Humans and variational autoencoders agree: attractive faces are average and feminine, not symmetric and young

Francisco M. López (lopez@fias.uni-frankfurt.de)
Frankfurt Institute for Advanced Studies
Frankfurt am Main, Germany

Jochen Triesch (triesch@fias.uni-frankfurt.de)
Frankfurt Institute for Advanced Studies
Frankfurt am Main, Germany

Abstract

Why do we find some faces more attractive than others? Four properties are frequently identified as sources of beauty: symmetry, youth, femininity, and averageness. Recent experiments show that only the latter two determine facial attractiveness. Here, we test whether a neural network trained with unsupervised learning can reproduce and explain this phenomenon. We train a variational autoencoder (VAE) on face images and estimate its preference judgments as the compression of a face in the latent space. We find that, like humans, the VAE's preferences are significantly correlated with the averageness and the femininity of the faces but not their symmetry or youth. Furthermore, the VAE correlates with human attractiveness judgments. In sum, this work suggests that human aesthetic face preferences can be explained by the efficiency with which a face can be encoded.

Keywords: face perception; facial attractiveness; perceptual fluency; efficient coding; unsupervised learning

Introduction

Why do we find some faces more attractive than others? The perceptual fluency theory posits that faces are preferred when they are easy to process (Reber et al., 2004), leading to more efficient visual representations, namely by exploiting redundancies that facilitate compression (López et al., 2024). Two facial features are frequently identified as sources of redundancy related to beauty: symmetry and averageness (Locher & Nodine, 1989; Ryali et al., 2020; Burton et al., 2005; Rhodes et al., 1999). Both are shown to correlate with human aesthetic judgments independently (Jones et al., 2007; Bertamini et al., 2019; Trujillo et al., 2014). Additionally, the femininity of a face, regardless of its gender, also correlates with beauty scores (Hoss et al., 2005). However, experiments that control for all three properties reveal that averageness and femininity, but not symmetry, determines facial attractiveness (Lee et al., 2025; Kleisner et al., 2024). It is currently unclear whether machine learning models trained to represent human faces can reproduce, and help to explain, this behavior. That is the question that we aim to answer in this work.

(a) Original images

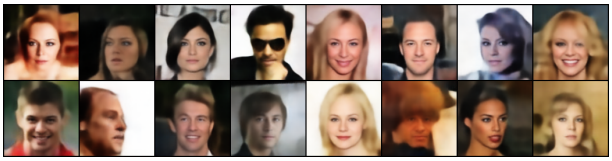

(b) VAE reconstructions

Figure 1: Samples from the CelebA dataset used for training.

