Supplementary Material: Brain-Supervised Image Editing

Keith M. Davis III[1], Carlos de la Torre-Ortiz[1], Tuukka Ruotsalo[1,2]
first.last@helsinki.fi
[1]University of Helsinki, Helsinki, Finland
[2]University of Copenhagen, Copenhagen, Denmark

1 Datasets and Models

To protect the privacy of the participants (see Ethics Statement below), raw EEG data collected from the neurophysiological experiment will not be made publicly available. However, these data are available from the corresponding author upon reasonable request.

The processed EEG data, which have been epoched and averaged into 100ms windows such that the privacy of participants has been preserved, *are available to the public*. These data, as well as the GAN model, labeled facial stimuli, latent vectors used to produce the stimuli, and corresponding code used to produce the results found in this manuscript, can all be found at:
https://github.com/Cognitive-Computing-Group/Brain-Supervised-Image-Editing/.

2 Neurophysiological Experiment

Here we provide additional information on exactly how the neurophysiological experiment was conducted, including equipment type and EEG configuration.

An LCD monitor, running at 60 Hz with a resolution of 1680 by 1050 pixels and positioned 60 cm from the participants, was used to present stimuli. The stimulus presentation software used was Psychology Software Tools E-Prime 3.0.3.60, which assisted with optimizing the timing of the display and EEG amplifier trigger control.

The EEG data were recorded using 32 Ag/AgCl electrodes with EasyCap elastic caps (EasyCap GmbH, Herrsching, Germany). The positions of the electrodes were placed equidistant around the scalp according to the 10-10 system. A diagram of the electrode positions is shown in Figure 1. An electrode at AFz was used as ground. A QuickAmp USB amplifier (BrainProducts GmbH, Gilching, Germany) running at 2000 Hz was used for hardware amplification, filtering, and digitization of the signal.

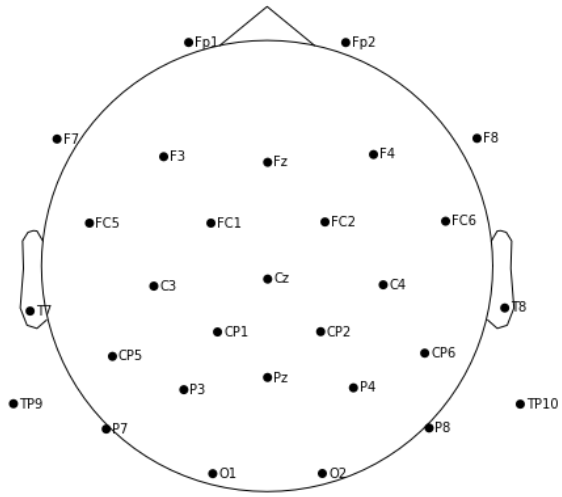

Figure 1: Electrode position diagram (not to scale)

3 Evaluation Experiment

For each of the 30 participants, 2 source images were randomly sampled from the stimuli used in each of the 8 saliency recognition tasks. Each source image was transformed by the respective participant's brain model, the explicit label model, and the two random control models. Thus, a total of 1,920 sets of transformed images were produced and evaluated by the primary assessor. We randomized the order in which each set of images was presented. Within each set, the order of the images was also randomized. Although the results of the 4 separate models (the brain model and 3 baselines) were being evaluated, the assessors were not made aware of this at any time to ensure this knowledge did not influence their assessments. Examples of the graphical user interface (GUI) used by the assessors are shown in Figure 2 and 3.

