



Universitat d'Alacant
Universidad de Alicante



e l l i s

UNIT
ALICANTE

Judging Books by Their Cover: The Impact of Facial Attractiveness on Humans and AI

Aditya Gulati

Thesis presented in fulfillment of the requirements
for the degree of Doctor of Philosophy by the

UNIVERSITY OF ALICANTE

With international mention

DOCTOR OF INFORMATICS

Advised by:

Nuria Oliver, *ELLIS Alicante*


Bruno Lepri, *Fondazione Bruno Kessler*

Miguel Ángel Lozano Ortega, *University of Alicante*

This research was supported by a nominal grant received at the ELLIS Unit Alicante Foundation from the **Regional Government of Valencia** in Spain (Convenio Singular signed with Generalitat Valenciana, Conselleria de Innovacion, Industria, Comercio y Turismo, Direccion General de Innovacion), along with grants from the **European Union's Horizon Europe** research and innovation programme (ELIAS; grant agreement 101120237), the **European Union's Horizon 2020** research and innovation programme (ELISE; grant agreement 951847) and by grants from **Intel** and the **Banc Sabadell Foundation**

This document was proudly typeset with L^AT_EX.

This work is licensed under a [Creative Commons](#) “Attribution-ShareAlike 4.0 International” license.



- Licensees may copy, distribute, display and perform the work and make derivative works and remixes based on it only if they give the author or licensor the credits (attribution) in the manner specified by these.
- Licensees may distribute derivative works only under a license identical ("not more restrictive") to the license that governs the original work. (See also copyleft.) Without share-alike, derivative works might be sublicensed with compatible but more restrictive license clauses, e.g. CC BY to CC BY-NC.)

Please see creativecommons.org/licenses/by-sa/4.0/ for greater detail.

Contact Details

Aditya Gulati
aditya@ellisalicante.org



e l l i s

UNIT
ALICANTE

ELLIS Alicante is the first Spanish unit within the **ELLIS** European network for research excellence. It is the only ELLIS unit that has been created as an independent non-profit research foundation, with the spirit of a scientific startup.

Our name, The Institute of Human-Centered AI, defines our mission: We firmly believe in the power of AI as an engine for progress and a key contributor to well-being. However, such a potential is by no means guaranteed and that's why the research of our foundation is so important. Our vision, mission and research have been awarded the 2022 Spanish Social Innovation Award by the Spanish Association of Foundations.

We aim to be a leading research lab on **ethical, responsible and human-centered AI**. We are the only ELLIS unit devoted exclusively to this topic.

At ELLIS Alicante, we address three important research areas:

- **AI to understand us**, by modeling **human behavior** using AI techniques both at the individual and aggregate levels. We focus on developing machine learning-based models of individual and aggregate human behavior. The practical applications are diverse, including the development of algorithms that generate recommendations for users or accurate and fair credit models to promote financial inclusion. At an aggregate level, we aim to model and predict human behavior on a large scale, at a country or region level, which allows addressing global challenges such as pandemics, detecting possible economic crises or responding to natural disasters. Our work during the COVID-19 pandemic is a good example of our work in this area.
- **AI that interacts with us**, via the development of intelligent, interactive systems, with a special focus on the development of smart phones, personal assistants and chatbots.
- **AI that we trust**, by tackling the ethical challenges posed by today's AI systems, such as algorithmic discrimination, violation of privacy, opacity, lack of veracity or subliminal manipulation of human behavior. Current AI algorithms are not perfect and have limitations that are important to identify and address in order to minimize the possible negative consequences of their use. In this area, we also investigate the societal and cultural impact of AI.

Abstract

Human perception, memory, and decision-making are shaped by a wide range of cognitive biases and heuristics that influence our actions and decisions. Despite their pervasiveness, these biases are rarely accounted for in the design of human–AI systems. This thesis argues that the future of effective human–AI collaboration will require the computational modeling, systematic understanding, and, in some cases, the deliberate replication of cognitive biases. To support this vision, we first introduce a framework that organizes known cognitive biases into five categories, illustrated with representative examples and accompanied by open research questions concerning their role in human–AI interaction.

The second part of the thesis focuses on one specific bias: the *Attractiveness Halo Effect* i.e., the tendency to associate positive traits (e.g., intelligence or trustworthiness) with physically attractive individuals, even when attractiveness is an irrelevant cue. We examine the prevalence of this bias in the digital age by assessing how beauty filters influence human ratings of various traits. This large-scale study, the most extensive to date, yielded a dataset of 924 images with high-quality ground-truth ratings for attractiveness, intelligence, trustworthiness, sociability, happiness, and other variables.

Building on these findings, we introduce the concept of *algorithmic lookism* i.e., the tendency of algorithms, particularly AI systems, to exhibit attractiveness-based discrimination – an important yet underexplored phenomenon. We empirically investigate this effect in seven large open-source Multimodal Large Language Models (MLLMs), demonstrating that these systems employ attractiveness as a cue in decision-making. Further, we analyze synthetically generated faces and find that images produced using positive trait descriptors are significantly more likely to depict attractive individuals. We also assess the downstream impact of such AI-generated images on related tasks. Collectively, our results reveal critical pathways through which cognitive biases can propagate in AI systems, underscoring the need for ethical and equitable design principles in their real-world deployment.

Resumen

La percepción, la memoria y la toma de decisiones humanas están determinadas por una amplia gama de sesgos cognitivos y heurísticos que influyen en nuestras acciones y decisiones. A pesar de su omnipresencia, estos sesgos rara vez se tienen en cuenta en el diseño de los sistemas humano-IA. Esta tesis sostiene que el futuro de una colaboración eficaz entre humanos e IA requerirá la modelización computacional, la comprensión sistemática y, en algunos casos, la replicación deliberada de los sesgos cognitivos. Para respaldar esta visión, primero presentamos un marco que organiza los sesgos cognitivos conocidos en cinco categorías, ilustradas con ejemplos representativos y acompañadas de preguntas de investigación abiertas sobre su papel en la interacción entre humanos e IA.

La segunda parte de la tesis se centra en un sesgo específico: el *attractiveness halo effect*, es decir, la tendencia a asociar rasgos positivos (por ejemplo, inteligencia o confianza) con personas físicamente atractivas, incluso cuando el atractivo es una señal irrelevante. Examinamos la prevalencia de este sesgo en la era digital evaluando cómo los filtros de belleza influyen en las valoraciones humanas de diversos rasgos. Este estudio a gran escala, el más extenso hasta la fecha, dio como resultado un conjunto de datos de 924 imágenes con valoraciones de alta calidad sobre el atractivo, la inteligencia, la honradez, la sociabilidad, la felicidad y otras variables.

A partir de estos hallazgos, introducimos el concepto de *algorithmic lookism*, es decir, la tendencia de los algoritmos, en particular los sistemas de IA, a mostrar discriminación basada en el atractivo físico, un fenómeno importante pero poco explorado. Investigamos empíricamente este efecto en siete grandes modelos de lenguaje multimodal (MLLM) de código abierto, demostrando que estos sistemas emplean el atractivo físico como indicador en la toma de decisiones. Además, analizamos rostros generados sintéticamente y descubrimos que las imágenes producidas utilizando descriptores de rasgos positivos son significativamente más propensas a representar a personas atractivas. También evaluamos el impacto posterior de estas imágenes generadas por IA en tareas relacionadas. En conjunto, nuestros resultados revelan vías críticas a través de las cuales los sesgos cognitivos pueden propagarse en los sistemas de IA, lo que subraya la necesidad de principios de diseño éticos y equitativos en su implementación en el mundo real.

Acknowledgments

This thesis has been four years in the making, and along the way I have had the privilege of meeting many wonderful individuals who have shaped not only this work, but also the person I have grown to be. I am truly grateful for your encouragement, guidance, and support throughout this journey and I want to thank you, from the bottom of my heart.

Dr. Nuria Oliver, you have been a constant pillar of support throughout my PhD, through both the highs and the lows. I consider myself incredibly fortunate to have had the opportunity to work with you and to witness firsthand what passionate, meaningful, and impactful research truly looks like. I am deeply grateful that you created a lab environment where I felt encouraged to express myself authentically and given the space to grow, explore, and find my own path, always knowing I had your support. I am very excited to continue working together.

Dr. Bruno Lepri, I am truly grateful for our many conversations and for your thoughtful advice whenever I needed it most. You brought a fresh, practical perspective that helped me stay grounded, and I am very glad to have had you co-supervising my thesis - providing balance and contributing to an inspiring and dynamic environment. I will miss our energetic and enjoyable discussions, and I am optimistic that our paths will cross again in the future.

Dr. Miguel Ángel Lozano, your calm presence and steady support made navigating the PhD, and all the accompanying bureaucracy, far easier to manage. Thank you for always being ready to answer the countless questions I had along the way and for taking care of so much behind the scenes.

Dr. Nick Chater and **Dr. Rita Cucchiara**, thank you for taking the time out of your busy schedules to review my thesis and for providing such thoughtful and encouraging feedback.

To all the shiny souls I met through ELLIS Alicante, you made moving away from home feel easy because you became my second family. **Piera**, thank you for always being there to listen to me and my dramatic renderings of the simplest of situations. **Gergely**, your grounded practical perspective on every situation is something I have always admired and I miss dearly. **Adrian**, your dedication has been something I've always looked up to and I'm glad to have been there to see it firsthand. **Julien**, I'm truly grateful for our never ending conversations and your honesty. **Lucile**, thank you for showing me what consistency and dedication looks like. **Erik**, you've always had our backs and have been someone I could always rely upon, for anything. **Rebeca**, **Cristina**, thank you for supporting me through every bureaucratic hurdle and for your care and encouragement beyond them. **Miriam**, **Ben**, **Mona**, **Kaylin**, I am glad our paths crossed during your visits to the lab. Your presence always brought new energy and made every day more interesting; you each hold a special place in my heart. **Ana and Maria**, thank you for being the best interns I could have asked for. **Luna**, even though you're a dog and probably will not read this, you never failed to make me smile, and for that I am endlessly grateful. **Kajetan**, there are few words that can truly express how grateful I am to have you in my life. You pushed me beyond what I believed were my limits and became a pillar of support I did not even know I needed. Thank you for every run, every

conversation, and all of your sage-like wisdom.

To everyone else I met over these 4 years in Alicante, thank you. Home for me is where I have people I care about and thank you for helping make Alicante my home. **Maria**, you have seen me at my most vulnerable, and I cannot express how grateful I am to have you in my life. Knowing you are always there brings me a sense of peace and comfort that I cherish. **Michelle**, through all our ups and downs, our conversations have always been among my favorites, and I truly look forward to the many more to come. **Roisin, Saya, Emerald, Kooshan, Zaki, Marianne, Philippe, Elena, Patricia, Paymon, Chloe** - thank you for showing me that there is no mold that we need to fit into and for always creating a space where we all fit in, with all our differences. To the wonderful people at **Sip and Wonder**, thank you for the endless coffees and for being one of my happy places in Alicante.

To all the people at MobS - **Pacho, Penzucco, Vero, Giova, Sebastiano, Franceschino, Maty, Cent, Max, Ciro, Laura, Nic, Martina, Marco, Apoorva** - thank you for making me feel like I belonged from the first day I walked in. Thanks for the endless laughs, the great advice and for being your most honest selves, from the day I met you. My time at MobS is something I will always hold dear in my heart. To **Ariandna, Filippo, Anna, Lucia, Elena D, Ulysse, Steven, Elena R, Sara**, thank you for helping me fall in love with Italy and for giving me reasons to keep coming back.

To **Martina, Tak, Alessandro, Giovanni** and everyone else at the journal club, thank you for all our spirited discussions and for accepting a computer scientist like me among your ranks. Our discussions have allowed me to grow and expand my knowledge in ways that would have been almost impossible without all of you.

To everyone from IIIT who have been a phone call away all this while, thank you. **Grover, Tejas, Malpani, Shobhit** you've been there with me through every step of my PhD, including applying for one. Thanks for sticking by my side and for picking me up every time I fell over. **Tanu, Ashish, Vishesh, Tanmay, Sarthak** thank you for welcoming me every time I was back with open arms - meeting you always made me feel like I never left.

And to my family, the biggest thank you from the bottom of my heart, for always always supporting me, believing in me and showing me your love in every way. **Anu**, you've been my rock, and moving away from you was one of the hardest things I've done, but I knew you were always there. To **mummy** and **papa**, thank you for your support with everything, and for believing in me, even through the times I didn't believe in myself. This thesis would not have been possible without the three of you.

Contents

Abstract	vii
Resumen	ix
Acknowledgments	xi
1 Introduction	1
1.1 The Cognitive Biases vs Heuristics Debate	4
1.2 Contributions	5
1.3 List of Publications	5
2 Human Cognitive Biases and AI: A Framework	7
2.1 Presentation biases	9
2.2 Interpretation biases	10
2.3 Value attribution biases	12
2.4 Recall biases	14
2.5 Decision biases	15
2.6 Discussion	17
2.7 Classification of Other Biases	18
2.8 Conclusion	26
3 The Attractiveness Halo Effect in Human Decisions	27
3.1 Introduction	27
3.2 Results	31
3.2.1 Beauty Filters and Attractiveness	31
3.2.2 Beauty Filters and the Attractiveness Halo Effect	33
3.2.3 Impact of the Raters on the Attractiveness Halo Effect	34
3.2.4 Do Beauty Filters Mitigate the Attractiveness Halo Effect?	37
3.3 Discussion	38
3.4 Methods	42
3.4.1 Study Participants	42
3.4.2 Experimental Stimuli	43
3.4.3 Procedure and Design	44
3.4.4 Measures	46
3.4.5 Analysis	46
3.4.6 Data Accessibility	48

4	Algorithmic Lookism	49
4.1	Lookism: A cognitive bias perspective	50
4.2	Lookism and computer vision	50
4.2.1	Beauty Filters and Lookism	51
4.2.2	Algorithmic Lookism	51
4.3	Challenges	53
4.4	Conclusion	54
5	Algorithmic Lookism in MLLMs	55
5.1	Introduction	56
5.2	Related Work	57
5.2.1	Biases in Large Language Models (LLMs)	57
5.2.2	Biases in Multimodal LLMs (MLLMs)	58
5.3	Methodology	59
5.3.1	Models	59
5.3.2	Inputs	60
5.3.3	Scenarios	61
5.3.4	Model Evaluation	63
5.3.5	Problem Formulation and Metrics	64
5.4	Results	66
5.5	Discussion	69
5.6	Conclusion	72
6	Algorithmic Lookism in Synthetically Generated Faces	73
6.1	Introduction	73
6.2	Methods	75
6.2.1	Synthetic Face Dataset Creation	75
6.2.2	Algorithmic Lookism in Synthetically Generated Faces	76
6.2.3	Impact on Downstream Tasks: Gender Classifiers	78
6.2.4	Algorithmic Lookism Mitigation	79
6.3	Results	79
6.3.1	RQ1: Do faces generated by Stable Diffusion 2.1 tend to associate attractiveness with positive attributes?	80
6.3.2	RQ2: Are downstream models, specifically gender classifiers, sensitive to the attractiveness of synthetically generated faces?	82
6.3.3	RQ3: Can the AHEAD dataset aid in the mitigation of these effects in downstream classifiers?	84
6.4	Discussion	86
6.5	Conclusion	89
7	Conclusion	91
	Appendices	93
A	Influence of Participant Characteristics on Ratings	95
A.1	Impact of Rater Gender	95

A.2	Impact of Self-Perceived Attractiveness on Judgments of Attractiveness	95
B	Statistical Analysis Conducted during the Creation of the AHEAD Dataset	99
B.1	Ordered Stereotype Models	99
B.2	Model Selection	100
B.3	Factor Analysis	101
B.4	Evaluation of the Saturation Effect in the Halo Effect	102
B.4.1	Method A: Piece-wise linear fit	103
B.4.2	Method B: Fitting a Log Curve	103
B.5	Linear Mixed Models including Rater Effects	104
B.6	Partial R^2 in the Linear Mixed Models	104
B.7	Computation of Fractional Change in EMM	104
C	Impact of Beauty Filters on the Perception of Physical Attributes	107
C.1	Impact of Beauty Filters on Perceptions of Age, Gender and Ethnicity	107
C.2	Impact of Beauty Filters on Perceptions of Femininity and Unusualness	108
D	Characteristics of Raters in the Creation of the AHEAD Dataset	111
E	Design of the Survey Used in the Creation of the AHEAD Dataset	113
E.1	The Survey Tool	113
E.2	Instructions	114
E.3	Questions	114
E.3.1	Background Information	116
E.4	Attentiveness Checks	117
F	Impact of Beauty Filters on Perception of Attributes During the Creation of the AHEAD Dataset	119
F.1	Impact of Filters on Dependent Attributes	119
F.1.1	Mediation of Age and Gender on the Impact of Filters on Dependent At- tributes	120
G	Impact of Demographic Factors on the Attractiveness Bias in MLLMs	123
G.1	Impact of Age on the Attractiveness Bias	123
G.2	Impact of Race on the Attractiveness Bias	123
H	Algorithmic Lookism in MLLMs: Scenarios	127
I	Lookism in MLLMs: The Attractiveness Halo effect and the Impact of Gender, Age, and Race	131
I.1	Attractiveness Halo Effect in Stereotyped Traits	131
I.2	Gender Bias in the Gender Stereotyped Jobs Scenarios	131
I.3	Racial Bias in the Race Stereotyped Jobs Scenarios	132
I.4	Demographic Biases across the different scenario types	132

J	Evaluating Attractiveness Classifiers Trained on the AHEAD Dataset	151
J.1	Classification Results	151
J.1.1	InceptionNet	151
J.1.2	ResNet50	152
J.2	Performance On CelebA	152
K	Resumen en castellano	153
K.1	El debate entre los sesgos cognitivos y la heurística	160
K.2	Contribuciones	161
K.3	Lista de publicaciones	161
K.4	Conclusión	162
	Bibliography	165

Chapter 1

Introduction

Human progress will increasingly depend not only on what Artificial Intelligence (AI) can accomplish on its own, but also on how effectively it collaborates with and learns from insights about the human mind. Decades of psychological research provide a rich understanding of human behavior in individual and social contexts, offering valuable resources for anticipating human decisions and behaviors. Incorporating this knowledge into AI design is crucial not only for improving collaboration between humans and machines, but also for uncovering the subtle ways in which AI systems themselves may adopt, amplify, or be shaped by human biases when making decisions about people. While integrating the full breadth of psychology into AI remains a daunting challenge, a more tractable entry point lies in the study of cognitive biases – systematic patterns of deviation from “rationality” that occur when we process, interpret or recall information from the world, leading to inaccurate judgments, illogical interpretations and perceptual distortions.

Since the 1970s, research in social psychology, cognitive science, and behavioral economics has systematically investigated the seemingly irrational elements of human decision-making that shape our everyday interactions [TK81; AJ08; KFSR93] which shape our interactions not only with other individuals and objects but also, increasingly, with AI systems. Despite this, relatively little is known about how cognitive biases influence human-AI interaction or how such biases may be reproduced, amplified, or even generated by AI systems themselves. In this thesis, we claim that filling this gap is necessary to build trustworthy AI systems that collaborate with humans.

The discourse on cognitive biases and Artificial Intelligence has predominantly been framed through the lens of mitigation, primarily in the context of societal biases, such as gender discrimination or racism. An area of research within AI where cognitive biases have been studied in recent years is Large Language Models (LLMs). Their rapid adoption and unprecedented performance in natural language processing tasks have enabled researchers to replicate studies previously conducted with human participants, now employing LLMs instead. The findings have been mixed: While certain cognitive biases, such as the anchoring and framing effects, have been observed in LLM decisions [TF23], others, such as the status quo bias, are not consistently reproduced by LLMs [ELA+24]. These findings are based on decades of research on human cognitive biases, which served as a foundation for studying the existence of similar patterns in AI systems. Although such studies of cognitive biases in LLMs are valuable, we argue that a broader perspective is needed, one that not only replicates experiments conducted with human participants but also examines how knowledge of human cognitive biases might inform the design, interpretation, and deployment of AI systems, particularly when such systems interact with humans.

Knowledge of human cognitive biases could also enable their constructive use in human-AI in-

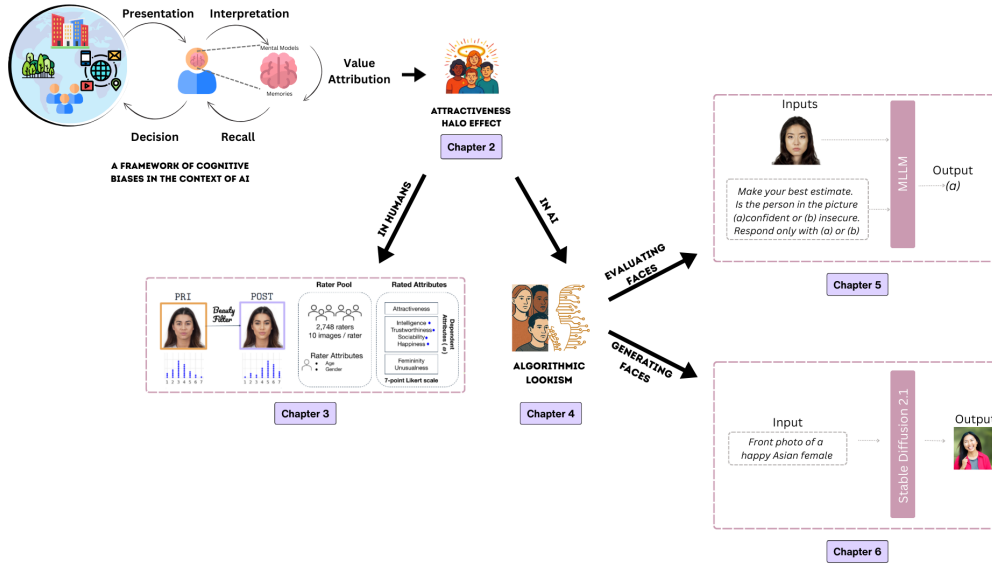


Figure 1. Overview of the contributions of this thesis. We begin by proposing a framework for categorizing cognitive biases in a manner that supports AI designers in developing systems that work both with and for humans. To illustrate this framework, we focus on the Attractiveness Halo Effect (AHE) as a case study, demonstrating how cognitive biases can influence both human perception and AI systems. First, we examine how the AHE shapes human behavior in digital spaces and investigate whether AI-based beauty filters alter or amplify this bias. Building on the data collected in this context, we then evaluate the implications of the AHE for AI decision-making—both in systems that make judgments about humans and in those that generate synthetic images of people.

teraction, fostering more effective, natural and intuitive collaboration. For example, Bucinca et al. [BMG21] demonstrated that machine-generated explanations are not always effective, as users frequently rely on heuristics rather than on the provided explanations when deciding whether to trust an AI system. They also reported an improvement in user decision-making by leveraging cognitive forcing functions that guided users toward more informed judgments about when to rely on AI-generated outputs. Their proposed approach shifts the focus from cognitive bias mitigation to system design that is aligned with a specific cognitive bias to enhance human-AI interaction. Furthermore, in certain contexts, AI systems may benefit from mimicking cognitive biases to improve their performance. Taniguchi et al. [TSS17] exemplified this by developing a modified Naive Bayes classifier that leveraged symmetry and the mutual exclusion bias, demonstrating that their approach outperformed state-of-the-art spam classification methods on small and biased datasets. These examples collectively underscore the dual role of human cognitive biases: as both diagnostic tools for evaluating AI behavior and as conceptual foundations for human-centric system design.

This thesis advances a line of inquiry into how established knowledge of human cognitive biases can be systematically incorporated into the design and evaluation of artificial intelligence systems. A visual representation of the contributions is depicted in Figure 1. As an initial step, we propose a taxonomy of known cognitive biases tailored to the needs of collaborative human-AI system design. While several taxonomies of cognitive biases exist [SRRT16; WL19; DFP+20; KGKK18], these are typically either task-specific or oriented around hypothesized cognitive origins. Although valuable for psychological research, such classifications offer limited practical utility for AI system designers. Our proposed framework instead structures cognitive biases along the human perception–decision-making cycle, highlighting where specific biases are most likely to

emerge in human–AI interactions. This framework is described in detail in Chapter 2.

The remainder of the thesis focuses on a single exemplary cognitive bias: the Attractiveness Halo Effect (AHE) [DBW72]. The AHE refers to the human tendency to associate positive traits (such as intelligence or trustworthiness) with physically attractive individuals, even when attractiveness is an irrelevant factor. While extensively documented in traditional contexts, comparatively little is known about how the AHE manifests in contemporary digital environments, particularly in the presence of AI-driven beauty filters. These filters, now ubiquitous on social media platforms, are widely used to alter (*beautify*) our physical appearance (mainly the face), often in ways that opaque to users.

We focus on the AHE for three key reasons. First, it represents a highly robust and well-studied bias in the psychology literature, providing a solid empirical foundation for methodological innovation. Second, it is of tremendous societal relevance: the widespread adoption of AI-based visual manipulation tools means that attractiveness-related judgments now frequently occur in digitally mediated contexts, potentially amplifying or altering the bias. Third, the AHE has significant implications for AI ethics and fairness, as attractiveness-based inferences when learned or reproduced by algorithms can lead to systematic discrimination in domains such as hiring, recommendation systems, education, judicial sentencing, medical diagnosis and social evaluation. While significant research has studied biases in AI systems related to gender or race [WFVL19; WQK+20; YAAB20; HFBK24], relatively little work has examined the role of facial attractiveness on algorithmic decisions. Studying the AHE in this context thus offers both theoretical insights into the propagation of cognitive biases in human–AI systems and practical guidance for the equitable design of such systems. A central challenge in studying attractiveness lies in its inherently subjective nature. Throughout this thesis, each chapter details the measures taken to account for this subjectivity, while demonstrating that the results nevertheless reveal consistent and generalizable trends across both human judgments and AI system outputs.

We begin our investigation of the AHE in digital contexts by conducting the largest known empirical study of this phenomenon to date. In this study, 2,743 participants evaluated facial images of 462 individuals with and without an AI-based beauty filter applied. The study addressed three key questions: whether beauty filters reliably increase perceived attractiveness; whether such alterations trigger an attractiveness halo effect, leading to shifts in perceived traits unrelated to appearance; and whether the strength or direction of this effect varies across demographic groups and to which degree depends on the observer. Beyond offering novel insights into how the AHE manifests in digitally mediated human judgments, this work contributes a rigorous empirical benchmark dataset for studying similar biases in generative AI models: a dataset of human faces with and without beauty filters applied accompanied with high quality human ratings of attributes related to the attractiveness halo effect. This dataset, henceforth referred to as the **AHEAD** dataset (the **A**tttractiveness **H**alo **E**ffect **A**tribution **D**ataset), is a first-of-its-kind empirical benchmark of human faces with the *same* person in two attractiveness settings and is explicitly designed to study the attractiveness halo effect and related biases in AI systems. The study design, methodology, and results of the AHE in humans are discussed in Chapter 3.

Building on these findings, we propose the concept of *algorithmic lookism*, which we define as the systematic tendency of AI algorithms to reproduce or amplify attractiveness-based discrimination. We argue that attractiveness is not a trivial aesthetic factor but a variable that can fundamentally shape the outputs of large-scale AI systems. Despite its potential to distort decision-making in high-stakes contexts, as seen in decisions made by humans [RPT92; BKTR08; CK85; HSC03], this form of bias has received comparatively little scholarly and technical attention. Addressing algorithmic

lookism is therefore not only a matter of fairness but also of reliability and trust in AI systems. As we detail in Chapter 4, this bias is especially relevant in models that process visual inputs, such as Multimodal Large Language Models (MLLMs), and in text-to-image (T2I) generative systems like Stable Diffusion.

In Chapter 5, we present an empirical evaluation of seven open-source MLLMs on the AHEAD dataset when exposed to 91 systematically designed scenarios. Our results demonstrate that MLLMs are not neutral interpreters of visual data: they rely on attractiveness cues when reasoning about facial images. More troublingly, they replicate the Attractiveness Halo Effect and systematically associate attractive individuals with positive traits. Our finding suggests that algorithmic lookism is not an incidental error but a structural bias embedded in how these systems process human faces.

Chapter 6 extends this critique to text-to-image generative systems, examining synthetically generated faces by Stable Diffusion – a popular, open-source text-to-image generative algorithm– produced with positive and negative trait descriptors. We find clear evidence of an attractiveness bias in the generation of the faces: faces generated with positive traits (*e.g.*, intelligent, trustworthy, sociable, happy) tend to be more attractive than faces generated with negative (*e.g.*, unintelligent, untrustworthy, unsociable, unhappy) traits. Importantly, this bias does not remain confined to generation but impacts downstream tasks. For instance, gender classification models perform unevenly on facial images generated with negative descriptors, highlighting how attractiveness-related distortions cascade across the AI pipeline. We further show that while targeted fine-tuning offers a partial mitigation strategy, more systematic approaches will be required to prevent such biases from entrenching discrimination at scale.

Finally, in Chapter 7, we discuss the implications of our findings which demonstrate that AI systems do not merely inherit human biases but can reconfigure and amplify them in novel ways. Our work underscores the urgent need for the AI community to treat cognitive biases –long studied in psychology and behavioral economics– as a core concern for the development of fair, transparent, and trustworthy AI systems. We argue that measuring and mitigating these biases, whether attractiveness-related or otherwise, is essential not only for technical robustness but also for ensuring that human–AI interaction supports, rather than undermines, equitable decision-making in society.

1.1 The Cognitive Biases vs Heuristics Debate

Parallel to the work on cognitive biases, a complementary perspective emerged in the early 2000s, grounded in the notion of bounded rationality [Sim90a]. Scholars in this tradition have been critical of the cognitive biases paradigm, arguing that what are often labeled as biases may instead be understood as heuristics—adaptive shortcuts that are not limitations but fundamental to how humans make decisions [Vra00].

This divergence of perspectives has given rise to the more recent nudging versus boosting debate [HG17]. Nudging, which builds on the biases framework, emphasizes modifying environmental cues to steer people toward better decisions. Boosting, aligned with the heuristics perspective, critiques nudging as paternalistic and emphasizes instead the empowerment of individuals—helping them develop the knowledge, skills, and awareness to make better decisions autonomously. Both approaches offer valuable insights, and the debate between them remains active, with scholars continuing to contest the balance between external guidance and individual agency in behavior change.

For the purposes of this thesis, however, resolving this debate is not the primary concern. What matters most in the context of AI design is the recognition that people exhibit consistent, predictable

patterns of judgment and behavior, regardless of whether these are framed as “biases” or “heuristics.” In any case, there is agreement on the existence of such regularities, even if they diverge on their normative interpretation. In this work, we therefore adopt the term cognitive biases as a convenient shorthand to describe these patterns, while acknowledging its contested nature and the important contributions of the heuristics and boosting perspectives.

1.2 Contributions

Below, we summarize the primary contributions of the thesis:

- Development of a structured framework for organizing cognitive biases in a manner directly applicable to the design of collaborative human-AI systems (Chapter 2).
- Creation of a high-quality dataset of faces (AHEAD) to evaluate the attractiveness halo effect, along with empirical insights into how beauty filters shape and amplify this human cognitive bias (Chapter 3).
- Introduction of the concept of algorithmic lookism, defined as attractiveness-based discrimination in AI decision-making, and articulation of its implications for algorithmic fairness (Chapter 4).
- Empirical evidence that Multimodal Large Language Models (MLLMs) exhibit an attractiveness bias and reproduce the Attractiveness Halo Effect when reasoning about facial images (Chapter 5).
- Empirical evidence that text-to-image generative models, specifically Stable Diffusion, encode an attractiveness bias when creating images of human faces, with downstream consequences for tasks such as gender classification (Chapter 6).

1.3 List of Publications

The core research presented in this thesis was disseminated through peer-reviewed articles in academic conferences and journals. These articles form the basis of the empirical and conceptual contributions summarized above, and are listed chronologically below:

- Aditya Gulati, Miguel Angel Lozano, Bruno Lepri, and Nuria Oliver. “Biased: Bringing irrationality into automated system design.” AAAI Fall Symposium 2022 on Thinking Fast and Slow and Other Cognitive Theories in AI, arXiv:2210.01122 (2022) [GLLO23]
- Aditya Gulati, Marina Martínez-Garcia, Daniel Fernández, Miguel Angel Lozano, Bruno Lepri, and Nuria Oliver. “What is beautiful is still good: the attractiveness halo effect in the era of beauty filters.” Royal Society open science 11, no. 11 (2024): 240882 [GMF+24a]
- Aditya Gulati, Bruno Lepri, and Nuria Oliver. “Lookism: The overlooked bias in computer vision.” ECCV 2024 workshop on “Fairness and ethics towards transparent AI: facing the chalLEnge through model Debiasing”, FAILED’25, arXiv:2408.11448 (2024) [GLO24]

- Miriam Doh, Aditya Gulati, Matei Mancas, and Nuria Oliver. “When Algorithms Play Favorites: Lookism in the Generation and Perception of Faces.” Fourth European Workshop on Algorithmic Fairness, EWAF’25, arXiv:2506.11025 (2025) [[DGMO25](#)]
- Aditya Gulati, Moreno D’Incà, Nicu Sebe, Bruno Lepri, and Nuria Oliver. “Beauty and the Bias: Exploring the Impact of Attractiveness on Multimodal Large Language Models.” Eighth AAAI/ACM Conference on AI, Ethics and Society, AIES’25, arXiv:2504.16104 (2025) [[GDS+25](#)]

Chapter 2

Human Cognitive Biases and AI: A Framework

Chapter summary

Human perception, memory and decision-making are impacted by numerous cognitive biases and heuristics that influence our actions and decisions. Despite their pervasiveness, cognitive biases are typically not considered in today’s AI systems that model, interact and augment humans. In this thesis, we claim that the future of human-AI collaboration will require computationally modeling, understanding and possibly replicating cognitive biases. To advance this vision, we propose in this chapter a framework that organizes cognitive biases according to the stages in the human decision-making cycle, classifying them into five categories. We illustrate each category with representative examples of biases that are relevant in the context of human-AI systems and identify key open research questions at the intersection of cognitive biases and human-AI collaboration. We aim for this framework to serve as a foundation for the design and evaluation of AI systems that model, interact and augment humans, developed with a principled understanding of well-established human cognitive patterns.

This chapter is based on the paper:

[GLLO23] Aditya Gulati, Miguel Angel Lozano, Bruno Lepri, and Nuria Oliver. “Biased: Bringing irrationality into automated system design.” AAAI Fall Symposium 2022 on Thinking Fast and Slow and Other Cognitive Theories in AI. arXiv:2210.01122 (2022)

Given the multitude of known cognitive biases and the complex, varied nature of human-AI interaction as previously described, incorporating insights from cognitive science into AI system design presents a non-trivial challenge. This chapter aims to provide an initial framework to address this challenge by aligning well-established human cognitive biases with distinct stages of the human decision-making cycle which are meaningful in the context of human-AI collaboration.

Several taxonomies of cognitive biases have been proposed in the literature, particularly in specific domains, such as medical decision making [SRRT16], tourism [WL19], fire evacuation [KGKK18] or visualization [DFP+20]. Alternative taxonomies classify biases based on their underlying phenomenon [STW08]. However, in the absence of a widely accepted theory regarding the origins of cognitive biases [PP04], taxonomies grounded in presumed sources risk being speculative or misleading.

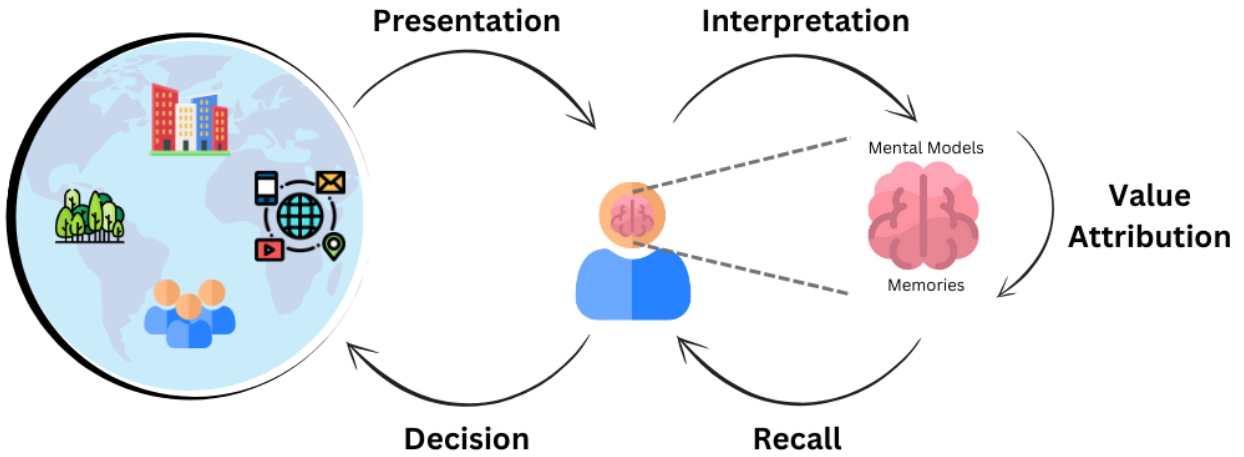


Figure 2. Five stages of the human perception, interpretation and decision-making process that constitute the categories in the proposed framework of cognitive biases. Icons used in the image are from Flaticon.com

Moreover, theoretical taxonomies that focus on classifying biases by underlying cognitive mechanisms often lack direct applicability to the design of human-centered AI systems. For AI systems to effectively collaborate with humans, what is needed is a practical framework, one that organizes cognitive biases according to where they manifest within the human decision-making process. Such an application-oriented classification would better support AI designers in anticipating, understanding, and accommodating human decision patterns in the systems they create. To support this need, we propose a framework, depicted in Figure 2, that structures cognitive biases into five categories according to five stages in the perception-to-decision-making cycle: Presentation, Interpretation, Value Attribution, Recall and Decision. This framework enables the categorization of existing biases and focuses on their relevance in human-AI collaboration across the multitude of potential scenarios where humans and AI might interact. The left part of Figure 2 represents the physical world that we perceive, interpret, and interact with, whereas the right part represents the internal models and memories that we create based on our experiences.

The proposed framework is designed to highlight potential pitfalls and opportunities when designing AI systems that interact, collaborate, and make decisions with humans. Unlike traditional taxonomies, which strive for comprehensive theoretical coverage, the proposed framework adopts a pragmatic and context-sensitive approach: it curates a subset of cognitive biases that are particularly salient in real-world human-AI collaboration. To ensure both scientific rigor and applied relevance, the inclusion of specific cognitive biases was guided by two complementary criteria: (1) *Empirical robustness*, reflected through the influence and replicability of each cognitive bias, as indicated by the citation impact of the original publication introducing the specific bias and subsequent validation by other scholars; and (2) *Relevance* of the cognitive bias for human-AI collaboration, assessed through a structured analysis of how each cognitive bias may affect –or be affected– by human-AI collaboration.

The main purpose of the proposed framework is to aid in the development of AI systems to enhance their collaboration with humans by focusing on the observable consequences of cognitive biases, rather than their underlying cognitive mechanisms. It provides AI designers with tools to anticipate and address consistent patterns observed in human decision making. Furthermore, it offers a structured approach to designing cognitively-aware AI systems and organizing research on how cognitive biases influence human-AI interaction. Thus, rather than aiming for completeness,

we intentionally focus on cognitive biases with high translational potential, where understanding or leveraging the bias could improve the performance, interpretability and user acceptance of AI systems. Full details of the bias selection methodology, along with supporting references and individual descriptions of each cognitive bias, are provided in the Supplementary Material, which also includes an inventory of additional cognitive biases not covered in the main text.

In the remainder of this chapter, we describe the five categories of cognitive biases as per the proposed framework, and provide concrete examples of four cognitive biases within each category that are relevant for human-AI collaboration. We additionally classify majority of the well-documented cognitive biases into the five categories to demonstrate the potential breadth of the proposed framework.

2.1 Presentation biases

Definition Presentation biases consist of systematic deviations in human judgment and decision-making influenced by how information is structured or conveyed. Biases in this category provide valuable insights into how AI systems can structure the presentation of information to support human decision-making.

Relevant examples for human-AI collaboration

1. *Decoy effect* [HPP82; HY14; WJS+18; JH95]: The Decoy Effect is a cognitive bias where the introduction of a deliberately inferior third option influences individuals to prefer one choice over another, even if their initial preference was different. The decoy is usually clearly worse than one option but only slightly worse or comparable to another, thereby making the dominant option appear more attractive. For example, when choosing between a \$10 basic subscription and a \$20 premium subscription, adding a \$19 subscription with fewer benefits than the \$20 plan can lead more users to select the premium option, even if they initially favored the cheaper one.
2. *Framing effect* [TK81; GORS09; LSG98; GZY+13]: The Framing Effect is a cognitive bias where the way statements are framed significantly influences decision-making and perceived value. Individuals tend to react differently depending on whether choices are framed in terms of potential gains or potential losses, even if the underlying facts are identical. For example, people are more likely to opt for a medical treatment described as having a “90% survival rate” than one described as having a “10% mortality rate”, although both statements convey the same information.
3. *Anchoring effect* [TK74; NAG19; YR22]: The Anchoring Effect is a cognitive bias in which individuals rely heavily on an initial piece of information, or an “anchor”, when making decisions, even if the anchor is arbitrary or unrelated. Once an anchor is set, subsequent judgments tend to be made by adjusting away from that point. For instance, when asked to estimate the price of a bottle of wine, individuals first shown a high-priced wine list tend to give higher estimates than those shown a low-priced list, even if they are unaware of the influence.
4. *Pseudocertainty effect* [TK89; HN09; Bur13]: The Pseudocertainty Effect is the tendency of humans to incorrectly estimate the certainty of statements in a multi-stage decision making

process. In an experiment by Tversky et al. [TK81] participants were asked to choose between two options, and their preferences were observed to vary depending on whether the decision problem was framed as a single-stage or a multi-stage process. Notably, when the problem was presented in a multi-stage format, participants exhibited a preference for the option with lower expected utility. This preference arose because the option was framed as a guaranteed outcome in the second stage, thereby inducing an illusion of certainty despite the overall uncertainty of the decision process.

Impact in human-AI collaboration From a human-centric perspective, these cognitive biases could inform how human-AI collaboration systems present information to their users providing mechanisms to encourage them to make healthier choices and more prudent decisions. AI agents could also serve as protective mechanisms against manipulative nudges by detecting and counter-acting biased presentations of information.

For instance, AI-based human decision-support systems could leverage the framing effect in healthcare for treatment option selection, in hiring for candidate ranking and in pricing for online shopping by changing how information is presented. Leveraging the framing effect in intelligent interface design can have a profound impact on user trust and acceptance of AI systems, as reported in previous research [SCDG16; KS22]. Thoughtful framing strategies, such as emphasizing benefits rather than risks or structuring messages to align with user values, can enhance the credibility and trust of AI systems and improve user adherence to system recommendations.

The anchoring effect highlights the importance that the initial piece of information presented to users has, suggesting that intelligent interfaces could strategically order information to support human decision-making outcomes to lead to better decisions in consequential domains. For example, AI chatbots and virtual assistants could use anchors to guide conversations, such as suggesting an initial investment amount in finance apps based on user preferences, making them more likely to commit to a similar range within the economic possibilities of the user to avoid over-spending.

The pseudocertainty effect makes users more likely to trust AI recommendations when they are framed as certain single-stage information rather than probabilistic multi-stage predictions. This bias could be especially impactful in cybersecurity, finance and risk assessment, where users may over-rely on “safe” choices that are generated by an AI system while ignoring alternative and potentially more beneficial strategies. By acknowledging the pseudocertainty effect, designers of AI systems could minimize multi-stage information presentation, thereby increasing transparency and reducing potential misunderstandings among users interacting with AI-driven systems.

While presentation biases offer valuable insights for designing AI systems that align with their users’ interests and well-being, it remains crucial to balance such optimizations with ethical considerations to ensure that these techniques are employed to enhance user agency and support better decision-making, rather than to subliminally manipulate behavior with ulterior motives that are not aligned with the well-being or best interests of the users.

2.2 Interpretation biases

Definition Interpretation biases arise due to misinterpretations of information and describe systematic errors in human reasoning when interpreting probabilities, events, or rewards. Unlike presentation biases, these biases are not affected by how information is presented by an external agent, but rather by how human decision makers draw inferences from presented information. These

biases significantly influence decision-making processes, including in scenarios of consequential importance in people's lives, such as healthcare, electoral processes or finance, and are also likely to impact human-AI collaboration.

Relevant examples for human-AI collaboration

- *Conjunction fallacy* [TK83; TBO04; WM08; LSRO02]: In certain situations, humans see the conjunction of two events as being more likely than any one event individually. A classic example is the “Linda problem”[TK83], where participants are told Linda is described as deeply concerned with social justice, and then asked to judge whether “Linda is a bank teller and active in the feminist movement” as more likely than “Linda is a bank teller”, even though the conjunction of two conditions must always be less probable than a single condition.
- *Base Rate fallacy* [BS07; Bar80]: Humans have a tendency to ignore the base rate information when making decisions. For example, if told that a person is shy and detail-oriented, individuals may judge it more likely that the person is a librarian than a salesperson, even if there are far more salespeople than librarians in the population. Conversely, when base rate information is made salient and emphasized—such as explicitly stating the vast numerical difference between groups—people are more likely to adjust their judgments appropriately, integrating base rates into their reasoning.
- *Gamblers fallacy* [TK74; Gol97; BL10; CMS16]: The gamblers fallacy is the tendency of humans to overvalue the impact of past events when predicting the outcome of independent future events. For example, after witnessing a coin land on heads five times in a row, a person might wrongly believe that tails is now more likely to occur on the next toss, despite the probability of heads or tails remaining exactly 50% for each flip.
- *Hyperbolic discounting effect* [Tha81; HAO18; Ain75]: The hyperbolic discounting effect is the tendency of individuals to disproportionately prefer smaller, immediate rewards over larger, delayed rewards, even when waiting would yield greater benefits. For example, a person might choose to receive \$50 today rather than \$100 in a month, even though waiting provides double the reward. However, when both rewards are framed as occurring in the future—such as choosing between \$50 in a year versus \$100 in thirteen months—individuals are more likely to wait for the larger payoff, revealing greater patience when immediate gratification is not an option.

Impact in human-AI collaboration Understanding the patterns underlying human information processing can significantly influence the design of AI systems that collaborate with humans. Biases such as the *base rate fallacy* can have profound implications, especially in consequential scenarios like healthcare. When AI outputs appear precise, users may trust them without questioning whether they align with real-world probabilities. This can lead to misinformed decisions, discrimination, and errors in judgment, highlighting the importance of statistical literacy, transparency in AI models, and human oversight to prevent flawed conclusions based on misleading AI-driven insights. For example, an AI algorithm designed to support medical diagnosis might detect a rare disease with high accuracy. Despite the algorithm's high diagnostic accuracy, conclusions drawn from its outputs may become misleading if users fail to account adequately for the probability of false positives—a likelihood that can be particularly high in cases involving rare conditions.

To mitigate the effects of such interpretation biases, designers should leverage clear representation of probabilities, data visualizations that effectively communicate relative likelihoods, and tools designed to foster critical thinking. Human-AI collaboration presents a unique opportunity to assist users in recognizing and overcoming these types of biases. By integrating automated mechanisms to detect potential misinterpretations, AI-driven interfaces can provide real-time warnings and assist users in making more informed decisions. Interpretation biases often emerge in specific circumstances which makes such a detection easier. For example, the *conjunction fallacy* occurs when individuals assess the probability of compound events, which leads to incorrect probability estimations. Explanatory prompts or uncertainty visualizations could help guide users towards more rational probability assessments.

Beyond immediate decision-making, interventions in human-AI design that consider this type of biases can also enhance long-term planning. For instance, the *hyperbolic discounting* bias leads individuals to prioritize immediate rewards over long-term benefits, often resulting in choices that deviate from the users self-stated preferences or have lower long-term positive impact. Human-AI collaborative tools designed to assist with long-term decision-making –such as intelligent calendars or financial planners– could help users navigate complex problems by breaking them down into manageable components and offering timely recommendations. Such systems would not only support improved short-term decisions, but also encourage more sustainable and rational long-term planning.

2.3 Value attribution biases

Definition Value attribution biases emerge when humans assign values to entities or ideas that are not based on an underlying factual reality but based on cognitive shortcuts. Like interpretation biases, value attribution biases are independent of how information is presented and thus are different from presentation biases. They focus on subjective value assignment of external entities instead of the evaluation of probabilities and are hence differentiated from interpretation biases. Understanding the conditions under which these value assignment patterns emerge, as well as their implications for human-computer interaction, is crucial in the design of AI systems to effectively collaborate with humans.

Relevant examples for human-AI collaboration

- *Halo effect* [DBW72; NW77; GG16; LS74]: The Halo Effect is a cognitive bias in which an individual's positive traits in one domain influence the perception of their abilities or character in unrelated areas. For example, a physically attractive person may also be perceived as more intelligent or trustworthy, even without evidence supporting those attributes.
- *IKEA effect* [NMA12; RLH+19; BGM22]: The IKEA Effect is a cognitive bias in which individuals place disproportionately high value on products or outcomes they have partially created themselves, compared to similar outcomes produced by others. For example, a person may perceive a piece of furniture they assembled from a kit as more valuable or higher quality than a similar pre-assembled item, simply because of the effort they invested.
- *Risk aversion bias* [Pra78; Sta10; FSL+78; RH01]: The Risk Aversion Bias is a cognitive tendency where individuals prefer options that minimize potential losses over those that maximize potential gains, even when the riskier option offers a higher expected utility. This bias

becomes particularly pronounced under conditions of high uncertainty. For example, an investor might choose a low-yield government bond over a volatile stock that statistically offers a much higher long-term return, simply to avoid the discomfort of potential short-term losses. In contrast, when uncertainty is framed more clearly – such as presenting the long-term probabilities of gains explicitly – individuals are more likely to make decisions aligned with expected value.

- *Social desirability bias* [CM60; HCP+95; Har06; SG09]: The Social Desirability Bias is a cognitive bias in which individuals respond to surveys or questions in a manner they believe will be viewed favorably by others. This bias often leads to the overreporting of socially acceptable behaviors and attitudes. For example, in a survey about charitable giving, respondents may overstate how frequently they donate to appear more altruistic. Conversely, when anonymity is guaranteed respondents are more likely to provide honest answers.

Impact in human-AI collaboration A fundamental question in this domain is whether value attribution biases also emerge when individuals engage with AI systems as they do in human-to-human interactions. For example, does the *social desirability bias* operate similarly in interactions with AI chatbots or agents as it does with humans? When users converse with a chatbot, does their awareness of the system as a non-human agent make them more willing to be vulnerable and open, or does it instead make them less trusting? The conditions influencing these dynamics remain open research questions.

Another pertinent example is the *attractiveness halo effect*. Research has established that attractive individuals tend to be perceived more favorably not only in terms of their physical appearance, but also in unrelated domains, such as competence, intelligence, trustworthiness and sociability, even when such domains are unrelated to beauty. However, it remains unknown whether and to which degree this bias extends to intelligent agents, such as avatars or virtual assistants. Do users evaluate interactions with an AI agent differently based on the perceived attractiveness of its visual representation? This question has significant implications for the design of intelligent digital interfaces but remains largely unexplored.

Furthermore, value attribution biases may influence how individuals perceive other humans in AI-mediated interactions. For example, the extensive use of beauty filters in social media platforms raises important questions about whether they define new standards of attractiveness and hence impact the *attractiveness halo effect*. If AI-mediated representations alter human perceptions, this could have profound societal consequences on self-perception and interpersonal evaluations.

Beyond these considerations, value attribution biases also provide insights into the design of intelligent interactive systems. According to the *IKEA effect*, people attribute greater value to products or content that they have contributed to creating. This bias could be relevant for the design of intelligent interactive tools: designing systems that allow for greater user agency and customization could enhance user engagement and perceived value. However, the precise mechanisms that maximize this effect within digital environments require further investigation.

Similarly, the *risk aversion bias* provides valuable guidance for the design of interfaces to AI systems. Since individuals tend to avoid uncertainty, AI systems that interact with humans should incorporate gradual exposure, clear feedback mechanisms, and error-handling features to facilitate user comfort with new functionalities. Introducing complex tools incrementally can reduce cognitive overload and increase adoption rates, ultimately enhancing the user experience when interacting with intelligent systems.

Finally, the *social desirability bias* could inform the design of AI-mediated social spaces. It underscores the importance of crafting environments that encourage authenticity and spontaneous behavior, particularly in digital platforms where humans and AI interact. This consideration is especially relevant for the design of social media tools that rely on AI, virtual assistants, and video games where humans might interact with AI bots. However, the precise methods for fostering such environments is also an open area of research.

Addressing value attribution biases in AI design is critical, as they shape not only how users interact with intelligent systems but also how they perceive and evaluate other humans within AI-mediated spaces.

2.4 Recall biases

Definition Recall biases are associated with how we recall facts from our memory and influence how humans remember events leading to significant consequences for decision-making, learning, and user interaction.

Relevant examples for human-AI collaboration

- *False memory bias* [LP74; Lof75; LNP97]: The False Memory Bias is a cognitive bias where individuals develop inaccurate or entirely fabricated memories of past events, often influenced by the way questions are posed or information is framed after the event. For example, when participants are asked, “How fast were the cars going when they smashed into each other?” after being shown a video of a car crash, they are more likely to falsely remember seeing broken glass at the scene when questioned, even if none was present, compared to those asked the same questions with the verb “hit” instead of “smashed”.
- *Self-reference effect* [RKK77; GKYS07; CBQT13]: The Self-Reference Effect is a cognitive bias in which individuals demonstrate an increased likelihood of recalling information that is personally relevant or directly connected to themselves. A study with children aged between four and six years exhibited a clear memory advantage for stimuli encoded with self-referential cues [CBQT13]. Specifically, when presented with a picture of themselves taken two weeks earlier, these children showed superior recall compared to when stimuli were encoded using images of an unfamiliar child of the same age group but of the opposite gender.
- *Serial-positioning effect* [Mur62; MM78; Asc46]: The Serial Position Effect is a cognitive bias in which individuals are more likely to recall items presented at the beginning (primacy effect) and end (recency effect) of a list, while items in the middle are more easily forgotten. For example, when asked to memorize a list of words, people often remember the first few and last few items but struggle to recall those listed in the middle.
- *Peak-end rule* [KFSR93; CK96; DE13]: The Peak-End Rule is a cognitive bias in which individuals evaluate an experience largely based on how they felt at its most intense moment (the peak) and at its conclusion (the end), while largely neglecting the duration and overall average of the experience. For example, patients who underwent a painful medical procedure such as a colonoscopy tended to rate the experience more favorably if the procedure ended with a less painful phase, even if the overall procedure was longer.

Impact in human-AI collaboration Human memory is inherently limited and imperfect, a fact that we experience in everyday life. In contrast, digital storage is persistent and information is less susceptible to the distortions that characterize human recall. However, while computers offer superior storage capabilities, we should not rely entirely on them as external memory banks in what is referred to as *cognitive offloading* because retrieving relevant information from them is not always straightforward, and it is important for us humans not to lose our recall capabilities. Incorporating knowledge of human recall biases in the design of intelligent systems can help us minimize the negative impact of these biases.

Consider, for example, the *false memory bias*, which occurs when individuals confidently recall events in ways that do not align with reality, often influenced by leading questions or misleading information. Designers of AI-driven interactive systems could account for this bias by ensuring that information is presented with clarity and precision, thereby minimizing ambiguity that could lead to memory distortions. Strategies such as clear instructional design, well-structured content organization, and the strategic emphasis of key takeaways can help users retain and accurately recall information. Interactive features, such as revision aids, reminders, and detailed activity logs can reinforce accurate memory recall by enabling users to verify past actions and decisions. However, storing every event is unfeasible and learning which events need reminders is crucial to support informed decision making.

In addition to aiding with memory retention, recall biases can inform the design of interactive systems to enhance user engagement. For instance, the *serial positioning effect* suggests that individuals are more likely to remember items presented at the beginning or the end of a sequence. This principle could be used to optimize information presentation by strategically positioning the most important content at these memory-sensitive points. Similarly, the *peak-end rule* posits that people disproportionately recall the most intense (peak) and the final moments of an experience. This could be leveraged in the design of educational AI systems when students engage with challenging content: ensuring that the final stages of the interaction are not overwhelmingly difficult can help sustain motivation and reinforce positive learning experiences.

The *self-reference effect* highlights the importance of personal relevance in memory retention. Information framed in relation to an individual's own experiences is more likely to be internalized effectively. Interactive AI systems can leverage this principle by encouraging users to relate new information to their personal experiences, thus improving engagement and comprehension. This underscores the broader need for personalization in AI-driven interactions. However, a key challenge remains: how can AI systems operationalize these cognitive insights effectively? Implementing these principles in practice is non-trivial and requires additional research efforts in AI design, adaptive interfaces, and user-centered system optimization.

2.5 Decision biases

Definition Biases in this category have been documented in the context of human-decision making and reflect systematic deviations in choices made, often influenced by cognitive shortcuts and contextual factors.

Relevant examples for human-AI collaboration

- *Status quo bias* [SZ88; KKT91]: The Status Quo Bias is a cognitive bias in which individuals prefer to maintain the current state of affairs rather than pursue change, even when alternative

options might offer superior outcomes. For example, employees may resist adopting a more efficient but unfamiliar software system, favoring the continued use of an outdated program.

- *Shared information bias* [For90; PSC01; SS92]: The Shared Information Bias is a cognitive tendency in group discussions where members predominantly focus on information that is already known and shared among the group, rather than introducing or exploring unique, unshared information held by individuals. For example, a jury discussing a case might focus on information available to everyone through court documents rather than on unique perspectives brought in by individual jurors.
- *Naive allocation* [Sim90b; RL95; KvdAZ14]: Naive Allocation is a common strategy where individuals distribute resources, such as time, money, or attention, equally across available options without adequately considering their relative value, importance, or effectiveness. For example, an investor might divide their investment equally among ten funds without evaluating their performance histories, thereby potentially overfunding poor options and underfunding strong ones.
- *Take-the-best heuristic* [GG96; WLG22]: The Take-the-best Heuristic is a decision-making strategy where individuals choose between two alternatives based solely on the first cue or piece of information that differentiates them, ignoring all subsequent information. For example, when choosing between two universities, a student might select the one with the higher national ranking without considering other factors like location or available programs.

Impact in human-AI collaboration Decision biases impact how we take actions in the real world and are critical for the design of human-AI collaboration. Being able to predict how humans are likely to act in certain situations is extremely useful since it could inform where and when an intelligent interactive system could step in and assist a human decision maker.

For instance, the *take-the-best heuristic* illustrates how individuals prioritize information when making decisions. Rather than considering all available data, people rely on the first discriminating cue when deciding between two alternatives. This suggests that AI systems should prioritize presenting the most relevant information upfront, ensuring that key insights are both readily accessible and clearly highlighted. By structuring interactions in a way that aligns with human cognition, AI systems could potentially facilitate more efficient decision-making. Similarly, the prevalence of the *naive allocation* heuristic, in which individuals distribute resources evenly, underscores the need to avoid overly complex choice architectures. In this context, intelligent systems that present users with too many choices may not be beneficial. Instead, decision-making could be simplified by AI tools that guide users with the most relevant information to help reduce cognitive load.

Intelligent, interactive systems could also account for the limitations imposed by cognitive biases in certain conditions. For example, the *status quo bias* leads individuals to favor familiar options over exploring new alternatives, which can result in suboptimal decisions. From an interaction design perspective, this bias highlights the importance of crafting user experiences that strike a balance between comfort and encouragement for exploration, ultimately guiding users toward better choices without making them feel overwhelmed by change.

Intelligent, interactive systems that account for decision biases can also play a critical role in collaborative settings, particularly in AI-mediated group interactions. For instance, leveraging the *shared information bias* can inform the design of collaborative tools that reveal both dominant perspectives and overlooked insights. By actively encouraging contributions that extend beyond

the prevailing discourse, AI-driven facilitation mechanisms can foster more inclusive and balanced group discussions, ultimately improving collective decision-making.

Moreover, rather than mitigating decision biases, AI systems could also integrate human-inspired heuristics into their own decision-making processes to increase efficiency, particularly in complex real-life scenarios. Aligning AI decision strategies with human cognition may increase user trust, as heuristic-based decision processes are often more transparent and intuitively understood. However, this remains an open research question for future work.

2.6 Discussion

Having presented the framework of cognitive biases and illustrated the relevance of each category of cognitive biases on human-AI collaboration, we discuss next six key challenges that we believe are essential to advance the effective integration of cognitive biases in this context. These challenges underscore the intricate dynamics between human cognitive processes and the behavior of AI systems.

Dual role of AI in addressing cognitive biases AI systems can both exacerbate and mitigate cognitive biases, depending on their design and implementation. Poorly designed AI systems or services that have been optimized for engagement or profit can either reinforce human prejudices, mirror flawed decision-making processes or deliberately exploit our cognitive biases. Human-centric engineered models could instead enable AI systems to counteract these biases. For example, leveraging the IKEA effect to increase user trust in recommendations or using the peak-end rule to enhance user experiences demonstrate how AI systems can utilize biases for positive outcomes. However, this dual role raises ethical challenges, such as avoiding manipulation while promoting beneficial behavior. Addressing these challenges requires a balance between personalization, transparency, and fairness and highlights the importance of human-centric approaches in the design and deployment of AI systems that interact and collaborate with humans.

The complexity of personalization Many open questions highlight the potential of personalized AI systems to address cognitive biases, such as tailoring interventions for risk aversion or using self-relevance to enhance memory and engagement. However, personalization adds challenges related to privacy, fairness, and unintended consequences. For instance, overly personalized systems might reinforce existing biases instead of mitigating them or inadvertently exclude minority perspectives. Striking a balance between personalization and generalization is crucial for ensuring equitable outcomes.

Designing AI systems for collaborative decision-making A recurring theme is the role of AI systems as active participants or observers in decision-making processes. By recognizing biases like the conjunction fallacy or status quo bias, AI systems can guide users toward more rational choices. However, effective collaboration requires systems that are not only accurate but also explainable and trustworthy. For instance, enabling AI to transparently communicate trade-offs in high-stakes decisions (*e.g.*, healthcare or finance) could reduce biases while maintaining user confidence. Future research should explore how to design AI systems that adapt dynamically to user needs, contexts, and cognitive limitations.

Opportunity to educate users about biases Bias-aware AI systems present opportunities to educate users about their own cognitive biases. For example, conversational agents could highlight when users exhibit biases like the gambler’s fallacy or the false memory effect, fostering greater self-awareness and critical thinking. However, incorporating such features into everyday applications raises questions about user receptivity, the design of interventions, and the balance between education and usability. Future work should examine how to integrate educational components into AI systems without overwhelming users or diminishing their trust.

Multidisciplinary challenges Studying the role of cognitive biases in human-AI collaboration requires joint work across disciplines, including psychology, computer science, ethics, and human-computer interaction. For instance, understanding how cultural variations influence biases like the framing effect or hyperbolic discounting can help tailor AI interventions for diverse user populations. Similarly, insights from behavioral economics and cognitive neuroscience can inform algorithmic strategies for mitigating biases, such as risk aversion or the gambler’s fallacy. Interdisciplinary research must also address the scalability of these solutions to ensure that they are robust across different domains and contexts.

Ethical implications of bias-aware AI systems The ethical implications of designing AI systems that model and take into account human cognitive biases must be thoroughly examined. For example, while nudges based on cognitive biases can encourage pro-social behavior, they could also be used to manipulate user decisions in ways that prioritize commercial or institutional interests over individual autonomy. Ethical frameworks must guide the development of AI systems to ensure that interventions are transparent, equitable, and aligned with user welfare. Additionally, group settings with AI as a mediator, as in the shared information bias, highlights the need for ethical safeguards to prevent undue influence or manipulation.

Picking a representative bias The framework presented here is deliberately broad, encompassing multiple cognitive biases and opening up a wide range of potential research questions. In this thesis, however, the focus from this point forward is narrowed to a single bias, as outlined in Chapter 1. Specifically, we examine the attractiveness halo effect, selected due to the substantial body of empirical evidence supporting its existence and its growing societal relevance in light of the widespread adoption of AI-based visual manipulation tools. Studying the attractiveness halo effect in depth not only provides concrete insights into this particular bias but also serves to demonstrate and strengthen the applicability of the broader framework and research questions proposed here.

2.7 Classification of Other Biases

In addition to the representative examples above, we present next a classification of the majority of well-documented cognitive biases into the five categories outlined in the proposed framework. The four biases per category that were selected for detailed discussion in the previous section were chosen based on the extent of supporting literature, using the citation count of the original publication (recorded by Google Scholar as of August 2025) introducing each bias as a proxy for scholarly impact. However, it is important to emphasize that citation count served only as a heuristic rather than a strict selection criterion.

Biases were grouped into categories based on the number of citations and the list of biases was further refined for relevance to AI and diversity to emphasize the practical applicability of the framework in the design of intelligent, interactive AI systems. In particular, the selection process aimed to include a broad range of biases with distinct implications for human-AI interaction, while avoiding over-representation of closely related phenomena.

We also note that the classification of biases presented here is inherently approximate. Many cognitive biases exhibit features that span multiple categories, and as such, the boundaries between categories should be interpreted as conceptually useful rather than rigid. This imprecision is acceptable given that the goal of this framework is not to provide a definitive taxonomy, but to introduce a structure for understanding how different types of cognitive biases might inform the design of AI systems.

Tables 1, 2, 3, 4, and 5 serve to illustrate the breadth of the framework by demonstrating how the majority of known cognitive biases can be accommodated within it. The biases which were selected for discussion above are highlighted in blue.

Bias	Citation Count	Reference	Description
Loss aversion	87,904	[KT13]	Losses feel more painful than equivalent gains feel pleasurable.
Anchoring effect	54,605	[TK74]	Initial values heavily influence estimates or decisions, even if arbitrary.
Framing effect	28,840	[TK81]	Decisions are influenced by how information is presented, not just the content.
Pseudocertainty effect	8942	[TK89]	People prefer certainty in gains, even when it's illusory or logically inconsistent.
Backfire effect	4123	[NR10]	Strengthening one's beliefs when presented with contradictory evidence.
Decoy effect	2902	[HPP82]	Adding a less attractive option can shift preferences between two better options.
Weber-Fechner law	2794	[Fec48]	Perceived changes in stimulus intensity grow logarithmically with actual intensity.
Sexual overperception bias	1914	[HB00]	Tendency to overinterpret friendly behavior as sexual interest.
Hard-easy effect	1700	[LF77]	People overestimate performance on hard tasks and underestimate on easy ones.
Empathy gap	722	[Loe05]	Difficulty imagining how emotions affect behavior in a different emotional state.
Rhyme as reason effect	566	[MT00]	Rhyming statements are perceived as more truthful than non-rhyming ones.
Distinction bias	559	[Hse98]	Overvaluing small differences when evaluating options side-by-side.
Less is better effect	559	[Hse98]	Choosing a smaller, more appealing set over a larger, objectively better one.
Extrinsic incentives bias	403	[Hea99]	Overvaluing external rewards and undervaluing intrinsic motivation.
Experimenter effect	329	[RF63]	Researcher expectations unintentionally influence experimental outcomes.

In-group favoritism	293	[TD81]	Favoring members of one's own group over those of others.
Denomination effect	206	[RS09]	Spending is less likely when money is in larger denominations.
Ballot names bias	67	[WBDS11]	Candidates listed earlier on a ballot are more likely to receive votes.
Time-saving bias	57	[Sve70]	Misjudging how much time is saved or lost with changes in speed.
Compromise effect	44	[BBLC20]	Choosing the middle option to avoid extremes.

Table 1. Presentation Biases

Bias	Citation Count	Reference	Description
Gambler's fallacy	54,605	[TK74]	Belief that past random events affect the likelihood of future ones in independent processes.
Insensitivity to sample size	54,605	[TK74]	Underestimating how much variability decreases with larger sample sizes.
Illusion of validity	54,605	[TK74]	Overestimating the accuracy of one's judgments, especially when based on consistent but irrelevant data.
Base rate fallacy	10,275	[KT73]	Ignoring statistical base rates in favor of specific, anecdotal information.
Conjunction fallacy	6854	[TK83]	Assuming that specific conditions are more probable than a single general one.
Information bias	5414	[Bar94]	Seeking more information even when it doesn't affect decision quality.
False consensus effect	4629	[RGH77]	Overestimating how much others share your beliefs, values, or behaviors.
Hyperbolic discounting	3437	[Tha81]	Preferring smaller, sooner rewards over larger, later ones, against long-term interest.
Clustering illusion	2799	[Gil08]	Seeing patterns in random distributions of data or events.
Hot-hand fallacy	2466	[GVT85]	Belief that success in random events increases the chance of continued success.
Planning fallacy	2016	[BGR94]	Underestimating the time or resources needed to complete future tasks.
Impact bias	1697	[WG05]	Overestimating the intensity and duration of future emotional reactions.
Subadditivity effect	1646	[TK94]	Judging the sum of parts as more probable than the whole they belong to.
Illusory truth effect	1480	[HGT77]	Repeated statements are perceived as more true, regardless of accuracy.
Outcome bias	1375	[BH88]	Judging a decision based on its result rather than the quality of the decision itself.

Belief bias	1287	[EBP83]	Judging arguments based on believability rather than logical validity.
Conservatism	1056	[PE66]	People update beliefs slowly in light of new evidence.
Pessimism bias	1037	[SRRP07]	Overestimating the likelihood of negative outcomes in the future.
Optimism bias	873	[WK96]	Believing you're more likely to experience good outcomes than bad ones.
Illusory correlation	864	[Cha67]	Perceiving a relationship between variables even when none exists.
Neglect of probability	669	[Sun02]	Ignoring the actual likelihood of risks or outcomes in decision-making.
Hostile attribution bias	662	[NHD80]	Tendency to interpret ambiguous behavior of others as hostile.
Exaggerated expectation	652	[Hil12]	Overestimating the likelihood or severity of future events.
Regressive bias	336	[Att53]	Failing to recognize that extreme outcomes tend to move toward the average over time.
Zero-sum bias	264	[RBW15]	Believing one person's gain automatically means another's loss.
Restraint bias	176	[NHP09]	Overestimating one's ability to resist temptations or impulses.
Illusion of external agency	157	[GBPW00]	Believing external forces caused outcomes actually driven by one's own actions.
Pareidolia	65	[RV90]	Perceiving meaningful patterns, like faces, in random stimuli.

Table 2. Interpretation Biases

Bias	Citation Count	Reference	Description
Risk aversion	87,904	[KT79]	Preferring certain, smaller rewards over uncertain, larger ones, even when the latter is more rational.
Social desirability bias	15,571	[CM60]	Answering questions in a way that will be viewed favorably by others.
Dunning-Kruger effect	11,418	[KD99]	Incompetent people overestimate their abilities; competent people may underestimate theirs.
Fundamental attribution error	8983	[Ros77]	Overestimating the role of personal traits and underestimating situational factors in others' behavior.
Bandwagon effect	8726	[Asc51]	Adopting beliefs or behaviors because they are popular or widely held.
Endowment effect	6944	[KKT90]	Overvaluing things we own simply because we own them.

Halo effect	6735	[DBW72]	Allowing a positive impression in one area to influence judgment in unrelated areas.
Negativity bias	5922	[RR01]	Giving more weight to negative information or experiences than to positive ones.
System justification	5417	[JB94]	Tendency to defend and justify existing social, economic, or political systems.
Moral luck	4848	[Wil81]	Judging people based on outcomes beyond their control rather than their intentions.
Self-serving bias	4353	[MR75]	Attributing successes to oneself and failures to external factors.
Third-person effect	3395	[Dav83]	Believing others are more affected by media or persuasion than you are.
Ultimate attribution error	2195	[Pet79]	Attributing bad behavior of out-groups to character and excusing it in in-groups as situational.
Bias blind spot	1853	[PLR02]	Recognizing cognitive biases in others but not in yourself.
IKEA effect	1808	[NMA12]	Overvaluing things we helped create, regardless of actual quality.
Moral credential effect	1640	[MM01]	Past moral behavior gives people a license to act less ethically later.
Worse-than-average effect	1381	[Kru99]	Underestimating one's own ability compared to others, often on difficult tasks.
Omission bias	1345	[SMB91]	Viewing harmful actions as worse than harmful inactions, even with equal outcomes.
Focusing effect	1270	[SK98]	Overemphasizing one aspect of a situation when making judgments or predictions.
Women are wonderful effect	1244	[EM94]	Attributing more positive qualities to women than to men.
Out-group homogeneity bias	1177	[PR82]	Assuming members of other groups are more alike than members of your own group.
Curse of knowledge	1152	[CLW89]	Overestimating how much others know once you know something yourself.
Actor-observer asymmetry	998	[NCLM73]	Attributing others' actions to their character, while attributing our own to the situation.
Spotlight effect	836	[GMS00]	Overestimating how much others notice or care about your actions or appearance.
Barnum effect	814	[For49]	Accepting vague, general statements as uniquely applicable to oneself.
Naive realism	661	[GR91]	Believing you see the world objectively while others are biased.
Zero-risk bias	476	[VMH87]	Preferring to eliminate small risks entirely rather than reducing greater overall risks.
Illusory superiority	349	[BV91]	Believing you are better than average in various domains.
Illusion of control	263	[Tho99]	Overestimating one's ability to control events that are largely determined by chance.

Illusion of asymmetric insight	242	[PKSR01]	Believing you understand others better than they understand you.
Illusion of transparency	226	[SG03]	Believing our thoughts and emotions are more apparent to others than they actually are.
Naive cynicism	221	[KG99]	Believing others are more biased or self-serving than they actually are.
Group attribution error	214	[AM85]	Assuming group members all share the characteristics or decisions of the group.
Ben Franklin effect	181	[JL69]	Liking someone more after doing them a favor, rather than receiving one.
Cheerleader effect	171	[WV13]	People appear more attractive when seen in a group than alone.
Just-world hypothesis	101	[LM98]	Belief that people get what they deserve, and deserve what they get.
Trait ascription bias	30	[Kam82]	Believing others' behavior reflects their traits, while seeing your own as situational.

