



# The Emergence of Convergence in Different Levels of Biology and AI

**Brian Cheung**

**cheungb@mit.edu**

Minyoung Huh (Jacob)



Tongzhou Wang



Phillip Isola



Erin Grant



Sophie Wang



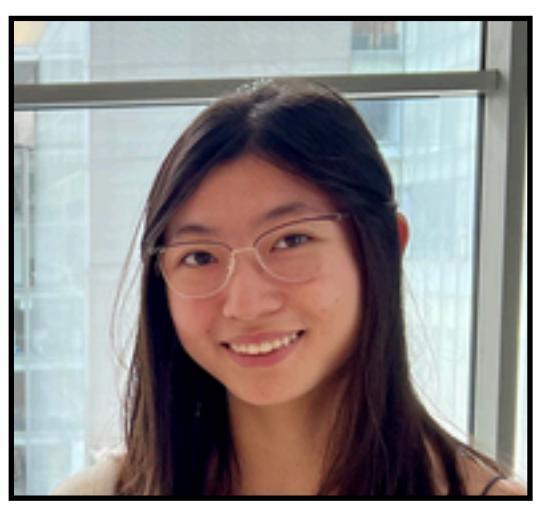
Yena Han



Eric Weiss



Helen Yang



Kushagra Tiwary



Aaron Young



Pulkit Agrawal



Bruno Olshausen



Boris Katz



Tomaso Poggio



Jascha Sohl-Dickstein



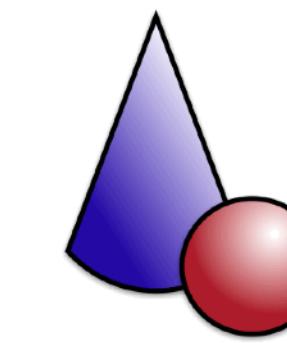
# REDWOOD CENTER

for Theoretical Neuroscience



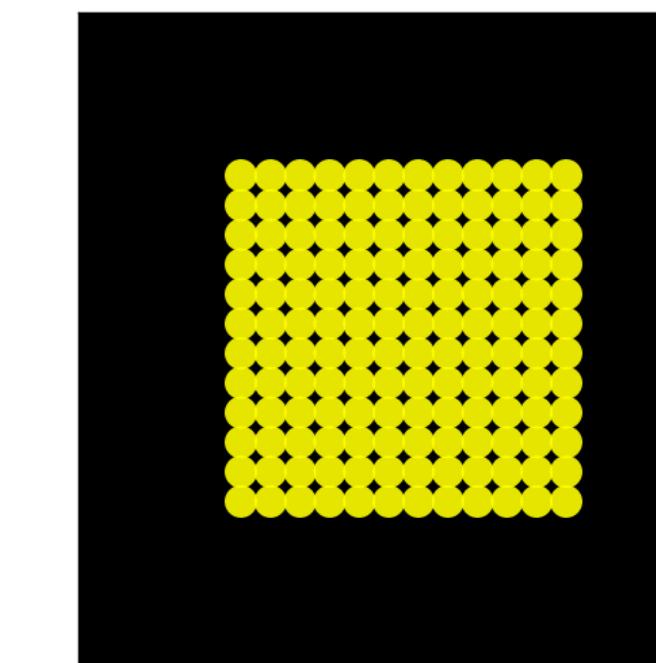
Google DeepMind

Observing convergence across many forms of intelligence



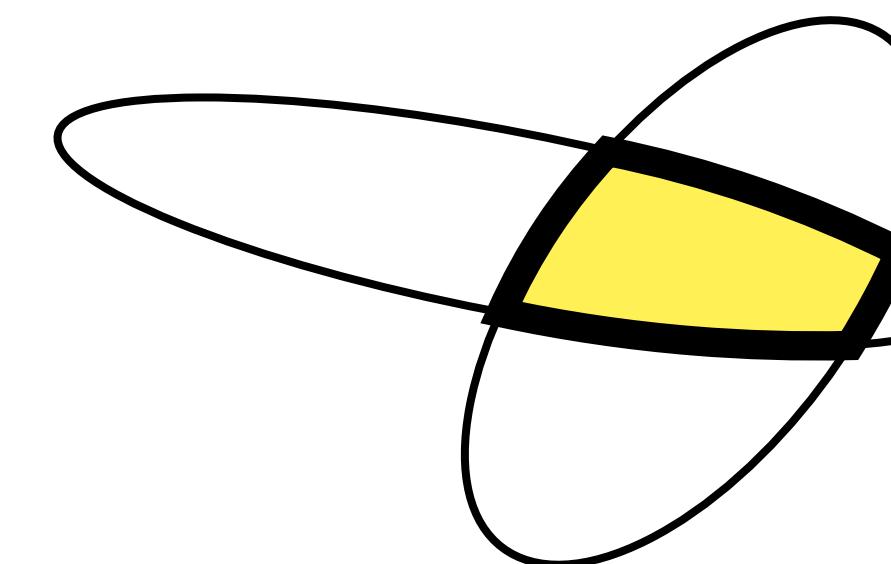
---

Controlling convergence to explain biology

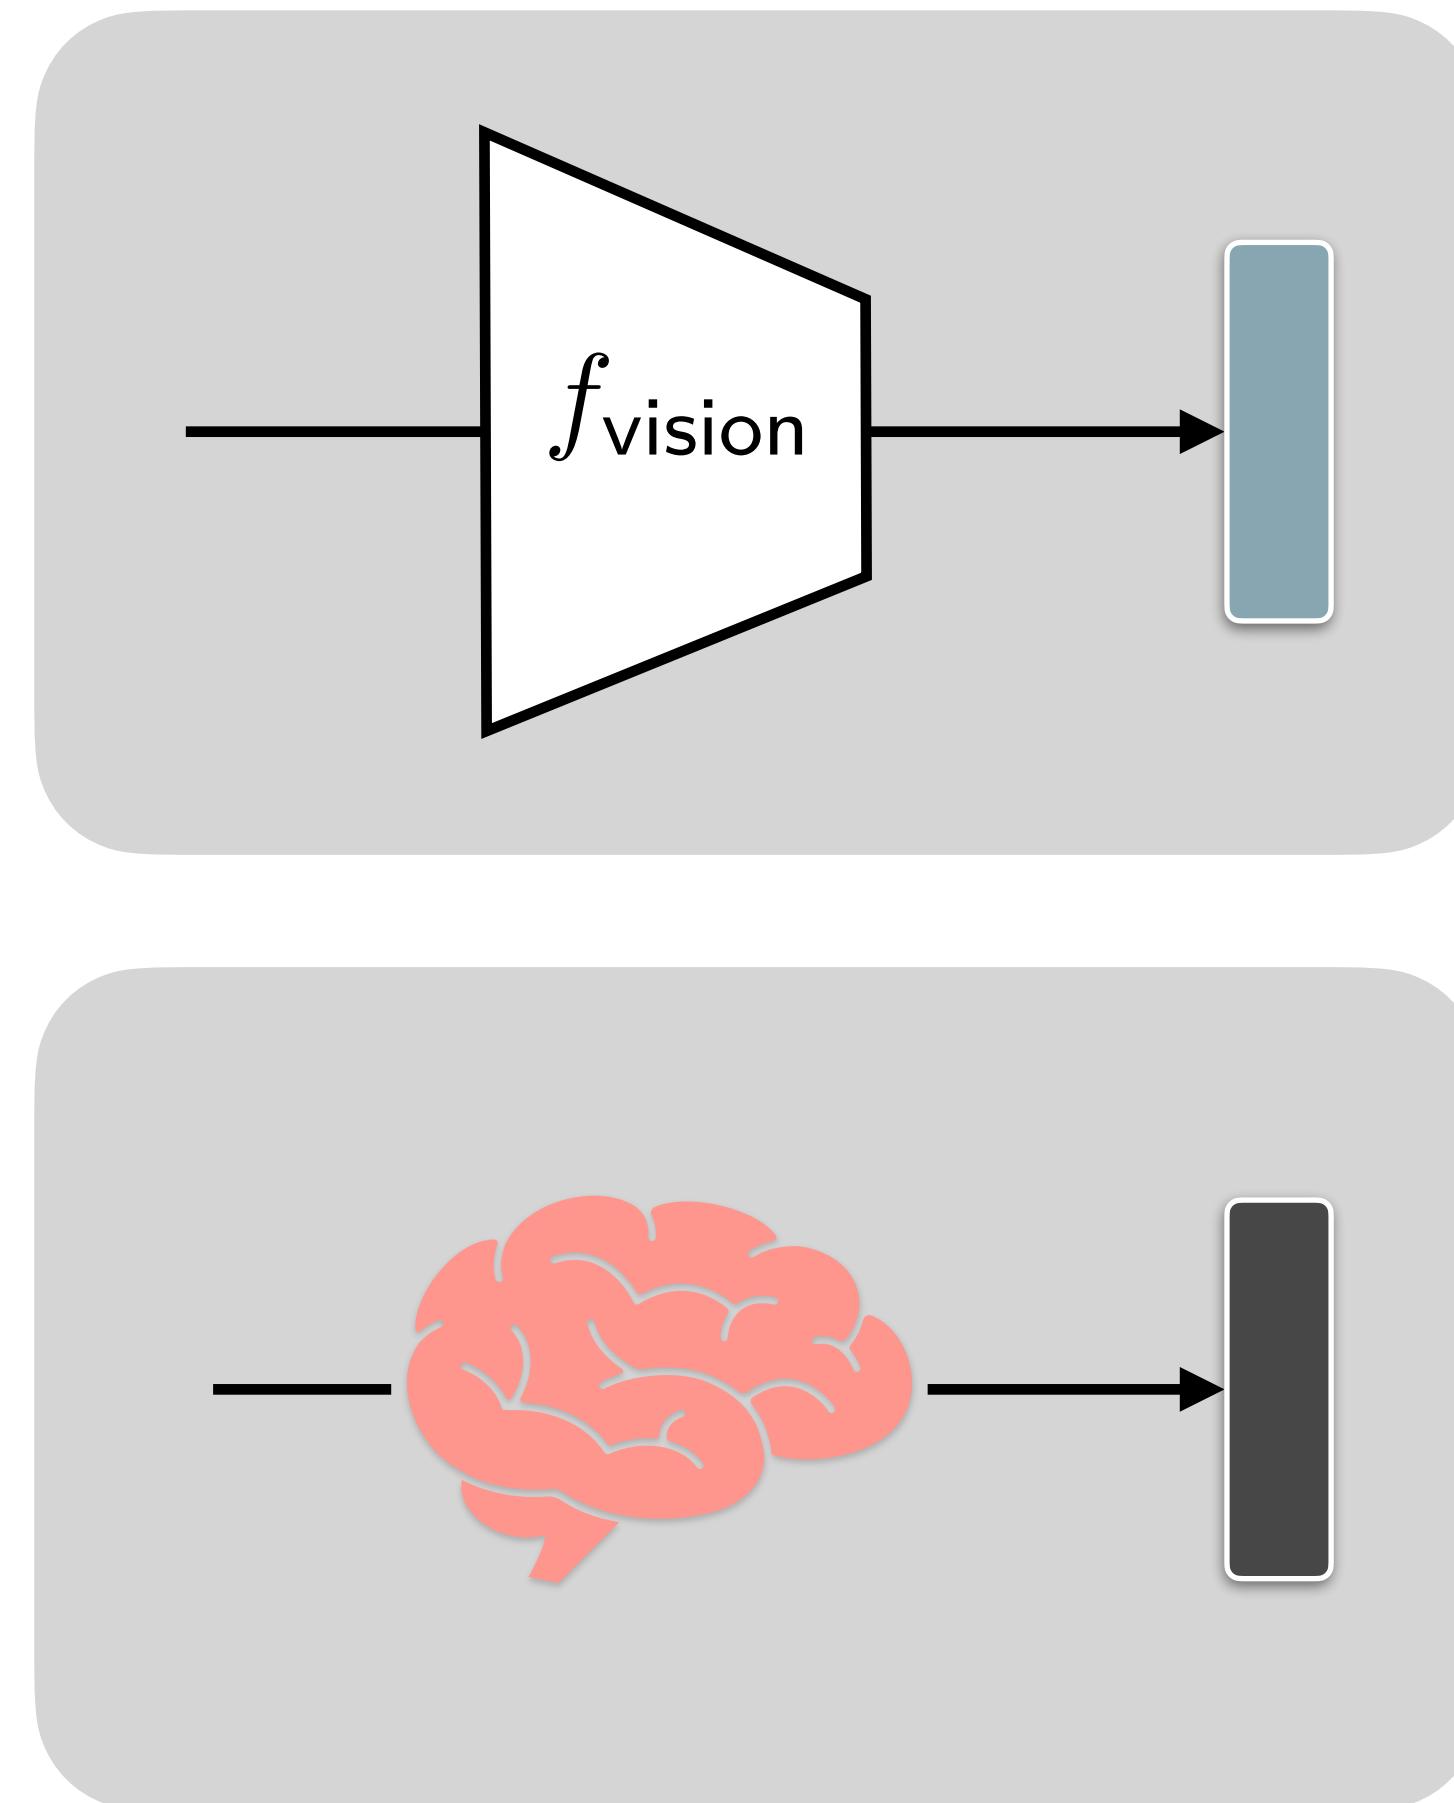
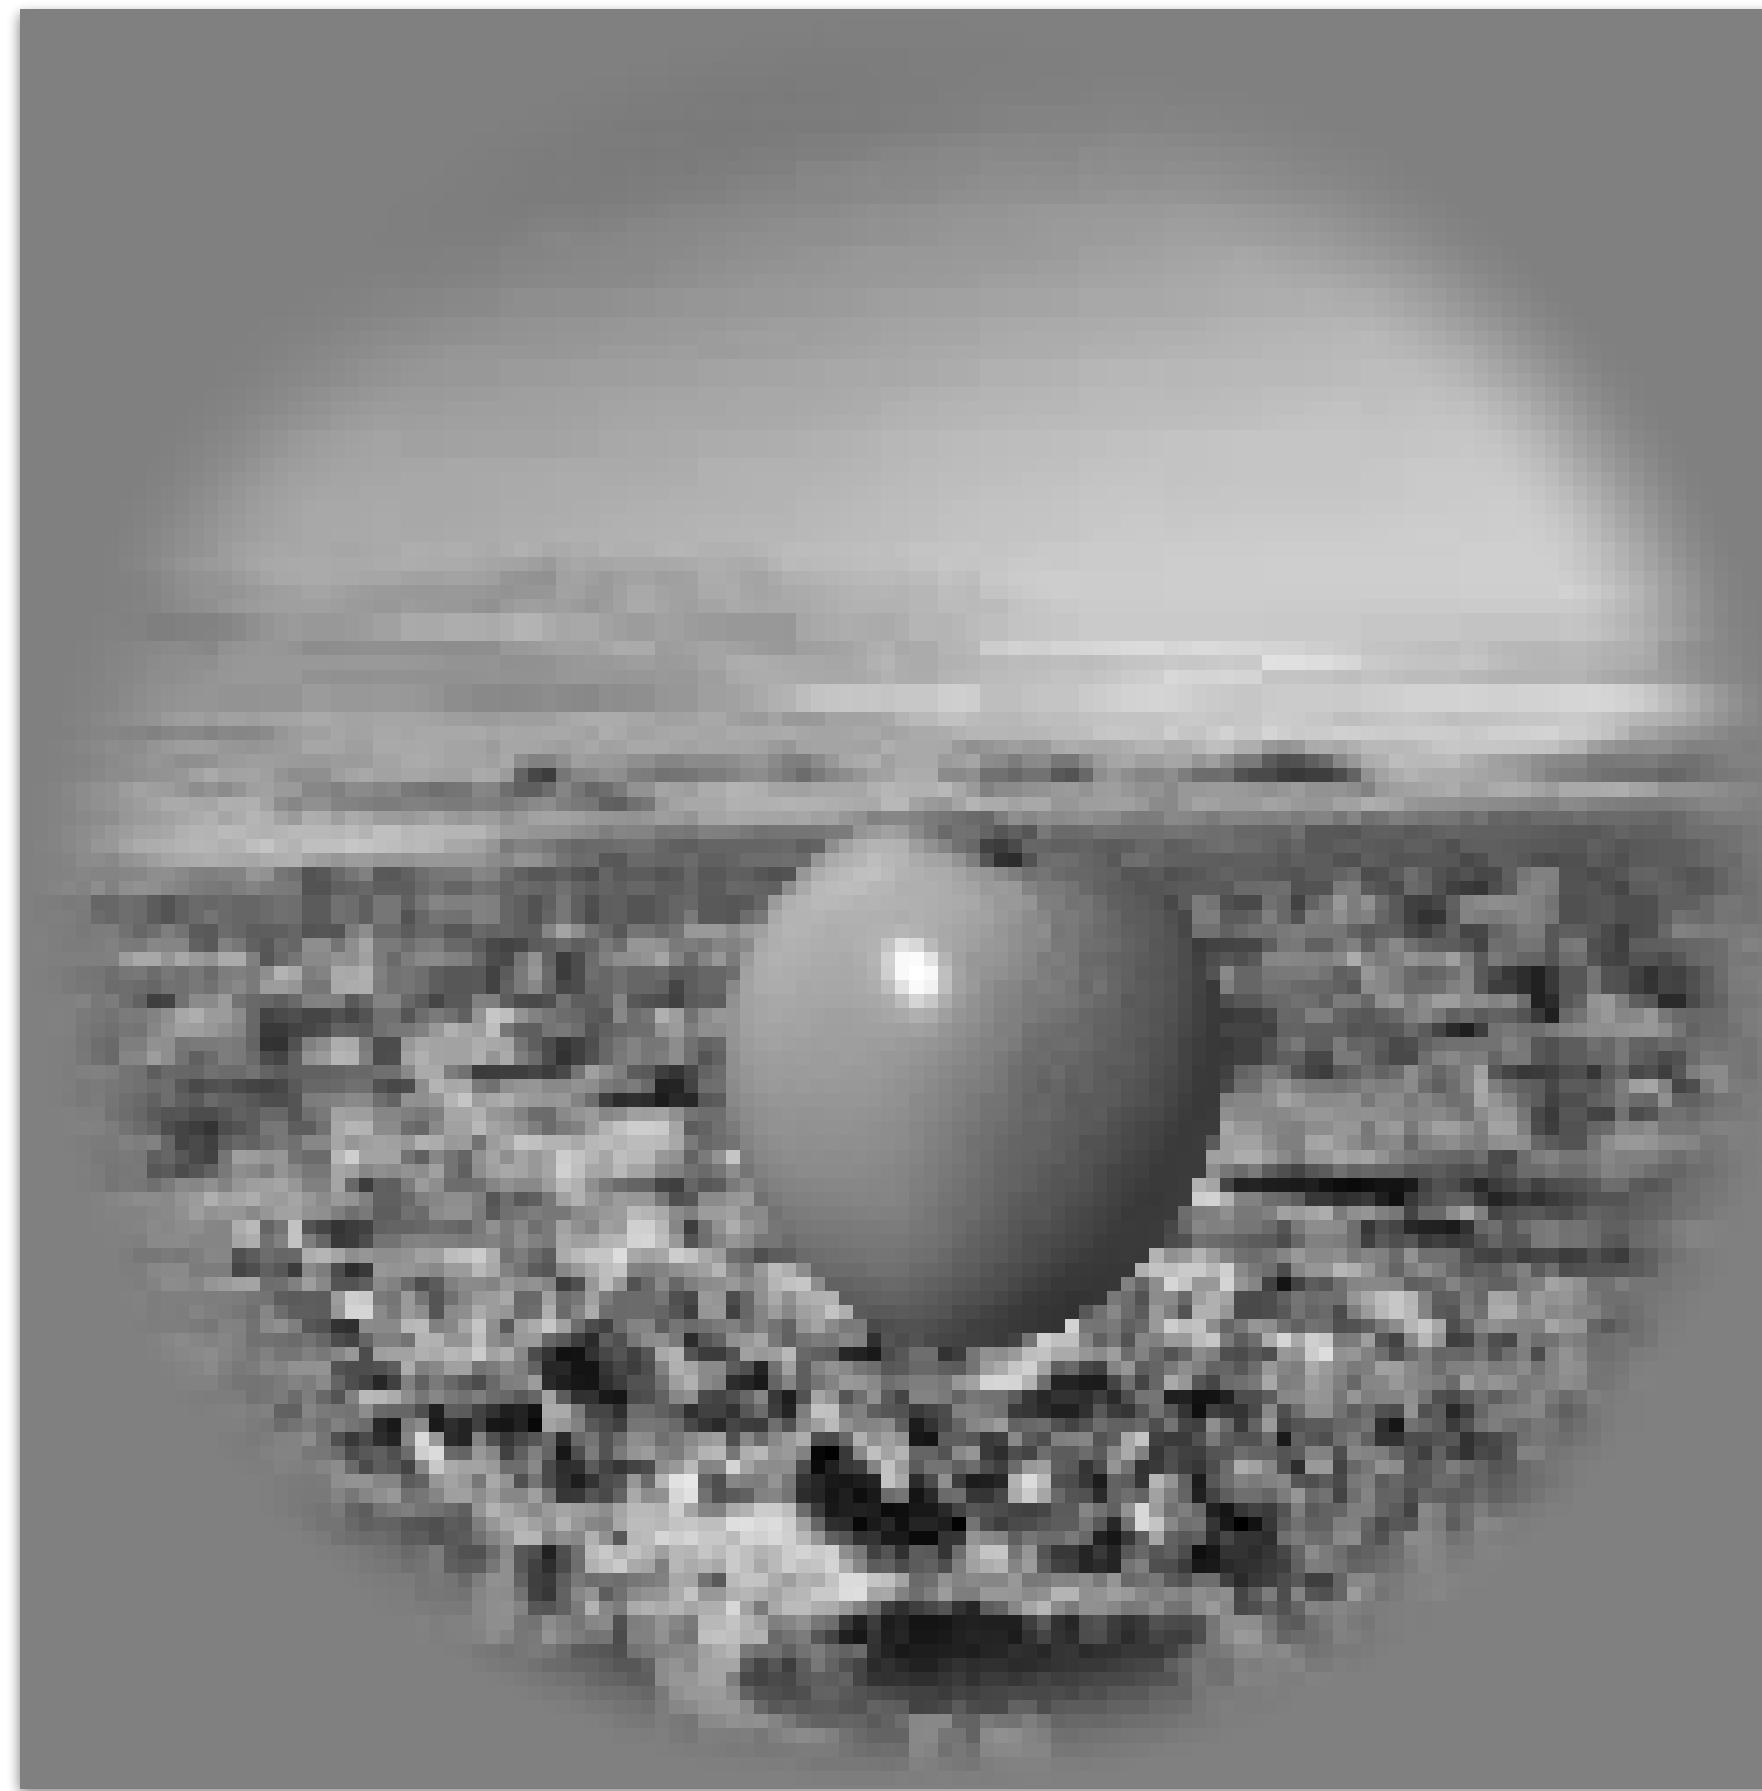


---

Future Work: Optimizing for convergence



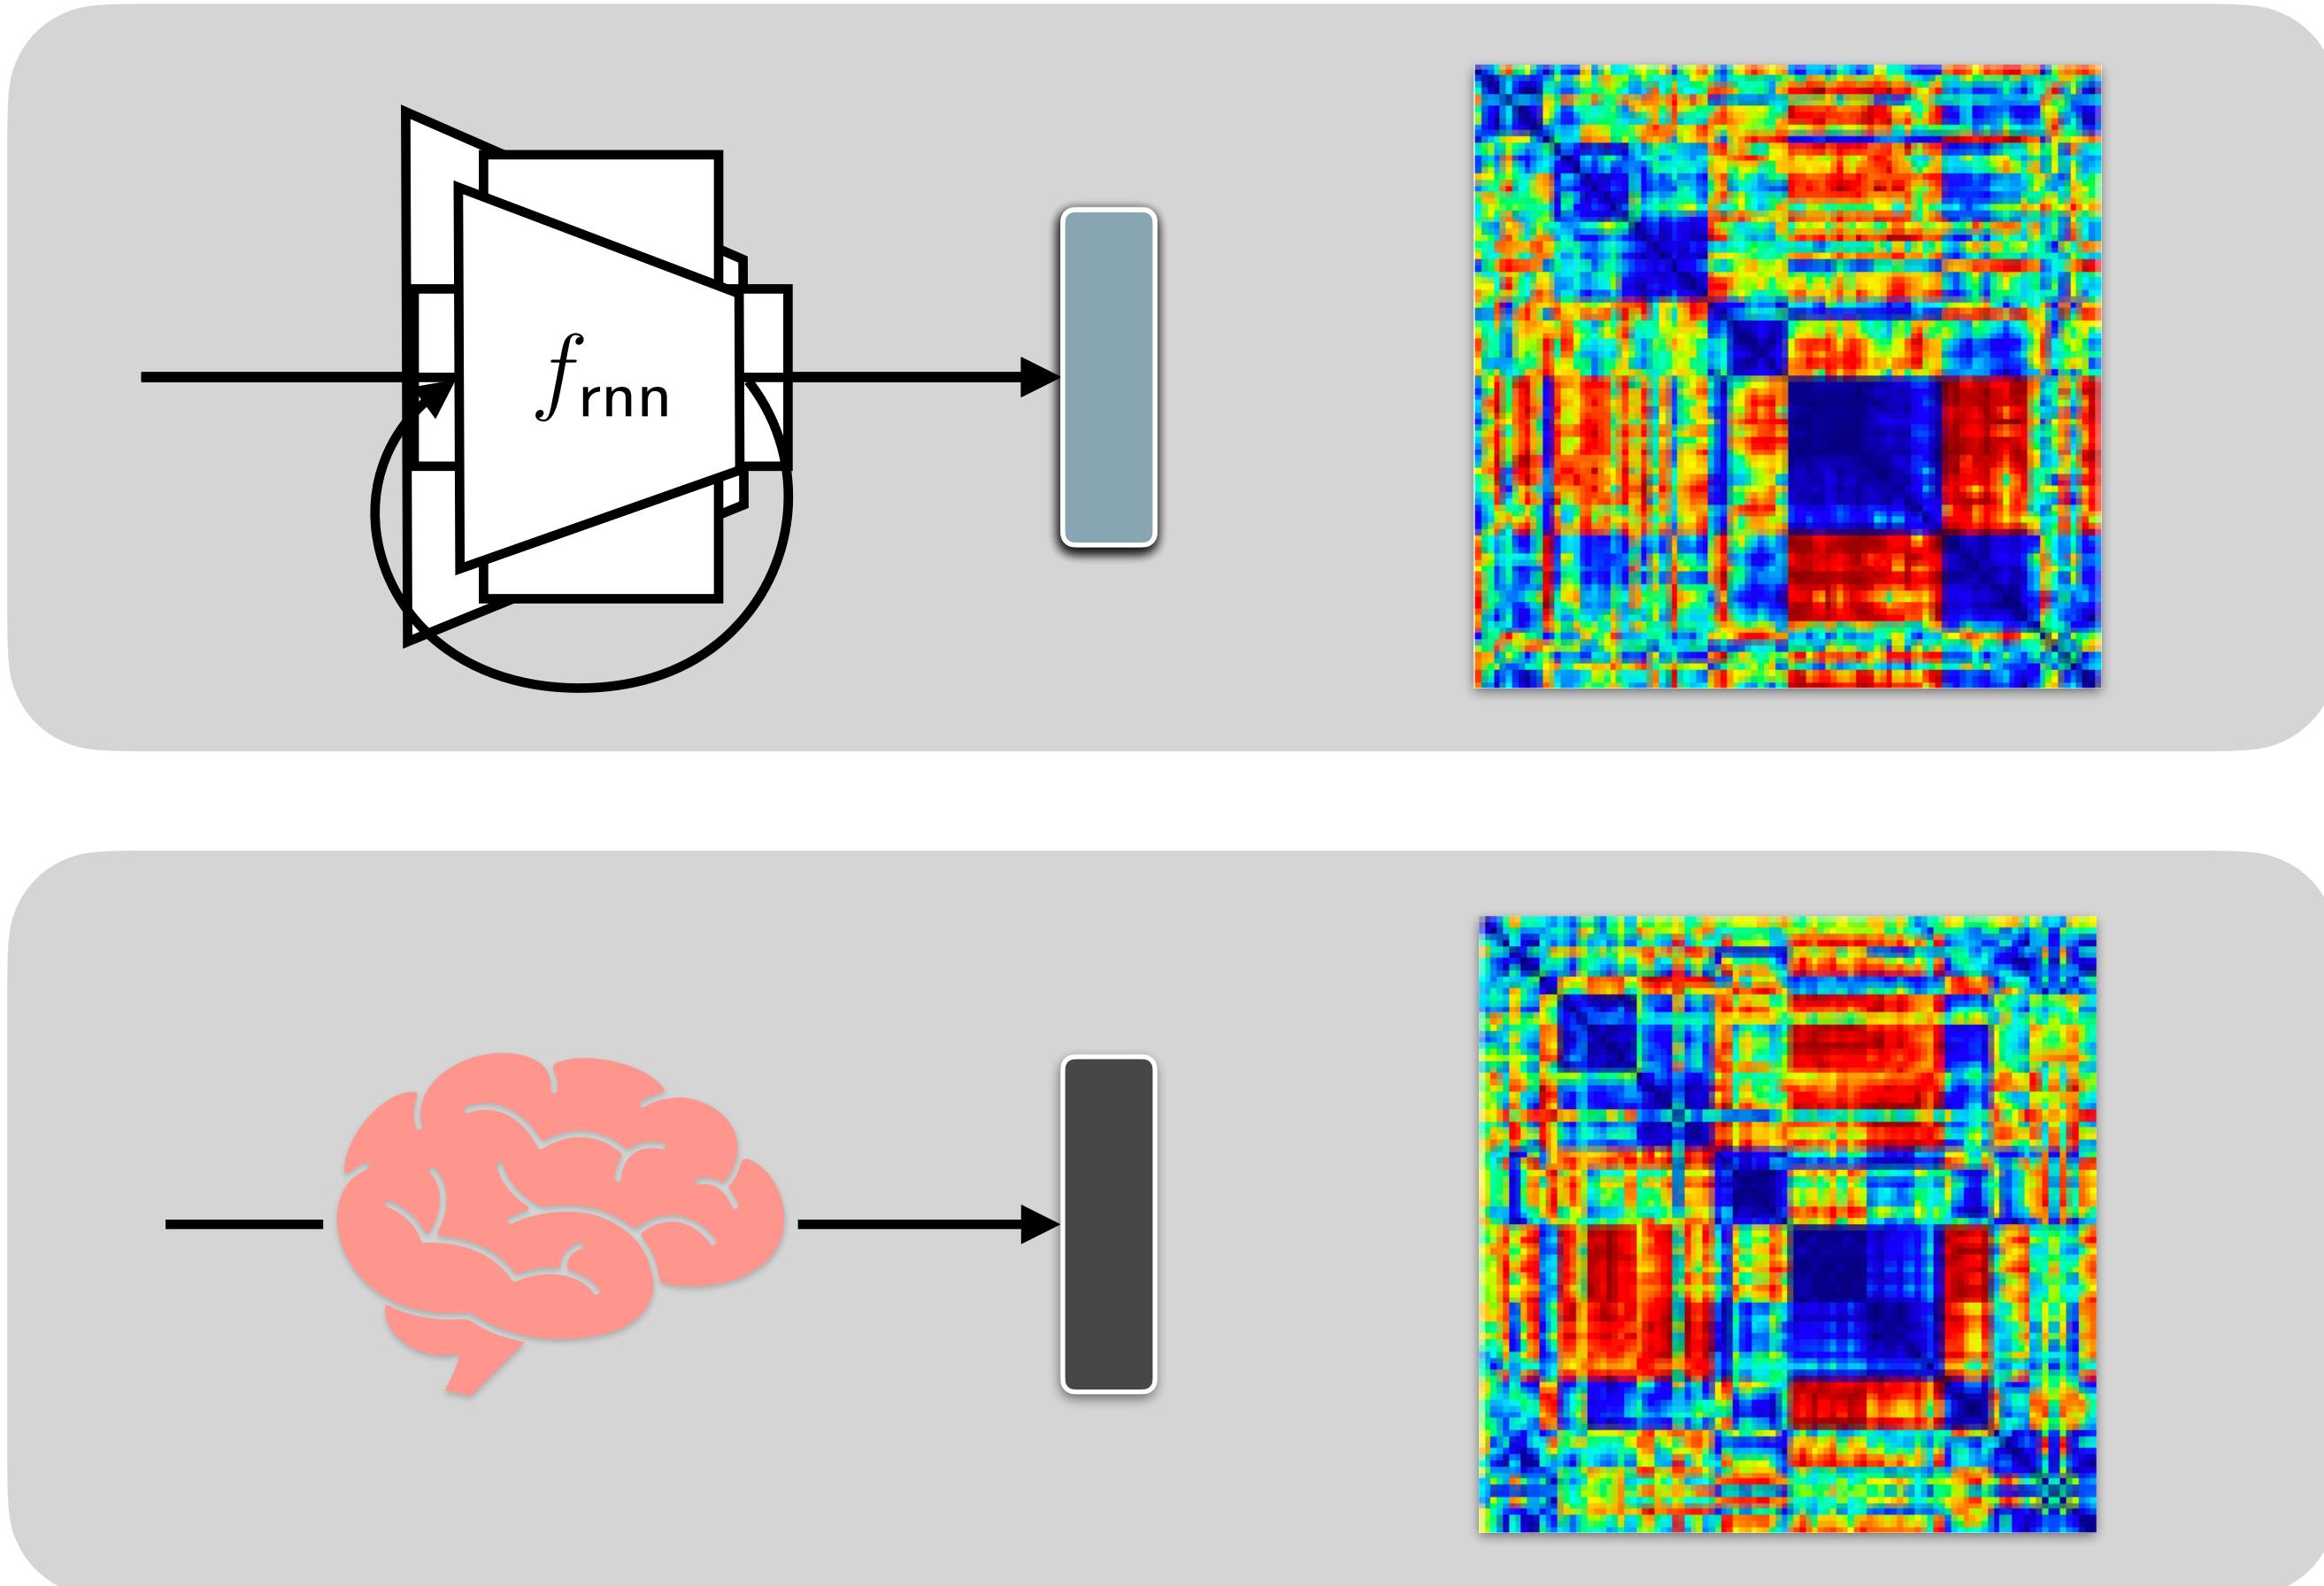
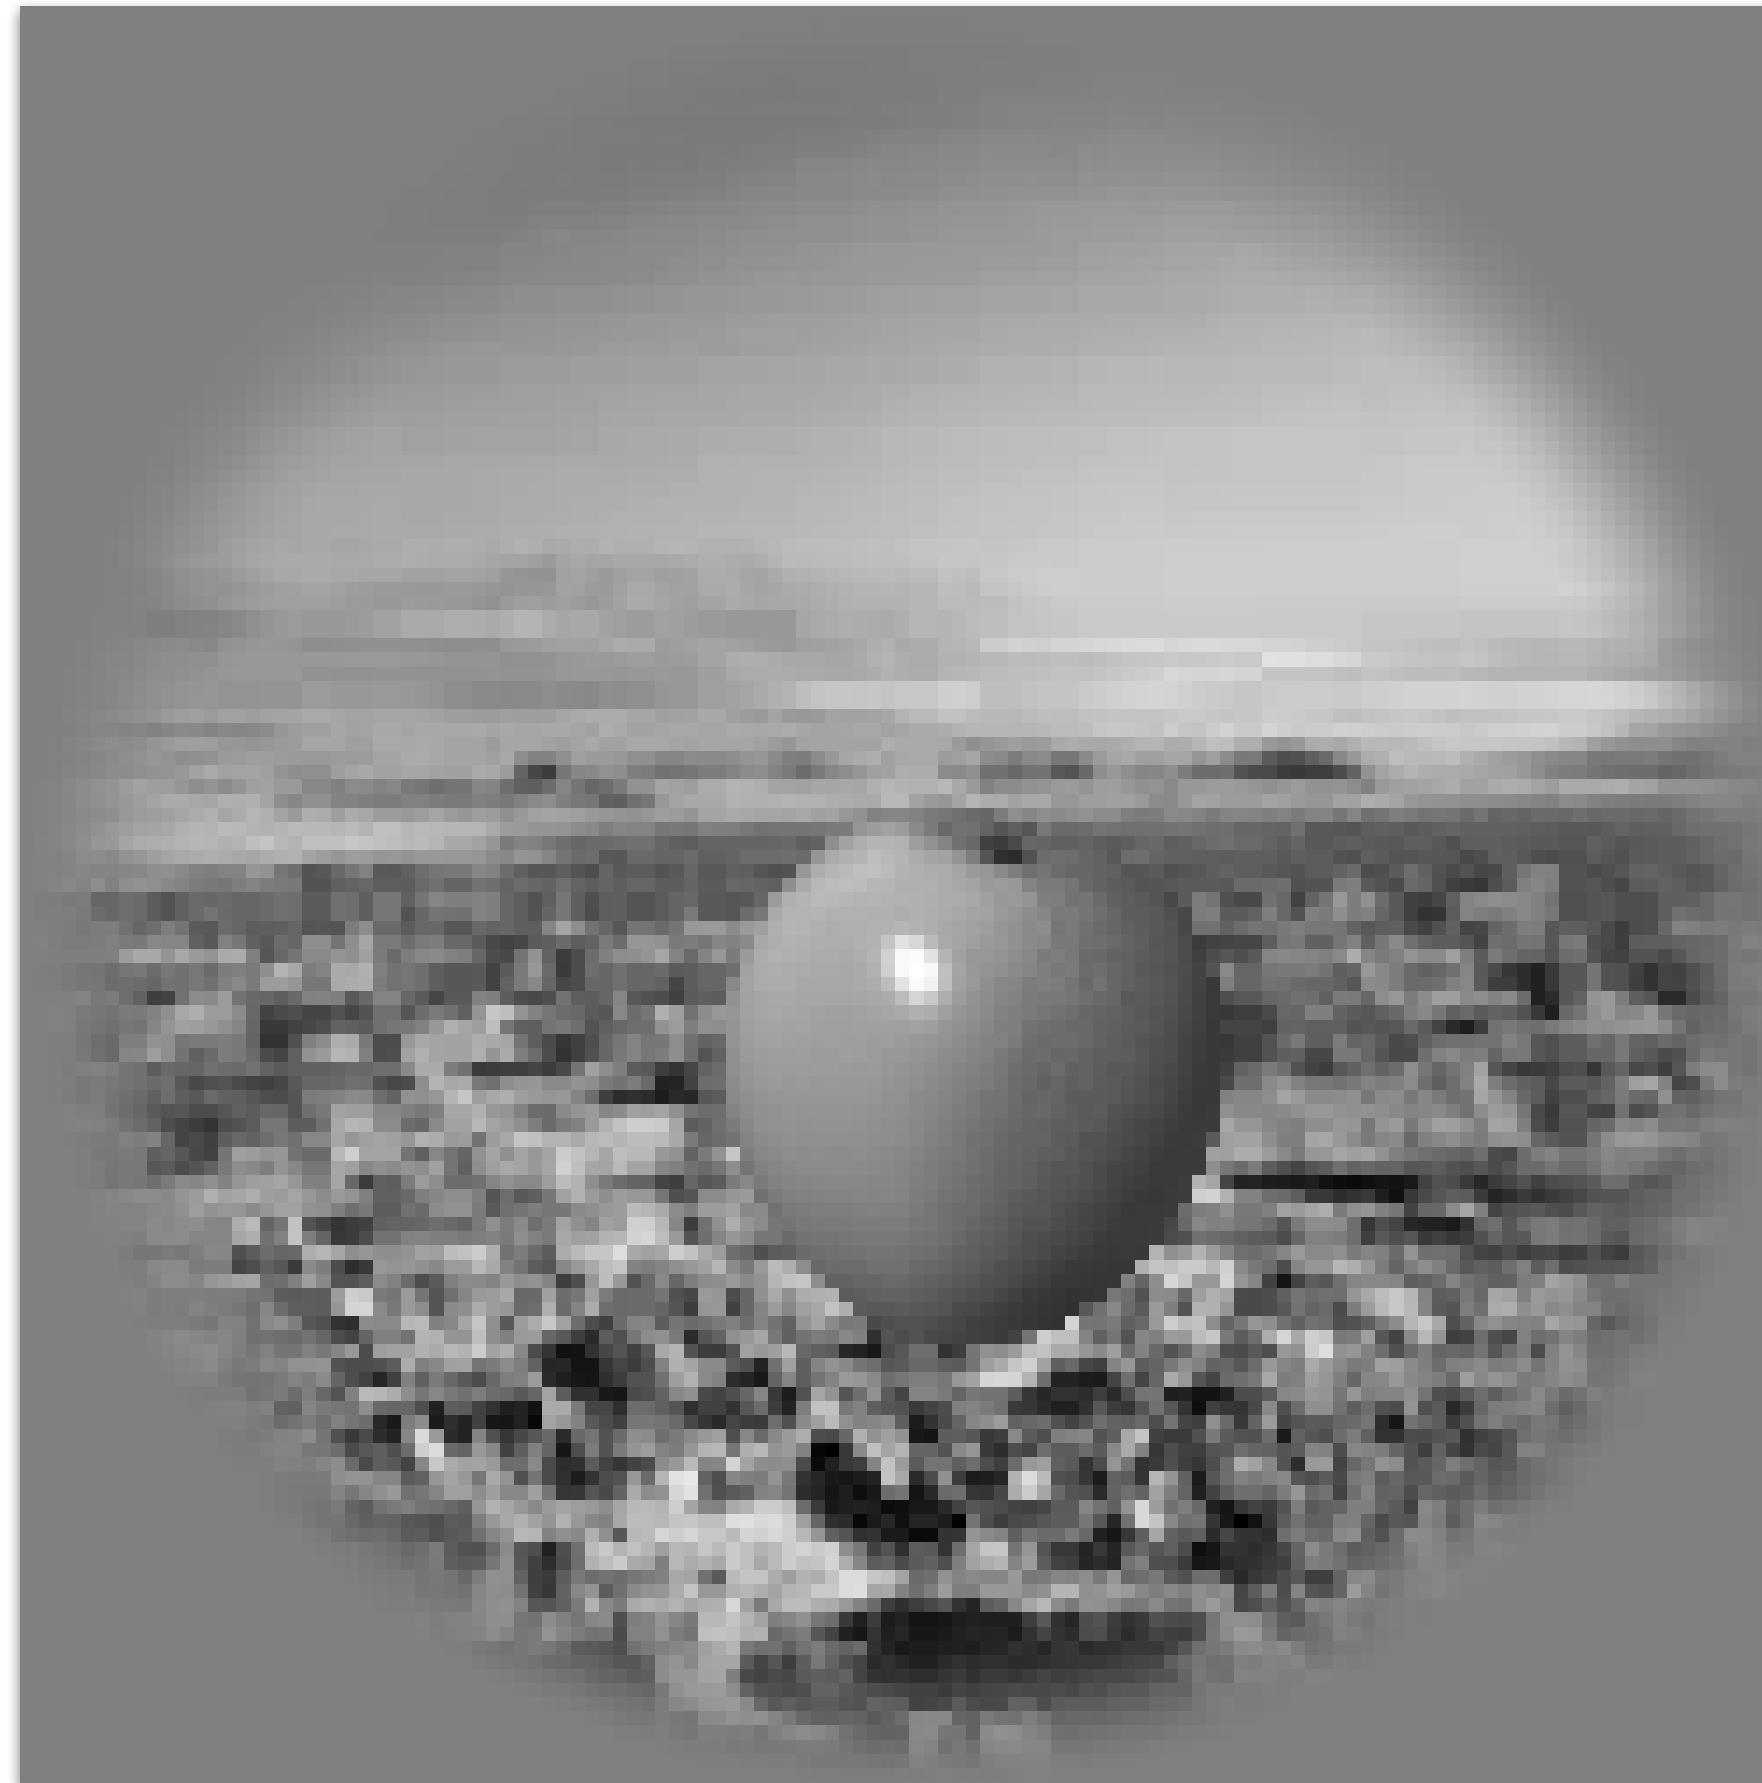
# Representational similarities between Brains and Machines



[Han, Poggio, **Cheung** 2023]

[Yamins, Hong, Cadieu, Solomon, Seibert, DiCarlo 2014]

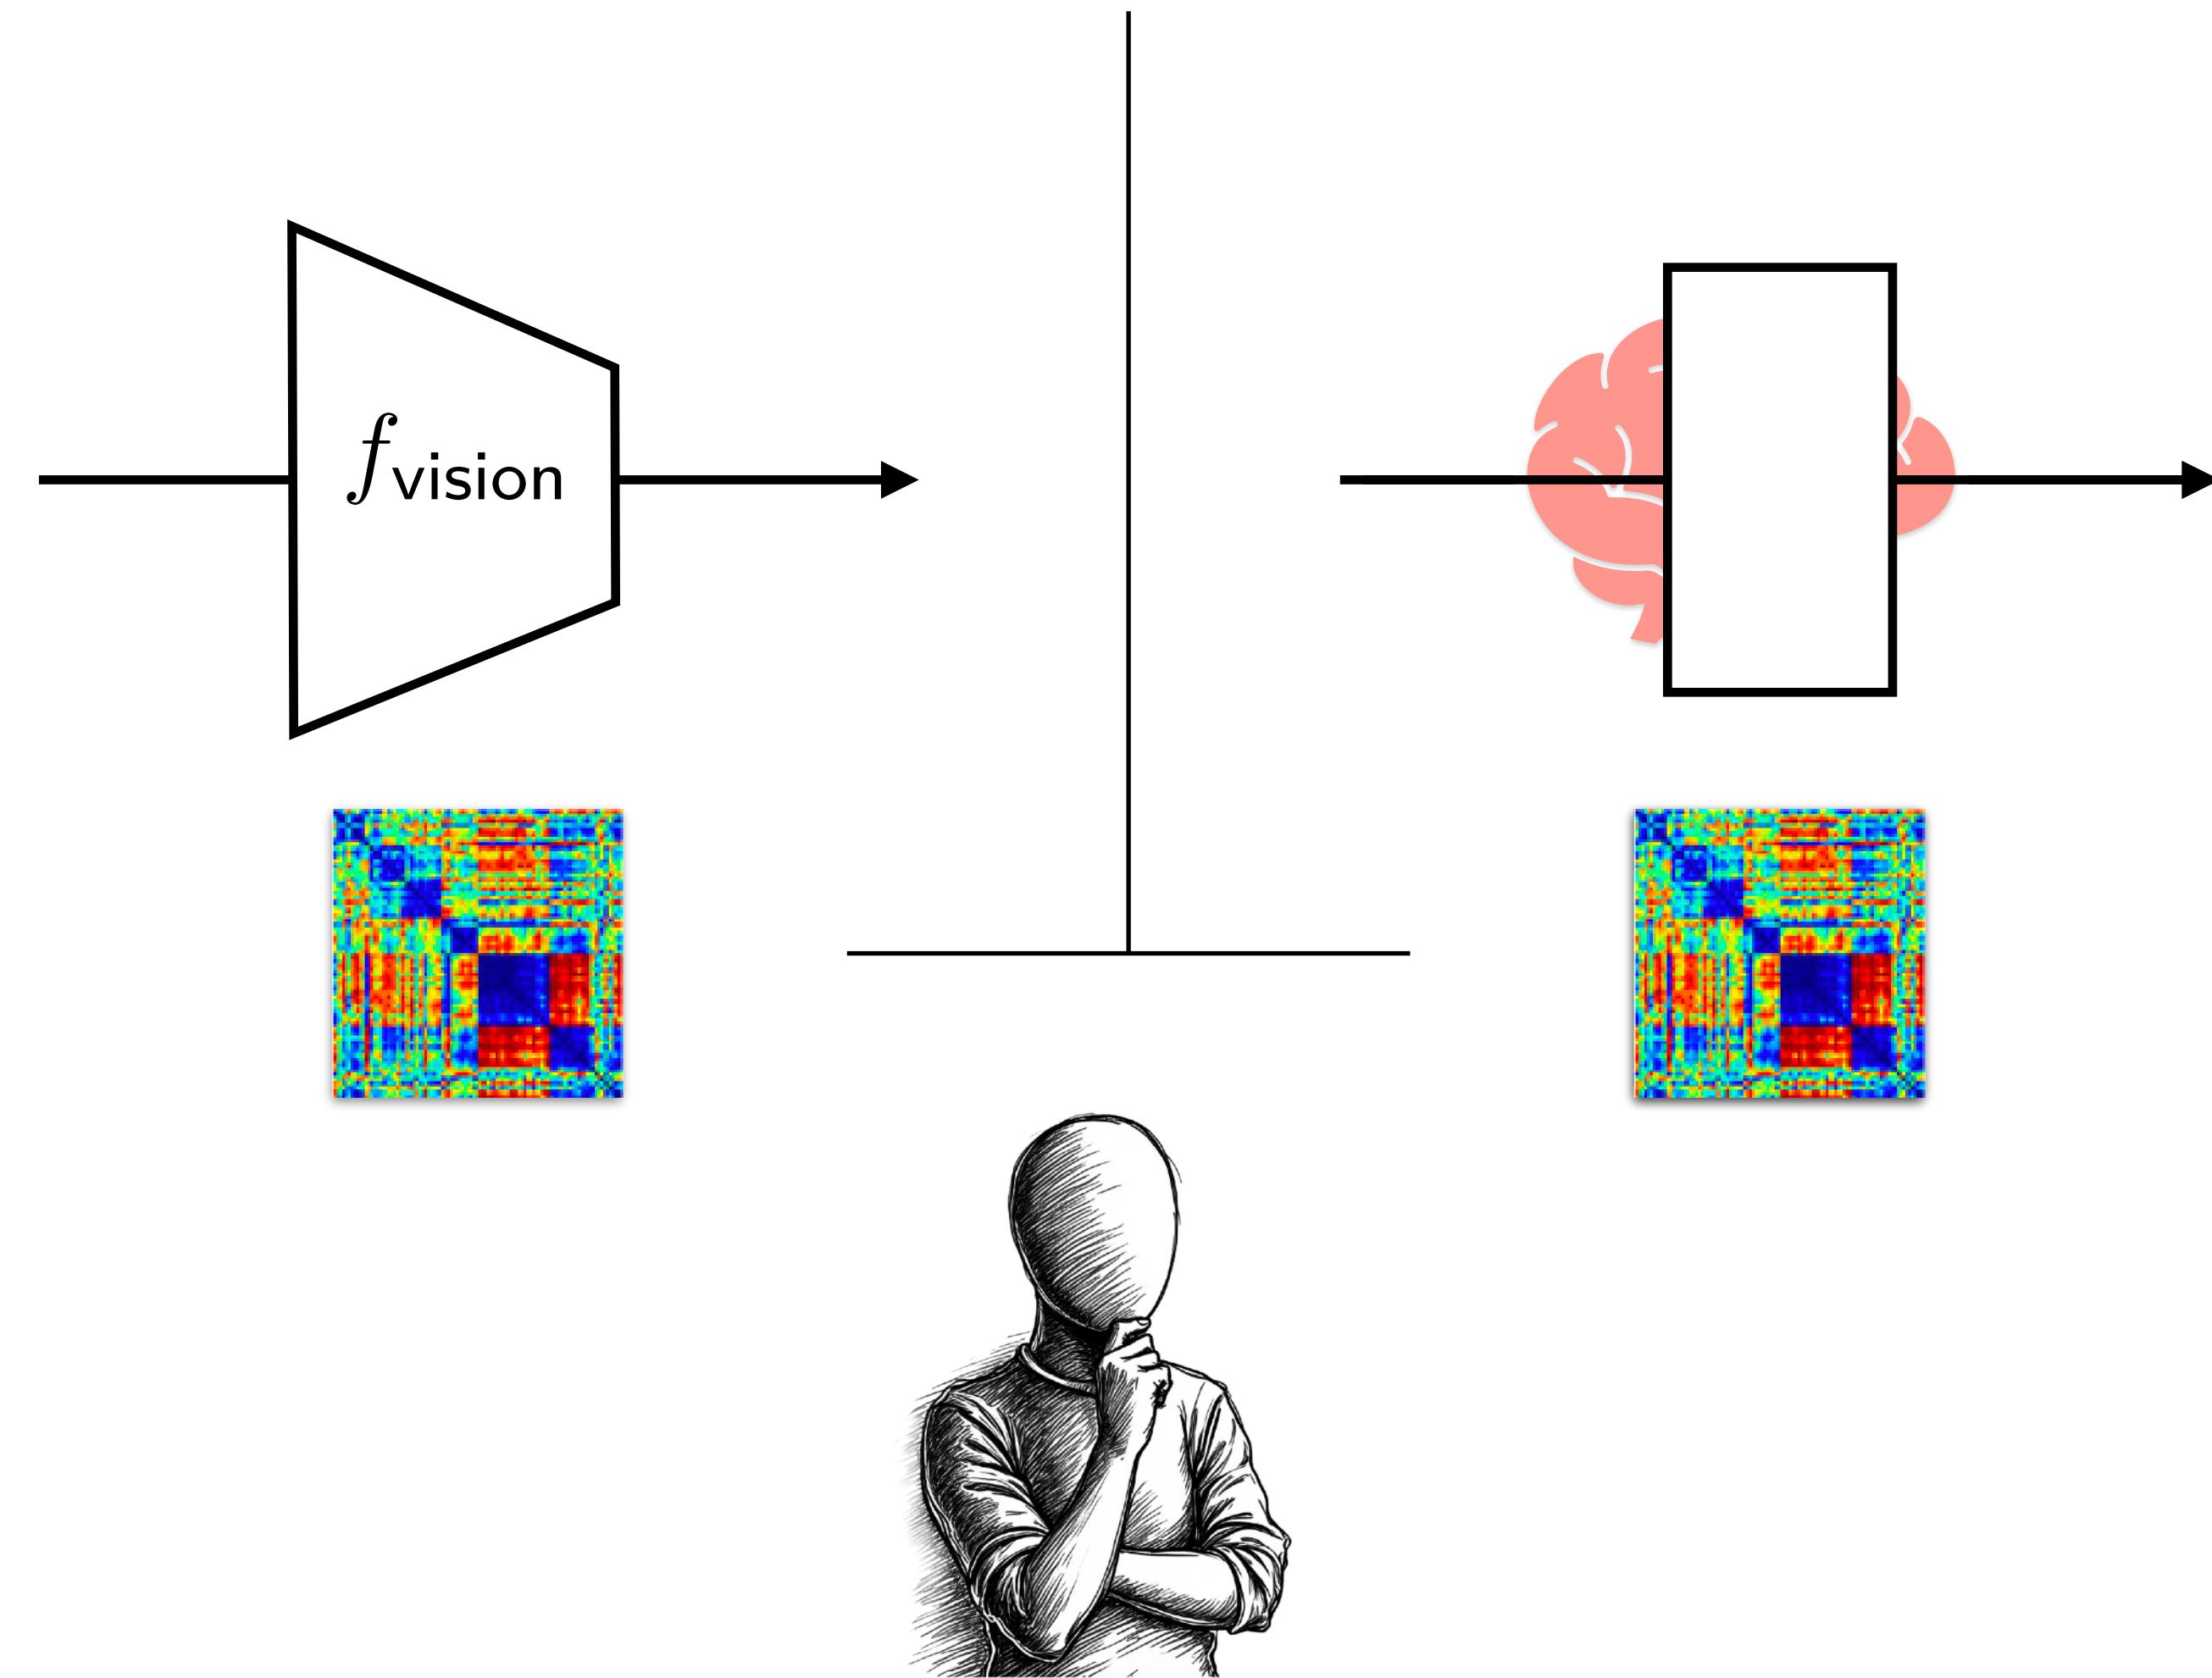
# Representational similarities between Brains and Machines



[Han, Poggio, **Cheung** 2023]

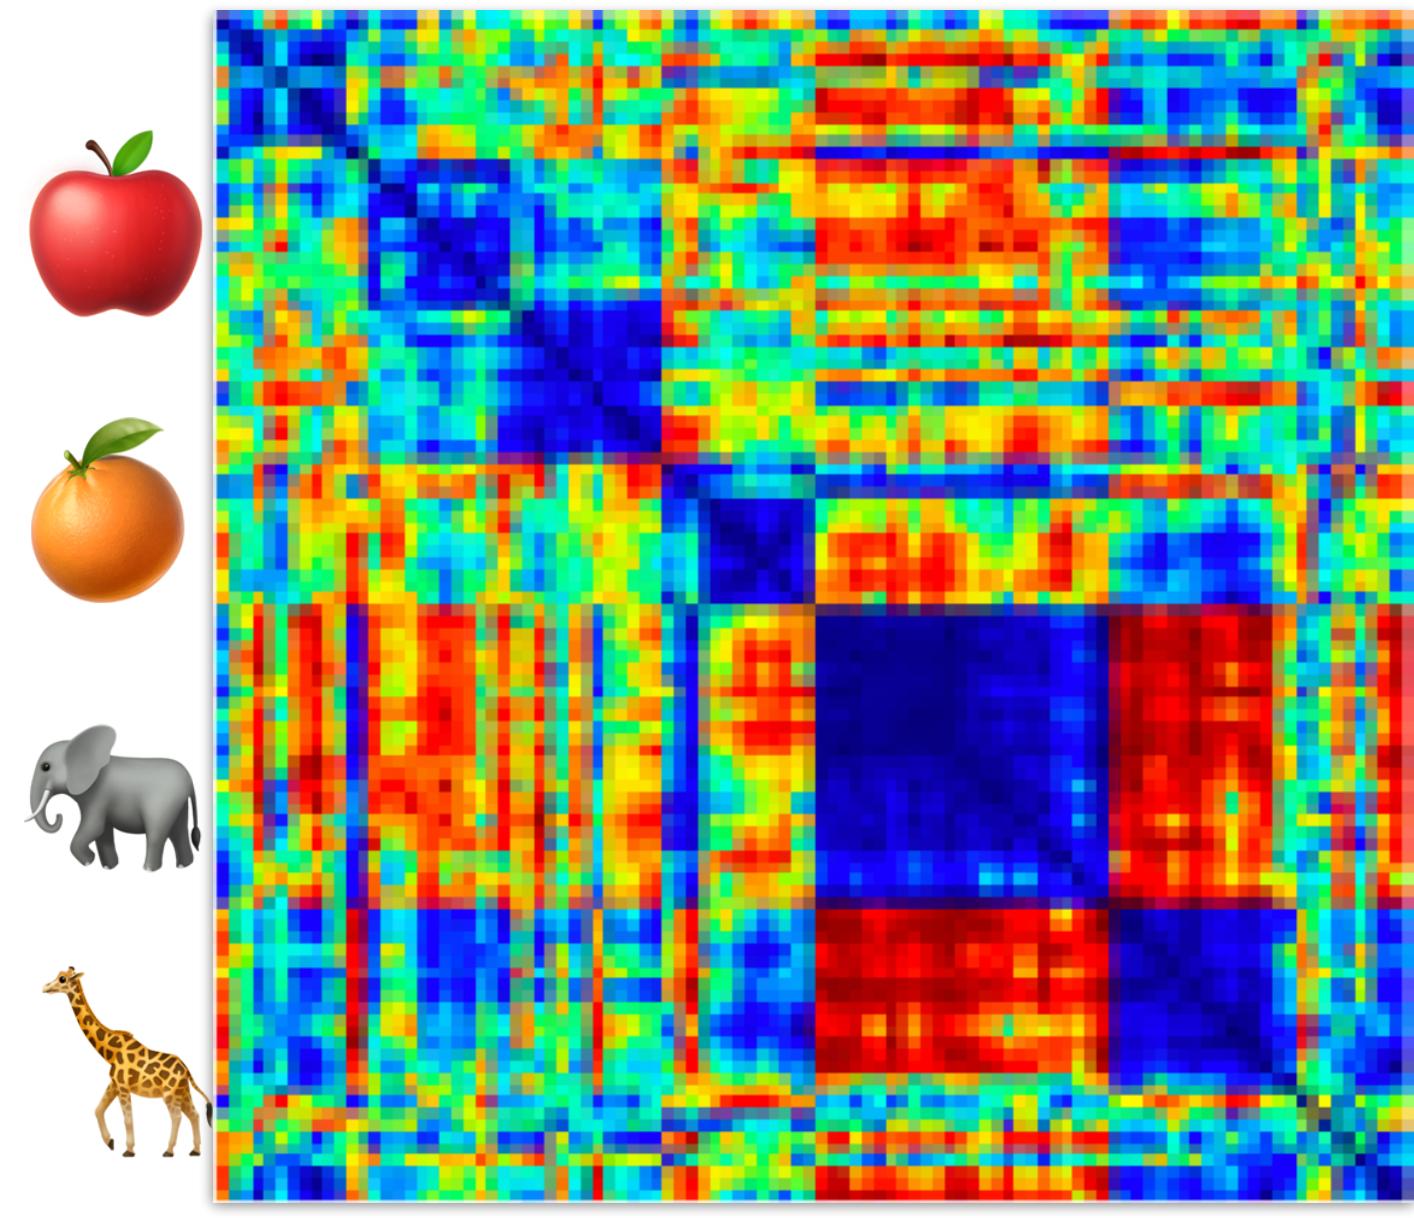
[Yamins, Hong, Cadieu, Solomon, Seibert, DiCarlo 2014]

# Representational Turing Test



# Characterizing representations using kernels

$K_{\text{vision}}$



Restrict our attention to **vector embeddings**

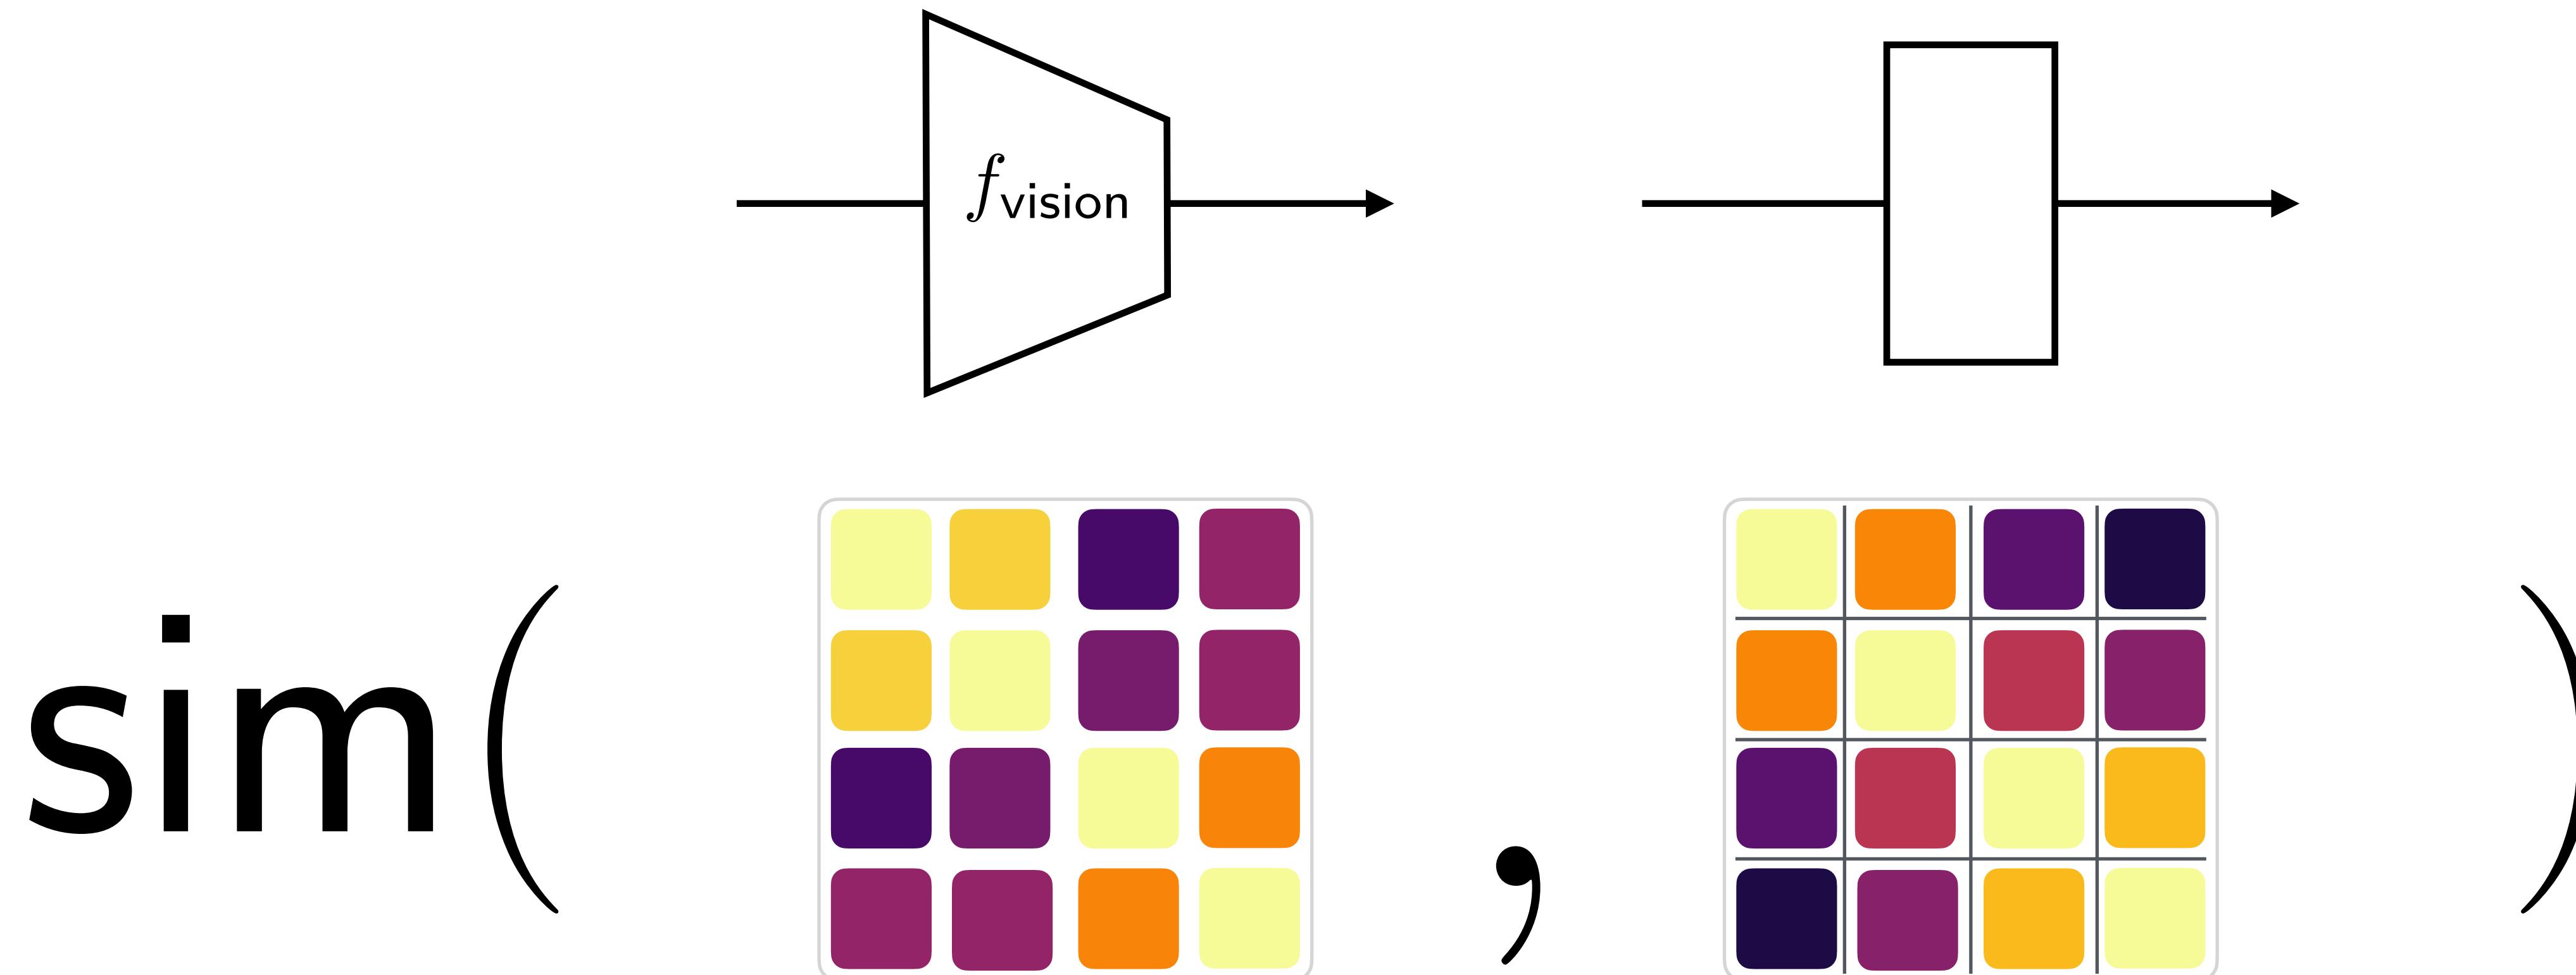
$$f: \mathcal{X} \rightarrow \mathbb{R}^n$$

Characterize a representation  
in terms of its **kernel**

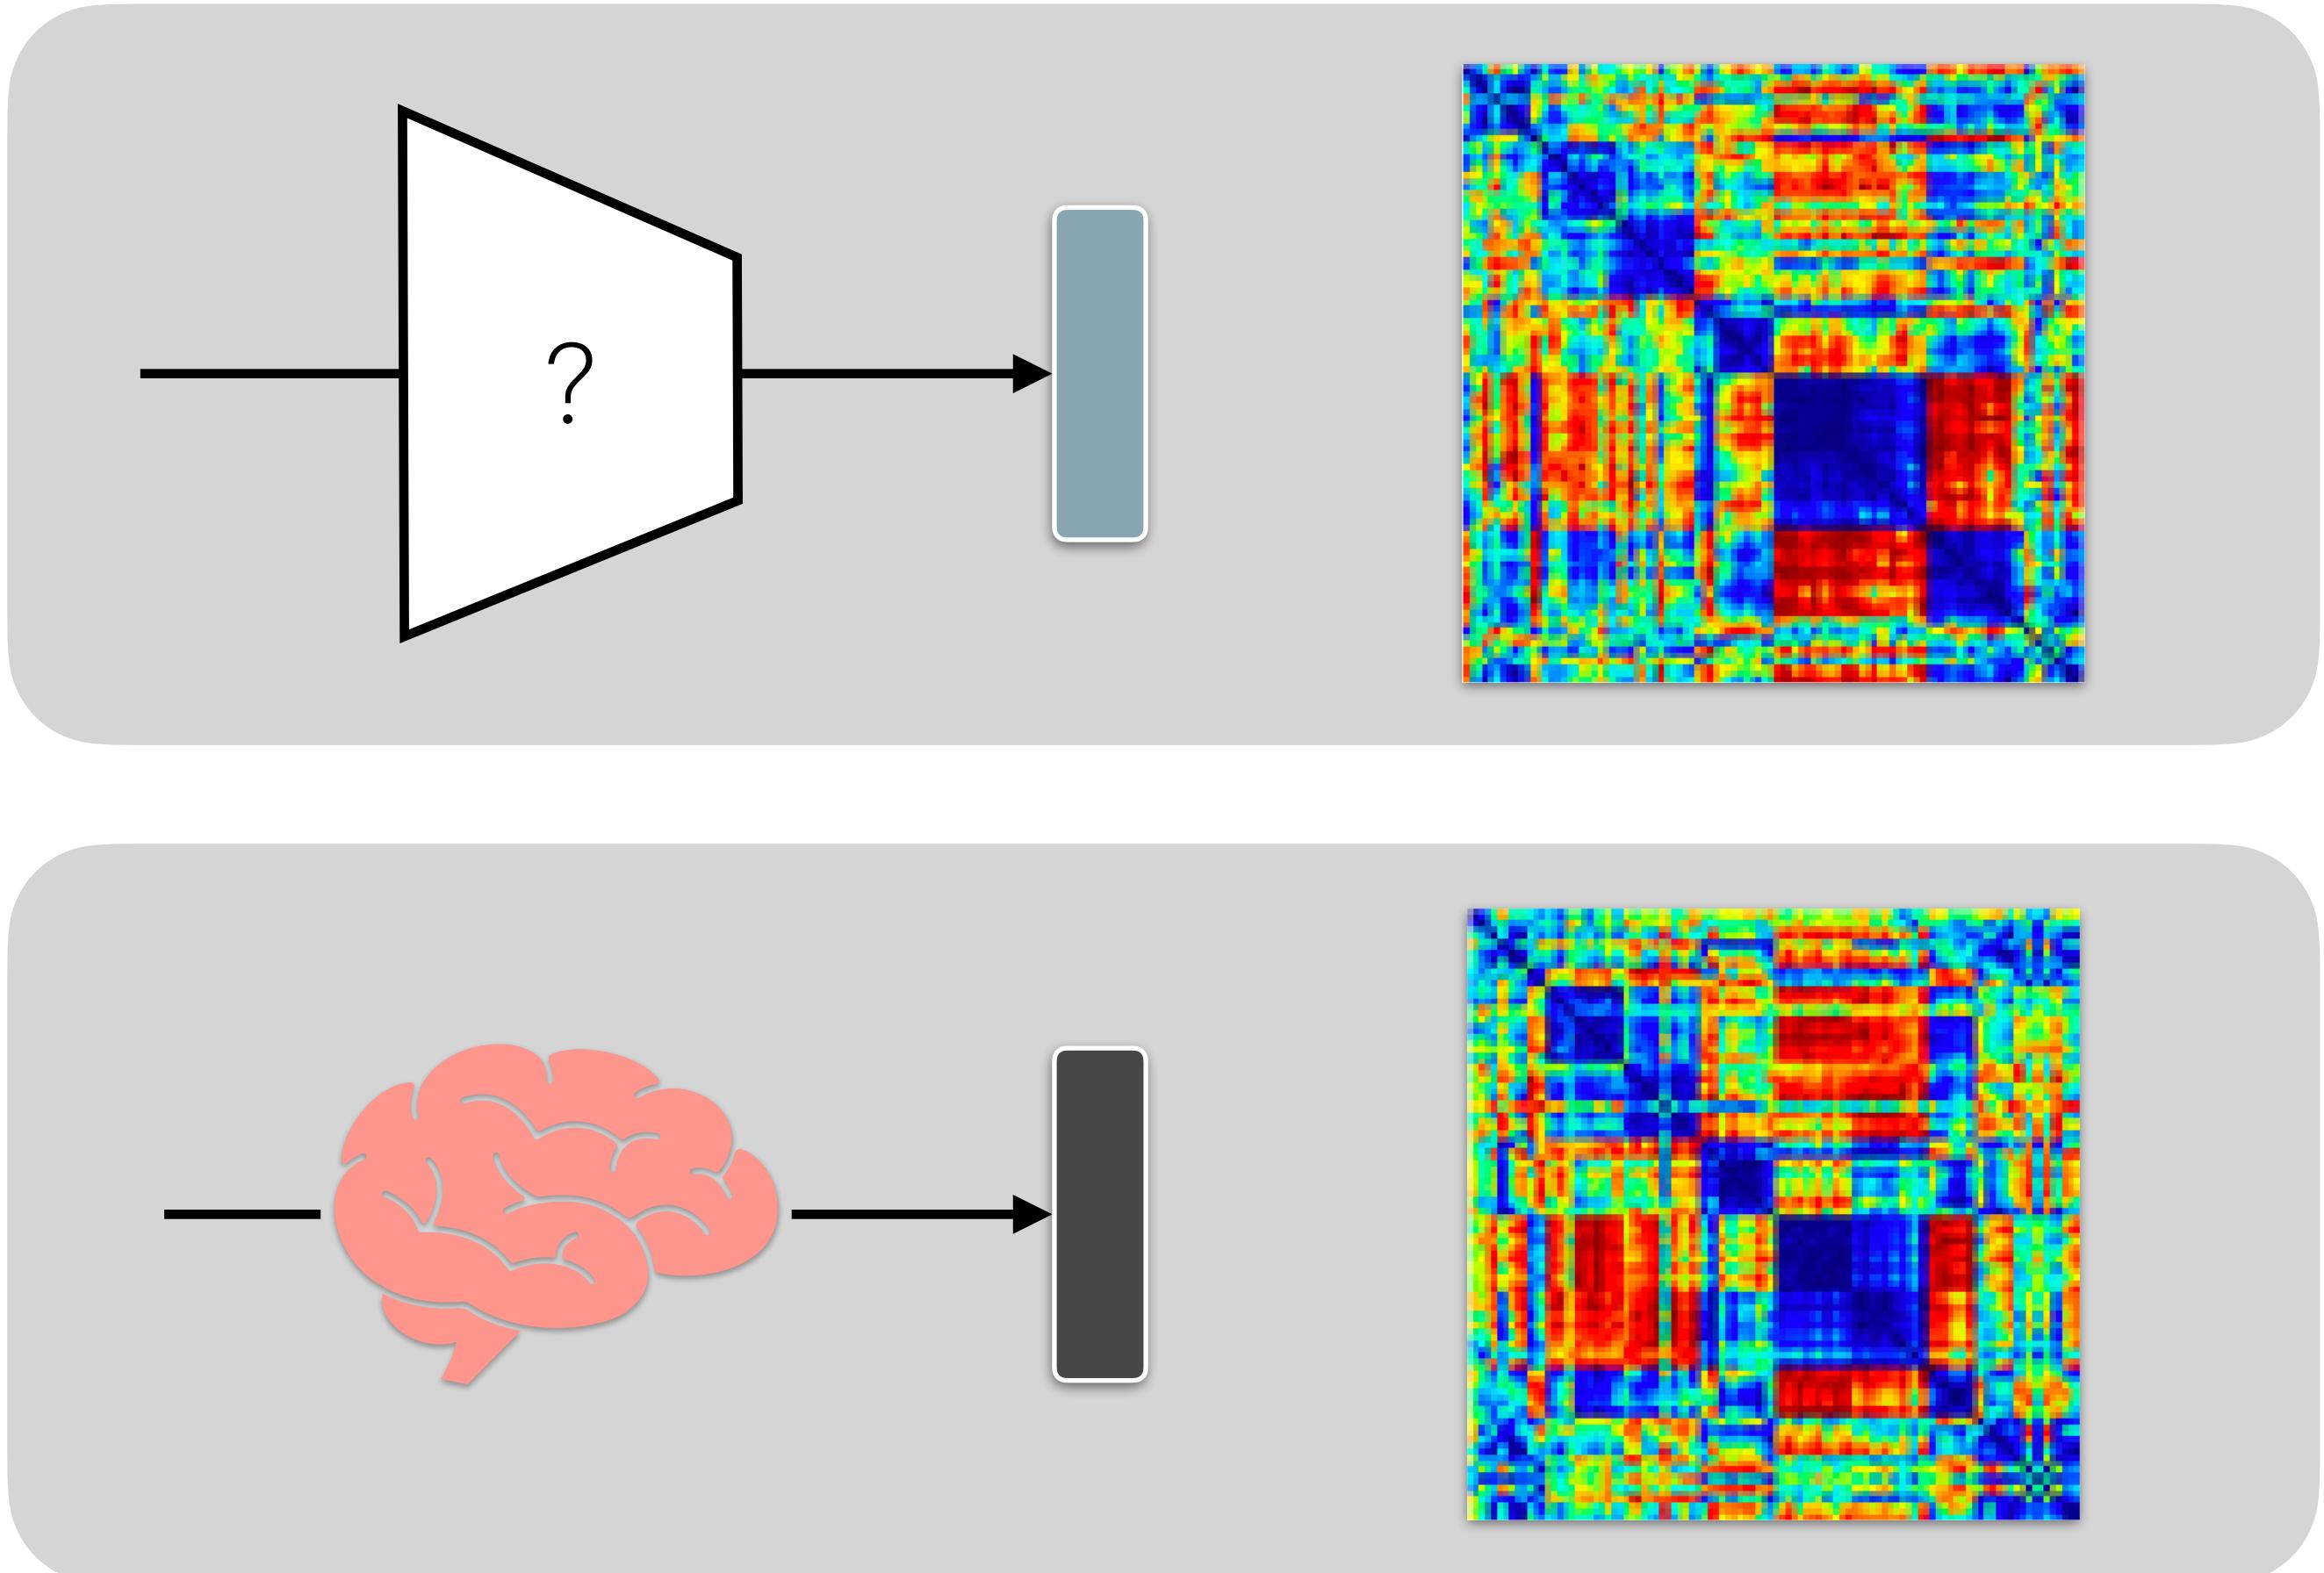
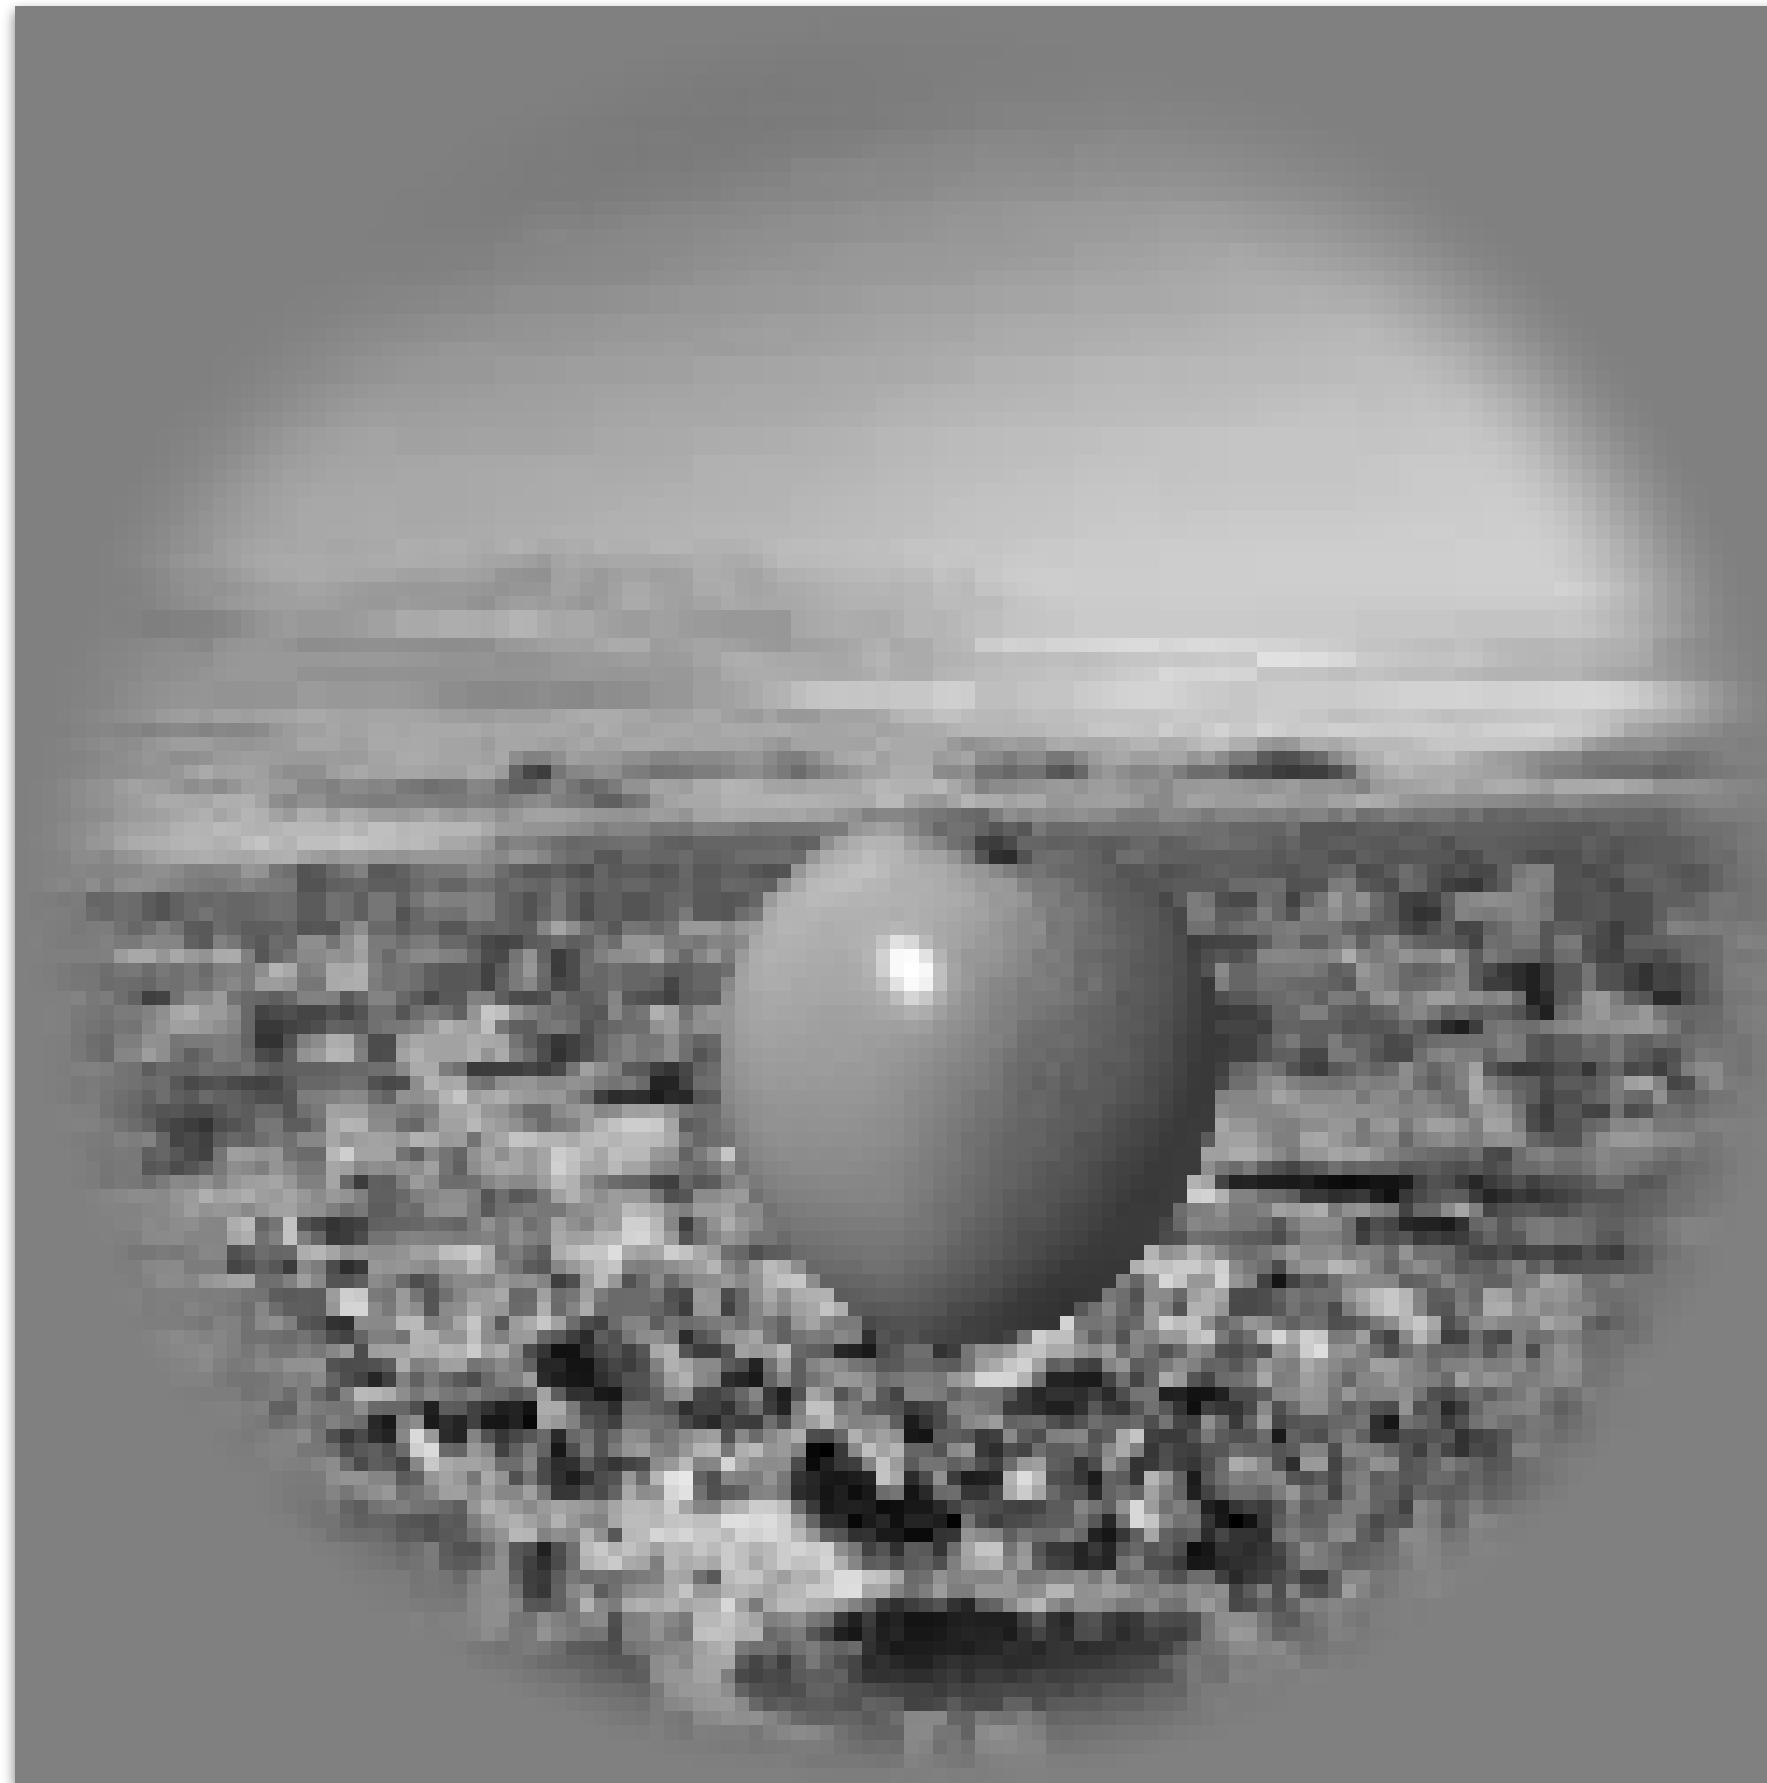
$$K: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$$

