BrainBits: How Much of the Brain are Generative Reconstruction Methods Using?

David Mayo*, Christopher Wang*, Asa Harbin*,
Abdul Alabdulkareem, Albert Shaw, Boris Katz, Andrei Barbu

Diffusion models enable surprisingly good reconstruction from the brain!

Stimuli shown to fMRI subject

Reconstruction

Ozcelik, Furkan, and Rufin VanRullen. "Natural scene reconstruction from fMRI signals using generative latent diffusion." *Scientific Reports* 13.1 (2023): 15666.

Background: Image reconstruction from the brain

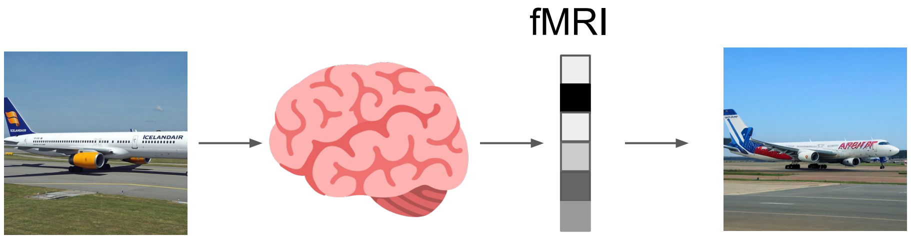

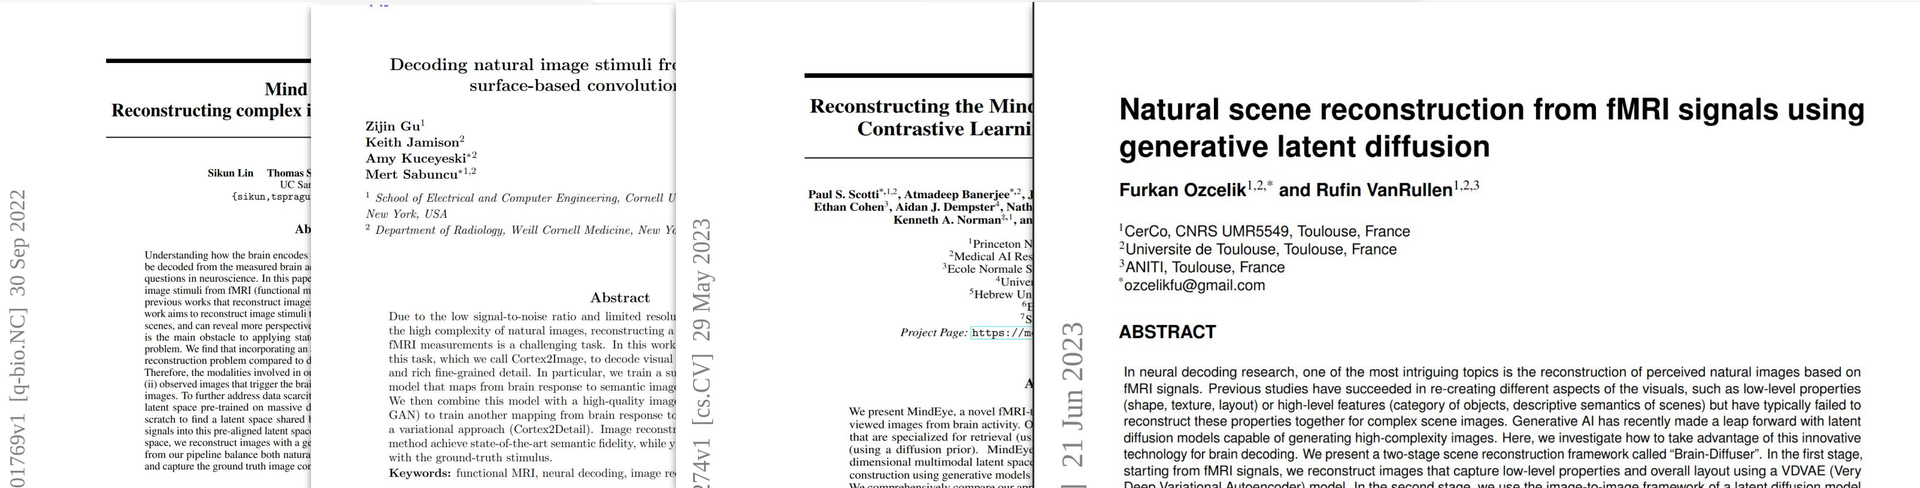

01769v1 [q-bio.NC] 30 Sep 2022

Mind
Reconstructing complex i

Sikun Lin Thomas S
UC San
{sikun,tspragu

Ab

Understanding how the brain encodes be decoded from the measured brain a questions in neuroscience. In this pape image stimuli from fMRI (functional m previous works that reconstruct image work aims to reconstruct image stimuli t scenes, and can reveal more perspective is the main obstacle to applying stat problem. We find that incorporating an reconstruction problem compared to d Therefore, the modalities involved in ou (ii) observed images that trigger the brai images. To further address data scarcit latent space pre-trained on massive d scratch to find a latent space shared l signals into this pre-aligned latent spac space, we reconstruct images with a ge from our pipeline balance both natura and capture the ground truth image co

Decoding natural image stimuli fr
surface-based convolution

Zijin Gu[1]
Keith Jamison[2]
Amy Kuceyeski[*2]
Mert Sabuncu[*1,2]

[1] *School of Electrical and Computer Engineering, Cornell U*
New York, USA
[2] *Department of Radiology, Weill Cornell Medicine, New Yo*

Abstract

Due to the low signal-to-noise ratio and limited resolu the high complexity of natural images, reconstructing a fMRI measurements is a challenging task. In this work this task, which we call Cortex2Image, to decode visual and rich fine-grained detail. In particular, we train a su model that maps from brain response to semantic image We then combine this model with a high-quality image GAN) to train another mapping from brain response to a variational approach (Cortex2Detail). Image reconstr method achieve state-of-the-art semantic fidelity, while y with the ground-truth stimulus.