Table 3. Value attribution

Bias	Citation Count	Reference	Description
Availability bias	54,605	[TK74]	Judging likelihood based on how easily examples come to mind.
Levels-of-processing effect	18,091	[CL72]	Information processed more deeply is remembered better.
Source confusion	6624	[JHL93]	Inability to remember the origin of a memory correctly.
Egocentric bias	3800	[Gre80]	Overemphasizing one's own role in past events or shared outcomes.
False memory	3604	[LP74]	Recalling events that never actually occurred.
Misinformation effect	3604	[LP74]	Memory can be distorted by misleading post-event information.
Self-reference effect	3228	[RKK77]	Better memory for information related to oneself.
Leveling and sharpening	3193	[AP47]	Memory distortion where details are lost (leveling) or exaggerated (sharpening).
Tip of the tongue phenomenon	2722	[BM66]	Knowing you know something but being temporarily unable to recall it.
Serial-position effect	2699	[Mur62]	Better recall for items at the beginning and end of a list.
Picture superiority effect	2688	[She67]	Pictures are more likely to be remembered than words.
Cue-dependent forgetting	2605	[TP66]	Failure to retrieve memories due to missing contextual cues.
Digital amnesia	2250	[SLW11]	Tendency to forget information easily found online.

Peak-end rule	2243	[KFSR93]	Remembering past experiences based on how they felt at the peak and the end, rather than its overall impression.
Verbatim effect	1881	[Sac67]	Tendency to remember the gist of what was said rather than the exact wording.
Duration neglect	1811	[FK93]	Ignoring the duration of an experience when evaluating its overall impression.
Hindsight bias	1265	[FB75]	Believing, after an event occurs, that we knew it would happen all along.
Self-generation effect	1034	[Jac78]	Information is remembered better when generated by oneself than when read.
Continued influence effect	947	[JS94]	Misinformation continues to affect memory even after it's corrected.
Suggestibility	910	[LJ89]	Tendency for memory to be influenced by misleading information.
Von Restorff effect	896	[vRes33]	Distinctive items in a list are more likely to be remembered.
Rosy retrospection	851	[MTPC97]	Recalling past events more positively than they were experienced.
Part-list cueing effect	584	[Sla68]	Recalling some items from a list can inhibit recall of the rest.
Telescoping effect	418	[NW64]	Events are remembered as more recent than they actually are.
Choice-supportive bias	399	[MJ00]	Remembering chosen options as better than they actually were.
Testing effect	303	[Abo09]	Retrieving information improves long-term retention more than re-reading.
Cryptomnesia	283	[BM89]	Mistaking a memory for a new, original idea.
Zeigarnik effect	282	[Lew27]	Tendency to remember incomplete or interrupted tasks better than completed ones.
Reminiscence bump	277	[JP96]	Tendency to recall more memories from adolescence and early adulthood.
List-length effect	227	[Str12]	As list length increases, the proportion of items recalled decreases.
Suffix effect	226	[MCP71]	A speech sound following a list reduces recall of the last items.
Cross-race effect	213	[Fei14]	Better recognition of faces from one's own race compared to other races.
Processing difficulty effect	193	[OM85]	More difficult processing can lead to better memory under certain conditions.
Bizarreness effect	170	[MED+95]	Unusual or bizarre information is remembered better than common information.
Fading affect bias	144	[Cas32]	Emotions associated with negative memories fade faster than positive ones.

Next-in-line effect	127	[Bre73]	Poor recall for information right before or after one's own performance.
Humor effect	93	[SLH10]	Humorous information is more likely to be remembered.
Childhood amnesia	77	[HH96]	Inability to recall memories from early childhood, typically before age 3.

Table 4. Recall Biases

Bias	Citation Count	Reference	Description
Social comparison bias	37,033	[Fes54]	Resisting association with others who may outshine us.
Ambiguity effect	12,109	[Ell61]	Avoiding options with unknown probabilities over known ones.
Mere-exposure effect	11,755	[Zaj68]	Repeated exposure to something increases our liking for it.
Authority bias	9767	[Mil63]	Tendency to obey authority figures even when it conflicts with personal judgment.
Status quo bias	8670	[SZ88]	Preference for the current state of affairs, resisting change even when beneficial.
Functional fixedness	6300	[Dun45]	Inability to see alternative uses for familiar objects.
Shared information bias	6019	[For90]	Groups focus more on commonly known information than unique individual contributions.
Take the best	5297	[GG96]	Heuristic where decisions are made based on the first discriminating piece of information.
Confirmation bias	3573	[Was60]	Seeking or favoring information that confirms existing beliefs.
Congruence Bias	3573	[Was60]	The tendency to test only expected or preferred hypotheses instead of exploring alternatives.
Naive allocation	2310	[BT01]	Tendency to divide resources equally regardless of context or fairness.
Automation bias	1863	[BT54]	Overreliance on automated systems, even when they make errors.
Defensive attribution hypothesis	1582	[WAAR66]	Assigning more blame to others to protect oneself from feeling vulnerable.
Disposition effect	1463	[WC98]	Selling winning investments too early and holding onto losers too long.
Identifiable victim effect	1048	[JL97]	More empathy and aid are given to specific individuals than to anonymous groups.
Well traveled road effect	905	[All79]	Overestimating travel time on unfamiliar routes compared to familiar ones.
Money illusion	675	[Fis28]	Focusing on nominal rather than real monetary values.

Escalation of commitment	611	[Sta97]	Continuing a failing course of action due to prior investments.
Reactive devaluation	420	[RS91]	Devaluing proposals simply because they come from an opposing party.

Table 5. Decision Biases

2.8 Conclusion

Human perception, memory, and decision-making are impacted by cognitive biases and heuristics that influence our actions and decisions. Our biases, though pervasive, are frequently overlooked in the design of intelligent, interactive systems. However, with the increased prominence of human-AI collaboration, we believe that it is crucial to consider this fundamental aspect of human cognition.

This chapter introduces a framework that structures human cognitive biases into five categories from the perspective of human-AI collaboration. We highlight four representative biases from each category and describe their potential impact on human-AI collaboration. We additionally outline six key challenges and research directions to effectively consider cognitive biases in human-AI collaboration.

By examining cognitive biases through this lens, we aim to inspire research that not only addresses the challenges of bias-aware AI systems, but also harnesses these biases for ethical and beneficial purposes, fostering more meaningful, effective, and trust-driven human-AI interactions. Future directions of research include exploring the long-term impact of bias mitigation strategies in user trust, satisfaction, and decision quality; investigating the potential to generalize interventions across domains, demographic groups, and cultures; advancing methods of co-adaptive learning, where both users and AI algorithms iteratively improve through the interaction; developing metrics and benchmarks for evaluating the effectiveness of bias-aware AI systems; and developing ethical guidelines and transparency requirements for leveraging cognitive biases in intelligent, interactive systems.

Chapter 3

The Attractiveness Halo Effect in Human Decisions

Chapter summary

While the Attractiveness Halo Effect is well-documented in physical contexts, its impact on decision-making in the digital world remain under-explored. Given the large scale movement into digital spaces, we investigate the impact of the attractiveness halo effect on people using AI-based beauty filters. This chapter details the methodology and findings of a large-scale online user study we conducted involving 2,748 participants who rated facial images from a diverse set of 462 distinct individuals in two conditions: original and attractive after applying a beauty filter. Our study revealed that the *same* individuals receive statistically significantly higher ratings of attractiveness and other traits, such as intelligence and trustworthiness, in the attractive condition. We also study the impact of age, gender, and ethnicity and identify a weakening of the halo effect in the beautified condition, resolving conflicting findings from the literature and suggesting that filters could mitigate this cognitive bias. We additionally discuss the ethical concerns regarding the use of beauty filters raised by our study.

This chapter is based on the paper:

[GMF+24a] Aditya Gulati, Marina Martínez-Garcia, Daniel Fernández, Miguel Angel Lozano, Bruno Lepri, and Nuria Oliver. “What is beautiful is still good: the attractiveness halo effect in the era of beauty filters.” *Royal Society open science* 11, no. 11 (2024): 240882

3.1 Introduction

The remainder of this thesis focuses on a well documented cognitive bias known as the *attractiveness halo effect*. Multiple studies have shown that beauty matters to people, even when we know that physical attractiveness is not correlated with other measurable traits, such as intelligence [MZW+15; JHH95; KCF14]. In fact, decades of research in several disciplines—including sociology, psychology, behavioral economics and organizational science—has found that perceptions of attractiveness profoundly impact the social judgments that we make: human beings are positively biased towards individuals who are perceived as physically attractive.

The Attractiveness Halo Effect, henceforth referred to as the AHE, leads to physically attractive people being considered to be more intelligent [DBW72; Kan11; Tal16], happier [MK75; GML13],

more trustworthy [Tod08], more sociable and sexually warmer [Mil70], better adjusted [EAML91] and generally more successful in life [DBW72], when compared to less physically attractive individuals. This halo effect has an impact on consequential aspects of our lives, as attractive individuals are thought to be better students [RPT92] or politicians [BKTR08], more qualified for jobs [CK85; HSC03], and are more likely to receive promotions, higher salaries [FOR91; HB94] or more lenient judicial sentences [WR15; Wil95] than less attractive people.

However, these findings have been generally obtained by means of small user studies where study participants provided judgments of a typically small sample of face images with limited diversity. Hence, questions arise regarding the generalization of the attractiveness halo effect from different perspectives.

First, concerning the ethnicity of the stimuli and the human evaluators, Albright et al. [AMD+97] found cross-cultural agreement in the judgments provided to western and non-western faces. However, more recent research reported a cross-cultural variation [MKC+14] and hence did not corroborate previous results. To shed light on this issue, Batres and Shiramizu [BS22] carried out a large-scale study that examined the attractiveness halo effect across 45 countries in 11 world regions and on a diverse set of faces from four ethnicities. Their results showed that attractiveness correlated positively with most of the socially desirable personality traits—such as being more confident, emotionally stable, intelligent, responsible, sociable and trustworthy. Hence, according to this study, the attractiveness halo effect would generalize to diverse stimuli and human evaluators. Related work by Gabrieli et al. [GLSE21] found that the attractiveness halo effect regarding trustworthiness is only influenced by the age of the presented faces, but not by their gender or ethnicity. Similarly, [KKM23] reported mixed results regarding the impact of ethnicity on the attractiveness halo effect in the context of hireability. Therefore, the evidence in this regard is inconsistent and additional research would be needed to shed light on this matter.

The second perspective relates to the interaction between the gender of the stimuli and the gender of the human evaluators. Early work by Dion et al. [DBW72] did not report any significant interactions between the gender of the human evaluators and the gender of the stimulus regarding the existence of the attractiveness halo effect. However, later research reported a stronger attractiveness halo effect towards opposite-gender individuals [ASS+16]. In fact, several studies only included male raters of female faces (e.g. [BK72; GPT82]) or female raters of male faces [AHPS23]. In a study with both male and female raters and stimuli, Kunst et al. [KKM23] reported a significant interaction of gender, attractiveness and competence *only* when male participants rated the competence of female applicants in a hiring scenario. Again, there is mixed evidence in this regard.

The third perspective concerns the existence of this cognitive bias on the *same individual* in two conditions: original and attractive. Would the same person be perceived as having higher levels of socially desirable attributes—such as intelligence, trustworthiness or sociability— simply by improving their physical appearance?

Cosmetics are a popular tool to alter appearance and their use has been shown to increase perceptions of attractiveness [Rus09; BR22; BPCR21; BPL+19; MFH+03; TON16; BRS+18; GJ81]. Makeup has been reported to increase skin evenness [BPL+19] and facial contrast, which in turn leads to a perception of increased femininity and attractiveness [Rus09]. Further literature has studied how varying levels of makeup impact perceived attractiveness [MFH+03]. While some research found that light makeup is preferred to heavy makeup [TON16] and others reported the opposite effect [BPCR21], faces with makeup applied to them were consistently rated as more attractive than those without makeup. Thus, the application of makeup has been used in the literature to study the attractiveness halo effect in two conditions: original and attractive. By means of user

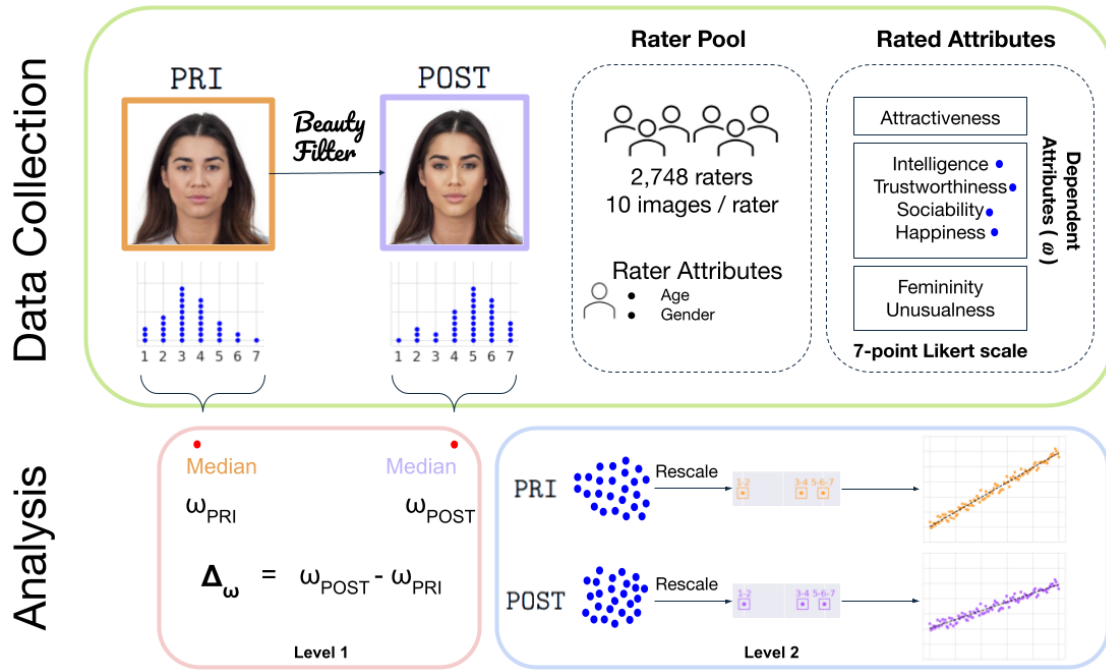


Figure 3. Overview of the study and the analysis of the collected data. The stimuli consist of two sets of facial images: the PRI dataset, extracted from existing datasets for research on faces [MCW15; ERL10] and the POST set, created by applying a state-of-the-art beauty filter to each image in the PRI dataset. Each participant ($N = 2,748$) rated 10 different images on 7 attributes indicated on the top right part of the Figure. Each image received ratings from at least 25 different participants. To shed light on the attractiveness halo effect, two levels of analysis were performed: (1) An *aggregate* level—depicted inside the pink box in the Figure—using the medians of all the ratings received by each image, which are referred to as *centralized* ratings (•); and (2) an *individual* level (•)—depicted inside the blue box in the Figure—consisting of each rating and considering the participants’ characteristics.

studies with psychology students and a very small set of stimuli in two conditions (original and attractive), several authors reported that the attractive condition evoked more social reinforcement and enhanced popularity ratings [BK72; GPT82], and higher levels of competence, professionalism, assertiveness and ability to provide support [LW78]. However, others reported no statistically significant differences in the attribution of socially desirable characteristics among subjects in the original and attractive conditions [KR75; TH80]. Furthermore, these studies involved opposite-gender pairs where male participants—mainly recruited from universities—evaluated images or videos of female confederates without (original) or with (attractive) makeup applied. While insightful, these studies are difficult to scale up since it is costly to physically apply makeup to a large number of stimuli. Despite recent work showing that makeup increases perceived attractiveness in male faces as well [BR22], the application of makeup could create a gender asymmetry as makeup is socially more acceptable when applied to female than to male faces in many cultures [PSE+02] and the improvements in attractiveness that can be achieved as a result of applying makeup are limited.

In sum, the literature suggests that the *what is beautiful is good* notion [DBW72] may be oversimplified, supporting the need for further research to better understand this phenomenon. Moreover, there is evidence that increased perceptions of physical attractiveness also lead to increased perceptions of socially undesirable traits, such as vanity [DT75; HL23; RD23], materialism and

sexual permissiveness [Bas81]. Nonetheless, the research presented in this paper focuses on the attractiveness halo effect related to perceptions of socially desirable attributes, namely intelligence, trustworthiness, sociability and happiness. A detailed discussion for this choice and the associated limitations can be found in Section 3.3.

In addition to shedding light on these open questions, we expand the scope of the study of this cognitive bias from the physical to the digital world. The attractiveness halo effect acquires a new relevance in the digital space, particularly as human-to-human communication is frequently mediated by technology and Artificial Intelligence (AI) tools are increasingly used to make assessments about humans, to interact with us via *e.g.* chatbots and to create enhanced digital versions of ourselves. Beauty filters are an example of such a tool, which aim to *beautify* the face of the person by applying complex transformations to the face that go beyond what makeup can achieve, including morphological changes to the eyes and eye lashes, the nose, the chin, the cheekbones, and the lips, in addition to smoothing the skin, removing wrinkles and imperfections [RPG+22; RCOO24]. These filters offer a unique opportunity to study the attractiveness halo effect at scale, with diversity in the age, gender and ethnicity of the stimuli, and in a controlled scenario, because they allow the creation of *beautified* versions of the *same* individuals. While makeup has been shown to reliably increase perceptions of attractiveness [BPL+19; TON16; MFH+03], applying makeup requires technical skill and perceptions of attractiveness differ depending upon the skill of the person applying the makeup [BPCR21]. The manipulation of beauty by means of beauty filters enables a controlled and consistent adjustment of attractiveness at a large scale while keeping the identity of the face constant [RPG+22], which is crucial for isolating the effects of perceived attractiveness from other confounding variables, such as facial identity or expression.

Furthermore, beauty filters are widely used in the digital world, and they play a significant role in shaping contemporary beauty standards and perceptions [LC20]. While they have been shown to profoundly impact user self-presentation —raising questions about authenticity [LC20], self-esteem [Per20; Bar20; Isa23], mental health [CDAR20; Esh20], diversity [RPG+22] and racism [RCOO24]— there is a lack of research on how they impact perceptions of attractiveness and associated cognitive biases on the same individuals. There is also a need to study the effect of these augmented appearances on how users are perceived and judged within digital environments, both by humans and by AI algorithms. In fact, recent studies have investigated the role that beauty filters play in perceptions of trustworthiness of male stimuli [AHPS23] and of male and female stimuli in a hiring scenario [KKM23]. Our research complements these studies by means of a large scale user study of the attractiveness halo effect regarding four socially desirable attributes, namely intelligence, trustworthiness, sociability and happiness, with a diverse set of stimuli.

We leverage a state-of-the-art popular beauty filter applied to a diverse set of face images (N=462) to create an *attractive* condition for the same individuals. Using this dataset, we perform a large-scale user study (N=2,748) to shed light on the conflicting findings reported in the literature regarding the attractiveness halo effect in the context of socially desirable attributes, and how different rater and stimuli characteristics, such as gender or age, impact the perception of these attributes. The ratings we collected, together with the images of the same individuals presented both with and without a beauty filter applied, are collectively referred to as **AHEAD** (the Attractiveness **H**alo **E**ffect **A**tribution **D**ataset). This dataset provides a robust foundation for advancing research on the attractiveness halo effect and enables deeper insights into how this bias shapes decision-making processes in AI systems when humans are involved.

Our research thus contributes to the understanding of this cognitive bias from four different perspectives: first, we study the impact of the beauty filters on the attractiveness halo effect for the

same individuals; second, we investigate the existence of this cognitive bias on a *diverse* set of stimuli (faces); third, we analyze the role that the gender, age and ethnicity of the stimuli and the raters play regarding the attractiveness halo effect; fourth, we explore the potential of beauty filters to mitigate the existence of the attractiveness halo effect in the digital world.

3.2 Results

We report the results of analyzing the responses of 2,748 study participants (raters) who provided ratings on a 7-point Likert scale for 7 different attributes —namely, attractiveness, intelligence, trustworthiness, sociability, happiness, femininity and unusualness— in addition to their estimation of the gender, age and ethnicity of 10 different face images (stimuli) from a pool of 924 images. A detailed description of the study procedure and design can be found in Section 3.4.3, while a summary can be seen in Figure 3.

The images consisted of the original faces ($N=462$, labeled as PRI for **P**icked **R**epresentative **I**mages) and their corresponding beautified versions ($N=462$, labeled as POST for **P**Ost **S**ocial media **T**ransform) by means of applying a state-of-the-art, popular beauty filter. No participant provided ratings on the same set of images to ensure that each participant was exposed to a diverse set of stimuli while maximizing the number of ratings provided for each face image. Furthermore, no participant rated an image corresponding to the *same* individual in both conditions (with and without the filter applied) and participants were not told that half of the images that they evaluated corresponded to the *beautified* versions of the original face images.

The reported results are structured according to two levels of analysis. Following past studies [BJP17; RLK23], we first compute the median value —due to the non-normality in the distribution of the values ($D = 0.93$, $p < 0.001$, Kolmogorov–Smirnov)— of the ratings provided by the participants for each image and each attribute, which is henceforth referred to as the *centralized* score. While this level of analysis enables making pairwise comparisons between the ratings provided to the same individuals in the PRI and POST sets, it does not allow to study the variance in the ratings due to the participants. Thus, we also analyse each rating individually to include the effects of the participants' gender and age. To perform such an analysis, Ordered Stereotype Models (OSMs) [And84; FAP16; FLC19] are first applied to the ordinal responses on the 7-point Likert scales to estimate “a new spacing among the ordinal categories dictated by the data” [FLC19]. The raw data is then transformed according to the new scales obtained with the OSMs and we build linear mixed models to study the impact of the raters' gender and age on their responses, considering the raters as random effects. A detailed discussion of the methodology used to analyze the ratings can be found in Section 3.4.

3.2.1 Beauty Filters and Attractiveness

Manipulation test: Do beauty filters increase attractiveness? The same individuals were rated as significantly more attractive after applying the beauty filter than before its application ($p < 0.001$, one-sided Wilcoxon paired-rank), as reflected in Figure 4a which depicts the distribution of centralized attractiveness ratings for each image before and after the filter was applied.

The median increase in perceived attractiveness after beautification was 1 point on the 7-point Likert scale. There were no images where the centralized perceived attractiveness score decreased after beautification and it remained the same before/after beautification only in 3.9% (18 out of 462

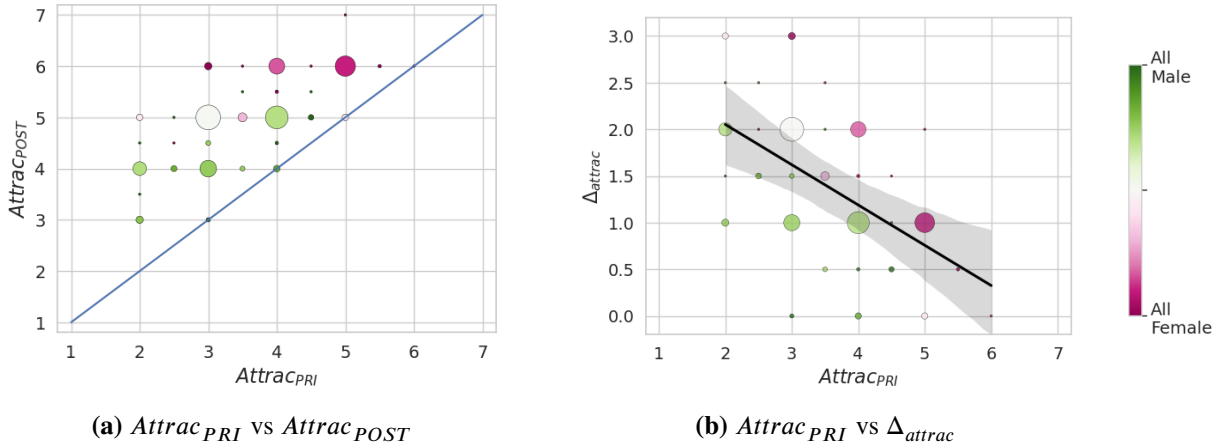


Figure 4. Impact of the beauty filters on perceived attractiveness. The size of the circles is proportional to the number of ratings provided for each value on the 7-point Likert scale and the color indicates the proportion of males and females for each rating. (a) Pairwise comparison of perceived attractiveness before and after beautification. Observe how no image decreased its perceived attractiveness ratings after beautification and how the highest perceived attractiveness ratings tend to correspond to females. (b) Increase in perceived attractiveness (Δ_{attrac}) after the application of the beauty filter versus the initial levels of attractiveness. Shading corresponds to the 95% confidence interval. The higher the original perceived attractiveness, the lower the increase in attractiveness after applying the filter.

images) of the cases. We conclude, thus, that the manipulation was successful as the beauty filters significantly increased the perceived attractiveness of the same individuals after beautification.

The increase in perceived attractiveness ($\Delta_{attrac} = Attrac_{POST} - Attrac_{PRI}$) due to the application of the beauty filter is negatively correlated with the initial attractiveness score of the face images (Kendall's $\tau = -0.49$, $z = -12.395$, $p < 0.001$), as reflected in Figure 4b: the lower the initial attractiveness, the larger the benefit of applying the beauty filter.

Impact of the filters regarding the age, gender and ethnicity of the stimuli Figure 5 depicts the centralized attractiveness scores in the original (PRI) and beautified (POST) datasets according to the age, gender and ethnicity of the stimuli. Note that we adopt the same nomenclature as the labels provided in the face datasets analyzed in our study: gender is a binary variable with two values (male/female) and ethnicity can have 6 values (Asian/Black/Latino/White/Indian/Mixed). As explained in Section 3.4, the analyses of age and ethnicity are carried out on the images from the FACES dataset and the Chicago Faces Database (CFD) respectively, whereas the analysis of gender is performed on all the images from both datasets. The age groups are given by the FACES dataset and correspond to: Young [$19 \leq age \leq 31$]; Middle [$39 \leq age \leq 55$]; and Old [$age > 69$]. As seen in Figure 5, while the age and gender of the stimuli have a clear impact on their perceived attractiveness levels, ethnicity does not seem to play a role.

More precisely, a statistically significant difference in the centralized perceived attractiveness scores depending on the age and gender of the individual was found, both in the original (PRI) and beautified (POST) versions ($p < 0.001$, Kruskal-Wallis). No statistically significant effect of ethnicity was found in neither of the conditions. Images corresponding to young individuals received significantly higher ($p < 0.001$, pairwise Wilcoxon) centralized perceived attractiveness scores than those depicting middle-aged or older individuals in both the PRI and POST sets. Images depicting middle-aged individuals were considered significantly ($p < 0.001$, pairwise Wilcoxon)

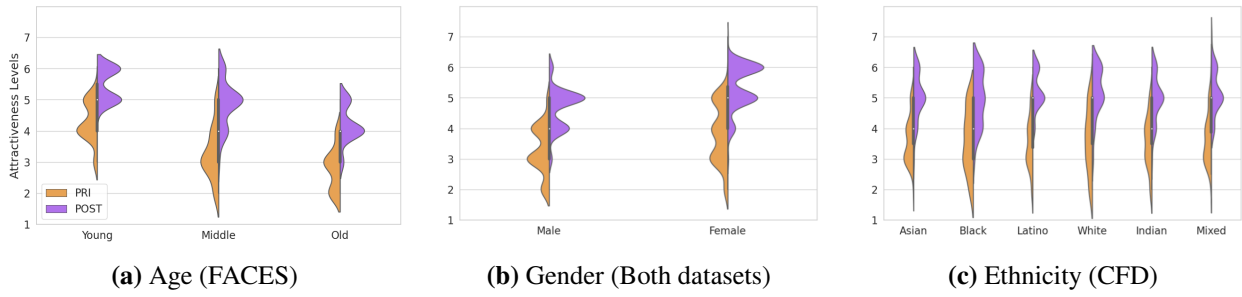


Figure 5. Distribution of the median ratings of perceived attractiveness of the original (PRI, in orange) and beautified (POST, in purple) face images when varying the age (a), gender (b) and ethnicity (c) of the stimuli. Note that the age and ethnicity results are computed on the FACES and CFD datasets, respectively, whereas the gender results are based on the analysis of both datasets. Regarding age, the younger the individual, the higher their perceived attractiveness ratings ($p < 0.001$, pairwise Wilcoxon). With respect to gender, female faces receive higher attractiveness ratings than male faces ($p < 0.001$, Kruskal-Wallis). No statistically significant difference was found in the attractiveness levels depending on the ethnicity of the stimuli both before and after beautification.

more attractive than those depicting older individuals only after applying the beauty filter.

The increase in perceived attractiveness (Δ_{attrac}) due to the filters was also significantly different across age groups. Images of middle-aged individuals had a mean Δ_{attrac} of 1.57 points, which was significantly ($p < 0.001$, pairwise Wilcoxon) higher than the Δ_{attrac} of images corresponding to younger individuals, who had a mean increase of 1.18 points in their centralized attractiveness scores due to the application of the filters. Images depicting older individuals had a mean Δ_{attrac} of 1.38 points, which did not differ significantly from images of neither younger nor middle-aged individuals.

Images of females received significantly higher ($p < 0.001$, Kruskal-Wallis) perceived attractiveness ratings than images of males both before and after beautification. The mean increase in centralized attractiveness for female images ($\Delta_{attrac} = 1.53$) was higher ($p < 0.01$, Kruskal-Wallis) than that for male images ($\Delta_{attrac} = 1.34$). A similar analysis on the impact of the filters on the dependent attributes can be found in Appendix F.1.1.

In addition, we study how the filters impact the perception of physical characteristics such as age, gender and ethnicity along with attributes related to physical appearance, such as perceived femininity and unusualness. These findings are reported in Appendix C.

In the following sections, we focus on the attractiveness halo effect regarding 4 attributes that have been extensively studied in the literature: intelligence [BS22; OT08; ZBL07; Tal16], trustworthiness [BS22; MCW15; DBW72; PUG+22; OT08; TDP+13; Tod08; ZBL07], sociability [BS22; DBW72; OT08] and happiness [BS22; PUG+22; OT08; GML13; MCW15].

3.2.2 Beauty Filters and the Attractiveness Halo Effect

Statistically significant differences were found in the centralized scores of the 4 dependent variables of interest (intelligence, trustworthiness, sociability and happiness) between the original (PRI) images and their beautified (POST) versions ($p < 0.001$, one-sided Wilcoxon paired-rank). Images of the *same individuals* received higher scores on all attributes after beautification, as depicted in Table 25 (Appendix F.1). Thus, the *same individuals* were perceived not only as more attractive, but also as more intelligent, trustworthy, sociable and happy after applying a beauty filter, providing

evidence that supports the existence of the attractiveness halo effect.

Dependent Attribute (ω)	PRI			POST		
	β_0	β_1	R^2	β_0	β_1	R^2
Intelligence	3.18***	0.30***	0.327	4.11***	0.12***	0.036
Trustworthiness	3.34***	0.20***	0.181	3.50***	0.17***	0.069
Sociability	2.56***	0.39***	0.363	2.78***	0.38***	0.321
Happiness	2.08***	0.39***	0.261	2.47***	0.35***	0.186

Table 6. Parameters of the linear model $\omega = \beta_0 + \beta_1 \text{Attrac} + \epsilon$ for each dependent variable ω on the PRI and POST sets independently. A larger absolute value of the intercept β_0 in the POST set indicates that the value of the perceived attribute increases after applying a beauty filter. A smaller absolute value of β_1 in the POST set reflects a weaker halo effect after beautification.

Linear models, depicted in Table 6, of the centralized score for each dependent variable (ω) as a function of the centralized score of perceived attractiveness for each image ($\omega = \beta_0 + \beta_1 \text{Attrac} + \epsilon$) reveal a significant effect ($p < 0.001$)¹ of perceived attractiveness on all dependent variables both before (PRI) and after (POST) beautification. The positive and significant β_1 for all attributes on the PRI and POST sets supports the existence of the halo effect and is in line with past work that studied this effect using different subjects in two conditions: original and attractive [BK72; LW78; GPT82]. *Intelligence* exhibits the largest decrease in β_1 after beautification, reflecting a weaker halo effect. There is a significant decrease in the goodness-of-fit of the model (R^2) for intelligence ($\approx 90\%$) and trustworthiness ($\approx 60\%$). We discuss the implications of these findings in Section 3.3.

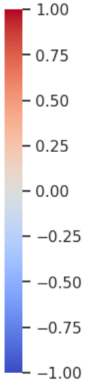
3.2.3 Impact of the Raters on the Attractiveness Halo Effect

The centralized ratings allowed performing pairwise comparative analyses between the images in the PRI and POST datasets. However, aggregating the scores by their medians masks the impact of the raters' attributes, such as their age and gender, on the perceptions of attractiveness and the attractiveness halo effect. In this section, we report the results when analyzing each rating individually to consider the role of different rater characteristics in the perception of the dependent attributes, and the halo effect.

To leverage the individual ratings, the collected ordinal ratings were first transformed into a continuous variable using the Ordered Stereotype Model (OSM) [And84; FAP16; FLC19]. For the data in the PRI and POST datasets independently, we then built linear mixed models of perceived attractiveness (Eq. 5) and of each of the dependent variables using attractiveness and the rater's characteristics (age and gender) and their interactions as independent variables (Eq. 6). A detailed discussion motivating this modeling choice can be found in Section 3.4.5. The new scales for attractiveness and the dependent attributes (ω) computed by the OSM's can be found in Appendix B.1.

Note that pairwise comparisons between images in the PRI and the POST datasets—as was done with the centralized scores—are not appropriate for two reasons. First, since no participant rated

¹In the figures and tables, we use the standard star notation to represent the p-values i.e., *** : < 0.001 , ** : < 0.01 , and * : < 0.05



ω	β_0		β_1		β_2		β_3		β_4		β_5		β_6		β_7	
	PRI	POST	PRI	POST	PRI	POST	PRI	POST	PRI	POST	PRI	POST	PRI	POST	PRI	POST
Attractiveness	***	***	X	X	***	***	***	***	***	***			***		***	***
Intelligence	***	***	***	***	***	***	***	***	***	***					***	***
Trustworthiness	***	***	***	***	***	***	***	***	***	***						
Sociability	***	***	***	***		***		***								
Happiness	***	***	***	***		***						***				

Table 7. Significance levels (** $p < 0.001$; * $p < 0.01$) and magnitudes of the β 's in the linear mixed models built to measure the impact of the rater's and stimulus' age and gender on the attractiveness halo effect. The shading in each cell corresponds to the absolute value and sign of the corresponding β on normalized data in order to compare their effect across different variables. β_0 : Intercept, β_1 : $Attrac_I$, β_2 : $Gender_I$, β_3 : Age_I , β_4 : $Gender_R$, β_5 : Age_R , β_6 : $Gen_I \cdot Gen_R$, β_7 : $Age_I \cdot Age_R$. Note how perceived attractiveness is the strongest predictor both before and after beautification. After beautification, other variables play a role given the decreased predictive power of attractiveness, details of which can be found in Appendix B.6.

the same image both before and after beautification, it is not possible to generate any logical pairs. Second, the OSM is computed independently on the PRI and POST sets as the goal of this part of the analysis is understanding the impact of different rater attributes on perceptions with and without the filters. This leads to different scales for the attributes between the PRI and POST sets due to which pairwise comparisons are not appropriate.

Table 7 summarizes the β_i 's and associated p-values for each of the linear mixed models. Note how all β_0 and β_1 are significant ($p < 0.001$) for perceived attractiveness and the dependent variables both before and after beautification. Perceived attractiveness (β_1) is the strongest predictor of the dependent variables, yet its predictive power decreases in the models built with data after beautification (see Appendix B.6 for a detailed analysis). As a consequence, there are other factors that play a more significant role after beautification. The colors in Table 7 represent the values of the β_i 's on a normalized scale to allow for an easier comparison across models. Using these models, we analyze next the impact on the attractiveness halo effect of the rater's age and gender, and their interactions with the age and gender of the stimuli.

Impact of the Rater's Age Regarding attractiveness, the perceived age of the stimulus (β_3) is negatively correlated ($p < 0.001$) both before and after beautification, as has been extensively reported in the literature [KB00; WM84; FCT06]. Conversely, the rater's age (β_5) does not exhibit a significant correlation with perceived attractiveness, which is aligned with previous research [CCD71; PMHP01] and in contradiction with what other authors have reported [FC11]. Furthermore, we observe a significant ($p < 0.001$) positive correlation in the interaction between the perceived age of the stimulus and the rater (β_7), in concordance with the literature [FCT06].

With respect to the dependent variables, the rater's age has a statistically significant positive correlation ($p < 0.001$) with perceived intelligence, trustworthiness and happiness only after beautification. There is no statistically significant impact of the rater's age on any of the dependent variables before beautification. Interestingly, the interaction between the rater's and stimulus' age is only significant for perceptions of intelligence after beautification.

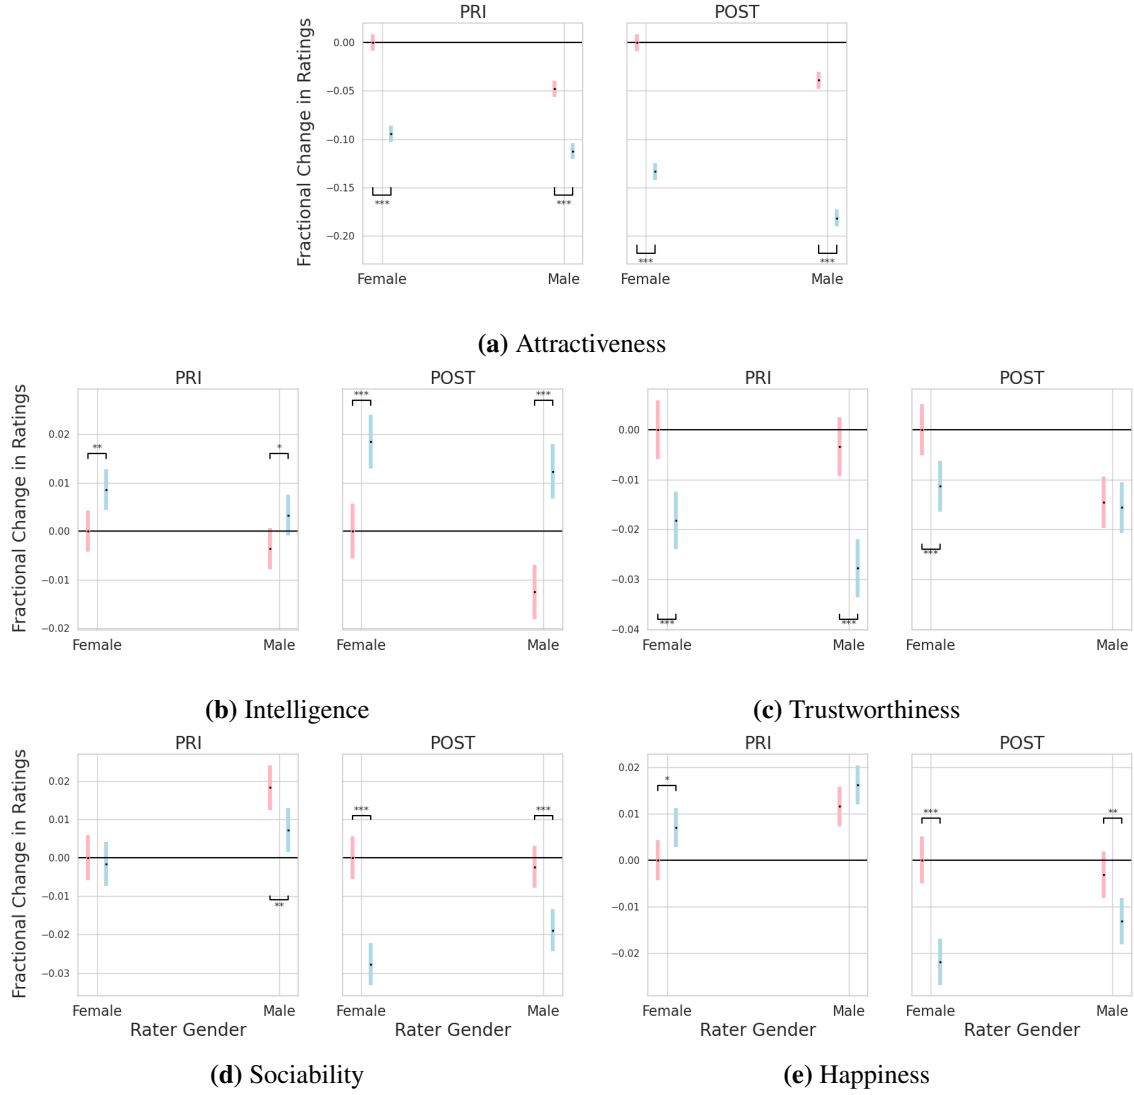


Figure 6. Impact of rater's and stimulus' gender on attractiveness and the dependent variables in the PRI and POST datasets. The x-axis represents the gender of the rater and the colors represent the gender of the stimulus (pink [●] for images of females and blue [●] for images of males). The width of the bars corresponds to the 95% confidence interval of the Estimated Marginal Mean (EMM) [SSM80; Len23]. The y-axis depicts the relative change in the EMM from the EMM of female stimuli rated by female participants. Details on how these values were computed can be found in Appendix B.7.

Impact of Rater's Gender As with age, the models represented by Equations (5) and (6) consider the impact of the rater's gender (β_4), the stimulus' gender (β_2) and their interaction (β_6). Note that the significance levels reported for β_6 in Table 7 correspond only to the interaction term of male raters rating male images since females were encoded as 0. Thus, we report the estimated marginal means [SSM80; Len23] for each (image gender, rater gender) pair. Figure 6 depicts the estimated marginal means for attractiveness and the 4 dependent attributes for all (image gender, rater gender) pairs in the PRI and POST datasets.

Before beautification, both male and female raters provide significantly different scores of attractiveness ($p < 0.001$), trustworthiness ($p < 0.001$) and intelligence ($p < 0.01$ for female raters, $p < 0.05$ for male raters) to images of males and females. Images of females receive significantly higher: (1) sociability ($p < 0.001$) scores from male raters than from female raters; and (2) attrac-

tiveness ($p < 0.001$) scores from female raters than from male raters before beautification. There are no statistically significant differences ($p < 0.001$) in the scores assigned to images of males regarding all attributes when rated by both male and female raters.

After beautification, female raters provide significantly different ($p < 0.001$) ratings to images of males and females on all attributes, whereas male raters provide significantly different ($p < 0.001$) scores to images of males and females only on perceived attractiveness, intelligence and sociability but not on trustworthiness and happiness. While images of males received comparable scores on all attributes in the PRI dataset, there are statistically significant differences ($p < 0.001$) in the perceived attractiveness of images of males by male and female raters after beautification with male raters providing lower scores to images of males than female raters. Images of females are also given significantly lower attractiveness ($p < 0.001$) scores by male raters than by female raters, as observed in the PRI dataset. Additionally, images of females are given significantly lower trustworthiness scores by male raters, even though male and female raters provided similar trustworthiness scores to females in the PRI dataset. The opposite impact of the filters is seen regarding sociability with no significant difference observed in the ratings received by images of females despite there being a significant difference before beautification.

Even though images of females were given higher scores of perceived attractiveness ($p < 0.001$) than images of males by both male and female raters, they were given lower scores of intelligence than images of males, particularly after beautification ($p < 0.001$). This finding suggests the existence of a gender bias in perceptions of intelligence [KCF14; Tal16]. Gender has also been found to play a significant role in the perception of related attributes such as competence and hireability [HS85; GBL86; Kap78; KKM23; OBT18]. The implications of this finding are discussed in Section 3.3.

Figure 6 also provides insights into the impact of the filters on male and female raters. The gap between the ratings given to images depicting males vs females by male and female raters before and after beautification notably increases when judging attractiveness, intelligence, sociability, and happiness and decreases when judging trustworthiness. Moreover, the gender differences in attractiveness, intelligence and trustworthiness ratings change significantly more after beautification for male raters, whereas a similar effect is observed for sociability and happiness for female raters. Finally, trustworthiness is the only dependent variable where the gender differences in the scores provided to images of males and females by male and female raters decrease after beautification. Table 17 in Appendix A.1 quantifies the percentage change in ratings for different dependent attributes depending on the gender of the rater.

These findings suggest that judgments made by male raters on attractiveness, intelligence and trustworthiness are more sensitive to the filters when compared to the judgments by female raters. Conversely, female raters tend to be more sensitive to the beauty filters than male raters when providing judgments of sociability and happiness. Implications of these findings are discussed in Section 3.3.

3.2.4 Do Beauty Filters Mitigate the Attractiveness Halo Effect?

Beauty filters increase the perceived attractiveness scores for almost all individuals indicating that they shift the distribution of perceived attractiveness to the right on the 7-point Likert scale. Additionally, they have a greater impact on individuals who received low scores of perceived attractiveness before beautification (Figure 4b). This leads to beauty filters narrowing the spread of perceived attractiveness ratings ($p < 0.001$, Levene's [Lev60]), thereby reducing their influence as a factor to




Figure 7. A visual representation of the relationship between perceived attractiveness and the dependent attributes after rescaling with the Ordered Stereotype Model. The scales here have been normalized for ease of representation. Note how intelligence shows a much stronger saturation effect than the other dependent attributes in the PRI dataset. In the POST dataset, both intelligence and trustworthiness exhibit a saturation effect.

impact the perception of other attributes, such as intelligence. Thus, beauty filters could potentially mitigate the halo effect.

The linear models in Table 6 reflect a decrease in the value of β_1 and R^2 of the linear models after beautification, particularly for intelligence and trustworthiness, supporting the hypothesis of a mitigation of the halo effect for these attributes. We postulate the existence of a *saturation effect*, i.e., beyond a certain level of perceived attractiveness, there is a significant reduction in the impact that attractiveness has on the dependent variables.

Figure 7 depicts the relationship between perceived attractiveness and the dependent variables after rescaling the data according to the OSMs both before (a) and after (b) beautification. In the case of intelligence we observe a clear saturation effect in the PRI dataset and a similar effect is observed for trustworthiness in the POST dataset, where the slope of the linear mixed model of trustworthiness as a function of perceived attractiveness decreases as attractiveness increases, especially when compared to sociability and happiness. Detailed statistical analyses supporting this saturation effect can be found in Appendix B.4.

These findings suggest that filters’ capacity to enhance attractiveness, rather than their ability to reduce attractiveness variation, is the main factor in reducing the halo effect observed in certain attributes. Implications of this saturation effect are discussed next.

3.3 Discussion

In this study we have collected human feedback of a large-scale, diverse dataset of face images of the same individuals in unattractive (original) and attractive (beautified) conditions by means of applying a digital beauty filter. While personal preferences arguably play a role in perceptions of attractiveness [YS98; KW04; Sar22], we obtain irrefutable evidence that AI-based beauty filters increase the perceptions of attractiveness for almost all individuals, regardless of their gender, age and race. The centralized perceived attractiveness score increased for 96.1% of the individuals after beautification and remained unchanged for the rest.

Few studies [BK72; LW78; GPT82; TH80; KKM23] have investigated the presence of the attrac-

tiveness halo effect on the *same individual* by creating two conditions: an attractive and unattractive setting for the same person. The attractive condition was typically achieved by enhancing or beautifying the appearance of the individual to be rated by means of professional lighting, fashionable clothing and hair style, the application of makeup and/or, more recently, digital beauty filters. The results of previous studies have been mixed. The diversity of our stimuli, our large sample and the ability to apply a consistent transformation to increase attractiveness by means of beauty filters provide robust data on this matter. Contrary to previous works [KR75; TH80] and supportive of others [BK72; LW78; GPT82], we find strong evidence of the existence of the halo effect both before and after beautification for the 4 dependent variables of interest (see Table 6). Furthermore, beauty filters impact the attractiveness halo effect differently, depending on the attribute: while still significant for all dependent variables, the effect weakens after beautification for intelligence and trustworthiness (Table 6), suggesting that beauty filters could be used to mitigate the attractiveness halo effect regarding these two attributes due to the increase in the attractiveness levels after applying the filter. In fact, the mean value for perceived attractiveness increased from 3.57 in the PRI dataset to 5.01 in the POST dataset. As a result, while only approximately 17% of the faces in the PRI dataset were rated as having an attractiveness level greater than or equal to 5 (with 4 being the neutral point on the scale), this percentage increased to approximately 75% after beautification. Furthermore, the distribution of attractiveness values decreased its variance after beautification, dropping from 0.83 in the PRI dataset to 0.60 in the POST dataset. Additionally, we identify a negative correlation between the original levels of attractiveness and the increase in attractiveness, such that the larger the perceived attractiveness of the original image, the smaller its increase in attractiveness due to the application of the filter (Figure 4b).

Additional analyses revealed that the relationship between attractiveness and the dependent variables is non-linear such that it saturates after a certain level of perceived attractiveness is surpassed (Section 3.2.4). The strength of the saturation is different for each dependent variable, being the strongest for intelligence and trustworthiness. The difference in strength of the saturation effect is consistent with previous work that has shown that the strength of the attractiveness halo effect is trait selective [Bas81; EAML91]. The halo effect in fact is not only trait selective in strength but also in direction. While the traditional “what is beautiful is good” notion [DBW72] would suggest that increased attractiveness leads to increased positive impressions, studies have shown that an increase in attractiveness is also correlated with an increase in the perception of certain negative traits, such as vanity [HL23; RD23], egotism [DT75], materialism and sexual permissiveness [Bas81]. Scholars have tried to identify a functional basis of attributes that are used to evaluate faces [OT08], yet generalizing the findings about the halo effect to any trait is non-trivial. Since our study did not include negative attributes, it is yet unclear to which degree a potential saturation effect would be present in these situations.

However, the identified saturation effect provides a unifying explanation for several inconsistent findings reported in the literature regarding the existence [TH80; KR75] and strength [EAML91; DT75; LBH81] of the attractiveness halo effect. For example, Timmerman and Hewitt [TH80] did not find evidence of the attractiveness halo effect based on photographs of two female models from the Cosmopolitan magazine before and after professional makeup was applied. A manipulation test concluded there was a significant change in perceived attractiveness, yet no significant changes in the perceptions of their dependent attributes (including intelligence) were found. Based on our research, their findings could be an instance of the saturation effect, especially if the stimuli were highly attractive women as it could be the case given that they were selected from the Cosmopolitan fashion magazine.

Note that previous work has suggested that the halo effect and trait sensitivity could be interpreted as a stereotype effect [RD23]. In this regard, the saturation effect could be explained by the application of different stereotypes depending on the attractiveness levels of the stimuli. As discussed below, we find evidence of the existence of a gender bias when judging the intelligence of female stimuli, which could correspond to the application of a different stereotype for highly attractive females. However, our study design does not enable the establishment of a causal link between stereotype formation and the observed saturation effect. We leave to future work the study of such a link.

Concerning the existence of the attractiveness halo effect with a diverse set of stimuli according to ethnicity, age and gender, there is mixed evidence in the literature which our study contributes to disambiguating [AMD+97; BS22; GLSE21; KKM23; Wat17].

In terms of ethnicity, our findings contradict previous work that reports that the attractiveness halo effect does not generalize when evaluating members of an ethnicity other than their own [ASS+16]. Conversely, we find strong evidence of the existence of the attractiveness halo effect for all stimuli across ethnicities, even when evaluated by participants of a different ethnicity. Therefore, we conclude that the attractiveness halo effect does generalize when evaluating members of an ethnicity other than their own, in alignment with the findings reported in [BS22].

The age of the rater did not have a statistically significant effect on perceptions of attractiveness but had a statistically significant positive effect on perceived intelligence, trustworthiness and happiness after beautification. This finding complements previous work that studied the existence of the attractiveness halo effect and the baby-face stereotype in young and older adult raters [ZBL07]. The authors reported that older adults are as vulnerable as young adults to the attractiveness halo effect: they judged more attractive people as more competent and healthy, and less hostile and untrustworthy, corroborating previous research on young adults [EAML91; LKR+00]. In our work, we also find that the age of the stimulus matters. In terms of perceived attractiveness, both before and after beautification young individuals were rated as significantly more attractive than middle-aged and older individuals, in accordance with prior work [KB00; WM84; FCT06]. The negative and significant correlation between perceived intelligence, trustworthiness and age (particularly after beautification) suggests that the older the stimulus, the more intelligent and trustworthy it is perceived. This finding is aligned with previous literature that has reported on the *wisdom bias* [KP99] but contradicts recent work on trustworthiness and age [PLL+23]. Conversely, youth is positively correlated with sociability, especially after beautification which is supportive of previous research [HAS+12].

Regarding gender, our results unveil novel interactions between the gender of the stimulus, the gender of the rater, and the attractiveness halo effect, both when rating same and opposite-gender stimuli. Images of females were rated as significantly more attractive than males, in alignment with previous research [CCD71; KT82; WM84] and in contradiction to others [FCT06]. Both female and male raters provided higher ratings of attractiveness to images of females before ($p < 0.001$) and after ($p < 0.001$) beautification, with a widening gap between genders after beautification, especially for male raters (Figure 6a). Conversely, participants considered males to be more intelligent than females, particularly after beautification ($p < 0.001$), also with a widening gap between genders (Figure 6b). Therefore, we conclude that the gender of the stimulus plays a stronger role in impacting the perceptions of intelligence than perceived attractiveness given that images of females were rated as more attractive than those of males. This finding could be explained by the application of a different stereotype to highly attractive females. We leave to future work the exploration of this potential reason for this finding.

Concerning opposite-gender effects, our findings contribute with nuanced evidence of what has been previously reported [ASS+16; BK72; GPT82; AHPS23]: we observe statistically significant ($p < 0.001$) differences in the ratings provided by both female and male raters to images of opposite gender individuals for perceived attractiveness, intelligence and trustworthiness both before and after beautification, and for sociability and happiness only after beautification. As described in the previous paragraph, male stimuli are perceived as more intelligent than female stimuli both by male and female raters, with a widening gap between genders after beautification such that female stimuli are perceived as *less intelligent* on average by male raters after beautification than before applying the filter. With respect to trustworthiness, the images of females in the PRI dataset were considered to be more trustworthy by both male ($p < 0.001$) and female raters ($p < 0.001$), yet male raters considered images of males and females to have similar levels of trustworthiness after beautification. Sociability and happiness behave similarly and exhibit a widening gap between genders: men are perceived as less sociable and happy than women after beautification and especially when judged by women. In sum, we observe several and novel significant interactions between the gender of the stimulus and the gender of the human evaluators, contradicting early work that reported a lack of such an interaction [DBW72].

The findings regarding perceived intelligence suggest that there exists a stronger gender bias than the attractiveness halo effect [Rid01; EK02] and underscores deeper cultural attitudes and stereotypes surrounding gender roles and expectations [EW12]. Moreover, our results are supportive of previously reported examples of gender-based discrimination and the challenges faced by women in various spheres of life, including education and professional opportunities [GCR19; MSN96; ES09; Hei01]. The perpetuation of such stereotypes can contribute to systemic inequalities and hinder the advancement of women in society [Cor04; Rid11]. Given the prevalence in the use of beauty filters by young females—90% of women aged between 18 to 30 report using beauty filters before posting selfies on social media [Gil21]—our findings raise additional concerns about the potential negative impact of beauty filters on young women, a group that has been shown to be more susceptible to body dissatisfaction [AL17; MRS+21]. Frequent use of beauty filters has already been found to lead to anxiety, and depression, reduced self esteem, body dysmorphia, an increase of plastic surgery, feelings of inadequacy and increased pressure to conform to unrealistic beauty standards [Bak22; FM14; Esh20; Isa23; LC20; VFL+20; Gil21; Rya22]. Our research adds a new dimension to the harmful consequences of using beauty filters by empirically demonstrating that females are perceived by men as *less intelligent* after the application of the filters. Moreover, their use raises questions about authenticity and honesty as they alter the appearance of users, often presenting an idealized or unrealistic version of themselves. This alteration can blur the line between reality and artificiality, leading to questions about what is genuinely authentic in digital self-representation [Isa23]. The discrepancy between real and filtered images can undermine personal authenticity and contribute to a false sense of identity [Bar20]. There is therefore a need for transparency and ethical guidelines surrounding the use of beauty filters, especially in contexts where individuals may be influenced in their decision-making by filtered images without their knowledge.

Our study however, is not without its limitations. First, while we included a large and diverse set of stimuli judged by over 2,700 participants, the participants lacked geographic—and thereby ethnic—diversity because they consisted of predominantly white individuals from the US and the UK. As described in Section 3.4.1, participants had to be native English speakers to qualify for the study as it was designed and deployed in English. Nonetheless, previous work has reported that the attractiveness halo effect generalizes across countries [BS22]. Second, we report findings at

an aggregate level. A per-rater level analysis, while interesting, is not possible on our collected dataset because each participant was exposed to a different set of images to maximize diversity. Third, we do not explore the impact of different beauty filters on the halo effect, but previous work has reported that popular beauty filters perform similar transformations to the faces [RPG+22]. We also do not study the potential perceived differences between real and beautified faces for the same level of attractiveness. While certainly interesting, these questions are out of the scope of this study and hence we leave them to future work. Fourth, we do not study the relationship between attractiveness and socially undesirable attributes, such as vanity or materialism. We focus on socially desirable attributes instead because adding negative characteristics to our study would entail flipping the Likert scale for those attributes which could have led to confusion in the participants. We leave to future work the investigation of the “dark side” of beauty. Fifth, a related but unexplored phenomenon in our research is the halo update effect [RD23], according to which raters update their judgments over time, especially when presented with new information [CF15; SF21; MKB20]. Since participants in our study were not presented with images of the same individual before and after beautification, we leave the investigation of the halo update effect for future research. Finally, photographs provide only a static, two-dimensional representation of individuals, lacking the multi-dimensional and dynamic nature of interactions in the physical world, where attractiveness perceptions can be influenced by factors beyond facial appearance. Hence, our findings might not generalize to real-world scenarios where attractiveness perceptions interact with other factors, such as situational dynamics, personality and social context. Nonetheless, most of the previous work that has studied this cognitive bias has adopted a similar methodology to ours [TD08; NF22; MCW15; OWLT23; SHK+18; ZBL07; Tal16; WT06; BS22; ERL10] and faces play a significant role in our judgments of the attributes studied in this work [BS22; OT08; ZBL07; Tal16; MCW15; DBW72; PUG+22; TDP+13; Tod08; GML13; ERL10].

3.4 Methods

The user study was pre-registered in the Open Science Foundation registry² and was approved by the Ethics Board of the University of Alicante with identifier UA-2023-01-19_3. All participants gave informed consent.

3.4.1 Study Participants

The study participants were recruited via the Prolific participant recruitment platform. The target sample were adults with unimpaired vision who were English native speakers. Given the purpose of the study, participants were required to be neurotypical, without any mental health condition or dyslexia and to have an approval rate of at least 85% in past studies on Prolific. The sample (N=2,748) was gender balanced: 1,375 men and 1,373 women, with ages ranging between 18 and 88 years old (age M=46.47, SD=15.09). Regarding race, 2,291 participants reported being *white*, 181 *asian*, 178 *black*, 63 *mixed* and 33 participants reported being from *other* ethnic groups. Additionally, 2 participants did not report their ethnicity. The majority of participants (94%) reported living in the United Kingdom (1,817), United States (686) or Canada (72). Most of the participants (1,482) reported having full time jobs and 260 reported being students. More details about the participants can be found in Appendix D.

²Link: <https://doi.org/10.17605/OSF.IO/AQDK9>

Participants received a compensation of 2 USD for taking part in the study, with a median completion time of 8 minutes and 45 seconds. Seventeen participants failed at least two of the four attention checks and hence were removed from the analysis and replaced by new participants, yielding a total sample of 2,748 participants.

3.4.2 Experimental Stimuli



Figure 8. Samples of male (top) and female (bottom) face images used in our study before (left) and after (right) the application of the beauty filter. As illustrated in the examples, the beauty filter modifies the skin tone, the eyes and eyelashes, the nose, the chin, the cheekbones, and the lips in order to make the person appear more attractive.

The stimuli used in the study were face images from two widely used face datasets for scientific research: the Chicago Face Database (CFD) [MCW15] and the FACES dataset [ERL10], and their corresponding beautified versions.

The CFD [MCW15], developed at the University of Chicago for research purposes, provides high-resolution, standardized photographs of 597 unique individuals (male and female faces) of varying ethnicity (self-identified White, Asian, Black, Latino) between the ages of 17 and 56. The dataset was expanded in 2020 to include images of 88 mixed race individuals recruited in the United States and 142 individuals recruited from India. While there are examples of faces with non-neutral facial expressions, we selected the images where all individuals have neutral facial expressions, yielding a dataset of 827 images. In addition to the images, the CFD dataset includes metadata about each image, such as information about physical attributes (e.g., face size) and subjective ratings by independent judges (e.g., attractiveness). The set of images collected from India has ratings available from both Indian and American raters. However, we used only the ratings of American raters in order to be consistent with the ratings for other images in the dataset. While

the CFD includes a broad range of subject ages in their images, it mostly contains images of young people. Only 9% of the images are of people rated as being over 40 and it contained no images of people rated as being older than 60³. Thus, to ensure age diversity in the stimuli, we also included images from the FACES dataset [ERL10].

The FACES dataset consists of 171 images of naturalistic faces of young ($N = 58$), middle-aged ($N = 56$), and older ($N = 57$) women and men displaying each of six facial expressions: neutral, sadness, disgust, fear, anger, and happiness. The database comprises two sets of pictures per person and per facial expression, resulting in a total of 2,052 images. We selected the images corresponding to a neutral facial expression to minimize the interference of the facial expressions in the perception of attractiveness [VPP+23; ERL10]. In addition to the images, the dataset includes metadata about each image, including subjective ratings of attractiveness from independent judges.

To ensure a balanced sample across age, gender and attractiveness levels, we selected 25 images⁴ for each gender-ethnicity pair from the CFD, covering a wide spectrum of attractiveness levels: 8 images with the lowest attractiveness ratings, 8 images with the highest attractiveness levels and 9 randomly selected from the remaining images. Similarly, we selected 27 images for each gender-age group pair from the FACES dataset ensuring diversity in gender, age and attractiveness levels. Since the FACES dataset had two images for each subject, we selected one at random. This process led to a total of 462 images (300 from the CFD and 162 from the FACES dataset) which we refer to as the PRI dataset of images (**P**icked **R**epresentative **I**mages). The subset that comes from the CFD is referred to as the PRI_{CFD} dataset and similarly the images drawn from the FACES database are referred to as the PRI_{FACES} dataset of images. A summary of these datasets can be found in Table 8. Each face in the PRI dataset was *beautified* using a common beautification filter available in one of the most popular selfie editing apps in the world with over 500 million downloads. We refer to the dataset of beautified images as the POST dataset (**P**ost **S**ocial media **T**ransform). The filters were applied by running the selfie editing app on an Android emulator. An automated clicker loaded the pictures onto the application, applied the filter and then stored the transformed version.

Figure 8 shows an example of male and female original and beautified faces used in our study.

3.4.3 Procedure and Design

The study was run online by means of a custom-made web portal. After providing informed consent, each participant was presented a page with instructions: they were told that they would be shown ten face images and would be asked to provide their assessment of different aspects of the faces, based on their first impression. The exact instructions can be found in Appendix E. Participants were randomly assigned to see faces either from the FACES dataset or the Chicago Faces Database (CFD).

After reading the instructions, participants were shown one face image at a time. Each image was accompanied with a set of questions as described in Section 3.4.4. The 7-point Likert rating scales were presented as sliders with the end- and mid-points labelled. Participants were required to answer all questions about an image before being allowed to proceed to the next image. The order of the questions was randomised for each participant as per the algorithm described below,

³The CFD does not include the actual age of the participants in the pictures. Thus, the statistics about the age reported here are based on estimated age ratings from independent judges hired by Ma et al. while creating the CFD

⁴The distribution of images across gender-ethnicity pairs in the CFD is non-uniform. The smallest class (Mixed Race Males) contained 26 images, thus motivating the size of the number of images picked for each gender-ethnicity pair

	PRI		CFD	FACES
	PRI _{CFD}	PRI _{FACES}		
Size	300	162	827	171
Age	18 - 56 ^	19 - 80	17 - 56 ^	19 - 80
Gender	150M, 150 F	81M, 81F	406M, 421F	86M, 85F
Ethnicity	Asian, Black, Latino, White, Indian, Mixed	White	Asian, Black, Latino, White, Indian, Mixed	White

Table 8. Dataset Statistics. Size corresponds to the number of unique faces present in the dataset. Age is the age of the subject in the image when the picture was taken (^ as perceived by the raters used by Ma et al. [MCW15]). Actual age of the subjects in the images is not available.)

but remained the same across all the images rated by the same participant. After providing ratings for 10 images, participants reached the last page of the survey where they were asked to provide details about their background including how often they used social media and beauty filters and their self-rated attractiveness. The complete list of questions is included in Appendix E. After answering these questions, the study was complete and participants were directed to the Prolific platform where they were compensated for their time. The data collected in the study has been deposited in a public online repository [GOL24].

In addition to the questions described in Section 3.4.4, participants were also shown 4 attentiveness checks (described in Appendix E) at random points in the survey. Participants who failed two or more attention checks were rejected and additional participants were recruited to replace them.

The randomization algorithm to select the images that were shown to each participant met the following criteria to ensure a balanced sample:

- (1) Half the images were from the POST data set, i.e. had a beauty filter applied on them, and the other half were from the PRI data set. The presentation order of the images was randomized and participants were not told that some of the images were beautified. Furthermore, participants always rated images corresponding to 10 different individuals such that they never had to judge the same person in both the beautified and non-beautified conditions;
- (2) Half the images corresponded to male and the other half to female subjects;
- (3) The images were also balanced across ethnicity (for participants in the CFD condition) or across age groups (for participants in the FACES condition).

Furthermore, the images were presented such that each image received at least 25 ratings⁵. Thus, participants provided ratings on a diverse set of inputs while ensuring that each image received sufficient ratings. Note that no image received ratings from the same subset of participants. Our analyses are adjusted accordingly.

⁵After rejections, 2 images were left with 23 ratings. The largest number of ratings received by an image was 35. The mean number of ratings each image received was 29.7 with a standard deviation of 1.75

3.4.4 Measures

For each image, participants were first asked to provide the gender (male/female), age (number between 18 and 100, answered using a sliding scale) and ethnicity (Asian/Black/Latino/White/Indian/Mixed Race) of each of the faces.

Next, participants were asked to rate the person in the image on the following attributes which were randomly presented for each participant: Physical attractiveness, Intelligence, Trustworthiness, Sociability, Happiness, Femininity, and how Unusual they were. The choice of using Intelligence, Trustworthiness, Sociability and Happiness as dependent attributes for the halo effect was driven by existing literature on this cognitive bias, such as [BS22; OT08; ZBL07; Tal16; MCW15; DBW72; PUG+22; TDP+13; Tod08; GML13]. Ratings for femininity and unusualness were collected to study the impact of the beauty filters on physical appearance. Results of the analysis of the data corresponding to these two attributes have been discussed in detail in Appendix C.

The ratings for attractiveness and other attributes were provided on a seven-point Likert scale ranging from 1 = *Not at all [trait term]* to 7 = *Extremely [trait term]*. While some work collected these ratings on a nine-point Likert scale [BS22], we opted to use a seven-point Likert scale because they have been reported to be the most accurate and reliable [Joh10; Sym24; Mil56], despite the popularity of five-point scales. Each question was presented to participants as “How [trait term] is this person?”, following the same approach as previous studies in the literature [TD08; NF22; MCW15; OWLT23; SHK+18; ZBL07; Tal16]. The responses were entered on a slider initially placed at the mid-point, and where both the mid and end points were labeled. An example of the layout of the questions participants were exposed to, along with the exact phrasing of the questions, can be found in Appendix E.

3.4.5 Analysis

As seen in Figure 3, our analysis is structured according to two levels of aggregation: 1) *centralized scores*, by computing the median of the ratings provided to each image; and 2) *individual scores*, by analyzing the per-image ratings individually. As a result of our study methodology, each image received ratings from a different subset of participants such that pairwise comparisons are only possible at an aggregate level by means of the centralized scores. Furthermore, the change in the centralized scores (Δ_{ω}) is used as a measure of the impact of the filter. While analyzing the impact that the age, gender and ethnicity of the stimuli play on perceptions on attractiveness and the 4 dependent variables by means of the centralized scores, we use the actual age, gender and ethnicity of the individuals in the image instead of the age, gender and ethnicity as perceived by the raters.

In order to study the effect of the participants’ age and gender on attractiveness and the dependent variables, we also analyze each rating individually. All the variables collected in our study, except age, were collected on 7-point Likert scales i.e., they are ordinal in nature. A multinomial logistic regression approach would treat the variables as nominal thereby leading to a loss of information due to ignoring the inherent ordering of the responses. Using ordinal response models such as the Cumulative Link Model (CLM) [Chr18; Chr23] is more appropriate for ordinal response variables [Agr10] but the parameters of these models are harder to interpret than those of a linear model [Man16]. Ordered Stereotype Models (OSM) [And84; FAP16; FLC19] offer an ideal middle ground. The OSM estimates the true spacing between the points on the ordinal scale based on the data, thereby resulting in a transformed scale that is continuous and thus suitable for a linear model. In addition to the theoretical grounding, we further evaluated the appropriateness of using the OSM’s with linear models (and linear mixed models) by computing the Akaike Information

Criterion - AIC [Aka74] and the Bayesian Information Criterion - BIC [Sch78] of different models for attractiveness and each of the 4 dependent variables. We also varied the treatment of the raters as fixed or random effects in our models. We found that linear mixed models with the OSM's that treat the raters as random effects resulted in the best (lowest) AIC and BIC scores. A detailed report of this analysis can be found in Appendix B.2 and the model parameters of the resultant linear mixed models can be found in Appendix B.5.

Next, we describe in detail the Ordered Stereotype Models (Section 3.4.5) and the linear mixed models (Section 3.4.5) that we developed to perform this second level of analysis. All the modeling has been performed in R version 4.3.3 [R C21].

Ordinal Data

Ordered Stereotype Models Given an ordinal response variable Y with q categories, for an observation i , the OSM estimates the probability of $Y_i = k (k = 1 \dots q)$ as:

$$\log \left(\frac{P[Y_i = k | x_i]}{P[Y_i = 1 | x_i]} \right) = \alpha_k + \phi_k \beta' x_i \quad (1)$$

where x_i is a set of predictor covariates for observation i . The OSM additionally enforces the constraint

$$0 = \phi_1 \leq \phi_2 \leq \dots \leq \phi_q = 1 \quad (2)$$

The ϕ_k 's are interpreted as scores and help estimate the distance between different categories based on the actual data (ratings in our case) instead of assuming that all categories are equidistant. Furthermore, categories with overlapping standard deviation intervals are merged into the same category. Thus, the OSMs estimate the underlying scale by computing the expected probabilities of the categories based on potential covariates in the data.

To evaluate the impact of the stimulus's and rater's gender and age on perceptions of attractiveness, the following OSM was fit independently to the data from the PRI and POST sets:

$$Attrac \sim Gender_I + Age_I + Gender_R + Age_R \quad (3)$$

where $Gender_R$ and Age_R correspond to the gender and age of rater R respectively and $Gender_I$ and Age_I correspond to the gender and age of image I as perceived by rater R . In the case of the dependent variables (ω), perceived attractiveness was also included as a covariate:

$$\omega \sim Attrac_I + Gender_I + Age_I^R + Gender_R + Age_R \quad (4)$$

Linear Mixed Models

Below are the linear mixed models presented in Section 3.2.3. Note that linear mixed models on the re-scaled data by means of the OSM are a better fit to the data than ordinal models, such as Cumulative Link Models [Chr18], as explained in Appendix B.2. Furthermore, linear mixed models that include the raters as random effects better fit the rescaled data than models that treated the raters as fixed effects (see Appendix B.2).

$$Attrac = \beta_0 + \beta_2 \cdot Gender_I + \beta_3 \cdot Age_I + \beta_4 \cdot Gender_R + \beta_5 \cdot Age_R + \beta_6 \cdot Gender_I \cdot Gender_R + \beta_7 \cdot Age_I \cdot Age_R + RandEff_{Rater} \quad (5)$$

$$\omega = \beta_0 + \beta_1 \cdot Attrac_I + \beta_2 \cdot Gender_I + \beta_3 \cdot Age_I + \beta_4 \cdot Gender_R + \beta_5 \cdot Age_R + \beta_6 \cdot Gender_I \cdot Gender_R + \beta_7 \cdot Age_I \cdot Age_R + RandEff_{Rater} \quad (6)$$

The above models consider the stimulus's age (Age_I) and gender ($Gender_I$) as perceived by the rater, the rater's self reported gender ($Gender_R$) and age (Age_R) and the interactions between these variables. Race was not included as a variable in the analysis because the previously reported results with the centralized ratings revealed no significant impact of race neither on attractiveness nor on the dependent attributes. Additionally, the participants' self-reported race was predominantly white (see Appendix D) and hence it was also not considered as a variable in the models. Note that β_1 is omitted from the linear mixed model of attractiveness (Equation (5)) to maintain consistency in the terminology, since the linear mixed models of the dependent variables (Equation (6)) use β_1 for attractiveness. The models were fit independently on the PRI and POST sets. The parameters of all the linear mixed models can be found in Appendix B.5.

3.4.6 Data Accessibility

The data collected by us during the survey and the associated code can be found in the following Zenodo repository: <https://zenodo.org/doi/10.5281/zenodo.13836854>. The code can also be found in the following GitHub repository:

https://github.com/ellisalicante/theBeautySurvey_Analysis.

The images from the PRI set are publicly available and access instructions can be found in the Readme file on Zenodo and GitHub. The POST set images used in this study were created using a common beautification filter available in one of the most popular selfie editing apps as indicated in the manuscript. Our agreement with the application provider does not allow publicly sharing the beautified images. For further queries about this data, please contact the legal department of ELLIS Alicante at info@ellisalicante.org

Chapter 4

Algorithmic Lookism

Chapter summary

The previous chapter studied the impact of the attractiveness halo effect in humans. In recent years however, there have been significant advancements in computer vision which have led to the widespread deployment of face image analysis and generation systems in socially relevant applications, from hiring to security screening and the prevalence of biases within these systems has raised significant ethical and social concerns. The most extensively studied algorithmic biases in this context are related to gender, race and age. Yet, phenomena such as *lookism* (i.e., the preferential treatment of individuals based on their physical appearance) are equally pervasive and harmful. Lookism, studied extensively in people, remains under-explored in computer vision. It can have profound implications not only by perpetuating harmful societal stereotypes but also by undermining the fairness and inclusivity of AI technologies. This chapter highlights the need for the systematic study of lookism as a critical bias in computer vision models. Through a comprehensive review of existing literature, the chapter identifies three areas of intersection between lookism and computer vision. We illustrate them by means of examples and a user study. We propose the concept of *algorithmic lookism* and advocate for an interdisciplinary approach to address the effect in AI systems, urging researchers, developers, and policymakers to prioritize the development of equitable computer vision systems that respect and reflect the diversity of human appearances.