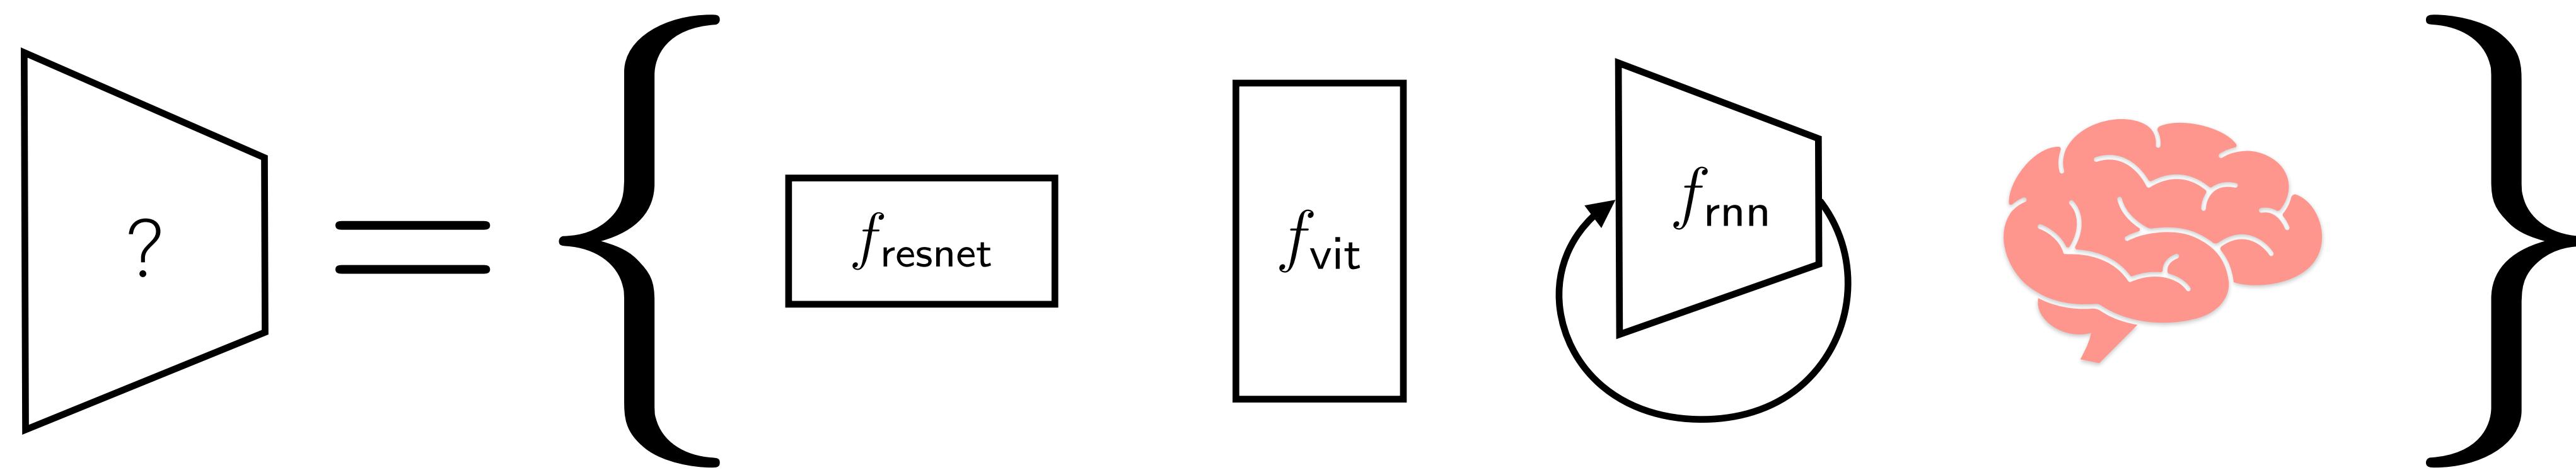
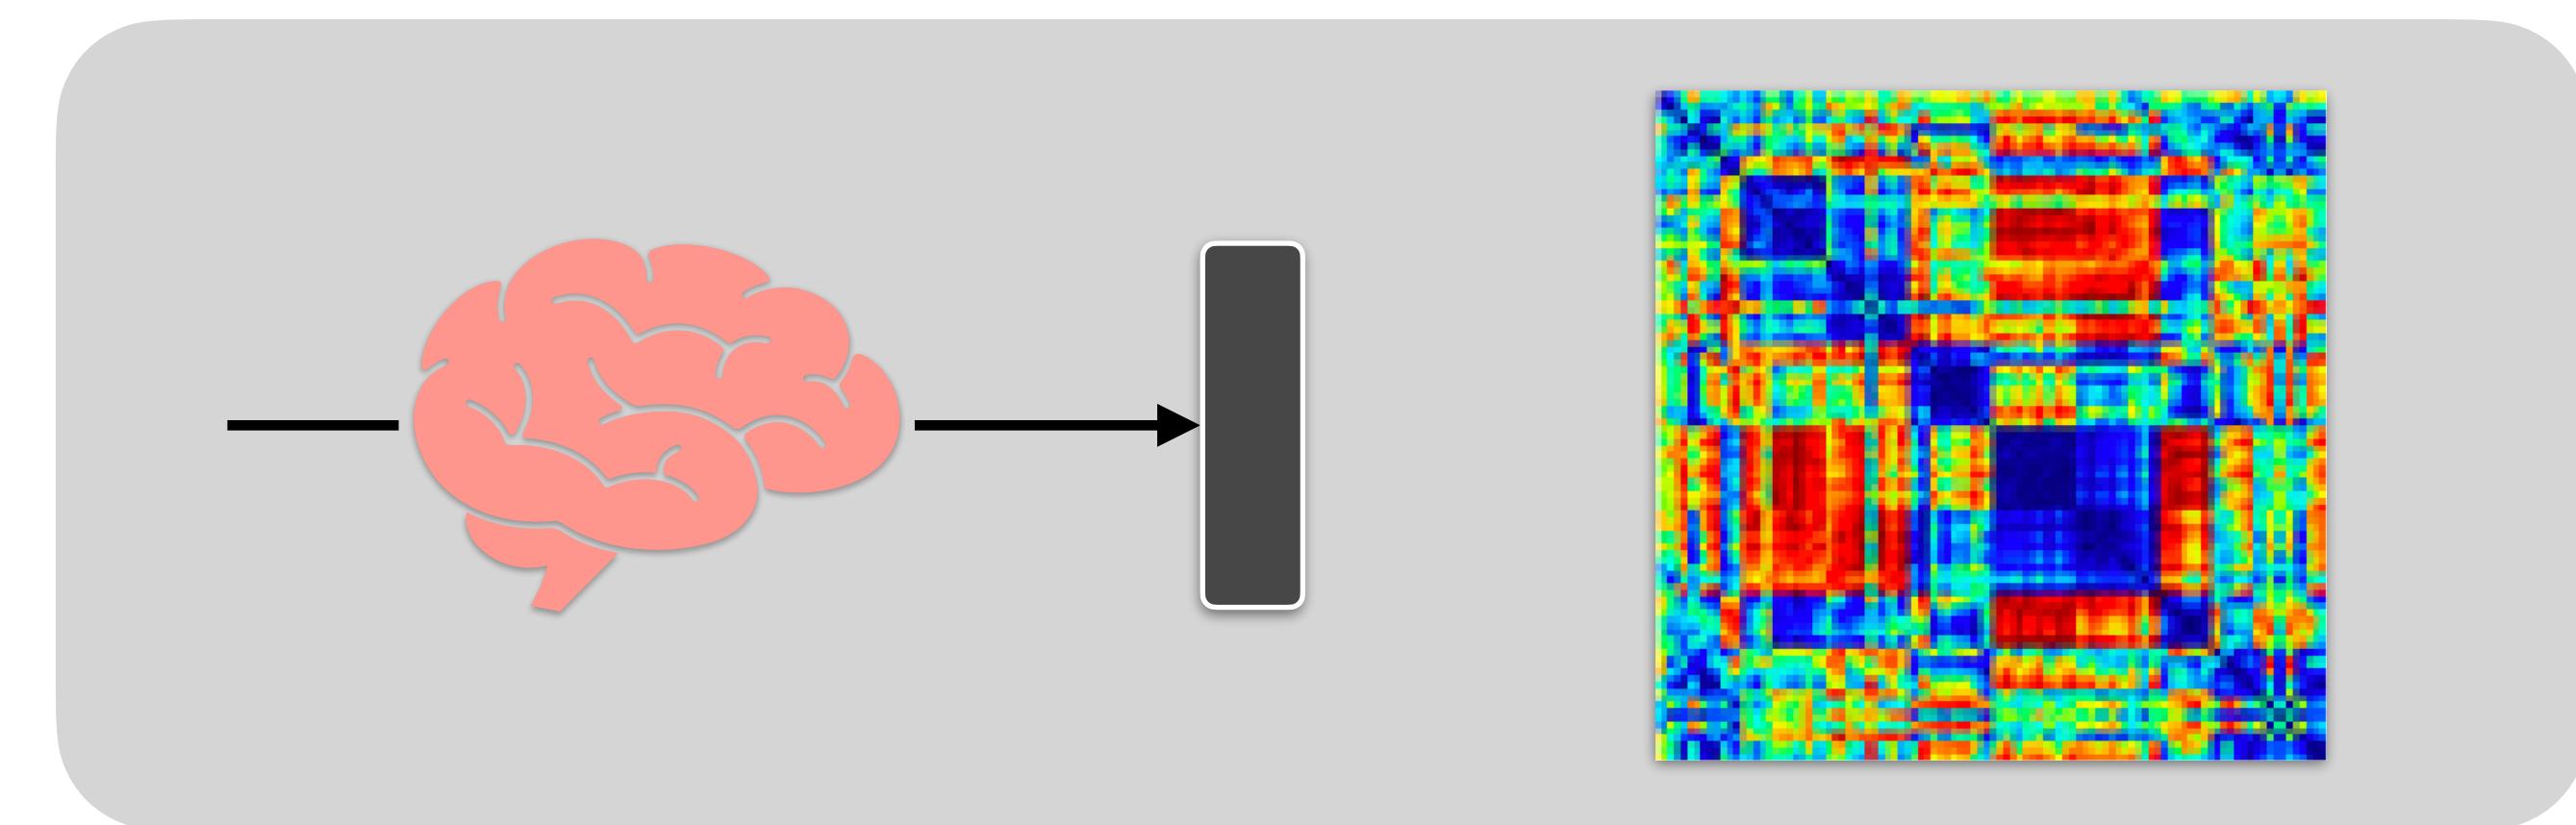
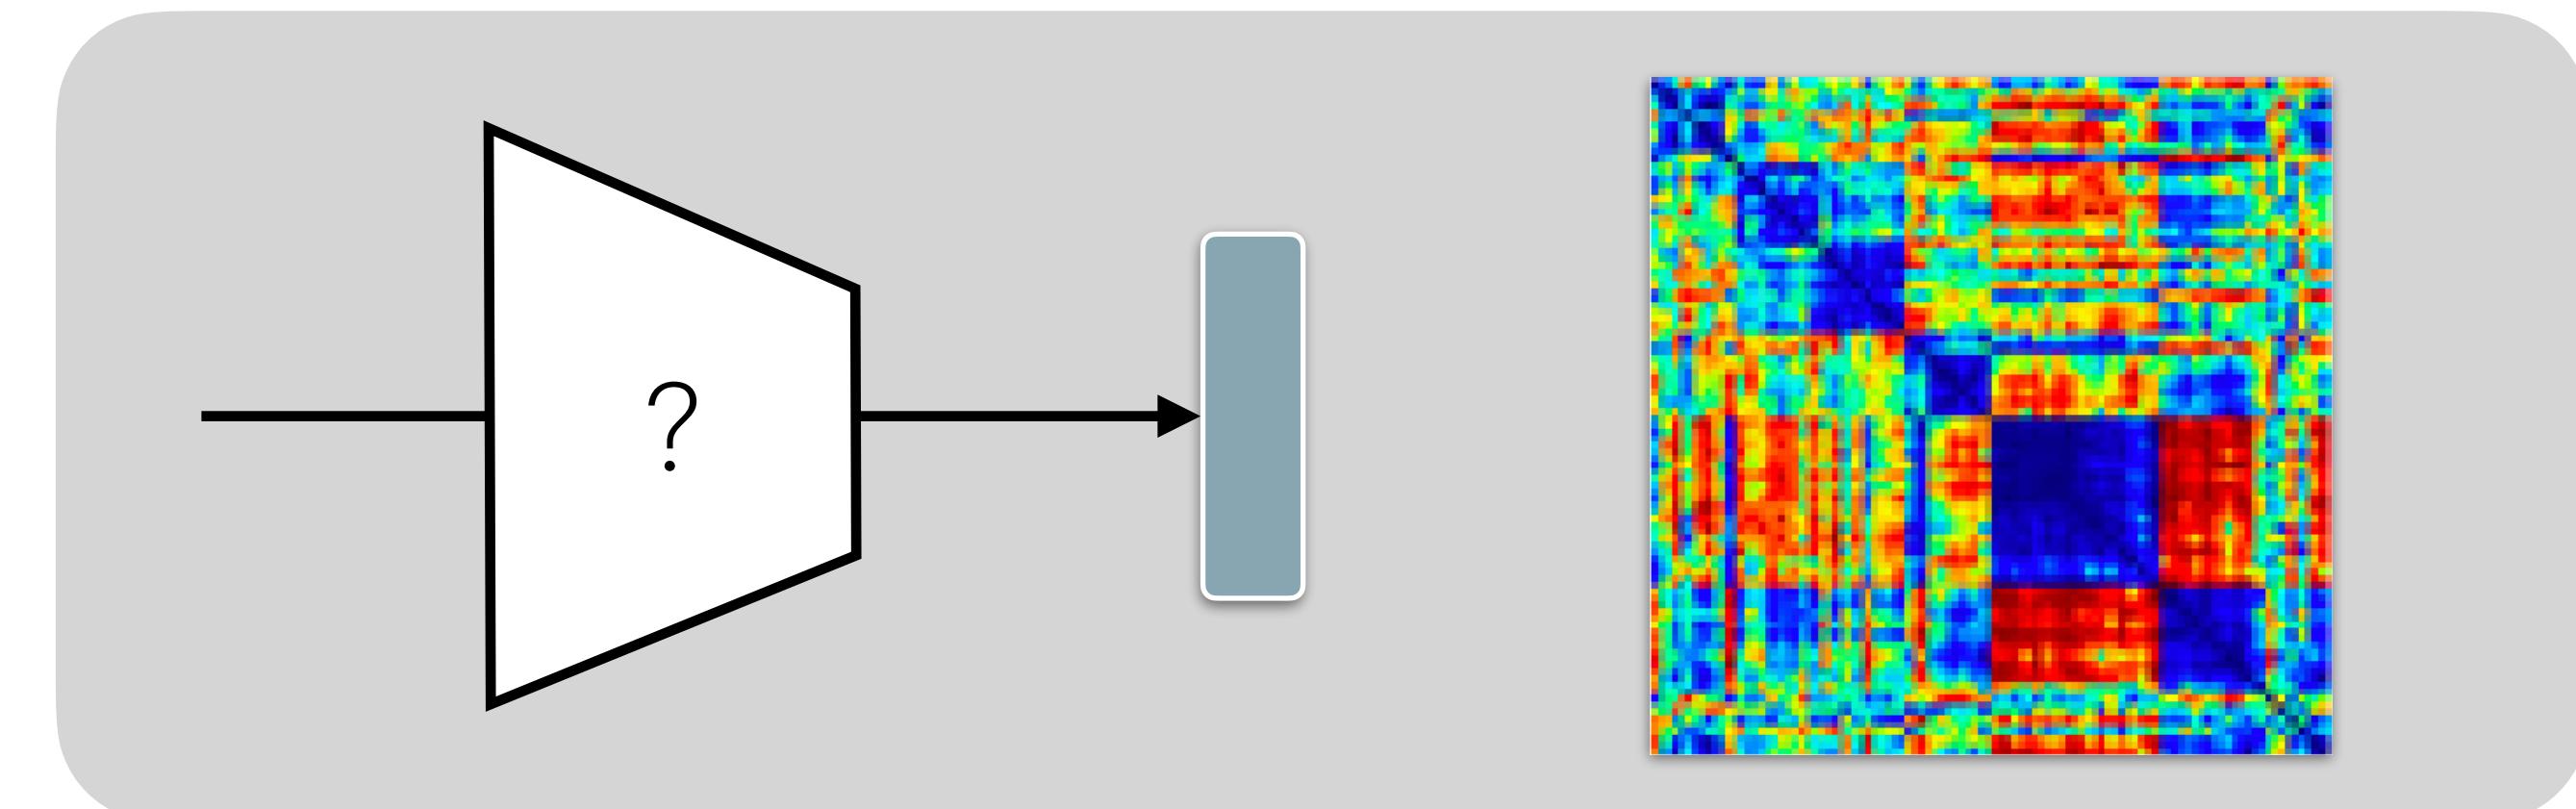
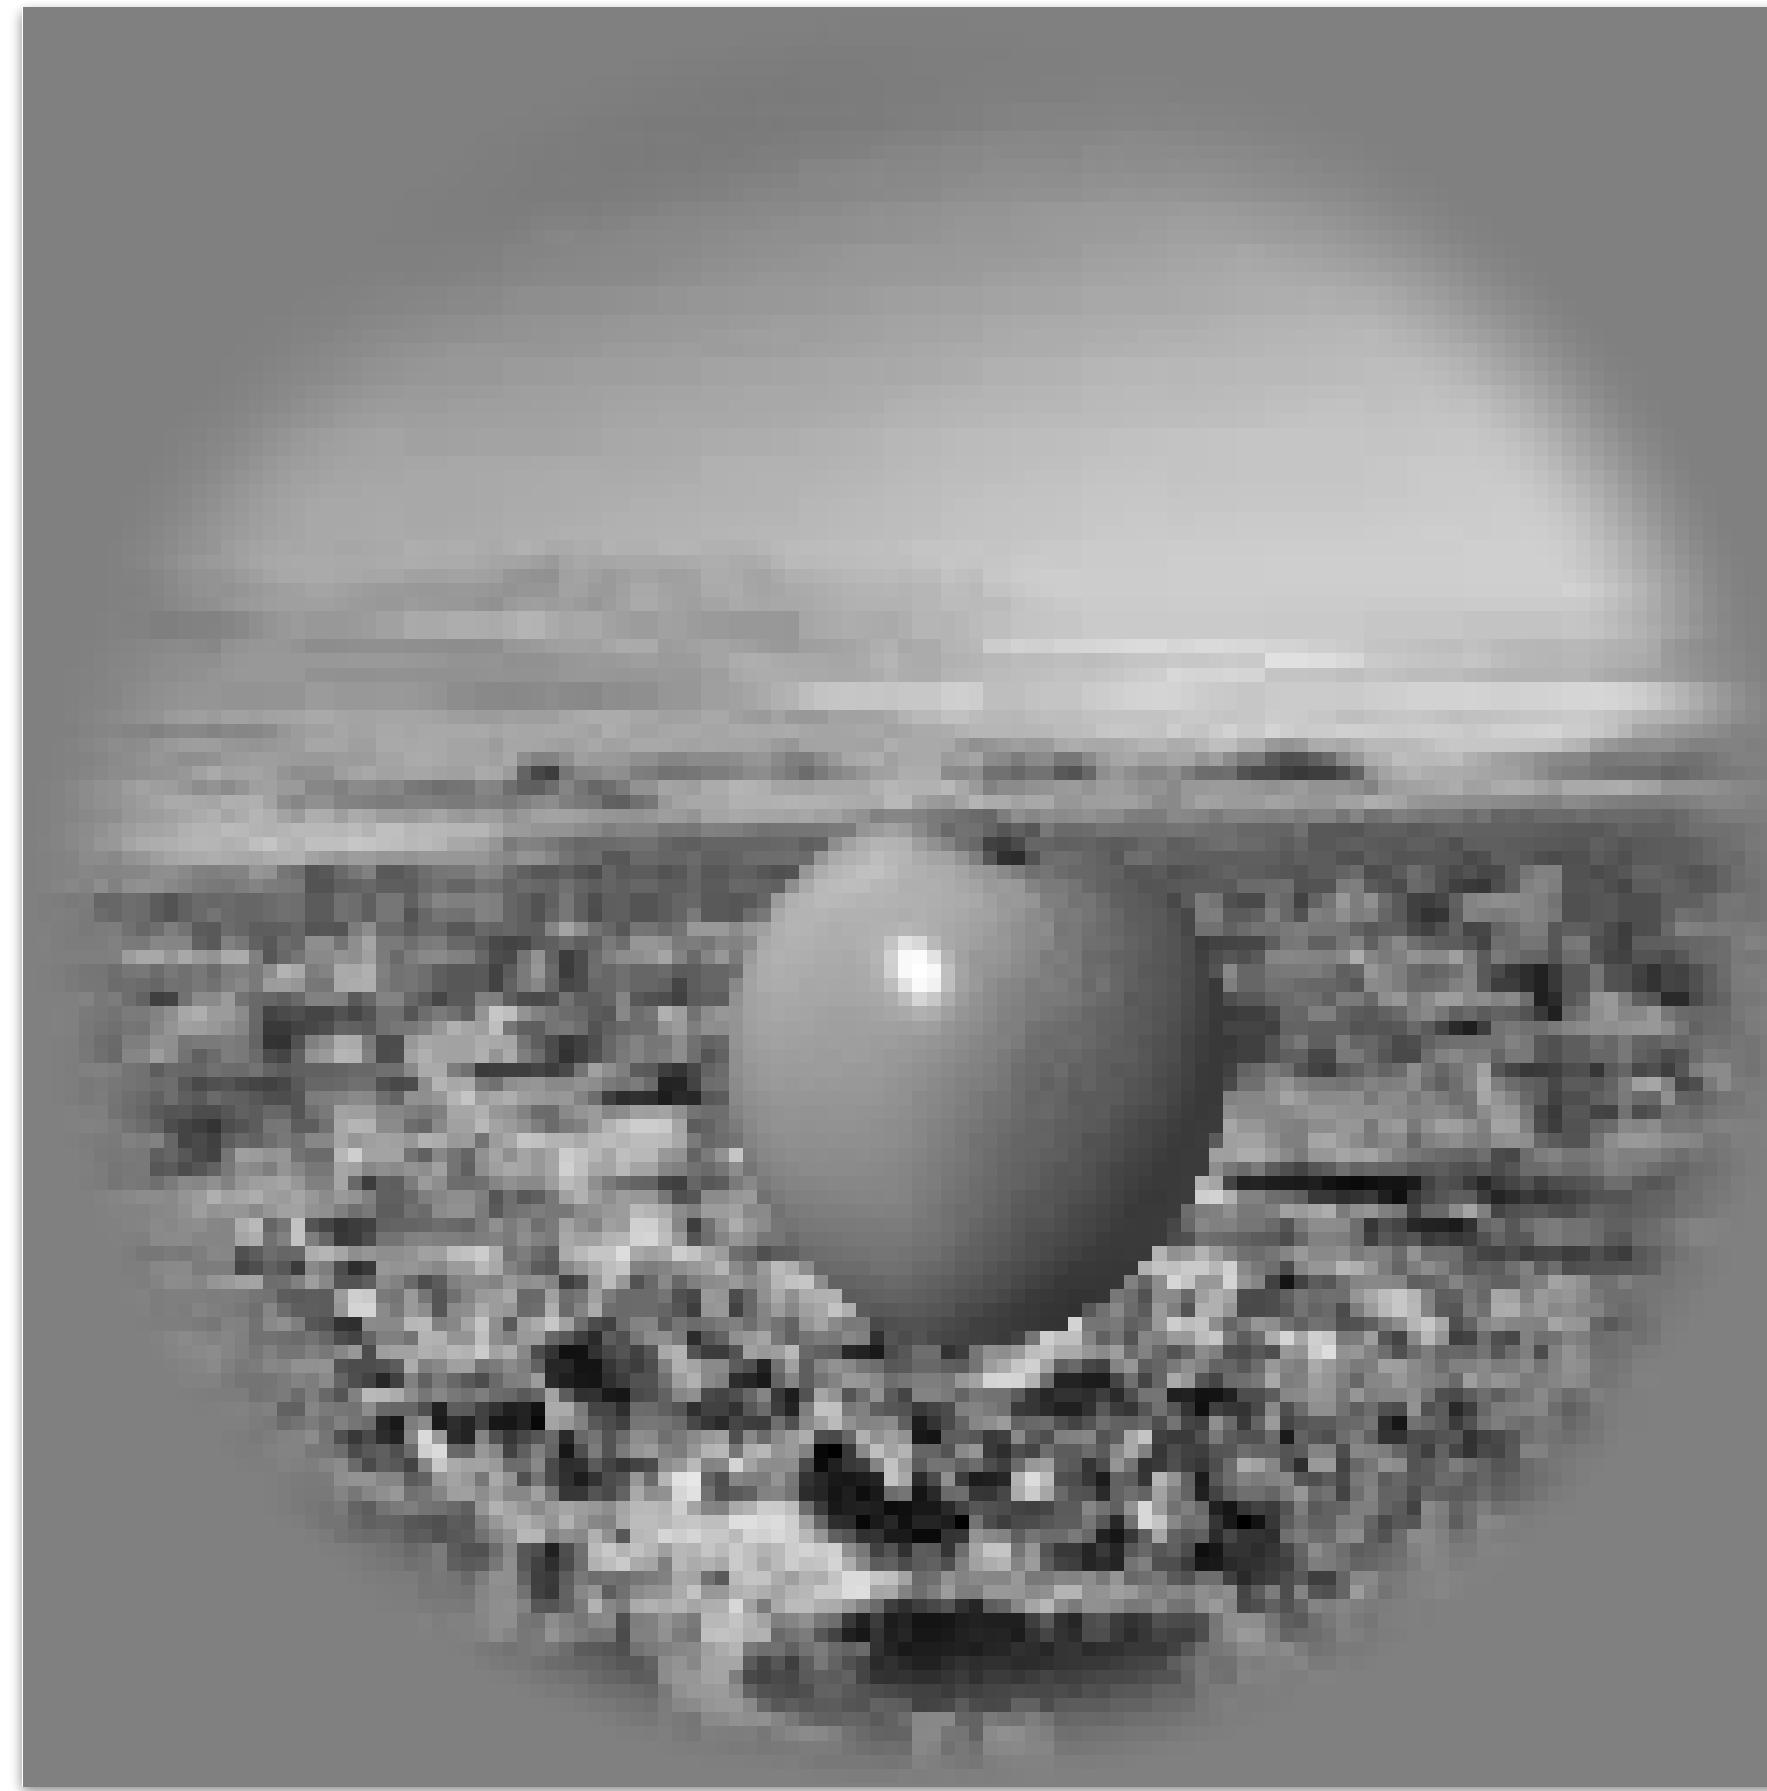
$$K(x_i, x_j) = \langle f(\text{apple}), f(\text{orange}) \rangle$$

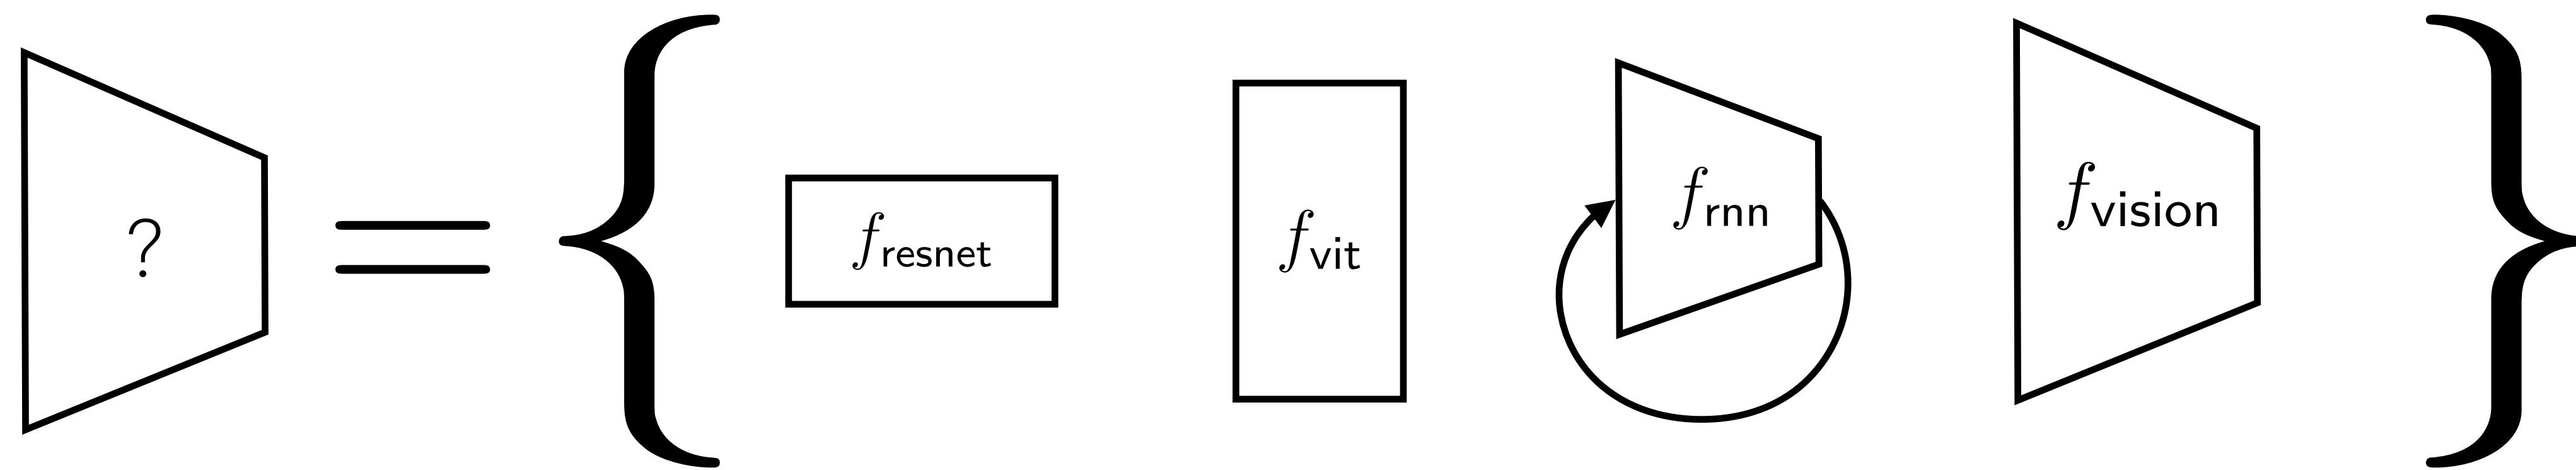
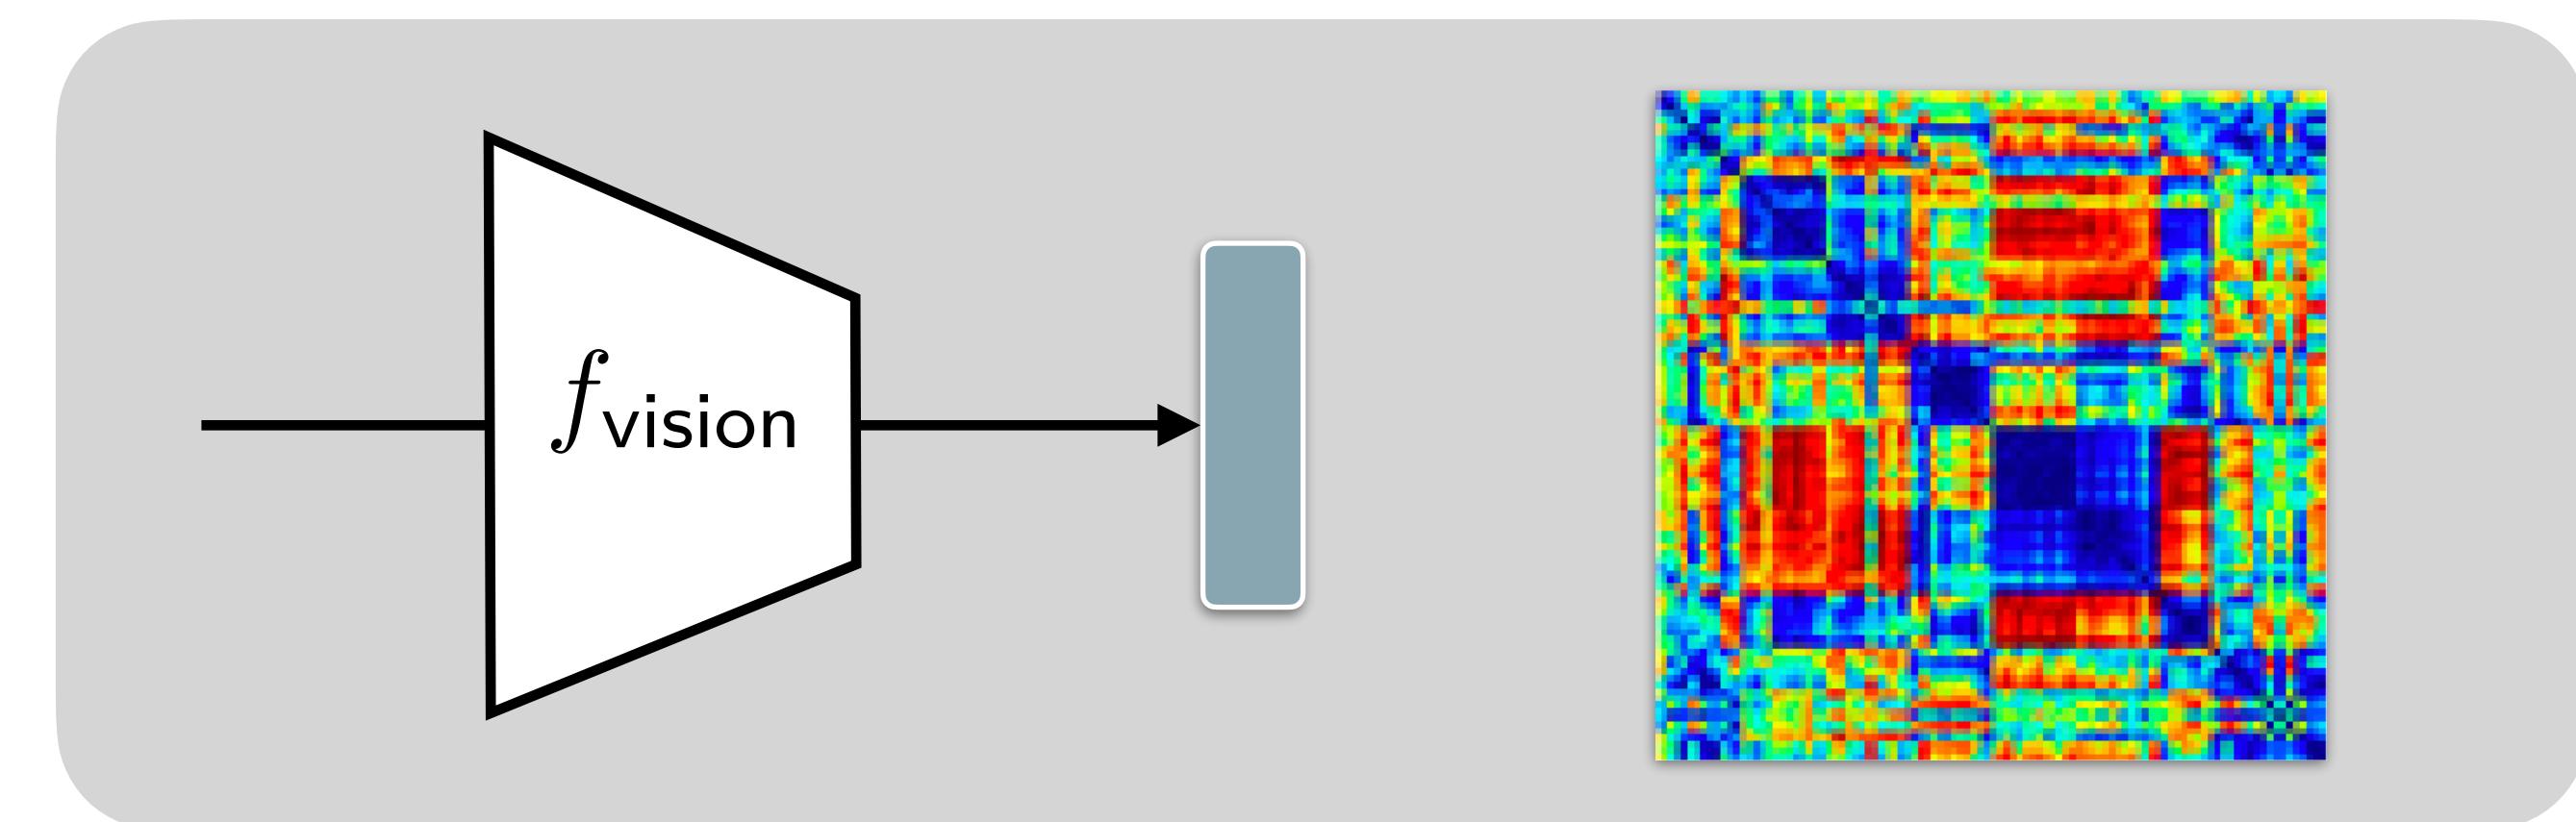
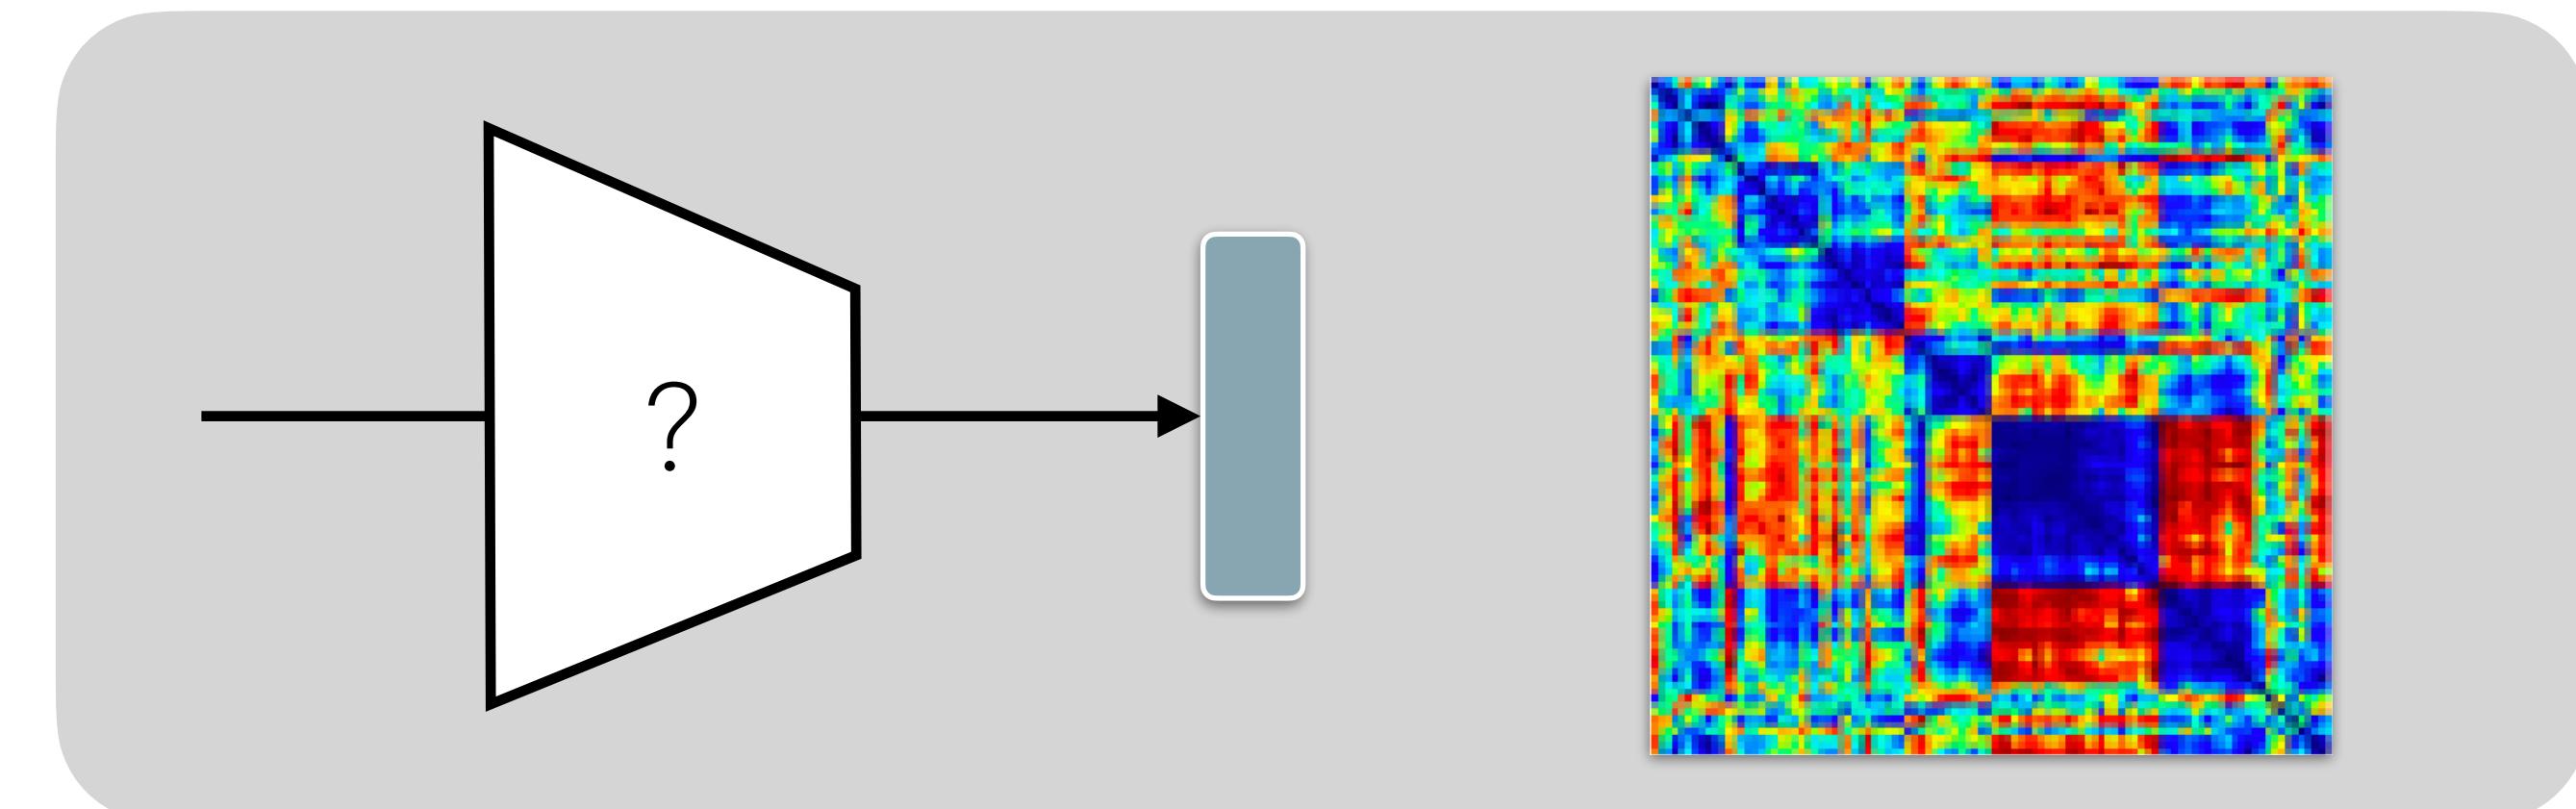
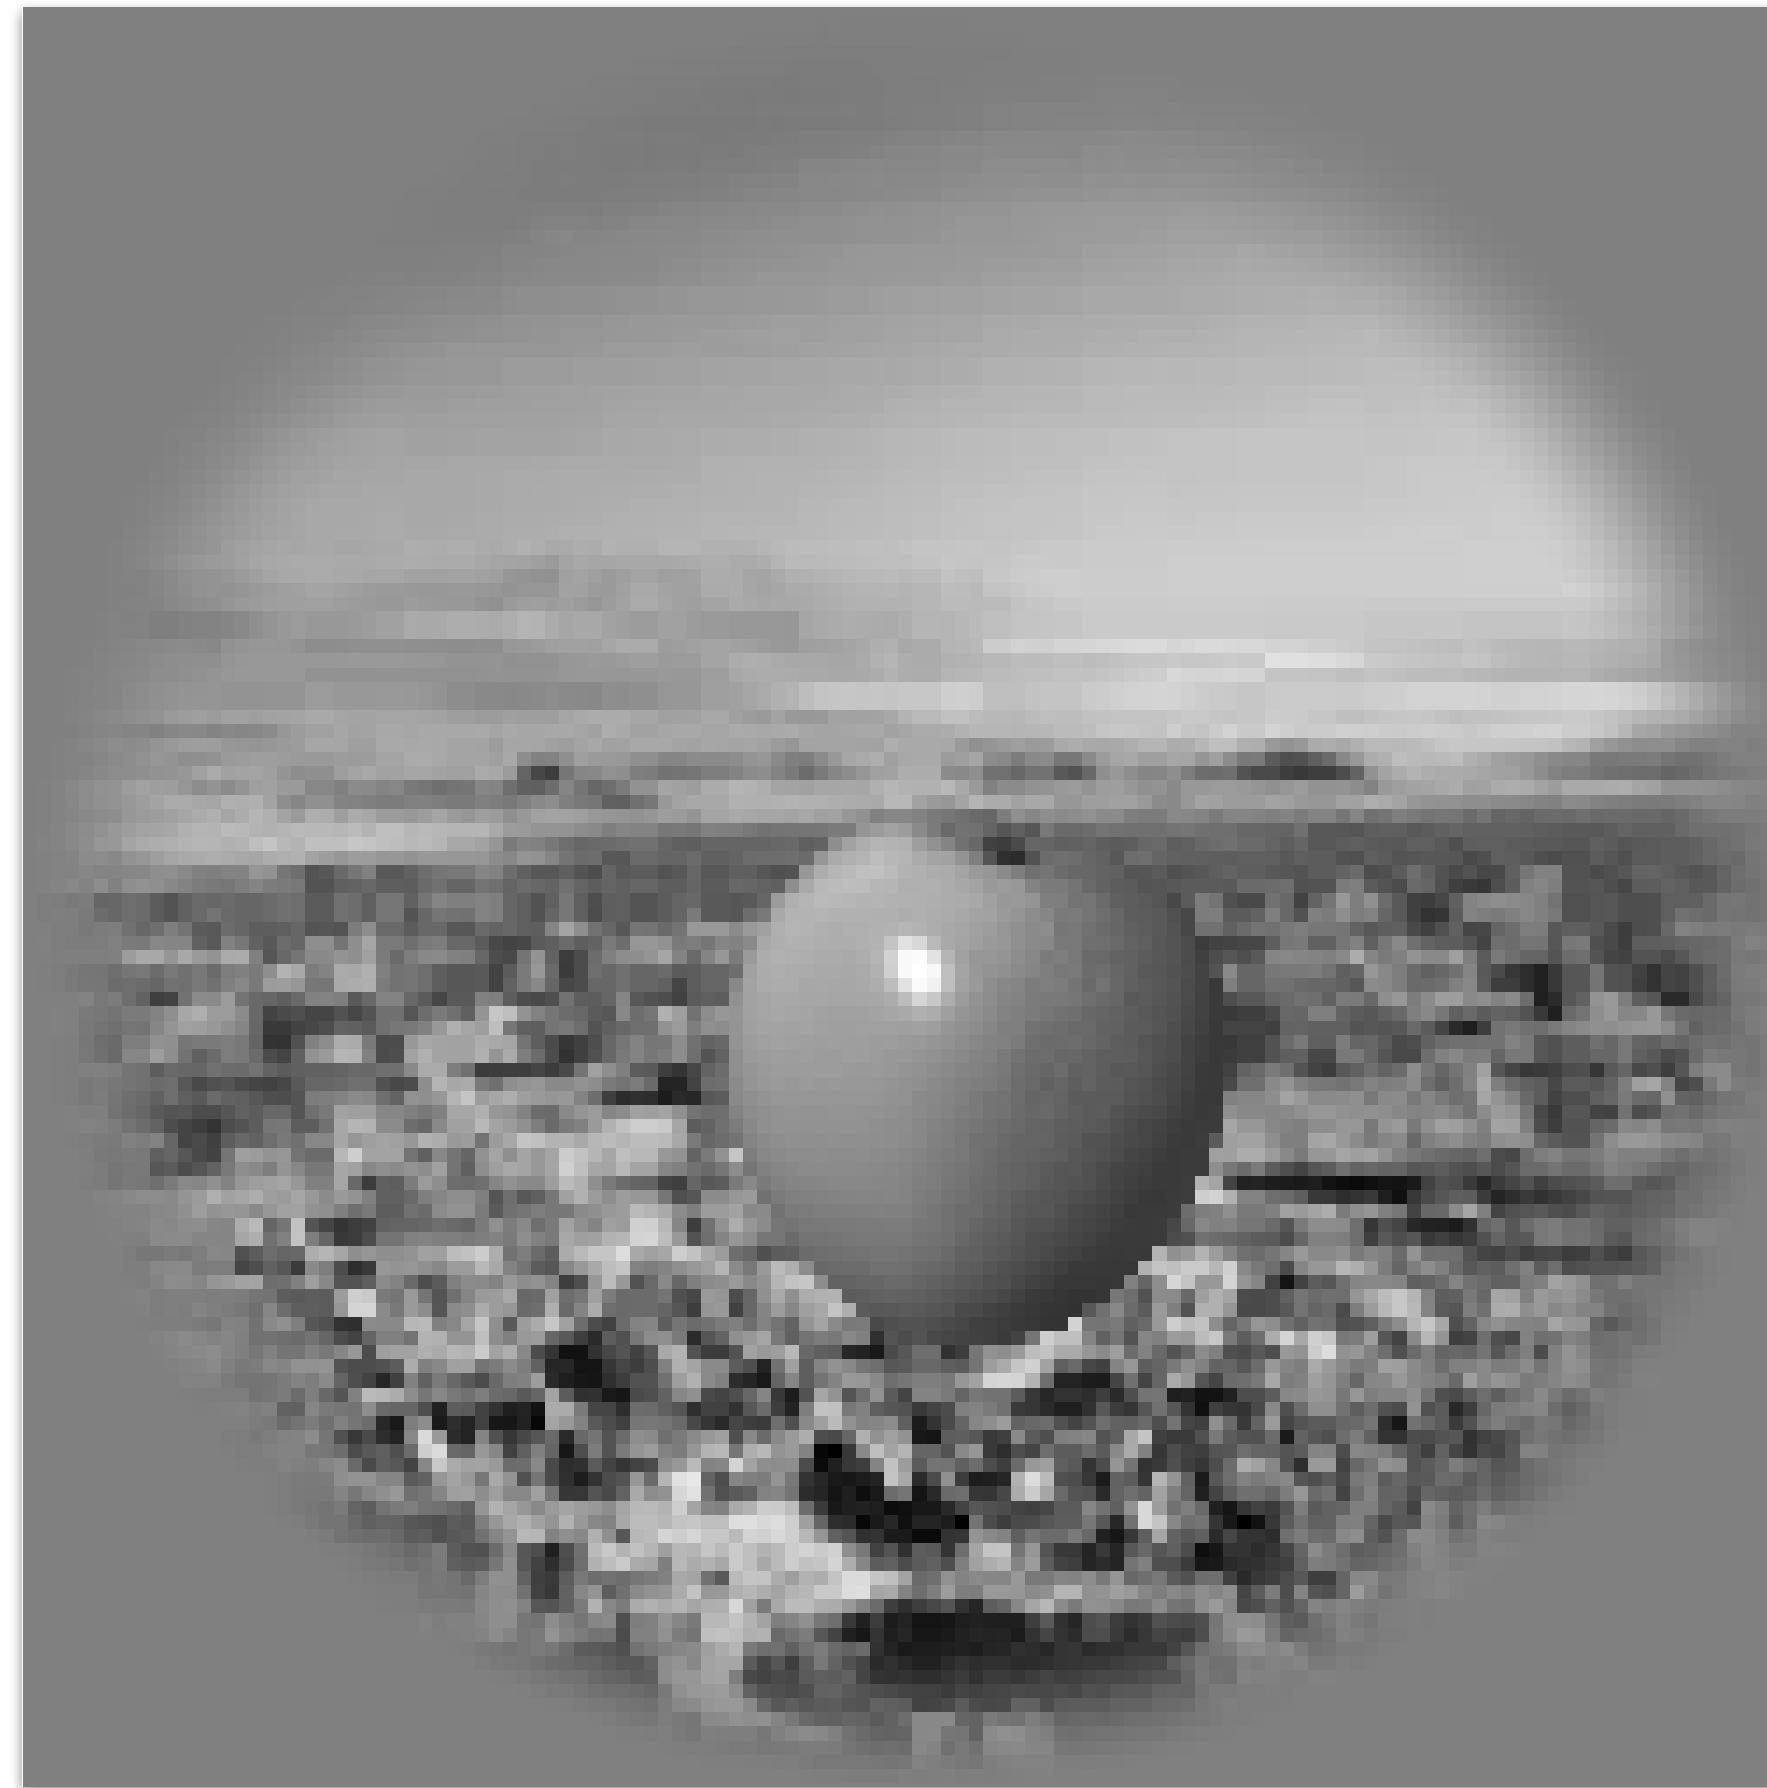
# Do models measure similarity in **similar** ways?

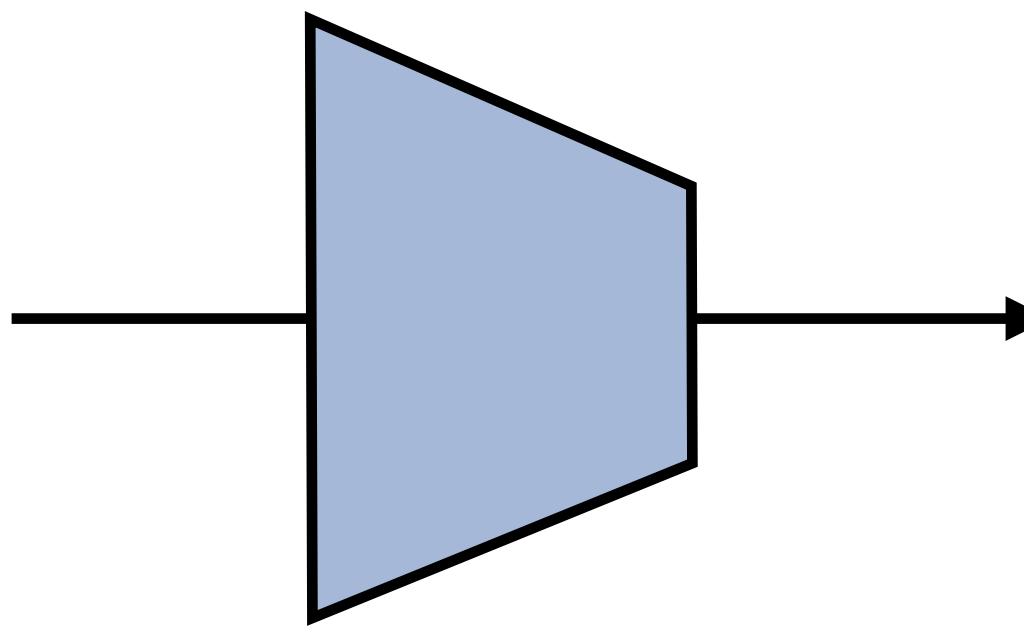
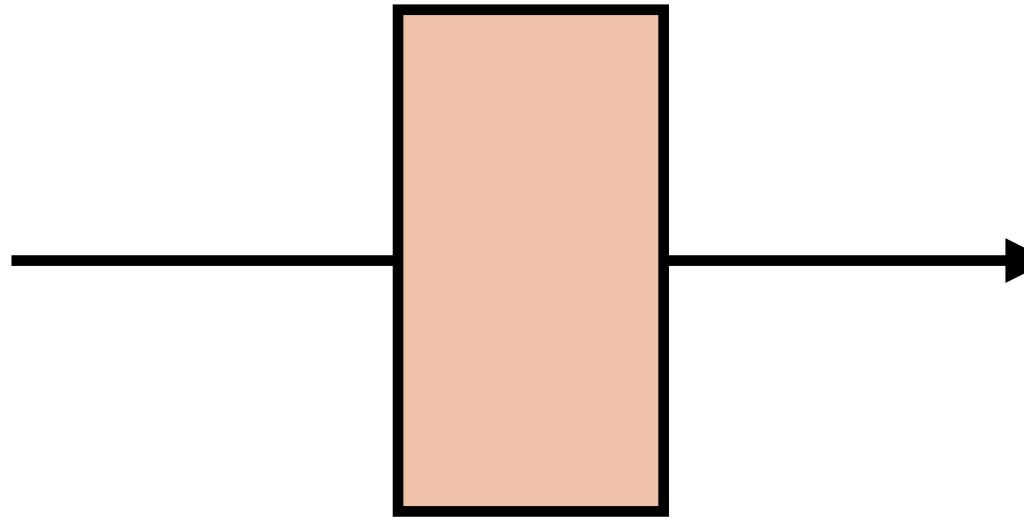


# Representational Turing Test





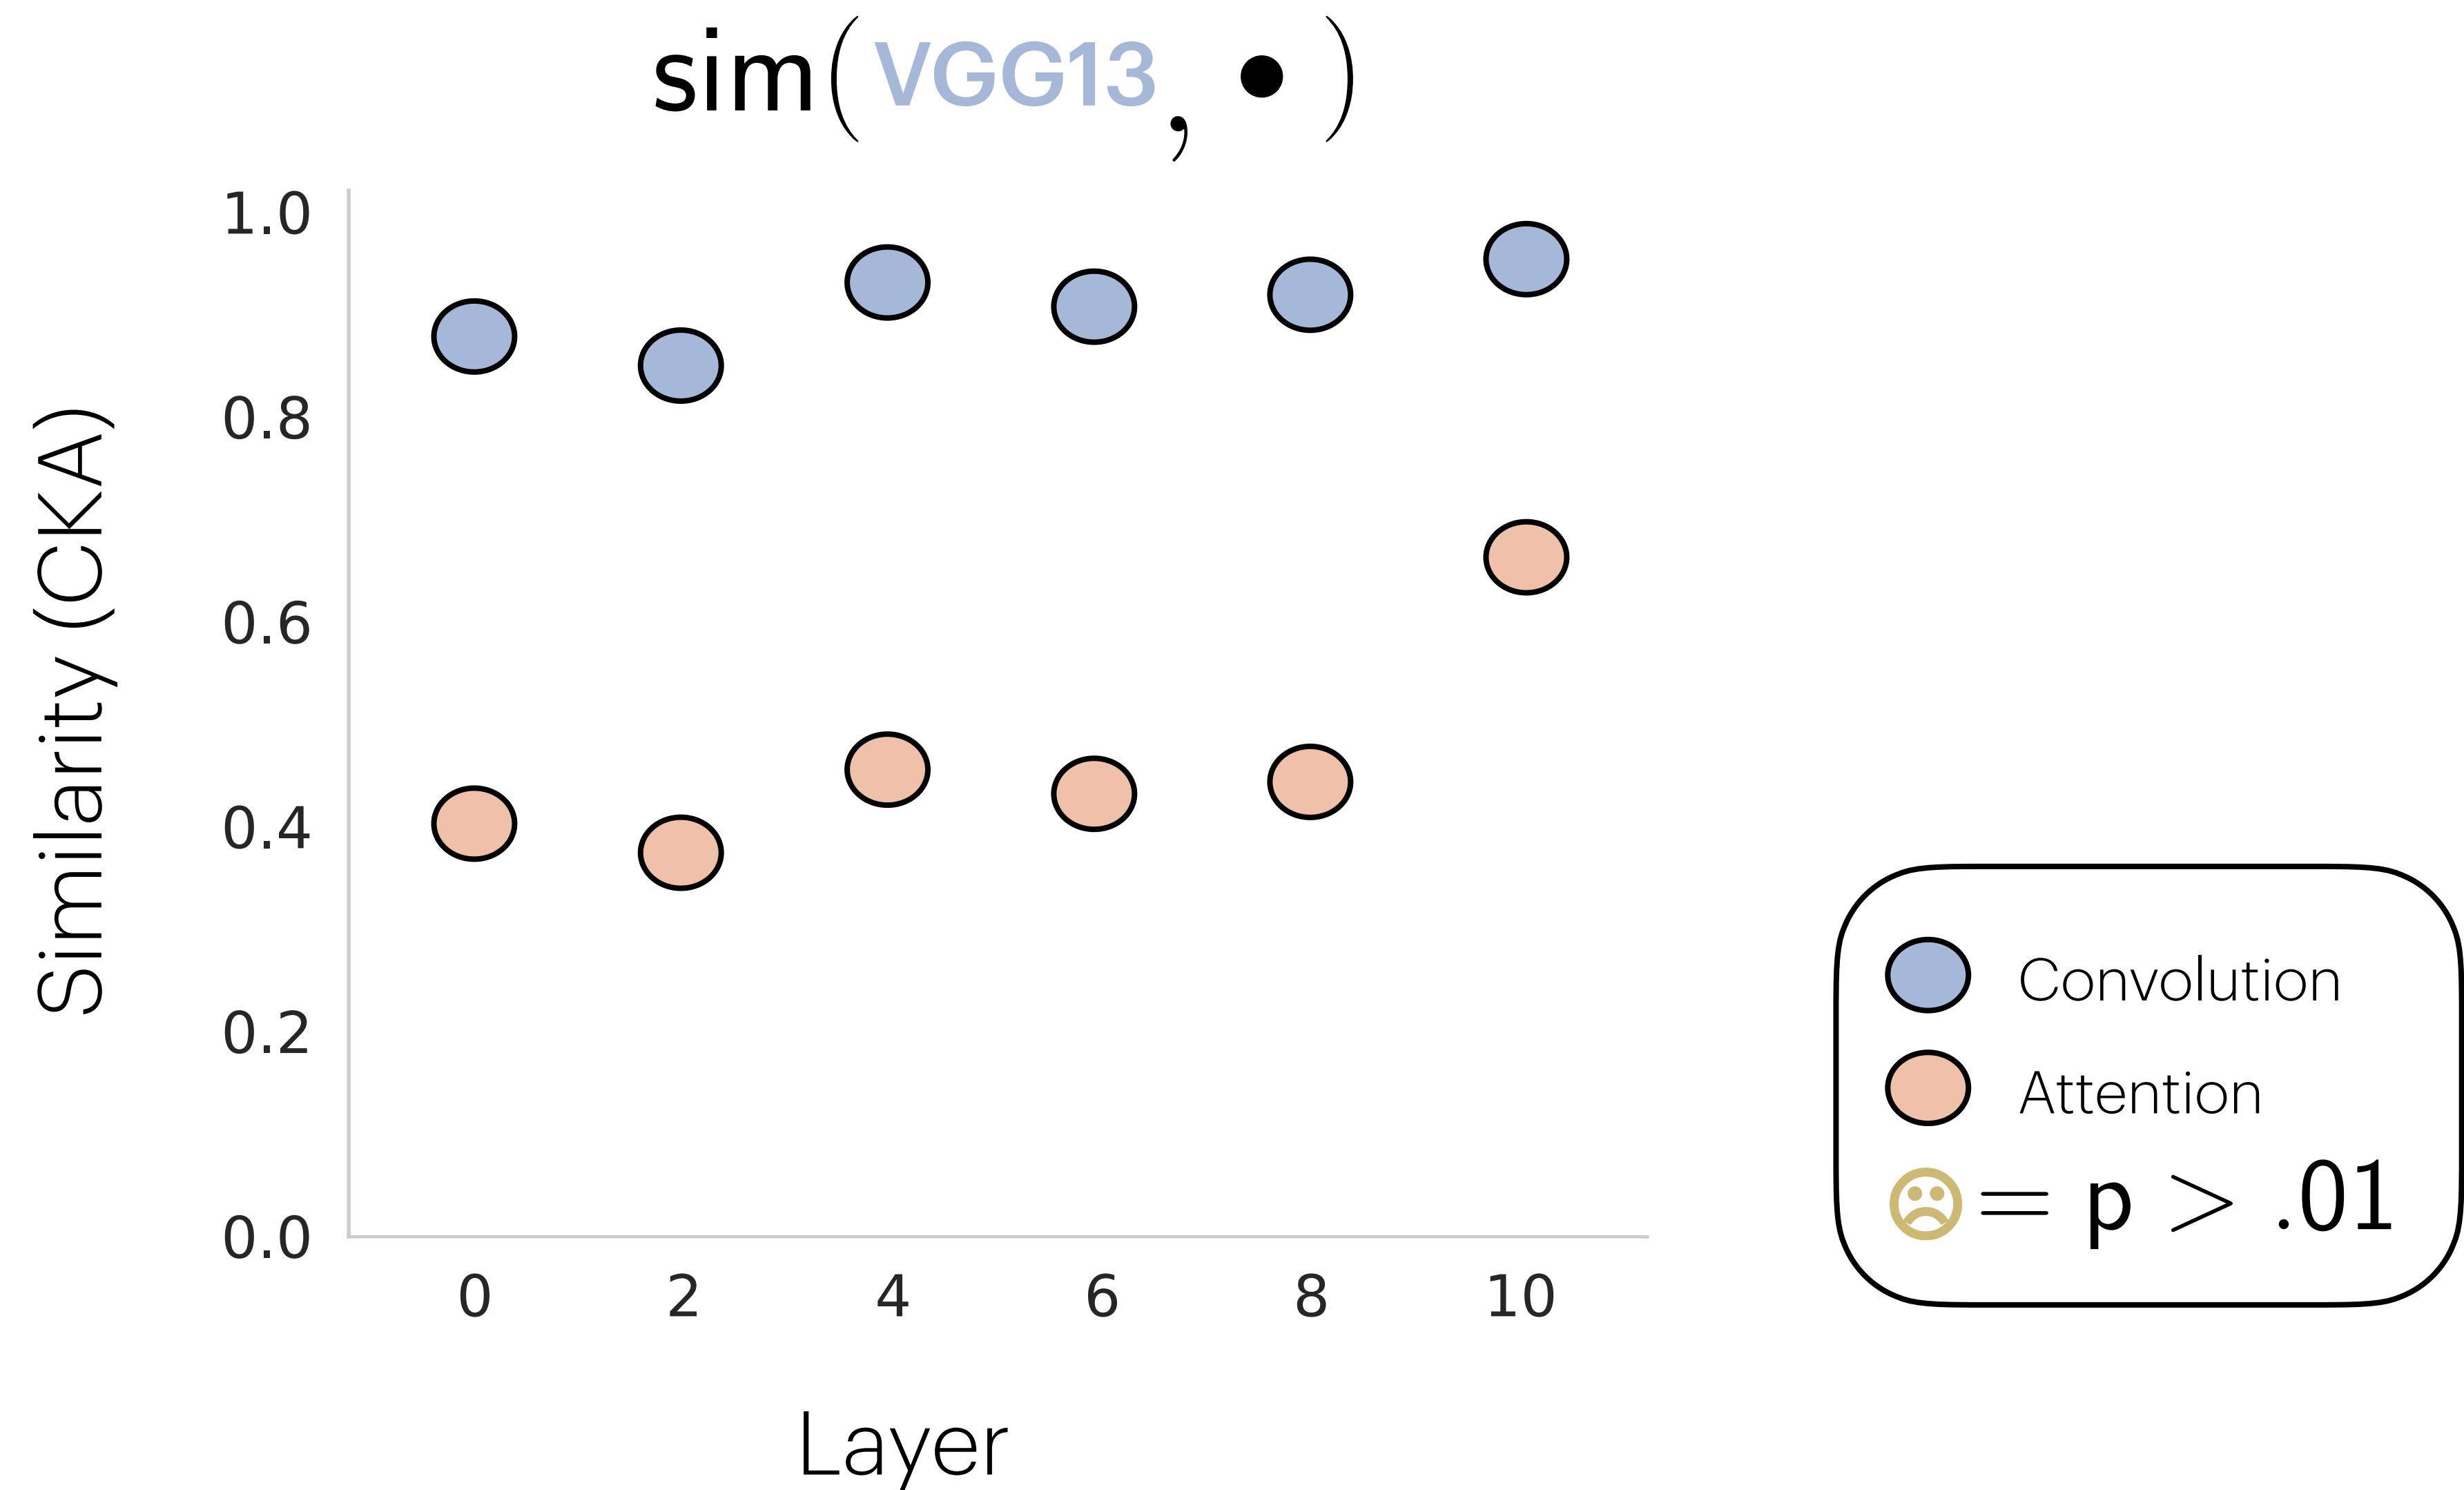


sim(  ,  )

Convolution

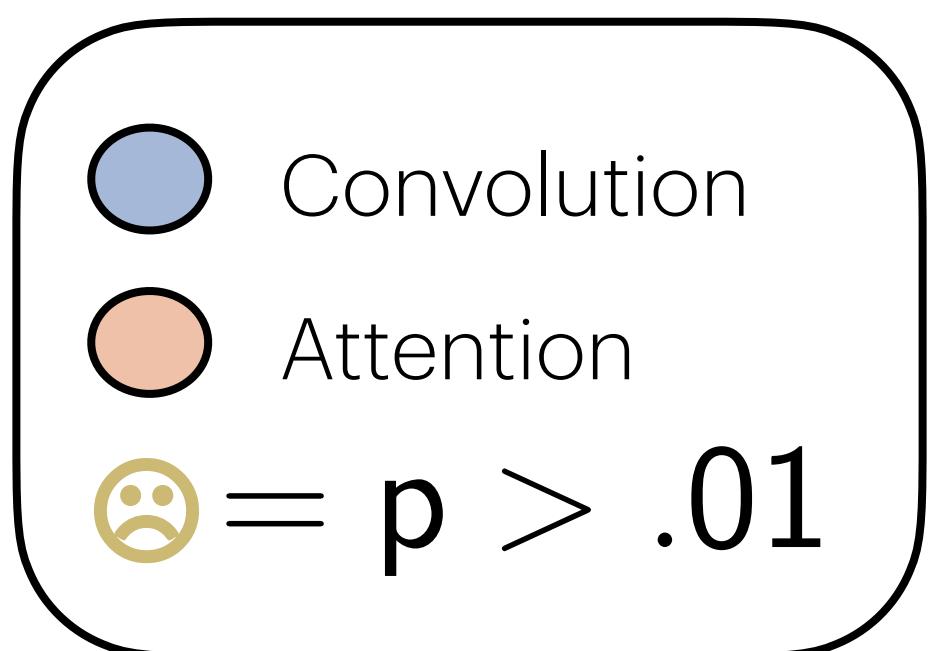
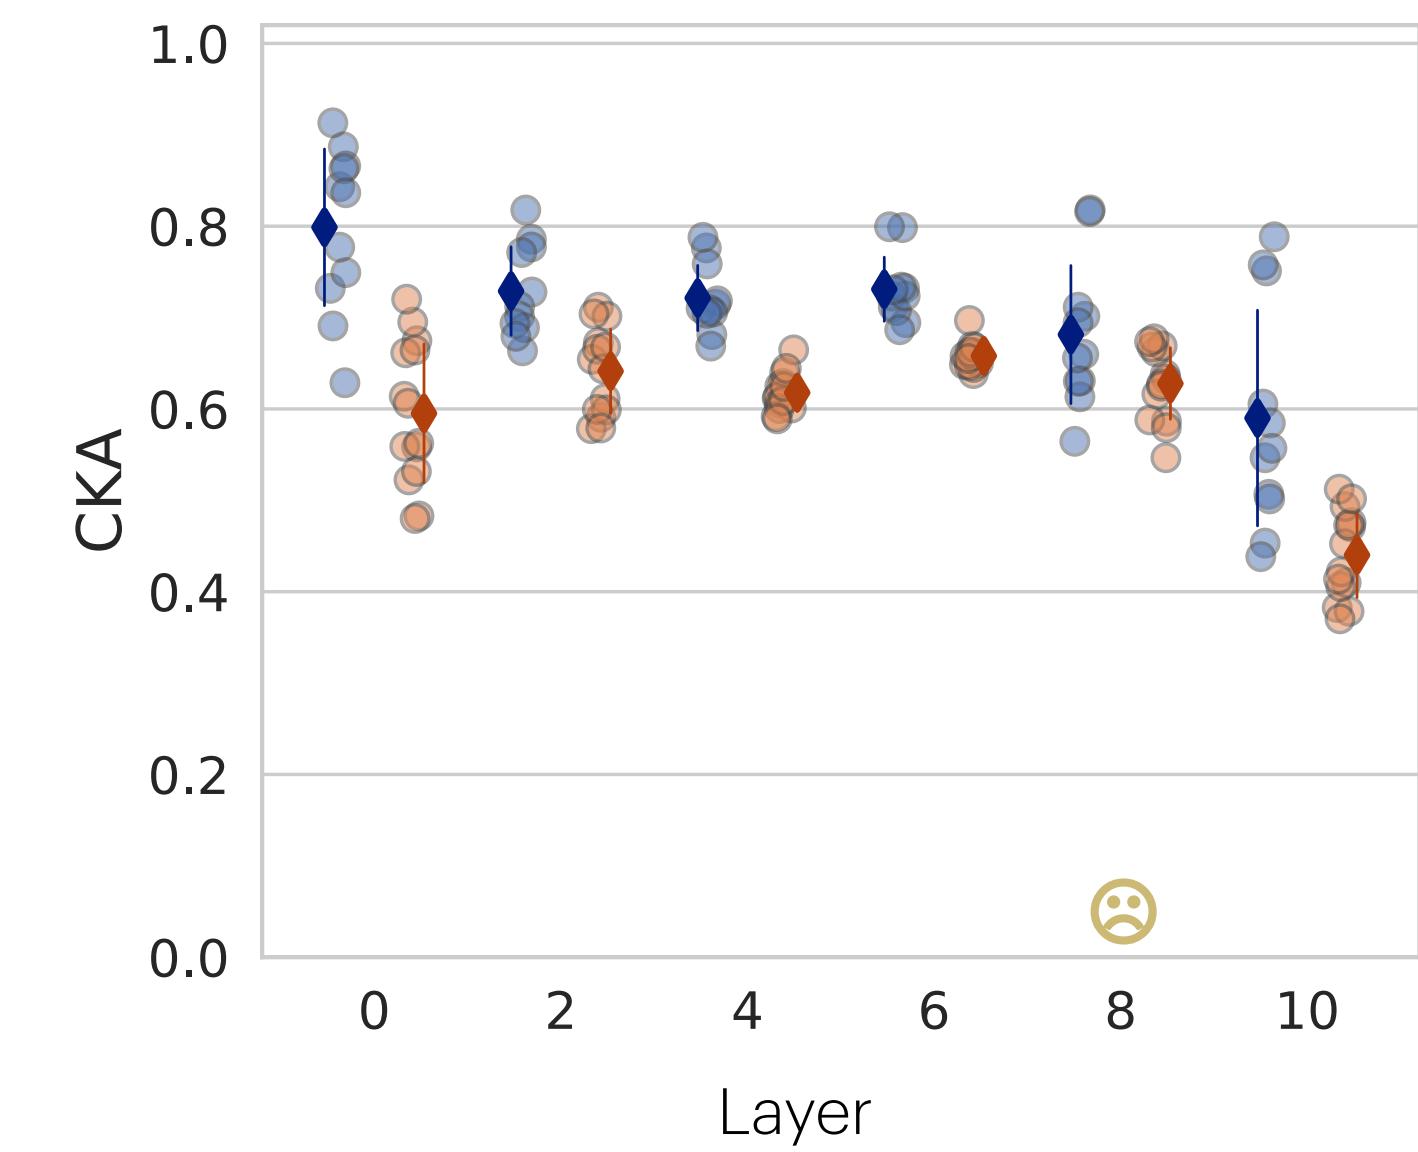
Attention

Easy to tell Convolution apart from Attention



# Easy to tell Convolution apart from Attention?

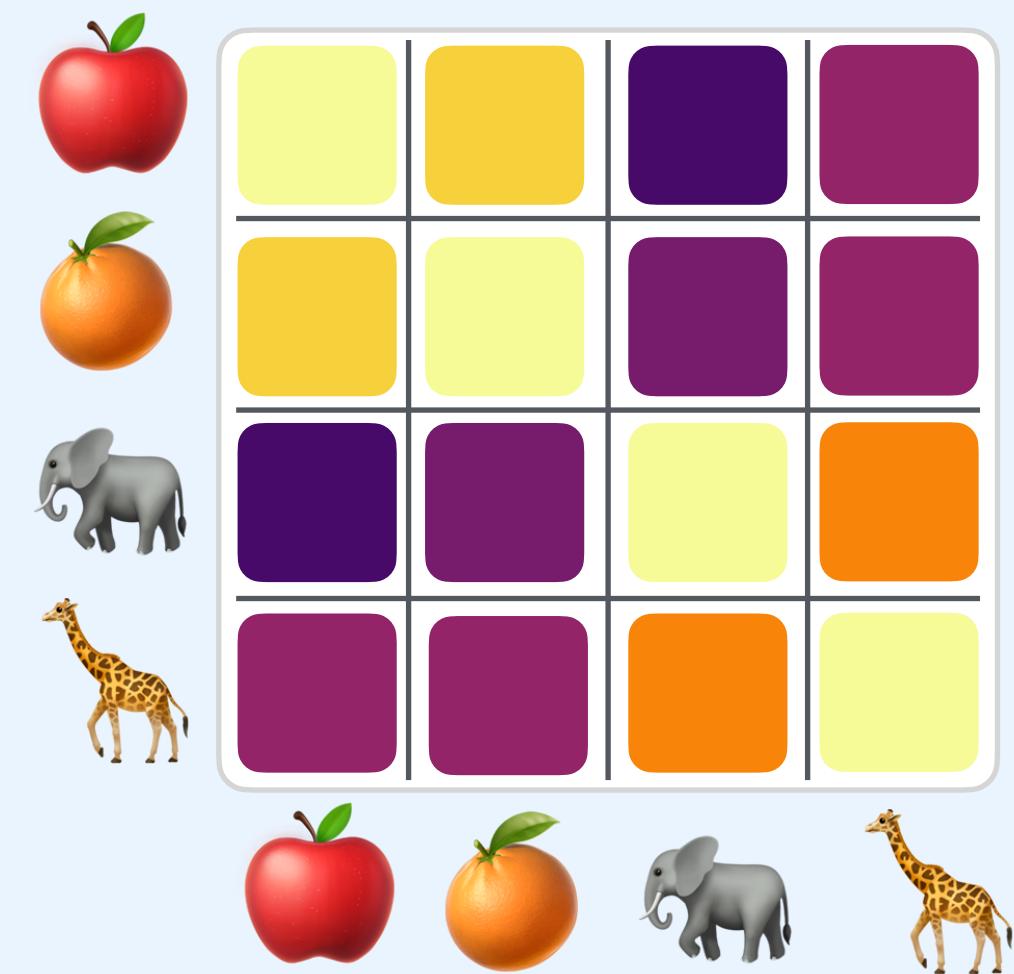
$\text{sim}(\text{VGG13}, \bullet)$



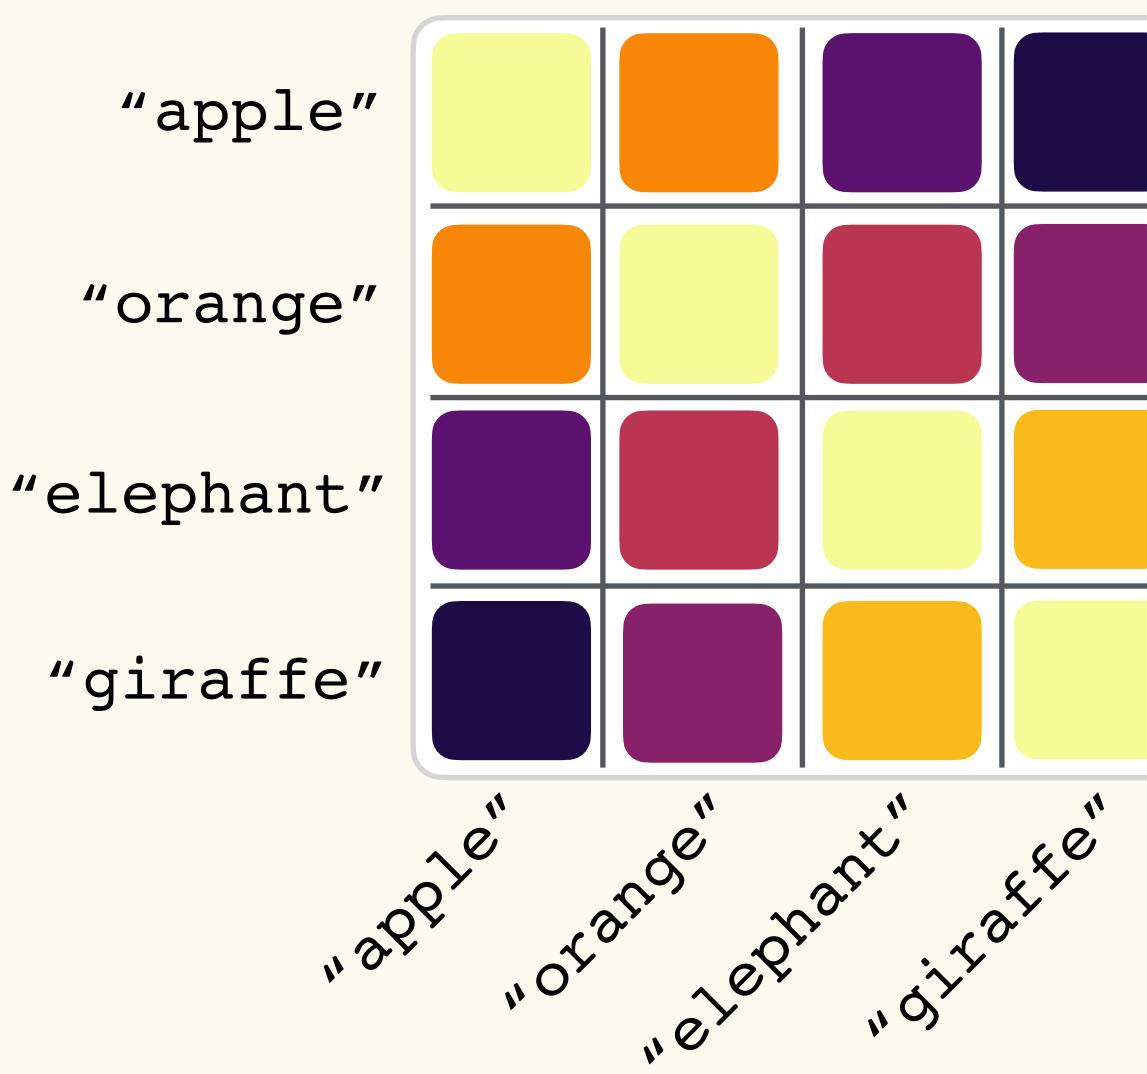
Do **different** models represent the world in **different** ways?  
Or are they somehow **all becoming alike**?

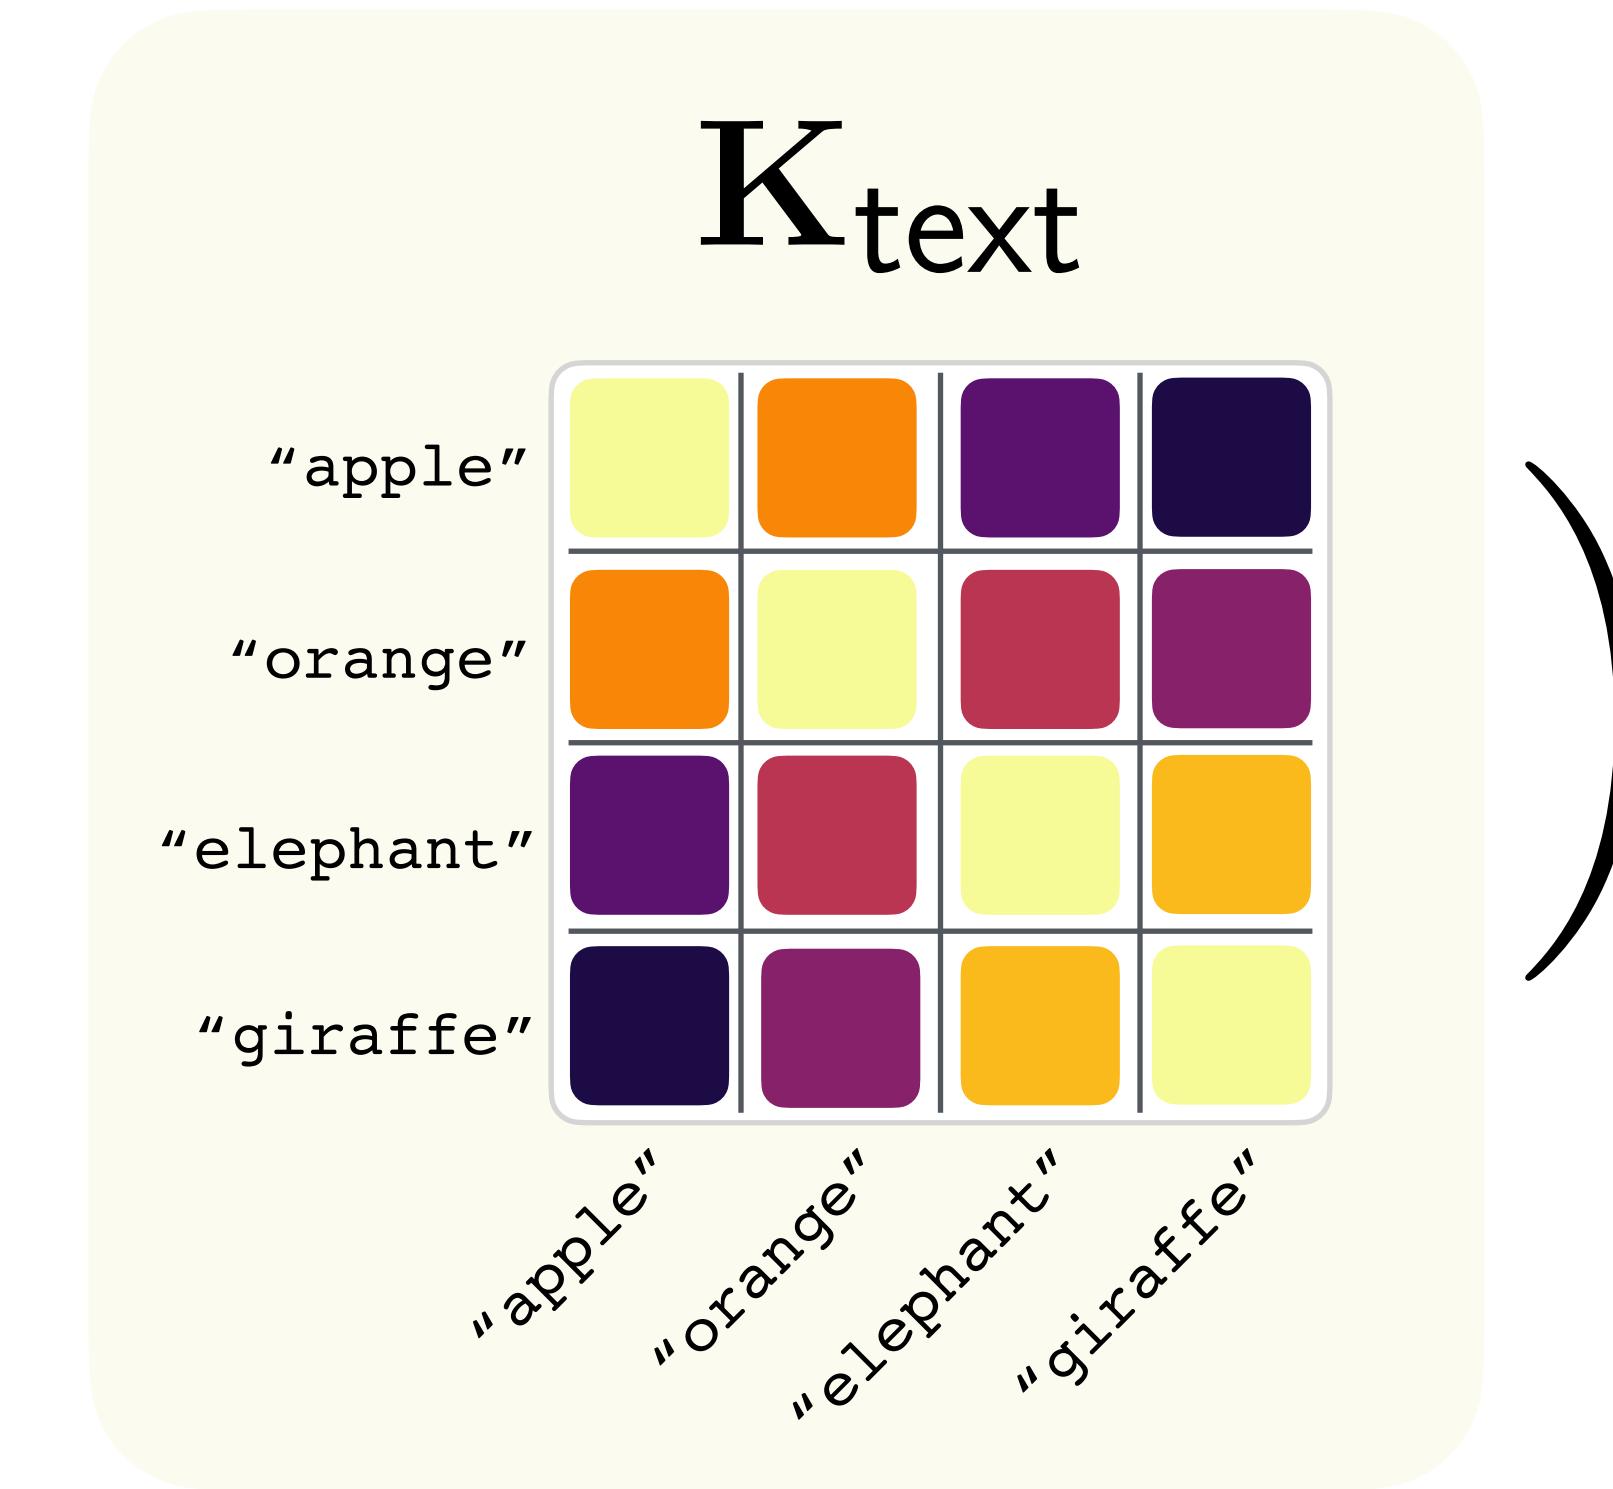
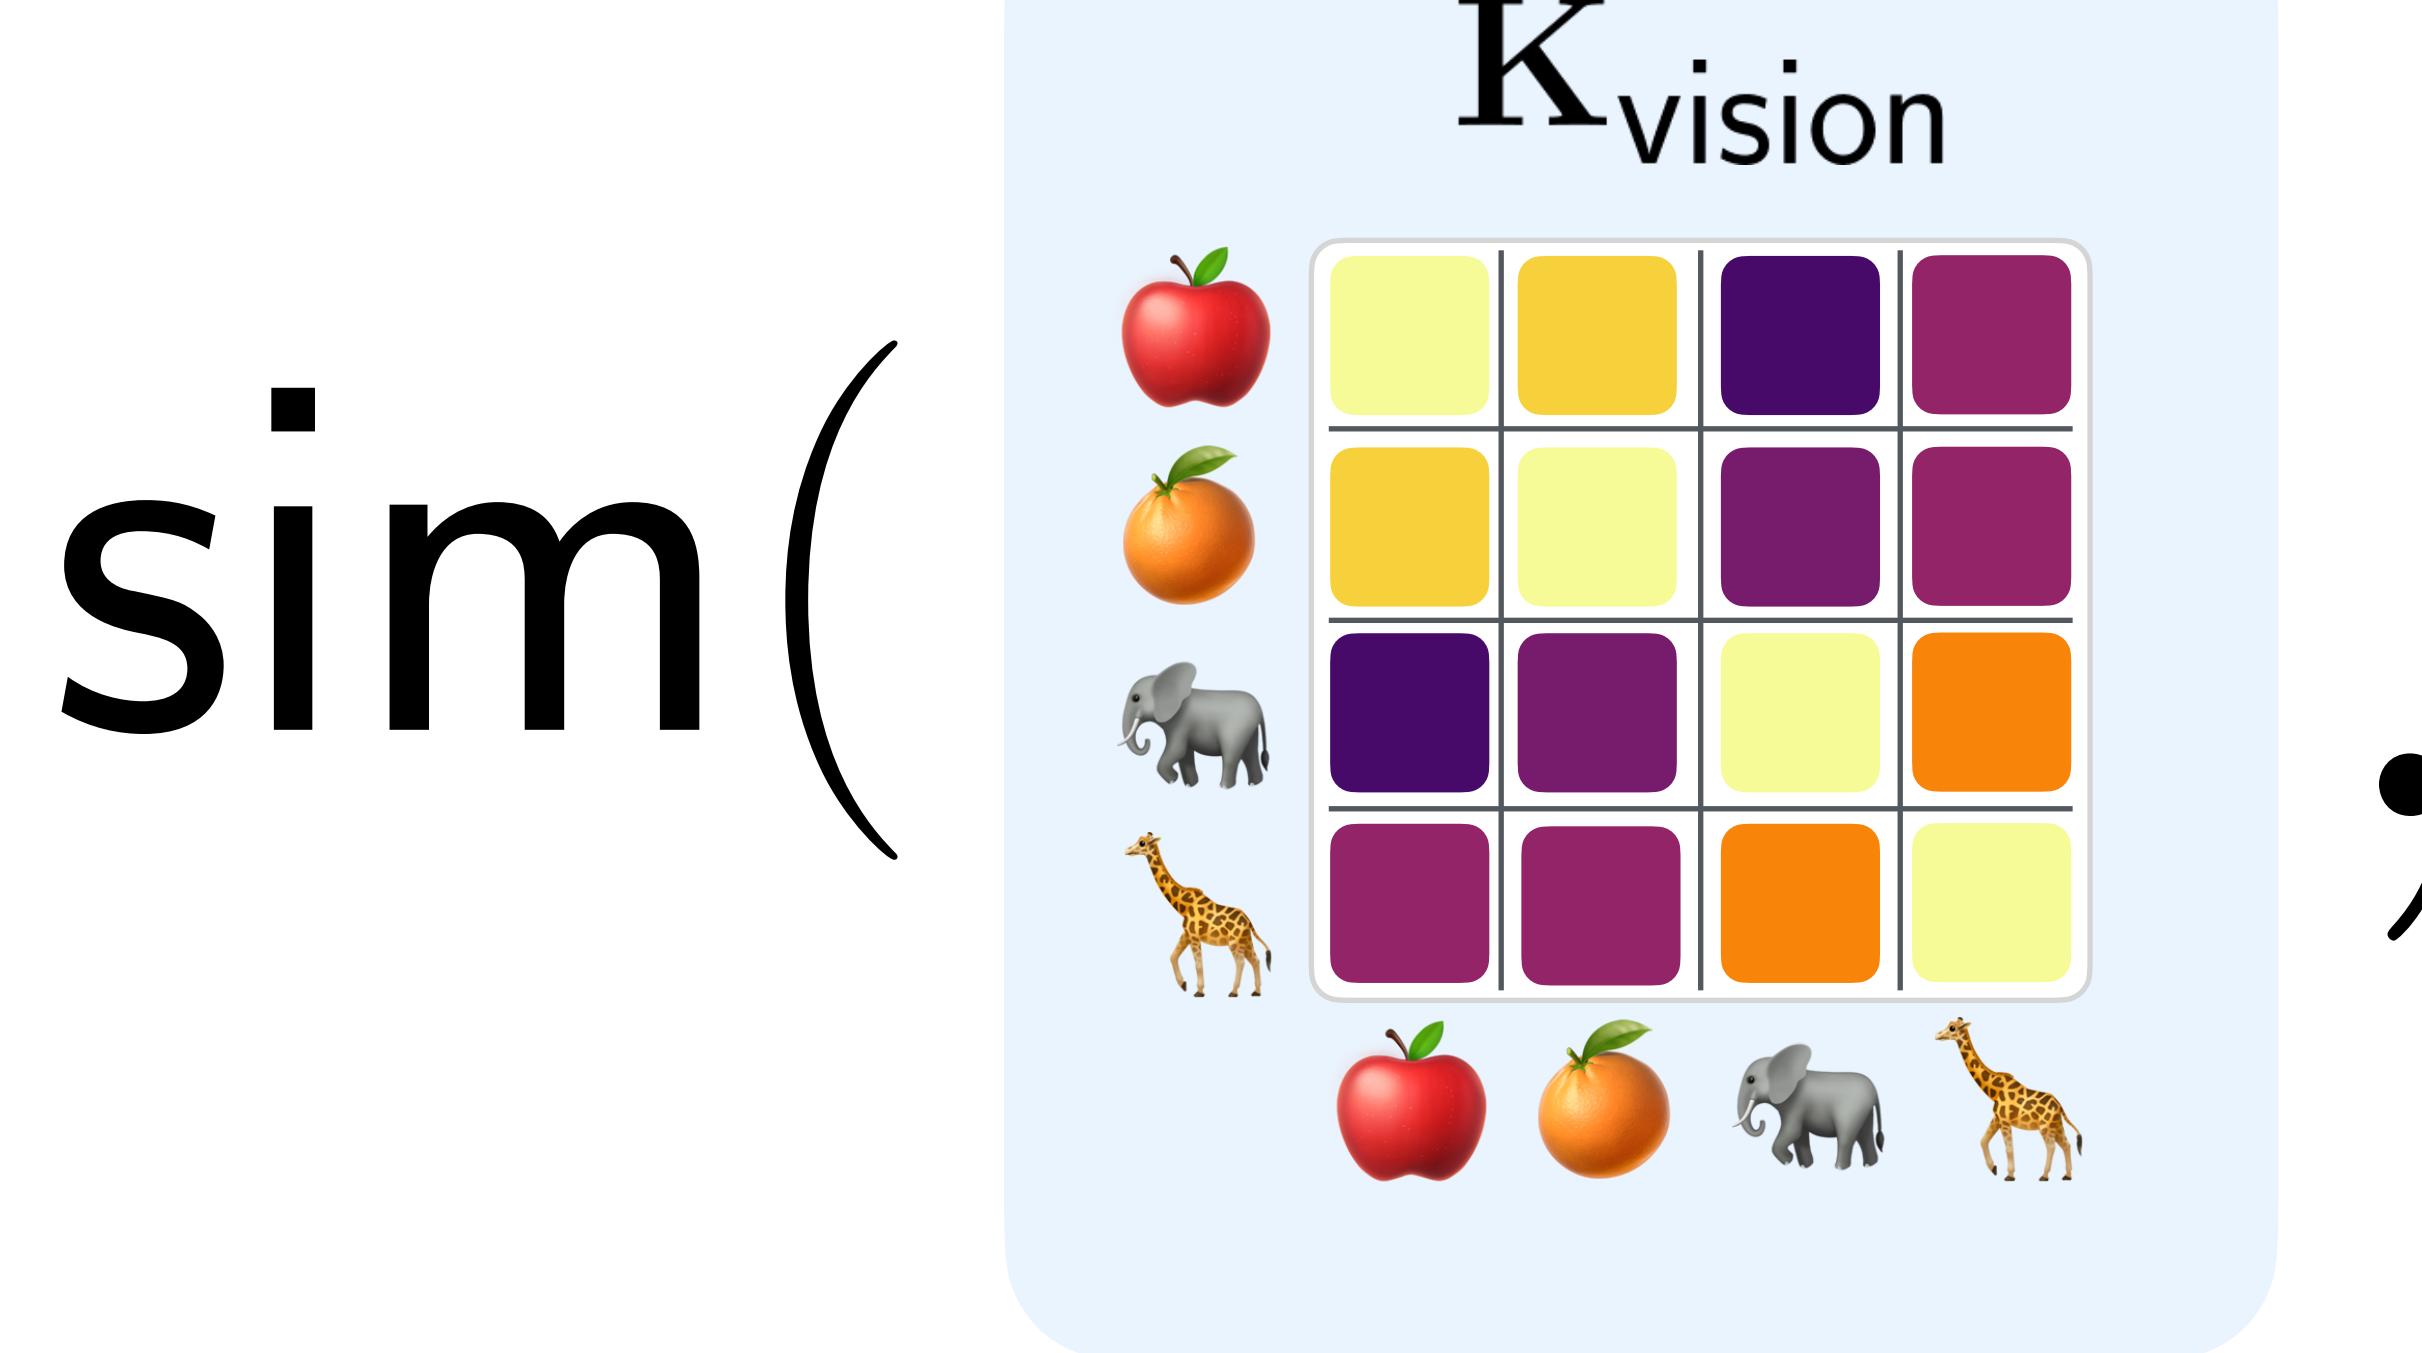
# Do models measure similarity in **similar** ways?