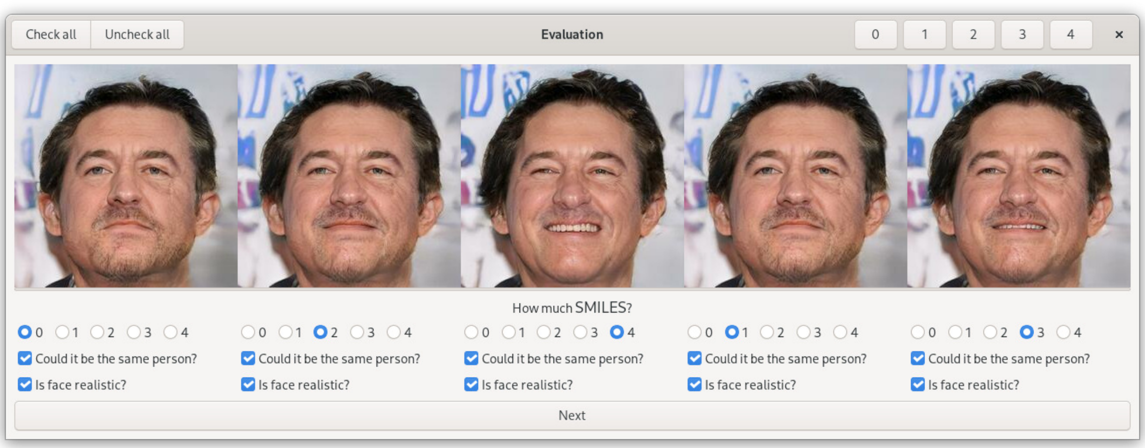

Figure 2: GUI prompt for tasks smile and no smile. The radio buttons from 0 to 4 under each image are used to annotate the saliency of the semantic feature of interest, with 0 indicating minimal salience and 4 indicating maximum salience. The checkboxes *"Could it be the same person?"* and *"Is face realistic?"* are used to indicate identity preservation and realism, respectively. For *"Could it be the same person?"*, the assessors were instructed to make a guess as to which image was the original and compare the other images to it.

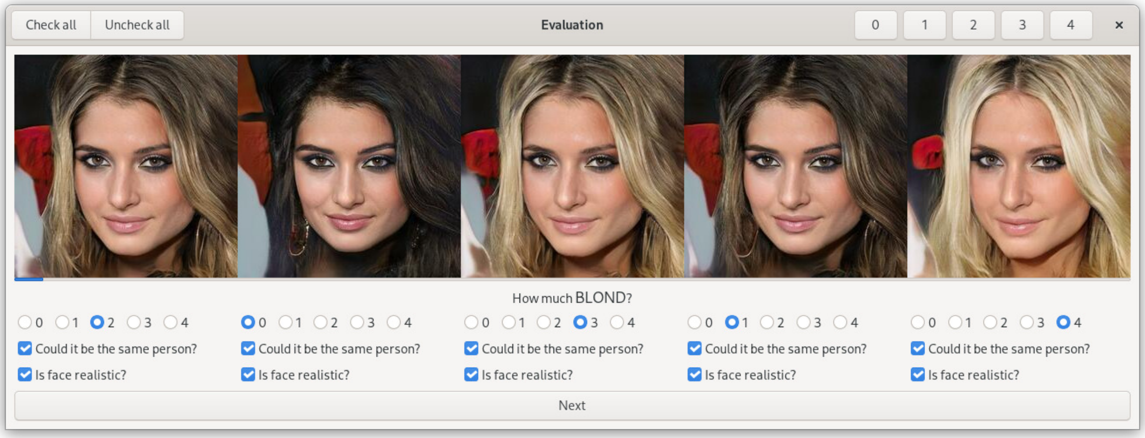

Figure 3: GUI prompt for tasks blond and dark-haired.

4 Results

Here, provide additional examples of our results that could not fit within the paper. Additional results for tasks *blond*, *dark hair*, *female*, *male*, *no smile*, *smile*, *old* and *young*, are found respectively in Figures 4, 5, 6, 7, 8, 9, 10, and 11. While it is certainly worth showing additional examples of the brain model completing the semantic editing task correctly, the authors also believe the reader may benefit from viewing some exceptions where the brain model did not fully succeed in the editing task. In the overwhelming majority of these cases, the explicit model also performed poorly. This suggests that other factors may be negatively affecting the results, such as a small training set (only a few hundred examples were used to learn features within a 512-dimensional space for each semantic saliency task) or flaws directly attributed to the GAN model itself.

