Human-in-the-Loop Optimization for Deep Stimulus Encoding in Visual Prostheses

Jacob Granley[1], Tristan Fauvel[2], Matthew Chalk[2], Michael Beyeler[1]
[1]University of California, Santa Barbara
[2]Institut de la Vision, Sorbonne Université

BIONIC VISION LAB

Overview

Goal:
Black-box optimization of stimulus parameters for prosthesis patients

Challenges:
- Limited, noisy patient feedback
- High-dimensional search space (100-1000's params)
- Large patient-to-patient variations in perception

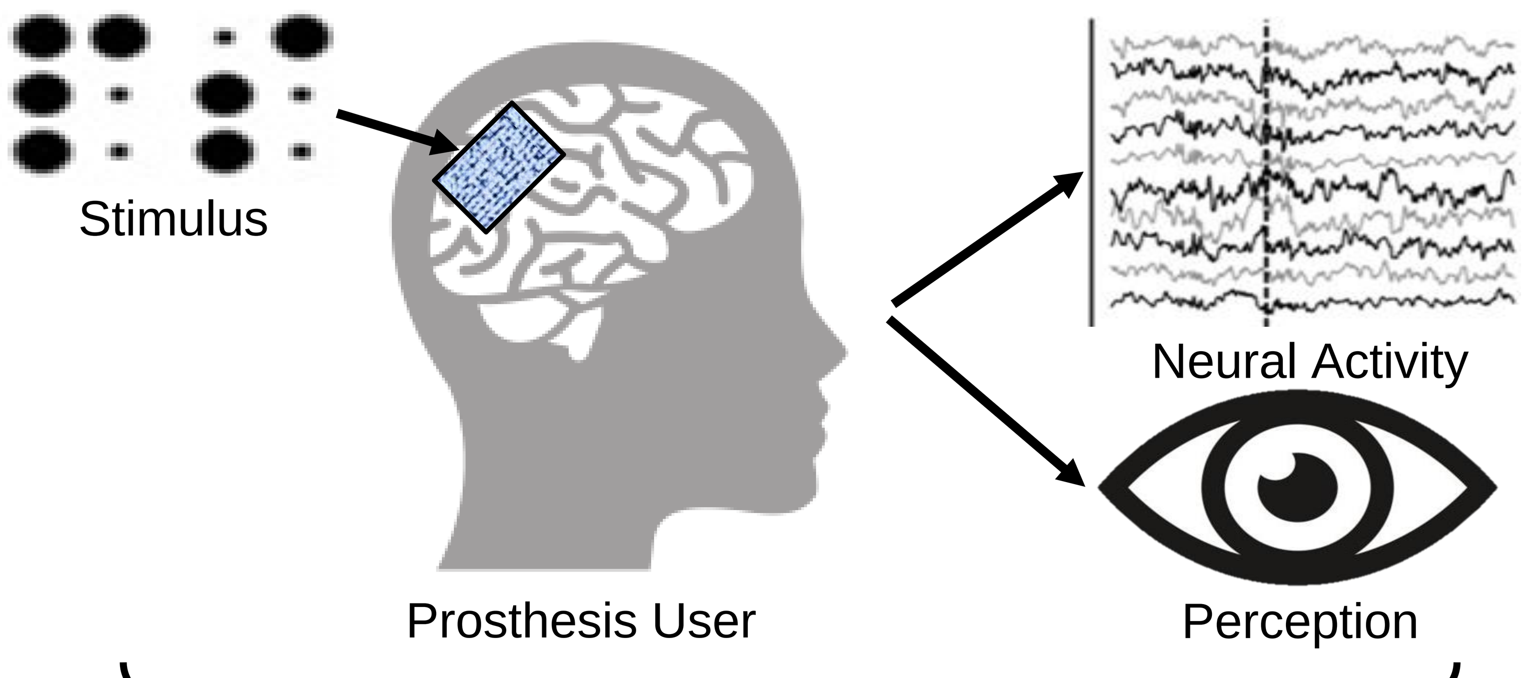

Forward Model
- Predicts visual perception (phosphenes) from stimulus
- Captures individual variations with patient-specific vector φ
- SOTA model

