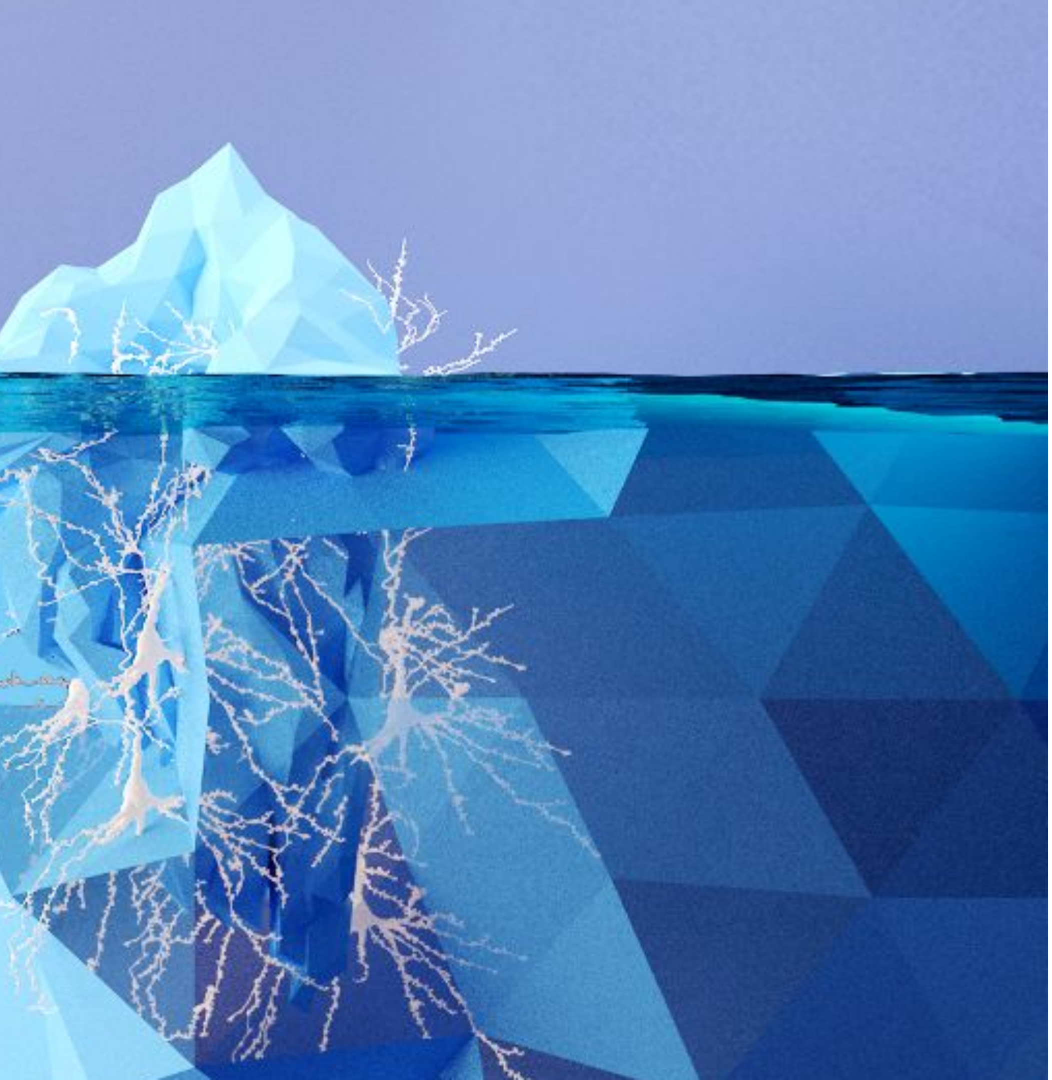

Section: Walker et al. 2019

Background Recap:

Section W10

What are neurons in V1 encoding?

Orientation selectivity (Hubel & Wiesel)

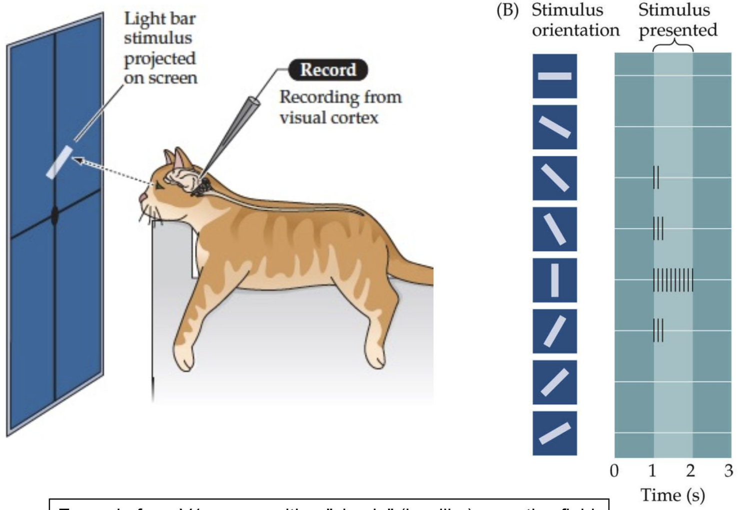

Example for a V1 neuron with a "simple" (bar-like) receptive field

Purves Fig. 12.8

Spike triggered average

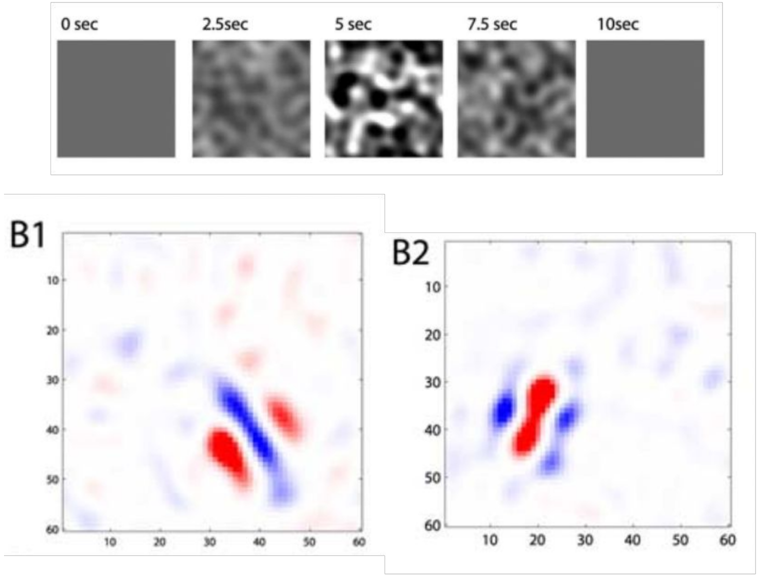

Niell and Stryker, 2008

V1 RFs resemble Gabor filters and neural response is sparse

Gabor filters

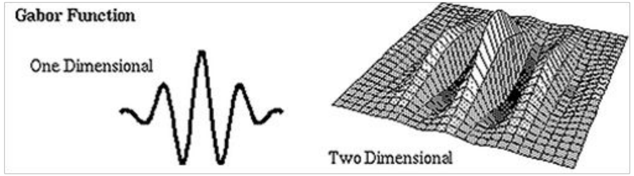

edge detector neurons can be explained with sparse autoencoding

$$I(x,y) = \sum_i a_i\, \phi_i(x,y) + \epsilon(x,y)$$

Adapted from A. Mathis

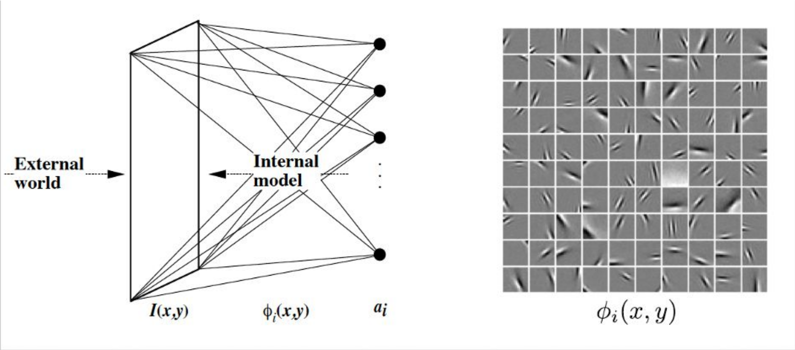

Olshausen & Field, 1996 Nature

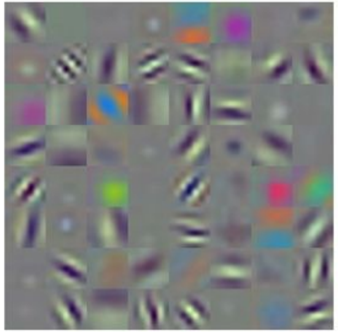

Representations in ImageNet trained (CNN)

