Real-time reconstruction of human visual perception from fMRI

Rishab S. Iyer, Jiaxin Cindy Tu†, Anish Mahishi†, Cesar Kadir Torrico Villanueva†, Ross P. Kempner, Jacob S. Prince, Hritik Arasu, Amaar Chughtai, Elizabeth A. McDevitt*, Paul S. Scotti*, Kenneth A. Norman*

Workshop on Foundation Models for the Brain and Body @ NeurIPS 2025

Supported by: Princeton University ODFR Innovation Fund for New Industrial Collaborations, Sophont, and NIMH Grant RF1MH125318

Bridging Foundation Models to Closed-loop Applications

Foundation/decoding models in fMRI are enabling rich latents & fine-grained decoding

- NeuroMamba[1]: 11.9M parameters
- BrainJEPA[2]: 307M parameters
- BrainLM[3]: 650M parameters
- Dynadiff[4]: 400M parameters
- MindEye2[5]: 722M parameters

Closed-loop neurofeedback can facilitate:

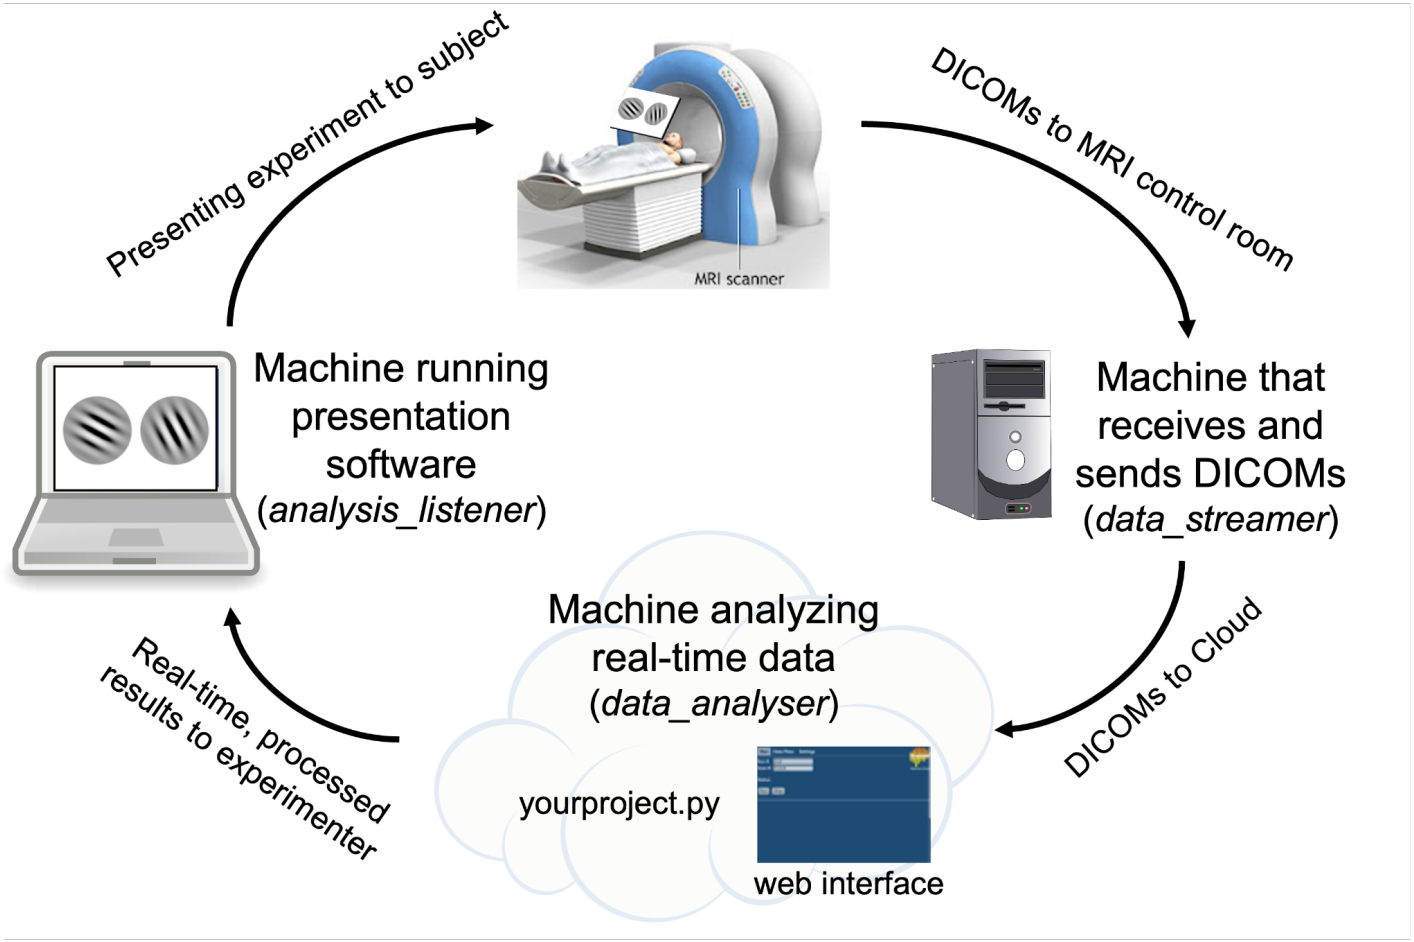
- Causal manipulations of neural activity[6,7]
- Clinical treatment via non-invasive neuromodulation[8,9]