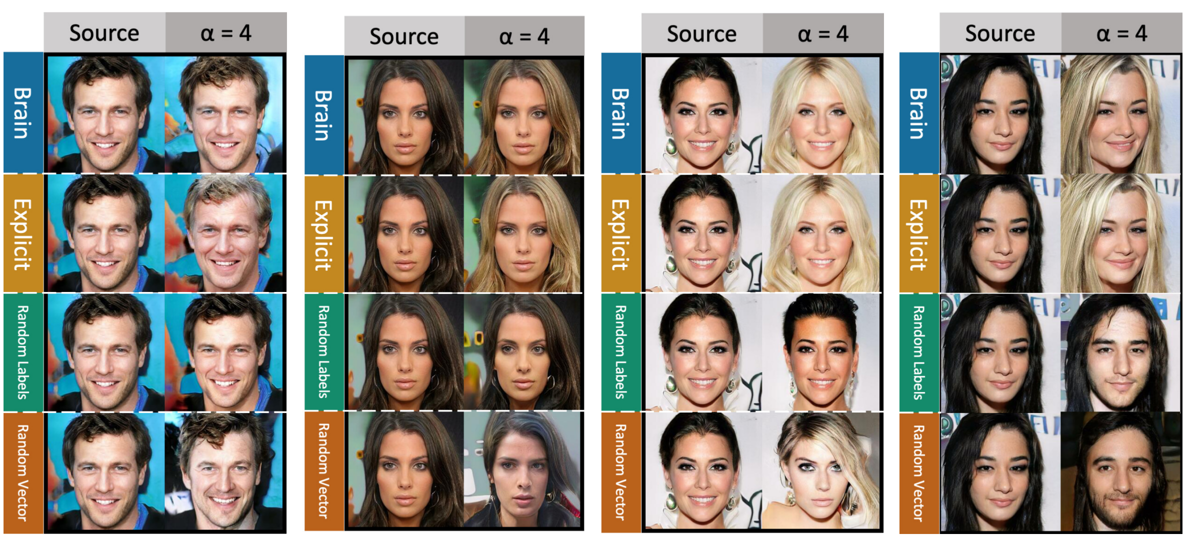

(a) Delta = 0 (b) Delta = 1 (c) Delta = 4 (d) Delta = 4

Figure 4: Example results for task *blond*, selected to illustrate variations in model performance. Recall that the goal of the task is to transform an image of a person into one with blond hair. In 4a and 4b, the final images produced by the brain model has only marginally lighter hair, whereas the explicit model produces images with more visibly blond hair given the same α. In 4c and 4d, both models produce results that are nearly indistinguishable from each other.

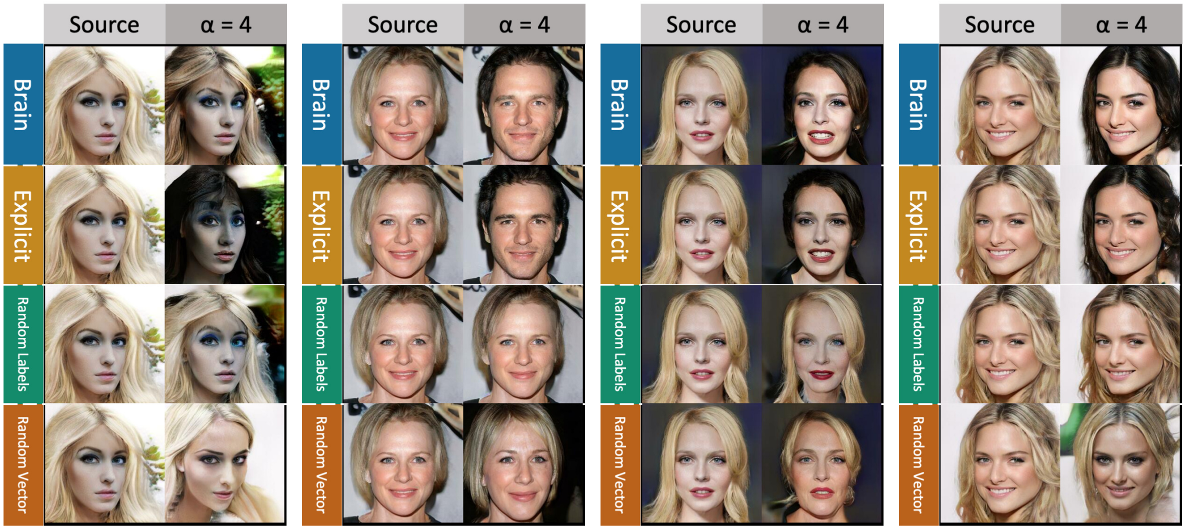

(a) Delta = 1 (b) Delta = 3 (c) Delta = 4 (d) Delta = 4

Figure 5: Example results for task *dark-haired*, selected to illustrate variations in model performance. Recall that the goal of the task is to transform an image of a person with light hair (blond) into one with dark hair. In 5a, the brain model only produces slightly darker hair compared to the explicit model. However, the explicit model also significantly darkens the rest of the image, and ultimately produces an image that is not very realistic. As in tasks *old* and *young*, occasionally the facial structure and hair length of the source images were also transformed, as in 5b, In general, both models performed the task quite well and produced similar results, as in 5c and 5d.

