

Revealing an Unattractivity Bias in Mental Reconstruction of Occluded Faces using Generative Image Models

Frederik Riedmann (frederik.riedmann@fau.de)

Departement Informatik

Bernhard Egger* (bernhard.egger@fau.de)

Departement Informatik

Tim Rohe* (tim.rohe@fau.de)

Institut für Psychologie

Friedrich-Alexander-Universität Erlangen-Nürnberg
Erlangen, Germany

Abstract

Previous studies have shown that faces are rated as more attractive when they are partially occluded. The cause of this observation remains unclear. One explanation is a mental reconstruction of the occluded face parts which is biased towards a more attractive percept as shown in face-attractiveness rating tasks. We aimed to test for this hypothesis by using a delayed matching-to-sample task, which directly requires mental reconstruction. In two online experiments, we presented observers with unattractive, neutral or attractive synthetic reconstructions of the occluded face parts using a state-of-the-art diffusion-based image generator. Our experiments do not support the initial hypothesis and reveal an unattractiveness bias for occluded faces instead. This suggests that facial attractiveness rating tasks do not prompt reconstructions. Rather, the attractivity bias may arise from global image features, and faces may actually be reconstructed with unattractive properties when mental reconstruction is applied.

Keywords: Human Face Perception; Face Occlusion; Mental Reconstruction; Generative Models; Stable Diffusion

Intro and Research Question

During the COVID-19 pandemic, medical masks not only helped to prevent infections, but also increased the wearer's perceived attractiveness: Numerous studies have shown that observers rate human faces as more attractive when they are partially occluded (Orghian & Hidalgo, 2020; Patel, Mazzaferro, Sarwer, & Bartlett, 2020; Kamatani, Ito, Miyazaki, & Kawahara, 2021; Pazhoohi & Kingstone, 2022; Hies & Lewis, 2022). A recent hypothesis is that observers reconstruct occluded face parts using an averaged face template based on the faces they have seen previously in their life. Especially, for unfamiliar faces from other races than their own, this template would be skewed towards attractive faces (e.g. by celebrities), resulting in an attractivity bias (Kamatani, Miyazaki, & Kawahara, 2023).

Previous studies used a simple task to measure the attractiveness bias: Participants were presented with a face with partial occlusion (e.g., superimposed mask or cropped image) or a full face and rated its overall attractiveness. However, this face-attractiveness rating (FAR) task does not require observers to mentally reconstruct the occluded face part, but observers provide holistic attractiveness ratings which may be biased by a general positive judgment of face masks (e.g., indicating responsible health behavior). In our study, we designed a delayed matching-to-sample (DMTS, (Paule et al., 1998)) task that requires observers to mentally reconstruct the occluded face part based on the visible face features: After a brief delay, observers matched one masked or unmasked face (sample face) to unmasked comparison faces with either attractive, neutral or unattractive mouth area of the face (match faces). The delay required observers to recollect the sample face from memory including a mental reconstruction of the occluded face part to perform the matching operation. Even though we could replicate the attractivity bias in the standard FAR task, our results from the DMTS task show that observers rather reconstruct the missing face part with an unattractive version, revealing an unattractivity bias in mental reconstruction of the missing face parts.

Methods

For our experiments, we created five different images for 100 face identities: **i)** Original image **ii)** Masked version, i.e. original image with a medical face mask covering mouth and nose **iii)** Unattractive version with replaced oral area of low attractiveness **iv)** Neutral version with replaced oral area of neutral attractiveness **v)** Attractive version with replaced oral area of high attractiveness. Examples can be seen in Figure 1.

The original and masked versions were taken from the MaskedFace-Net dataset (Cabani et al., 2020). To create the unattractive, neutral and attractive version, the same area that was occluded in the masked version was inpainted using the “stable-diffusion-2-inpainting” model (Rombach, Blattmann, Lorenz, Esser, & Ommer, 2022). Multiple prompts for faces with a neutral expression and low or high attractiveness were embedded, and a linear extrapolation between these vectors

*These two authors contributed equally to this work.

