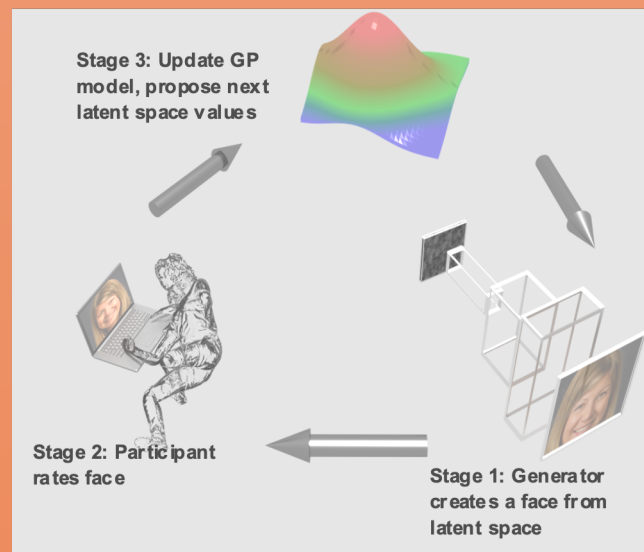


Bayesian Optimization for real-time, automatic design of face stimuli in human-centred research

F da Costa, P.^{1,2}, Lorenz, R.^{3,4}, Pio Monti, R.⁵, Jones, E.², Leech, R.¹

¹ Centre for Neuroimaging Sciences, Institute of Psychiatry, Psychology and Neuroscience, King's College London, UK, ² Centre for Brain and Cognitive Development, Birkbeck, University of London, UK, ³ MRC Cognition and Brain Sciences Unit, University of Cambridge, UK, ⁴ Department of Psychology, Stanford University, USA, ⁵ Gatsby Computational Neuroscience, University College London, UK

We built a framework that optimizes face stimuli according to any measured target, by leveraging the flexibility of GANs and the efficiency of Bayesian optimization.



A. Framework Overview.

You can try our proof of concept:




Introduction

Most research in neuroscience and in psychology overlooks inter-individual differences in face processing by relying on pre-selected stimuli. To address this shortcoming we built a closed-loop system that tailors face stimuli to maximise an individual subject's response (e.g., neural, behavioural or subjective) in a short number of iterations. The stimulus is generated by the *StyleGAN* [1] and the space of faces is obtained by manipulating the latent space according to previously learned semantic directions.

Proof of Concept

We conducted a study of 30 participants, in which they had to rate how much a manipulated photo of themselves looked like them. The space of faces was composed of two directions: emotion and age. Results showed a higher intra-subject spatial correlation than the mean inter-subject correlation suggesting inter-individual response patterns.



B. Participants performance along the space of faces.