Hierarchical visual processing

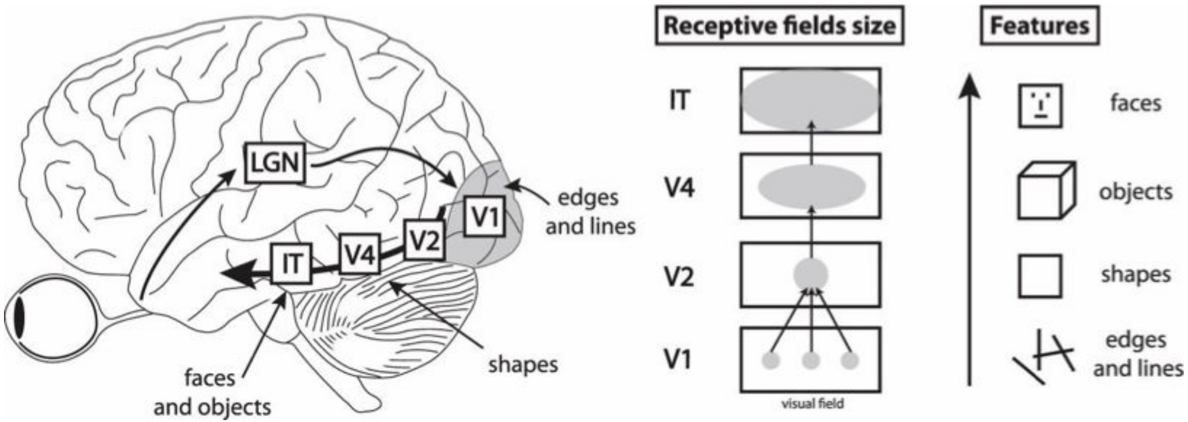

From: When crowding of crowding leads to uncrowding
Journal of Vision. 2013;13(13):10. doi:10.1167/13.13.10