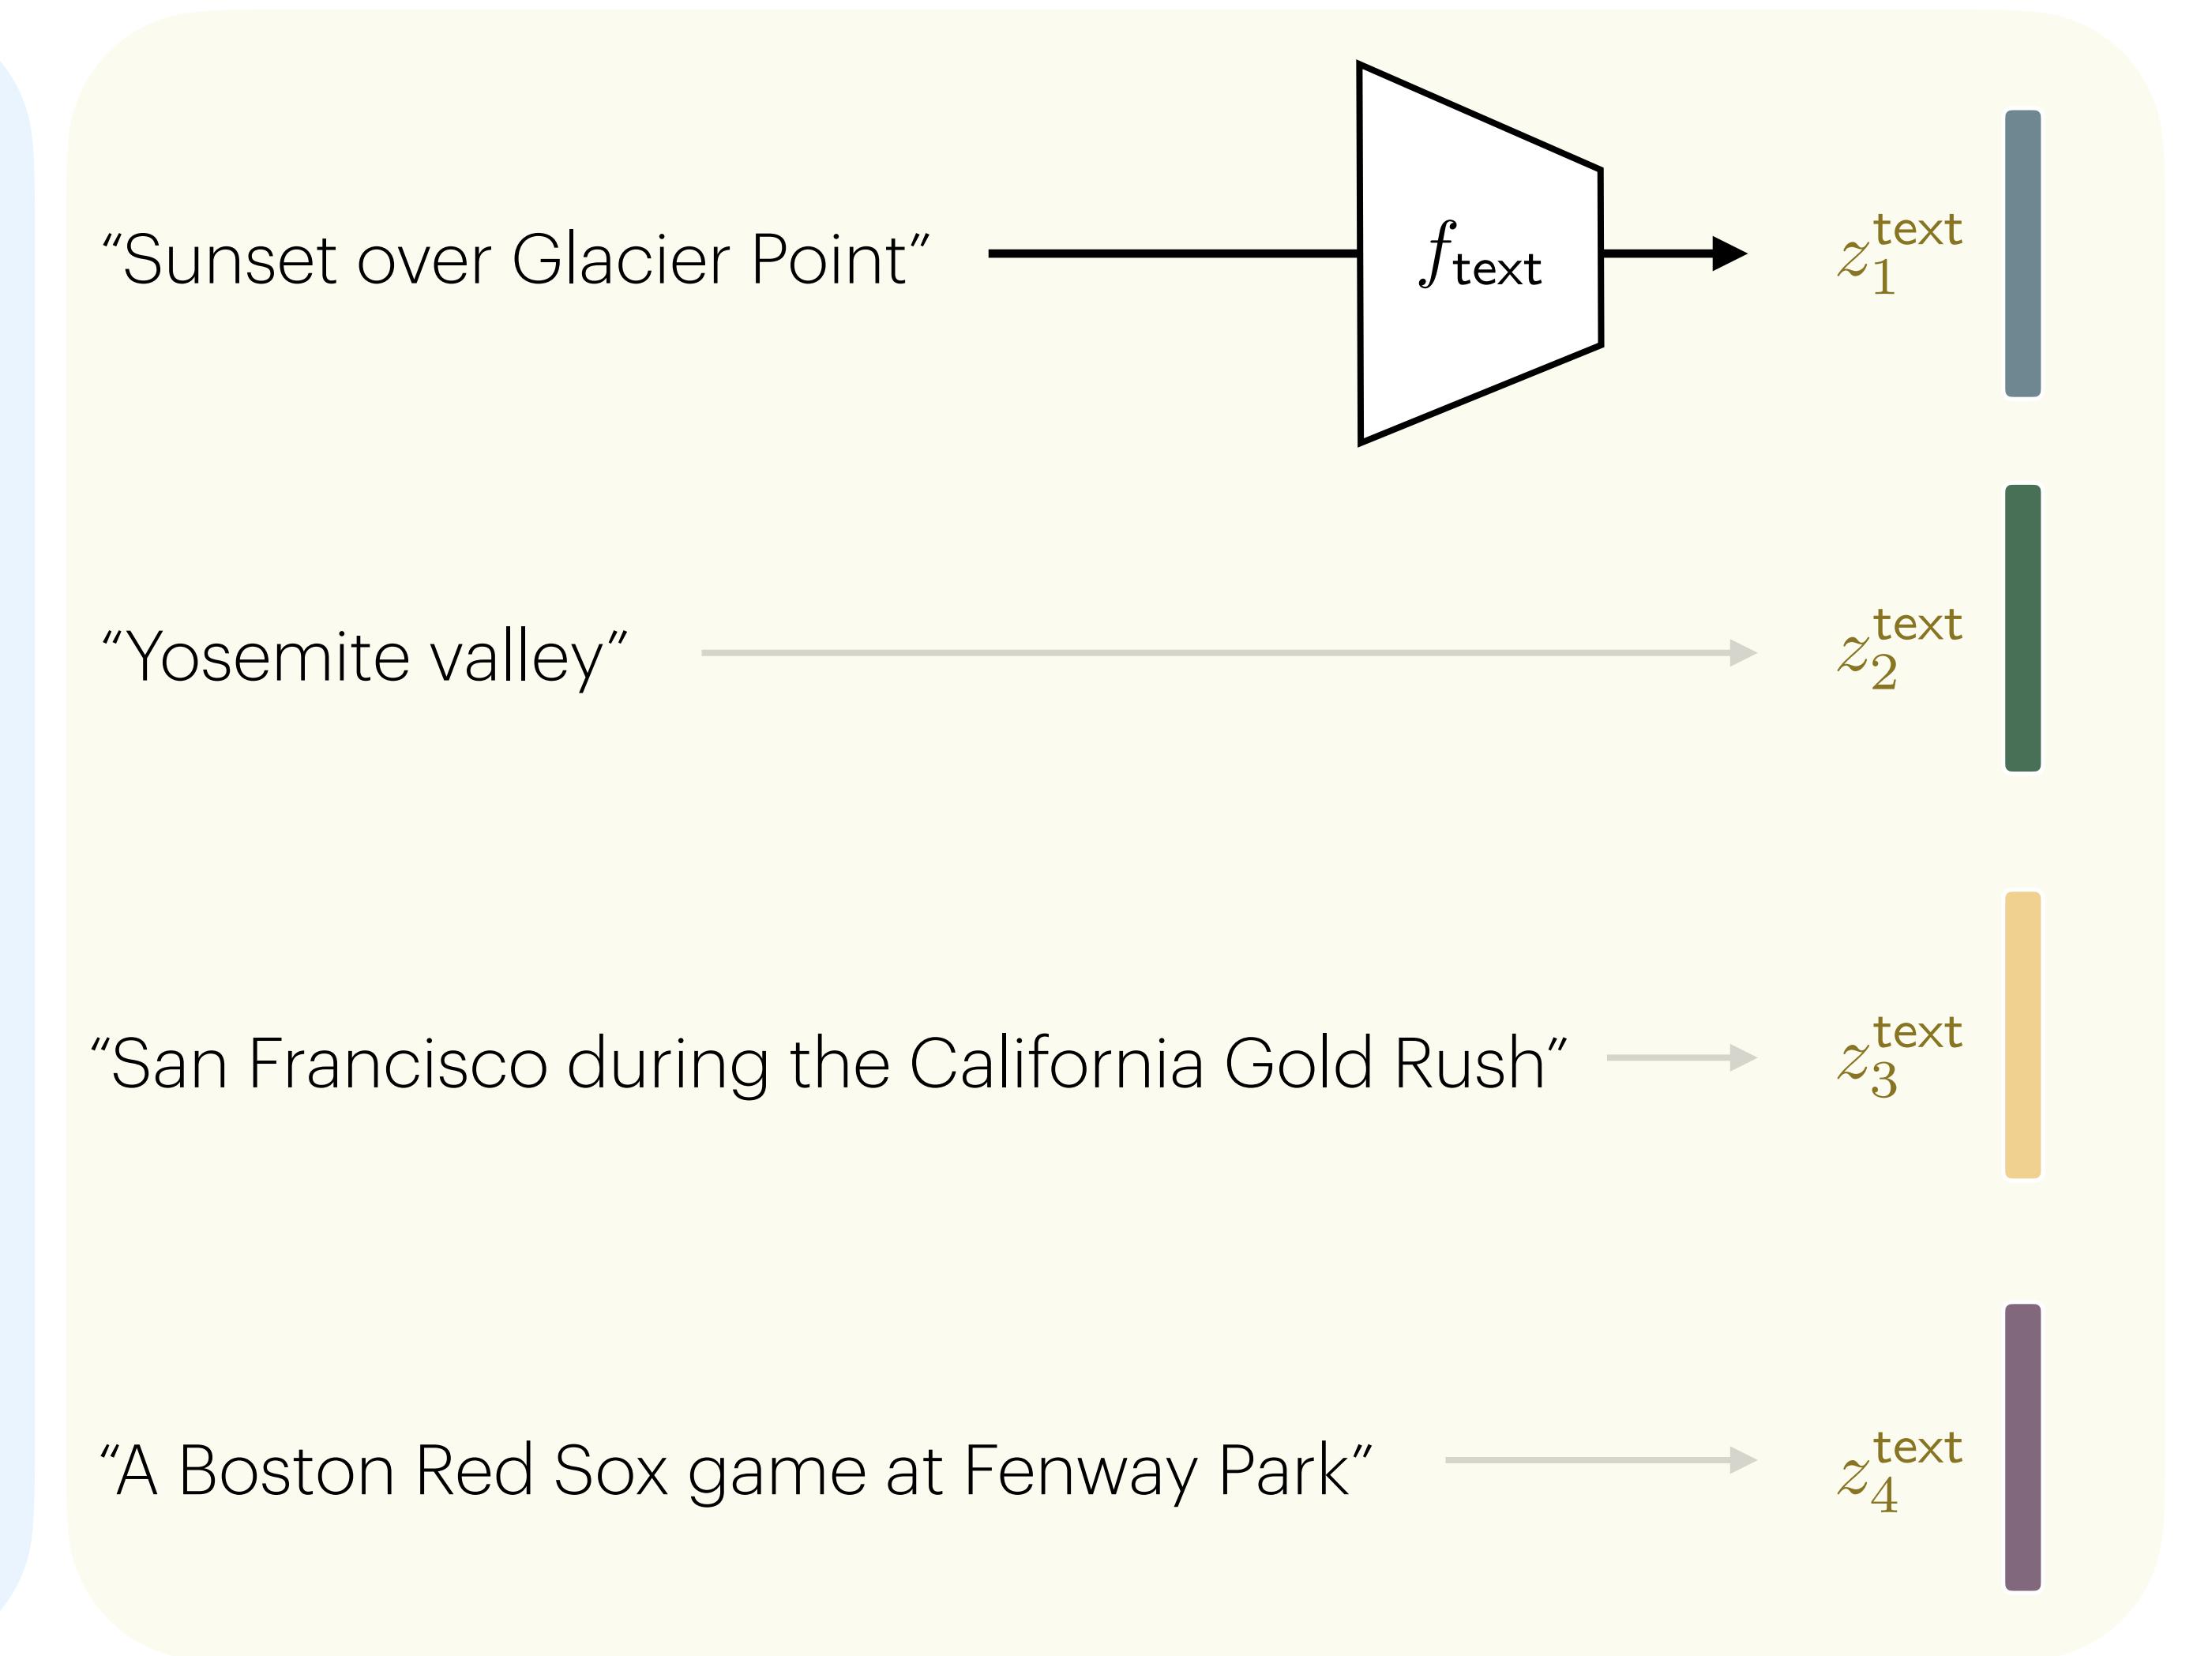
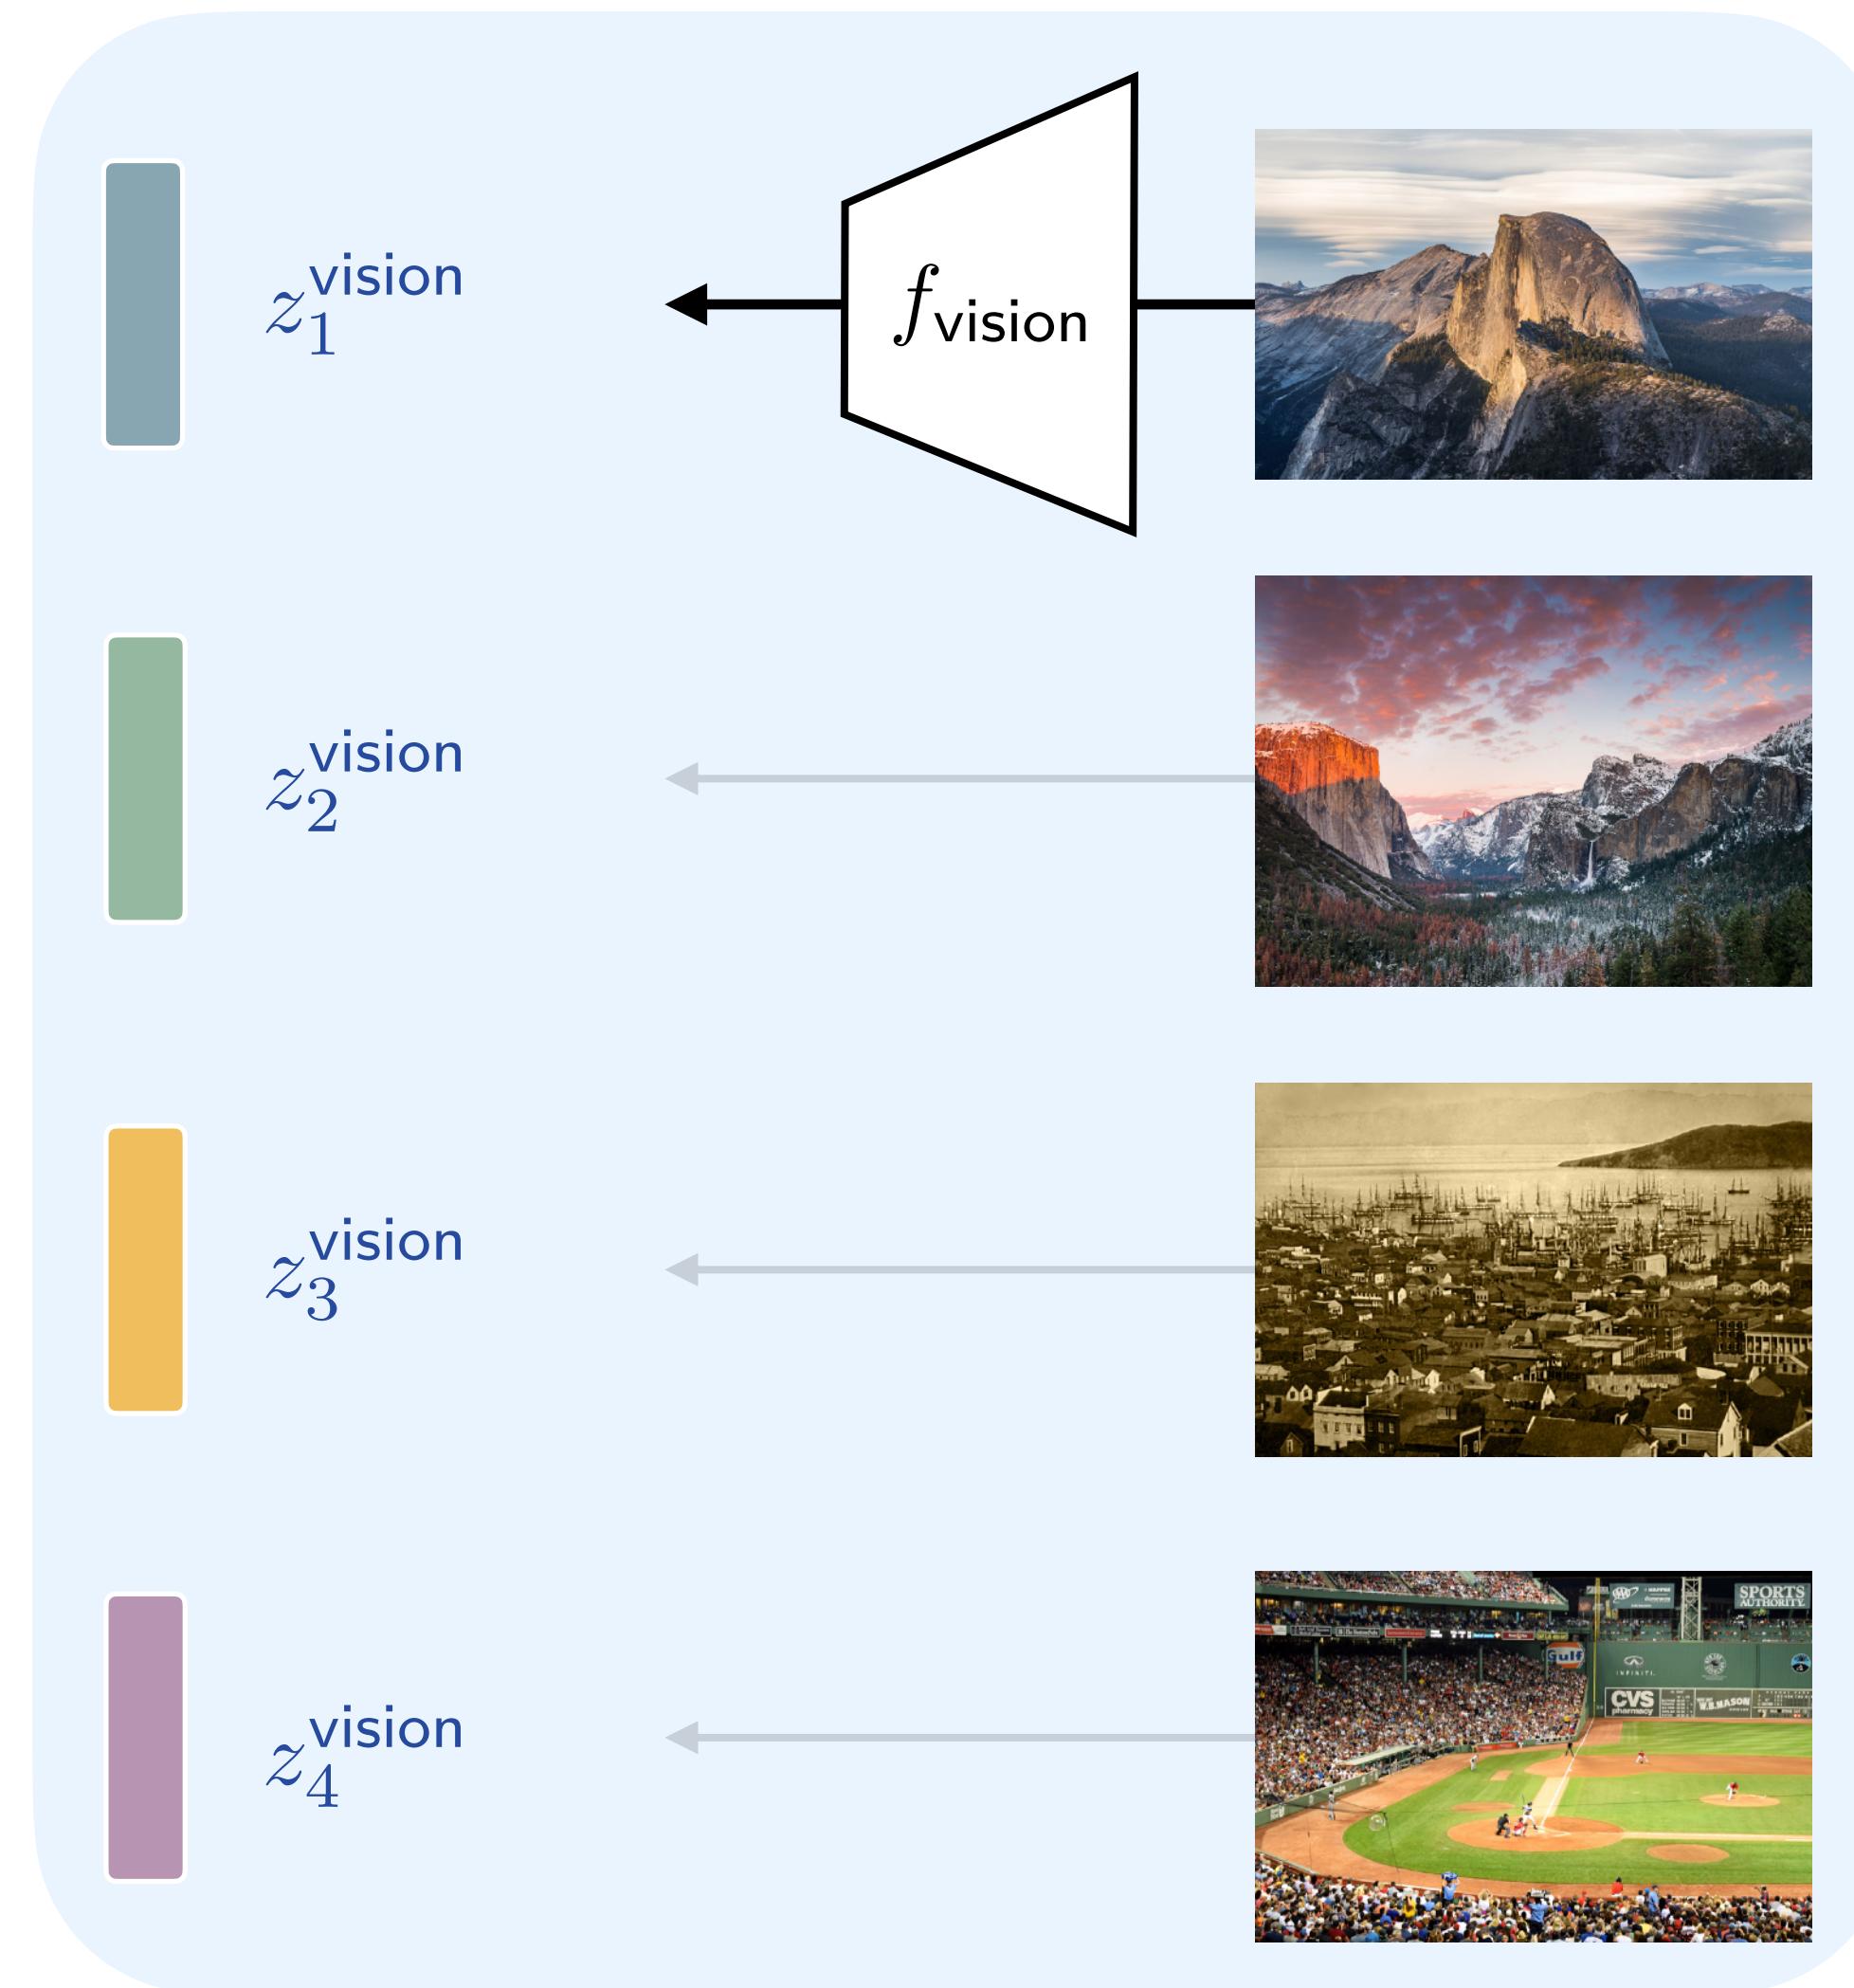
$K_{\text{vision}}$



$K_{\text{text}}$



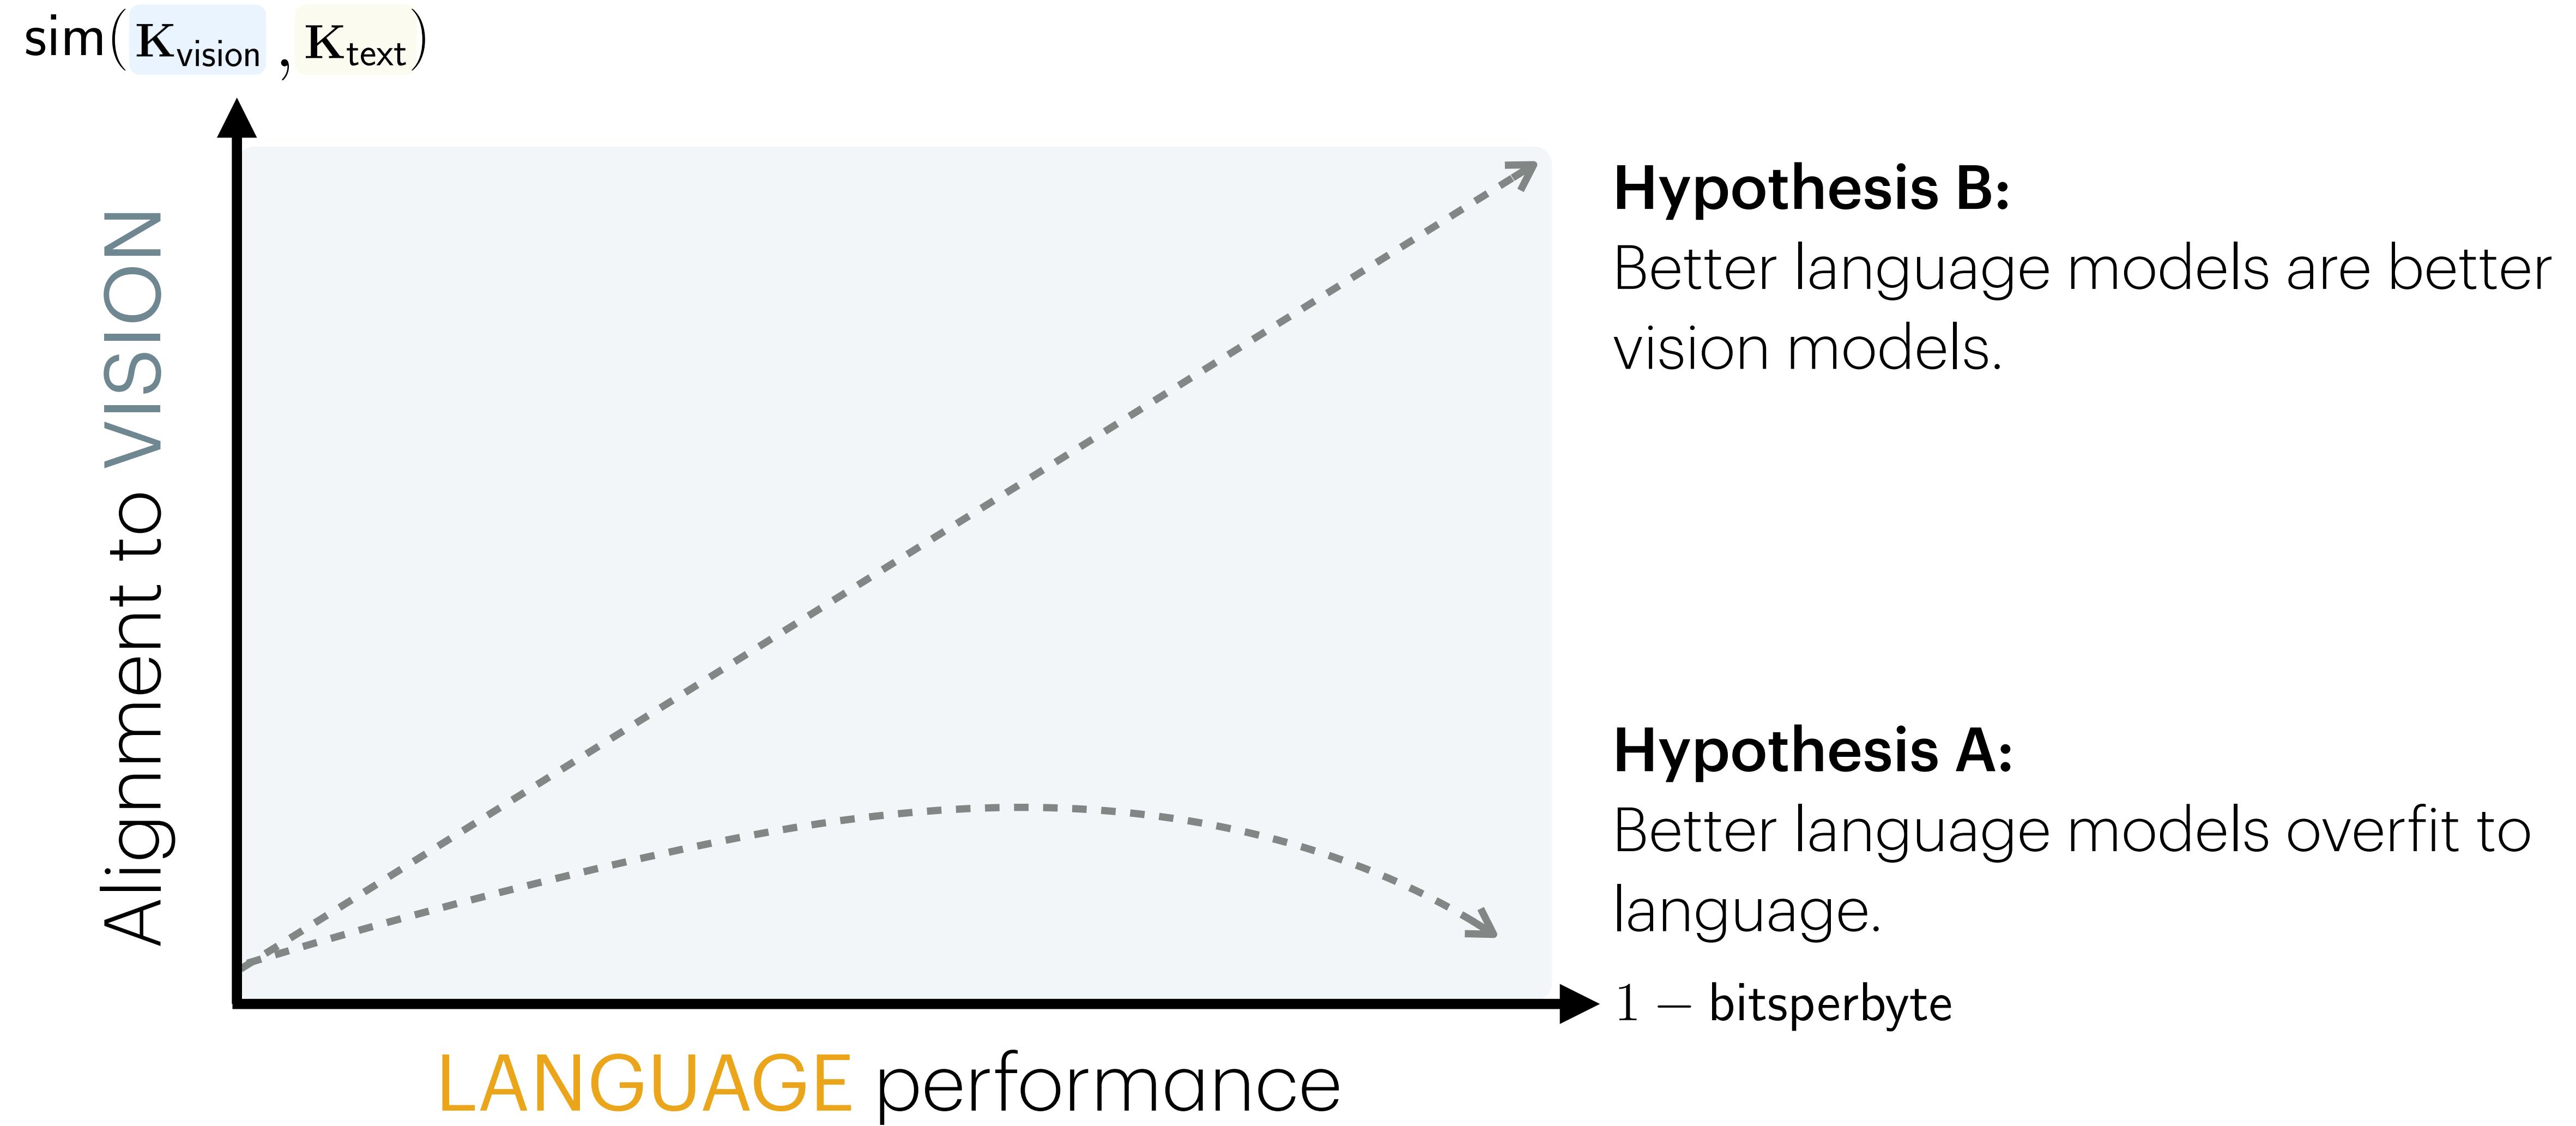




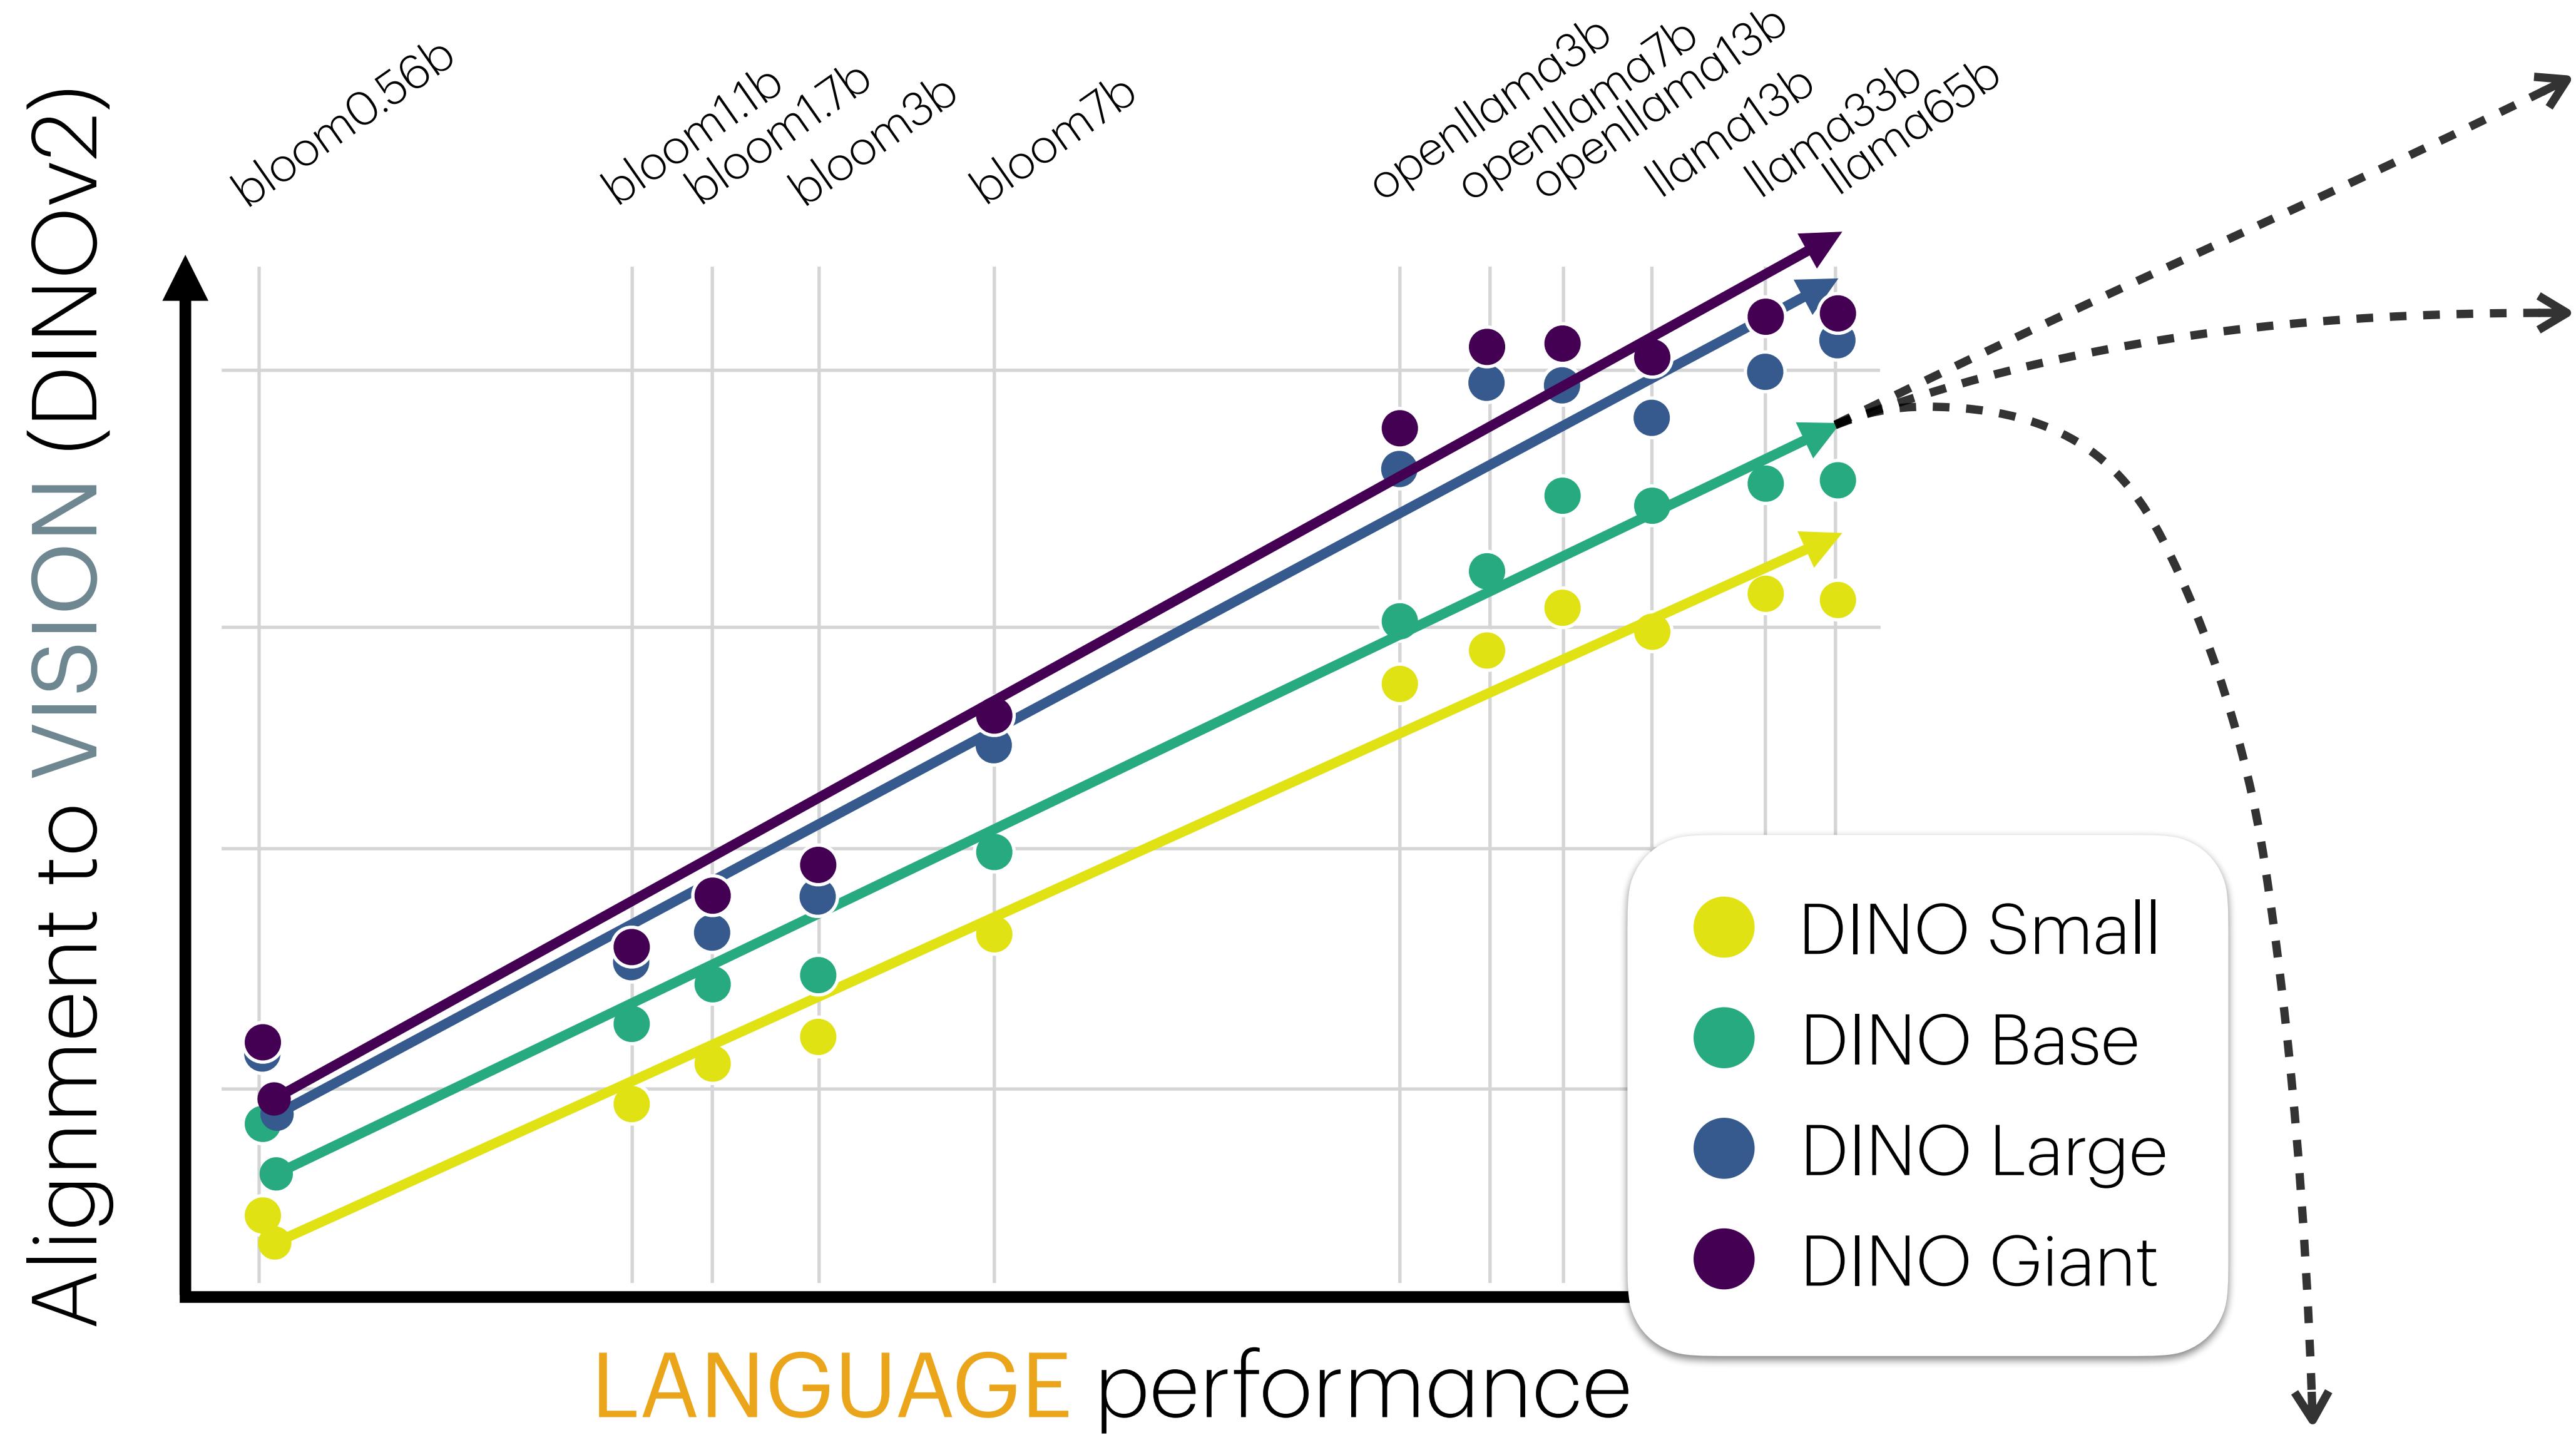
## Wikipedia Image Text Dataset

[Srinivasan, Raman, Chen, Bendersky, Najork 2021]

# Do strong language models align better with vision models?



# Strong models converge in representation



There's a large space of work that has shown similar results supporting our hypothesis

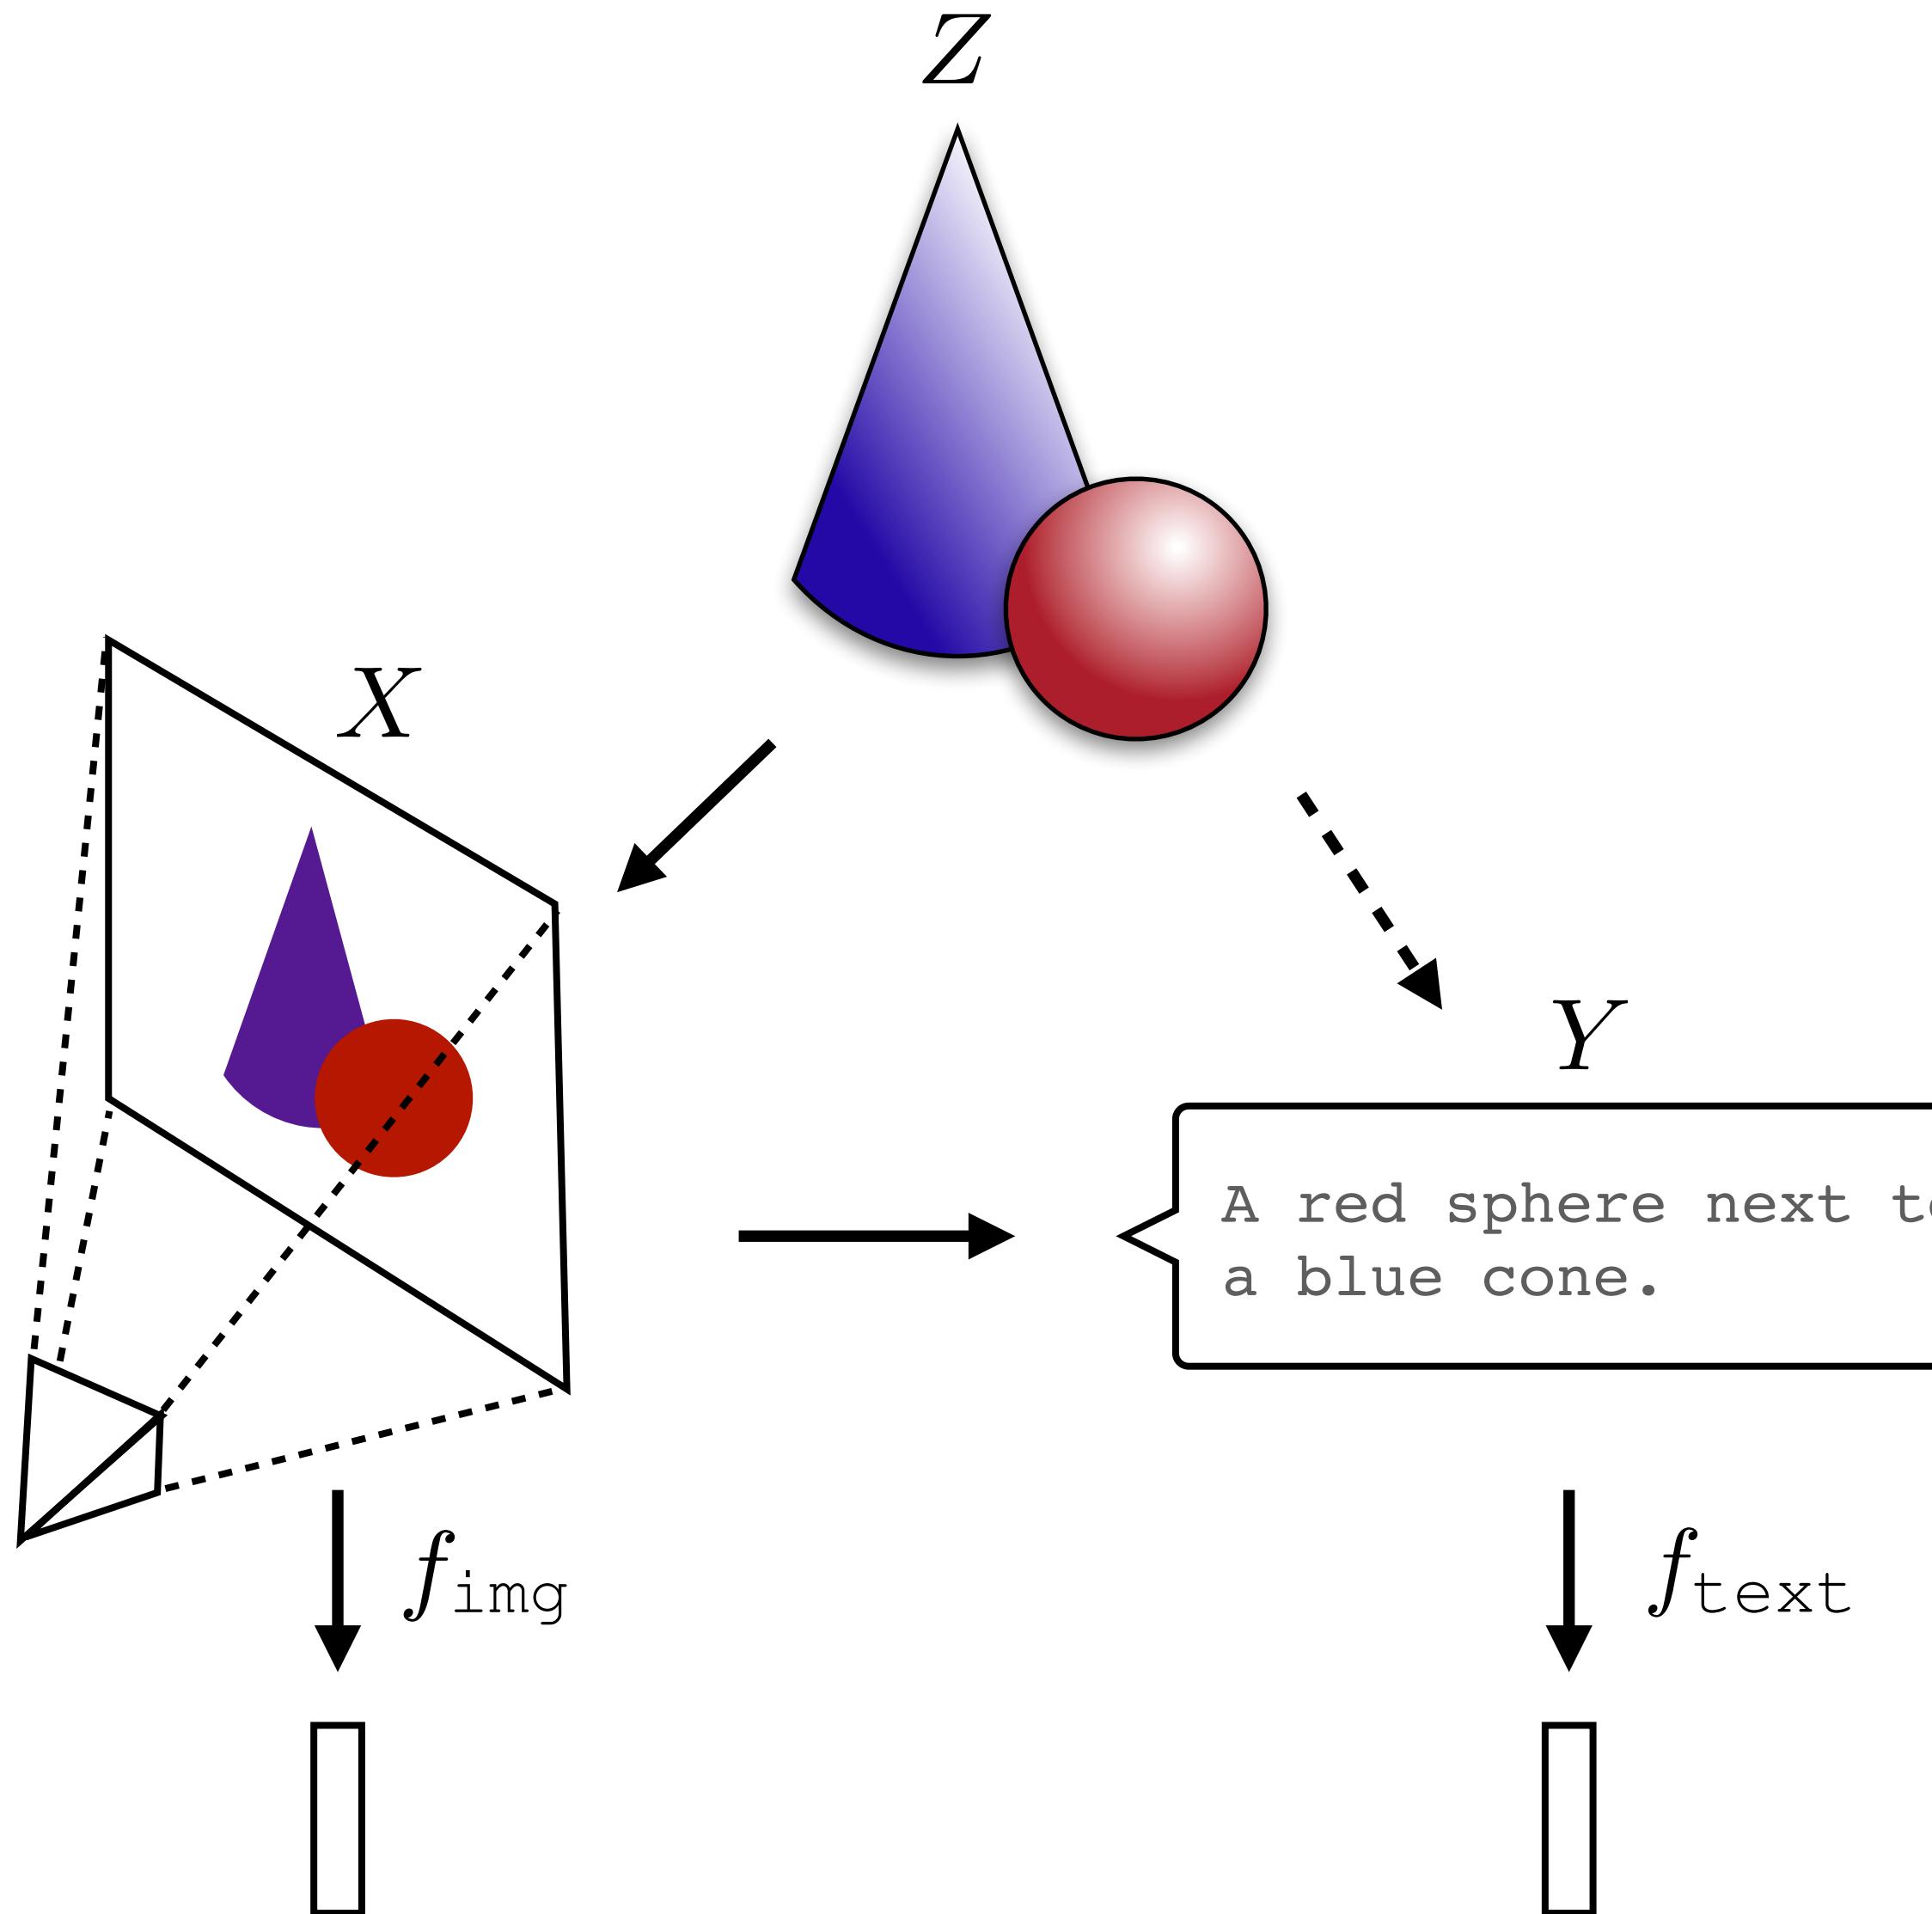
A collage of academic paper snippets and a journal cover. On the left, a snippet from a paper titled 'RELATIVE REPRESENTATIONS ENABLE ZERO-SHOT LATENT SPACE COMMUNICATION' by Luca Moschella and Valentino Maiorca is shown, with a note that it was 'Published as a conference paper at ICLR 2023'. In the center, a snippet from a paper titled 'Emergence of simple-cell receptive field properties by learning a sparse code for natural images' by Bruno A. Olshausen and David J. Field is shown. On the right, a snippet from a paper titled 'Explanatory models in neuroscience' by Rosa Cao and Daniel Yamins is shown, with a table of contents for the journal issue. The table of contents includes: 'Opinion' (2), 'Up Intelligibility' (4), 'Aren't Mere Curve-Fits' (7), 'the no-miracles argument' (10), 'tation' (12), and 'ciple' (17). The journal cover features a large image of a grid of simple cell receptive fields.

# Why is this happening?

- ~~It's all about the data!~~
- ~~It's all about transformers!~~
- *It's all about the world*

[Plato 375 BC]





*Platonic representation hypothesis*

---

Different neural networks are converging toward the **same** way of representing the world (same kernel).

# How many words is a picture worth?



5 words

"Illuminated escalators in indoor mall."

10 words

"Illuminated escalators in a lush indoor mall with greenery and polished tiles."

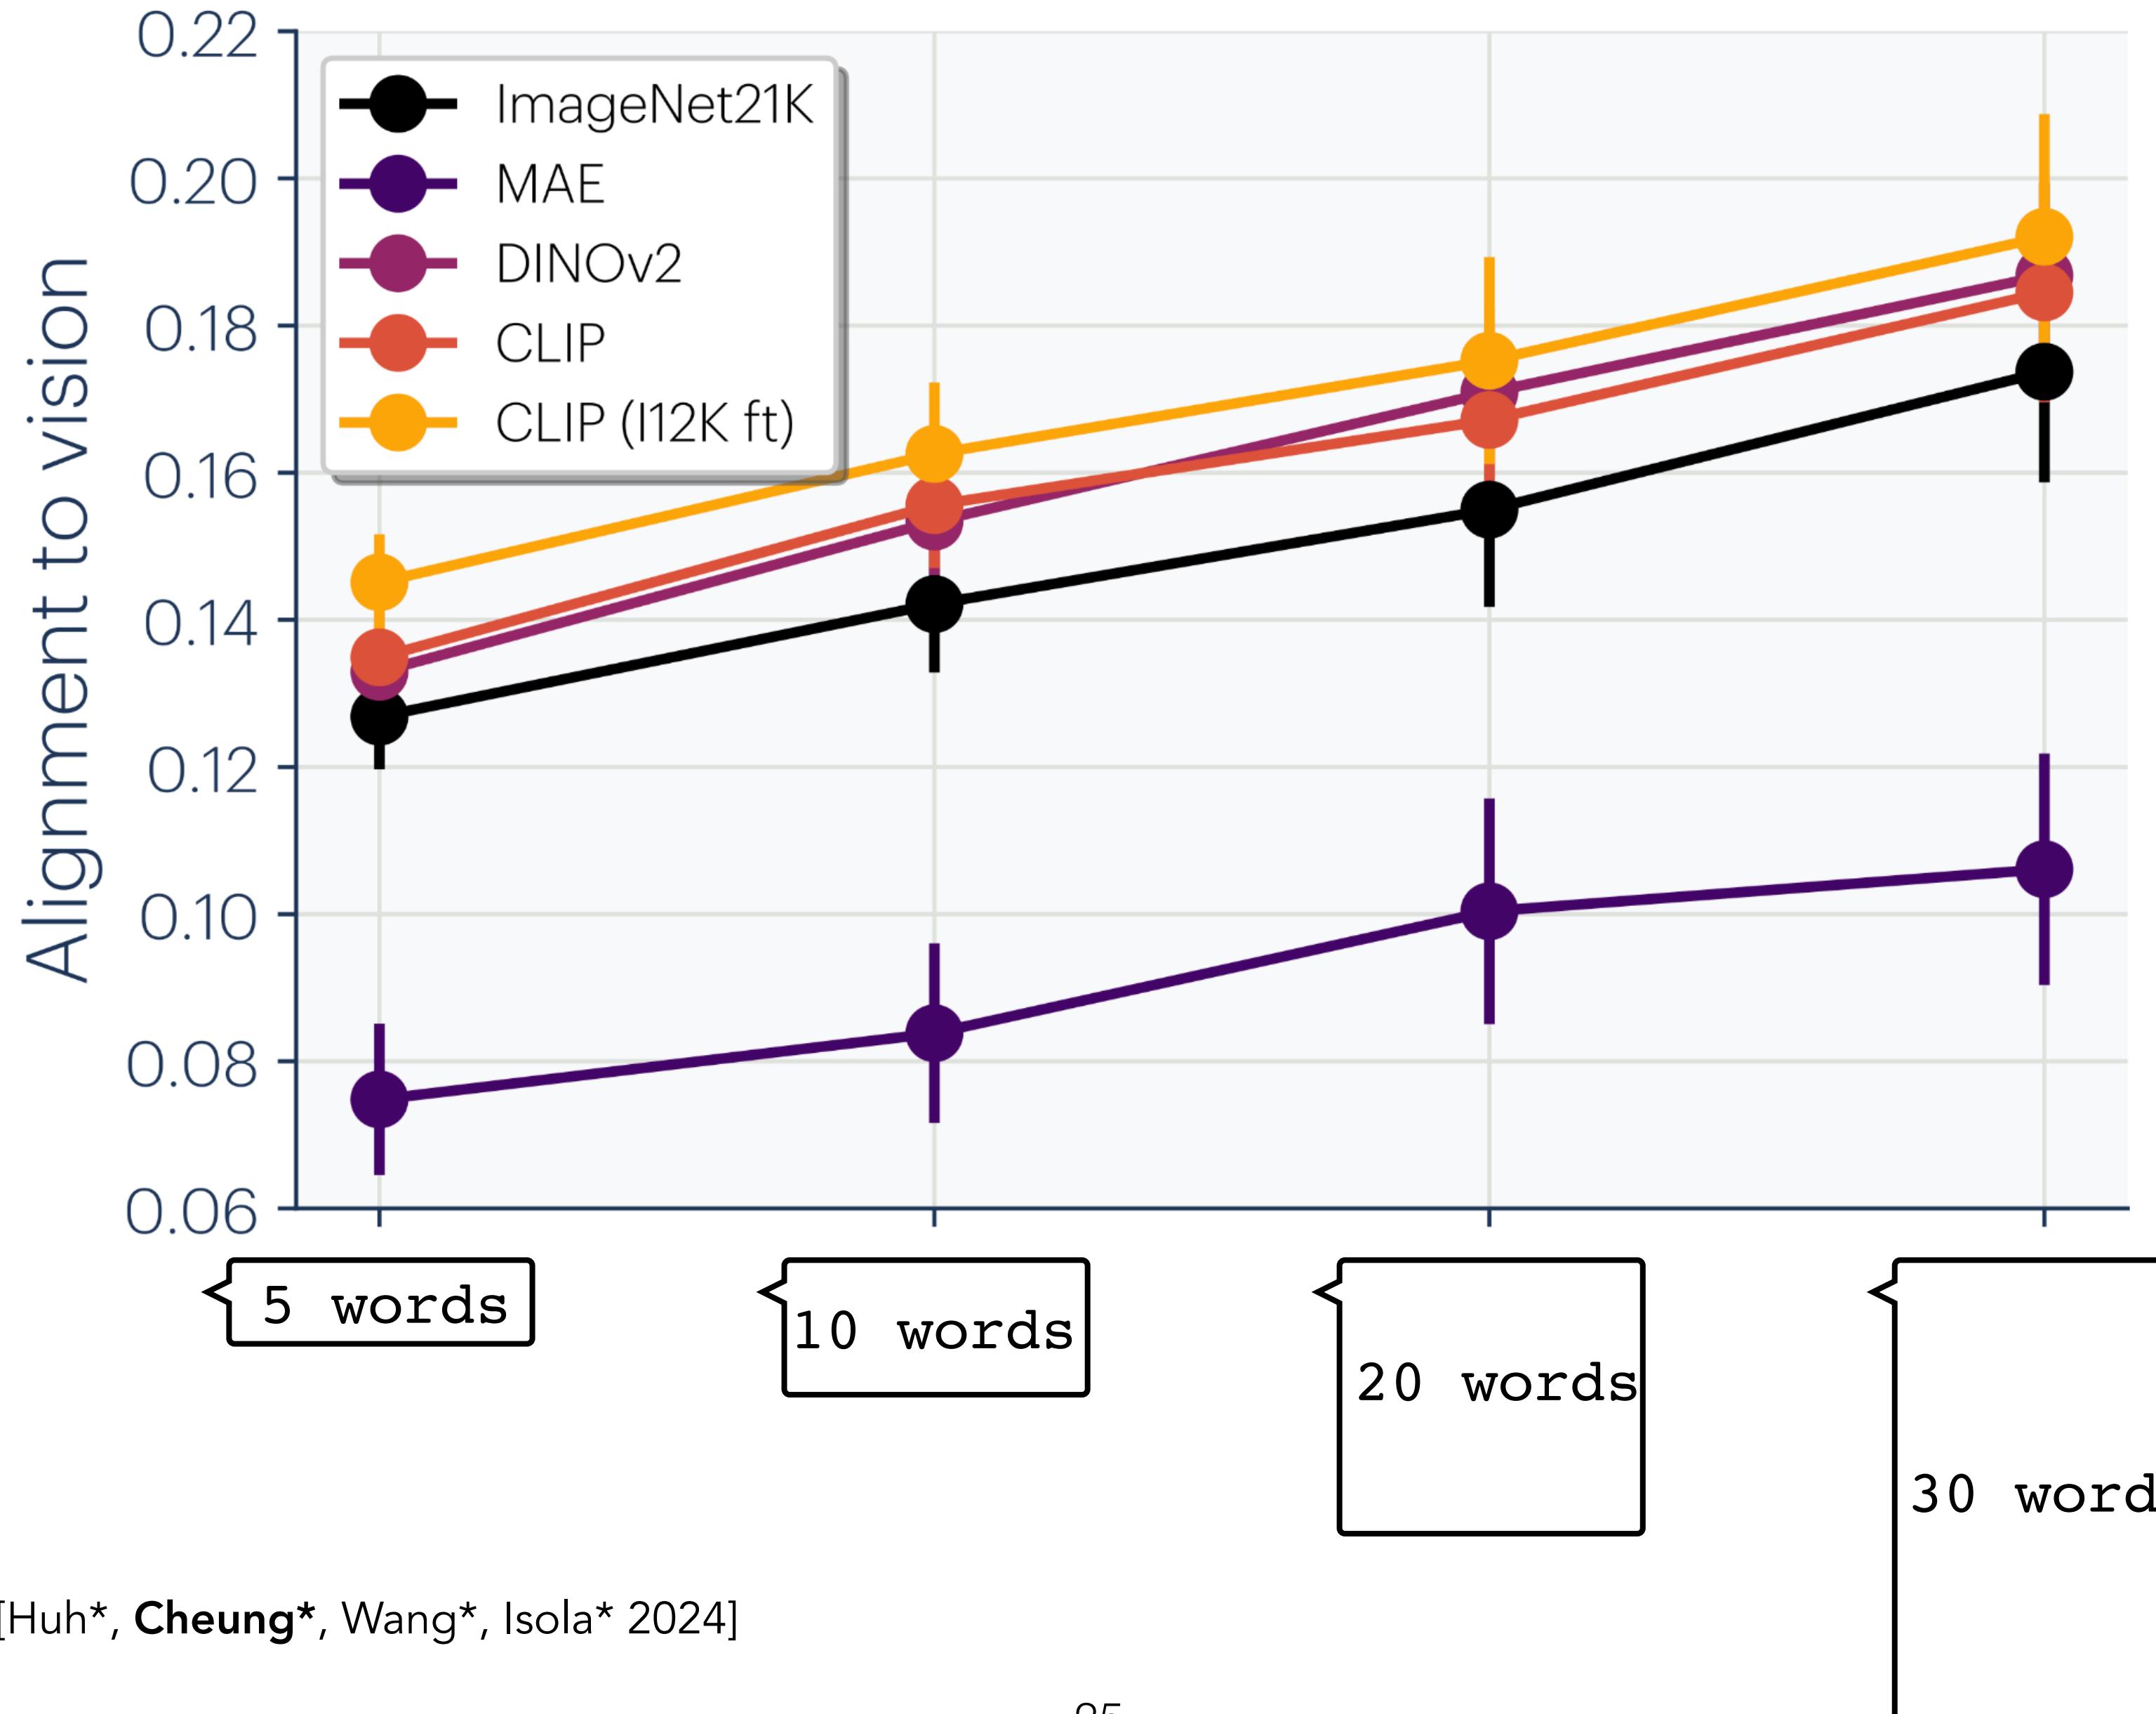
20 words

"Illuminated escalators in a vibrant indoor mall with lush greenery, polished tiles, and tropical plants, surrounded by stone columns and recessed lighting."

30 words

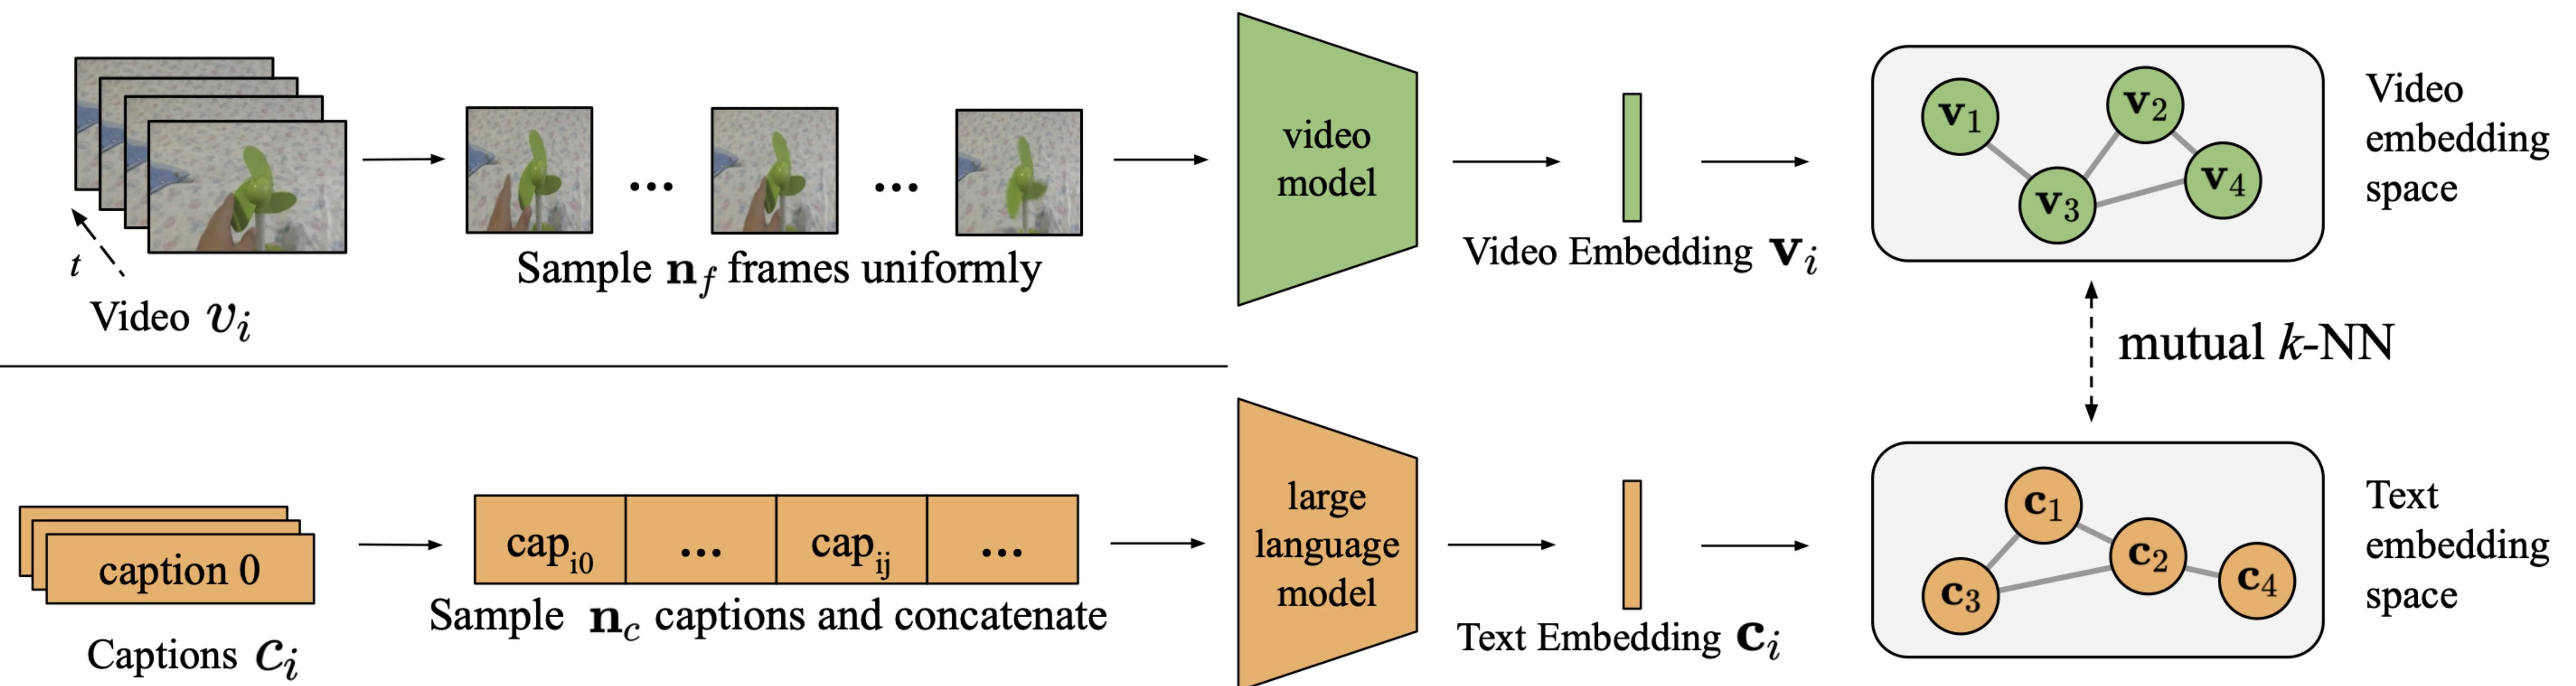
"An indoor shopping mall with three illuminated escalators, surrounded by lush greenery and polished colored tiles, featuring a man ascending one escalator and various shops and plants on the upper level."

# How many words is a picture worth?

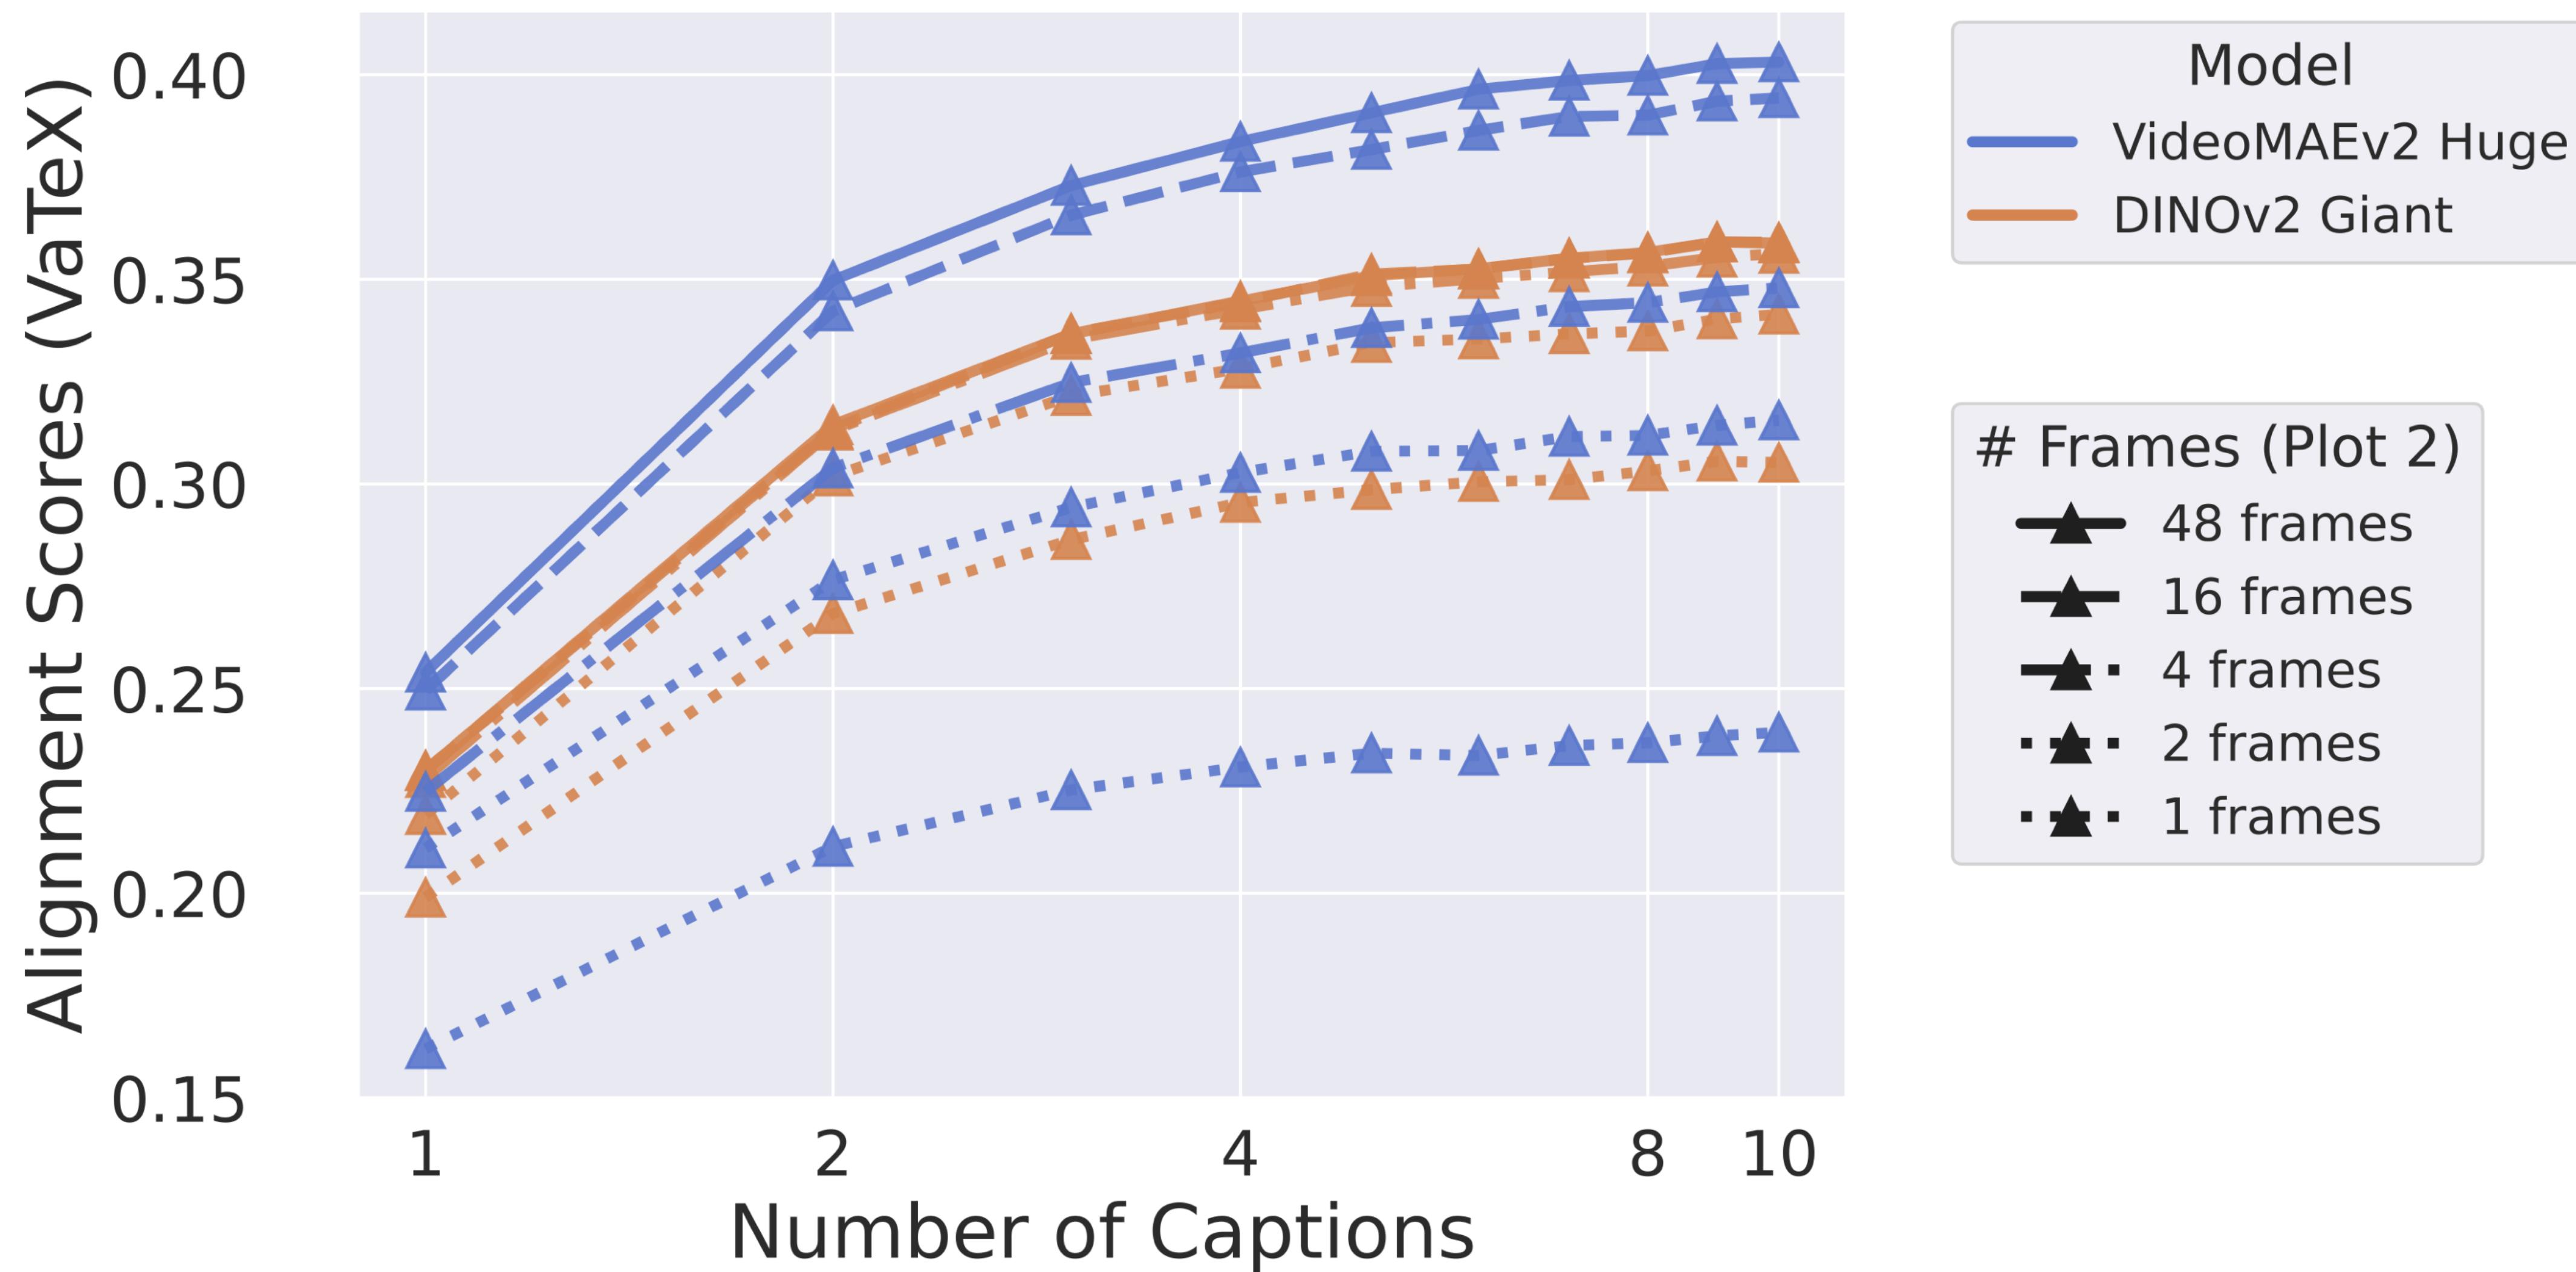


# DYNAMIC REFLECTIONS: PROBING VIDEO REPRESENTATIONS WITH TEXT ALIGNMENT

**Tyler Zhu<sup>†\*</sup> Tengda Han<sup>‡</sup> Leonidas Guibas<sup>‡</sup> Viorica Pătrăucean<sup>‡</sup> Maks Ovsjanikov<sup>‡</sup>**  
Princeton University<sup>†</sup> Google DeepMind<sup>‡</sup>



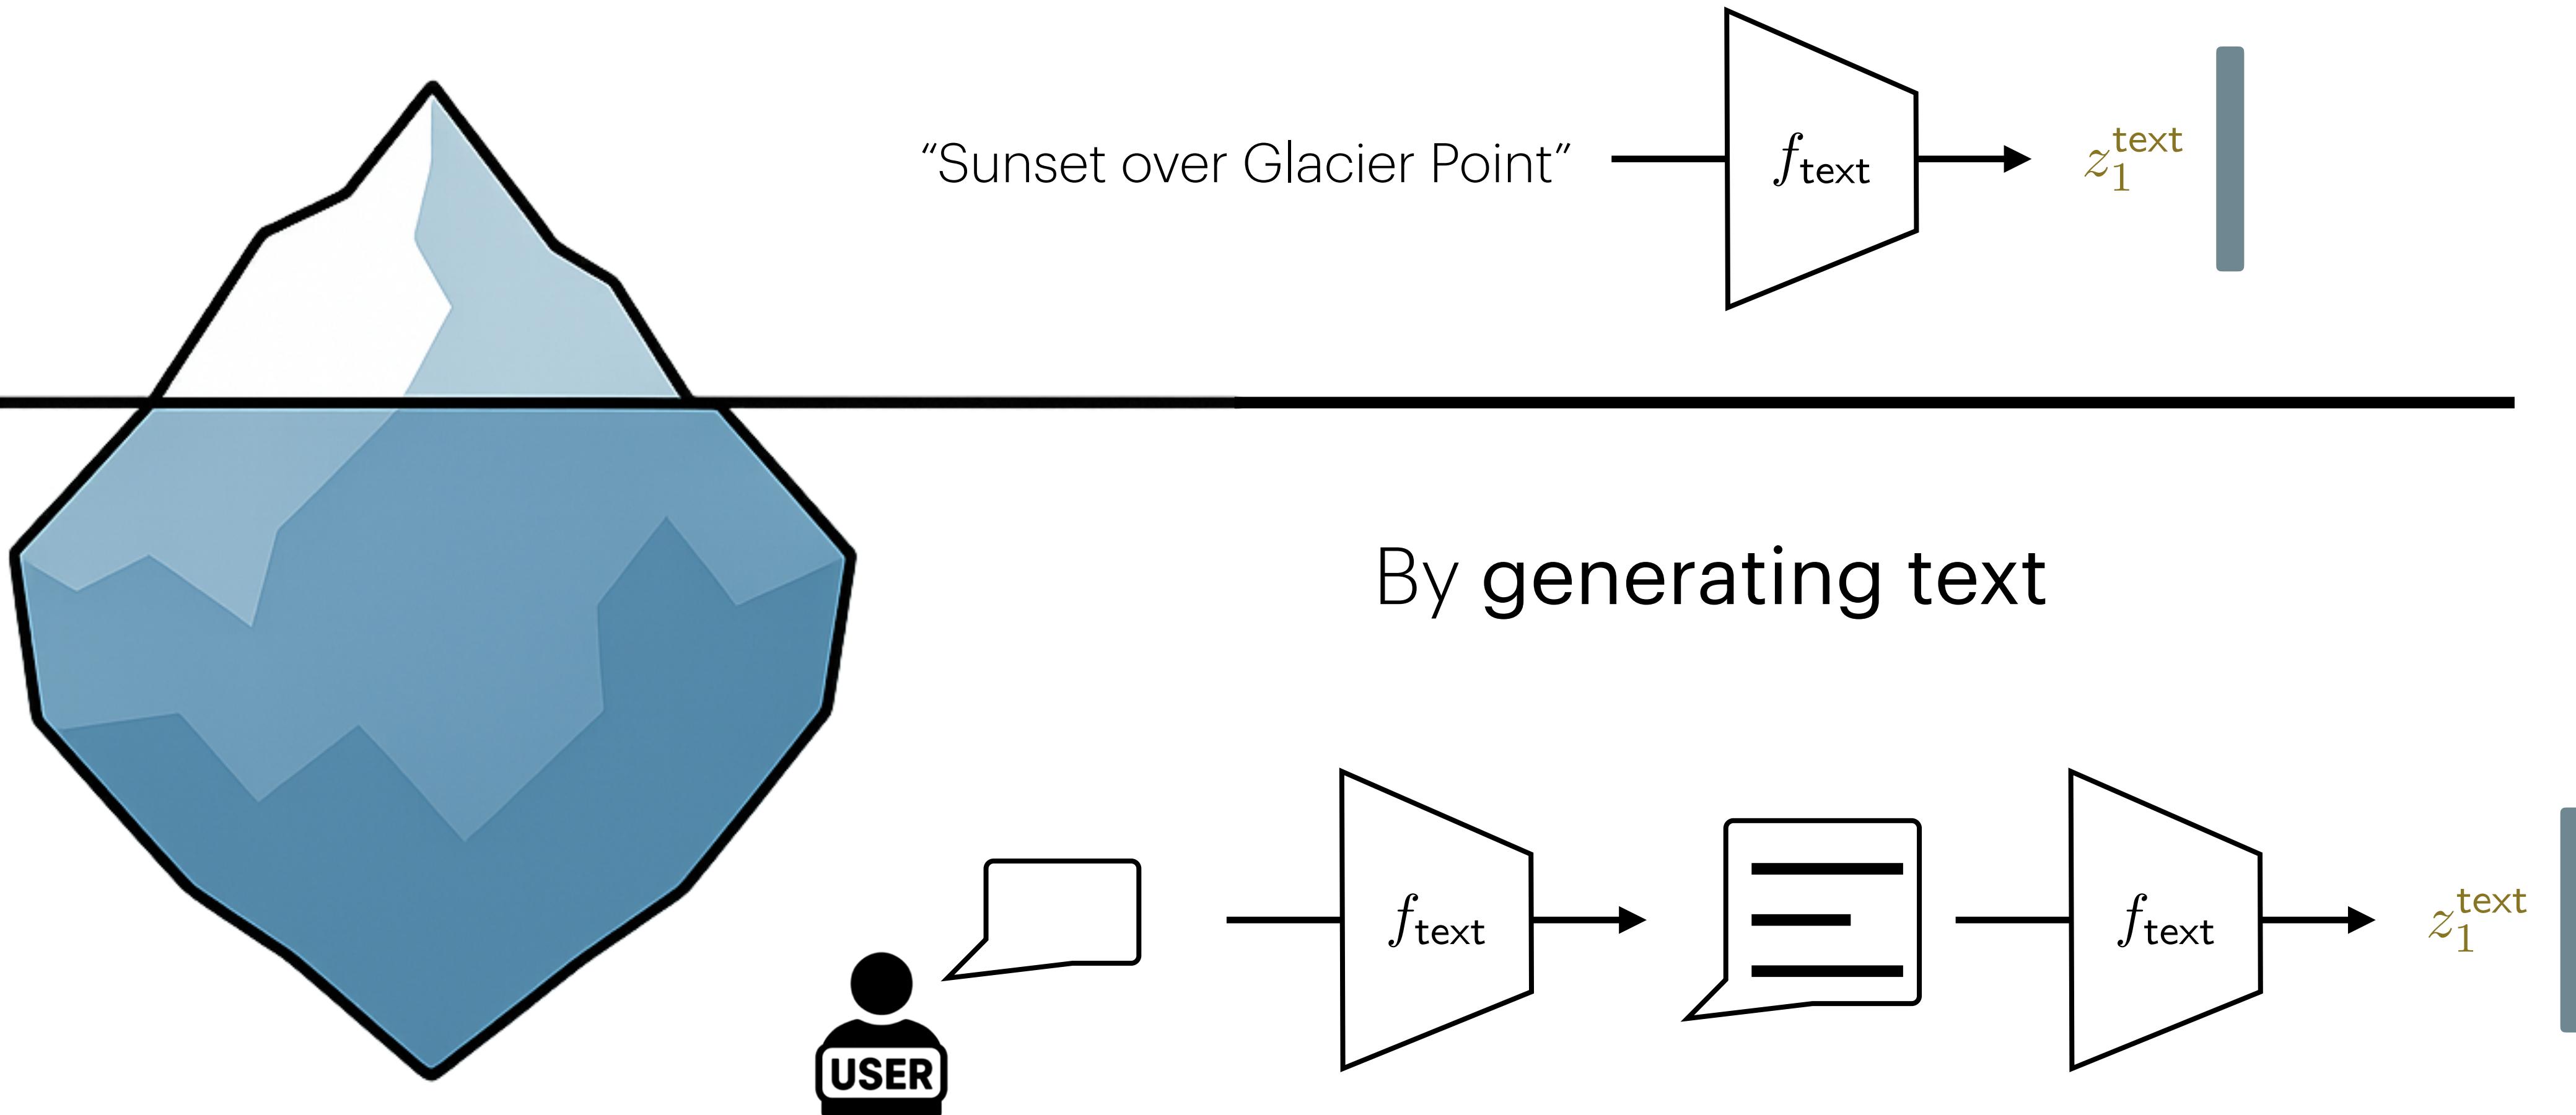
# How many captions is a video worth?



[Zhu et al. 2025]

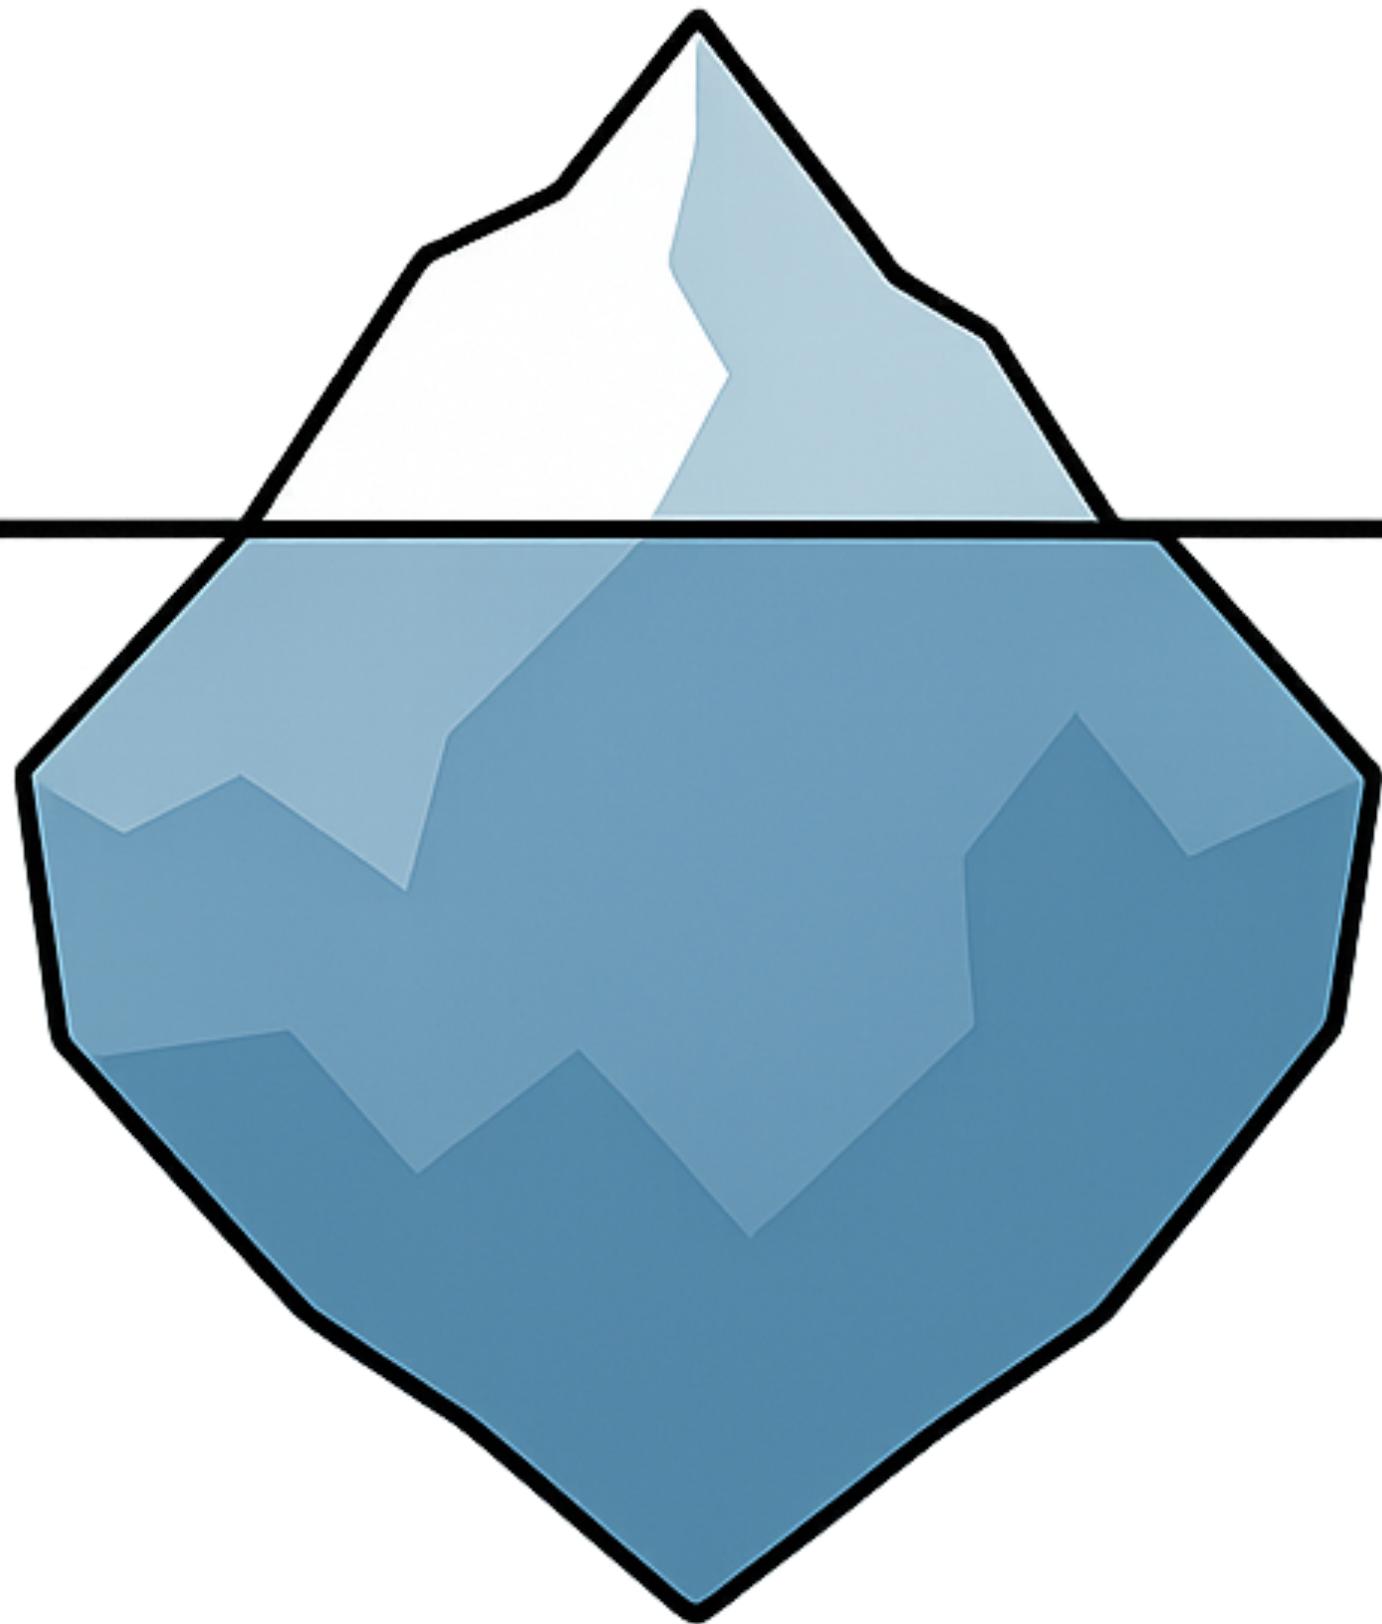
# Probing what language models know

By embedding text

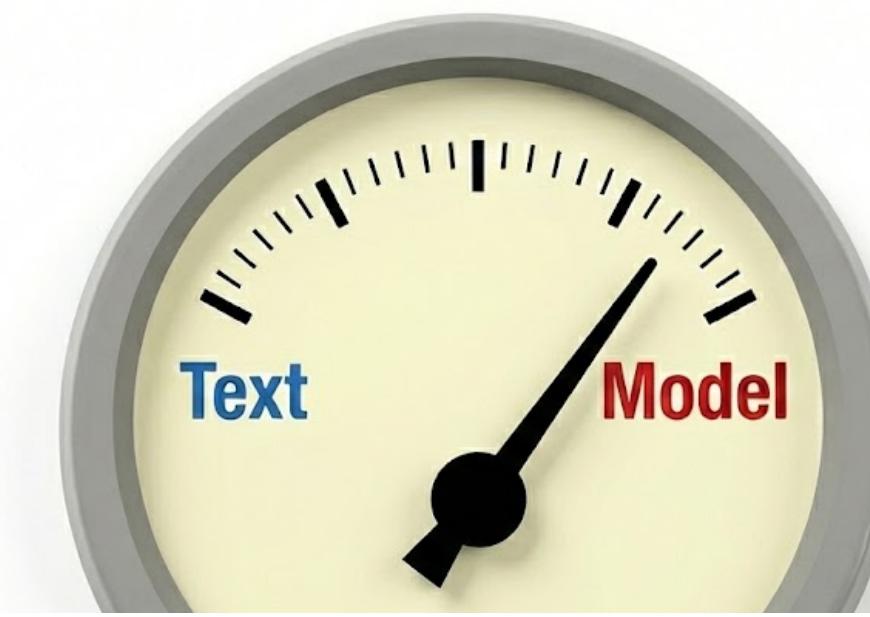


# Influence on model representation

By embedding text



By generating text



“You can just do things.”

—Sam Altman, CEO of OpenAI



You can just ask for things.

# Chain-of-Thought Prompting

## Zero-shot

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A: The answer (arabic numerals) is

---

## Zero-shot-CoT

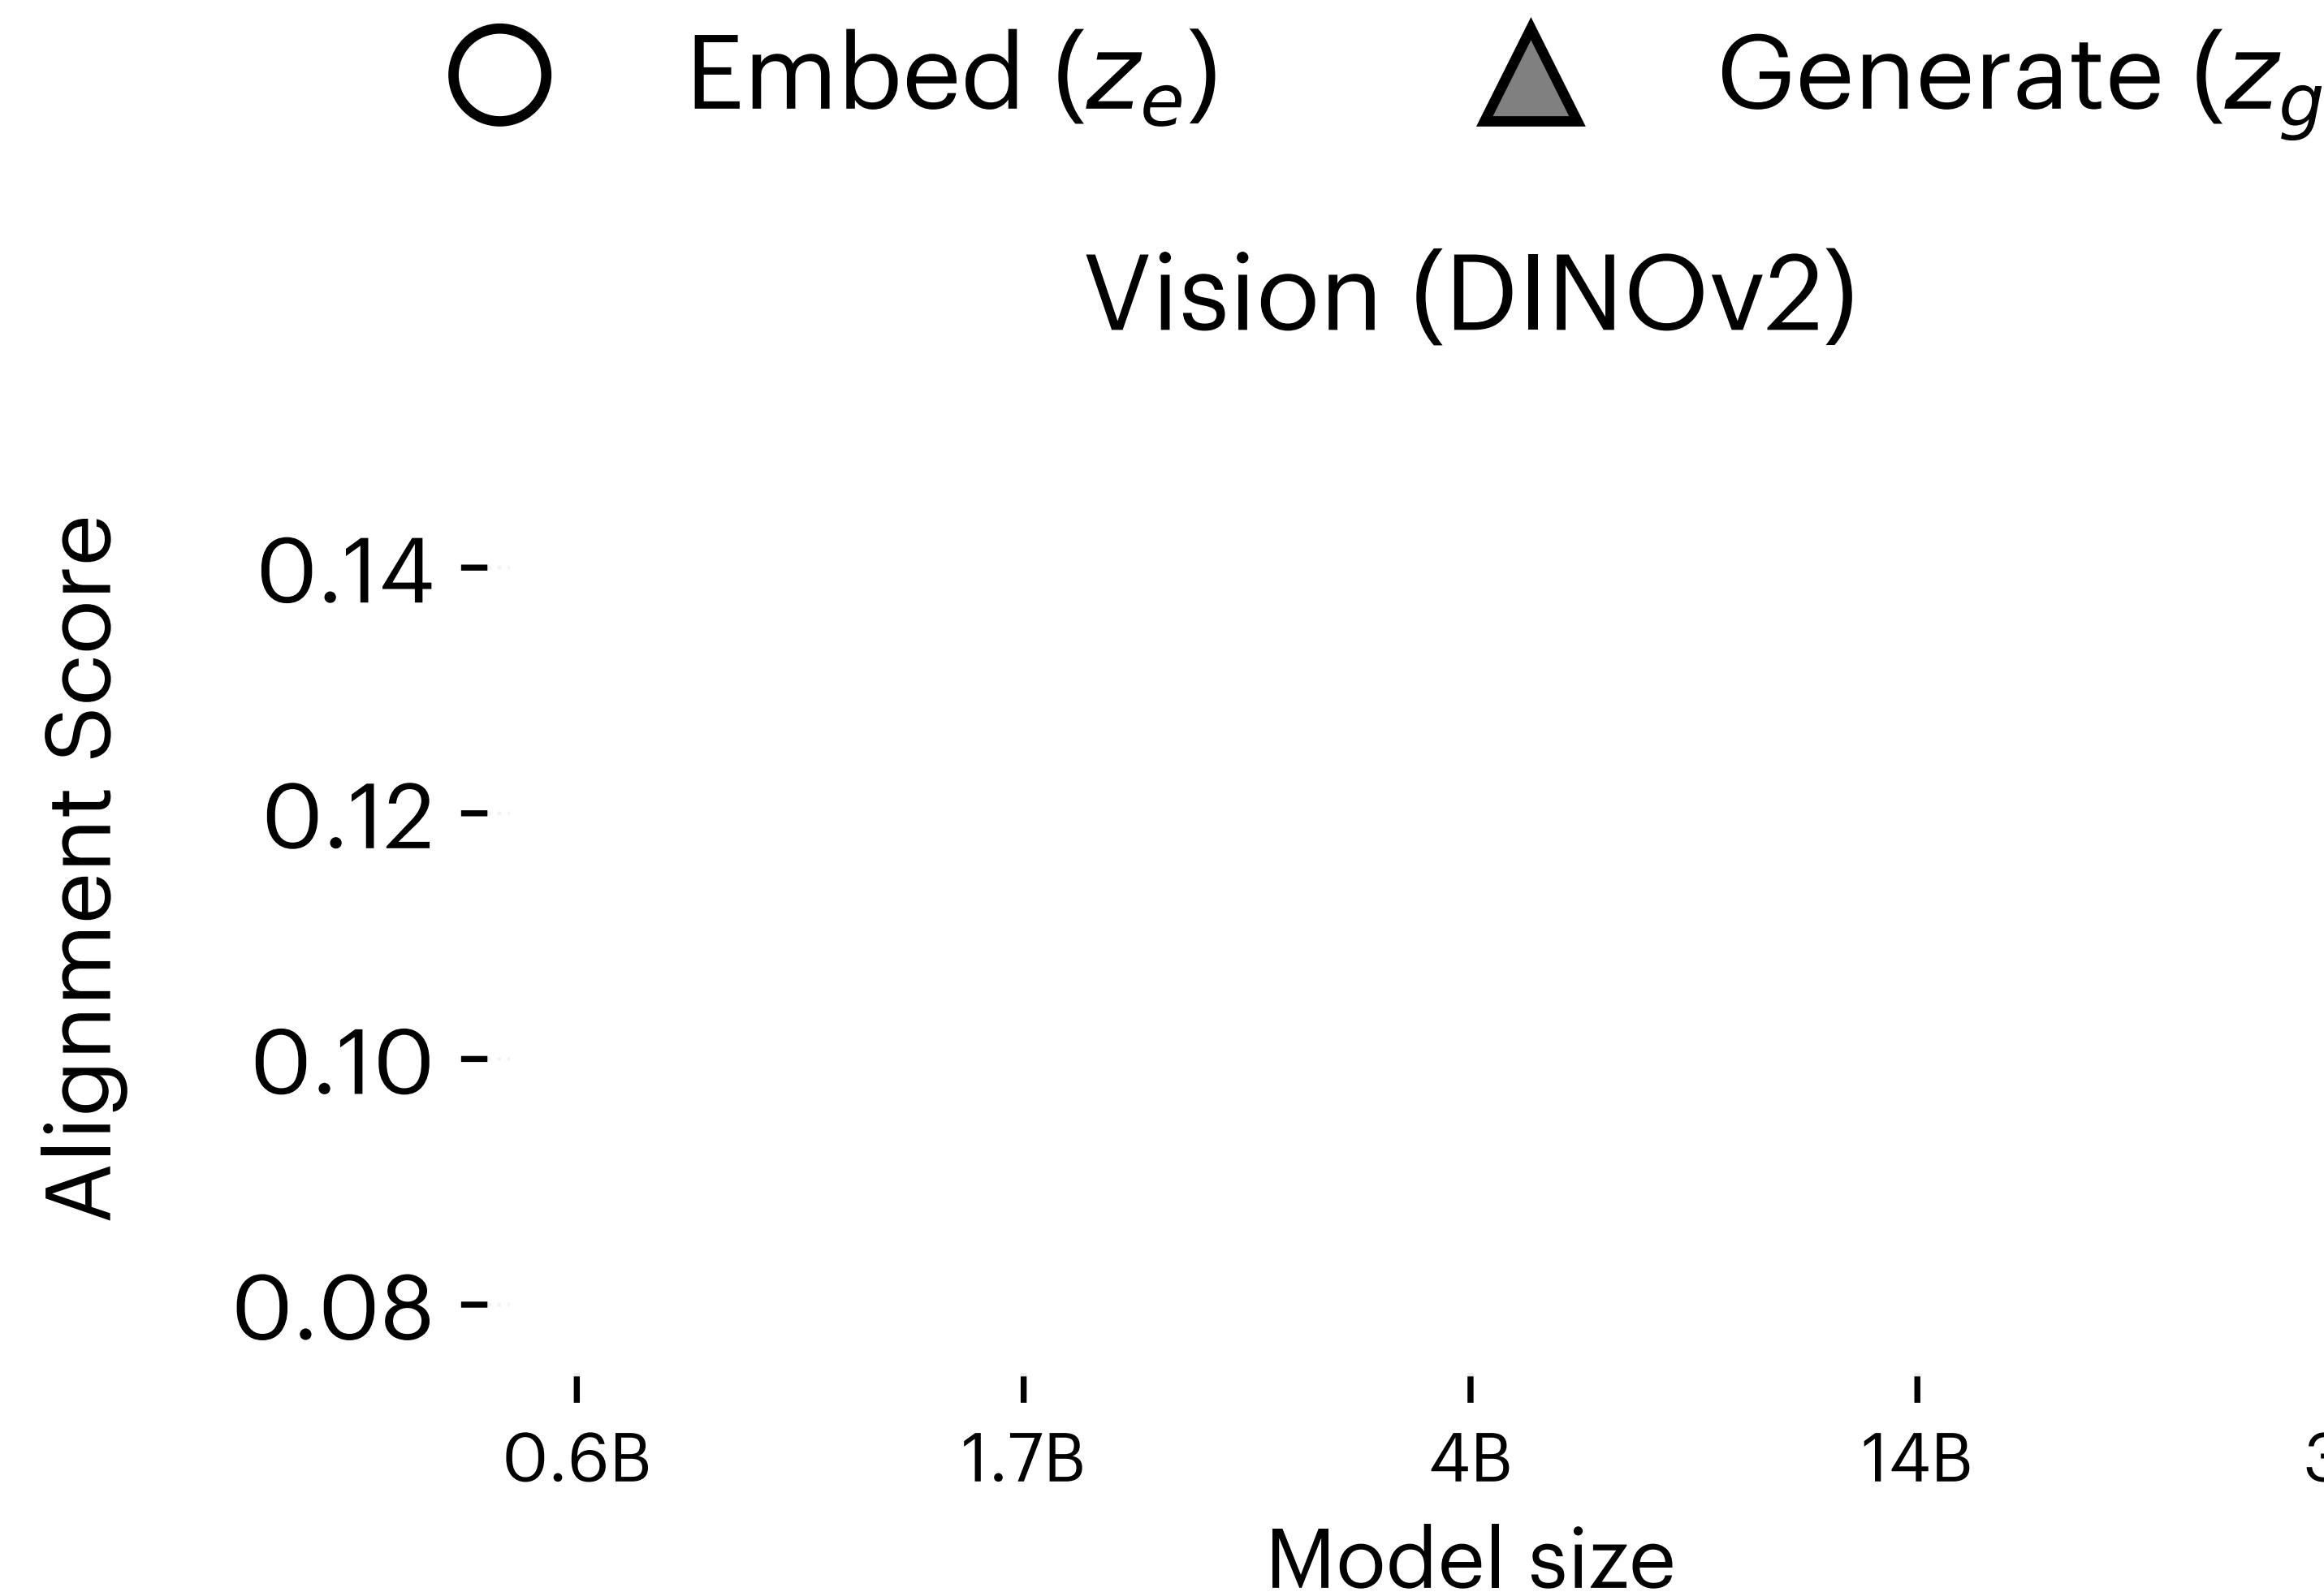
Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A:

---

[Kojima et al. 2022]

Asking the model to Imagine: {caption}.

