AI-Generated Faces Could Be Making Real People Less Employable

Originally published Mon 19 Oct 2026 at https://www.unite.ai/ai-generated-faces-could-be-making-real-people-less-employable/



You don't really need scientists to know that personal attractiveness opens doors, since you've likely been the victim or the beneficiary of this to some extent in your life; nor to discover that AI-generated faces tend to be more attractive than real people (for a variety of reasons, many involving human-made preferences and processes).

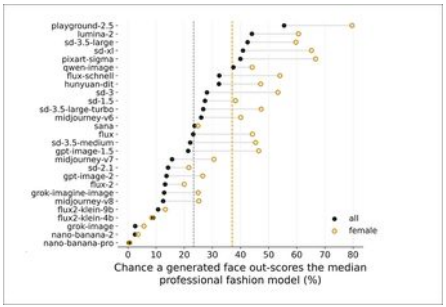
Nonetheless, a related, new academic collaboration from the US caught my eye today, citing evidence that AI-generated beauty may be costing real people job interviews. In a study involving 570 people, those previously exposed to unusually attractive AI-generated faces subsequently judged real job applicants as less employable and less trustworthy; and their willingness to offer interviews fell from 77.4% to 72.5%, even though they were assessing *exactly the same applicants*.



Examples of AI-generated faces from the study, showing how widely their attractiveness scores varied. Researchers tested 27 image generators using ordinary prompts that did not mention beauty. *The top row shows faces given some of the highest attractiveness scores by automated rating systems; the bottom row shows some of the lowest. All eight people pictured are AI-generated, not real.* Source

The researchers also investigated how attractive AI-generated faces appear compared to real ones, using images produced by diverse commercial and open-weight models. Asking 550 people to rate photographs of real people alongside AI-generated faces, it was found that the synthetic faces consistently received higher scores; and in an additional comparison, apposite metrics found that 37% of AI-generated female faces were rated more attractive *than the typical professional fashion model*.

© Martin Anderson and/or commissioning entities



Comparison of AI-generated faces with professional fashion models, across 27 image generators. The chart shows the percentage of generated faces scoring higher than the median professional model, using automated attractiveness measures. Black dots represent all generated faces, while yellow circles represent female faces only. Overall, 23% of generated faces exceeded this benchmark, rising to 37% for female faces. The dashed lines mark these two averages.

The paper† states:

> *'We audited 27 text-to-image generators released between 2022 and 2026 using simple prompts that never mention attractiveness or beauty, analyzing 51,300 generated faces and 160,196 photographs of real people across primary and supplementary comparisons.*
>
> *'We used three automated measures and collected ratings from 550 human raters on a subset of images.*
>
> *'Nearly all generators produced faces scoring above photographed faces under automated measures, and human raters judged the outputs of all 23 generators as significantly more attractive than photographed faces, even after correction for multiple comparisons.*
>
> *'When focusing our analysis on gender, we find that 37% of female-prompted generated faces outscored the median professional fashion model. '*

The authors also note that while experts – both human and AI – may disagree and debate around who is top of the 'beautiful' quadrant, there is no such granularity in the *'not attractive'* or '*minimally attractive'* ratings range, which is, by comparison, a kind of generic dustbin*:

> *'The greater cross-measure agreement over who is valued least than over who is valued most resembles an important feature of aesthetic classification in settings like fashion modeling: people and institutions do not have to agree on a single ideal in order to produce recurring inequalities in whose appearance receives less value.'*

Face Values

Many of the study's finding reiterate the common human experience of *contextual and relative perception*: for instance, if you were the smartest or most attractive person in your school, entering an Ivy League university may tend to greatly raise the ambient standards, and lower your relative position in the new pool you've just entered.

However, the world is not full of AI-generated people actually walking about and elbowing flesh-and-blood folk aside at job interviews; so how can the cultural phenomenon of genAI beauty aesthetics make any difference?

Well, besides the alarming possibility that people may resort to plastic surgery in order to more closely adhere to the 'AI standard', the effect described in the new study is similar to the pre-AI culture of beauty in movies, TV, fashion, and other visible and affecting sectors of human culture: an environment wherein we were/are saturated with unrealistic physical ideals; and where we learn to deal with it as a kind of signal-to-noise issue. We dream of movie stars, but marry who we can.

The paper describes the growing influence of AI's own visual aesthetics in a similar context*:

> *'If generated faces are both more attractive than real ones and unequally so, [the] question follows: what are the consequences for how real people are judged? Exposure to highly selected or edited imagery alters appearance comparison and subsequent judgment, and classic contrast experiments show that exposure to highly attractive people can lower subsequent attractiveness judgments.*
>
> *'Generative AI can deliver repeated streams of distinct faces shaped by the same aesthetic processes. Whether exposure to AI-generated faces alters consequential evaluations of real people remains unclear.'*

The authors propose no specific remedy, instead calling for further research to establish when AI-generated beauty affects our judgments of real people; what aspects of synthetic faces drive this effect; and whether the results would hold outside the simulated hiring experiment.