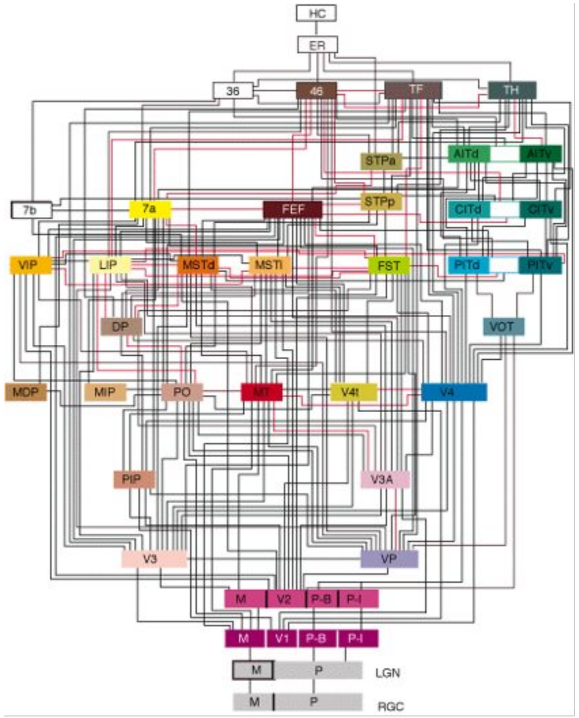

Felleman and Van Essen *Cerebral Cortex* 1991

Hierarchical visual processing in the mouse

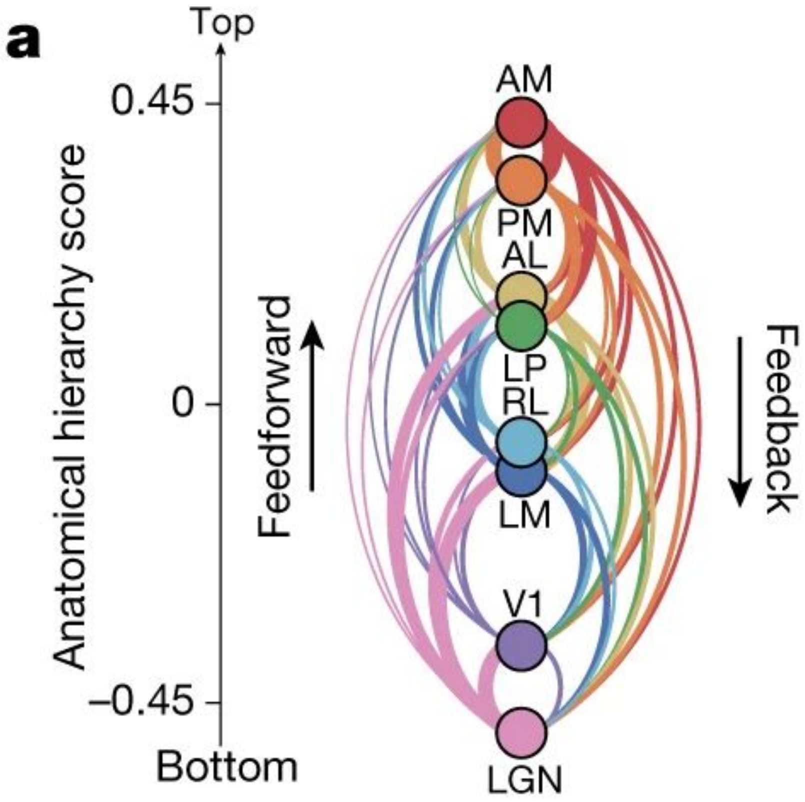

Siegle et al. 2021

Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex

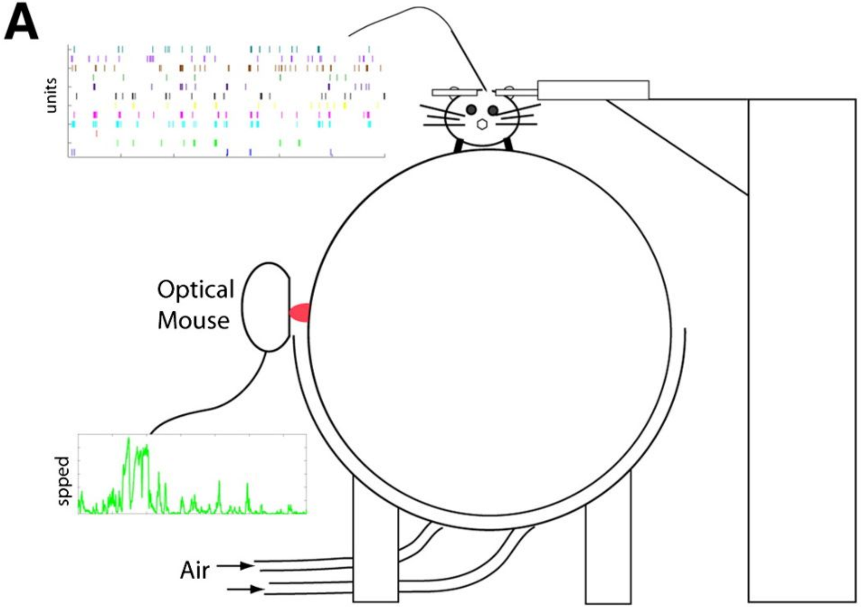

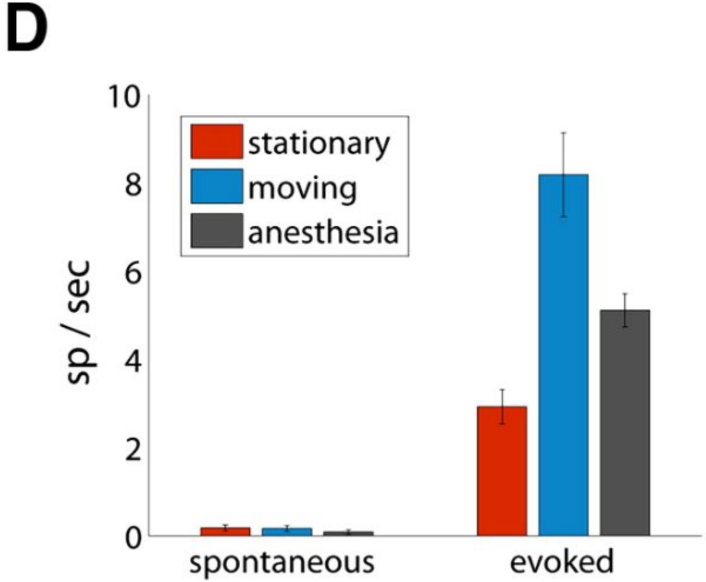

Niell & Stryker, 2010

Binocular vision in humans and mice

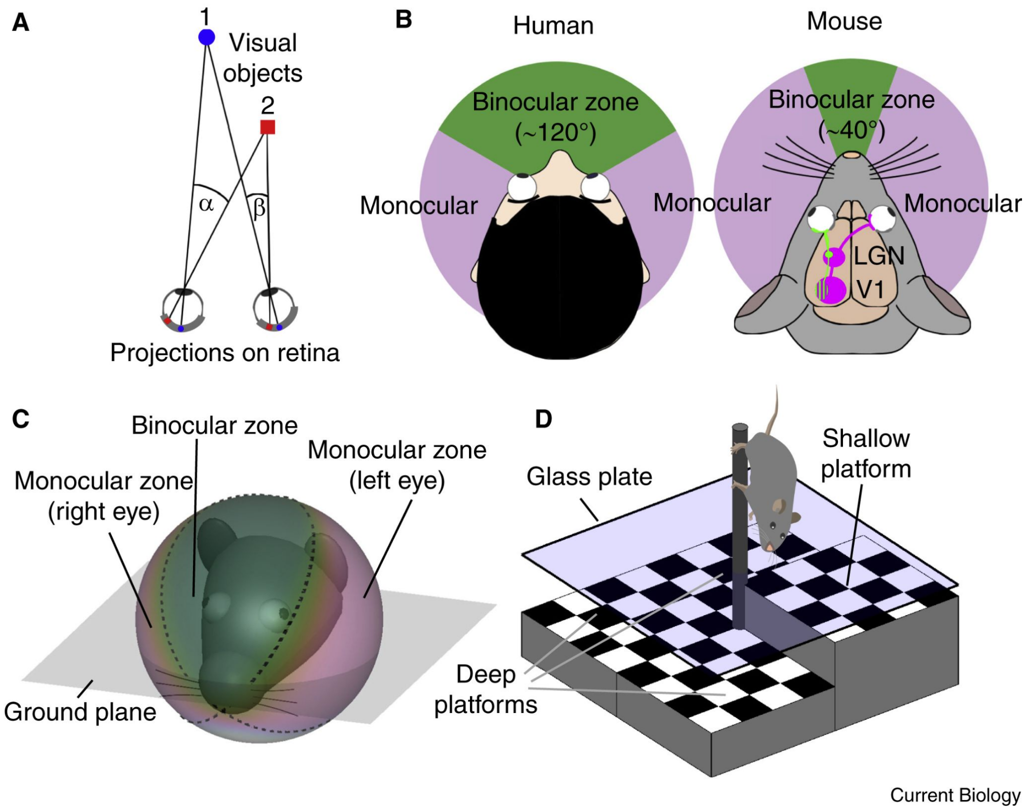

Poort & Meyer, 2021

From week 1:
many unsolved challenge…

How Do Neurons Interact?

7. How Can the Brain Be So Fast?
8. What Is the Neural Code?
9. Are Single Cortical Neurons Soloists or Are They Obedient Members of a Huge Orchestra?
10. What Is the Other 85 Percent of V1 Doing?

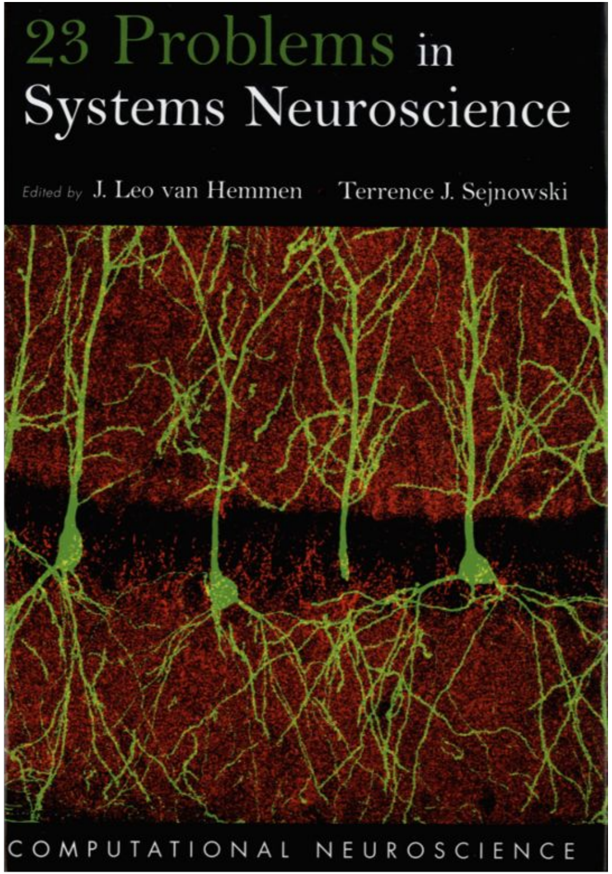

From this week: What other tasks? What other stimuli is the brain (visual) encoding?

Biological: recordings in visual system

Building "digital twins", NN models of the system

vision

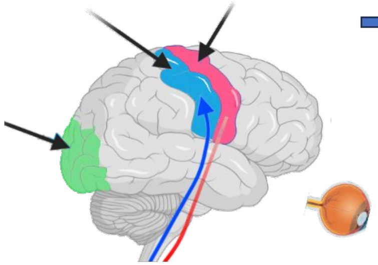

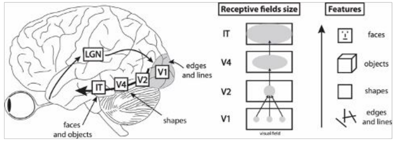

Hubel & Wiesel discoveries in cat V1 inspired convolutional neural networks

We now now a lot more (faces, motion, value coding) but we never can give enough stimuli ....
What would the ideal stimulus be for a given neuron?

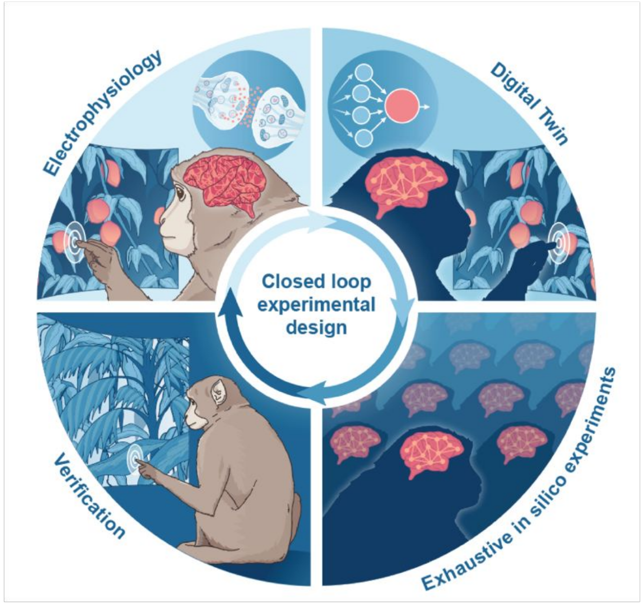

Mathis, Perez Rotondo, Chang, Tolias, Mathis (unpublished)

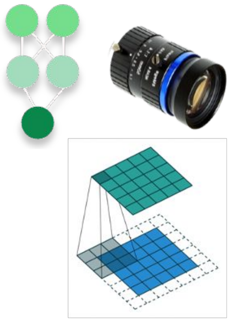

Can we use our NN to produce predictions of optimal stimuli? Can this help reveal anew computational principle, or validate a discovered rule?

→ Cortical sensory processings is **nonlinear**, inputs are **high dimensional**.

What are the other neural responses in V1 representing?

Aim: Develop a deep predictive model for causal testing of visual processing.