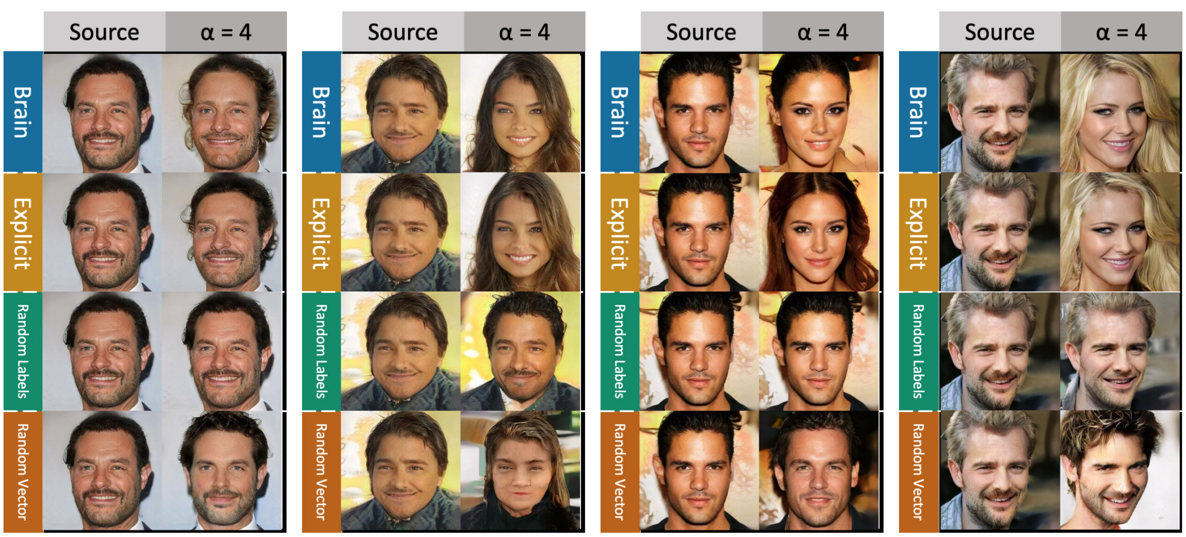

(a) Delta = 0 (b) Delta = 3 (c) Delta = 4 (d) Delta = 4

Figure 6: Example results for task *female*, selected to illustrate variations in model performance. Recall that the goal of the task is to transform an image of a person into one who appears more female. While in general this was one of the easiest tasks for the models to perform successfully as can be seen in 6c and 6d, for 2 of the 60 source images transformed, the figure remained unchanged as in 6a. When clothing was visible, the models did not always preserve these features, as can be seen in 6b.