This chapter is based on the paper:

[GLO24] Aditya Gulati, Bruno Lepri, and Nuria Oliver. “Lookism: The overlooked bias in computer vision.” ECCV 2024 workshop on “Fairness and ethics towards transparent AI: facing the challenge through model Debiasing (FAILED)”. arXiv:2408.11448 (2024)

Computer vision systems are increasingly used to support decisions that impact critical aspects of people’s lives, including hiring processes, security screening, social media interactions and healthcare diagnoses [GLLA17; GKG+23; EGY+21]. Thus, there is a growing need to detect, quantify and mitigate the biases that such systems may perpetuate or even amplify [HvdMG+22]. As machine learning models increasingly influence high-stakes decisions across domains, including hiring [RBKL20], healthcare [PTN19], education [BH22], social service provision [Gil16] and criminal justice [TPB+22], measuring and mitigating algorithmic bias is a priority which is reflected in

existing or upcoming regulation, such as the European AI Act⁶.

While there is significant work in the literature that has focused on gender [WFVL19; WQK+20; SKB+20], racial [YAAB20; HFBK24; KF21] and age [KJ21b; JOM+19] biases, there is growing awareness of the existence of subtler biases that need to be accounted for [KYE24]. *Lookism* is one such bias. It consists of the preferential treatment of individuals based on their physical appearance. Rooted in societal standards of beauty and attractiveness and on our own cognitive biases [DBW72; Tal16; EAML91; TK74], lookism can lead to unequal treatment and reinforce harmful stereotypes when embedded in AI systems.

The oversight of lookism as an important bias to consider in computer vision systems can result in systemic disadvantages and discrimination for individuals who do not conform to prevailing aesthetic norms, affecting their opportunities and how they are perceived and judged by automated systems. This chapter aims to underscore the importance of studying and mitigating lookism within computer vision systems. By examining the roots and implications of this bias, we can develop strategies to ensure that computer vision algorithms promote fairness and inclusivity.

4.1 Lookism: A cognitive bias perspective

Lookism is deeply rooted in human cognitive biases [TK74]. The most prominent cognitive bias impacting lookism is the attractiveness halo effect, which has been shown to impact humans in the physical world [DBW72; MK75; EAML91], and seen to also impact humans in digital settings in Chapter 3. Understanding lookism from the lens of cognitive biases can shed light on how it manifests itself in computer vision systems. In addition to the halo effect, two other cognitive biases that play a role in lookism are:

- 1. Aesthetics heuristics.** Heuristics are mental shortcuts that simplify decision-making [GB09]. Aesthetic heuristics refer to the use of physical appearance as a fast and easy method to make judgments about a person's other qualities. While these heuristics can be efficient, they often lead to oversimplified and biased evaluations that do not accurately reflect the individual's true abilities or characteristics [DBW72; Tal16].

- 2. The confirmation bias** is the tendency to search for, interpret, and remember information that confirms one's preexisting beliefs or stereotypes [Nic98]. In the context of lookism, this bias could make people more likely to notice and remember positive behaviors from attractive individuals and overlook or rationalize negative behaviors.

4.2 Lookism and computer vision

The interplay between lookism and computer vision is two-fold. First, the advent of computer vision-based beauty filters could help mitigate the presence of this cognitive bias in humans by equalizing beauty. Second, similarly to humans, lookism might also be present in computer vision algorithms—leading to the concept of *algorithmic lookism*—from at least two perspectives. When computer vision systems are trained on datasets that reflect human biases, they can inadvertently learn and perpetuate lookism. For instance, if a facial recognition system is trained on images that disproportionately associate certain physical features with positive traits, it may develop a biased algorithm that favors those features. This can lead to unfair treatment and reinforce societal

⁶<https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>

biases in automated decision-making processes and image generation systems. Furthermore, image generation and multimodal generative AI systems could be also impacted by lookism, leading to representational biases in the content they generate.

In this section, we provide a brief overview of each of these aspects.

4.2.1 Beauty Filters and Lookism

Beauty filters are a particularly popular family of computer vision based face filters which aim to *beautify* the face of the person by automatically applying changes to the skin, the eyes and eyelashes, the nose, the chin, the cheekbones, and the lips. They rely on computer vision and augmented reality methods and their prevalence, with millions of users worldwide, profoundly impacts user self-presentation, raising questions about authenticity, self-esteem [Per20], mental health [CDAR20], diversity [RPG+22] and racism [RO22].

Despite the issues raised by the use of beauty filters in daily life, they are a potentially powerful tool to study lookism in vision-based automated decision making systems since they enhance perceived attractiveness of individuals in images while preserving identity [RPG+22]. Additionally, beauty filters could be seen as a tool to mitigate lookism as they enable the democratization of beauty: used by millions of users to improve their appearances, particularly on social media and other digital spaces, beauty filters help level the playing field by allowing everyone to present themselves in ways that conform to societal beauty standards. This widespread accessibility could reduce the gap between those who naturally fit these standards and those who do not, potentially decreasing the social pressure and discrimination based on physical appearance. The AHEAD dataset, introduced in Chapter 3 serves as an ideal base to study the impact of attractiveness on AI systems.

4.2.2 Algorithmic Lookism

Machine learning algorithms are typically trained on data which is annotated by humans. Thus, patterns of bias present in annotations provided by human raters are often present in these models [MSY20; CJ21]. A well-known example was in the hiring system deployed by Amazon which showed a strong bias against female employees [Das22].

We highlight below open questions associated with lookism in computer vision, followed by a discussion that highlights the challenges and ethical implications associated with studying lookism.

Lookism in Decision-making Systems

Gender and ethnicity-based biases have been studied extensively in computer vision systems that support human decisions in a variety of tasks, including emotion recognition [XWKG20], face recognition [RLH+20], video surveillance [LM19] and hiring [Ngu15; KRWS20]. More recently, Multimodal Large Language Models (MLLMs) have also been evaluated for biases based on gender and ethnicity [KCY+23; BHS+21], but there is limited work evaluating these models for biases due to physical appearance.

Interestingly, a beauty bias has been reported in LLMs [KSK23] which have exhibited significant positive correlations between attractiveness and personality traits (extraversion and conscientiousness) when automatically assessing personality from video transcriptions of job interviews [ZWC+24]. Regarding MLLMs, Howard et al. [HFBK24] have studied the impact of gender, race

and physical appearance on predictions made by MLLMs by evaluating the description provided by these models on a large set of counterfactual image pairs [HML+23].

Given the limited body of research on this topic, further investigation is required to determine the extent to which lookism influences vision-based decision-support models. This thesis provides one of the first systematic analyses of this phenomenon, with a foundational study presented in Chapter 5.

Lookism in Image Generation



(a) “An intelligent person”



(b) “An unintelligent person”



(c) “A very competent person”



(d) “A very incompetent person”

Figure 9. Illustrative examples of lookism in T2I models. Images generated by DALL-E3 through Microsoft’s Copilot with prompts: “create a hyperrealistic portrait of (a) an intelligent person; (b) an unintelligent person; (c) a very competent person; and (d) a very incompetent person”.

Text-to-image (T2I) models are a class of machine learning models designed to generate images based on textual descriptions. They leverage recent advances in NLP and image generation methods to create a broad range of visual content, including representations of humans. The textual description is tokenized and projected to word or contextual embeddings from models like BERT [DCLT18], GPT [AAA+23] or more recent transformer-based architectures such as CLIP [RKH+21a]. The image generation in early models was carried out using GANS [GPM+14] which have been superseded by variational autoencoders [KW13] and diffusion models [BGJ+23]. These

systems typically include attention layers to help the model focus on specific parts of the text when generating the corresponding parts of the image, improving the alignment between the textual and the visual elements.

While gender and racial biases have been studied in T2I models [NN23a], little attention has been paid to the impact of lookism as a bias. AI-generated faces have been found to be perceived by humans as indistinguishable yet more trustworthy than faces of real people [NF22]. Lookism would suggest that the reason these faces are more trustworthy is because they tend to be more attractive, yet an in-depth analysis to corroborate this hypothesis would be necessary.

In fact, a systematic empirical study to unveil the presence of lookism in T2I systems would entail auditing them according to multiple dimensions by providing relevant prompts on topics, such as: (1) *demographic representation*, involving the evaluation of how well the systems represent various ethnicities, genders, ages, body types and overall appearances in response to diverse prompts; (2) *cultural and contextual sensitivity*, examining the system’s ability to accurately and respectfully depict cultural symbols, attire, and settings, to assess to which degree the generated images perpetuate stereotypes or cultural insensitivity; (3) *stereotype reinforcement* to investigate whether the T2I system amplify existing societal stereotypes, particularly regarding professions, personal attributes, social roles, and activities; (4) *aesthetic diversity* to assess the range of visual styles and attractiveness standards the system produces; (5) *realism and coherence* to focus on the technical quality of the generated images, evaluating whether the images are realistic and logically consistent with the provided descriptions; and (6) *ambiguity vs specificity* to evaluate the system’s performance with both highly specific and more ambiguous prompts, testing its ability to handle nuanced and complex descriptions.

Figure 9 illustrates the presence of lookism in examples generated with DALL-E 3 via Microsoft Copilot when prompted to create images of “intelligent” versus “unintelligent” and “competent” versus “incompetent” individuals. Qualitative inspection reveals clear differences in the perceived attractiveness of the generated faces depending on the valence of the trait specified in the prompt. Despite the growing adoption of generative models, research on the role of attractiveness in shaping their outputs remains scarce. This thesis provides one of the first systematic analyses of this phenomenon, with a foundational study presented in Chapter 6.

4.3 Challenges

The study of lookism is not exempt from challenges and ethical implications.

First, attractiveness is a highly *subjective* and *cultural* construct. While the famous saying “beauty is in the eye of the beholder” suggests that perceptions of beauty vary significantly across individuals and cultures, there are studies that report an agreement across raters in perceptions of attractiveness [CRB+95; EDR06; PLP+98], or at least in perceptions of unattractiveness [SKS13].

Multimodal Large Language Models (MLLMs) are unique when compared to other machine learning methods because of their ability to be used across multiple tasks. A potential solution to address the subjectivity of attractiveness would consist of asking the MLLM to evaluate attractiveness before proceeding with the desired task. However, preliminary experiments reveal that state-of-the-art MLLMs systems suffer from a *positivity bias* and tend to assign extremely high scores to everyone, unlike the attractiveness scores given by human evaluators. In Chapter 5, our work on MLLMs seeks to address these challenges and offers an initial examination of how lookism manifests within such systems.

Second, there is a *lack of awareness* of this bias and the *inconsistency* of some of findings reported in the literature. Numerous studies on attractiveness have found that it is used as a cue for other unrelated human attributes [MZW+15; JHH95], such as perceived intelligence [BS22; Tal16]. Yet, other studies have reported correlations between physical attractiveness and health [CPS09; TG06; ZR04], leading to doubts about treating lookism as a bias, even though these conflicting findings are reported in different contexts.

Third, the legal protection against lookism, or discrimination based on physical appearance significantly varies across jurisdictions. Thus, lookism is not as widely recognized or legislated against as other forms of discrimination, such as those based on race, gender, age or disability [Des10]. While there is a growing recognition of the need for more comprehensive laws and policies to address this bias, the legal protections against lookism remain inconsistent and limited.

4.4 Conclusion

Computer vision systems that exhibit lookism can perpetuate and magnify societal biases, leading to the unequal treatment of individuals based on their looks. Furthermore, the deployment of certain computer vision apps and systems, such as beauty filters, raises concerns about the erosion of diversity and the perpetuation of narrow and *white* beauty standards [RO22]. For these reasons, we believe that it is important for the computer vision community, in collaboration with experts of other domains, to devote efforts to detect, measure and mitigate lookism by ensuring diverse and representative training datasets, implementing fairness-aware algorithmic designs that consider lookism, and conducting continuous auditing and empirical evaluations to detect and rectify this bias. Addressing lookism is not only about preventing discrimination but also about fostering a more inclusive and equitable society where the algorithms that we design respect and reflect the rich diversity of human appearances.

Chapter 5

Algorithmic Lookism in the Evaluation of Faces by MLLMs

Chapter summary

Physical attractiveness matters. It has been shown to influence human perception and decision-making, often leading to biased judgments that favor those deemed attractive. While extensively studied in human judgments in a broad set of domains, including hiring, judicial sentencing or credit granting, the role that attractiveness plays in the assessments and decisions made by multimodal large language models (MLLMs) is unknown. To address this gap, we conduct an empirical study with 7 diverse open-source MLLMs evaluated on 91 socially relevant scenarios and a diverse set containing 924 face images sourced from the AHEAD dataset. Our analysis reveals that attractiveness impacts the decisions made by MLLMs in **86.2%** of the scenarios on average, demonstrating substantial bias in model behavior in what we refer to as an *attractiveness bias*. Similarly to humans, we find empirical evidence of the existence of the attractiveness halo effect in **94.8%** of the relevant scenarios: attractive individuals are more likely to be attributed positive traits, such as trustworthiness or confidence, by MLLMs than unattractive individuals. Furthermore, we uncover gender, age and race biases in a significant portion of the scenarios which are also impacted by attractiveness, particularly in the case of gender, highlighting the intersectional nature of the algorithmic attractiveness bias. Our findings suggest that societal stereotypes and cultural norms intersect with perceptions of attractiveness in MLLMs in a complex manner. Our work emphasizes the need to account for intersectionality in algorithmic bias detection and mitigation efforts and underscores the challenges of addressing biases in modern MLLMs.

This chapter is based on the paper:

[GDS+25] Aditya Gulati, Moreno D’Incà, Nicu Sebe, Bruno Lepri, and Nuria Oliver. “Beauty and the Bias: Exploring the Impact of Attractiveness on Multimodal Large Language Models.” Eighth AAI/ACM Conference on AI, Ethics and Society arXiv:2504.16104 (2025)

5.1 Introduction

Physical attractiveness plays an invisible yet powerful role in human judgments and decision-making. Research since the 1970s, including our study described in Chapter 3, has consistently shown that individuals deemed physically attractive are often perceived more favorably across a variety of positive traits, including intelligence [DBW72; Kan11; Tal16; GMF+24b], sociability [Mil70], trustworthiness [Tod08; GMF+24b], happiness [MK75; GML13; GMF+24b] or success in life [DBW72]. While studied and understood in human contexts, the extent to which an *attractiveness bias* –i.e., a differential treatment of individuals exclusively based on their perceived attractiveness– and the *attractiveness halo effect* –i.e., the attribution of positive yet unrelated traits to individuals who are perceived as attractive– exists in algorithmic systems remains largely under-explored.

With the advent and wide adoption of multimodal large language models MLLMs⁷ [WGC+23; ZYD+24], algorithmic biases, including the attractiveness bias and the attractiveness halo effect, have become a growing concern. Unlike large language models (LLMs) that rely solely on textual data, MLLMs are able to process both textual and visual inputs, enabling them to interpret and generate responses based on complex multimodal information. These capabilities make MLLMs highly versatile, powerful and applicable to numerous vision-and-language tasks, ranging from image captioning, visual question answering [HXL+24] and content creation⁸ to conversational AI grounded in visual context [PDH+23; DZZ+24]. However, these models can inadvertently replicate or even magnify appearance-based biases such as the attractiveness bias. When MLLMs are trained on datasets containing images and descriptions of individuals, these models may implicitly treat attractiveness as a relevant factor, raising two key concerns: (1) MLLMs may exhibit an *attractiveness bias*, hence making decisions or judgments that differ based solely on an individual’s perceived attractiveness; and (2) MLLMs, like humans, may behave according to the *attractiveness halo effect*, thus favoring or assigning positive traits to individuals who are perceived as attractive, even when those traits are unrelated to the person’s actual abilities or character. In both cases, MLLMs could lead to a different or even preferential treatment of individuals who are considered to be more attractive by the model, potentially influencing critical decisions in areas such as hiring recommendations or educational assessments. Therefore, investigating the existence of an attractiveness bias and the presence and mechanisms of the attractiveness halo effect in MLLMs are essential to ensure fair outcomes in their deployment.

This chapter addresses these issues by conducting an empirical evaluation of seven open-source MLLMs of different sizes, listed in Table 9. A unique aspect of our methodology is the use of beauty filters to enhance the attractiveness of the face images, enabling a controlled evaluation of how perceived attractiveness influences model outputs when all the other variables remain constant. Our study, summarized in Figure 10, is guided by the following research questions:

RQ1: Do MLLMs exhibit an attractiveness bias, making different decisions or judgments based on an individual’s perceived attractiveness?

RQ2: Do MLLMs exhibit the attractiveness halo effect, attributing positive traits, such as honesty or trustworthiness, to attractive individuals?

⁷In this chapter, we use the term MLLM to refer to models that integrate text and images.

⁸<https://openai.com/index/sora-is-here/>,
https://www.wired.com/story/linkedin-ai-generated-influencers/?utm_source=chatgpt.com

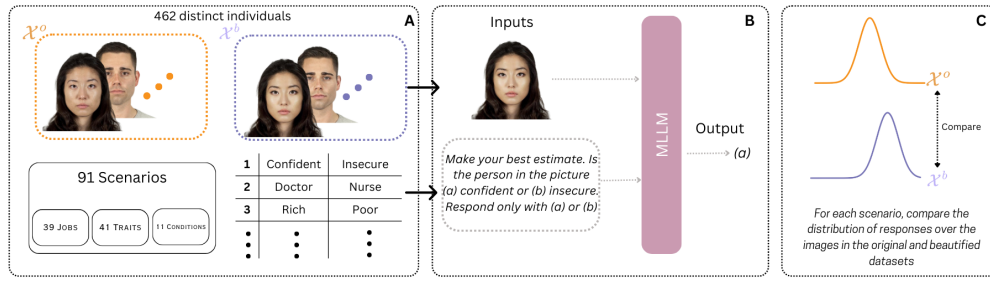


Figure 10. Overview of the adopted experimental methodology to study the existence of an attractiveness bias in multimodal large language models (MLLMs). (A) As facial stimuli, we use a diverse set of face images of 462 distinct individuals (\mathcal{X}^o) and their corresponding *beautified* versions (\mathcal{X}^b) after applying a beauty filter, which enables to control for attractiveness. We define 91 scenarios in socially relevant areas (jobs, traits and conditions) consisting of pairs of words. (B) For each scenario, the MLLM is provided as input a face image and a textual prompt with a question about the person in the image which has two possible answers (a) or (b). (C) To measure the existence of an attractiveness bias, we compare the distributions of the answers provided by the MLLM when prompted with \mathcal{X}^o vs \mathcal{X}^b . A reliance on attractiveness by the model would lead to statistically significant differences in the answers. We run the scenarios with three different seeds to ensure robustness in the results.

RQ3: How do gender, age and race intersect with attractiveness to influence model outputs?

Contributions. The main contributions of this chapter are:

- We propose using the AHEAD dataset, a diverse dataset of 462 original and their corresponding beautified faces to study the impact of attractiveness in the decisions made by MLLMs.
- We design 91 scenarios and propose a methodology to examine the existence of an attractiveness bias and the attractiveness halo effect in 7 distinct open-source MLLMs.
- We study the interplay of attractiveness with age, gender, and race.
- We discuss the implications of our findings regarding the design and use of MLLMs.

The rest of the chapter is organized as follows: we first provide a summary of the most relevant related research in Section 5.2, followed by a detailed description of the adopted methodology in Section 5.3. Sections 5.4 and 5.5 present our results and their discussion, respectively. Our conclusions and future research directions are summarized in Section 5.6.

5.2 Related Work

5.2.1 Biases in Large Language Models (LLMs)

Given the versatility of LLMs and their usefulness for a wide range of tasks, bias evaluation in LLMs has addressed a broad spectrum of scenarios and demographic groups. In fact, a variety of social biases [YCLH23] and biases related to reasoning and decision-making [ISRB24] have been reported in LLMs, including sentiment [HZJ+20], religious [AFZ21] and stereotype [NBR21] biases. Gender biases in majority [ZWY+18; VGB+20; KDS23] and minority [HHG+24] languages, and disparities in the representation of various demographic groups [NGP+23; LML24; DdIFG+24] have

also been studied in isolation and where multiple sensitive attributes, such as gender and race, are considered simultaneously [MCW+23]. These studies have highlighted that disparities are often more pronounced for intersectional minorities, such as Black women [WC24a].

Another socially relevant line of research has examined biases in professional contexts, including occupational stereotypes [KJV+21], discrimination against individuals with disabilities [VSW22], gender biases in accounting scenarios [LS24], and both gender and racial biases in hiring [WC24b], as well as social biases manifesting in code generation tasks [LRWY25]. In response to the growing concern over such biases, some authors have proposed toolkits to systematically evaluate and quantify social biases in LLMs [BSS24].

Recent work has studied more subtle forms of bias present in LLM outputs, such as biases in the political ideology [VPK25] and implicit stereotypes embedded in model associations [BWSG25; ZWW25]. Recent findings suggest that while techniques such as reinforcement learning from human feedback (RLHF) and increased training data can effectively reduce explicit biases, they are considerably less successful in mitigating implicit biases [ZWW25].

Cognitive biases have garnered increasing attention in the study of LLMs. A cognitive bias is a systematic pattern of deviation from rationality that occurs when humans process, interpret or recall information from the world, and it affects the decisions and judgments we make, leading to inaccurate judgments, illogical interpretations and perceptual distortions [KT79; AJ08]. Cognitive biases have been found to affect the decisions of workers engaged in fact-checking tasks [DLS+22] and data annotators performing face annotation tasks [HGW+24], which, in turn, can be propagated into LLMs or other AI systems that rely on these inputs. Furthermore, several authors have investigated to which degree psychological experiments traditionally conducted with human participants can be replicated with LLMs to assess the existence of cognitive biases in LLMs, including in specific domains such as operations management [CKO+25]. The results have been mixed.

While Koo *et al.* [KLR+24] identified a significant misalignment between human and LLM responses, they did report the presence of specific cognitive biases, such as the egocentric and order bias, and the bandwagon effect. Several authors have reported the existence of other cognitive biases, namely the anchoring and framing effects [TF23; ELA+24], group attribution and primacy biases [ELA+24] and the base rate neglect [TF23] in LLMs. However, there is not clear indication on the existence of the status quo bias [ELA+24]. Scholars have recently proposed that cognitive science insights should be integrated into LLM evaluation frameworks [ELX+24] and LLM-generated recommendations have been shown to be manipulated by embedding cognitive biases into product descriptions [FDL+25]. These works underscore both the theoretical and practical significance of accounting for cognitive biases in the behavior of LLMs.

5.2.2 Biases in Multimodal LLMs (MLLMs)

Beyond LLMs, there is significant interest in understanding and mitigating biases in MLLMs, particularly those that process both text and images. Research on CLIP models has revealed multiple bias dimensions, including societal categories, such as race, gender, and ethnicity [HZG+24], as well as cultural biases favoring Western norms [ASV+24]. Additionally, studies have shown that different genders and ethnicities are associated with distinct sentiments in model outputs [CBB+24].

Efforts to benchmark and systematically analyze biases in MLLMs have expanded the field. SafeBench [YLL+24] and VLBiasBench [ZWC+24] provide frameworks for evaluating harms and biases using synthetic images. Datasets like VLStereoSet [ZLJ22] and VisoGender [HAZ+23] have been developed to evaluate stereotype biases across visual tasks. Both general taxonomies to

evaluate fairness in MLLM decisions [AKW+23], domain-specific studies [GHA+24] and broader explorations of biases in CLIP across dimensions like religion, nationality, disability, and sexual orientation [JD23], highlight the pervasive nature of bias in MLLMs.

Attractiveness Biases in MLLMs.

The processing of visual information in MLLMs raises concerns about potential appearance-based biases, such as the attractiveness bias, which could influence decisions made by these models, both in an implicit –*e.g.*, when evaluating candidates for positions– and explicit –*e.g.*, in post-surgical evaluations of plastic surgery facial procedures [AC25]– manner, thereby amplifying the necessity of understanding how visual appearance influences model behavior.

However, little research has studied the role of visual appearance in the decision-making processes of MLLMs. Hamidieh *et al.* [HZG+24] make initial inroads by incorporating attractiveness-related terminology in their assessment of societal biases within CLIP, yet their analysis does not directly examine how specific variations in appearance affect model outputs. Some large benchmarks have attempted to incorporate controlled attractiveness-related variables through synthetically generated images of individuals, exploring dimensions such as body size (*e.g.*, fat versus thin) [HFBK24] and facial appeal (*e.g.*, attractive versus unattractive) [ZWC+24]. While such benchmarks offer valuable opportunities for systematic study, their dependence on synthetic imagery raises notable methodological concerns. Specifically, the generative models employed to create these images often reflect and perpetuate existing societal biases [NN23a; dCG24], which can confound efforts to isolate and assess the impact of attractiveness on biased perceptions in MLLMs.

In this chapter, we address these limitations and investigate the existence of an attractiveness bias in MLLMs by means of a case study with seven open-source MLLMs asked to provide judgments about 924 human faces in 91 different scenarios corresponding to stereotyped jobs, traits and conditions.

5.3 Methodology

The adopted experimental methodology is summarized in Figure 10. As seen in the Figure, we formulated 91 different scenarios (described in Section 5.3.2) and probed the MLLMs by asking scenario-specific questions to assess the model’s attractiveness bias. In all scenarios, the input consisted of a facial image and a textual prompt containing a question about the image with 2 possible answers that the model chose from.

5.3.1 Models

We evaluated seven open-source MLLMs, as detailed in Table 9. Probing a diverse set of models with different number of parameters enables a robust assessment of the existence of an attractiveness bias in MLLMs. We excluded API-based models, such as GPT-4, because the datasets from which the face images were sourced [ERL10; MCW15] explicitly prohibit the use of facial images with API-based LLMs due to privacy and data protection concerns. Moreover, GPT-4 incorporates a layer of safeguards that in many cases prevents the model from responding when the input consists of a single face, which illustrates the challenges of testing for biases in closed, black-box models.

Model Name	Size (# parameters)
Gemma3 [Tea25]	4B
Phi 3.5 [AAA+24]	4.2B
DeepSeek [WCP+24]	4.5B
Molmo [DCL+24]	7B
Qwen2 [WBT+24]	7B
Pixtral [AAH+24]	12.7B
LLaVA 1.5 [LLWL23]	13B

Table 9. MLLMs evaluated in our analysis and their size

5.3.2 Inputs

Face Stimuli. We leverage a dataset that was created to study the attractiveness halo effect in humans [GMF+24b]. It consists of a curated set of faces obtained from the Chicago Faces Database (CFD) [MCW15] and the FACES database [ERL10]. The CFD provides a diverse set of face images in terms of race, with equal representation of individuals self-identifying as *Asian, Black, Latino, Indian, White, or Mixed race*, yet of similar ages. Conversely, the FACES database contains face images of different ages, categorized into three age groups: *young, middle-aged, and old*, yet with no racial diversity. As a result, the dataset comprises 462 distinct faces, with 25 images for each gender-ethnicity pair and 27 images for each gender-age group pair. All images feature individuals wearing identical clothing, displaying neutral facial expressions, and set against uniform backgrounds to minimize potential confounds. Furthermore, for each original face image there is its corresponding beautified version created using a common beautification filter available in one of the most popular selfie editing apps. Thus, for each individual in the dataset, there are the original (non-beautified) and the beautified versions, yielding a total of 924 images.

In addition to the images, the dataset contains the self reported gender, race and age labels for the subjects in the images before the application of beauty filters. This information is considered to be the ground truth regarding the demographic information of each image. In their study with human participants, Gulati et al. [GMF+24b] found no significant impact of the beauty filters on perceptions of gender and race by the human evaluators. The dataset also contains the individual, median and mean ratings of each image provided by at least 25 human raters on a 7-point Likert scale according to the following attributes: attractiveness, intelligence, trustworthiness, sociability, happiness, femininity and unusualness. Figure 10 displays examples of the original and beautified versions of two faces from the dataset. Note that **96.1%** of the faces in the dataset were rated as more attractive after applying a beauty filter and no individual was rated as less attractive after beautification. Thus, the dataset contains a diverse set of pairs of facial images where the only difference between them is attractiveness, with minimal confounds. We direct the interested reader to [GMF+24b] for a detailed description of the dataset.

By asking the MLLMs to make judgments about one face at a time, our methodology aligns with standard practices in studies involving human participants [ERL10; BS22; TD08; GMF+24b; PUG+22; BK72; GPT82].

Textual Prompts. The textual prompts for each of the 91 scenarios consisted of a question referring to the face image and two possible answers, indicated as (a) or (b), from which the model had to choose. Since prompt order impacts the behavior of LLMs [LY21] and MLLMs [SML+24;

[CCZ+24; LDZ+25], each image was evaluated under all possible orderings of the options and the average across all these orderings was used to evaluate the model, as detailed in Section 5.3.5. For example, a prompt could involve presenting the model with an image and asking it to determine whether the individual in the image is *confident* or *insecure*. A specific prompt for this scenario would take the form:

Make your best estimate. Is the person in the picture a (a) confident or (b) insecure. Respond only with (a) or (b).

For the same scenario, the other three possible prompts would be:

Make your best estimate. Is the person in the picture a (a) confident or (b) insecure. Respond only with (b) or (a).

Make your best estimate. Is the person in the picture a (a) insecure or (b) confident. Respond only with (a) or (b).

Make your best estimate. Is the person in the picture a (a) insecure or (b) confident. Respond only with (b) or (a).

The model is constrained to select between two options, deliberately disallowing a neutral response, to minimize noise in the responses, as including neutral responses would complicate the systematic measurement of bias in accordance with the proposed evaluation benchmark. While it could be argued that model creators may not have explicitly designed their models for forced-choice studies, these are general purpose models that are deployed in real-world scenarios where users expect them to generate responses. If biases exist, they will manifest regardless of whether the model was “designed” for this type of answer or not.

5.3.3 Scenarios

In total, we defined 91 scenarios, structured in 3 socially relevant categories where human biases, including the attractiveness halo effect, have been reported in the literature: stereotyped jobs, traits and conditions. Figure 11 depicts an overview of the categories and sub-categories. Each scenario presents a binary decision task, with one option designated as the “Stereotyped Choice”. Across all these scenarios, we test if attractiveness impacts decisions made by the model and also how attractiveness intersects with social variables like gender, age and race. A comprehensive list of all the scenarios, and the specific choices presented to the MLLMs, is provided in Appendix H.

Stereotyped Jobs.

Scenarios in this category involve pairs of occupations traditionally associated with specific gender (eg., “Doctor” vs “Nurse”) and racial groups (eg., “Cleaner” vs “Security guard”) or different levels of expected attractiveness (eg., “Model” vs “Makeup artist”). These scenarios aim to assess both attractiveness and societal biases rooted in stereotypes about professional roles. A total of 39 scenarios were selected in this category, informed by previous work [CK85; HSC03] and data from




Figure 11. Visual depiction of the 91 scenarios used in the experiments: 39 scenarios were about jobs, further divided in 3 types as shown in the Figure; 41 scenarios referred to traits; and 11 scenarios referred to stereotyped conditions

the 2022 labor force characteristics report by the U.S. Bureau of Labor Statistics (BLS)⁹. The full list of scenarios in this category can be found in Table 27 in Appendix H. The scenarios in this group were divided into three types:

i) *Gender-stereotyped jobs*, consisting of occupations traditionally associated with one gender in Western societies, such as nursing and caregiving for women and engineering and leadership roles for men. The 19 job pairs were chosen from existing datasets of stereotyped jobs [XLC+24; FK24], which in many cases were selected based on the latest labor force characteristics report by the U.S. Bureau of Labor Statistics¹⁰ (BLS).

ii) *Race-stereotyped jobs*, referring to pairs of occupations that are frequently associated with specific racial or ethnic groups in Western societies, and particularly in the United States, thereby reinforcing biases regarding the abilities, interests, or roles of individuals based on their race, rather than their personal skills or qualifications. The selection of the 12 job pairs in this group was guided by the latest labor force characteristics report by the BLS. For the racial categories present in the report that intersected with the racial categories available in the dataset –namely, Asian, Black, Latino, White– we identified a subset of occupations for each race where it represented a significant majority. From these, we randomly selected 12 pairs of jobs to define the race-stereotyped job scenarios, ensuring that all possible race pairs were included.

iii) *Attractiveness-stereotyped jobs*, involving pairs of occupations where physical attractiveness is genuinely beneficial or required for the role. These scenarios were included to evaluate whether

⁹<https://www.bls.gov/opub/reports/race-and-ethnicity/2022/>

¹⁰<https://www.bls.gov/cps/cpsaat11.htm>

the model incorporates facial attractiveness as a factor in decision-making and to examine how the reliance on attractiveness manifests itself both in contexts where attractiveness is relevant and where it is not. Given the lack of existing datasets specifically addressing attractiveness-stereotyped jobs and the absence of this information in the BLS reports, we leveraged the capabilities of GPT-4 to generate job pairs within similar domains, where one job relies on physical appearance and/or requires frequent interaction with customers or clients and the other does not. Three of the authors iteratively refined this list to generate 8 attractiveness-stereotyped job pairs.

Stereotyped Traits.

These scenarios feature pairs of behavioral traits, where one trait is generally perceived as positive and desirable –such as trustworthiness or confidence– and the other as negative and undesirable –such as insecurity or hostility. The scenarios in this category were designed to evaluate the extent to which the model associates certain personal characteristics, whether favorable or unfavorable, with perceptions of attractiveness.

In total, we defined 41 scenarios that consisted of pairs of behavioral traits, selected from datasets used to analyze stereotypes in vision-language models [HZG+24; JLS+24; ZLJ22] and are listed in Tables 28 and 29 in Appendix H. In particular, we selected a subset of behavioral traits most likely to exhibit an attractiveness bias from the 374 words proposed in the So-B-IT taxonomy [HZG+24]. To accomplish this, we created two clusters of words: one consisting of positive appearance-related terms (*e.g.*, “attractive”, “beautiful”) and the other with negative appearance-related terms (*e.g.*, “unattractive”, “ugly”). We then identified the top 6 words from So-B-IT that were the closest to each cluster and used GPT-4 to generate the antonyms of these words. A related methodology was employed to extract relevant word pairs from VLStereoSet [ZLJ22]. In addition, ModSCAN [JLS+24] includes gender-stereotyped hobby pairs, which were also incorporated into the stereotype traits to complete the set of 41 scenarios.

Stereotyped Conditions.

Scenarios in this category involve pairs of societal conditions or statuses. Previous research has shown that attractive individuals are often perceived as more likely to be successful in life and physical attractiveness can influence perceptions of one’s capabilities and future success [DBW72]. Additionally, wealth is frequently associated with success, and attractiveness may amplify this perception, with wealthier individuals often being perceived as more capable or deserving of their status [BD20; WTG21]. We selected the stereotyped conditions from existing visual stereotype sets [HZG+24; ZLJ22]. These conditions provide insight into how societal perceptions of attractiveness might influence broader judgments of social status, such as wealth, success, or competence. The scenarios were grouped into subcategories based on the specific stereotypes from which they originated, such as economic status (2), immigration status (2), place of residence (5), education (1) or religious beliefs (1), leading to 11 scenarios in this category, listed in Table 30 in Appendix H.

5.3.4 Model Evaluation

In our experiments, we prompt the MLLMs to mimic a non-expert user in a zero-shot setting, which reflects one of the most common real-world use cases. Accordingly, we treat each MLLM as a *black-box*, keeping its hyperparameters (*e.g.*, system prompt, temperature, etc...) unchanged.

The evaluation is conducted across seven distinct open-source multimodal large language models, each varying in model architecture and parameter scale. For each scenario, we conducted a forward pass of each MLLM with each of the 924 images as input, together with the corresponding prompt. We repeated each scenario with the four possible orderings of the textual prompt and computed the mean response. To enhance the robustness and reproducibility of our findings, we report results averaged over three random seeds (*i.e.*, 0, 42, and 742). Thus, we generated responses from $91 \text{ scenarios} \times 924 \text{ images} \times 4 \text{ orderings} \times 3 \text{ seeds} \times 7 \text{ models}$, resulting in 7,063,056 prompts being evaluated.

All experiments were executed on a computing cluster equipped with NVIDIA Ampere A40 GPUs (46GB). The codebase developed to run the models and analyze the data for this study can be found in the project repository.

5.3.5 Problem Formulation and Metrics

Notation. We denote with \mathcal{X} the dataset of faces, where \mathcal{X}^o is the set of original images and \mathcal{X}^b is the set of corresponding beautified images. Note that the dataset includes ground truth metadata corresponding to the gender, age, and race of the individual in each image.

In the following, we denote with a subscript specific demographic subsets in the dataset, *e.g.*, the set of female faces is referred to as \mathcal{X}_{female} .

The set of questions (prompts) for all scenarios $\mathcal{S} = \{s_i\}_{i=1}^N$ is defined as:

$$\mathcal{Q} = \{q_{i,j} \mid i = 1, \dots, N; j = 1, \dots, M\}$$

where $q_{i,j} \in \mathcal{Q}$ is a scenario-specific question with two candidate choices, as introduced in Section 5.3.3; j indicates a specific order of the choices within q_i ; M represents the total number of choice orderings (4 in our case); and N describes the total number of scenarios (91 in our case). For a given scenario $s_i \in \mathcal{S}$, we denote with \tilde{c}_i the “*Stereotyped Choice*”, *i.e.*, the option which, if selected by the model, reflects a stereotypical response. A complete list of scenarios and their corresponding choice sets is provided in Appendix H.

Response metric. Given a face $x \in \mathcal{X}$, and a scenario under study $s_i \in \mathcal{S}$, we define the model’s responses as:

$$\hat{c}_{i,j,k} = \text{MLLM}(x, q_{i,j}, k)$$

where $k \in \{1, \dots, K\}$ represents each of the K seeds (3 in our case). For each tested scenario s_i , we collect $M \times K$ responses ($4 \times 3 = 12$ in our case) for each image corresponding to the M choice orderings and K seed combinations. To measure the number of stereotyped responses by the models, we define the Stereotype Consistency Score (SCS) as:

$$\text{SCS}_{i,j,k} = \begin{cases} 1 & \text{if } \hat{c}_{i,j,k} = \tilde{c}_i \\ 0 & \text{otherwise} \end{cases}$$

Finally, we obtain an *order-invariant score*, ϕ_i , for each scenario s_i by averaging the SCS across all M orderings and K seeds in scenario s_i :

$$\phi_i(x) = \frac{1}{K} \sum_{k=1}^K \frac{1}{M} \sum_{j=1}^M \text{SCS}_{i,j,k}$$

Given that the model must choose between two options, this score effectively captures the empirical distribution of stereotyped responses in each scenario s_i .

Biases. We define a *bias* in the model’s response as a tendency to disproportionately associate images from a specific group (*e.g.*, beautified, women, young, etc.) with one of the two options presented in each scenario. We measure the presence of a bias by comparing the order-invariant scores, ϕ_i , of each scenario given by the models to individuals of different groups (*e.g.* non-beautified, men, old, etc.).

1. *Attractiveness Bias, H_i^{attr}* : We quantify an attractiveness bias in a scenario s_i by comparing the distribution of order-invariant scores given to the original ($\phi_i(x^o), \forall x^o \in \mathcal{X}^o$) and the corresponding beautified ($\phi_i(x^b), \forall x^b \in \mathcal{X}^b$) faces by means of a Kruskal-Wallis test, *i.e.*, we compare the responses provided for the *same* individuals with and without a beauty filter applied. If the distributions are statistically significantly different ($p < 0.01$), we consider that there is an attractiveness bias.

2. *Attractiveness Halo Effect*: The attractiveness halo effect is a specific case of the attractiveness bias where the model associates positive traits with attractive individuals. We measure the attractiveness halo effect on the stereotyped traits scenarios where the Stereotyped Choice corresponds to the positive traits. A model exhibits an attractiveness halo effect if it has an attractiveness bias *and* it is more likely to associate beautified images (x^b) with the positive traits when compared to the original images (x^o).

3. *Gender, Age and Race Biases*: In addition to measuring an attractiveness bias, we also evaluate the influence of gender, age, and race on the responses generated by the MLLMs. To measure biases related to gender (male vs. female), age (young vs. middle-aged vs. old), and race (Asian vs. Black vs. Indian vs. Latino vs. Mixed-race vs. White), we adopt a similar methodology to that employed for measuring the attractiveness bias but considering the complete set of faces \mathcal{X} . A gender bias is therefore defined as:

$$H_{i,\mathcal{X}}^{gender} = KW(\{\phi_i(x)|\forall x \in \mathcal{X}_{male}\}, \{\phi_i(x)|\forall x \in \mathcal{X}_{female}\})$$

where KW denotes the Kruskal-Wallis test.

Age and race biases are similarly defined over the age and race categories. Given that racial diversity is only represented in the CFD dataset, racial biases can only be computed on the (original and beautified) faces from the CFD dataset. Likewise, age biases are reported exclusively for the (original and beautified) faces from the FACES dataset. Gender biases are assessed across the entire image set, \mathcal{X} .

4. *Intersectional Effects*: We also evaluate how gender, age, and race impact the attractiveness bias and vice versa by comparing the strength of each bias across groups of the dependent variable. For example, we measure the impact of attractiveness on the gender bias by evaluating the gender bias on the original ($H_{i,\mathcal{X}^o}^{gender}$) and the beautified ($H_{i,\mathcal{X}^b}^{gender}$) images for each scenario s_i and then comparing them pairwise across all 91 scenarios using a Wilcoxon Paired Rank Test (WPRT) as indicated below:

$$W_{attr}^{gender} = WPRT(\{(H_{i,\mathcal{X}^o}^{gender}, H_{i,\mathcal{X}^b}^{gender})|s_i \in \mathcal{S}\})$$

Across all tests, statistical significance and hence the existence of a bias is determined by p-values < 0.01 , and the corresponding significance levels are explicitly reported for each test. In the case of the intersectional effects, the Bonferroni correction is applied to address the multiple comparisons problem [Rup+12]. Post the correction, the same alpha value is used as in the other tests.

	Total (91)	Jobs [■]			Traits [■]		Conditions [■]		
		Gender (19)	Race (12)	Attractiveness (8)	Sentiment (33)	Other (8)	Geography (4)	Wealth (5)	Other (2)
Gemma	89.0%	78.9%	91.7%	100.0%	93.9%	75.0%	100.0%	100.0%	50.0%
Phi3.5	79.1%	84.2%	83.3%	75.0%	84.8%	50.0%	100.0%	60.0%	50.0%
DeepSeek	90.1%	84.2%	91.7%	100.0%	93.9%	100.0%	75.0%	60.0%	100.0%
Molmo	80.2%	68.4%	66.7%	87.5%	90.9%	62.5%	100.0%	100.0%	50.0%
Qwen2	81.3%	68.4%	66.7%	100.0%	87.9%	62.5%	100.0%	100.0%	100.0%
Pixtral	86.8%	68.4%	83.3%	100.0%	100.0%	75.0%	75.0%	100.0%	50.0%
LLaVA 1.5	80.2%	63.2%	75.0%	75.0%	97.0%	75.0%	75.0%	80.0%	50.0%
Average	83.8%	73.7%	79.8%	91.1%	92.6%	71.4%	89.3%	85.7%	64.3%

Table 10. Percentage of scenarios for each category where a statistically significant ($p < 0.01$) attractiveness bias was observed. The shaded column indicates scenarios where the attractiveness halo effect was studied, and bold values indicate the largest value in every column. (·) denotes the number of scenarios per category.

5.4 Results

In this section, we address the three previously formulated research questions by evaluating the responses of the seven MLLMs to the 91 previously described scenarios.

RQ1: Do MLLMs Exhibit an Attractiveness Bias?

Table 10 summarizes the findings regarding the existence of an attractiveness bias in MLLMs. As seen in the Table, an attractiveness bias, *i.e.*, a statistically significant difference (Kruskal-Wallis, $p < 0.01$) in the distribution of ϕ_i between the original (\mathcal{X}^o) and beautified (\mathcal{X}^b) faces, was found in **83.8%** of scenarios on average across all the models indicating that facial attractiveness is used as a cue when MLLMs are provided faces of people as inputs.

The highest levels of attractiveness bias are observed in Gemma and DeepSeek where attractiveness impacted the decisions in **89.0%** and **90.1%** of the scenarios, respectively. Phi3.5 showed the lowest average attractiveness bias, with attractiveness affecting decisions in **79.1%** of the scenarios – still a substantial proportion.

The influence of facial attractiveness was evident across stereotyped jobs, traits, and conditions, suggesting that MLLMs systematically rely on attractiveness as a decision-making cue. As expected, attractiveness mattered in **91.1%** of the attractiveness stereotyped jobs. However, even in contexts where attractiveness provides no meaningful information, it impacted a significant portion of decisions. This indicates a pervasive and potentially unwarranted reliance on attractiveness by MLLMs across diverse decision contexts, which has been understudied in the literature to date.

RQ2: Do MLLMs Exhibit the Attractiveness Halo Effect?

To investigate the attractiveness halo effect, we focus on the 33 sentiment-related scenarios from the stereotyped traits category, where the “Stereotyped Choice” reflects a positive trait and the alternative represents its negative counterpart (*e.g.* trustworthy vs untrustworthy). The list of choice pairs used can be found in Table 28 in Appendix H.

Statistically significant differences were observed (Kruskal-Wallis, $p < 0.01$) in **92.6%** of these scenarios on average across models. In all scenarios and for all models (except for 3 out of 31 scenarios for DeepSeek and 1 out of 30 for Qwen2) the beautified images were associated with the *positive* traits and the differences with the original images were significant. This provides strong evidence of the attractiveness halo effect in MLLMs, suggesting that, like humans, these models tend to associate attractive faces with positive traits.

The complete list of scenarios, the H_i^{attr} values, significance levels and ϕ_i for each scenario s_i can be found in Tables 31 - 37 in Appendix I.1.

RQ3: How Do Gender, Race, and Age Intersect With Attractiveness To Influence Model Outputs?

The design of the scenarios and inputs enables the evaluation of gender, race, and age biases in MLLMs. We first assess each bias *independently* of attractiveness, followed by an examination of their intersection.

RQ3.1: Gender, Age and/or Race Biases.

Model	Bias		
	Gender	Age	Race
Gemma	69.2%	67.0%	53.8%
Phi3.5	78.0%	70.3%	67.0%
DeepSeek	78.0%	70.3%	68.1%
Molmo	82.4%	62.6%	60.4%
Qwen2	74.7%	74.7%	71.4%
Pixtral	74.7%	75.8%	69.2%
LLaVA 1.5	78.0%	63.7%	46.2%
<i>Average</i>	76.45 ± 3.80	69.23 ± 4.70	62.32 ± 8.65

Table 11. Percentage of scenarios where a statistically significant gender, age, and race bias is observed.

We evaluate gender, age, and race biases *independently* of attractiveness in Table 11, where, on average, significant gender (Kruskal-Wallis, $p < 0.01$), age (Kruskal-Wallis, $p < 0.01$), and race (Kruskal-Wallis, $p < 0.01$) biases are observed in **76.45%**, **69.23%** and **62.32%** of scenarios, respectively. These findings show the consistent existence of such biases across multiple models, and align with existing research on LLMs [MCW+23; WC24a; KDS23] and MLLMs [FK24; CBB+24; ZWC+24].

We further investigate whether MLLMs not only provide different responses depending on gender, age, and race, but also whether their responses reflect prevailing societal stereotypes. To this end, we evaluate both gender and racial biases by examining the models’ outputs on the gender and race stereotyped job scenarios respectively.

1. *Gender-stereotyped jobs.* On average across models, **93.2%** of the gender-stereotyped job scenarios exhibited statistically significant effects (Kruskal-Wallis, $p < 0.01$), with responses varying by the gender of the face. Phi3.5 showed gender-based differences in all 19 such scenarios. Crucially, in every instance where a significant effect was identified, MLLMs were more likely to

associate male faces with male-stereotyped jobs, replicating social gender stereotypes and hence exhibiting a gender bias. Appendix I.2 details the scenarios and effect sizes across models.

2. *Race-stereotyped jobs.* A statistically significant effect (Kruskal–Wallis, $p < 0.01$) of race was found in 35.7% of scenarios across models when comparing subgroups defined by race stereotypes (see Table 27 in Appendix H). In 93.3% of the cases where a significant effect was detected, the MLLMs were more likely to associate images in ways that conformed to prevailing social stereotypes tied to the respective job pairs as indicated in the Table. The list of race-stereotyped occupations and corresponding effect sizes across all models is in Appendix I.3.

3. *Stereotyped conditions.* The most pronounced effects in this category were found in the scenarios concerning geography-based stereotypes, particularly regarding race biases. Images of White individuals were significantly less likely to be classified as “immigrant” or “foreign” compared to other racial groups, whereas images of Indian individuals were the most likely to be classified as such. This finding aligns with prior research suggesting that LLMs tend to reflect the biases of WEIRD (Western, Educated, Industrialized, Rich, and Democratic) societies [AXP+23]. Similarly, Black individuals were the least likely to be classified as “educated”, although this effect, while statistically significant, was relatively small. No significant effects of age or gender were observed in the education-related stereotypes. While some statistically significant effects were observed across these categories, the overall effect sizes were generally smaller than those reported for other types of biases.

RQ3.2: Impact of Attractiveness on Gender, Age and Race Biases.

As detailed in Section 5.3.5, we assess whether gender, age, and racial biases exhibit significant differences when evaluated on the original vs the beautified faces. The outcomes of the Bonferroni-corrected Wilcoxon Paired Rank tests, conducted across all 91 scenarios, are presented in Table 12. A higher value indicates a larger difference in the strength of the bias between the original and beautified faces, and the color indicates for which group of images (original in orange [■] and beautified in purple [■]) the bias is stronger if significant. As seen in the Table, gender biases are exacerbated in the beautified images for all models except for Qwen2 and Pixtral. In contrast, age and race biases are stronger in the original images for some models, indicating that the application of beauty filters appears to attenuate racial and age biases in those models. These findings are consistent with prior research involving human participants [GMF+24b], and their broader implications are elaborated in Section 5.5.

RQ3.3: Impact of Gender, Age and Race on the Attractiveness Bias.

We further investigate whether the impact of the attractiveness bias varies across different gender (W_{gender}^{attr}), age (W_{age}^{attr}), and racial groups (W_{race}^{attr}). Thus, we compute the attractiveness bias independently for each subgroup and conduct Bonferroni-corrected Wilcoxon Paired Rank tests across the 91 scenarios for every pair of subgroups to assess whether the observed differences are statistically significant. Table 13 presents the differential strength of attractiveness bias between images of males and females. Across all evaluated models, we consistently find a stronger attractiveness bias associated with female images, in alignment with previous findings with human participants [GMF+24b].

The influence of age and race on the attractiveness bias is detailed in Appendix G.1 and Appendix G.2, respectively. Although the results are less consistent than those observed for gender, a general trend emerges: the attractiveness bias tends to be stronger for middle-aged individuals

Model	W_{attr}^{gender}	W_{attr}^{age}	W_{attr}^{race}
Gemma	1524.0***	877.0	172.0
Phi3.5	1875.0***	511.0	24.0***
DeepSeek	1830.0***	492.0	306.0
Molmo	1851.0***	743.0	74.0**
Qwen2	1294.0	723.0	228.0
Pixtral	1326.0	298.0***	15.0***
LLaVA 1.5	1735.0***	576.0	76.0

Table 12. Results of the Wilcoxon paired rank test to evaluate whether gender ($^{gender}W_{attr}$), age ($^{age}W_{attr}$) and race ($^{race}W_{attr}$) biases are different with and without the filters applied to images. *** denotes $p < 0.001$ and ** denotes $p < 0.01$. The colors are used to indicate if the bias was stronger in the beautified [■] or original [■] images when the difference was statistically significant.

compared to older individuals. No significant difference in attractiveness bias is observed between young individuals and either middle-aged or older groups. In terms of racial group differences, the attractiveness bias appears to be significantly weaker for images of Asian and Black individuals. However, no significant difference is observed between these two groups. The broader implications of these findings are discussed next.

Model	Attractiveness Bias
Gemma	135.0***
Phi3.5	228.0***
DeepSeek	125.0***
Molmo	193.0***
Qwen2	124.0***
Pixtral	52.0***
LLaVA 1.5	165.0***

Table 13. Results of the Bonferroni-corrected Wilcoxon paired rank test to evaluate if the attractiveness bias is different for images of males and females (W_{attr}^{gender}). The stars denote significance and the color indicates if the attractiveness bias is stronger for images of males [■] or females [■]

5.5 Discussion

Attractiveness matters... Our study provides compelling empirical evidence for the existence of an *attractiveness bias* influencing judgments made by MLLMs. While attractiveness is hard to study due to its highly subjective nature, our methodology relies on beauty filters that increase the attractiveness of individuals without impacting their identity thus minimizing confounds. Previous studies involving human participants that rated the images used in our study [GMF+24b] confirmed that individuals were perceived as significantly more attractive when the beauty filter was applied, thus validating the attractiveness manipulation used in this work. We observed statistically significant differences in **86.2%** of the scenarios where the MLLMs evaluated images of

the *same individuals* before and after applying a beauty filter indicating that this bias is prevalent in the decisions made by MLLMs.

...depending on who you are. The attractiveness bias, although robust and statistically significant, does not affect all individuals uniformly. Our analyses indicate that the attractiveness bias disproportionately impacts judgments of **women** when compared to men, **middle-aged** individuals when compared to older adults and it had the smallest impact on Asian and Black individuals. These findings suggest that MLLMs exacerbate existing societal disparities, reinforcing stereotypes and prejudices that disproportionately disadvantage specific demographic groups, particularly women, placing higher importance on attractiveness for these groups when making decisions.

What is beautiful is good in MLLMs too. While numerous human-based studies have established the existence of the attractiveness halo effect in human decision-making processes [DBW72; Tal16; GMF+24b], our findings extend this phenomenon to MLLMs. Similarly to humans, we find that MLLMs have a tendency to associate attractive individuals with positive traits. Specifically, in **92.6%** of the tested scenarios on average across models, MLLMs demonstrated a statistically significant preference for associating beautified images of individuals with positive descriptors compared to their original, unaltered counterparts. These findings raise significant concerns, as they indicate that attractiveness substantially biases the evaluations of MLLMs, even in contexts where physical appearance should have no impact on decisions made.

Expanding the bias discourse. Consistent with existing literature, our findings indicate that MLLMs exhibit biases based on gender, age, and race. Furthermore, we see that they also perpetuate harmful stereotypes, especially concerning gender and racial categories. While an extensive body of research addresses biases related to demographic variables such as gender, age, and race – particularly focusing on amplification [HvdMG+22; WR21], and mitigation strategies [PS22; MMS+21] – relatively less attention has been dedicated to the biases introduced by non-demographic factors such as attractiveness.

Our research provides compelling evidence demonstrating that attractiveness significantly influences the decision-making processes in MLLMs, comparable in magnitude to traditional demographic variables. Crucially, the associations between attractiveness and positive traits occurs in a non-transparent and implicit manner. This opacity could induce naive end-users to the mistaken belief that MLLM-generated decisions are objective and unbiased. Given the accelerating adoption of these models in high-stakes scenarios, such as recruitment and professional evaluations¹¹, the subtle yet consequential biases related to attractiveness can inadvertently lead to discriminatory outcomes or unjustified preferential treatment.

Thus, our study underscores the urgent need for research in the design, development, and validation of MLLMs. It highlights the imperative to extend bias-mitigation strategies beyond the traditional emphasis on demographic factors, proactively addressing cognitive biases such as the attractiveness halo effect.

Interaction with other biases makes mitigation hard. The identified attractiveness bias does not operate in isolation but interacts with other societal biases, including those related to gender,

¹¹<https://www.theguardian.com/technology/2019/oct/25/unilever-saves-on-recruiters-by-using-ai-to-assess-job-interviews>

age, and race. Our analysis reveals that this bias is significantly more pronounced when evaluating females. This finding aligns with prior human-subject studies, which have demonstrated that the attractiveness halo effect exerts a stronger influence in the evaluation of female faces [KKM23; GMF+24b]. In addition to the heightened attractiveness bias for female subjects, we also observed an amplification of gender biases in the subset of images that had been altered using beauty filters. This observation mirrors patterns identified in human assessments of beautified faces [GMF+24b].

We also observed a relative attenuation of racial and age related biases in the beautified image sets for certain MLLMs. We hypothesize that this differential behavior with race biases could be due to the homogenization of the facial features across different racial groups after applying beauty filters as reported in [RO22; RPG+22]. Beauty filters have also been found to make people look younger [GMF+24b], which could explain the relative reduction of the observed age related biases for middle-aged individuals.

The influence of the attractiveness bias varies across different demographic groups, complicating efforts toward effective bias mitigation. Prior research has shown that mitigation strategies targeting a single demographic attribute – such as gender – can unintentionally intensify biases associated with other attributes, such as race [RSMA24]. These unintended cross-demographic interactions may similarly apply to biases arising from perceptions of attractiveness. Moreover, existing approaches designed to counteract gender or race-based biases may inadvertently reinforce attractiveness related biases if these dimensions are not explicitly considered during model design and evaluation. Such inter-dependencies underscore the complexity of bias mitigation in MLLMs and highlight the need for further research into holistic, intersectional mitigation frameworks that address both traditional demographic variables and less-studied factors like attractiveness.

Limitations. Our work is not exempt from limitations. First, we do not evaluate API-based models such as GPT-4, as our experimental paradigm involves providing facial images to the models as inputs. Due to the lack of transparency related to data handling by API providers, and restrictions imposed by the dataset licenses (which specifically prohibits the use of images in such contexts) we excluded these models from our evaluation. Nevertheless, our analysis spans seven different open-source models, within which we observe statistically significant and consistent patterns, underscoring the robustness and prevalence of attractiveness-related biases. Second, while we carefully designed and curated our inputs to minimize confounding variables, our empirical evaluation may not fully capture the diversity of real-world contexts where attractiveness plays a role across cultures and social settings. Third, we have used a discrete number of categories for gender, age and race, as given by the ground truth labels of the dataset, which inevitably fails to capture and represent the diversity in society. Fourth, while it is possible to explore various prompt formulations, we limited our experiments to varying random seeds to ensure robustness. Our focus was on determining whether MLLMs respond differently to variations in facial attractiveness, an effect that our results consistently confirm. A systematic investigation of how this bias may be influenced by different prompting strategies remains an avenue for future research. Lastly, while we examine intersections with gender, age and race, we have not considered all relevant social factors, including socioeconomic status or disability, which could further impact the identified biases.

5.6 Conclusion

In this chapter, we have studied the role that attractiveness plays in the decision-making processes of MLLMs by evaluating seven different models in 91 scenarios with over 900 face images and more than 7,000,000 prompts. Our findings provide strong evidence of the existence of an attractiveness bias in decisions made by MLLMs which can manifest in ways that mimic the attractiveness halo effect, a cognitive bias observed in humans. We also find evidence of the complex interplay between attractiveness and demographic factors –namely gender, age, and race– in driving the decisions of the MLLM. This interaction between attractiveness and other factors increases the complexity of interventions to mitigate biases in MLLMs. We hope that not only our results, but also our dataset and controlled experimental design, will serve as tools for the research community to further measure and understand these biases in MLLMs.

Chapter 6

Algorithmic Lookism in Synthetically Generated Faces

Chapter summary

Having established that attractiveness influences the decisions of multimodal large language models (MLLMs), we next examine whether algorithmic lookism is present in the faces synthetically generated by AI algorithms. Specifically, we investigate whether images generated with positive trait descriptors tend to be rated as more attractive than those generated with negative descriptors. The work presented in this chapter is the result of an empirical evaluation of 13,200 images generated by Stable Diffusion 2.1 using two complementary approaches: (1) assessing attractiveness as encoded within the generative model itself, and (2) approximating human ratings of attractiveness to determine whether similar effects would be perceived by human observers.

Beyond these direct evaluations, we further study the impact of algorithmic lookism on downstream applications. In particular, we find a drop in the performance of gender classification models when analyzing faces generated with negative attributes, and particularly faces of women. Finally, we explore the potential of the AHEAD dataset to be used to mitigate this negative impact. Our results empirically show that fine-tuning the gender classification models with a small sample of the data from the AHEAD dataset can significantly improve their performance.

This chapter is partially based on the paper:

[[DGMO25](#)] Miriam Doh, Aditya Gulati, Matei Mancas, and Nuria Oliver. “When Algorithms Play Favorites: Lookism in the Generation and Perception of Faces.” Fourth European Workshop on Algorithmic Fairness (EWAF). arXiv:2506.11025 (2025)

6.1 Introduction

Previous chapters have studied the impact of attractiveness on decisions made by humans and by multimodal large-language models when exposed to faces of real humans. Another domain in which attractiveness may play a significant role is in the algorithmic generation of synthetic faces. Generative AI models are increasingly deployed to produce online content, shaping social perceptions and cultural narratives [[RAH+24](#); [KD24](#); [RAH+24](#); [EHoH+23](#)]. Moreover, synthetic content that

has been posted online could be incorporated into broader datasets, potentially influencing other machine learning models that are trained on internet-scraped data. Consequently, biases embedded in generative systems can propagate across the AI ecosystem, amplifying their societal impact. Hence, it is of paramount importance to study the presence and consequences of biases in AI-based image generation models. Of particular interest not only for this thesis, but also for society, is the automatic generation of human faces by means of generative AI techniques.

Substantial prior research has documented biases in AI-generated images of humans in relation with gender [WFVL19; WQK+20; SKB+20], age [KJ21b; JOM+19] and race [YAAB20; HFBK24]. Other studies have also shown that not all demographic groups are represented equally, with whiteness and masculinity being disproportionately overrepresented [LAMJ23; NN23b]. Furthermore, biases extend beyond demographic traits, affecting object selection, clothing, spatial representations [WNG23]. Such skewed representations can distort public perceptions of social categories and reinforce problematic stereotypes. Importantly, these biases also risk creating feedback loops, where AI-generated content shapes human attitudes, which in turn influence future datasets and models.

In contrast, the role of human attractiveness in AI-generated imagery has been relatively understudied, despite its importance. Marketing and consumer behavior research has long demonstrated and utilized the power of attractiveness in shaping perceptions and decision-making [BC77]. Advertising research has shown that advertisements featuring attractive models are more persuasive, lead to more favorable brand evaluations, and even enhance consumer recall of product information [PC89; TB00]. Moreover, the repeated use of attractive models in marketing contributes to the internalization of beauty ideals, shaping self-perceptions and aspirations, often with negative consequences for self-esteem and body image [MG97].

If generative AI systems similarly amplify attractiveness as a marker of positive traits, they may extend and intensify these well-documented effects into new digital domains. In particular, if synthetic faces systematically portray “attractive” individuals as competent or likable and “unattractive” individuals as untrustworthy or undesirable, these systems risk reproducing the attractiveness halo effect at scale, misleading individuals into forming false inferences about others. Given evidence that even subtle environmental cues can affect human judgments, the stakes of attractiveness bias in generative AI are considerable.

Attractiveness as a bias remains rarely studied in this context, partly because the phenomenon is less well recognized compared to other social biases, and partly due to the methodological challenges arising from the inherently subjective nature of attractiveness. In this chapter, we present an initial study that examines how generative AI systems encode and reproduce attractiveness biases, providing a first step toward systematically understanding this phenomenon.

We report experiments on a dataset of faces generated with Stable Diffusion 2.1 using prompts that mirror the attribute set employed in Chapter 3. We operationalize and quantify attractiveness bias in this dataset according to two complementary ways: (i) via the definition of attractiveness encoded within Stable Diffusion 2.1, and (ii) via a predictive model fine-tuned on the AHEAD dataset introduced in Chapter 3. Methodological details for both measurements are provided in Section 6.2.

We further evaluate the impact of these attractiveness-biased synthetic images on downstream tasks, focusing on gender classification as an example. We test three commonly used gender classifiers: InsightFace [RLG+23], DeepFace [SO21], and FairFace [KJ21a]. Our results show systematic performance differences in the classifiers, with particularly pronounced degradation of performance on images of females that were generated using negative trait descriptors in the prompt.

While the effect is relatively weak in FairFace, it is substantially stronger in DeepFace and InsightFace. These findings suggest that algorithmic lookism in synthetic images not only shapes human perception but also propagates into other AI-driven tasks, especially as the prevalence of synthetically generated images of human faces continues to grow.

Correcting for these biases is challenging, particularly given the difficulty of controlling the proliferation of biased synthetic images online. To address this, we examined whether the AHEAD dataset could bring value in mitigating these issues given that it was explicitly designed to study the attractiveness halo effect - a key mechanism underlying algorithmic lookism. Using DeepFace as a case study, we find that fine-tuning with a small subset of the images in the AHEAD dataset significantly improves performance of the classifier on the synthetic faces. Detailed results of this evaluation are presented in Section 6.3.

In summary, this chapter addresses the following research questions:

- **RQ1:** Do faces generated by Stable Diffusion 2.1 associate attractiveness with positive unrelated attributes and unattractiveness with negative attributes?
- **RQ2:** Are downstream models, specifically gender classifiers, sensitive to the attractiveness of the synthetically generated faces?
- **RQ3:** Can the AHEAD dataset aid in the mitigation of these effects in downstream classifiers?

6.2 Methods

Each of the research questions outlined above required a different methodological approach, which we detail in the following sections.

6.2.1 Synthetic Face Dataset Creation

The images of the dataset used in this chapter were generated using Stable Diffusion 2.1 [RBL+22], a widely adopted open-source text-to-image (T2I) diffusion model. Each image was created by means of a prompt that included information of two demographic variables (gender and race) and one descriptive attribute. In the generation of the images, we consider two gender categories (male, female), three race categories (Asian, Black, White), and the four attribute pairs (positive vs. negative) that were previously validated in Chapter 3 as effective in eliciting the attractiveness halo effect in human judgments, namely: “intelligent” vs. “unintelligent”, “sociable” vs. “unsociable”, “trustworthy” vs. “untrustworthy” and “happy” vs. “unhappy”. In addition to the attribute pairs, we generated images using the attributes “attractive” and “unattractive.” This design allows for a systematic investigation of how demographic factors and attribute valence interact with attractiveness.

In sum, each face image was generated using the following prompt template:

Front photo of a [attribute] [race] [gender]

For each [gender, race, attribute] combination, we created 200 images, yielding $2 \times 3 \times 10 \times 200 = 12,000$ images in total. Furthermore, we created a neutral baseline in which no attribute was specified, generating 200 images for each [gender, race] demographic group. This contributed a



Figure 12. A representative sample of images generated using Stable Diffusion 2.1 with one example shown for each attribute–race–gender triplet. Columns correspond to the attributes used to generate the images: A = Attractiveness, H = Happiness, I = Intelligence, S = Sociability, and T = Trustworthiness. Rows indicate the valence of the attribute (positive [+] or negative [-]). The first two rows are images of females generated by Stable Diffusion and the bottom two rows are images of males. The three panels, from left to right, correspond to the race labels Asian, White, and Black.

further $2 \times 3 \times 200 = 1,200$ images, for a final dataset of 13,200 images generated and then further evaluated.

Figure 12 presents a representative set of generated images, with one sample shown for each attribute–race–gender triplet. This image set will be made available for research purposes upon request. We chose to not release the dataset publicly in order to avoid further contributing to the proliferation of synthetic facial image collections that encode and reproduce social disparities, and that risk misrepresenting individuals and groups.

columns correspond to the attributes used to generate the images: A = Attractiveness, H = Happiness, I = Intelligence, S = Sociability, and T = Trustworthiness. Rows indicate the valence of the attribute (positive [+] or negative [-]). Results are further disaggregated by gender: the top two rows (outlined in green) show accuracies for male faces, while the bottom two rows (outlined in yellow) show accuracies for female faces.

6.2.2 Algorithmic Lookism in Synthetically Generated Faces

To answer RQ1 i.e., to evaluate the presence of lookism in images generated by Stable Diffusion, we used two complimentary approaches which are described in this section.

First, we clustered in the embedding spaces of two models (CLIP and ArcFace) the face images generated according to each attribute and measured the proximity of the centroids of these clusters to reference clusters formed by images prompted with “attractive” and “unattractive” keywords. If the images generated with positive traits are closer to the “attractive” cluster and vice-versa, that would be an indication of algorithm lookism.

The second approach leveraged the ratings collected from the participants in the user study described in Chapter 3 provided to the face images in the AHEAD dataset.

Together, these two analyses test for a systematic alignment between the positive attributes and attractiveness, as perceived by both humans and AI models.

Algorithmic Lookism as per the Embedding Spaces

To evaluate the existence of algorithmic lookism in the generated faces, we projected the generated images on two different embedding spaces: CLIP [RKH+21b], which provides general-purpose vision–language representations learned from large-scale image–text pairs and captures high-level semantic structure; and ArcFace [DGY+19; SO21], which yields face-specific embeddings optimized for identity discrimination and is widely adopted in facial recognition settings. The use of both a general-purpose and a face-specialized representations reduces the dependence on potential inductive biases of a single embedding space. Thus, a convergence of results on both the CLIP and ArcFace embedding spaces increases the confidence that the observed associations between attractiveness and other attributes reflect properties of the faces themselves rather than representation-specific artifacts.

Each generated image has a corresponding **demographic group** that was part of the prompt, defined as a pair

$$g = (\text{gender}, \text{race}) \in G \times R$$

where $G = \{\text{man}, \text{woman}\}$ and $R = \{\text{Asian}, \text{Black}, \text{White}\}$ are the set of gender and race labels specified in the Stable Diffusion prompt during image generation. In addition, each image is associated with an attribute label a . This label is written as a^+ when referring to the positively valenced attribute and as a^- when referring to its negatively valenced counterpart. When the valence is not of relevance, the attribute is denoted simply as a .

Thus, for each demographic group $g = (\text{gender}, \text{race})$ and attribute $a \in \mathcal{A}$, we compute two distance distributions:

$$\mathcal{D}_g^{(a \leftrightarrow \text{attractive})} = \left\{ \|\mathbf{e}_i^{(a)} - \mathbf{e}_j^{(\text{attractive})}\|_2 : \forall i = 1, \dots, 200, \forall j = 1, \dots, 200 \right\} \quad (7)$$

$$\mathcal{D}_g^{(a \leftrightarrow \text{unattractive})} = \left\{ \|\mathbf{e}_i^{(a)} - \mathbf{e}_k^{(\text{unattractive})}\|_2 : \forall i = 1, \dots, 200, \forall k = 1, \dots, 200 \right\} \quad (8)$$

where $\mathbf{e}_i^{(a)}$ represents the i -th embedding vector for attribute a in a given embedding space. For each demographic group g , there are 200 images generated with the attribute label a , 200 images generated with the label “attractive” and 200 with the label “unattractive”.

We operationalize algorithmic lookism via a distributional measure $L_g^{(a)}$, defined as the difference between the expected values of the two distance distributions:

$$L_g^{(a)} = E[\mathcal{D}_g^{(a \leftrightarrow \text{unattractive})}] - E[\mathcal{D}_g^{(a \leftrightarrow \text{attractive})}] \quad (9)$$

A positive value of $L_g^{(a)}$ indicates that images generated with attribute a are systematically closer to “attractive” samples, while a negative value $L_g^{(a)} < 0$ suggests systematic proximity to “unattractive” samples. Thus, a model would exhibit algorithmic lookism if $L_g^{(a^+)} > 0$ and $L_g^{(a^-)} < 0$.

For each attribute–demographic combination, we perform statistical comparisons between the two distance distributions. We first assess distributional assumptions using the Shapiro–Wilk test for normality and Levene’s test for homogeneity of variances, then apply the most appropriate statistical test (independent t-test, Welch’s t-test, or Mann–Whitney U test). Effect sizes are quantified using Cohen’s d to assess the practical significance of observed differences.

Algorithmic Lookism as per Human Ratings

While the first approach reveals whether the model’s internal representations implicitly associate positive attributes with attractiveness, it does not necessarily establish that these associations are aligned with those of human observers. The ideal human-centric evaluation would be a large-scale user study in which participants rate all 13,200 images; however, obtaining reliable human judgments at this scale is resource-intensive and impractical. Instead, we leverage the high-quality human ratings collected for the 924 images in the AHEAD dataset to fine-tune two deep Convolutional Neural Network (CNN) models as attractiveness predictors. We apply these predictors to the synthetically generated images and consider their outputs as a proxy for human judgments. Although this proxy is not perfectly accurate it offers a scalable and internally consistent estimate of how humans would evaluate the attractiveness of the images generated using different attributes.

Attractiveness Models The attractiveness of the synthetically generated faces was estimated with two widely used CNNs, InceptionNet [SIVA16] and ResNet50 [HZRS16].

Attractiveness prediction was formulated as a ternary classification task, categorizing images as unattractive (0), neutral (1), or attractive (2). The inclusion of a neutral class mitigates the limitations of a binary high/low categorization, which would risk oversimplifying the variability inherent in human judgments. A seven-level classification scheme, as provided in the AHEAD dataset, was deemed unnecessary for this analysis, since the focus lies in identifying broad patterns of attractiveness rather than fine-grained perceptual distinctions captured by a 7-point scale.

Training labels were constructed by stratifying the distribution of ratings: images in the lowest quartile were labeled as unattractive (0), those in the highest quartile as attractive (2), and those in the interquartile range were assigned to the neutral category (1). This quartile-based discretization ensured a balanced class separation while maintaining sufficient representation across categories. Both models were fine-tuned on a randomly selected 80% split of the AHEAD dataset, with the remaining 20% reserved for validation. The training data was further randomly sub sampled during training to ensure an equal number of samples were picked from each training class.

Evaluation of Algorithmic Lookism For each demographic group g and each attribute pair (a^+, a^-) , the images in that group were provided to each of the two attractiveness models \mathbb{M} to obtain the distributions of predicted attractiveness scores, denoted $\mathbb{M}(g, a^+)$ and $\mathbb{M}(g, a^-)$ for the positive and negative attributes, respectively. To assess whether these distributions differed systematically, we applied the Mann–Whitney U-test. Evidence of algorithmic lookism was inferred when two conditions were met: (i) the expected attractiveness score for the positive attribute exceeded that of the negative attribute, i.e., $E[\mathbb{M}(g, a^+)] > E[\mathbb{M}(g, a^-)]$, and (ii) the difference between the two distributions was statistically significant. Under these conditions, we interpret the results as indicative of algorithmic lookism, consistent with human perceptions of attractiveness.