Section Paper:

Inception loops discover what excites neurons most using deep predictive models

Edgar Y. Walker[1,2,8*], Fabian H. Sinz[1,2,3,4,8*], Erick Cobos[1,2], Taliah Muhammad[1,2], Emmanouil Froudarakis[1,2], Paul G. Fahey[1,2], Alexander S. Ecker[1,3,5,6], Jacob Reimer[1,2], Xaq Pitkow[1,2,7] and Andreas S. Tolias[1,2,7*]

Finding sensory stimuli that drive neurons optimally is central to understanding information processing in the brain. However, optimizing sensory input is difficult due to the predominantly nonlinear nature of sensory processing and high dimensionality of the input. We developed 'inception loops', a closed-loop experimental paradigm combining in vivo recordings from thousands of neurons with in silico nonlinear response modeling. **Our end-to-end trained, deep-learning-based model predicted thousands of neuronal responses to arbitrary, new natural input with high accuracy** and was used to **synthesize optimal stimuli**—most exciting inputs (MEIs). For mouse primary visual cortex (V1), MEIs exhibited complex spatial features that occurred frequently in natural scenes but deviated strikingly from the common notion that Gabor-like stimuli are optimal for V1. When presented back to the same neurons in vivo, MEIs drove responses significantly better than control stimuli. **Inception loops represent a widely applicable technique for dissecting the neural mechanisms of sensation.**

Figure 1. Experimental paradigm and model

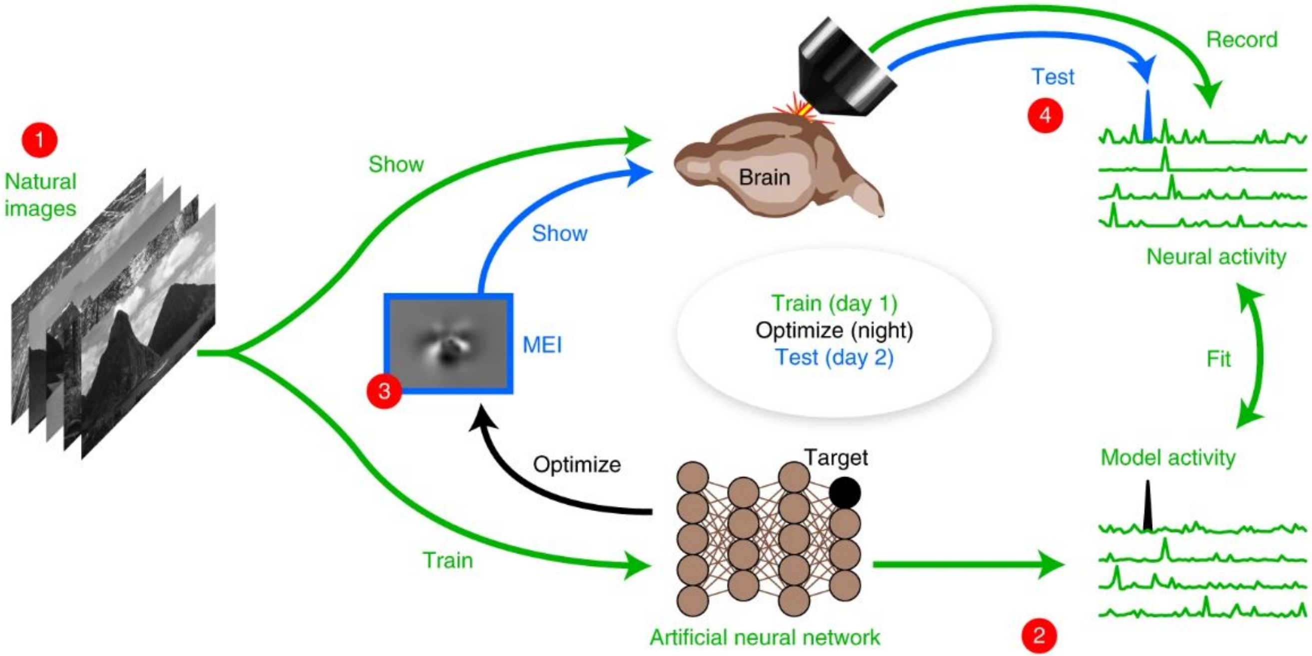

Explain the experimental setup. What is the goal?

closed-loop experimental paradigm combining in silico modeling (data-driven models) and in vivo neural recordings (causal testing) to synthesize stimuli that evoke a desired response confirmed in vivo. Causal testing of visual processing.

Figure 1. Experimental paradigm and model

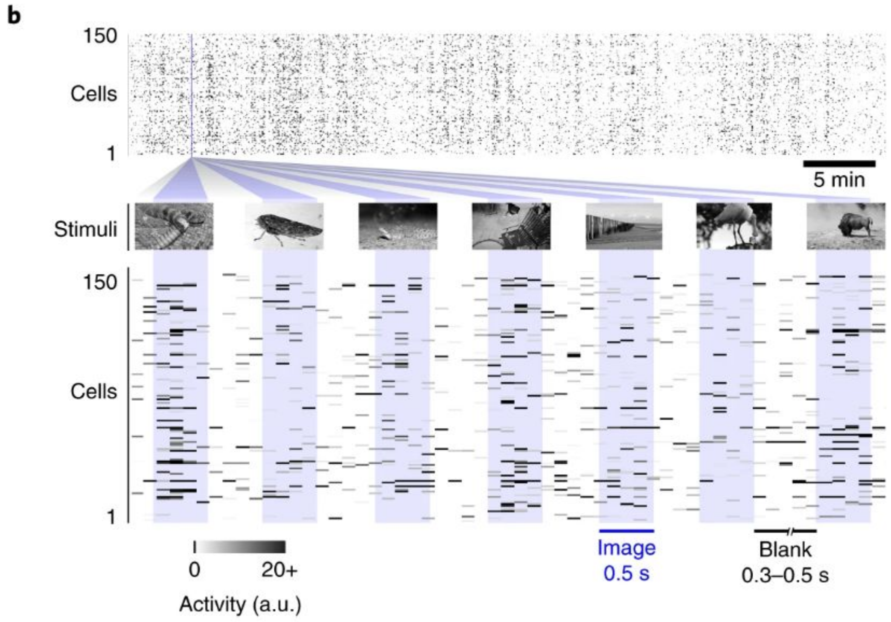

What is Panel (b.) showing?

Dataset: 5'100 unique natural images (ImageNet). 500ms each. 100 images x 10 repeat. random ITI. Mouse head-fixed, cylindrical treadmill. Recording: ~2000 neurons per animal, 5 animals, V1 L2/3, wide-field 2P microscope.

Why do they allow the mice to run during the experiment?

Movements have been shown to be an important modulator in visual responses in mice (Niell & Stryker 2010).

Why do they consider neural activity on a time window of 50-550 ms after image onset?

Time window 50 ms after image onset to consider signal latency to reach the visual cortex.

Figure 1. Experimental paradigm and model

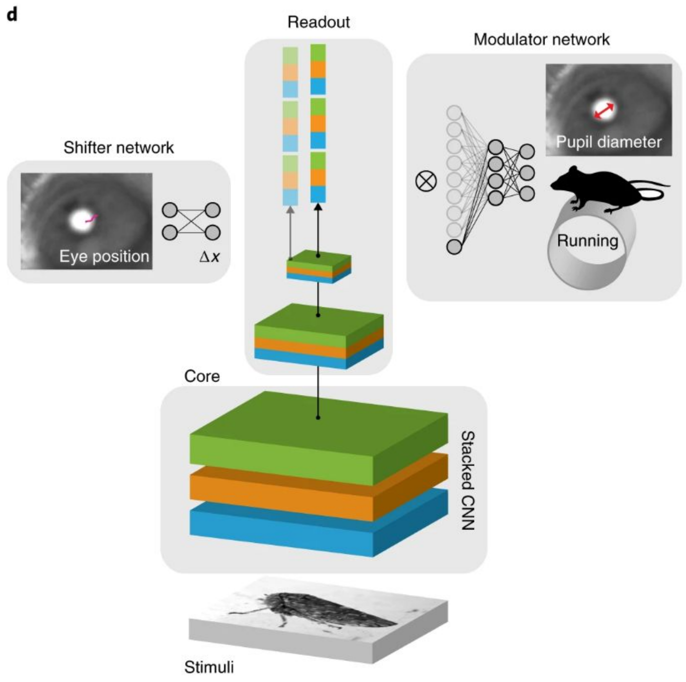

Why do the authors added the shifter network?

Account for receptive fields shifts from pupil position changes.

And the modulator network?

Attenuate the effect of movement (running state) and arousal (pupil dilation) on visual responses.