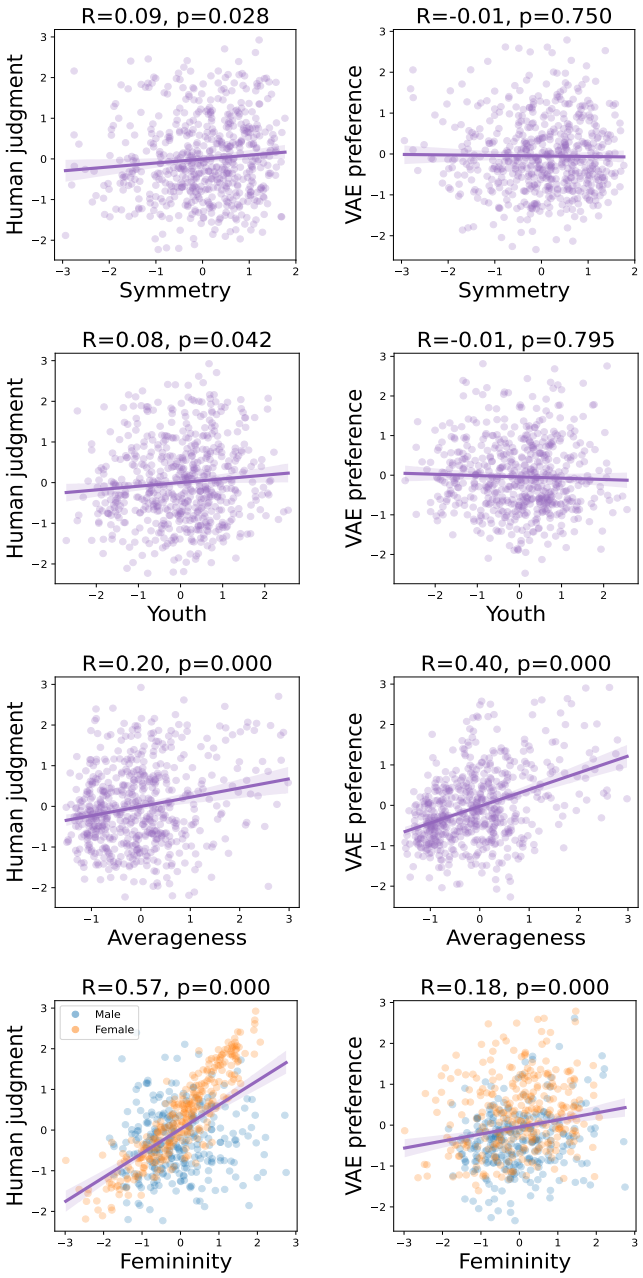

Figure 2: Facial attractiveness scores by humans (left) and VAE (right) as functions of symmetry, youth, averageness, and femininity. Plot titles show the results of the linear regressions (correlation coefficient and significance). The femininity plots show the dots colored by gender but the linear regressions are performed over both sets of faces.

Methods

We reproduce human attractiveness judgment experiments using a variational autoencoder (VAE) trained on the CelebA dataset of celebrity faces (see Fig. 1). The model's preference for a face is given by its processing fluency: the more easily a face's pixel image can be compressed into a compact representation, the more it preferred. We calculate

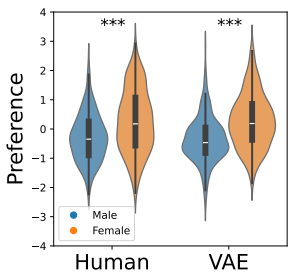

Figure 3: Distributions of preference judgments split by gender.

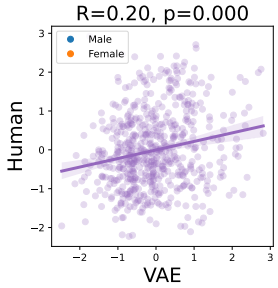

Figure 4: Correlation between human judgments and VAE preferences.

the compression $c(\mathbf{x})$ of an image $\mathbf{x}$ as defined by: $c(\mathbf{x}) = -\mathcal{L}_{VAE}(\mathbf{x})/E(\mathbf{x})$, where $E(\cdot)$ is an energy function that sums the total local intensity changes in the image and $\mathcal{L}_{VAE}(\cdot)$ is the loss function of the VAE (see López et al., 2024). This compression function favors images with high information and low error, following the efficient coding principles.

Using the Chicago Face Database (Ma et al., 2015), which includes calibrated face images and corresponding human judgments of attractiveness, femininity, and age, we compare the preferences of humans and VAE. Symmetry is computed as the negative average distance between corresponding left and right facial markers detected by a MediaPipe face mesh (Lugaresi et al., 2019). We compute a face's averageness as its negative mean squared error in pixels with the corresponding average per category (ethnicity and gender). All values are z-scored over all faces, except femininity, which is z-scored separately for males and females.

Results

Averageness and femininity, but not symmetry or youth, align with human and VAE preferences

The main result of this contribution is shown in Fig. 2. We explore the dependence of perceptual preferences for humans (attractiveness judgments) and VAE (compression) on symmetry, youth, averageness, and femininity. First, we find that symmetry and youth are only weakly correlated with attractiveness in the human judgments, and uncorrelated with the VAE's preference. On the other hand, there are significant dependences of both human and model preferences on the facial averageness and femininity. This confirms the recent results of human judgments reported by Lee et al. (2025). While the VAE's preference for average faces may be expected, the preference for feminine faces but not for symmetric ones is a surprising finding that provides new insights about what makes a face more compressible, i.e. easy to process.

VAE preferences correlate with human judgments

Given the results of Fig. 2, we investigate whether the perceptual preferences of humans and VAE are themselves correlated. The results are shown in Fig. 4. We find a significant correlation, indicating that the VAE aligns with the human judgments. Additionally, we perform separate correlation analyses for each ethnicity, which returns significant results for all cases (not shown). The outcomes of this analysis provides further evidence that human aesthetic preferences are based on processing fluency, as captured by the VAE's compression.

Humans favor femininity, VAE favors averageness

Finally, we investigate the relative contributions of averageness and femininity to perceptual preferences. No significant correlations were found between these two properties for this dataset. We then perform a multiple linear regression for human attractiveness judgments against averageness and femininity, and repeat it for VAE compression preferences. We find significant values for both features in both cases. However, as visualized in Fig. 5 and quantified by the regression coefficients, human judgments depend most on femininity whereas VAE preferences favor averageness. Nevertheless, both humans and VAE have significantly higher preferences for female faces than male faces (Fig. 3), in line with previous human preference reports (e.g. Lee et al., 2025).