Method and Data Curation

To establish how widely AI image generators favor attractive faces, the researchers tested 27 systems released between 2022 and 2026, namely Stable Diffusion 1.5, 2.1, XL, 3, 3.5-medium, 3.5-large and 3.5-large-turbo; FLUX.1 dev and schnell, FLUX.2, FLUX.2-klein 9B, and FLUX.2-klein 4B; GPT Image 1.5, and GPT Image 2; Grok Image, and Grok Imagine; Nano Banana 2 and Nano Banana Pro; Midjourney v6, v7 and v8; PixArt-Sigma; Playground 2.5; HunyuanDiT; Qwen-Image; Sana; and Lumina-2.

The same simple prompt was used across all systems to generate portraits representing 14 demographic categories across male and female categories and eight occupations (*doctor*; *lawyer*; *veterinarian*; *pharmacist*; *cook*; *construction worker*; *nurse and retail worker*), without mentioning beauty or attractiveness.

Prompts omitting demographic, sex or occupation details were also tested, producing a total of 51,300 images for analysis.

© Martin Anderson and/or commissioning entities

For comparison, 94,249 photographs were drawn from FairFace; IdenProf; SCUT-FBP5500; MEBeauty; the Chicago Face Database and its multiracial and Indian extensions; and the Face Research Lab London Set. Usable faces were detected in 89,435 photographs.



Diverse western and Asian faces from the SCUT-FBP5500 dataset, one of six used in the new study. Source

Additional comparisons were made using the 10k US Adult Faces Database; a collection of 15,182 professional fashion-model photographs; LookismBench-SyntheticFaces; FFHQ; and StyleGAN2. The fashion-model photographs were used as a 'high-attractiveness' benchmark, while the remaining collections were analyzed separately to assess the attractiveness measures and provide further comparisons.

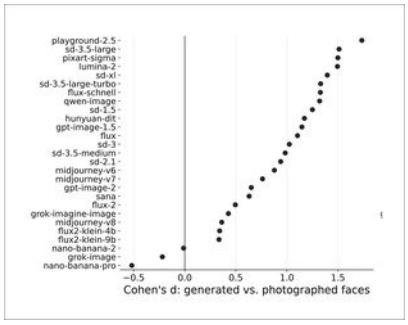
Faces were detected and cropped using MTCNN, then standardized. Attractiveness was assessed using ResNet-50 models trained on SCUT-FBP5500 and MEBeauty; and SwinFace, trained on CelebA annotations.

To check whether the automated attractiveness scores matched human opinions, a selection of 1,374 real and AI-generated faces was assessed by 550 US-based adults.

In a separate experiment, 570 US adults were divided into three groups and shown job applicants with either real photographs, retouched photographs or AI-generated faces. Afterwards, the same real applicants were assessed by all three groups, allowing differences in hiring decisions to be measured. The study's main analyses were preregistered.

Tests

Under both main attractiveness metrics, AI-generated faces scored significantly higher than real photographs for 24 of the 27 generators tested, a finding also supported by the CelebA/SwinFace metric, as shown below:

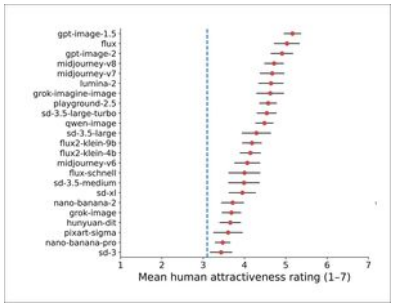


Attractiveness differences between AI-generated faces and real photographs across 27 image generators, expressed as average Cohen's d scores across three automated metrics. Positive values indicate higher attractiveness scores for generated faces.

Across all three metrics, the paper reports, 13 generators produced faces rated more than one standard deviation above real photographs – though no advantage was found for Nano Banana Pro, Grok Image or Nano Banana 2.

The tendency to produce unusually attractive faces was less pronounced in more recent Midjourney and OpenAI models, while the effect was roughly twice as strong in open-weight models as in commercial offerings from OpenAI, Google, Midjourney and xAI.