Keywords: functional MRI, neural decoding, image re

274v1 [cs.CV] 29 May 2023

Reconstructing the Min
Contrastive Learni

Paul S. Scotti[*,1,2], Atmadeep Banerjee[*,2], J
Ethan Cohen[3], Aidan J. Dempster[4], Nath
Kenneth A. Norman[‡,1], an

[1]Princeton N
[2]Medical AI Res
[3]Ecole Normale S
[4]Univer
[5]Hebrew Un

Project Page: https://m

A

We present MindEye, a novel fMRI-t viewed images from brain activity. O that are specialized for retrieval (us (using a diffusion prior). MindEye dimensional multimodal latent spac construction using generative models

21 Jun 2023

Natural scene reconstruction from fMRI signals using generative latent diffusion

Furkan Ozcelik[1,2,*] **and Rufin VanRullen**[1,2,3]

[1]CerCo, CNRS UMR5549, Toulouse, France
[2]Universite de Toulouse, Toulouse, France
[3]ANITI, Toulouse, France
[*]ozcelikfu@gmail.com

ABSTRACT

In neural decoding research, one of the most intriguing topics is the reconstruction of perceived natural images based on fMRI signals. Previous studies have succeeded in re-creating different aspects of the visuals, such as low-level properties (shape, texture, layout) or high-level features (category of objects, descriptive semantics of scenes) but have typically failed to reconstruct these properties together for complex scene images. Generative AI has recently made a leap forward with latent diffusion models capable of generating high-complexity images. Here, we investigate how to take advantage of this innovative technology for brain decoding. We present a two-stage scene reconstruction framework called "Brain-Diffuser". In the first stage, starting from fMRI signals, we reconstruct images that capture low-level properties and overall layout using a VDVAE (Very

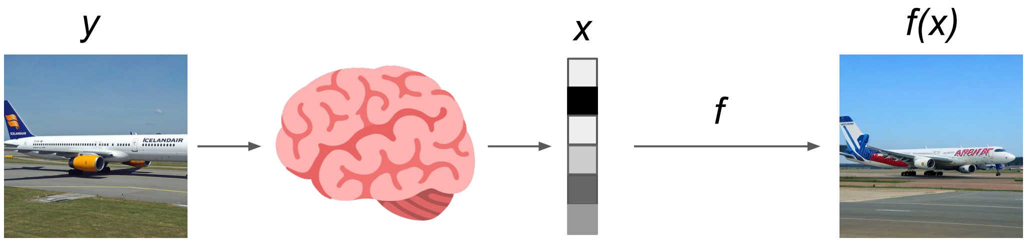

According to image similarity metrics, the images keep getting better and better…

But are we really getting better at the brain or just getting better at image generation?