Reminder: Field of view

Field of view:

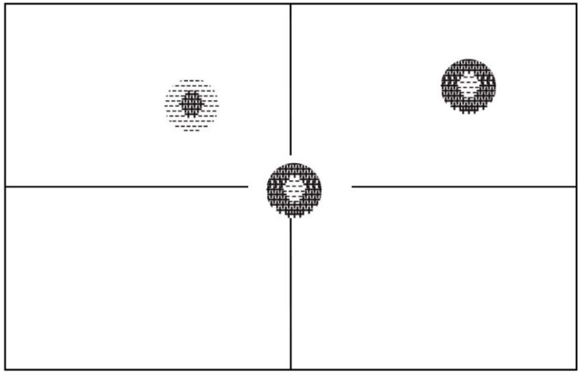

If you record from a retinal ganglion cell (RGC).
They fire APs with generally two types of responses:

On-center
ganglion cell

"ON-center":
↑ AP- frequency
in the center of the receptive field (RF)

Off-center
ganglion cell

"OFF center"
↓ AP-frequency
in the center of the RF

Single-cell recordings of sensory neurons at any position on the retina (sensory epithelium)

Figure 1. Experimental paradigm and model

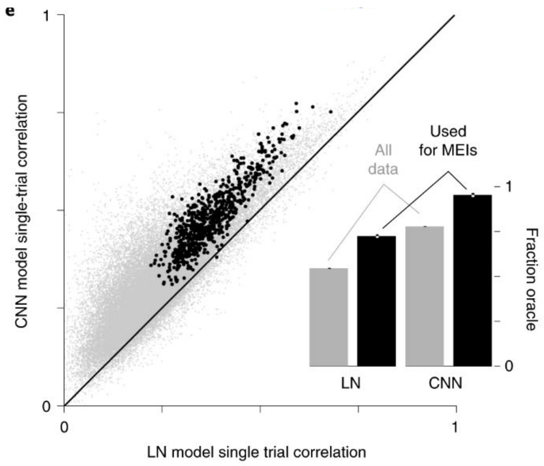

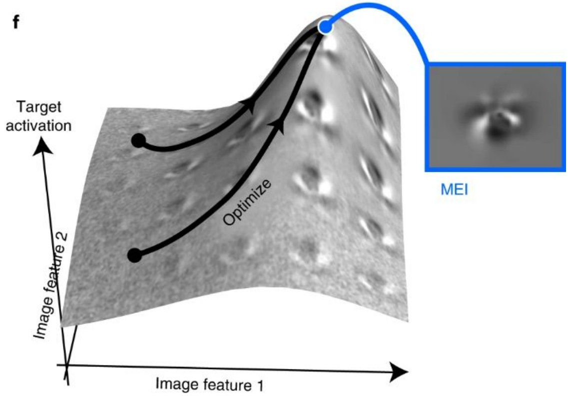

Why are they comparing a nonlinear (CNN) to a linear (LN) model?

CNN outperforms LN. Gabor-features are linear. The fact that a nonlinear model captures a bigger variability of responses suggest that Gabor filters might not express the full diversity of V1 neural responses.

How is the MEI optimized?

- *150 neurons reliable and reasonably well-predicted by both the CNN and the LN model.*
- *regularized gradient ascent from random image.*

Figure 2. Most Exciting Inputs (MEIs)

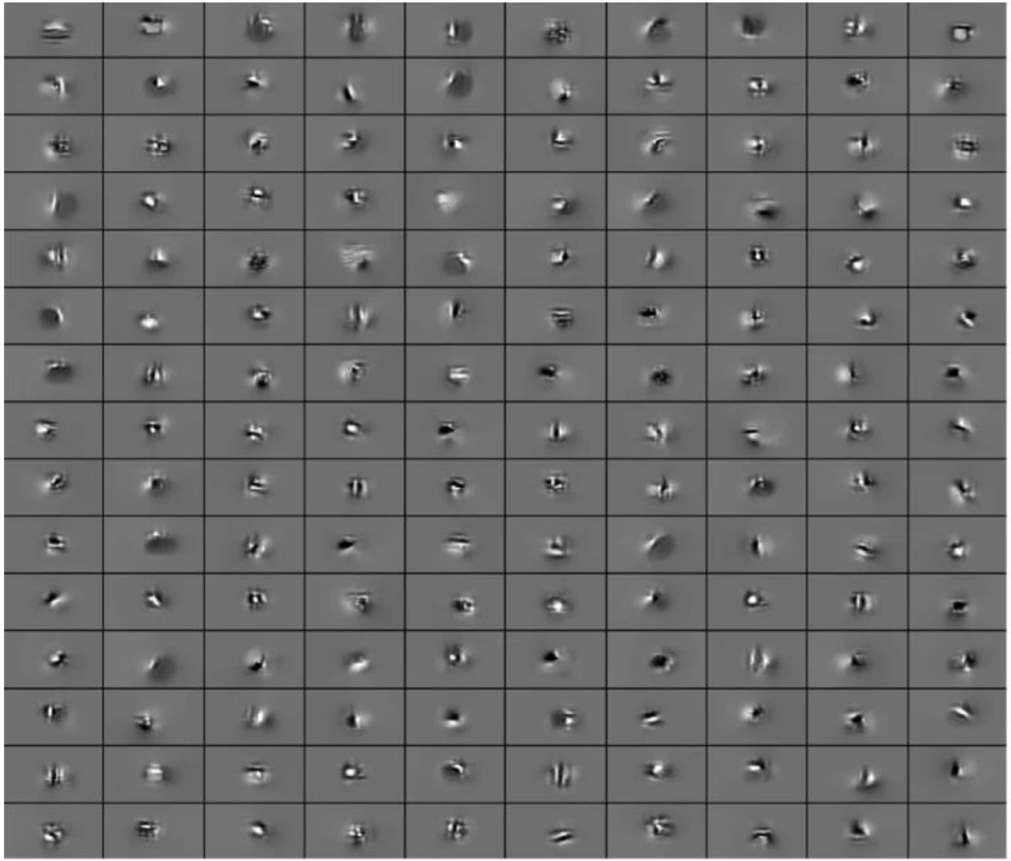

From the way RFs are traditionally thought about in V1, what can you say about the obtained MEIs? Is it what you would expect?

Some resulting MEIs are Gabor-like. Lots of them deviate substantially from Gabor-shaped V1 RFs → sharp corners, checkerboard patterns, irregular pointillist textures and a variety of curved strokes.

Figure 2. Most Exciting Inputs (MEIs)

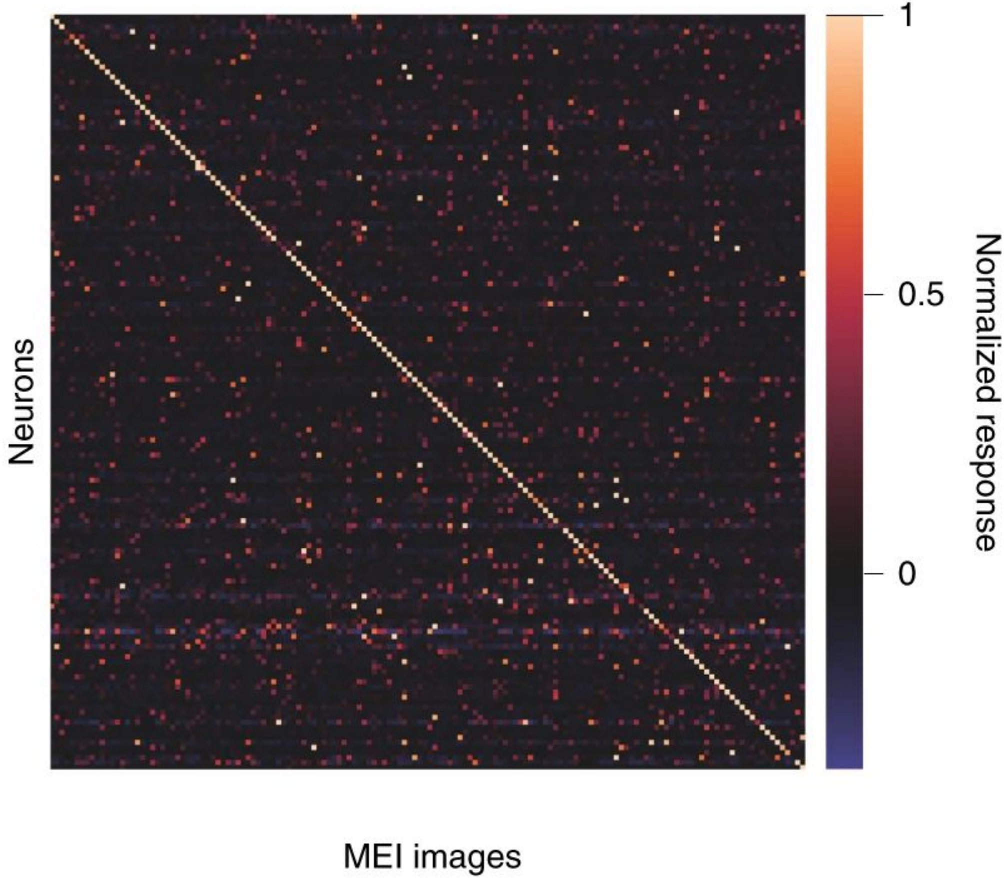

Why do we see sparsity?