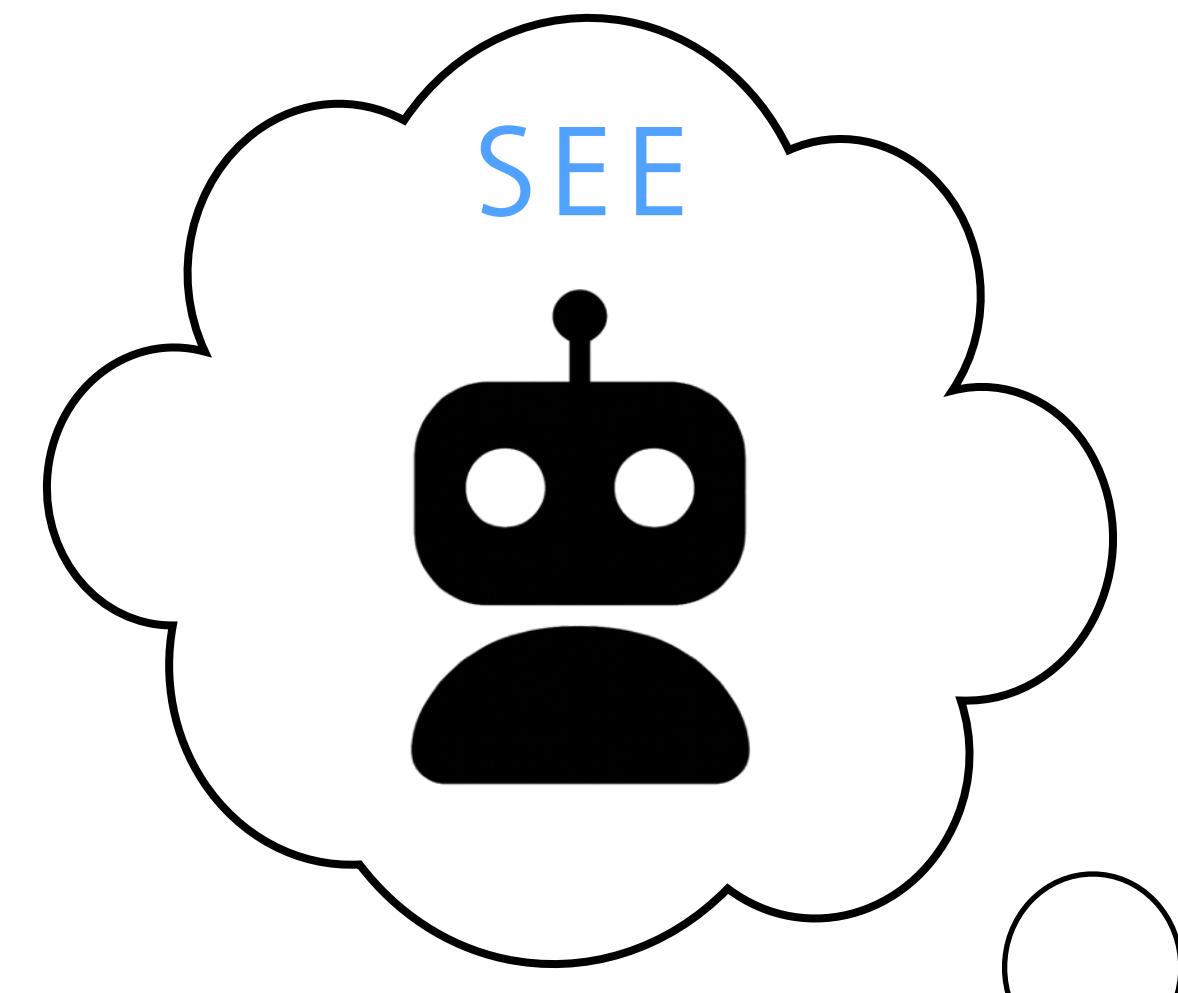


[Wang, Isola, **Cheung\*** 2025]

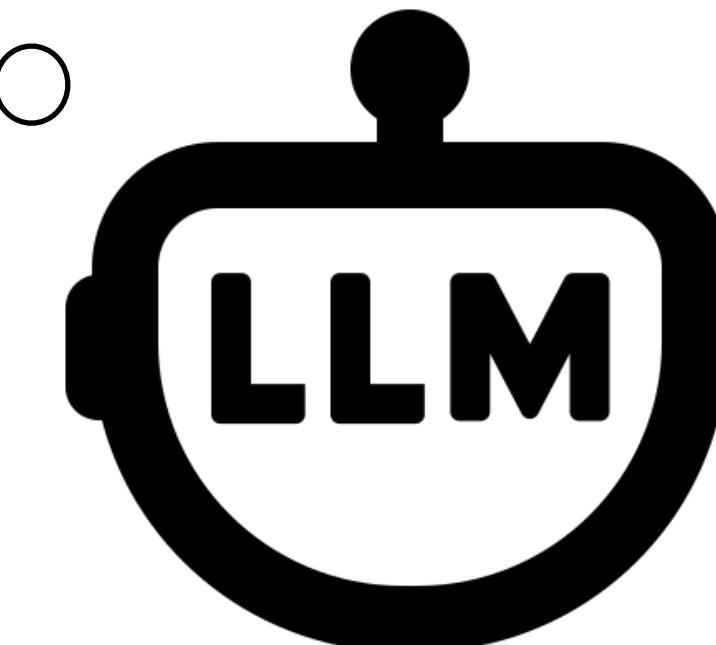
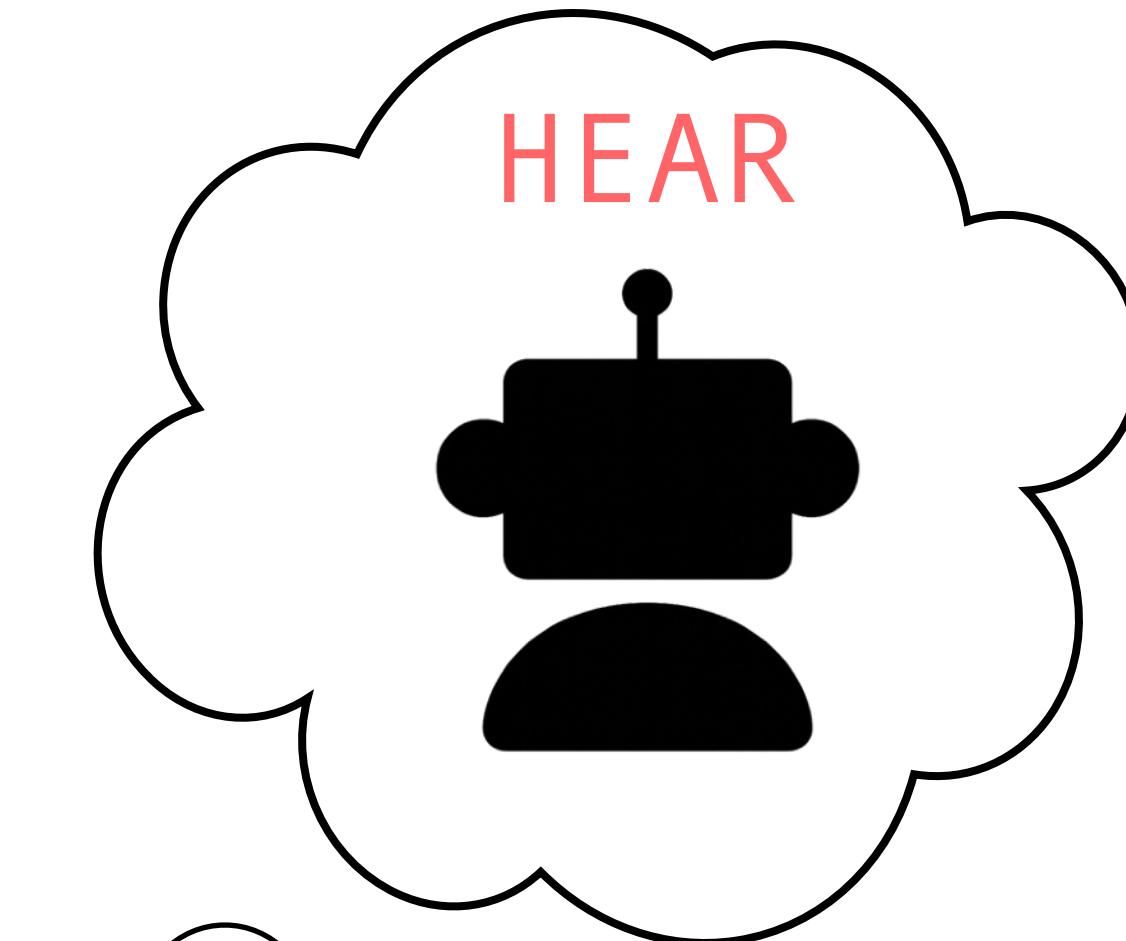
# What happens when you ask a language model to imagine senses it never experienced?



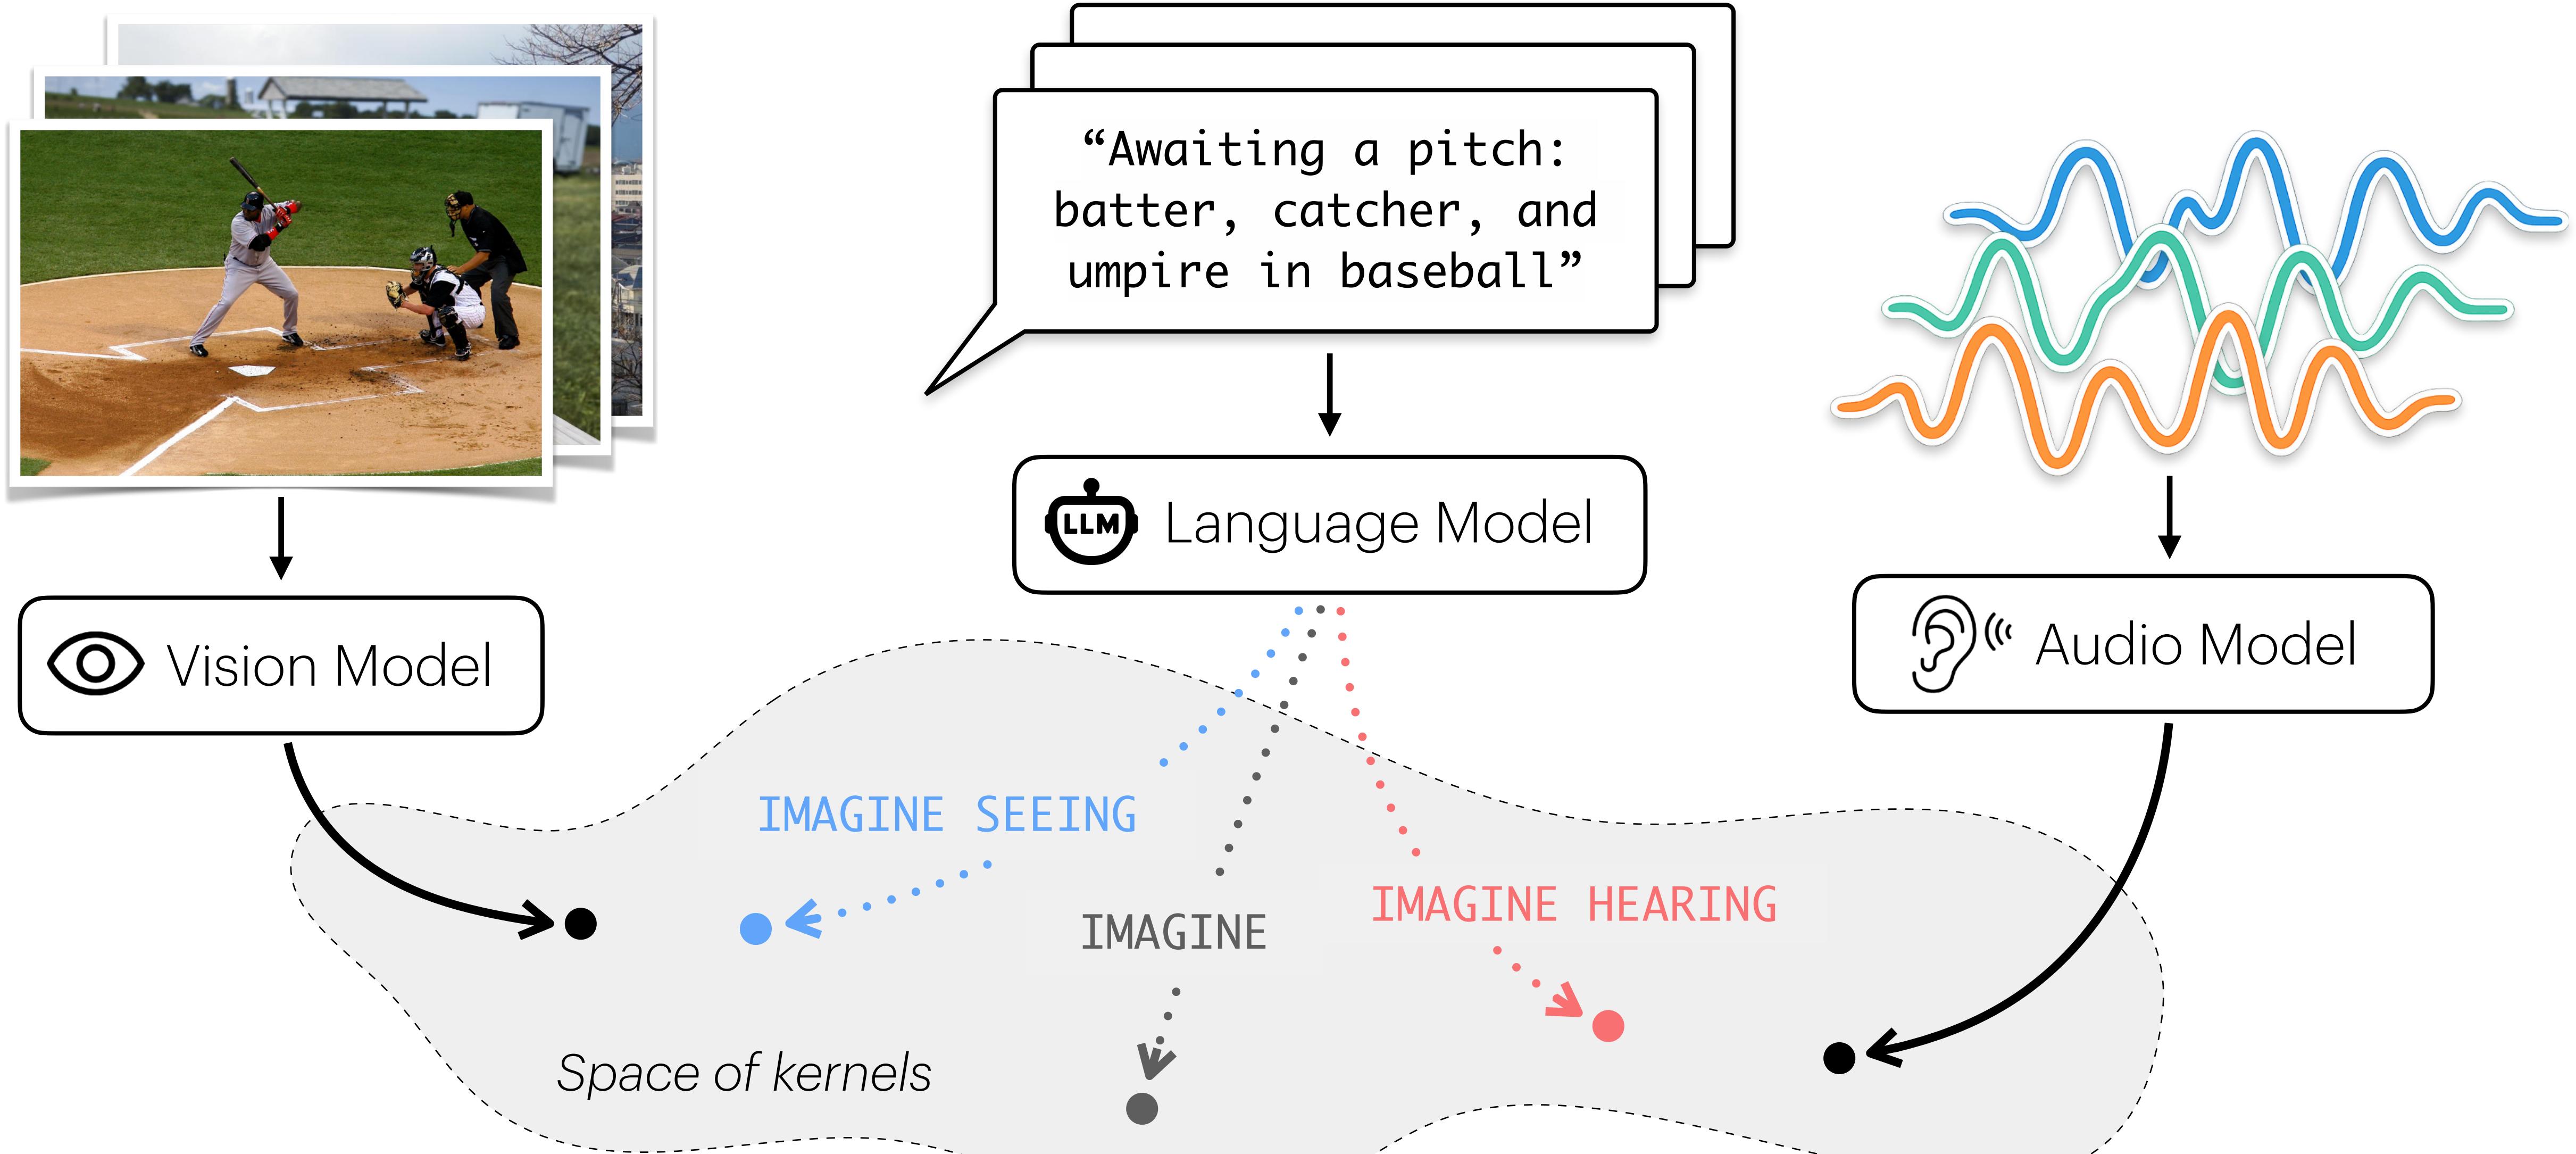
Imagine what it would **look** like to **see** {caption}.

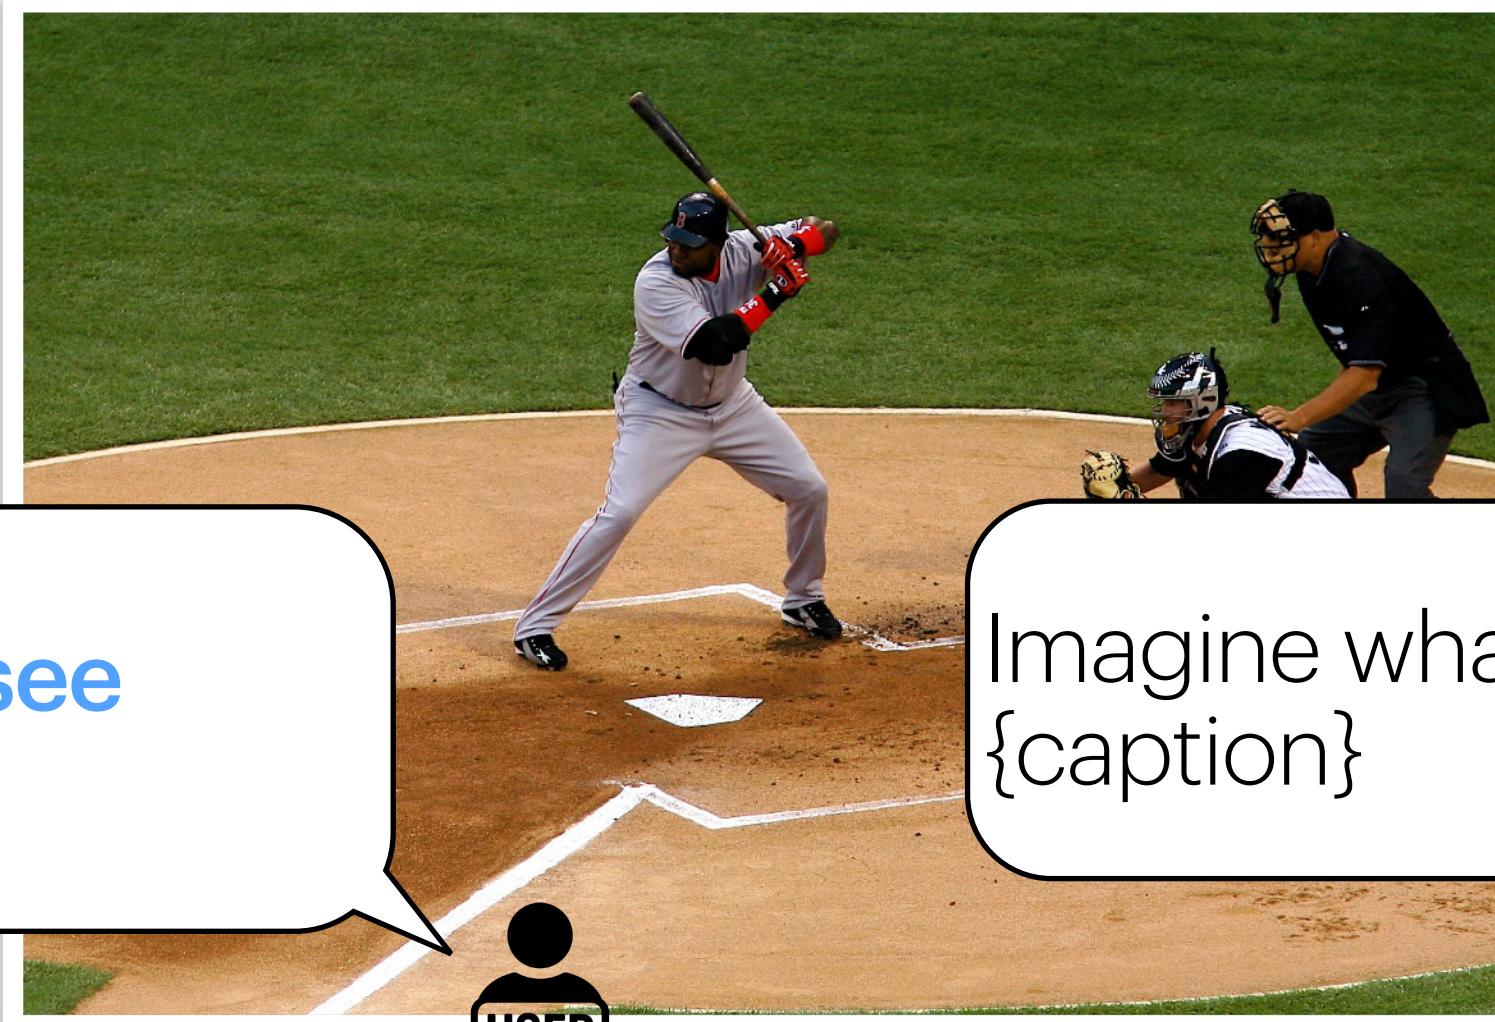


Imagine what it would **sound** like to **hear** {caption}.



# Sensory Prompting





Imagine what it would **look** like to **see**  
{caption}

Imagine what it would **sound** like to **hear**  
{caption}

...I need to **visualize** a baseball field. The batter stands at the home plate, right? They're **holding the bat**, probably in a ready position, feet shoulder-width apart. Their **body language** should show they're focused...

USER

USER

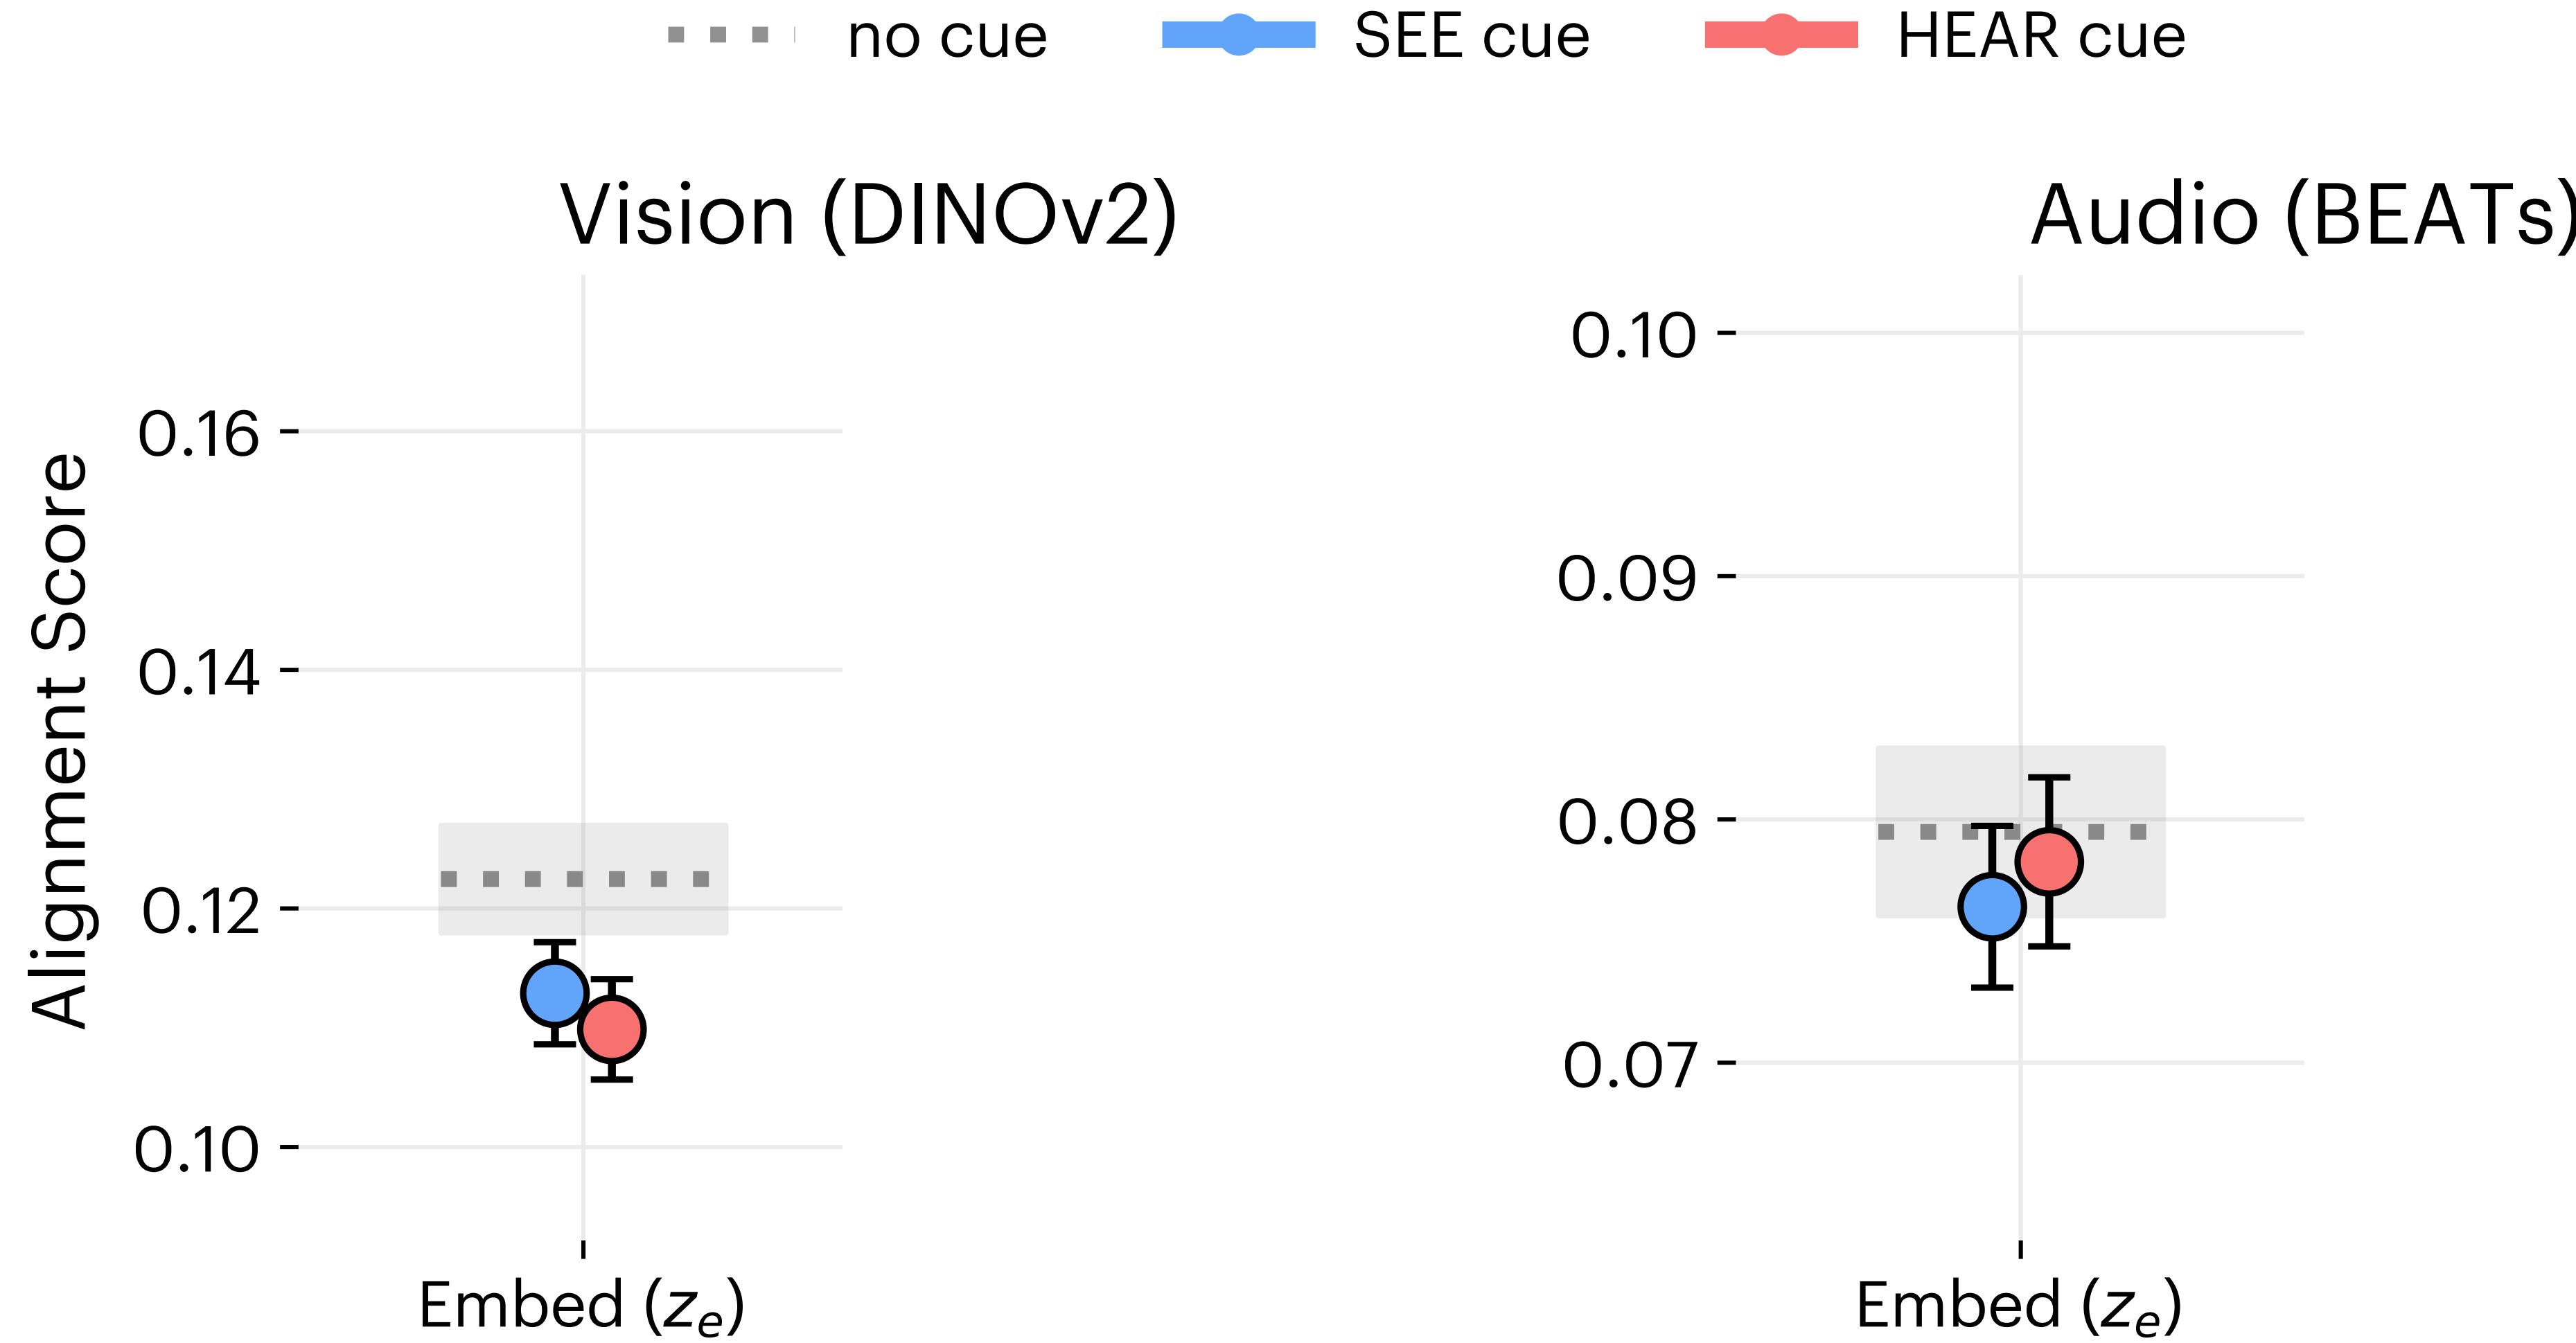
batter, catcher

...There's the general **ambient noise** of the stadium: **crowd murmurs**, maybe some distant **cheers or chatter**. The **sound of players moving**, like cleats scraping on the dirt, gloves being adjusted...

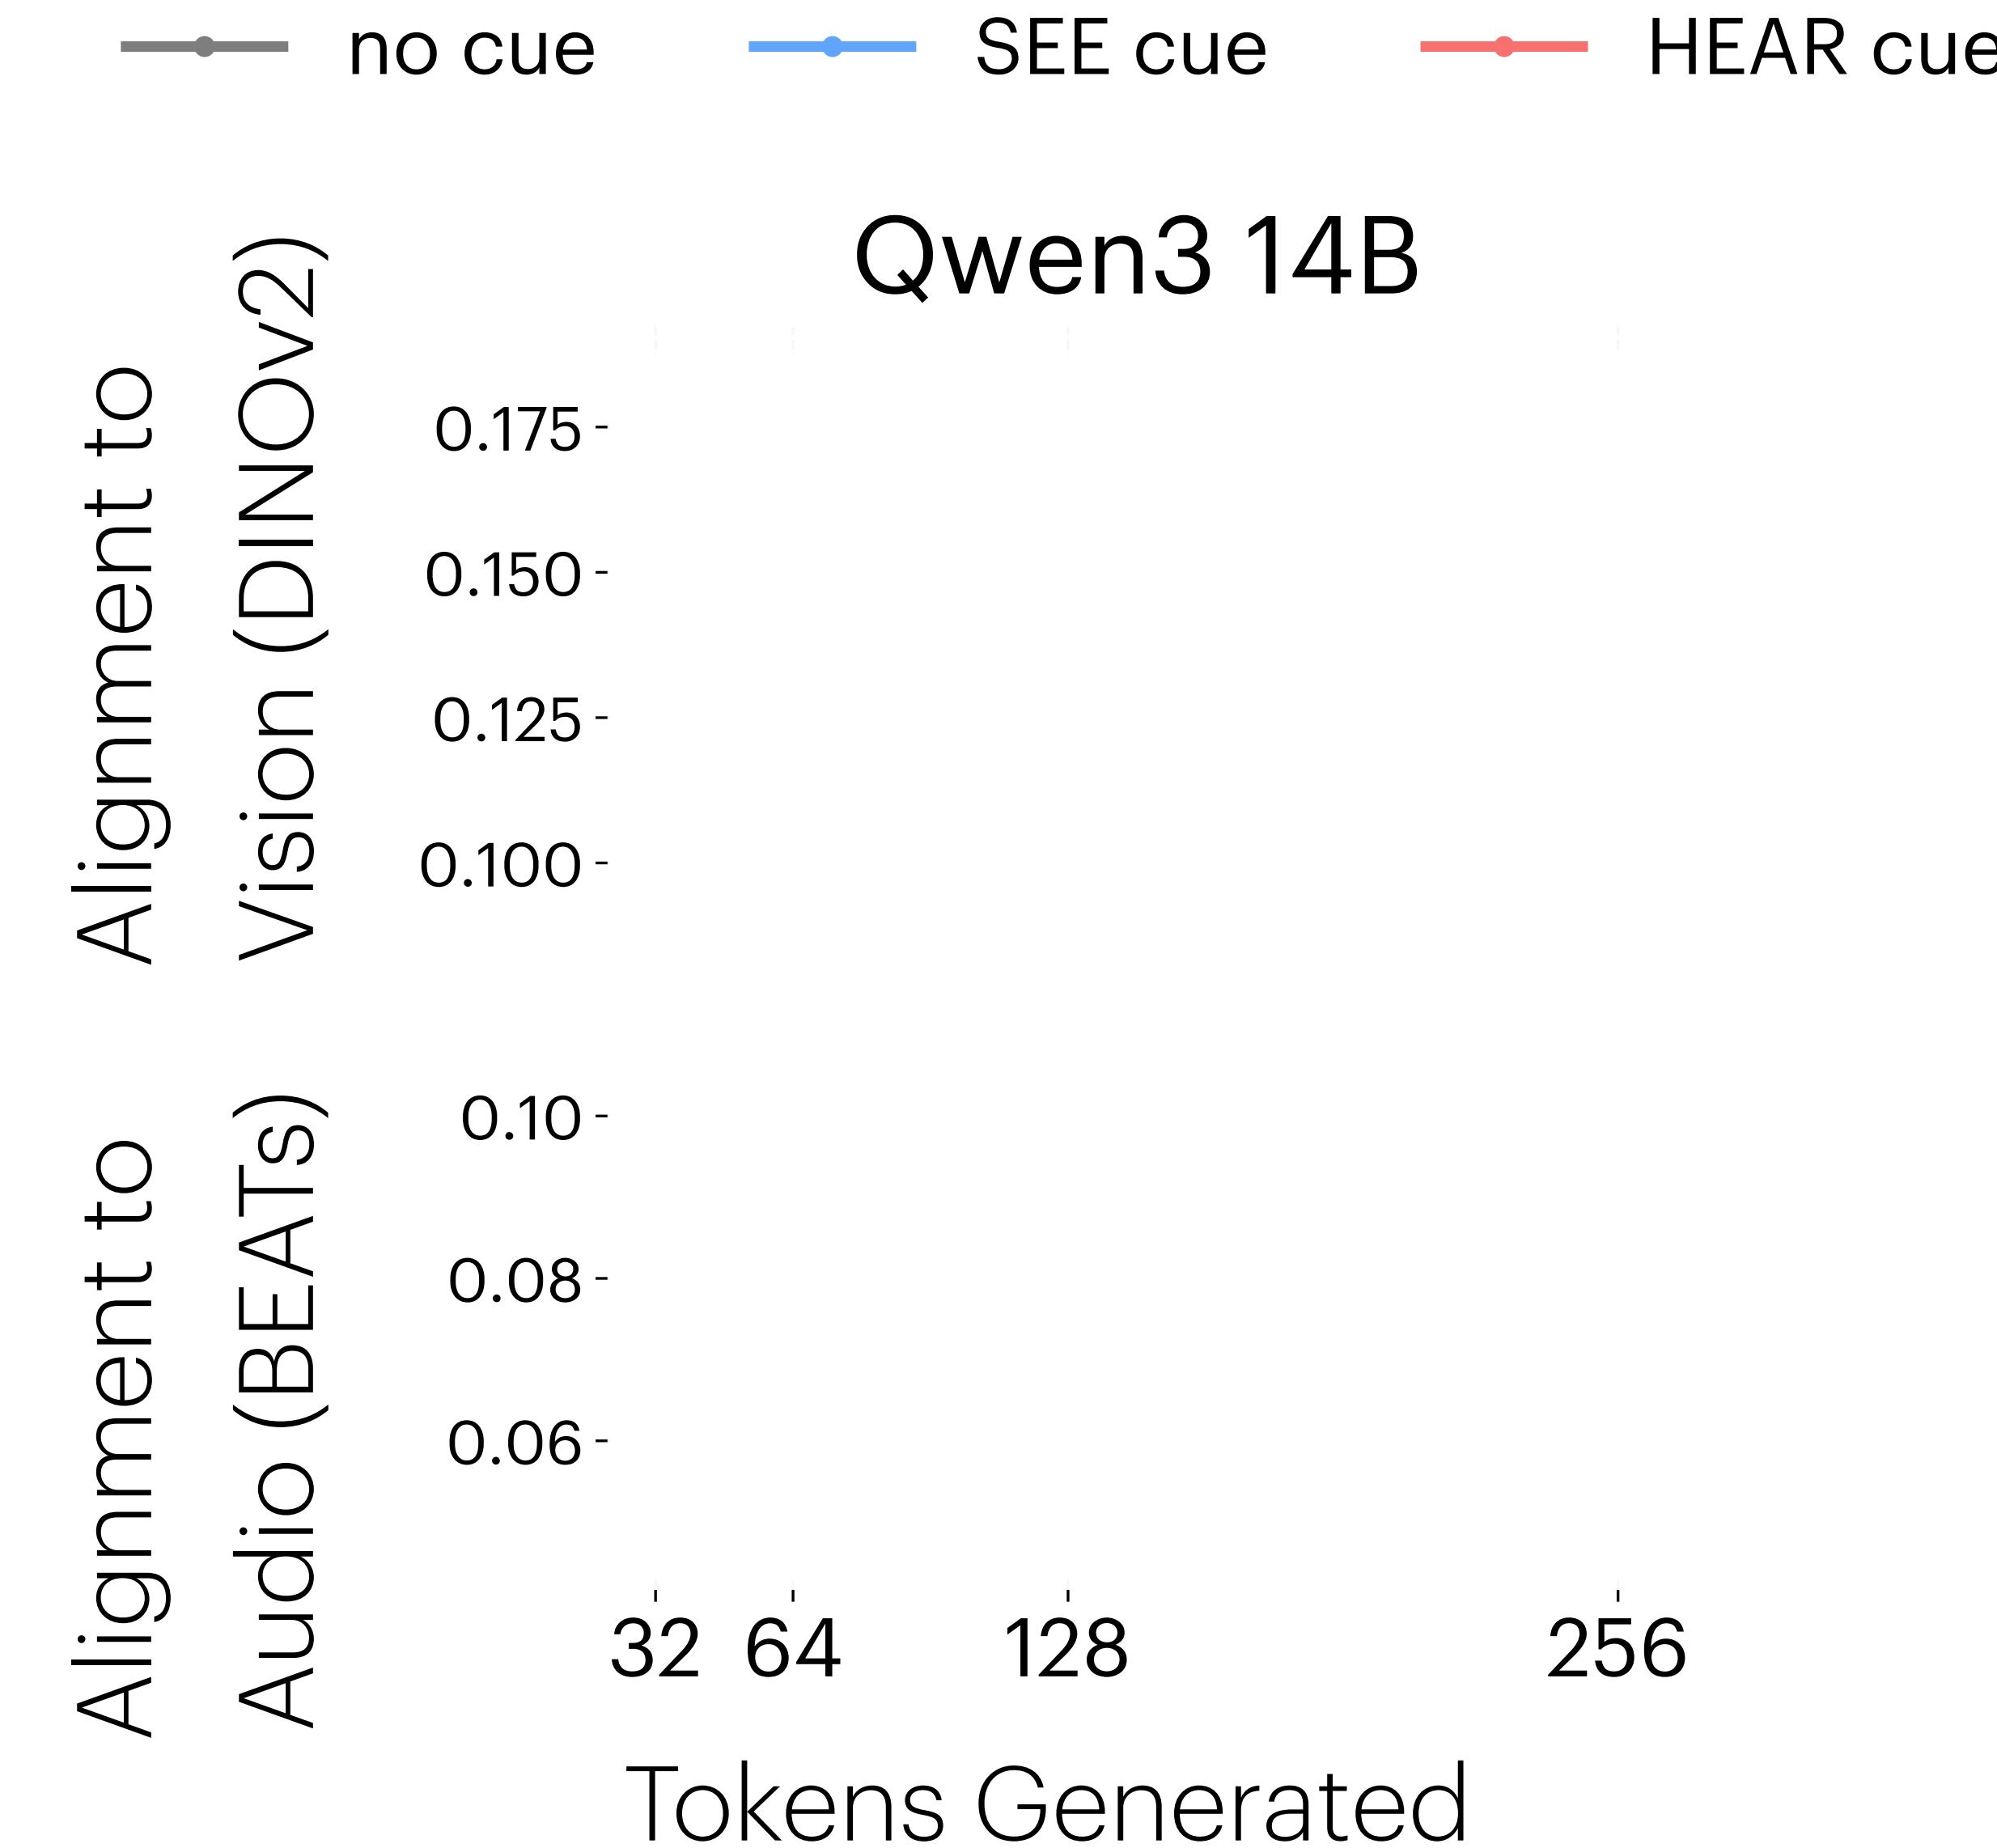
LLM

LLM

# Sensory Embedding vs Sensory Generation



# How many spoken words is a picture worth?



# Redirecting Sensory Cues

## Seeing to Hearing

Please rewrite the following text to describe what the scene would sound like **audibly**, rather than how it **looks**: {text}

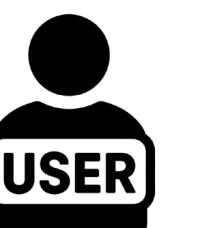


The scene **hums** with anticipation, the only sounds the **soft rustle of grass** underfoot and the occasional **murmur** of the **crowd** in the distance. A tense **silence** hangs in the **air**...



## Hearing to Seeing

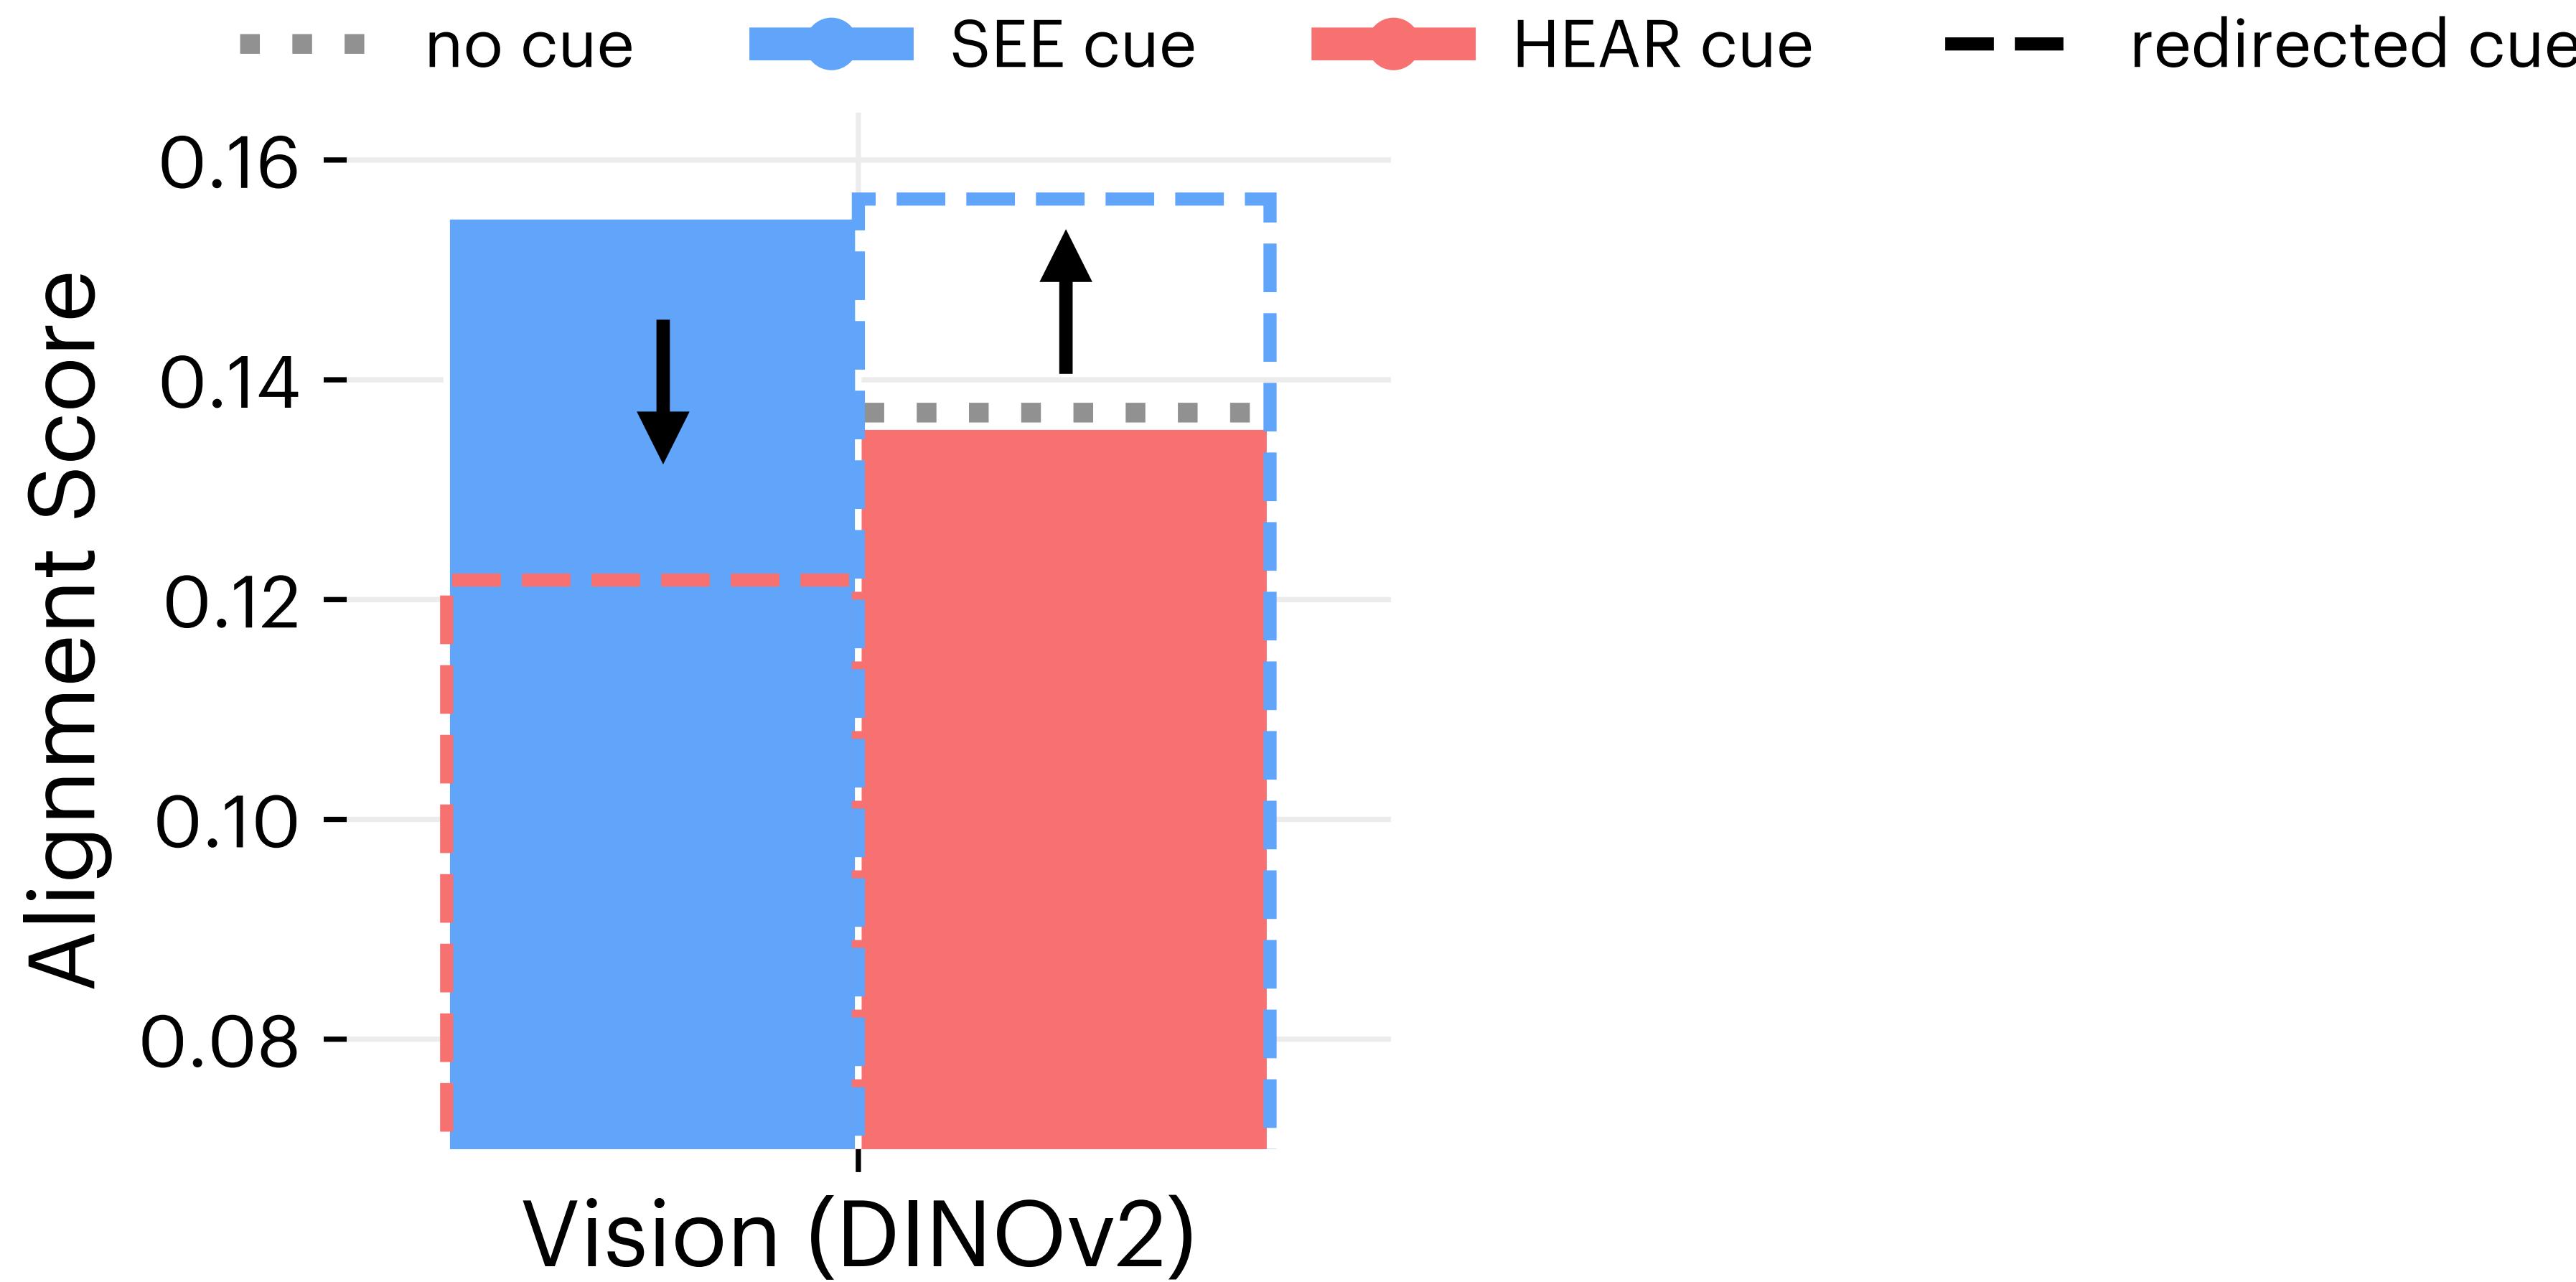
Please rewrite the following text to describe what the scene would look like **visually**, rather than how it **sounds**: {text}



...The batter stands at the plate, gripping the bat tightly, his muscles coiled with anticipation. His eyes are **locked onto** the pitcher, scanning for the slightest movement...

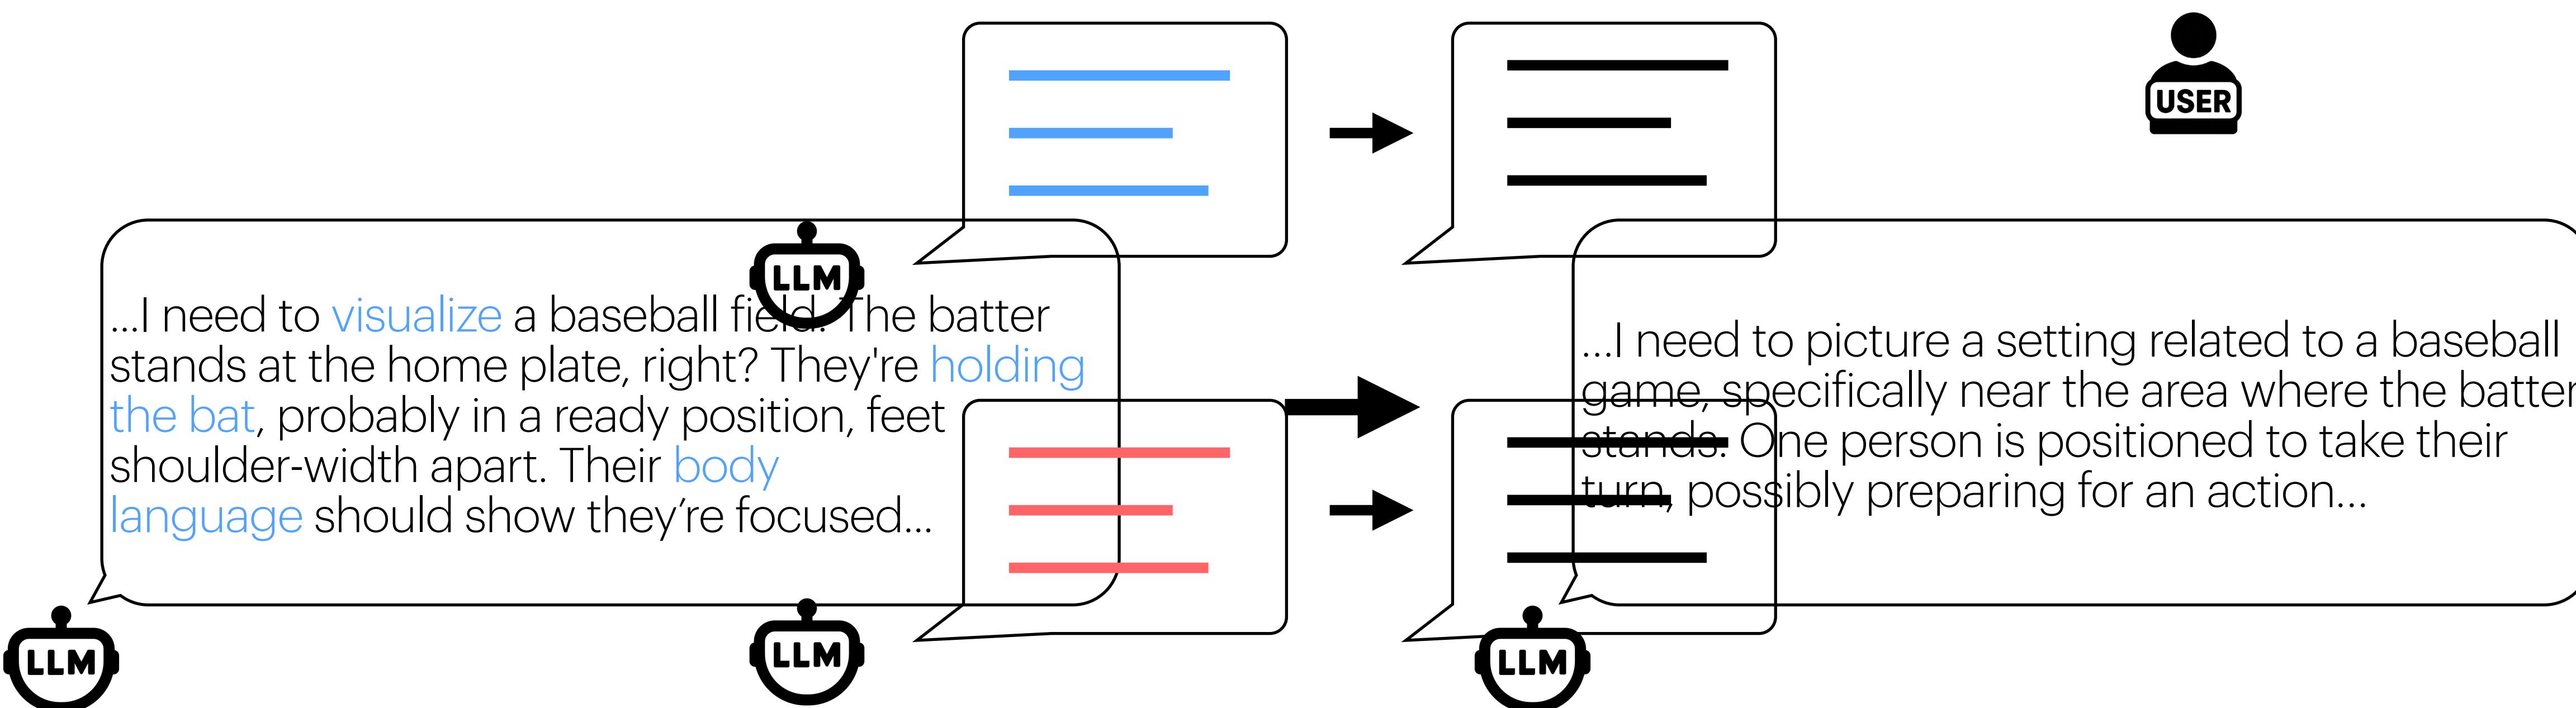


# Redirecting sensory cues preserves modality-specific alignment

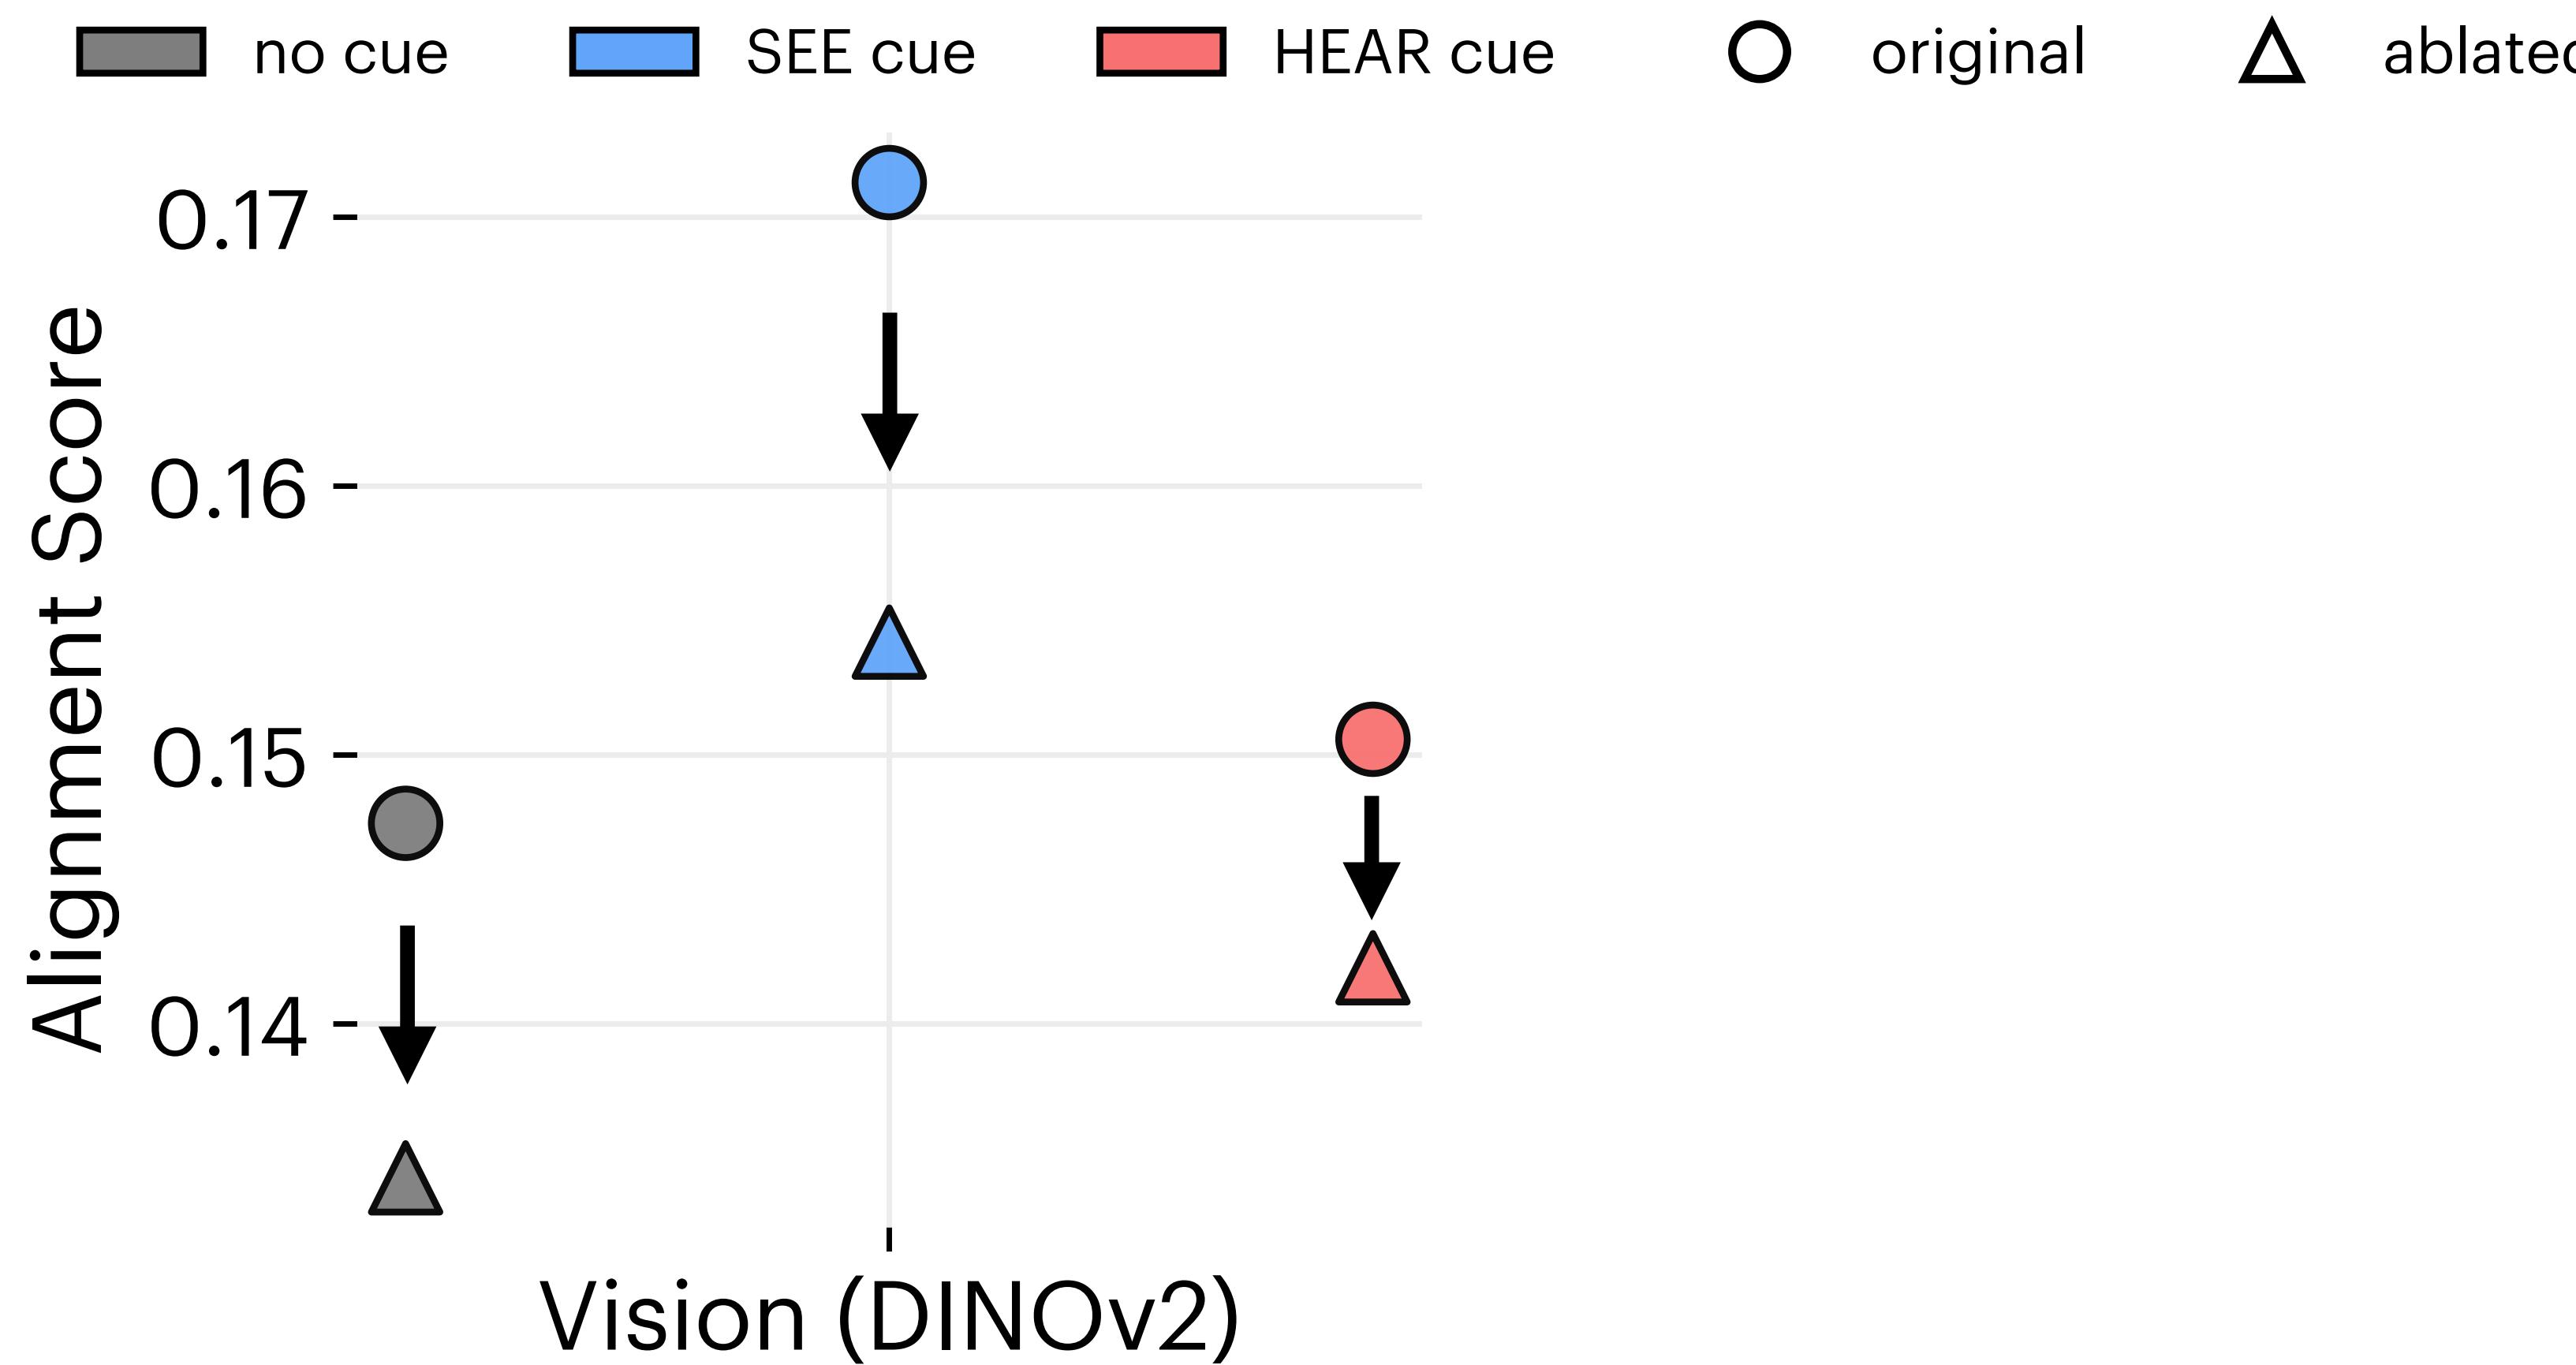


# Removing sensory references

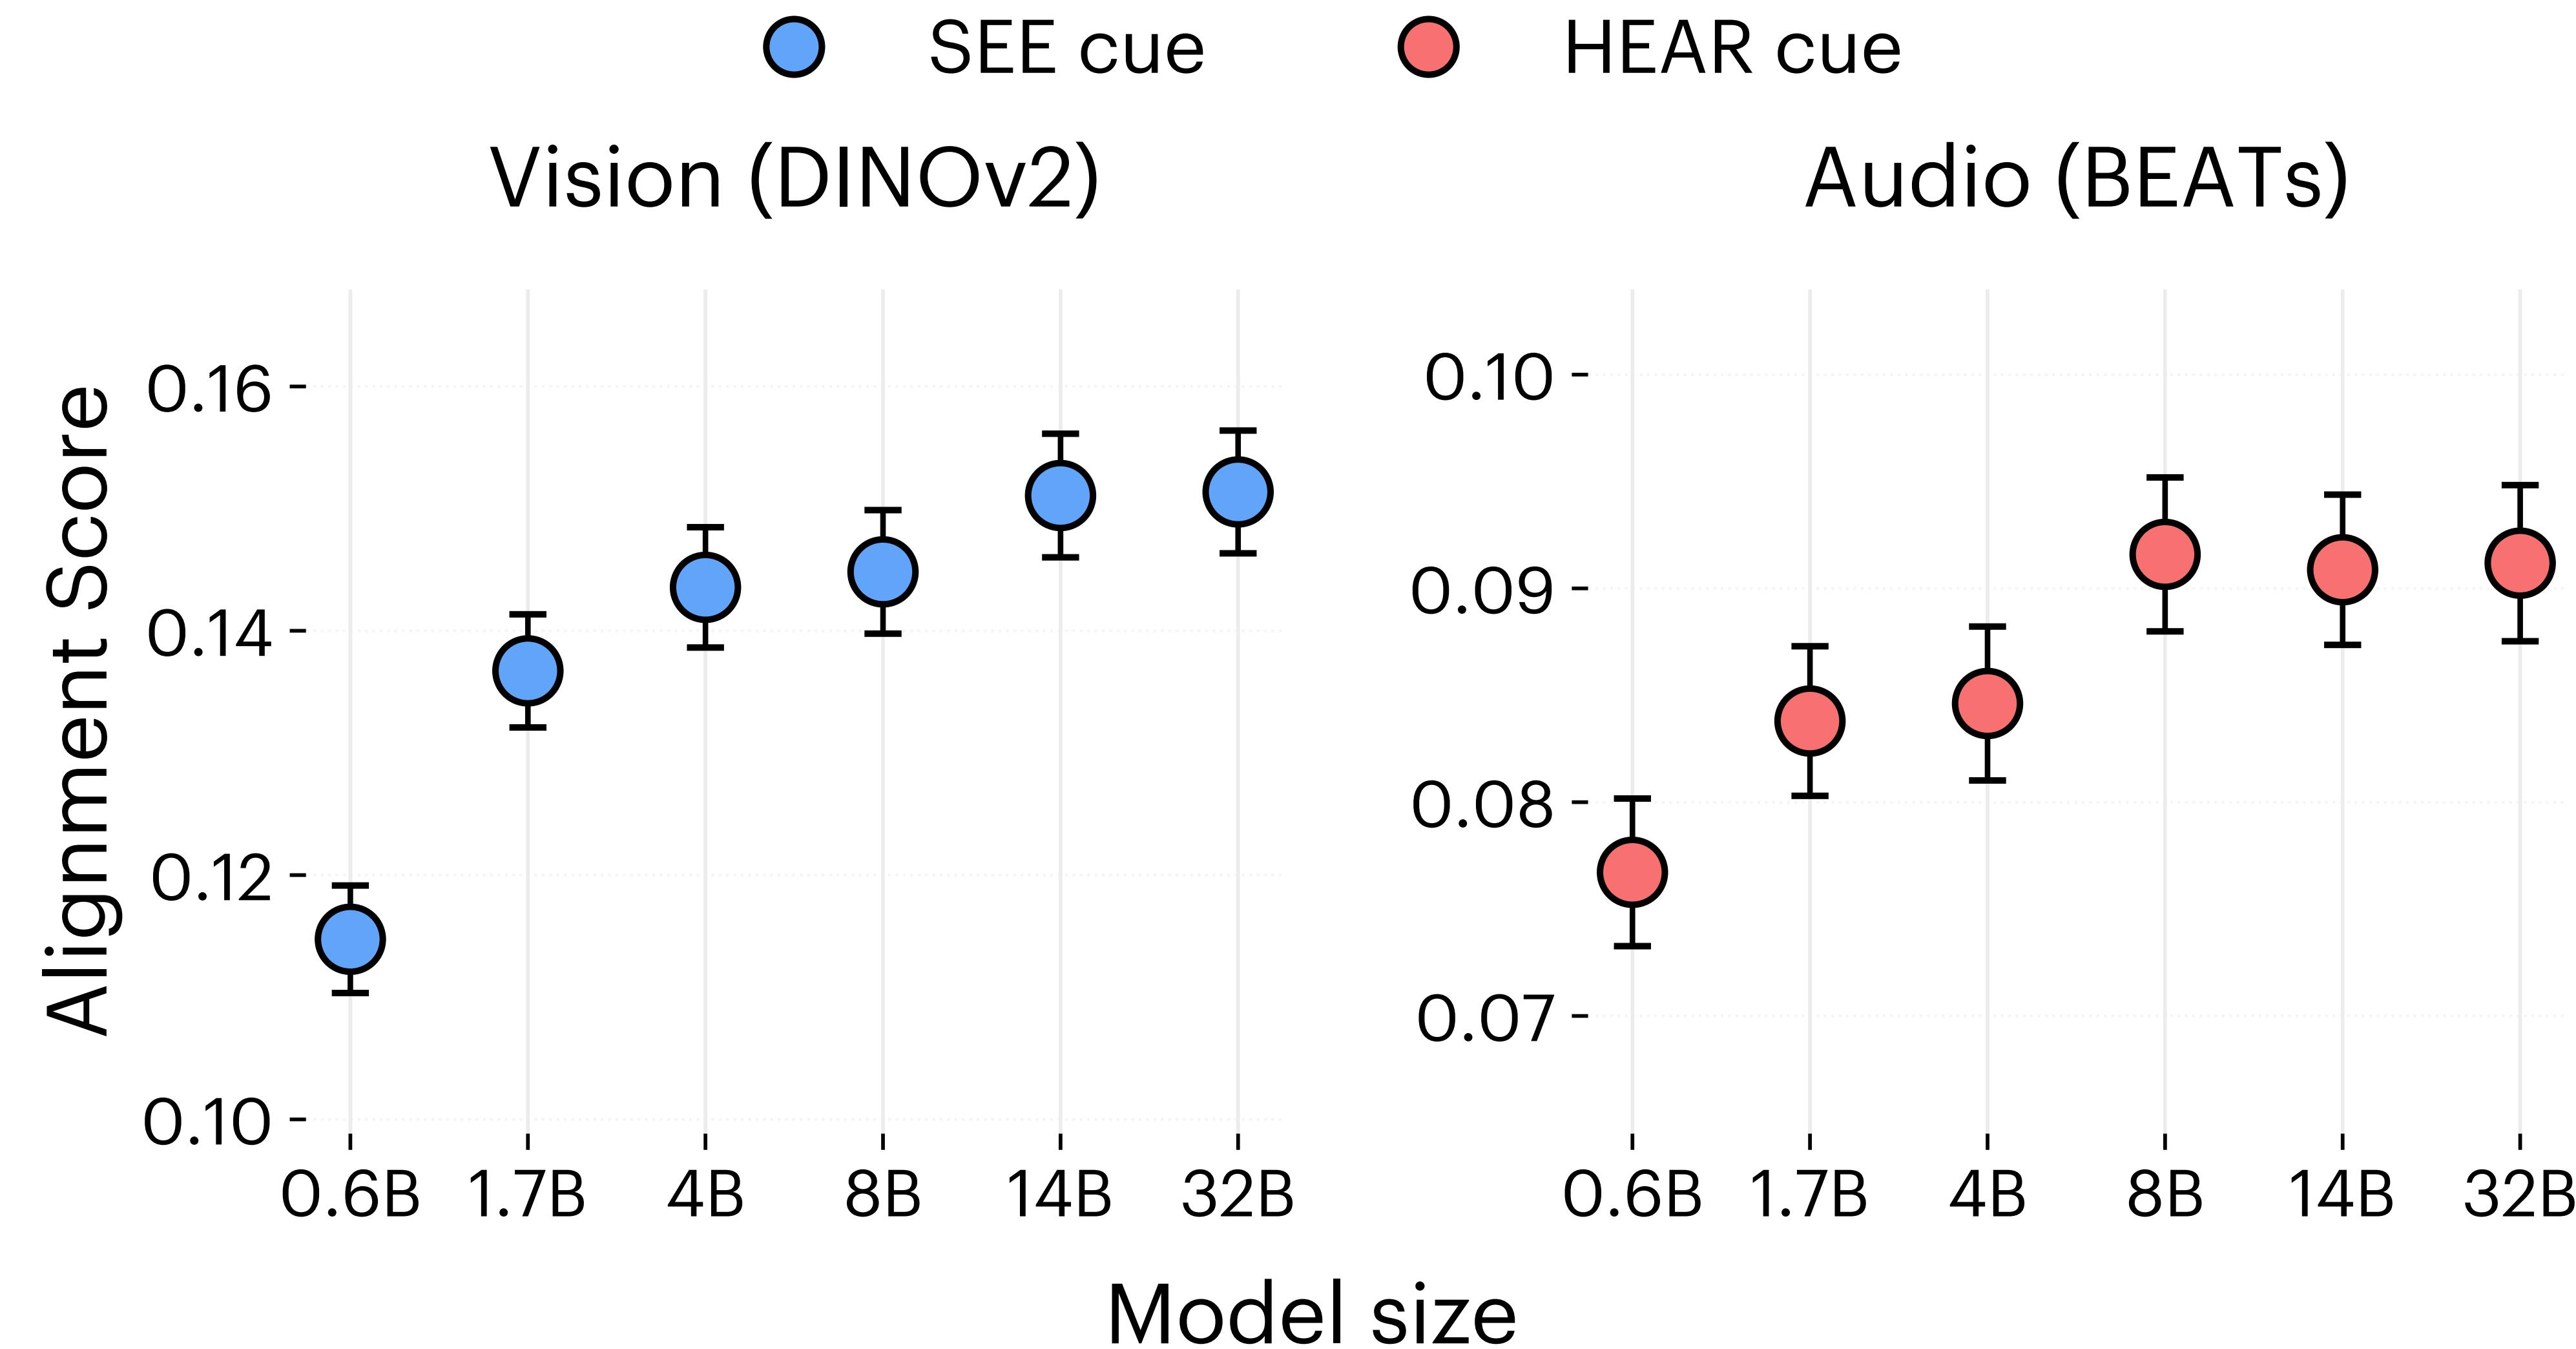
Please rewrite the following text by **removing all sensory-specific words** or descriptions (e.g., related to sound, sight, smell, touch, taste), and **replace them with neutral, non-sensory words**. The result should preserve the event or action described but remove explicit sensory grounding: {text}



# Use of sensory language strengthens modality-specific representations



# Alignment is higher in larger models



# What about senses humans don't have?

## X-ray Crystallography



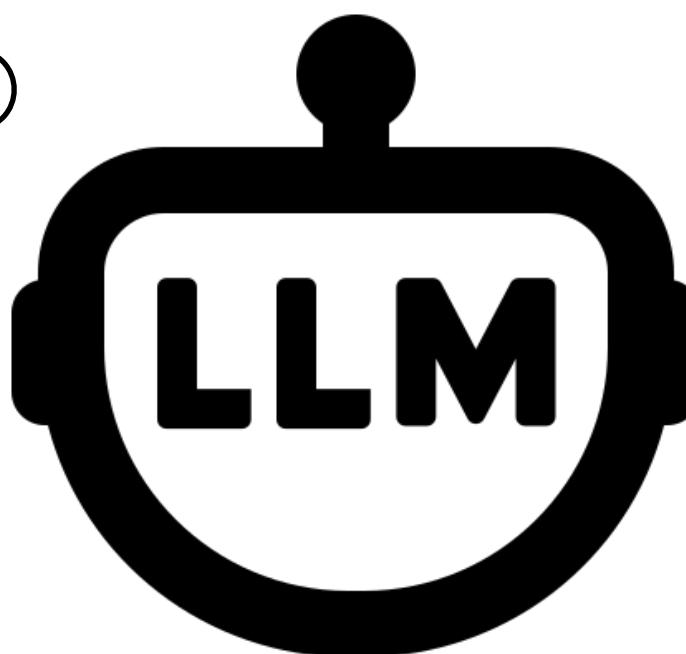
<https://msg.ucsf.edu/rigaku-ru-200>

## Cryo-EM



<https://emcore.ucsf.edu/krios-g2-cryo-tem>

# What about senses humans don't have?



# Bio Sensory Prompting

SEQUENCE:  
KKAVINGEQIRSIISDLHQTLKKEALALPEYYGENLDALWDCLTGWV  
EYPLVLEWRQFEQSKQLTENGAESVLQVFREAKAEGCDITIILS



Provide a thorough summary of {protein name}.  
Include its gene name, protein family, molecular weight, known structural domains, function in the cell, binding sites, any known interactions or pathways it participates in.



RIBONUCLEASE SA COMPLEX

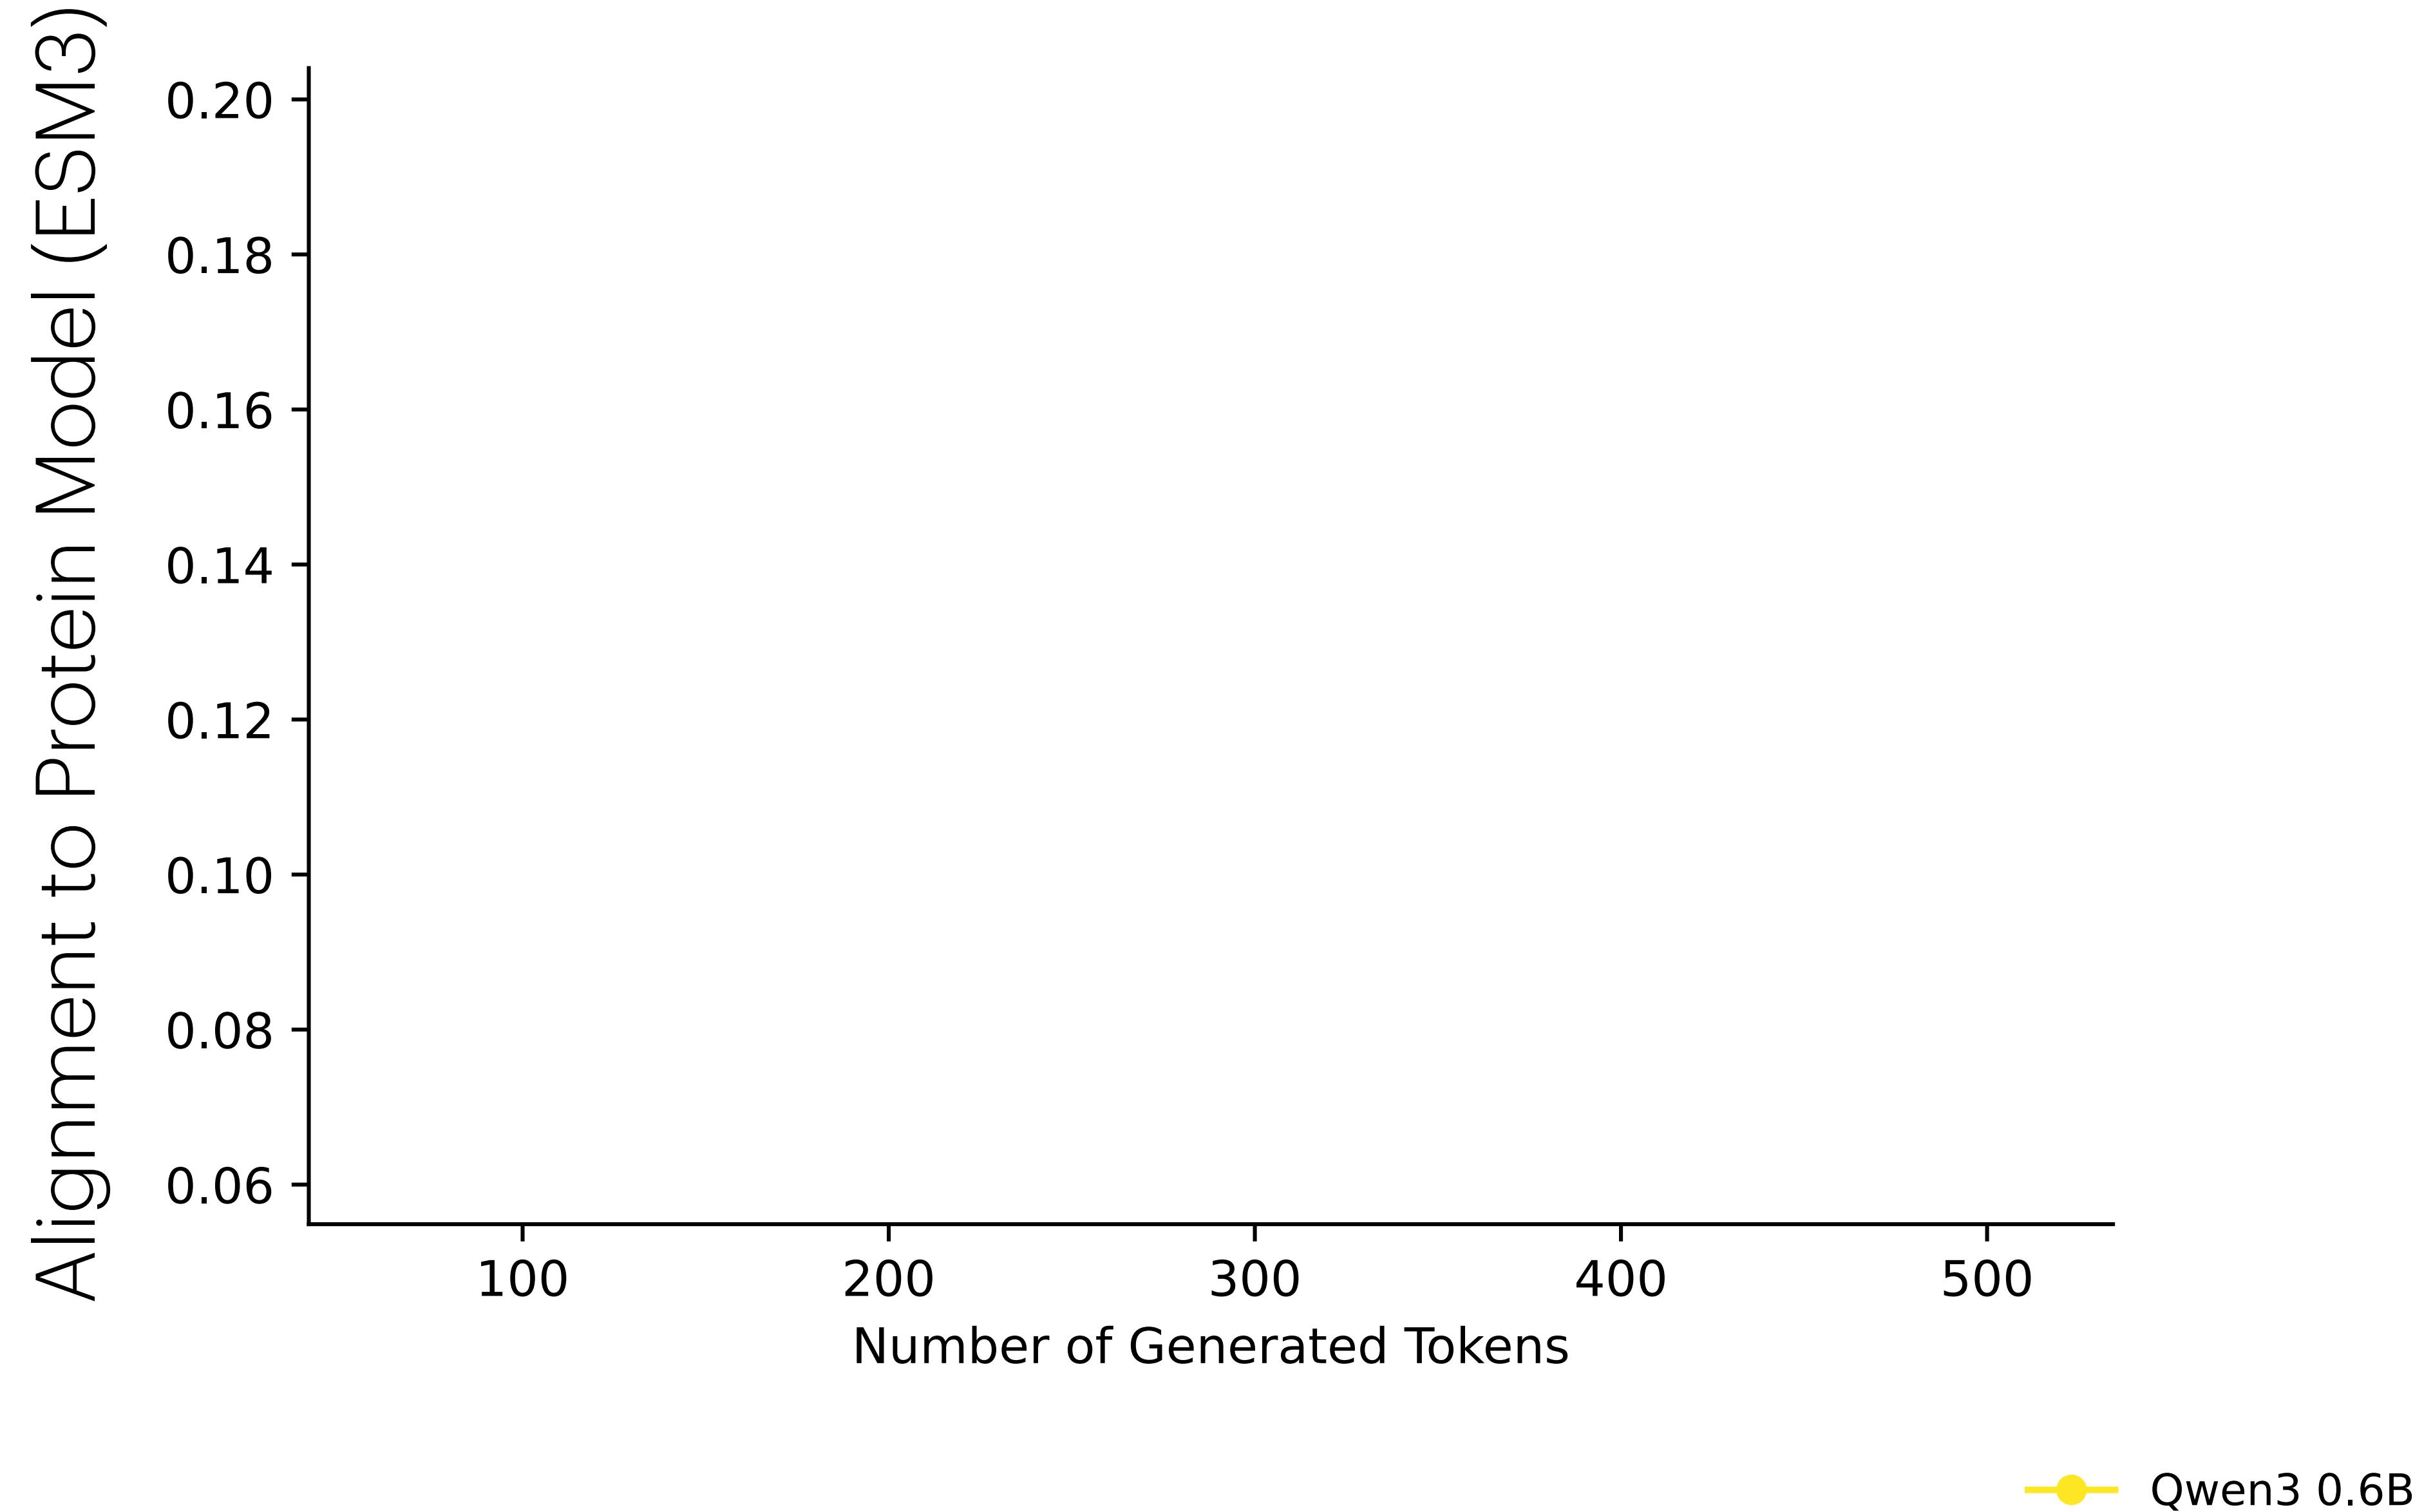


Okay, I need to provide a thorough summary of the RIBONUCLEASE SA COMPLEX WITH BARSTAR. Let me start by recalling what I know about these proteins.

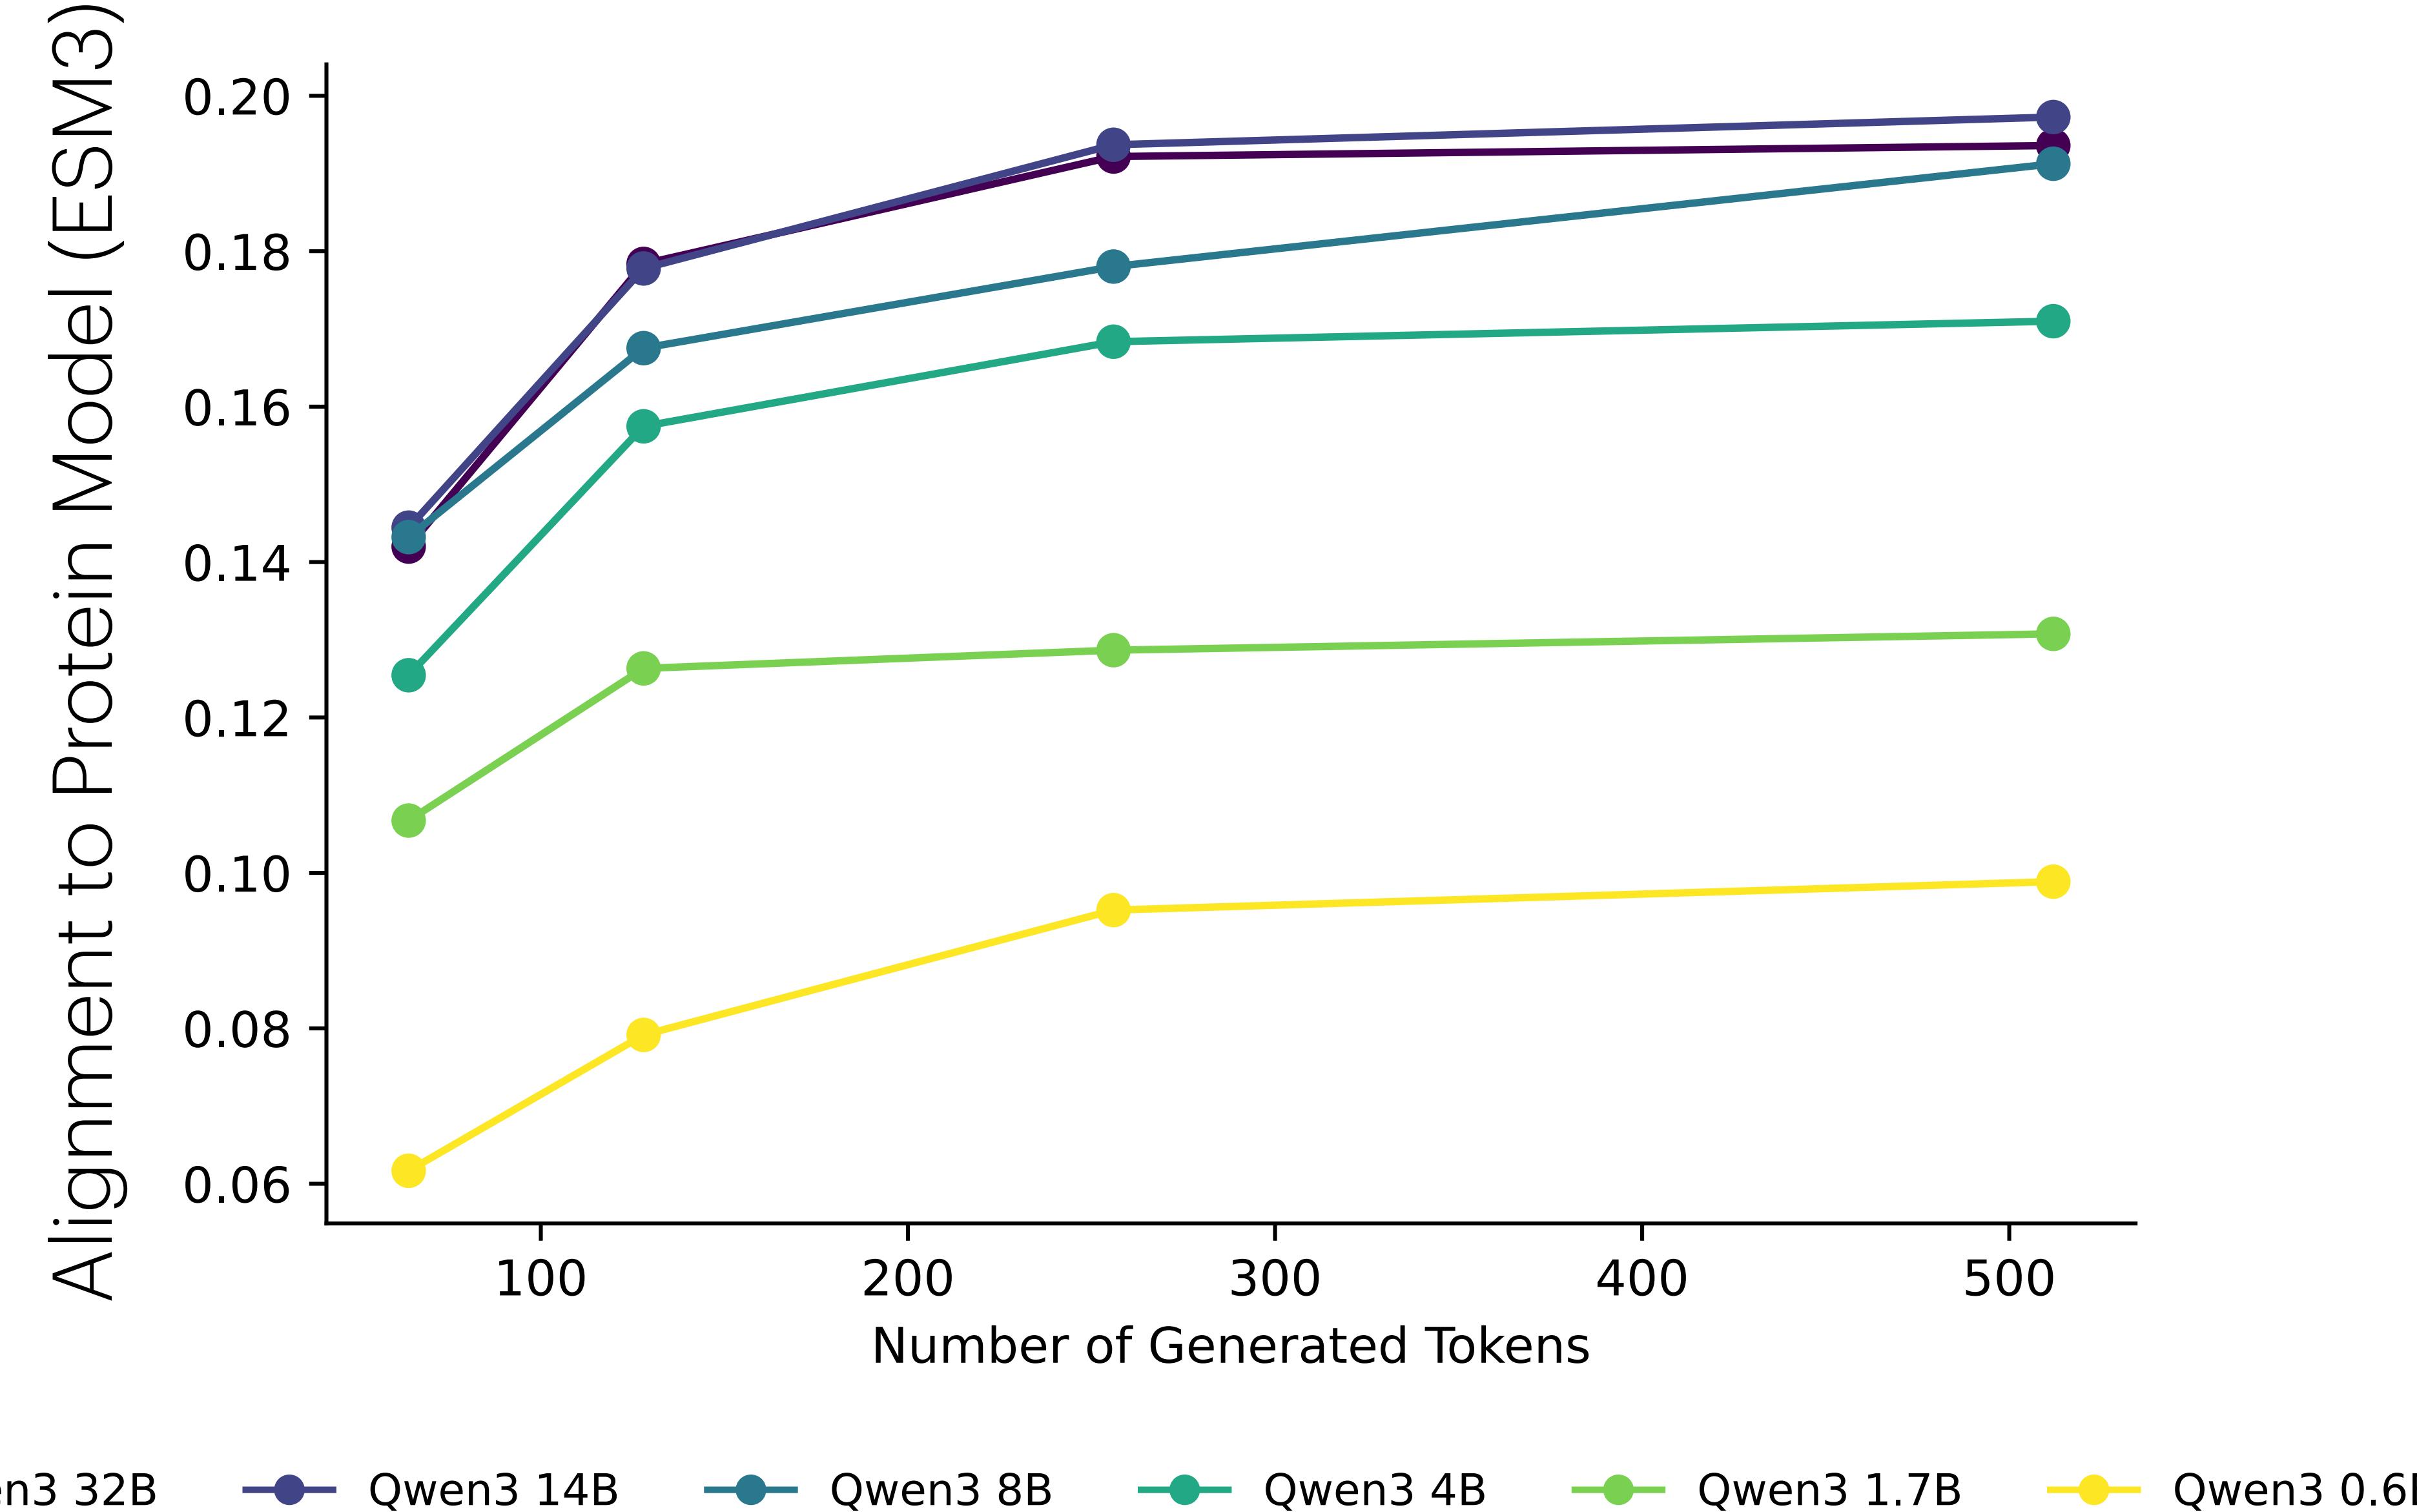
First, the gene names. I remember that RNase SA is from the bacterium *Staphylococcus aureus*. The gene might be named rnsa, but I should verify that. Barstar is the inhibitor of RNase SA, so its gene might be bar...