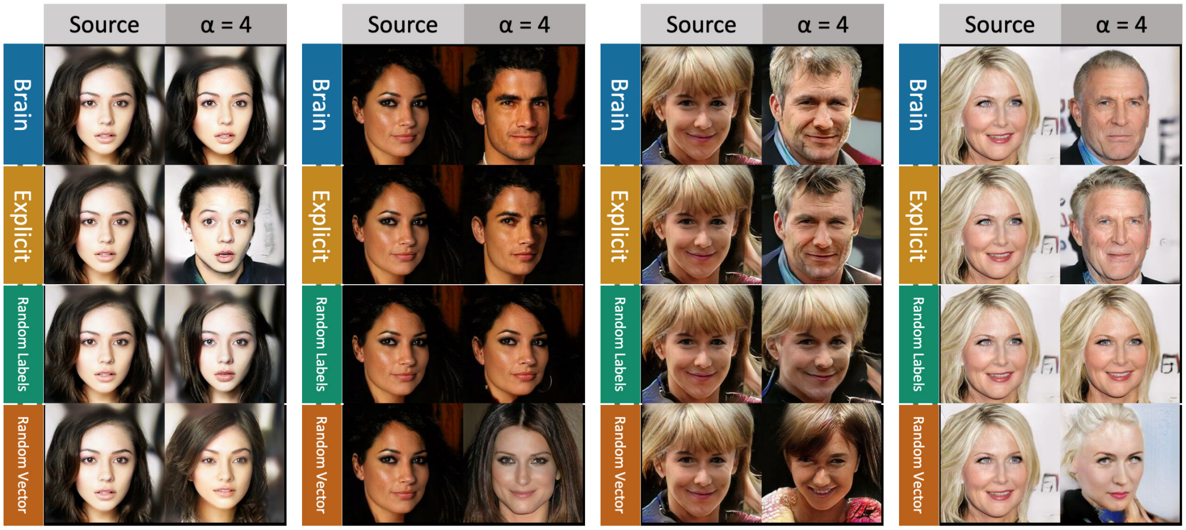

(a) Delta = 0 (b) Delta = 3 (c) Delta = 4 (d) Delta = 4

Figure 7: Example results for task *male*, selected to illustrate variations in model performance. Recall that the goal of the task is to transform an image of a person who appears female into one who appears more male. As with task *female*, while the vast majority of images were successfully transformed to ones that appeared more female, occasionally the model did not adequately modify a source, such as in 7a. In 7b, the explicit model performs slightly better in preserving other facial features, such as nose shape, although the result has a less stereotypically male facial structure. In general, the results for the brain model and explicit model were very similar, as seen in 7c and 7d.

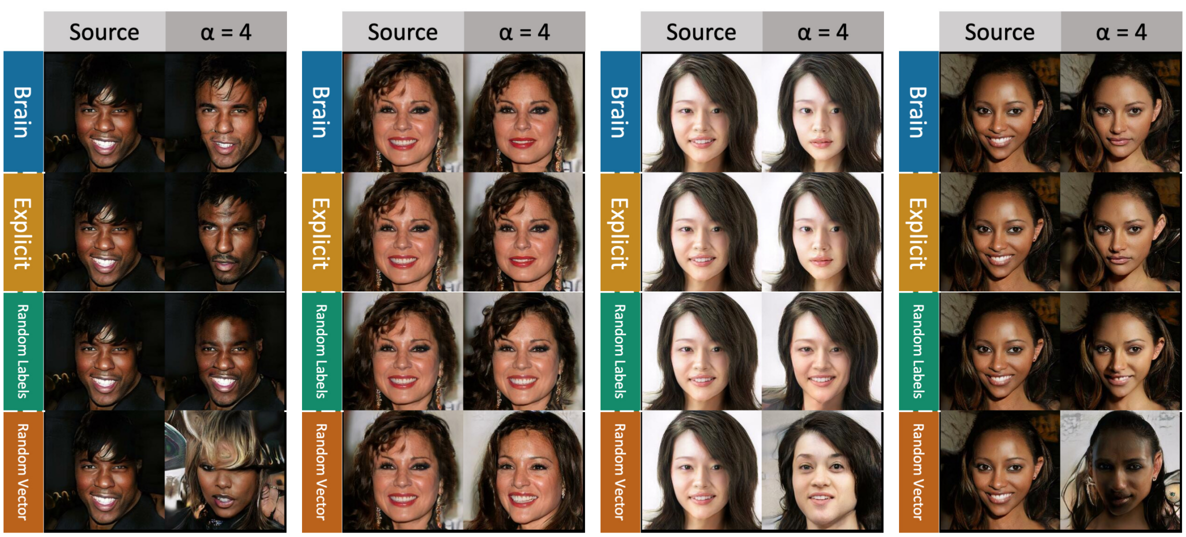

(a) Delta = 0 (b) Delta = 2 (c) Delta = 3 (d) Delta = 4

Figure 8: Example results for task *no smile*, selected to illustrate variations in model performance. Recall that the goal of the task is to transform an image of a smiling person into one where they are not smiling. For a small number of source images, the brain model produced outputs where the smile was only moderately reduced, as in 8a and 8b. Both brain and explicit models occasionally produced marginally lighter skin tones, as can be seen in 8c and 8d, but otherwise maintained most other features.