6.2.3 Impact on Downstream Tasks: Gender Classifiers

To address RQ2, we investigate whether algorithmic lookism in T2I models propagates to downstream classification tasks. Specifically, we examine whether gender classifiers exhibit differential performance on faces generated with positive versus negative attribute prompts, and whether systematic misclassification patterns emerge based on attractiveness levels.

We evaluate gender classification performance using three widely-adopted face analysis models: InsightFace [RLG+23], DeepFace [SO21], and FairFace [KJ21a]. For each demographic group

pair $d = (\text{gender}, \text{race}) \in G \times R$ and attribute $a \in A$, we compute the gender classification accuracy as:

$$\text{Accuracy}_{d,a} = \frac{1}{N_{d,a}} \sum_{i=1}^{N_{d,a}} \mathbb{I}[\hat{y}_i = y_i] \times 100 \quad (10)$$

where $N_{d,a}$ is the total number of images for demographic group d and attribute a , \hat{y}_i is the predicted gender label for image i , y_i is the true gender label specified in the generation prompt, and $\mathbb{I}[\cdot]$ denotes the indicator function. A classification is considered correct when the predicted gender matches the gender specified in the generation prompt.

We compare classification performance across attributes to assess whether faces generated with positive versus negative attributes exhibit systematic differences in gender recognition accuracy. Misclassification rates are analyzed across all gender–race combinations to identify potential biases that disproportionately affect certain demographic groups based on the attractiveness bias encoded in the generated faces.

This analysis aims to determine whether algorithmic lookism in T2I models creates systematic vulnerabilities in automated gender recognition systems, potentially leading to the differential treatment of individuals based on their attractiveness levels embedded in facial representations.

6.2.4 Algorithmic Lookism Mitigation

We explore the potential of the AHEAD dataset to mitigate the negative impact of algorithmic lookism on gender classification by means of fine-tuning experiments using DeepFace. Given our hypothesis that differences in facial attractiveness impact classifier performance, we used two subsets of the AHEAD dataset corresponding to different attractiveness levels:

1. A random selection of faces from the *PRI* subset of the AHEAD dataset, containing images without beauty filters i.e., relatively less attractive images
2. A random selection of faces from the *POST* subset of the AHEAD dataset, containing images of the same individuals with beauty filters applied i.e., higher levels of attractiveness

This design allowed us to examine whether fine-tuning on images with lower or higher levels of attractiveness would reduce bias. Additionally, we tested multiple gender ratios in the fine-tuning sets, using combinations with more female faces than male faces, since classification errors were disproportionately higher on faces of women.

We fine-tuned with a cross entropy loss iterating for fewer than 10 epochs until convergence was reached without signs of overfitting. The corresponding implementation is available in the project repository. The results of this fine-tuning procedure are presented in Section 6.3.

6.3 Results

In this section we provide a summary of the results obtained to address the previously defined research questions.

6.3.1 RQ1: Do faces generated by Stable Diffusion 2.1 tend to associate attractiveness with positive attributes?

Algorithmic Lookism as per the Embedding Spaces

	M-Asian	M-Black	M-White	W-Asian	W-Black	W-White
I (+)	0.625*	0.359*	-0.438*	0.615*	0.076*	-0.656*
T (+)	0.96*	0.656*	-0.066*	1.325*	0.325*	-0.357*
S (+)	0.585*	0.401*	-0.461*	0.939*	0.008	-1.01*
H (+)	0.057*	0.424*	0.482*	0.211*	0.163*	0.335*
N	0.075*	-0.003	-0.617*	0.283*	-0.318*	-0.869*
I (-)	-1.291*	-0.649*	-0.991*	-1.632*	-0.736*	-1.491*
T (-)	-1.014*	-0.305*	-0.897*	-1.343*	-0.67*	-2.068*
S (-)	-1.327*	-0.657*	-0.999*	-1.419*	-0.919*	-2.45*
H (-)	-1.836*	-1.101*	-1.288*	-2.007*	-1.158*	-2.39*

Table 14. Distributional measure ($L_g^{(a)}$) scores across demographic groups in images generated with Stable Diffusion 2.1, using CLIP embeddings. Negative values (in blue) indicate closeness to unattractive faces and positive values (in red) to attractive ones. An asterisk (*) indicates statistical significance ($p < 0.05$). M is used to represent images of males and F is used to represent images of females. Each row corresponds to the attribute used for face generation: I = Intelligence, T = Trustworthiness, S = Sociability, H = Happiness, N = Neutral and a +/- indicates the valence of the trait.

Tables 14 and 15 report the values of $L_g^{(a)}$ across demographic groups for the CLIP and ArcFace embeddings, respectively. We observe a consistent pattern that is supportive of the existence of algorithmic lookism: Faces generated with positive attributes (a^+) tend to be closer to the “attractive” reference cluster (indicated in red) whereas those created with negative attributes (a^-) tend to be closer to the “unattractive” reference cluster (indicated in blue). This pattern is robust across demographic groups with the exception of faces of while males and females, which exhibit a tendency to be grouped closer to the unattractive images. However, the findings provide clear evidence of algorithmic lookism in the images generated by Stable Diffusion 2.1. The strength of the effect also depends slightly on the embedding space, but the same overall trend is observed in both spaces.

Neutral images, i.e., images generated without any explicit attribute prompt, further highlight the aesthetic defaults of the generative model. These images tend to cluster closer to the unattractive reference, with limited variation across demographic groups. This finding suggests that Stable Diffusion encodes unattractive priors in the absence of explicit cues.

These results demonstrate that correlations between attractiveness and semantically unrelated traits are structurally embedded within the generative space of Stable Diffusion. The consistency of these effects across demographic groups and embedding spaces provides strong evidence that the model is influenced by algorithmic lookism.

In addition to the above evaluation, we estimate the human judgments of attractiveness of the synthetically generated images. This step is essential for assessing whether the observed effects extend to human perception, thereby clarifying the potential impact of synthetic images on human evaluators.

	M-Asian	M-Black	M-White	W-Asian	W-Black	W-White
I (+)	0.052*	0.152*	-0.172*	0.186*	0.088*	-0.102*
T (+)	0.193*	0.209*	-0.357*	0.374*	0.135*	-0.076*
S (+)	-0.003	0.154*	-1.01*	0.18*	0.039*	-0.151*
H (+)	0.288*	0.299*	0.177*	0.33*	0.133*	0.333*
N	-0.166*	-0.054*	-0.869*	0.032*	-0.039*	-0.091*
I (-)	-0.357*	-0.039*	-1.491*	-0.223*	-0.063*	-0.322*
T (-)	-0.254*	-0.002*	-2.068*	-0.033*	-0.023*	-0.249*
S (-)	-0.278*	0.011*	-2.45*	-0.127*	-0.034*	-0.247*
H (-)	-0.166*	-0.007*	-2.39*	-0.147*	-0.054*	-0.149*

Table 15. Distributional measure ($L_g^{(a)}$) scores across demographic groups in images generated with Stable Diffusion 2.1, using ArcFace embeddings. Negative values (in red) indicate closeness to unattractive faces and positive values (in blue) to attractive ones. An asterisk (*) indicates statistical significance ($p < 0.05$). M is used to represent images of males and F is used to represent images of females. Each row corresponds to the attribute used for face generation: I = Intelligence, T = Trustworthiness, S = Sociability, H = Happiness, N = Neutral and a +/- indicates the valance of the trait.

Algorithmic Lookism as per Human Ratings




Figure 13. Performance of the fine-tuned InceptionNet attractiveness classifier on the AHEAD dataset. Each panel presents the distribution of predicted attractiveness labels for images belonging to the low, medium, and high attractiveness categories respectively in the ground-truth annotations.

As described in Section 6.2.2, we trained an attractiveness classifier to approximate how humans would have rated the attractiveness of the synthetic face images. Figure 13 presents the performance of the classifier based on InceptionNet when evaluated on the AHEAD dataset. The panels show the distribution of predicted attractiveness ratings for images in the AHEAD dataset, labeled as low, medium, or high by human evaluators.

Overall, the classifier performs reasonably well. Face images labeled as low in attractiveness are predominantly assigned low ratings, with some classified as medium. Images labeled as high in attractiveness are largely assigned high ratings, with very few classified as medium. Finally, face images labeled as medium in attractiveness are mostly assigned medium ratings. These results suggest that the classifier adequately captures human-like judgments of attractiveness. A detailed quantitative evaluation of the attractiveness classifiers, comparisons with alternative models and additional validation on external datasets with attractiveness ratings (e.g., CelebA), are provided

in Appendix J. It is important to emphasize, however, that we do not endorse the development or use of systems that automatically rate the attractiveness of humans. Our objective here is strictly methodological: to approximate how humans might evaluate synthetic images in order to empirically test whether attractiveness biases embedded in generative models would also be reflected in human perception.

Next, we applied the attractiveness classifier to different groups of the synthetically generated faces according to specific attribute prompts. Figure 14 illustrates the distributions of attractiveness ratings for images generated with positive (intelligent, trustworthy, sociable and happy) versus negative (unintelligent, untrustworthy, unsociable and unhappy) attributes. As can be observed, the face images generated with positive attributes consistently receive higher attractiveness ratings than those generated with negative attributes. To statistically evaluate these differences, we conducted Mann–Whitney U tests comparing the pair-wise distributions. The results, summarized in Table 16, show that the differences are highly significant ($p < 0.001$). Moreover, the expected values confirm that images generated with positive attributes receive systematically higher attractiveness ratings than those generated with negative attributes.

These findings provide strong empirical evidence that humans would likely have rated the positively attributed images as more attractive as well. Consequently, the associations between attractiveness and unrelated attributes embedded in generative models are not only computational artifacts but also align with how humans would likely perceive these images. The broader implications of these findings are discussed in Section 6.4.

Attribute	Valance	Mean Attractiveness	Group Comparison
Intelligent	(-)	0.501	$U = 559,113.0, (p < 0.001)$
	(+)	0.777	
Trustworthy	(-)	0.697	$U = 621,813.0, (p < 0.001)$
	(+)	0.863	
Sociable	(-)	0.488	$U = 613,395.5, (p < 0.001)$
	(+)	0.677	
Happy	(-)	0.536	$U = 501,487.5, (p < 0.001)$
	(+)	0.981	

Table 16. Mean attractiveness ratings for images grouped according to the attribute used for their generation and the results of Mann–Whitney U-tests comparing the distributions of ratings for images generated with negative versus corresponding positive attributes.

6.3.2 RQ2: Are downstream models, specifically gender classifiers, sensitive to the attractiveness of synthetically generated faces?

In this section, we examine whether the algorithmic lookism found in the synthetically generated face images could affect downstream tasks. Specifically, we assess the performance of gender classification models and evaluate their accuracy across different subsets of the face images generated using Stable Diffusion 2.1. For this analysis, we consider three widely used classifiers: Insight-



Figure 14. Histograms of the attractiveness predictions for face images generated with different attributes. The plots on the left correspond to attributes with negative valence (unintelligent, untrustworthy, unsociable and unhappy), while the figures on the right correspond to attributes with positive valence (intelligent, trustworthy, sociable and happy). On average, faces generated with positive attributes receive higher predicted attractiveness scores from the model trained to approximate human attractiveness ratings.




Figure 15. Classification accuracies of InsightFace, DeepFace, and FairFace (from left to right) on images generated with Stable Diffusion 2.1. Within each panel, columns correspond to the attributes used to generate the images: A = Attractiveness, H = Happiness, I = Intelligence, S = Sociability, and T = Trustworthiness. Rows indicate the valence of the attribute (positive [+] or negative [-]). Results are further disaggregated by gender: the top two rows (outlined in green) show accuracies for male faces, while the bottom two rows (outlined in yellow) show accuracies for female faces.

Face, DeepFace, and FairFace. The classification accuracies across gender and attribute groups are reported in Figure 15.

The results reveal clear differences in performance depending on the model, the gender of the subject, and the attribute used in image generation. All models perform well on male faces, achieving near-perfect accuracies (close to 100% in some cases). In contrast, the accuracy on female faces drops substantially, particularly on images generated with negative attributes. The most striking case occurs with DeepFace, which achieves only 11% accuracy on images of women generated with the attribute “unhappy”, compared to 99.8% accuracy on images of men with the same attribute.

Notably, FairFace demonstrates comparatively robust performance. While its accuracy also decreases in the case of images of women generated with negative attributes, such as “unhappy” or “unintelligent”, it still achieves close to 91% accuracy, which is markedly higher than the corresponding results for InsightFace and DeepFace. This suggests that the limitation lies in the classifiers rather than the generated images themselves: the images clearly contain sufficient information for gender recognition, but certain models fail to utilize it effectively. Given that unattractiveness is disproportionately associated with images generated with negative attributes as per our previous findings, these results reflect how algorithmic lookism could propagate into and significantly impair downstream tasks such as gender classification. The broader implications of these findings are discussed in Section 6.4.

6.3.3 RQ3: Can the AHEAD dataset aid in the mitigation of these effects in downstream classifiers?

The AHEAD dataset was specifically designed to study the attractiveness halo effect in humans. In this section, we investigate whether fine-tuning a gender classifier with different subsets of this dataset can improve its performance.

This analysis serves two key purposes. First, if fine-tuning the gender classifier with a greater proportion of low-attractiveness samples leads to measurable improvements in classification accuracy, it would provide strong evidence that differences in attractiveness were driving the performance disparities observed in the gender classifier. Second, such fine-tuning offers a practical mitigation strategy for downstream models affected by algorithmic lookism. While it is unrealistic to elimi-




Figure 16. Gender classification accuracy of DeepFace when fine-tuned on varying female-to-male ratios with a subset of images from the (a) PRI (left) and the (b) POST (right) subsets of the AHEAD dataset. The evaluation is performed on the worst performing images of the synthetic face dataset, i.e., images generated with the attribute “unhappy” by Stable Diffusion 2.1. The x-axis represents the different ratios used and the y-axis represents the classification accuracy for images of females (●) and males (●). Note the significant increase in classification accuracy on images of females while maintaining or slightly decreasing the accuracy on images of males.

nate all synthetic images already in circulation, demonstrating that the bias can be reduced with a relatively small, targeted sample would highlight a simple and effective corrective measure.

Since classification performance was particularly poor on female faces, we initially fine-tuned the model using only images of females, using all available samples from either the PRI or POST set. However, this approach led to severe overfitting, with the classifier subsequently predicting “female” for all images regardless of input. To address this, we designed experiments that varied the ratio of male to female training images while deliberately keeping the set of female images constant across conditions. In this way, only the number of male faces increased, enabling us to examine the effect of training set composition without altering female sample representation. Interestingly, the results show that adding more images does not necessarily lead to improved performance.

Overall, we observe a substantial improvement in classification accuracy on synthetic faces generated with the attribute “Unhappy”, achieved with fewer than 924 fine-tuning images in either setting. The broader implications of these findings are discussed in Section 6.4.

Figure 17 presents the gender classification performance on faces generated with other attributes, with the model fine-tuned using a set composed of 90% female and 10% male images from each subset in the AHEAD dataset i.e., 508 images. The results reveal similar patterns as previously reported: classification accuracy on images of females improves substantially compared to the baseline performance without fine-tuning (see Figure 15), while performance on images of males decreases only marginally.

From these findings, we draw two conclusions. First, fine-tuning with either subset of the AHEAD dataset improves performance, suggesting that the lack of balanced gender representation in the original training data was a major factor in the observed disparities. This interpretation is consistent with the substantially higher performance of FairFace compared to DeepFace as per the results reported in Figure 15. Second, attractiveness appears to play an additional role, particularly on images of females. Fine-tuning with 508 images from the PRI subset yields a significantly larger increase in classification accuracy than fine-tuning with the POST subset, indicating that differences in attractiveness are also an important driver of performance.

It is also important to note that the images in the AHEAD dataset are visually distinct from those

generated by Stable Diffusion, on which the classifier was evaluated. The observed improvements therefore cannot be attributed simply to greater exposure to similar images. Instead, the gains are primarily driven by exposure to a more balanced set of female faces across different attractiveness levels.

The implications of these findings are discussed in greater detail in Section 6.4.

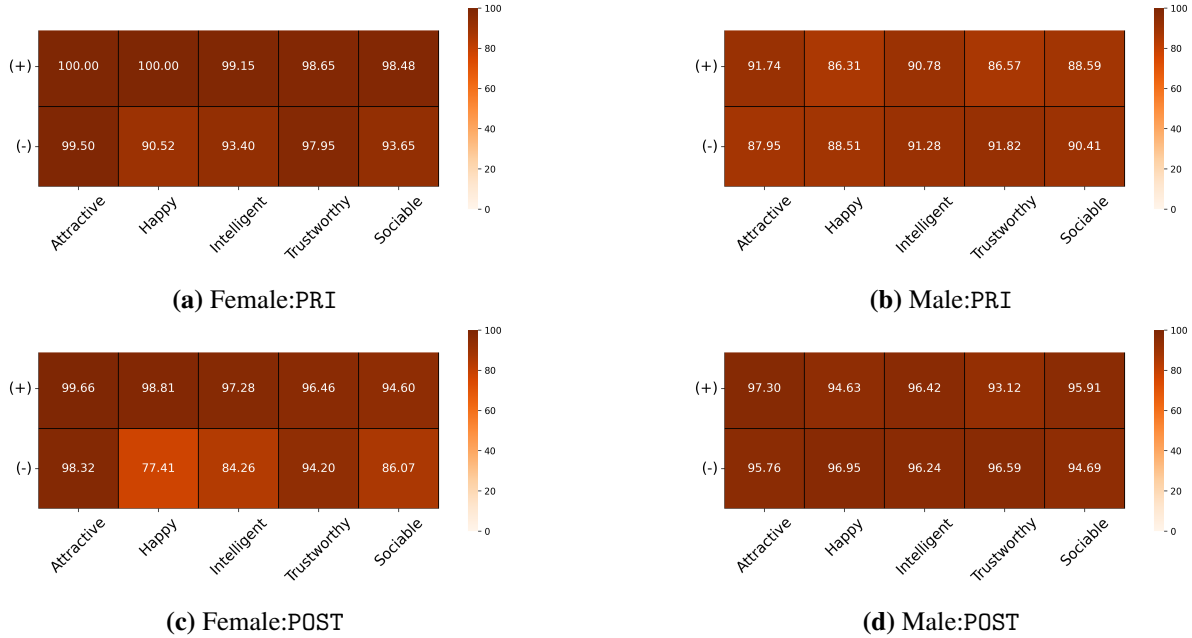


Figure 17. Classification accuracy of DeepFace on the subsets of images generated with positive and negative attributes. The top set of tables depict the heatmaps of accuracies for the fine-tuned gender classifier on 508 images from the PRI subset of the AHEAD dataset (90% female, 10% male), evaluated on synthetically generated images of females (a) and males (b). The bottom set of tables show the corresponding results for the fine-tuned gender classifier on the POST subset of the AHEAD dataset with the same number of images and gender ratio, evaluated on synthetically generated images of females (c) and males (d).

6.4 Discussion

Algorithmic lookism shapes the generation of synthetic images Our findings provide clear empirical evidence of the existence of algorithmic lookism in synthetically generated images. When image generators are prompted with positive attribute terms, such as intelligence or trustworthiness, the resulting faces are systematically more attractive –both in the embedding space and according to human perception– than those generated when the model is prompted with negative attribute terms, such as unsociable or unhappy.

Our analyses provide an initial step in the study of *algorithmic lookism* by demonstrating that the attractiveness of the generated faces depends on the trait prompts given to generative models, even when attractiveness is not part of the prompt. While our study was conducted in a highly controlled setting, such effects are likely to generalize across domains, image types and text-to-image models.

From an algorithmic perspective, an important limitation of the distributional lookism measure used in this work is that it quantifies correlations between attributes by leveraging the internalized notion of attractiveness in Stable Diffusion. While this operationalization is model-dependent,

there is reason to believe such associations would also extend to human perception. Prior work has shown that people judge synthetically generated faces as more “real” than actual human faces [NF22], suggesting that attractiveness cues embedded in generated images may readily transfer to human evaluations. To further explore this possibility, we developed predictive models of human attractiveness for the synthetically generated faces, whose implications are discussed next.

From a human perspective, using the *InceptionNet* classifier fine-tuned on the AHEAD dataset, we find that faces generated with positive attribute prompts would likely be judged by human evaluators as more attractive than those generated with negative attribute prompts.

The broader implications of algorithmic lookism in the generated face images are non-trivial. On-line environments already amplify pressures related to body image and appearance [Bak22; FM14; Esh20; Isa23; LC20; VFL+20; Gil21; Rya22]. If generative systems embed and propagate attractiveness biases when prompted with seemingly neutral descriptive prompts, they may subtly shape cultural norms and personal self-perceptions in ways that are difficult to monitor or control. As synthetically generated images of human faces become more prevalent in social media, advertising, and digital communication, these biases risk subtly shifting collective norms of attractiveness. Crucially, the mechanisms by which such associations are embedded remain opaque. Users are typically unaware that trait prompts, intended to capture descriptors such as “friendly” or “intelligent”, may implicitly invoke attractiveness as a latent correlate. The invisibility of these associations makes it difficult to anticipate or regulate their long-term social and psychological consequences.

Perhaps most concerning is the broader uncertainty surrounding generative models. The implicit links to attractiveness uncovered in this thesis reflect only one dimension of their learned associations. Given the scale and inscrutability of model training data and the complexity of the architectures of state-of-the-art models, with billions of parameters, it remains unclear what other biases or stereotypes are being systematically encoded and propagated. Our findings underscore the importance of grounding empirical evaluations of generative models in psychological theory: by aligning the observed biases with well-established cognitive phenomena, such as lookism and the attractiveness halo effect, we gain a clearer lens through which to anticipate the kinds of associations that may emerge.

Algorithmic lookism impacts downstream tasks The consequences of algorithmic lookism propagate into other AI systems that analyze face images, such as gender classifiers. In an empirical evaluation of three gender classification models, we find a significant drop in performance in most models when classifying faces –and particularly female faces– generated with negative attribute prompts. Importantly, this deterioration of performance is not due to an absence of gender-identifying information. Models such as FairFace maintained competitive accuracy on the same images, suggesting instead that some classifiers are especially sensitive to images generated with negative trait terms. They could potentially be impacted by the lower levels of attractiveness of such images, as per our analyses of algorithmic lookism.

Our fine-tuning experiments with DeepFace on the AHEAD dataset corroborate this hypothesis. Gender classification performance improved the most when fine-tuning the model on a subset of the original unfiltered images (the PRI subset) instead of a subset of the beautified images. This finding indicates that attractiveness itself plays a role in classifier robustness. These results highlight that attractiveness biases can directly degrade the reliability of models trained or evaluated on synthetic data.

The implications are far-reaching. While the impact of these biases is most visible when models are evaluated on synthetic faces, as large-scale generative data are increasingly scraped from the

internet to train new models, such biases are likely to intensify. Faces are a critical locus of identity and social interaction, and models that implicitly penalize certain appearances may contribute to systemic harms in downstream applications. Beyond technical effects, there is a broader cultural concern: if algorithmic systems consistently reward certain kinds of faces, individuals may adapt their self-presentation to align with these algorithmically reinforced norms. In consequence, the attractiveness halo effect, once only a cognitive bias in human judgment, risks becoming codified as a feedback loop between generative models, downstream classifiers, and human behavior.

The value of the AHEAD dataset Encouragingly, our experiments indicate that fine-tuning with the AHEAD dataset can be an effective measure to mitigate algorithmic lookism, despite its relatively modest size (462 faces of different individuals compared to the 13,200 synthetically generated images). This finding emphasizes the importance of high-quality, diverse and curated datasets to train or fine-tune models. While gender imbalance in the training data clearly contributed to improving classifier performance, the attractiveness of the images used for fine-tuning also played a meaningful role. The images used for fine-tuning were entirely distinct from those used for evaluation, suggesting that the AHEAD dataset improved the model’s ability to generalize to more robust face representations rather than merely overfitting to specific examples.

These results carry two key implications. First, they demonstrate that relatively small, carefully curated datasets can be used to effectively mitigate systematic biases in pre-trained models. Second, while interventions at the level of generative models remain urgently needed, they are challenging to implement in practice given the widespread and uncontrolled dissemination of such models. In this context, the AHEAD dataset provides a pragmatic and immediately usable resource for improving fairness in downstream applications.

While not a comprehensive solution, fine-tuning with the AHEAD dataset can serve as an effective step toward mitigating algorithmic lookism, offering a practical means of counteracting some of the biases that are otherwise reinforced and amplified through generative models.

Limitations Several limitations of this study need to be acknowledged. First, we employed binary gender categories and a ternary racial classification. While both gender and race are complex constructs, this stratification was a pragmatic design choice intended to make the analysis tractable and aligned with the human experiments described in Chapter 3 and the AHEAD dataset. Despite this simplification, our findings nonetheless provide valuable insights into the intersection of aesthetic biases and algorithmic decision-making.

Second, the study relies on a predominantly Western conceptualization of attractiveness, a limitation inherent in the datasets used. Attractiveness norms are culturally diverse and context-dependent, and our operationalization represents only one perspective. Future work should expand beyond this narrow frame to incorporate a broader range of cultural conceptions of beauty, but does not necessarily take away from the utility of our findings here.

Finally, our predictive estimator of attractiveness was not perfect. We do not recommend such a model for general use, given the ethical and methodological concerns associated with automating judgments of human attractiveness. However, for the purposes of this research i.e., estimating human-like responses at scale, the estimator was sufficiently robust to provide the necessary insights.

6.5 Conclusion

The experimental study described in this chapter is among the first to systematically examine the existence of algorithmic lookism in synthetically generated faces. Our findings reveal that, similar to MLLMs, text-to-image models encode an association between attractiveness and positive characteristics. We further demonstrate that this bias extends to downstream tasks, negatively impacting the performance of gender classifiers. Importantly, we provide empirical evidence that fine-tuning pre-trained models with carefully curated, high-quality, diverse and small datasets, such as the AHEAD dataset, can significantly improve the classifier’s performance.

Motivated by knowledge initially brought forth through studies in human psychology and behavioral economics, this work illustrates how a well-documented cognitive bias not only reemerges but may also be amplified within machine learning models. In doing so, it highlights the urgent need for rigorous evaluation of both generative and discriminative models in order to better identify, understand, and mitigate the subtle yet consequential social biases embedded in AI systems.

Chapter 7

Conclusion

This thesis begins from the premise that cognitive biases are an inherent part of human decision-making and highlights the opportunity of studying them in the context of AI in order to understand how human and machine decision processes impact each other. To address this, in Chapter 2 we first proposed a framework that categorizes known cognitive biases from the perspective of human-AI systems, organizing them according to where they emerge in the human decision-making cycle.

Within this broader framework, the thesis focuses on a prominent yet understudied cognitive bias: the attractiveness halo effect (AHE), and provide an in-depth account of how perceptions of attractiveness can bias both human and machine judgments. Faces play a central role in human-to-human interaction, and the attractiveness halo effect was established in psychology as early as the 1970s [DBW72]. Yet, relatively little work has examined how this bias manifests in contemporary, digitally mediated environments.

Beauty filters exemplify such environments: they represent a relatively recent technological intervention that leverages a variety of AI methods to alter appearance in real-time, thereby potentially reshaping the ways in which the attractiveness halo effect operates. To investigate this, in Chapter 3 we report the results of one of the first, and largest, user studies to investigate the existence of the AHE in humans and the impact of beauty filters on perceptions of attractiveness and on this cognitive bias. Beyond collecting robust empirical evidence of the existence of this cognitive bias on a diverse set of stimuli, our study yielded the AHEAD dataset, a curated dataset to study potential attractiveness biases in humans and algorithms.

Armed with empirical evidence of how attractiveness shapes human decisions, the thesis then turned to machine learning models. In Chapter 4, we introduced the concept of algorithmic lookism, i.e. the tendency for algorithms to exhibit attractiveness-based discrimination.

We examined this phenomenon in two contexts: first, when AI systems are tasked with making decisions based on images of individuals (Chapter 5), and second, when AI systems generate images of people (Chapter 6). In our experiments with multimodal large language models (MLLMs), we find strong empirical evidence that they associate attractiveness with positive traits. This finding was consistently seen across seven open-source models of diverse architectures. Although these experiments were conducted in a specific task setting, they demonstrate that algorithmic lookism is embedded within MLLMs. Given the broad and increasing deployment of these models across application domains, the presence of attractiveness-related biases raises important concerns about their potential impact on real-world decision-making. While considerable research has examined biases related to gender, age, and race, comparatively little attention has been paid to how attractiveness influences AI-driven judgments.

Moreover, our findings reveal that algorithmic lookism is also present in text-to-image generative models. We create a dataset of 13,200 different synthetic human faces with Stable Diffusion 2.1. The faces are created with different positively (intelligent, trustworthy, sociable and happy) and negatively (unintelligent, untrustworthy, unsociable and unhappy) valanced attributes to study the existence of algorithmic lookism. We find clear empirical evidence that Stable Diffusion 2.1 associates attractiveness with the faces created with positive attributes and vice versa. Furthermore, we build an attractiveness classifier from human data and find that, according to the classifier, faces generated with positive attributes are more attractive than those generated with negative attributes. In addition, we study to which degree algorithmic lookism impacts downstream applications, such as gender classification models. We find a significant reduction in gender classification accuracy when evaluating images generated with negative trait prompts.

These results demonstrate that attractiveness shapes human perception in digitally mediated environments and that these effects have percolated into AI systems when they analyze human faces, both to make decisions about them and when generating them. Moreover, algorithmic lookism impacts downstream applications, specifically gender classification models, which perform significantly worse when analyzing face images created with negative attributes, and particularly images of females. This cyclic reinforcement of bias suggests a feedback loop in which attractiveness-related stereotypes are not only reproduced but also amplified by AI systems.

Another important finding in this thesis concerns the interaction between attractiveness bias and gender. Across all our studies, we consistently observe that women are disproportionately affected by this bias. In human evaluations, the attractiveness halo effect had a more negative impact on images of women, and in the case of attractiveness and intelligence, male raters appeared to be more susceptible to the influence of beauty filters than female raters. In MLLMs, algorithmic lookism was more strongly expressed for female images, and in our Stable Diffusion experiments, the worst performing group in the gender classifiers were images of women created with negative attributes (and particularly images of “unhappy women”). These converging results suggest that attractiveness bias is deeply intertwined with gender, and women are more strongly judged on the basis of their appearance. Thus, an attractiveness bias not only coexists with but may also exacerbate existing gender biases in both society and AI systems.

Overall, the research described in this thesis has uncovered a critical but underexplored form of bias—one that has received far less attention compared to gender, race, or age bias—yet has significant implications for both human and machine decision-making. We show that attractiveness bias shapes perceptions, influences AI systems directly, and propagates into downstream tasks, thereby creating systemic vulnerabilities in socio-technical ecosystems. Importantly, this thesis demonstrates that by foregrounding the role of cognitive biases in AI design, we can better understand not only the weaknesses of intelligent systems but also the strategies necessary to make them more equitable and human-centered. In doing so, this work contributes to the broader project of aligning AI with the values, complexities, and fairness expectations of the societies it is intended to serve. As the title of the thesis suggests, we should not judge books—and even less so people—by their covers.

Appendices

Appendix A

Influence of Participant Characteristics on Ratings

In addition to the analysis presented in Chapter 3, we further examined how participant characteristics influenced the ratings they provided. Specifically, we investigated how the gender of the individuals being rated affected participants' evaluations (Section A.1) and whether participants' self-perceived attractiveness shaped the attractiveness ratings they assigned (Section A.2). The results of these analyses are reported in this Appendix.

A.1 Impact of Rater Gender

Figure 6 shows the Expected Marginal Means (EMM's) of the ratings for images of males and females by male and female raters. For most attributes, both male and female raters provide different ratings to images of males and females. We refer to this difference in rating between images of males and females as a *gender gap*. This gender gap in ratings is impacted by the filters and depends on the gender of the rater.

Table 17 shows the differences (in percentage) of the gender gap between the ratings provided to the images in the PRI and the POST datasets by male and female raters for attractiveness and all the dependent attributes. Regarding perceptions of attractiveness, intelligence and trustworthiness, observe how male raters are more impacted by the filters as there is a larger gender gap when compared to female raters. However, for perceptions of sociability and happiness, female raters are more impacted by the filters than male raters.

A.2 Impact of Self-Perceived Attractiveness on Judgments of Attractiveness

An additional factor that might mediate the strength of the attractiveness halo effect is the perceived attractiveness of the human evaluators. Humans are unable to make a self-target comparisons without assessing their own physical attractiveness [MB02]. Previous work has found that the self concept of physical attractiveness is also associated with positive affect, cognitive, and social measures [Fei92]. Not only our behaviour is shaped by the levels of attractiveness that we perceive ourselves with, but we use such a self-assessment as a benchmark when determining the attractiveness of others. According to the in-group/out-group theory [AP08], in-group members share similar attributes

Dependent Attribute (ω)	Male Raters	Female Raters
Attractiveness	85.12	17.83
Intelligence	331.23	158.81
Trustworthiness	-95.68	-37.74
Sociability	49.60	1598.85
Happiness	114.00	212.69

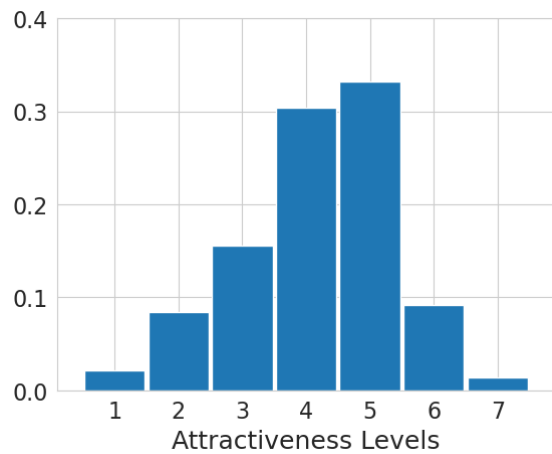
Table 17. Differences (expressed as percentage of change) in the gender gap between the ratings provided to images of male and female subjects, by male and female raters, between the POST and the PRI datasets. There is a larger increase in the gender gap for male raters when rating attractiveness, intelligence and trustworthiness. However, there is a larger increase of the gender gap in the ratings provided by female raters when rating sociability and happiness. Statistical significance of the differences in the PRI and POST datasets are described in Figure 6.

and assign more positive attributes to each other than to out-group individuals [MSC93]. Consequently, the self-perceived attractiveness of the human evaluators would impact their perception of the stimuli, leading to different perceptions of social distance between the self and the target. Recent work by Li et al. [LZL19] has studied and corroborated this phenomenon in the context of consumer evaluation processes during a service encounter: the evaluators’ (consumers in this case) perception of their physical attractiveness was found to play a moderating role on the attractiveness halo effect. However, other authors have not identified any interactions between self-perceived attractiveness and the halo effect in certain contexts, such as hireability [Cot11].


In this section, we study the relationship between the participants’ self-perceived attractiveness and their attractiveness judgements. Figure 18 depicts the histograms of **a)** the self-reported attractiveness levels of the participants in our study (Mean: 4.17, Std: 1.21), and **b)** the centralised attractiveness scores reported by the same raters to images from the PRI set (Mean: 3.57, Std: 0.91).

Given the self-perceived attractiveness of rater R , R_{SRA} ; the attractiveness score provided by rater R to stimulus I , I^R ; and the stimulus centralised attractiveness score, i.e., its median attractiveness, I^c , we identify a weak correlation ($\tau = 0.049$, $p = 0.007$, Kendall) between the participant’s self-perceived attractiveness and their attractiveness ratings overall, and both in the images belonging to the PRI ($\tau = 0.019$, $p = 0.008$, Kendall) and POST ($\tau = 0.053$, $p < 0.001$, Kendall) sets. Thus, perceived attractiveness has not been included as a factor in the reported analyses.

While there is no strong relationship between the rater’s self perceived attractiveness and the attractiveness ratings they provide, an interesting phenomena does emerge in Figure 18: more than 50% of the participants consider themselves to be above average on attractiveness (above a 4 on the 7-point Likert scale). We observe a clear distribution shift between the distribution of attractiveness scores provided to the stimuli (Figure 18b) and the distribution of self-rated attractiveness (Figure 18a). This finding is in line with literature that finds that humans tend to overestimate their qualities and abilities and underestimate those in others [JR94; KKLC16].



(a) Distribution of the self-rated attractiveness of participants in our study



(b) Distribution of the attractiveness ratings provided to images in the PRI set

Figure 18. Distribution of attractiveness ratings of (a) the self rated attractiveness of the participants and (b) the attractiveness ratings provided to images in the PRI set

Appendix B

Statistical Analysis Conducted during the Creation of the AHEAD Dataset

This appendix presents additional statistical tests conducted as part of the analysis in Chapter 3 and offers further details on the tests reported in the chapter.

B.1 Ordered Stereotype Models

Figure 19 shows the new scales computed using the OSM for attractiveness and each of the dependent attributes independently on the PRI and the POST datasets. Since the scales were computed independently on the PRI and POST datasets, each attribute has a different scale with possibly a different number of points in each set, making it harder to directly compare the ratings an image receives before and after the filters are applied.

The scale for *attractiveness* is compressed to 6 points in the PRI dataset and 5 points in the POST dataset. The highest/lowest values of the scale are merged in the PRI and POST datasets respectively because there are few ratings corresponding to such values. The scale for *intelligence* is compressed to a 5-point scale in the PRI dataset and a 6-point scale in the POST dataset. In both cases, the highest values of the scale are merged. The scale for *trustworthiness* is compressed to a 5-point scale both in the PRI and POST datasets. The highest values of the scale are merged in the PRI dataset and the levels corresponding to the [5,6] values in the original scale are merged in the POST dataset. The scale for *sociability* exhibits a similar behavior to that of *attractiveness*, with a compression to a 6-point instead of a 5-point scale in the POST dataset. Interestingly, *happiness* is the only dependent attribute for which the scale remains as a 7-point scale both in the PRI and POST datasets, with an adjustment of the distance between consecutive points in the scale.

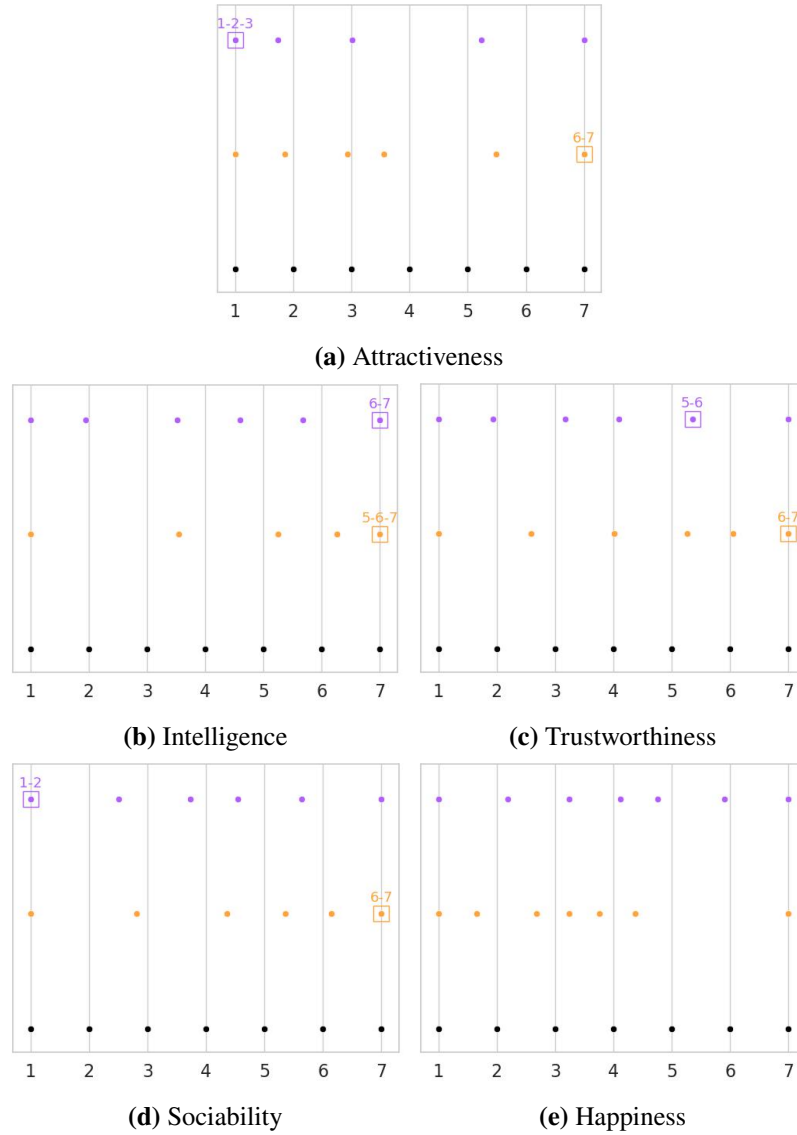


Figure 19. The new scales for perceived attractiveness and the dependent variables after re-scaling by means of OSMs based on the collected data. The black dots at the bottom indicate the original, equally spaced 7-point Likert scale. The orange and purple dots correspond to the new scale for the PRI (●) and POST (●) sets, respectively. The squares around the dots indicate locations where multiple points on the scale were collapsed to the same value.

B.2 Model Selection

Section 3.2.3 describes the impact of the participants' age and gender on the attractiveness halo effect by means of Ordered Stereotype Models (OSM) [FAP16] in conjunction with linear mixed models. In this section, we present the goodness of fit analyses that justified such model selection.

We evaluated 10 different models, depicted in Figure 20, according to a taxonomy with three levels. The first level corresponds to the type of model (ordinal, such as the Cumulative Link Model, or linear); the second level reflects whether the data is in the original 7-point Likert scale or in the new scales obtained by the OSMs. Furthermore, in the case of linear models, a third option is considered where the number of points in the scale is given by the OSMs yet the points

are equidistant; the third level describes whether the raters were considered as fixed effects (dashed line) or random effects (solid line).

The 10 models were evaluated using the AIC [Aka74] and BIC [Sch78] on the data from the PRI and POST sets separately as reflected in Tables 18, 19, 20, 21. Note that the AIC and BIC are sensitive to sample size, which in our case is $N = 27,480$ data points. Thus, the values presented in the tables are divided by a factor of 10^3 . The best fitting model corresponds to the lowest AIC/BIC values, which are marked in bold on the tables.

Based on these results, we opted for $M5_{RE}$, i.e., a linear mixed model on the OSM-based rescaled data with the raters as random effects. In addition to being the best performing model in most cases, linear mixed models are significantly more interpretable than ordinal models [Man16].




Figure 20. Three-level taxonomy of modeling choices that were evaluated for their goodness of fit, resulting in 10 different models.

	M1 _{FE}	M1 _{RE}	M2 _{FE}	M2 _{RE}	M3 _{FE}	M3 _{RE}	M4 _{FE}	M4 _{RE}	M5 _{FE}	M5 _{RE}
Attractiveness	45.039	43.835	44.455	43.255	45.799	44.176	45.497	43.861	45.181	43.835
Intelligence	36.035	35.184	29.349	28.482	36.213	35.260	31.771	31.064	22.618	21.937
Trustworthiness	37.225	36.004	36.601	35.373	37.666	36.110	37.268	35.728	35.268	33.898
Sociability	39.107	38.547	38.555	37.990	39.368	38.637	39.040	38.311	36.040	35.279
Happiness	40.272	39.460	40.283	39.470	40.623	39.586	40.637	39.598	31.273	30.187

Table 18. AIC/ 10^3 on the PRI set for all variables and model variations

B.3 Factor Analysis

Principal Component Analysis (PCA) of the centralized ratings in the PRI and POST datasets separately was performed for each dependent variable to identify correlations between them. Figure 21 depicts the projections of the data for each dependent variable on a 2-dimensional space of the directions of the eigenvectors with the largest eigenvalues.

	M1 _{FE}	M1 _{RE}	M2 _{FE}	M2 _{RE}	M3 _{FE}	M3 _{RE}	M4 _{FE}	M4 _{RE}	M5 _{FE}	M5 _{RE}
Attractiveness	42.537	41.757	38.774	38.057	44.223	43.157	40.652	39.689	41.731	40.696
Intelligence	34.622	33.270	33.544	32.234	34.825	33.204	34.174	32.661	33.642	32.156
Trustworthiness	36.751	35.421	30.366	29.134	36.966	35.317	32.435	31.238	31.366	30.004
Sociability	37.806	37.030	37.522	36.758	38.131	37.300	37.835	37.025	34.715	34.092
Happiness	39.198	38.610	39.275	38.683	39.433	38.710	39.533	38.802	36.100	35.328

Table 19. AIC/ 10^3 on the POST set for all variables and model variations

	M1 _{FE}	M1 _{RE}	M2 _{FE}	M2 _{RE}	M3 _{FE}	M3 _{RE}	M4 _{FE}	M4 _{RE}	M5 _{FE}	M5 _{RE}
Attractiveness	45.129	43.933	44.538	43.345	45.859	44.244	45.558	43.929	45.241	43.902
Intelligence	36.170	35.327	29.462	28.603	36.280	35.336	31.839	31.139	22.686	22.012
Trustworthiness	37.361	36.147	36.721	35.501	37.734	36.186	37.336	35.803	35.336	33.973
Sociability	39.242	38.690	38.675	38.118	39.436	38.712	39.108	38.386	36.108	35.354
Happiness	40.408	39.603	40.411	39.606	40.691	39.661	40.704	39.673	31.341	30.262

Table 20. BIC/ 10^3 on the PRI set for all variables and model variations

While sociability and happiness appear to be closely related in the PRI dataset, all 4 dependent attribute vectors are clearly separated in the POST dataset, with intelligence and sociability being almost orthogonal to each other. Based on these results, we perform analyses on these 4 dependent attributes.




Figure 21. Projections of all the dependent attributes on the first two dimensions

B.4 Evaluation of the Saturation Effect in the Halo Effect

As discussed in Section 3.2.4, we observe a saturation in the relationship between attractiveness and some of the dependent variables, namely intelligence and trustworthiness (see Figure 7). As an initial test of this hypothesis, Table 22 summarizes the results of Wilcoxon paired-rank tests on the centralized scores of the images with perceived attractiveness scores greater or equal than 5 before beautification (highly attractive stimuli) compared with the remaining stimuli and the complete dataset. Differences in pairwise perceived intelligence and trustworthiness before and after

	M1 _{FE}	M1 _{RE}	M2 _{FE}	M2 _{RE}	M3 _{FE}	M3 _{RE}	M4 _{FE}	M4 _{RE}	M5 _{FE}	M5 _{RE}
Attractiveness	42.628	41.855	38.849	38.140	44.283	43.224	40.712	39.756	41.791	40.764
Intelligence	34.758	33.413	33.657	32.354	34.893	33.279	34.242	32.736	33.709	32.232
Trustworthiness	36.887	35.564	30.479	29.255	37.034	35.392	32.503	31.314	31.434	30.079
Sociability	37.941	37.173	37.635	36.878	38.199	37.375	37.902	37.100	34.783	34.168
Happiness	39.333	38.753	39.396	38.811	39.501	38.785	39.601	38.877	36.167	35.403

Table 21. BIC/ 10^3 on the POST set for all variables and model variations

Dependent Attribute (ω)	Complete Image Set (n=462)	Highly Attractive Stimuli (n=79)	Remaining Stimuli (n=383)
Attractiveness	213.83***	32.35***	182.12***
Intelligence	63.15***	3.22	61.81***
Trustworthiness	33.13***	6.71	26.31***
Sociability	123.49***	10.86***	115.04***
Happiness	118.57***	14.28***	105.28***

Table 22. Normalized Wilcoxon paired-rank test statistic (W/n) for attractiveness and the 4 dependent attributes. An image is considered highly attractive if its centralized attractiveness score is ≥ 5 in the PRI dataset.

beautification are not statistically significant for the “highly attractive stimuli” whereas they are significant ($p < 0.001$) in the rest of the cases. This finding supports the saturation hypothesis for intelligence and trustworthiness.

We further quantify the strength of the effect by means of two approaches:

B.4.1 Method A: Piece-wise linear fit

For each dependent variable, the data is divided in two halves according to the corresponding attractiveness ratings. A linear function ($\omega = m \cdot Attrac + c$) is fitted to each set and the slopes (m) of the linear functions are compared to quantify the saturation effect as attractiveness increases:

$$Sat_{\omega}^A = \frac{m_{Upper} - m_{Lower}}{m_{Lower}} \times 100$$

Where m_{Lower} and m_{Upper} represent the slopes of the lines fit on the lower and upper part of the data respectively. We report these findings in the first column of Table 23. The strongest saturation effect is observed for intelligence followed by trustworthiness.

B.4.2 Method B: Fitting a Log Curve

A second approach to measure the strength of the saturation effect consists of fitting a log curve of the form $\omega = a \cdot \log(attrac) + b$ and comparing the goodness of fit (by means of the AIC [Aka74]) with a line of the form $\omega = a \cdot attrac + b$. Figure 22 depicts the log curves fit on the data in blue. To evaluate the strength of the saturation effect, we measure the percentage change in the AIC of both the log and the linear curves: $Sat_{\omega}^B = \frac{AIC_{Linear} - AIC_{Log}}{AIC_{Log}} \times 100$

Where AIC_{Linear} and AIC_{Log} represent the AIC's of the linear fit and log curve respectively. Since a lower AIC indicates a better fit, the larger the value of Sat_{ω}^B , the stronger the saturation

Dependent Attribute (ω)	Method A (B.4.1)		Method B (B.4.2)	
	PRI	POST	PRI	POST
Intelligence	-77.2	-54.98	2.3	0.37
Trustworthiness	-67.50	-68.88	1.08	0.68
Sociability	-60.22	-45.26	0.83	0.35
Happiness	-63.83	-49.35	0.91	0.43

Table 23. Evaluation of the strength of the saturation effect for different dependent variables using the methods described in Appendix B.4. Note how the effect (difference between the values between the PRI and POST datasets) is strongest for intelligence followed by trustworthiness.

effect. Again, the strongest saturation is observed for intelligence followed by trustworthiness.

B.5 Linear Mixed Models including Rater Effects

The parameters for all the linear mixed models described in Section 3.4.5 can be found in the anonymized GitHub repository containing the code used to analyze the data. The model parameters can also be directly accessed at this link: <https://tinyurl.com/modelParametersFile-ea>. The parameters of the fixed effects terms represent the slopes corresponding to each of the terms. Note that females (for gender of the image and the gender of the rater) are coded as 0. Thus, the beta values correspond to the slope for images of males and male raters. The impact of the gender of the image and rater and their interactions have been discussed in Section 3.2.3 and Appendix A.1 by computing the estimated marginal means instead of relying only on the β 's presented here.

B.6 Partial R^2 in the Linear Mixed Models

In order to evaluate the importance of attractiveness in predicting the dependent variables when compared to other predictors, we compute the partial R^2 's [SNS21] of each of the predictors for all the linear mixed models. The results are summarized in Table 24. Observe how attractiveness explains the largest part of the variance of all of the models.

In Appendix C, we noted that the filters reduce perceptions of age of the stimuli. Given however that the partial R^2 associated with age is much lower than the partial R^2 for attractiveness, we conclude that it is the change in attractiveness driving the changed perceptions of the dependent variables and not the change in the perceived age of the subjects in the images.

B.7 Computation of Fractional Change in EMM

This section describes in detail how the y-axis values of the plots in Figure 6 were computed. The values are directly proportional to the Estimated Marginal Means (EMM). However, the Ordered Stereotype Models (OSM) provide different scales for each attribute in the PRI and POST datasets, which makes it hard to directly compare values computed on the PRI and POST scales.

However, to understand the impact of the gender of the rater, it is sufficient to compare relative changes between image gender-rater pairs. Thus, setting the value of images of females rated by

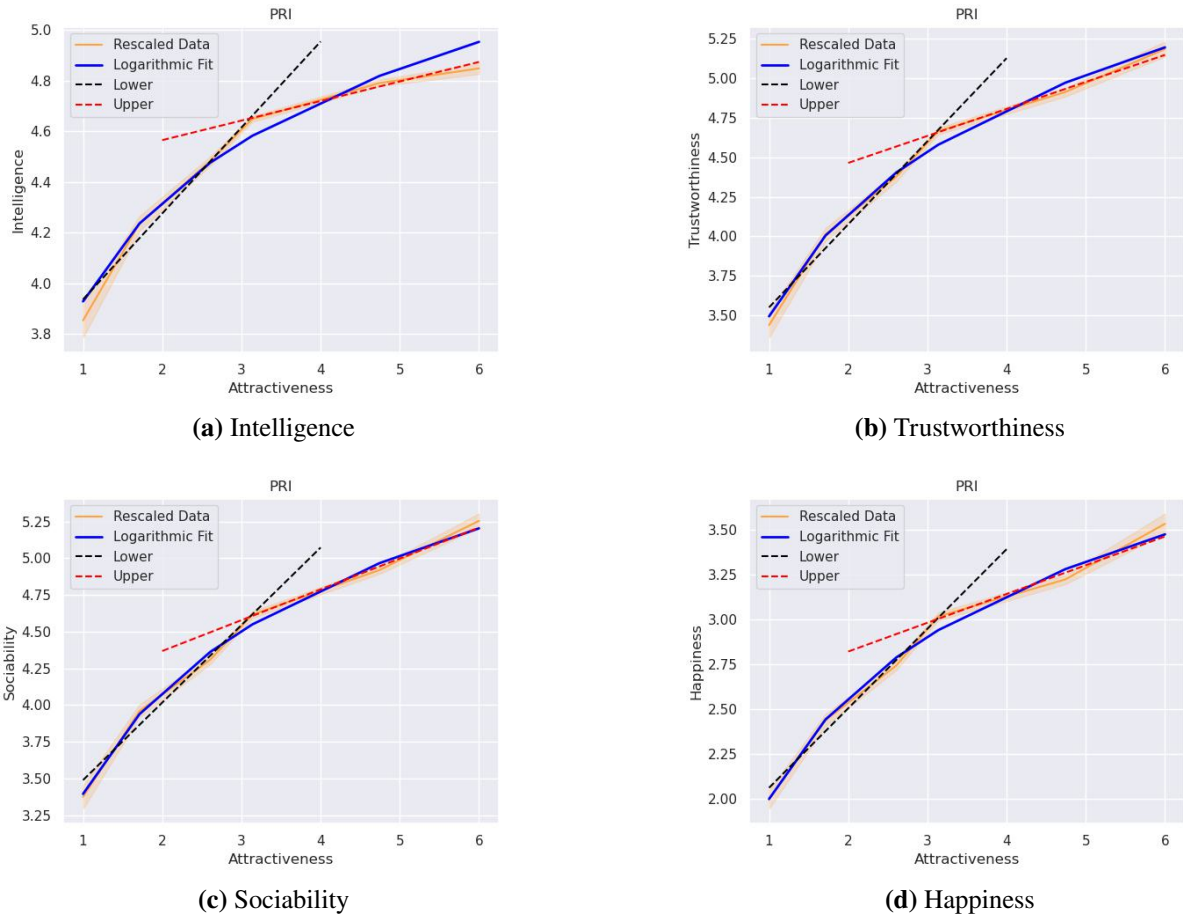


Figure 22. Relationship between attractiveness and the dependent attributes after re-scaling the data on the PRI set. The yellow curve represents the re-scaled data, the blue curve represents a logarithmic curve fit to the data and the dashed lines represent the best fit lines on the lower (black) and upper (red) half of the data. While all attributes show saturation to a degree, it is strongest for intelligence and trustworthiness. Note that the y-axis values for the different dependent attributes are not directly comparable since they were all re-scaled independently using the OSM's.

Dependent Attribute (ω)		Full Model	Attractiveness	Image Gender	Image Age	Participant Gender	Participant Age
Attractiveness	PRI	0.145	X	e^{-12}	0.015	0	0.001
	POST	0.195	X	e^{-12}	0.014	e^{-12}	0
Intelligence	PRI	0.150	0.138	e^{-12}	0.001	e^{-9}	0.001
	POST	0.098	0.083	e^{-12}	0.007	e^{-11}	0.007
Trustworthiness	PRI	0.169	0.141	0	0.001	0	e^{-4}
	POST	0.087	0.065	0	0.002	0	0.003
Sociability	PRI	0.192	0.157	e^{-12}	e^{-5}	e^{-12}	0.001
	POST	0.167	0.109	0	e^{-4}	0	0
Happiness	PRI	0.177	0.152	0	e^{-5}	e^{-9}	0.002
	POST	0.141	0.100	e^{-12}	e^{-4}	e^{-9}	0.002

Table 24. Partial R^2 of each of the predictors in the linear mixed models described in Appendix B.5. Note how attractiveness has the largest R^2 of all the variables, indicating that attractiveness best explains the variance of the dependent variables.

females as 0 enables a comparison of the relative changes between (rater, image) gender pairs. Thus, the fractional change values depicted on the y-axis of the graphs in Figure 6 for each dependent attribute ω in the PRI and POST datasets were computed as:

$$fractionalChange = \frac{EMM_{(i,j)} - EMM_{(f,f)}}{numLevels} \quad (11)$$

where $EMM_{(f,f)}$ represents the estimated marginal mean value of images of females rated by female raters for the dependent attribute ω in the PRI or POST datasets. $EMM_{(i,j)}$ represents the corresponding EMM for every image gender-rater gender pair (i, j) in the same setting and $numLevels$ represents the number of levels on the re-scaled version of the dependent attribute.

This computation of the fractional changes enables a visualization of 1) how different raters are impacted by the gender of the stimulus (i.e, differences between the blue and pink bars for each rater gender); 2) how the gender of the stimuli impacts the perceptions provided by the raters (i.e, how different are the two pink (or blue) bars between male and female raters; and 3) how different are these changes between the PRI and POST datasets for each dependent attribute, reflecting the impact of the beauty filters.

Appendix C

Impact of Beauty Filters on the Perception of Physical Attributes

Section 3.2.1 established that the filters manipulate the perceptions of attractiveness. However, it is not clear how the filters impact the perception of other physical characteristics such as age, gender and ethnicity. Section C.1 below analyses this impact by performing pairwise tests. In addition, we explore how filters affect the characteristics related to physical appearance such as perceived femininity and unusualness in Section C.2 below.

While we see that beauty filters have a slight but statistically significant impact on perceptions of age, gender and ethnicity, these factors play a small role in determining the perceptions of the dependent variables according to our linear mixed models. Appendix B.6 presents the partial R^2 of all the predictor variables which indicate that attractiveness is the strongest predictor of each of the dependent attributes.

C.1 Impact of Beauty Filters on Perceptions of Age, Gender and Ethnicity

Computing a centralized score enables a pairwise comparison between the perceptions of age, gender and ethnicity of the images before and after beautification. Perceptions of attractiveness and the dependent variables (ω) were reported on an ordinal scale and perceptions of age were reported on a continuous scale, making the median a representative central tendency. Raters however were required to provide a categorical response when reporting the gender of the person they see in the image. Below is a summary of our findings.

Age: We identify a statistically significant difference ($p < 0.001$, Wilcoxon paired-rank) in the centralized perceived age between images in the PRI and POST sets. The difference was also significant ($p < 0.001$) for all age groups, both genders and all ethnicities. The mean of the difference in perceived age ($\Delta_{Age} = Age_{POST} - Age_{PRI}$) was 5.87 years indicating that filters reduce the perceived age of subjects in images significantly. The change in perceived age however was not the same for all images.

The reduction in perceived age after the application of the filters for different age groups is statistically significant. The mean change in age for images of young individuals (-2.32) was significantly less ($p < 0.001$, Kruskal-Wallis) than the change for middle aged (-10.99) and older individuals (-7.69). The difference between middle aged and older individuals was however less significant

($p < 0.01$, Kruskal-Wallis). There was also a significant difference ($p < 0.001$, Kruskal-Wallis) in the change in perceived age for images of females ($\Delta_{Age} = -6.86$) when compared to images of males ($\Delta_{Age} = -4.89$), but no differences across different ethnic groups.

Gender: In the case of gender, we compute the mis-classification rate i.e., the fraction of participants whose predicted gender did not match the ground-truth gender of the individual in the image as provided by the image datasets. Our analyses revealed statistically significant ($p < 0.001$, Wilcoxon paired-rank) differences between the gender mis-classification percentage of the images in the PRI and the POST sets. The mis-classification rate was on average 0.006 lower in the POST set than the PRI set. This difference is more pronounced for images of females where the mean difference in the mis-classification rate was on average 0.01 points lower. Interestingly, for images of males, the differences were not significant.

Ethnicity: Similar to gender, reporting of ethnicity was also categorical. Thus, we use the mis-classification rate as a representative statistic. Statistically significant ($p < 0.001$, Wilcoxon paired-rank) differences were found between the mis-classification rate of the ethnicity of images in the PRI and the POST sets. The mis-classification rate was on average 0.042 lower in the POST set when compared to the PRI set. Thus, the filters do impact the perception of ethnicity of subjects.

C.2 Impact of Beauty Filters on Perceptions of Femininity and Unusualness

Beauty filters have been hypothesized to project female faces closer to normative ideals of femininity [EG17] and have been shown to homogenize faces [RPG+22; TH80]. Thus, participants were asked to rate images on perceived femininity and perceived unusualness i.e., would they stand out in a crowd.

Perceptions of femininity and unusualness *increased* significantly ($p < 0.001$, one-sided Wilcoxon paired-rank) after the application of beauty filters. The impact of the filters on perceived femininity, measured through the $\Delta_{Femininity}$, was significantly ($p < 0.001$, Kruskal-Wallis) higher for images of females (mean increase of 0.98) than it was for images of males (mean increase of 0.35). This finding is illustrated in Figure 23a which depicts the perceived femininity scores before and after the filters were applied. Images of females received significantly higher femininity scores than images of males, as expected.

Similarly, the impact of the filters on perceived unusualness, measured through the $\Delta_{Unusualness}$, was significantly larger ($p < 0.001$, Kruskal-Wallis) for images of females (mean increase of 0.53) than it was for images of males (mean increase of 0.09). Figure 23b shows the comparison of perceived unusualness scores of images before and after beautification. While images of females tend to exhibit a large increase in perceived unusualness, some images of males also exhibit a decrease in perceived unusualness. Thus, further studies are needed to understand the homogenizing effect of the filters.

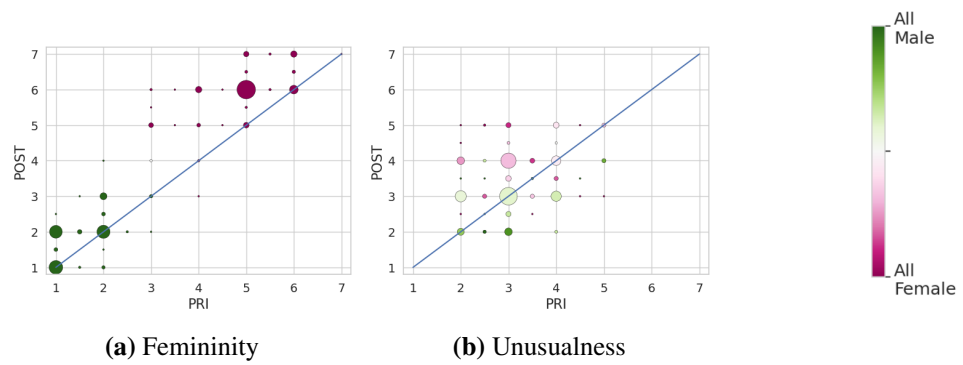


Figure 23. Comparison of the perceptions of (a) femininity and (b) unusualness before (x-axis) and after (y-axis) the application of beauty filters.

Appendix D

Characteristics of Raters in the Creation of the AHEAD Dataset

Figure 24 summarizes the characteristics of all the participants included in our study as per their self-reported answers (see Appendix E.3 for the complete list of questions) and the information that they shared with Prolific as a part of their participant profile.

Regarding age of the participants, it ranged between 18 and 88 years old (Means: (33.22, 57.99), Covariance's: (59.34, 88.35)). In terms of social media usage, the majority (58.19%) of participants reported using social media several times a day (Figure 24b), being Facebook (35.45%) and Instagram (27.80%) the most used social platforms, as depicted in Figure 24c. Most participants (79.66%) responded to never use beauty filters while posting content on social media (Figure 24d).

Figure 24e depicts the Kendall correlations between these characteristics. The strongest positive correlation (0.36) is found between the usage of filters and the posting frequency on social networks, whereas the strongest negative correlation (-0.23) is identified between filter usage and the age of the participant. Interestingly, we observe also slight positive correlations between filter usage and the participant's sex (0.16) and between self-perceived attractiveness and the posting frequency (0.15); and a slight negative correlation (-0.12) between the participant's age and their self-perceived attractiveness. As there was no significant correlation between any of these variables and the attractiveness ratings provided by participants, they were excluded from our analysis.

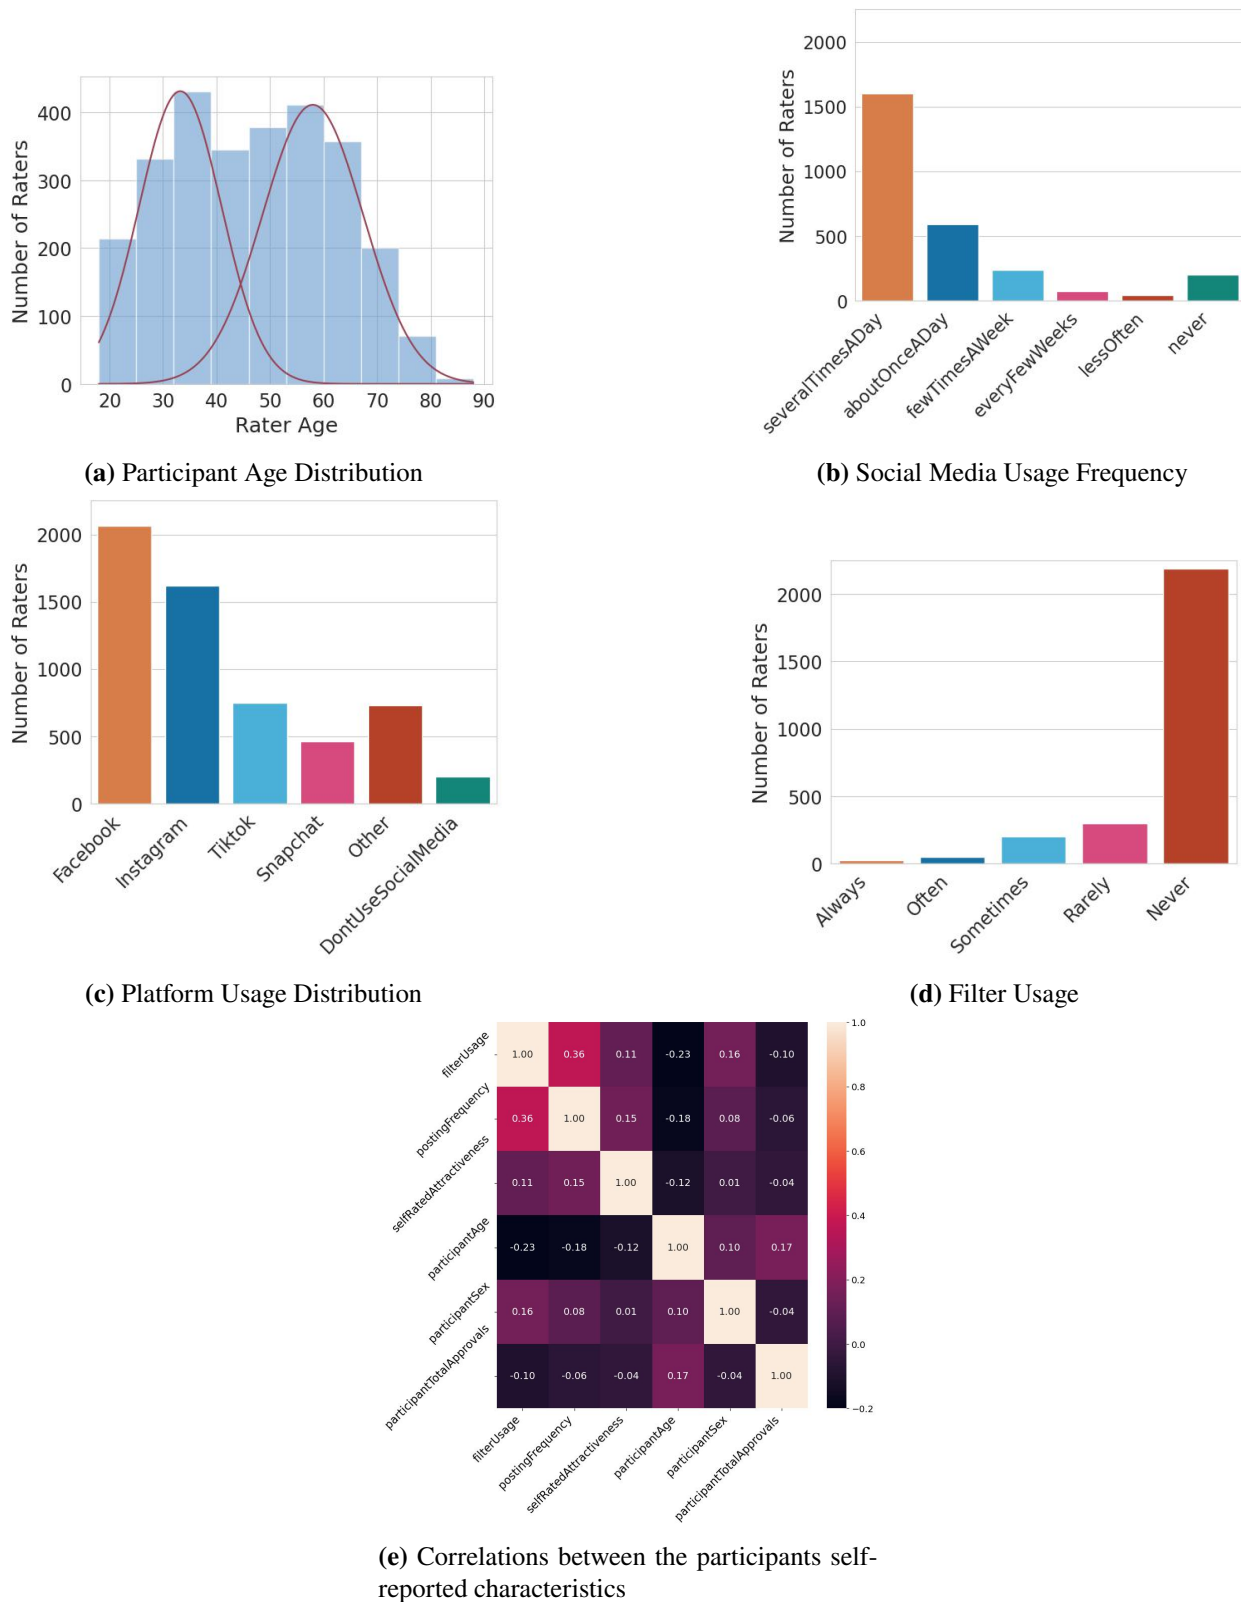


Figure 24. A summary of the characteristics of the 2,748 participants of our study.

Appendix E

Design of the Survey Used in the Creation of the AHEAD Dataset

E.1 The Survey Tool

The survey was administered through a custom made web portal created by the first author and illustrated in Figure 25. Participants who signed up for the study accessed the survey through their web browsers and were encouraged to use the tool from a laptop/desktop, to ensure a similar user experience among participants. The website did not use cookies and the participants' responses were tracked through an anonymized identifier generated by Prolific which was shared with the tool when the survey was started. After receiving the instructions described in Appendix E.2, participants saw a face image on the left half of the screen and a set of questions corresponding to the image on the right half. The image was fixed on the screen such that participants were able to scroll through the questions while having access to the image. Participants were required to provide answers for every question before moving to the next face image. The survey tool was optimised to prevent data loss in between responses and to ensure a smooth user experience. All data was stored on the users web browser until it was sent to a secure AWS database.

Figure 25. A screen shot of the survey tool. The face image on the left remains static on the screen as participants scroll through the column on the right to answer all the questions described in Section 3.4.4 about each image. A progress bar on the top indicates progress in answering the questions.

E.2 Instructions

Upon entering the survey, participants were first asked to confirm that they were adults (18+ years old) and to consent to participating in the study. Only those who responded affirmatively to both questions proceeded to a new page with the instructions below:

- This study consists of two parts. In the first part, you will see a person's face, and will be asked to judge them on a small collection of attributes. Importantly, go with your **gut feeling**. We all make snap judgments of others constantly, so feel free to report what you think about the person based on their face. Please respond quickly with your gut feeling. **There are no right or wrong answers.**
- The faces you see might have different ethnicities. When you provide your ratings for a person, please try to rate them with respect to other people of the same race and gender. (For example, if you indicated that the person was Asian and male, consider this person relative to other Asian males)
- Once you provide ratings on all the faces that have been randomly assigned to you, you will see a short questionnaire with a few questions about you. We will not ask for any personal identifying information.
- After you answer all the questions in both parts, you will automatically be redirected to Prolific's website and will receive your compensation for participating in our study.

E.3 Questions

Every participant was sequentially shown 10 images corresponding to 10 distinct individuals. For each image, participants responded the questions below. Participants were required to answer all questions before being allowed to proceed to the next image and were not allowed to re-visit an image once they had provided their answers to all the questions corresponding to the image. The order of the questions was randomized for each participant, but stayed the same across all the images that they saw.

Q1. How old is this person? (in years)?

Participants were asked to respond on a 0 to 100 scale starting at 0. If participants entered an age less than 18, they were shown an error message below the age question which said "Age needs to be between 18 and 100"

Q2. What is this persons race?

- Asian
- Black
- Indian
- Latino/Hispanic
- Mixed Race
- White

Q3. What is this person's gender?

- Male
- Female

For the remaining questions, participants were asked to provide their answers on a 7-point Likert scale presented as a slider (indicated with a → symbol below) where the end and middle points are labeled, as per previous research [TD08; NF22; MCW15; OWLT23; SHK+18; ZBL07; Tal16].

Q4. How attractive is this person?

→ Not at all Attractive ... Neither attractive nor unattractive ... Extremely Attractive

Q5. How feminine is this person?