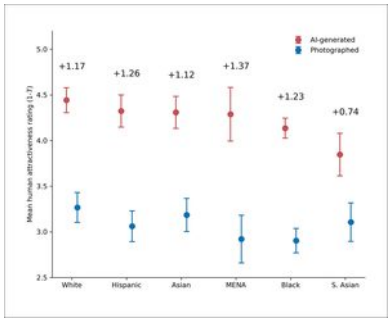
In human evaluations, faces produced by all 23 generators tested with enough images for a reliable comparison were judged more attractive than real photographs. On a scale of 1 (least attractive) to 7 (most attractive), real faces received an average score of 3.10, while AI-generated faces scored between 0.34 and 2.06 points higher, depending on the generator, as indicated below:



Human attractiveness ratings for faces produced by 23 AI image generators, on a scale of 1 to 7. Generators scored above average for real photographs, indicated by blue dashed line. Red dots indicate average ratings, horizontal bars the uncertainty around each average.

Higher ratings for AI-generated faces were recorded across all six racial and ethnic categories (see graph below), even for Nano Banana Pro, Grok Image and Nano Banana 2 – whose faces had scored no higher than real photographs in the tests:

© Martin Anderson and/or commissioning entities



Human attractiveness ratings for AI-generated faces (red) and real photographs (blue) across six racial and ethnic categories. AI-generated faces received higher average scores in every category, with differences ranging from 0.74 to 1.37 points on a seven-point scale. Vertical bars indicate the uncertainty around each average.

Human ratings also differed from the automated results for newer models. Little change was recorded between OpenAI's GPT Image 1.5 and GPT Image 2, while progressively higher attractiveness ratings were recorded for Midjourney versions 6, 7 and 8.

In other words, newer AI models were still judged (by people) to produce unusually attractive faces, even when the automated metrics suggested that this tendency had diminished.

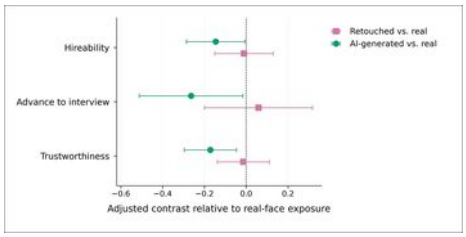
The Unfair Advantage

The final experiment investigated whether exposure to attractive AI-generated faces could affect how real job applicants were judged.

For a simulated 'Customer Success Associate' vacancy, 570 US adults were divided into three groups. After assessing four identical applicants, each group was shown eight profiles featuring either real photographs; retouched photographs; or AI-generated faces. The AI faces had previously scored 5.75 out of seven for attractiveness, compared to a score of 2.99 for real photographs. The retouched images had been smoothed and brightened using an effect based on Facetune's One-tap Enhance feature.

Twelve further applicants, with identical photographs and qualifications across all three groups, were included, to test whether exposure to the different images affected hiring decisions.

Applicants with AI-generated faces received slightly higher employability and trustworthiness ratings, though no significant increase in interview invitations was found. After exposure to these faces, however, the *same real applicants were rated less employable and trustworthy*, with interview invitations falling from 77.4% to 72.5%:



Hiring test results showing changes in real applicants' perceived employability, trustworthiness and likelihood of receiving an interview after exposure to different face types. Lower evaluations followed exposure to AI-generated faces, but not retouched photographs. Points show differences from the real-photo control group; vertical lines indicate 95% confidence intervals.

No comparable decline was recorded after exposure to retouched photographs, meaning that real humans supposedly looking their best did not 'raise the bar' as the AI images had done.

The authors comment:

> *'Candidates paired with AI-generated faces received only slightly better direct ratings, but participants who saw those faces subsequently rated the same real candidates more negatively on all three outcomes.*
>
> *'The AI-generated candidates themselves received slightly higher direct trustworthiness ratings, which makes a simple generalized dislike of the AI-faced profiles less likely.*
>
> *'Still, the experiment tells us that exposure mattered, not why.'*

Conclusion

Anecdotally, AI-generated characters have certainly come to dominate street-level advertising in the city where I live, but that's 2D art. Whether or not the use of diffusion-based generative *video* in TV and movies will overcome current legal and ethical restraints, remains to be seen; if it does, we can perhaps expect a far wider population of AI-generated and/or AI-augmented actors 'raising the aesthetic standard', and bringing to pass the scenario of 'ambient AI beauty' that the new paper outlines.

© Martin Anderson and/or commissioning entities

The principle is nothing new; but the potential for automation is. In such a case, one could distill the essence of current or former real-world celebrities, models and public representatives, propagate or 'improve' them, and diffuse them broadly, and without any expensive contract negotiations. But that's yet to come, if come it will.

† *Titled* Generative AI Produces and Transmits Beauty Bias*, from eight researchers across Cornell, the University of Florida, and the University of Pennsylvania.*

* *My conversion of the authors' inline citations to hyperlinks.*

First published Monday, October 19, 2026

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.
Website: martinanderson.ai
Contact: martin@martinanderson.ai