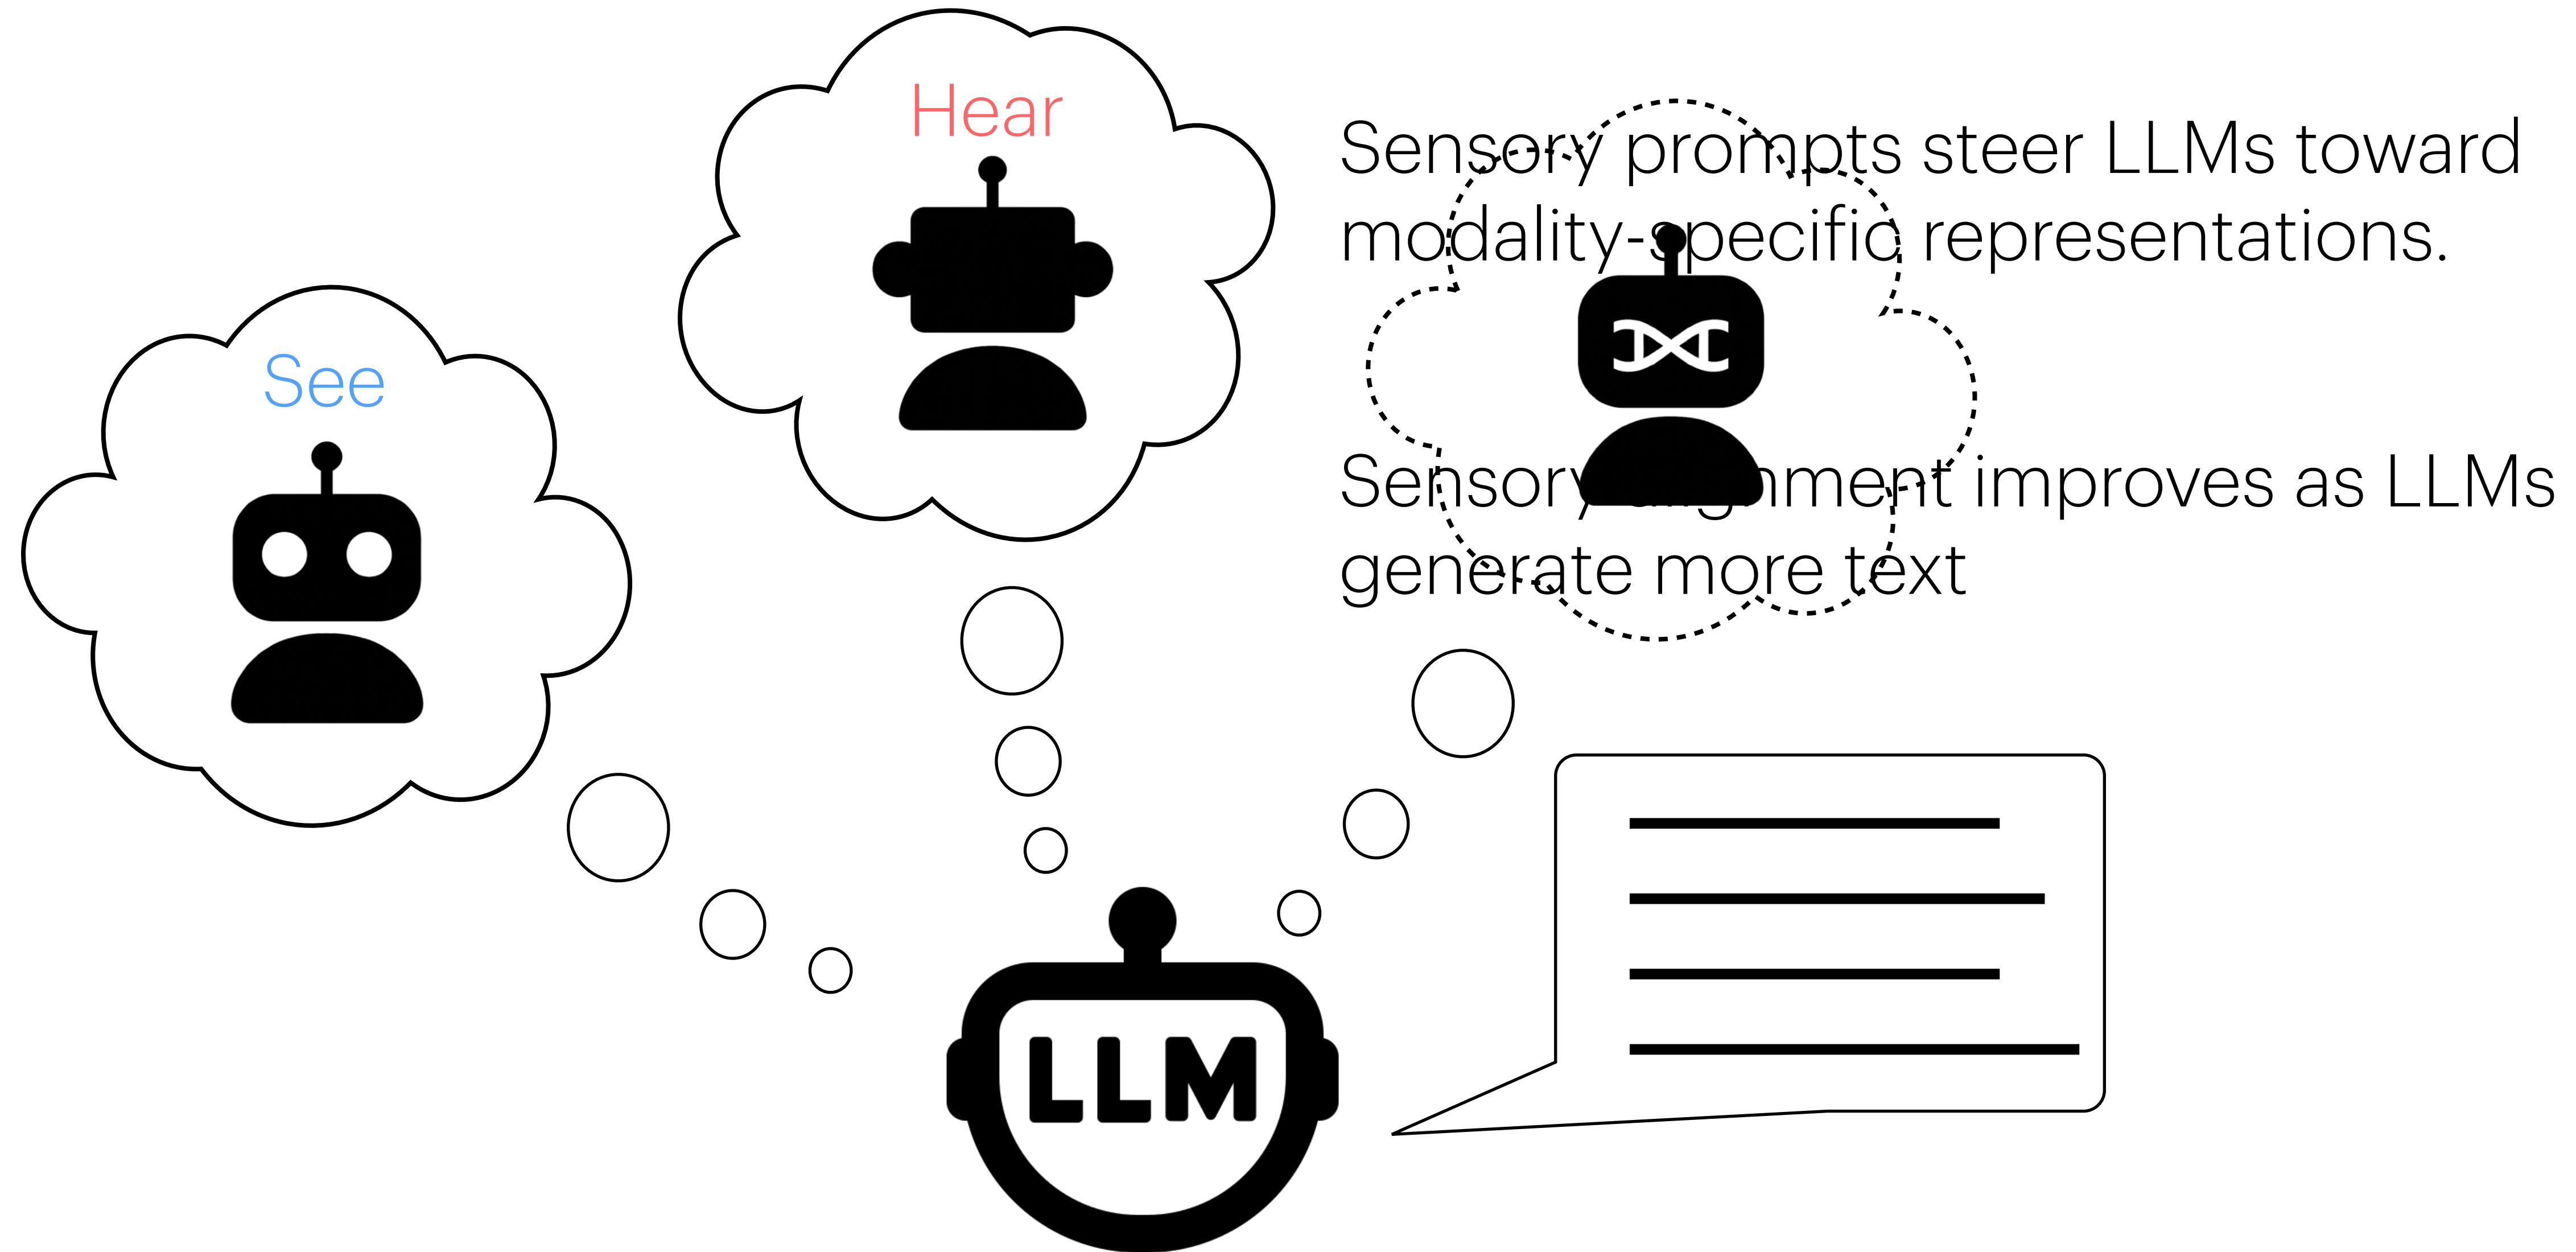
# Bio Sensory Prompting



# Bio Sensory Prompting

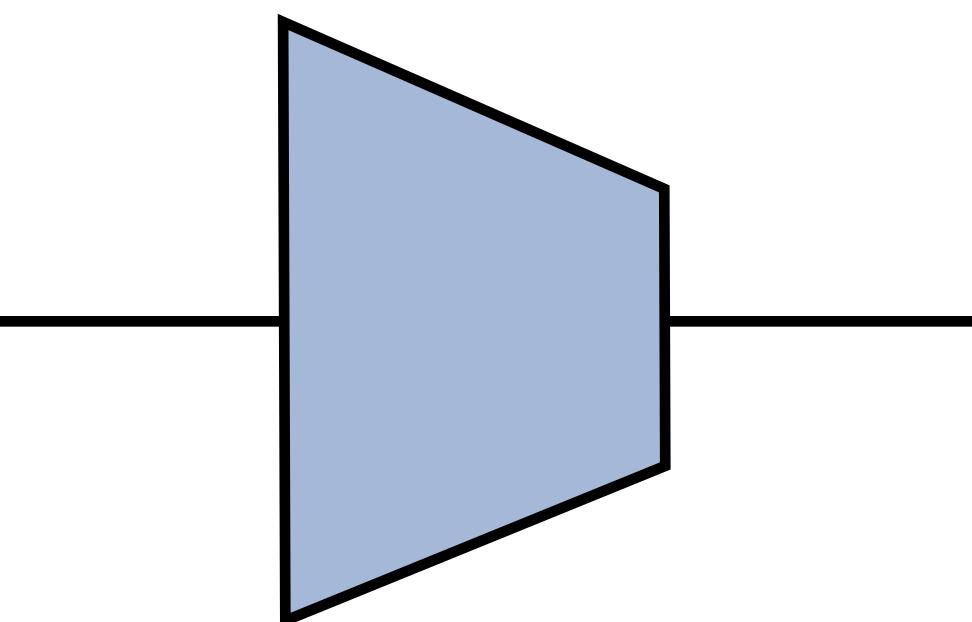
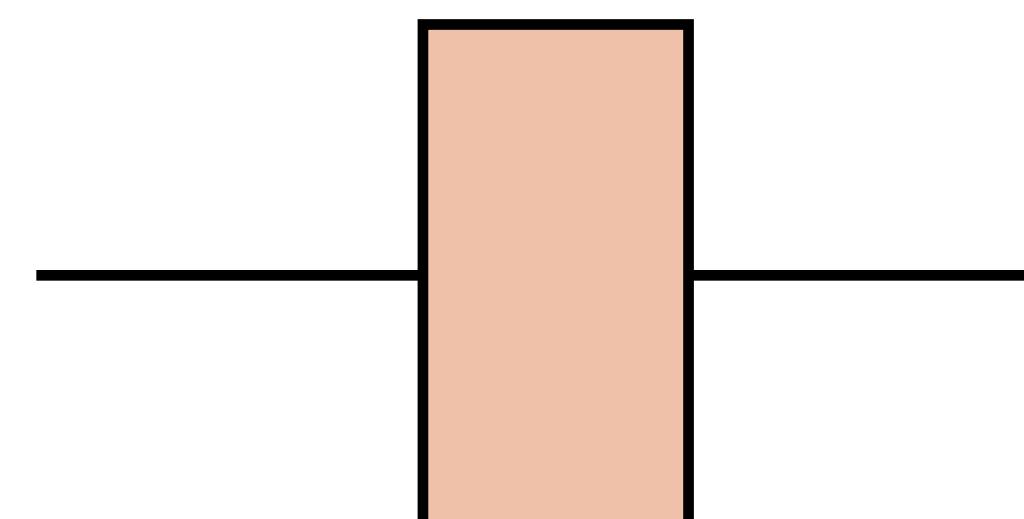


# You can just ask for perception

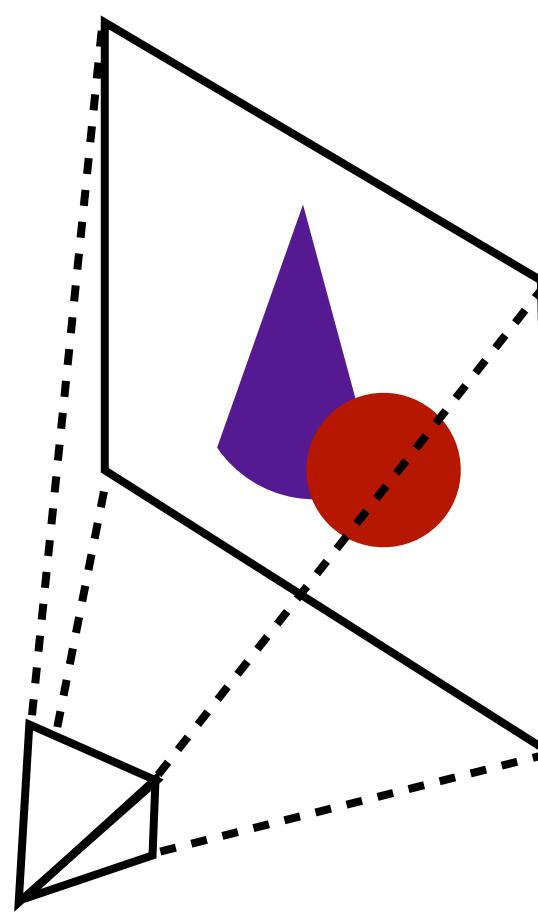


Representational alignment predicts downstream performance



sim(  ,  )

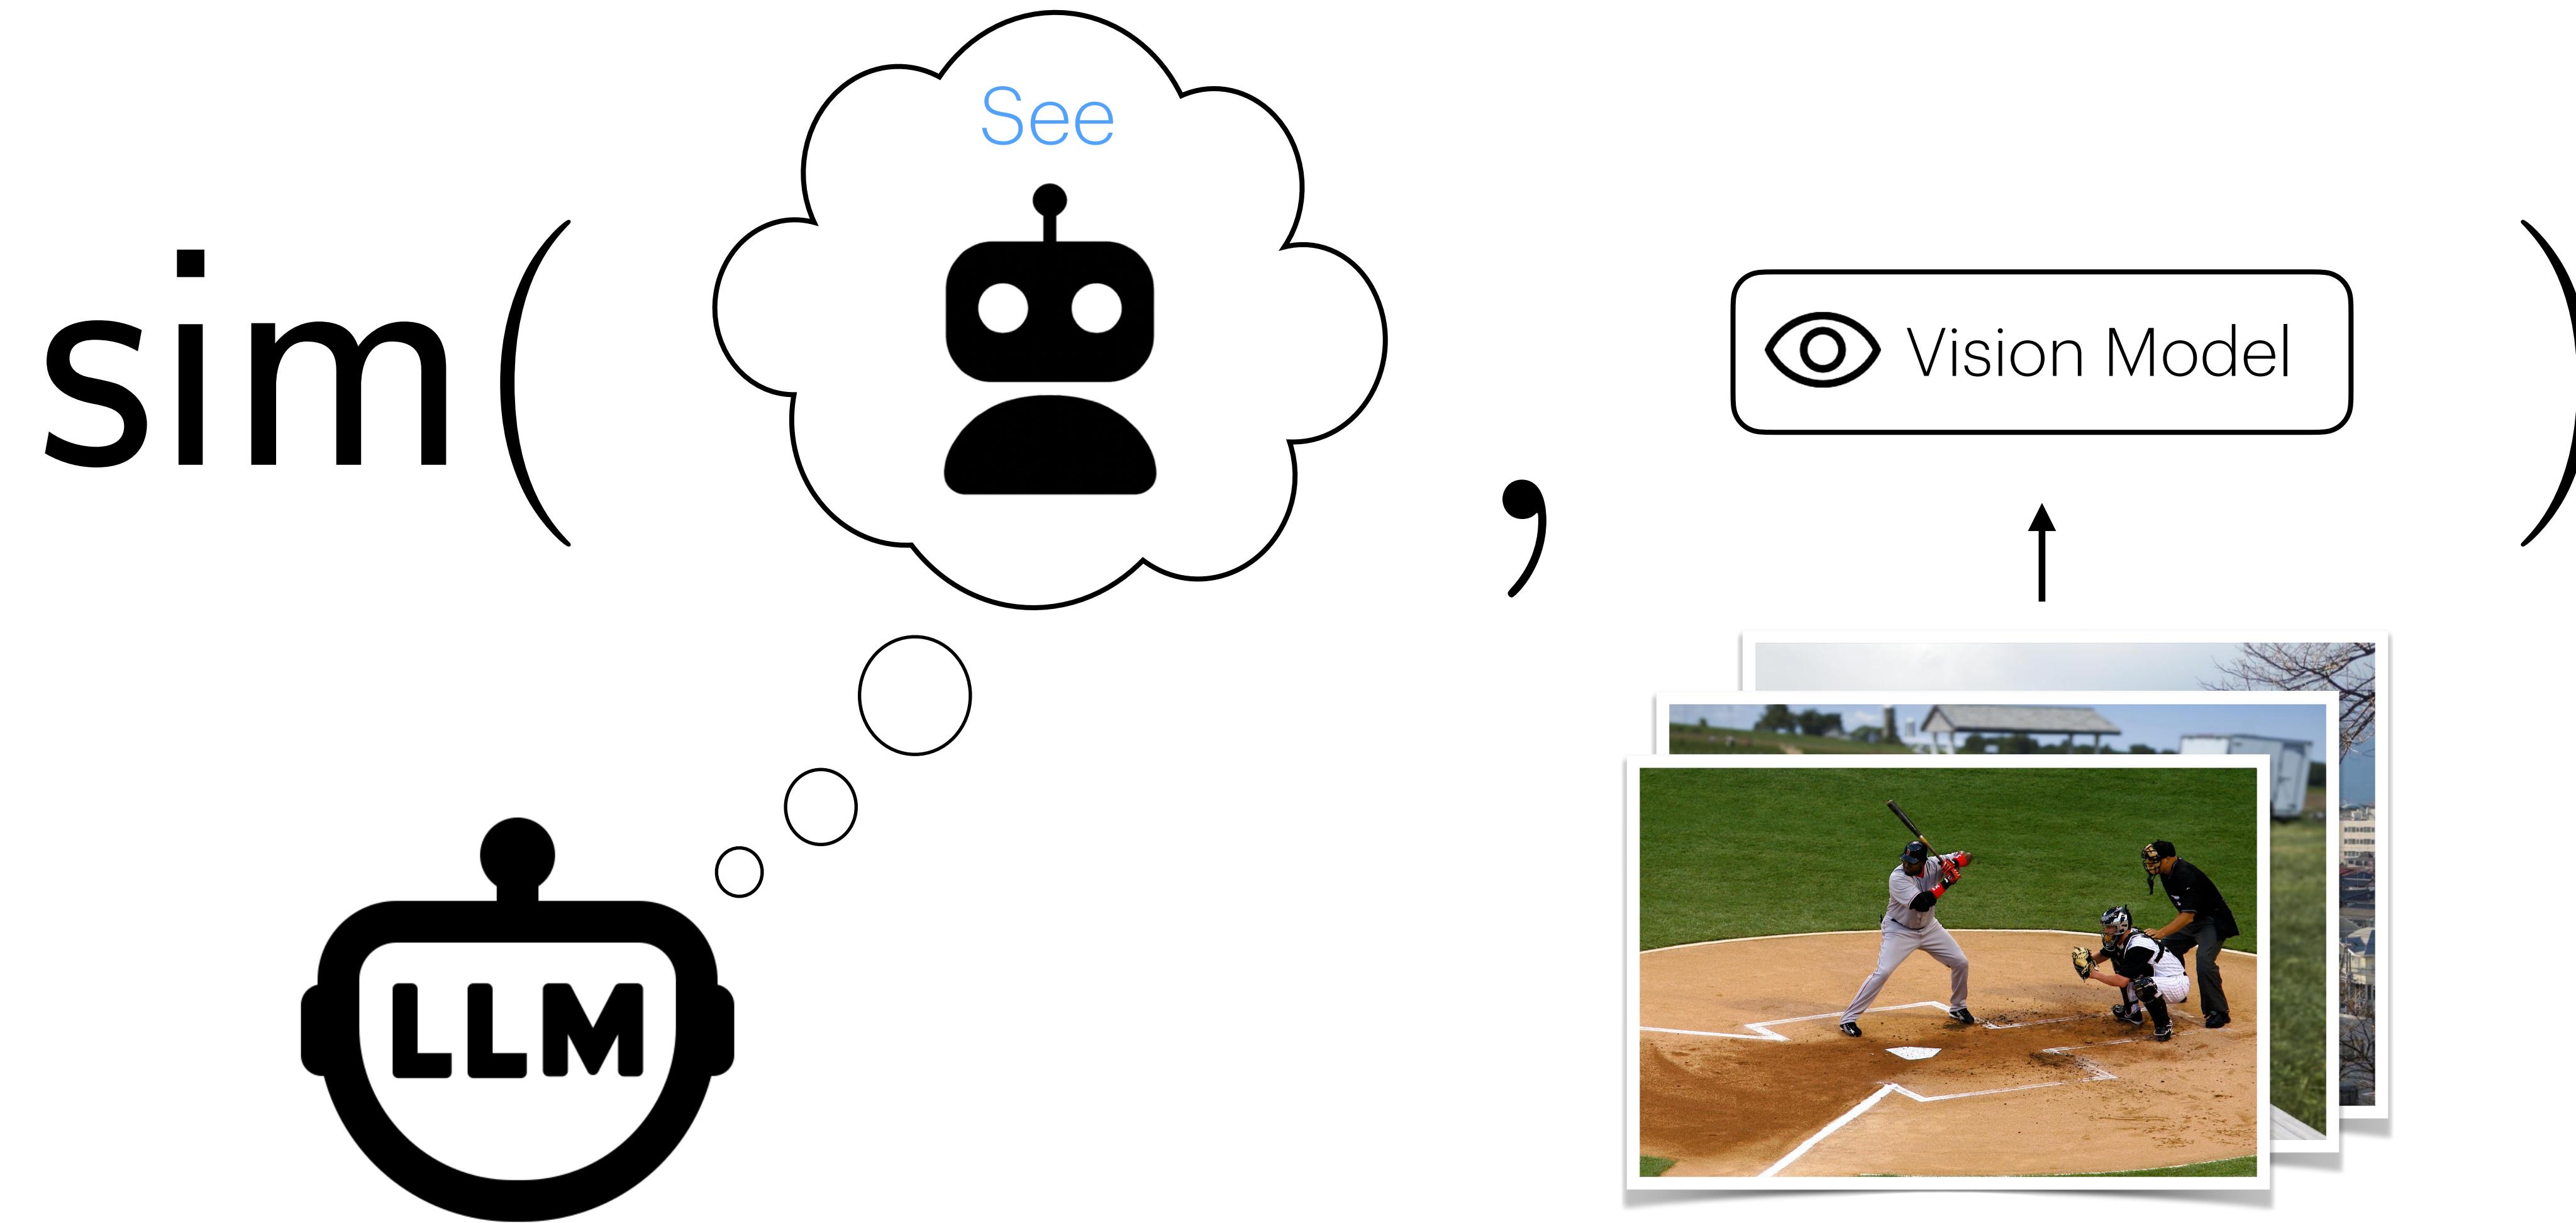
sim(



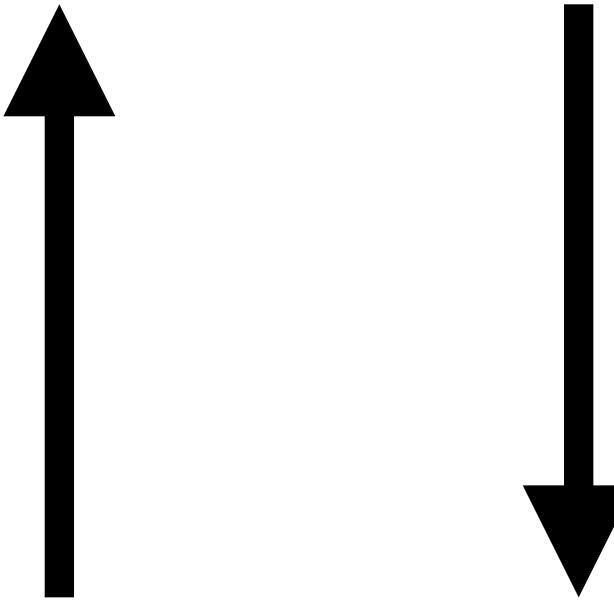
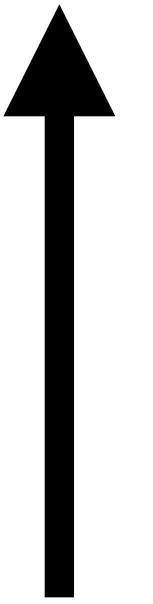
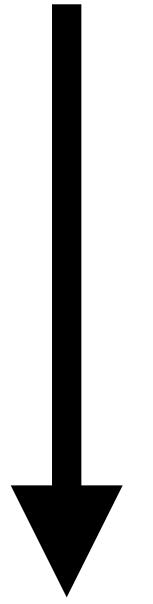
,

A red sphere next  
to a blue cone.

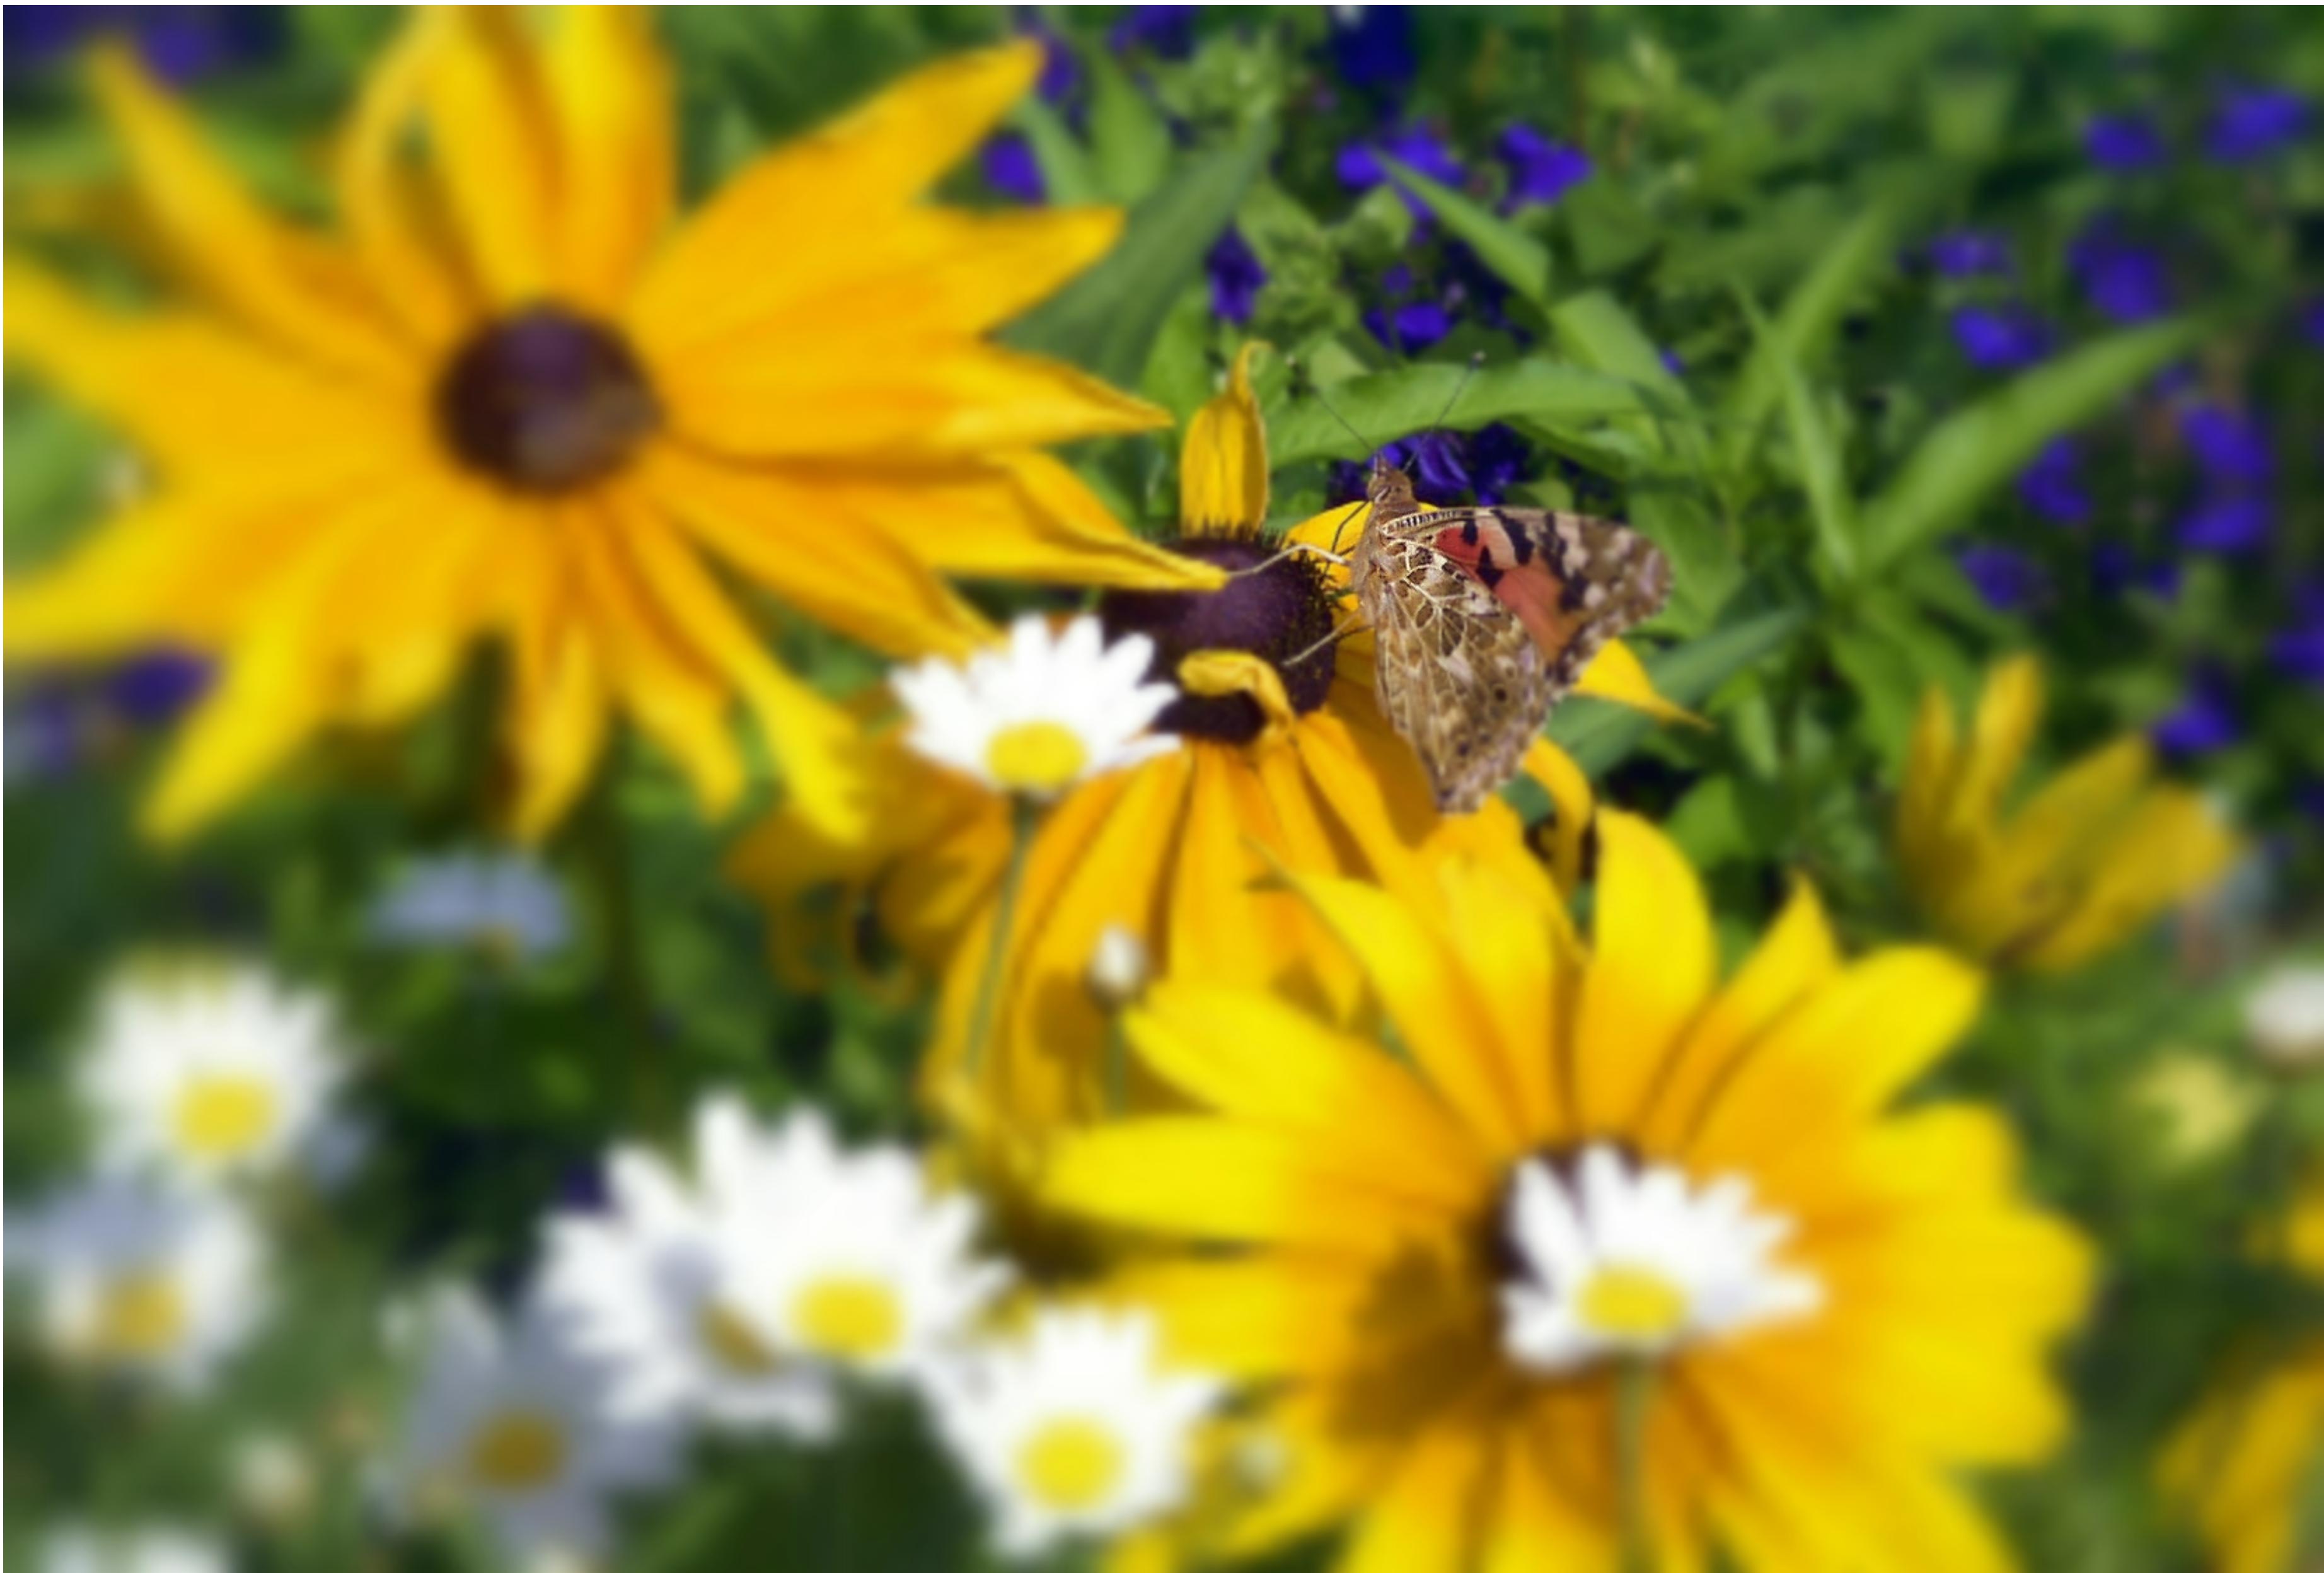
)



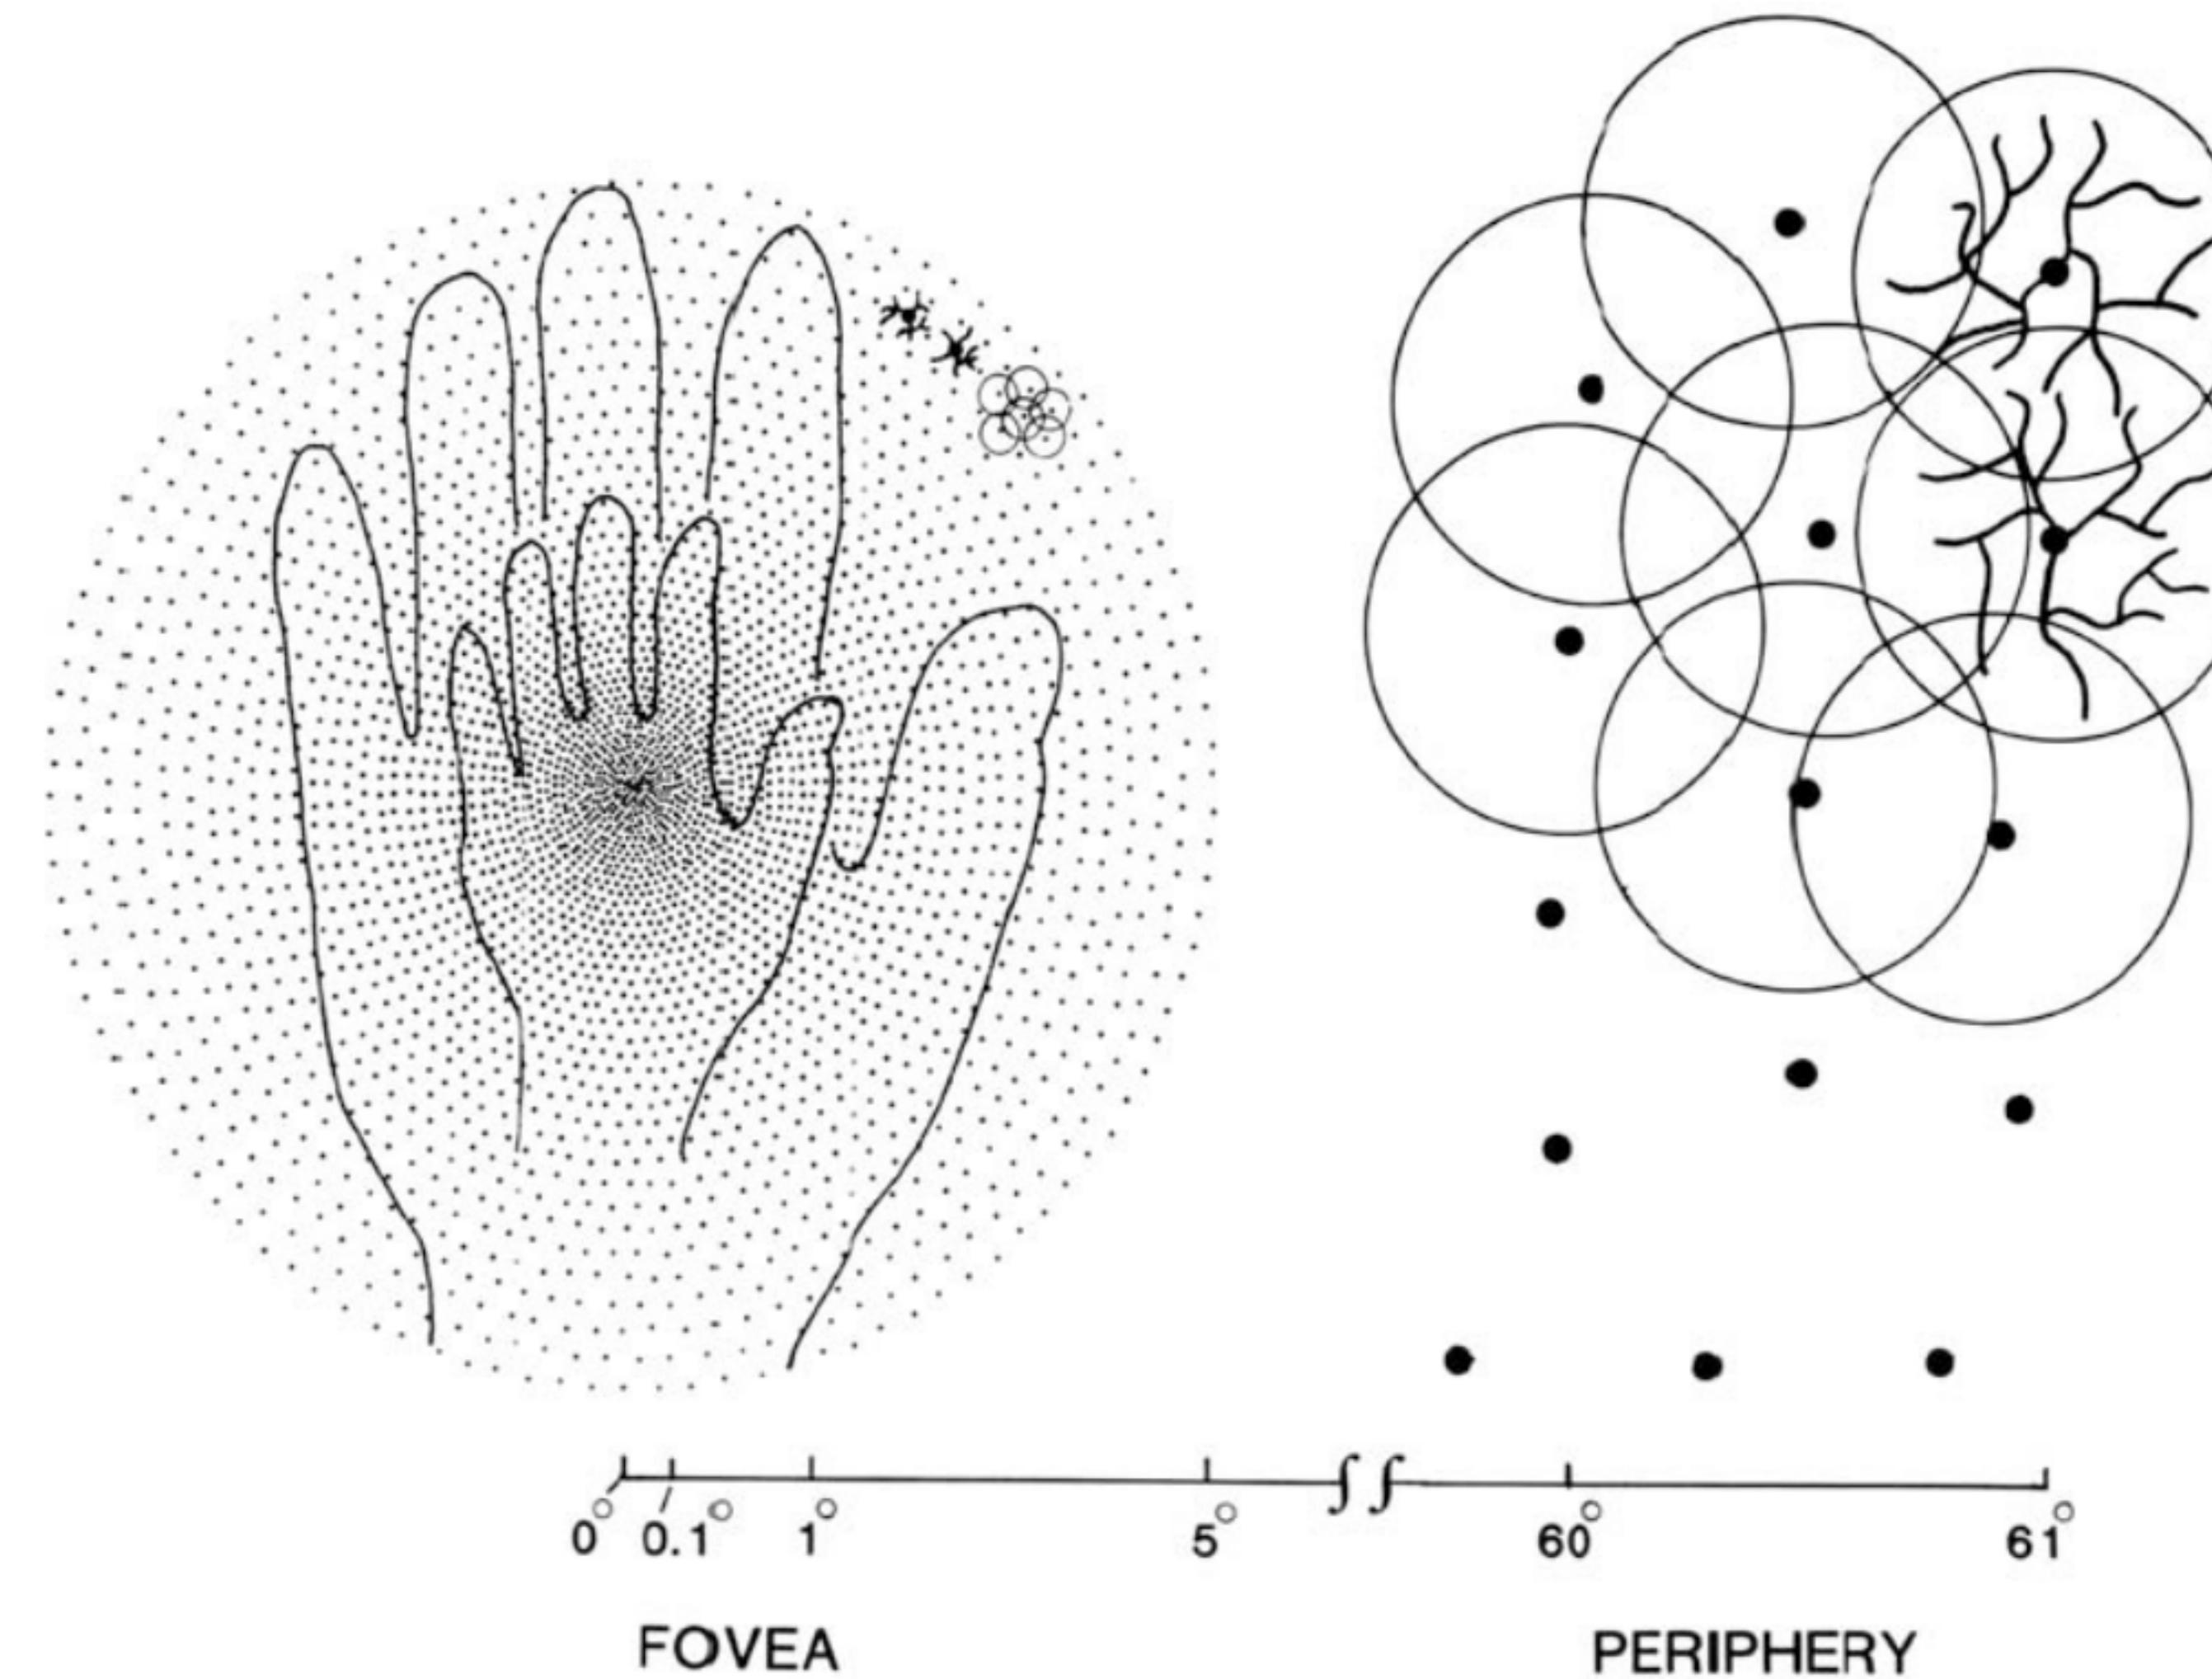
With the convergence of artificial and biological systems,  
what questions about biology can we answer?

**sim**( in-vivo , in-silico ) =   

Why do we have a fovea?

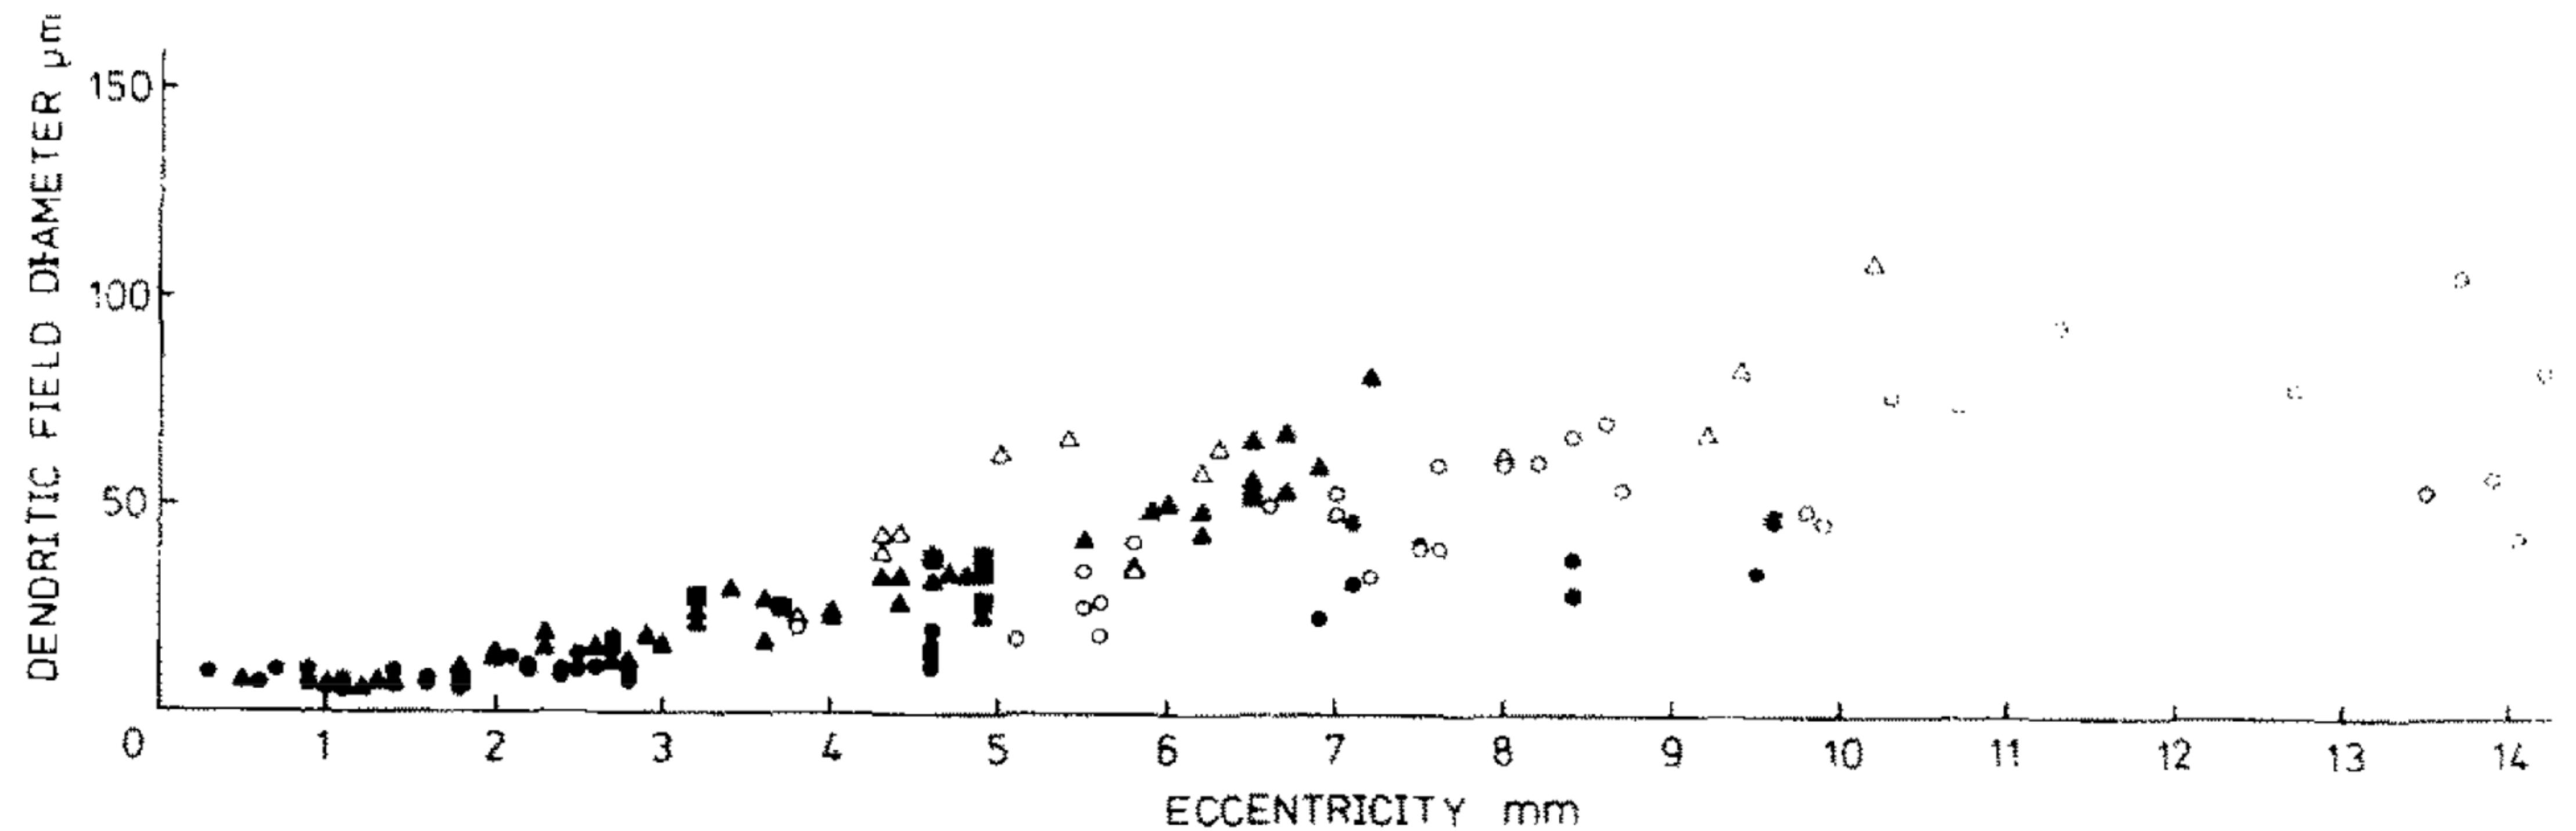


# What is a fovea?



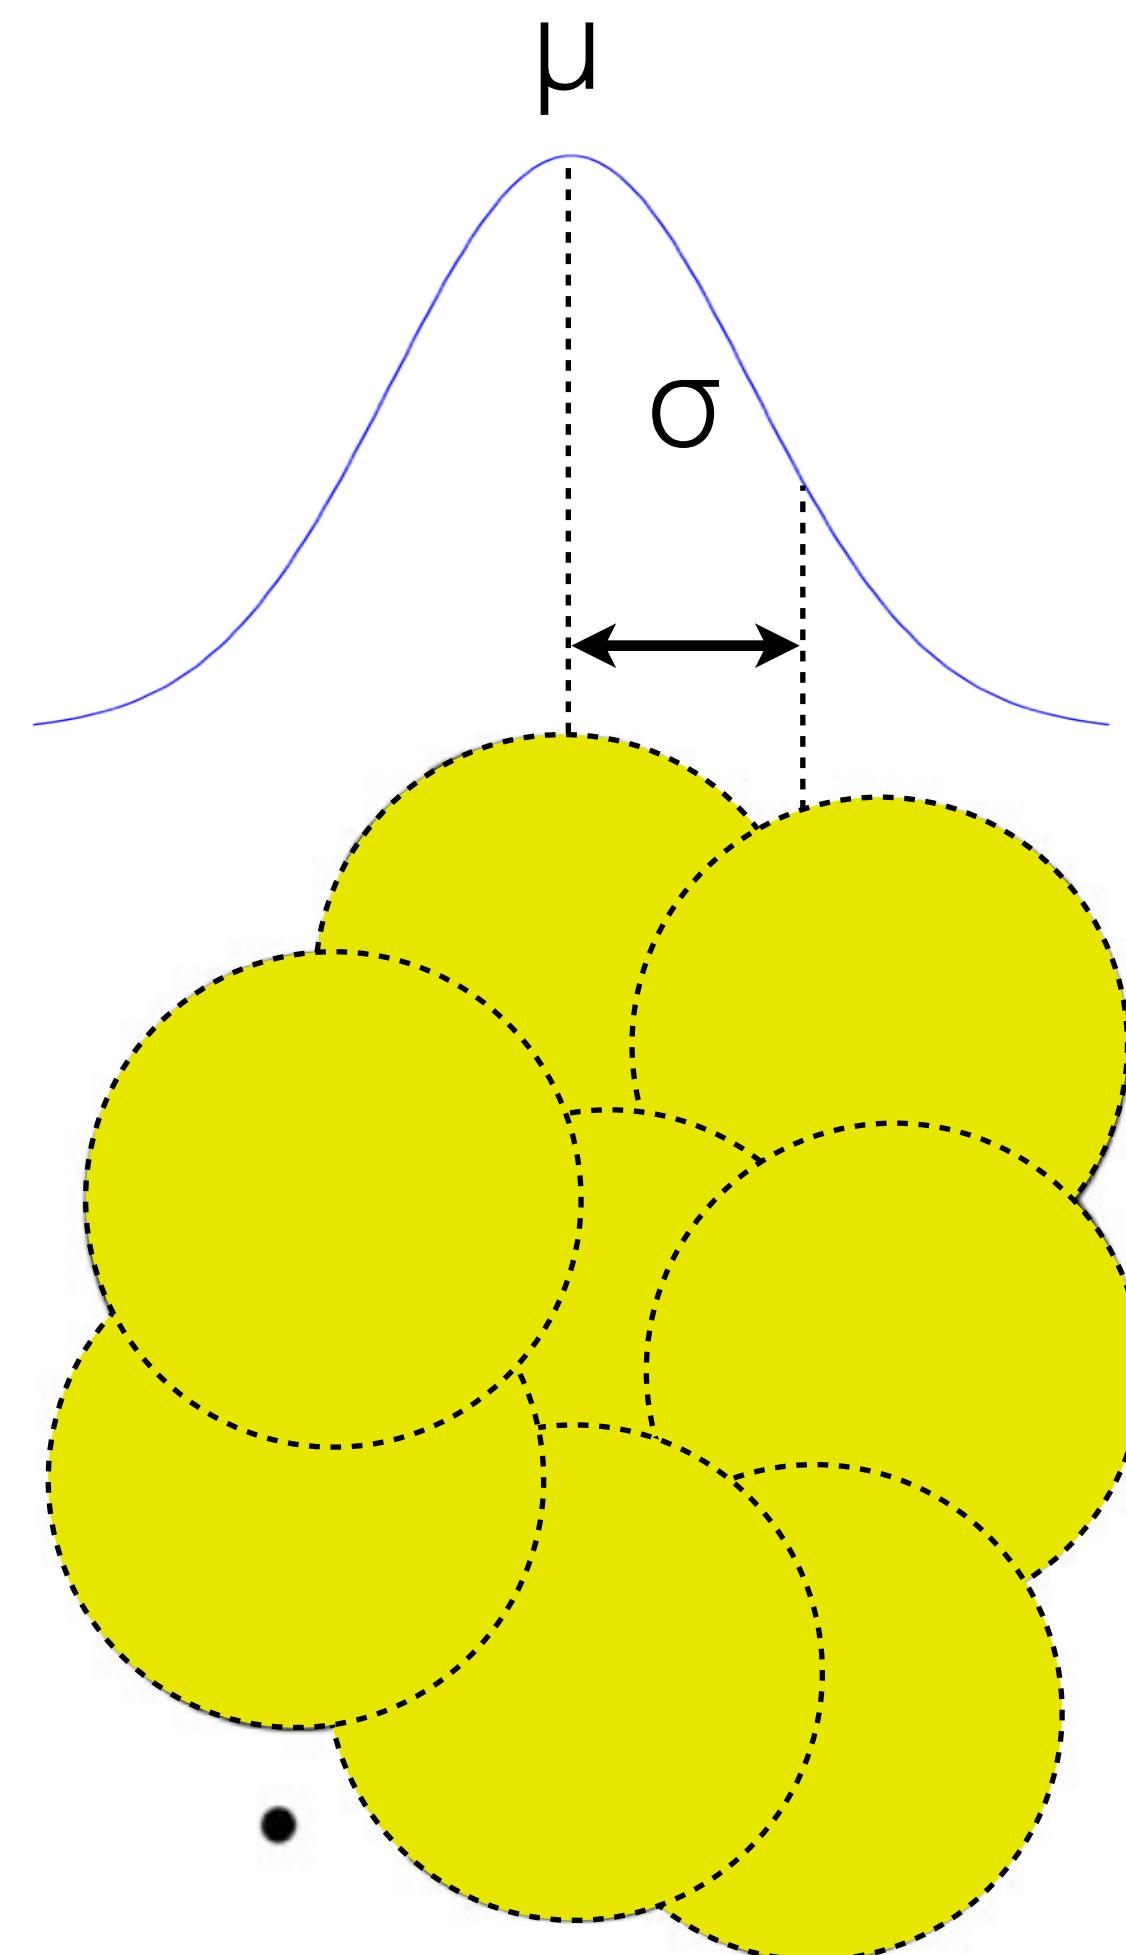
[Van Essen and Anderson 1995]

# Scaling Law for Structure

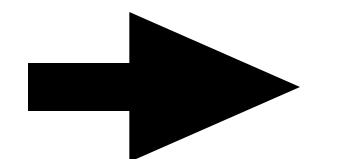
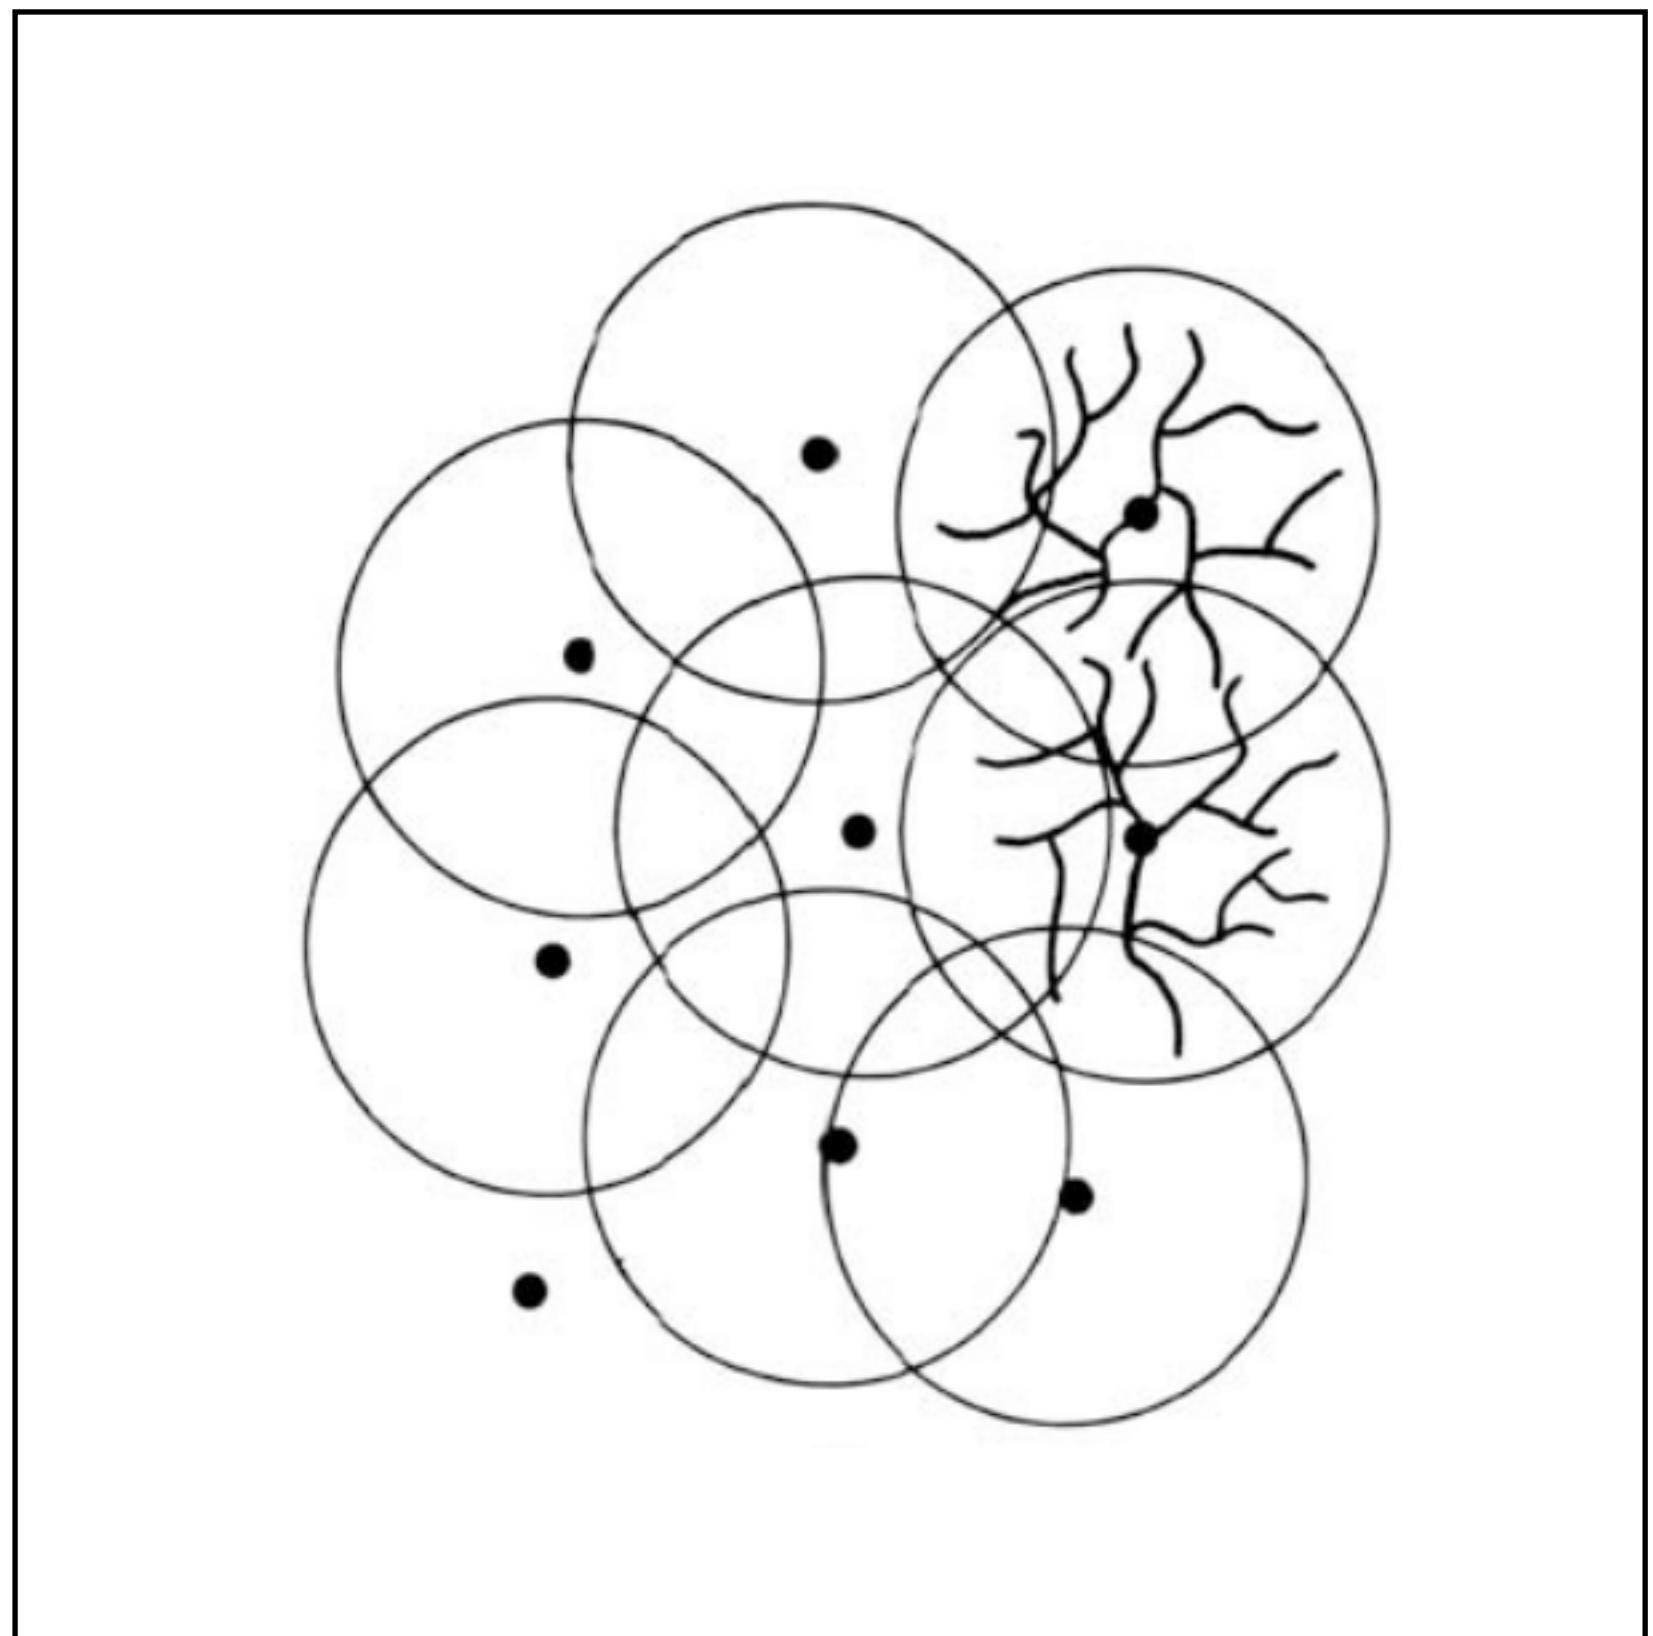


[Perry, Oehler, Cowey 1984]

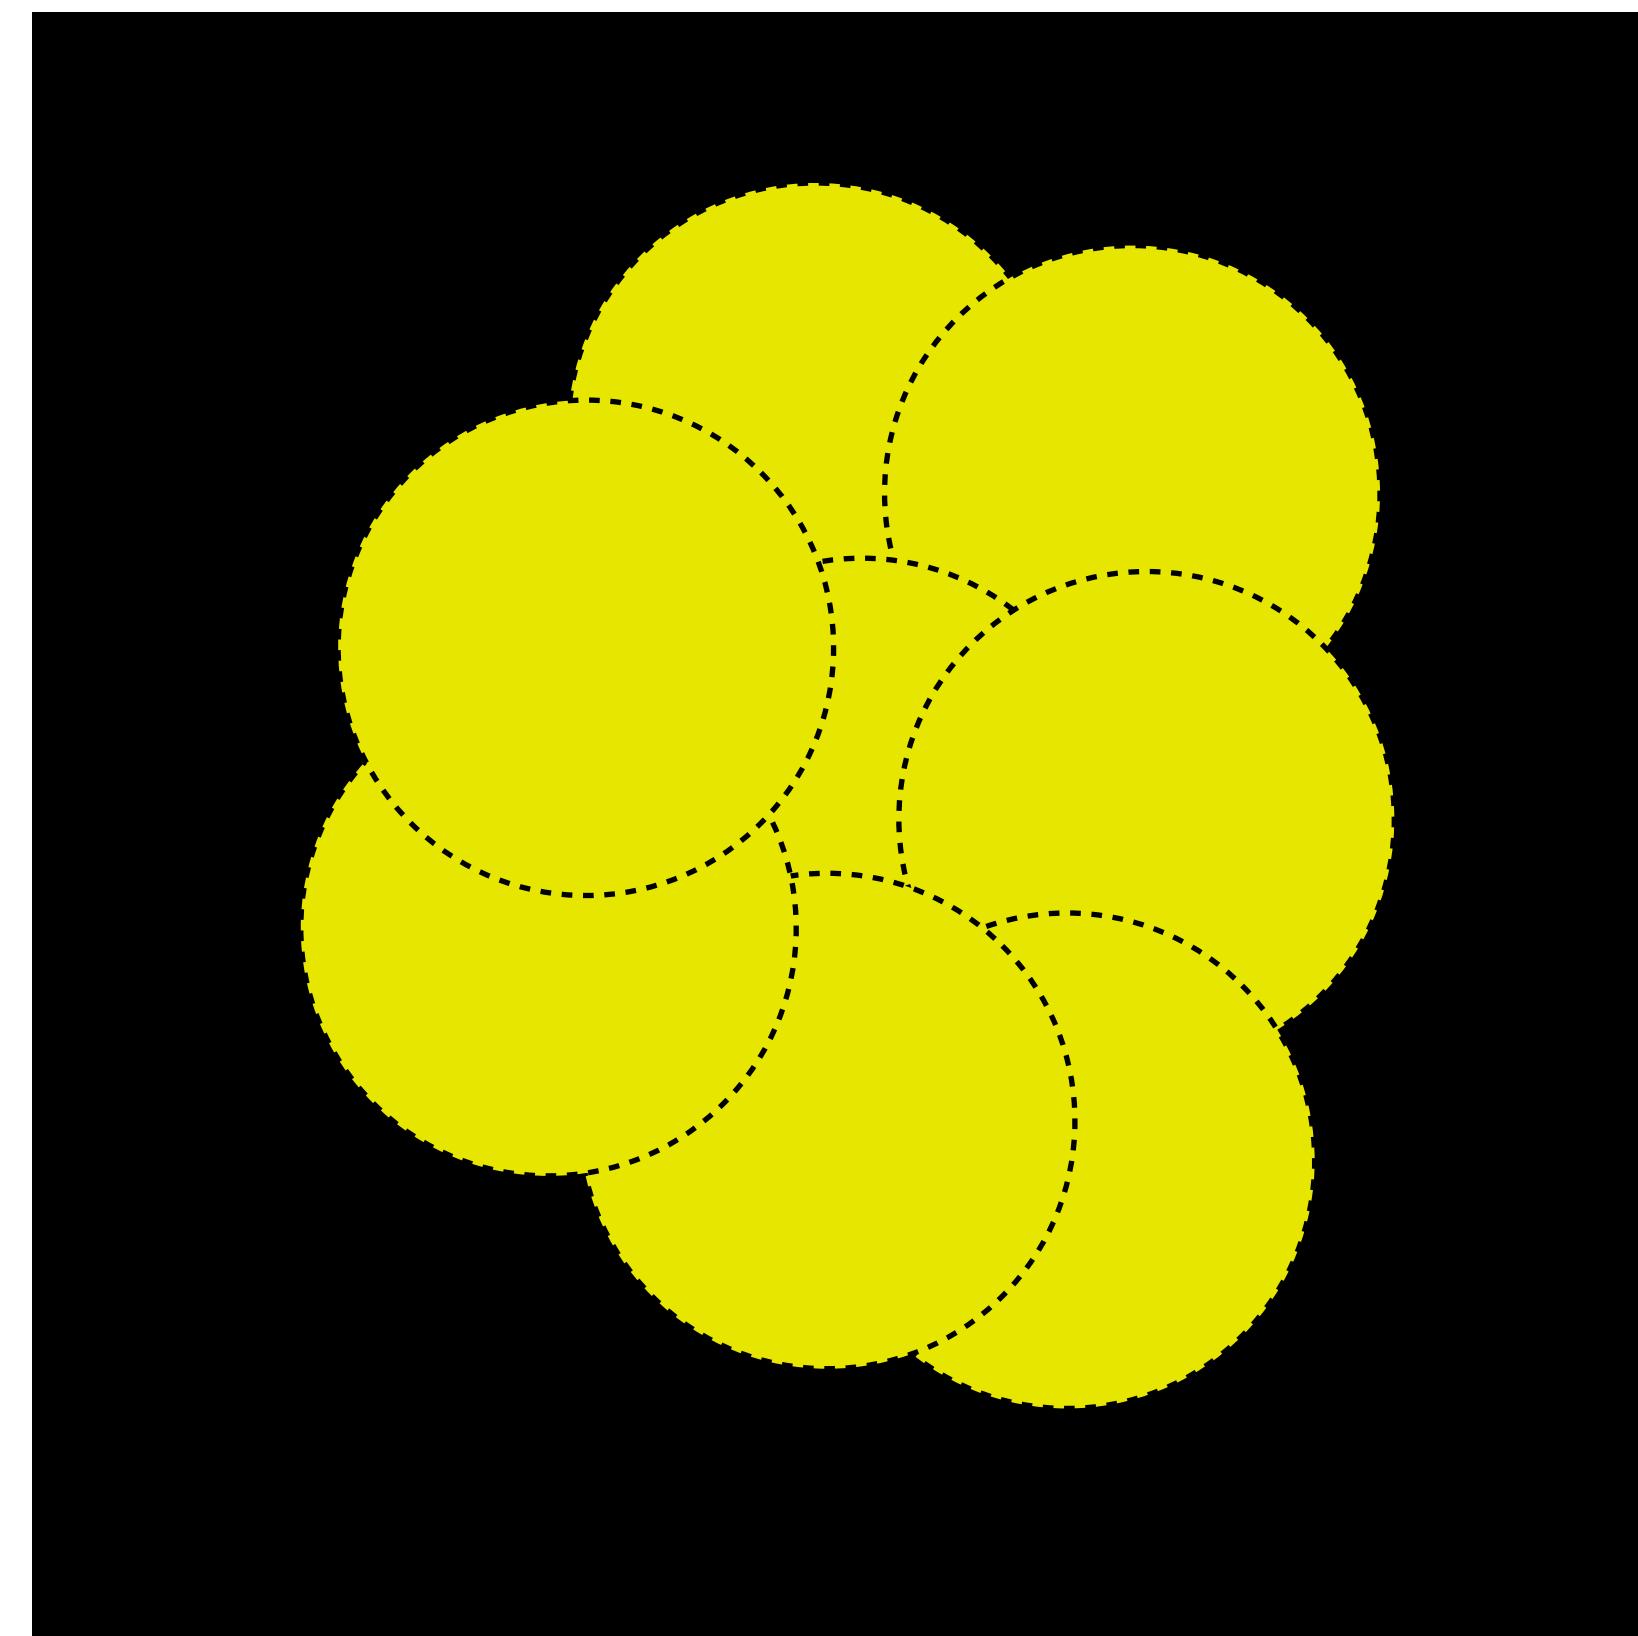
# Model of a Retinal Ganglion Cell



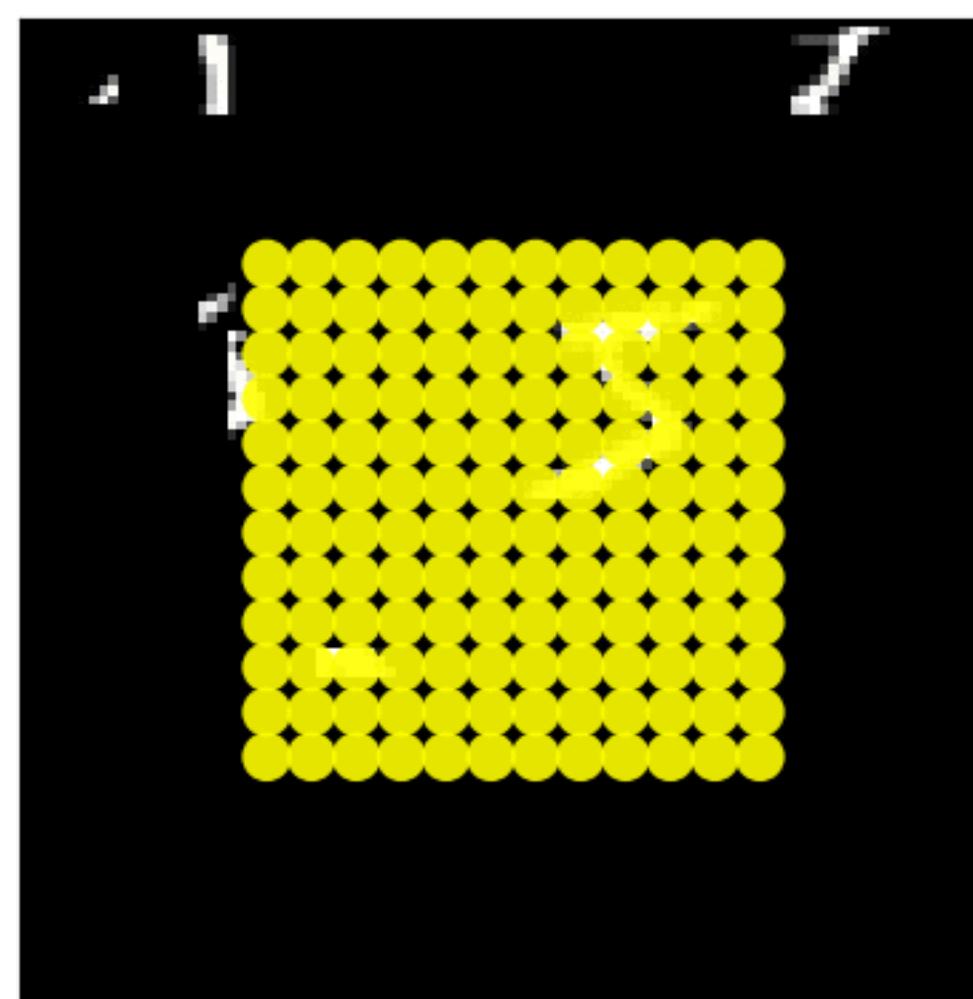
in-vivo



in-silico

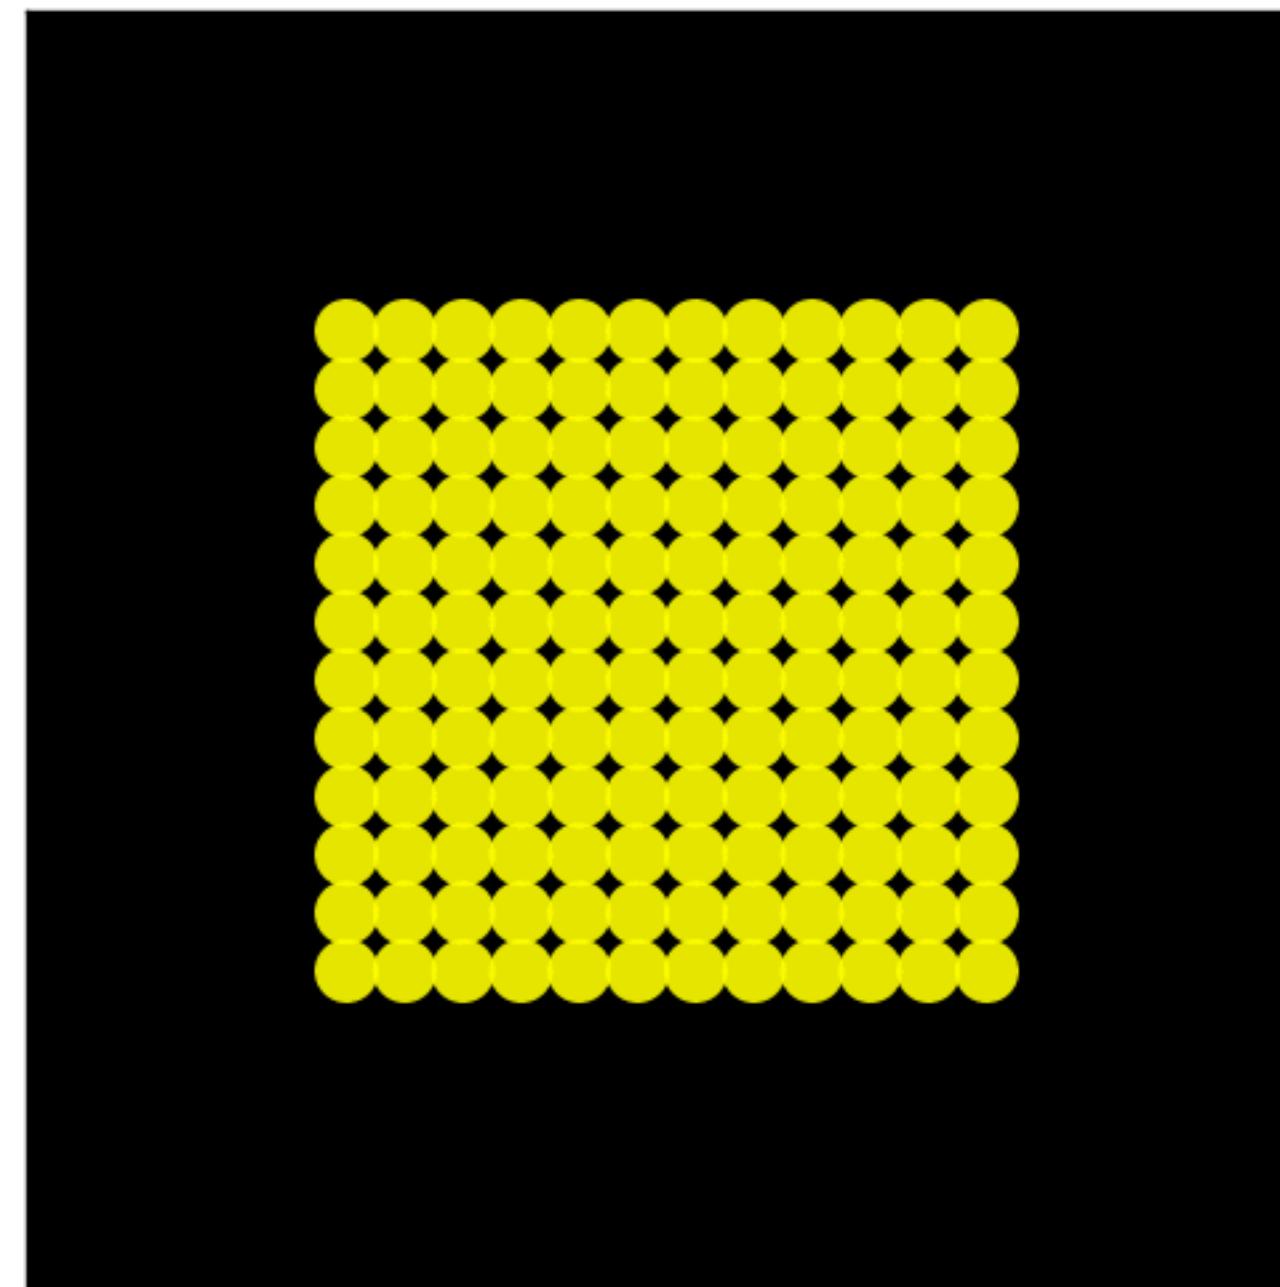


# Attention with Saccades



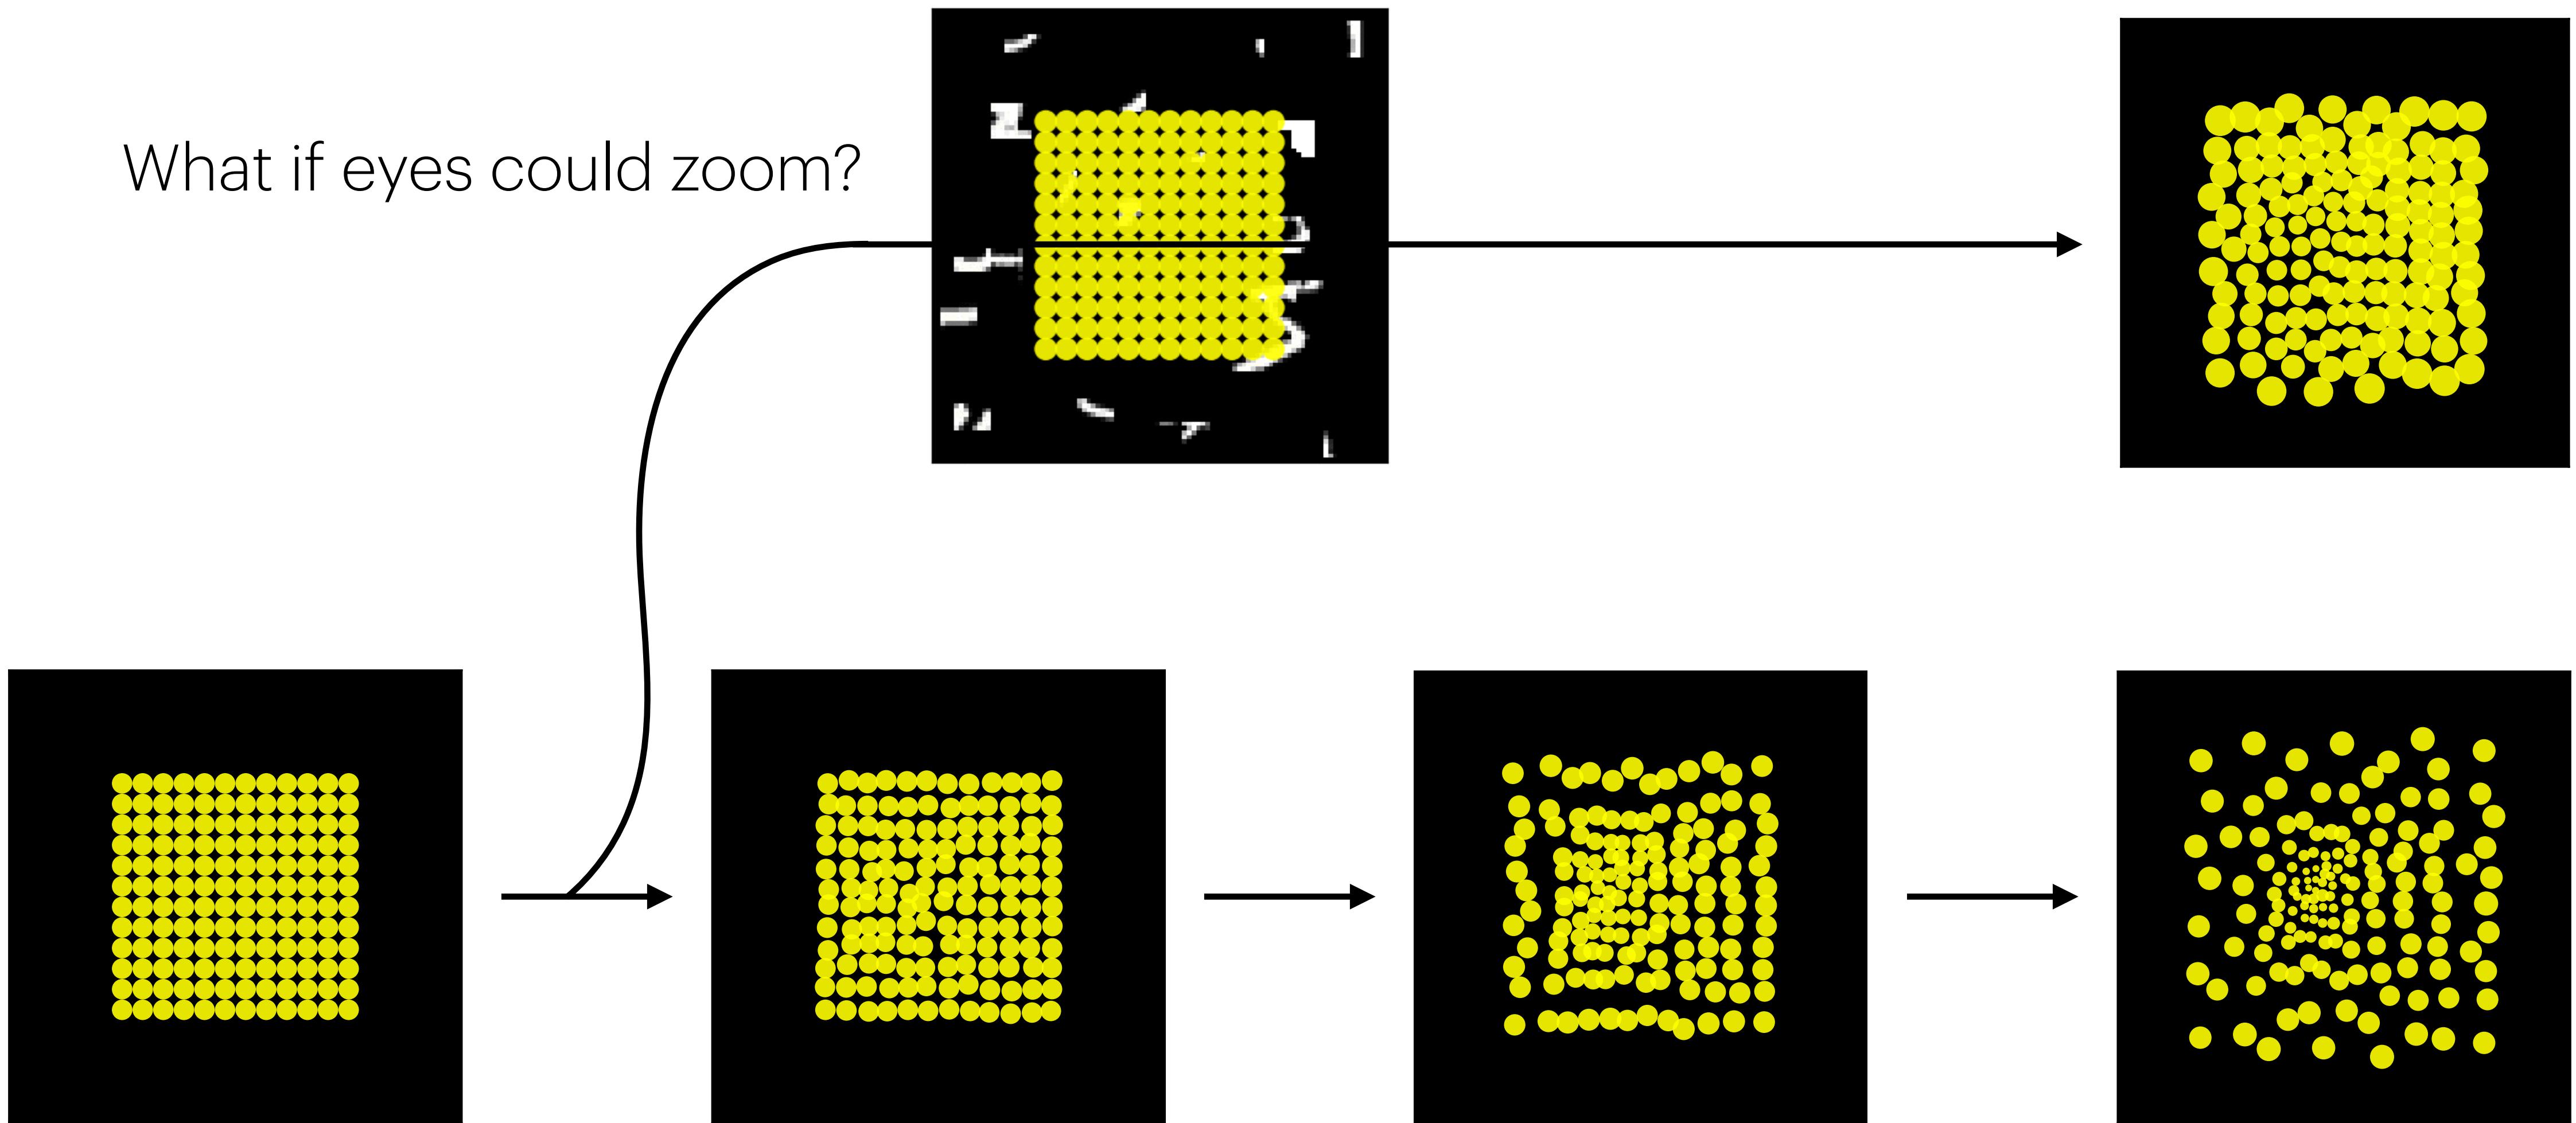
# Evolution of a fovea

Low Acuity      High Acuity



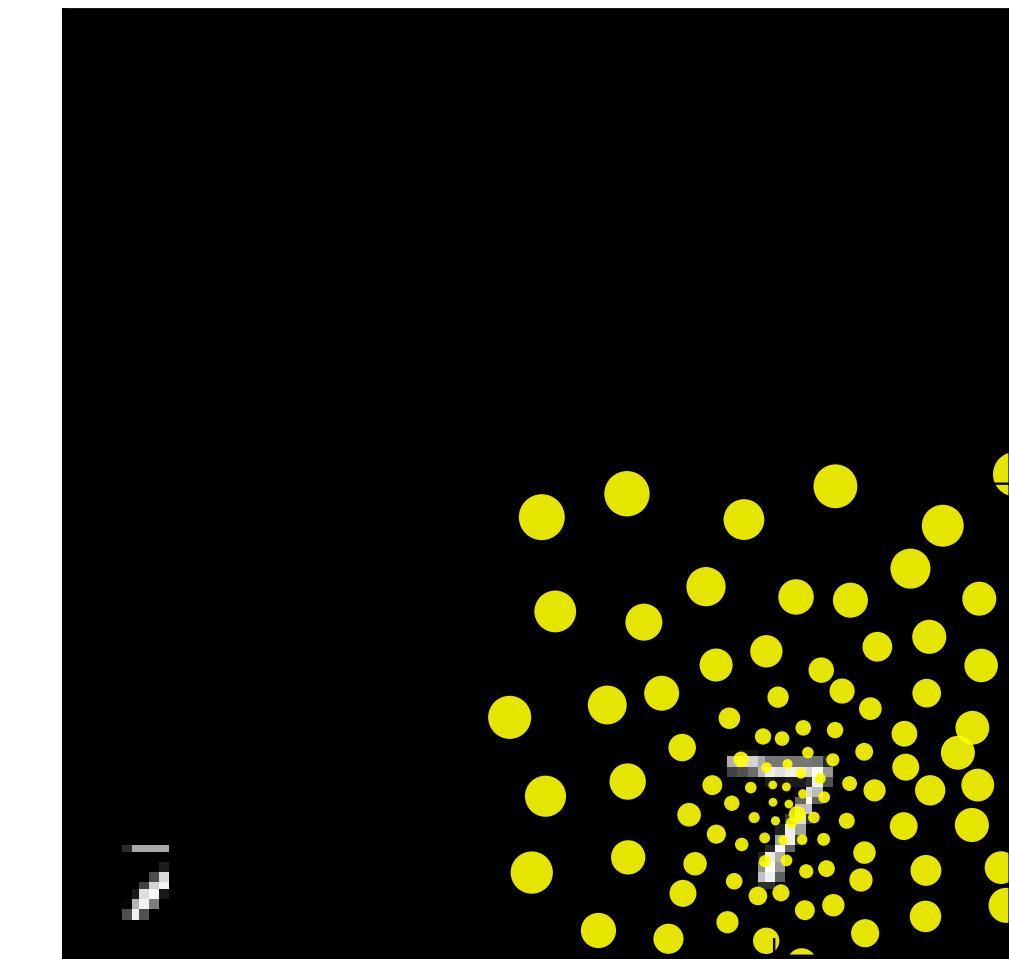
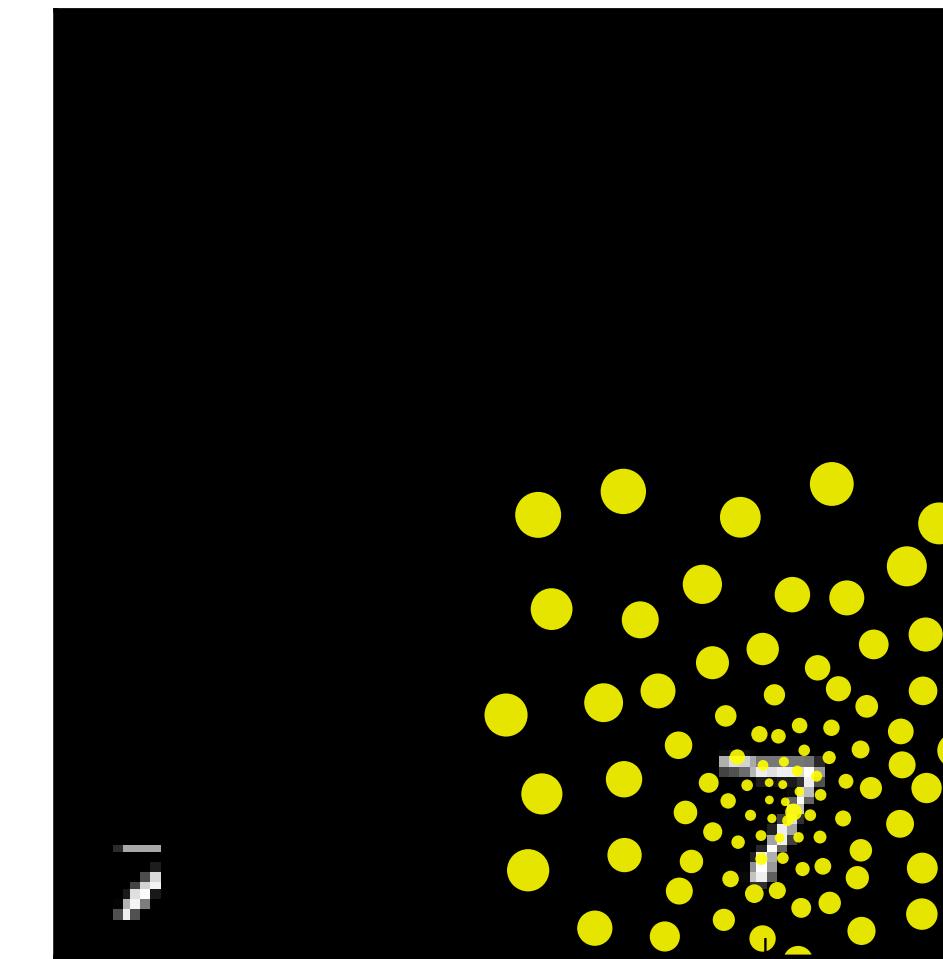
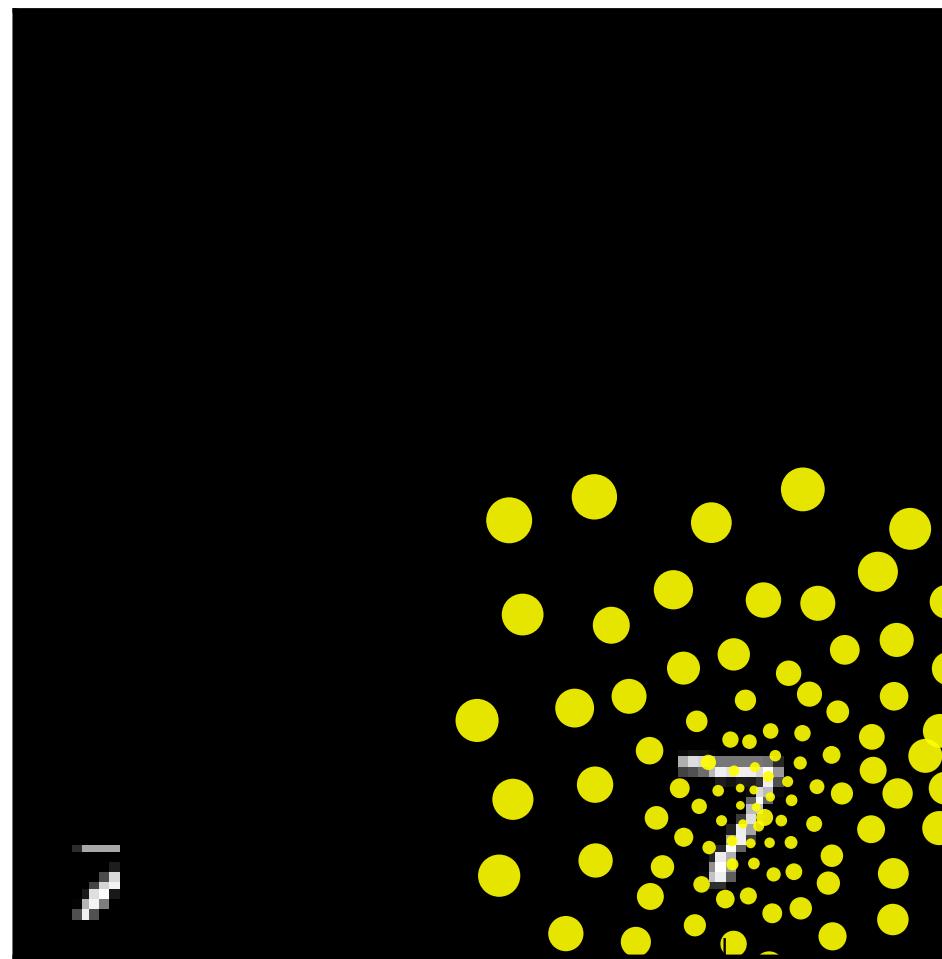
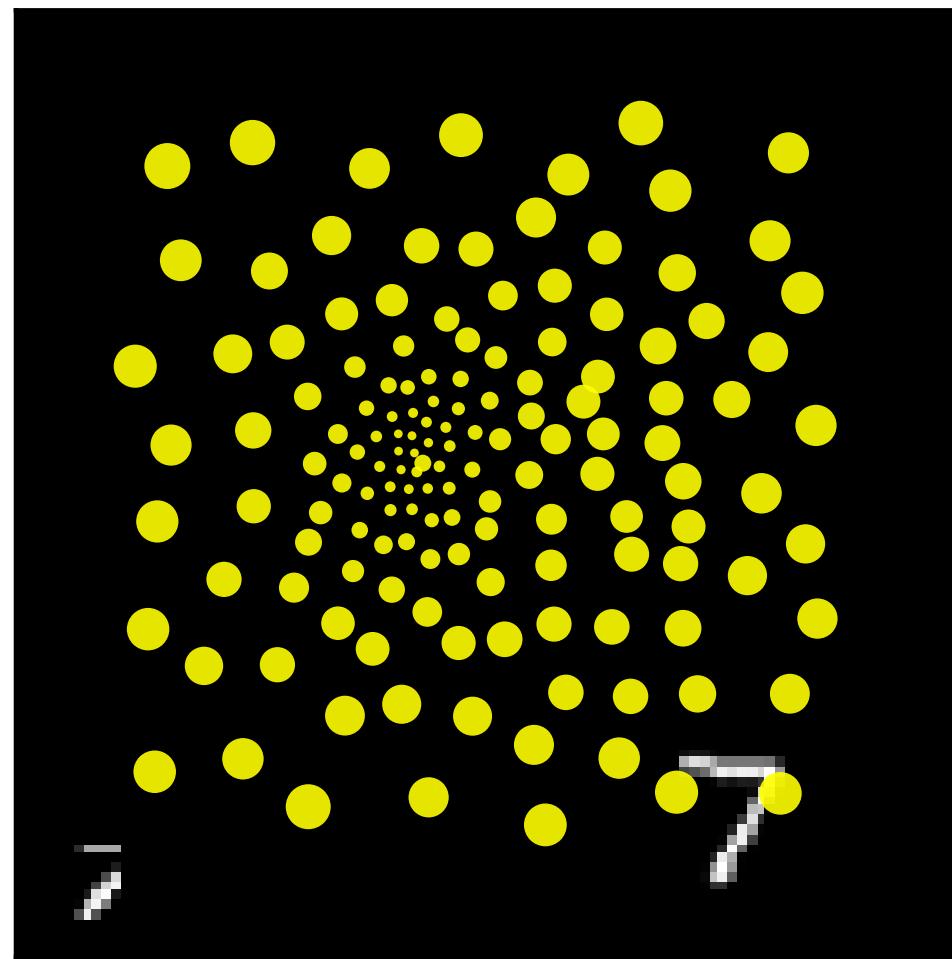
[Cheung, Weiss, Olshausen 2017]

What if eyes could zoom?