However, current real-time pipelines are typically limited to computationally light analyses

- Univariate activity/functional connectivity
- Simple multivoxel pattern classifiers

[1]Choi et al., 2025, [2]Dong et al., 2024, [3]Caro et al., 2024, [4]Careil et al., 2025, [5]Scotti et al., 2024, [6]Peng et al., 2024, [7]Iordan et al., 2024, [8]Mennen et al., 2021, [9]Khaleghi et al., 2025

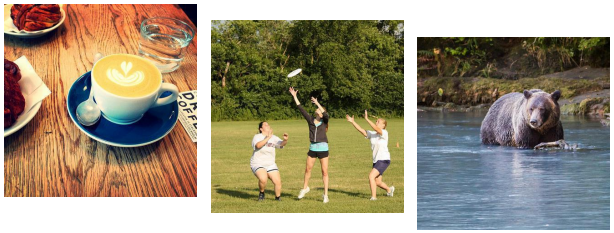
RT-Cloud



[10]Wallace et al., 2022

Training:
Pre-train on multiple subjects (NSD: 7T fMRI)
Fine-tune on 1 hour of data from new subject

Inference:



Preprocess + deconvolve HRF using GLM
-> single-trial response estimates

MindEye2[5] inference

Predicted CLIP[12] embeddings

Retrieval: pick the top-k nearest neighbor image CLIP embeddings

Reconstruction: generate an image based on the predicted CLIP embedding

Image retrievals/reconstructions

1. Baseline: offline 7T
Fine-tune on held-out subject (NSD; 7T fMRI)

2. Generalize to new dataset: offline 3T
Fine-tune on subject from novel dataset (3T fMRI)

3. Fastest real-time