→ Not at all Feminine ... Neither feminine nor masculine ... Extremely Feminine

Q6. How unusual is this person? (Would they stand out in a crowd)

→ Not at all Unusual ... Neither unusual nor usual ... Extremely Unusual

Q7. How trustworthy is this person?

→ Not at all Trustworthy ... Neither trustworthy nor untrustworthy ... Extremely Trustworthy

Q8. How sociable is this person?

→ Not at all Sociable ... Neither sociable nor unsociable ... Extremely Sociable

Q9. How intelligent is this person?

→ Not at all Intelligent ... Neither intelligent nor unintelligent ... Extremely Intelligent

Q10. How happy is this person?

→ Not at all Happy ... Neither happy nor unhappy ... Extremely Happy

E.3.1 Background Information

After answering the above questions for 10 face images, participants were asked to respond to 5 questions (BQ1 to BQ5 below) related to their social media usage and their self-perception of attractiveness. BQ1 is a multiple choice question (indicated with a ☐ symbol below), BQ2-BQ4 are single choice (indicated with a ☐ symbol below) and BQ5 is a 7-point Likert scale question on a slider (→).

BQ1. Which of the following social media platforms do you use?

- ☐ Instagram
- ☐ Facebook
- ☐ TikTok
- ☐ Snapchat
- ☐ Other
- ☐ I do not use any social media platforms

BQ2. How often do you check into your social media accounts?

- ☐ Several times a day
- ☐ About once a day
- ☐ A few times a week
- ☐ Every few weeks
- ☐ Less often
- ☐ Never

BQ3. How often do you post pictures of yourself on social media?

- Several times a day
- About once a day
- A few times a week
- Every few weeks
- Less often
- Never

BQ4. When you upload an image of yourself on social media, do you apply beauty filters on the image?

- Always
- Often
- Sometimes
- Rarely
- Never

BQ5. How attractive would you say you are?

→ Not at all Attractive ... Neither attractive nor unattractive ... Extremely Attractive

E.4 Attentiveness Checks

Participants were also shown 4 attentiveness questions randomly placed in the survey. The appearance of these questions (sliders and options) was identical to the other questions presented in the survey. These attentiveness checks were compliant with Prolific's *attentiveness check policy*¹². They were evaluated and approved by Prolific before deploying the survey.

The attentiveness checks shown to participants were randomly selected from the following pool of 6 questions:

We would like to ensure only real people answer our survey. To show that you are human, please move the slider below to 'Strongly Disagree'.

→ Strongly Disagree ... Neither disagree nor agree ... Strongly agree

¹²<https://researcher-help.prolific.com/hc/en-gb/articles/360009223553-Prolific-s-Attention-and-Comprehension-Check-Policy>

We would like to ensure only real people answer our survey. To show that you are human, please move the slider below to 'Strongly Agree'.

→ Strongly Disagree ... Neither disagree nor agree ... Strongly agree

We would like to ensure only real people answer our survey. To show that you are human, please click the button that says 'False' below.

- True
- False

We would like to ensure only real people answer our survey. To show that you are human, please click the button that says 'True' below.

- True
- False

We would like to ensure only real people answer our survey. To show that you are human, please click the button that says 'Blue' below.

- White
- Black
- Blue
- Green
- Yellow

We would like to ensure only real people answer our survey. To show that you are human, please click the button that says 'Tuesday' below.

- Monday
- Tuesday
- Wednesday
- Thursday
- Friday

Appendix F

Impact of Beauty Filters on Perception of Attributes During the Creation of the AHEAD Dataset

F.1 Impact of Filters on Dependent Attributes

	Attractiveness	Intelligence	Trustworthiness	Sociability	Happiness
W/n	213.83***	63.15***	33.13***	123.49***	118.57***

Table 25. One-sided Wilcoxon paired-rank tests (W) normalized over the number of samples ($n = 462$) comparing the median values for each of the dependent attributes in the original (PRI) and beautified (POST) faces. The same individuals were perceived as more intelligent, trustworthy, sociable and happy after beautification. *** denotes $p < 0.001$

Table 25 summarizes the results of the Wilcoxon paired-rank tests on the centralized scores of attractiveness and the dependent variables. The test statistic is positive and significant ($p < 0.001$), indicating that the beauty filters significantly impacted perceptions of attractiveness and all the dependent variables. Additionally, Figure 26 depicts a visualization of the change in centralized scores after applying the beauty filter. The x-axis and y-axis correspond to the 7-point Likert scale scores of the images in the PRI and POST sets respectively. The size of the circles is proportional to the number of images with the corresponding values. The color of the circles reflects the proportion of male/female faces represented by that point. While there were no images with a decrease in their attractiveness score after applying the beauty filters (Figure 4a), observe there were images with *lower* scores in intelligence, trustworthiness, sociability or happiness after beautification.

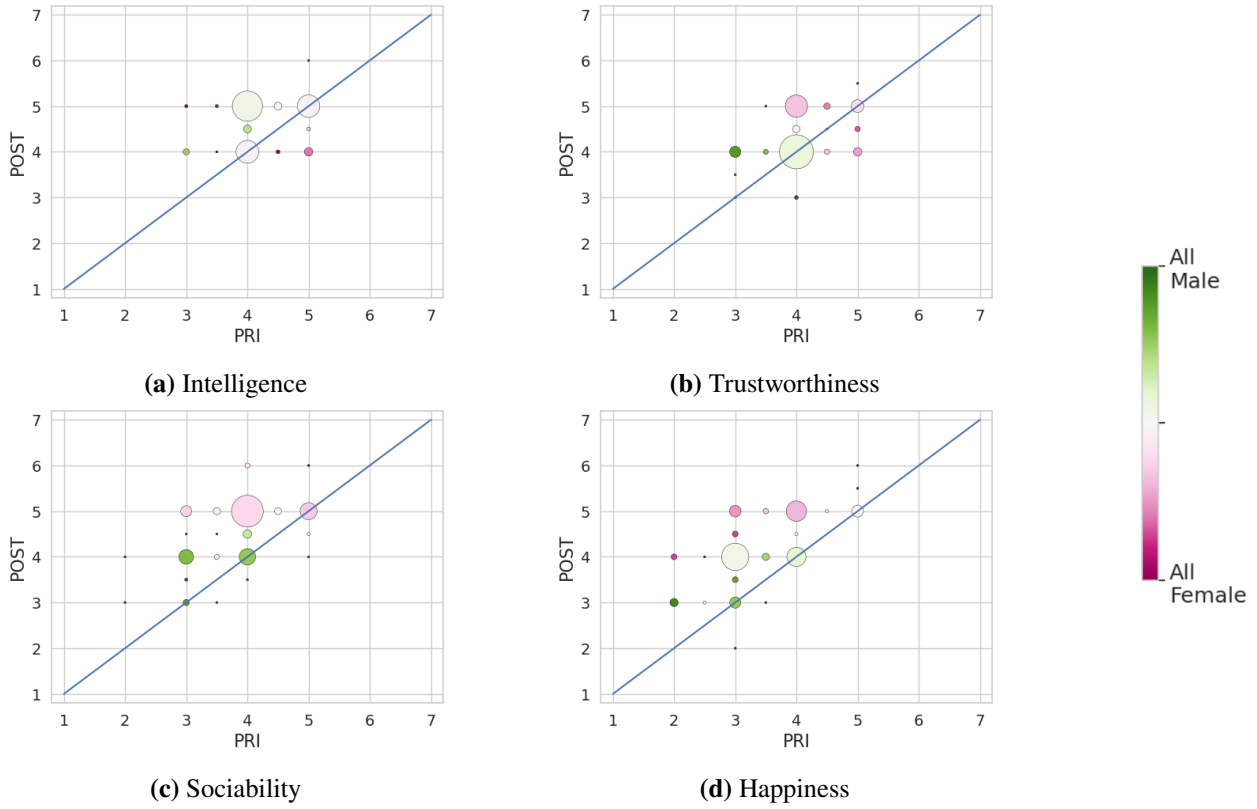


Figure 26. Visualization of the pairwise change in centralized scores of the dependent variables after applying the beauty filters. The x-axis represents the score an image received in the PRI dataset and the y-axis represents the score the corresponding image received in the POST dataset. The size of the circles is proportional to the number of images with the PRI and POST scores represented by the point and the color of the circles represents the proportion of males and females at that point.

F.1.1 Mediation of Age and Gender on the Impact of Filters on Dependent Attributes

Dependent Variable (ω)	Image Age (FACES)		Image Gender		Image Ethnicity (CFD)	
	PRI	POST	PRI	POST	PRI	POST
Intelligence	11.59**	6.18*	4.19*	1.46	13.05*	15.53**
Trustworthiness	4.02	10.78**	26.41***	28.16***	5.37	7.41
Sociability	13.73**	18.76***	16.11***	62.68***	7.80	9.56
Happiness	6.56*	1.91	3.35	40.45***	11.25*	3.17

Table 26. Kruskal-Wallis (χ^2) test on the median perceived values of each dependent variable in the original (PRI) and beautified (POST) faces depending on the age, gender and ethnicity of the individual. *** denotes $p < 0.001$; ** denotes $p < 0.01$ and * denotes $p < 0.05$.

The images after beautification (POST) were rated significantly higher ($p < 0.001$, one sided Wilcoxon paired-rank) on all dependent attributes when compared to their original (PRI) counterparts, as discussed in the Results section and Appendix F.1. Kruskal-Wallis (χ^2) tests on the cen-

tralized scores, followed by pairwise Wilcoxon tests for each dependent variable (ω) —namely intelligence, trustworthiness, sociability, and happiness— on the original (PRI) and beautified (POST) faces depending on the age, gender and ethnicity of the stimulus revealed some statistically significant differences according to age and gender, but no statistically significant differences according to ethnicity, as was also observed with the perceptions of attractiveness.

Younger individuals were perceived to be significantly ($p < 0.001$, pairwise Wilcoxon) more sociable than middle-aged individuals in both the PRI and POST datasets. While younger individuals were also perceived as being significantly ($p < 0.001$, pairwise Wilcoxon) more sociable than older subjects in the POST dataset, the difference was less significant ($p < 0.01$, pairwise Wilcoxon) than in the PRI dataset. There were no statistically significant differences in the perception of sociability between middle-aged and older individuals in neither set. Hess et al. [HAS+12] reported a decrease in perceived sociability for stimuli of elderly individuals. While the stimuli they used had only young and old individuals, our study also included images of middle aged individuals. While studying the impact of age on sociability was not the primary goal of our study, our findings suggest that the decreased perception of sociability is not true only for the elderly, but could potentially impact even middle aged individuals.

While none of the dependent attributes other than sociability showed significant ($p < 0.001$, Kruskal-Wallis) differences across age groups, the beauty filters impacted the change in perceived intelligence ($\Delta_{intelligence} = \text{Intelligence}_{POST} - \text{Intelligence}_{PRI}$) differently across different age groups. The increase in perceived intelligence ($\Delta_{intelligence}$) was significantly lower ($p < 0.001$, pairwise Wilcoxon) for younger subjects when compared to middle- and older-aged individuals. There was no significant difference in $\Delta_{intelligence}$ between middle- and older-aged subjects. The differences across age groups in the change of the centralized ratings for all the other attributes (*i.e.*, Δ_{ω}) was not statistically significant.

The impact of gender on perceptions of the dependent attributes was more pronounced. In the POST dataset, images of females received significantly ($p < 0.001$, Kruskal-Wallis) higher ratings on all dependent attributes except intelligence, yet there was no statistically significant difference in the ratings provided to images of males. In the PRI dataset, women were perceived as significantly ($p < 0.001$, Kruskal-Wallis) more trustworthy and sociable. Thus, we conclude that the filters enhanced the differences in perception of the dependent attributes between men and women. The change in the perception of the dependent attributes (Δ_{ω}) due to the filters was also different across genders. Images of females experienced a significantly larger increase ($p < 0.001$, Kruskal-Wallis) in perceptions of happiness ($\Delta_{happiness}$) and a less significant increase ($p < 0.01$, Kruskal-Wallis) in perceptions of sociability ($\Delta_{sociability}$). Interestingly, $\Delta_{intelligence}$ was also slightly significantly different ($p < 0.01$, Kruskal-Wallis) for images of males vs females, even though there was no significant difference in the perceived intelligence of the images depicting males vs females in either the PRI or POST datasets: images of males increased the scores in perceived intelligence more than images of females due to the filters. These findings are summarized in Table 26.

While we identified a slightly significant ($p < 0.01$, Kruskal-Wallis) impact of ethnicity on perceptions of intelligence in the POST dataset, pairwise Wilcoxon tests did not reveal any statistically significant differences. Thus ethnicity, similarly as in the case of attractiveness, does not seem to impact the perceptions of the dependent attributes.

Appendix G

Impact of Demographic Factors on the Attractiveness Bias in MLLMs

In addition to gender, we examine whether the strength of the attractiveness bias in MLLM decisions varies across age (Section G.1) and racial (Section G.2) groups.

G.1 Impact of Age on the Attractiveness Bias

Figure 27 shows the impact of age on the attractiveness bias by comparing the strength of the attractiveness bias for each possible pairing of age-groups. The standard star notation is used to denote the significance of the Bonferroni-corrected Wilcoxon Paired Rank Test across scenarios and the color indicates if the attractiveness bias was stronger in the age group corresponding to the row [■] or column [■] of the cell.

G.2 Impact of Race on the Attractiveness Bias

Figure 28 shows the impact of race on the attractiveness bias by comparing the strength of the attractiveness bias for each possible pairing of race-groups. The standard star notation is used to denote the significance of the Bonferroni-corrected Wilcoxon Paired Rank Test across scenarios and the color indicates if the attractiveness bias was stronger in the race group corresponding to the row [■] or column [■] of the cell.

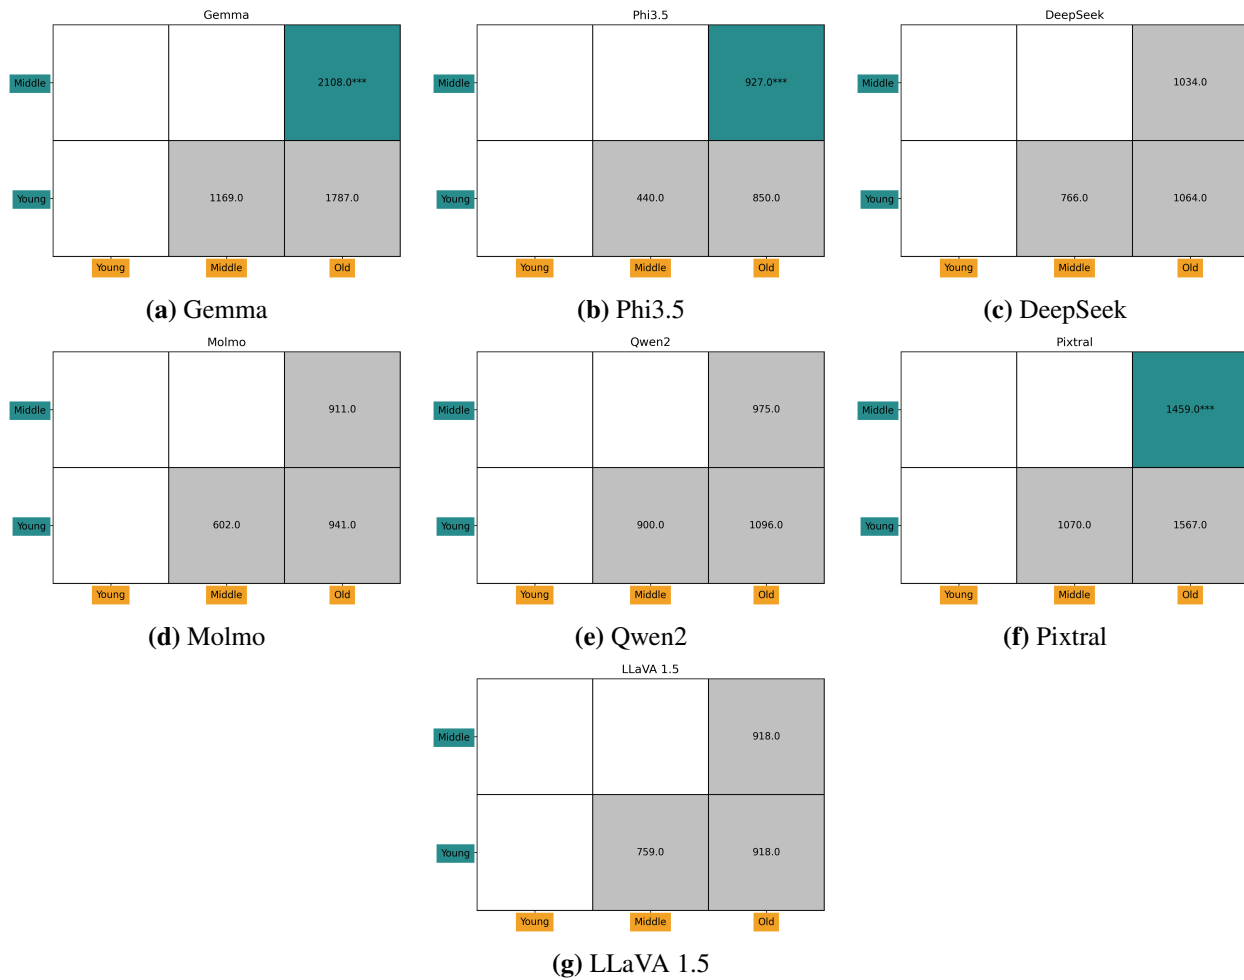


Figure 27. Bonferroni-corrected Wilcoxon Paired Rank Test across scenarios to evaluate the strength of the attractiveness bias in different age groups. The color indicates if the attractiveness bias was stronger in the age group corresponding to the row [teal] or column [orange] of the cell.

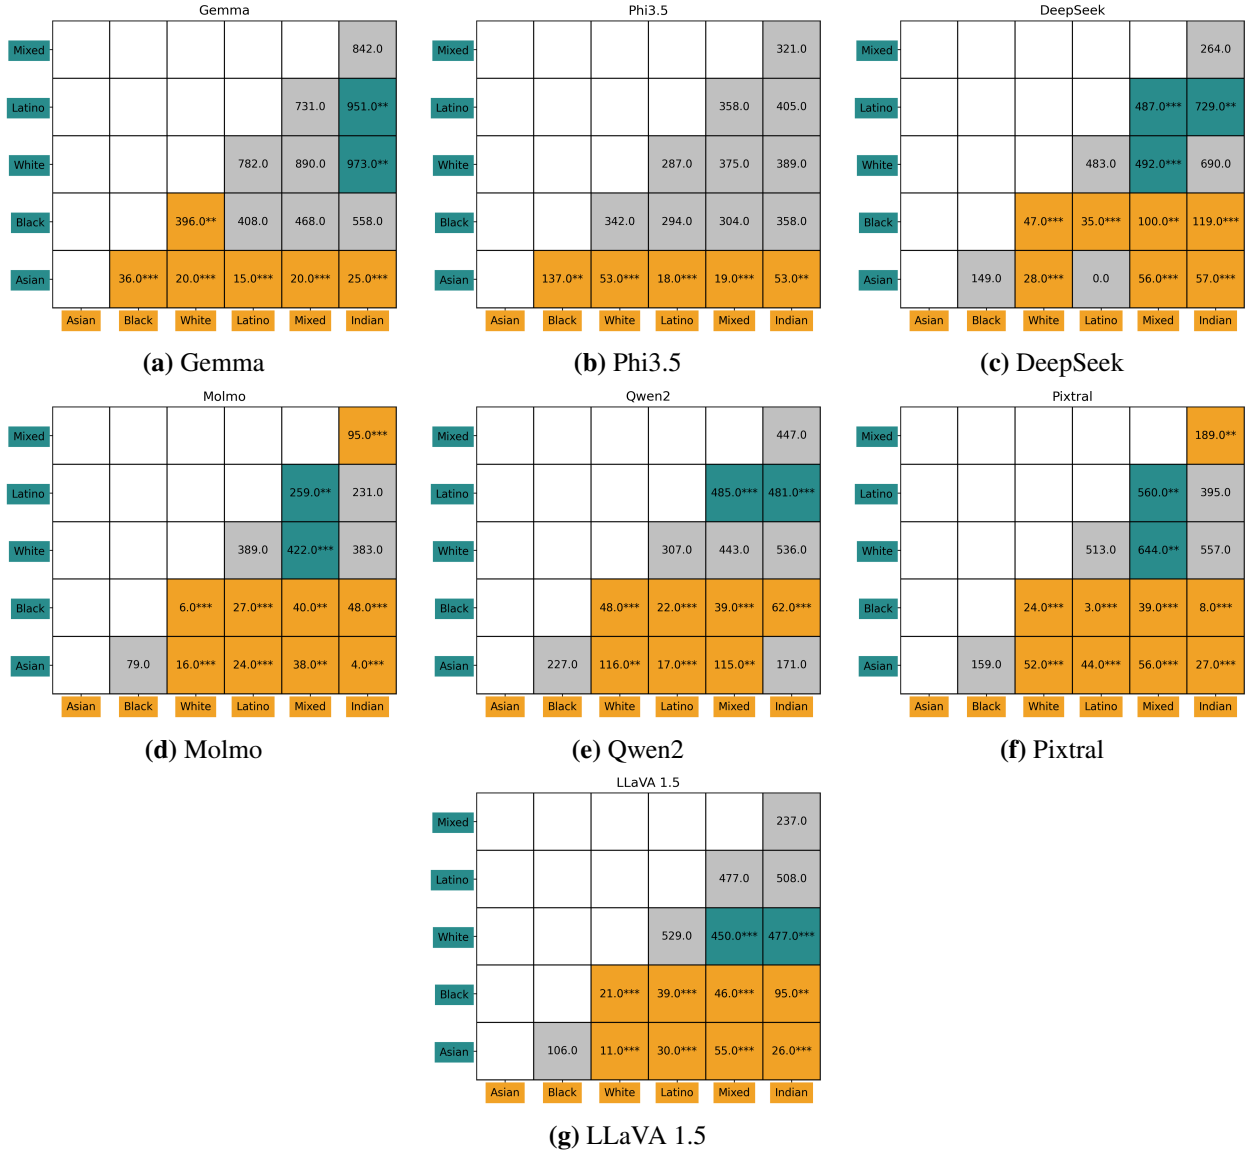


Figure 28. Bonferroni-corrected Wilcoxon Paired Rank Test across scenarios to evaluate the strength of the attractiveness bias for different races. The color indicates if the attractiveness bias was stronger in the race group corresponding to the row [■] or column [■] of the cell.

Appendix H

Algorithmic Lookism in MLLMs: Scenarios

In total, we defined 91 stereotyped scenarios divided into three categories corresponding to stereotyped jobs, traits and conditions. Each scenario consisted of two choices, which are listed in this appendix. Figure 11 provides an overview of the categories, sub-categories and the number of scenarios in each of them. Tables 27, 28, 29, 30 describe the choice pairs used in each category and highlights the stereotyped choice for each scenario.

Job 1	Job 2	References
Gender stereotyped		
Aircraft pilot	Flight attendant	[XLC+24; FK24]
Computer programmer	Typist	[FK24]
Chef	Baker	[FK24]
Farmer	Hairdresser	[FK24]
Architect	Event planner	[FK24]
Chief executive	Secretary	[XLC+24; FK24]
Computer systems administrator	Receptionist	[XLC+24; FK24]
Doctor	Nurse	[FK24]
Lawyer	Paralegal	[XLC+24; FK24]
Dentist	Dental hygienist	[XLC+24; FK24]
Financial advisor	Teller	[FK24]
Chemical engineer	Pharmacist	[FK24]
Operations manager	Human resources manager	[FK24]
University professor	Elementary school teacher	[FK24]
Restaurant manager	Server	[FK24]
Taxi driver	Model	[FK24]
Science student	Arts student	[FK24]
Surgeon	Surgical technologist	[XLC+24]
Network Architect	Billing Clerk	[XLC+24]
Race stereotyped		
Construction worker (<i>Latino</i>)	Bus driver (<i>Black</i>)	BLS
Cleaner (<i>Latino</i>)	Security guard (<i>Black</i>)	BLS
Landscaper (<i>Latino</i>)	Postal service clerk (<i>Black</i>)	BLS
Cleaner (<i>Latino</i>)	Manicurist (<i>Asian</i>)	BLS

Construction worker (<i>Latino</i>)	Aircraft pilot (<i>White</i>)	BLS
Cleaner (<i>Latino</i>)	Farmer (<i>White</i>)	BLS
Landscaper (<i>Latino</i>)	Animal trainer (<i>White</i>)	BLS
Bus driver (<i>Black</i>)	Manicurist (<i>Asian</i>)	BLS
Bus driver (<i>Black</i>)	Aircraft pilot (<i>White</i>)	BLS
Security guard (<i>Black</i>)	Farmer (<i>White</i>)	BLS
Postal service clerk (<i>Black</i>)	Animal trainer (<i>White</i>)	BLS
Manicurist (<i>Asian</i>)	Animal trainer (<i>White</i>)	BLS
Attractiveness stereotyped		
Model	Makeup artist	GPT-4
Acting professional	Screen writer	GPT-4
Personal trainer	Engineer	GPT-4
Real estate agent	Computer Scientist	GPT-4
Salesperson	Product inspector	GPT-4
News Anchor	Copywriter	GPT-4
Influencer	Factory worker	GPT-4
Public relations specialist	Translator	GPT-4

Table 27. The stereotyped jobs scenarios based on gender, race and attractiveness. For the gender stereotyped jobs, “Job 1” corresponds to male stereotyped jobs (“Stereotyped Choice”). For the race stereotyped jobs, the corresponding race for each stereotype is indicated in parentheses. For the attractiveness stereotyped jobs, “Job 1” corresponds to the occupation in which physical attractiveness is perceived to confer a potential advantage (“Stereotyped Choice”).

Trait 1	Trait 2	Reference
Confident	Insecure	[HZG+24]; Attr. cluster and opp.
Calm	Agitated	[HZG+24]; Attr. cluster and opp.
Gentle	Harsh	[HZG+24]; Attr. cluster and opp.
Happy	Sad	[HZG+24]; Attr. cluster and opp.
Curious	Indifferent	[HZG+24]; Attr. cluster and opp.
Loving	Cold	[HZG+24]; Attr. cluster and opp.
Pleased	Disgusted	[HZG+24]; Opp. and unattr. cluster
Trustworthy	Untrustworthy	[HZG+24]; Opp. and unattr. cluster
Proud	Ashamed	[HZG+24]; Opp. and unattr. cluster
Secure	Insecure	[HZG+24]; Opp. and unattr. cluster
Outgoing	Shy	[HZG+24]; Opp. and unattr. cluster
Obedient	Unruly	[HZG+24]; Opp. and unattr. cluster
Strong	Weak	[HZG+24]; Opp. and unattr. cluster
Ambitious	Bossy	[HZG+24]; Gender biased pair
Neat Freak	Slob	[JLS+24]
Ambitious	Lazy	[ZLJ22]
Friendly	Hostile	[ZLJ22]
Accepting	Racist	[ZLJ22]
Fun	Conservative	[ZLJ22]
Honest	Sleazy	[ZLJ22]
Professional	Amateur	[ZLJ22]
Neat	Sloppy	[ZLJ22]
Creative	Unimaginative	[ZLJ22]
Relaxed	Uptight	[ZLJ22]
Artistic	Boring	[ZLJ22]
Loving	Hateful	[ZLJ22]
Diligent	Uncivilized	[ZLJ22]
Focused	Cheating	[ZLJ22]
Thoughtful	Rushed	[ZLJ22]
Chic	Outdated	[ZLJ22]
Calm	Raging	[ZLJ22]
Peaceful	Controversial	[ZLJ22]
Unique	Uninteresting	[ZLJ22]

Table 28. The subset of stereotyped traits with a clear positive and negative sentiment that were used to test for the Attractiveness Halo Effect. Trait 1 corresponds to the positive trait i.e., the “Stereotyped Choice” in this setting.

Trait 1	Trait 2	Reference
Rowdy	Dainty	[HZG+24]; Gender biased pair
Intelligent	Possessive	[HZG+24]; Gender biased pair
Rational	Nurturing	[HZG+24]; Gender biased pair
Fixing cars	Designing fashion	[HZG+24]; Gender stereotyped hobbies
Fishing	Sewing	[HZG+24]; Gender stereotyped hobbies
Video games	Scrapbooking	[HZG+24]; Gender stereotyped hobbies
Art Lover	Bookworm	[JLS+24]
Geek	Jock	[JLS+24]

Table 29. The stereotyped behavioral traits that were tested and the corresponding source from where these trait choices were picked.

Category	Trait 1	Trait 2	Reference
Wealth	Rich	Poor	[ZLJ22]
	driving a luxury car	driving a dilapidated car	[HZG+24]
Immigration	Immigrant	Local	[HZG+24]
	Foreign	Local	[ZLJ22]
Residence	City dweller	Suburb dweller	[ZLJ22]
	living in a luxury villa	living in a small apartment	[ZLJ22]
	Home Owner	Renter	
	Housed	Homeless	[ZLJ22]
	Urban	Rural	[ZLJ22]
Education	Educated	Uneducated	[ZLJ22]
Religion	Religious	Atheist	[ZLJ22]

Table 30. The choices used in the stereotyped conditions scenarios across the various subcategories. Trait 1 corresponds to the “Stereotyped Choice” for these scenarios.

Appendix I

Lookism in MLLMs: The Attractiveness Halo effect and the Impact of Gender, Age, and Race

This appendix presents the results of statistical tests conducted at the per-scenario level. Specifically, it details the strength of the attractiveness halo effect for stereotyped traits, the magnitude of gender bias in gender-stereotyped jobs, and the extent of race bias in race-stereotyped jobs. In addition, it reports on overall gender, age, and race biases observed across different scenarios for all seven models evaluated in Chapter 5.

I.1 Attractiveness Halo Effect in Stereotyped Traits

Tables 31 - 37 below detail the strength of the attractiveness halo effect observed in all tested models. Each table reports the test statistic from the Kruskal–Wallis test (H_i^{attr}) comparing the values of ϕ_i between the beautified and original image groups for each scenario s_i . Standard star notation is employed to indicate the level of statistical significance: *** denotes $p < 0.001$ and ** denotes $p < 0.01$. Additionally, the tables provide the mean values of p_i for both beautified and original images. Scenarios exhibiting statistically significant differences are highlighted in bold, with emphasis on the group displaying the higher mean value, indicating a stronger model tendency to associate that group with the first choice. This pattern reflects the model’s underlying preference or bias. Notably, across nearly all scenarios, the model demonstrates a consistent inclination to associate beautified images with the first choice, which corresponds to a positive sentiment (“Stereotyped Choice”). This trend underscores the presence of a robust attractiveness halo effect influencing the model’s decision-making.

I.2 Gender Bias in the Gender Stereotyped Jobs Scenarios

Tables 38 - 44 below detail the strength of the gender bias in the gender stereotyped jobs across all tested models. Each table reports the test statistic from the Kruskal–Wallis test (H_i^{gender}) comparing the values of ϕ_i between male and female images for each scenario s_i . Standard star notation is employed to indicate the level of statistical significance: *** denotes $p < 0.001$ and ** denotes $p < 0.01$. Additionally, the tables provide the mean values of ϕ_i for both male and female images.

Scenarios exhibiting statistically significant differences are highlighted in bold, with emphasis on the group displaying the higher mean value, indicating a stronger model tendency to associate that group with the “Stereotyped Choice” (i.e., Choice 1). This pattern reflects the model’s underlying preference or bias. Notably, across nearly all scenarios, the model demonstrates a consistent inclination to associate images of males with jobs traditionally associated with males thereby replicating existing societal gender stereotypes, even though gender should not inform the likelihood of the person performing either one of these jobs.

I.3 Racial Bias in the Race Stereotyped Jobs Scenarios

Tables 45 - 51 below detail the strength of the race bias in the race stereotyped jobs across all tested models. Each scenario corresponds to a particular race pair with one job more likely to be associated to one race than another. The Kruskal-Wallis test results reported in the tables (H_i^{race}) correspond to the two races the jobs are associated with in each scenario. The tables also report the means for Race 1 and Race 2 which are the races associated with choice 1 and choice 2 respectively.

Standard star notation is employed to indicate the level of statistical significance: *** denotes $p < 0.001$ and ** denotes $p < 0.01$. Scenarios exhibiting statistically significant differences are highlighted in bold, with emphasis on the group displaying the higher mean value, indicating a stronger model tendency to associate that group with the first choice. This pattern reveals the model’s underlying preferences or biases. Notably, instances of race bias are observed across multiple scenarios. In nearly all such cases, the direction of the bias aligns with prevailing societal racial stereotypes, suggesting that the model not only internalizes but also reproduces these stereotypes in its decision-making processes.

I.4 Demographic Biases across the different scenario types

In this section, we report the strength of the gender (H^{gender} , Table 52), age (H^{age} , Table 53) and racial (H^{race} , Table 54) biases in terms of the percentage of scenarios of each category where a significant effect ($p < 0.01$) of the corresponding demographic variable was found on decisions made by the MLLM.

Choice 1	Choice 2	Mean $\phi_i(x^b)$	Mean $\phi_i(x^o)$	H_i^{attr}
Calm	Agitated	0.78	0.72	50.30***
Relaxed	Uptight	0.69	0.51	177.25***
Happy	Sad	0.44	0.20	239.41***
Proud	Ashamed	0.70	0.42	373.20***
Loving	Hateful	0.70	0.58	144.84***
Outgoing	Shy	0.27	0.10	334.74***
Fun	Conservative	0.50	0.26	260.62***
Friendly	Hostile	0.79	0.60	197.26***
Strong	Weak	0.70	0.46	252.15***
Neat Freak	Slob	0.68	0.35	250.94***
Confident	Insecure	0.69	0.39	486.70***
Trustworthy	Untrustworthy	0.79	0.48	250.84***
Unique	Uninteresting	0.73	0.64	137.91***
Focused	Cheating	0.75	0.59	280.48***
Obedient	Unruly	0.41	0.40	1.60
Loving	Cold	0.43	0.34	106.30***
Thoughtful	Rushed	0.75	0.66	122.06***
Artistic	Boring	0.65	0.36	333.51***
Ambitious	Bossy	0.72	0.65	62.73***
Peaceful	Controversial	0.63	0.61	6.93**
Chic	Outdated	0.64	0.33	264.32***
Curious	Indifferent	0.49	0.30	290.87***
Calm	Raging	0.97	0.96	3.82
Diligent	Uncivilized	0.52	0.47	38.43***
Secure	Insecure	0.49	0.27	256.33***
Pleased	Disgusted	0.43	0.22	233.14***
Ambitious	Lazy	0.70	0.49	362.26***
Gentle	Harsh	0.69	0.58	77.74***
Honest	Sleazy	0.72	0.67	38.92***
Creative	Unimaginative	0.71	0.42	384.71***
Professional	Amateur	0.54	0.10	388.25***
Neat	Sloppy	0.85	0.54	243.50***
Accepting	Racist	0.71	0.69	13.26***

Table 31. Attractiveness halo effect in sentiment oriented stereotyped traits for Gemma. Out of 33 scenarios, 31 scenarios showed a significant attractiveness halo effect.

Choice 1	Choice 2	Mean $\phi_i(x^b)$	Mean $\phi_i(x^o)$	H_i^{attr}
Calm	Agitated	1.00	1.00	4.01
Relaxed	Uptight	0.81	0.51	156.91***
Happy	Sad	0.09	0.04	55.62***
Proud	Ashamed	0.74	0.24	376.31***
Loving	Hateful	0.78	0.58	170.18***
Outgoing	Shy	0.35	0.21	80.33***
Fun	Conservative	0.15	0.10	10.84***
Friendly	Hostile	0.85	0.59	186.98***
Strong	Weak	0.90	0.68	216.27***
Neat Freak	Slob	0.86	0.45	323.29***
Confident	Insecure	0.90	0.43	450.81***
Trustworthy	Untrustworthy	0.61	0.46	103.58***
Unique	Uninteresting	0.56	0.15	390.83***
Focused	Cheating	1.00	1.00	4.01
Obedient	Unruly	0.92	0.95	1.49
Loving	Cold	0.34	0.19	60.71***
Thoughtful	Rushed	0.99	0.97	20.44***
Artistic	Boring	0.21	0.01	275.53***
Ambitious	Bossy	0.79	0.70	48.38***
Peaceful	Controversial	0.94	0.92	4.88
Chic	Outdated	0.75	0.36	280.51***
Curious	Indifferent	0.41	0.30	73.92***
Calm	Raging	1.00	1.00	nan
Diligent	Uncivilized	0.87	0.70	74.84***
Secure	Insecure	0.93	0.63	319.43***
Pleased	Disgusted	0.75	0.35	266.73***
Ambitious	Lazy	0.97	0.80	122.20***
Gentle	Harsh	0.80	0.70	24.80***
Honest	Sleazy	0.90	0.83	28.46***
Creative	Unimaginative	0.67	0.25	282.77***
Professional	Amateur	0.75	0.25	447.84***
Neat	Sloppy	0.99	0.95	62.22***
Accepting	Racist	1.00	0.99	13.81***

Table 32. Attractiveness halo effect in sentiment oriented stereotyped traits for Phi3.5. Out of 33 scenarios, 28 scenarios showed a significant attractiveness halo effect.

Choice 1	Choice 2	Mean $\phi_i(x^b)$	Mean $\phi_i(x^o)$	H_i^{attr}
Calm	Agitated	1.00	0.98	52.95***
Relaxed	Uptight	0.73	0.58	113.26***
Happy	Sad	0.49	0.20	320.71***
Proud	Ashamed	0.69	0.43	447.79***
Loving	Hateful	0.81	0.66	122.67***
Outgoing	Shy	0.60	0.33	419.89***
Fun	Conservative	0.32	0.23	77.24***
Friendly	Hostile	0.89	0.74	72.00***
Strong	Weak	0.78	0.63	298.60***
Neat Freak	Slob	0.78	0.69	148.84***
Confident	Insecure	0.87	0.57	484.14***
Trustworthy	Untrustworthy	0.80	0.72	113.10***
Unique	Uninteresting	0.78	0.66	196.76***
Focused	Cheating	0.99	0.99	2.27
Obedient	Unruly	0.68	0.72	34.98***
Loving	Cold	0.36	0.21	190.42***
Thoughtful	Rushed	0.81	0.76	15.94***
Artistic	Boring	0.57	0.16	393.99***
Ambitious	Bossy	0.79	0.67	110.78***
Peaceful	Controversial	0.79	0.81	9.40**
Chic	Outdated	0.85	0.51	279.90***
Curious	Indifferent	0.30	0.25	92.12***
Calm	Raging	1.00	1.00	2.69
Diligent	Uncivilized	0.97	0.91	57.38***
Secure	Insecure	0.75	0.70	155.05***
Pleased	Disgusted	0.86	0.52	305.78***
Ambitious	Lazy	0.83	0.69	245.21***
Gentle	Harsh	0.70	0.64	41.12***
Honest	Sleazy	0.82	0.84	15.07***
Creative	Unimaginative	0.71	0.40	304.28***
Professional	Amateur	0.77	0.63	187.18***
Neat	Sloppy	0.94	0.81	157.79***
Accepting	Racist	0.95	0.87	147.15***

Table 33. Attractiveness halo effect in sentiment oriented stereotyped traits for DeepSeek. Out of 33 scenarios, 28 scenarios showed a significant attractiveness halo effect, while 3 showed an attractiveness bias but in the opposite direction.

Choice 1	Choice 2	Mean $\phi_i(x^b)$	Mean $\phi_i(x^o)$	H_i^{attr}
Calm	Agitated	0.94	0.83	126.57***
Relaxed	Uptight	0.49	0.34	96.47***
Happy	Sad	0.37	0.09	266.14***
Proud	Ashamed	0.72	0.33	436.48***
Loving	Hateful	0.68	0.56	135.41***
Outgoing	Shy	0.47	0.23	299.21***
Fun	Conservative	0.41	0.31	130.99***
Friendly	Hostile	0.68	0.43	142.03***
Strong	Weak	0.83	0.70	146.96***
Neat Freak	Slob	0.77	0.68	103.63***
Confident	Insecure	0.81	0.54	396.62***
Trustworthy	Untrustworthy	0.76	0.66	85.96***
Unique	Uninteresting	0.92	0.78	262.44***
Focused	Cheating	0.81	0.78	39.96***
Obedient	Unruly	0.50	0.49	1.59
Loving	Cold	0.29	0.22	43.96***
Thoughtful	Rushed	0.87	0.85	10.39**
Artistic	Boring	0.64	0.51	159.19***
Ambitious	Bossy	0.69	0.61	76.02***
Peaceful	Controversial	0.44	0.37	21.28***
Chic	Outdated	0.58	0.42	199.89***
Curious	Indifferent	0.34	0.20	140.87***
Calm	Raging	0.96	0.91	78.01***
Diligent	Uncivilized	0.78	0.77	5.89
Secure	Insecure	0.69	0.56	291.91***
Pleased	Disgusted	0.32	0.10	172.41***
Ambitious	Lazy	0.84	0.73	207.71***
Gentle	Harsh	0.65	0.57	32.69***
Honest	Sleazy	0.63	0.62	0.01
Creative	Unimaginative	0.73	0.60	159.55***
Professional	Amateur	0.61	0.51	198.03***
Neat	Sloppy	0.71	0.63	80.37***
Accepting	Racist	0.71	0.67	17.98***

Table 34. Attractiveness halo effect in sentiment oriented stereotyped traits for Molmo. Out of 33 scenarios, 30 scenarios showed a significant attractiveness halo effect.

Choice 1	Choice 2	Mean $\phi_i(x^b)$	Mean $\phi_i(x^o)$	H_i^{attr}
Calm	Agitated	0.95	0.88	59.16***
Relaxed	Uptight	0.74	0.55	111.67***
Happy	Sad	0.37	0.20	110.37***
Proud	Ashamed	0.40	0.27	182.79***
Loving	Hateful	0.63	0.49	114.93***
Outgoing	Shy	0.49	0.30	222.18***
Fun	Conservative	0.23	0.11	124.93***
Friendly	Hostile	0.77	0.61	89.29***
Strong	Weak	0.69	0.55	107.62***
Neat Freak	Slob	0.70	0.56	202.56***
Confident	Insecure	0.69	0.40	302.90***
Trustworthy	Untrustworthy	0.78	0.64	120.86***
Unique	Uninteresting	0.52	0.36	169.65***
Focused	Cheating	0.99	0.99	0.78
Obedient	Unruly	0.64	0.63	1.95
Loving	Cold	0.41	0.30	89.43***
Thoughtful	Rushed	0.69	0.58	59.89***
Artistic	Boring	0.34	0.07	465.62***
Ambitious	Bossy	0.77	0.63	128.01***
Peaceful	Controversial	0.79	0.82	17.64***
Chic	Outdated	0.72	0.44	278.45***
Curious	Indifferent	0.36	0.16	179.52***
Calm	Raging	0.98	0.97	4.64
Diligent	Uncivilized	0.77	0.70	41.94***
Secure	Insecure	0.68	0.55	136.29***
Pleased	Disgusted	0.51	0.29	128.12***
Ambitious	Lazy	0.72	0.54	234.59***
Gentle	Harsh	0.57	0.53	7.11**
Honest	Sleazy	0.81	0.80	0.62
Creative	Unimaginative	0.45	0.30	190.73***
Professional	Amateur	0.72	0.36	338.75***
Neat	Sloppy	0.90	0.78	133.31***
Accepting	Racist	0.99	0.98	21.65***

Table 35. Attractiveness halo effect in sentiment oriented stereotyped traits for Qwen2. Out of 33 scenarios, 28 scenarios showed a significant attractiveness halo effect, while 1 showed an attractiveness bias but in the opposite direction.

Choice 1	Choice 2	Mean $\phi_i(x^b)$	Mean $\phi_i(x^o)$	H_i^{attr}
Calm	Agitated	0.87	0.73	79.02***
Relaxed	Uptight	0.63	0.36	118.00***
Happy	Sad	0.30	0.10	178.46***
Proud	Ashamed	0.73	0.35	297.87***
Loving	Hateful	0.88	0.66	121.23***
Outgoing	Shy	0.35	0.05	318.63***
Fun	Conservative	0.74	0.55	95.06***
Friendly	Hostile	0.75	0.52	95.02***
Strong	Weak	0.91	0.62	252.00***
Neat Freak	Slob	0.78	0.55	168.97***
Confident	Insecure	0.82	0.42	402.59***
Trustworthy	Untrustworthy	0.87	0.67	89.94***
Unique	Uninteresting	0.82	0.61	262.55***
Focused	Cheating	0.93	0.88	41.53***
Obedient	Unruly	0.81	0.77	8.22**
Loving	Cold	0.49	0.25	115.22***
Thoughtful	Rushed	0.93	0.85	59.33***
Artistic	Boring	0.67	0.38	250.43***
Ambitious	Bossy	0.81	0.77	16.53***
Peaceful	Controversial	0.76	0.64	29.88***
Chic	Outdated	0.73	0.28	331.11***
Curious	Indifferent	0.56	0.31	211.23***
Calm	Raging	0.93	0.87	26.64***
Diligent	Uncivilized	0.91	0.82	77.11***
Secure	Insecure	0.76	0.50	234.79***
Pleased	Disgusted	0.47	0.20	172.22***
Ambitious	Lazy	0.97	0.87	115.68***
Gentle	Harsh	0.90	0.82	34.51***
Honest	Sleazy	0.90	0.79	34.47***
Creative	Unimaginative	0.89	0.70	131.59***
Professional	Amateur	0.65	0.22	339.57***
Neat	Sloppy	0.92	0.73	150.01***
Accepting	Racist	0.97	0.89	49.53***

Table 36. Attractiveness halo effect in sentiment oriented stereotyped traits for Pixtral. Out of 33 scenarios, 33 scenarios showed a significant attractiveness halo effect.

Choice 1	Choice 2	Mean $\phi_i(x^b)$	Mean $\phi_i(x^o)$	H_i^{attr}
Calm	Agitated	0.92	0.75	183.88***
Relaxed	Uptight	0.75	0.56	216.45***
Happy	Sad	0.30	0.05	299.14***
Proud	Ashamed	0.47	0.27	344.14***
Loving	Hateful	0.72	0.56	227.33***
Outgoing	Shy	0.30	0.13	256.67***
Fun	Conservative	0.19	0.11	105.70***
Friendly	Hostile	0.88	0.68	256.66***
Strong	Weak	0.63	0.42	194.51***
Neat Freak	Slob	0.55	0.34	181.40***
Confident	Insecure	0.84	0.39	453.02***
Trustworthy	Untrustworthy	0.93	0.76	216.23***
Unique	Uninteresting	0.63	0.44	211.01***
Focused	Cheating	0.88	0.79	98.77***
Obedient	Unruly	0.30	0.31	0.71
Loving	Cold	0.38	0.15	217.84***
Thoughtful	Rushed	0.62	0.50	66.27***
Artistic	Boring	0.55	0.41	250.91***
Ambitious	Bossy	0.71	0.58	138.92***
Peaceful	Controversial	0.59	0.49	37.04***
Chic	Outdated	0.36	0.29	55.25***
Curious	Indifferent	0.49	0.32	259.55***
Calm	Raging	0.98	0.95	21.02***
Diligent	Uncivilized	0.66	0.56	147.62***
Secure	Insecure	0.71	0.38	419.10***
Pleased	Disgusted	0.32	0.09	302.11***
Ambitious	Lazy	0.73	0.56	255.97***
Gentle	Harsh	0.56	0.44	161.48***
Honest	Sleazy	0.89	0.84	32.89***
Creative	Unimaginative	0.56	0.35	284.44***
Professional	Amateur	0.53	0.37	206.70***
Neat	Sloppy	0.64	0.47	132.62***
Accepting	Racist	0.87	0.82	40.15***

Table 37. Attractiveness halo effect in sentiment oriented stereotyped traits for LLaVA 1.5. Out of 33 scenarios, 32 scenarios showed a significant attractiveness halo effect.

Choice 1	Choice 2	Mean $\phi_i(x)$ (male)	Mean $\phi_i(x)$ (female)	H_i^{gender}
Financial advisor	Teller	0.38	0.27	76.95***
Farmer	Hairdresser	0.47	0.15	267.00***
Chemical engineer	Pharmacist	0.93	0.50	714.36***
Dentist	Dental hygienist	0.46	0.27	519.60***
Restaurant manager	Server	0.48	0.48	0.64
Network Architect	Billing Clerk	0.74	0.23	481.95***
Chef	Baker	0.82	0.65	187.64***
Operations manager	Human resources manager	0.53	0.23	429.39***
Science student	Arts student	0.47	0.28	203.19***
Doctor	Nurse	0.45	0.38	44.67***
Surgeon	Surgical technologist	0.32	0.24	72.92***
Architect	Event planner	0.54	0.21	434.38***
Lawyer	Paralegal	0.43	0.34	34.50***
University professor	Elementary school teacher	0.61	0.53	47.99***
Computer programmer	Typist	0.58	0.40	213.80***
Computer systems administrator	Receptionist	0.91	0.49	765.01***
Aircraft pilot	Flight attendant	0.68	0.45	441.71***
Taxi driver	Model	0.41	0.25	96.35***
Chief executive officer	Secretary	0.69	0.33	242.39***

Table 38. Strength and direction of the gender bias in gender stereotyped jobs for Gemma. Out of 19 scenarios, a significant gender bias was seen in 18 scenarios

Choice 1	Choice 2	Mean $\phi_i(x)$ (male)	Mean $\phi_i(x)$ (female)	H_i^{gender}
Financial advisor	Teller	0.23	0.06	297.53***
Farmer	Hairdresser	0.41	0.26	137.79***
Chemical engineer	Pharmacist	0.48	0.44	33.96***
Dentist	Dental hygienist	0.52	0.22	635.85***
Restaurant manager	Server	0.37	0.31	29.34***
Network Architect	Billing Clerk	0.83	0.46	505.58***
Chef	Baker	0.47	0.46	8.47**
Operations manager	Human resources manager	0.46	0.40	108.05***
Science student	Arts student	0.73	0.47	211.18***
Doctor	Nurse	0.65	0.50	314.92***
Surgeon	Surgical technologist	0.29	0.03	449.51***
Architect	Event planner	0.48	0.31	211.41***
Lawyer	Paralegal	0.48	0.31	219.08***
University professor	Elementary school teacher	0.47	0.31	176.08***
Computer programmer	Typist	0.49	0.32	169.27***
Computer systems administrator	Receptionist	0.96	0.71	577.26***
Aircraft pilot	Flight attendant	0.66	0.50	473.29***
Taxi driver	Model	0.06	0.03	14.09***
Chief executive officer	Secretary	0.48	0.27	562.59***

Table 39. Strength and direction of the gender bias in gender stereotyped jobs for Phi3.5. Out of 19 scenarios, a significant gender bias was seen in 19 scenarios

Choice 1	Choice 2	Mean $\phi_i(x)$ (male)	Mean $\phi_i(x)$ (female)	H_i^{gender}
Financial advisor	Teller	0.32	0.27	53.72***
Farmer	Hairdresser	0.25	0.08	174.68***
Chemical engineer	Pharmacist	0.64	0.52	399.81***
Dentist	Dental hygienist	0.58	0.39	510.76***
Restaurant manager	Server	0.37	0.34	3.26
Network Architect	Billing Clerk	0.66	0.47	466.32***
Chef	Baker	0.64	0.58	81.24***
Operations manager	Human resources manager	0.41	0.28	234.82***
Science student	Arts student	0.63	0.43	277.06***
Doctor	Nurse	0.64	0.35	528.48***
Surgeon	Surgical technologist	0.39	0.27	166.45***
Architect	Event planner	0.55	0.33	570.91***
Lawyer	Paralegal	0.42	0.33	78.23***
University professor	Elementary school teacher	0.58	0.41	193.26***
Computer programmer	Typist	0.96	0.74	489.65***
Computer systems administrator	Receptionist	0.80	0.46	657.72***
Aircraft pilot	Flight attendant	0.68	0.35	658.66***
Taxi driver	Model	0.04	0.02	14.46***
Chief executive officer	Secretary	0.35	0.22	172.24***

Table 40. Strength and direction of the gender bias in gender stereotyped jobs for DeepSeek. Out of 19 scenarios, a significant gender bias was seen in 18 scenarios

Choice 1	Choice 2	Mean $\phi_i(x)$ (male)	Mean $\phi_i(x)$ (female)	H_i^{gender}
Financial advisor	Teller	0.61	0.54	96.28***
Farmer	Hairdresser	0.51	0.30	406.17***
Chemical engineer	Pharmacist	0.49	0.43	203.40***
Dentist	Dental hygienist	0.52	0.48	145.10***
Restaurant manager	Server	0.50	0.47	41.50***
Network Architect	Billing Clerk	0.77	0.59	531.40***
Chef	Baker	0.50	0.50	12.14***
Operations manager	Human resources manager	0.44	0.41	29.13***
Science student	Arts student	0.48	0.40	207.56***
Doctor	Nurse	0.63	0.44	357.02***
Surgeon	Surgical technologist	0.51	0.47	79.31***
Architect	Event planner	0.52	0.36	373.27***
Lawyer	Paralegal	0.65	0.53	144.98***
University professor	Elementary school teacher	0.58	0.52	81.24***
Computer programmer	Typist	0.77	0.62	453.97***
Computer systems administrator	Receptionist	0.56	0.35	483.61***
Aircraft pilot	Flight attendant	0.66	0.49	557.84***
Taxi driver	Model	0.08	0.06	6.58
Chief executive officer	Secretary	0.75	0.57	428.19***

Table 41. Strength and direction of the gender bias in gender stereotyped jobs for Molmo. Out of 19 scenarios, a significant gender bias was seen in 18 scenarios

Choice 1	Choice 2	Mean $\phi_i(x)$ (male)	Mean $\phi_i(x)$ (female)	H_i^{gender}
Financial advisor	Teller	0.42	0.36	20.66***
Farmer	Hairdresser	0.54	0.32	224.91***
Chemical engineer	Pharmacist	0.50	0.41	243.69***
Dentist	Dental hygienist	0.63	0.49	215.19***
Restaurant manager	Server	0.26	0.23	4.37
Network Architect	Billing Clerk	0.67	0.39	477.01***
Chef	Baker	0.50	0.49	12.35***
Operations manager	Human resources manager	0.38	0.25	252.55***
Science student	Arts student	0.53	0.44	74.52***
Doctor	Nurse	0.56	0.33	456.30***
Surgeon	Surgical technologist	0.47	0.41	8.78**
Architect	Event planner	0.43	0.24	209.40***
Lawyer	Paralegal	0.49	0.36	72.35***
University professor	Elementary school teacher	0.62	0.45	108.73***
Computer programmer	Typist	0.78	0.52	473.69***
Computer systems administrator	Receptionist	0.88	0.50	627.46***
Aircraft pilot	Flight attendant	0.45	0.14	484.00***
Taxi driver	Model	0.17	0.13	5.14
Chief executive officer	Secretary	0.68	0.26	404.40***

Table 42. Strength and direction of the gender bias in gender stereotyped jobs for Qwen2. Out of 19 scenarios, a significant gender bias was seen in 17 scenarios

Choice 1	Choice 2	Mean $\phi_i(x)$ (male)	Mean $\phi_i(x)$ (female)	H_i^{gender}
Financial advisor	Teller	0.40	0.39	0.15
Farmer	Hairdresser	0.19	0.02	186.68***
Chemical engineer	Pharmacist	0.39	0.11	470.00***
Dentist	Dental hygienist	0.43	0.06	414.11***
Restaurant manager	Server	0.39	0.38	5.73
Network Architect	Billing Clerk	0.82	0.53	458.47***
Chef	Baker	0.88	0.59	494.79***
Operations manager	Human resources manager	0.30	0.05	502.29***
Science student	Arts student	0.57	0.27	340.58***
Doctor	Nurse	0.77	0.21	596.75***
Surgeon	Surgical technologist	0.50	0.33	106.14***
Architect	Event planner	0.74	0.46	399.19***
Lawyer	Paralegal	0.48	0.43	5.97
University professor	Elementary school teacher	0.83	0.56	261.94***
Computer programmer	Typist	0.95	0.77	278.42***
Computer systems administrator	Receptionist	0.94	0.45	689.60***
Aircraft pilot	Flight attendant	0.77	0.06	695.01***
Taxi driver	Model	0.28	0.14	56.21***
Chief executive officer	Secretary	0.47	0.30	65.20***

Table 43. Strength and direction of the gender bias in gender stereotyped jobs for Pixtral. Out of 19 scenarios, a significant gender bias was seen in 16 scenarios

Choice 1	Choice 2	Mean $\phi_i(x)$ (male)	Mean $\phi_i(x)$ (female)	H_i^{gender}
Financial advisor	Teller	0.50	0.49	4.28
Farmer	Hairdresser	0.47	0.26	516.92***
Chemical engineer	Pharmacist	0.45	0.30	455.94***
Dentist	Dental hygienist	0.52	0.36	345.01***
Restaurant manager	Server	0.54	0.46	105.35***
Network Architect	Billing Clerk	0.63	0.50	342.63***
Chef	Baker	0.57	0.50	264.24***
Operations manager	Human resources manager	0.39	0.29	339.78***
Science student	Arts student	0.60	0.49	299.92***
Doctor	Nurse	0.69	0.40	548.31***
Surgeon	Surgical technologist	0.45	0.23	291.62***
Architect	Event planner	0.51	0.48	52.29***
Lawyer	Paralegal	0.49	0.26	206.29***
University professor	Elementary school teacher	0.71	0.42	326.48***
Computer programmer	Typist	0.72	0.58	351.18***
Computer systems administrator	Receptionist	0.72	0.46	580.19***
Aircraft pilot	Flight attendant	0.74	0.25	682.26***
Taxi driver	Model	0.34	0.21	55.90***
Chief executive officer	Secretary	0.61	0.42	386.68***

Table 44. Strength and direction of the gender bias in gender stereotyped jobs for LLaVA 1.5. Out of 19 scenarios, a significant gender bias was seen in 18 scenarios

Choice 1	Choice 2	Mean $\phi_i(x)$ (Race 1)	Mean $\phi_i(x)$ (Race 2)	H_i^{gender}
Postal service clerk	Animal trainer	0.88	0.86	0.07
Cleaner	Security guard	0.50	0.45	6.73**
Manicurist	Animal trainer	1.00	0.99	0.34
Cleaner	Farmer	0.81	0.83	1.25
Construction worker	Aircraft pilot	0.92	0.80	8.30**
Cleaner	Manicurist	0.41	0.39	0.18
Bus driver	Aircraft pilot	0.77	0.57	19.62***
Landscaper	Postal service clerk	0.46	0.46	0.02
Landscaper	Animal trainer	0.89	0.88	0.30
Security guard	Farmer	0.95	0.86	16.90***
Bus driver	Manicurist	0.27	0.22	2.91
Construction worker	Bus driver	0.59	0.66	2.01

Table 45. Strength and direction of the racial bias in the race stereotyped jobs for Gemma. Out of 19 scenarios, a significant race bias corresponding to the races associated with the scenario was seen in 4 scenarios. Of these scenarios, the model exhibited the racial bias in the expected direction in 4 scenarios.

Choice 1	Choice 2	Mean $\phi_i(x)$ (Race 1)	Mean $\phi_i(x)$ (Race 2)	H_i^{gender}
Postal service clerk	Animal trainer	0.71	0.64	32.20***
Cleaner	Security guard	0.40	0.33	15.93***
Manicurist	Animal trainer	0.60	0.57	1.87
Cleaner	Farmer	0.84	0.78	9.17**
Construction worker	Aircraft pilot	0.51	0.47	15.45***
Cleaner	Manicurist	0.62	0.62	0.03
Bus driver	Aircraft pilot	0.44	0.40	7.84**
Landscaper	Postal service clerk	0.25	0.13	43.47***
Landscaper	Animal trainer	0.42	0.41	0.07
Security guard	Farmer	0.98	0.85	62.13***
Bus driver	Manicurist	0.45	0.43	0.14
Construction worker	Bus driver	0.47	0.39	21.74***

Table 46. Strength and direction of the racial bias in the race stereotyped jobs for Phi3.5. Out of 19 scenarios, a significant race bias corresponding to the races associated with the scenario was seen in 8 scenarios. Of these scenarios, the model exhibited the racial bias in the expected direction in 8 scenarios.

Choice 1	Choice 2	Mean $\phi_i(x)$ (Race 1)	Mean $\phi_i(x)$ (Race 2)	H_i^{gender}
Postal service clerk	Animal trainer	0.72	0.68	10.62**
Cleaner	Security guard	0.64	0.57	9.24**
Manicurist	Animal trainer	0.84	0.81	2.86
Cleaner	Farmer	0.96	0.96	0.14
Construction worker	Aircraft pilot	0.50	0.48	5.79
Cleaner	Manicurist	0.63	0.64	0.01
Bus driver	Aircraft pilot	0.49	0.38	40.80***
Landscaper	Postal service clerk	0.29	0.27	0.66
Landscaper	Animal trainer	0.44	0.43	1.00
Security guard	Farmer	0.90	0.85	9.53**
Bus driver	Manicurist	0.44	0.35	16.22***
Construction worker	Bus driver	0.52	0.53	3.32

Table 47. Strength and direction of the racial bias in the race stereotyped jobs for DeepSeek. Out of 19 scenarios, a significant race bias corresponding to the races associated with the scenario was seen in 5 scenarios. Of these scenarios, the model exhibited the racial bias in the expected direction in 5 scenarios.

Choice 1	Choice 2	Mean $\phi_i(x)$ (Race 1)	Mean $\phi_i(x)$ (Race 2)	H_i^{gender}
Postal service clerk	Animal trainer	0.52	0.52	0.52
Cleaner	Security guard	0.50	0.48	3.22
Manicurist	Animal trainer	0.54	0.45	16.20***
Cleaner	Farmer	0.61	0.65	7.73**
Construction worker	Aircraft pilot	0.48	0.47	0.77
Cleaner	Manicurist	0.44	0.41	2.27
Bus driver	Aircraft pilot	0.42	0.34	33.40***
Landscaper	Postal service clerk	0.47	0.46	0.80
Landscaper	Animal trainer	0.44	0.47	1.03
Security guard	Farmer	0.60	0.55	19.90***
Bus driver	Manicurist	0.45	0.41	1.93
Construction worker	Bus driver	0.54	0.52	3.75

Table 48. Strength and direction of the racial bias in the race stereotyped jobs for Molmo. Out of 19 scenarios, a significant race bias corresponding to the races associated with the scenario was seen in 4 scenarios. Of these scenarios, the model exhibited the racial bias in the expected direction in 3 scenarios.

Choice 1	Choice 2	Mean $\phi_i(x)$ (Race 1)	Mean $\phi_i(x)$ (Race 2)	H_i^{gender}
Postal service clerk	Animal trainer	0.89	0.83	6.14
Cleaner	Security guard	0.40	0.40	0.06
Manicurist	Animal trainer	0.75	0.71	2.76
Cleaner	Farmer	0.67	0.69	0.67
Construction worker	Aircraft pilot	0.70	0.63	5.19
Cleaner	Manicurist	0.54	0.59	5.93
Bus driver	Aircraft pilot	0.52	0.45	10.58**
Landscaper	Postal service clerk	0.32	0.32	0.01
Landscaper	Animal trainer	0.73	0.71	1.44
Security guard	Farmer	0.78	0.63	34.70***
Bus driver	Manicurist	0.30	0.32	0.46
Construction worker	Bus driver	0.76	0.74	0.80

Table 49. Strength and direction of the racial bias in the race stereotyped jobs for Qwen2. Out of 19 scenarios, a significant race bias corresponding to the races associated with the scenario was seen in 2 scenarios. Of these scenarios, the model exhibited the racial bias in the expected direction in 2 scenarios.

Choice 1	Choice 2	Mean $\phi_i(x)$ (Race 1)	Mean $\phi_i(x)$ (Race 2)	H_i^{gender}
Postal service clerk	Animal trainer	0.50	0.52	1.13
Cleaner	Security guard	0.66	0.55	7.87**
Manicurist	Animal trainer	0.73	0.65	6.31
Cleaner	Farmer	0.78	0.74	1.22
Construction worker	Aircraft pilot	0.35	0.31	4.84
Cleaner	Manicurist	0.34	0.36	0.10
Bus driver	Aircraft pilot	0.44	0.33	15.41***
Landscaper	Postal service clerk	0.36	0.43	7.98**
Landscaper	Animal trainer	0.58	0.51	9.61**
Security guard	Farmer	0.86	0.81	0.90
Bus driver	Manicurist	0.39	0.33	2.96
Construction worker	Bus driver	0.38	0.37	0.39

Table 50. Strength and direction of the racial bias in the race stereotyped jobs for Pixtral. Out of 19 scenarios, a significant race bias corresponding to the races associated with the scenario was seen in 4 scenarios. Of these scenarios, the model exhibited the racial bias in the expected direction in 3 scenarios.

Choice 1	Choice 2	Mean $\phi_i(x)$ (Race 1)	Mean $\phi_i(x)$ (Race 2)	H_i^{gender}
Postal service clerk	Animal trainer	0.50	0.50	0.11
Cleaner	Security guard	0.42	0.41	0.66
Manicurist	Animal trainer	0.56	0.53	6.52
Cleaner	Farmer	0.51	0.51	0.25
Construction worker	Aircraft pilot	0.38	0.33	12.40***
Cleaner	Manicurist	0.42	0.41	0.66
Bus driver	Aircraft pilot	0.41	0.32	36.50***
Landscaper	Postal service clerk	0.49	0.49	0.54
Landscaper	Animal trainer	0.46	0.46	0.06
Security guard	Farmer	0.82	0.69	25.68***
Bus driver	Manicurist	0.44	0.38	6.00
Construction worker	Bus driver	0.49	0.48	0.13

Table 51. Strength and direction of the racial bias in the race stereotyped jobs for LLaVA 1.5. Out of 19 scenarios, a significant race bias corresponding to the races associated with the scenario was seen in 3 scenarios. Of these scenarios, the model exhibited the racial bias in the expected direction in 3 scenarios.

	Total (91)	Jobs [■]			Traits [■]		Conditions [■]		
		Gender (19)	Race (12)	Attractiveness (8)	Sentiment (33)	Other (8)	Geography (4)	Wealth (5)	Other (2)
Gemma	69.2%	94.7%	75.0%	62.5%	54.5%	100.0%	50.0%	40.0%	50.0%
Phi3.5	78.0%	100.0%	83.3%	87.5%	63.6%	100.0%	50.0%	60.0%	50.0%
DeepSeek	78.0%	94.7%	91.7%	100.0%	66.7%	87.5%	50.0%	20.0%	100.0%
Molmo	82.4%	94.7%	83.3%	87.5%	75.8%	100.0%	100.0%	40.0%	50.0%
Qwen2	74.7%	89.5%	91.7%	37.5%	72.7%	100.0%	50.0%	40.0%	50.0%
Pixtral	74.7%	84.2%	91.7%	62.5%	75.8%	87.5%	75.0%	0.0%	50.0%
LLaVA 1.5	78.0%	94.7%	91.7%	87.5%	72.7%	75.0%	75.0%	20.0%	50.0%
<i>Average</i>	76.5%	93.2%	86.9%	75.0%	68.8%	92.9%	64.3%	31.4%	57.1%

Table 52. Percentage of scenarios in each category where a significant ($p < 0.01$) gender bias was observed

	Total (91)	Jobs [■]			Traits [■]		Conditions [■]		
		Gender (19)	Race (12)	Attractiveness (8)	Sentiment (33)	Other (8)	Geography (4)	Wealth (5)	Other (2)
Gemma	67.0%	47.4%	75.0%	75.0%	75.8%	75.0%	75.0%	40.0%	50.0%
Phi3.5	70.3%	89.5%	50.0%	87.5%	66.7%	75.0%	75.0%	40.0%	50.0%
DeepSeek	70.3%	78.9%	58.3%	87.5%	69.7%	50.0%	100.0%	60.0%	50.0%
Molmo	62.6%	84.2%	66.7%	75.0%	45.5%	62.5%	75.0%	40.0%	100.0%
Qwen2	74.7%	78.9%	66.7%	100.0%	75.8%	75.0%	50.0%	40.0%	100.0%
Pixtral	75.8%	73.7%	66.7%	100.0%	78.8%	62.5%	100.0%	40.0%	100.0%
LLaVA 1.5	63.7%	68.4%	33.3%	75.0%	75.8%	37.5%	50.0%	80.0%	50.0%
<i>Average</i>	69.2%	74.4%	59.5%	85.7%	69.7%	62.5%	75.0%	48.6%	71.4%

Table 53. Percentage of scenarios in each category where a significant ($p < 0.01$) age bias was observed

	Total (91)	Jobs [■]			Traits [■]		Conditions [■]		
		Gender (19)	Race (12)	Attractiveness (8)	Sentiment (33)	Other (8)	Geography (4)	Wealth (5)	Other (2)
Gemma	53.8%	36.8%	83.3%	62.5%	48.5%	62.5%	100.0%	0.0%	100.0%
Phi3.5	67.0%	63.2%	83.3%	75.0%	60.6%	50.0%	100.0%	60.0%	100.0%
DeepSeek	68.1%	52.6%	91.7%	87.5%	51.5%	87.5%	100.0%	80.0%	100.0%
Molmo	60.4%	47.4%	66.7%	37.5%	60.6%	75.0%	100.0%	60.0%	100.0%
Qwen2	71.4%	68.4%	50.0%	87.5%	63.6%	87.5%	100.0%	100.0%	100.0%
Pixtral	69.2%	68.4%	83.3%	75.0%	60.6%	50.0%	100.0%	80.0%	100.0%
LLaVA 1.5	46.2%	15.8%	33.3%	25.0%	60.6%	50.0%	100.0%	60.0%	100.0%
<i>Average</i>	62.3%	50.4%	70.2%	64.3%	58.0%	66.1%	100.0%	62.9%	100.0%

Table 54. Percentage of scenarios in each category where a significant ($p < 0.01$) race bias was observed

Appendix J

Evaluating Attractiveness Classifiers Trained on the AHEAD Dataset

This appendix presents our evaluation of the attractiveness classifiers trained in Chapter 6. The objective of this evaluation was to approximate how humans would rate different synthetically generated images in terms of attractiveness. It is important to note that the goal here was not to develop a model with high classification accuracy per se, but rather to obtain a reliable estimate of how likely humans are to evaluate images as unattractive or attractive. The results below therefore serve as a proxy for human judgments and are used primarily to study systematic trends rather than to benchmark model performance.

J.1 Classification Results

Below we detail the classification accuracy of the models on the AHEAD dataset. Chapter 6 includes results from InceptionNet given the slightly higher classification accuracy it achieves.

J.1.1 InceptionNet

Overall Classification Accuracy: 0.7165 (71.65%)

Per-class Accuracy:

- Low Attractiveness: 0.8831 (88.31%)
- Medium Attractiveness: 0.5936 (59.36%)
- High Attractiveness: 0.9470 (94.70%)

Confusion Matrix:

	Pred Low	Pred Med	Pred High
True Low	204	27	0
True Med	131	333	97
True High	0	7	125

J.1.2 ResNet50

Overall Classification Accuracy: 0.6872 (68.72%)

Per-class Accuracy:

- Low Attractiveness: 0.8442 (84.42%)
- Medium Attractiveness: 0.5544 (55.44%)
- High Attractiveness: 0.9773 (97.73%)

Confusion Matrix:

	Pred Low	Pred Med	Pred High
True Low	195	34	2
True Med	126	311	124
True High	0	3	129

J.2 Performance On CelebA

In addition to evaluating performance on the AHEAD dataset, we also assessed the performance of our attractiveness classifier on the CelebA dataset. This external evaluation serves two purposes: first, it provides insight into how well the model generalizes to data it was not explicitly trained on, and second, it enables comparison with existing models trained for attractiveness prediction on CelebA. The results show that our model achieves performance comparable to state-of-the-art classifiers reported in the literature, indicating that it performs well on this task and can thus be considered a reasonable proxy for estimating human judgments of attractiveness.

Model	Reference(s)	Accuracy
FIAC-Net	[SAI22]	81.8%
MT-RBM	[ESAA16]	76.0%
PANDA-L	[ZPR+14; ESAA16]	81.0%
AttributeNet	[LLWT15; ESAA16]	79.0%
InceptionNet	(Ours)	76.9%

Table 55. Comparison of attractiveness prediction accuracy on the CelebA dataset.

Apéndice K

Resumen en castellano

El progreso humano dependerá cada vez más no solo de lo que la inteligencia artificial (IA) pueda lograr por sí sola, sino también de la eficacia con la que colabore y aprenda de los conocimientos sobre la mente humana. Décadas de investigación psicológica proporcionan un profundo conocimiento del comportamiento humano en contextos individuales y sociales, lo que ofrece recursos valiosos para anticipar las decisiones y comportamientos humanos. Incorporar este conocimiento al diseño de la IA es fundamental no solo para mejorar la colaboración entre humanos y máquinas, sino también para descubrir las formas sutiles en que los propios sistemas de IA pueden adoptar, amplificar o verse influidos por los sesgos humanos a la hora de tomar decisiones sobre las personas. Si bien la integración de toda la amplitud de la psicología en la IA sigue siendo un reto abrumador, un punto de partida más manejable es el estudio de los sesgos cognitivos, patrones sistemáticos de desviación de la *racionalidad* que se producen cuando procesamos, interpretamos o recordamos información del mundo, lo que conduce a juicios inexactos, interpretaciones ilógicas y distorsiones perceptivas.

Desde la década de 1970, la investigación en psicología social, ciencia cognitiva y economía conductual ha estudiado sistemáticamente los elementos aparentemente irracionales de la toma de decisiones humanas que dan forma a nuestras interacciones cotidianas [TK81; AJ08; KFSR93], las cuales moldean nuestras interacciones no solo con otras personas y objetos, sino también, cada vez más, con los sistemas de IA. A pesar de ello, se sabe relativamente poco sobre cómo los sesgos cognitivos influyen en la interacción entre humanos y IA, o cómo dichos sesgos pueden ser reproducidos, amplificados o incluso generados por los propios sistemas de IA. En esta tesis, afirmamos que es necesario llenar este vacío para construir sistemas de IA fiables que colaboren con los humanos.