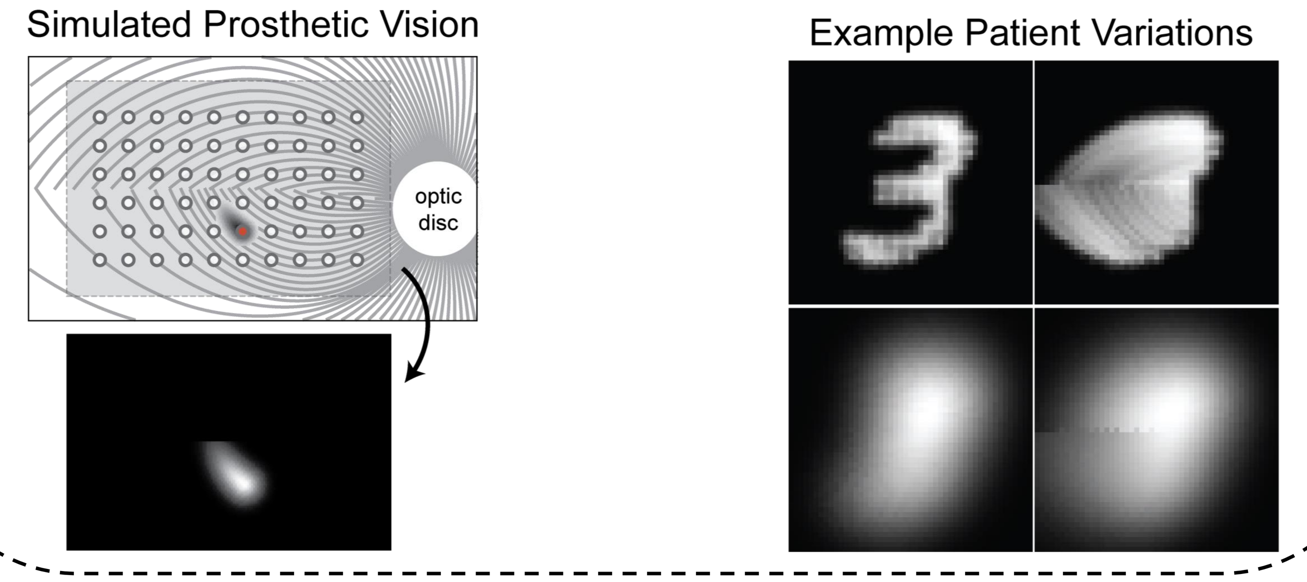

Proposed Solution:
Use Bayesian optimization (BO) to tune a deep neural network stimulus encoder (DSE) based on patient feedback

End-to-end framework:
1. Create a forward model of sensory system
2. Backpropagate through model to optimize stimuli, dependent on latent patient-specific parameters φ
3. Use Bayesian optimization to learn optimal φ for new patients

Human-in-the-Loop Optimization (HILO)