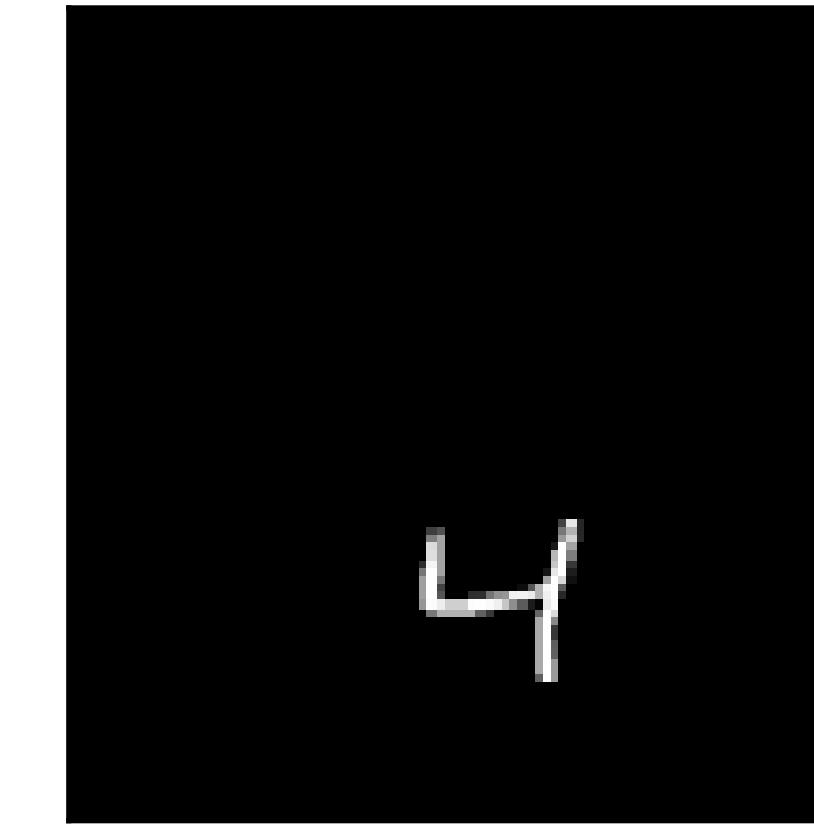
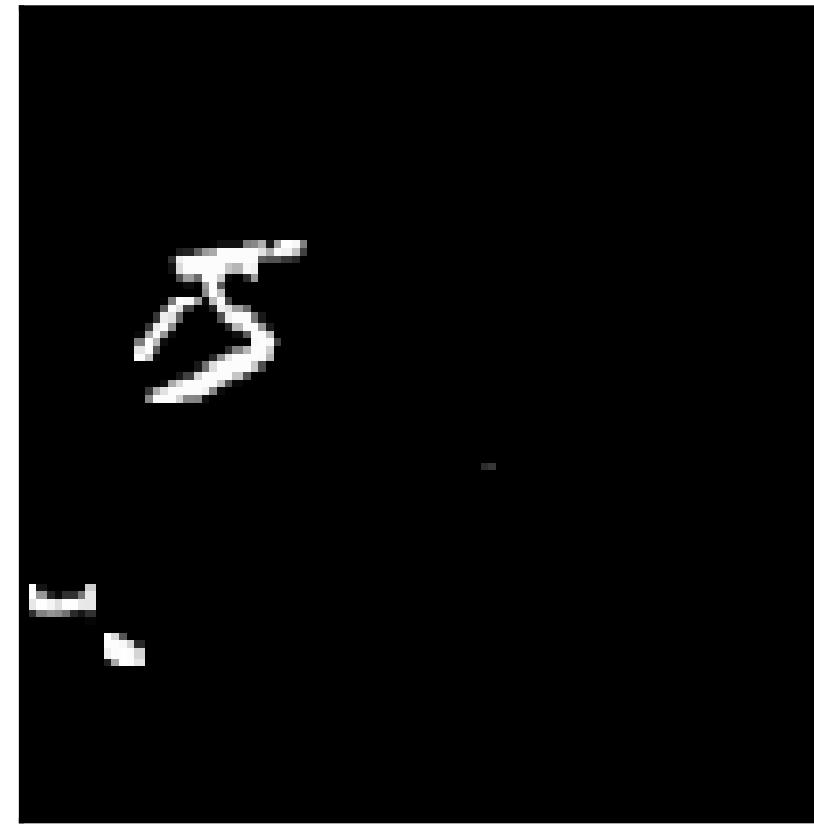
[Cheung, Weiss, Olshausen 2017]

## How to use a fovea

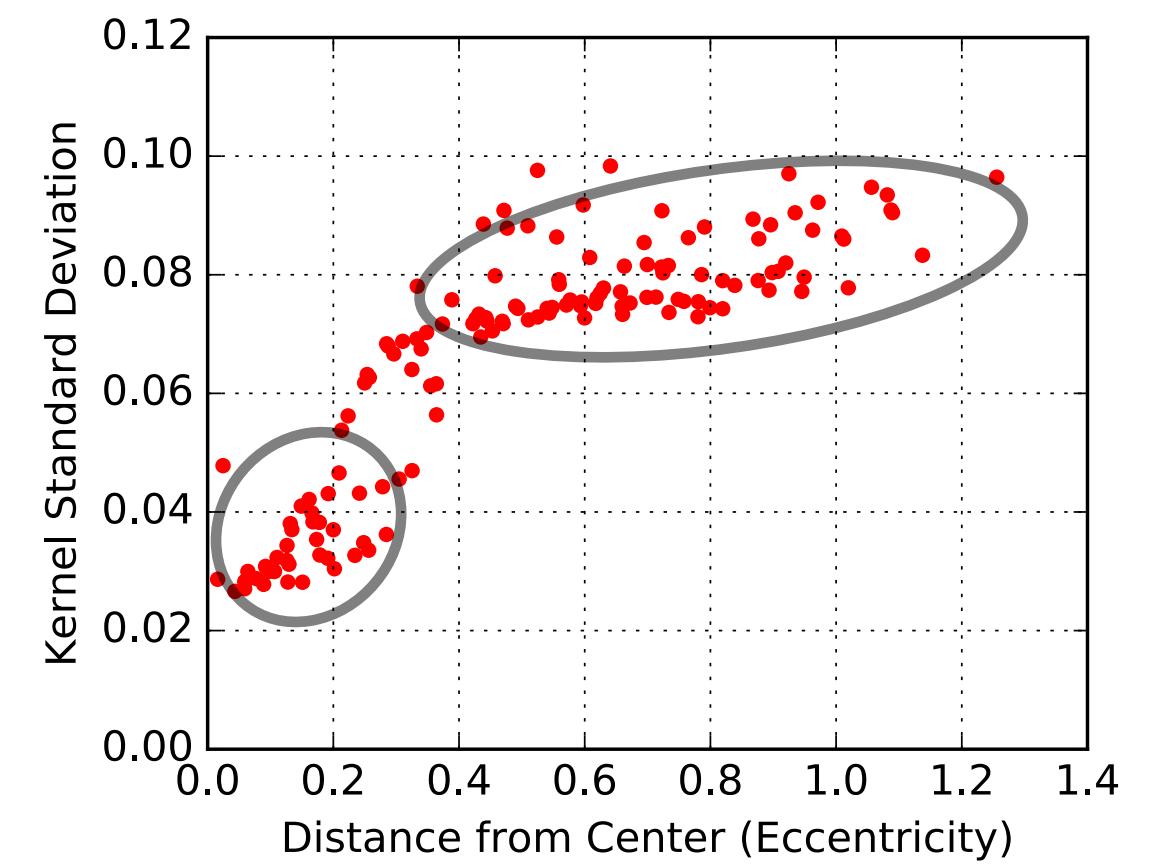
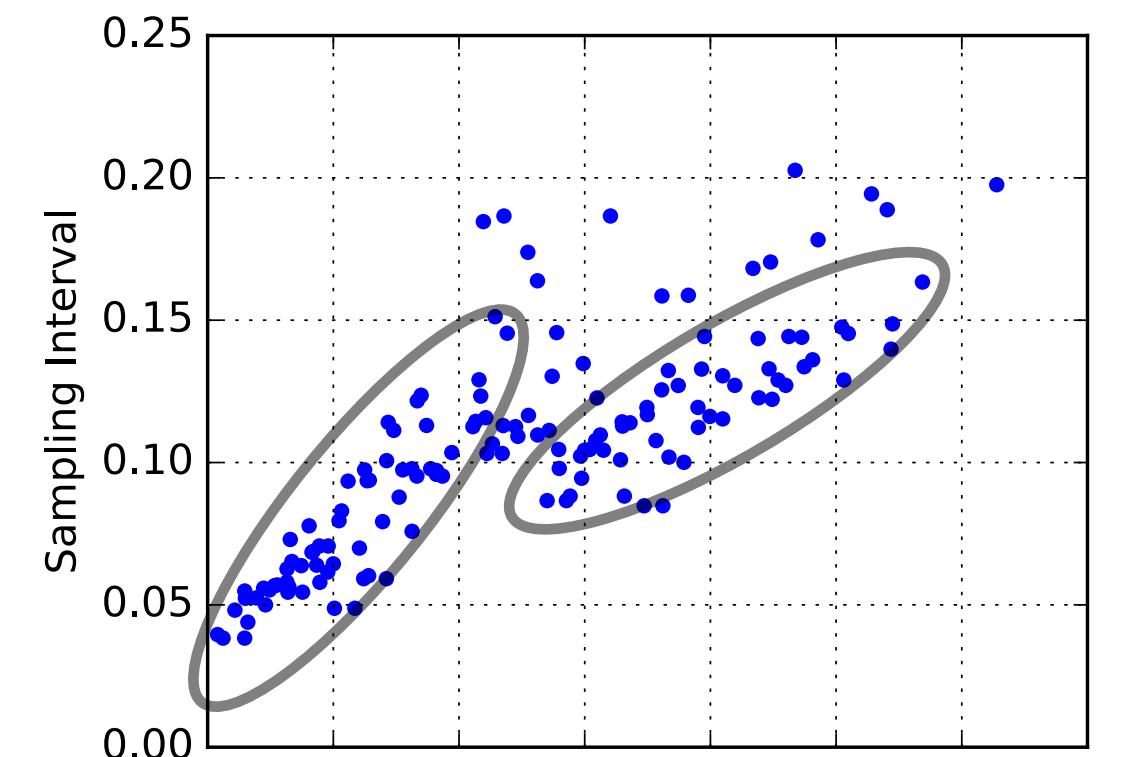
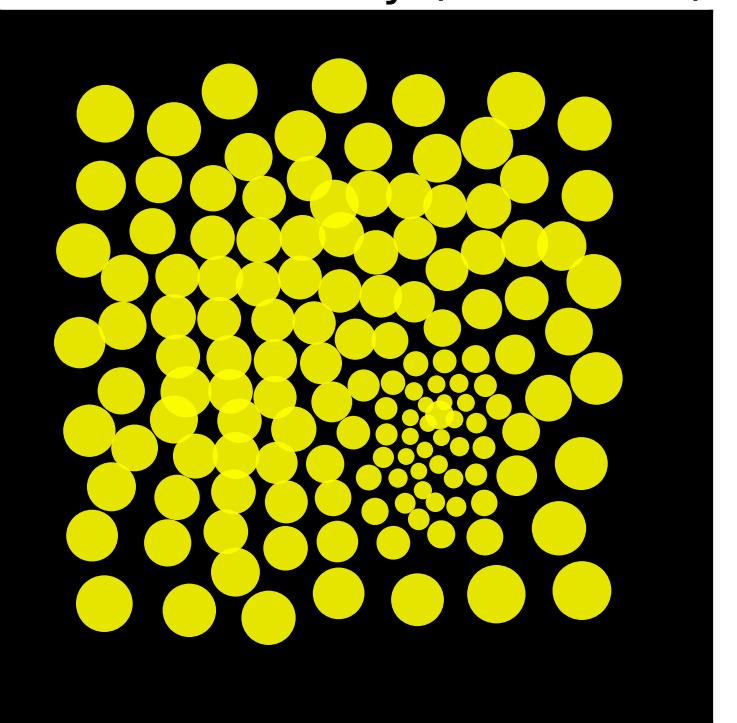


What if the data were different?

Dataset 1

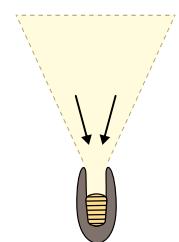


Translation Only (Dataset 1)

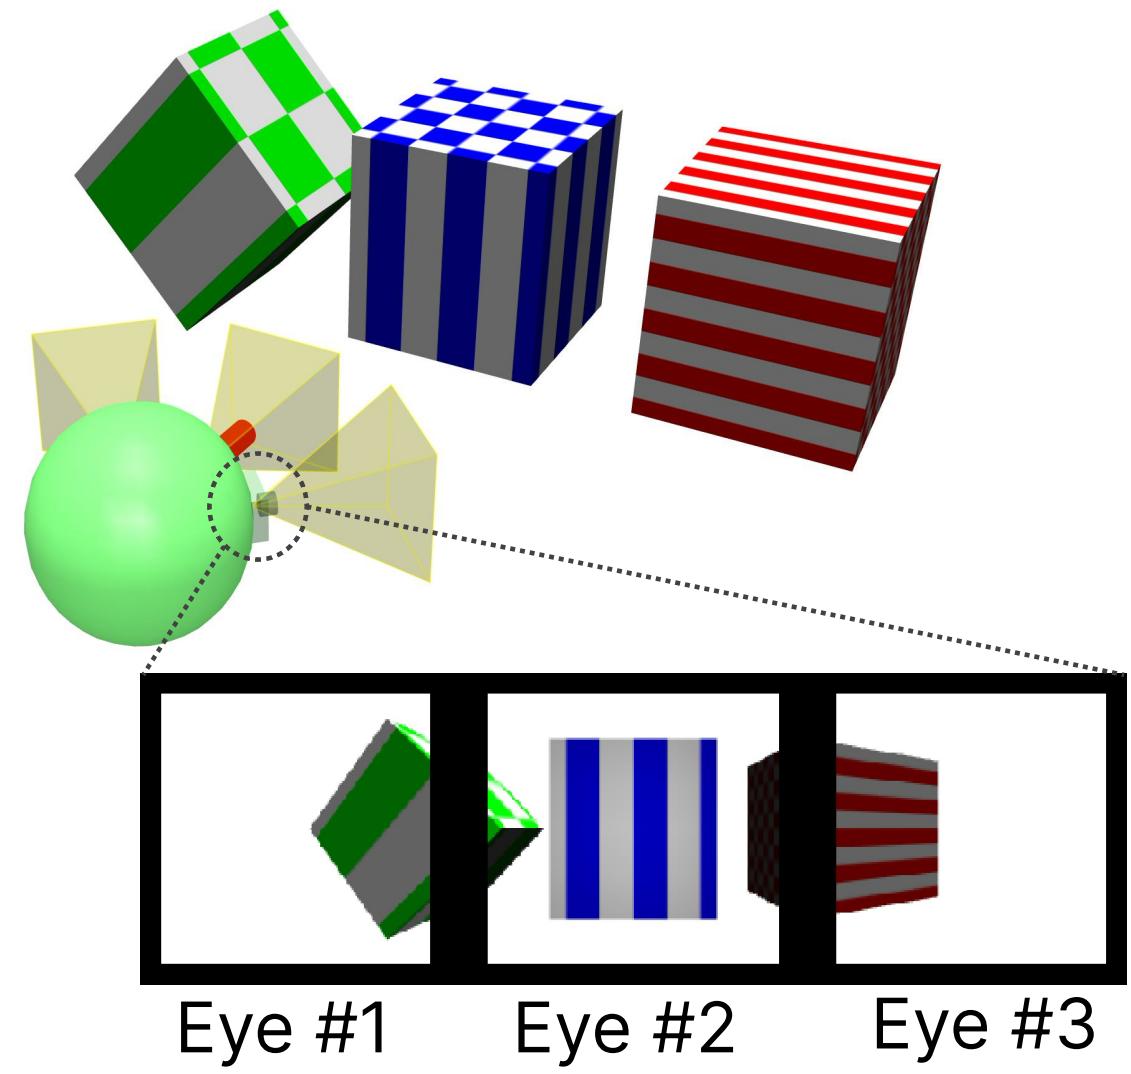


# Building a '**What if**' machine for vision

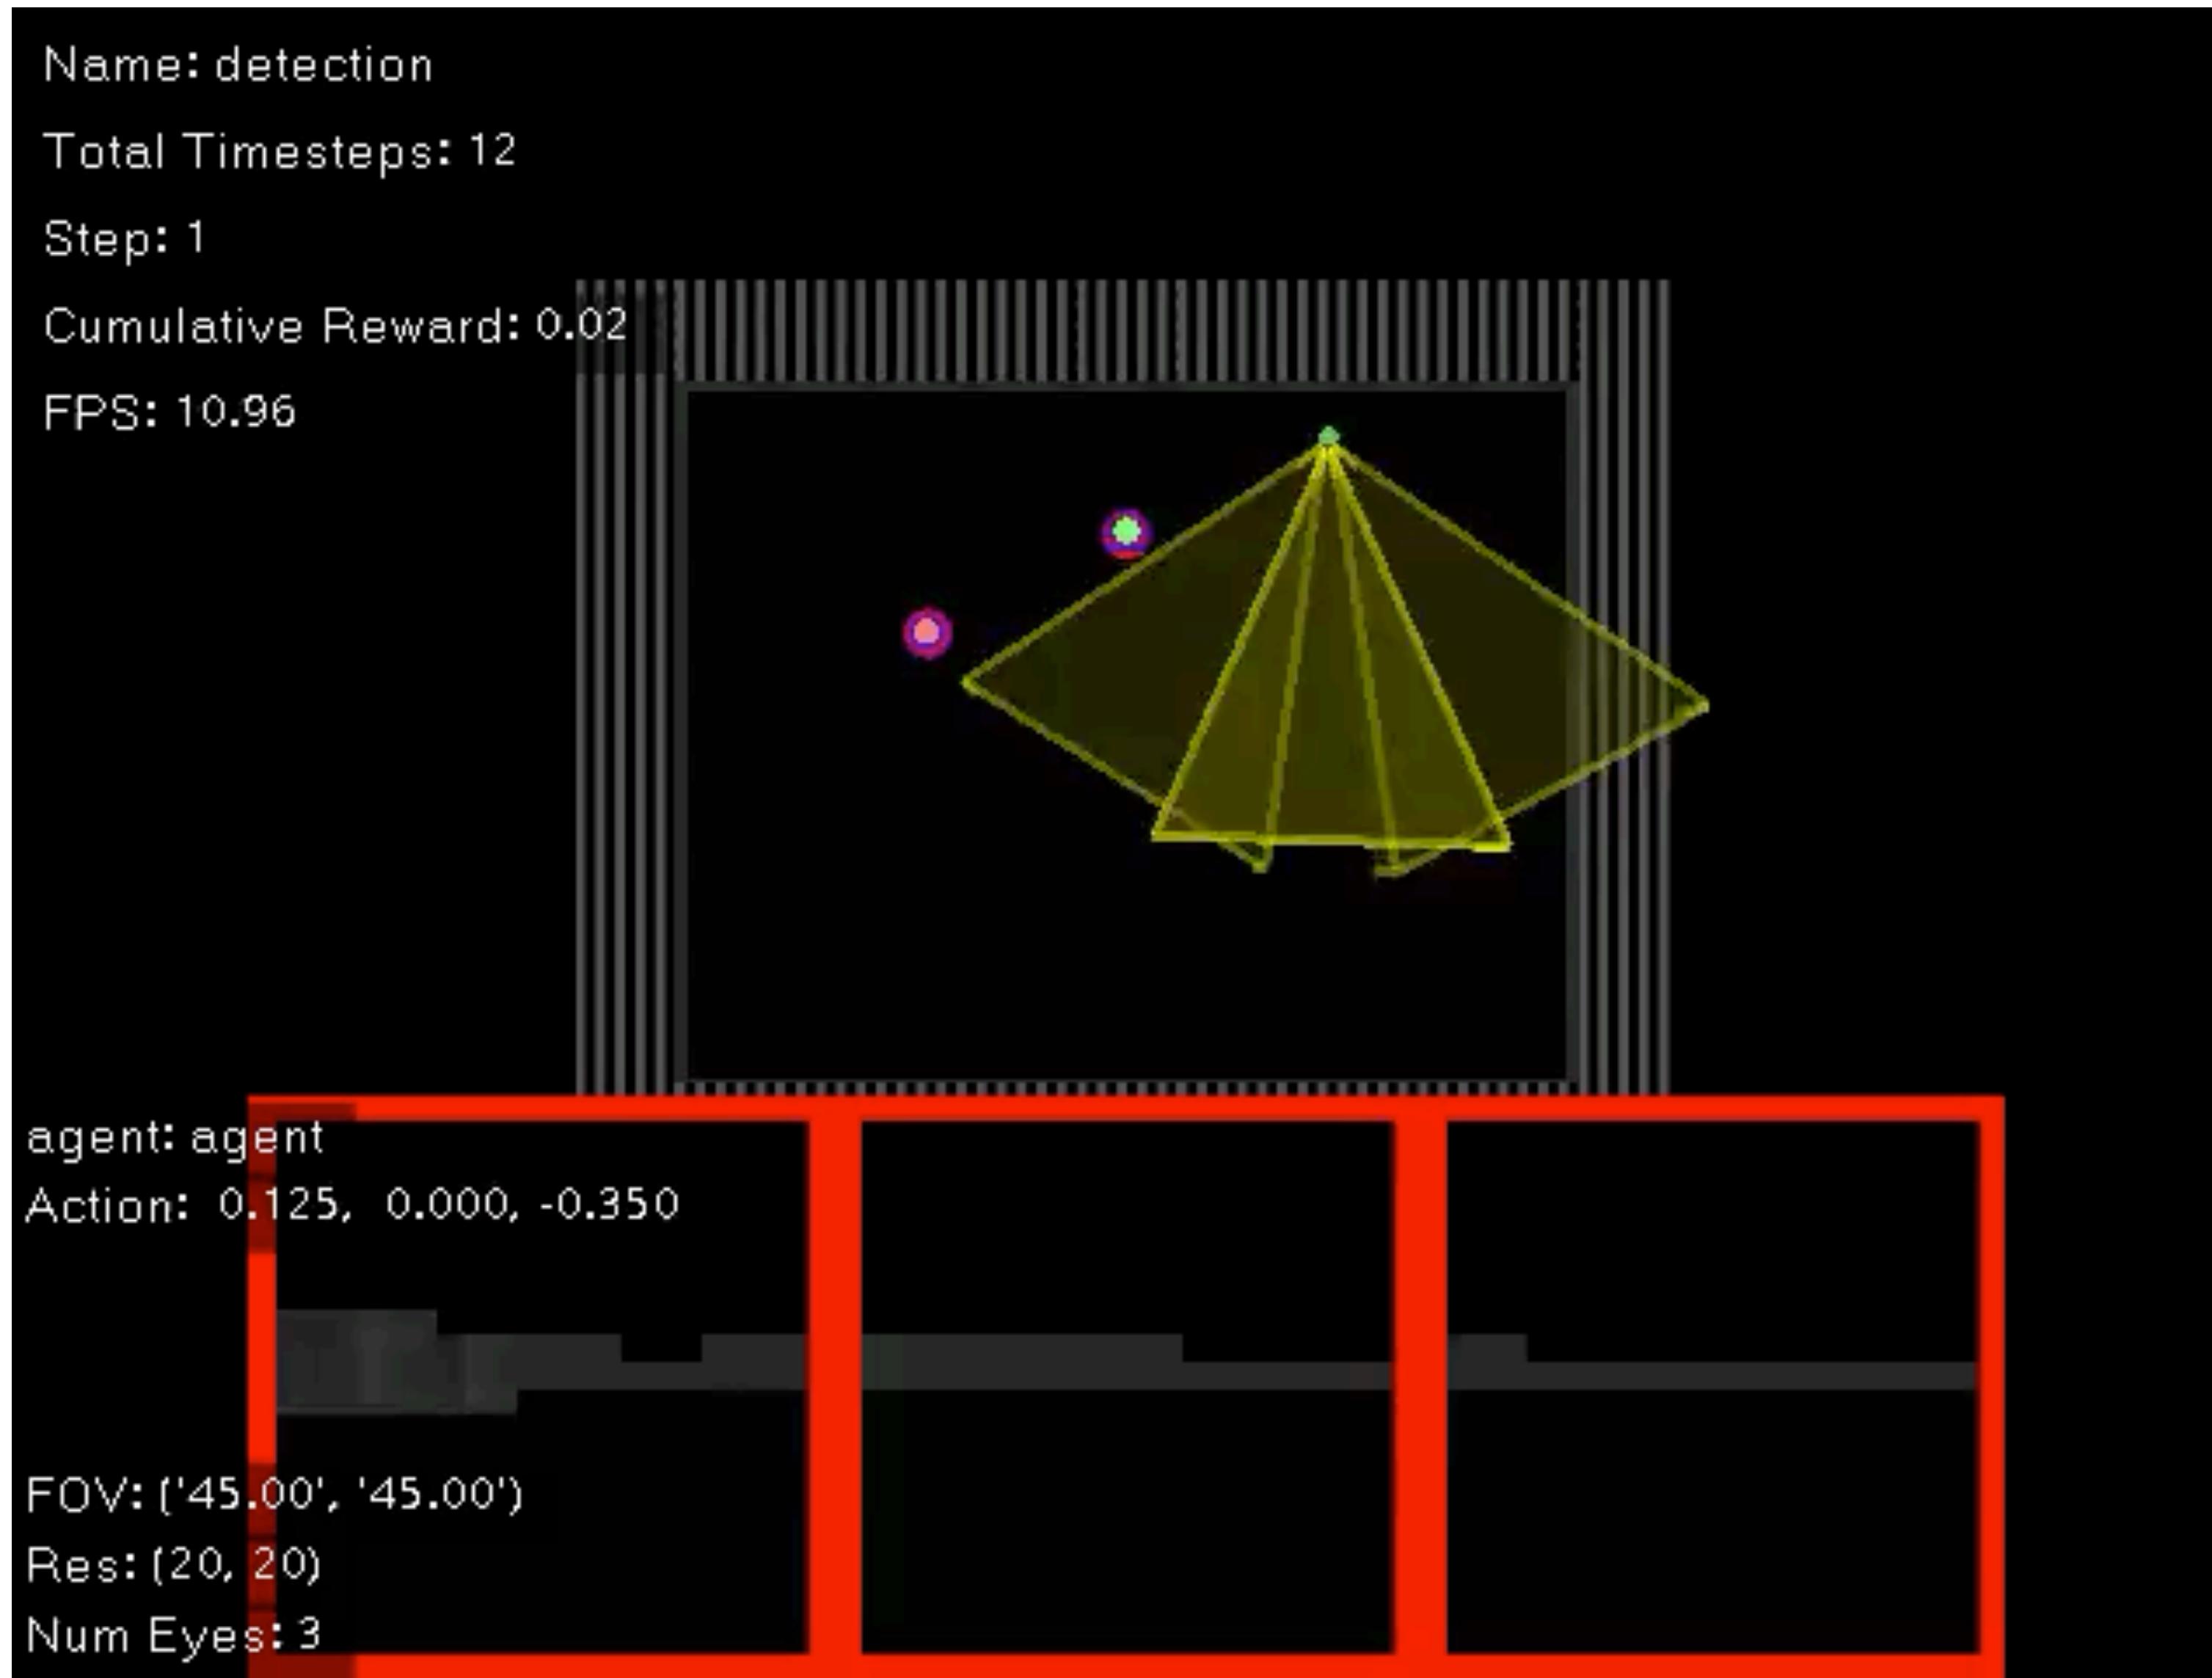
Photoreceptor



# Evolution of an “AI-ball”



# Reinforcement Learning Environments

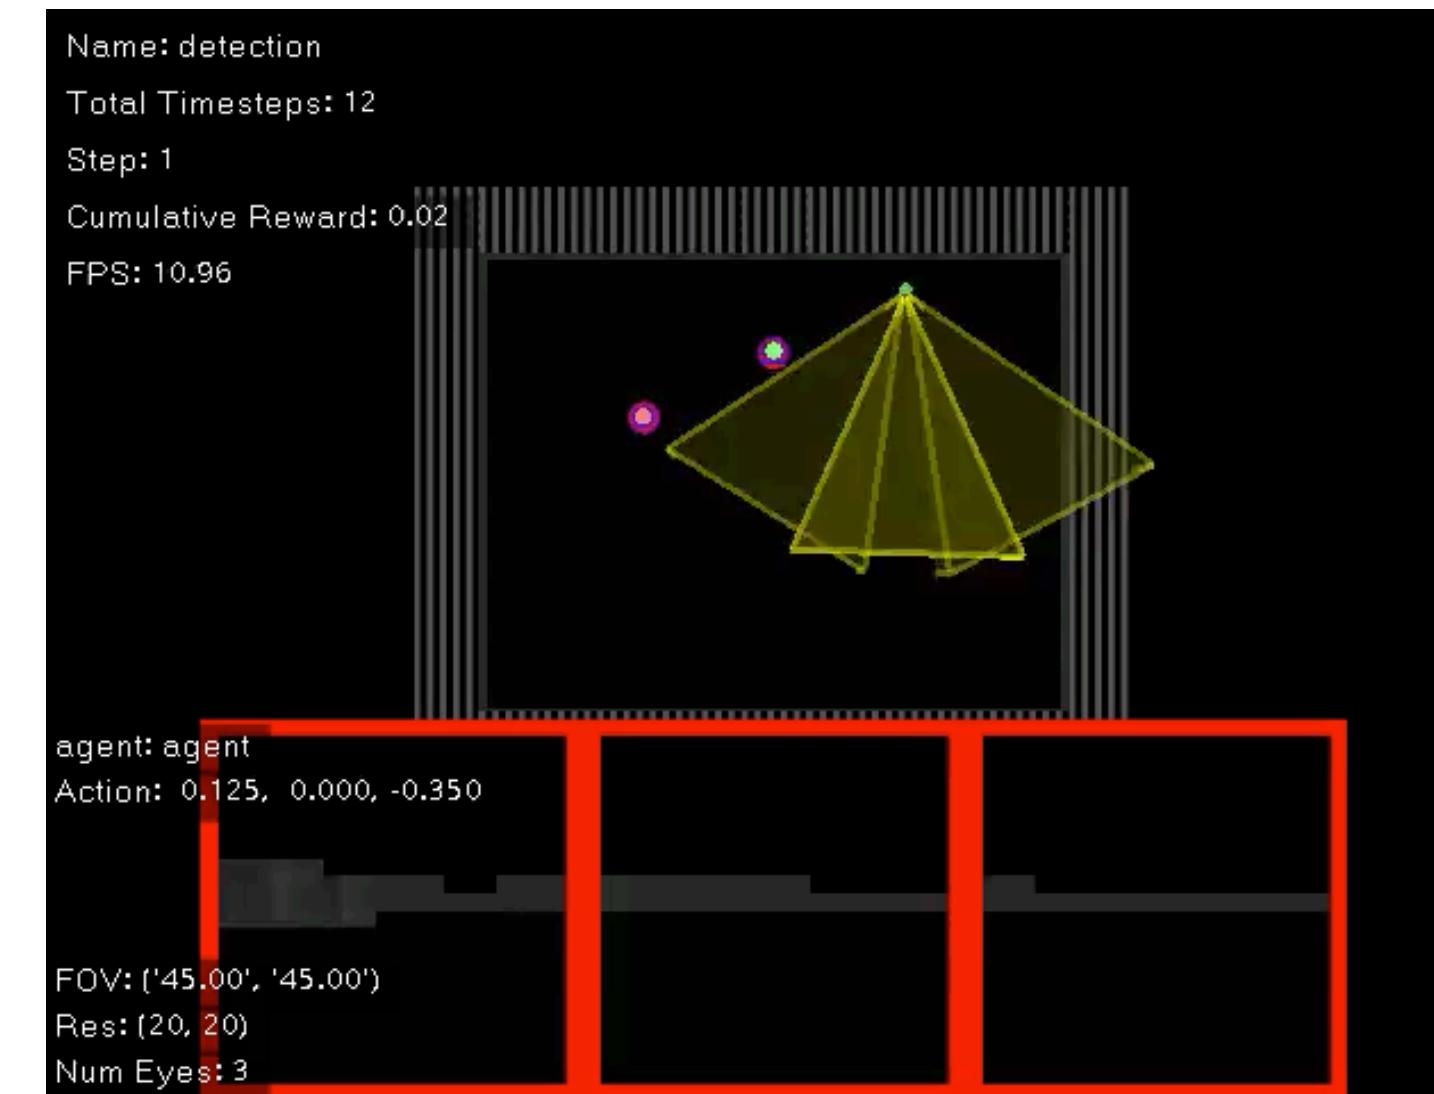


# Reinforcement Learning Environments

## Navigation



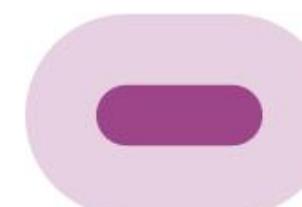
## Detection



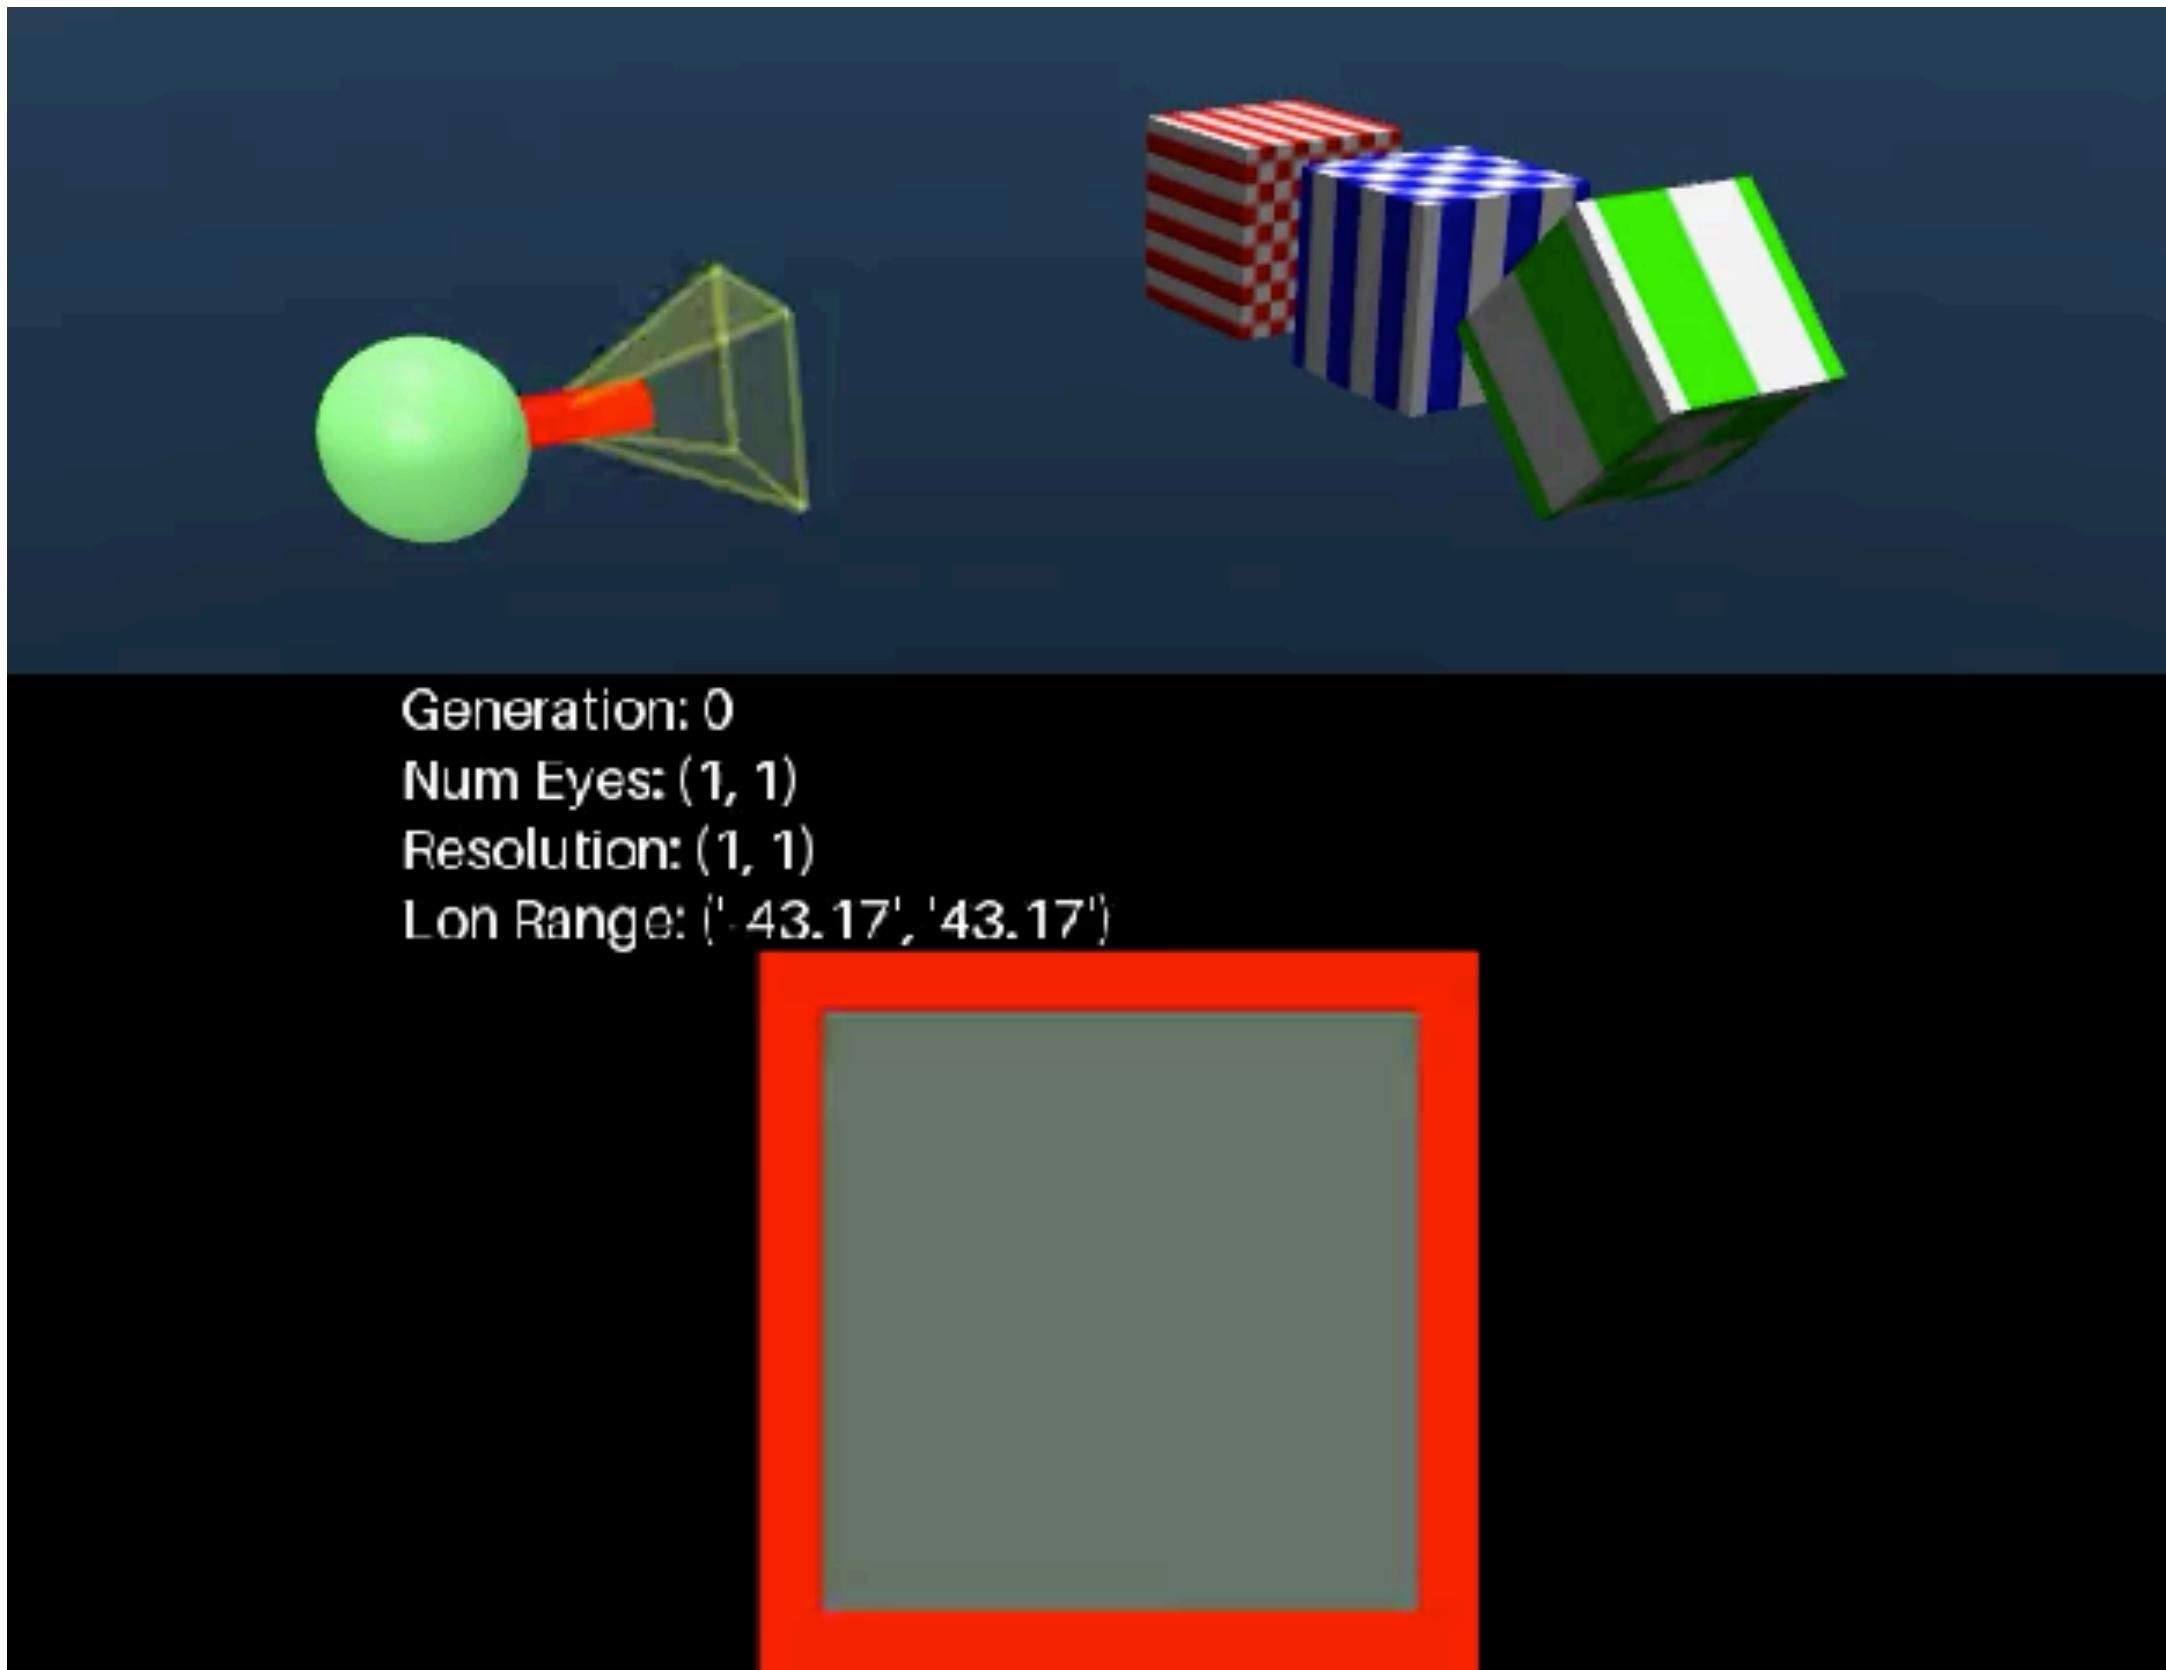
## Tracking



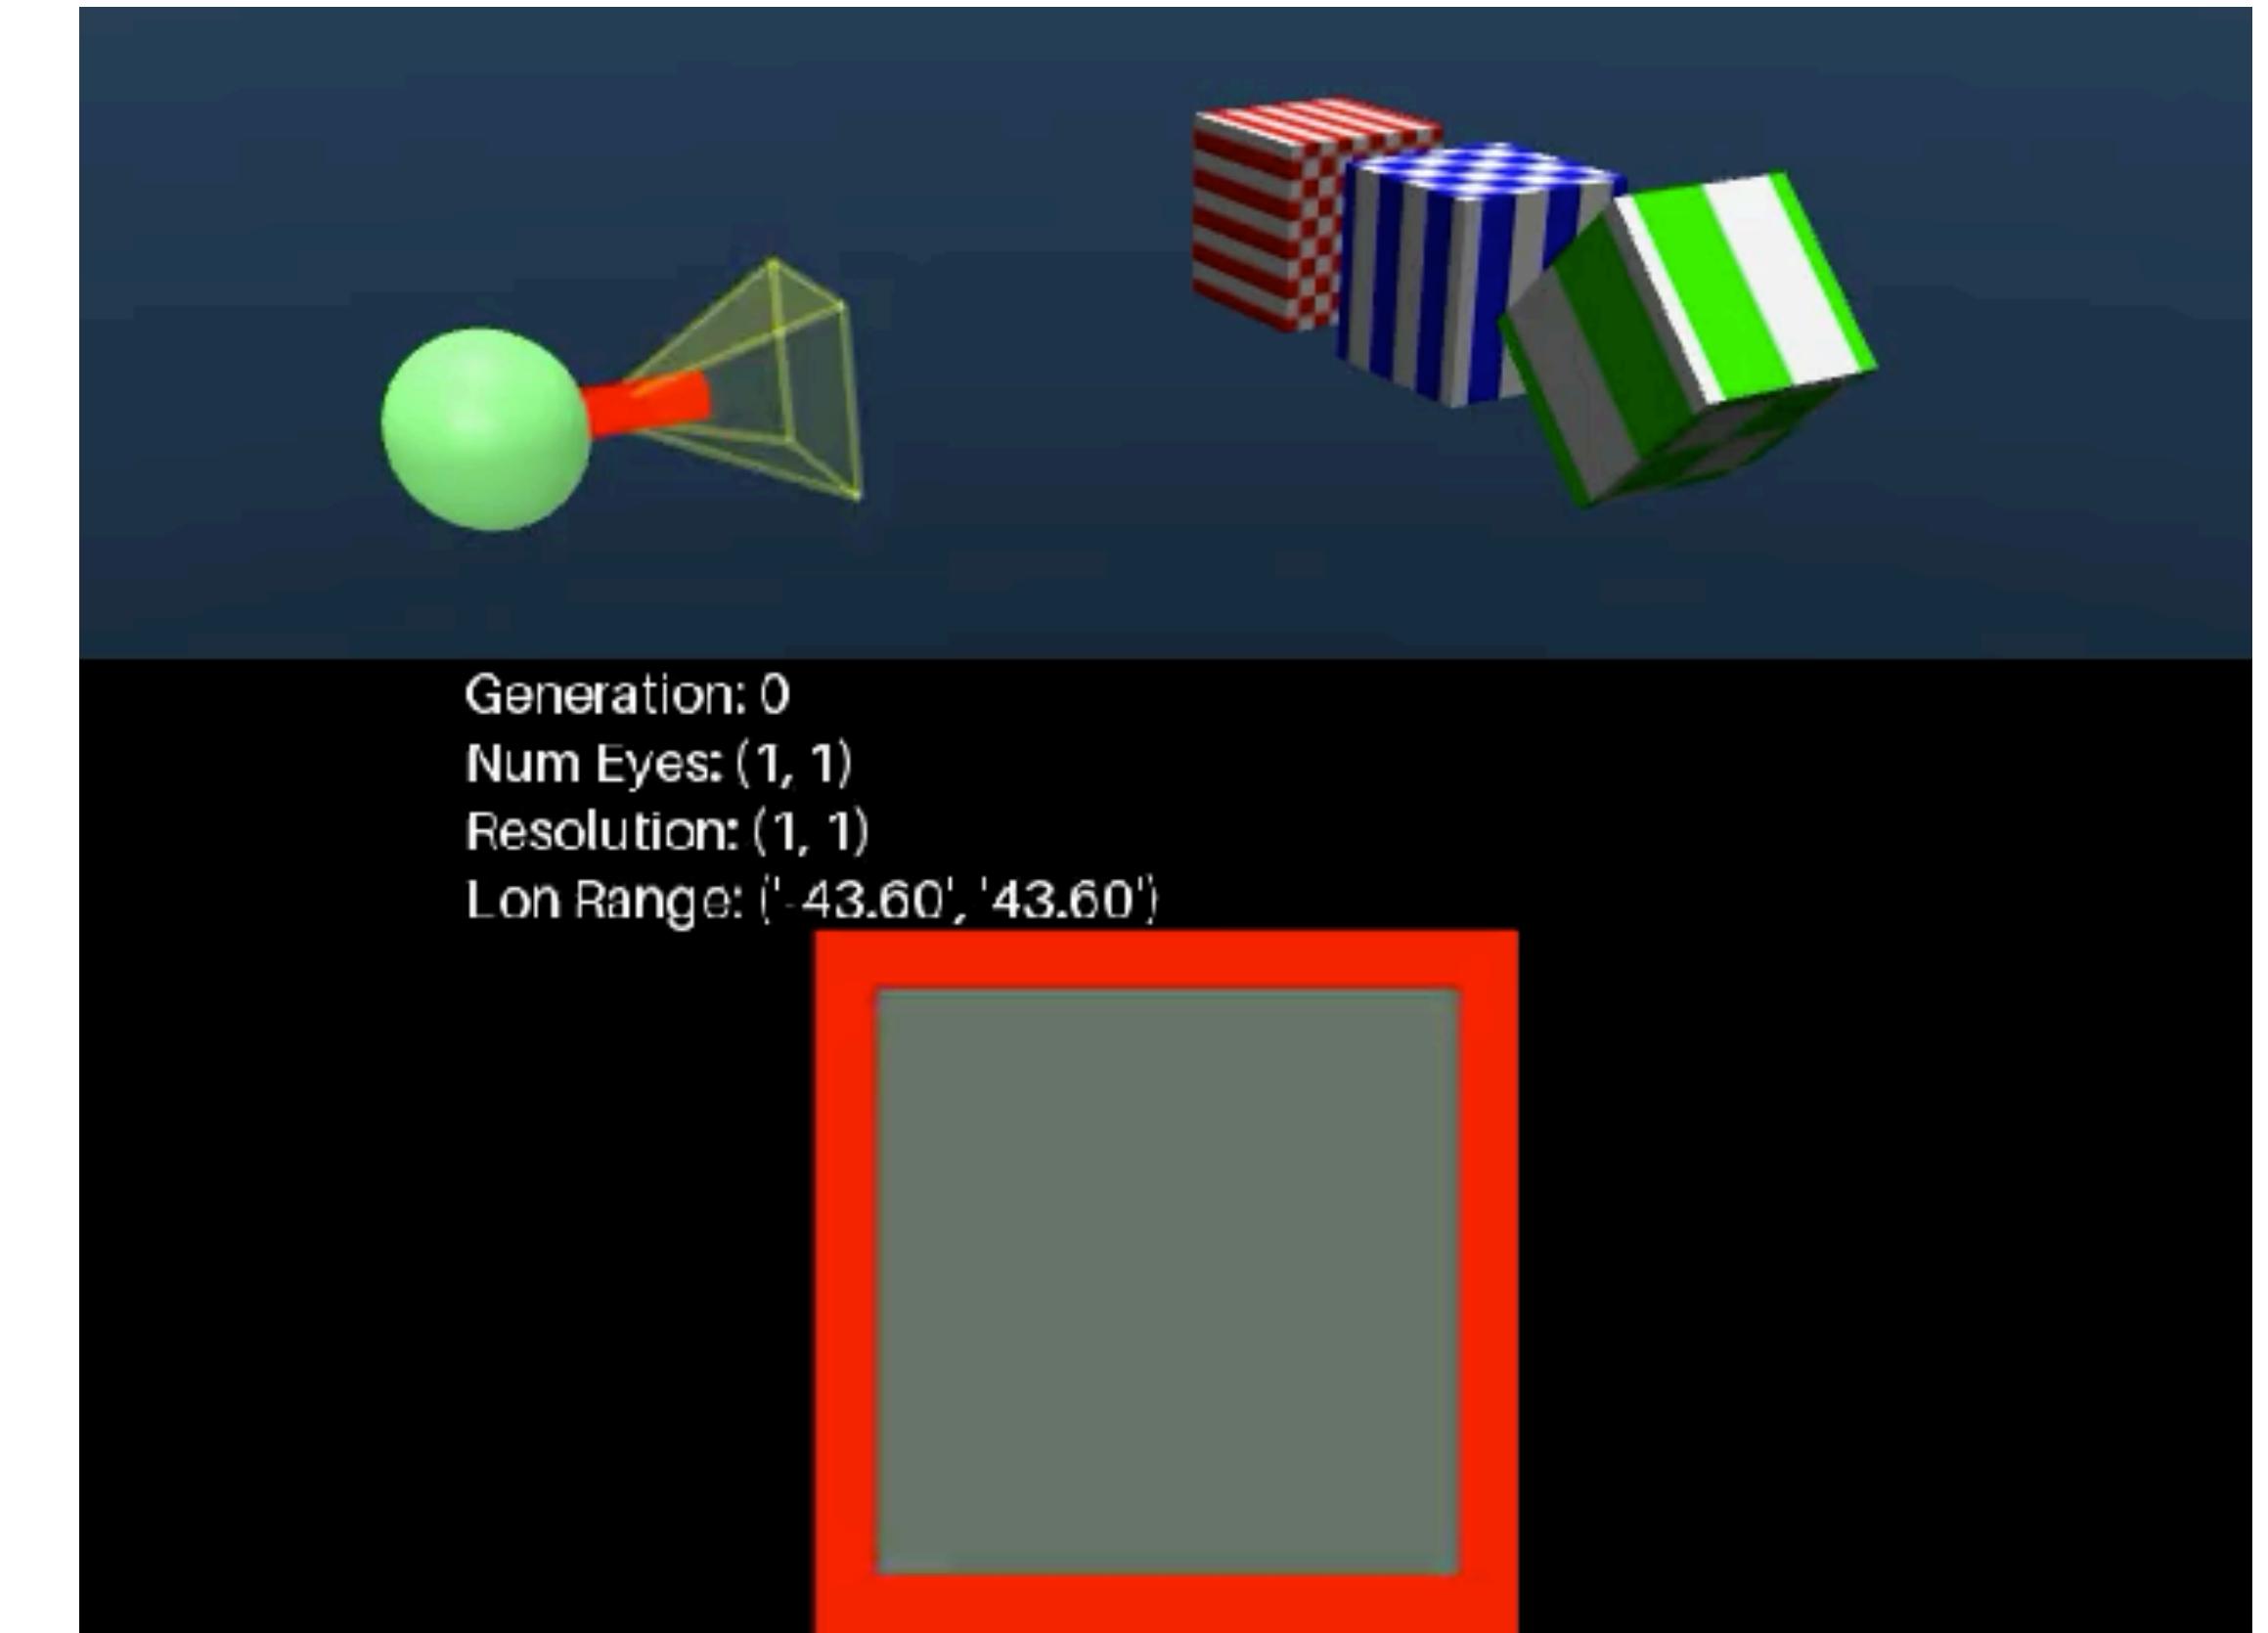
# What if the ~~data~~ goals were different?



## Navigation

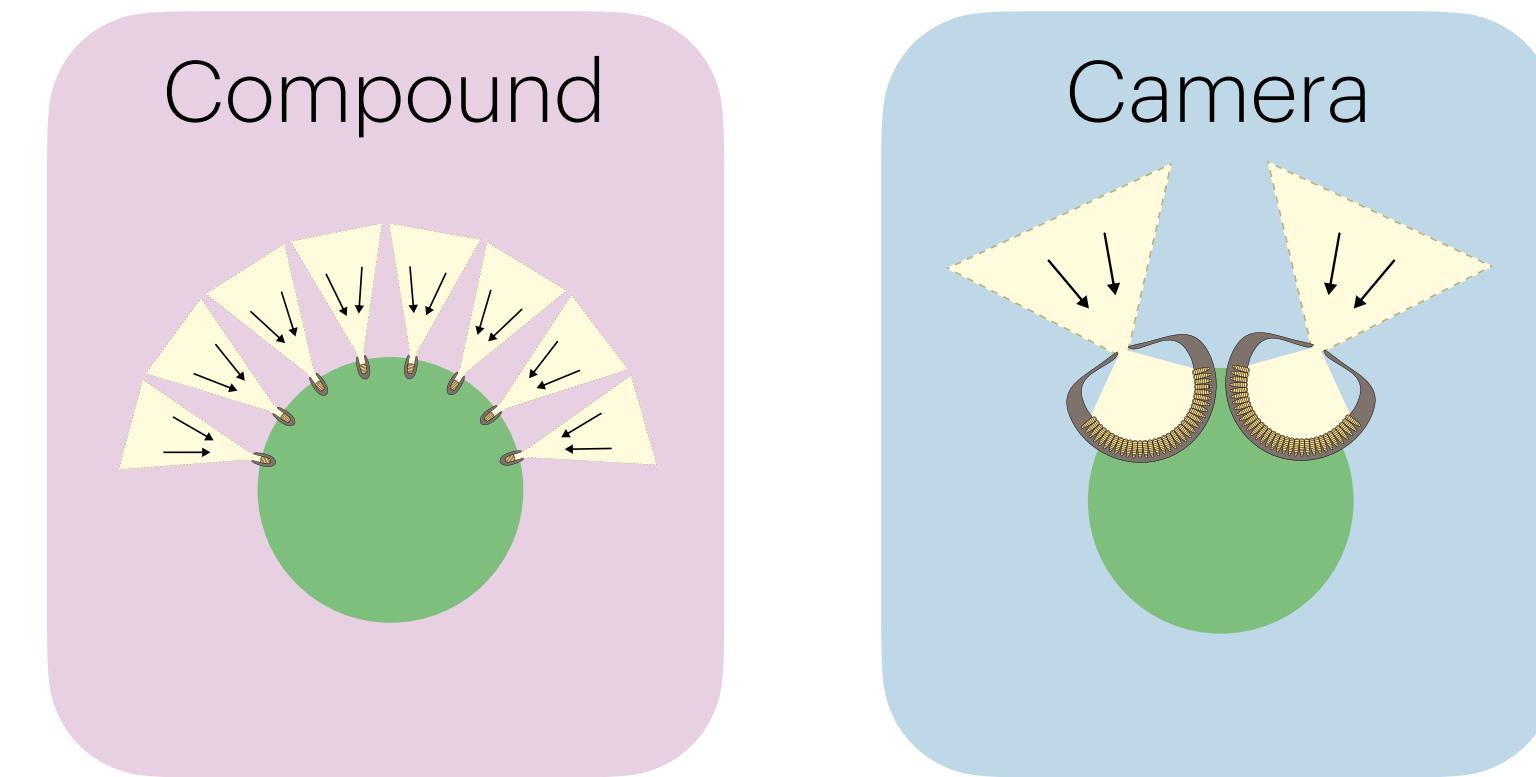


## Detection



[Tiwary\*, Young\*, Klinghoffer, Tasneem, Dave, Poggio, Nilsson, **Cheung\*\***, Raskar\*\* 2025]

What if the ~~data~~ goals were different?



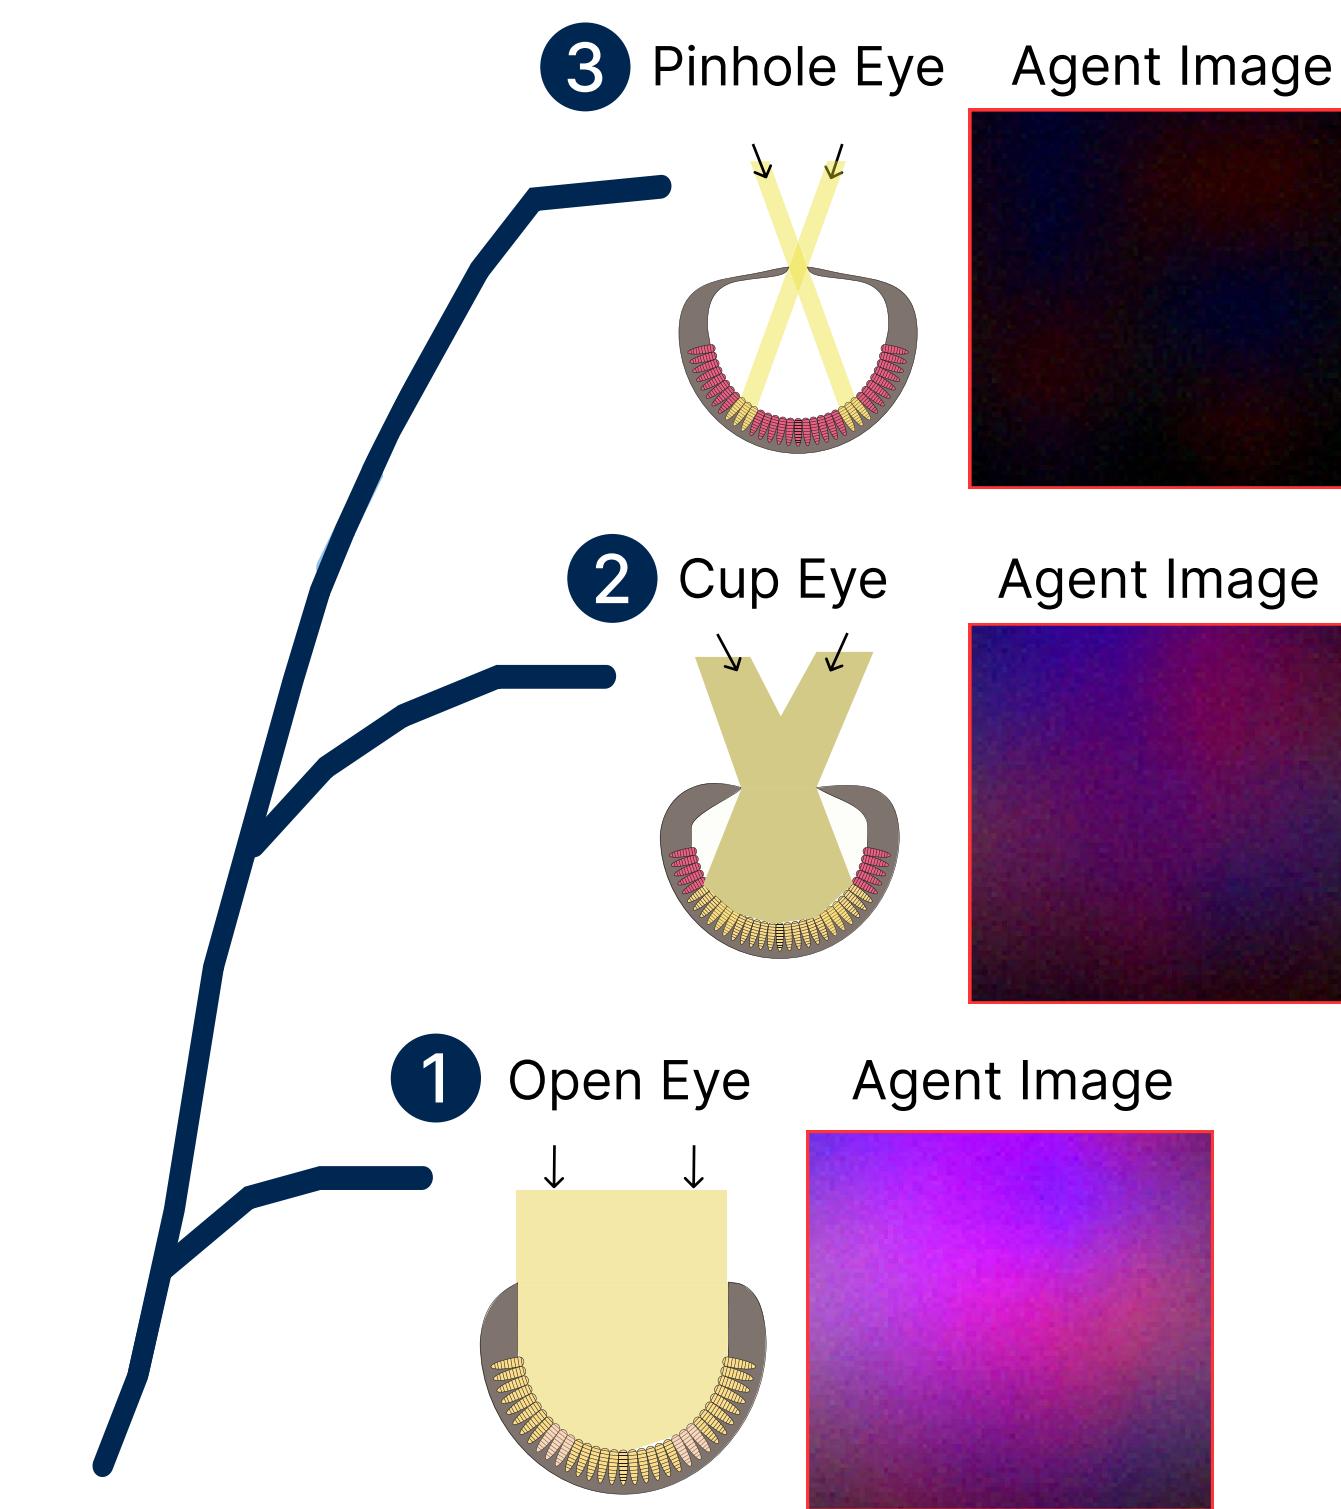
# Navigation at Generation 0



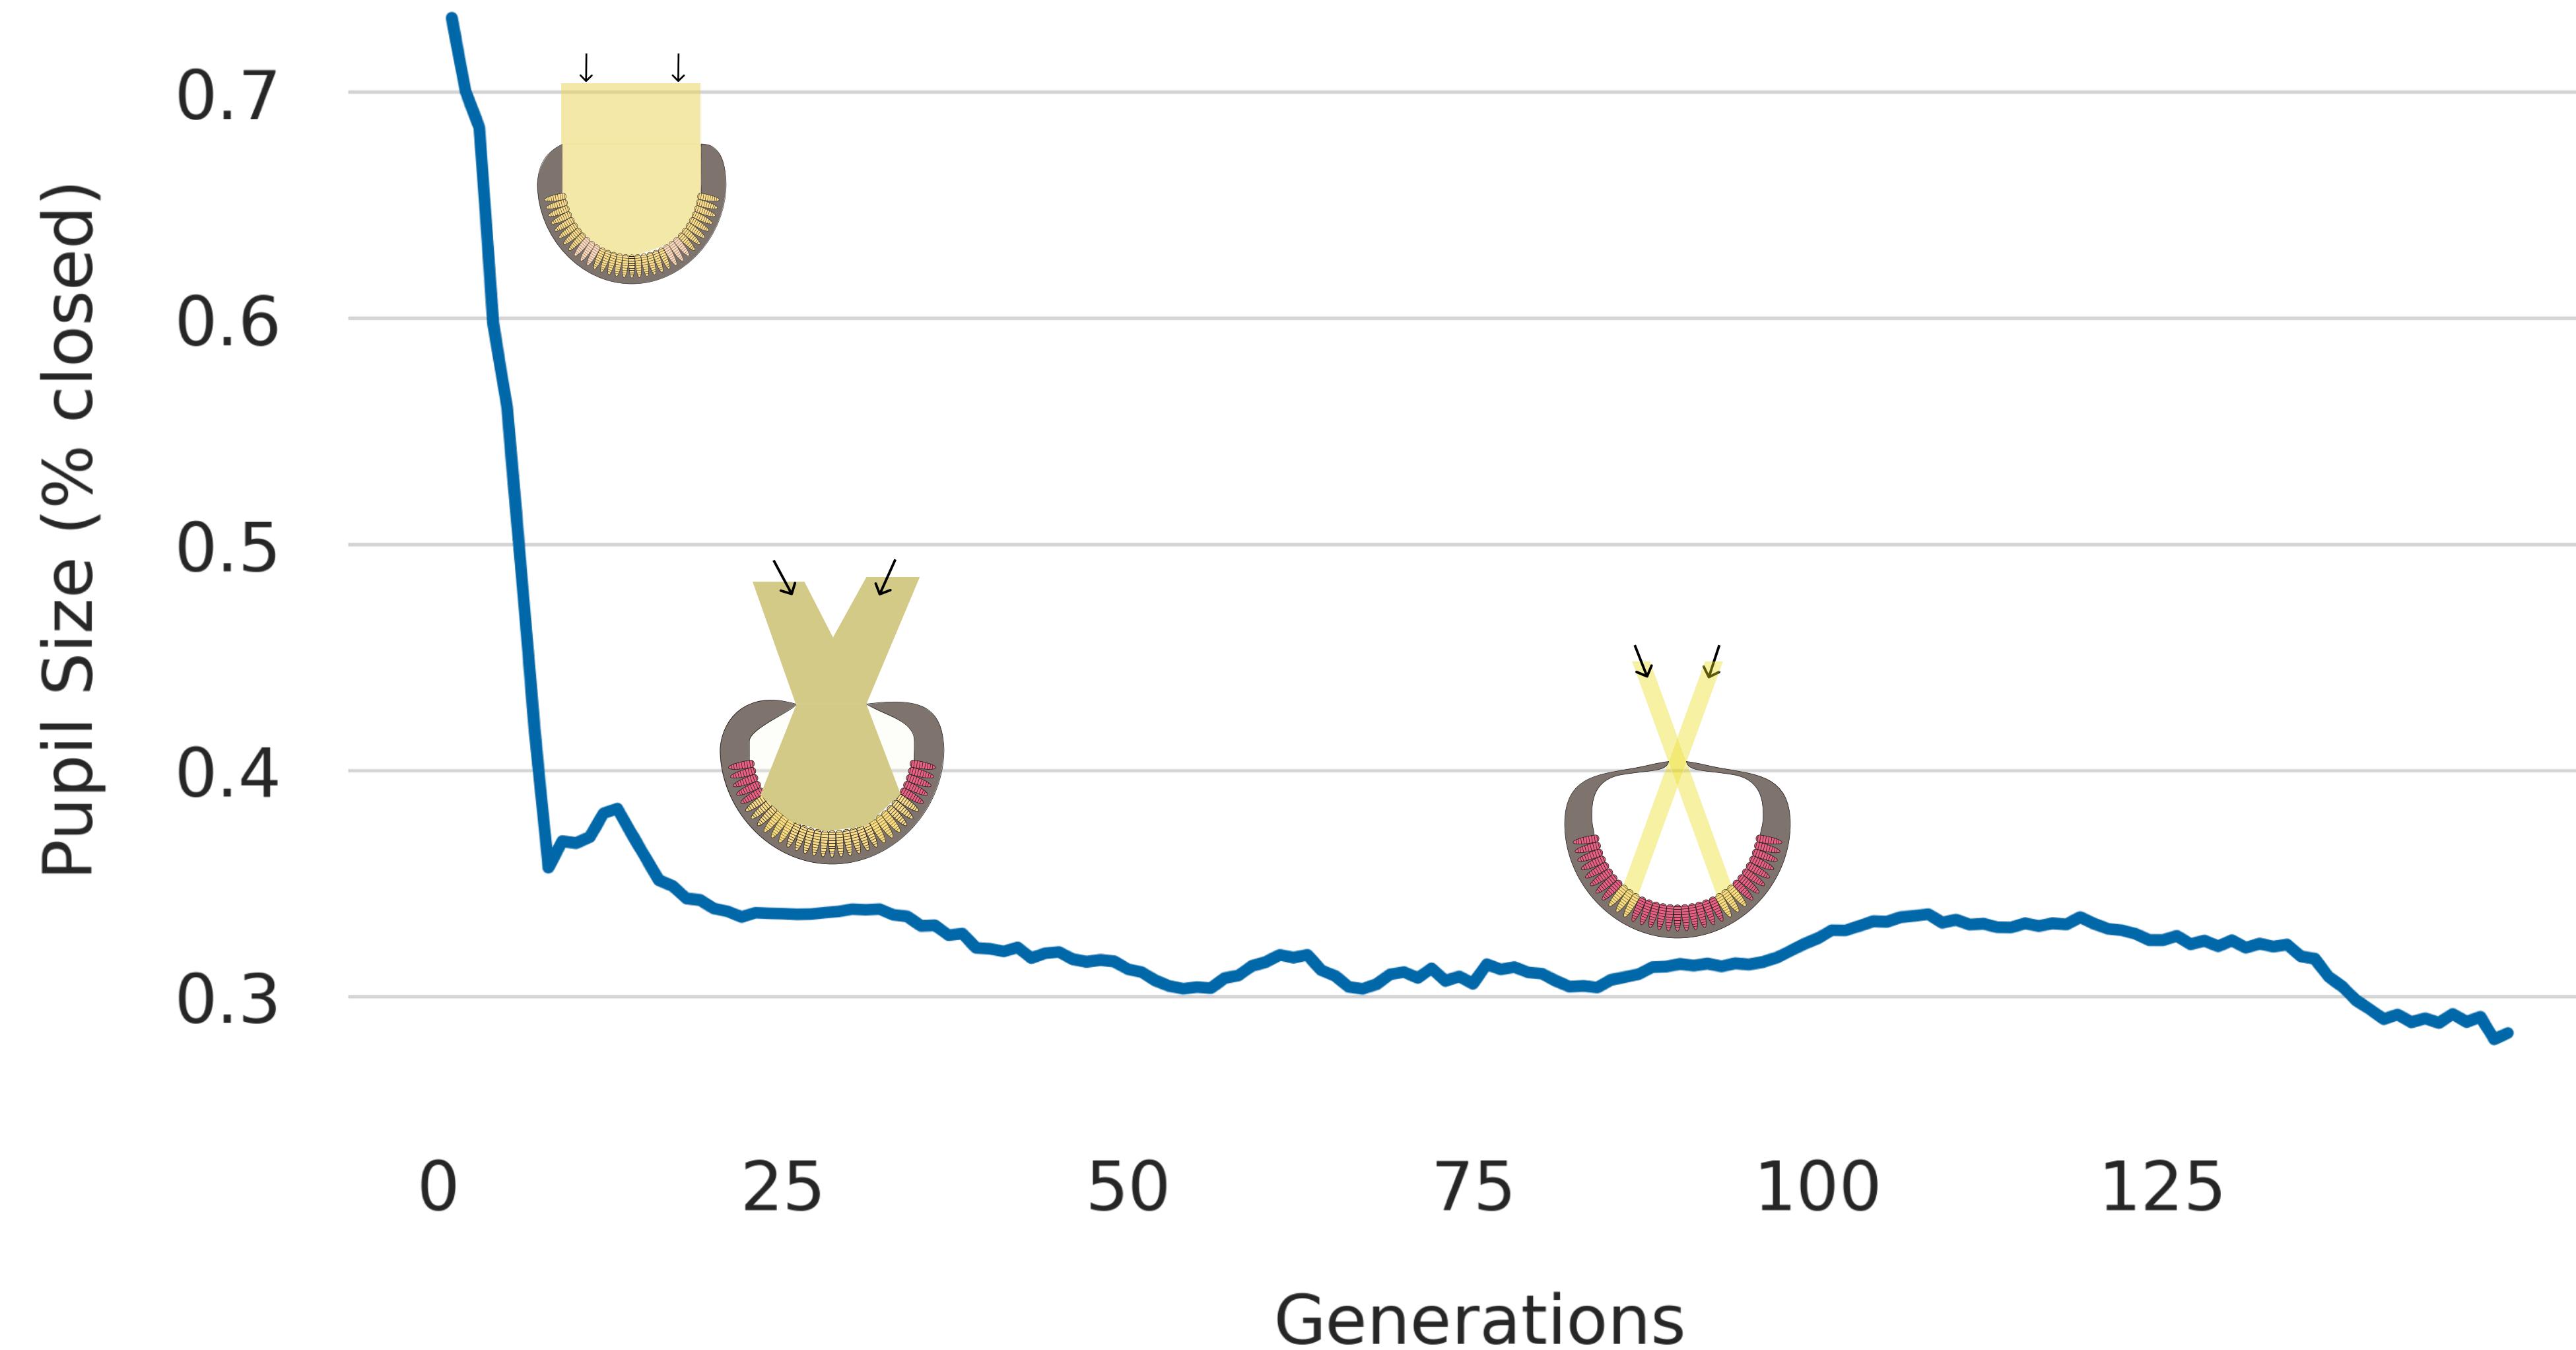
# Navigation at Generation 50



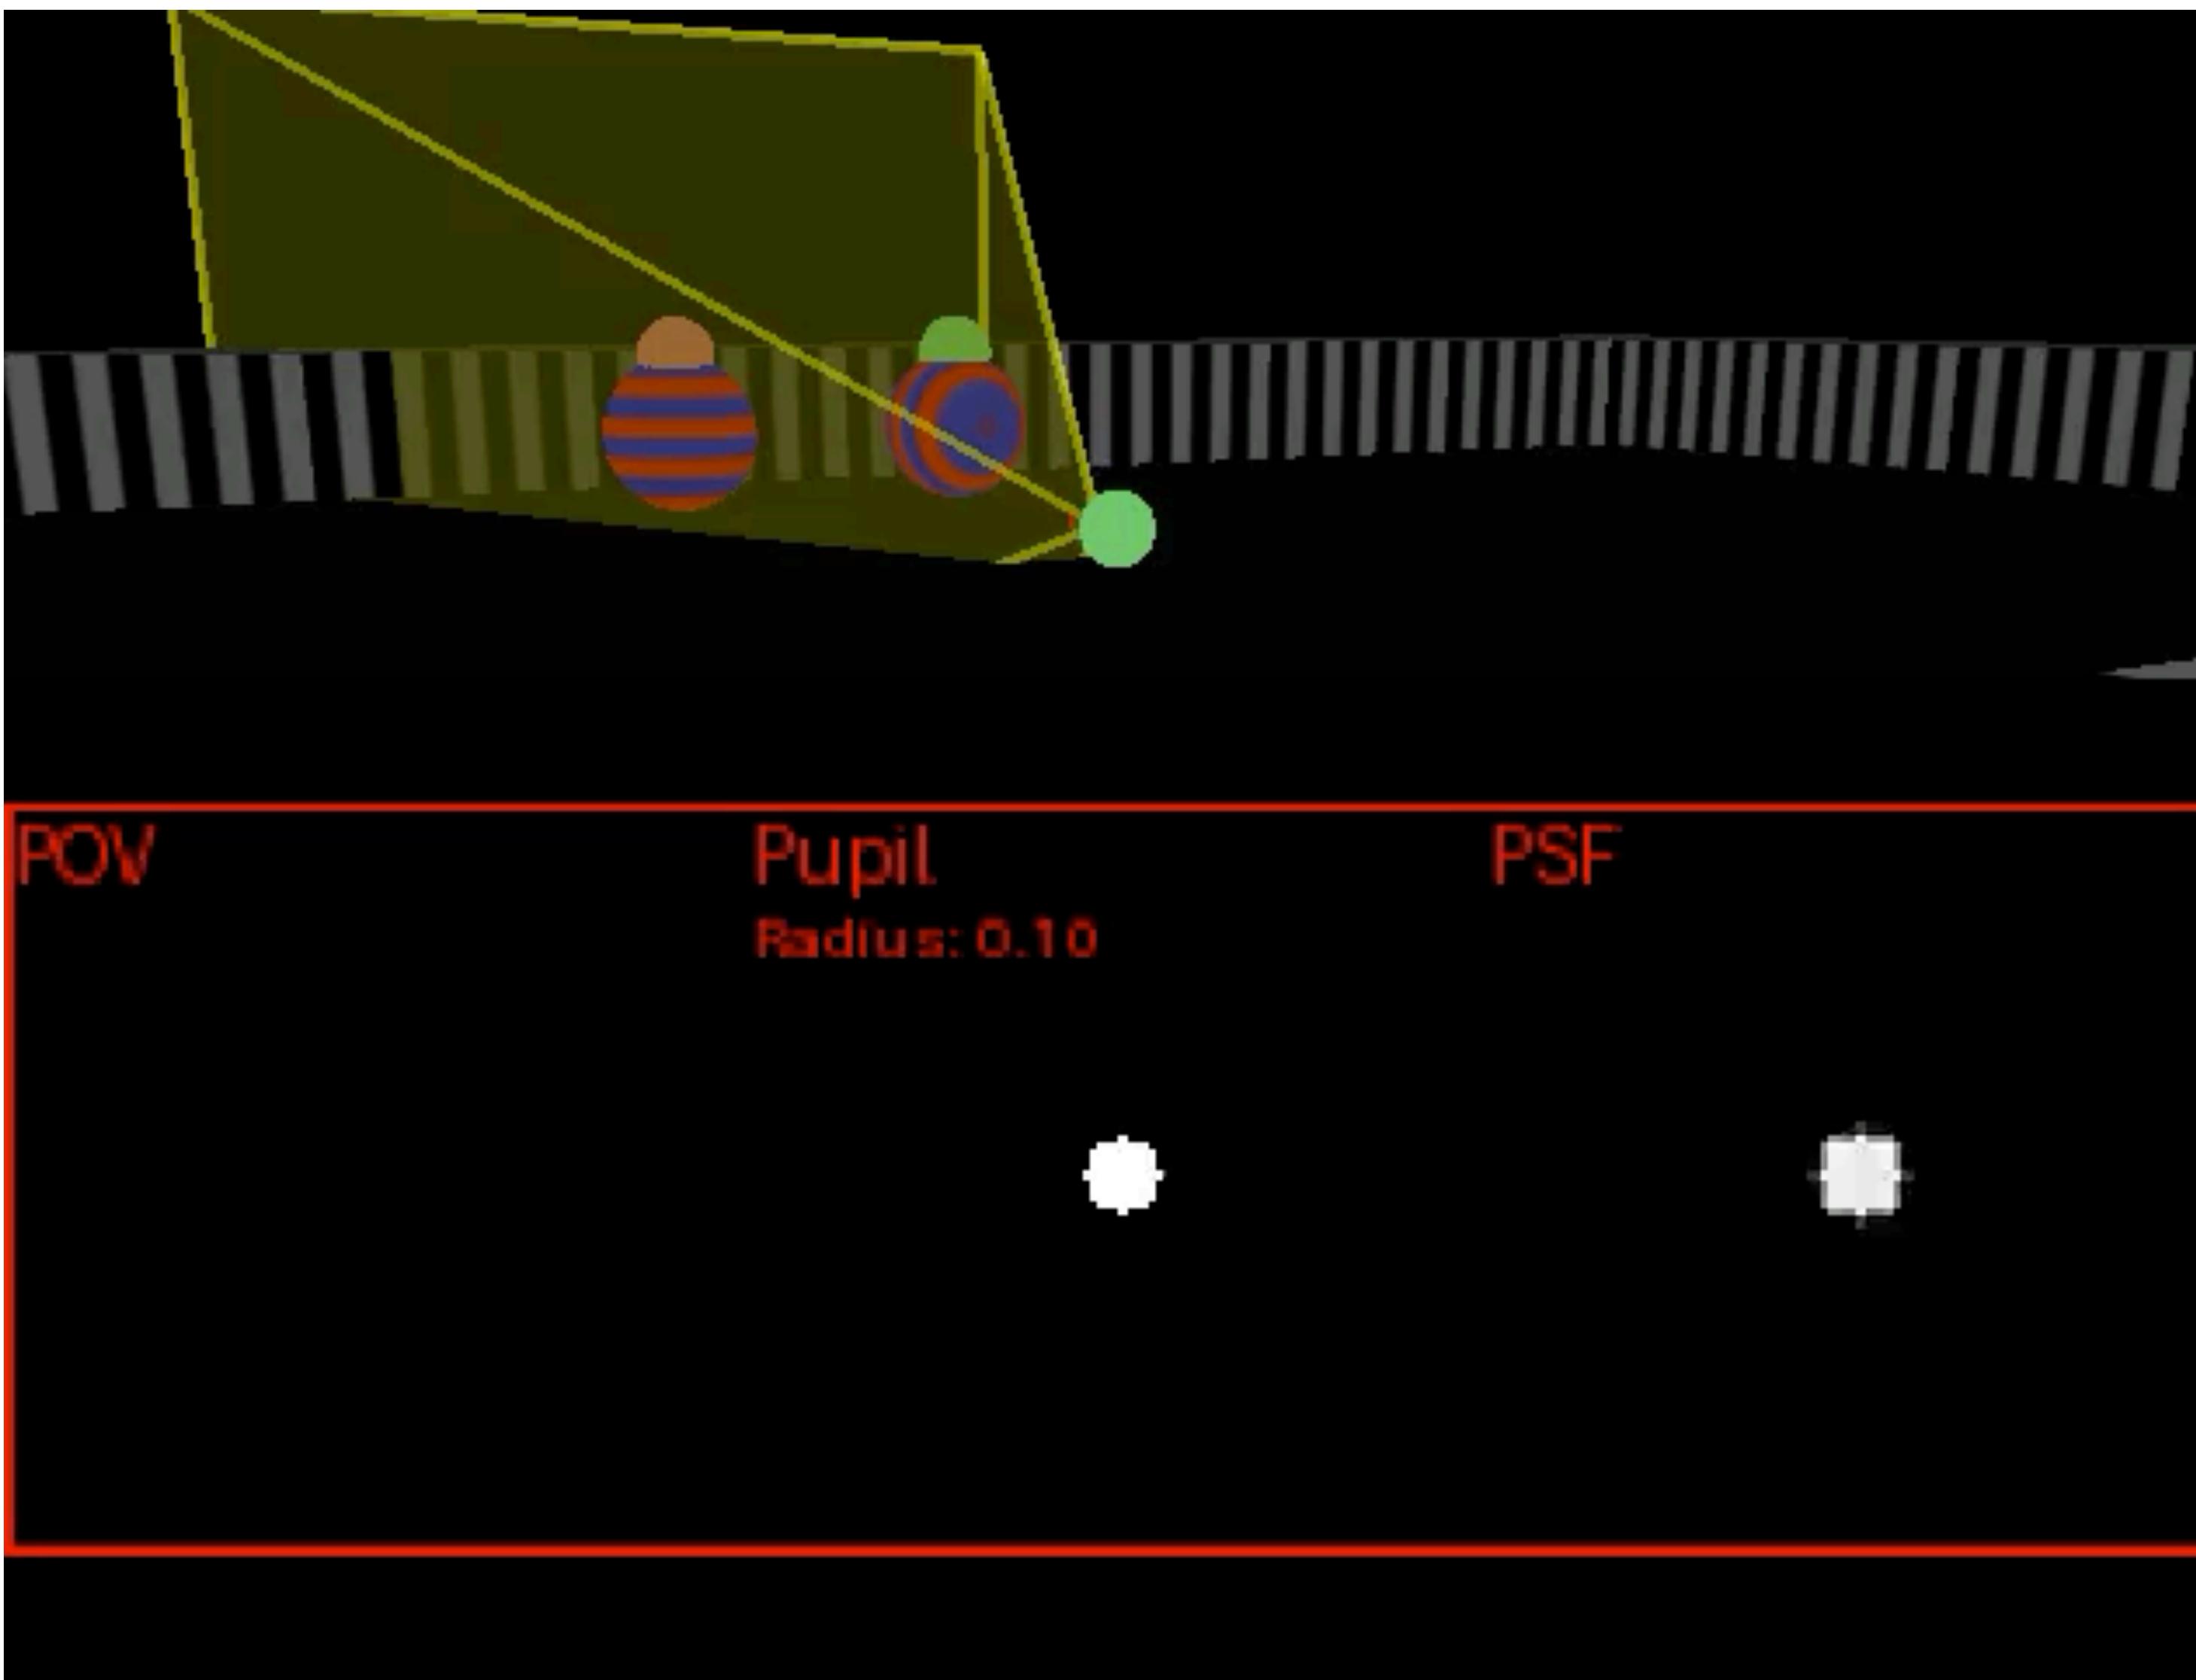
# Emergence of Pinhole Eyes



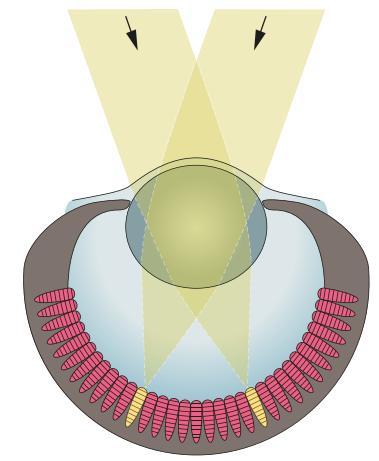
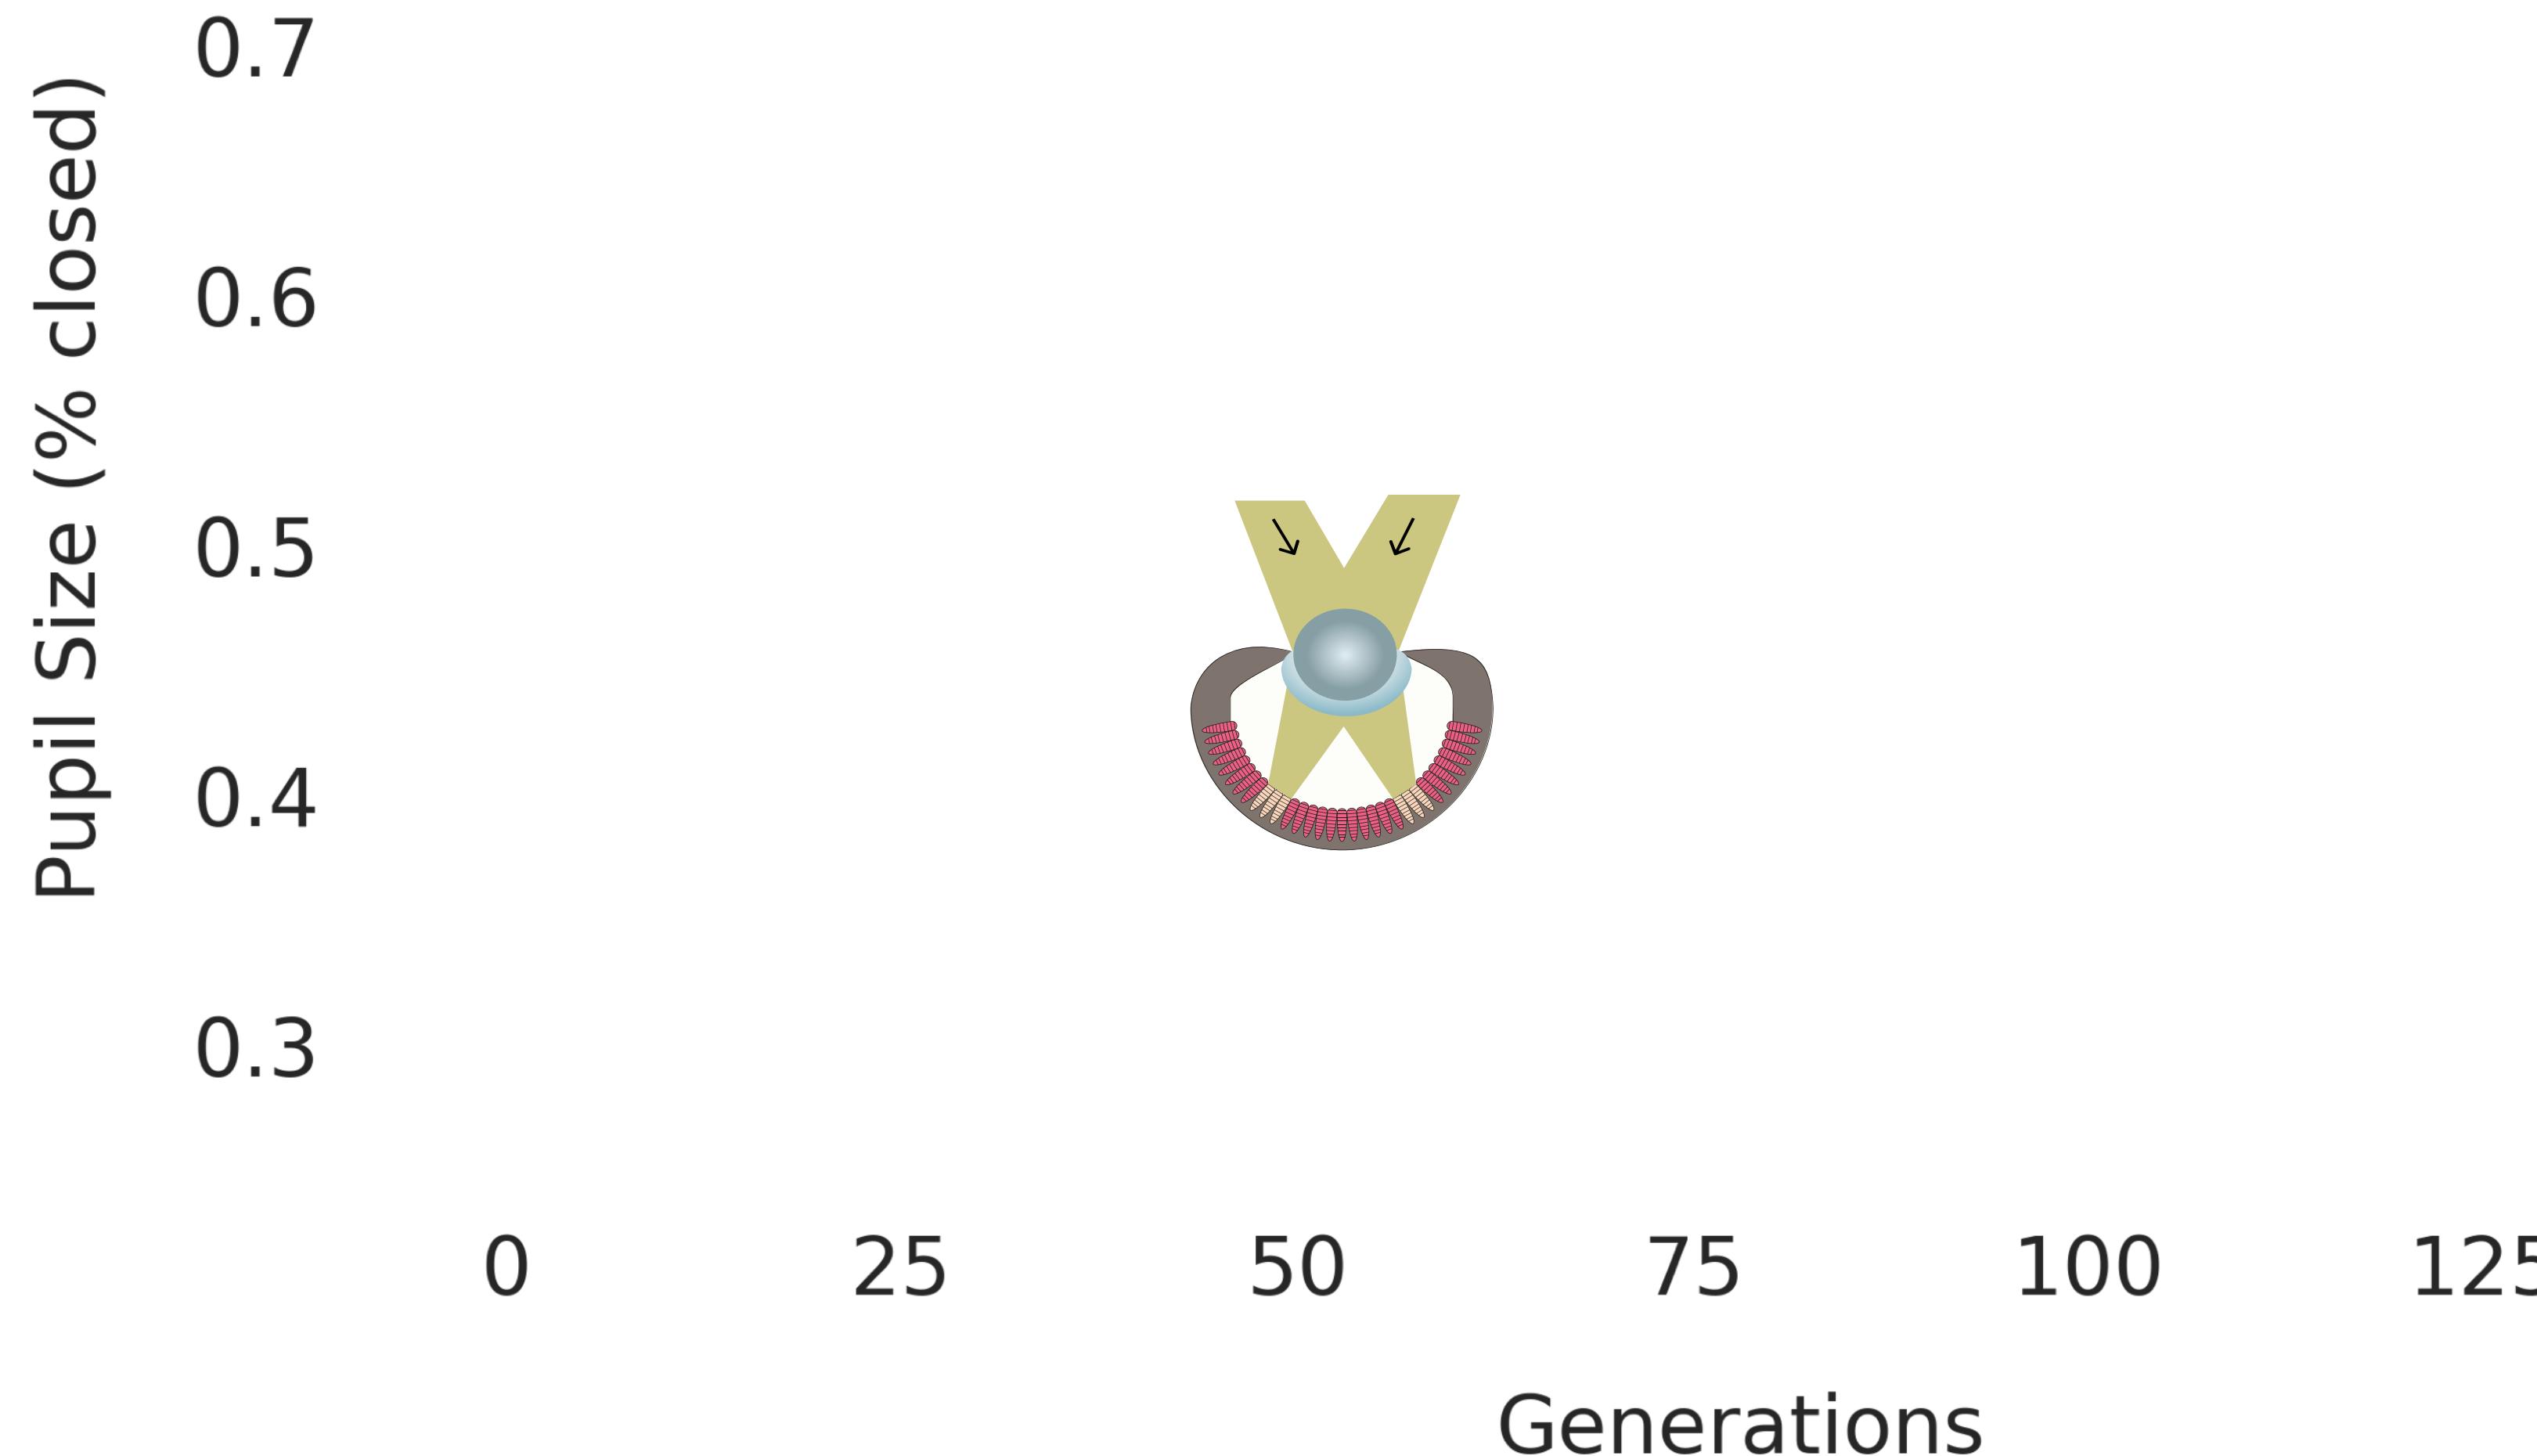
# Emergence of Pinhole Eyes