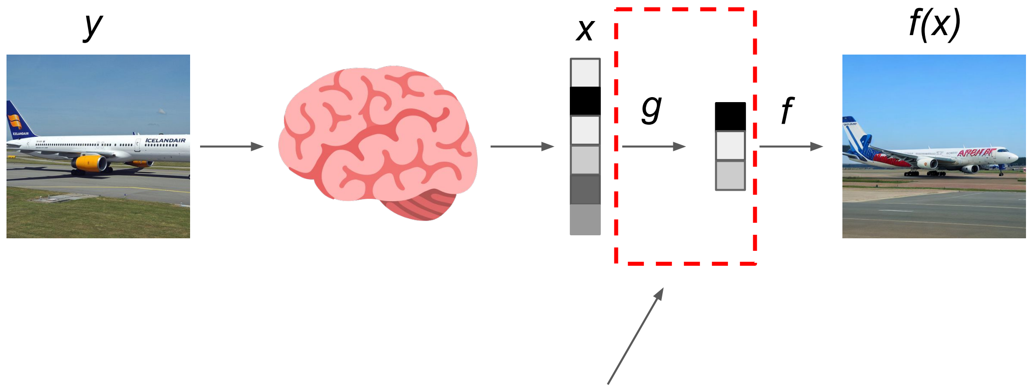

BrainBits: Insert an information bottleneck in the reconstruction pipeline.

How much information is needed to get the reconstructions we're seeing?

This is a general method that could be applied to any stimuli reconstruction task from neural data.

Finding 1: We can restrict the information available at reconstruction time to a great degree and still get “good looking” images

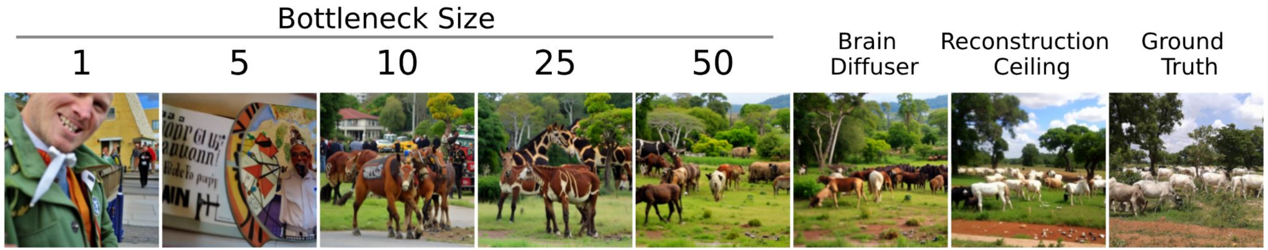

Finding 1: We can restrict the information available at reconstruction time to a great degree and still get “good looking” images