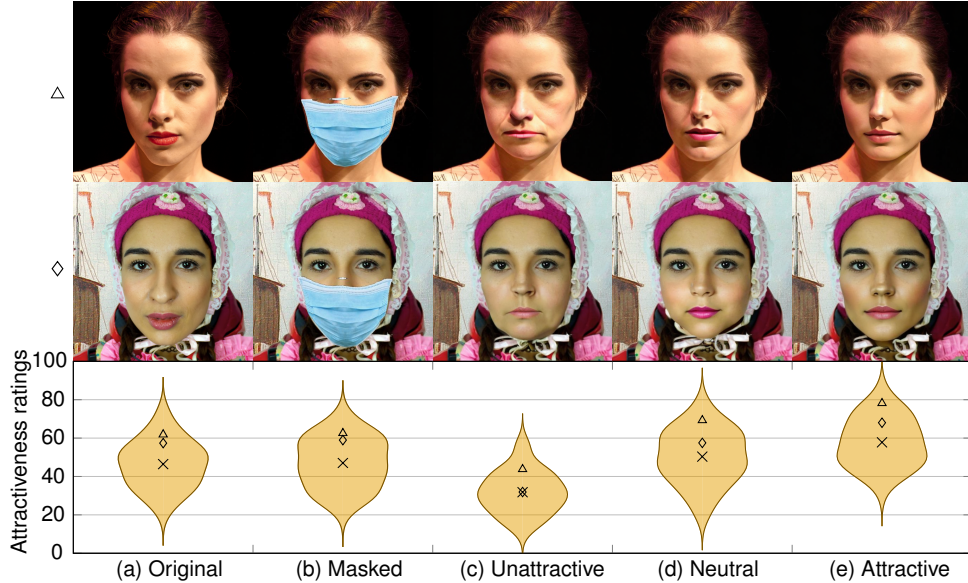


Figure 1: All five variations of exemplary identities. Violin plots display the distribution of all identities' mean attractiveness ratings. The two examples are highlighted in the violin plots and \times marks the mean value. Original and masked images provided by the MaskedFace-Net (Cabani et al., 2020) (CC BY-NC-SA 4.0).

was used to generate images of different levels of attractiveness. All prompts were structurally similar to following examples, but words were replaced with synonyms or increases: **i)** “an attractive beautiful person with neutral facial expression” **ii)** “a good looking person with a perfect face and neutral facial expression” **iii)** “an ugly person with neutral facial expression” **iv)** “a hideous appalling person with neutral facial expression”. In addition to common negative prompts like “injuries, unhealthy, mutated body, bad anatomy, deformed, bad composition, disfigured”, we also discouraged “open mouth, emotion, smile” and “frown”.

Using these images, we performed two online surveys on SoSci Survey to conduct the following experiments (Leiner, 2024): **i)** In Experiment 1, we implemented a standard FAR task as a replication and control: Participants were presented with all five versions of each image and rated the faces' overall attractiveness ratings using a slider (mapped to a 100-point scale). **ii)** In Experiment 2, we implemented a DMTS task using two experimental conditions: The participants were shown the masked version or the unmasked original image as a sample face for 1.25 seconds. This image disappeared and after a delay of 0.5 seconds, the three unattractive, neutral or attractive inpainted face images appeared in random spatial order. Participants were asked to select one face as a match for which they believed it was most similar to the sample face. For the masked faces, this task required participants to mentally reconstruct the missing face part. The unmasked sample face served as a control condition which also required recall of the sample face from memory, but no mental reconstruction of the missing part because the face was fully visible. Thus, comparison to this condition controlled for memory processes and response biases when choosing the sample face. Further, each individual observer was randomly presented with each individual face ($n = 99$ trials in total) in either masked or unmasked version to avoid retesting effects.

After giving written informed consent, 126 participants com-

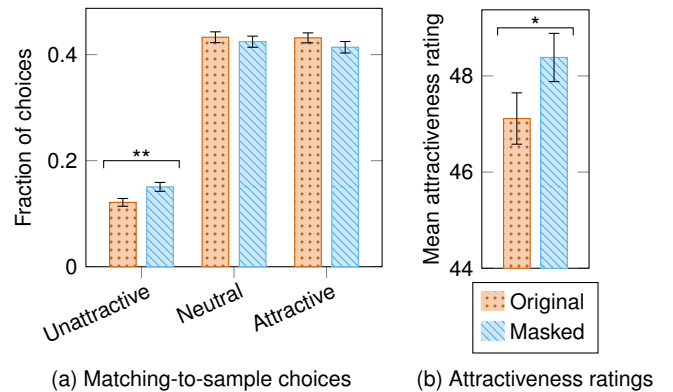


Figure 2: Survey results: Error bars indicate standard error.

pleted survey 1 (81 male, 43 female, 2 other; mean age 25.2 ± 8.2 STD years) and 41 participants completed survey 2 (21 female, 20 male; mean age 34.4 ± 17.3 STD years). All images were re-categorized into the unattractive, neutral and attractive categories based on participants' attractiveness ratings in all subsequent analyses.

