Self-Calibrating BCIs: Ranking and Recovery of Mental Targets Without Labels

Jonathan Grizou*, **Carlos de la Torre-Ortiz***, Tuukka Ruotsalo (*Equal contribution)

Paper: https://arxiv.org/pdf/2506.11151
Code: https://jgrizou.github.io/neurips-self-calibrating-bci
Dataset: https://huggingface.co/datasets/ctorre/self-calibrating-bci

We recover the face image a participant holds in mind, using only unlabeled EEG responses

The self-calibration (SC) problem

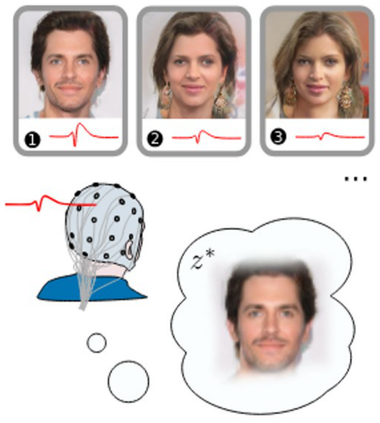

1) We give you unlabeled image and EEG pairs

2) You must recover the (face) image I had in mind

3) AND do it with:

- No labels
- No pre-trained EEG→image decoders
- No pre-recorded data
- No per-user calibration phase

"We test candidate images of faces (hypothetical targets); choosing the correct one yields the strongest EEG-distance predictability."

We propose the continuous SC-BCI framework and solve it

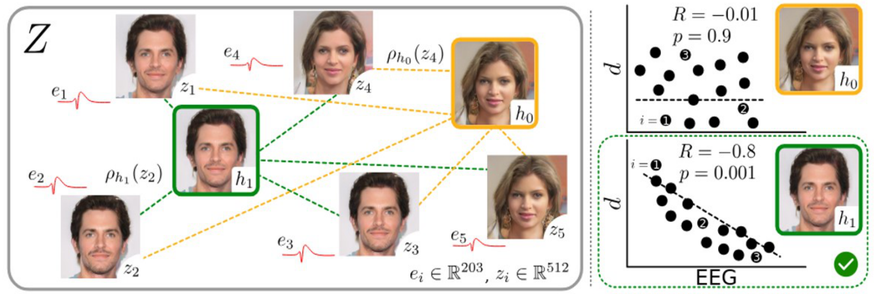

We contribute a new EEG dataset

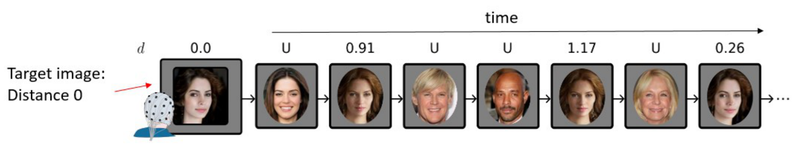

Limitations

- This study is restricted to face stimuli from one pre-trained GAN.
- We relied on a predefined similarity function in the stimuli latent space.
- The optimization was run in a lower-dimensional latent subspace.
- Real-time performance remains untested.

Ethical considerations are critical as inferring intent without explicit consent raises privacy concerns.

Our scores correlate with human judgement

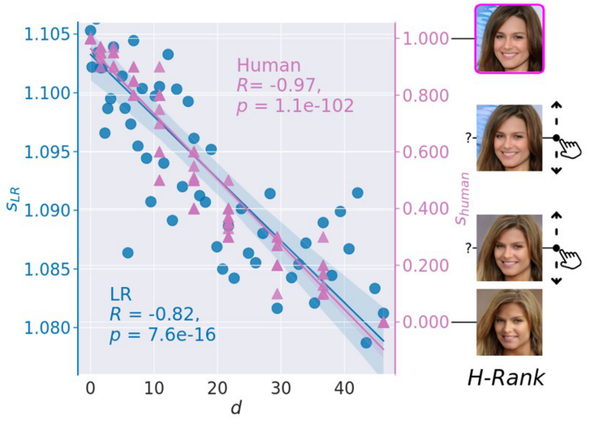

Top-1 retrieved faces are near the target

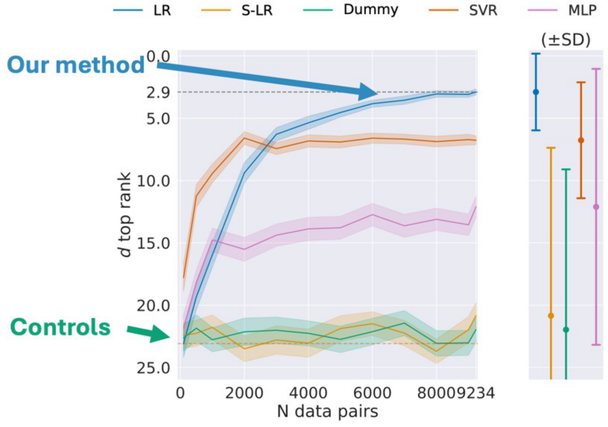

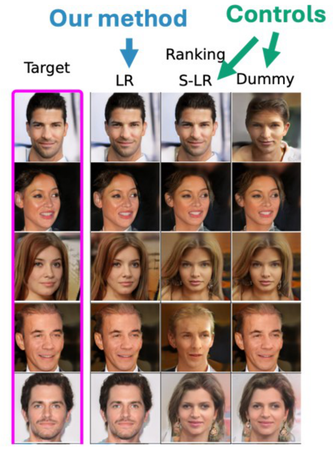

We can optimize and generate an *unknown* target