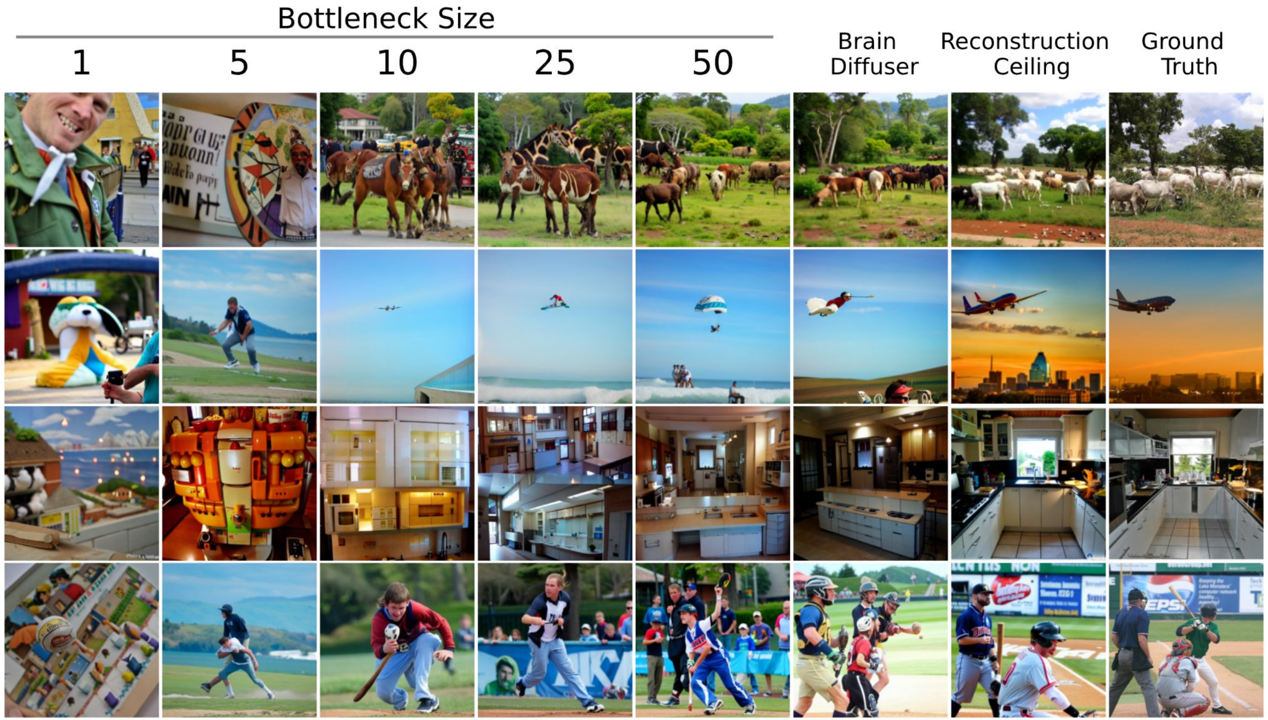

How good?

Finding 2: According to existing metrics, reconstruction performance quickly plateaus

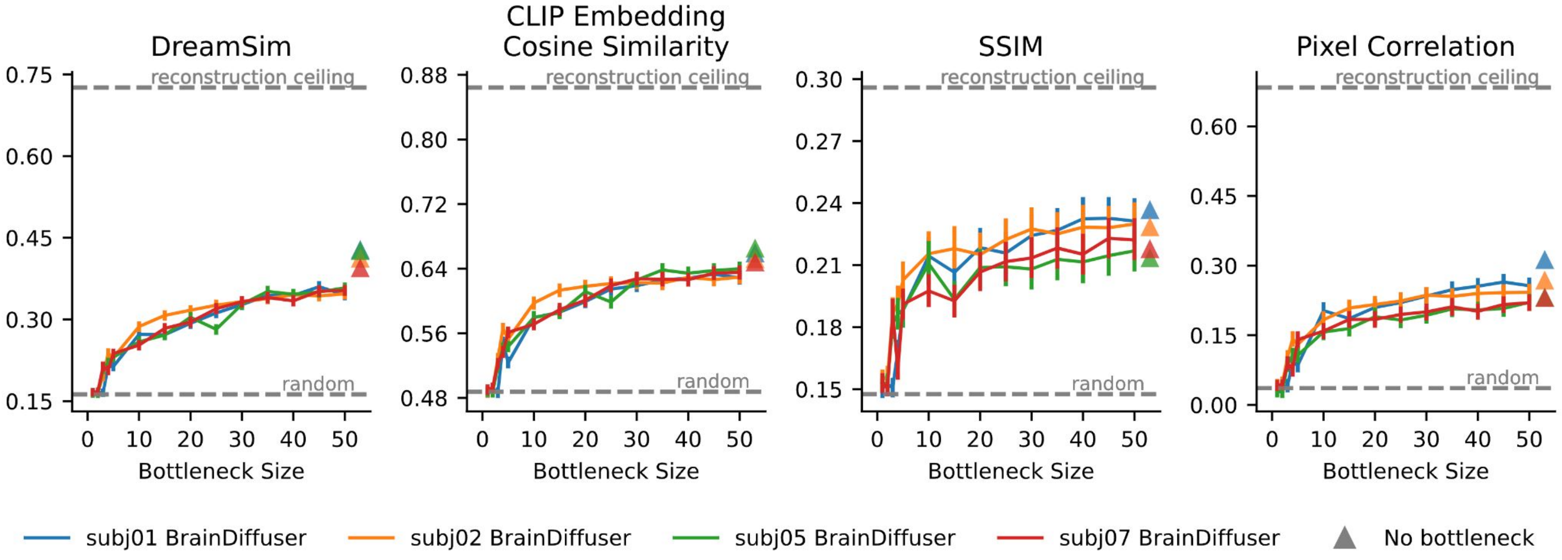

What does this mean?