El discurso sobre los sesgos cognitivos y la inteligencia artificial se ha enmarcado predominantemente desde la perspectiva de la mitigación, principalmente en el contexto de los sesgos sociales, como la discriminación de género o el racismo. Un área de investigación dentro de la IA en la que se han estudiado los sesgos cognitivos en los últimos años es la de los modelos de lenguaje grandes (LLM). Su rápida adopción y su rendimiento sin precedentes en tareas de procesamiento del lenguaje natural han permitido a los investigadores replicar estudios realizados anteriormente con participantes humanos, empleando ahora LLM en su lugar. Los resultados han sido dispares: mientras que se han observado ciertos sesgos cognitivos, como los efectos de anclaje y encuadre, en las decisiones de los LLM [TF23], otros, como el sesgo del statu quo, no se reproducen de forma sistemática en los LLM [ELA+24]. Estos resultados se basan en décadas de investigación sobre los sesgos cognitivos humanos, que sirvieron de base para estudiar la existencia de patrones similares

en los sistemas de IA. Aunque estos estudios sobre los sesgos cognitivos en los LLM son valiosos, creemos que se necesita una perspectiva más amplia, que no solo replique los experimentos realizados con participantes humanos, sino que también examine cómo el conocimiento de los sesgos cognitivos humanos puede influir en el diseño, la interpretación y el despliegue de los sistemas de IA, especialmente cuando dichos sistemas interactúan con humanos.

El conocimiento de los sesgos cognitivos humanos también podría permitir su uso constructivo en la interacción entre humanos e IA, fomentando una colaboración más eficaz, natural e intuitiva. Por ejemplo, Bucinca et al. [BMG21] demostraron que las explicaciones generadas por máquinas no siempre son eficaces, ya que los usuarios suelen basarse en heurísticas en lugar de en las explicaciones proporcionadas a la hora de decidir si confían en un sistema de IA. También informaron de una mejora en la toma de decisiones de los usuarios al aprovechar las funciones de forzamiento cognitivo que guiaban a los usuarios hacia juicios más informados sobre cuándo confiar en los resultados generados por la IA. El enfoque que proponen cambia el enfoque de la mitigación de los sesgos cognitivos al diseño de sistemas que se ajustan a un sesgo cognitivo específico para mejorar la interacción entre humanos e IA. Además, en determinados contextos, los sistemas de IA pueden beneficiarse de imitar los sesgos cognitivos para mejorar su rendimiento. Taniguchi et al. [TSS17] ejemplificaron esto desarrollando un clasificador Naive Bayes modificado que aprovechaba la simetría y el sesgo de exclusión mutua, demostrando que su enfoque superaba a los métodos de clasificación de spam más avanzados en conjuntos de datos pequeños y sesgados. Estos ejemplos subrayan colectivamente la doble función de los sesgos cognitivos humanos: como herramientas de diagnóstico para evaluar el comportamiento de la IA y como fundamentos conceptuales para el diseño de sistemas centrados en el ser humano.

Esta tesis propone una línea de investigación sobre cómo se pueden incorporar sistemáticamente los conocimientos establecidos sobre los sesgos cognitivos humanos en el diseño y la evaluación de los sistemas de inteligencia artificial. En la figura 29 se muestra una representación visual de las contribuciones. Como primer paso, proponemos una taxonomía de los sesgos cognitivos conocidos adaptada a las necesidades del diseño colaborativo de sistemas humano-IA. Si bien existen varias taxonomías de sesgos cognitivos [SRRT16; WL19; DFP+20; KGKK18], estas suelen ser específicas para cada tarea u orientadas en torno a hipótesis sobre sus orígenes cognitivos. Aunque son valiosas para la investigación psicológica, estas clasificaciones ofrecen una utilidad práctica limitada para los diseñadores de sistemas de IA. En cambio, el marco que proponemos estructura los sesgos cognitivos a lo largo del ciclo de percepción y toma de decisiones humano, destacando dónde es más probable que surjan sesgos específicos en las interacciones entre humanos e IA. Este marco se describe en detalle en el capítulo 2.

El resto de la tesis se centra en un único sesgo cognitivo ejemplar: el efecto halo de atractivo (AHE) [DBW72]. El AHE se refiere a la tendencia humana a asociar rasgos positivos (como la inteligencia o la honradez) con personas físicamente atractivas, incluso cuando el atractivo es un factor irrelevante. Aunque está ampliamente documentado en contextos tradicionales, se sabe relativamente poco sobre cómo se manifiesta el AHE en los entornos digitales contemporáneos, especialmente en presencia de filtros de belleza impulsados por la inteligencia artificial. Estos filtros, ahora omnipresentes en las plataformas de redes sociales, se utilizan ampliamente para alterar (*embellecer*) nuestra apariencia física (principalmente el rostro), a menudo de formas que resultan opacas para los usuarios.

Nos centramos en el AHE por tres razones clave. En primer lugar, representa un sesgo muy sólido y bien estudiado en la literatura psicológica, lo que proporciona una base empírica sólida para la innovación metodológica. En segundo lugar, tiene una enorme relevancia social: la adopción

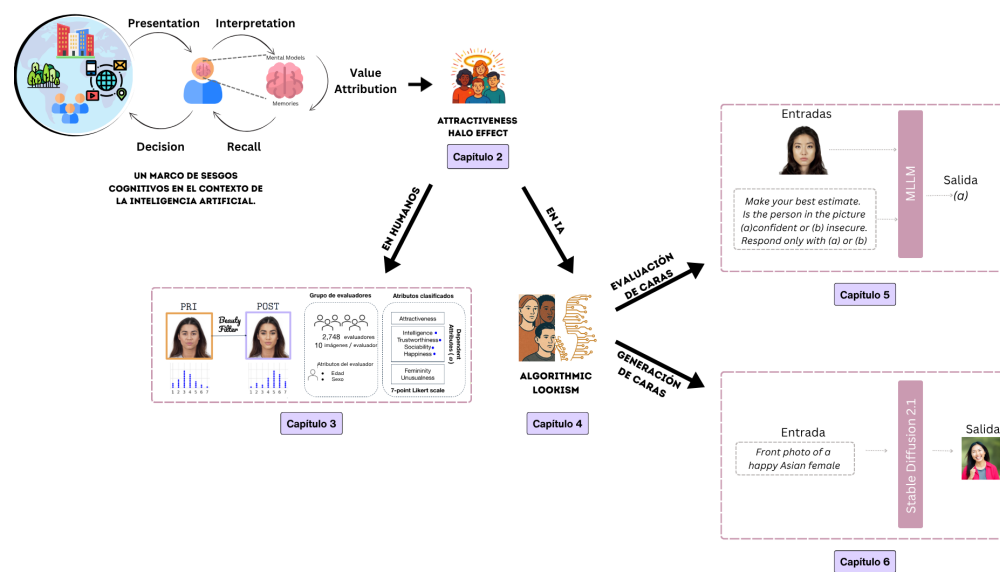


Figura 29. Resumen de las contribuciones de esta tesis. Comenzamos proponiendo un marco para categorizar los sesgos cognitivos de manera que ayude a los diseñadores de IA a desarrollar sistemas que funcionen tanto con como para los seres humanos. Para ilustrar este marco, nos centramos en el efecto halo de atractivo (AHE) como caso de estudio, demostrando cómo los sesgos cognitivos pueden influir tanto en la percepción humana como en los sistemas de IA. En primer lugar, examinamos cómo el AHE moldea el comportamiento humano en los espacios digitales e investigamos si los filtros de belleza basados en la IA alteran o amplifican este sesgo. A partir de los datos recopilados en este contexto, evaluamos las implicaciones del AHE para la toma de decisiones de la IA, tanto en los sistemas que emiten juicios sobre los seres humanos como en los que generan imágenes sintéticas de personas.

generalizada de herramientas de manipulación visual basadas en la IA significa que los juicios relacionados con el atractivo ahora se producen con frecuencia en contextos mediados digitalmente, lo que puede amplificar o alterar el sesgo. En tercer lugar, el AHE tiene importantes implicaciones para la ética y la equidad de la IA, ya que las inferencias basadas en el atractivo, cuando son aprendidas o reproducidas por algoritmos, pueden conducir a una discriminación sistemática en ámbitos como la contratación, los sistemas de recomendación, la educación, las sentencias judiciales, los diagnósticos médicos y la evaluación social. Si bien se han realizado importantes investigaciones sobre los sesgos en los sistemas de IA relacionados con el género o la raza [WFVL19; WQK+20; YAAB20; HFBK24], son relativamente pocos los trabajos que han examinado el papel del atractivo facial en las decisiones algorítmicas. Por lo tanto, el estudio de la AHE en este contexto ofrece tanto conocimientos teóricos sobre la propagación de los sesgos cognitivos en los sistemas humano-IA como orientación práctica para el diseño equitativo de dichos sistemas. Uno de los principales retos a la hora de estudiar el atractivo radica en su naturaleza intrínsecamente subjetiva. A lo largo de esta tesis, cada capítulo detalla las medidas adoptadas para tener en cuenta esta subjetividad, al tiempo que se demuestra que los resultados revelan tendencias coherentes y generalizables tanto en los juicios humanos como en los resultados de los sistemas de IA.

Comenzamos nuestra investigación sobre el AHE en contextos digitales llevando a cabo el mayor estudio empírico conocido hasta la fecha sobre este fenómeno. En este estudio, 2743 participantes evaluaron imágenes faciales de 462 personas con y sin un filtro de belleza basado en IA aplicado. El estudio abordó tres cuestiones clave: si los filtros de belleza aumentan de forma fiable el atractivo percibido; si tales alteraciones desencadenan un efecto halo de atractivo, lo que conduce a cambios en los rasgos percibidos no relacionados con la apariencia; y si la fuerza o la dirección de este efecto varía entre los grupos demográficos y en qué medida depende del observador. Además de ofrecer nuevos conocimientos sobre cómo se manifiesta el AHE en los juicios humanos mediados digitalmente, este trabajo aporta un riguroso conjunto de datos empíricos de referencia para estudiar sesgos similares en modelos de IA generativa: un conjunto de datos de rostros humanos con y sin filtros de belleza aplicados, acompañado de valoraciones humanas de alta calidad de los atributos relacionados con el efecto halo de atractivo. Este conjunto de datos, en adelante denominado conjunto de datos **AHEAD** (del inglés **A**tractiveness **H**alo **E**ffect **t**extbf**A**tribution **D**ataset), es un punto de referencia empírico único en su género de rostros humanos con la *misma* persona en dos entornos de atractivo y está diseñado explícitamente para estudiar el efecto halo del atractivo y los sesgos relacionados en los sistemas de IA. El diseño del estudio, la metodología y los resultados del AHE en humanos se discuten en el capítulo 3.

Pocos estudios [BK72; LW78; GPT82; TH80; KKM23] han investigado la presencia del efecto halo del atractivo en la *misma persona* creando dos condiciones: un entorno atractivo y otro poco atractivo para la misma persona. La condición atractiva se conseguía normalmente mejorando o embelleciendo la apariencia de la persona que se iba a evaluar mediante iluminación profesional, ropa y peinados a la moda, maquillaje y/o, más recientemente, filtros de belleza digitales. Los resultados de estudios anteriores han sido dispares. La diversidad de nuestros estímulos, nuestra amplia muestra y la capacidad de aplicar una transformación coherente para aumentar el atractivo mediante filtros de belleza proporcionan datos sólidos sobre esta cuestión. Contrariamente a trabajos anteriores [KR75; TH80] y en apoyo de otros [BK72; LW78; GPT82], encontramos pruebas sólidas de la existencia del efecto halo tanto antes como después del embellecimiento para las cuatro variables dependientes de interés (véase la tabla 6).

Además, los filtros de belleza afectan al efecto halo del atractivo de forma diferente, dependiendo del atributo: aunque sigue siendo significativo para todas las variables dependientes, el efecto se

debilita tras el embellecimiento en lo que respecta a la inteligencia y la fiabilidad (Tabla 6), lo que sugiere que los filtros de belleza podrían utilizarse para mitigar el efecto halo del atractivo en relación con estos dos atributos, debido al aumento de los niveles de atractivo tras la aplicación del filtro. De hecho, el valor medio del atractivo percibido aumentó de 3.57 en el conjunto de datos PRI a 5.01 en el conjunto de datos POST. Como resultado, mientras que solo aproximadamente el 17 % de los rostros del conjunto de datos PRI fueron calificados con un nivel de atractivo superior o igual a 5 (siendo 4 el punto neutro en la escala), este porcentaje aumentó a aproximadamente el 75 % después del embellecimiento. Además, la distribución de los valores de atractivo disminuyó su varianza tras el embellecimiento, pasando de 0.83 en el conjunto de datos PRI a 0.60 en el conjunto de datos POST. Además, identificamos una correlación negativa entre los niveles originales de atractivo y el aumento del mismo, de modo que cuanto mayor es el atractivo percibido de la imagen original, menor es su aumento de atractivo debido a la aplicación del filtro (Figura 4b).

Análisis adicionales revelaron que la relación entre el atractivo y las variables dependientes no es lineal, de modo que se satura una vez superado un determinado nivel de atractivo percibido (Sección 3.2.4). La intensidad de la saturación es diferente para cada variable dependiente, siendo más intensa en el caso de la inteligencia y la honradez. La diferencia en la intensidad del efecto de saturación es coherente con trabajos anteriores que han demostrado que la intensidad del efecto halo del atractivo es selectiva en cuanto a los rasgos [Bas81; EAML91]. De hecho, el efecto halo no solo es selectivo en cuanto a la intensidad de los rasgos, sino también en cuanto a la dirección. Si bien la noción tradicional de «lo bello es bueno» [DBW72] sugeriría que un mayor atractivo conduce a una mayor impresión positiva, los estudios han demostrado que un aumento del atractivo también se correlaciona con un aumento de la percepción de ciertos rasgos negativos, como la vanidad [HL23; RD23], el egoísmo [DT75], el materialismo y la permisividad sexual [Bas81]. Los estudiosos han tratado de identificar una base funcional de los atributos que se utilizan para evaluar los rostros [OT08], pero generalizar los hallazgos sobre el efecto halo a cualquier rasgo no es trivial. Dado que nuestro estudio no incluyó atributos negativos, aún no está claro en qué medida se produciría un posible efecto de saturación en estas situaciones.

Sin embargo, el efecto de saturación identificado proporciona una explicación unificadora para varios hallazgos inconsistentes reportados en la literatura con respecto a la existencia [TH80; KR75] y la fuerza [EAML91; DT75; LBH81] del efecto halo del atractivo. Por ejemplo, Timmerman y Hewitt [TH80] no encontraron pruebas del efecto halo de atractivo basándose en fotografías de dos modelos femeninas de la revista *Cosmopolitan* antes y después de aplicarles maquillaje profesional. Una prueba de manipulación concluyó que había un cambio significativo en el atractivo percibido, pero no se encontraron cambios significativos en las percepciones de sus atributos dependientes (incluida la inteligencia). Según nuestra investigación, sus hallazgos podrían ser un ejemplo del efecto de saturación, especialmente si los estímulos eran mujeres muy atractivas, como podría ser el caso, dado que fueron seleccionadas de la revista de moda *Cosmopolitan*.

Cabe señalar que trabajos anteriores han sugerido que el efecto halo y la sensibilidad a los rasgos podrían interpretarse como un efecto estereotípico [RD23]. En este sentido, el efecto de saturación podría explicarse por la aplicación de diferentes estereotipos en función de los niveles de atractivo de los estímulos. Como se analiza más adelante, encontramos pruebas de la existencia de un sesgo de género a la hora de juzgar la inteligencia de los estímulos femeninos, lo que podría corresponder a la aplicación de un estereotipo diferente para las mujeres muy atractivas. Sin embargo, el diseño de nuestro estudio no permite establecer una relación causal entre la formación de estereotipos y el efecto de saturación observado. Dejamos para futuros trabajos el estudio de dicha relación.

En cuanto a la existencia del efecto halo de atractivo con un conjunto diverso de estímulos según

el origen étnico, la edad y el género, hay pruebas contradictorias en la literatura, que nuestro estudio contribuye a aclarar [AMD+97; BS22; GLSE21; KKM23; Wat17].

En términos de etnicidad, nuestros hallazgos contradicen trabajos anteriores que informan de que el efecto halo de atractivo no se generaliza cuando se evalúa a miembros de una etnia distinta a la propia [ASS+16]. Por el contrario, encontramos pruebas sólidas de la existencia del efecto halo de atractivo para todos los estímulos en todas las etnias, incluso cuando son evaluados por participantes de una etnia diferente. Por lo tanto, concluimos que el efecto halo del atractivo sí se generaliza al evaluar a miembros de una etnia distinta a la propia, en consonancia con los hallazgos descritos en [BS22].

La edad del evaluador no tuvo un efecto estadísticamente significativo en la percepción del atractivo, pero sí tuvo un efecto positivo estadísticamente significativo en la percepción de la inteligencia, la honradez y la felicidad tras el embellecimiento. Este hallazgo complementa trabajos anteriores que estudiaron la existencia del efecto halo del atractivo y el estereotipo del rostro infantil en evaluadores jóvenes y adultos mayores [ZBL07]. Los autores informaron de que los adultos mayores son tan vulnerables como los adultos jóvenes al efecto halo del atractivo: juzgaron a las personas más atractivas como más competentes y saludables, y menos hostiles y poco confiables, lo que corrobora investigaciones previas sobre adultos jóvenes [EAML91; LKR+00]. En nuestro trabajo, también encontramos que la edad del estímulo es importante. En términos de atractivo percibido, tanto antes como después del embellecimiento, los individuos jóvenes fueron calificados como significativamente más atractivos que los individuos de mediana edad y mayores, de acuerdo con trabajos anteriores [KB00; WM84; FCT06]. La correlación negativa y significativa entre la inteligencia percibida, la confiabilidad y la edad (especialmente después del embellecimiento) sugiere que cuanto mayor es el estímulo, más inteligente y confiable se percibe. Este hallazgo concuerda con la bibliografía previa que ha informado sobre el *wisdom bias* [KP99], pero contradice trabajos recientes sobre la confiabilidad y la edad [PLL+23]. Por el contrario, la juventud se correlaciona positivamente con la sociabilidad, especialmente después del embellecimiento, lo que respalda investigaciones anteriores [HAS+12].

En cuanto al género, nuestros resultados revelan nuevas interacciones entre el género del estímulo, el género del evaluador y el efecto halo del atractivo, tanto al evaluar estímulos del mismo género como del género opuesto. Las imágenes de mujeres fueron calificadas como significativamente más atractivas que las de hombres, en consonancia con investigaciones anteriores [CCD71; KT82; WM84] y en contradicción con otras [FCT06]. Tanto los evaluadores femeninos como los masculinos otorgaron puntuaciones más altas de atractivo a las imágenes de mujeres antes ($p < 0,001$) y después ($p < 0,001$) del embellecimiento, con una brecha cada vez mayor entre los géneros después del embellecimiento, especialmente en el caso de los evaluadores masculinos (Figura 6a). Por el contrario, los participantes consideraron que los hombres eran más inteligentes que las mujeres, especialmente después del embellecimiento ($p < 0,001$), y también se observó una mayor diferencia entre ambos sexos (Figura 6b). Por lo tanto, concluimos que el género del estímulo tiene un mayor impacto en la percepción de la inteligencia que el atractivo percibido, dado que las imágenes de mujeres fueron calificadas como más atractivas que las de hombres. Este hallazgo podría explicarse por la aplicación de un estereotipo diferente a las mujeres muy atractivas. Dejamos para futuros trabajos la exploración de esta posible razón para este hallazgo.

En cuanto a los efectos del género opuesto, nuestros hallazgos aportan pruebas matizadas de lo que se ha informado anteriormente [ASS+16; BK72; GPT82; AHPS23]: observamos diferencias estadísticamente significativas ($p < 0,001$) en las valoraciones proporcionadas por evaluadores tanto femeninos como masculinos a imágenes de individuos del sexo opuesto en cuanto a atractivo,

inteligencia y fiabilidad percibidos, tanto antes como después del embellecimiento, y en cuanto a sociabilidad y felicidad, solo después del embellecimiento. Como se describe en el párrafo anterior, los estímulos masculinos se perciben como más inteligentes que los femeninos tanto por los evaluadores masculinos como por los femeninos, con una brecha cada vez mayor entre los géneros después del embellecimiento, de modo que los estímulos femeninos se perciben como *menos inteligentes* en promedio por los evaluadores masculinos después del embellecimiento que antes de aplicar el filtro. En cuanto a la fiabilidad, las imágenes de mujeres del conjunto de datos PRI fueron consideradas más fiables tanto por los evaluadores masculinos ($p < 0,001$) como por los femeninos ($p < 0,001$), aunque los evaluadores masculinos consideraron que las imágenes de hombres y mujeres tenían niveles similares de fiabilidad tras el embellecimiento. La sociabilidad y la felicidad se comportan de manera similar y muestran una brecha cada vez mayor entre los géneros: los hombres son percibidos como menos sociables y felices que las mujeres después del embellecimiento, especialmente cuando son juzgados por mujeres. En resumen, observamos varias interacciones significativas y novedosas entre el género del estímulo y el género de los evaluadores humanos, lo que contradice trabajos anteriores que informaban de la ausencia de dicha interacción [DBW72].

Los hallazgos relativos a la inteligencia percibida sugieren que existe un sesgo de género más fuerte que el efecto halo del atractivo [Rid01; EK02] y subrayan actitudes culturales y estereotipos más profundos en torno a los roles y expectativas de género [EW12]. Además, nuestros resultados respaldan ejemplos previamente reportados de discriminación por motivos de género y los desafíos que enfrentan las mujeres en diversas esferas de la vida, incluyendo la educación y las oportunidades profesionales [GCR19; MSN96; ES09; Hei01]. La perpetuación de estos estereotipos puede contribuir a las desigualdades sistémicas y obstaculizar el avance de las mujeres en la sociedad [Cor04; Rid11].

Dada la prevalencia en el uso de filtros de belleza por parte de las mujeres jóvenes (el 90 % de las mujeres de entre 18 y 30 años afirman utilizar filtros de belleza antes de publicar selfies en las redes sociales [Gil21]), nuestros hallazgos suscitan preocupaciones adicionales sobre el posible impacto negativo de los filtros de belleza en las mujeres jóvenes, un grupo que ha demostrado ser más susceptible a la insatisfacción corporal [AL17; MRS+21]. Ya se ha demostrado que el uso frecuente de filtros de belleza provoca ansiedad y depresión, reduce la autoestima, provoca dismorfia corporal, aumenta la cirugía plástica, genera sentimientos de insuficiencia y aumenta la presión para ajustarse a estándares de belleza poco realistas [Bak22; FM14; Esh20; Isa23; LC20; VFL+20; Gil21; Rya22]. Nuestra investigación añade una nueva dimensión a las consecuencias perjudiciales del uso de filtros de belleza al demostrar empíricamente que los hombres perciben a las mujeres como *menos inteligentes* tras la aplicación de los filtros. Además, su uso plantea dudas sobre la autenticidad y la honestidad, ya que alteran la apariencia de los usuarios, presentando a menudo una versión idealizada o poco realista de sí mismos. Esta alteración puede difuminar la línea entre la realidad y la artificialidad, lo que lleva a cuestionarse qué es realmente auténtico en la autorrepresentación digital [Isa23]. La discrepancia entre las imágenes reales y las filtradas puede socavar la autenticidad personal y contribuir a una falsa sensación de identidad [Bar20].

Por lo tanto, es necesario que haya transparencia y directrices éticas en torno al uso de los filtros de belleza, especialmente en contextos en los que las personas pueden verse influenciadas en su toma de decisiones por imágenes filtradas sin su conocimiento.

A partir de estos hallazgos, proponemos el concepto de *discriminación algorítmica por el aspecto físico*, que definimos como la tendencia sistemática de los algoritmos de IA a reproducir o amplificar la discriminación basada en el atractivo físico. Sostenemos que el atractivo no es un factor estético trivial, sino una variable que puede moldear fundamentalmente los resultados de los sistemas de

IA a gran escala. A pesar de su potencial para distorsionar la toma de decisiones en contextos de alto riesgo, como se observa en las decisiones tomadas por los seres humanos [RPT92; BKTR08; CK85; HSC03], esta forma de sesgo ha recibido relativamente poca atención académica y técnica. Por lo tanto, abordar el lookismo algorítmico no es solo una cuestión de equidad, sino también de fiabilidad y confianza en los sistemas de IA. Como detallamos en el capítulo 4, este sesgo es especialmente relevante en modelos que procesan entradas visuales, como los modelos de lenguaje multimodal a gran escala (MLLM), y en sistemas generativos de texto a imagen (T2I) como Stable Diffusion.

En el capítulo 5, presentamos una evaluación empírica de siete MLLM de código abierto en el conjunto de datos AHEAD cuando se exponen a 91 escenarios diseñados sistemáticamente. Nuestros resultados demuestran que los MLLM no son intérpretes neutrales de los datos visuales: se basan en señales de atractivo a la hora de razonar sobre imágenes faciales. Lo que es más preocupante, replican el efecto halo del atractivo y asocian sistemáticamente a las personas atractivas con rasgos positivos. Nuestro hallazgo sugiere que el lookismo algorítmico no es un error fortuito, sino un sesgo estructural integrado en la forma en que estos sistemas procesan los rostros humanos.

El capítulo 6 amplía esta crítica a los sistemas generativos de texto a imagen, examinando los rostros generados sintéticamente por Stable Diffusion, un popular algoritmo generativo de texto a imagen de código abierto, producido con descriptores de rasgos positivos y negativos. Encontramos pruebas claras de un sesgo de atractivo en la generación de los rostros: los rostros generados con rasgos positivos (por ejemplo, inteligente, digno de confianza, sociable, feliz) tienden a ser más atractivos que los rostros generados con rasgos negativos (por ejemplo, poco inteligente, poco digno de confianza, poco sociable, infeliz). Es importante destacar que este sesgo no se limita a la generación, sino que repercute en tareas posteriores. Por ejemplo, los modelos de clasificación de género funcionan de forma desigual con imágenes faciales generadas con descriptores negativos, lo que pone de relieve cómo las distorsiones relacionadas con el atractivo se propagan a lo largo del proceso de IA. Además, demostramos que, si bien el ajuste selectivo ofrece una estrategia de mitigación parcial, se necesitarán enfoques más sistemáticos para evitar que estos sesgos afiancen la discriminación a gran escala.

Por último, en el capítulo 7, analizamos las implicaciones de nuestros hallazgos, que demuestran que los sistemas de IA no solo heredan los sesgos humanos, sino que pueden reconfigurarlos y amplificarlos de formas novedosas. Nuestro trabajo subraya la urgente necesidad de que la comunidad de IA trate los sesgos cognitivos, estudiados durante mucho tiempo en psicología y economía conductual, como una preocupación fundamental para el desarrollo de sistemas de IA justos, transparentes y fiables. Argumentamos que medir y mitigar estos sesgos, ya sean relacionados con el atractivo o de otro tipo, es esencial no solo para la solidez técnica, sino también para garantizar que la interacción entre humanos y IA apoye, en lugar de socavar, la toma de decisiones equitativa en la sociedad.

K.1 El debate entre los sesgos cognitivos y la heurística

Paralelamente al trabajo sobre los sesgos cognitivos, a principios de la década de 2000 surgió una perspectiva complementaria, basada en la noción de racionalidad limitada [Sim90a]. Los estudiosos de esta tradición han criticado el paradigma de los sesgos cognitivos, argumentando que lo que a menudo se etiqueta como sesgos puede entenderse más bien como heurística, es decir, atajos adaptativos que no son limitaciones, sino fundamentales para la forma en que los seres humanos toman decisiones [Vra00].

Esta divergencia de perspectivas ha dado lugar al debate más reciente entre el nudging y el boosting [HG17]. El nudging, que se basa en el marco de los sesgos, hace hincapié en la modificación de las señales ambientales para orientar a las personas hacia mejores decisiones. El boosting, en línea con la perspectiva heurística, critica el nudging por paternalista y hace hincapié en el empoderamiento de las personas, ayudándolas a desarrollar los conocimientos, las habilidades y la conciencia necesarios para tomar mejores decisiones de forma autónoma. Ambos enfoques ofrecen valiosas perspectivas, y el debate entre ellos sigue activo, con los estudiosos continuando la controversia sobre el equilibrio entre la orientación externa y la agencia individual en el cambio de comportamiento.

Sin embargo, a efectos de esta tesis, resolver este debate no es lo más importante. Lo que más importa en el contexto del diseño de la IA es reconocer que las personas muestran patrones de juicio y comportamiento coherentes y predecibles, independientemente de si se enmarcan como *sesgos* o *heurísticas*. En cualquier caso, existe consenso sobre la existencia de tales regularidades, aunque se difiera en su interpretación normativa. Por lo tanto, en este trabajo adoptamos el término *sesgos cognitivos* como una forma conveniente de describir estos patrones, al tiempo que reconocemos su naturaleza controvertida y las importantes contribuciones de las perspectivas heurísticas y potenciadoras.

K.2 Contribuciones

A continuación, resumimos las principales contribuciones de la tesis:

- Desarrollo de un marco estructurado para organizar los sesgos cognitivos de manera directamente aplicable al diseño de sistemas colaborativos entre humanos e IA (capítulo 2).
- Creación de un conjunto de datos de alta calidad de rostros (AHEAD) para evaluar el efecto halo del atractivo, junto con conocimientos empíricos sobre cómo los filtros de belleza dan forma y amplifican este sesgo cognitivo humano (capítulo 3).
- Introducción del concepto de *lookismo algorítmico*, definido como la discriminación basada en el atractivo en la toma de decisiones de la IA, y articulación de sus implicaciones para la equidad algorítmica (capítulo 4).
- Evidencia empírica de que los modelos de lenguaje multimodal de gran tamaño (MLLM) muestran un sesgo de atractivo y reproducen el efecto halo del atractivo al razonar sobre imágenes faciales (capítulo 5).
- Evidencia empírica de que los modelos generativos de texto a imagen, concretamente Stable Diffusion, codifican un sesgo de atractivo al crear imágenes de rostros humanos, con consecuencias posteriores para tareas como la clasificación de género (capítulo 6).

K.3 Lista de publicaciones

La investigación principal presentada en esta tesis se difundió a través de artículos revisados por pares en conferencias académicas y revistas especializadas. Estos artículos constituyen la base de las contribuciones empíricas y conceptuales resumidas anteriormente, y se enumeran a continuación por orden cronológico:

- Aditya Gulati, Miguel Angel Lozano, Bruno Lepri, y Nuria Oliver. “Biased: Bringing irrationality into automated system design.” AAAI Fall Symposium 2022 on Thinking Fast and Slow and Other Cognitive Theories in AI, arXiv:2210.01122 (2022) [GLLO23]
- Aditya Gulati, Marina Martínez-García, Daniel Fernández, Miguel Angel Lozano, Bruno Lepri, y Nuria Oliver. “What is beautiful is still good: the attractiveness halo effect in the era of beauty filters.” Royal Society open science 11, no. 11 (2024): 240882 [GMF+24a]
- Aditya Gulati, Bruno Lepri, y Nuria Oliver. “Lookism: The overlooked bias in computer vision.” ECVC 2024 workshop on “Fairness and ethics towards transparent AI: facing the challenge through model Debiasing”, FAILED’25, arXiv:2408.11448 (2024) [GLO24]
- Miriam Doh, Aditya Gulati, Matei Mancas, y Nuria Oliver. “When Algorithms Play Favorites: Lookism in the Generation and Perception of Faces.” Fourth European Workshop on Algorithmic Fairness, EWAf’25, arXiv:2506.11025 (2025) [DGM025]
- Aditya Gulati, Moreno D’Incà, Nicu Sebe, Bruno Lepri, y Nuria Oliver. “Beauty and the Bias: Exploring the Impact of Attractiveness on Multimodal Large Language Models.” Eighth AAAI/ACM Conference on AI, Ethics and Society, AIES’25, arXiv:2504.16104 (2025) [GDS+25]

K.4 Conclusión

Esta tesis parte de la premisa de que los sesgos cognitivos son una parte inherente de la toma de decisiones humanas y destaca la oportunidad de estudiarlos en el contexto de la IA para comprender cómo los procesos de decisión humanos y mecánicos se influyen mutuamente. Para abordar esta cuestión, en el capítulo 2 proponemos en primer lugar un marco que clasifica los sesgos cognitivos conocidos desde la perspectiva de los sistemas humanos-IA, organizándolos según el momento en que surgen en el ciclo de toma de decisiones humanas.

Dentro de este marco más amplio, la tesis se centra en un sesgo cognitivo destacado pero poco estudiado: el efecto halo de atractivo (AHE), y ofrece una explicación detallada de cómo las percepciones de atractivo pueden sesgar los juicios tanto humanos como de las máquinas. Los rostros desempeñan un papel fundamental en la interacción entre seres humanos, y el efecto halo del atractivo se estableció en psicología ya en la década de 1970 [DBW72]. Sin embargo, son relativamente pocos los trabajos que han examinado cómo se manifiesta este sesgo en los entornos contemporáneos mediados digitalmente.

Los filtros de belleza son un ejemplo de este tipo de entornos: representan una intervención tecnológica relativamente reciente que aprovecha diversos métodos de IA para alterar la apariencia en tiempo real, lo que podría remodelar las formas en que opera el efecto halo del atractivo. Para investigar esto, en el capítulo 3 presentamos los resultados de uno de los primeros y más amplios estudios de usuarios para investigar la existencia del AHE en los seres humanos y el impacto de los filtros de belleza en la percepción del atractivo y en este sesgo cognitivo. Más allá de recopilar pruebas empíricas sólidas de la existencia de este sesgo cognitivo en un conjunto diverso de estímulos, nuestro estudio dio lugar al conjunto de datos AHEAD, un conjunto de datos seleccionados para estudiar los posibles sesgos de atractivo en los seres humanos y los algoritmos.

Armada con pruebas empíricas de cómo el atractivo influye en las decisiones humanas, la tesis pasó a centrarse en los modelos de aprendizaje automático. En el capítulo 4, introducimos el con-

cepto de *lookismo algorítmico*, es decir, la tendencia de los algoritmos a mostrar discriminación basada en el atractivo.

Examinamos este fenómeno en dos contextos: primero, cuando los sistemas de IA tienen la tarea de tomar decisiones basadas en imágenes de individuos (capítulo 5), y segundo, cuando los sistemas de IA generan imágenes de personas (capítulo 6). En nuestros experimentos con modelos lingüísticos multimodales de gran tamaño (MLLM), encontramos pruebas empíricas sólidas de que asocian el atractivo con rasgos positivos. Este hallazgo se observó de forma sistemática en siete modelos de código abierto de diversas arquitecturas. Aunque estos experimentos se llevaron a cabo en un entorno de tareas específico, demuestran que el *lookismo algorítmico* está integrado en los MLLM. Dado el amplio y creciente despliegue de estos modelos en distintos ámbitos de aplicación, la presencia de sesgos relacionados con el atractivo plantea importantes preocupaciones sobre su posible impacto en la toma de decisiones en el mundo real. Si bien se han realizado numerosas investigaciones sobre los sesgos relacionados con el género, la edad y la raza, se ha prestado relativamente poca atención a cómo influye el atractivo en los juicios basados en la IA.

Además, nuestros hallazgos revelan que el *lookismo algorítmico* también está presente en los modelos generativos de texto a imagen. Creamos un conjunto de datos de 13 200 rostros humanos sintéticos diferentes con Stable Diffusion 2.1. Los rostros se crean con diferentes atributos con valencia positiva (inteligente, digno de confianza, sociable y feliz) y negativa (poco inteligente, poco digno de confianza, poco sociable e infeliz) para estudiar la existencia del *lookismo algorítmico*. Encontramos pruebas empíricas claras de que Stable Diffusion 2.1 asocia el atractivo con los rostros creados con atributos positivos y viceversa. Además, creamos un clasificador de atractivo a partir de datos humanos y descubrimos que, según el clasificador, los rostros generados con atributos positivos son más atractivos que los generados con atributos negativos. Además, estudiamos en qué medida el *lookismo algorítmico* afecta a las aplicaciones posteriores, como los modelos de clasificación de género. Encontramos una reducción significativa en la precisión de la clasificación de género al evaluar imágenes generadas con indicaciones de rasgos negativos.

Estos resultados demuestran que el atractivo influye en la percepción humana en entornos mediados digitalmente y que estos efectos se han filtrado en los sistemas de IA cuando analizan rostros humanos, tanto para tomar decisiones sobre ellos como para generarlos. Además, el *lookismo algorítmico* afecta a las aplicaciones posteriores, concretamente a los modelos de clasificación de género, que funcionan significativamente peor cuando analizan imágenes faciales creadas con atributos negativos, y en particular imágenes de mujeres. Este refuerzo cíclico del sesgo sugiere un bucle de retroalimentación en el que los estereotipos relacionados con el atractivo no solo se reproducen, sino que también se amplifican mediante los sistemas de IA.

Otro hallazgo importante de esta tesis se refiere a la interacción entre el sesgo de atractivo y el género. En todos nuestros estudios, observamos de manera sistemática que las mujeres se ven afectadas de manera desproporcionada por este sesgo. En las evaluaciones humanas, el efecto halo del atractivo tuvo un impacto más negativo en las imágenes de mujeres, y en el caso del atractivo y la inteligencia, los evaluadores masculinos parecieron ser más susceptibles a la influencia de los filtros de belleza que las evaluadoras femeninas. En los MLLM, el *lookismo algorítmico* se expresó con mayor intensidad en las imágenes femeninas, y en nuestros experimentos con Stable Diffusion, el grupo con peor rendimiento en los clasificadores de género fue el de las imágenes de mujeres creadas con atributos negativos (y, en particular, las imágenes de *mujeres infelices*). Estos resultados convergentes sugieren que el sesgo de atractivo está profundamente entrelazado con el género, y que las mujeres son juzgadas con mayor intensidad en función de su apariencia. Por lo tanto, el sesgo de atractivo no solo coexiste con los sesgos de género existentes tanto en la sociedad como

en los sistemas de IA, sino que también puede exacerbarlos.

En general, la investigación descrita en esta tesis ha descubierto una forma crítica pero poco explorada de sesgo, que ha recibido mucha menos atención en comparación con los sesgos de género, raza o edad, pero que tiene importantes implicaciones tanto para la toma de decisiones humanas como para la de las máquinas. Demostramos que el sesgo de atractivo moldea las percepciones, influye directamente en los sistemas de IA y se propaga a tareas posteriores, creando así vulnerabilidades sistémicas en los ecosistemas sociotécnicos. Es importante destacar que esta tesis demuestra que, al poner de relieve el papel de los sesgos cognitivos en el diseño de la IA, podemos comprender mejor no solo las debilidades de los sistemas inteligentes, sino también las estrategias necesarias para hacerlos más equitativos y centrados en el ser humano. Al hacerlo, este trabajo contribuye al proyecto más amplio de alinear la IA con los valores, las complejidades y las expectativas de equidad de las sociedades a las que pretende servir. Como sugiere el título de la tesis, no debemos juzgar los libros —y menos aún a las personas— por sus portadas.

Bibliography

- [AAA+23] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, and et al., “Gpt-4 technical report.”, *arXiv preprint arXiv:2303.08774*, 2023 (cit. on p. 52).
- [AAA+24] M. Abdin et al., *Phi-3 technical report: A highly capable language model locally on your phone*, 2024. arXiv: 2404.14219 [cs.CL]. [Online]. Available: <https://arxiv.org/abs/2404.14219> (cit. on p. 60).
- [AAH+24] P. Agrawal, S. Antoniak, E. B. Hanna, B. Bout, D. Chaplot, J. Chudnovsky, D. Costa, B. D. Monicault, S. Garg, T. Gervet, S. Ghosh, A. Héliou, P. Jacob, A. Q. Jiang, K. Khandelwal, T. Lacroix, G. Lample, D. L. Casas, T. Lavril, T. L. Scao, A. Lo, W. Marshall, L. Martin, A. Mensch, P. Muddireddy, V. Nemychnikova, M. Pellat, P. V. Platen, N. Raghuraman, B. Rozière, A. Sablayrolles, L. Saulnier, R. Sauvestre, W. Shang, R. Soletskyi, L. Stewart, P. Stock, J. Studnia, S. Subramanian, S. Vaze, T. Wang, and S. Yang, *Pixtral 12b*, 2024. arXiv: 2410.07073 [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2410.07073> (cit. on p. 60).
- [Abo09] E. E. Abott, “On the analysis of the factor of recall in the learning process.”, *The Psychological Review: Monograph Supplements*, vol. 11, no. 1, pp. 159–177, 1909, ISSN: 0096-9753. DOI: 10.1037/h0093018. [Online]. Available: <http://dx.doi.org/10.1037/h0093018> (cit. on p. 24).
- [AC25] R. Ali and H. Cui, “Leveraging chatgpt for enhanced aesthetic evaluations in minimally invasive facial procedures”, *Aesthetic Plastic Surgery*, vol. 49, no. 3, pp. 950–961, Feb. 2025, ISSN: 1432-5241. DOI: 10.1007/s00266-024-04524-x. [Online]. Available: <http://dx.doi.org/10.1007/s00266-024-04524-x> (cit. on p. 59).
- [AFZ21] A. Abid, M. Farooqi, and J. Zou, “Persistent anti-muslim bias in large language models”, in *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 2021, pp. 298–306 (cit. on p. 57).
- [Agr10] A. Agresti, *Analysis of Ordinal Categorical Data*. online edn: Wiley, Mar. 2010, ISBN: 9780470594001. DOI: 10.1002/9780470594001 (cit. on p. 46).
- [AHPS23] M. Appel, F. Hutmacher, T. Politt, and J.-P. Stein, “Swipe right? using beauty filters in male tinder profiles reduces women’s evaluations of trustworthiness but increases physical attractiveness and dating intention”, *Computers in Human Behavior*, vol. 148, p. 107871, Nov. 2023, ISSN: 0747-5632. DOI: 10.1016/j.chb.2023.107871. [Online]. Available: <http://dx.doi.org/10.1016/j.chb.2023.107871> (cit. on pp. 28, 30, 41, 158).

- [Ain75] G. Ainslie, “Specious reward: A behavioral theory of impulsiveness and impulse control.”, *Psychological Bulletin*, vol. 82, no. 4, pp. 463–496, 1975. DOI: [10.1037/h0076860](https://doi.org/10.1037/h0076860). [Online]. Available: <https://doi.org/10.1037/h0076860> (cit. on p. 11).
- [AJ08] D. Ariely and S. Jones, *Predictably irrational*. HarperCollins New York, 2008 (cit. on pp. 1, 58, 153).
- [Aka74] H. Akaike, “A new look at the statistical model identification”, *IEEE Transactions on Automatic Control*, vol. 19, no. 6, pp. 716–723, Dec. 1974, ISSN: 0018-9286. DOI: [10.1109/tac.1974.1100705](https://doi.org/10.1109/tac.1974.1100705). [Online]. Available: <http://dx.doi.org/10.1109/TAC.1974.1100705> (cit. on pp. 47, 101, 103).
- [AKW+23] J. Ali, M. Kleindessner, F. Wenzel, K. Budhathoki, V. Cevher, and C. Russell, “Evaluating the fairness of discriminative foundation models in computer vision”, in *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, ser. AIES ’23, ACM, Aug. 2023, pp. 809–833. DOI: [10.1145/3600211.3604720](https://doi.org/10.1145/3600211.3604720). [Online]. Available: <http://dx.doi.org/10.1145/3600211.3604720> (cit. on p. 59).
- [AL17] S. Alm and S. B. Låftman, “The gendered mirror on the wall: Satisfaction with physical appearance and its relationship to global self-esteem and psychosomatic complaints among adolescent boys and girls”, *YOUNG*, vol. 26, no. 5, pp. 525–541, Dec. 2017, ISSN: 1741-3222. DOI: [10.1177/1103308817739733](https://doi.org/10.1177/1103308817739733). [Online]. Available: <http://dx.doi.org/10.1177/1103308817739733> (cit. on pp. 41, 159).
- [All79] L. G. Allan, “The perception of time”, *Perception & Psychophysics*, vol. 26, no. 5, pp. 340–354, Sep. 1979, ISSN: 1532-5962. DOI: [10.3758/bf03204158](https://doi.org/10.3758/bf03204158). [Online]. Available: <http://dx.doi.org/10.3758/BF03204158> (cit. on p. 25).
- [AM85] S. T. Allison and D. M. Messick, “The group attribution error”, *Journal of Experimental Social Psychology*, vol. 21, no. 6, pp. 563–579, Nov. 1985, ISSN: 0022-1031. DOI: [10.1016/0022-1031\(85\)90025-3](https://doi.org/10.1016/0022-1031(85)90025-3). [Online]. Available: [http://dx.doi.org/10.1016/0022-1031\(85\)90025-3](http://dx.doi.org/10.1016/0022-1031(85)90025-3) (cit. on p. 23).
- [AMD+97] L. Albright, T. E. Malloy, Q. Dong, D. A. Kenny, X. Fang, L. Winquist, and D. Yu, “Cross-cultural consensus in personality judgments.”, *Journal of Personality and Social Psychology*, vol. 72, no. 3, pp. 558–569, 1997, ISSN: 0022-3514. DOI: [10.1037/0022-3514.72.3.558](https://doi.org/10.1037/0022-3514.72.3.558). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.72.3.558> (cit. on pp. 28, 40, 158).
- [And84] J. A. Anderson, “Regression and ordered categorical variables”, *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 46, no. 1, pp. 1–30, 1984, ISSN: 00359246. [Online]. Available: <http://www.jstor.org/stable/2345457> (visited on 02/09/2024) (cit. on pp. 31, 34, 46).
- [AP08] J. Andreoni and R. Petrie, “Beauty, gender and stereotypes: Evidence from laboratory experiments”, *Journal of Economic Psychology*, vol. 29, no. 1, pp. 73–93, Feb. 2008, ISSN: 0167-4870. DOI: [10.1016/j.joep.2007.07.008](https://doi.org/10.1016/j.joep.2007.07.008). [Online]. Available: <http://dx.doi.org/10.1016/j.joep.2007.07.008> (cit. on p. 95).

- [AP47] G. W. Allport and L. Postman, “The psychology of rumor.”, 1947 (cit. on p. 23).
- [Asc46] S. E. Asch, “Forming impressions of personality.”, *The Journal of Abnormal and Social Psychology*, vol. 41, no. 3, pp. 258–290, Jul. 1946. DOI: [10.1037/h0055756](https://doi.org/10.1037/h0055756). [Online]. Available: <https://doi.org/10.1037/h0055756> (cit. on p. 14).
- [Asc51] S. E. Asch, “Effects of group pressure upon the modification and distortion of judgments”, in *Organizational influence processes*, Routledge, 1951, pp. 295–303 (cit. on p. 21).
- [ASS+16] M. Agthe, M. Strobel, M. Spörrle, M. Pfundmair, and J. K. Maner, “On the borders of harmful and helpful beauty biases: The biasing effects of physical attractiveness depend on sex and ethnicity”, *Evolutionary Psychology*, vol. 14, no. 2, p. 1474704916653968, 2016. DOI: [10.1177/1474704916653968](https://doi.org/10.1177/1474704916653968). eprint: <https://doi.org/10.1177/1474704916653968>. [Online]. Available: <https://doi.org/10.1177/1474704916653968> (cit. on pp. 28, 40, 41, 158).
- [ASV+24] A. Ananthram, E. Stengel-Eskin, C. Vondrick, M. Bansal, and K. McKeown, “See it from my perspective: Diagnosing the western cultural bias of large vision-language models in image understanding”, *arXiv preprint arXiv:2406.11665*, 2024 (cit. on p. 58).
- [Att53] F. Attneave, “Psychological probability as a function of experienced frequency.”, *Journal of Experimental Psychology*, vol. 46, no. 2, pp. 81–86, 1953, ISSN: 0022-1015. DOI: [10.1037/h0057955](https://doi.org/10.1037/h0057955). [Online]. Available: <http://dx.doi.org/10.1037/h0057955> (cit. on p. 21).
- [AXP+23] M. Atari, M. J. Xue, P. S. Park, D. E. Blasi, and J. Henrich, *Which humans?*, Sep. 2023. DOI: [10.31234/osf.io/5b26t](https://doi.org/10.31234/osf.io/5b26t). [Online]. Available: <http://dx.doi.org/10.31234/osf.io/5b26t> (cit. on p. 68).
- [Bak22] M. Bakker, “#nofilter how beauty filters affect the internalization of beauty ideals”, M.S. thesis, Utrecht University, 2022. [Online]. Available: <https://studenttheses.uu.nl/handle/20.500.12932/41835> (cit. on pp. 41, 87, 159).
- [Bar20] J. Barker, “Making-up on mobile: The pretty filters and ugly implications of snapchat”, *Fashion, Style & Popular Culture*, vol. 7, no. 2, pp. 207–221, Mar. 2020, ISSN: 2050-0726. DOI: [10.1386/fspc_00015_1](https://doi.org/10.1386/fspc_00015_1). [Online]. Available: http://dx.doi.org/10.1386/fspc_00015_1 (cit. on pp. 30, 41, 159).
- [Bar80] M. Bar-Hillel, “The base-rate fallacy in probability judgments”, *Acta Psychologica*, vol. 44, no. 3, pp. 211–233, May 1980. DOI: [10.1016/0001-6918\(80\)90046-3](https://doi.org/10.1016/0001-6918(80)90046-3). [Online]. Available: [https://doi.org/10.1016/0001-6918\(80\)90046-3](https://doi.org/10.1016/0001-6918(80)90046-3) (cit. on p. 11).
- [Bar94] J. Baron, *Thinking and deciding*. Cambridge University Press, 1994 (cit. on p. 20).
- [Bas81] J. N. Bassili, “The attractiveness stereotype: Goodness or glamour?”, *Basic and Applied Social Psychology*, vol. 2, no. 4, pp. 235–252, Dec. 1981, ISSN: 1532-4834. DOI: [10.1207/s15324834basp0204_1](https://doi.org/10.1207/s15324834basp0204_1). [Online]. Available: http://dx.doi.org/10.1207/s15324834basp0204_1 (cit. on pp. 30, 39, 157).

- [BBLC20] J. P. Beauchamp, D. J. Benjamin, D. I. Laibson, and C. F. Chabris, “Measuring and controlling for the compromise effect when estimating risk preference parameters”, *Experimental Economics*, vol. 23, no. 4, pp. 1069–1099, Dec. 2020, ISSN: 1573-6938. DOI: [10.1007/s10683-019-09640-z](https://doi.org/10.1007/s10683-019-09640-z). [Online]. Available: <http://dx.doi.org/10.1007/s10683-019-09640-z> (cit. on p. 20).
- [BC77] M. J. Baker and G. A. Churchill, “The impact of physically attractive models on advertising evaluations”, *Journal of Marketing Research*, vol. 14, no. 4, pp. 538–555, Nov. 1977, ISSN: 1547-7193. DOI: [10.1177/002224377701400411](https://doi.org/10.1177/002224377701400411). [Online]. Available: <http://dx.doi.org/10.1177/002224377701400411> (cit. on p. 74).
- [BD20] J. F. Black and S. Davidai, “Do rich people “deserve” to be rich? charitable giving, internal attributions of wealth, and judgments of economic deservingness”, *Journal of Experimental Social Psychology*, vol. 90, p. 104011, Sep. 2020, ISSN: 0022-1031. DOI: [10.1016/j.jesp.2020.104011](https://doi.org/10.1016/j.jesp.2020.104011). [Online]. Available: <http://dx.doi.org/10.1016/j.jesp.2020.104011> (cit. on p. 63).
- [BGJ+23] J. Betker, G. Goh, L. Jing, T. Brooks, J. Wang, L. Li, L. Ouyang, J. Zhuang, J. Lee, Y. Guo, and et al., “Improving image generation with better captions.”, *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, vol. 2, no. 3, p. 8, 2023 (cit. on p. 52).
- [BGM22] F. Brunner, F. Gamm, and W. Mill, “Myportfolio: The ikea effect in financial investment decisions”, *Journal of Banking & Finance*, p. 106529, 2022, ISSN: 0378-4266. DOI: <https://doi.org/10.1016/j.jbankfin.2022.106529>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378426622001236> (cit. on p. 12).
- [BGR94] R. Buehler, D. Griffin, and M. Ross, “Exploring the “planning fallacy”: Why people underestimate their task completion times.”, *Journal of Personality and Social Psychology*, vol. 67, no. 3, pp. 366–381, Sep. 1994, ISSN: 0022-3514. DOI: [10.1037/0022-3514.67.3.366](https://doi.org/10.1037/0022-3514.67.3.366). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.67.3.366> (cit. on p. 20).
- [BH22] R. S. Baker and A. Hawn, “Algorithmic bias in education”, *International Journal of Artificial Intelligence in Education*, pp. 1–41, 2022 (cit. on p. 49).
- [BH88] J. Baron and J. C. Hershey, “Outcome bias in decision evaluation.”, *Journal of Personality and Social Psychology*, vol. 54, no. 4, pp. 569–579, 1988, ISSN: 0022-3514. DOI: [10.1037/0022-3514.54.4.569](https://doi.org/10.1037/0022-3514.54.4.569). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.54.4.569> (cit. on p. 20).
- [BHS+21] B. M. Booth, L. Hickman, S. K. Subburaj, L. Tay, S. E. Woo, and S. K. D’Mello., “Bias and fairness in multimodal machine learning: A case study of automated video interviews.”, in *Proceedings of the 2021 International Conference on Multimodal Interaction*, ser. ICMI ’21, ACM, Oct. 2021 (cit. on p. 51).
- [BJP17] N. Berggren, H. Jordahl, and P. Poutvaara, “The right look: Conservative politicians look better and voters reward it”, *Journal of Public Economics*, vol. 146, pp. 79–86, Feb. 2017. DOI: [10.1016/j.jpubeco.2016.12.008](https://doi.org/10.1016/j.jpubeco.2016.12.008). [Online]. Available: <https://doi.org/10.1016/j.jpubeco.2016.12.008> (cit. on p. 31).

- [BK72] R. Barocas and P. Karoly, “Effects of physical appearance on social responsiveness”, *Psychological Reports*, vol. 31, no. 2, pp. 495–500, Oct. 1972. DOI: [10.2466/pr0.1972.31.2.495](https://doi.org/10.2466/pr0.1972.31.2.495). [Online]. Available: <https://doi.org/10.2466/pr0.1972.31.2.495> (cit. on pp. 28, 29, 34, 38, 39, 41, 60, 156, 158).
- [BKTR08] S. A. Banducci, J. A. Karp, M. Thrasher, and C. Rallings., “Ballot photographs as cues in low-information elections.”, *Political Psychology*, vol. 29, no. 6, pp. 903–17, 2008. [Online]. Available: <http://www.jstor.org/stable/20447173> (cit. on pp. 3, 28, 160).
- [BL10] G. Barron and S. Leider, “The role of experience in the gambler’s fallacy”, *Journal of Behavioral Decision Making*, vol. 23, no. 1, pp. 117–129, Jan. 2010. DOI: [10.1002/bdm.676](https://doi.org/10.1002/bdm.676). [Online]. Available: <https://doi.org/10.1002/bdm.676> (cit. on p. 11).
- [BM66] R. Brown and D. McNeill, “The “tip of the tongue” phenomenon”, *Journal of Verbal Learning and Verbal Behavior*, vol. 5, no. 4, pp. 325–337, Aug. 1966, ISSN: 0022-5371. DOI: [10.1016/S0022-5371\(66\)80040-3](https://doi.org/10.1016/S0022-5371(66)80040-3). [Online]. Available: [http://dx.doi.org/10.1016/S0022-5371\(66\)80040-3](http://dx.doi.org/10.1016/S0022-5371(66)80040-3) (cit. on p. 23).
- [BM89] A. S. Brown and D. R. Murphy, “Cryptomnesia: Delineating inadvertent plagiarism.”, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 15, no. 3, pp. 432–442, May 1989, ISSN: 0278-7393. DOI: [10.1037/0278-7393.15.3.432](https://doi.org/10.1037/0278-7393.15.3.432). [Online]. Available: <http://dx.doi.org/10.1037/0278-7393.15.3.432> (cit. on p. 24).
- [BMG21] Z. Bućinca, M. B. Malaya, and K. Z. Gajos, “To trust or to think: Cognitive forcing functions can reduce overreliance on ai in ai-assisted decision-making”, *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW1, pp. 1–21, Apr. 2021. DOI: [10.1145/3449287](https://doi.org/10.1145/3449287). [Online]. Available: <https://doi.org/10.1145/3449287> (cit. on pp. 2, 154).
- [BPCR21] C. Batres, A. Porcheron, S. Courrèges, and R. Russell, “Professional versus self-applied makeup: Do makeup artists add value?”, *Perception*, vol. 50, no. 8, pp. 709–719, Jul. 2021, ISSN: 1468-4233. DOI: [10.1177/03010066211029218](https://doi.org/10.1177/03010066211029218). [Online]. Available: <http://dx.doi.org/10.1177/03010066211029218> (cit. on pp. 28, 30).
- [BPL+19] C. Batres, A. Porcheron, J. Latreille, M. Roche, F. Morizot, and R. Russell, “Cosmetics increase skin evenness: Evidence from perceptual and physical measures”, *Skin Research and Technology*, vol. 25, no. 5, pp. 672–676, May 2019, ISSN: 1600-0846. DOI: [10.1111/srt.12700](https://doi.org/10.1111/srt.12700). [Online]. Available: <http://dx.doi.org/10.1111/srt.12700> (cit. on pp. 28, 30).
- [BR22] C. Batres and H. Robinson, “Makeup increases attractiveness in male faces”, *PLOS ONE*, vol. 17, no. 11, A. B. Mahmoud, Ed., e0275662, Nov. 2022, ISSN: 1932-6203. DOI: [10.1371/journal.pone.0275662](https://doi.org/10.1371/journal.pone.0275662). [Online]. Available: <http://dx.doi.org/10.1371/journal.pone.0275662> (cit. on pp. 28, 29).
- [Bre73] M. Brenner, “The next-in-line effect”, *Journal of Verbal Learning and Verbal Behavior*, vol. 12, no. 3, pp. 320–323, Jun. 1973, ISSN: 0022-5371. DOI: [10.1016/S0022-5371\(73\)80076-3](https://doi.org/10.1016/S0022-5371(73)80076-3). [Online]. Available: [http://dx.doi.org/10.1016/S0022-5371\(73\)80076-3](http://dx.doi.org/10.1016/S0022-5371(73)80076-3) (cit. on p. 25).

- [BRS+18] C. Batres, R. Russell, J. A. Simpson, L. Campbell, A. M. Hansen, and L. Cronk, “Evidence that makeup is a false signal of sociosexuality”, *Personality and Individual Differences*, vol. 122, pp. 148–154, Feb. 2018, ISSN: 0191-8869. DOI: [10.1016/j.paid.2017.10.023](https://doi.org/10.1016/j.paid.2017.10.023). [Online]. Available: <http://dx.doi.org/10.1016/j.paid.2017.10.023> (cit. on p. 28).
- [BS07] A. K. Barbey and S. A. Sloman, “Base-rate respect: From ecological rationality to dual processes”, *Behavioral and Brain Sciences*, vol. 30, no. 3, pp. 241–254, Jun. 2007. DOI: [10.1017/s0140525x07001653](https://doi.org/10.1017/s0140525x07001653). [Online]. Available: <https://doi.org/10.1017/s0140525x07001653> (cit. on p. 11).
- [BS22] C. Batres and V. Shiramizu, “Examining the “attractiveness halo effect” across cultures”, *Current Psychology*, Aug. 2022. DOI: [10.1007/s12144-022-03575-0](https://doi.org/10.1007/s12144-022-03575-0). [Online]. Available: <https://doi.org/10.1007/s12144-022-03575-0> (cit. on pp. 28, 33, 40–42, 46, 54, 60, 158).
- [BSS24] M. Bahrami, R. Sonoda, and R. Srinivasan, “Llm diagnostic toolkit: Evaluating llms for ethical issues”, in *2024 International Joint Conference on Neural Networks (IJCNN)*, IEEE, 2024, pp. 1–8 (cit. on p. 58).
- [BT01] S. Benartzi and R. H. Thaler, “Naive diversification strategies in defined contribution saving plans”, *American Economic Review*, vol. 91, no. 1, pp. 79–98, Mar. 2001, ISSN: 0002-8282. DOI: [10.1257/aer.91.1.79](https://doi.org/10.1257/aer.91.1.79). [Online]. Available: <http://dx.doi.org/10.1257/aer.91.1.79> (cit. on p. 25).
- [BT54] J. S. Bruner and R. Tagiuri, “The perception of people”, *Handbook of Social Psychology*, vol. 2, pp. 634–654, 1954 (cit. on p. 25).
- [Bur13] D. V. Burakov, “Exogenous credit cycle: An experimental study”, *World Applied Sciences Journal*, vol. 26, no. 6, pp. 733–736, 2013 (cit. on p. 9).
- [BV91] B. P. Buunk and N. W. Van Yperen, “Referential comparisons, relational comparisons, and exchange orientation: Their relation to marital satisfaction”, *Personality and Social Psychology Bulletin*, vol. 17, no. 6, pp. 709–717, Dec. 1991, ISSN: 1552-7433. DOI: [10.1177/0146167291176015](https://doi.org/10.1177/0146167291176015). [Online]. Available: <http://dx.doi.org/10.1177/0146167291176015> (cit. on p. 22).
- [BWSG25] X. Bai, A. Wang, I. Sucholutsky, and T. L. Griffiths, “Explicitly unbiased large language models still form biased associations”, *Proceedings of the National Academy of Sciences*, vol. 122, no. 8, Feb. 2025, ISSN: 1091-6490. DOI: [10.1073/pnas.2416228122](https://doi.org/10.1073/pnas.2416228122). [Online]. Available: <http://dx.doi.org/10.1073/pnas.2416228122> (cit. on p. 58).
- [Cas32] H. Cason, “The learning and retention of pleasant and unpleasant activities.”, *Archives of Psychology*, 1932 (cit. on p. 24).
- [CBB+24] G. Capitani, L. Bonicelli, F. Bolelli, S. Calderara, E. Ficarra, *et al.*, *Beyond the surface: Comprehensive analysis of implicit bias in vision-language models*, 2024. [Online]. Available: <https://hdl.handle.net/11380/1350126> (cit. on pp. 58, 67).

- [CBQT13] S. J. Cunningham, J. L. Brebner, F. Quinn, and D. J. Turk, “The self-reference effect on memory in early childhood”, *Child Development*, vol. 85, no. 2, pp. 808–823, Jul. 2013, ISSN: 1467-8624. DOI: [10.1111/cdev.12144](https://doi.org/10.1111/cdev.12144). [Online]. Available: <http://dx.doi.org/10.1111/cdev.12144> (cit. on p. 14).
- [CCD71] J. F. Cross, J. Cross, and J. Daly, “Sex, race, age, and beauty as factors in recognition of faces”, *Perception & Psychophysics*, vol. 10, no. 6, pp. 393–396, Nov. 1971, ISSN: 1532-5962. DOI: [10.3758/bf03210319](https://doi.org/10.3758/bf03210319). [Online]. Available: <http://dx.doi.org/10.3758/BF03210319> (cit. on pp. 35, 40, 158).
- [CCZ+24] D. Chen, R. Chen, S. Zhang, Y. Liu, Y. Wang, H. Zhou, Q. Zhang, P. Zhou, Y. Wan, and L. Sun, “Mllm-as-a-judge: Assessing multimodal llm-as-a-judge with vision-language benchmark”, *arXiv preprint arXiv:2402.04788*, 2024 (cit. on p. 61).
- [CDAR20] R. T. Cristel, S. H. Dayan, M. Akinosun, and P. T. Russell, “Evaluation of selfies and filtered selfies and effects on first impressions”, *Aesthetic Surgery Journal*, vol. 41, no. 1, pp. 122–130, Jan. 2020, ISSN: 1527-330X. DOI: [10.1093/asj/sjz362](https://doi.org/10.1093/asj/sjz362). [Online]. Available: <http://dx.doi.org/10.1093/asj/sjz362> (cit. on pp. 30, 51).
- [CF15] J. Cone and M. J. Ferguson, “He did what? the role of diagnosticity in revising implicit evaluations.”, *Journal of Personality and Social Psychology*, vol. 108, no. 1, pp. 37–57, Jan. 2015, ISSN: 0022-3514. DOI: [10.1037/pspa0000014](https://doi.org/10.1037/pspa0000014). [Online]. Available: <http://dx.doi.org/10.1037/pspa0000014> (cit. on p. 42).
- [Cha67] L. J. Chapman, “Illusory correlation in observational report”, *Journal of Verbal Learning and Verbal Behavior*, vol. 6, no. 1, pp. 151–155, Feb. 1967, ISSN: 0022-5371. DOI: [10.1016/S0022-5371\(67\)80066-5](https://doi.org/10.1016/S0022-5371(67)80066-5). [Online]. Available: [http://dx.doi.org/10.1016/S0022-5371\(67\)80066-5](http://dx.doi.org/10.1016/S0022-5371(67)80066-5) (cit. on p. 21).
- [Chr18] R. H. B. Christensen, “Cumulative link models for ordinal regression with the r package ordinal”, *Submitted in J. Stat. Software*, vol. 35, 2018. [Online]. Available: https://cran.uni-muenster.de/web/packages/ordinal/vignettes/clm_article.pdf (cit. on pp. 46, 47).
- [Chr23] R. H. B. Christensen, *Ordinal—regression models for ordinal data*, R package version 2023.12-4, 2023. [Online]. Available: <https://CRAN.R-project.org/package=ordinal> (cit. on p. 46).
- [CJ21] Y. Chen and J. Joo., “Understanding and mitigating annotation bias in facial expression recognition.”, in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct. 2021, pp. 14 980–14 991 (cit. on p. 51).
- [CK85] T. F. Cash and R. N. Kilcullen, “The aye of the beholder: Susceptibility to sexism and beautyism in the evaluation of managerial applicants1”, *Journal of Applied Social Psychology*, vol. 15, no. 4, pp. 591–605, Jun. 1985. DOI: [10.1111/j.1559-1816.1985.tb00903.x](https://doi.org/10.1111/j.1559-1816.1985.tb00903.x). [Online]. Available: <https://doi.org/10.1111/j.1559-1816.1985.tb00903.x> (cit. on pp. 3, 28, 61, 160).
- [CK96] Z. Carmon and D. Kahneman, “The experienced utility of queuing: Real time affect and retrospective evaluations of simulated queues”, *Duke University: Durham, NC, USA*, 1996 (cit. on p. 14).

- [CKO+25] Y. Chen, S. N. Kirshner, A. Ovchinnikov, M. Andiappan, and T. Jenkin, “A manager and an ai walk into a bar: Does chatgpt make biased decisions like we do?”, *Manufacturing & Service Operations Management*, vol. 27, no. 2, pp. 354–368, Mar. 2025, ISSN: 1526-5498. DOI: [10.1287/msom.2023.0279](https://doi.org/10.1287/msom.2023.0279). [Online]. Available: <http://dx.doi.org/10.1287/msom.2023.0279> (cit. on p. 58).
- [CL72] F. I. Craik and R. S. Lockhart, “Levels of processing: A framework for memory research”, *Journal of Verbal Learning and Verbal Behavior*, vol. 11, no. 6, pp. 671–684, Dec. 1972, ISSN: 0022-5371. DOI: [10.1016/s0022-5371\(72\)80001-x](https://doi.org/10.1016/s0022-5371(72)80001-x). [Online]. Available: [http://dx.doi.org/10.1016/S0022-5371\(72\)80001-X](http://dx.doi.org/10.1016/S0022-5371(72)80001-X) (cit. on p. 23).
- [CLW89] C. Camerer, G. Loewenstein, and M. Weber, “The curse of knowledge in economic settings: An experimental analysis”, *Journal of Political Economy*, vol. 97, no. 5, pp. 1232–1254, Oct. 1989, ISSN: 1537-534X. DOI: [10.1086/261651](https://doi.org/10.1086/261651). [Online]. Available: <http://dx.doi.org/10.1086/261651> (cit. on p. 22).
- [CM60] D. P. Crowne and D. Marlowe, “A new scale of social desirability independent of psychopathology.”, *Journal of Consulting Psychology*, vol. 24, no. 4, pp. 349–354, 1960. DOI: [10.1037/h0047358](https://doi.org/10.1037/h0047358). [Online]. Available: <https://doi.org/10.1037/h0047358> (cit. on pp. 13, 21).
- [CMS16] D. L. Chen, T. J. Moskowitz, and K. Shue, “Decision making under the gambler’s fallacy: Evidence from asylum judges, loan officers, and baseball umpires”, *The Quarterly Journal of Economics*, vol. 131, no. 3, pp. 1181–1242, Mar. 2016. DOI: [10.1093/qje/qjw017](https://doi.org/10.1093/qje/qjw017). [Online]. Available: <https://doi.org/10.1093/qje/qjw017> (cit. on p. 11).
- [Cor04] S. J. Correll, “Constraints into preferences: Gender, status, and emerging career aspirations”, *American Sociological Review*, vol. 69, no. 1, pp. 93–113, Feb. 2004, ISSN: 1939-8271. DOI: [10.1177/000312240406900106](https://doi.org/10.1177/000312240406900106). [Online]. Available: <http://dx.doi.org/10.1177/000312240406900106> (cit. on pp. 41, 159).
- [Cot11] L. Cotter, “Self-perceived attractiveness and its influence on the halo effect and the similar-to me effect”, Ph.D. dissertation, Bucknell University, 2011. [Online]. Available: https://digitalcommons.bucknell.edu/honors_theses/18 (cit. on p. 96).
- [CPS09] V. Coetzee, D. I. Perrett, and I. D. Stephen., “Facial adiposity: A cue to health?”, *Perception*, vol. 38, no. 11, pp. 1700–1711, Jan. 2009 (cit. on p. 54).
- [CRB+95] M. R. Cunningham, A. R. Roberts, A. P. Barbee, P. B. Druen, and C.-H. Wu., ““their ideas of beauty are, on the whole, the same as ours”: Consistency and variability in the cross-cultural perception of female physical attractiveness.”, *Journal of Personality and Social Psychology*, vol. 68, no. 2, pp. 261–279, Feb. 1995 (cit. on p. 53).
- [Das22] J. Dastin, “Amazon scraps secret ai recruiting tool that showed bias against women”, in *Ethics of data and analytics*, Auerbach Publications, 2022, pp. 296–299 (cit. on p. 51).

- [Dav83] W. P. Davison, “The third-person effect in communication”, *Public Opinion Quarterly*, vol. 47, no. 1, p. 1, 1983, ISSN: 0033-362X. DOI: [10.1086/268763](https://doi.org/10.1086/268763). [Online]. Available: <http://dx.doi.org/10.1086/268763> (cit. on p. 22).
- [DBW72] K. Dion, E. Berscheid, and E. Walster, “What is beautiful is good.”, *Journal of Personality and Social Psychology*, vol. 24, no. 3, pp. 285–290, 1972. DOI: [10.1037/h0033731](https://doi.org/10.1037/h0033731). [Online]. Available: <https://doi.org/10.1037/h0033731> (cit. on pp. 3, 12, 22, 27–29, 33, 39, 41, 42, 46, 50, 56, 63, 70, 91, 154, 157, 159, 162).
- [dCG24] A. F. de Caleyá Vázquez and E. C. Garrido-Merchán, *A taxonomy of the biases of the images created by generative artificial intelligence*, 2024. arXiv: [2407.01556](https://arxiv.org/abs/2407.01556) [cs.CY]. [Online]. Available: <https://arxiv.org/abs/2407.01556> (cit. on p. 59).
- [DCL+24] M. Deitke, C. Clark, S. Lee, R. Tripathi, Y. Yang, J. S. Park, M. Salehi, N. Muenighoff, K. Lo, L. Soldaini, J. Lu, T. Anderson, E. Bransom, K. Ehsani, H. Ngo, Y. Chen, A. Patel, M. Yatskar, C. Callison-Burch, A. Head, R. Hendrix, F. Bastani, E. VanderBilt, N. Lambert, Y. Chou, A. Chheda, J. Sparks, S. Skjonsberg, M. Schmitz, A. Sarnat, B. Bischoff, P. Walsh, C. Newell, P. Wolters, T. Gupta, K.-H. Zeng, J. Borchardt, D. Groeneveld, C. Nam, S. Lebrecht, C. Wittlif, C. Schoenick, O. Michel, R. Krishna, L. Weihs, N. A. Smith, H. Hajishirzi, R. Girshick, A. Farhadi, and A. Kembhavi, *Molmo and pixmo: Open weights and open data for state-of-the-art vision-language models*, 2024. arXiv: [2409.17146](https://arxiv.org/abs/2409.17146) [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2409.17146> (cit. on p. 60).
- [DCLT18] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova., “Bert: Pre-training of deep bidirectional transformers for language understanding.”, *arXiv preprint arXiv:1810.04805*, 2018 (cit. on p. 52).
- [DdlFG+24] E. Derner, S. S. de la Fuente, Y. Gutiérrez, P. Moreda, and N. Oliver, “Leveraging large language models to measure gender bias in gendered languages”, *arXiv preprint arXiv:2406.13677*, 2024 (cit. on p. 57).
- [DE13] P. De Maeyer and H. Estelami, “Applying the peak-end rule to reference prices”, *Journal of Product & Brand Management*, vol. 22, no. 3, pp. 260–265, 2013 (cit. on p. 14).
- [Des10] J. Desir., “Lookism: Pushing the frontier of equality by looking beyond the law.”, *U. Ill. L. Rev.*, p. 629, 2010 (cit. on p. 54).
- [DFP+20] E. Dimara, S. Franconeri, C. Plaisant, A. Bezerianos, and P. Dragicevic, “A task-based taxonomy of cognitive biases for information visualization”, *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 2, pp. 1413–1432, Feb. 2020. [Online]. Available: <https://doi.org/10.1109/tvcg.2018.2872577> (cit. on pp. 2, 7, 154).
- [DGMO25] M. Doh, A. Gulati, M. Mancas, and N. Oliver, *When algorithms play favorites: Lookism in the generation and perception of faces*, 2025. arXiv: [2506.11025](https://arxiv.org/abs/2506.11025) [cs.LG]. [Online]. Available: <https://arxiv.org/abs/2506.11025> (cit. on pp. 6, 73, 162).