- *Neuron activation with high specificity for each MEI.*
- *Sparse response (few images activating a neuron).*

Sparse coding in the visual cortex.

Figure 2. Most Exciting Inputs (MEIs)

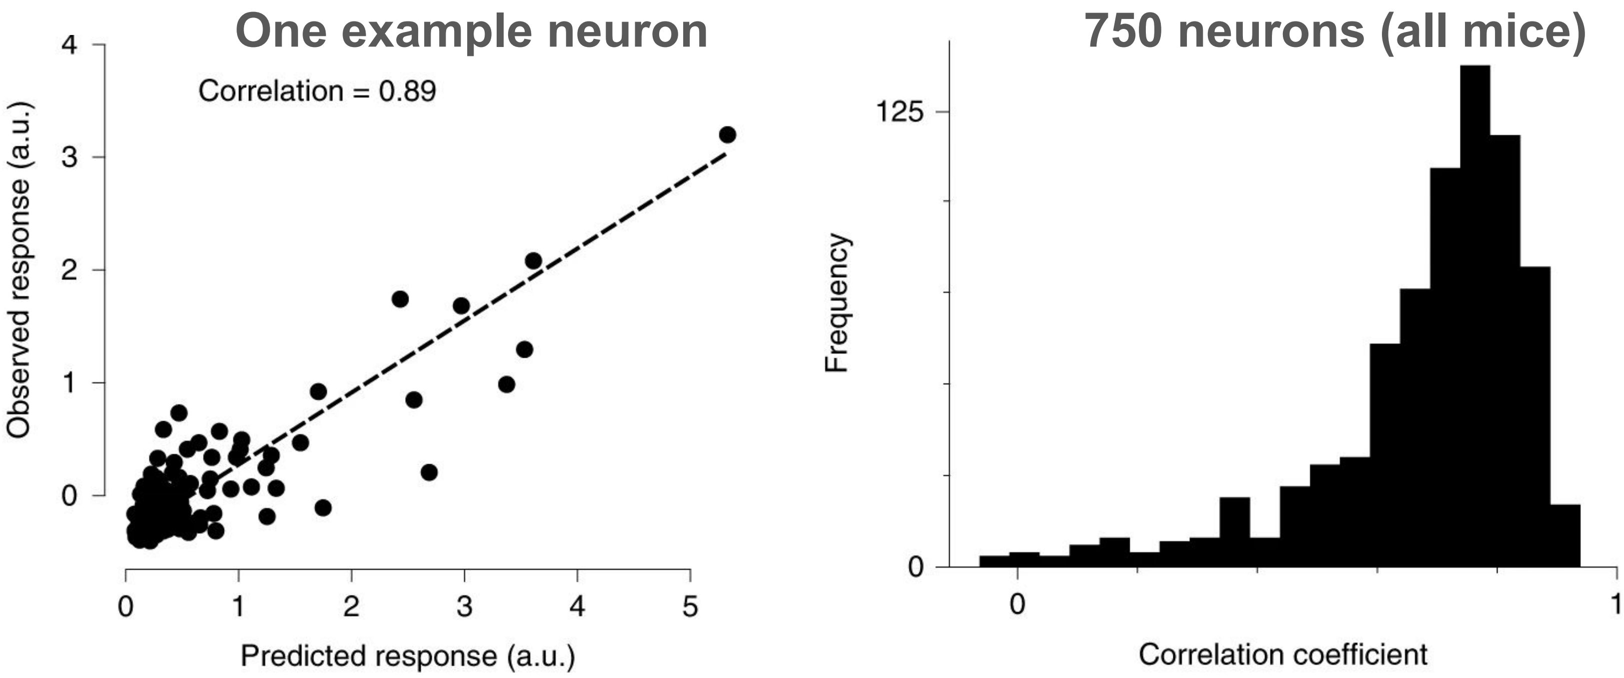

Model predictions correlates highly with observed responses.

Recap:

They manage to obtain MEIs that are specific to a neuron in V1.

Now, how can we understand the specific visual features encoded by those neurons?

What else should we compare the MEIs to?

Figure 3/4. Comparison of MEIs and linear RFs

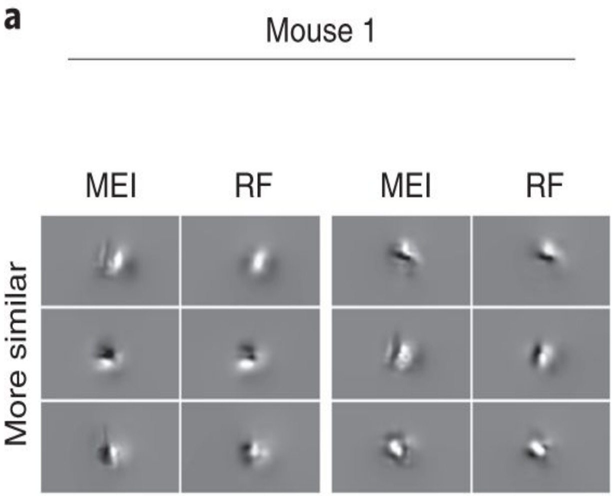

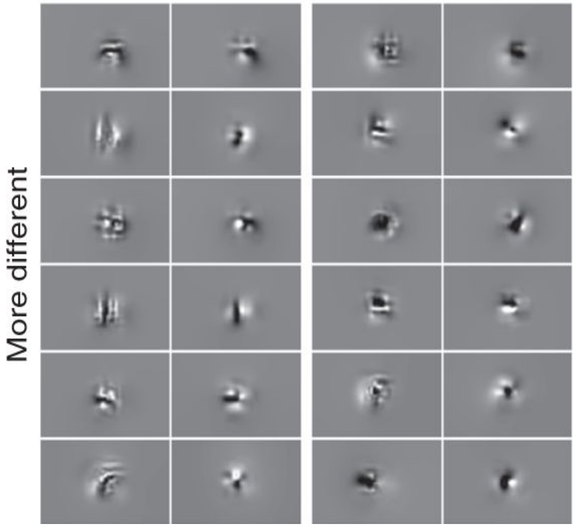

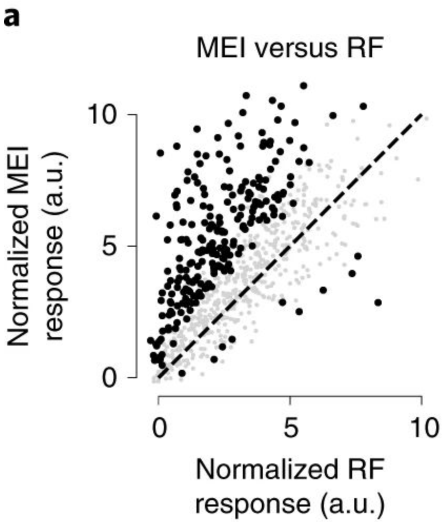

black: significant difference in mean response ($P < 0.05$).

Observation?

MEIs drive stronger neural response than linear RFs from the LN.

Figure 3/4. Comparison of MEIs and Gabor filters

Observation?

MEIs drive stronger neural response than Gabor filters.

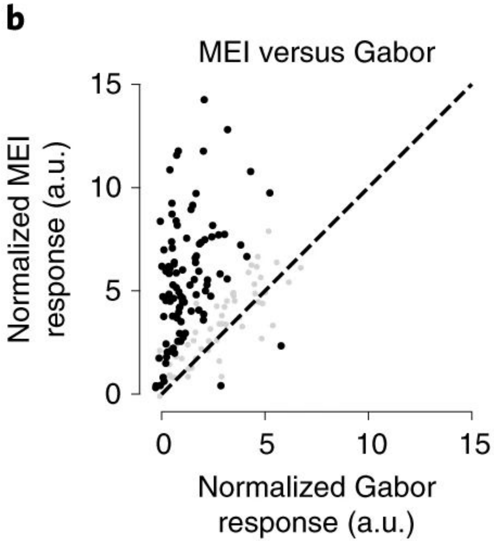

Figure 3/4. Comparison of MEIs and best masked natural images

Observation?

MEIs drive stronger neural response than masked natural images + resemblance.