The message: if we care about studying visual processing in the brain, we need to measure signal extracted, not image quality

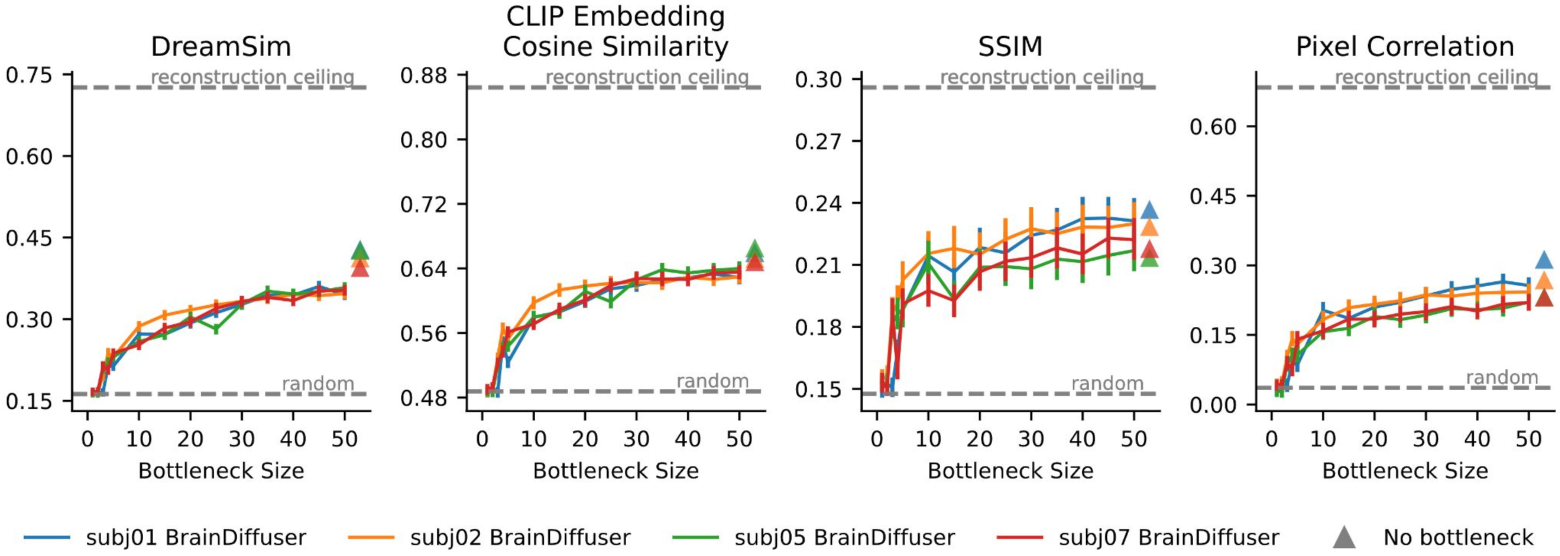

What information do bottlenecks contain?

Edge energy, brightness and contrast are mostly exhausted early. Larger bottlenecks are needed to extract more object class information above chance.