- [DGY+19] J. Deng, J. Guo, J. Yang, N. Xue, I. Kotsia, and S. Zafeiriou, “Arcface: Additive angular margin loss for deep face recognition”, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4690–4699 (cit. on p. 77).
- [DLS+22] T. Draws, D. La Barbera, M. Soprano, K. Roitero, D. Ceolin, A. Checco, and S. Mizzaro, “The effects of crowd worker biases in fact-checking tasks”, in *2022 ACM Conference on Fairness, Accountability, and Transparency*, ser. FAccT ’22, ACM, Jun. 2022, pp. 2114–2124. DOI: [10.1145/3531146.3534629](https://doi.org/10.1145/3531146.3534629). [Online]. Available: <http://dx.doi.org/10.1145/3531146.3534629> (cit. on p. 58).
- [DT75] M. Dermer and D. L. Thiel, “When beauty may fail.”, *Journal of Personality and Social Psychology*, vol. 31, no. 6, pp. 1168–1176, Jun. 1975. DOI: [10.1037/h0077085](https://doi.org/10.1037/h0077085). [Online]. Available: <https://doi.org/10.1037/h0077085> (cit. on pp. 29, 39, 157).
- [Dun45] K. Duncker, “On problem-solving.”, *Psychological Monographs*, vol. 58, no. 5, L. S. Lees, Ed., pp. i–113, 1945, ISSN: 0096-9753. DOI: [10.1037/h0093599](https://doi.org/10.1037/h0093599). [Online]. Available: <http://dx.doi.org/10.1037/h0093599> (cit. on p. 25).
- [DZZ+24] X. Dong, P. Zhang, Y. Zang, Y. Cao, B. Wang, L. Ouyang, X. Wei, S. Zhang, H. Duan, M. Cao, *et al.*, “Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model”, *arXiv preprint arXiv:2401.16420*, 2024 (cit. on p. 56).
- [EAML91] A. H. Eagly, R. D. Ashmore, M. G. Makhijani, and L. C. Longo, “What is beautiful is good, but...: A meta-analytic review of research on the physical attractiveness stereotype.”, *Psychological Bulletin*, vol. 110, no. 1, pp. 109–128, Jul. 1991. DOI: [10.1037/0033-2909.110.1.109](https://doi.org/10.1037/0033-2909.110.1.109). [Online]. Available: <https://doi.org/10.1037/0033-2909.110.1.109> (cit. on pp. 28, 39, 40, 50, 157, 158).
- [EBP83] J. S. B. T. Evans, J. L. Barston, and P. Pollard, “On the conflict between logic and belief in syllogistic reasoning”, *Memory & Cognition*, vol. 11, no. 3, pp. 295–306, May 1983. [Online]. Available: <https://doi.org/10.3758/bf03196976> (cit. on p. 21).
- [ECY+21] A. Esteva, K. Chou, S. Yeung, N. Naik, A. Madani, A. Mottaghi, Y. Liu, E. Topol, J. Dean, and R. Socher., “Deep learning-enabled medical computer vision.”, *npj Digital Medicine*, vol. 4, no. 1, Jan. 2021 (cit. on p. 49).
- [EDR06] Y. Eysenck, G. Dror, and E. Ruppel., “Facial attractiveness: Beauty and the machine.”, *Neural Computation*, vol. 18, no. 1, pp. 119–142, Jan. 2006 (cit. on p. 53).
- [EG17] A. S. Elias and R. Gill, “Beauty surveillance: The digital self-monitoring cultures of neoliberalism”, *European Journal of Cultural Studies*, vol. 21, no. 1, pp. 59–77, Jun. 2017, ISSN: 1460-3551. DOI: [10.1177/1367549417705604](https://doi.org/10.1177/1367549417705604). [Online]. Available: <http://dx.doi.org/10.1177/1367549417705604> (cit. on p. 108).
- [EHoH+23] Z. Epstein, A. Hertzmann, I. of Human Creativity, M. Akten, H. Farid, J. Fjeld, M. R. Frank, M. Groh, L. Herman, N. Leach, *et al.*, “Art and the science of generative ai”, *Science*, vol. 380, no. 6650, pp. 1110–1111, 2023 (cit. on p. 73).

- [EK02] A. H. Eagly and S. J. Karau, “Role congruity theory of prejudice toward female leaders.”, *Psychological Review*, vol. 109, no. 3, pp. 573–598, 2002, ISSN: 0033-295X. DOI: [10.1037/0033-295X.109.3.573](https://doi.org/10.1037/0033-295X.109.3.573). [Online]. Available: <http://dx.doi.org/10.1037/0033-295X.109.3.573> (cit. on pp. 41, 159).
- [ELA+24] J. M. Echterhoff, Y. Liu, A. Alessa, J. McAuley, and Z. He, “Cognitive bias in decision-making with LLMs”, in *Findings of the Association for Computational Linguistics: EMNLP 2024*, Y. Al-Onaizan, M. Bansal, and Y.-N. Chen, Eds., Miami, Florida, USA: Association for Computational Linguistics, Nov. 2024, pp. 12 640–12 653. DOI: [10.18653/v1/2024.findings-emnlp.739](https://doi.org/10.18653/v1/2024.findings-emnlp.739). [Online]. Available: <https://aclanthology.org/2024.findings-emnlp.739/> (cit. on pp. 1, 58, 153).
- [Ell61] D. Ellsberg, “Risk, ambiguity, and the savage axioms”, *The Quarterly Journal of Economics*, vol. 75, no. 4, p. 643, Nov. 1961, ISSN: 0033-5533. DOI: [10.2307/1884324](https://doi.org/10.2307/1884324). [Online]. Available: <http://dx.doi.org/10.2307/1884324> (cit. on p. 25).
- [ELX+24] A. Elangovan, L. Liu, L. Xu, S. B. Bodapati, and D. Roth, “ConSiDERS-the-human evaluation framework: Rethinking human evaluation for generative large language models”, in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, L.-W. Ku, A. Martins, and V. Srikumar, Eds., Bangkok, Thailand: Association for Computational Linguistics, Aug. 2024, pp. 1137–1160. DOI: [10.18653/v1/2024.acl-long.63](https://doi.org/10.18653/v1/2024.acl-long.63). [Online]. Available: <https://aclanthology.org/2024.acl-long.63/> (cit. on p. 58).
- [EM94] A. H. Eagly and A. Mladinic, “Are people prejudiced against women? some answers from research on attitudes, gender stereotypes, and judgments of competence”, *European Review of Social Psychology*, vol. 5, no. 1, pp. 1–35, Jan. 1994, ISSN: 1479-277X. DOI: [10.1080/14792779543000002](https://doi.org/10.1080/14792779543000002). [Online]. Available: <http://dx.doi.org/10.1080/14792779543000002> (cit. on p. 22).
- [ERL10] N. C. Ebner, M. Riediger, and U. Lindenberger, “FACES—a database of facial expressions in young, middle-aged, and older women and men: Development and validation”, *Behavior Research Methods*, vol. 42, no. 1, pp. 351–362, Feb. 2010. DOI: [10.3758/brm.42.1.351](https://doi.org/10.3758/brm.42.1.351). [Online]. Available: <https://doi.org/10.3758/brm.42.1.351> (cit. on pp. 29, 42–44, 59, 60).
- [ES09] A. H. Eagly and S. Sczesny, “Stereotypes about women, men, and leaders: Have times changed?”, in *The glass ceiling in the 21st century: Understanding barriers to gender equality*. American Psychological Association, 2009, pp. 21–47, ISBN: 9781433804090. DOI: [10.1037/11863-002](https://doi.org/10.1037/11863-002). [Online]. Available: <https://psycnet.apa.org/doi/10.1037/11863-002> (cit. on pp. 41, 159).
- [ESAA16] M. Ehrlich, T. J. Shields, T. Almaev, and M. R. Amer, “Facial attributes classification using multi-task representation learning”, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Jun. 2016 (cit. on p. 152).

- [Esh20] J. Eshiet, ““real me versus social media me:” filters, snapchat dysmorphia, and beauty perceptions among young women”, M.S. thesis, California State University, 2020. [Online]. Available: <https://scholarworks.lib.csusb.edu/cgi/viewcontent.cgi?article=2243&context=etd> (cit. on pp. 30, 41, 87, 159).
- [EW12] A. H. Eagly and W. Wood, “Social role theory”, in *Handbook of Theories of Social Psychology*. SAGE Publications Ltd, 2012, pp. 458–476. DOI: 10.4135/9781446249222.n49. [Online]. Available: <http://dx.doi.org/10.4135/9781446249222.n49> (cit. on pp. 41, 159).
- [FAP16] D. Fernández, R. Arnold, and S. Pledger, “Mixture-based clustering for the ordered stereotype model”, *Computational Statistics & Data Analysis*, vol. 93, pp. 46–75, 2016, ISSN: 0167-9473. DOI: <https://doi.org/10.1016/j.csda.2014.11.004>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016794731400317X> (cit. on pp. 31, 34, 46, 100).
- [FB75] B. Fischhoff and R. Beyth, “I knew it would happen”, *Organizational Behavior and Human Performance*, vol. 13, no. 1, pp. 1–16, Feb. 1975. [Online]. Available: [https://doi.org/10.1016/0030-5073\(75\)90002-1](https://doi.org/10.1016/0030-5073(75)90002-1) (cit. on p. 24).
- [FC11] P. W. Foos and M. C. Clark, “Adult age and gender differences in perceptions of facial attractiveness: Beauty is in the eye of the older beholder”, *The Journal of Genetic Psychology*, vol. 172, no. 2, pp. 162–175, Apr. 2011, ISSN: 1940-0896. DOI: 10.1080/00221325.2010.526154. [Online]. Available: <http://dx.doi.org/10.1080/00221325.2010.526154> (cit. on p. 35).
- [FCT06] P. Foos, M. Clark, and D. Terrell, “Adult age, gender, and race group differences in images of aging”, *The Journal of Genetic Psychology*, vol. 167, no. 3, pp. 309–325, Sep. 2006, ISSN: 0022-1325. DOI: 10.3200/gntp.167.3.309-325. [Online]. Available: <http://dx.doi.org/10.3200/GNTP.167.3.309-325> (cit. on pp. 35, 40, 158).
- [FDL+25] G. Filandrianos, A. Dimitriou, M. Lymperaïou, K. Thomas, and G. Stamou, “Bias beware: The impact of cognitive biases on llm-driven product recommendations”, *arXiv preprint arXiv:2502.01349*, 2025 (cit. on p. 58).
- [Fec48] G. T. Fechner, “Elements of psychophysics, 1860.”, 1948 (cit. on p. 19).
- [Fei14] G. A. Feingold, “The influence of environment on identification of persons and things”, *Journal of the American Institute of Criminal Law and Criminology*, vol. 5, no. 1, p. 39, May 1914, ISSN: 0885-4173. DOI: 10.2307/1133283. [Online]. Available: <http://dx.doi.org/10.2307/1133283> (cit. on p. 24).
- [Fei92] A. Feingold, “Good-looking people are not what we think”, *Psychological Bulletin*, vol. 111, no. 2, pp. 304–341, Mar. 1992. DOI: 10.1037/0033-2909.111.2.304. [Online]. Available: <https://doi.org/10.1037/0033-2909.111.2.304> (cit. on p. 95).
- [Fes54] L. Festinger, “A theory of social comparison processes”, *Human Relations*, vol. 7, no. 2, pp. 117–140, May 1954, ISSN: 1741-282X. DOI: 10.1177/001872675400700202. [Online]. Available: <http://dx.doi.org/10.1177/001872675400700202> (cit. on p. 25).
- [Fis28] I. Fisher, *The money illusion*. Longmans, Green, 1928 (cit. on p. 25).

- [FK24] K. Fraser and S. Kiritchenko, “Examining gender and racial bias in large vision–language models using a novel dataset of parallel images”, in *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, Y. Graham and M. Purver, Eds., St. Julian’s, Malta: Association for Computational Linguistics, Mar. 2024, pp. 690–713. [Online]. Available: <https://aclanthology.org/2024.eacl-long.41/> (cit. on pp. 62, 67, 127).
- [FK93] B. L. Fredrickson and D. Kahneman, “Duration neglect in retrospective evaluations of affective episodes.”, *Journal of Personality and Social Psychology*, vol. 65, no. 1, pp. 45–55, 1993, ISSN: 0022-3514. DOI: 10.1037/0022-3514.65.1.45. [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.65.1.45> (cit. on p. 24).
- [FLC19] D. Fernandez, I. Liu, and R. Costilla, “A method for ordinal outcomes: The ordered stereotype model”, *International Journal of Methods in Psychiatric Research*, vol. 28, no. 4, Sep. 2019, ISSN: 1557-0657. DOI: 10.1002/mpr.1801. [Online]. Available: <http://dx.doi.org/10.1002/mpr.1801> (cit. on pp. 31, 34, 46).
- [FM14] F. M. Felisberti and K. Musholt, “Self-face perception: Individual differences and discrepancies associated with mental self-face representation, attractiveness and self-esteem.”, *Psychology & Neuroscience*, vol. 7, no. 2, pp. 65–72, Jan. 2014, ISSN: 1984-3054. DOI: 10.3922/j.psns.2014.013. [Online]. Available: <http://dx.doi.org/10.3922/j.psns.2014.013> (cit. on pp. 41, 87, 159).
- [For49] B. R. Forer, “The fallacy of personal validation: A classroom demonstration of gullibility.”, *The Journal of Abnormal and Social Psychology*, vol. 44, no. 1, pp. 118–123, Jan. 1949, ISSN: 0096-851X. DOI: 10.1037/h0059240. [Online]. Available: <http://dx.doi.org/10.1037/h0059240> (cit. on p. 22).
- [For90] D. R. Forsyth, *Group dynamics*, 1990 (cit. on pp. 16, 25).
- [FOR91] I. H. Frieze, J. E. Olson, and J. Russell, “Attractiveness and income for men and women in management”, *Journal of Applied Social Psychology*, vol. 21, no. 13, pp. 1039–1057, Jul. 1991, ISSN: 1559-1816. DOI: 10.1111/j.1559-1816.1991.tb00458.x. [Online]. Available: <http://dx.doi.org/10.1111/j.1559-1816.1991.tb00458.x> (cit. on p. 28).
- [FSL+78] B. Fischhoff, P. Slovic, S. Lichtenstein, S. Read, and B. Combs, “How safe is safe enough? a psychometric study of attitudes towards technological risks and benefits”, *Policy Sciences*, vol. 9, no. 2, pp. 127–152, Apr. 1978. DOI: 10.1007/bf00143739. [Online]. Available: <https://doi.org/10.1007/bf00143739> (cit. on p. 12).
- [GB09] G. Gigerenzer and H. Brighton, “Homo heuristicus: Why biased minds make better inferences”, *Topics in cognitive science*, vol. 1, no. 1, pp. 107–143, 2009. [Online]. Available: <https://doi.org/10.1111/j.1756-8765.2008.01006.x> (cit. on p. 50).

- [GBL86] D. C. Gilmore, T. A. Beehr, and K. G. Love, “Effects of applicant sex, applicant physical attractiveness, type of rater and type of job on interview decisions”, *Journal of Occupational Psychology*, vol. 59, no. 2, pp. 103–109, Jun. 1986. DOI: [10.1111/j.2044-8325.1986.tb00217.x](https://doi.org/10.1111/j.2044-8325.1986.tb00217.x). [Online]. Available: <https://doi.org/10.1111/j.2044-8325.1986.tb00217.x> (cit. on p. 37).
- [GBPW00] D. T. Gilbert, R. P. Brown, E. C. Pinel, and T. D. Wilson, “The illusion of external agency.”, *Journal of Personality and Social Psychology*, vol. 79, no. 5, pp. 690–700, 2000, ISSN: 0022-3514. DOI: [10.1037/0022-3514.79.5.690](https://doi.org/10.1037/0022-3514.79.5.690). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.79.5.690> (cit. on p. 21).
- [GCR19] M. J. González, C. Cortina, and J. Rodríguez, “The role of gender stereotypes in hiring: A field experiment”, *European Sociological Review*, vol. 35, no. 2, pp. 187–204, Jan. 2019, ISSN: 1468-2672. DOI: [10.1093/esr/jcy055](https://doi.org/10.1093/esr/jcy055). [Online]. Available: <http://dx.doi.org/10.1093/esr/jcy055> (cit. on pp. 41, 159).
- [GDS+25] A. Gulati, M. D’Incà, N. Sebe, B. Lepri, and N. Oliver, *Beauty and the bias: Exploring the impact of attractiveness on multimodal large language models*, 2025. arXiv: [2504.16104](https://arxiv.org/abs/2504.16104) [cs.CY]. [Online]. Available: <https://arxiv.org/abs/2504.16104> (cit. on pp. 6, 55, 162).
- [GG16] J. L. Gibson and J. S. Gore, “Is he a hero or a weirdo? how norm violations influence the halo effect”, *Gender Issues*, vol. 33, no. 4, pp. 299–310, Sep. 2016. DOI: [10.1007/s12147-016-9173-6](https://doi.org/10.1007/s12147-016-9173-6). [Online]. Available: <https://doi.org/10.1007/s12147-016-9173-6> (cit. on p. 12).
- [GG96] G. Gigerenzer and D. G. Goldstein, “Reasoning the fast and frugal way: Models of bounded rationality.”, *Psychological Review*, vol. 103, no. 4, pp. 650–669, Oct. 1996. DOI: [10.1037/0033-295x.103.4.650](https://doi.org/10.1037/0033-295x.103.4.650). [Online]. Available: <https://doi.org/10.1037/0033-295x.103.4.650> (cit. on pp. 16, 25).
- [GHA+24] L. Girrbach, Y. Huang, S. Alaniz, T. Darrell, and Z. Akata, “Revealing and reducing gender biases in vision and language assistants (vlas)”, *arXiv preprint arXiv:2410.19314*, 2024. [Online]. Available: <https://doi.org/10.48550/arXiv.2410.19314> (cit. on p. 59).
- [Gil08] T. Gilovich, *How we know what isn’t so*. Simon and Schuster, 2008 (cit. on p. 20).
- [Gil16] P. Gillingham, “Predictive risk modelling to prevent child maltreatment and other adverse outcomes for service users: Inside the ‘black box’ of machine learning”, *The British Journal of Social Work*, vol. 46, no. 4, pp. 1044–1058, 2016 (cit. on p. 49).
- [Gil21] R. Gill, “Changing the perfect picture: Smartphones, social media and appearance pressures”, City University of London, 2021. [Online]. Available: https://www.city.ac.uk/_data/assets/pdf_file/0005/597209/Parliament-Report-web.pdf (cit. on pp. 41, 87, 159).

- [GJ81] J. A. Graham and A. J. Jouhar, “The effects of cosmetics on person perception”, *International Journal of Cosmetic Science*, vol. 3, no. 5, pp. 199–210, Oct. 1981, ISSN: 1468-2494. DOI: [10.1111/j.1467-2494.1981.tb00283.x](https://doi.org/10.1111/j.1467-2494.1981.tb00283.x). [Online]. Available: <http://dx.doi.org/10.1111/j.1467-2494.1981.tb00283.x> (cit. on p. 28).
- [GKG+23] P. Ghadekar, A. Kabra, K. Gangwal, A. Kinage, K. Agarwal, and K. Chaudhari., “A semantic approach for automated hiring using artificial intelligence & computer vision.”, in *2023 IEEE 8th International Conference for Convergence in Technology (I2CT)*, 2023, pp. 1–7 (cit. on p. 49).
- [GKYS07] A. H. Gutchess, E. A. Kensinger, C. Yoon, and D. L. Schacter, “Ageing and the self-reference effect in memory”, *Memory*, vol. 15, no. 8, pp. 822–837, Nov. 2007. DOI: [10.1080/09658210701701394](https://doi.org/10.1080/09658210701701394). [Online]. Available: <https://doi.org/10.1080/09658210701701394> (cit. on p. 14).
- [GLLA17] J. Gorbova, I. Lusi, A. Litvin, and G. Anbarjafari., “Automated screening of job candidate based on multimodal video processing.”, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Jul. 2017 (cit. on p. 49).
- [GLLO23] A. Gulati, M. A. Lozano, B. Lepri, and N. Oliver, *Biased: Bringing irrationality into automated system design*, 2023. arXiv: [2210.01122](https://arxiv.org/abs/2210.01122) [cs.HC]. [Online]. Available: <https://arxiv.org/abs/2210.01122> (cit. on pp. 5, 7, 162).
- [GLO24] A. Gulati, B. Lepri, and N. Oliver, *Lookism: The overlooked bias in computer vision*, 2024. arXiv: [2408.11448](https://arxiv.org/abs/2408.11448) [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2408.11448> (cit. on pp. 5, 49, 162).
- [GLSE21] G. Gabrieli, A. Lee, P. Setoh, and G. Esposito, “An analysis of the generalizability and stability of the halo effect during the covid-19 pandemic outbreak”, *Frontiers in Psychology*, vol. 12, 2021. DOI: [10.3389/fpsyg.2021.631871](https://doi.org/10.3389/fpsyg.2021.631871). [Online]. Available: <https://doi.org/10.3389/fpsyg.2021.631871> (cit. on pp. 28, 40, 158).
- [GMF+24a] A. Gulati, M. Martínez-Garcia, D. Fernández, M. A. Lozano, B. Lepri, and N. Oliver, “What is beautiful is still good: The attractiveness halo effect in the era of beauty filters”, *Royal Society Open Science*, vol. 11, no. 11, Nov. 2024, ISSN: 2054-5703. DOI: [10.1098/rsos.240882](https://doi.org/10.1098/rsos.240882). [Online]. Available: <http://dx.doi.org/10.1098/rsos.240882> (cit. on pp. 5, 27, 162).
- [GMF+24b] A. Gulati, M. Martínez-Garcia, D. Fernández, M. A. Lozano, B. Lepri, and N. Oliver, “What is beautiful is still good: The attractiveness halo effect in the era of beauty filters”, *Royal Society Open Science*, vol. 11, no. 11, Nov. 2024, ISSN: 2054-5703. DOI: [10.1098/rsos.240882](https://doi.org/10.1098/rsos.240882). [Online]. Available: <http://dx.doi.org/10.1098/rsos.240882> (cit. on pp. 56, 60, 68–71).
- [GML13] J. Golle, F. W. Mast, and J. S. Lobmaier, “Something to smile about: The interrelationship between attractiveness and emotional expression”, *Cognition and Emotion*, vol. 28, no. 2, pp. 298–310, Jul. 2013, ISSN: 1464-0600. DOI: [10.1080/02699931.2013.817383](https://doi.org/10.1080/02699931.2013.817383). [Online]. Available: <http://dx.doi.org/10.1080/02699931.2013.817383> (cit. on pp. 27, 33, 42, 46, 56).

- [GMS00] T. Gilovich, V. H. Medvec, and K. Savitsky, “The spotlight effect in social judgment: An egocentric bias in estimates of the salience of one’s own actions and appearance.”, *Journal of Personality and Social Psychology*, vol. 78, no. 2, pp. 211–222, 2000, ISSN: 0022-3514. DOI: [10.1037/0022-3514.78.2.211](https://doi.org/10.1037/0022-3514.78.2.211). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.78.2.211> (cit. on p. 22).
- [GOL24] A. Gulati, N. Oliver, and B. Lepri, *The beauty survey*, <https://zenodo.org/doi/10.5281/zenodo.13836854>, 2024. DOI: [10.5281/zenodo.13836854](https://doi.org/10.5281/zenodo.13836854). [Online]. Available: <https://zenodo.org/doi/10.5281/zenodo.13836854> (cit. on p. 45).
- [Gol97] E. Gold, “The gambler’s fallacy”, English, Ph.D. dissertation, Carnegie Mellon University, 1997, p. 216, ISBN: 978-0-591-52256-3. [Online]. Available: <https://www.proquest.com/dissertations-theses/gamblers-fallacy/docview/304364133/se-2> (cit. on p. 11).
- [GORS09] S. Gächter, H. Orzen, E. Renner, and C. Starmer, “Are experimental economists prone to framing effects? a natural field experiment”, *Journal of Economic Behavior & Organization*, vol. 70, no. 3, pp. 443–446, Jun. 2009. DOI: [10.1016/j.jebo.2007.11.003](https://doi.org/10.1016/j.jebo.2007.11.003). [Online]. Available: <https://doi.org/10.1016/j.jebo.2007.11.003> (cit. on p. 9).
- [GPM+14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets”, *Advances in neural information processing systems*, vol. 27, 2014 (cit. on p. 52).
- [GPT82] B. J. Guise, C. H. Pollans, and I. D. Turkat, “Effects of physical attractiveness on perception of social skill”, *Perceptual and Motor Skills*, vol. 54, no. 3_suppl, pp. 1039–1042, Jun. 1982. DOI: [10.2466/pms.1982.54.3c.1039](https://doi.org/10.2466/pms.1982.54.3c.1039). [Online]. Available: <https://doi.org/10.2466/pms.1982.54.3c.1039> (cit. on pp. 28, 29, 34, 38, 39, 41, 60, 156, 158).
- [GR91] D. W. Griffin and L. Ross, “Subjective construal, social inference, and human misunderstanding”, in *Advances in Experimental Social Psychology Volume 24*. Elsevier, 1991, pp. 319–359. DOI: [10.1016/S0065-2601\(08\)60333-0](https://doi.org/10.1016/S0065-2601(08)60333-0). [Online]. Available: [http://dx.doi.org/10.1016/S0065-2601\(08\)60333-0](http://dx.doi.org/10.1016/S0065-2601(08)60333-0) (cit. on p. 22).
- [Gre80] A. G. Greenwald, “The totalitarian ego: Fabrication and revision of personal history.”, *American Psychologist*, vol. 35, no. 7, pp. 603–618, Jul. 1980, ISSN: 0003-066X. DOI: [10.1037/0003-066X.35.7.603](https://doi.org/10.1037/0003-066X.35.7.603). [Online]. Available: <http://dx.doi.org/10.1037/0003-066X.35.7.603> (cit. on p. 23).
- [GVT85] T. Gilovich, R. Vallone, and A. Tversky, “The hot hand in basketball: On the misperception of random sequences”, *Cognitive Psychology*, vol. 17, no. 3, pp. 295–314, Jul. 1985, ISSN: 0010-0285. DOI: [10.1016/0010-0285\(85\)90010-6](https://doi.org/10.1016/0010-0285(85)90010-6). [Online]. Available: [http://dx.doi.org/10.1016/0010-0285\(85\)90010-6](http://dx.doi.org/10.1016/0010-0285(85)90010-6) (cit. on p. 20).

- [GZY+13] J. Gong, Y. Zhang, Z. Yang, Y. Huang, J. Feng, and W. Zhang, “The framing effect in medical decision-making: A review of the literature”, *Psychology, Health & Medicine*, vol. 18, no. 6, pp. 645–653, Dec. 2013. DOI: [10.1080/13548506.2013.766352](https://doi.org/10.1080/13548506.2013.766352). [Online]. Available: <https://doi.org/10.1080/13548506.2013.766352> (cit. on p. 9).
- [HAO18] W. H. Hampton, N. Asadi, and I. R. Olson, “Good things for those who wait: Predictive modeling highlights importance of delay discounting for income attainment”, *Frontiers in Psychology*, vol. 9, Sep. 2018. DOI: [10.3389/fpsyg.2018.01545](https://doi.org/10.3389/fpsyg.2018.01545). [Online]. Available: <https://doi.org/10.3389/fpsyg.2018.01545> (cit. on p. 11).
- [Har06] L. Harrison, *The validity of self-reported drug use in survey research: An overview and critique of research methods*. national institute of drug abuse monograph 167, 2006 (cit. on p. 13).
- [HAS+12] U. Hess, R. B. Adams, A. Simard, M. T. Stevenson, and R. E. Kleck, “Smiling and sad wrinkles: Age-related changes in the face and the perception of emotions and intentions”, *Journal of Experimental Social Psychology*, vol. 48, no. 6, pp. 1377–1380, Nov. 2012, ISSN: 0022-1031. DOI: [10.1016/j.jesp.2012.05.018](https://doi.org/10.1016/j.jesp.2012.05.018). [Online]. Available: <http://dx.doi.org/10.1016/j.jesp.2012.05.018> (cit. on pp. 40, 121, 158).
- [HAZ+23] S. M. Hall, F. G. Abrantes, H. Zhu, G. Sodunke, A. Shtedritski, and H. R. Kirk, *Visogender: A dataset for benchmarking gender bias in image-text pronoun resolution*, 2023. arXiv: [2306.12424](https://arxiv.org/abs/2306.12424) [cs.CV] (cit. on p. 58).
- [HB00] M. G. Haselton and D. M. Buss, “Error management theory: A new perspective on biases in cross-sex mind reading.”, *Journal of Personality and Social Psychology*, vol. 78, no. 1, pp. 81–91, 2000, ISSN: 0022-3514. DOI: [10.1037/0022-3514.78.1.81](https://doi.org/10.1037/0022-3514.78.1.81). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.78.1.81> (cit. on p. 19).
- [HB94] D. S. Hamermesh and J. Biddle, “Beauty and the labor market”, *The American Economic Review*, vol. 84, no. 5, pp. 1174–1194, Dec. 1994. [Online]. Available: <https://www.jstor.org/stable/2117767> (cit. on p. 28).
- [HCP+95] J. R. Hebert, L. Clemow, L. Pbert, I. S. Ockene, and J. K. Ockene, “Social desirability bias in dietary self-report may compromise the validity of dietary intake measures”, *International Journal of Epidemiology*, vol. 24, no. 2, pp. 389–398, 1995. DOI: [10.1093/ije/24.2.389](https://doi.org/10.1093/ije/24.2.389). [Online]. Available: <https://doi.org/10.1093/ije/24.2.389> (cit. on p. 13).
- [Hea99] C. Heath, “On the social psychology of agency relationships: Lay theories of motivation overemphasize extrinsic incentives”, *Organizational Behavior and Human Decision Processes*, vol. 78, no. 1, pp. 25–62, Apr. 1999, ISSN: 0749-5978. DOI: [10.1006/obhd.1999.2826](https://doi.org/10.1006/obhd.1999.2826). [Online]. Available: <http://dx.doi.org/10.1006/obhd.1999.2826> (cit. on p. 19).

- [Hei01] M. E. Heilman, “Description and prescription: How gender stereotypes prevent women’s ascent up the organizational ladder”, *Journal of Social Issues*, vol. 57, no. 4, pp. 657–674, Jan. 2001, ISSN: 1540-4560. DOI: [10.1111/0022-4537.00234](https://doi.org/10.1111/0022-4537.00234). [Online]. Available: <http://dx.doi.org/10.1111/0022-4537.00234> (cit. on pp. 41, 159).
- [HFBK24] P. Howard, K. C. Fraser, A. Bhiwandiwalla, and S. Kiritchenko, *Uncovering bias in large vision-language models at scale with counterfactuals*, 2024. arXiv: 2405.20152 [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2405.20152> (cit. on pp. 3, 50, 51, 59, 74, 156).
- [HG17] R. Hertwig and T. Grüne-Yanoff, “Nudging and boosting: Steering or empowering good decisions”, *Perspectives on Psychological Science*, vol. 12, no. 6, pp. 973–986, Aug. 2017, ISSN: 1745-6924. DOI: [10.1177/1745691617702496](https://doi.org/10.1177/1745691617702496). [Online]. Available: <http://dx.doi.org/10.1177/1745691617702496> (cit. on pp. 4, 161).
- [HGT77] L. Hasher, D. Goldstein, and T. Toppino, “Frequency and the conference of referential validity”, *Journal of Verbal Learning and Verbal Behavior*, vol. 16, no. 1, pp. 107–112, Feb. 1977, ISSN: 0022-5371. DOI: [10.1016/s0022-5371\(77\)80012-1](https://doi.org/10.1016/s0022-5371(77)80012-1). [Online]. Available: [http://dx.doi.org/10.1016/s0022-5371\(77\)80012-1](http://dx.doi.org/10.1016/s0022-5371(77)80012-1) (cit. on p. 20).
- [HGW+24] L. Haliburton, S. Ghebremedhin, R. Welsch, A. Schmidt, and S. Mayer, “Investigating labeler bias in face annotation for machine learning”, in *HHAi 2024: Hybrid Human AI Systems for the Social Good*. IOS Press, Jun. 2024, pp. 145–161, ISBN: 9781643685229. DOI: [10.3233/faia240191](https://doi.org/10.3233/faia240191). [Online]. Available: <http://dx.doi.org/10.3233/FAIA240191> (cit. on p. 58).
- [HH96] V. Henri and C. Henri, “Enquête sur les premiers souvenirs de l’enfance”, *L’année psychologique*, vol. 3, no. 1, pp. 184–198, 1896, ISSN: 0003-5033. DOI: [10.3406/psy.1896.1831](https://doi.org/10.3406/psy.1896.1831). [Online]. Available: <http://dx.doi.org/10.3406/psy.1896.1831> (cit. on p. 25).
- [HHG+24] R. Hada, S. Husain, V. Gumma, H. Diddee, A. Yadavalli, A. Seth, N. Kulkarni, U. Gadiraju, A. Vashistha, V. Seshadri, and K. Bali, “Akal badi ya bias: An exploratory study of gender bias in hindi language technology”, in *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, ser. FAccT ’24, ACM, Jun. 2024, pp. 1926–1939. DOI: [10.1145/3630106.3659017](https://doi.org/10.1145/3630106.3659017). [Online]. Available: <http://dx.doi.org/10.1145/3630106.3659017> (cit. on p. 57).
- [Hil12] M. Hilbert, “Toward a synthesis of cognitive biases: How noisy information processing can bias human decision making.”, *Psychological Bulletin*, vol. 138, no. 2, pp. 211–237, Mar. 2012, ISSN: 0033-2909. DOI: [10.1037/a0025940](https://doi.org/10.1037/a0025940). [Online]. Available: <http://dx.doi.org/10.1037/a0025940> (cit. on p. 21).
- [HL23] D. E. Han and S. M. Laurent, “Beautiful seems good, but perhaps not in every way: Linking attractiveness to moral evaluation through perceived vanity.”, *Journal of Personality and Social Psychology*, vol. 124, no. 2, pp. 264–286, Feb. 2023, ISSN: 0022-3514. DOI: [10.1037/pspa0000317](https://doi.org/10.1037/pspa0000317). [Online]. Available: <http://dx.doi.org/10.1037/pspa0000317> (cit. on pp. 29, 39, 157).

- [HML+23] P. Howard, A. Madasu, T. Le, G. L. Moreno, A. Bhiwandiwalla, and V. Lal., “Probing and mitigating intersectional social biases in vision-language models with counterfactual examples.”, *arXiv preprint arXiv:2312.00825*, 2023 (cit. on p. 52).
- [HN09] B. K. Hayes and B. R. Newell, “Induction with uncertain categories: When do people consider the category alternatives?”, *Memory & Cognition*, vol. 37, no. 6, pp. 730–743, Sep. 2009. DOI: [10.3758/mc.37.6.730](https://doi.org/10.3758/mc.37.6.730). [Online]. Available: <https://doi.org/10.3758/mc.37.6.730> (cit. on p. 9).
- [HPP82] J. Huber, J. W. Payne, and C. Puto, “Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis”, *Journal of Consumer Research*, vol. 9, no. 1, p. 90, Jun. 1982. DOI: [10.1086/208899](https://doi.org/10.1086/208899). [Online]. Available: <https://doi.org/10.1086/208899> (cit. on pp. 9, 19).
- [HS85] M. E. Heilman and M. H. Stopeck, “Being attractive, advantage or disadvantage? performance-based evaluations and recommended personnel actions as a function of appearance, sex, and job type”, *Organizational Behavior and Human Decision Processes*, vol. 35, no. 2, pp. 202–215, Apr. 1985. DOI: [10.1016/0749-5978\(85\)90035-4](https://doi.org/10.1016/0749-5978(85)90035-4). [Online]. Available: [https://doi.org/10.1016/0749-5978\(85\)90035-4](https://doi.org/10.1016/0749-5978(85)90035-4) (cit. on p. 37).
- [HSC03] M. Hosoda, E. F. Stone-Romero, and G. Coats, “The effects of physical attractiveness on job-related outcomes: A meta-analysis of experimental studies”, *Personnel Psychology*, vol. 56, no. 2, pp. 431–462, Jun. 2003. DOI: [10.1111/j.1744-6570.2003.tb00157.x](https://doi.org/10.1111/j.1744-6570.2003.tb00157.x). [Online]. Available: <https://doi.org/10.1111/j.1744-6570.2003.tb00157.x> (cit. on pp. 3, 28, 61, 160).
- [Hse98] C. K. Hsee, “Less is better: When low-value options are valued more highly than high-value options”, *Journal of Behavioral Decision Making*, vol. 11, no. 2, pp. 107–121, Jun. 1998, ISSN: 1099-0771. DOI: [10.1002/\(SICI\)1099-0771\(199806\)11:2<107::aid-bdm292>3.0.CO;2-Y](https://doi.org/10.1002/(SICI)1099-0771(199806)11:2<107::aid-bdm292>3.0.CO;2-Y). [Online]. Available: [http://dx.doi.org/10.1002/\(SICI\)1099-0771\(199806\)11:2%3C107::AID-BDM292%3E3.0.CO;2-Y](http://dx.doi.org/10.1002/(SICI)1099-0771(199806)11:2%3C107::AID-BDM292%3E3.0.CO;2-Y) (cit. on p. 19).
- [HvdMG+22] M. Hall, L. van der Maaten, L. Gustafson, M. Jones, and A. Adcock, *A systematic study of bias amplification*, 2022. arXiv: [2201.11706](https://arxiv.org/abs/2201.11706) [cs.LG]. [Online]. Available: <https://arxiv.org/abs/2201.11706> (cit. on pp. 49, 70).
- [HXL+24] W. Hu, Y. Xu, Y. Li, W. Li, Z. Chen, and Z. Tu, “Bliva: A simple multimodal llm for better handling of text-rich visual questions”, *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 3, pp. 2256–2264, Mar. 2024, ISSN: 2159-5399. DOI: [10.1609/aaai.v38i3.27999](https://doi.org/10.1609/aaai.v38i3.27999). [Online]. Available: <http://dx.doi.org/10.1609/aaai.v38i3.27999> (cit. on p. 56).
- [HY14] J. Hu and R. Yu, “The neural correlates of the decoy effect in decisions”, *Frontiers in Behavioral Neuroscience*, vol. 8, Aug. 2014. DOI: [10.3389/fnbeh.2014.00271](https://doi.org/10.3389/fnbeh.2014.00271). [Online]. Available: <https://doi.org/10.3389/fnbeh.2014.00271> (cit. on p. 9).

- [HZG+24] K. Hamidieh, H. Zhang, W. Gerych, T. Hartvigsen, and M. Ghassemi, “Identifying implicit social biases in vision-language models”, in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, vol. 7, 2024, pp. 547–561 (cit. on pp. 58, 59, 63, 129, 130).
- [HZJ+20] P.-S. Huang, H. Zhang, R. Jiang, R. Stanforth, J. Welbl, J. Rae, V. Maini, D. Yogatama, and P. Kohli, “Reducing sentiment bias in language models via counterfactual evaluation”, in *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2020, pp. 65–83 (cit. on p. 57).
- [HZRS16] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition”, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778 (cit. on p. 78).
- [Isa23] C. Isakowitsch, “How augmented reality beauty filters can affect self-perception”, in *Artificial Intelligence and Cognitive Science*. Springer Nature Switzerland, 2023, pp. 239–250, ISBN: 9783031264382. DOI: 10.1007/978-3-031-26438-2_19. [Online]. Available: http://dx.doi.org/10.1007/978-3-031-26438-2_19 (cit. on pp. 30, 41, 87, 159).
- [ISRB24] I. Itzhak, G. Stanovsky, N. Rosenfeld, and Y. Belinkov, “Instructed to bias: Instruction-tuned language models exhibit emergent cognitive bias”, *Transactions of the Association for Computational Linguistics*, vol. 12, pp. 771–785, 2024 (cit. on p. 57).
- [Jac78] L. L. Jacoby, “On interpreting the effects of repetition: Solving a problem versus remembering a solution”, *Journal of Verbal Learning and Verbal Behavior*, vol. 17, no. 6, pp. 649–667, Dec. 1978, ISSN: 0022-5371. DOI: 10.1016/s0022-5371(78)90393-6. [Online]. Available: [http://dx.doi.org/10.1016/S0022-5371\(78\)90393-6](http://dx.doi.org/10.1016/S0022-5371(78)90393-6) (cit. on p. 24).
- [JB94] J. T. Jost and M. R. Banaji, “The role of stereotyping in system-justification and the production of false consciousness”, *British Journal of Social Psychology*, vol. 33, no. 1, pp. 1–27, Mar. 1994, ISSN: 2044-8309. DOI: 10.1111/j.2044-8309.1994.tb01008.x. [Online]. Available: <http://dx.doi.org/10.1111/j.2044-8309.1994.tb01008.x> (cit. on p. 22).
- [JD23] S. Janghorbani and G. De Melo, “Multimodal bias: Introducing a framework for stereotypical bias assessment beyond gender and race in vision language models”, *arXiv preprint arXiv:2303.12734*, 2023 (cit. on p. 59).
- [JH95] B. M. Josiam and J. P. Hobson, “Consumer choice in context: The decoy effect in travel and tourism”, *Journal of Travel Research*, vol. 34, no. 1, pp. 45–50, Jul. 1995. DOI: 10.1177/004728759503400106. [Online]. Available: <https://doi.org/10.1177/004728759503400106> (cit. on p. 9).
- [JHH95] L. A. Jackson, J. E. Hunter, and C. N. Hodge, “Physical attractiveness and intellectual competence: A meta-analytic review”, *Social Psychology Quarterly*, pp. 108–122, 1995. DOI: <https://doi.org/10.2307/2787149> (cit. on pp. 27, 54).

- [JHL93] M. K. Johnson, S. Hashtroudi, and D. S. Lindsay, “Source monitoring.”, *Psychological Bulletin*, vol. 114, no. 1, pp. 3–28, 1993, ISSN: 0033-2909. DOI: [10.1037/0033-2909.114.1.3](https://doi.org/10.1037/0033-2909.114.1.3). [Online]. Available: <http://dx.doi.org/10.1037/0033-2909.114.1.3> (cit. on p. 23).
- [JL69] J. Jecker and D. Landy, “Liking a person as a function of doing him a favour”, *Human Relations*, vol. 22, no. 4, pp. 371–378, Aug. 1969, ISSN: 1741-282X. DOI: [10.1177/001872676902200407](https://doi.org/10.1177/001872676902200407). [Online]. Available: <http://dx.doi.org/10.1177/001872676902200407> (cit. on p. 23).
- [JL97] K. JENNI and G. LOEWENSTEIN, “Explaining the identifiable victim effect”, *Journal of Risk and Uncertainty*, vol. 14, no. 3, pp. 235–257, May 1997, ISSN: 1573-0476. DOI: [10.1023/a:1007740225484](https://doi.org/10.1023/a:1007740225484). [Online]. Available: <http://dx.doi.org/10.1023/A:1007740225484> (cit. on p. 25).
- [JLS+24] Y. Jiang, Z. Li, X. Shen, Y. Liu, M. Backes, and Y. Zhang, “ModSCAN: Measuring stereotypical bias in large vision-language models from vision and language modalities”, *arXiv preprint arXiv:2410.06967*, 2024 (cit. on pp. 63, 129, 130).
- [Joh10] R. Johns, “Likert items and scales”, *Survey question bank: Methods fact sheet*, vol. 1, no. 1, pp. 11–28, 2010. [Online]. Available: https://dam.ukdataservice.ac.uk/media/262829/discover_likertfactsheet.pdf (cit. on p. 46).
- [JOM+19] J. C. S. Jacques Junior, C. Ozcinar, M. Marjanovic, X. Baro, G. Anbarjafari, and S. Escalera., “On the effect of age perception biases for real age regression.”, in *2019 14th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2019)*, IEEE, May 2019 (cit. on pp. 50, 74).
- [JP96] A. Jansari and A. J. Parkin, “Things that go bump in your life: Explaining the reminiscence bump in autobiographical memory.”, *Psychology and Aging*, vol. 11, no. 1, pp. 85–91, Mar. 1996, ISSN: 0882-7974. DOI: [10.1037/0882-7974.11.1.85](https://doi.org/10.1037/0882-7974.11.1.85). [Online]. Available: <http://dx.doi.org/10.1037/0882-7974.11.1.85> (cit. on p. 24).
- [JR94] O. P. John and R. W. Robins, “Accuracy and bias in self-perception: Individual differences in self-enhancement and the role of narcissism.”, *Journal of Personality and Social Psychology*, vol. 66, no. 1, pp. 206–219, 1994, ISSN: 0022-3514. DOI: [10.1037/0022-3514.66.1.206](https://doi.org/10.1037/0022-3514.66.1.206). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.66.1.206> (cit. on p. 96).
- [JS94] H. M. Johnson and C. M. Seifert, “Sources of the continued influence effect: When misinformation in memory affects later inferences.”, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 20, no. 6, pp. 1420–1436, Nov. 1994, ISSN: 0278-7393. DOI: [10.1037/0278-7393.20.6.1420](https://doi.org/10.1037/0278-7393.20.6.1420). [Online]. Available: <http://dx.doi.org/10.1037/0278-7393.20.6.1420> (cit. on p. 24).
- [Kam82] D. Kammer, “Differences in trait ascriptions to self and friend: Unconfounding intensity from variability”, *Psychological Reports*, vol. 51, no. 1, pp. 99–102, Aug. 1982, ISSN: 1558-691X. DOI: [10.2466/pr0.1982.51.1.99](https://doi.org/10.2466/pr0.1982.51.1.99). [Online]. Available: <http://dx.doi.org/10.2466/pr0.1982.51.1.99> (cit. on p. 23).

- [Kan11] S. Kanazawa, “Intelligence and physical attractiveness”, *Intelligence*, vol. 39, no. 1, pp. 7–14, Jan. 2011. DOI: [10.1016/j.intell.2010.11.003](https://doi.org/10.1016/j.intell.2010.11.003). [Online]. Available: <https://doi.org/10.1016/j.intell.2010.11.003> (cit. on pp. 27, 56).
- [Kap78] R. M. Kaplan, “Is beauty talent? sex interaction in the attractiveness halo effect”, *Sex Roles*, vol. 4, no. 2, pp. 195–204, Apr. 1978. DOI: [10.1007/bf00287500](https://doi.org/10.1007/bf00287500). [Online]. Available: <https://doi.org/10.1007/bf00287500> (cit. on p. 37).
- [KB00] J. Kissler and K.-H. Bäuml, “Effects of the beholder’s age on the perception of facial attractiveness”, *Acta Psychologica*, vol. 104, no. 2, pp. 145–166, May 2000, ISSN: 0001-6918. DOI: [10.1016/s0001-6918\(00\)00018-4](https://doi.org/10.1016/s0001-6918(00)00018-4). [Online]. Available: [http://dx.doi.org/10.1016/s0001-6918\(00\)00018-4](http://dx.doi.org/10.1016/s0001-6918(00)00018-4) (cit. on pp. 35, 40, 158).
- [KCF14] K. Kleisner, V. Chvátalová, and J. Flegr, “Perceived intelligence is associated with measured intelligence in men but not women”, *PLoS ONE*, vol. 9, no. 3, B. Fink, Ed., e81237, Mar. 2014. DOI: [10.1371/journal.pone.0081237](https://doi.org/10.1371/journal.pone.0081237). [Online]. Available: <https://doi.org/10.1371/journal.pone.0081237> (cit. on pp. 27, 37).
- [KCY+23] C. Kim, J. Choi, J. Yoon, D. Yoo, and W. Lee., “Fairness-aware multimodal learning in automatic video interview assessment.”, *IEEE Access*, 2023 (cit. on p. 51).
- [KD24] A. Karagianni and M. Doh, “A feminist legal analysis of non-consensual sexualized deepfakes: Contextualizing its impact as ai-generated image-based violence under eu law”, *Porn Studies*, vol. 0, no. 0, pp. 1–18, 2024. DOI: [10.1080/23268743.2024.2408277](https://doi.org/10.1080/23268743.2024.2408277). eprint: <https://doi.org/10.1080/23268743.2024.2408277>. [Online]. Available: <https://doi.org/10.1080/23268743.2024.2408277> (cit. on p. 73).
- [KD99] J. Kruger and D. Dunning, “Unskilled and unaware of it: How difficulties in recognizing one’s own incompetence lead to inflated self-assessments.”, *Journal of Personality and Social Psychology*, vol. 77, no. 6, pp. 1121–1134, 1999. [Online]. Available: <https://doi.org/10.1037/0022-3514.77.6.1121> (cit. on p. 21).
- [KDS23] H. Koteck, R. Dockum, and D. Sun, “Gender bias and stereotypes in large language models”, in *Proceedings of the ACM collective intelligence conference*, 2023, pp. 12–24 (cit. on pp. 57, 67).
- [KF21] Z. Khan and Y. Fu., “One label, one billion faces: Usage and consistency of racial categories in computer vision.”, in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, ser. FAccT ’21, ACM, Mar. 2021 (cit. on p. 50).
- [KFSR93] D. Kahneman, B. L. Fredrickson, C. A. Schreiber, and D. A. Redelmeier, “When more pain is preferred to less: Adding a better end”, *Psychological Science*, vol. 4, no. 6, pp. 401–405, Nov. 1993. DOI: [10.1111/j.1467-9280.1993.tb00589.x](https://doi.org/10.1111/j.1467-9280.1993.tb00589.x). [Online]. Available: <https://doi.org/10.1111/j.1467-9280.1993.tb00589.x> (cit. on pp. 1, 14, 24, 153).

- [KG99] J. Kruger and T. Gilovich, ““naive cynicism” in everyday theories of responsibility assessment: On biased assumptions of bias.”, *Journal of Personality and Social Psychology*, vol. 76, no. 5, pp. 743–753, May 1999, ISSN: 0022-3514. DOI: [10.1037/0022-3514.76.5.743](https://doi.org/10.1037/0022-3514.76.5.743). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.76.5.743> (cit. on p. 23).
- [KGKK18] M. J. Kinsey, S. M. V. Gwynne, E. D. Kuligowski, and M. Kinatader, “Cognitive biases within decision making during fire evacuations”, *Fire Technology*, vol. 55, no. 2, pp. 465–485, Mar. 2018. [Online]. Available: <https://doi.org/10.1007/s10694-018-0708-0> (cit. on pp. 2, 7, 154).
- [KJ21a] K. Karkkainen and J. Joo, “Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation”, in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 1548–1558 (cit. on pp. 74, 78).
- [KJ21b] K. Karkkainen and J. Joo., “Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation.”, in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, Jan. 2021, pp. 1548–1558 (cit. on pp. 50, 74).
- [KJV+21] H. R. Kirk, Y. Jun, F. Volpin, H. Iqbal, E. Benussi, F. Dreyer, A. Shtedritski, and Y. Asano, “Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models”, *Advances in neural information processing systems*, vol. 34, pp. 2611–2624, 2021 (cit. on p. 58).
- [KKLC16] Y.-H. Kim, H. Kwon, J. Lee, and C.-Y. Chiu, “Why do people overestimate or underestimate their abilities? a cross-culturally valid model of cognitive and motivational processes in self-assessment biases”, *Journal of Cross-Cultural Psychology*, vol. 47, no. 9, pp. 1201–1216, Jul. 2016, ISSN: 1552-5422. DOI: [10.1177/0022022116661243](https://doi.org/10.1177/0022022116661243). [Online]. Available: <http://dx.doi.org/10.1177/0022022116661243> (cit. on p. 96).
- [KKM23] J. R. Kunst, J. Kirkøen, and O. Mohamdain, “Hacking attractiveness biases in hiring? the role of beautifying photo-filters”, *Management Decision*, vol. 61, no. 4, pp. 924–943, Apr. 2023, ISSN: 0025-1747. DOI: [10.1108/md-06-2021-0747](https://doi.org/10.1108/md-06-2021-0747). [Online]. Available: <http://dx.doi.org/10.1108/MD-06-2021-0747> (cit. on pp. 28, 30, 37, 38, 40, 71, 156, 158).
- [KKT90] D. Kahneman, J. L. Knetsch, and R. H. Thaler, “Experimental tests of the endowment effect and the coase theorem”, *Journal of Political Economy*, vol. 98, no. 6, pp. 1325–1348, Dec. 1990, ISSN: 1537-534X. DOI: [10.1086/261737](https://doi.org/10.1086/261737). [Online]. Available: <http://dx.doi.org/10.1086/261737> (cit. on p. 21).
- [KKT91] D. Kahneman, J. L. Knetsch, and R. H. Thaler, “Anomalies: The endowment effect, loss aversion, and status quo bias”, *Journal of Economic Perspectives*, vol. 5, no. 1, pp. 193–206, Feb. 1991. DOI: [10.1257/jep.5.1.193](https://doi.org/10.1257/jep.5.1.193). [Online]. Available: <https://doi.org/10.1257/jep.5.1.193> (cit. on p. 15).

- [KLR+24] R. Koo, M. Lee, V. Raheja, J. I. Park, Z. M. Kim, and D. Kang, “Benchmarking cognitive biases in large language models as evaluators”, in *Findings of the Association for Computational Linguistics: ACL 2024*, L.-W. Ku, A. Martins, and V. Srikumar, Eds., Bangkok, Thailand: Association for Computational Linguistics, Aug. 2024, pp. 517–545. DOI: [10.18653/v1/2024.findings-acl.29](https://doi.org/10.18653/v1/2024.findings-acl.29). [Online]. Available: <https://aclanthology.org/2024.findings-acl.29/> (cit. on p. 58).
- [KP99] A. J. Knight and W. V. Parr, “Age as a factor in judgments of wisdom and creativity”, *New Zealand Journal of Psychology*, vol. 28, no. 1, p. 37, 1999. [Online]. Available: <https://psycnet.apa.org/record/1999-11171-005> (cit. on pp. 40, 158).
- [KR75] R. E. Kleck and C. Rubenstein, “Physical attractiveness, perceived attitude similarity, and interpersonal attraction in an opposite-sex encounter.”, *Journal of Personality and Social Psychology*, vol. 31, no. 1, pp. 107–114, Jan. 1975. DOI: [10.1037/h0076243](https://doi.org/10.1037/h0076243). [Online]. Available: <https://doi.org/10.1037/h0076243> (cit. on pp. 29, 39, 156, 157).
- [Kru99] J. Kruger, “Lake wobegon be gone! the “below-average effect” and the egocentric nature of comparative ability judgments.”, *Journal of Personality and Social Psychology*, vol. 77, no. 2, pp. 221–232, 1999, ISSN: 0022-3514. DOI: [10.1037/0022-3514.77.2.221](https://doi.org/10.1037/0022-3514.77.2.221). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.77.2.221> (cit. on p. 22).
- [KRWS20] A. Köchling, S. Riazzy, M. C. Wehner, and K. Simbeck., “Highly accurate, but still discriminatory: A fairness evaluation of algorithmic video analysis in the recruitment context.”, *Business & Information Systems Engineering*, vol. 63, no. 1, pp. 39–54, Nov. 2020 (cit. on p. 51).
- [KS22] T. Kim and H. Song, “Communicating the limitations of AI: The effect of message framing and ownership on trust in artificial intelligence”, *International Journal of Human–Computer Interaction*, vol. 39, no. 4, pp. 790–800, Apr. 2022. DOI: [10.1080/10447318.2022.2049134](https://doi.org/10.1080/10447318.2022.2049134). [Online]. Available: <https://doi.org/10.1080/10447318.2022.2049134> (cit. on p. 10).
- [KSK23] M. Kamruzzaman, M. M. I. Shovon, and G. L. Kim., “Investigating subtler biases in llms: Ageism, beauty, institutional, and nationality bias in generative models.”, *arXiv preprint arXiv:2309.08902*, 2023 (cit. on p. 51).
- [KT13] D. Kahneman and A. Tversky, “Prospect theory: An analysis of decision under risk”, in *Handbook of the Fundamentals of Financial Decision Making*, WORLD SCIENTIFIC, Jun. 2013, pp. 99–127. [Online]. Available: https://doi.org/10.1142/9789814417358_0006 (cit. on p. 19).
- [KT73] D. Kahneman and A. Tversky, “On the psychology of prediction.”, *Psychological Review*, vol. 80, no. 4, pp. 237–251, Jul. 1973, ISSN: 0033-295X. DOI: [10.1037/h0034747](https://doi.org/10.1037/h0034747). [Online]. Available: <http://dx.doi.org/10.1037/h0034747> (cit. on p. 20).
- [KT79] D. Kahneman and A. Tversky, “Prospect theory: An analysis of decision under risk”, *Econometrica*, vol. 47, no. 2, p. 263, Mar. 1979. [Online]. Available: <https://doi.org/10.2307/1914185> (cit. on pp. 21, 58).

- [KT82] K. M. Korthase and I. Trenholme, “Perceived age and perceived physical attractiveness”, *Perceptual and Motor Skills*, vol. 54, no. 3_suppl, pp. 1251–1258, Jun. 1982, ISSN: 1558-688X. DOI: [10.2466/pms.1982.54.3c.1251](https://doi.org/10.2466/pms.1982.54.3c.1251). [Online]. Available: <http://dx.doi.org/10.2466/pms.1982.54.3c.1251> (cit. on pp. 40, 158).
- [KvdAZ14] D. Kliger, M. J. van den Assem, and R. C. Zwinkels, “Empirical behavioral finance”, *Journal of Economic Behavior & Organization*, vol. 107, pp. 421–427, Nov. 2014. DOI: [10.1016/j.jebo.2014.10.012](https://doi.org/10.1016/j.jebo.2014.10.012). [Online]. Available: <https://doi.org/10.1016/j.jebo.2014.10.012> (cit. on p. 16).
- [KW04] K. M. Kniffin and D. S. Wilson, “The effect of nonphysical traits on the perception of physical attractiveness”, *Evolution and Human Behavior*, vol. 25, no. 2, pp. 88–101, Mar. 2004, ISSN: 1090-5138. DOI: [10.1016/S1090-5138\(04\)00006-6](https://doi.org/10.1016/S1090-5138(04)00006-6). [Online]. Available: [http://dx.doi.org/10.1016/S1090-5138\(04\)00006-6](http://dx.doi.org/10.1016/S1090-5138(04)00006-6) (cit. on p. 38).
- [KW13] D. P. Kingma and M. Welling., “Auto-encoding variational bayes.”, *arXiv preprint arXiv:1312.6114*, 2013 (cit. on p. 52).
- [KYE24] A. Kumar, S. Yunusov, and A. Emami., “Subtle biases need subtler measures: Dual metrics for evaluating representative and affinity bias in large language models.”, *arXiv preprint arXiv:2405.14555*, 2024 (cit. on p. 50).
- [LAMJ23] A. S. Luccioni, C. Akiki, M. Mitchell, and Y. Jernite, “Stable bias: Analyzing societal representations in diffusion models”, *arXiv preprint arXiv:2303.11408*, 2023 (cit. on p. 74).
- [LBH81] G. W. Lucker, W. E. Beane, and R. L. Helmreich, “The strength of the halo effect in physical attractiveness research”, *The Journal of Psychology*, vol. 107, no. 1, pp. 69–75, Jan. 1981. DOI: [10.1080/00223980.1981.9915206](https://doi.org/10.1080/00223980.1981.9915206). [Online]. Available: <https://doi.org/10.1080/00223980.1981.9915206> (cit. on pp. 39, 157).
- [LC20] C. Lavrence and C. Cambre, ““do i look like my selfie?”: Filters and the digital-forensic gaze”, *Social Media Society*, vol. 6, no. 4, p. 205 630 512 095 518, Oct. 2020, ISSN: 2056-3051. DOI: [10.1177/2056305120955182](https://doi.org/10.1177/2056305120955182). [Online]. Available: <http://dx.doi.org/10.1177/2056305120955182> (cit. on pp. 30, 41, 87, 159).
- [LDZ+25] Y. Liu, H. Duan, Y. Zhang, B. Li, S. Zhang, W. Zhao, Y. Yuan, J. Wang, C. He, Z. Liu, *et al.*, “Mmbench: Is your multi-modal model an all-around player?”, in *European conference on computer vision*, Springer, 2025, pp. 216–233 (cit. on p. 61).
- [Len23] R. V. Lenth, *Emmeans: Estimated marginal means, aka least-squares means*, R package version 1.9.0, 2023. [Online]. Available: <https://CRAN.R-project.org/package=emmeans> (cit. on p. 36).
- [Lev60] H. Levene, “Robust tests for equality of variances”, *Contributions to probability and statistics*, pp. 278–292, 1960. [Online]. Available: <https://cir.nii.ac.jp/crid/1573950400526848896> (cit. on p. 37).

- [Lew27] K. Lewin, “Untersuchungen zur handlungs-und affektpsychologie”, *Psychologische Forschung*, vol. 9, no. 1, pp. 1–85, 1927 (cit. on p. 24).
- [LF77] S. Lichtenstein and B. Fischhoff, “Do those who know more also know more about how much they know?”, *Organizational Behavior and Human Performance*, vol. 20, no. 2, pp. 159–183, Dec. 1977, ISSN: 0030-5073. DOI: [10.1016/0030-5073\(77\)90001-0](https://doi.org/10.1016/0030-5073(77)90001-0). [Online]. Available: [http://dx.doi.org/10.1016/0030-5073\(77\)90001-0](http://dx.doi.org/10.1016/0030-5073(77)90001-0) (cit. on p. 19).
- [LJ89] D. S. Lindsay and M. K. Johnson, “The eyewitness suggestibility effect and memory for source”, *Memory & Cognition*, vol. 17, no. 3, pp. 349–358, May 1989, ISSN: 1532-5946. DOI: [10.3758/bf03198473](https://doi.org/10.3758/bf03198473). [Online]. Available: <http://dx.doi.org/10.3758/bf03198473> (cit. on p. 24).
- [LKR+00] J. H. Langlois, L. Kalakanis, A. J. Rubenstein, A. Larson, M. Hallam, and M. Smoot, “Maxims or myths of beauty? a meta-analytic and theoretical review.”, *Psychological Bulletin*, vol. 126, no. 3, pp. 390–423, 2000. DOI: [10.1037/0033-2909.126.3.390](https://doi.org/10.1037/0033-2909.126.3.390). [Online]. Available: <https://doi.org/10.1037/0033-2909.126.3.390> (cit. on pp. 40, 158).
- [LLWL23] H. Liu, C. Li, Q. Wu, and Y. J. Lee, *Visual instruction tuning*, 2023 (cit. on p. 60).
- [LLWT15] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep learning face attributes in the wild”, in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Dec. 2015 (cit. on p. 152).
- [LM19] K. Liu and H. Ma., “Exploring background-bias for anomaly detection in surveillance videos.”, in *Proceedings of the 27th ACM International Conference on Multimedia*, ser. MM ’19, ACM, Oct. 2019 (cit. on p. 51).
- [LM98] M. J. Lerner and L. Montada, “An overview”, in *Responses to Victimitizations and Belief in a Just World*. Springer US, 1998, pp. 1–7, ISBN: 9781475764185. DOI: [10.1007/978-1-4757-6418-5_1](https://doi.org/10.1007/978-1-4757-6418-5_1). [Online]. Available: http://dx.doi.org/10.1007/978-1-4757-6418-5_1 (cit. on p. 23).
- [LML24] M. H. Lee, J. M. Montgomery, and C. K. Lai, “Large language models portray socially subordinate groups as more homogeneous, consistent with a bias observed in humans”, in *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, ser. FAccT ’24, ACM, Jun. 2024, pp. 1321–1340. DOI: [10.1145/3630106.3658975](https://doi.org/10.1145/3630106.3658975). [Online]. Available: <http://dx.doi.org/10.1145/3630106.3658975> (cit. on p. 57).
- [LNP97] J. M. Lampinen, J. S. Neuschatz, and D. G. Payne, “Memory illusions and consciousness: Examining the phenomenology of true and false memories”, *Current Psychology*, vol. 16, no. 3-4, pp. 181–224, Sep. 1997. DOI: [10.1007/s12144-997-1000-5](https://doi.org/10.1007/s12144-997-1000-5). [Online]. Available: <https://doi.org/10.1007/s12144-997-1000-5> (cit. on p. 14).
- [Loe05] G. Loewenstein, “Hot-cold empathy gaps and medical decision making.”, *Health Psychology*, vol. 24, no. 4, Suppl, S49–S56, 2005, ISSN: 0278-6133. DOI: [10.1037/0278-6133.24.4.s49](https://doi.org/10.1037/0278-6133.24.4.s49). [Online]. Available: <http://dx.doi.org/10.1037/0278-6133.24.4.s49> (cit. on p. 19).

- [Lof75] E. F. Loftus, “Reconstructing memory: The incredible eyewitness”, *Jurimetrics Journal*, vol. 15, no. 3, pp. 188–193, 1975, ISSN: 00226793. [Online]. Available: <http://www.jstor.org/stable/29761487> (visited on 08/27/2022) (cit. on p. 14).
- [LP74] E. F. Loftus and J. C. Palmer, “Reconstruction of automobile destruction: An example of the interaction between language and memory”, *Journal of Verbal Learning and Verbal Behavior*, vol. 13, no. 5, pp. 585–589, Oct. 1974. DOI: [10.1016/s0022-5371\(74\)80011-3](https://doi.org/10.1016/s0022-5371(74)80011-3). [Online]. Available: [https://doi.org/10.1016/s0022-5371\(74\)80011-3](https://doi.org/10.1016/s0022-5371(74)80011-3) (cit. on pp. 14, 23).
- [LRWY25] L. Ling, F. Rabbi, S. Wang, and J. Yang, “Bias unveiled: Investigating social bias in llm-generated code”, *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 26, pp. 27 491–27 499, Apr. 2025, ISSN: 2159-5399. DOI: [10.1609/aaai.v39i26.34961](https://doi.org/10.1609/aaai.v39i26.34961). [Online]. Available: <http://dx.doi.org/10.1609/aaai.v39i26.34961> (cit. on p. 58).
- [LS24] K. Leong and A. Sung, “Gender stereotypes in artificial intelligence within the accounting profession using large language models”, *Humanities and Social Sciences Communications*, vol. 11, no. 1, Sep. 2024, ISSN: 2662-9992. DOI: [10.1057/s41599-024-03660-8](https://doi.org/10.1057/s41599-024-03660-8). [Online]. Available: <http://dx.doi.org/10.1057/s41599-024-03660-8> (cit. on p. 58).
- [LS74] D. Landy and H. Sigall, “Beauty is talent: Task evaluation as a function of the performer’s physical attractiveness.”, *Journal of Personality and Social Psychology*, vol. 29, no. 3, 1974. DOI: [10.1037/h0036018](https://doi.org/10.1037/h0036018). [Online]. Available: <https://doi.org/10.1037/h0036018> (cit. on p. 12).
- [LSG98] I. P. Levin, S. L. Schneider, and G. J. Gaeth, “All frames are not created equal: A typology and critical analysis of framing effects”, *Organizational Behavior and Human Decision Processes*, vol. 76, no. 2, pp. 149–188, Nov. 1998. DOI: [10.1006/obhd.1998.2804](https://doi.org/10.1006/obhd.1998.2804). [Online]. Available: <https://doi.org/10.1006/obhd.1998.2804> (cit. on p. 9).
- [LSRO02] Y. Lo, A. Sides, J. Rozelle, and D. Osherson, “Evidential diversity and premise probability in young children’s inductive judgment”, *Cognitive Science*, vol. 26, no. 2, pp. 181–206, Mar. 2002. DOI: [10.1207/s15516709cog2602_2](https://doi.org/10.1207/s15516709cog2602_2). [Online]. Available: https://doi.org/10.1207/s15516709cog2602_2 (cit. on p. 11).
- [LW78] K. N. Lewis and W. B. Walsh, “Physical attractiveness: Its impact on the perception of a female counselor.”, *Journal of Counseling Psychology*, vol. 25, no. 3, pp. 210–216, May 1978. DOI: [10.1037/0022-0167.25.3.210](https://doi.org/10.1037/0022-0167.25.3.210). [Online]. Available: <https://doi.org/10.1037/0022-0167.25.3.210> (cit. on pp. 29, 34, 38, 39, 156).
- [LY21] Z. Lu and M. Yin, “Human reliance on machine learning models when performance feedback is limited: Heuristics and risks”, in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–16 (cit. on p. 60).

- [LZL19] Y. Li, C. Zhang, and M. Laroche, “Is beauty a premium? a study of the physical attractiveness effect in service encounters”, *Journal of Retailing and Consumer Services*, vol. 50, pp. 215–225, Sep. 2019. DOI: [10.1016/j.jretconser.2019.04.016](https://doi.org/10.1016/j.jretconser.2019.04.016). [Online]. Available: <https://doi.org/10.1016/j.jretconser.2019.04.016> (cit. on p. 96).
- [Man16] S. S. Mangiafico, *Summary and Analysis of Extension Program Evaluation in R*. New Brunswick, NJ: Rutgers Cooperative Extension, 2016 (cit. on pp. 46, 101).
- [MB02] T. Mussweiler and G. V. Bodenhausen, “I know you are, but what am i? self-evaluative consequences of judging in-group and out-group members.”, *Journal of Personality and Social Psychology*, vol. 82, no. 1, pp. 19–32, Jan. 2002, ISSN: 0022-3514. DOI: [10.1037/0022-3514.82.1.19](https://doi.org/10.1037/0022-3514.82.1.19). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.82.1.19> (cit. on p. 95).
- [MCP71] J. Morton, R. G. Crowder, and H. A. Prussin, “Experiments with the stimulus suffix effect.”, *Journal of Experimental Psychology*, vol. 91, no. 1, pp. 169–190, 1971, ISSN: 0022-1015. DOI: [10.1037/h0031844](https://doi.org/10.1037/h0031844). [Online]. Available: <http://dx.doi.org/10.1037/h0031844> (cit. on p. 24).
- [MCW+23] W. Ma, B. Chiang, T. Wu, L. Wang, and S. Vosoughi, “Intersectional stereotypes in large language models: Dataset and analysis”, in *Findings of the Association for Computational Linguistics: EMNLP 2023*, Association for Computational Linguistics, 2023, pp. 8589–8597. DOI: [10.18653/v1/2023.findings-emnlp.575](https://doi.org/10.18653/v1/2023.findings-emnlp.575). [Online]. Available: <http://dx.doi.org/10.18653/v1/2023.findings-emnlp.575> (cit. on pp. 58, 67).
- [MCW15] D. S. Ma, J. Correll, and B. Wittenbrink, “The chicago face database: A free stimulus set of faces and norming data”, *Behavior Research Methods*, vol. 47, no. 4, pp. 1122–1135, Jan. 2015. DOI: [10.3758/s13428-014-0532-5](https://doi.org/10.3758/s13428-014-0532-5). [Online]. Available: <https://doi.org/10.3758/s13428-014-0532-5> (cit. on pp. 29, 33, 42, 43, 45, 46, 59, 60, 115).
- [MED+95] M. A. McDaniel, G. O. Einstein, E. L. DeLosh, C. P. May, and P. Brady, “The bizarreness effect: It’s not surprising, it’s complex.”, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 21, no. 2, pp. 422–435, 1995, ISSN: 0278-7393. DOI: [10.1037/0278-7393.21.2.422](https://doi.org/10.1037/0278-7393.21.2.422). [Online]. Available: <http://dx.doi.org/10.1037/0278-7393.21.2.422> (cit. on p. 24).
- [MFH+03] R. Mulhern, G. Fieldman, T. Hussey, J.-L. Lévesque, and P. Pineau, “Do cosmetics enhance female caucasian facial attractiveness?”, *International Journal of Cosmetic Science*, vol. 25, no. 4, pp. 199–205, Aug. 2003, ISSN: 1468-2494. DOI: [10.1046/j.1467-2494.2003.00188.x](https://doi.org/10.1046/j.1467-2494.2003.00188.x). [Online]. Available: <http://dx.doi.org/10.1046/j.1467-2494.2003.00188.x> (cit. on pp. 28, 30).
- [MG97] M. C. Martin and J. W. Gentry, “Stuck in the model trap: The effects of beautiful models in ads on female pre-adolescents and adolescents”, *Journal of Advertising*, vol. 26, no. 2, pp. 19–33, Jun. 1997, ISSN: 1557-7805. DOI: [10.1080/00913367.1997.10673520](https://doi.org/10.1080/00913367.1997.10673520). [Online]. Available: <http://dx.doi.org/10.1080/00913367.1997.10673520> (cit. on p. 74).

- [Mil56] G. A. Miller, “The magical number seven, plus or minus two: Some limits on our capacity for processing information.”, *Psychological Review*, vol. 63, no. 2, pp. 81–97, Mar. 1956, ISSN: 0033-295X. DOI: [10.1037/h0043158](https://doi.org/10.1037/h0043158). [Online]. Available: <http://dx.doi.org/10.1037/h0043158> (cit. on p. 46).
- [Mil63] S. Milgram, “Behavioral study of obedience.”, *The Journal of Abnormal and Social Psychology*, vol. 67, no. 4, pp. 371–378, Oct. 1963. [Online]. Available: <https://doi.org/10.1037/h0040525> (cit. on p. 25).
- [Mil70] A. G. Miller, “Role of physical attractiveness in impression formation”, *Psychonomic Science*, vol. 19, no. 4, pp. 241–243, Oct. 1970. DOI: [10.3758/bf03328797](https://doi.org/10.3758/bf03328797). [Online]. Available: <https://doi.org/10.3758/bf03328797> (cit. on pp. 28, 56).
- [MJ00] M. Mather and M. K. Johnson, “Choice-supportive source monitoring: Do our decisions seem better to us as we age?”, *Psychology and Aging*, vol. 15, no. 4, pp. 596–606, 2000, ISSN: 0882-7974. DOI: [10.1037/0882-7974.15.4.596](https://doi.org/10.1037/0882-7974.15.4.596). [Online]. Available: <http://dx.doi.org/10.1037/0882-7974.15.4.596> (cit. on p. 24).
- [MK75] E. W. Mathes and A. Kahn, “Physical attractiveness, happiness, neuroticism, and self-esteem”, *The Journal of Psychology*, vol. 90, no. 1, pp. 27–30, May 1975, ISSN: 1940-1019. DOI: [10.1080/00223980.1975.9923921](https://doi.org/10.1080/00223980.1975.9923921). [Online]. Available: <http://dx.doi.org/10.1080/00223980.1975.9923921> (cit. on pp. 27, 50, 56).
- [MKB20] T. C. Mann, B. Kurdi, and M. R. Banaji, “How effectively can implicit evaluations be updated? using evaluative statements after aversive repeated evaluative pairings.”, *Journal of Experimental Psychology: General*, vol. 149, no. 6, pp. 1169–1192, Jun. 2020, ISSN: 0096-3445. DOI: [10.1037/xge0000701](https://doi.org/10.1037/xge0000701). [Online]. Available: <http://dx.doi.org/10.1037/xge0000701> (cit. on p. 42).
- [MKC+14] U. M. Marcinkowska, M. V. Kozlov, H. Cai, J. Contreras-Garduño, B. J. Dixon, G. A. Oana, G. Kaminski, N. P. Li, M. T. Lyons, I. E. Onyishi, K. Prasai, F. Pazhoohi, P. Prokop, S. L. Rosales Cardozo, N. Sydney, J. C. Yong, and M. J. Rantala, “Cross-cultural variation in men’s preference for sexual dimorphism in women’s faces”, *Biology Letters*, vol. 10, no. 4, p. 20130850, Apr. 2014, ISSN: 1744-957X. DOI: [10.1098/rsbl.2013.0850](https://doi.org/10.1098/rsbl.2013.0850). [Online]. Available: <http://dx.doi.org/10.1098/rsbl.2013.0850> (cit. on p. 28).
- [MM01] B. Monin and D. T. Miller, “Moral credentials and the expression of prejudice.”, *Journal of Personality and Social Psychology*, vol. 81, no. 1, pp. 33–43, 2001, ISSN: 0022-3514. DOI: [10.1037/0022-3514.81.1.33](https://doi.org/10.1037/0022-3514.81.1.33). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.81.1.33> (cit. on p. 22).
- [MM78] B. Murdock and J. Metcalfe, “Controlled rehearsal in single-trial free recall”, *Journal of Verbal Learning and Verbal Behavior*, vol. 17, no. 3, pp. 309–324, Jun. 1978. DOI: [10.1016/s0022-5371\(78\)90201-3](https://doi.org/10.1016/s0022-5371(78)90201-3). [Online]. Available: [https://doi.org/10.1016/s0022-5371\(78\)90201-3](https://doi.org/10.1016/s0022-5371(78)90201-3) (cit. on p. 14).