Pretrain deep stimulus encoder (DSE) across patients

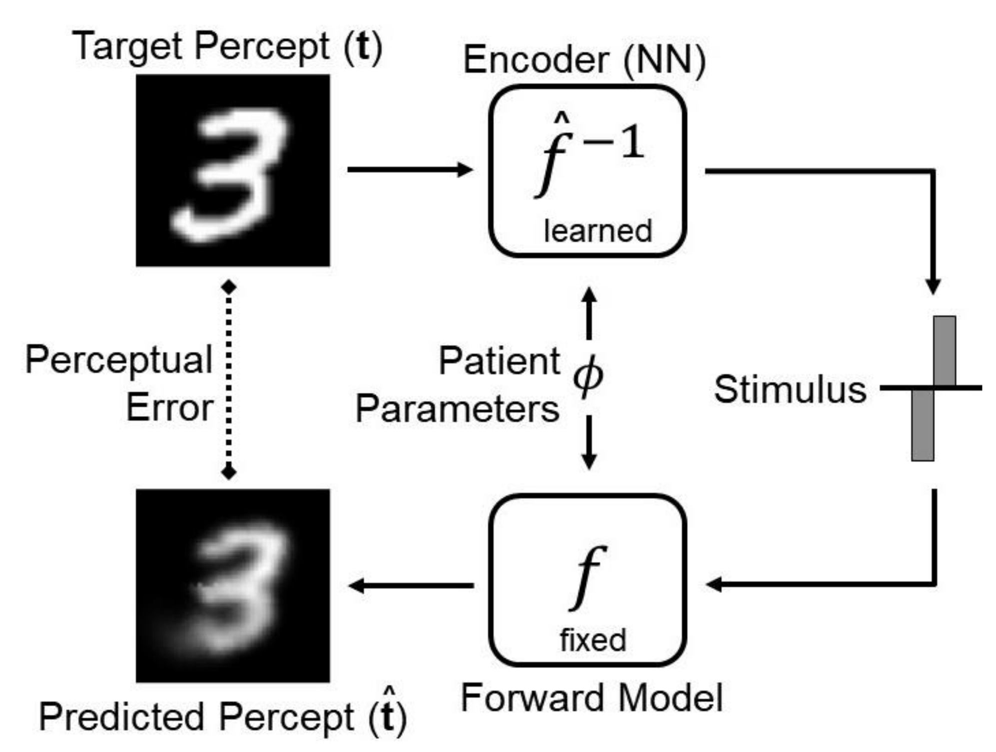

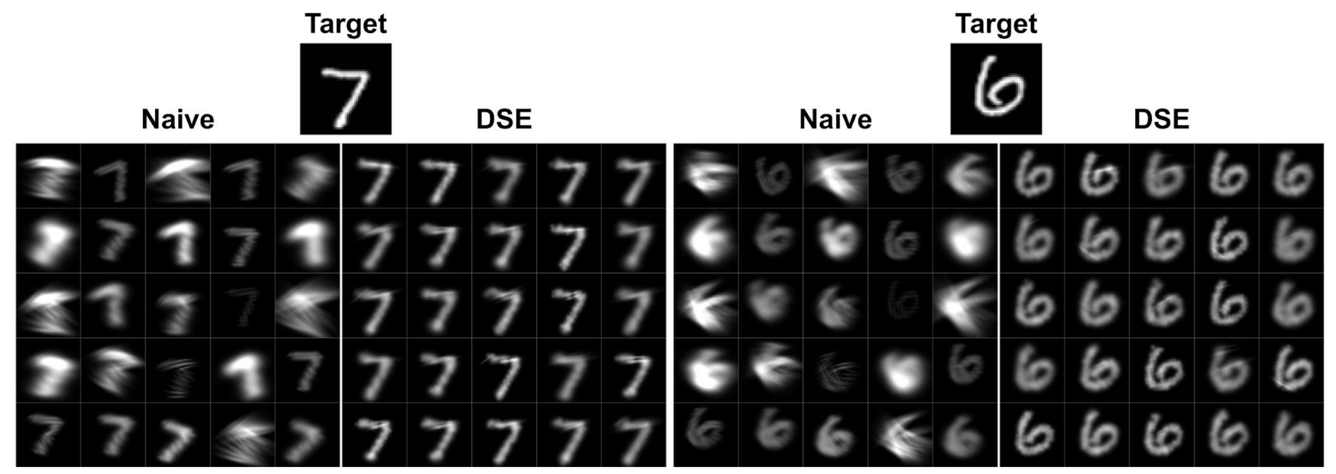

Use preferential Bayesian Optimization to learn φ for a new patient:

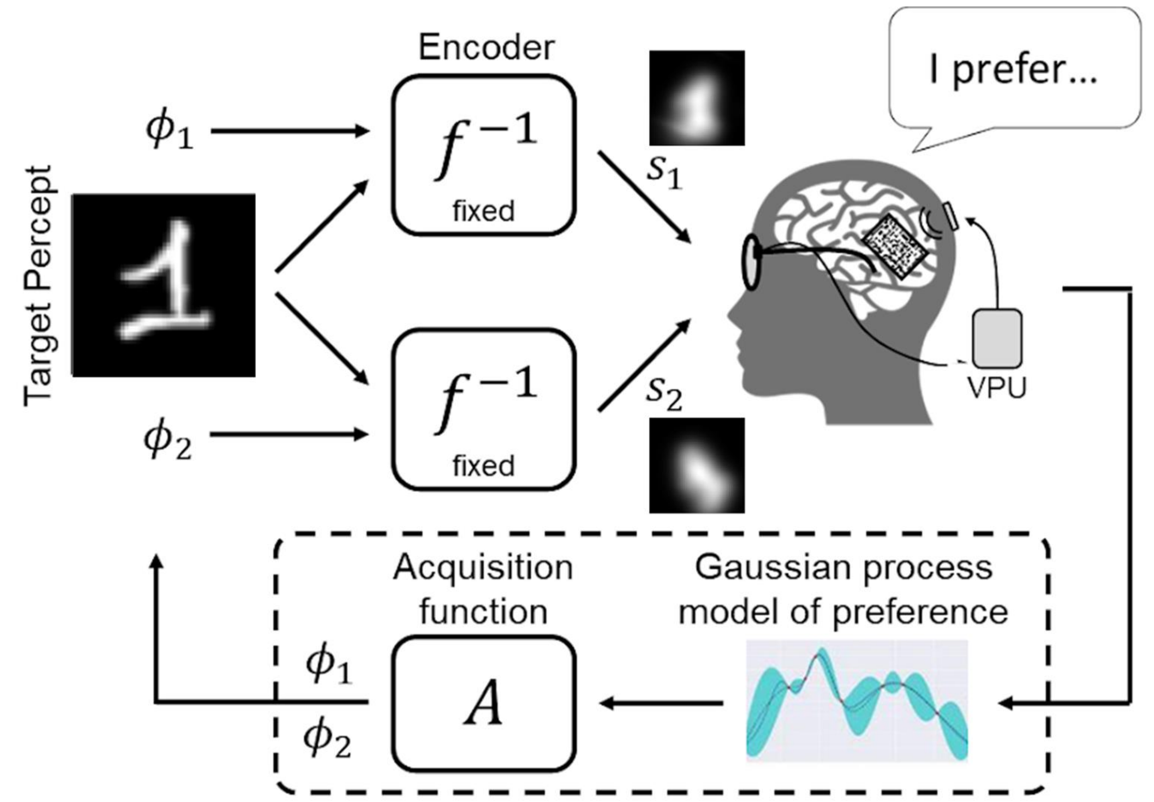

Gaussian process model:

$$P(\phi_1 \succ \phi_2 | g) = \Phi\big(g(\phi_1) - g(\phi_2)\big)$$

Acquisition function:

$$\phi_1 \to \arg\max_{\phi} \mathop{\mathbb{E}}_{p(g|\mathcal{D})}[g(\phi)]$$

$$\phi_2 \to \arg\max_{\phi} \mathop{\mathbb{V}}_{p(g|\mathcal{D})}[\Phi(g(\phi) - g(\phi_1))]$$

Robustness

Tested model variants where:
- the patient's decision was increasingly random
- the phosphene model was misspecified