Conclusions

We have shown that VAEs achieve higher compression for average and feminine faces rather than symmetrical or young ones, in line with human attractiveness judgments. What do these results reveal about face processing?

VAEs and other autoencoders learn to perform lossy compression of inputs into low-dimensional latent spaces. Doing so requires identifying and exploiting redundancies in the image statistics while neglecting uncommon details. Therefore, we propose that some faces are preferred not because they are inherently attractive but because they are prototypical and thus easier to compress. Our work can explain that human aesthetic preferences are driven by processing fluency and efficient coding.

Given the positive outcome of this study, it remains to be seen whether other aesthetic preferences in vision and other sensory modalities can also emerge in neural networks without additional input control, or whether the preference for compression needs to be enforced during training (López et al., 2024). This will be the topic of future research.

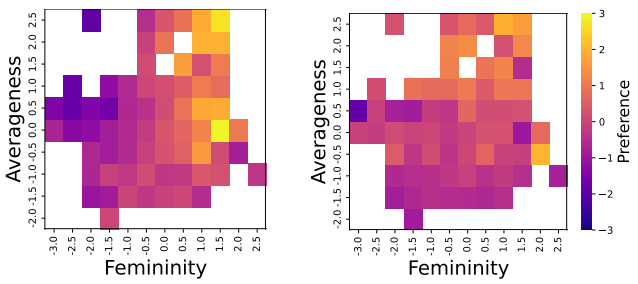

(a) Human judgments (b) VAE preference

Figure 5: Human attractiveness judgments and VAE preferences as functions of both facial averageness and femininity.

Acknowledgments

This research was supported by "The Adaptive Mind" funded by the Excellence Program of the Hessian Ministry of Higher Education, Research, Science and the Arts, Germany, and by the German Research Foundation (DFG) Project numbers 520617944, 520223571 ("Sensing LOOPS"). JT was supported by the Johanna Quandt foundation.

References

Bertamini, M., Rampone, G., Makin, A. D., & Jessop, A. (2019). Symmetry preference in shapes, faces, flowers and landscapes. *PeerJ*, *7*, e7078.

Burton, A. M., Jenkins, R., Hancock, P. J., & White, D. (2005). Robust representations for face recognition: The power of averages. *Cognitive psychology*, *51*(3), 256–284.

Hoss, R. A., Ramsey, J. L., Griffin, A. M., & Langlois, J. H. (2005). The role of facial attractiveness and facial masculinity/femininity in sex classification of faces. *Perception*, *34*(12), 1459–1474.

Jones, B. C., DeBruine, L. M., & Little, A. C. (2007). The role of symmetry in attraction to average faces. *Perception & psychophysics*, *69*, 1273–1277.

Kleisner, K., Tureček, P., Saribay, S. A., Pavlovič, O., Leongómez, J. D., Roberts, S. C., ... others (2024). Distinctiveness and femininity, rather than symmetry and masculinity, affect facial attractiveness across the world. *Evolution and Human Behavior*, *45*(1), 82–90.

Lee, P., Li, J., Rafiee, Y., Jones, B. C., & Shiramizu, V. K. (2025). Further evidence that averageness and femininity, rather than symmetry and masculinity, predict facial attractiveness judgments. *Scientific Reports*, *15*(1), 5498.

Locher, P., & Nodine, C. (1989). The perceptual value of symmetry. In *Symmetry 2* (pp. 475–484). Elsevier.

Lugaresi, C., Tang, J., Nash, H., McClanahan, C., Uboweja, E., Hays, M., ... others (2019). Mediapipe: A framework for perceiving and processing reality. In *Third workshop on computer vision for ar/vr at ieee computer vision and pattern recognition (cvpr)* (Vol. 2019).

López, F. M., Shi, B. E., & Triesch, J. (2024). Prioritizing compression explains human perceptual preferences. In *Intrinsically-motivated and open-ended learning workshop@ neurips2024.*

Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The chicago face database: A free stimulus set of faces and norming data. *Behavior research methods*, *47*, 1122–1135.

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and social psychology review*, *8*(4), 364–382.

Rhodes, G., Sumich, A., & Byatt, G. (1999). Are average facial configurations attractive only because of their symmetry? *Psychological science*, *10*(1), 52–58.

Ryali, C. K., Goffin, S., Winkielman, P., & Yu, A. J. (2020). From likely to likable: The role of statistical typicality in human social assessment of faces. *Proceedings of the National Academy of Sciences*, *117*(47), 29371–29380.

Trujillo, L. T., Jankowitsch, J. M., & Langlois, J. H. (2014). Beauty is in the ease of the beholding: A neurophysiological test of the averageness theory of facial attractiveness. *Cognitive, Affective, & Behavioral Neuroscience*, *14*, 1061–1076.