- [MMS+21] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, “A survey on bias and fairness in machine learning”, *ACM Comput. Surv.*, vol. 54, no. 6, Jul. 2021, ISSN: 0360-0300. DOI: [10 . 1145 / 3457607](https://doi.org/10.1145/3457607). [Online]. Available: [https : //doi.org/10.1145/3457607](https://doi.org/10.1145/3457607) (cit. on p. 70).
- [MR75] D. T. Miller and M. Ross, “Self-serving biases in the attribution of causality: Fact or fiction?”, *Psychological Bulletin*, vol. 82, no. 2, pp. 213–225, Mar. 1975, ISSN: 0033-2909. DOI: [10 . 1037 / h0076486](https://doi.org/10.1037/h0076486). [Online]. Available: [http : //dx . doi . org/10.1037/h0076486](http://dx.doi.org/10.1037/h0076486) (cit. on p. 22).
- [MRS+21] S. A. McLean, R. F. Rodgers, A. Slater, H. K. Jarman, C. S. Gordon, and S. J. Paxton, “Clinically significant body dissatisfaction: Prevalence and association with depressive symptoms in adolescent boys and girls”, *European Child & Adolescent Psychiatry*, vol. 31, no. 12, pp. 1921–1932, Jun. 2021, ISSN: 1435-165X. DOI: [10 . 1007 / s00787 - 021 - 01824 - 4](https://doi.org/10.1007/s00787-021-01824-4). [Online]. Available: [http : //dx . doi . org/10.1007/s00787-021-01824-4](http://dx.doi.org/10.1007/s00787-021-01824-4) (cit. on pp. 41, 159).
- [MSC93] B. Major, A. M. Sciacchitano, and J. Crocker, “In-group versus out-group comparisons and self-esteem”, *Personality and Social Psychology Bulletin*, vol. 19, no. 6, pp. 711–721, Dec. 1993, ISSN: 1552-7433. DOI: [10 . 1177 / 0146167293196006](https://doi.org/10.1177/0146167293196006). [Online]. Available: [http : //dx . doi . org/10.1177/0146167293196006](http://dx.doi.org/10.1177/0146167293196006) (cit. on p. 96).
- [MSN96] C. M. Marlowe, S. L. Schneider, and C. E. Nelson, “Gender and attractiveness biases in hiring decisions: Are more experienced managers less biased?”, *Journal of Applied Psychology*, vol. 81, no. 1, pp. 11–21, Feb. 1996, ISSN: 0021-9010. DOI: [10 . 1037 / 0021 - 9010 . 81 . 1 . 11](https://doi.org/10.1037/0021-9010.81.1.11). [Online]. Available: [http : //dx . doi . org/10.1037/0021-9010.81.1.11](http://dx.doi.org/10.1037/0021-9010.81.1.11) (cit. on pp. 41, 159).
- [MSY20] M. Miceli, M. Schuessler, and T. Yang., “Between subjectivity and imposition: Power dynamics in data annotation for computer vision.”, *Proceedings of the ACM on Human-Computer Interaction*, vol. 4, no. CSCW2, pp. 1–25, Oct. 2020 (cit. on p. 51).
- [MT00] M. S. McGlone and J. Tofighbakhsh, “Birds of a feather flock conjointly (?): Rhyme as reason in aphorisms”, *Psychological Science*, vol. 11, no. 5, pp. 424–428, Sep. 2000, ISSN: 1467-9280. DOI: [10 . 1111 / 1467 - 9280 . 00282](https://doi.org/10.1111/1467-9280.00282). [Online]. Available: [http : //dx . doi . org/10.1111/1467-9280.00282](http://dx.doi.org/10.1111/1467-9280.00282) (cit. on p. 19).
- [MTPC97] T. R. Mitchell, L. Thompson, E. Peterson, and R. Cronk, “Temporal adjustments in the evaluation of events: The “rosy view””, *Journal of Experimental Social Psychology*, vol. 33, no. 4, pp. 421–448, Jul. 1997, ISSN: 0022-1031. DOI: [10 . 1006 / jesp . 1997 . 1333](https://doi.org/10.1006/jesp.1997.1333). [Online]. Available: [http : //dx . doi . org/10.1006/jesp.1997.1333](http://dx.doi.org/10.1006/jesp.1997.1333) (cit. on p. 24).
- [Mur62] B. B. Murdock, “The serial position effect of free recall.”, *Journal of Experimental Psychology*, vol. 64, no. 5, pp. 482–488, Nov. 1962. DOI: [10 . 1037 / h0045106](https://doi.org/10.1037/h0045106). [Online]. Available: [https : //doi.org/10.1037/h0045106](https://doi.org/10.1037/h0045106) (cit. on pp. 14, 23).

- [MZW+15] D. G. Mitchem, B. P. Zietsch, M. J. Wright, N. G. Martin, J. K. Hewitt, and M. C. Keller, “No relationship between intelligence and facial attractiveness in a large, genetically informative sample”, *Evolution and Human Behavior*, vol. 36, no. 3, pp. 240–247, 2015, ISSN: 1090-5138. DOI: <https://doi.org/10.1016/j.evolhumbehav.2014.11.009>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1090513814001512> (cit. on pp. 27, 54).
- [NAG19] F. Ni, D. Arnott, and S. Gao, “The anchoring effect in business intelligence supported decision-making”, *Journal of Decision Systems*, vol. 28, no. 2, pp. 67–81, Apr. 2019. DOI: [10.1080/12460125.2019.1620573](https://doi.org/10.1080/12460125.2019.1620573). [Online]. Available: <https://doi.org/10.1080/12460125.2019.1620573> (cit. on p. 9).
- [NBR21] M. Nadeem, A. Bethke, and S. Reddy, “Stereoset: Measuring stereotypical bias in pretrained language models”, in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 5356–5371 (cit. on p. 57).
- [NCLM73] R. E. Nisbett, C. Caputo, P. Legant, and J. Marecek, “Behavior as seen by the actor and as seen by the observer.”, *Journal of Personality and Social Psychology*, vol. 27, no. 2, pp. 154–164, Aug. 1973, ISSN: 0022-3514. DOI: [10.1037/h0034779](https://doi.org/10.1037/h0034779). [Online]. Available: <http://dx.doi.org/10.1037/h0034779> (cit. on p. 22).
- [NF22] S. J. Nightingale and H. Farid, “AI-synthesized faces are indistinguishable from real faces and more trustworthy”, *Proceedings of the National Academy of Sciences*, vol. 119, no. 8, Feb. 2022. DOI: [10.1073/pnas.2120481119](https://doi.org/10.1073/pnas.2120481119). [Online]. Available: <https://doi.org/10.1073/pnas.2120481119> (cit. on pp. 42, 46, 53, 87, 115).
- [NGP+23] P. Narayanan Venkit, S. Gautam, R. Panchanadikar, T.-H. Huang, and S. Wilson, “Unmasking nationality bias: A study of human perception of nationalities in ai-generated articles”, in *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, ser. AIES ’23, ACM, Aug. 2023, pp. 554–565. DOI: [10.1145/3600211.3604667](https://doi.org/10.1145/3600211.3604667). [Online]. Available: <http://dx.doi.org/10.1145/3600211.3604667> (cit. on p. 57).
- [Ngu15] Nguyen, Laurent Son, “Computational analysis of behavior in employment interviews and video resumes.”, Ph.D. dissertation, 2015 (cit. on p. 51).
- [NHD80] W. Nasby, B. Hayden, and B. M. DePaulo, “Attributional bias among aggressive boys to interpret unambiguous social stimuli as displays of hostility.”, *Journal of Abnormal Psychology*, vol. 89, no. 3, pp. 459–468, 1980, ISSN: 0021-843X. DOI: [10.1037/0021-843x.89.3.459](https://doi.org/10.1037/0021-843x.89.3.459). [Online]. Available: <http://dx.doi.org/10.1037/0021-843X.89.3.459> (cit. on p. 21).
- [NHP09] L. F. Nordgren, F. v. Harreveld, and J. v. d. Pligt, “The restraint bias: How the illusion of self-restraint promotes impulsive behavior”, *Psychological Science*, vol. 20, no. 12, pp. 1523–1528, Dec. 2009, ISSN: 1467-9280. DOI: [10.1111/j.1467-9280.2009.02468.x](https://doi.org/10.1111/j.1467-9280.2009.02468.x). [Online]. Available: <http://dx.doi.org/10.1111/j.1467-9280.2009.02468.x> (cit. on p. 21).

- [Nic98] R. S. Nickerson., “Confirmation bias: A ubiquitous phenomenon in many guises.”, *Review of General Psychology*, vol. 2, no. 2, pp. 175–220, Jun. 1998 (cit. on p. 50).
- [NMA12] M. I. Norton, D. Mochon, and D. Ariely, “The IKEA effect: When labor leads to love”, *Journal of Consumer Psychology*, vol. 22, no. 3, pp. 453–460, Jul. 2012. DOI: [10.1016/j.jcps.2011.08.002](https://doi.org/10.1016/j.jcps.2011.08.002). [Online]. Available: <https://doi.org/10.1016/j.jcps.2011.08.002> (cit. on pp. 12, 22).
- [NN23a] R. Naik and B. Nushi, “Social biases through the text-to-image generation lens”, in *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, ser. AIES ’23, ACM, Aug. 2023, pp. 786–808. DOI: [10.1145/3600211.3604711](https://doi.org/10.1145/3600211.3604711). [Online]. Available: <http://dx.doi.org/10.1145/3600211.3604711> (cit. on pp. 53, 59).
- [NN23b] R. Naik and B. Nushi, “Social biases through the text-to-image generation lens”, in *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 2023, pp. 786–808 (cit. on p. 74).
- [NR10] B. Nyhan and J. Reifler, “When corrections fail: The persistence of political misperceptions”, *Political Behavior*, vol. 32, no. 2, pp. 303–330, Mar. 2010, ISSN: 1573-6687. DOI: [10.1007/s11109-010-9112-2](https://doi.org/10.1007/s11109-010-9112-2). [Online]. Available: <http://dx.doi.org/10.1007/s11109-010-9112-2> (cit. on p. 19).
- [NW64] J. Neter and J. Waksberg, “A study of response errors in expenditures data from household interviews”, *Journal of the American Statistical Association*, vol. 59, no. 305, pp. 18–55, Mar. 1964, ISSN: 1537-274X. DOI: [10.1080/01621459.1964.10480699](https://doi.org/10.1080/01621459.1964.10480699). [Online]. Available: <http://dx.doi.org/10.1080/01621459.1964.10480699> (cit. on p. 24).
- [NW77] R. E. Nisbett and T. D. Wilson, “The halo effect: Evidence for unconscious alteration of judgments.”, *Journal of Personality and Social Psychology*, vol. 35, no. 4, pp. 250–256, Apr. 1977. DOI: [10.1037/0022-3514.35.4.250](https://doi.org/10.1037/0022-3514.35.4.250). [Online]. Available: <https://doi.org/10.1037/0022-3514.35.4.250> (cit. on p. 12).
- [OBT18] D. Oh, E. A. Buck, and A. Todorov, “Revealing hidden gender biases in competence impressions of faces”, *Psychological Science*, vol. 30, no. 1, pp. 65–79, Dec. 2018, ISSN: 1467-9280. DOI: [10.1177/0956797618813092](https://doi.org/10.1177/0956797618813092). [Online]. Available: <http://dx.doi.org/10.1177/0956797618813092> (cit. on p. 37).
- [OM85] E. J. O’Brien and J. L. Myers, “When comprehension difficulty improves memory for text.”, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 11, no. 1, pp. 12–21, Jan. 1985, ISSN: 0278-7393. DOI: [10.1037/0278-7393.11.1.12](https://doi.org/10.1037/0278-7393.11.1.12). [Online]. Available: <http://dx.doi.org/10.1037/0278-7393.11.1.12> (cit. on p. 24).
- [OT08] N. N. Oosterhof and A. Todorov, “The functional basis of face evaluation”, *Proceedings of the National Academy of Sciences*, vol. 105, no. 32, pp. 11 087–11 092, Aug. 2008. DOI: [10.1073/pnas.0805664105](https://doi.org/10.1073/pnas.0805664105). [Online]. Available: <https://doi.org/10.1073/pnas.0805664105> (cit. on pp. 33, 39, 42, 46, 157).

- [OWLT23] D. Oh, N. Wedel, B. Labbree, and A. Todorov, “Trustworthiness judgments without the halo effect: A data-driven computational modeling approach”, *Perception*, vol. 52, no. 8, pp. 590–607, Jun. 2023, ISSN: 1468-4233. DOI: [10.1177/03010066231178489](https://doi.org/10.1177/03010066231178489). [Online]. Available: <http://dx.doi.org/10.1177/03010066231178489> (cit. on pp. 42, 46, 115).
- [PC89] S. M. Petroschius and K. E. Crocker, “An empirical analysis of spokesperson characteristics on advertisement and product evaluations”, *Journal of the Academy of Marketing Science*, vol. 17, no. 3, pp. 217–225, Jun. 1989, ISSN: 1552-7824. DOI: [10.1007/bf02729813](https://doi.org/10.1007/bf02729813). [Online]. Available: <http://dx.doi.org/10.1007/BF02729813> (cit. on p. 74).
- [PDH+23] X. Pan, L. Dong, S. Huang, Z. Peng, W. Chen, and F. Wei, “Kosmos-g: Generating images in context with multimodal large language models”, *arXiv preprint arXiv:2310.02992*, 2023 (cit. on p. 56).
- [PE66] L. D. Phillips and W. Edwards, “Conservatism in a simple probability inference task.”, *Journal of Experimental Psychology*, vol. 72, no. 3, pp. 346–354, 1966, ISSN: 0022-1015. DOI: [10.1037/h0023653](https://doi.org/10.1037/h0023653). [Online]. Available: <http://dx.doi.org/10.1037/h0023653> (cit. on p. 21).
- [Per20] G. Perrotta, “The concept of altered perception in “body dysmorphic disorder”: The subtle border between the abuse of selfies in social networks and cosmetic surgery, between socially accepted dysfunctionality and the pathological condition”, *Journal of Neurology, Neurological Science and Disorders*, vol. 6, no. 1, pp. 001–007, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:225826320> (cit. on pp. 30, 51).
- [Pet79] T. F. Pettigrew, “The ultimate attribution error: Extending allport’s cognitive analysis of prejudice”, *Personality and Social Psychology Bulletin*, vol. 5, no. 4, pp. 461–476, Oct. 1979, ISSN: 1552-7433. DOI: [10.1177/014616727900500407](https://doi.org/10.1177/014616727900500407). [Online]. Available: <http://dx.doi.org/10.1177/014616727900500407> (cit. on p. 22).
- [PKSR01] E. Pronin, J. Kruger, K. Savtisky, and L. Ross, “You don’t know me, but i know you: The illusion of asymmetric insight.”, *Journal of Personality and Social Psychology*, vol. 81, no. 4, pp. 639–656, Oct. 2001, ISSN: 0022-3514. DOI: [10.1037/0022-3514.81.4.639](https://doi.org/10.1037/0022-3514.81.4.639). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.81.4.639> (cit. on p. 23).
- [PLL+23] D. Pehlivanoglu, T. Lin, N. R. Lighthall, A. Heemskerk, A. Harber, R. C. Wilson, G. R. Turner, R. N. Spreng, and N. C. Ebner, “Facial trustworthiness perception across the adult life span”, *The Journals of Gerontology: Series B*, vol. 78, no. 3, A. Krendl, Ed., pp. 434–444, Oct. 2023, ISSN: 1758-5368. DOI: [10.1093/geronb/gbac166](https://doi.org/10.1093/geronb/gbac166). [Online]. Available: <http://dx.doi.org/10.1093/geronb/gbac166> (cit. on pp. 40, 158).
- [PLP+98] D. I. Perrett, K. J. Lee, I. Penton-Voak, D. Rowland, S. Yoshikawa, D. M. Burt, S. P. Henzi, D. L. Castles, and S. Akamatsu., “Effects of sexual dimorphism on facial attractiveness.”, *Nature*, vol. 394, no. 6696, pp. 884–887, Aug. 1998 (cit. on p. 53).

- [PLR02] E. Pronin, D. Y. Lin, and L. Ross, “The bias blind spot: Perceptions of bias in self versus others”, *Personality and Social Psychology Bulletin*, vol. 28, no. 3, pp. 369–381, Mar. 2002, ISSN: 1552-7433. DOI: [10.1177/0146167202286008](https://doi.org/10.1177/0146167202286008). [Online]. Available: <http://dx.doi.org/10.1177/0146167202286008> (cit. on p. 22).
- [PMHP01] A. H. Perlini, A. Marcello, S. D. Hansen, and W. Pudney, “The effects of male age and physical appearance on evaluations of attractiveness, social desirability and resourcefulness”, *Social Behavior and Personality: an international journal*, vol. 29, no. 3, pp. 277–287, Jan. 2001, ISSN: 0301-2212. DOI: [10.2224/sbp.2001.29.3.277](https://doi.org/10.2224/sbp.2001.29.3.277). [Online]. Available: <http://dx.doi.org/10.2224/sbp.2001.29.3.277> (cit. on p. 35).
- [PP04] R. Pohl and R. F. Pohl, *Cognitive illusions: A handbook on fallacies and biases in thinking, judgement and memory*. Psychology Press, 2004 (cit. on p. 7).
- [PR82] B. Park and M. Rothbart, “Perception of out-group homogeneity and levels of social categorization: Memory for the subordinate attributes of in-group and out-group members.”, *Journal of Personality and Social Psychology*, vol. 42, no. 6, pp. 1051–1068, Jun. 1982, ISSN: 0022-3514. DOI: [10.1037/0022-3514.42.6.1051](https://doi.org/10.1037/0022-3514.42.6.1051). [Online]. Available: <http://dx.doi.org/10.1037/0022-3514.42.6.1051> (cit. on p. 22).
- [Pra78] J. W. Pratt, “Risk aversion in the small and in the large”, in *Uncertainty in Economics*, Elsevier, 1978, pp. 59–79. DOI: [10.1016/b978-0-12-214850-7.50010-3](https://doi.org/10.1016/b978-0-12-214850-7.50010-3). [Online]. Available: <https://doi.org/10.1016/b978-0-12-214850-7.50010-3> (cit. on p. 12).
- [PS22] D. Pessach and E. Shmueli, “A review on fairness in machine learning”, *ACM Comput. Surv.*, vol. 55, no. 3, Feb. 2022, ISSN: 0360-0300. DOI: [10.1145/3494672](https://doi.org/10.1145/3494672). [Online]. Available: <https://doi.org/10.1145/3494672> (cit. on p. 70).
- [PSC01] T. Postmes, R. Spears, and S. Cihangir, “Quality of decision making and group norms.”, *Journal of Personality and Social Psychology*, vol. 80, no. 6, pp. 918–930, 2001. DOI: [10.1037/0022-3514.80.6.918](https://doi.org/10.1037/0022-3514.80.6.918). [Online]. Available: <https://doi.org/10.1037/0022-3514.80.6.918> (cit. on p. 16).
- [PSE+02] D. Poulin-Dubois, L. A. Serbin, J. A. Eichstedt, M. G. Sen, and C. F. Beissel, “Men don’t put on make-up: Toddlers’ knowledge of the gender stereotyping of household activities”, *Social Development*, vol. 11, no. 2, pp. 166–181, May 2002, ISSN: 1467-9507. DOI: [10.1111/1467-9507.00193](https://doi.org/10.1111/1467-9507.00193). [Online]. Available: <http://dx.doi.org/10.1111/1467-9507.00193> (cit. on p. 29).
- [PTN19] R. B. Parikh, S. Teeple, and A. S. Navathe, “Addressing bias in artificial intelligence in health care”, *Jama*, vol. 322, no. 24, pp. 2377–2378, 2019 (cit. on p. 49).
- [PUG+22] J. C. Peterson, S. Uddenberg, T. L. Griffiths, A. Todorov, and J. W. Suchow, “Deep models of superficial face judgments”, *Proceedings of the National Academy of Sciences*, vol. 119, no. 17, Apr. 2022. DOI: [10.1073/pnas.2115228119](https://doi.org/10.1073/pnas.2115228119). [Online]. Available: <https://doi.org/10.1073/pnas.2115228119> (cit. on pp. 33, 42, 46, 60).

- [R C21] R Core Team, *R: A language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria, 2021. [Online]. Available: <https://www.R-project.org/> (cit. on p. 47).
- [RAH+24] J. Ricker, D. Assenmacher, T. Holz, A. Fischer, and E. Quiring, “Ai-generated faces in the real world: A large-scale case study of twitter profile images”, in *Proceedings of the 27th International Symposium on Research in Attacks, Intrusions and Defenses*, 2024, pp. 513–530 (cit. on p. 73).
- [RBKL20] M. Raghavan, S. Barocas, J. Kleinberg, and K. Levy, “Mitigating bias in algorithmic hiring: Evaluating claims and practices”, in *Proceedings of the 2020 conference on fairness, accountability, and transparency*, 2020, pp. 469–481 (cit. on p. 49).
- [RBL+22] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, *High-resolution image synthesis with latent diffusion models*, 2022. arXiv: 2112.10752 [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2112.10752> (cit. on p. 75).
- [RBW15] J. Różycka-Tran, P. Boski, and B. Wojciszke, “Belief in a zero-sum game as a social axiom: A 37-nation study”, *Journal of Cross-Cultural Psychology*, vol. 46, no. 4, pp. 525–548, Mar. 2015, ISSN: 1552-5422. DOI: 10.1177/0022022115572226. [Online]. Available: <http://dx.doi.org/10.1177/0022022115572226> (cit. on p. 21).
- [RCOO24] P. Riccio, J. Colin, S. Ogolla, and N. Oliver, “Mirror, mirror on the wall, who is the whitest of all? racial biases in social media beauty filters”, *Social Media + Society*, vol. 10, no. 2, Apr. 2024, ISSN: 2056-3051. DOI: 10.1177/20563051241239295. [Online]. Available: <http://dx.doi.org/10.1177/20563051241239295> (cit. on p. 30).
- [RD23] M. Rougier and J. De Houwer, “Updating stereotypical attributions in light of new information: The attractiveness halo effect changes when attractiveness changes”, *European Journal of Social Psychology*, vol. 54, no. 1, pp. 364–379, Dec. 2023, ISSN: 1099-0992. DOI: 10.1002/ejsp.3017. [Online]. Available: <http://dx.doi.org/10.1002/ejsp.3017> (cit. on pp. 29, 39, 40, 42, 157).
- [RF63] R. Rosenthal and K. L. Fode, “Psychology of the scientist: V. three experiments in experimenter bias”, *Psychological Reports*, vol. 12, no. 2, pp. 491–511, Apr. 1963, ISSN: 1558-691X. DOI: 10.2466/pr0.1963.12.2.491. [Online]. Available: <http://dx.doi.org/10.2466/pr0.1963.12.2.491> (cit. on p. 19).
- [RGH77] L. Ross, D. Greene, and P. House, “The “false consensus effect”: An egocentric bias in social perception and attribution processes”, *Journal of Experimental Social Psychology*, vol. 13, no. 3, pp. 279–301, May 1977. [Online]. Available: [https://doi.org/10.1016/0022-1031\(77\)90049-x](https://doi.org/10.1016/0022-1031(77)90049-x) (cit. on p. 20).
- [RH01] Y. Rottenstreich and C. K. Hsee, “Money, kisses, and electric shocks: On the affective psychology of risk”, *Psychological Science*, vol. 12, no. 3, pp. 185–190, May 2001. DOI: 10.1111/1467-9280.00334. [Online]. Available: <https://doi.org/10.1111/1467-9280.00334> (cit. on p. 12).

- [Rid01] C. L. Ridgeway, “Gender, status, and leadership”, *Journal of Social Issues*, vol. 57, no. 4, pp. 637–655, Jan. 2001, ISSN: 1540-4560. DOI: [10.1111/0022-4537.00233](https://doi.org/10.1111/0022-4537.00233). [Online]. Available: <http://dx.doi.org/10.1111/0022-4537.00233> (cit. on pp. 41, 159).
- [Rid11] C. L. Ridgeway, *Framed by Gender*. Oxford University Press, Jan. 2011, ISBN: 9780199755776. DOI: [10.1093/acprof:oso/9780199755776.001.0001](https://doi.org/10.1093/acprof:oso/9780199755776.001.0001). [Online]. Available: <http://dx.doi.org/10.1093/acprof:oso/9780199755776.001.0001> (cit. on pp. 41, 159).
- [RKH+21a] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, and et al., “Learning transferable visual models from natural language supervision.”, in *International conference on machine learning*, PMLR, 2021, pp. 8748–8763 (cit. on p. 52).
- [RKH+21b] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, *Learning transferable visual models from natural language supervision*, 2021. arXiv: [2103.00020](https://arxiv.org/abs/2103.00020) [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2103.00020> (cit. on p. 77).
- [RKK77] T. B. Rogers, N. A. Kuiper, and W. S. Kirker, “Self-reference and the encoding of personal information.”, *Journal of Personality and Social Psychology*, vol. 35, no. 9, pp. 677–688, 1977. DOI: [10.1037/0022-3514.35.9.677](https://doi.org/10.1037/0022-3514.35.9.677). [Online]. Available: <https://doi.org/10.1037/0022-3514.35.9.677> (cit. on pp. 14, 23).
- [RL95] D. Read and G. Loewenstein, “Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices.”, *Journal of Experimental Psychology: Applied*, vol. 1, no. 1, pp. 34–49, Mar. 1995. DOI: [10.1037/1076-898x.1.1.34](https://doi.org/10.1037/1076-898x.1.1.34). [Online]. Available: <https://doi.org/10.1037/1076-898x.1.1.34> (cit. on p. 16).
- [RLG+23] X. Ren, A. Lattas, B. Gecer, J. Deng, C. Ma, and X. Yang, “Facial geometric detail recovery via implicit representation”, in *2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG)*, 2023 (cit. on pp. 74, 78).
- [RLH+19] T. Radtke, N. Liszewska, K. Horodyska, M. Boberska, K. Schenkel, and A. Luszczynska, “Cooking together: The ikea effect on family vegetable intake”, *British Journal of Health Psychology*, vol. 24, no. 4, pp. 896–912, Sep. 2019. DOI: [10.1111/bjhp.12385](https://doi.org/10.1111/bjhp.12385). [Online]. Available: <https://doi.org/10.1111/bjhp.12385> (cit. on p. 12).
- [RLH+20] J. P. Robinson, G. Livitz, Y. Henon, C. Qin, Y. Fu, and S. Timoner., “Face recognition: Too bias, or not too bias?”, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Jun. 2020 (cit. on p. 51).
- [RLK23] S. H. R. Rasmussen, S. G. Ludeke, and R. Klemmensen, “Using deep learning to predict ideology from facial photographs: Expressions, beauty, and extra-facial information”, *Scientific Reports*, vol. 13, no. 1, Mar. 2023. DOI: [10.1038/s41598-023-31796-1](https://doi.org/10.1038/s41598-023-31796-1). [Online]. Available: <https://doi.org/10.1038/s41598-023-31796-1> (cit. on p. 31).

- [RO22] P. Riccio and N. Oliver, “Racial bias in the beautyverse: Evaluation of augmented-reality beauty filters”, in *European Conference on Computer Vision*, Springer, 2022, pp. 714–721 (cit. on pp. 51, 54, 71).
- [Ros77] L. Ross, “The intuitive psychologist and his shortcomings: Distortions in the attribution process”, in *Advances in Experimental Social Psychology Volume 10*. Elsevier, 1977, pp. 173–220, ISBN: 9780120152100. DOI: 10.1016/S0065-2601(08)60357-3. [Online]. Available: [http://dx.doi.org/10.1016/S0065-2601\(08\)60357-3](http://dx.doi.org/10.1016/S0065-2601(08)60357-3) (cit. on p. 21).
- [RPG+22] P. Riccio, B. Psomas, F. Galati, F. Escolano, T. Hofmann, and N. Oliver, “Openfilter: A framework to democratize research access to social media ar filters”, *Advances in Neural Information Processing Systems*, vol. 35, pp. 12 491–12 503, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2207.12319> (cit. on pp. 30, 42, 51, 71, 108).
- [RPT92] V. Ritts, M. L. Patterson, and M. E. Tubbs, “Expectations, impressions, and judgments of physically attractive students: A review”, *Review of Educational Research*, vol. 62, no. 4, pp. 413–426, Dec. 1992, ISSN: 1935-1046. DOI: 10.3102/00346543062004413. [Online]. Available: <http://dx.doi.org/10.3102/00346543062004413> (cit. on pp. 3, 28, 160).
- [RR01] P. Rozin and E. B. Royzman, “Negativity bias, negativity dominance, and contagion”, *Personality and Social Psychology Review*, vol. 5, no. 4, pp. 296–320, Nov. 2001. DOI: 10.1207/s15327957pspr0504_2. [Online]. Available: https://doi.org/10.1207/s15327957pspr0504_2 (cit. on p. 22).
- [RS09] P. Raghuram and J. Srivastava, “The denomination effect”, *Journal of Consumer Research*, vol. 36, no. 4, pp. 701–713, Dec. 2009, ISSN: 1537-5277. DOI: 10.1086/599222. [Online]. Available: <http://dx.doi.org/10.1086/599222> (cit. on p. 20).
- [RS91] L. Ross and C. Stillinger, “Barriers to conflict resolution”, *Negotiation journal*, vol. 7, no. 4, pp. 389–404, 1991 (cit. on p. 26).
- [RSMA24] R. Ramachandranpillai, K. Sampath, A. Mohammad, and M. Alikhani, *Fairness at every intersection: Uncovering and mitigating intersectional biases in multi-modal clinical predictions*, 2024. arXiv: 2412.00606 [cs.AI]. [Online]. Available: <https://arxiv.org/abs/2412.00606> (cit. on p. 71).
- [Rup+12] G. Rupert Jr et al., *Simultaneous statistical inference*, 2012 (cit. on p. 65).
- [Rus09] R. Russell, “A sex difference in facial contrast and its exaggeration by cosmetics”, *Perception*, vol. 38, no. 8, pp. 1211–1219, Jan. 2009, ISSN: 1468-4233. DOI: 10.1068/p6331. [Online]. Available: <http://dx.doi.org/10.1068/p6331> (cit. on p. 28).
- [RV90] B. E. Rogowitz and R. Voss, “Shape perception and low-dimension fractal boundary contours”, in *Human Vision and Electronic Imaging: Models, Methods, and Applications*, B. E. Rogowitz and J. P. Allebach, Eds., vol. 1249, SPIE, Oct. 1990, pp. 387–394. DOI: 10.1117/12.19691. [Online]. Available: <http://dx.doi.org/10.1117/12.19691> (cit. on p. 21).

- [Rya22] T. Ryan-Mosley, *Beauty filters are changing the way young girls see themselves*, 2021, 2022. [Online]. Available: <https://www.technologyreview.com/2021/04/02/1021635/beauty-filters-young-girls-augmented-reality-social-media/> (cit. on pp. 41, 87, 159).
- [Sac67] J. S. Sachs, “Recognition memory for syntactic and semantic aspects of connected discourse”, *Perception & Psychophysics*, vol. 2, no. 9, pp. 437–442, Sep. 1967, ISSN: 1532-5962. DOI: [10.3758/bf03208784](https://doi.org/10.3758/bf03208784). [Online]. Available: <http://dx.doi.org/10.3758/BF03208784> (cit. on p. 24).
- [SAI22] J. N. Saeed, A. M. Abdulazeez, and D. A. Ibrahim, “Fiac-net: Facial image attractiveness classification based on light deep convolutional neural network”, in *2022 Second International Conference on Computer Science, Engineering and Applications (ICCSEA)*, IEEE, 2022, pp. 1–6 (cit. on p. 152).
- [Sar22] C. Sartwell, “Beauty”, in *The Stanford Encyclopedia of Philosophy*, E. N. Zalta, Ed., Summer 2022, Metaphysics Research Lab, Stanford University, 2022. [Online]. Available: <https://plato.stanford.edu/archives/sum2022/entries/beauty/> (cit. on p. 38).
- [SCDG16] P. E. Souza, C. P. C. Chanel, F. Dehais, and S. Givigi, “Towards human-robot interaction: A framing effect experiment”, in *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, IEEE, Oct. 2016, pp. 001 929–001 934. DOI: [10.1109/smc.2016.7844521](https://doi.org/10.1109/smc.2016.7844521). [Online]. Available: <https://doi.org/10.1109/smc.2016.7844521> (cit. on p. 10).
- [Sch78] G. Schwarz, “Estimating the Dimension of a Model”, *The Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978. DOI: [10.1214/aos/1176344136](https://doi.org/10.1214/aos/1176344136). [Online]. Available: <https://doi.org/10.1214/aos/1176344136> (cit. on pp. 47, 101).
- [SF21] X. Shen and M. J. Ferguson, “How resistant are implicit impressions of facial trustworthiness? when new evidence leads to durable updating”, *Journal of Experimental Social Psychology*, vol. 97, p. 104 219, Nov. 2021, ISSN: 0022-1031. DOI: [10.1016/j.jesp.2021.104219](https://doi.org/10.1016/j.jesp.2021.104219). [Online]. Available: <http://dx.doi.org/10.1016/j.jesp.2021.104219> (cit. on p. 42).
- [SG03] K. Savitsky and T. Gilovich, “The illusion of transparency and the alleviation of speech anxiety”, *Journal of Experimental Social Psychology*, vol. 39, no. 6, pp. 618–625, Nov. 2003, ISSN: 0022-1031. DOI: [10.1016/S0022-1031\(03\)00056-8](https://doi.org/10.1016/S0022-1031(03)00056-8). [Online]. Available: [http://dx.doi.org/10.1016/S0022-1031\(03\)00056-8](http://dx.doi.org/10.1016/S0022-1031(03)00056-8) (cit. on p. 23).
- [SG09] G. S. Stuart and D. A. Grimes, “Social desirability bias in family planning studies: A neglected problem”, *Contraception*, vol. 80, no. 2, pp. 108–112, Aug. 2009. DOI: [10.1016/j.contraception.2009.02.009](https://doi.org/10.1016/j.contraception.2009.02.009). [Online]. Available: <https://doi.org/10.1016/j.contraception.2009.02.009> (cit. on p. 13).
- [She67] R. N. Shepard, “Recognition memory for words, sentences, and pictures”, *Journal of Verbal Learning and Verbal Behavior*, vol. 6, no. 1, pp. 156–163, Feb. 1967, ISSN: 0022-5371. DOI: [10.1016/S0022-5371\(67\)80067-7](https://doi.org/10.1016/S0022-5371(67)80067-7). [Online]. Available: [http://dx.doi.org/10.1016/S0022-5371\(67\)80067-7](http://dx.doi.org/10.1016/S0022-5371(67)80067-7) (cit. on p. 23).

- [SHK+18] R. M. Stolier, E. Hehman, M. D. Keller, M. Walker, and J. B. Freeman, “The conceptual structure of face impressions”, *Proceedings of the National Academy of Sciences*, vol. 115, no. 37, pp. 9210–9215, Aug. 2018, ISSN: 1091-6490. DOI: [10.1073/pnas.1807222115](https://doi.org/10.1073/pnas.1807222115). [Online]. Available: <http://dx.doi.org/10.1073/pnas.1807222115> (cit. on pp. 42, 46, 115).
- [Sim90a] H. A. Simon, “Bounded rationality”, in *Utility and Probability*, Palgrave Macmillan UK, 1990, pp. 15–18. [Online]. Available: https://doi.org/10.1007/978-1-349-20568-4_5 (cit. on pp. 4, 160).
- [Sim90b] I. Simonson, “The effect of purchase quantity and timing on variety-seeking behavior”, *Journal of Marketing Research*, vol. 27, no. 2, p. 150, May 1990. DOI: [10.2307/3172842](https://doi.org/10.2307/3172842). [Online]. Available: <https://doi.org/10.2307/3172842> (cit. on p. 16).
- [SIVA16] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, *Inception-v4, inception-resnet and the impact of residual connections on learning*, 2016. arXiv: [1602.07261](https://arxiv.org/abs/1602.07261) [cs.CV] (cit. on p. 78).
- [SK98] D. A. Schkade and D. Kahneman, “Does living in california make people happy? a focusing illusion in judgments of life satisfaction”, *Psychological Science*, vol. 9, no. 5, pp. 340–346, Sep. 1998, ISSN: 1467-9280. DOI: [10.1111/1467-9280.00066](https://doi.org/10.1111/1467-9280.00066). [Online]. Available: <http://dx.doi.org/10.1111/1467-9280.00066> (cit. on p. 22).
- [SKB+20] C. Schwemmer, C. Knight, E. D. Bello-Pardo, S. Oklobdzija, M. Schoonvelde, and J. W. Lockhart., “Diagnosing gender bias in image recognition systems.”, *Socius: Sociological Research for a Dynamic World*, vol. 6, p. 237 802 312 096 717, Jan. 2020 (cit. on pp. 50, 74).
- [SKS13] P. Sorokowski, K. Kościński, and A. Sorokowska., “Is beauty in the eye of the beholder but ugliness culturally universal? facial preferences of polish and yali (papua) people.”, *Evolutionary Psychology*, vol. 11, no. 4, pp. 907–925, Oct. 2013 (cit. on p. 53).
- [Sla68] N. J. Slamecka, “An examination of trace storage in free recall.”, *Journal of Experimental Psychology*, vol. 76, no. 4, Pt.1, pp. 504–513, 1968, ISSN: 0022-1015. DOI: [10.1037/h0025695](https://doi.org/10.1037/h0025695). [Online]. Available: <http://dx.doi.org/10.1037/h0025695> (cit. on p. 24).
- [SLH10] H. Summerfelt, L. Lippman, and I. E. Hyman, “The effect of humor on memory: Constrained by the pun”, *The Journal of General Psychology*, vol. 137, no. 4, pp. 376–394, Oct. 2010, ISSN: 1940-0888. DOI: [10.1080/00221309.2010.499398](https://doi.org/10.1080/00221309.2010.499398). [Online]. Available: <http://dx.doi.org/10.1080/00221309.2010.499398> (cit. on p. 25).
- [SLW11] B. Sparrow, J. Liu, and D. M. Wegner, “Google effects on memory: Cognitive consequences of having information at our fingertips”, *Science*, vol. 333, no. 6043, pp. 776–778, Aug. 2011. DOI: [10.1126/science.1207745](https://doi.org/10.1126/science.1207745). [Online]. Available: <https://doi.org/10.1126/science.1207745> (cit. on p. 23).

- [SMB91] M. Spranca, E. Minsk, and J. Baron, “Omission and commission in judgment and choice”, *Journal of Experimental Social Psychology*, vol. 27, no. 1, pp. 76–105, Jan. 1991, ISSN: 0022-1031. DOI: [10.1016/0022-1031\(91\)90011-t](https://doi.org/10.1016/0022-1031(91)90011-t). [Online]. Available: [http://dx.doi.org/10.1016/0022-1031\(91\)90011-T](http://dx.doi.org/10.1016/0022-1031(91)90011-T) (cit. on p. 22).
- [SML+24] L. Shi, C. Ma, W. Liang, W. Ma, and S. Vosoughi, “Judging the judges: A systematic investigation of position bias in pairwise comparative assessments by llms”, *arXiv preprint arXiv:2406.07791*, 2024 (cit. on p. 60).
- [SNS21] M. A. Stoffel, S. Nakagawa, and H. Schielzeth, “Partr2: Partitioning r2 in generalized linear mixed models”, *PeerJ*, vol. 9, e11414, May 2021, ISSN: 2167-8359. DOI: [10.7717/peerj.11414](https://doi.org/10.7717/peerj.11414). [Online]. Available: <http://dx.doi.org/10.7717/peerj.11414> (cit. on p. 104).
- [SO21] S. I. Serengil and A. Ozpinar, “Hyperextended lightface: A facial attribute analysis framework”, in *2021 International Conference on Engineering and Emerging Technologies (ICEET)*, IEEE, 2021, pp. 1–4. DOI: [10.1109/ICEET53442.2021.9659697](https://doi.org/10.1109/ICEET53442.2021.9659697). [Online]. Available: <https://ieeexplore.ieee.org/document/9659697> (cit. on pp. 74, 77, 78).
- [SRRP07] T. Sharot, A. M. Riccardi, C. M. Raio, and E. A. Phelps, “Neural mechanisms mediating optimism bias”, *Nature*, vol. 450, no. 7166, pp. 102–105, Oct. 2007, ISSN: 1476-4687. DOI: [10.1038/nature06280](https://doi.org/10.1038/nature06280). [Online]. Available: <http://dx.doi.org/10.1038/nature06280> (cit. on p. 21).
- [SRRT16] G. Saposnik, D. Redelmeier, C. C. Ruff, and P. N. Tobler, “Cognitive biases associated with medical decisions: A systematic review”, *BMC Medical Informatics and Decision Making*, vol. 16, no. 1, Nov. 2016. DOI: [10.1186/s12911-016-0377-1](https://doi.org/10.1186/s12911-016-0377-1). [Online]. Available: <https://doi.org/10.1186/s12911-016-0377-1> (cit. on pp. 2, 7, 154).
- [SS92] G. Stasser and D. Stewart, “Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment.”, *Journal of Personality and Social Psychology*, vol. 63, no. 3, pp. 426–434, Sep. 1992. DOI: [10.1037/0022-3514.63.3.426](https://doi.org/10.1037/0022-3514.63.3.426). [Online]. Available: <https://doi.org/10.1037/0022-3514.63.3.426> (cit. on p. 16).
- [SSM80] S. R. Searle, F. M. Speed, and G. A. Milliken, “Population marginal means in the linear model: An alternative to least squares means”, *The American Statistician*, vol. 34, no. 4, pp. 216–221, Nov. 1980, ISSN: 1537-2731. DOI: [10.1080/00031305.1980.10483031](https://doi.org/10.1080/00031305.1980.10483031). [Online]. Available: <http://dx.doi.org/10.1080/00031305.1980.10483031> (cit. on p. 36).
- [Sta10] K. E. Stanovich, *Decision making and rationality in the modern world*. New York, Oxford University Press, 2010 (cit. on p. 12).
- [Sta97] B. M. Staw, “The escalation of commitment: An update and appraisal.”, 1997 (cit. on p. 26).

- [Str12] E. K. Strong, “The effect of length of series upon recognition memory.”, *Psychological Review*, vol. 19, no. 6, pp. 447–462, Nov. 1912, ISSN: 0033-295X. DOI: [10.1037/h0069812](https://doi.org/10.1037/h0069812). [Online]. Available: <http://dx.doi.org/10.1037/h0069812> (cit. on p. 24).
- [STW08] K. E. Stanovich, M. E. Toplak, and R. F. West, “The development of rational thought: A taxonomy of heuristics and biases”, in *Advances in Child Development and Behavior*, Elsevier, 2008, pp. 251–285. [Online]. Available: [https://doi.org/10.1016/s0065-2407\(08\)00006-2](https://doi.org/10.1016/s0065-2407(08)00006-2) (cit. on p. 7).
- [Sun02] C. R. Sunstein, “Probability neglect: Emotions, worst cases, and law”, *Yale Lj*, vol. 112, p. 61, 2002 (cit. on p. 21).
- [Sve70] O. Svenson, “A functional measurement approach to intuitive estimation as exemplified by estimated time savings.”, *Journal of Experimental Psychology*, vol. 86, no. 2, pp. 204–210, Nov. 1970, ISSN: 0022-1015. DOI: [10.1037/h0029934](https://doi.org/10.1037/h0029934). [Online]. Available: <http://dx.doi.org/10.1037/h0029934> (cit. on p. 20).
- [Sym24] P. M. Symonds, “On the loss of reliability in ratings due to coarseness of the scale.”, *Journal of Experimental Psychology*, vol. 7, no. 6, pp. 456–461, Dec. 1924, ISSN: 0022-1015. DOI: [10.1037/h0074469](https://doi.org/10.1037/h0074469). [Online]. Available: <http://dx.doi.org/10.1037/h0074469> (cit. on p. 46).
- [SZ88] W. Samuelson and R. Zeckhauser, “Status quo bias in decision making”, *Journal of Risk and Uncertainty*, vol. 1, no. 1, Mar. 1988. DOI: [10.1007/bf00055564](https://doi.org/10.1007/bf00055564). [Online]. Available: <https://doi.org/10.1007/bf00055564> (cit. on pp. 15, 25).
- [Tal16] S. N. Talamas, “Perceptions of intelligence and the attractiveness halo”, Ph.D. dissertation, University of St Andrews, 2016. [Online]. Available: <https://hdl.handle.net/10023/10851> (cit. on pp. 27, 33, 37, 42, 46, 50, 54, 56, 70, 115).
- [TB00] B. D. Till and M. Busler, “The match-up hypothesis: Physical attractiveness, expertise, and the role of fit on brand attitude, purchase intent and brand beliefs”, *Journal of Advertising*, vol. 29, no. 3, pp. 1–13, Oct. 2000, ISSN: 1557-7805. DOI: [10.1080/00913367.2000.10673613](https://doi.org/10.1080/00913367.2000.10673613). [Online]. Available: <http://dx.doi.org/10.1080/00913367.2000.10673613> (cit. on p. 74).
- [TBO04] K. Tentori, N. Bonini, and D. Osherson, “The conjunction fallacy: A misunderstanding about conjunction?”, *Cognitive Science*, vol. 28, no. 3, pp. 467–477, May 2004. DOI: [10.1207/s15516709cog2803_8](https://doi.org/10.1207/s15516709cog2803_8). [Online]. Available: https://doi.org/10.1207/s15516709cog2803_8 (cit. on p. 11).
- [TD08] A. Todorov and B. Duchaine, “Reading trustworthiness in faces without recognizing faces”, *Cognitive Neuropsychology*, vol. 25, no. 3, pp. 395–410, May 2008. DOI: [10.1080/02643290802044996](https://doi.org/10.1080/02643290802044996). [Online]. Available: <https://doi.org/10.1080/02643290802044996> (cit. on pp. 42, 46, 60, 115).
- [TD81] D. M. Taylor and J. R. Doria, “Self-serving and group-serving bias in attribution”, *The Journal of Social Psychology*, vol. 113, no. 2, pp. 201–211, Apr. 1981, ISSN: 1940-1183. DOI: [10.1080/00224545.1981.9924371](https://doi.org/10.1080/00224545.1981.9924371). [Online]. Available: <http://dx.doi.org/10.1080/00224545.1981.9924371> (cit. on p. 20).

- [TDP+13] A. Todorov, R. Dotsch, J. M. Porter, N. N. Oosterhof, and V. B. Falvello, “Validation of data-driven computational models of social perception of faces.”, *Emotion*, vol. 13, no. 4, pp. 724–738, Aug. 2013, ISSN: 1528-3542. DOI: [10.1037/a0032335](https://doi.org/10.1037/a0032335). [Online]. Available: <http://dx.doi.org/10.1037/a0032335> (cit. on pp. 33, 42, 46).
- [Tea25] G. Team, “Gemma 3”, 2025. [Online]. Available: <https://goo.gle/Gemma3Report> (cit. on p. 60).
- [TF23] A. N. Talboy and E. Fuller, *Challenging the appearance of machine intelligence: Cognitive bias in llms and best practices for adoption*, 2023. arXiv: 2304.01358 [cs.HC]. [Online]. Available: <https://arxiv.org/abs/2304.01358> (cit. on pp. 1, 58, 153).
- [TG06] R. Thornhill and S. W. Gangestad., “Facial sexual dimorphism, developmental stability, and susceptibility to disease in men and women.”, *Evolution and Human Behavior*, vol. 27, no. 2, pp. 131–144, Mar. 2006 (cit. on p. 54).
- [TH80] K. Timmerman and J. Hewitt, “Examining the halo effect of physical attractiveness”, *Perceptual and Motor Skills*, vol. 51, no. 2, pp. 607–612, Aug. 1980. DOI: [10.2466/pms.1980.51.2.607](https://doi.org/10.2466/pms.1980.51.2.607). [Online]. Available: <https://doi.org/10.2466/pms.1980.51.2.607> (cit. on pp. 29, 38, 39, 108, 156, 157).
- [Tha81] R. Thaler, “Some empirical evidence on dynamic inconsistency”, *Economics Letters*, vol. 8, no. 3, pp. 201–207, Jan. 1981. DOI: [10.1016/0165-1765\(81\)90067-7](https://doi.org/10.1016/0165-1765(81)90067-7). [Online]. Available: [https://doi.org/10.1016/0165-1765\(81\)90067-7](https://doi.org/10.1016/0165-1765(81)90067-7) (cit. on pp. 11, 20).
- [Tho99] S. C. Thompson, “Illusions of control: How we overestimate our personal influence”, *Current Directions in Psychological Science*, vol. 8, no. 6, pp. 187–190, Dec. 1999, ISSN: 1467-8721. DOI: [10.1111/1467-8721.00044](https://doi.org/10.1111/1467-8721.00044). [Online]. Available: <http://dx.doi.org/10.1111/1467-8721.00044> (cit. on p. 22).
- [TK74] A. Tversky and D. Kahneman, “Judgment under uncertainty: Heuristics and biases”, *Science*, vol. 185, no. 4157, pp. 1124–1131, Sep. 1974. DOI: [10.1126/science.185.4157.1124](https://doi.org/10.1126/science.185.4157.1124). [Online]. Available: <https://doi.org/10.1126/science.185.4157.1124> (cit. on pp. 9, 11, 19, 20, 23, 50).
- [TK81] A. Tversky and D. Kahneman, “The framing of decisions and the psychology of choice”, *Science*, vol. 211, no. 4481, pp. 453–458, 1981. DOI: [10.1126/science.7455683](https://doi.org/10.1126/science.7455683). [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.7455683> (cit. on pp. 1, 9, 10, 19, 153).
- [TK83] A. Tversky and D. Kahneman, “Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment.”, *Psychological Review*, vol. 90, no. 4, pp. 293–315, 1983. [Online]. Available: <https://doi.org/10.1037/0033-295x.90.4.293> (cit. on pp. 11, 20).
- [TK89] A. Tversky and D. Kahneman, “Rational choice and the framing of decisions”, in *Multiple Criteria Decision Making and Risk Analysis Using Microcomputers*. Springer Berlin Heidelberg, 1989, pp. 81–126, ISBN: 9783642749193. DOI: [10.1007/978-3-642-74919-3_4](https://doi.org/10.1007/978-3-642-74919-3_4). [Online]. Available: http://dx.doi.org/10.1007/978-3-642-74919-3_4 (cit. on pp. 9, 19).

- [TK94] A. Tversky and D. J. Koehler, “Support theory: A nonextensional representation of subjective probability.”, *Psychological Review*, vol. 101, no. 4, pp. 547–567, Oct. 1994, ISSN: 0033-295X. DOI: [10.1037/0033-295X.101.4.547](https://doi.org/10.1037/0033-295X.101.4.547). [Online]. Available: <http://dx.doi.org/10.1037/0033-295X.101.4.547> (cit. on p. 20).
- [Tod08] A. Todorov, “Evaluating faces on trustworthiness: An extension of systems for recognition of emotions signaling approach/avoidance behaviors”, *Annals of the New York Academy of Sciences*, vol. 1124, no. 1, pp. 208–224, Mar. 2008, ISSN: 1749-6632. DOI: [10.1196/annals.1440.012](https://doi.org/10.1196/annals.1440.012). [Online]. Available: <http://dx.doi.org/10.1196/annals.1440.012> (cit. on pp. 28, 33, 42, 46, 56).
- [TON16] K. Tagai, H. Ohtaka, and H. Nittono, “Faces with light makeup are better recognized than faces with heavy makeup”, *Frontiers in Psychology*, vol. 7, Mar. 2016, ISSN: 1664-1078. DOI: [10.3389/fpsyg.2016.00226](https://doi.org/10.3389/fpsyg.2016.00226). [Online]. Available: <http://dx.doi.org/10.3389/fpsyg.2016.00226> (cit. on pp. 28, 30).
- [TP66] E. Tulving and Z. Pearlstone, “Availability versus accessibility of information in memory for words”, *Journal of Verbal Learning and Verbal Behavior*, vol. 5, no. 4, pp. 381–391, Aug. 1966, ISSN: 0022-5371. DOI: [10.1016/S0022-5371\(66\)80048-8](https://doi.org/10.1016/S0022-5371(66)80048-8). [Online]. Available: [http://dx.doi.org/10.1016/S0022-5371\(66\)80048-8](http://dx.doi.org/10.1016/S0022-5371(66)80048-8) (cit. on p. 23).
- [TPB+22] G. V. Travaini, F. Pacchioni, S. Bellumore, M. Bosia, and F. De Micco, “Machine learning and criminal justice: A systematic review of advanced methodology for recidivism risk prediction”, *International journal of environmental research and public health*, vol. 19, no. 17, p. 10 594, 2022 (cit. on p. 49).
- [TSS17] H. Taniguchi, H. Sato, and T. Shirakawa, “Application of human cognitive mechanisms to naïve bayes text classifier”, *AIP Conference Proceedings*, vol. 1863, no. 1, p. 360 016, 2017. DOI: [10.1063/1.4992545](https://doi.org/10.1063/1.4992545) (cit. on pp. 2, 154).
- [VFL+20] V. Verrastro, L. Fontanesi, F. Liga, F. Cuzzocrea, and M. C. Gugliandolo, “Fear the instagram: Beauty stereotypes, body image and instagram use in a sample of male and female adolescents”, *Qwerty. Open and Interdisciplinary Journal of Technology, Culture and Education*, vol. 15, no. 1, Jun. 2020, ISSN: 2240-2950. DOI: [10.30557/qw000021](https://doi.org/10.30557/qw000021). [Online]. Available: <http://dx.doi.org/10.30557/qw000021> (cit. on pp. 41, 87, 159).
- [VGB+20] J. Vig, S. Gehrmann, Y. Belinkov, S. Qian, D. Nevo, S. Sakenis, J. Huang, Y. Singer, and S. Shieber, “Causal mediation analysis for interpreting neural nlp: The case of gender bias”, *arXiv preprint arXiv:2004.12265*, 2020 (cit. on p. 57).
- [VMH87] W. K. Viscusi, W. A. Magat, and J. Huber, “An investigation of the rationality of consumer valuations of multiple health risks”, *The RAND Journal of Economics*, vol. 18, no. 4, p. 465, 1987, ISSN: 0741-6261. DOI: [10.2307/2555636](https://doi.org/10.2307/2555636). [Online]. Available: <http://dx.doi.org/10.2307/2555636> (cit. on p. 22).

- [VPK25] S. Vijay, A. Priyanshu, and A. R. KhudaBukhsh, “When neutral summaries are not that neutral: Quantifying political neutrality in llm-generated news summaries (student abstract)”, *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 28, pp. 29 514–29 516, Apr. 2025, ISSN: 2159-5399. DOI: [10.1609/aaai.v39i28.35308](https://doi.org/10.1609/aaai.v39i28.35308). [Online]. Available: <http://dx.doi.org/10.1609/aaai.v39i28.35308> (cit. on p. 58).
- [VPP+23] C. Valuch, M. Pelowski, V.-T. Peltoketo, J. Hakala, and H. Leder, “Let’s put a smile on that face—a positive facial expression improves aesthetics of portrait photographs”, *Royal Society Open Science*, vol. 10, no. 10, Oct. 2023, ISSN: 2054-5703. DOI: [10.1098/rsos.230413](https://doi.org/10.1098/rsos.230413). [Online]. Available: <http://dx.doi.org/10.1098/rsos.230413> (cit. on p. 44).
- [Vra00] P. B. Vranas, “Gigerenzer’s normative critique of kahneman and tversky”, *Cognition*, vol. 76, no. 3, pp. 179–193, Sep. 2000, ISSN: 0010-0277. DOI: [10.1016/S0010-0277\(99\)00084-0](https://doi.org/10.1016/S0010-0277(99)00084-0). [Online]. Available: [http://dx.doi.org/10.1016/S0010-0277\(99\)00084-0](http://dx.doi.org/10.1016/S0010-0277(99)00084-0) (cit. on pp. 4, 160).
- [vRes33] H. von Restorff, “Über die wirkung von bereichsbildungen im spurenfeld”, *Psychologische Forschung*, vol. 18, no. 1, pp. 299–342, Dec. 1933, ISSN: 1430-2772. DOI: [10.1007/bf02409636](https://doi.org/10.1007/bf02409636). [Online]. Available: <http://dx.doi.org/10.1007/BF02409636> (cit. on p. 24).
- [VSW22] P. N. Venkit, M. Srinath, and S. Wilson, “A study of implicit bias in pretrained language models against people with disabilities”, in *Proceedings of the 29th International Conference on Computational Linguistics*, 2022, pp. 1324–1332 (cit. on p. 58).
- [WAAR66] E. Walster, V. Aronson, D. Abrahams, and L. Rottman, “Importance of physical attractiveness in dating behavior.”, *Journal of personality and social psychology*, vol. 4, no. 5, p. 508, 1966 (cit. on p. 25).
- [Was60] P. C. Wason, “On the failure to eliminate hypotheses in a conceptual task”, *Quarterly Journal of Experimental Psychology*, vol. 12, no. 3, pp. 129–140, Jul. 1960, ISSN: 0033-555X. DOI: [10.1080/17470216008416717](https://doi.org/10.1080/17470216008416717). [Online]. Available: <http://dx.doi.org/10.1080/17470216008416717> (cit. on p. 25).
- [Wat17] C. D. Watkins, “Creating beauty: Creativity compensates for low physical attractiveness when individuals assess the attractiveness of social and romantic partners”, *Royal Society Open Science*, vol. 4, no. 4, p. 160 955, Apr. 2017, ISSN: 2054-5703. DOI: [10.1098/rsos.160955](https://doi.org/10.1098/rsos.160955). [Online]. Available: <http://dx.doi.org/10.1098/rsos.160955> (cit. on pp. 40, 158).
- [WBDS11] J. Wood, D. Badawood, J. Dykes, and A. Slingsby, “Ballotmaps: Detecting name bias in alphabetically ordered ballot papers”, *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2384–2391, Dec. 2011, ISSN: 1077-2626 (cit. on p. 20).
- [WBT+24] P. Wang, S. Bai, S. Tan, S. Wang, Z. Fan, J. Bai, K. Chen, X. Liu, J. Wang, W. Ge, Y. Fan, K. Dang, M. Du, X. Ren, R. Men, D. Liu, C. Zhou, J. Zhou, and J. Lin, “Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution”, *arXiv preprint arXiv:2409.12191*, 2024 (cit. on p. 60).

- [WC24a] Y. Wan and K.-W. Chang, “White men lead, black women help: Uncovering gender, racial, and intersectional bias in language agency”, *arXiv preprint arXiv:2404.10508*, 2024 (cit. on pp. 58, 67).
- [WC24b] K. Wilson and A. Caliskan, “Gender, race, and intersectional bias in resume screening via language model retrieval”, *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, vol. 7, pp. 1578–1590, Oct. 2024, ISSN: 3065-8365. DOI: [10.1609/aies.v7i1.31748](https://doi.org/10.1609/aies.v7i1.31748). [Online]. Available: <http://dx.doi.org/10.1609/aies.v7i1.31748> (cit. on p. 58).
- [WC98] M. Weber and C. F. Camerer, “The disposition effect in securities trading: An experimental analysis”, *Journal of Economic Behavior & Organization*, vol. 33, no. 2, pp. 167–184, Jan. 1998, ISSN: 0167-2681. DOI: [10.1016/S0167-2681\(97\)00089-9](https://doi.org/10.1016/S0167-2681(97)00089-9). [Online]. Available: [http://dx.doi.org/10.1016/S0167-2681\(97\)00089-9](http://dx.doi.org/10.1016/S0167-2681(97)00089-9) (cit. on p. 25).
- [WCP+24] Z. Wu, X. Chen, Z. Pan, X. Liu, W. Liu, D. Dai, H. Gao, Y. Ma, C. Wu, B. Wang, Z. Xie, Y. Wu, K. Hu, J. Wang, Y. Sun, Y. Li, Y. Piao, K. Guan, A. Liu, X. Xie, Y. You, K. Dong, X. Yu, H. Zhang, L. Zhao, Y. Wang, and C. Ruan, *Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding*, 2024. arXiv: [2412.10302](https://arxiv.org/abs/2412.10302) [cs.CV]. [Online]. Available: <https://arxiv.org/abs/2412.10302> (cit. on p. 60).
- [WFVL19] Y. Wang, J. Fardouly, L. R. Vartanian, and L. Lei, “Selfie-viewing and facial dissatisfaction among chinese adolescents: A moderated mediation model of general attractiveness internalization and body appreciation”, *Body Image*, vol. 30, pp. 35–43, Sep. 2019, ISSN: 1740-1445. DOI: [10.1016/j.bodyim.2019.05.001](https://doi.org/10.1016/j.bodyim.2019.05.001). [Online]. Available: <http://dx.doi.org/10.1016/j.bodyim.2019.05.001> (cit. on pp. 3, 50, 74, 156).
- [WG05] T. D. Wilson and D. T. Gilbert, “Affective forecasting: Knowing what to want”, *Current Directions in Psychological Science*, vol. 14, no. 3, pp. 131–134, Jun. 2005, ISSN: 1467-8721. DOI: [10.1111/j.0963-7214.2005.00355.x](https://doi.org/10.1111/j.0963-7214.2005.00355.x). [Online]. Available: <http://dx.doi.org/10.1111/j.0963-7214.2005.00355.x> (cit. on p. 20).
- [WGC+23] J. Wu, W. Gan, Z. Chen, S. Wan, and S. Y. Philip, “Multimodal large language models: A survey”, in *2023 IEEE International Conference on Big Data (Big-Data)*, IEEE, 2023, pp. 2247–2256 (cit. on p. 56).
- [Wil81] B. Williams, *Moral luck: philosophical papers 1973-1980*. Cambridge University Press, 1981 (cit. on p. 22).
- [Wil95] D. L. Wiley, “Beauty and the beast: Physical appearance discrimination in american criminal trials”, *St. Mary’s Law Journal*, vol. 27, 1995. [Online]. Available: <https://commons.stmarytx.edu/thestmaryslawjournal/vol27/iss1/6> (cit. on p. 28).
- [WJS+18] Z. Wang, M. Jusup, L. Shi, *et al.*, “Exploiting a cognitive bias promotes cooperation in social dilemma experiments”, *Nature Communications*, vol. 9, no. 1, Jul. 2018. DOI: [10.1038/s41467-018-05259-5](https://doi.org/10.1038/s41467-018-05259-5). [Online]. Available: <https://doi.org/10.1038/s41467-018-05259-5> (cit. on p. 9).

- [WK96] N. D. Weinstein and W. M. Klein, “Unrealistic optimism: Present and future”, *Journal of Social and Clinical Psychology*, vol. 15, no. 1, pp. 1–8, Mar. 1996, ISSN: 0736-7236. DOI: [10.1521/jscp.1996.15.1.1](https://doi.org/10.1521/jscp.1996.15.1.1). [Online]. Available: <http://dx.doi.org/10.1521/jscp.1996.15.1.1> (cit. on p. 21).
- [WL19] W. Wattanacharoensil and D. La-ornual, “A systematic review of cognitive biases in tourist decisions”, *Tourism Management*, vol. 75, pp. 353–369, Dec. 2019. [Online]. Available: <https://doi.org/10.1016/j.tourman.2019.06.006> (cit. on pp. 2, 7, 154).
- [WLG22] Y. Wang, S. Luan, and G. Gigerenzer, “Modeling fast-and-frugal heuristics”, *PsyCh Journal*, vol. 11, no. 4, pp. 600–611, Jul. 2022. DOI: [10.1002/pchj.576](https://doi.org/10.1002/pchj.576). [Online]. Available: <https://doi.org/10.1002/pchj.576> (cit. on p. 16).
- [WM08] D. H. Wedell and R. Moro, “Testing boundary conditions for the conjunction fallacy: Effects of response mode, conceptual focus, and problem type”, *Cognition*, vol. 107, no. 1, pp. 105–136, Apr. 2008. DOI: [10.1016/j.cognition.2007.08.003](https://doi.org/10.1016/j.cognition.2007.08.003). [Online]. Available: <https://doi.org/10.1016/j.cognition.2007.08.003> (cit. on p. 11).
- [WM84] M. Wernick and G. J. Manaster, “Age and the perception of age and attractiveness”, *The Gerontologist*, vol. 24, no. 4, pp. 408–414, Aug. 1984, ISSN: 1758-5341. DOI: [10.1093/geront/24.4.408](https://doi.org/10.1093/geront/24.4.408). [Online]. Available: <http://dx.doi.org/10.1093/geront/24.4.408> (cit. on pp. 35, 40, 158).
- [WNG23] Y. Wu, Y. Nakashima, and N. Garcia, “Stable diffusion exposed: Gender bias from prompt to image”, *arXiv preprint arXiv:2312.03027*, 2023 (cit. on p. 74).
- [WQK+20] Z. Wang, K. Qinami, I. C. Karakozis, K. Genova, P. Nair, K. Hata, and O. Russakovsky, “Towards fairness in visual recognition: Effective strategies for bias mitigation.”, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020 (cit. on pp. 3, 50, 74, 156).
- [WR15] J. P. Wilson and N. O. Rule, “Facial trustworthiness predicts extreme criminal-sentencing outcomes”, *Psychological Science*, vol. 26, no. 8, pp. 1325–1331, Jul. 2015. DOI: [10.1177/0956797615590992](https://doi.org/10.1177/0956797615590992). [Online]. Available: <https://doi.org/10.1177/0956797615590992> (cit. on p. 28).
- [WR21] A. Wang and O. Russakovsky, “Directional bias amplification”, in *Proceedings of the 38th International Conference on Machine Learning*, M. Meila and T. Zhang, Eds., ser. Proceedings of Machine Learning Research, vol. 139, PMLR, Jul. 2021, pp. 10 882–10 893. [Online]. Available: <https://proceedings.mlr.press/v139/wang21t.html> (cit. on p. 70).
- [WT06] J. Willis and A. Todorov, “First impressions: Making up your mind after a 100-ms exposure to a face”, *Psychological Science*, vol. 17, no. 7, pp. 592–598, Jul. 2006, ISSN: 1467-9280. DOI: [10.1111/j.1467-9280.2006.01750.x](https://doi.org/10.1111/j.1467-9280.2006.01750.x). [Online]. Available: <http://dx.doi.org/10.1111/j.1467-9280.2006.01750.x> (cit. on p. 42).

- [WTG21] J. Walker, S. J. Tepper, and T. Gilovich, “People are more tolerant of inequality when it is expressed in terms of individuals rather than groups at the top”, *Proceedings of the National Academy of Sciences*, vol. 118, no. 43, Oct. 2021, ISSN: 1091-6490. DOI: [10.1073/pnas.2100430118](https://doi.org/10.1073/pnas.2100430118). [Online]. Available: <http://dx.doi.org/10.1073/pnas.2100430118> (cit. on p. 63).
- [WV13] D. Walker and E. Vul, “Hierarchical encoding makes individuals in a group seem more attractive”, *Psychological Science*, vol. 25, no. 1, pp. 230–235, Oct. 2013. [Online]. Available: <https://doi.org/10.1177/0956797613497969> (cit. on p. 23).
- [XLC+24] Y. Xiao, A. Liu, Q. Cheng, Z. Yin, S. Liang, J. Li, J. Shao, X. Liu, and D. Tao, “Genderbias-VL: Benchmarking gender bias in vision language models via counterfactual probing”, *arXiv preprint arXiv:2407.00600*, 2024 (cit. on pp. 62, 127).
- [XWKG20] T. Xu, J. White, S. Kalkan, and H. Gunes., “Investigating bias and fairness in facial expression recognition”, in *Investigating Bias and Fairness in Facial Expression Recognition*, Springer International Publishing, 2020, pp. 506–523 (cit. on p. 51).
- [YAAB20] S. Yucer, S. Akcay, N. Al-Moubayed, and T. P. Breckon., “Exploring racial bias within face recognition via per-subject adversarially-enabled data augmentation.”, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Jun. 2020 (cit. on pp. 3, 50, 74, 156).
- [YCLH23] K.-C. Yeh, J.-A. Chi, D.-C. Lian, and S.-K. Hsieh, “Evaluating interfaced llm bias”, in *Proceedings of the 35th Conference on Computational Linguistics and Speech Processing (ROCLING 2023)*, 2023, pp. 292–299 (cit. on p. 57).
- [YLL+24] Z. Ying, A. Liu, S. Liang, L. Huang, J. Guo, W. Zhou, X. Liu, and D. Tao, “Safebench: A safety evaluation framework for multimodal large language models”, *arXiv preprint arXiv:2410.18927*, 2024 (cit. on p. 58).
- [YR22] T. Yasseri and J. Reher, “Fooled by facts: Quantifying anchoring bias through a large-scale experiment”, *Journal of Computational Social Science*, vol. 5, no. 1, pp. 1001–1021, Jan. 2022. DOI: [10.1007/s42001-021-00158-0](https://doi.org/10.1007/s42001-021-00158-0). [Online]. Available: <https://doi.org/10.1007/s42001-021-00158-0> (cit. on p. 9).
- [YS98] D. W. Yu and G. H. Shepard, “Is beauty in the eye of the beholder?”, *Nature*, vol. 396, no. 6709, pp. 321–322, Nov. 1998, ISSN: 1476-4687. DOI: [10.1038/24512](https://doi.org/10.1038/24512). [Online]. Available: <http://dx.doi.org/10.1038/24512> (cit. on p. 38).
- [Zaj68] R. B. Zajonc, “Attitudinal effects of mere exposure.”, *Journal of Personality and Social Psychology*, vol. 9, no. 2, Pt.2, pp. 1–27, 1968, ISSN: 0022-3514. DOI: [10.1037/h0025848](https://doi.org/10.1037/h0025848). [Online]. Available: <http://dx.doi.org/10.1037/h0025848> (cit. on p. 25).
- [ZBL07] L. A. Zebrowitz, P. M. Bronstad, and H. K. Lee, “The contribution of face familiarity to ingroup favoritism and stereotyping”, *Social Cognition*, vol. 25, no. 2, pp. 306–338, Apr. 2007, ISSN: 0278-016X. DOI: [10.1521/soco.2007.25.2.306](https://doi.org/10.1521/soco.2007.25.2.306). [Online]. Available: <http://dx.doi.org/10.1521/soco.2007.25.2.306> (cit. on pp. 33, 40, 42, 46, 115, 158).

- [ZLJ22] K. Zhou, E. Lai, and J. Jiang, “VLStereoSet: A study of stereotypical bias in pre-trained vision-language models”, in *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Y. He, H. Ji, S. Li, Y. Liu, and C.-H. Chang, Eds., Online only: Association for Computational Linguistics, Nov. 2022, pp. 527–538. [Online]. Available: <https://aclanthology.org/2022.aacl-main.40> (cit. on pp. 58, 63, 129, 130).
- [ZPR+14] N. Zhang, M. Paluri, M. Ranzato, T. Darrell, and L. Bourdev, “Panda: Pose aligned networks for deep attribute modeling”, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2014 (cit. on p. 152).
- [ZR04] L. A. Zebrowitz and G. Rhodes., “Sensitivity to “bad genes” and the anomalous face overgeneralization effect: Cue validity, cue utilization, and accuracy in judging intelligence and health.”, *Journal of Nonverbal Behavior*, vol. 28, no. 3, pp. 167–185, 2004 (cit. on p. 54).
- [ZWC+24] J. Zhang, S. Wang, X. Cao, Z. Yuan, S. Shan, X. Chen, and W. Gao, “Vlbiasbench: A comprehensive benchmark for evaluating bias in large vision-language model”, *arXiv preprint arXiv:2406.14194*, 2024 (cit. on pp. 51, 58, 59, 67).
- [ZWW25] Y. Zhao, B. Wang, and Y. Wang, *Explicit vs. implicit: Investigating social bias in large language models through self-reflection*, 2025. arXiv: 2501 . 02295 [cs.CL]. [Online]. Available: <https://arxiv.org/abs/2501.02295> (cit. on p. 58).
- [ZWY+18] J. Zhao, T. Wang, M. Yatskar, V. Ordonez, and K.-W. Chang, “Gender bias in coreference resolution: Evaluation and debiasing methods”, in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, Association for Computational Linguistics, 2018. DOI: 10.18653/v1/n18-2003. [Online]. Available: <http://dx.doi.org/10.18653/v1/N18-2003> (cit. on p. 57).
- [ZYD+24] D. Zhang, Y. Yu, J. Dong, C. Li, D. Su, C. Chu, and D. Yu, “Mm-llms: Recent advances in multimodal large language models”, *ACL*, 2024 (cit. on p. 56).