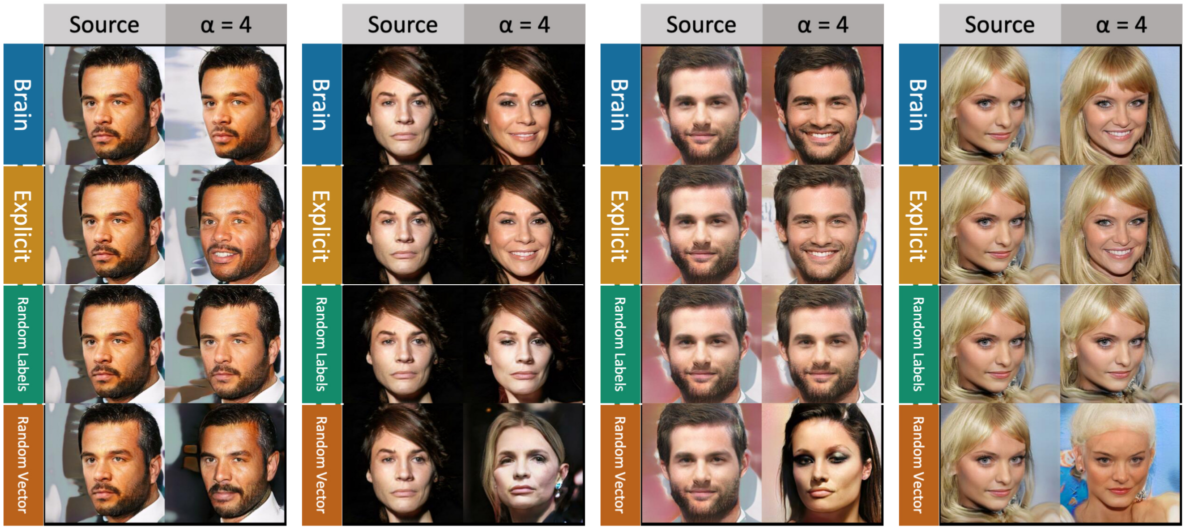

(a) Delta = 0 (b) Delta = 2 (c) Delta = 3 (d) Delta = 4

Figure 9: Example results for task *smile*, selected to illustrate variations in model performance. Recall that the goal of the task is to transform an image of a person who is not smiling into one where they are smiling. When a source image did not depict a stereotypically male or female person, both models tended to produce results that appeared more female, as in 9b. While the models did tend to preserve most other features, like hairstyle as in 9c, this was not always the case, as in 9d.