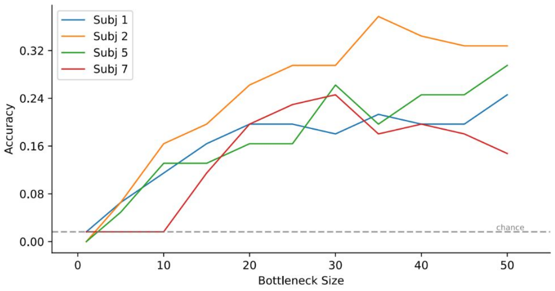

(a) Object class

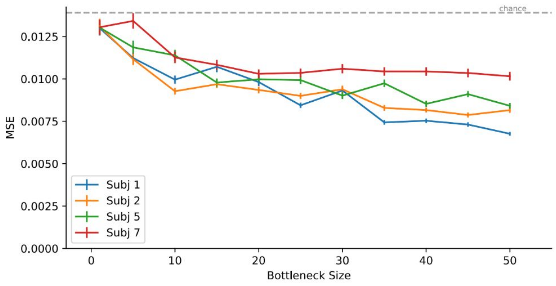

(b) Brightness

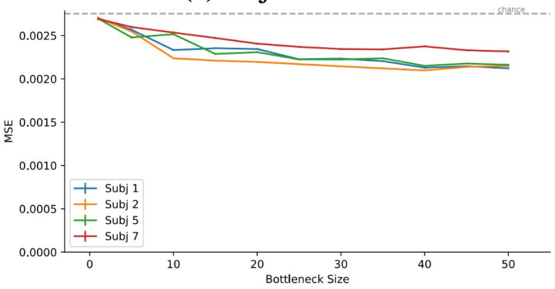

(c) RMS Contrast

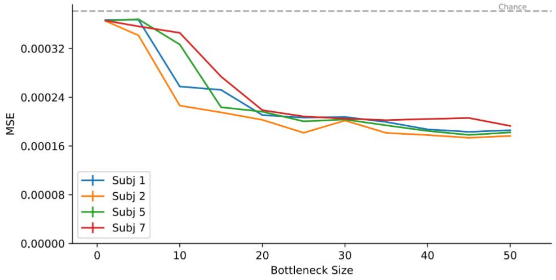

(d) Average gradient magnitude

What areas of the brain help reconstruction the most?

Models quickly zoom in on useful areas even at low bottleneck sizes.

As the bottleneck size goes up models exploit those original areas but do not meaningfully expand to new area

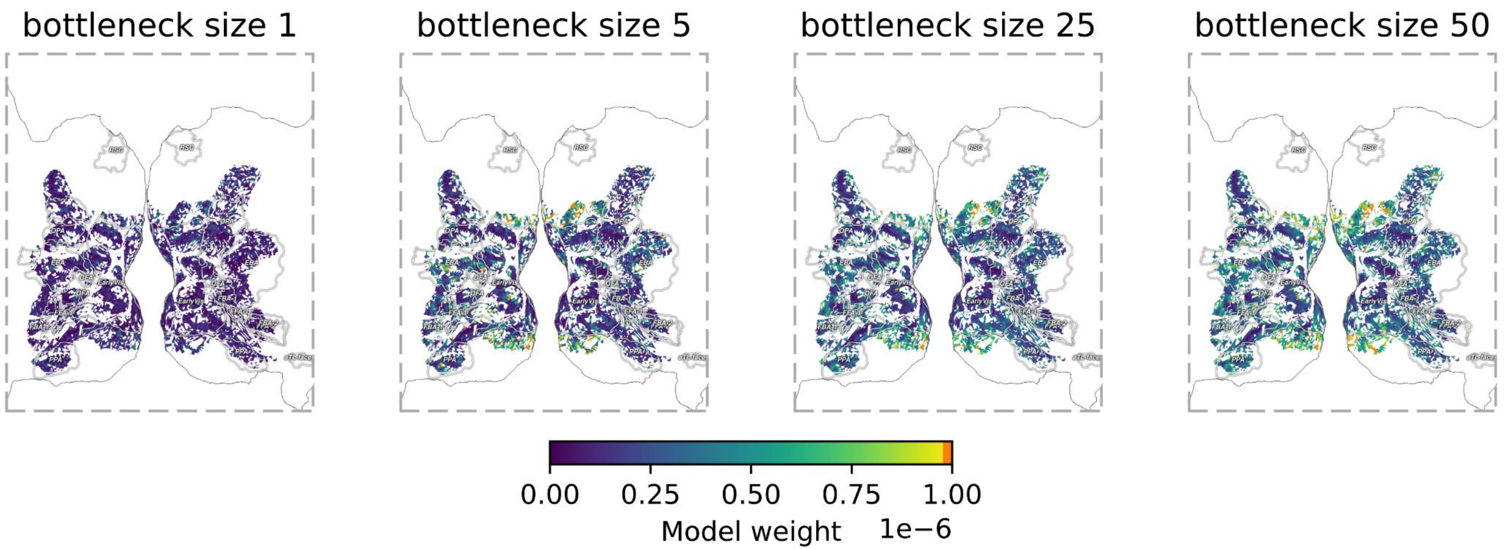

Check out our poster at NeurIPS 2024!

Paper: https://arxiv.org/abs/2411.02783
Github: https://github.com/czlwang/BrainBits