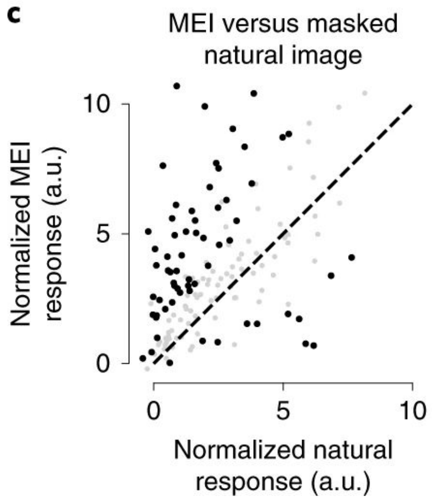

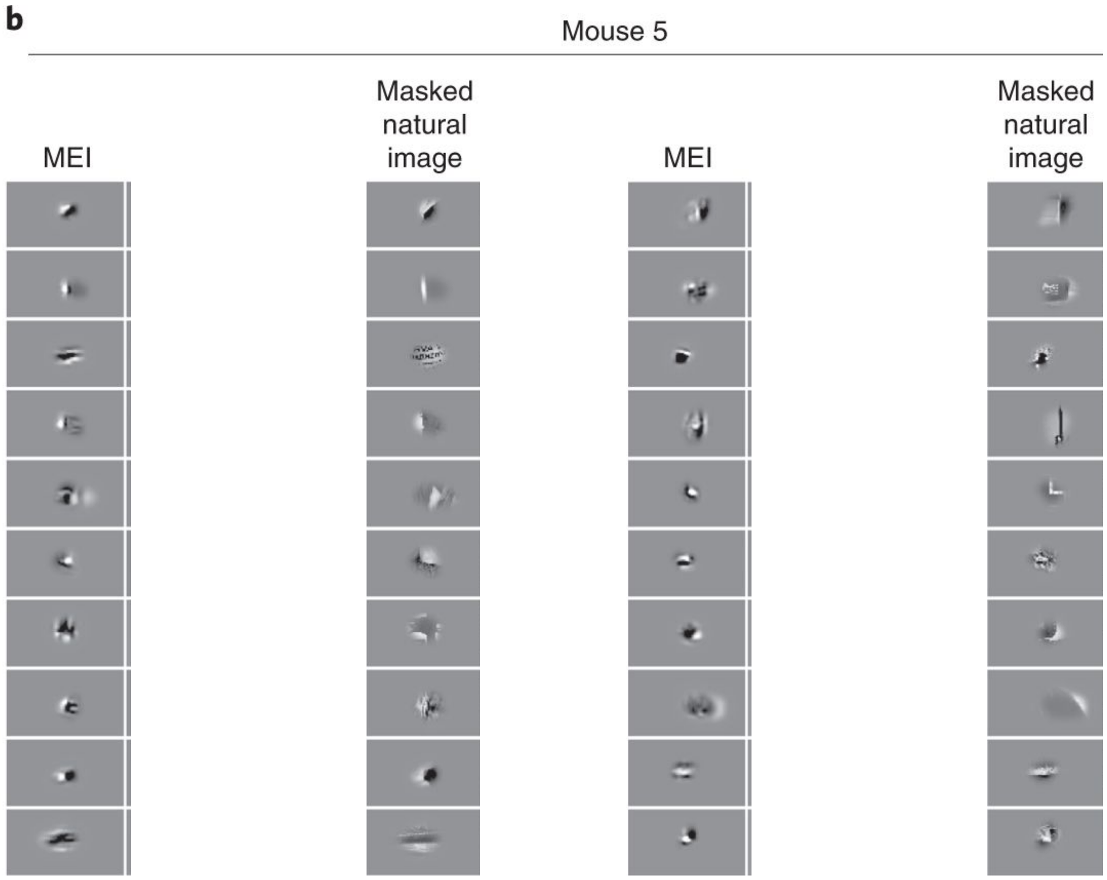

Figure 3/4. Comparison of MEIs and full-field natural images

Observation?

MEIs drive stronger neural response than full-field natural images.

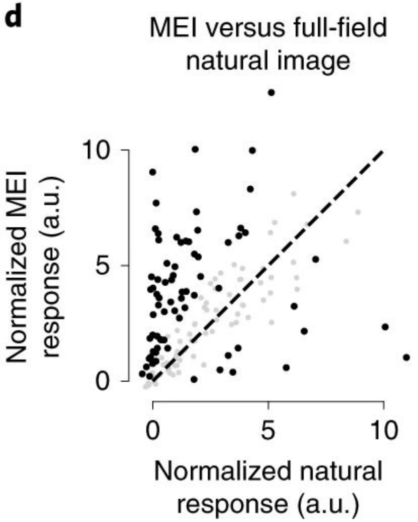

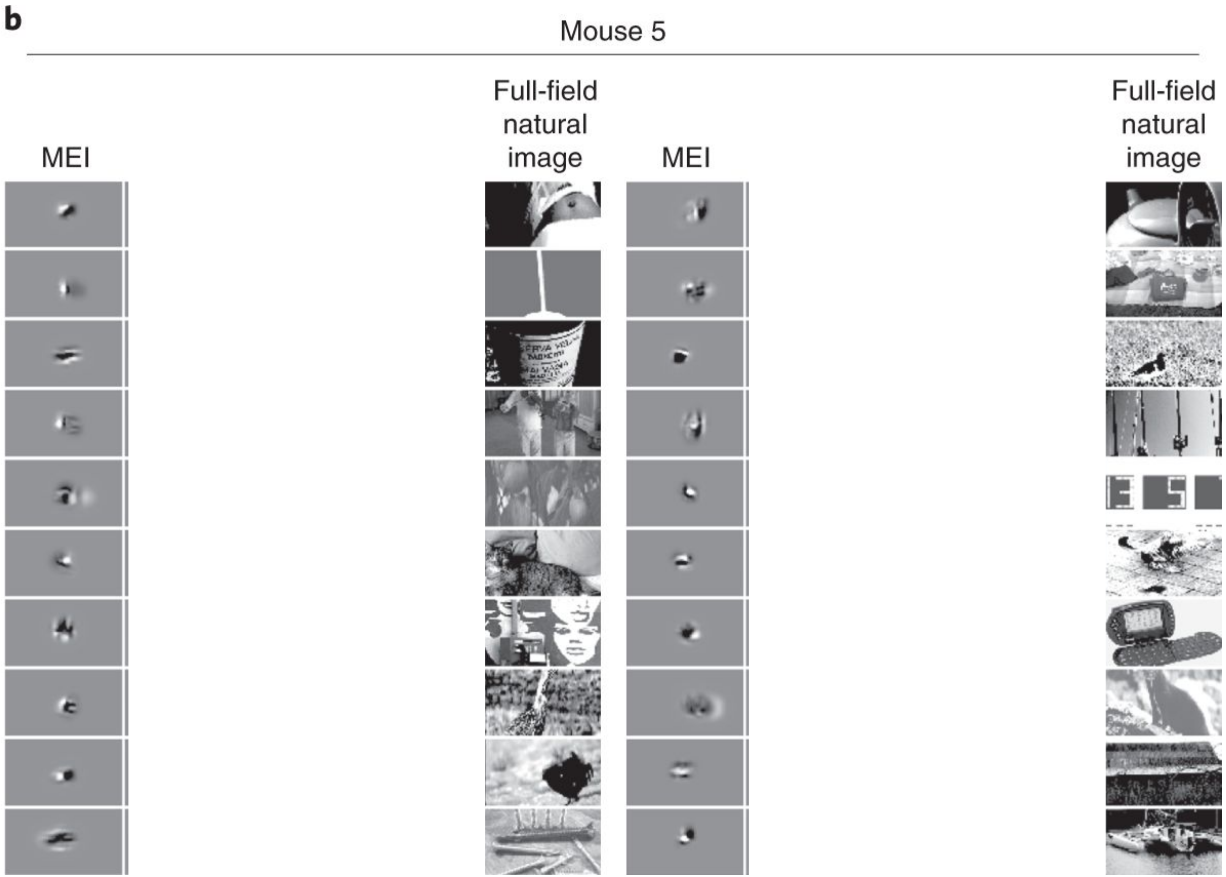

Figure 3/4. Comparison of MEIs and other control stimuli.