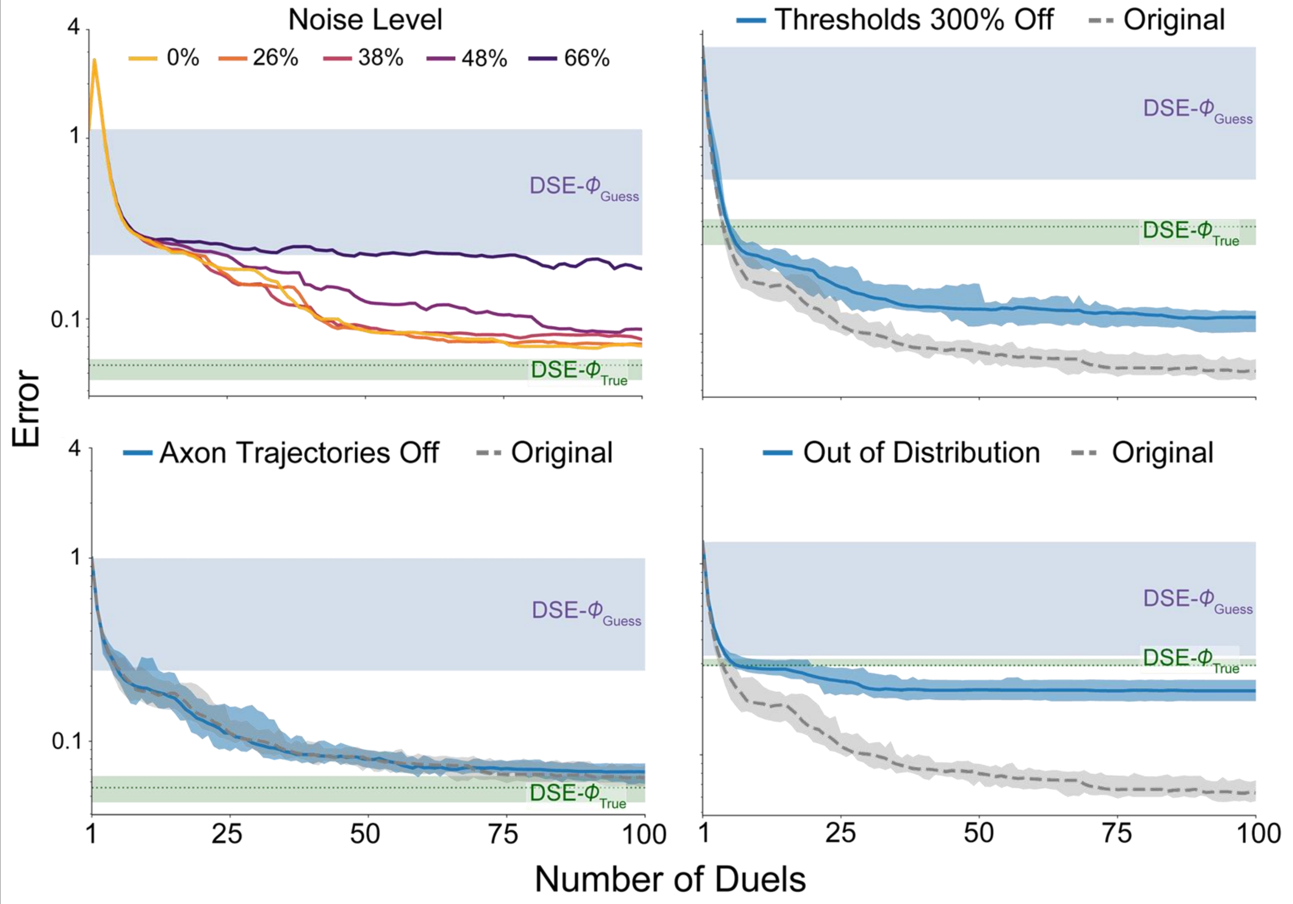

Optimizing based on user preferences is very robust!
- works even when user's decision is random 2/3 times
- works even when electrode thresholds are off by 300%
- works even when the phosphene model is incorrect

Results on Simulated Patients

Simulated patients:
- Randomly assigned φ
- Simulate percept according to forward model
- Choose duel outcome probabilistically, based on distance from each percept to the target

Baselines:
- Naïve encoder (used by current devices),
- DSE with guess for φ

Across 100 simulated patients:
- Better encoding after **~20 duels**
- **9x** reduction in reconstruction error
- **99%** of patients improved from HILO