Results

The FAR task in experiment 1 showed that we were able to replicate the previous attractiveness bias for masked over unmasked original faces (one-sided t-test, $t_{2027} = 1.893$, $p = 0.029$; see Figure 2b). Critically, our DMTS task in experiment 2 revealed the opposite result, that is an unattractiveness bias for masked faces (Figure 2a): For each participant and masked versus unmasked sample face, we computed the fraction of trials in which either the unattractive, neutral or attractive face was chosen as a match. A repeated-measures ANOVA with factors sample-face masking (yes vs. no) and match-face version (unattractive vs. attractive) showed that participants generally preferred the attractive

over the unattractive match face (main effect, $F_{1,125} = 659.73$, $p < 0.001$). Yet, for masked as compared to unmasked sample faces, participants more often chose the unattractive but less often chose the attractive match face (interaction effect, $F_{1,125} = 5.48$, $p = 0.02$). Two-sided t-tests confirmed that participants chose the unattractive version more often for a masked sample ($t_{125} = 2.98$, $p = 0.004$). This was not the case for the attractive versions ($t_{125} = -1.26$, $p = 0.208$) or the neutral match faces ($t_{125} = -0.61$, $p = 0.545$).

Conclusion

Our experiments replicated that face masks create an attractiveness bias in overall attractiveness ratings. In contrast to this established finding (Orghian & Hidalgo, 2020; Patel et al., 2020; Kamatani et al., 2021; Pazhoohi & Kingstone, 2022; Hies & Lewis, 2022), our DMTS task revealed an unattractiveness bias for masked faces. Because the DMTS task requires observers to mentally “fill-in” the masked part, this result contradicts the idea that observers mentally reconstruct occluded face parts by using an attractive average face template. Rather, the previously reported attractiveness bias may result from general factors that make a masked face appear more attractive in overall attractiveness rating tasks. For example, masks may increase perceived attractiveness by indicating responsible health behavior or high social status of medical staff. Further, an attractiveness rating task can be performed without mentally reconstructing occluded face parts as hypothetically assumed. In conclusion, by combining a DMTS task with a state-of-the-art inpainting model, we show that the classical attractiveness bias does not arise from mental reconstruction because this cognitive operation leads to an unattractiveness bias instead.

References

- Cabani, A., Hammoudi, K., Benhabiles, H., & Melkemi, M. (2020). Maskedface-net – a dataset of correctly/incorrectly masked face images in the context of covid-19. *Smart Health*. doi: <https://doi.org/10.1016/j.smhl.2020.100144>
- Hies, O., & Lewis, M. B. (2022, December). Beyond the beauty of occlusion: medical masks increase facial attractiveness more than other face coverings. *Cogn. Research*, 7(1), 1–6. doi: 10.1186/s41235-021-00351-9
- Kamatani, M., Ito, M., Miyazaki, Y., & Kawahara, J. I. (2021, May). Effects of Masks Worn to Protect Against COVID-19 on the Perception of Facial Attractiveness. *i-Perception*, 12(3), 20416695211027920. doi: 10.1177/20416695211027920
- Kamatani, M., Miyazaki, Y., & Kawahara, J. I. (2023, January). Occlusion of faces by sanitary masks improves facial attractiveness of other races. *Front. Psychol.*, 13, 953389. doi: 10.3389/fpsyg.2022.953389
- Leiner, D. J. (2024). *Sosci survey (version 3.5.01) [computer software]*. Retrieved 2024-02-15, from <https://www.soscisurvey.de>
- Orghian, D., & Hidalgo, C. A. (2020, January). Humans judge faces in incomplete photographs as physically more attractive. *Sci. Rep.*, 10(110), 1–12. doi: 10.1038/s41598-019-56437-4
- Patel, V., Mazzaferro, D. M., Sarwer, D. B., & Bartlett, S. P. (2020, August). Beauty and the Mask. *Plastic and Reconstructive Surgery Global Open*, 8(8). doi: 10.1097/GOX.0000000000003048
- Paule, M. G., Bushnell, P. J., Maurissen, J. P. J., Wenger, G. R., Buccafusco, J. J., Chelonis, J. J., & Elliott, R. (1998, September). Symposium Overview: The Use of Delayed Matching-to-Sample Procedures in Studies of Short-Term Memory in Animals and Humans. *Neurotoxicol. Teratol.*, 20(5), 493–502. doi: 10.1016/S0892-0362(98)00013-0
- Pazhoohi, F., & Kingstone, A. (2022, December). Unattractive faces are more attractive when the bottom-half is masked, an effect that reverses when the top-half is concealed. *Cogn. Research*, 7(1), 1–10. doi: 10.1186/s41235-022-00359-9
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022, June). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (cvpr)* (p. 10684-10695).