# Pinhole Eyes: Acuity at the cost of brightness

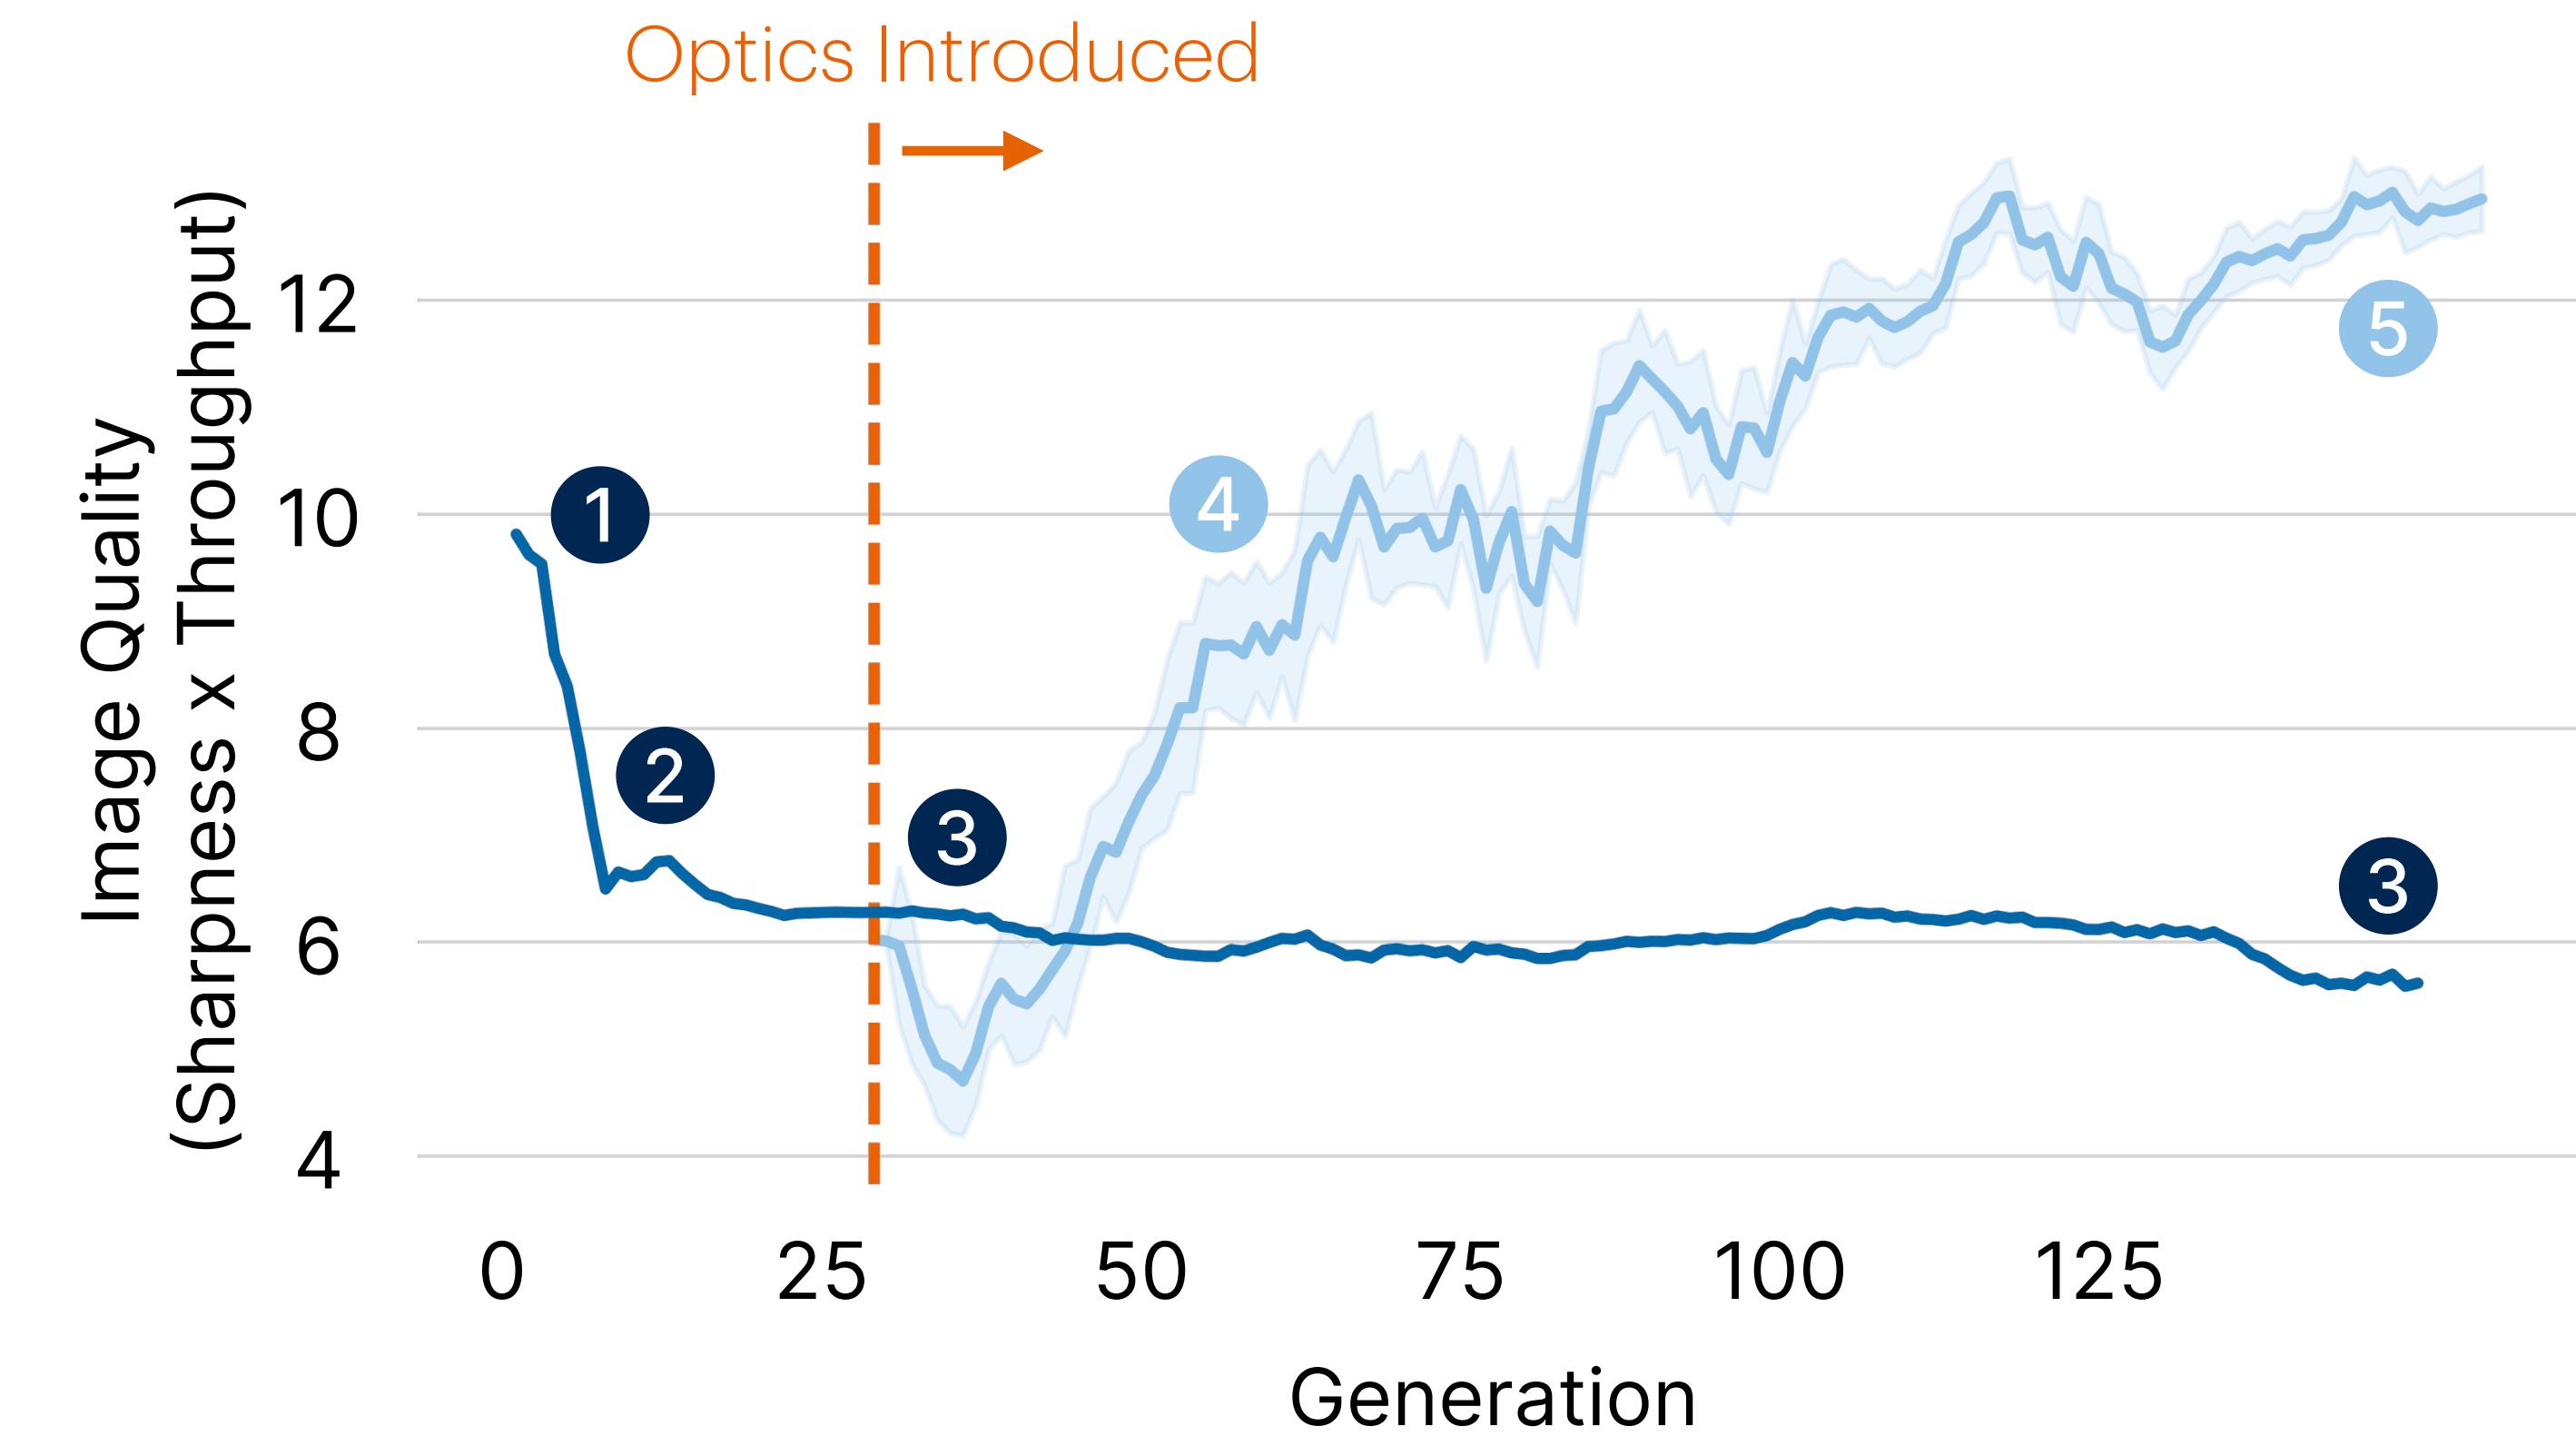
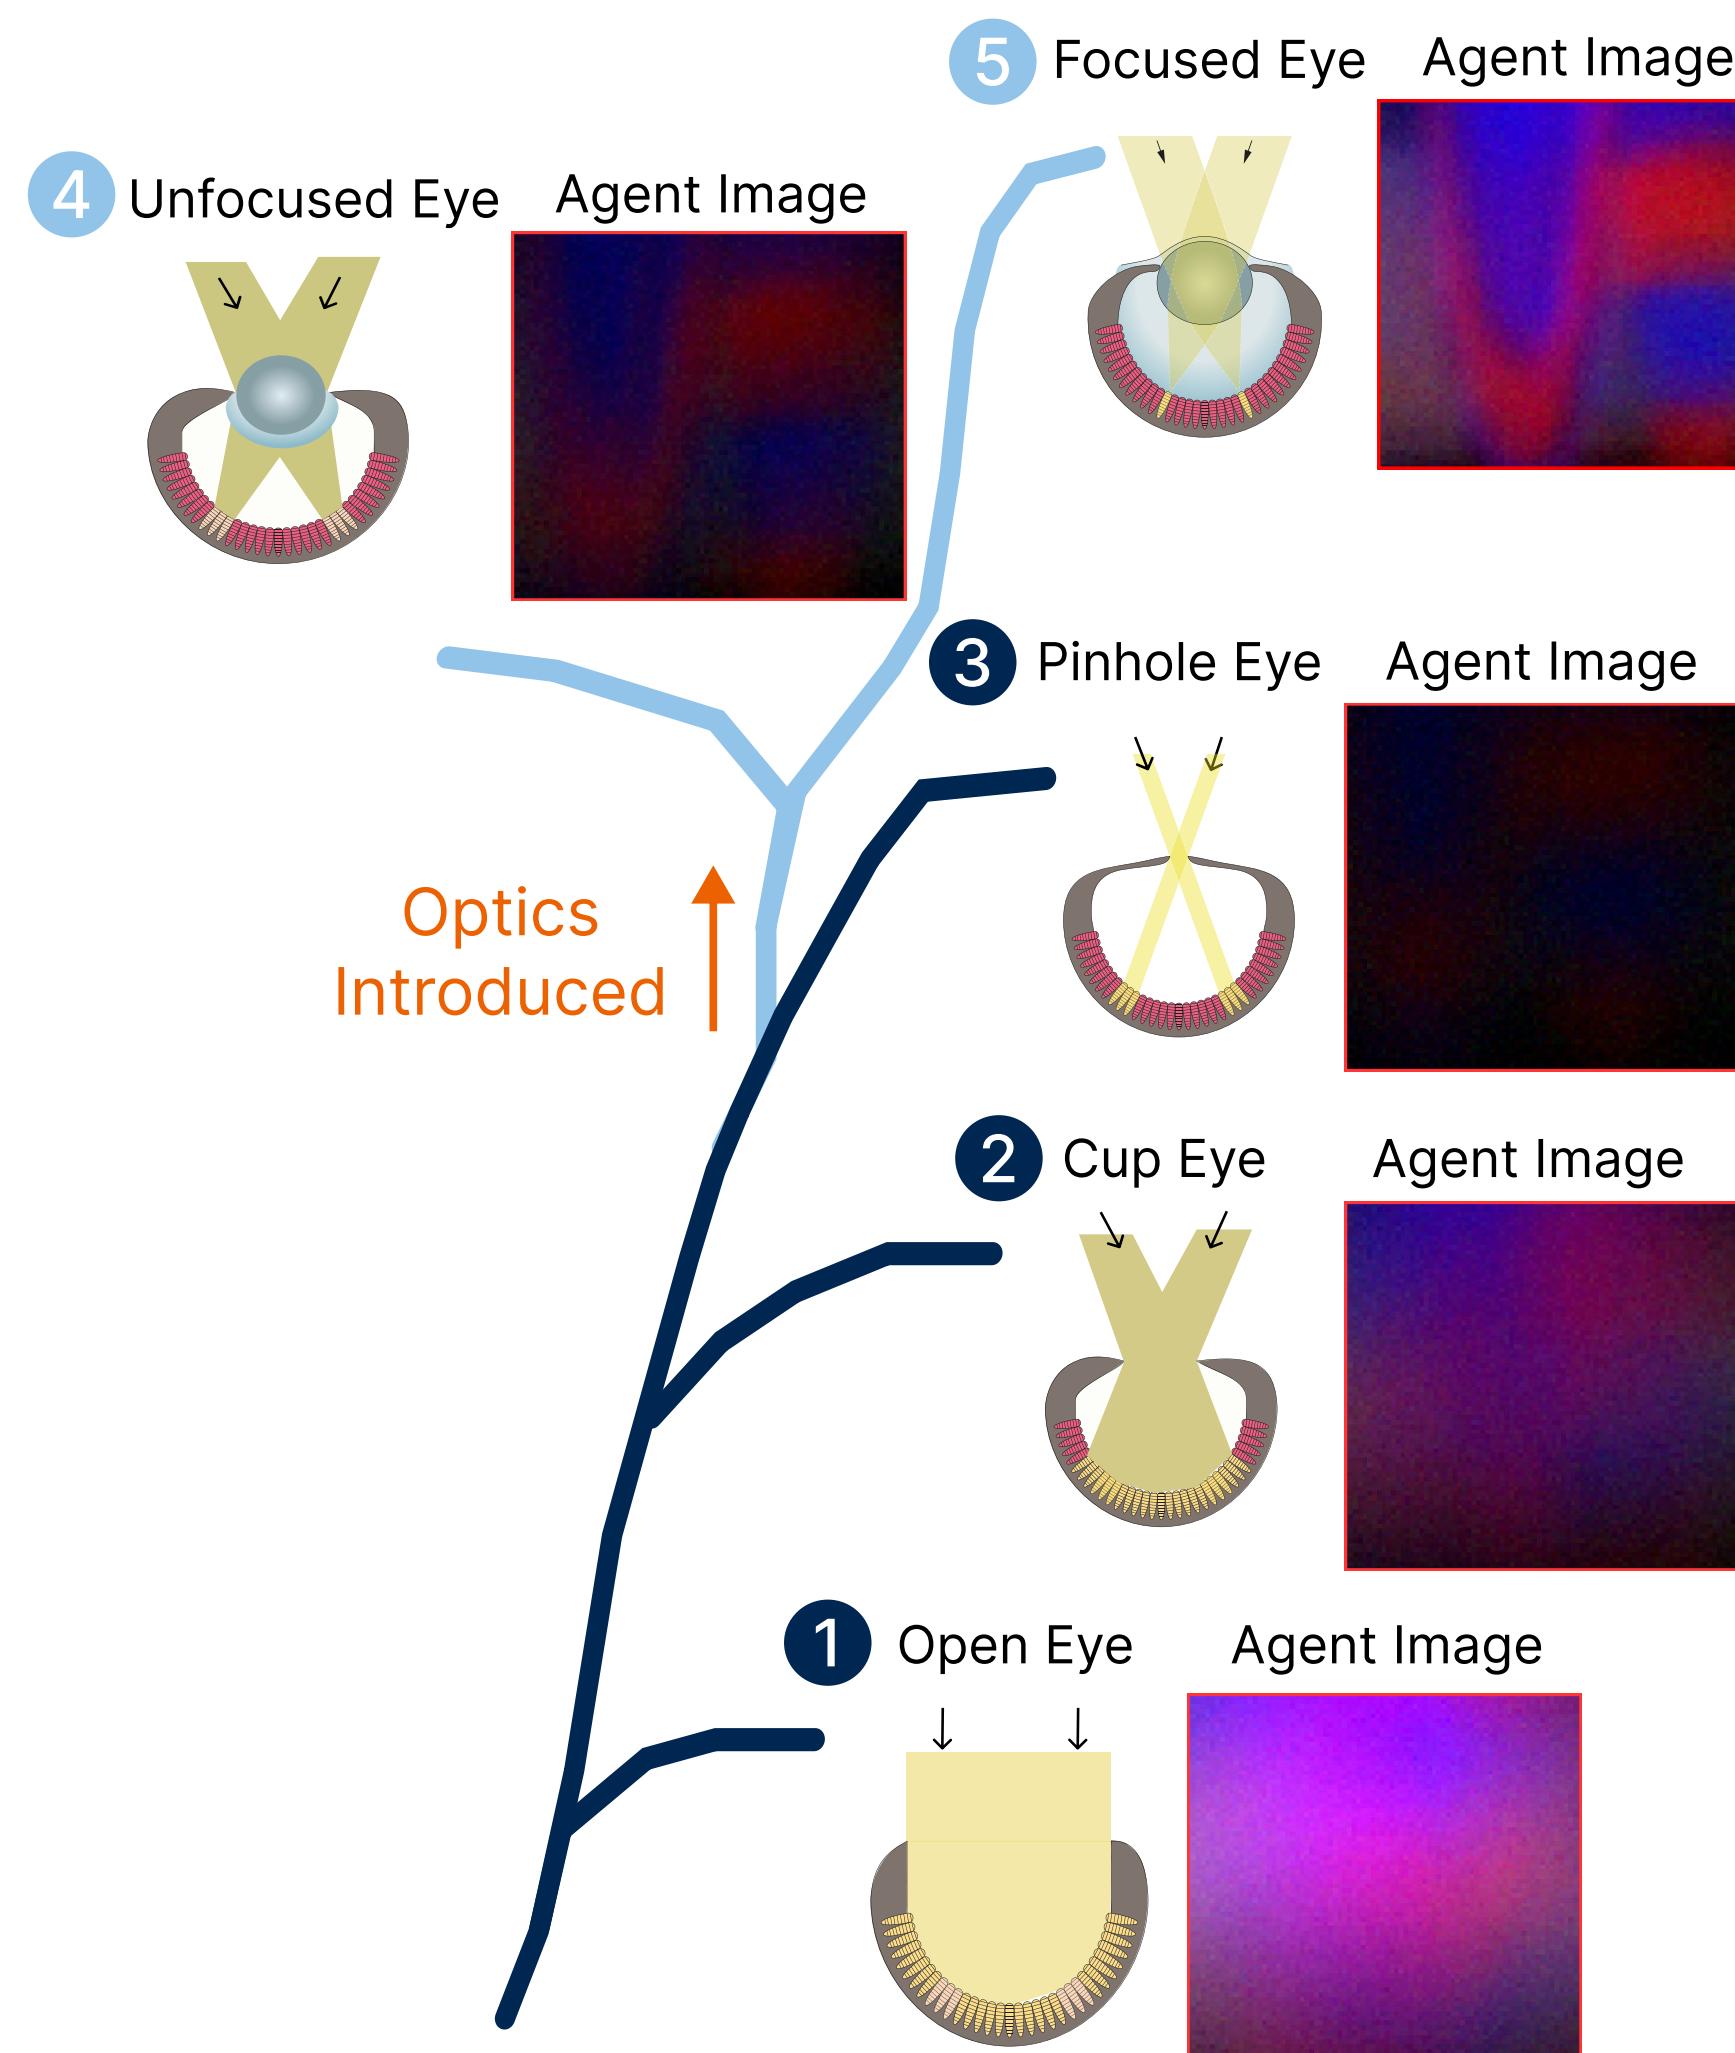


# What if eyes could **bend light**?

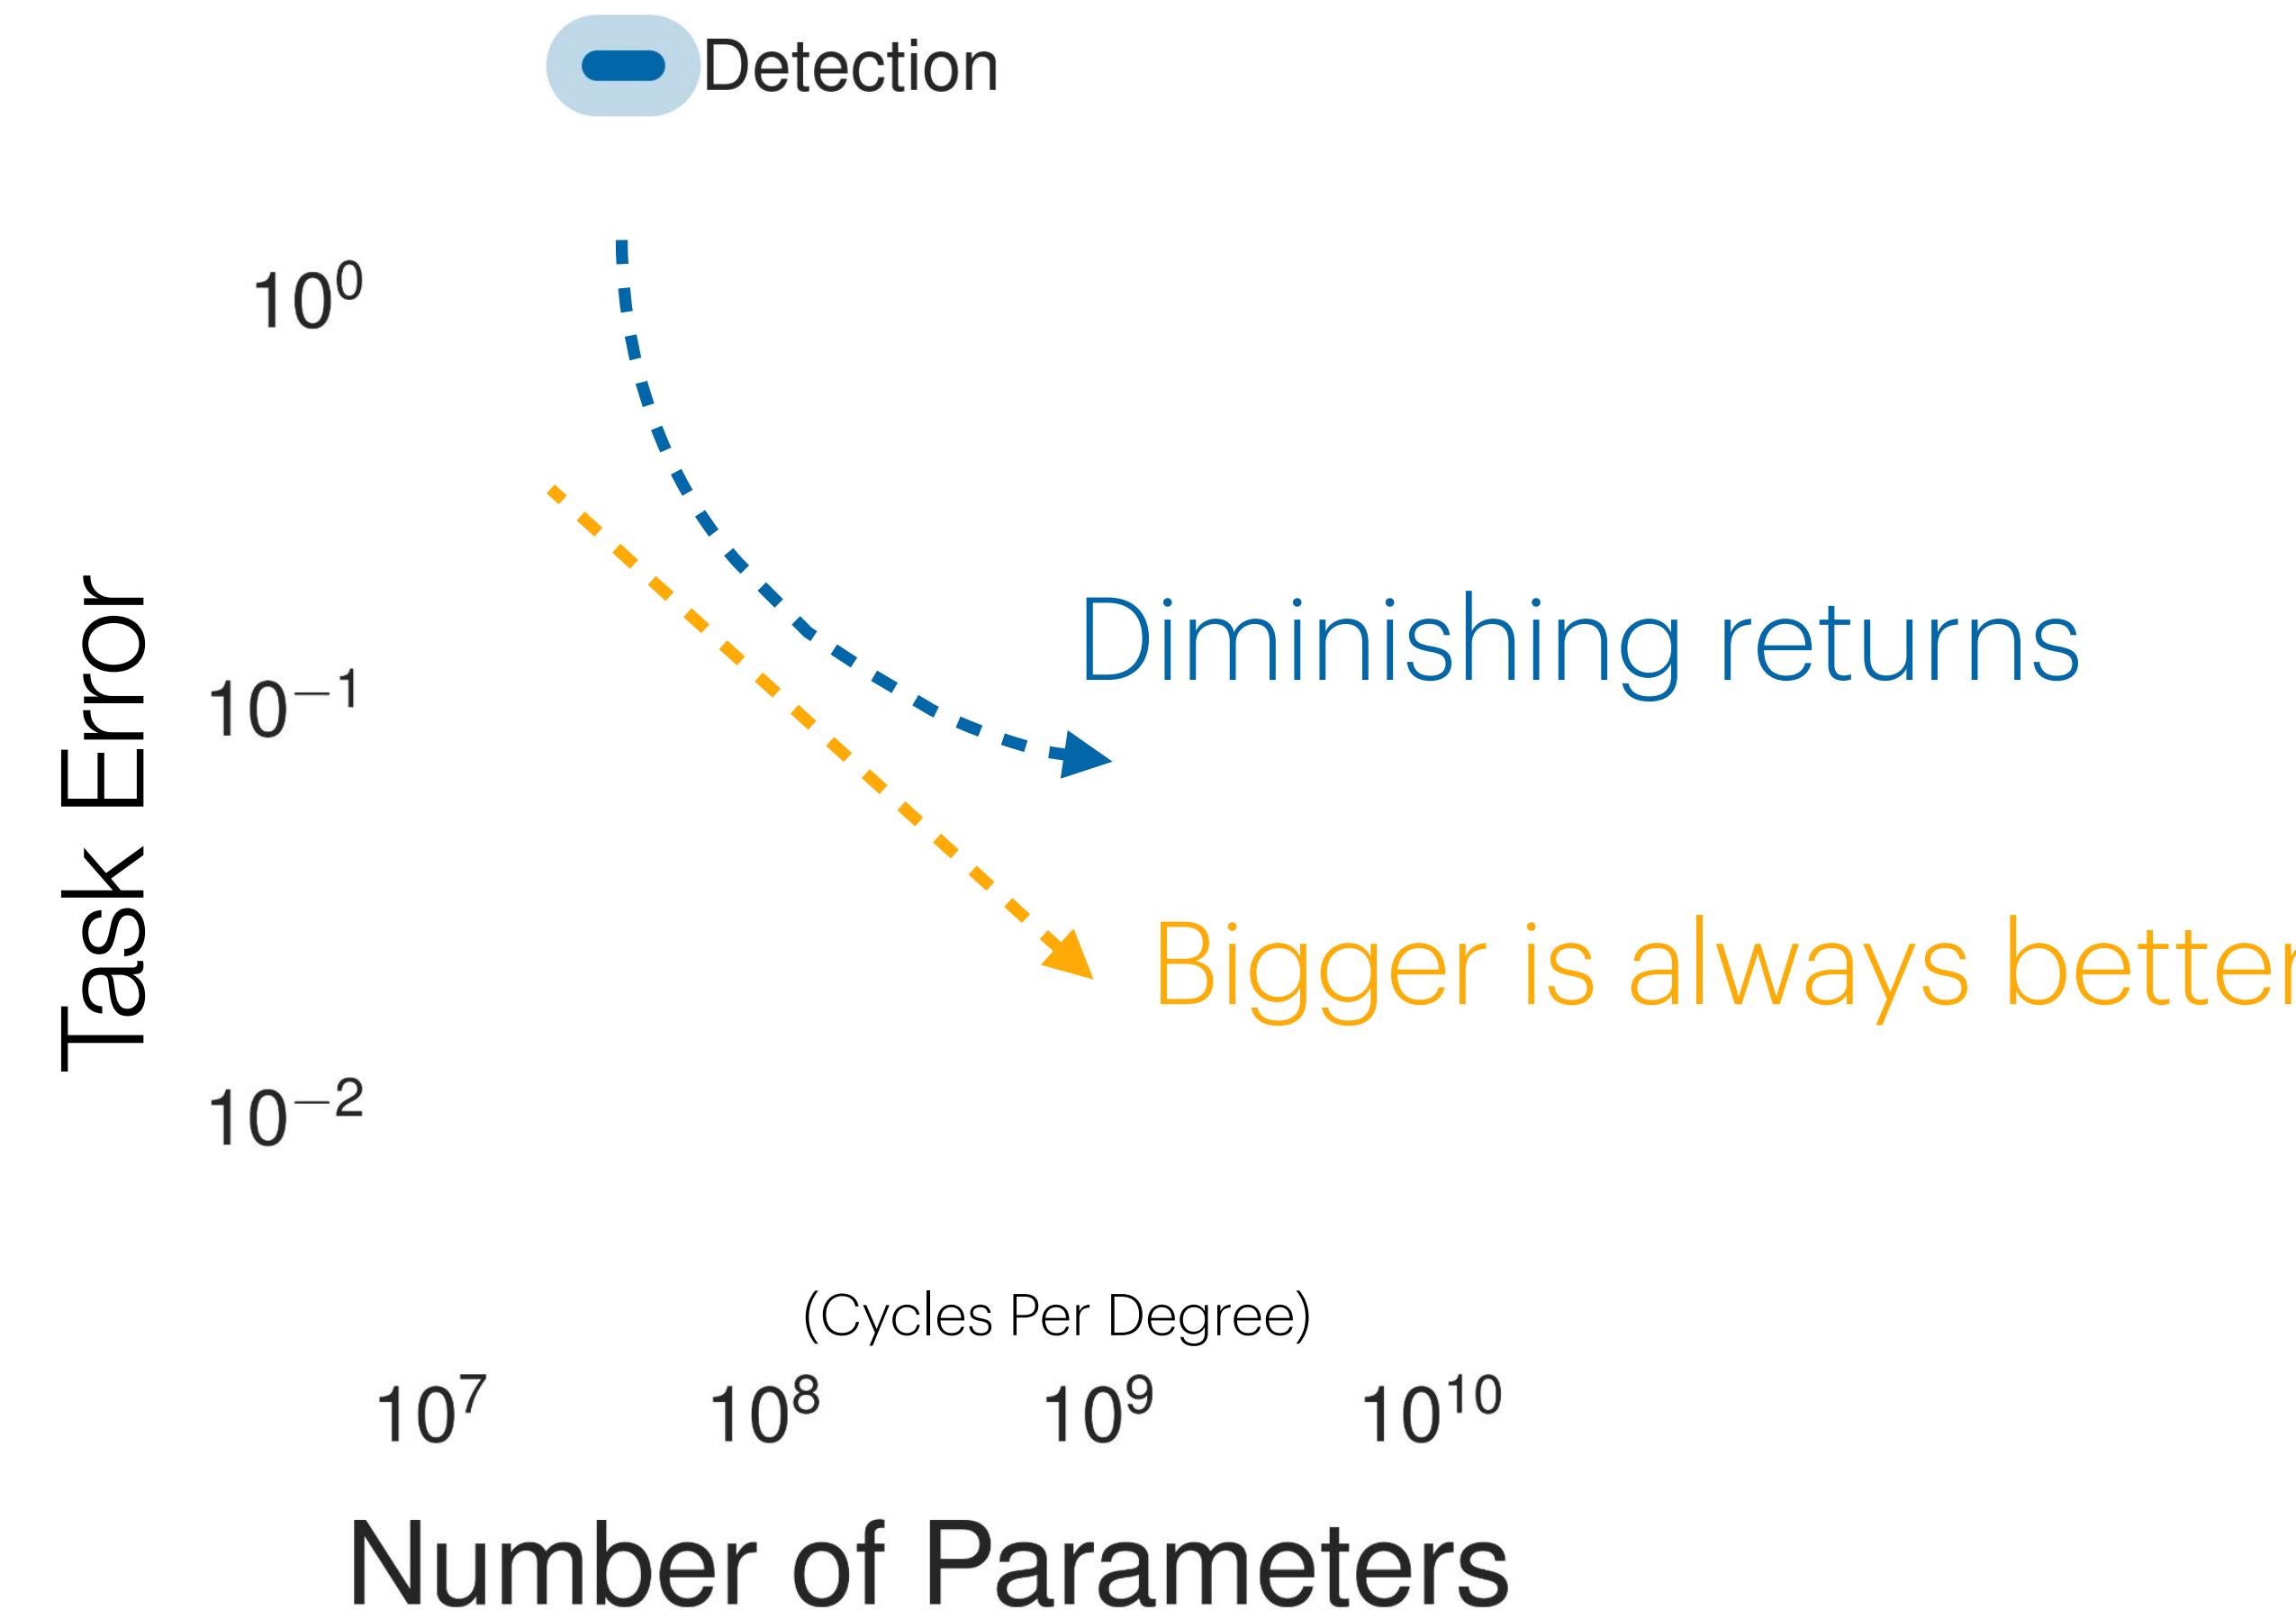


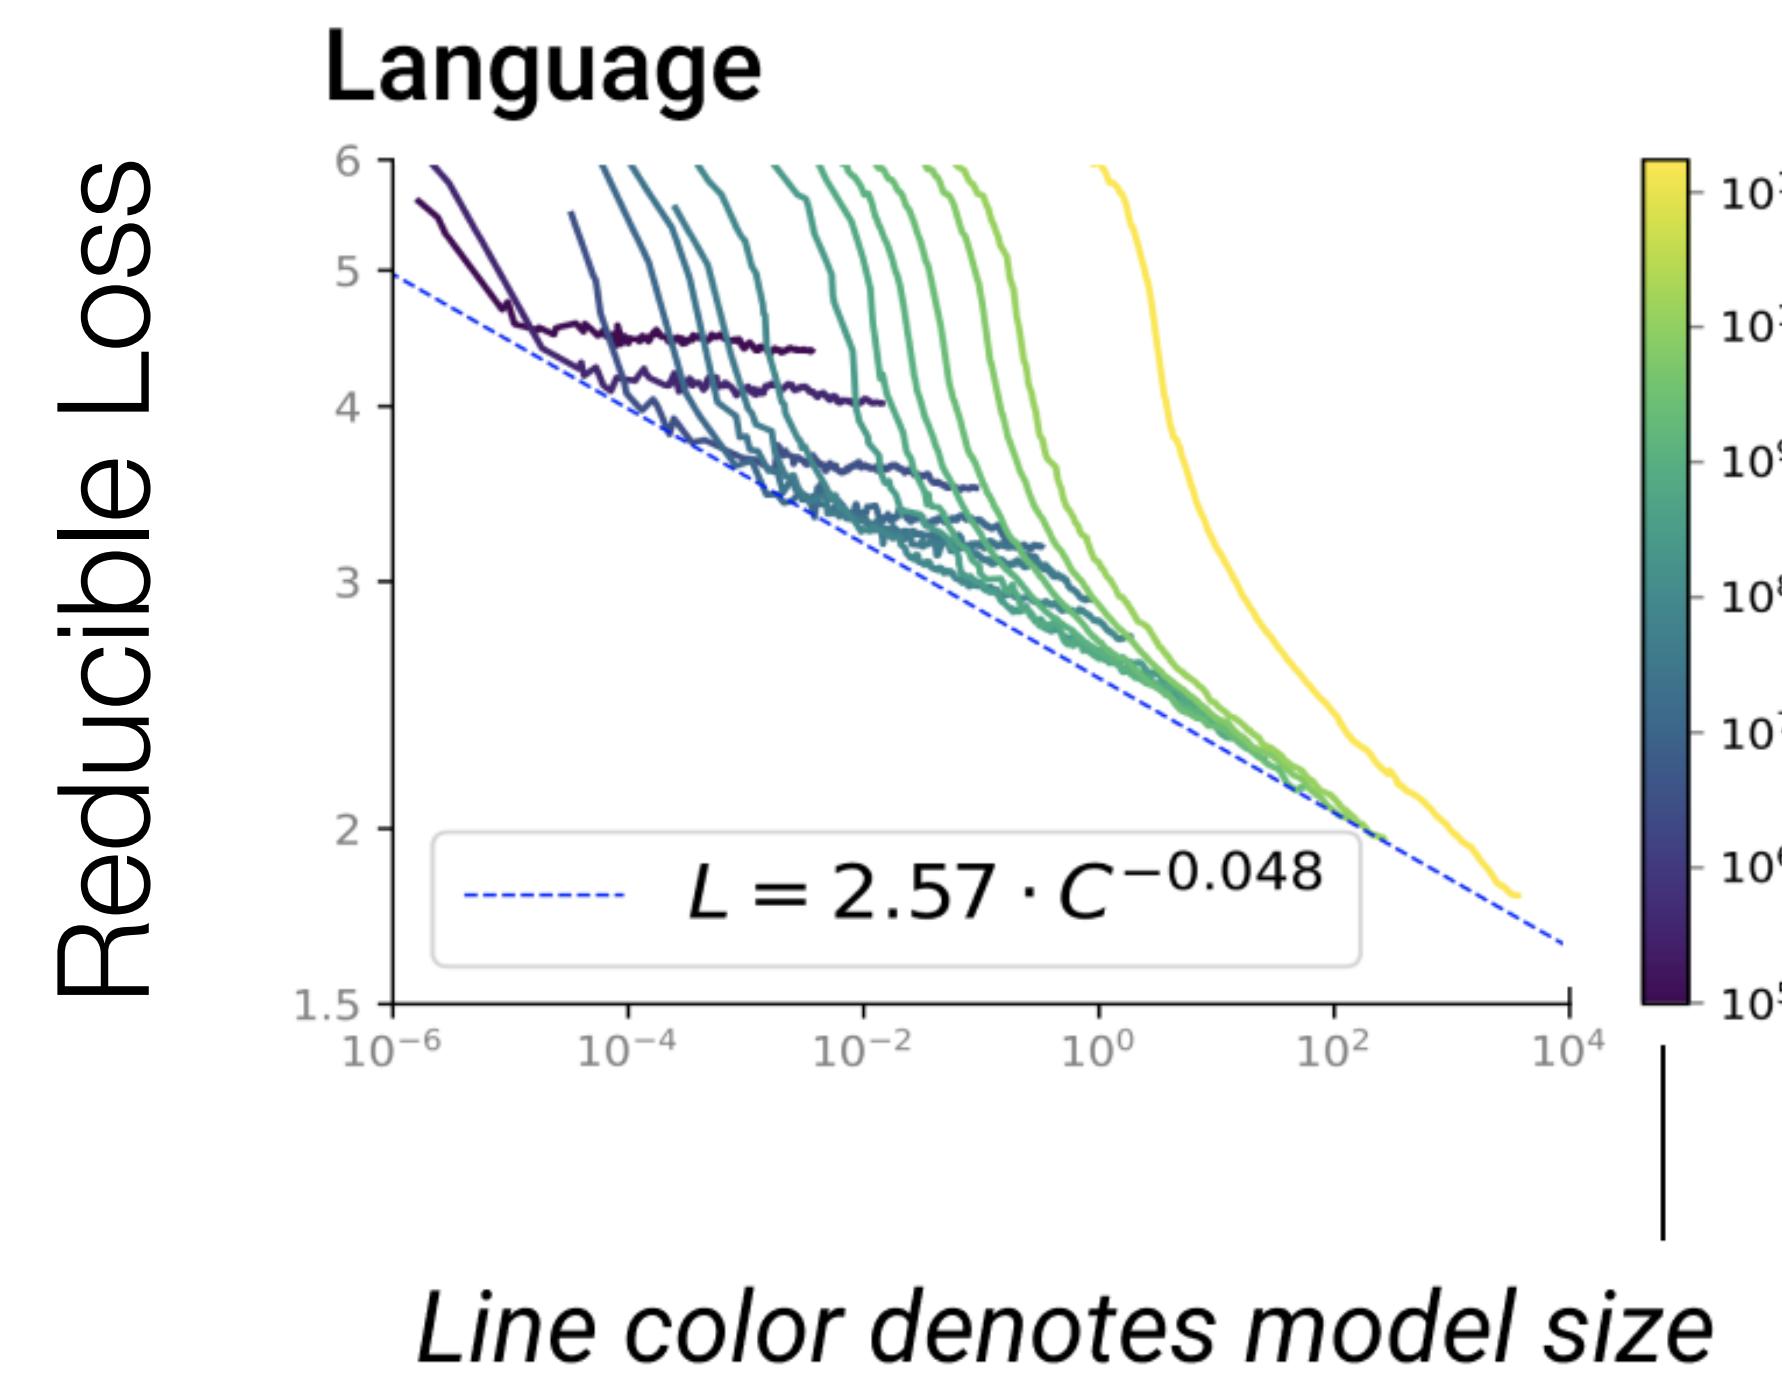
Lens Emerges

# What if eyes could **bend** light?



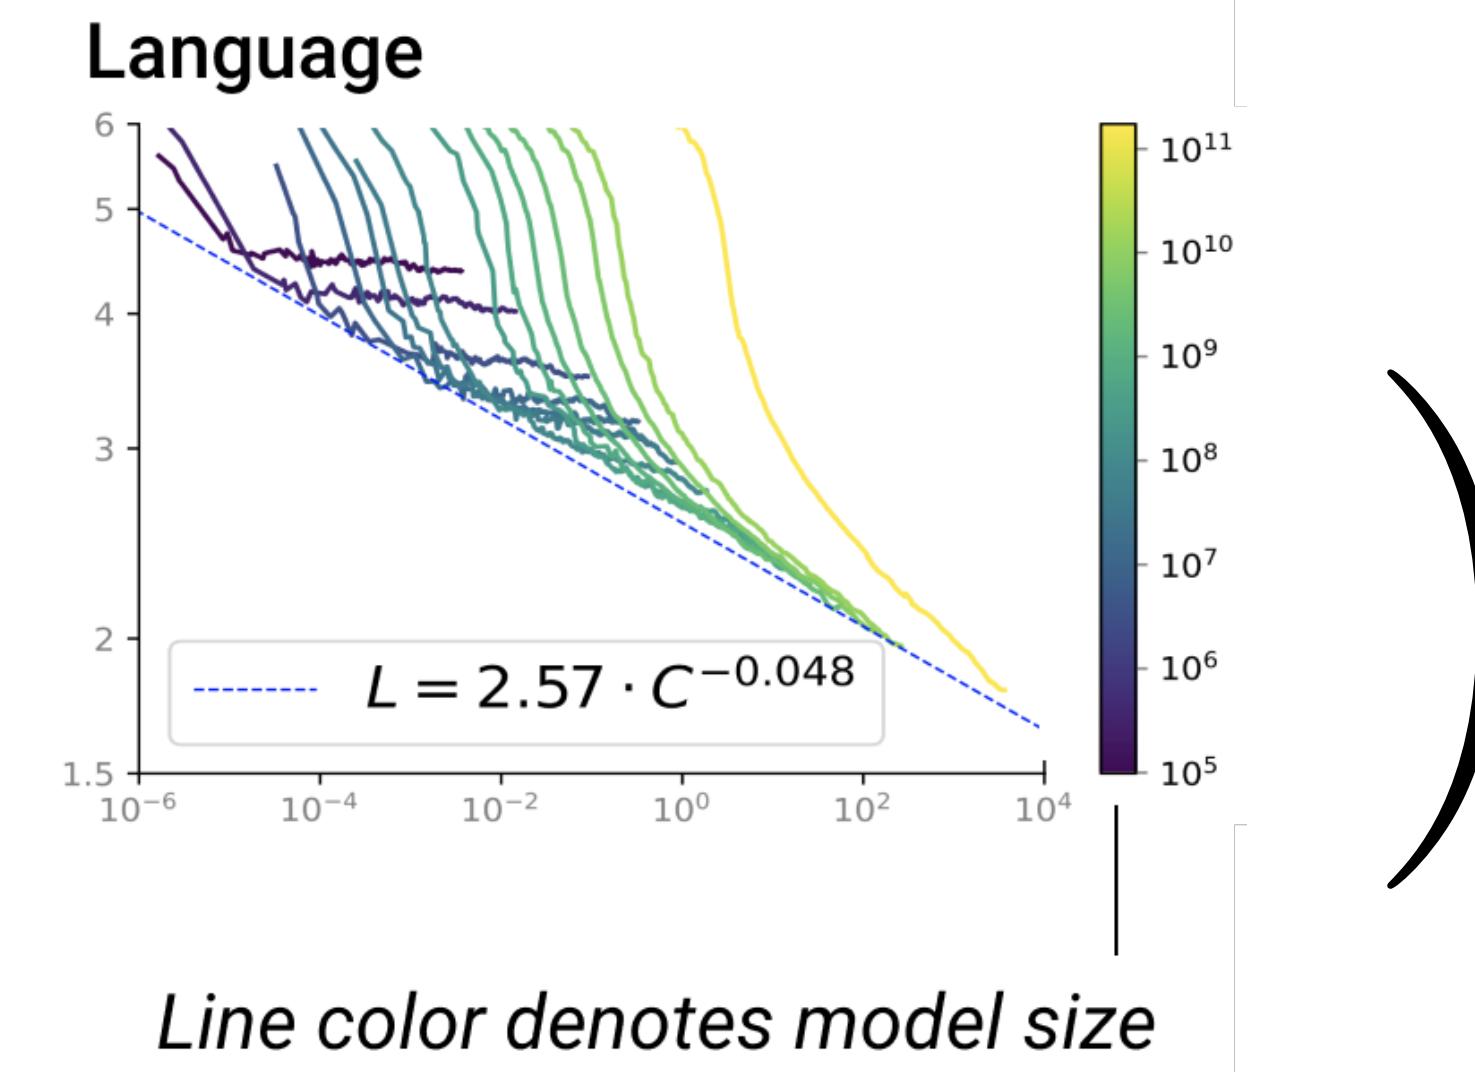
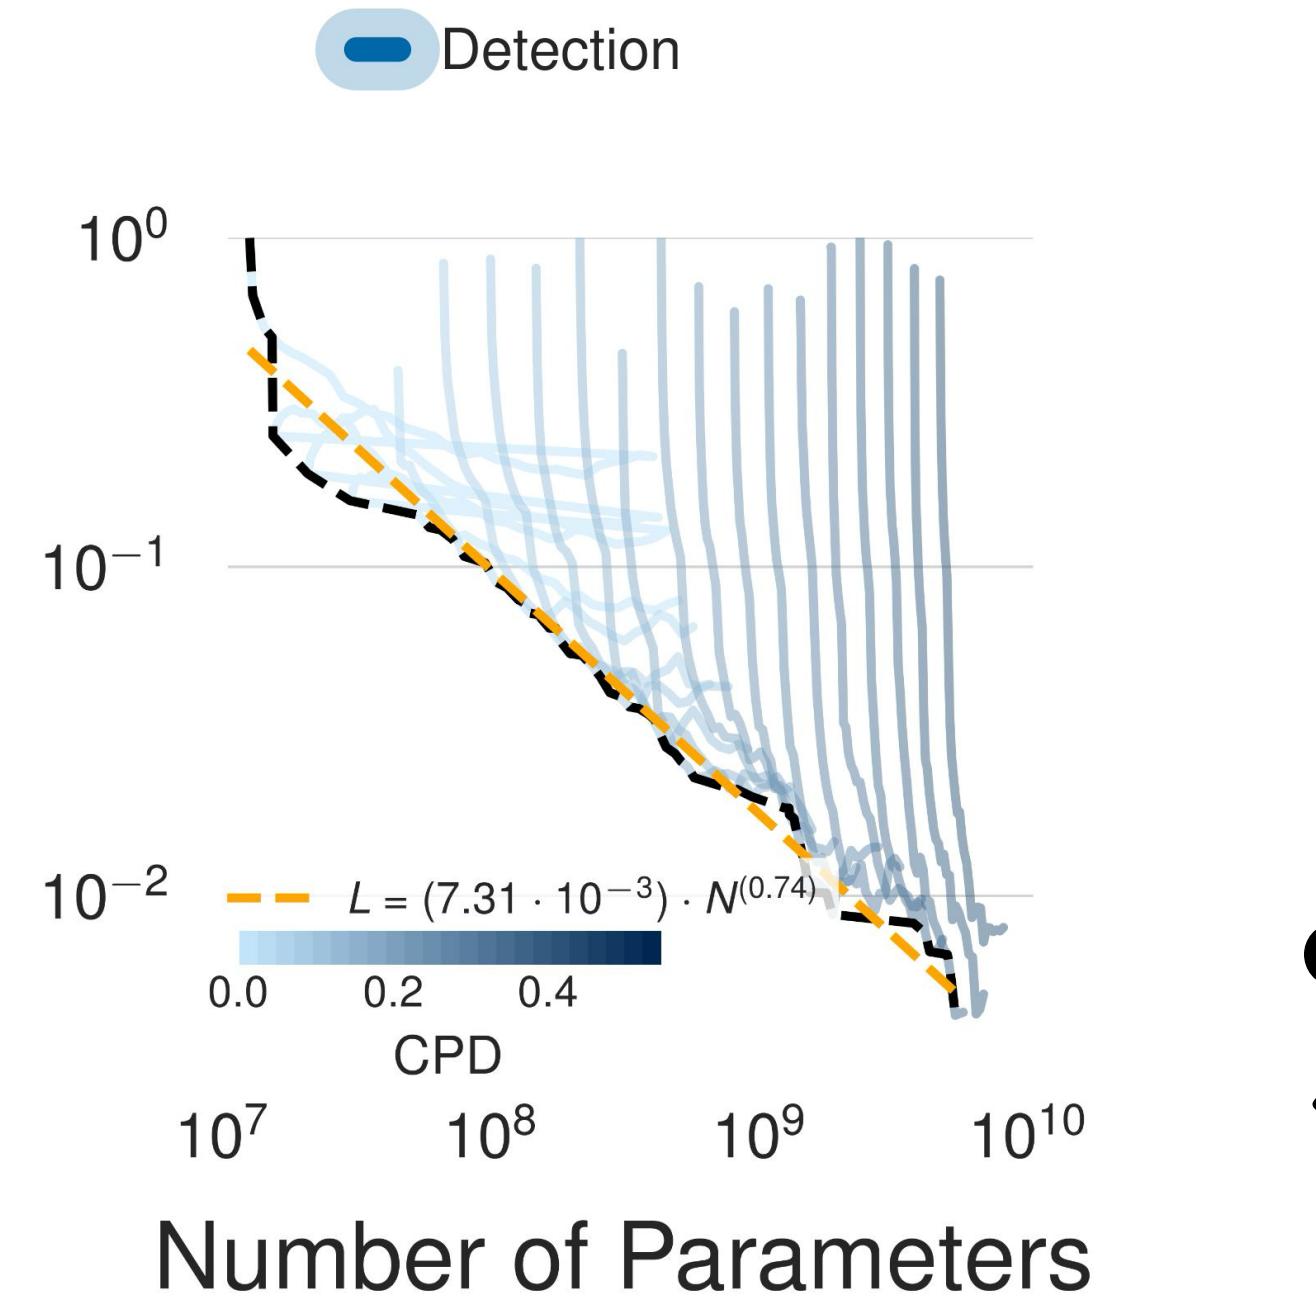
# What if the **brain** became larger?





# Scaling Laws

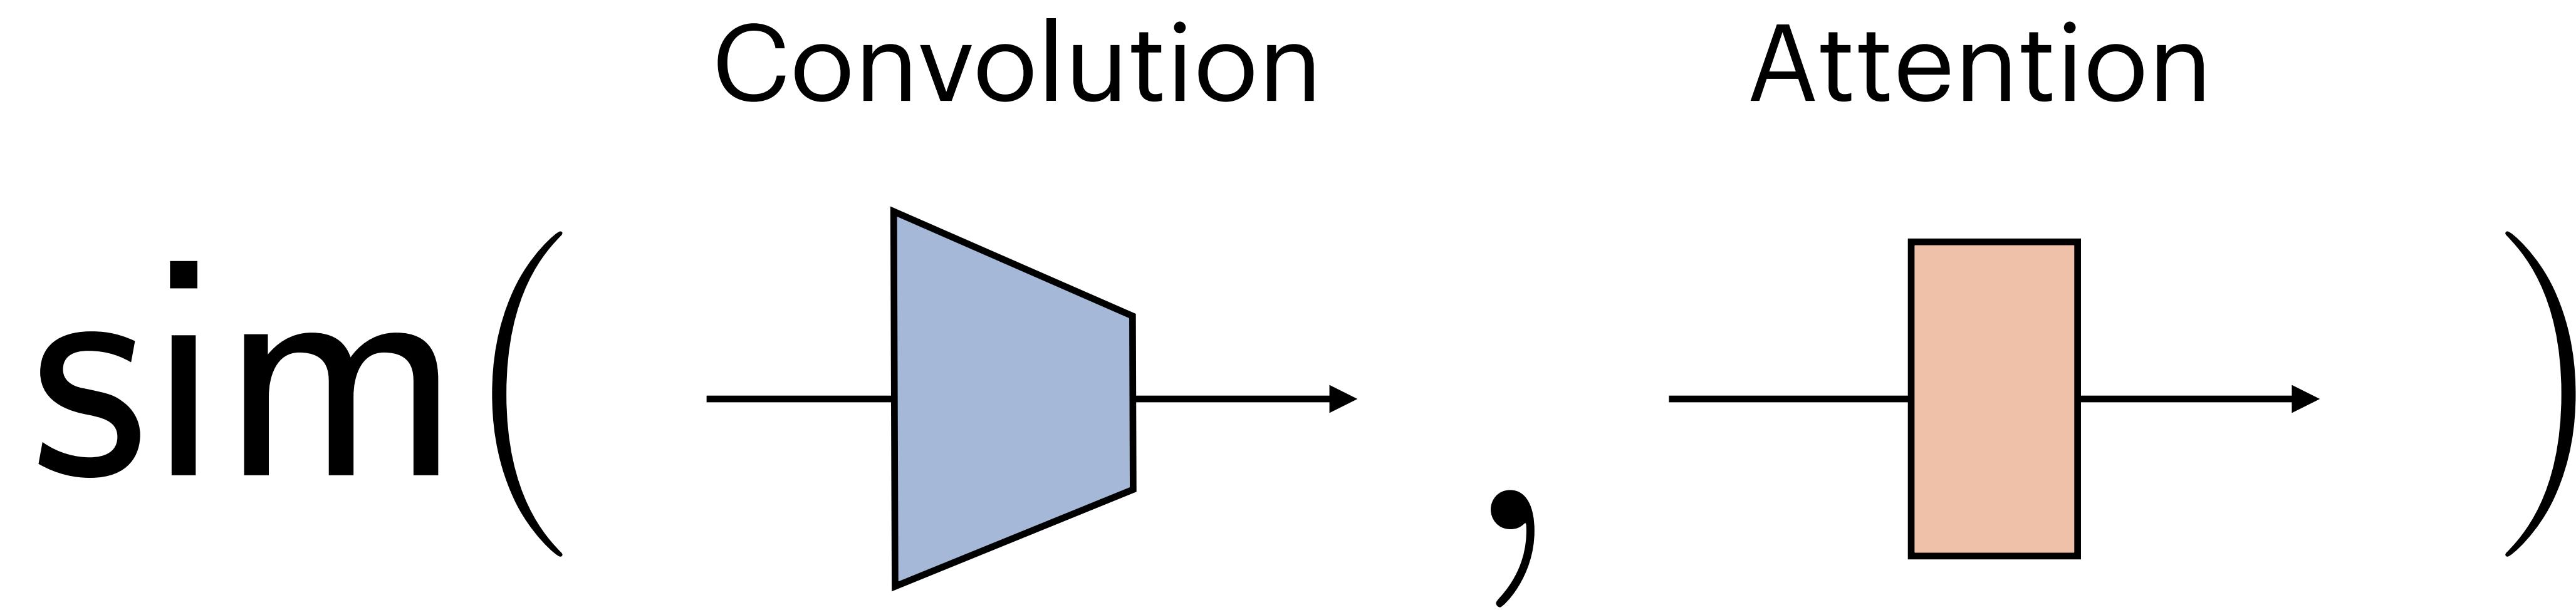
sim(



K. Tiwary\*, A. Young\*, T. Klinghoffer, Z. Tasneem, A. Dave, T. Poggio, D.-E. Nilsson

**B. Cheung\*\***, R. Raskar\*\*

“What if Eye...? Computationally Recreating Vision Evolution”

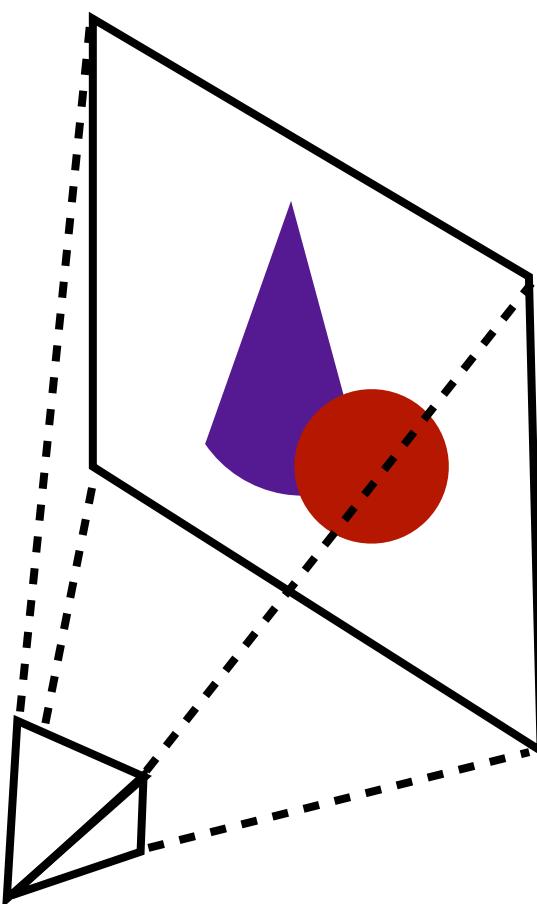


Y. Han, T. Poggio, **B. Cheung**

“System identification of neural systems: If we got it right, would we know?”

sim(

Vision



,

Language

A red sphere next  
to a blue cone.

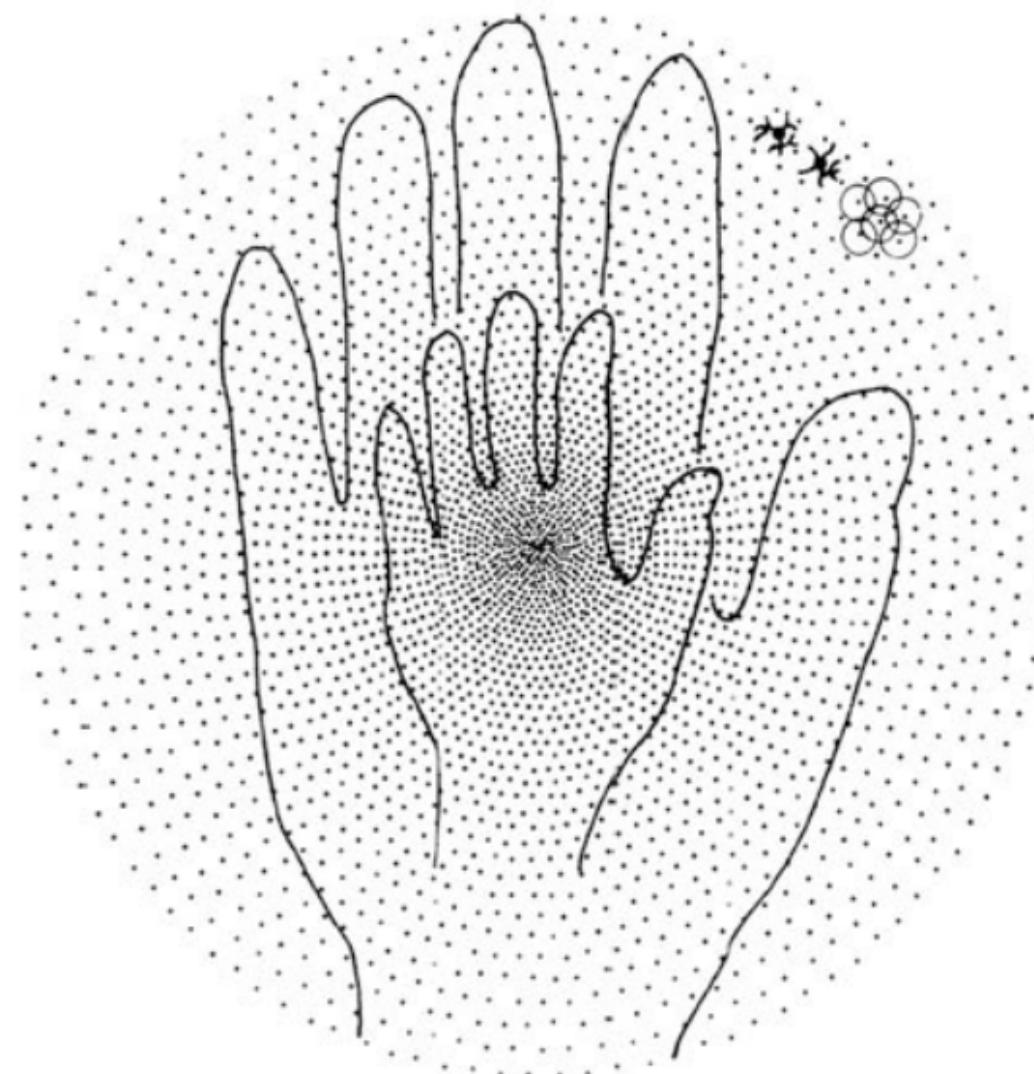
)

M. Huh\*, **B. Cheung\***, T. Wang\*, P. Isola\*

“The Platonic Representation Hypothesis”

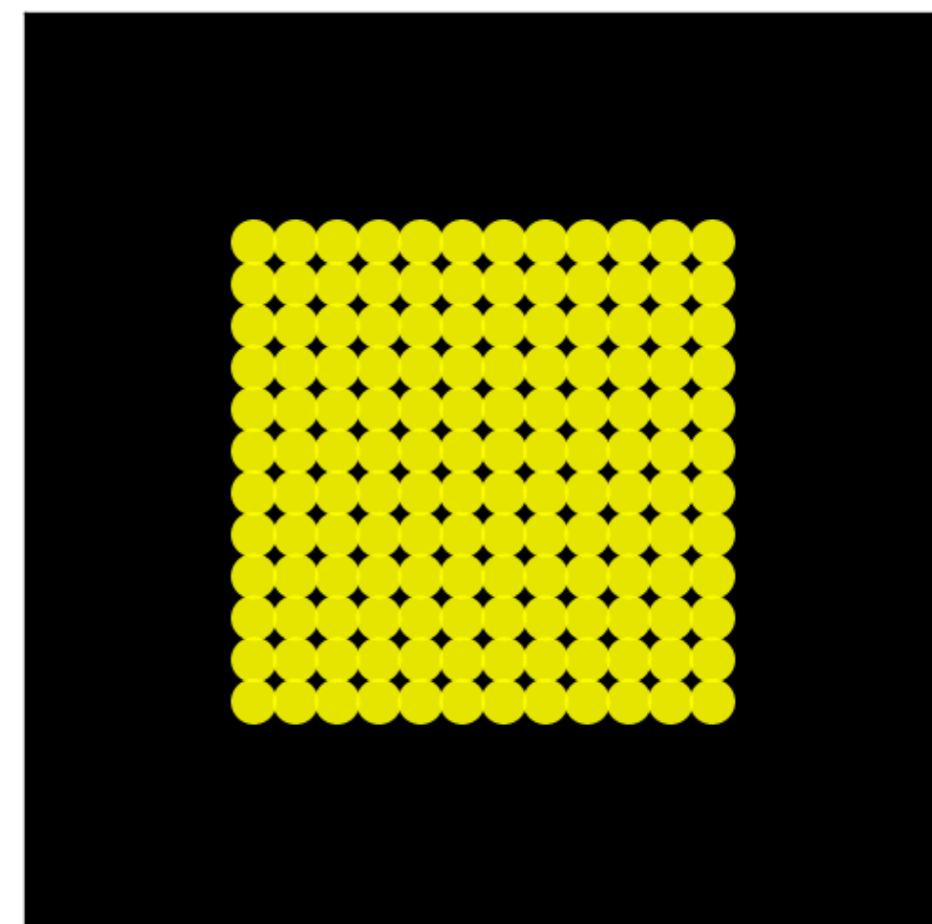
sim(

In-vivo



,

In-silico



)

**B. Cheung**, E.A. Weiss, B.A. Olshausen

“Emergence of foveal image sampling from learning to attend in visual scenes”

# Future Work

What is alignment?

A consistent representation across all ~~modalities~~ projections of the real world.

What is science?

The search for a **consistent** representation of the real world.

What if we **optimized for alignment** to find consistent representations?

# Plato's hypothesis for eyes



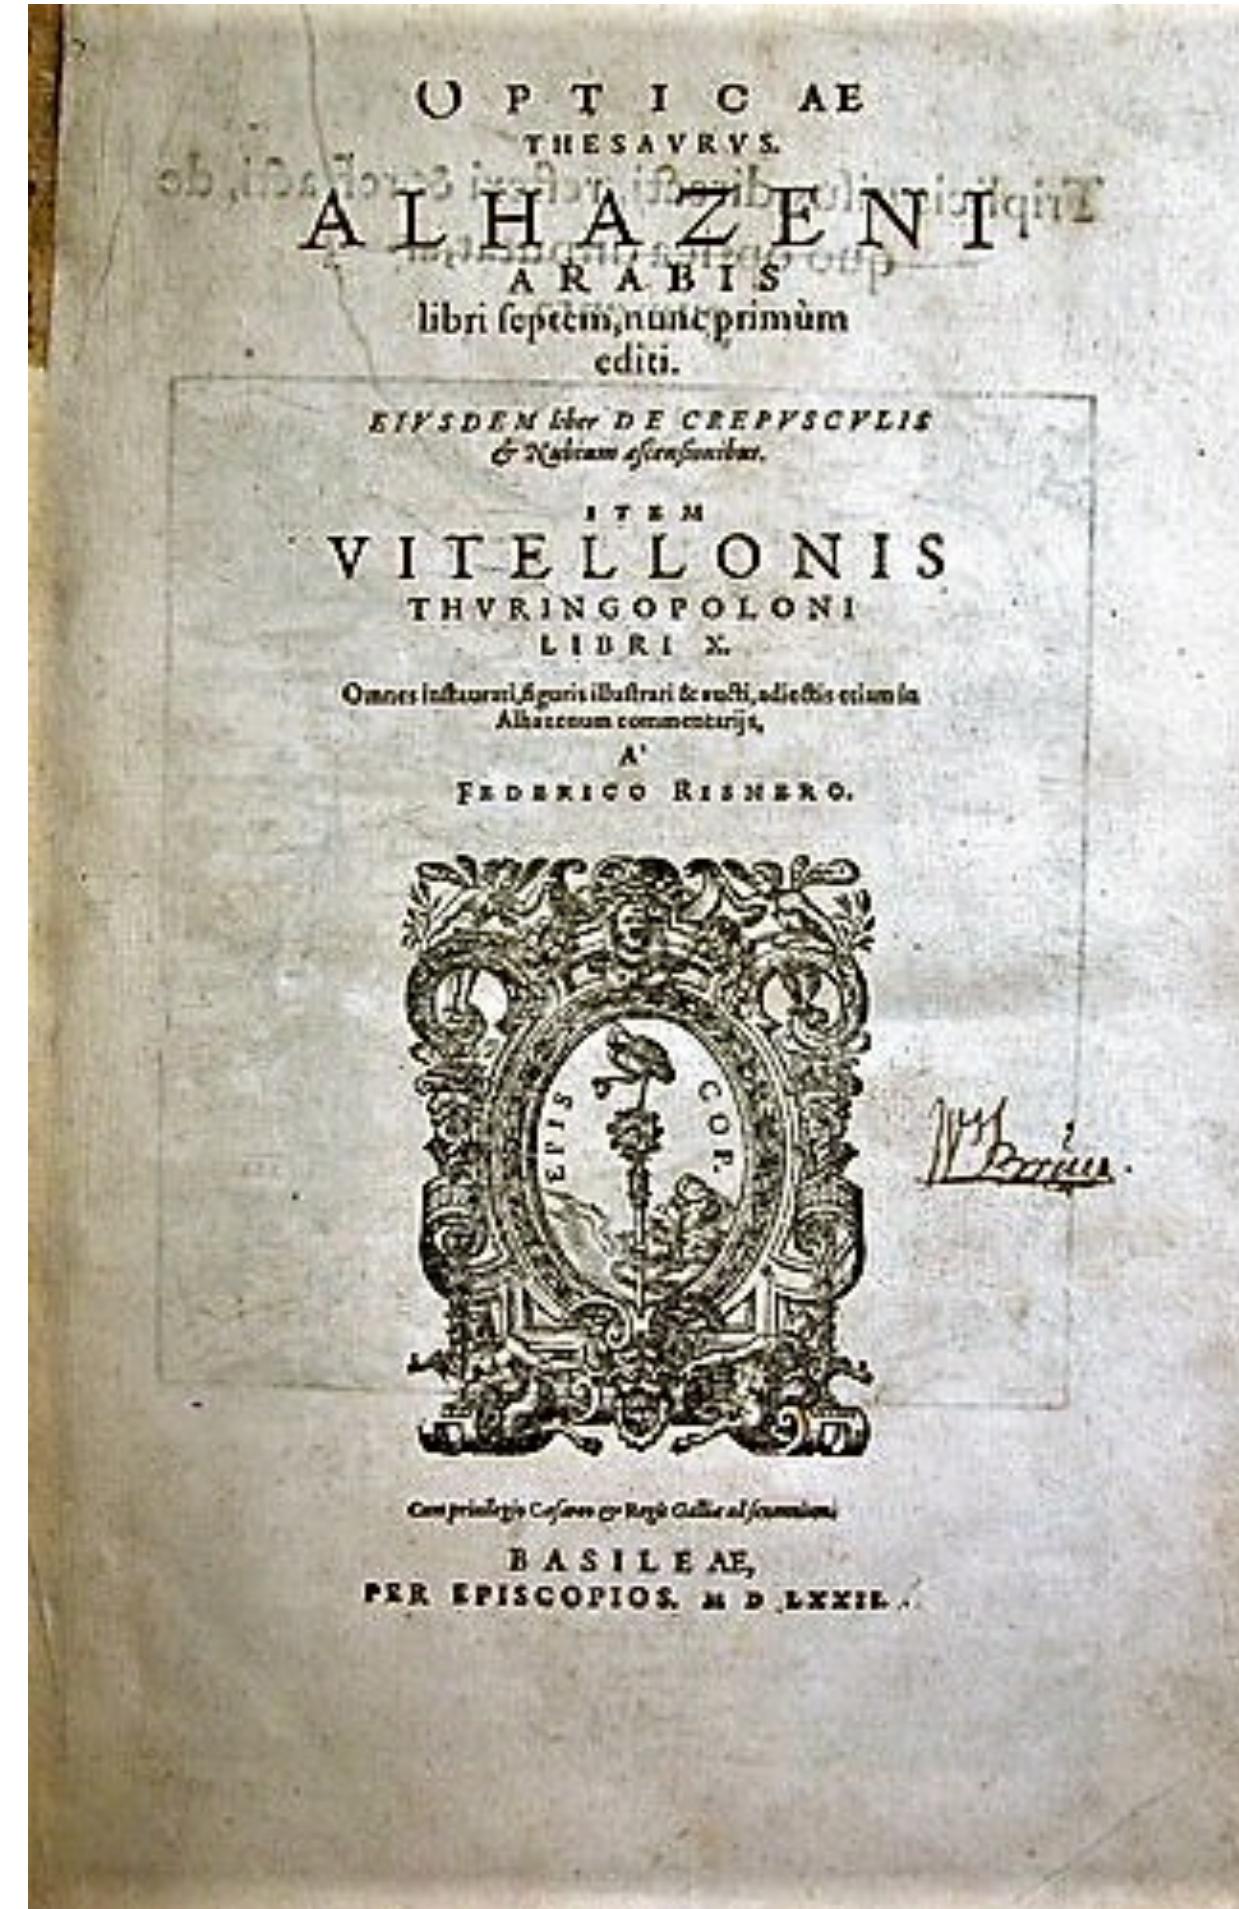
Quelle: Deutsche Fotothek

# Is my hypothesis **aligned** with my observation?

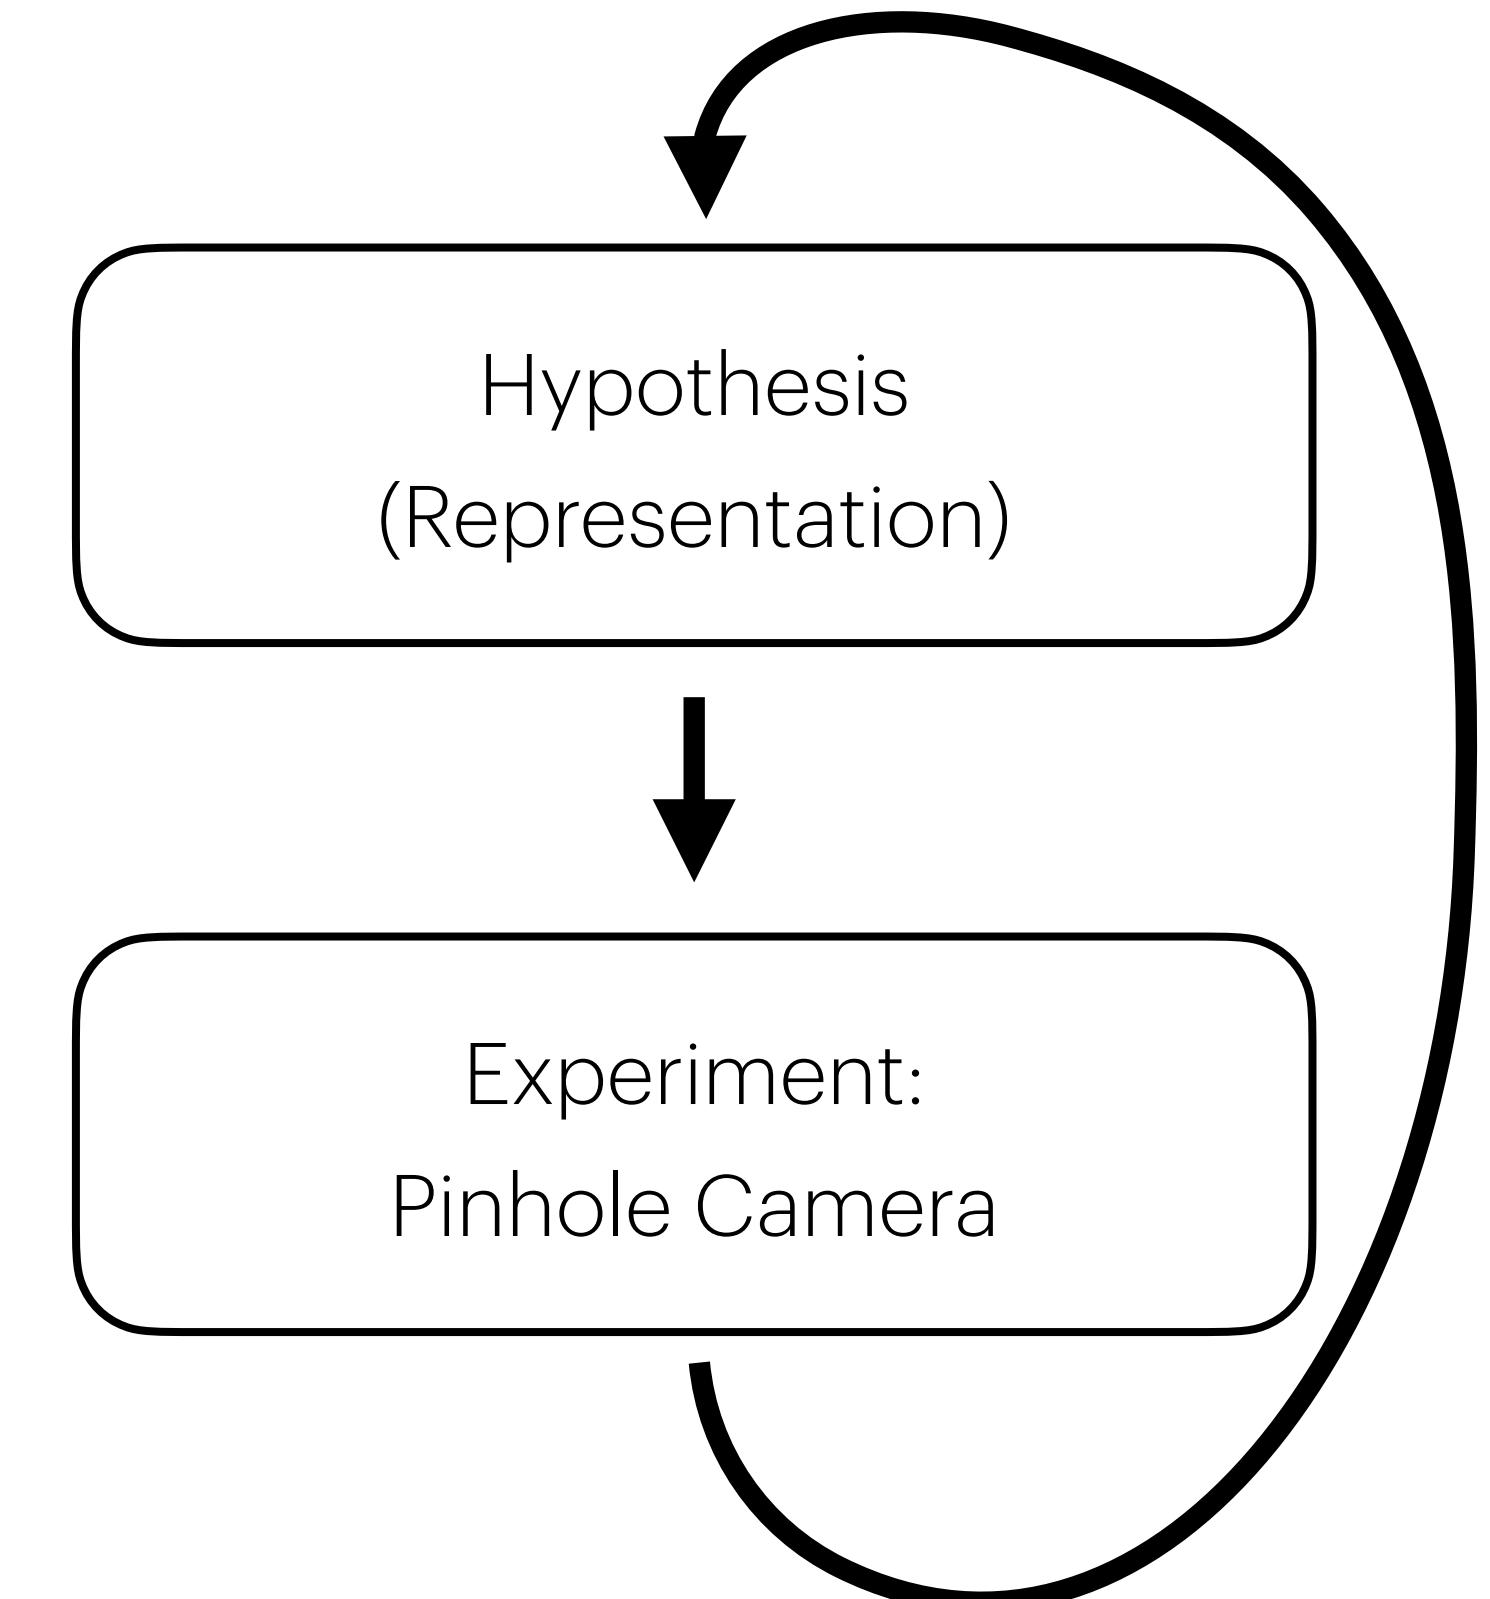
## Emission Theory



## Intromission Theory



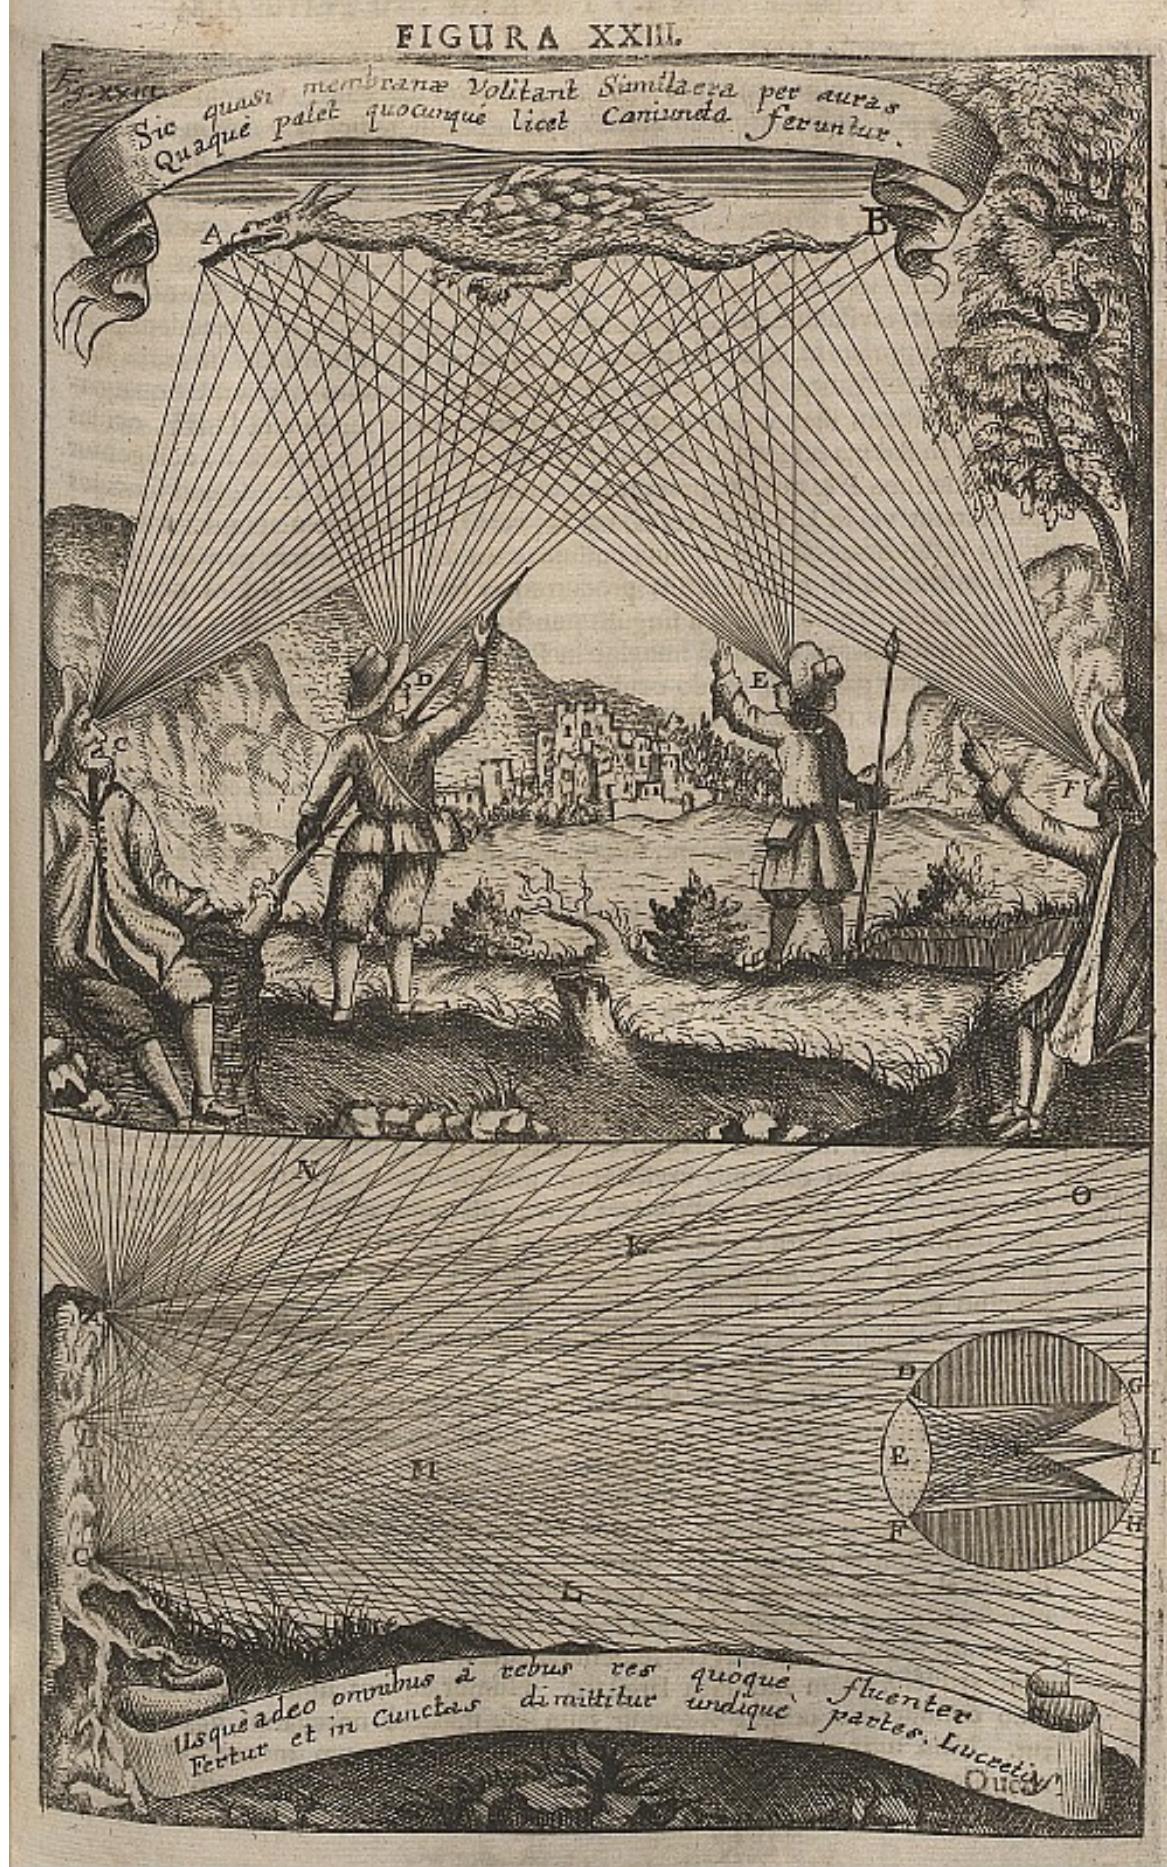
[Alhazen 1021]



Quelle: Deutsche Fotothek

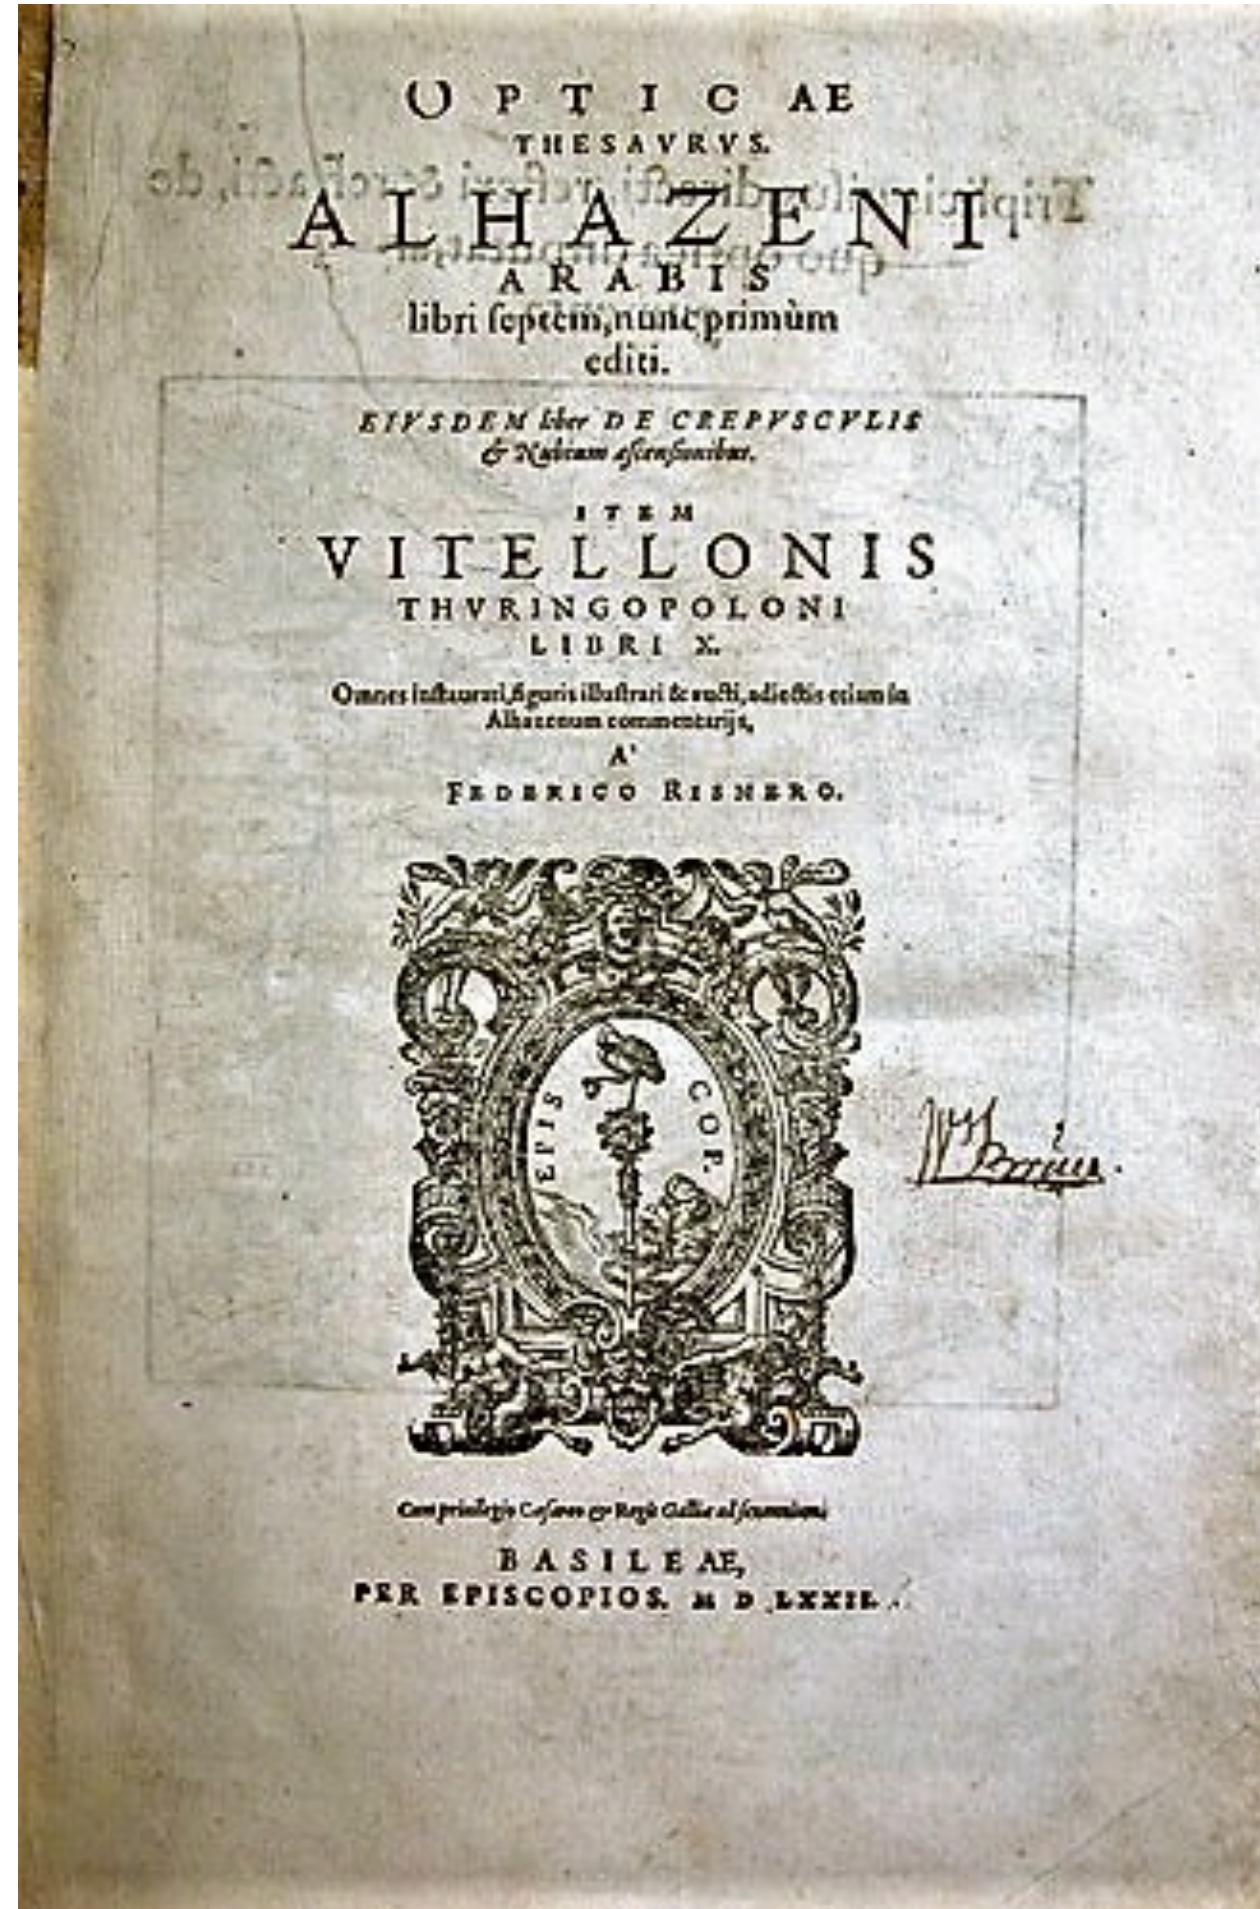
# Is my hypothesis **aligned** with my observation?