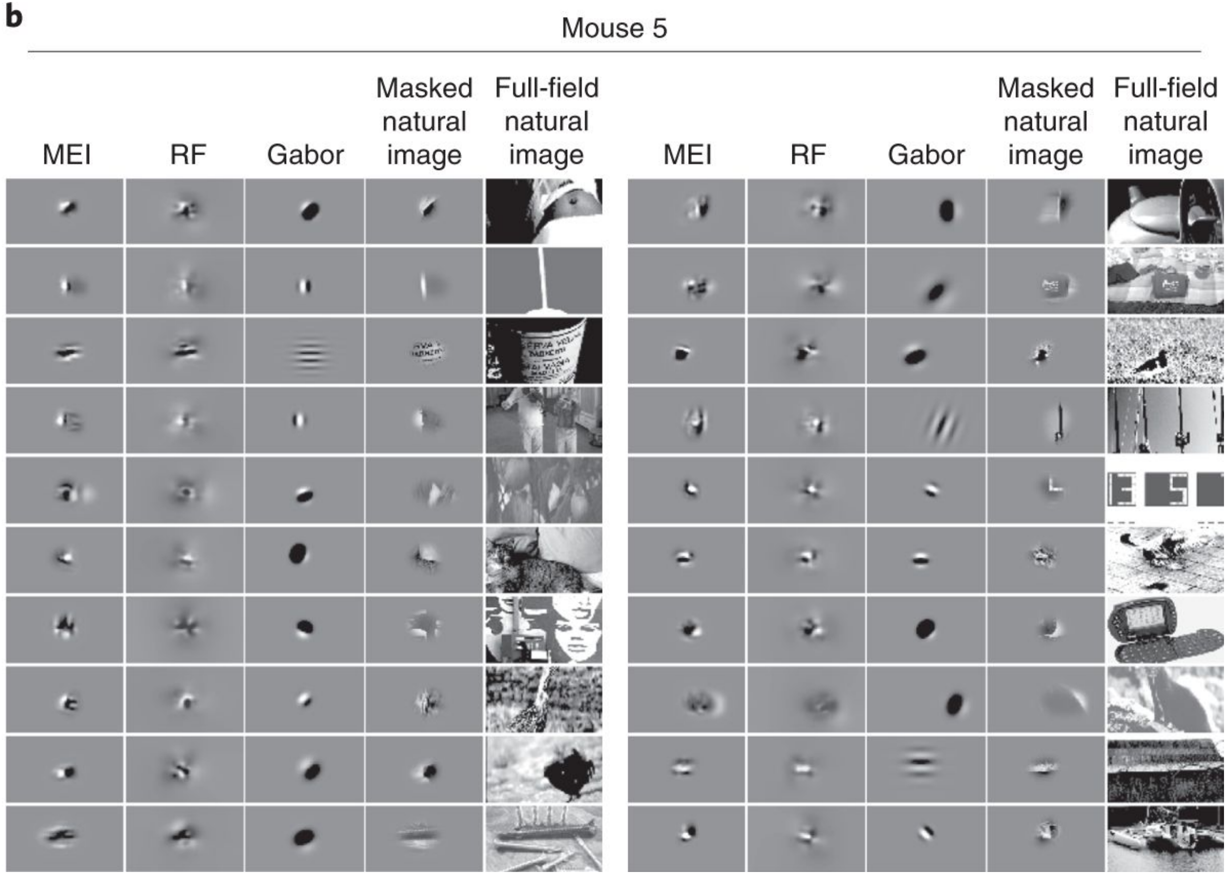

On average:

- 1.6% of images produced activations above 0.5 of MEI activation.
- 0.04% above 0.75. of the MEI activation.

→ more support for sparse encoding in V1.

Paper round-up

- They propose a closed-loop experimental paradigm combining in vivo recordings from thousands of neurons with in silico nonlinear response modeling.
- They show that high-performing, end-to-end trained, black-box models of the visual system generalize and can make in silico inferences about nontrivial computational properties of V1 neurons.
- They find that most MEIs deviate strikingly from Gabor-like stimuli, suggesting that even mouse V1 neurons prefer features that are more complex than the classical oriented edges (Gabor) described by Hubel and Wiesel.
- They show that the perceptual attributes of MEIs occur often in natural scenes.
- They propose their method to verify experimentally predictive models of optimized stimuli.

What did we learn? What questions do we have?

- **What points do they make in the discussion?**

- **Is anything unclear?**

- **What would you do next if you had to design an experiment?**
  - *Online causal testing rather than during the night? Build a foundation model of the visual cortex → Wang, 2023, preprint.*
  - *Apply to other visual cortex areas? → Wang, 2023, preprint.*
  - *Apply to other models → Bashivan, 2023, Science*
  - *Video rather than images → integration of motion to the RFs.*

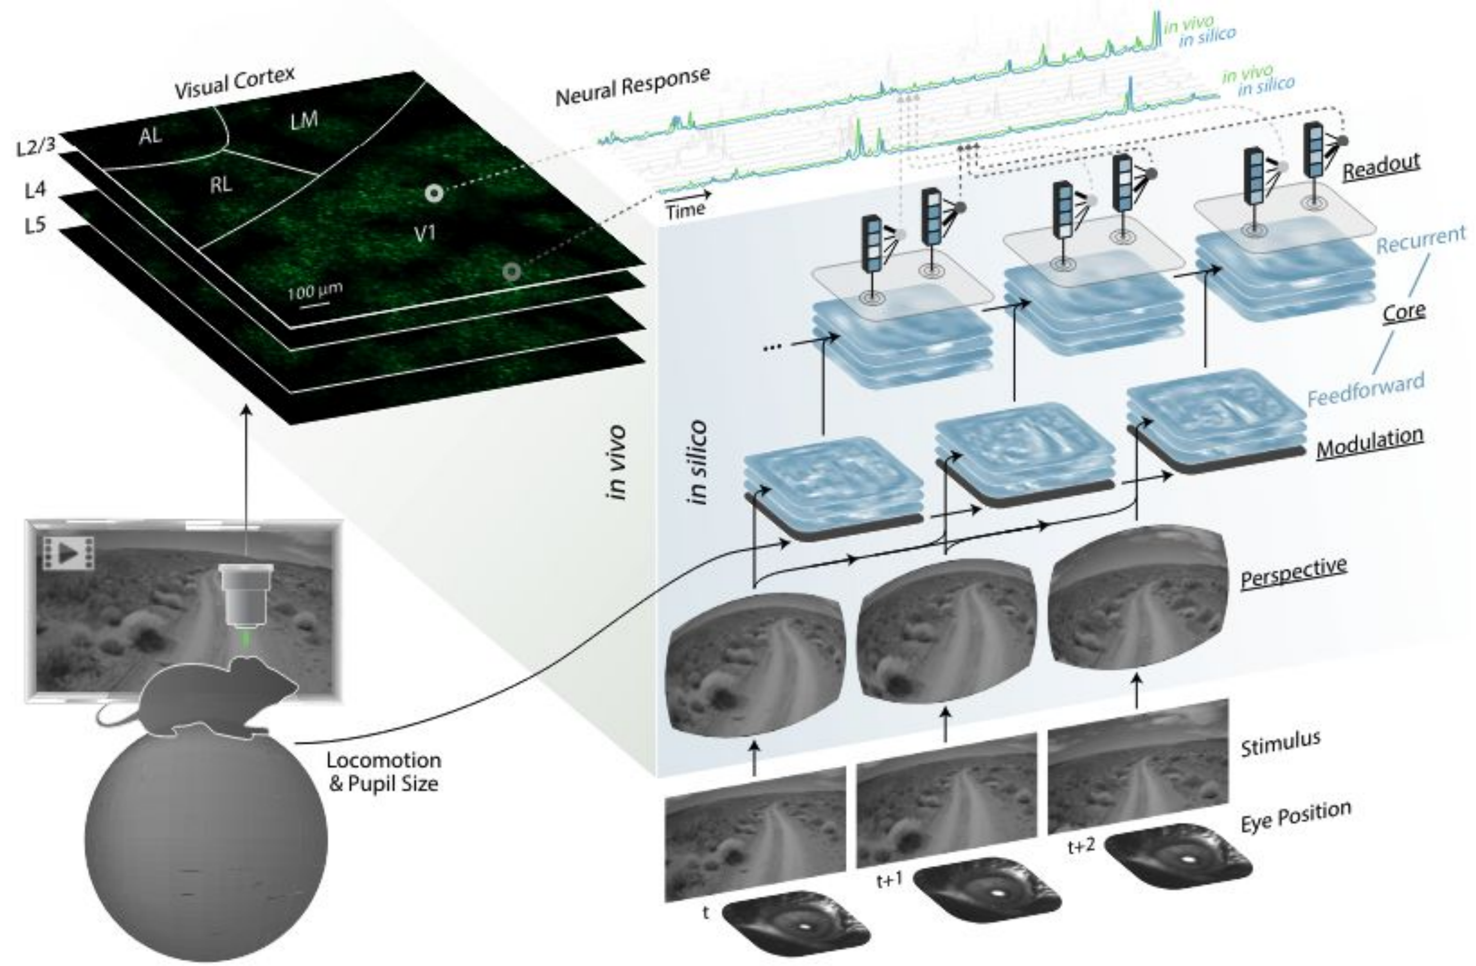

Wang, 2023, preprint