Fit GLM ~7.5s post stimulus onset (after HRF peak)

4. Slower real-time
Fit GLM after waiting for 5 more trials; ~27s delay

[5]Scotti et al., 2024, [11]Radford & Kim et al., 2021

Evaluations

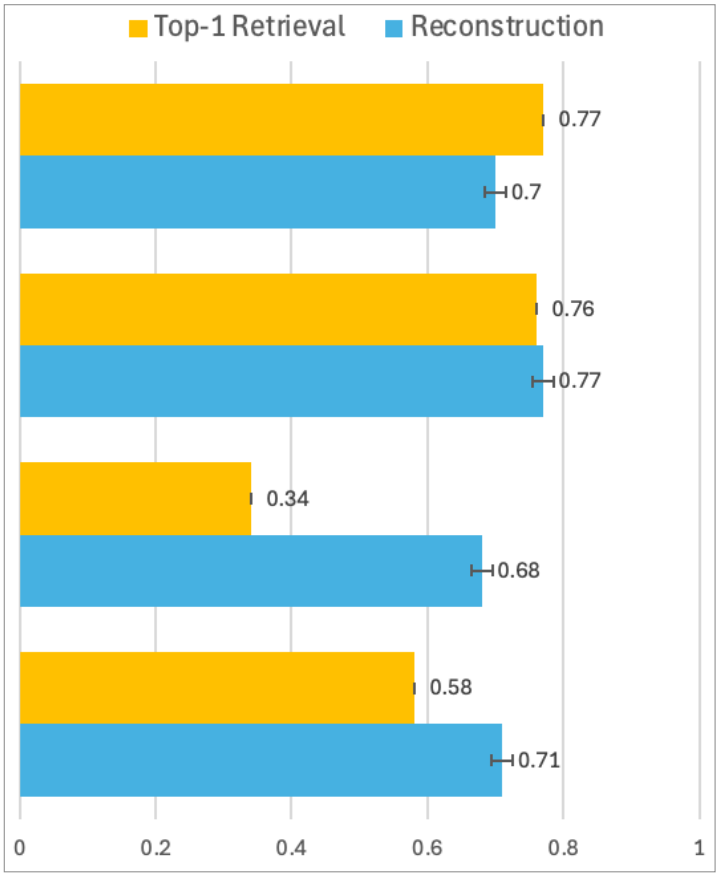
Single-trial decoding

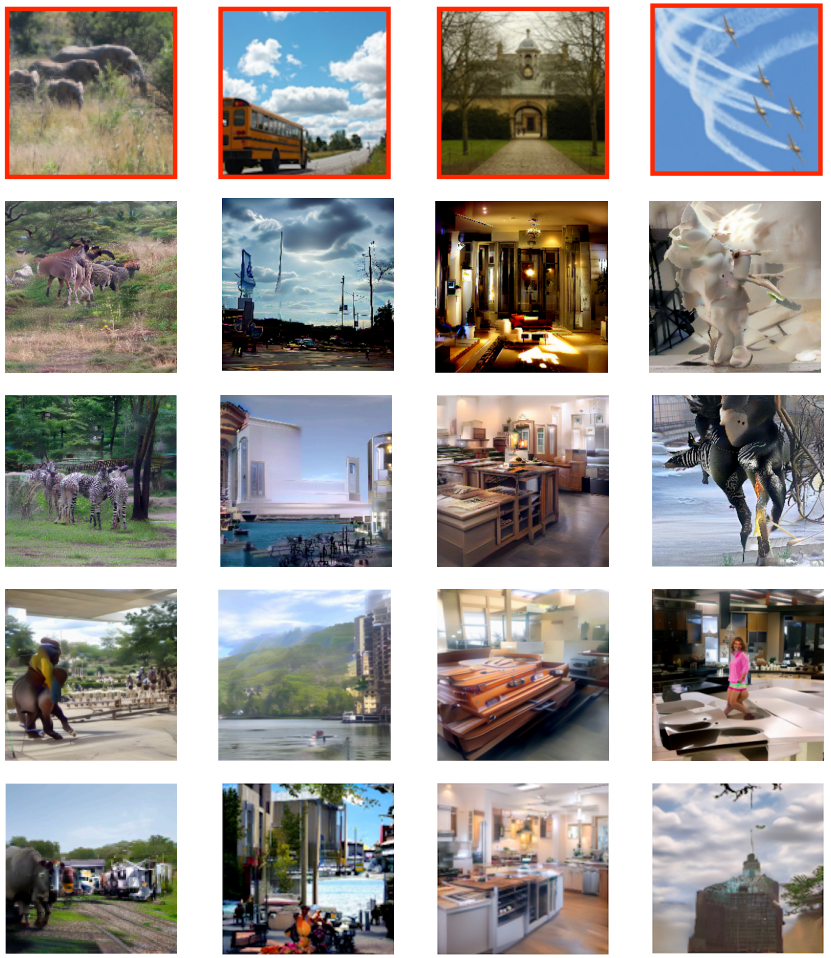
1. Baseline: offline 7T
Fine-tune on held-out subject (NSD: 7T fMRI)

2. Generalize to new dataset: offline 3T
Fine-tune on subject from novel dataset (3T fMRI)

3. Fastest real-time
Fit GLM ~7.5s post stimulus onset (after HRF peak)

4. Slower real-time
Fit GLM after waiting for 5 more trials (~27s delay)





Join us!

- RT-Cloud is fully open-source, enabling complex real-time fMRI workflows on any scanner, no specialized infrastructure needed
- We are building publicly + open-source on Discord (MedARC)
  - Open weekly meetings for anyone to join
  - Most co-authors are volunteers from Discord
- QR code: Notion homepage for real-time MindEye project, contact info, links to MedARC Discord & GitHub, these slides
- Open to collaborations and suggestions (rsiyer@princeton.edu)