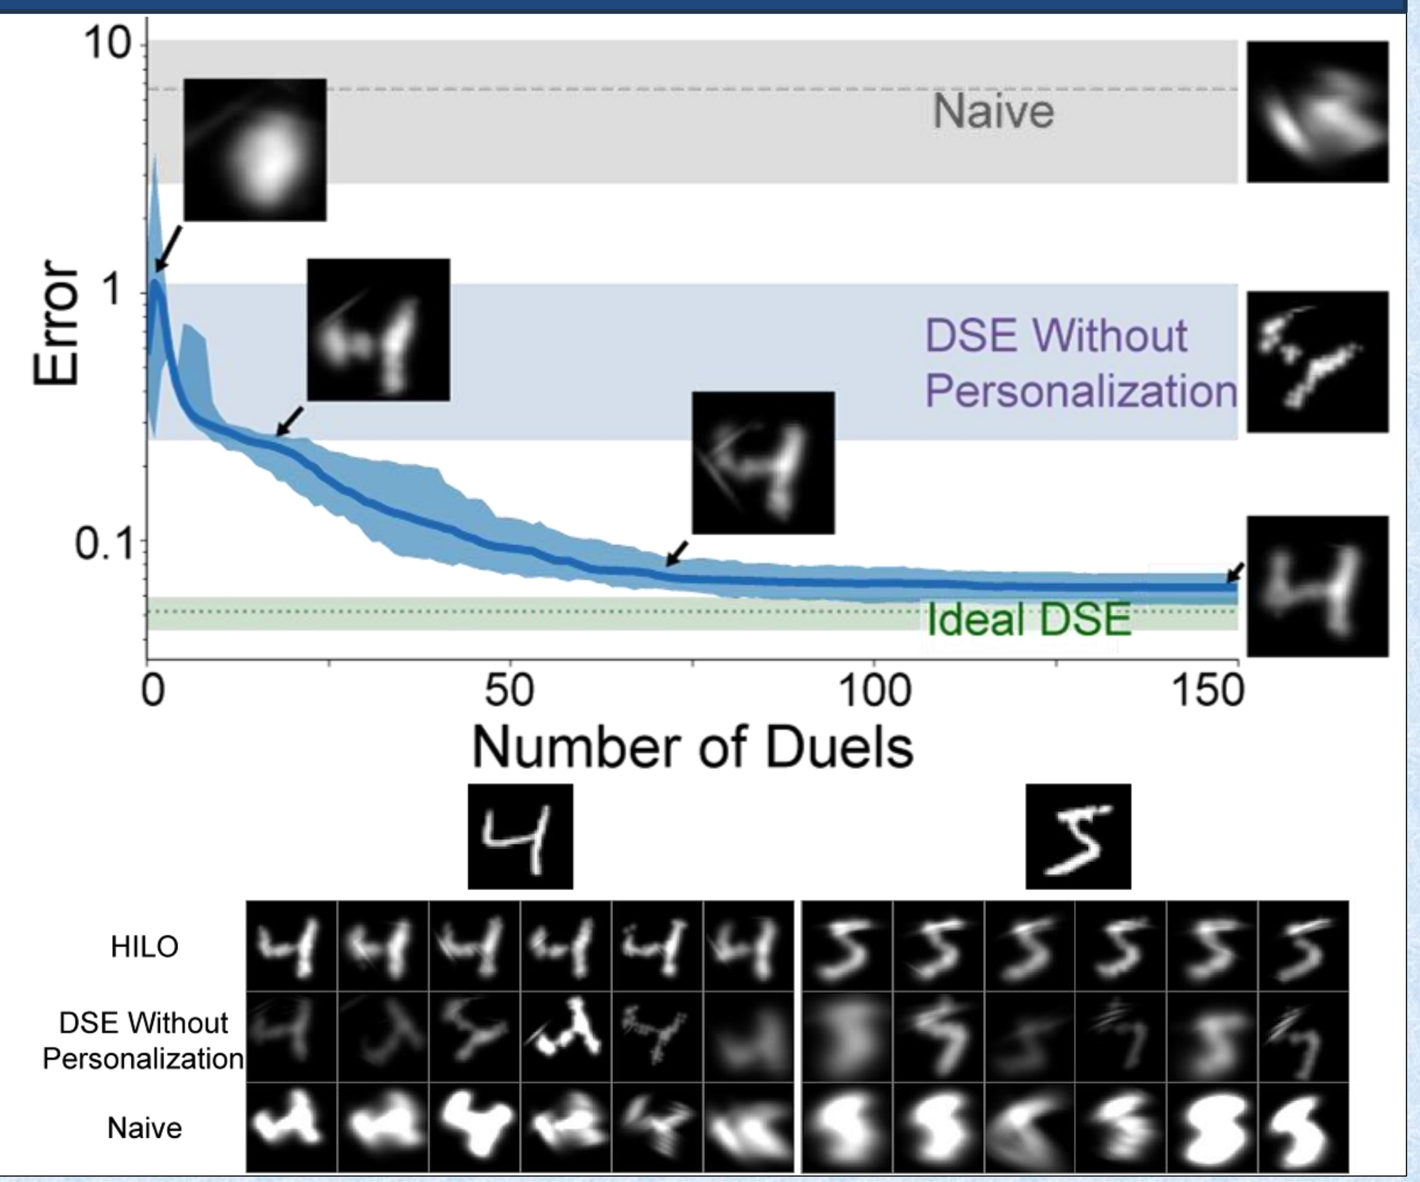

Conclusions

Bayesian optimization and deep learning are complimentary.

Bayesian optimization:
- Adapts to limited, noisy human feedback
- Robust against erroneous modeling assumptions

Deep learning:
- Allows for high dimensional, complex optimization

Could be deployed with commercial prostheses to periodically recalibrate without requiring a professional

Future work will test HILO on sighted and blind subjects

Broader Impacts:
- Forward models and DSEs have been successfully used across multiple sensory modalities and could potentially be adapted for personalization with HILO.
- The latent space of a DNN is a good target for integrating Bayesian optimization into a neural network.