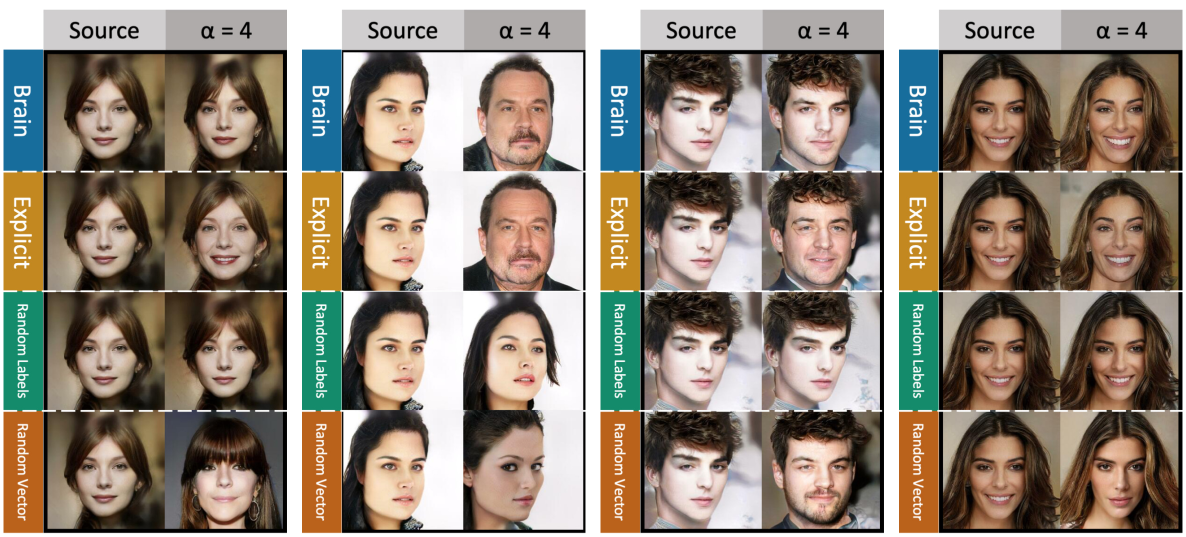

(a) Delta = 0 (b) Delta = 2 (c) Delta = 0 (d) Delta = 1

Figure 10: Example results for task *old*, selected to illustrate variations in model performance. Recall that the goal of the task is to transform an image of a person into one that looks much older. In 10a the brain model does not change the apparent age of the source image, while the explicit model does (although the changes are small). In 10b, both the brain and explicit models produce transformed images that not only appear older, but more male, than the source image. In 10c and 10d, while the assessed deltas are small, the modified images ultimately appear older than the source, revealing differences in the subjective perception of age.

5 Ethics Statement

As stated in the main text of our paper, all participants were fully informed of the nature of our study, their right to withdraw, and their right to have their data destroyed at any time. The experiment was

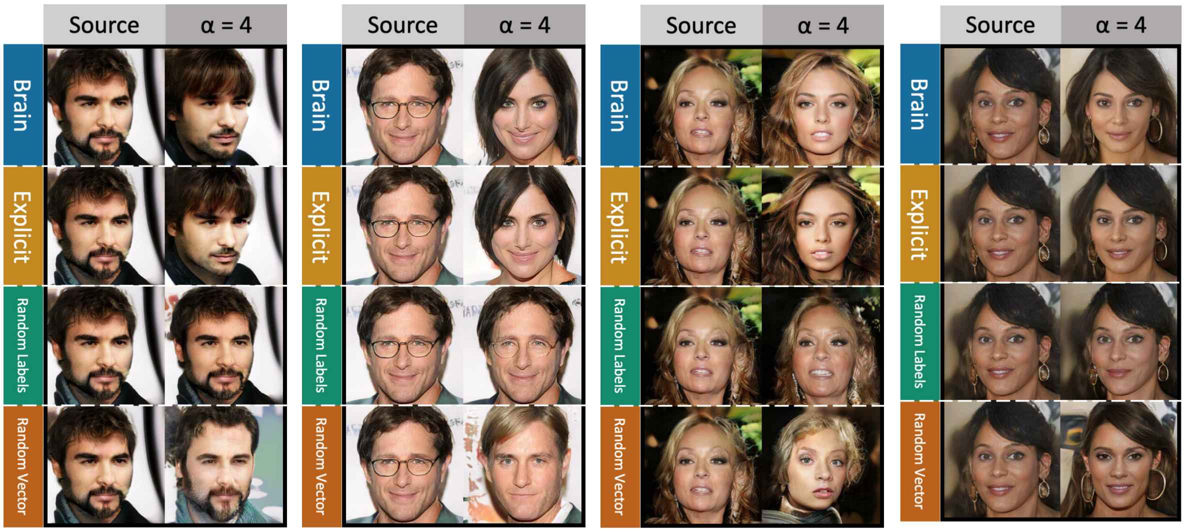

(a) Delta = 0 (b) Delta = 0 (c) Delta = 1 (d) Delta = 3

Figure 11: Example results for task *young*, selected to illustrate variations in model performance. Recall that the goal of the task is to transform an image of a person into one that looks much younger. In 11a, the changes produced by the brain and explicit models may be interpreted by some to create a slightly more youthful appearance (such as the change in hairstyle). However, they ultimately do not translate to a change in age as viewed by the assessors. In 11b, both brain and explicit models transform the male-appearing source into a female-appearing image of approximately the same age. In 11d and 11c, models achieve similar results, and the transformed images appear younger than the source.

approved by the ethical review board of the University of Helsinki.

While all information that could allow a participant to be easily identified was removed from the final dataset, the authors still exercise caution when handling and storing the data. EEG data can be used to detect and diagnose various neurophysiological diseases and conditions, and thus should be handled with the same care as medical information. Raw EEG data also contain signatures similar to a fingerprint, in that EEG data collected from an individual can potentially be linked to future EEG data collected.