## Emission Theory

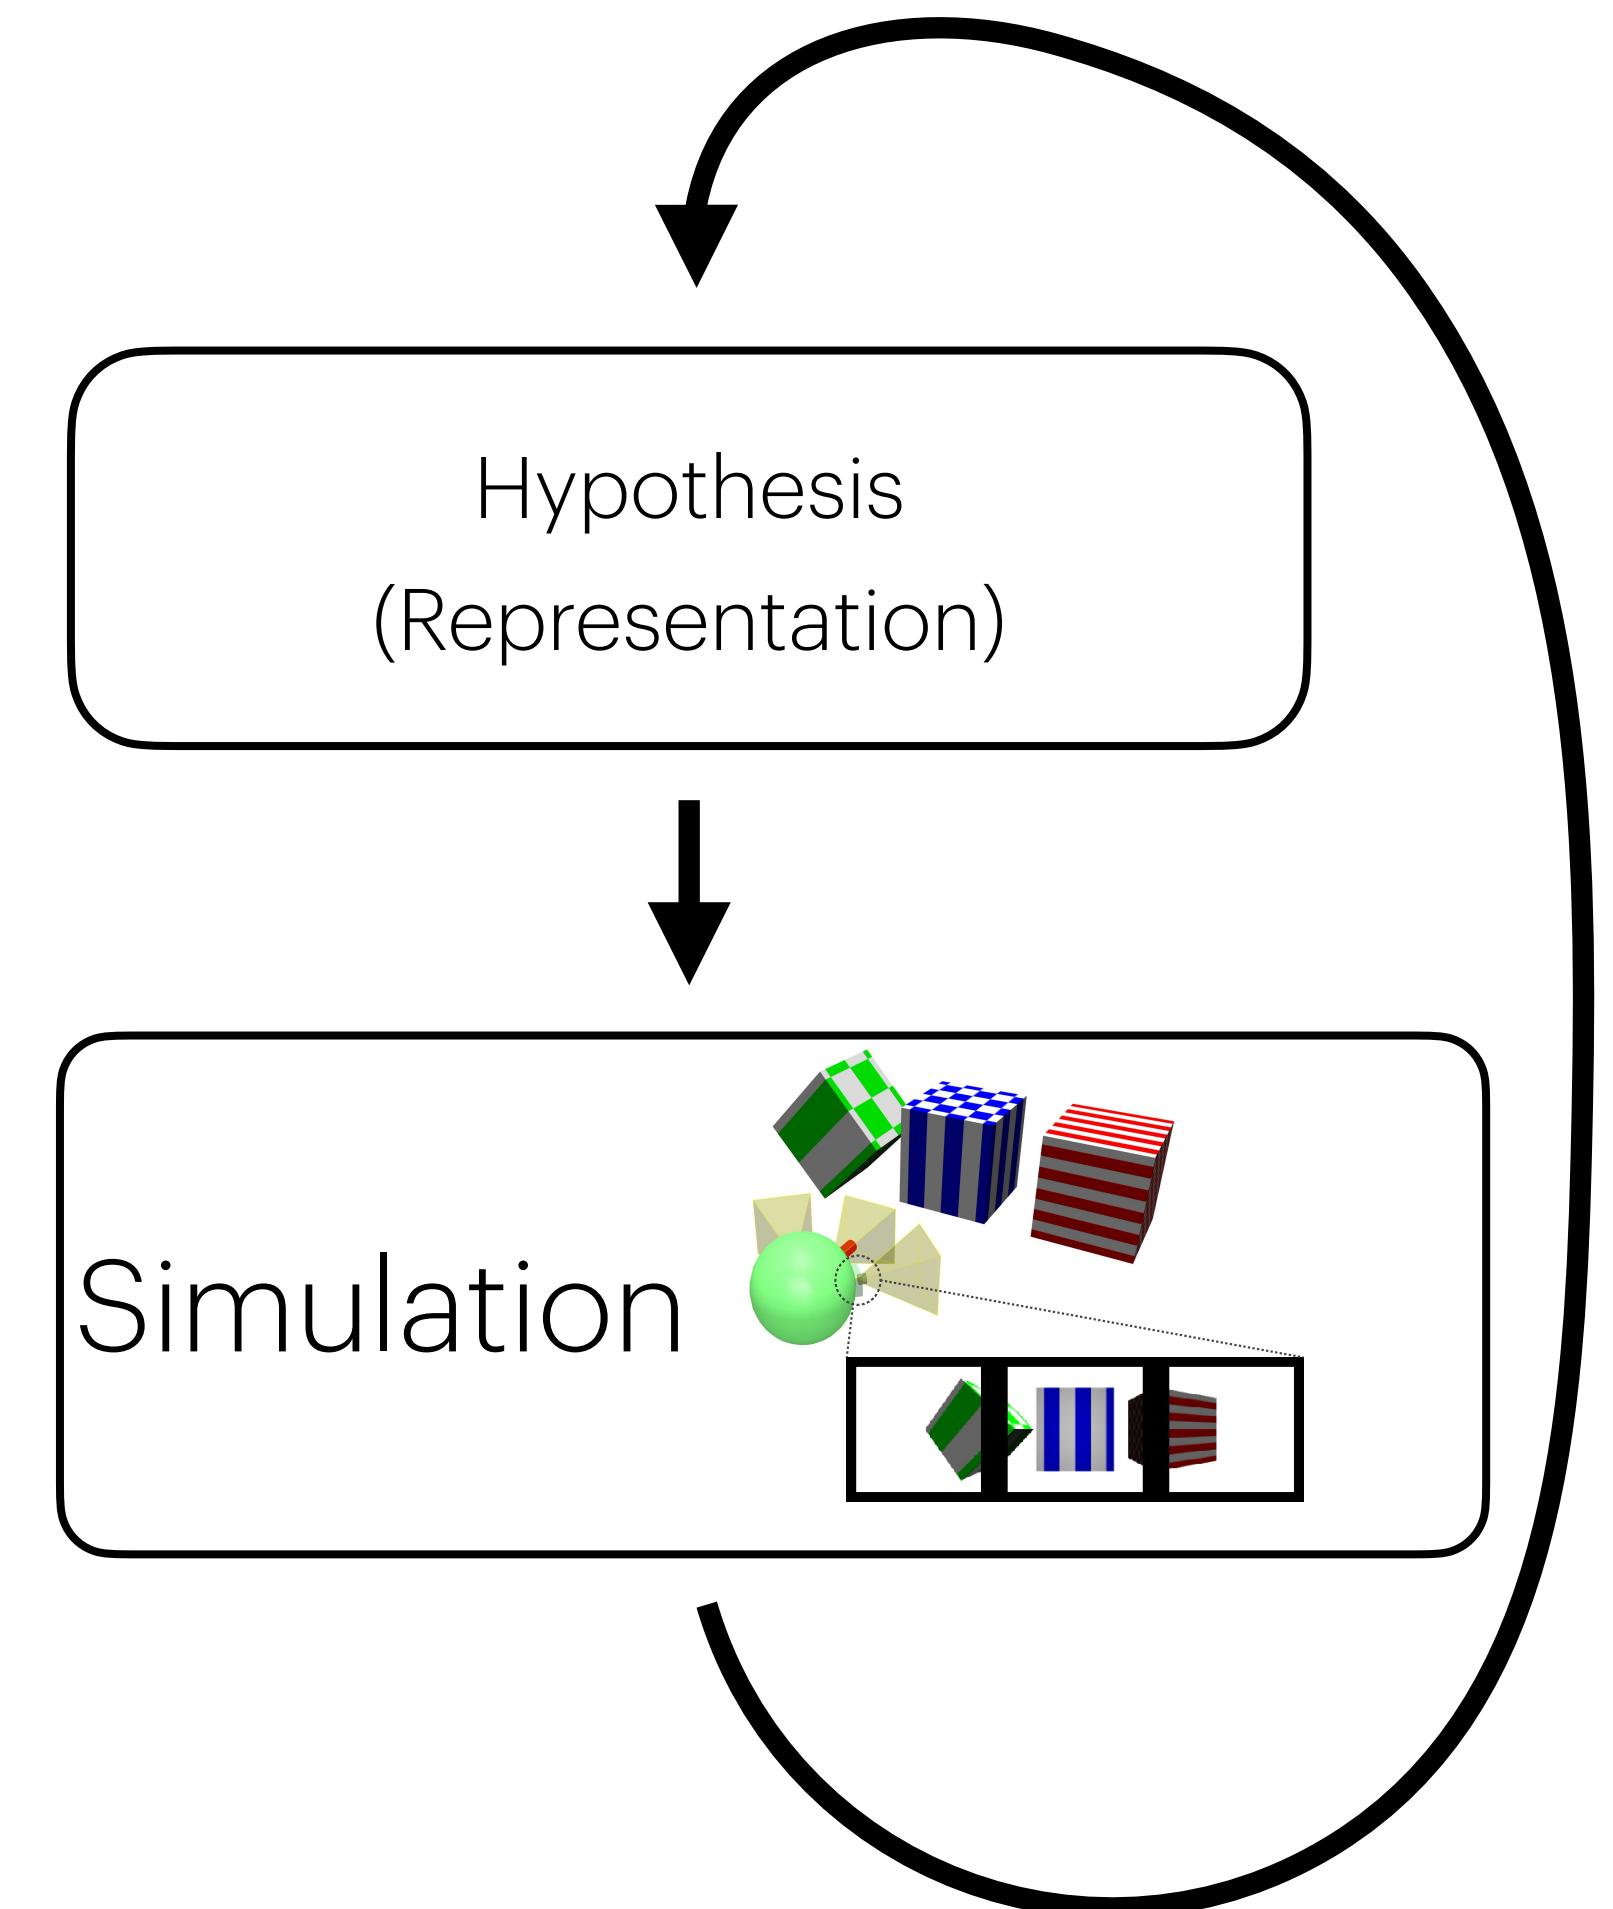


Quelle: Deutsche Fotothek

## Intromission Theory



[Alhazen 1021]





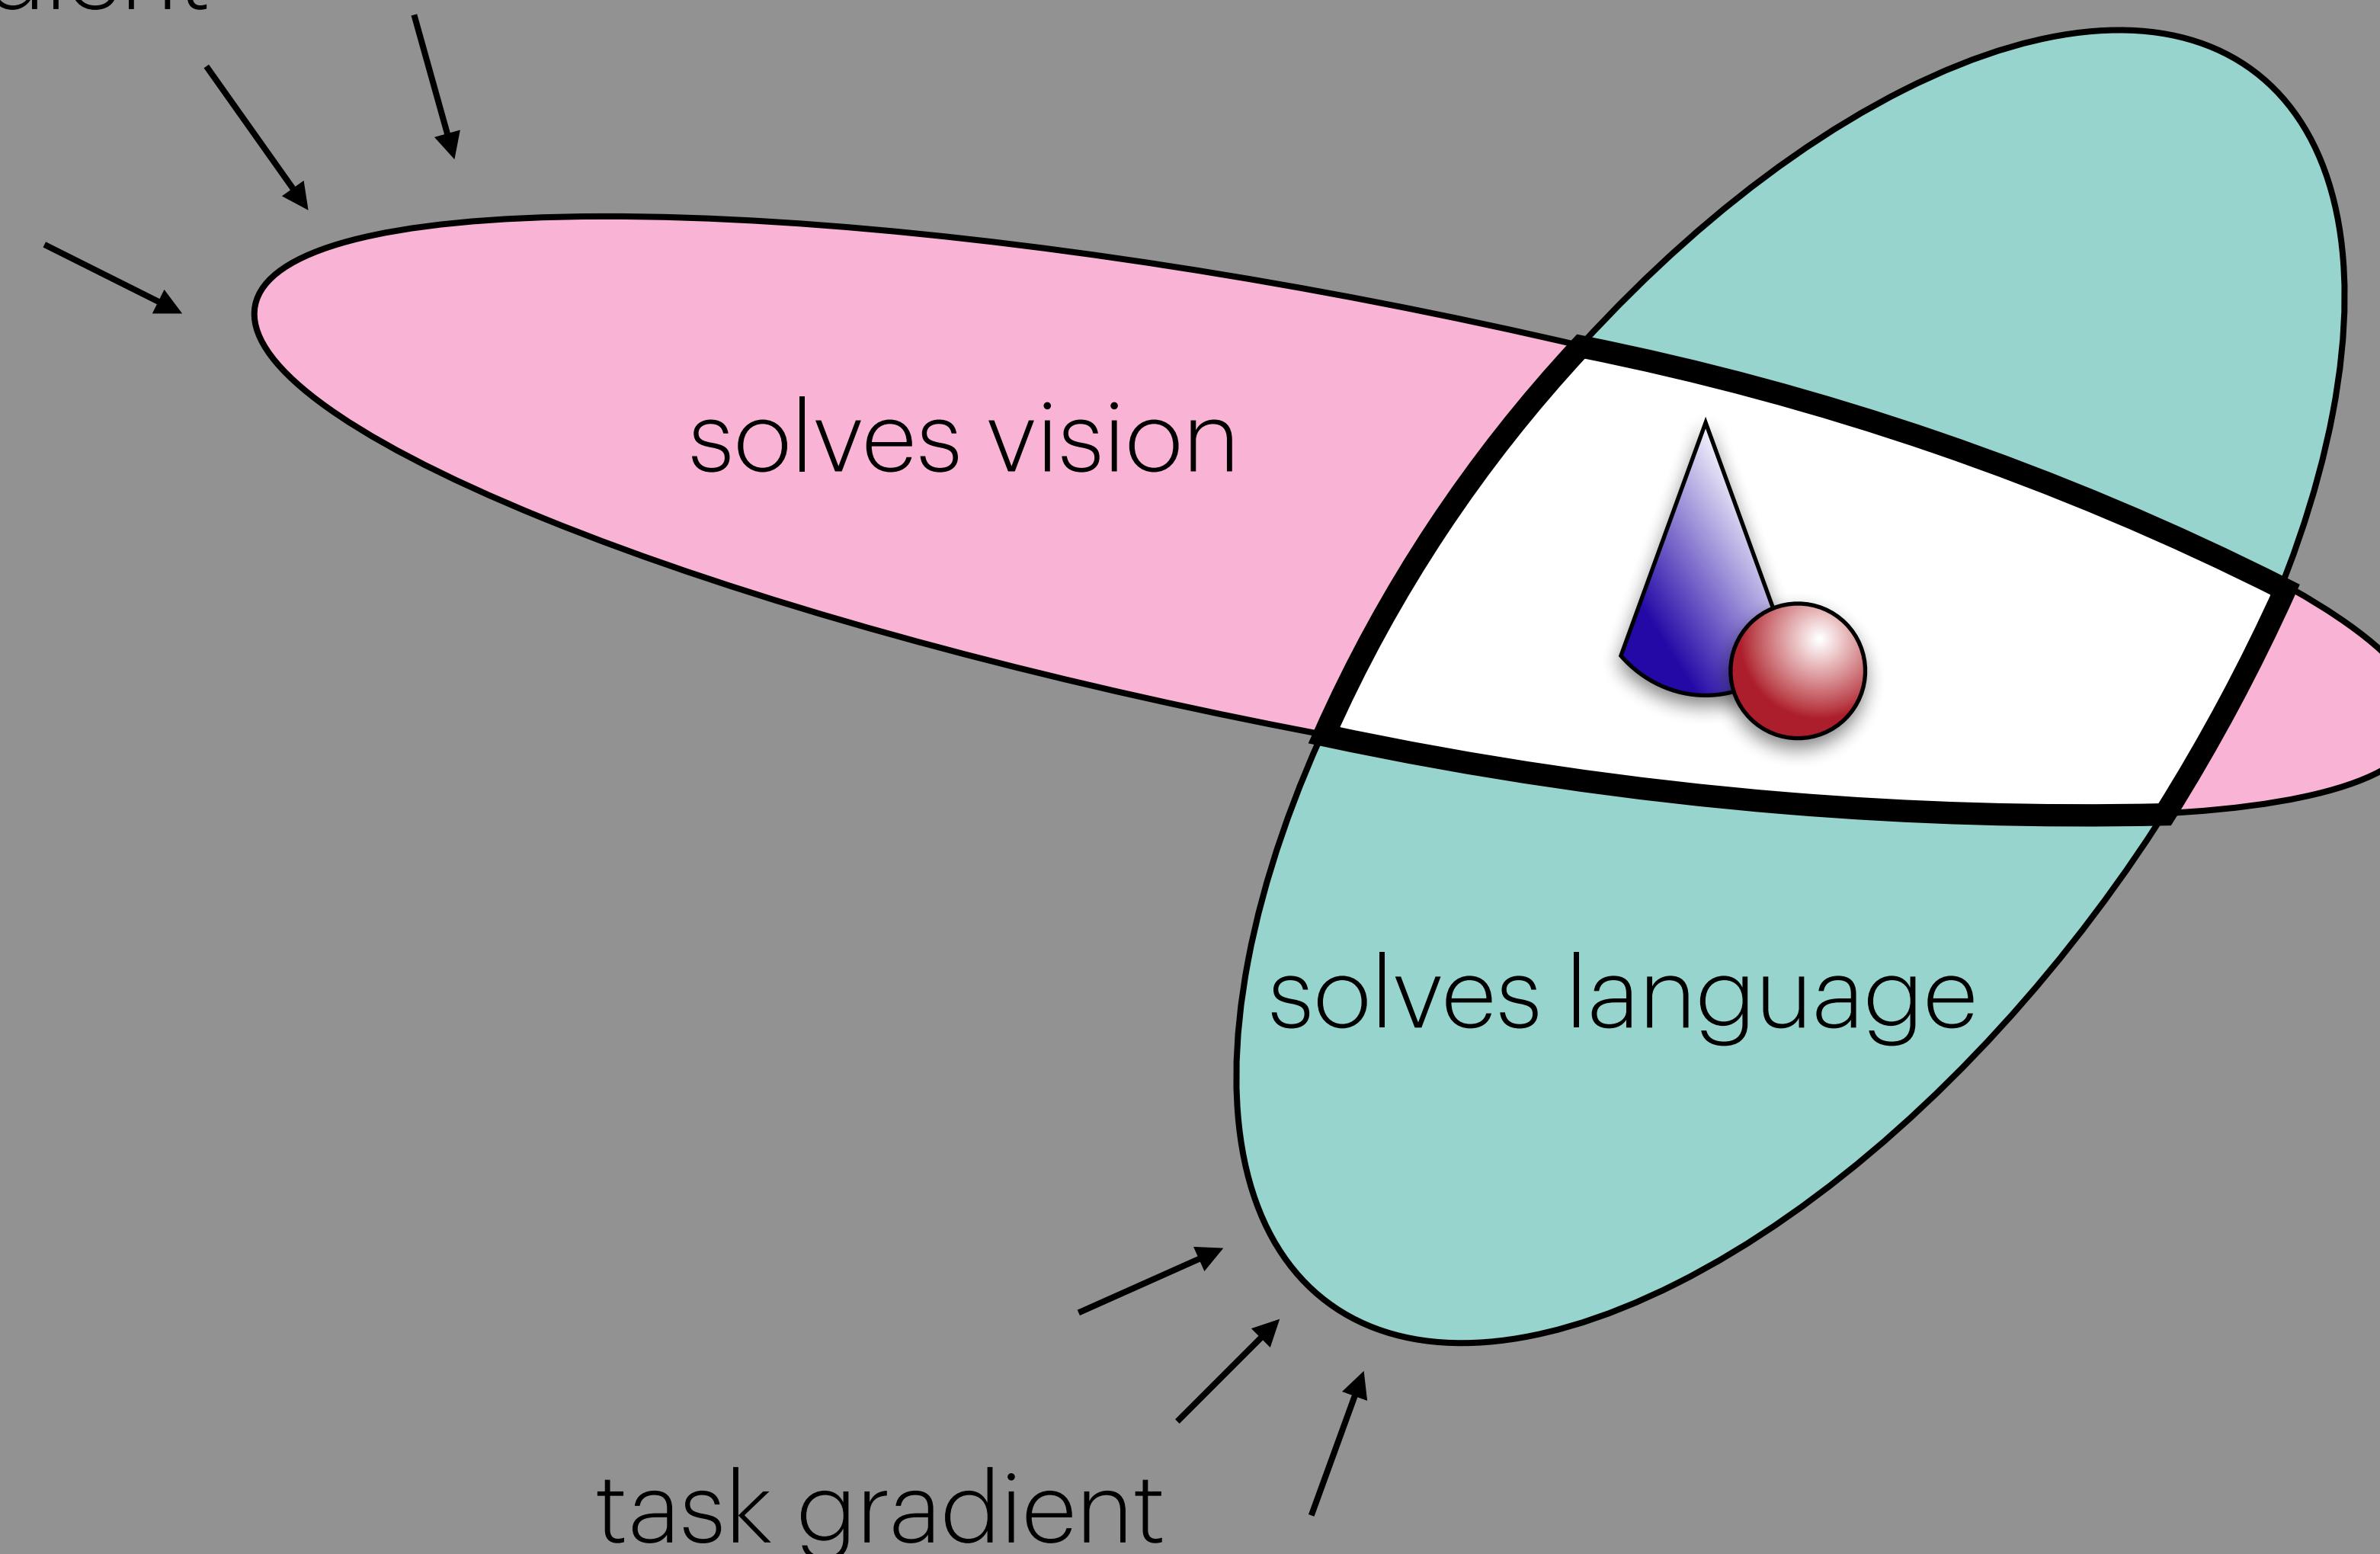
Alhazen (965 - 1040 CE)



“The duty of the man who investigates the writings of scientists, if learning the truth is his goal, is to make himself an enemy of all that he reads, and ... attack it from every side.”

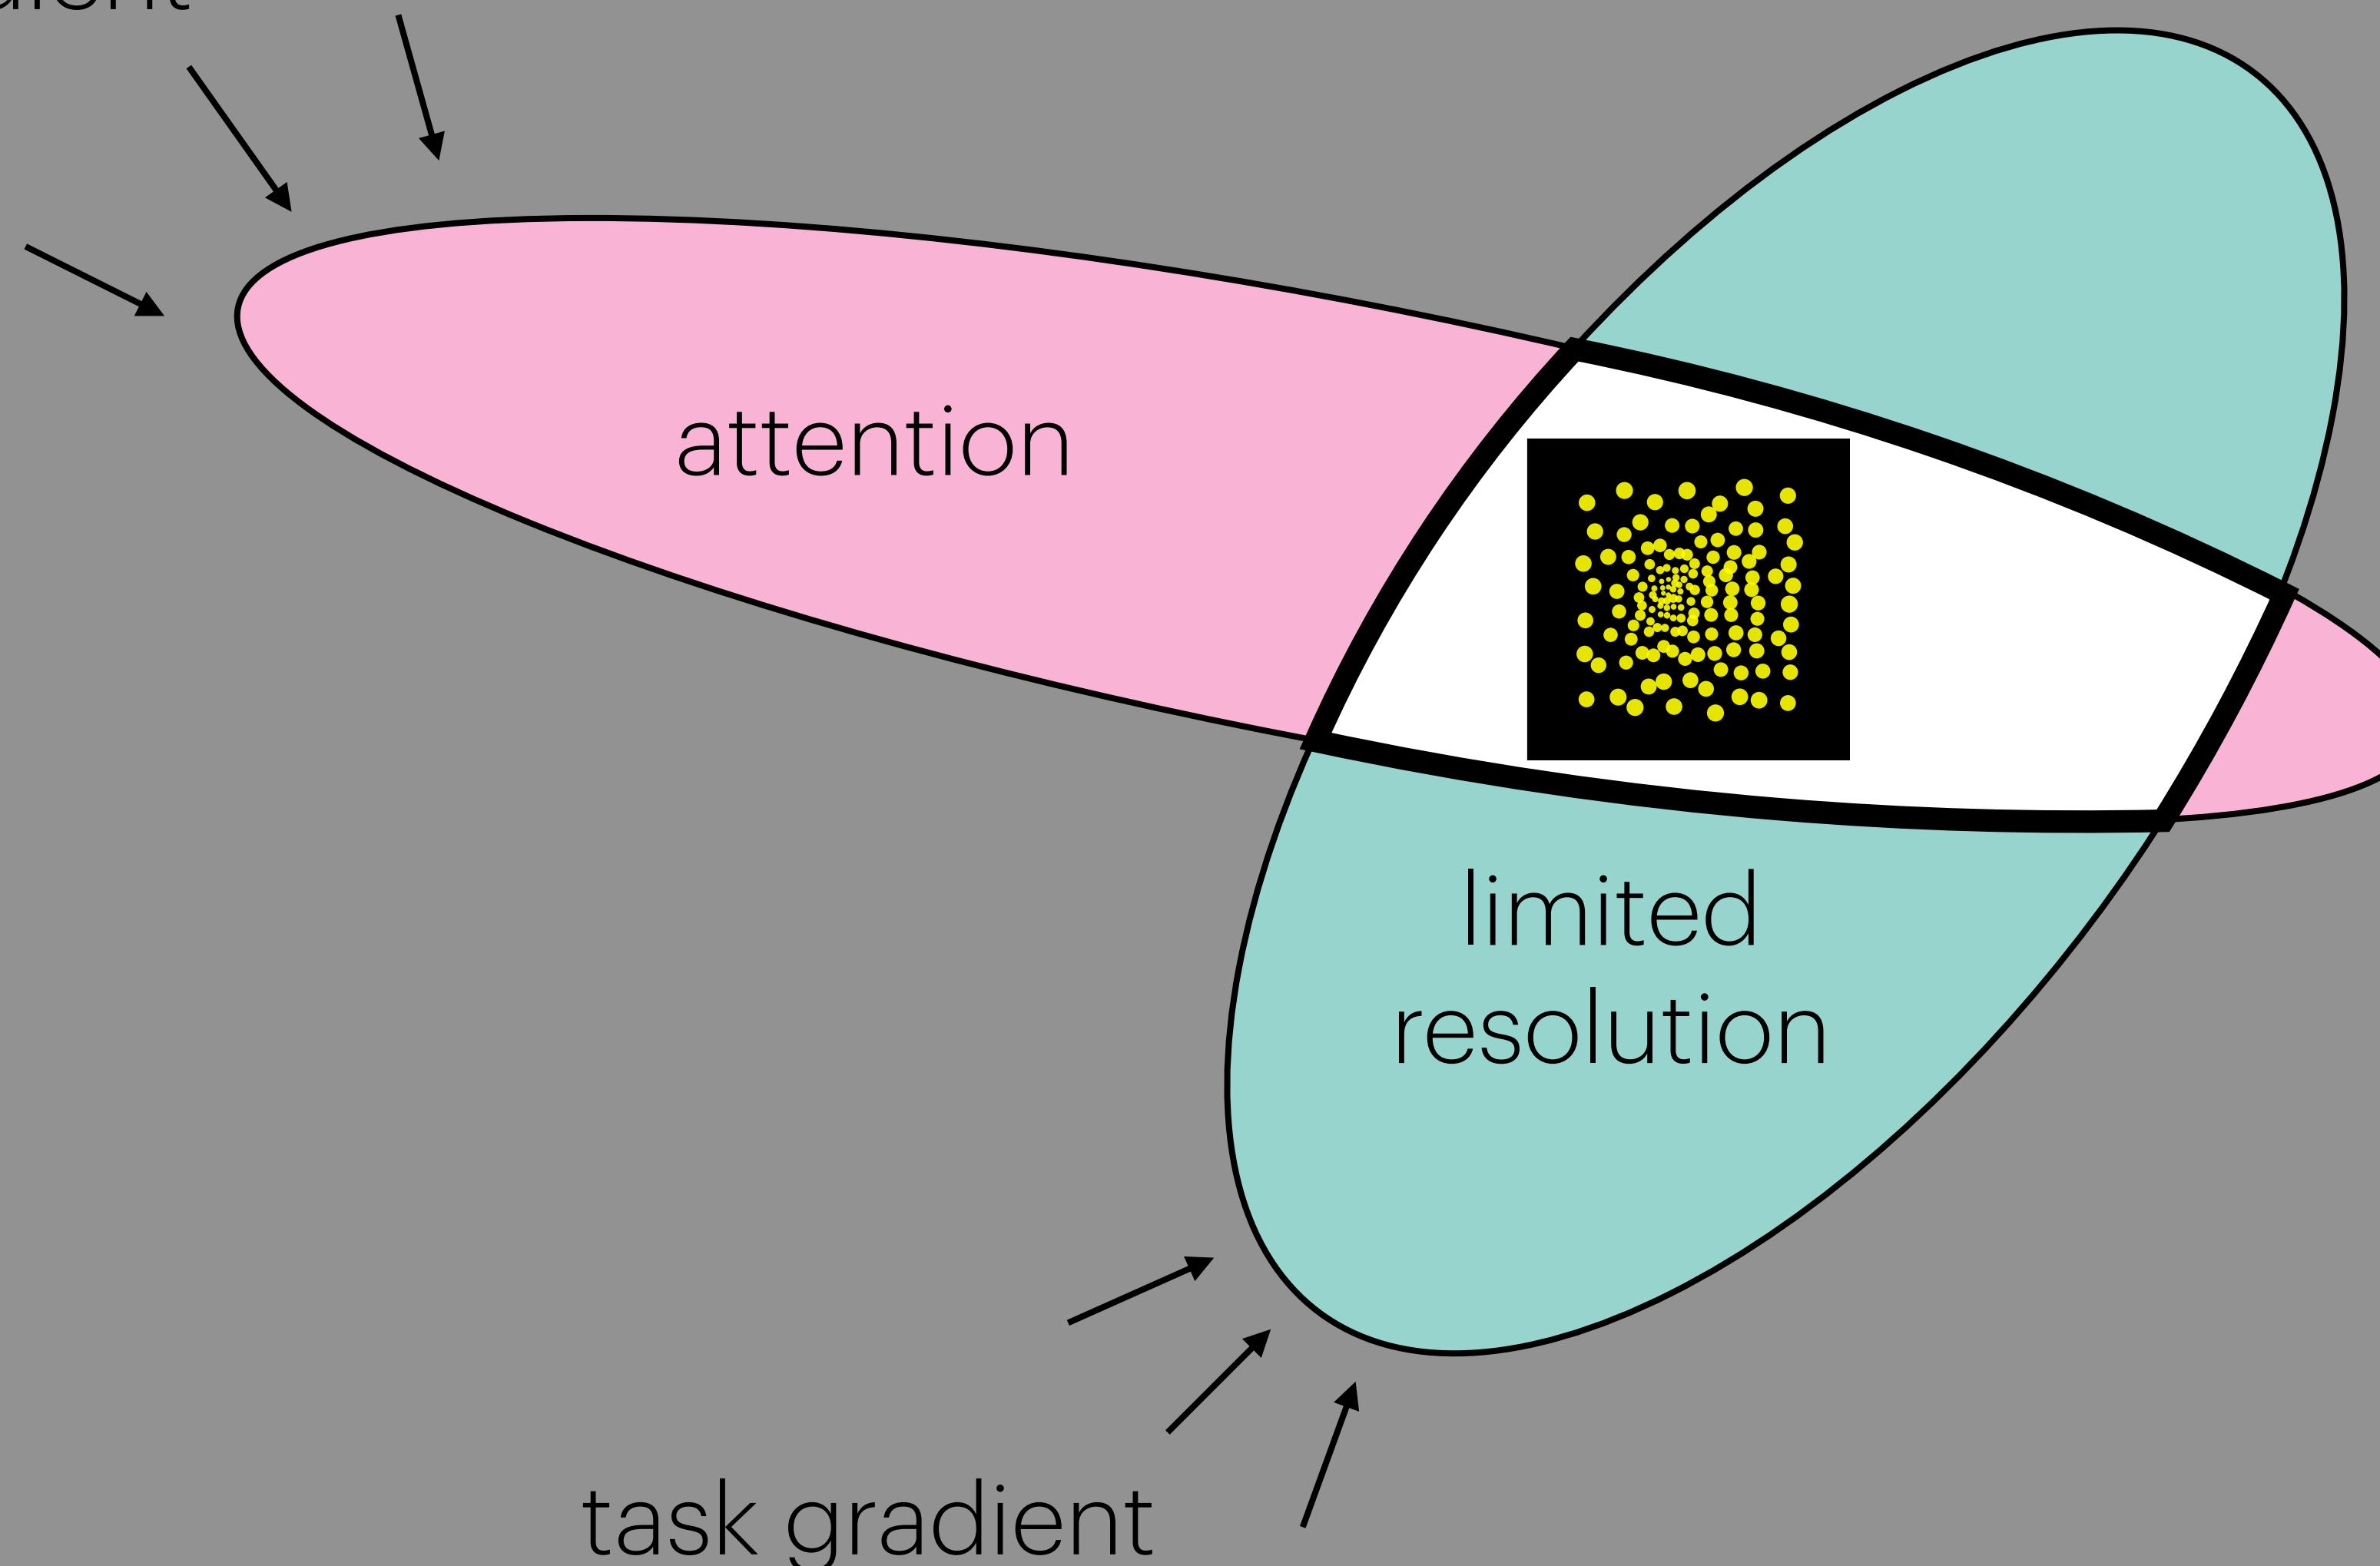
## Hypothesis Space

task gradient



# Hypothesis Space

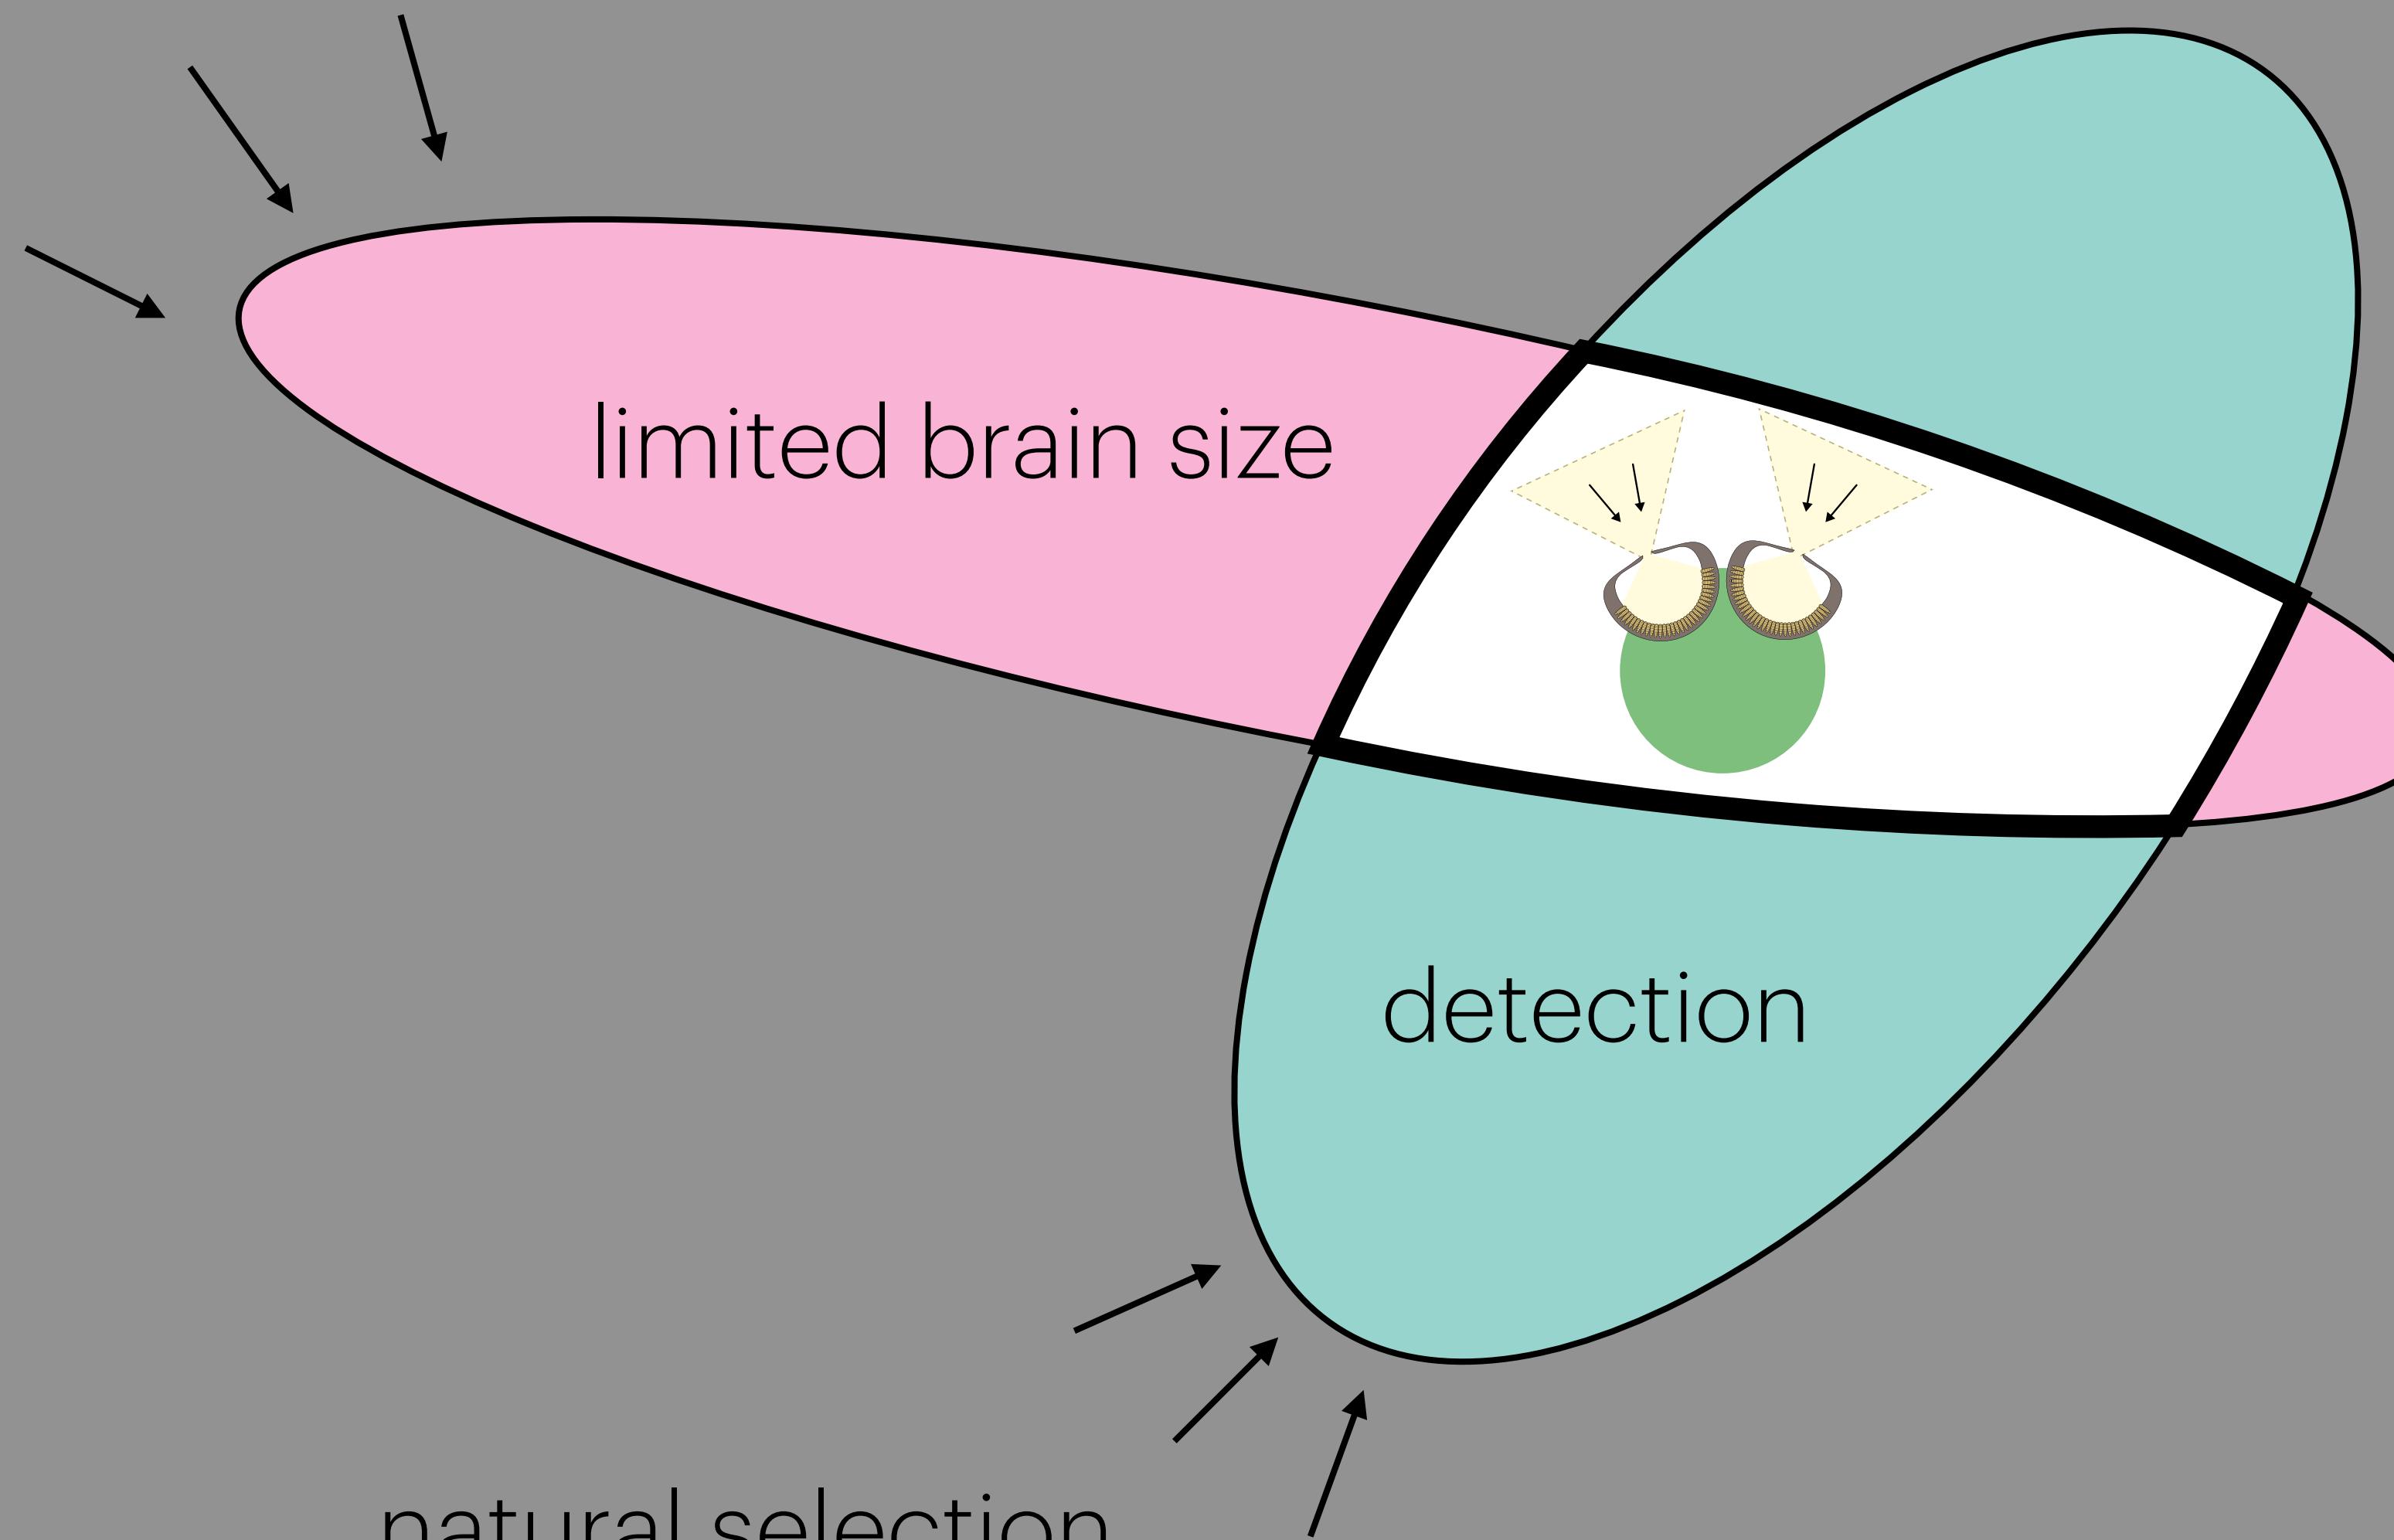
task gradient



task gradient

# Hypothesis Space

natural selection



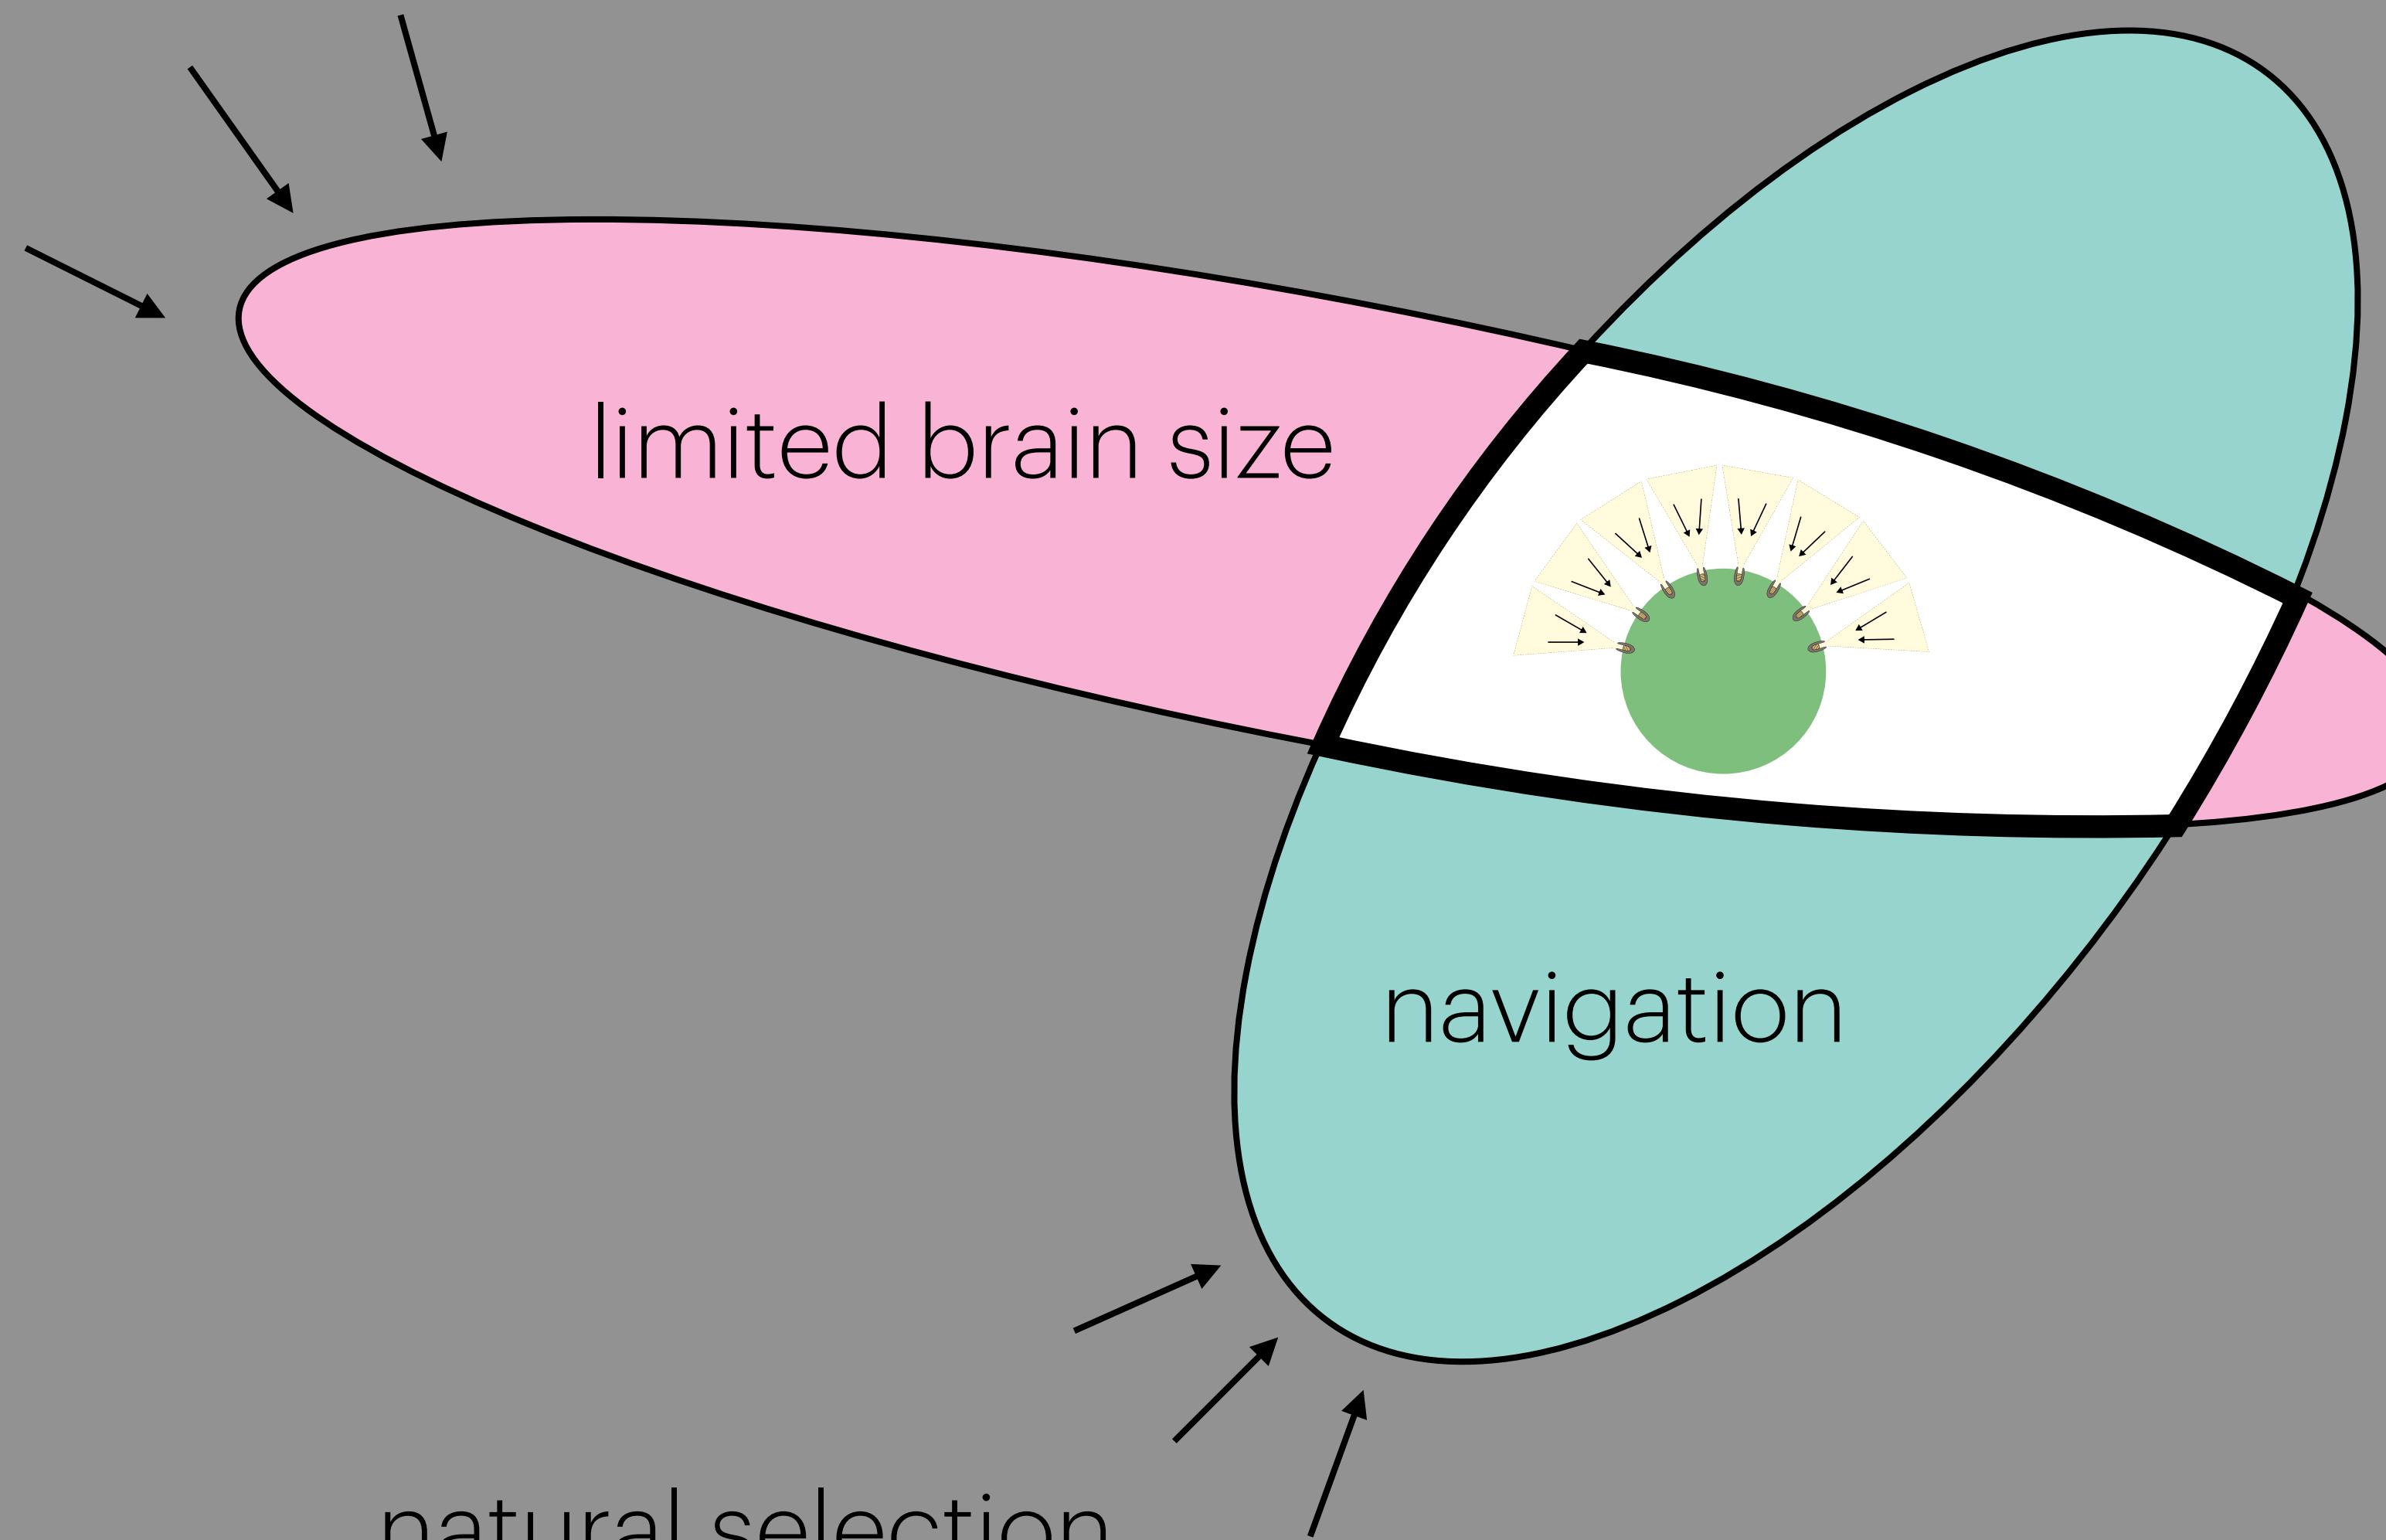
limited brain size

detection

natural selection

# Hypothesis Space

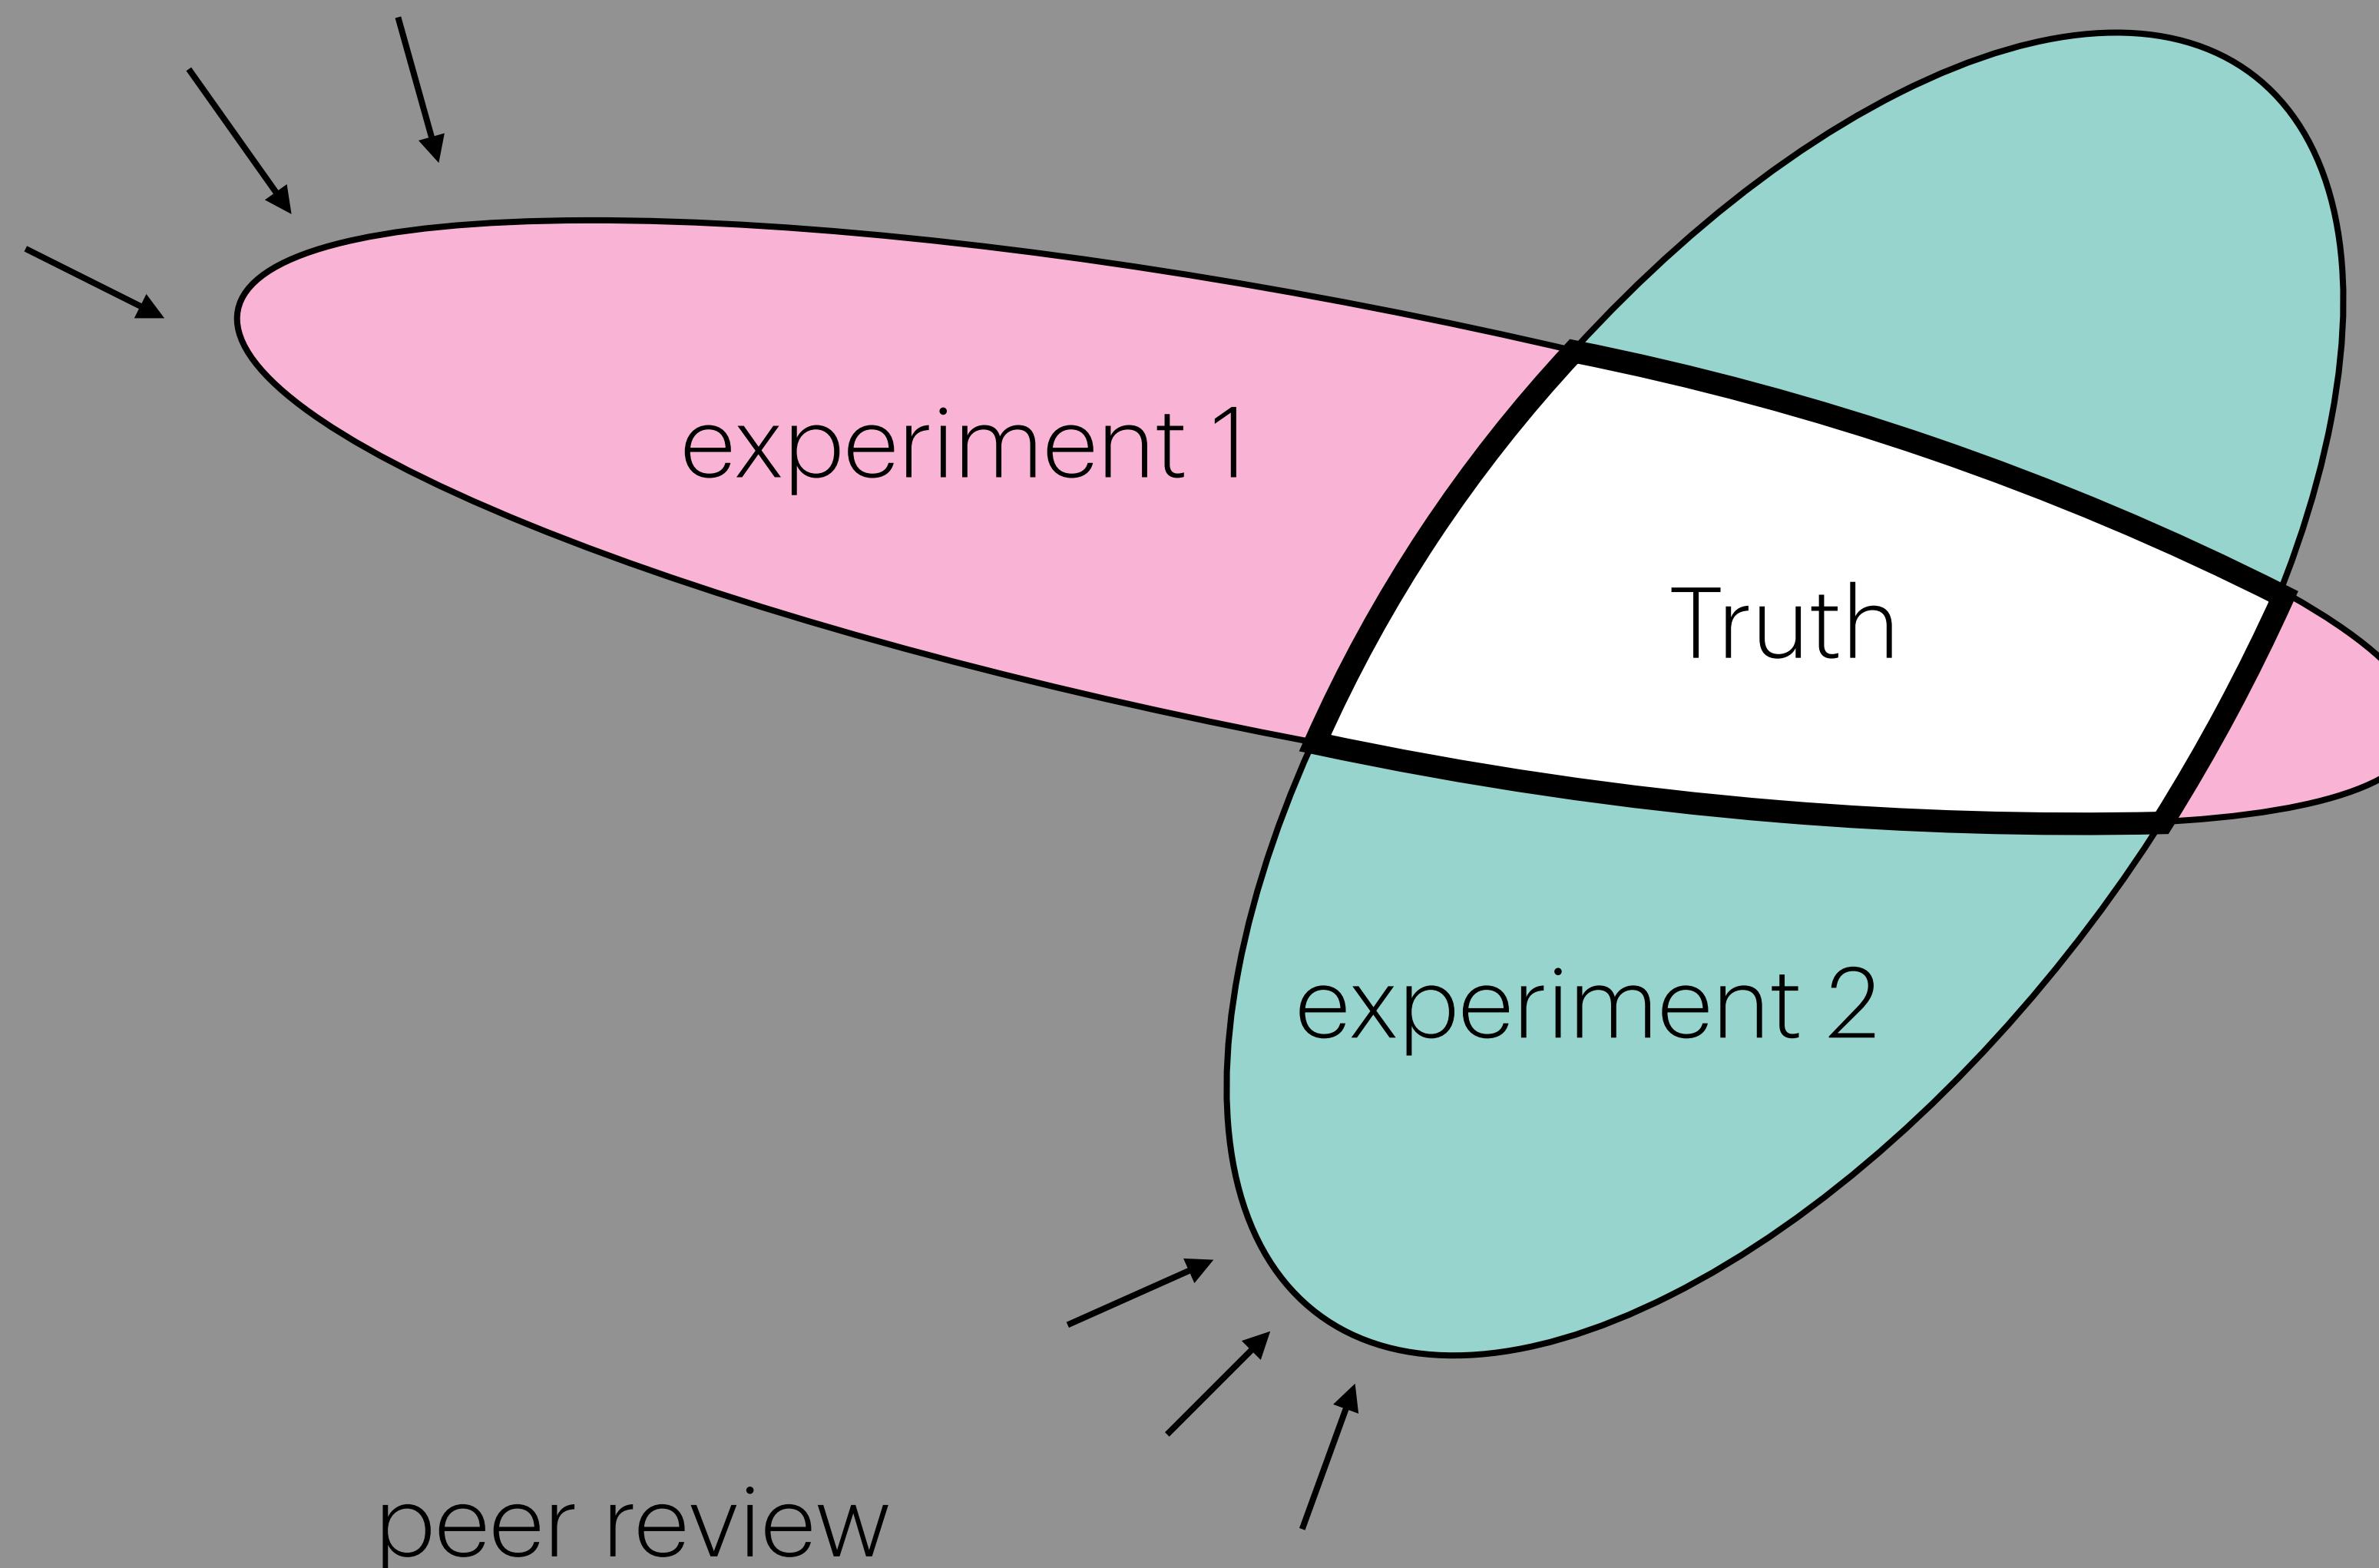
natural selection



natural selection

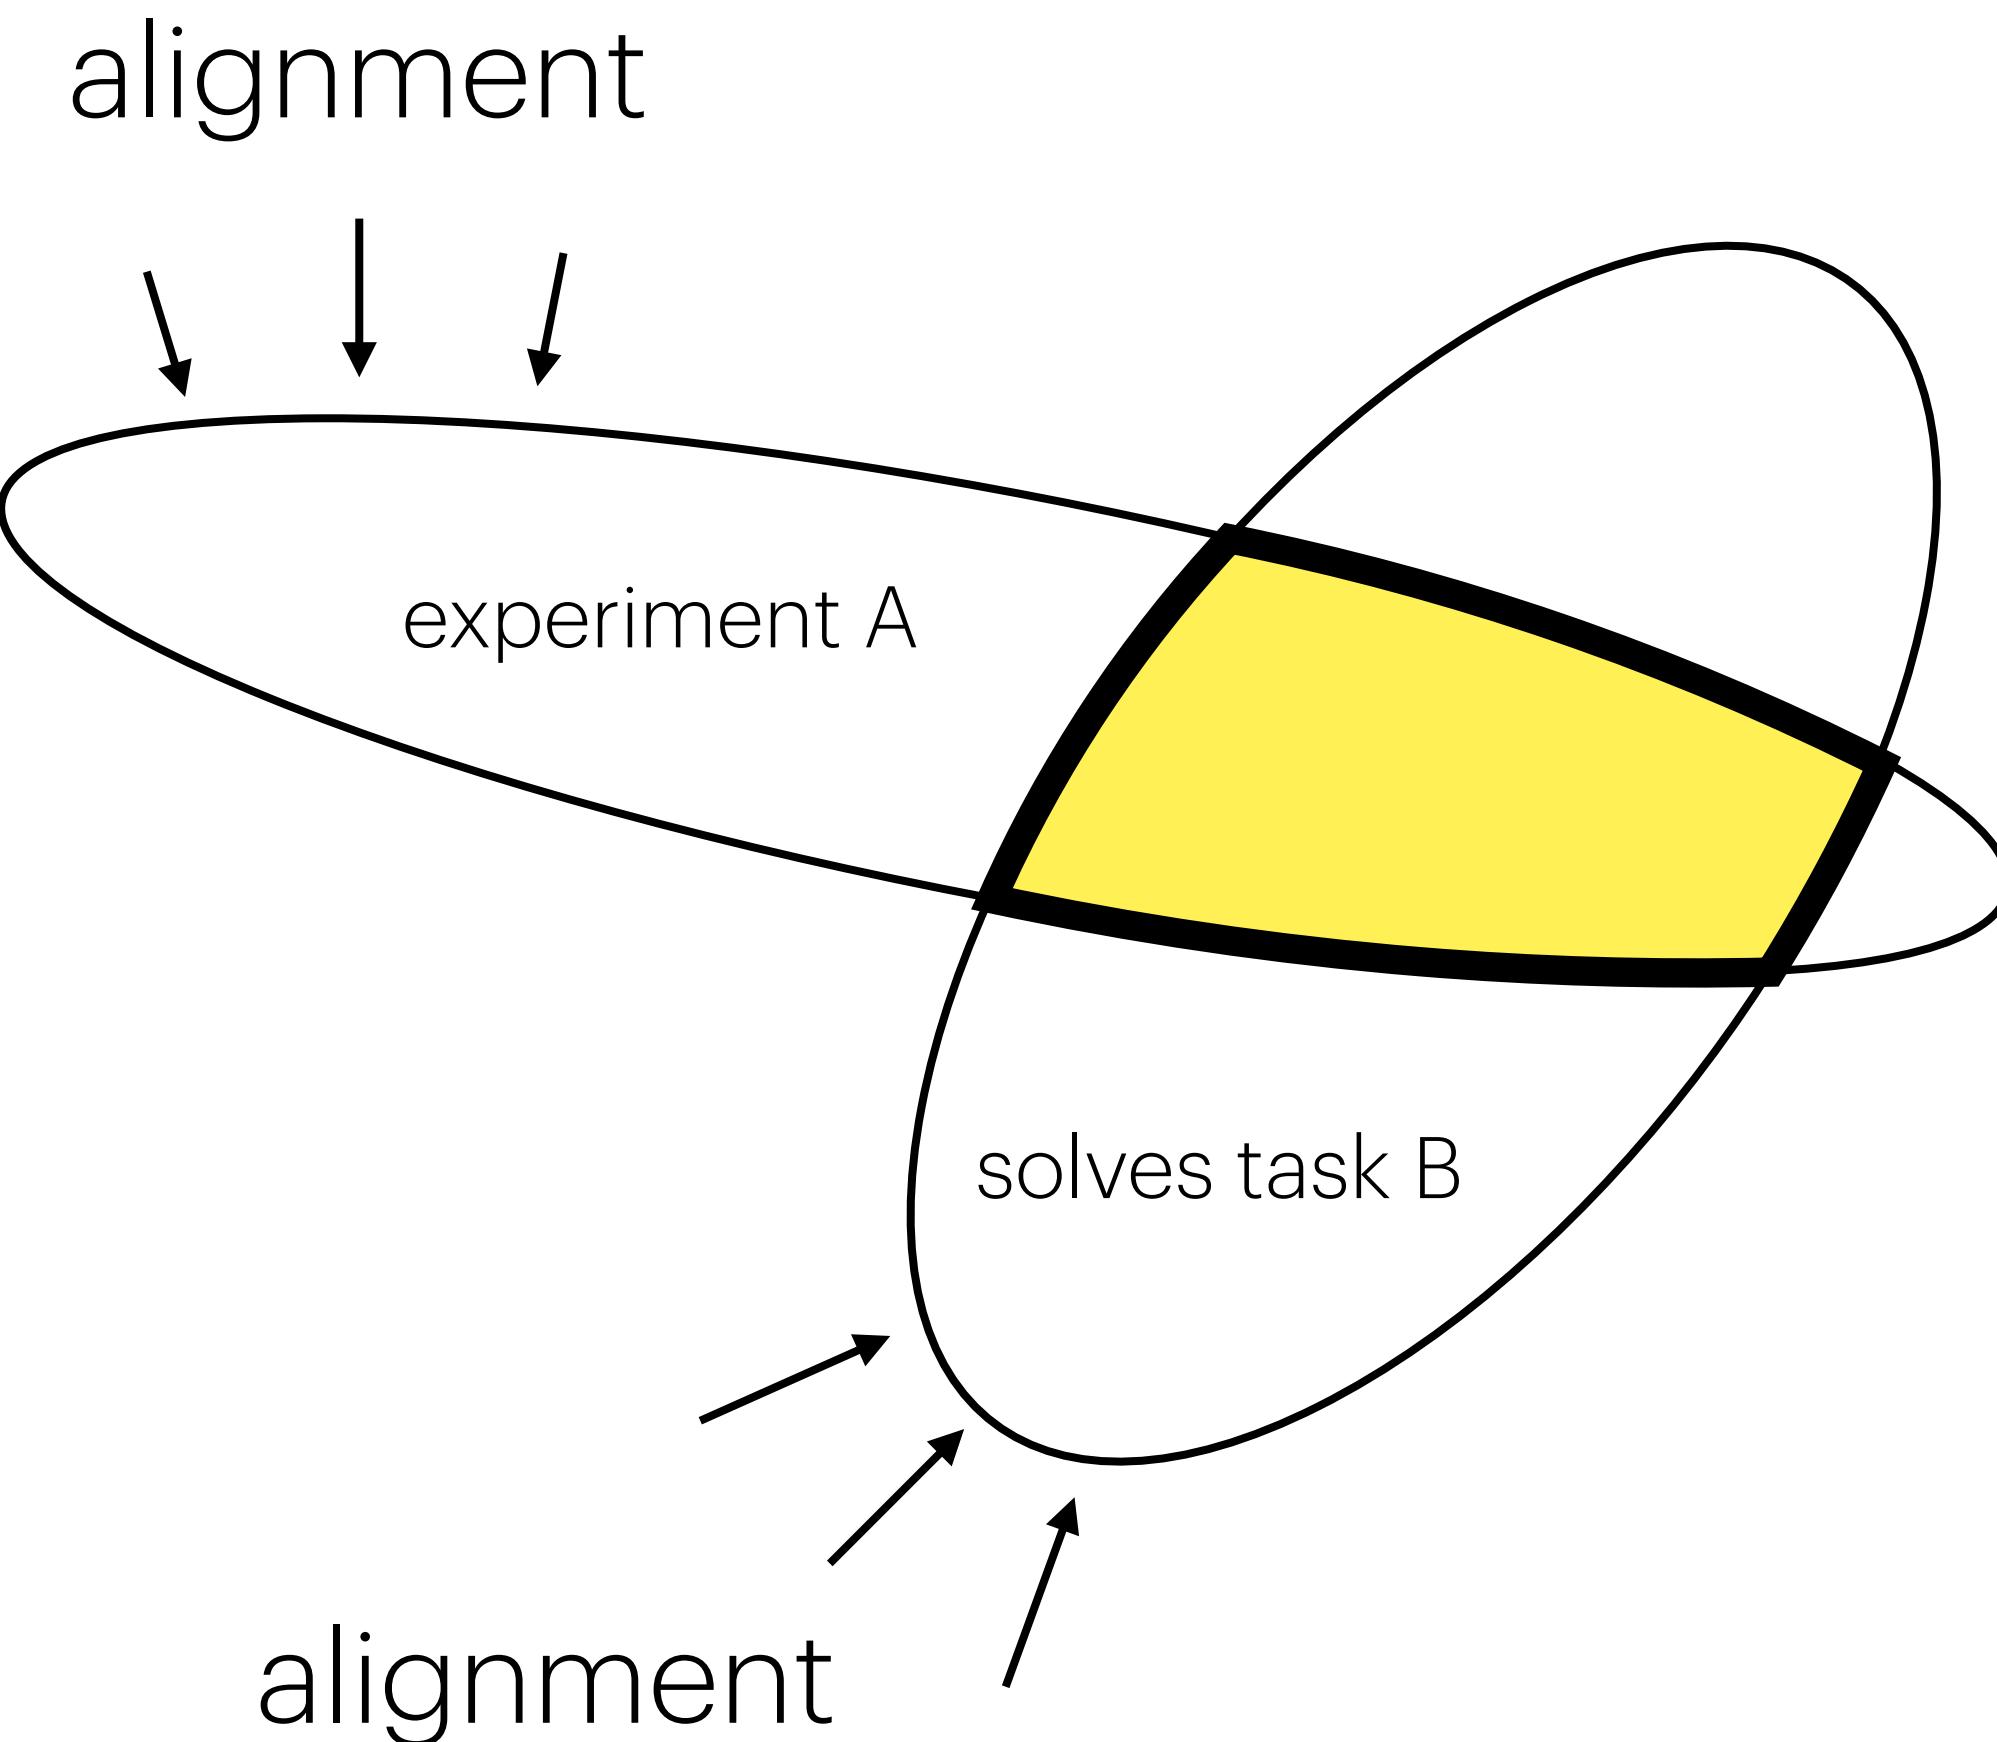
# Hypothesis Space

peer review



If you can **embed it**, you can **align it**!

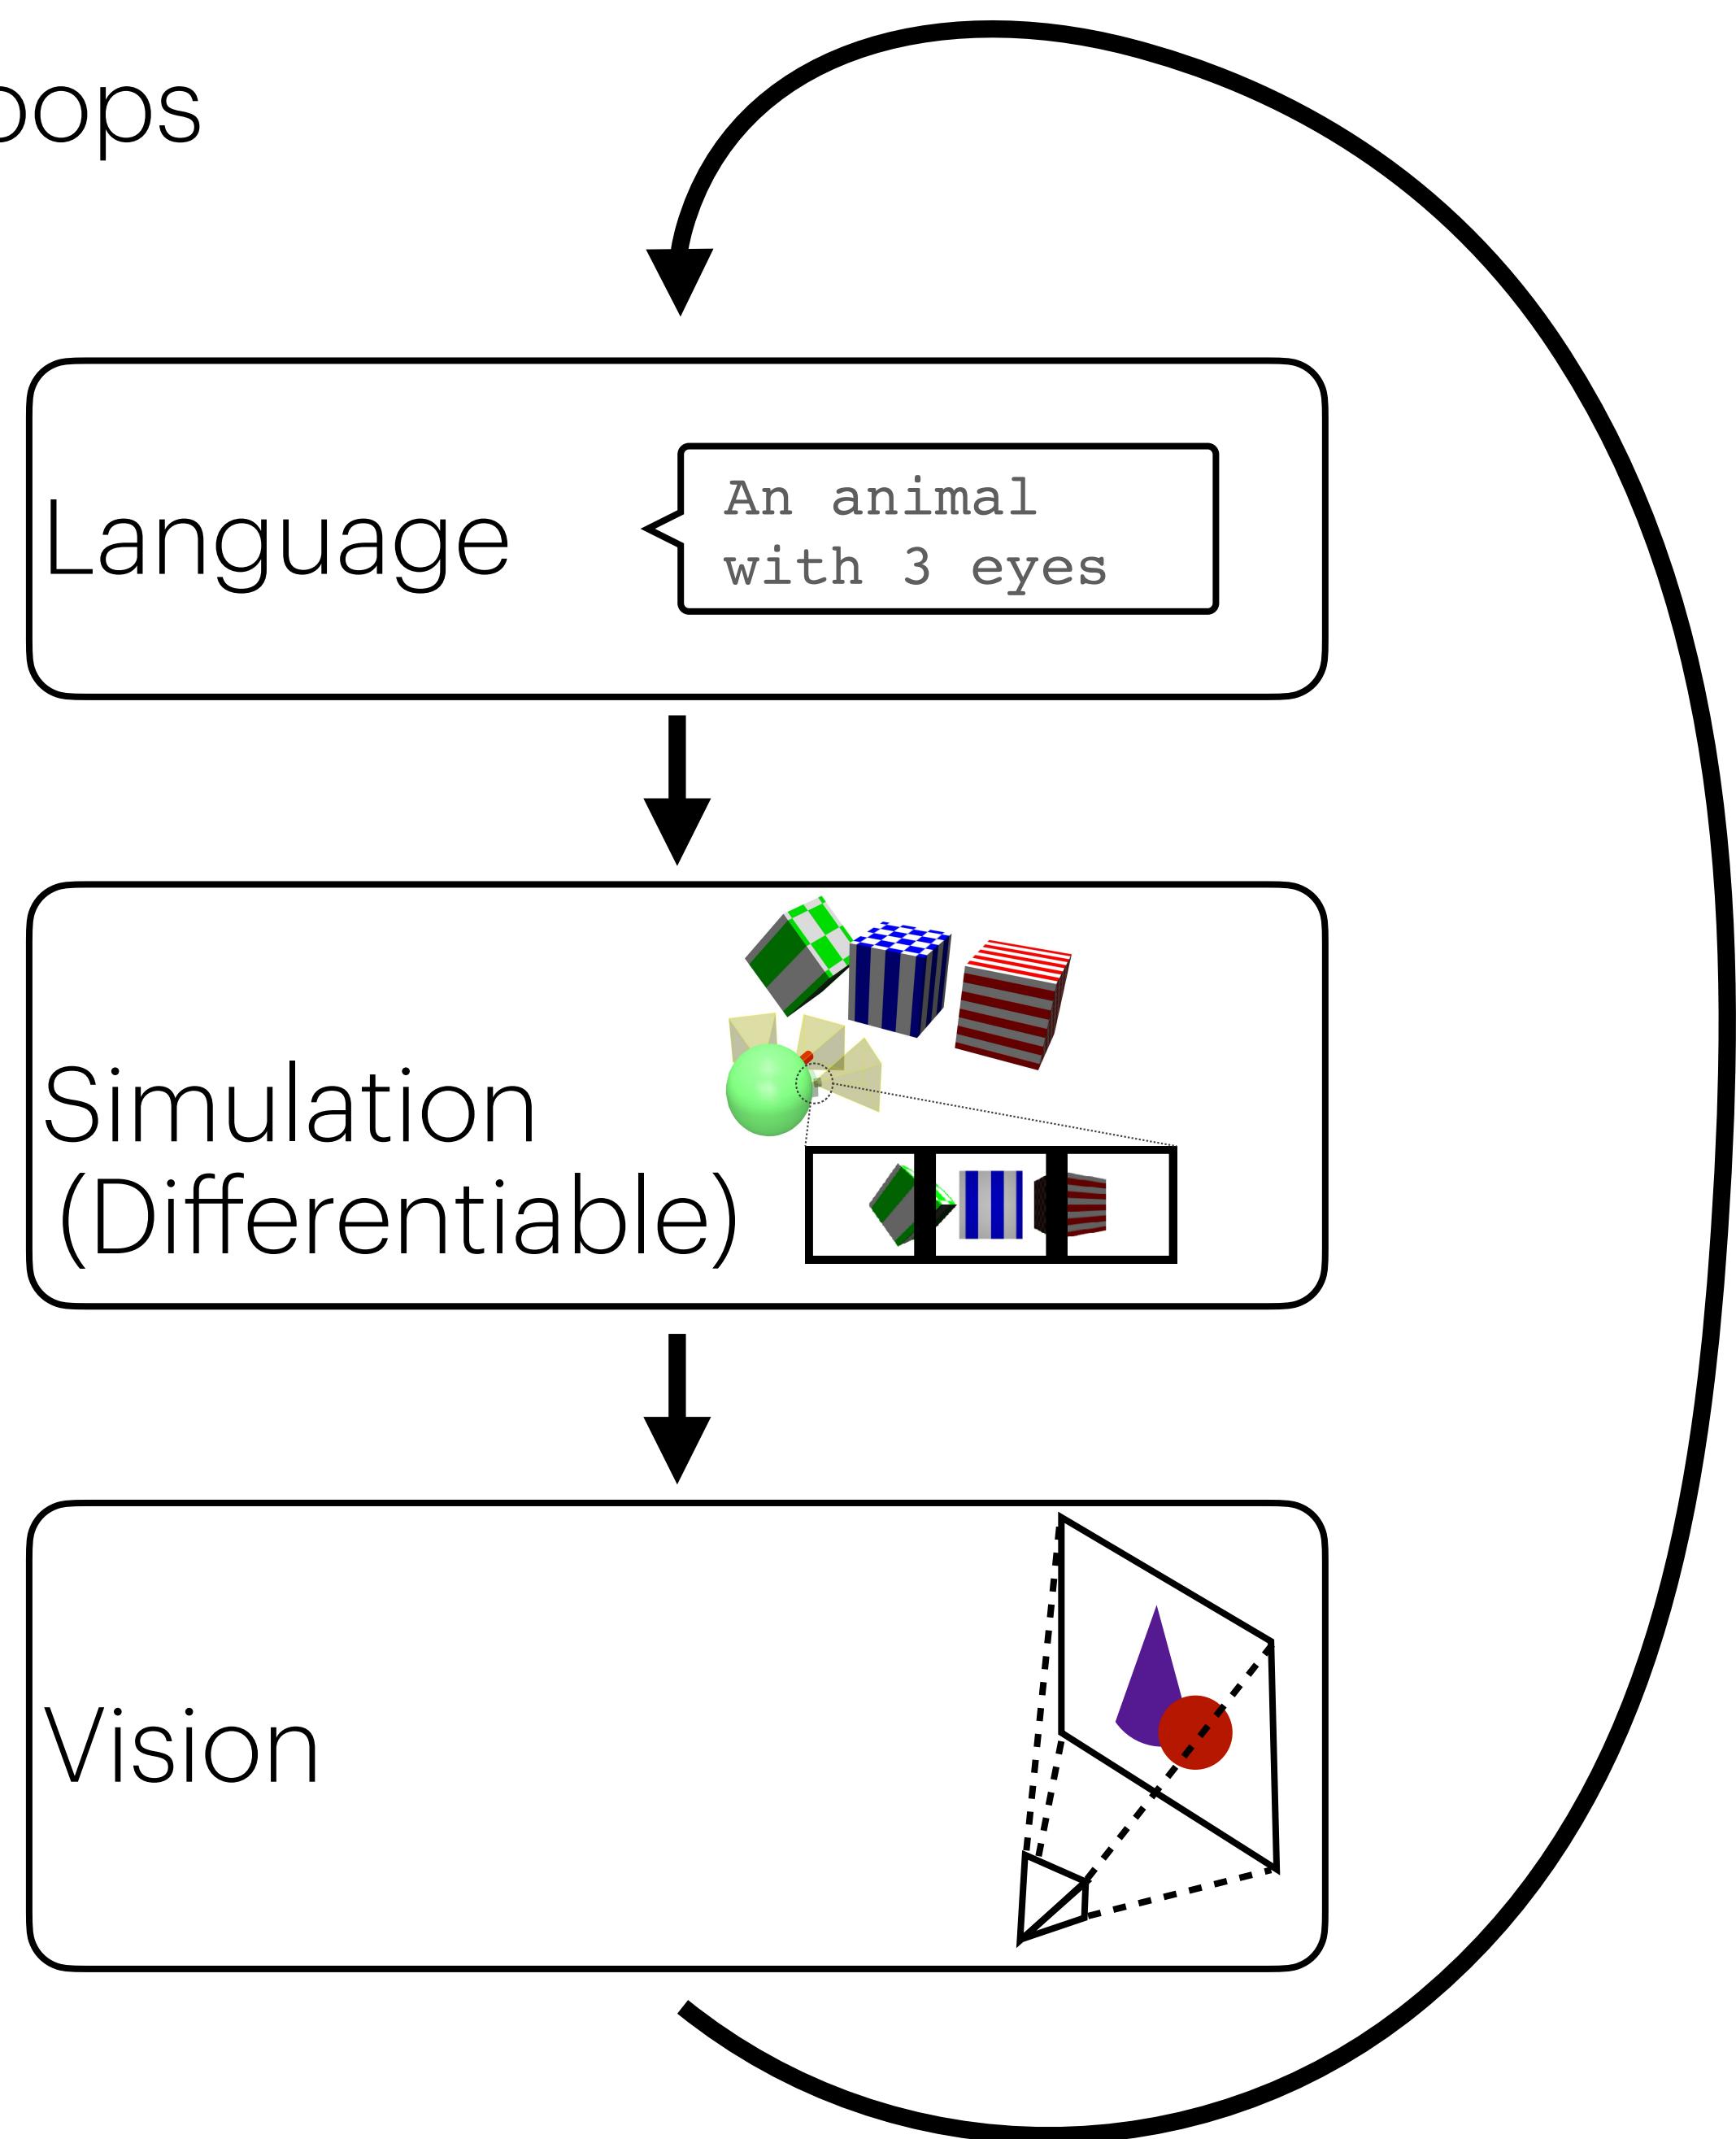
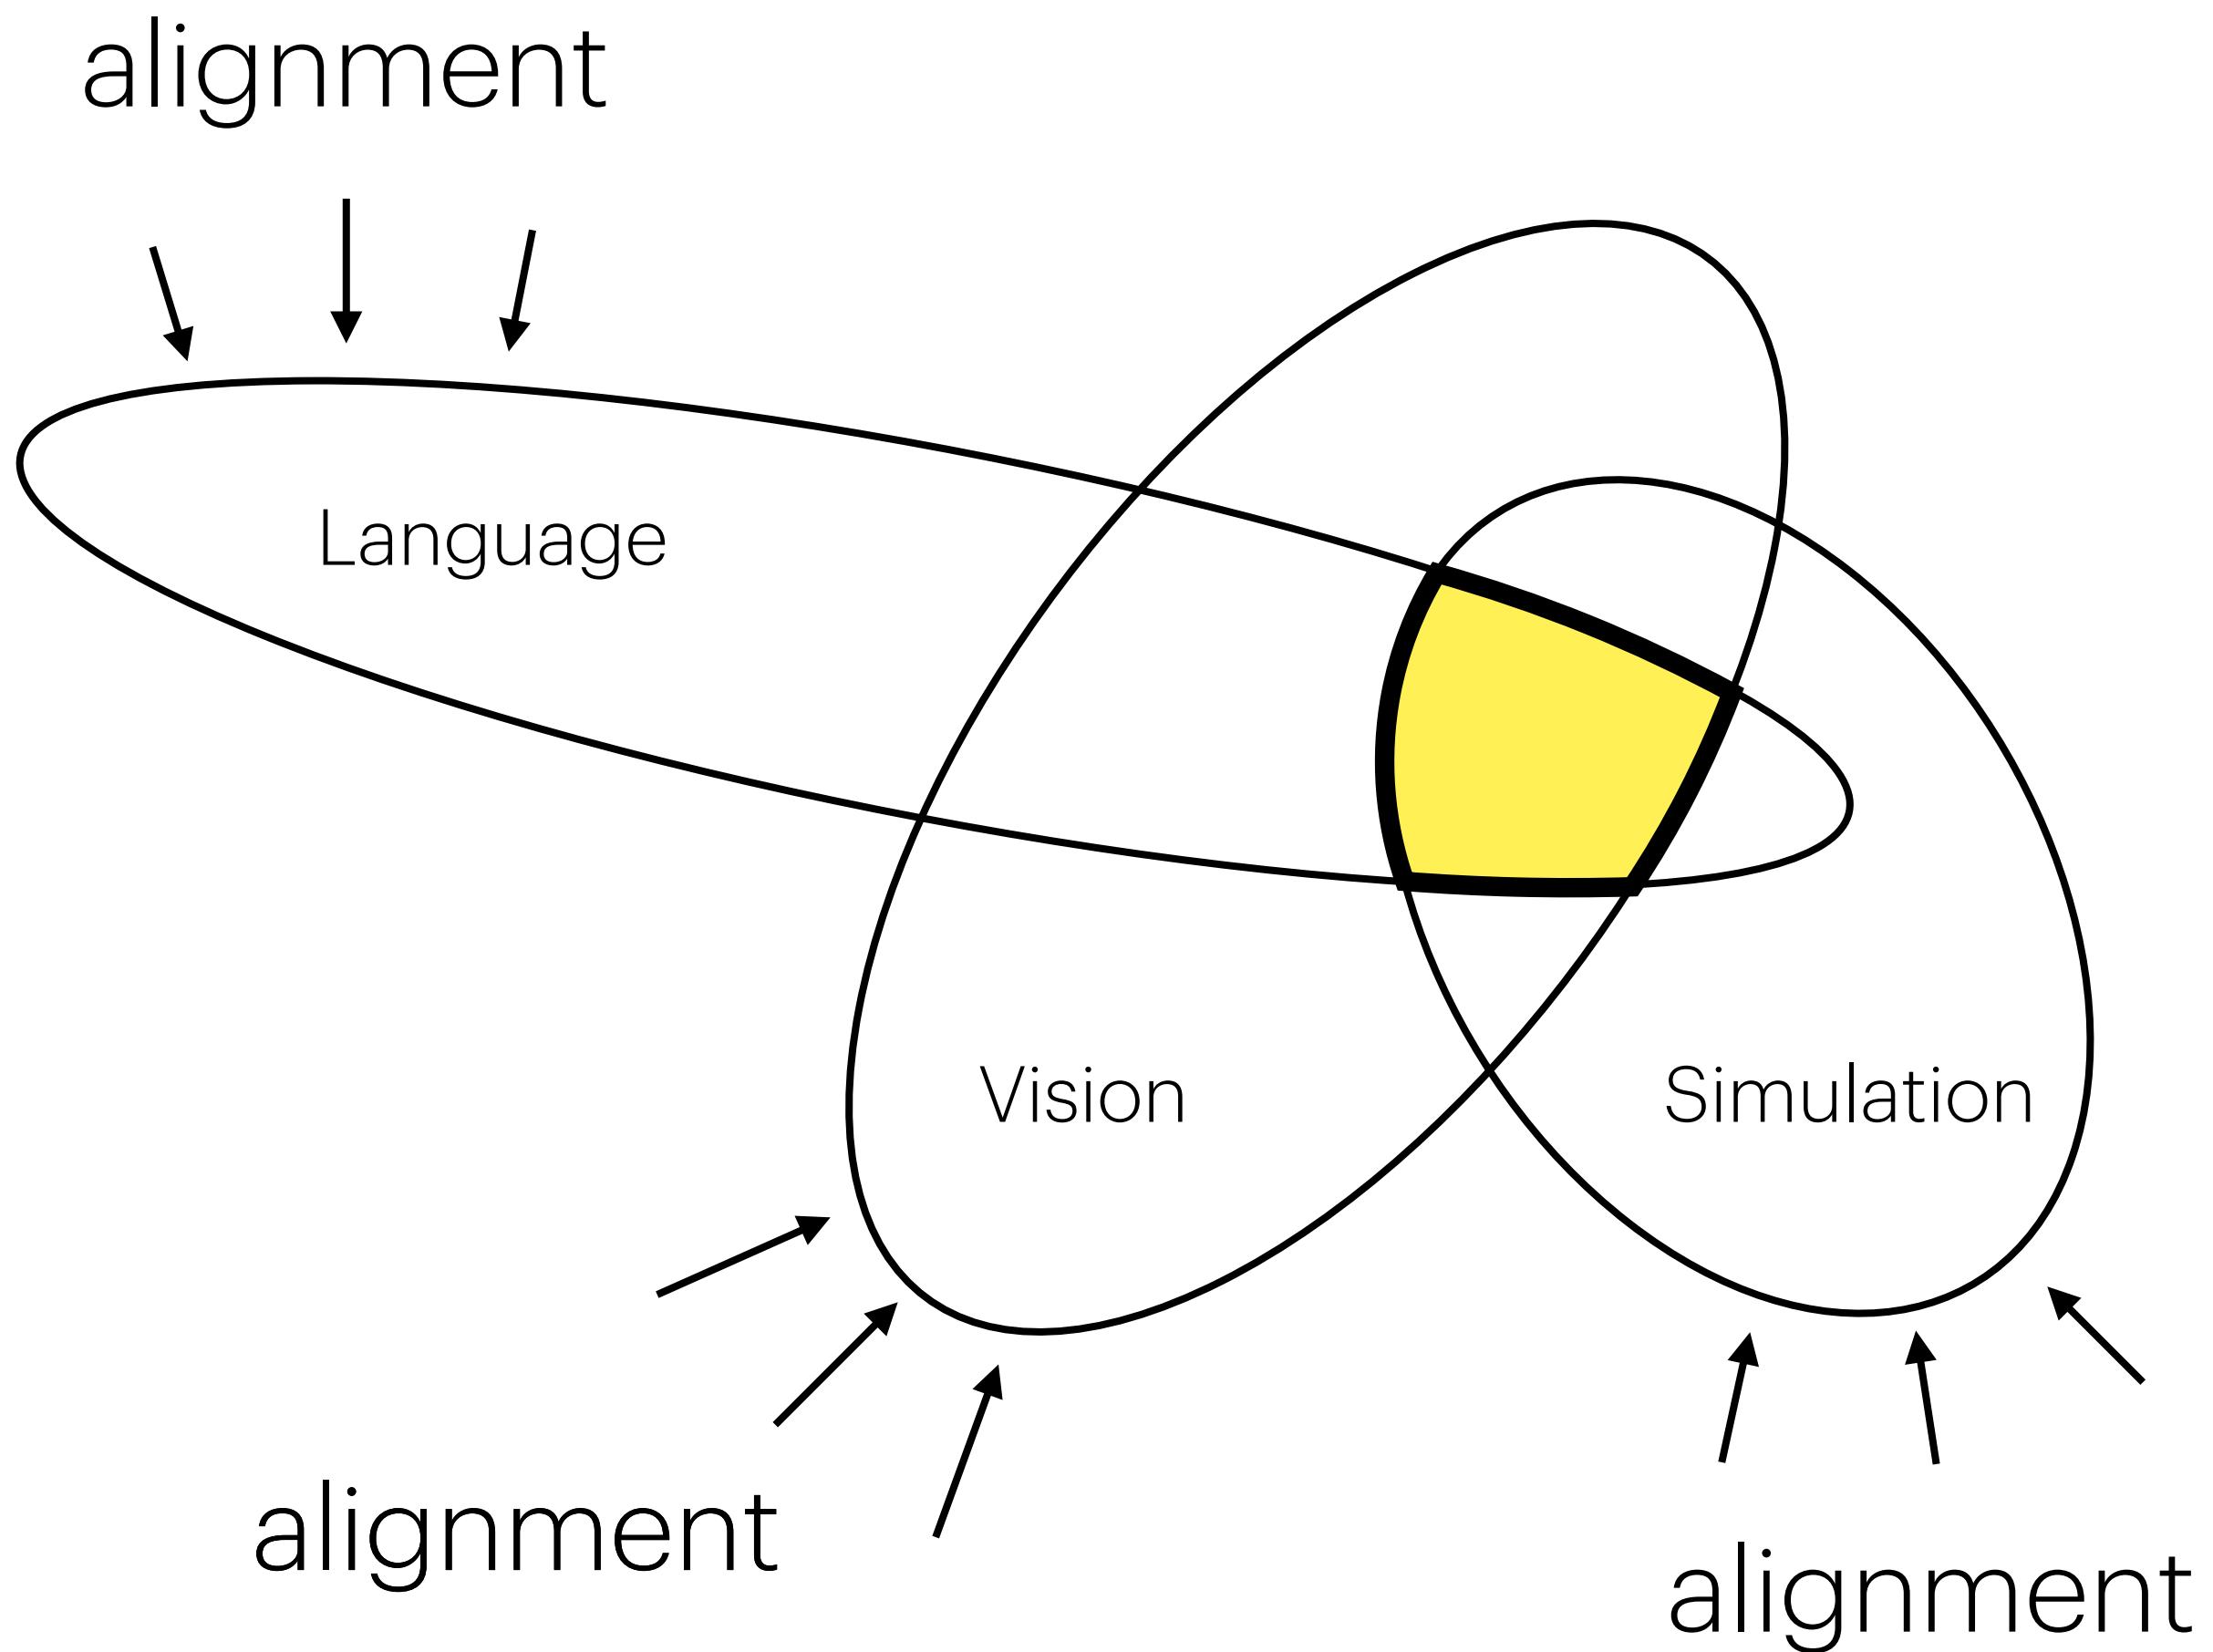
$$f : \mathcal{X} \rightarrow \mathbb{R}^n$$



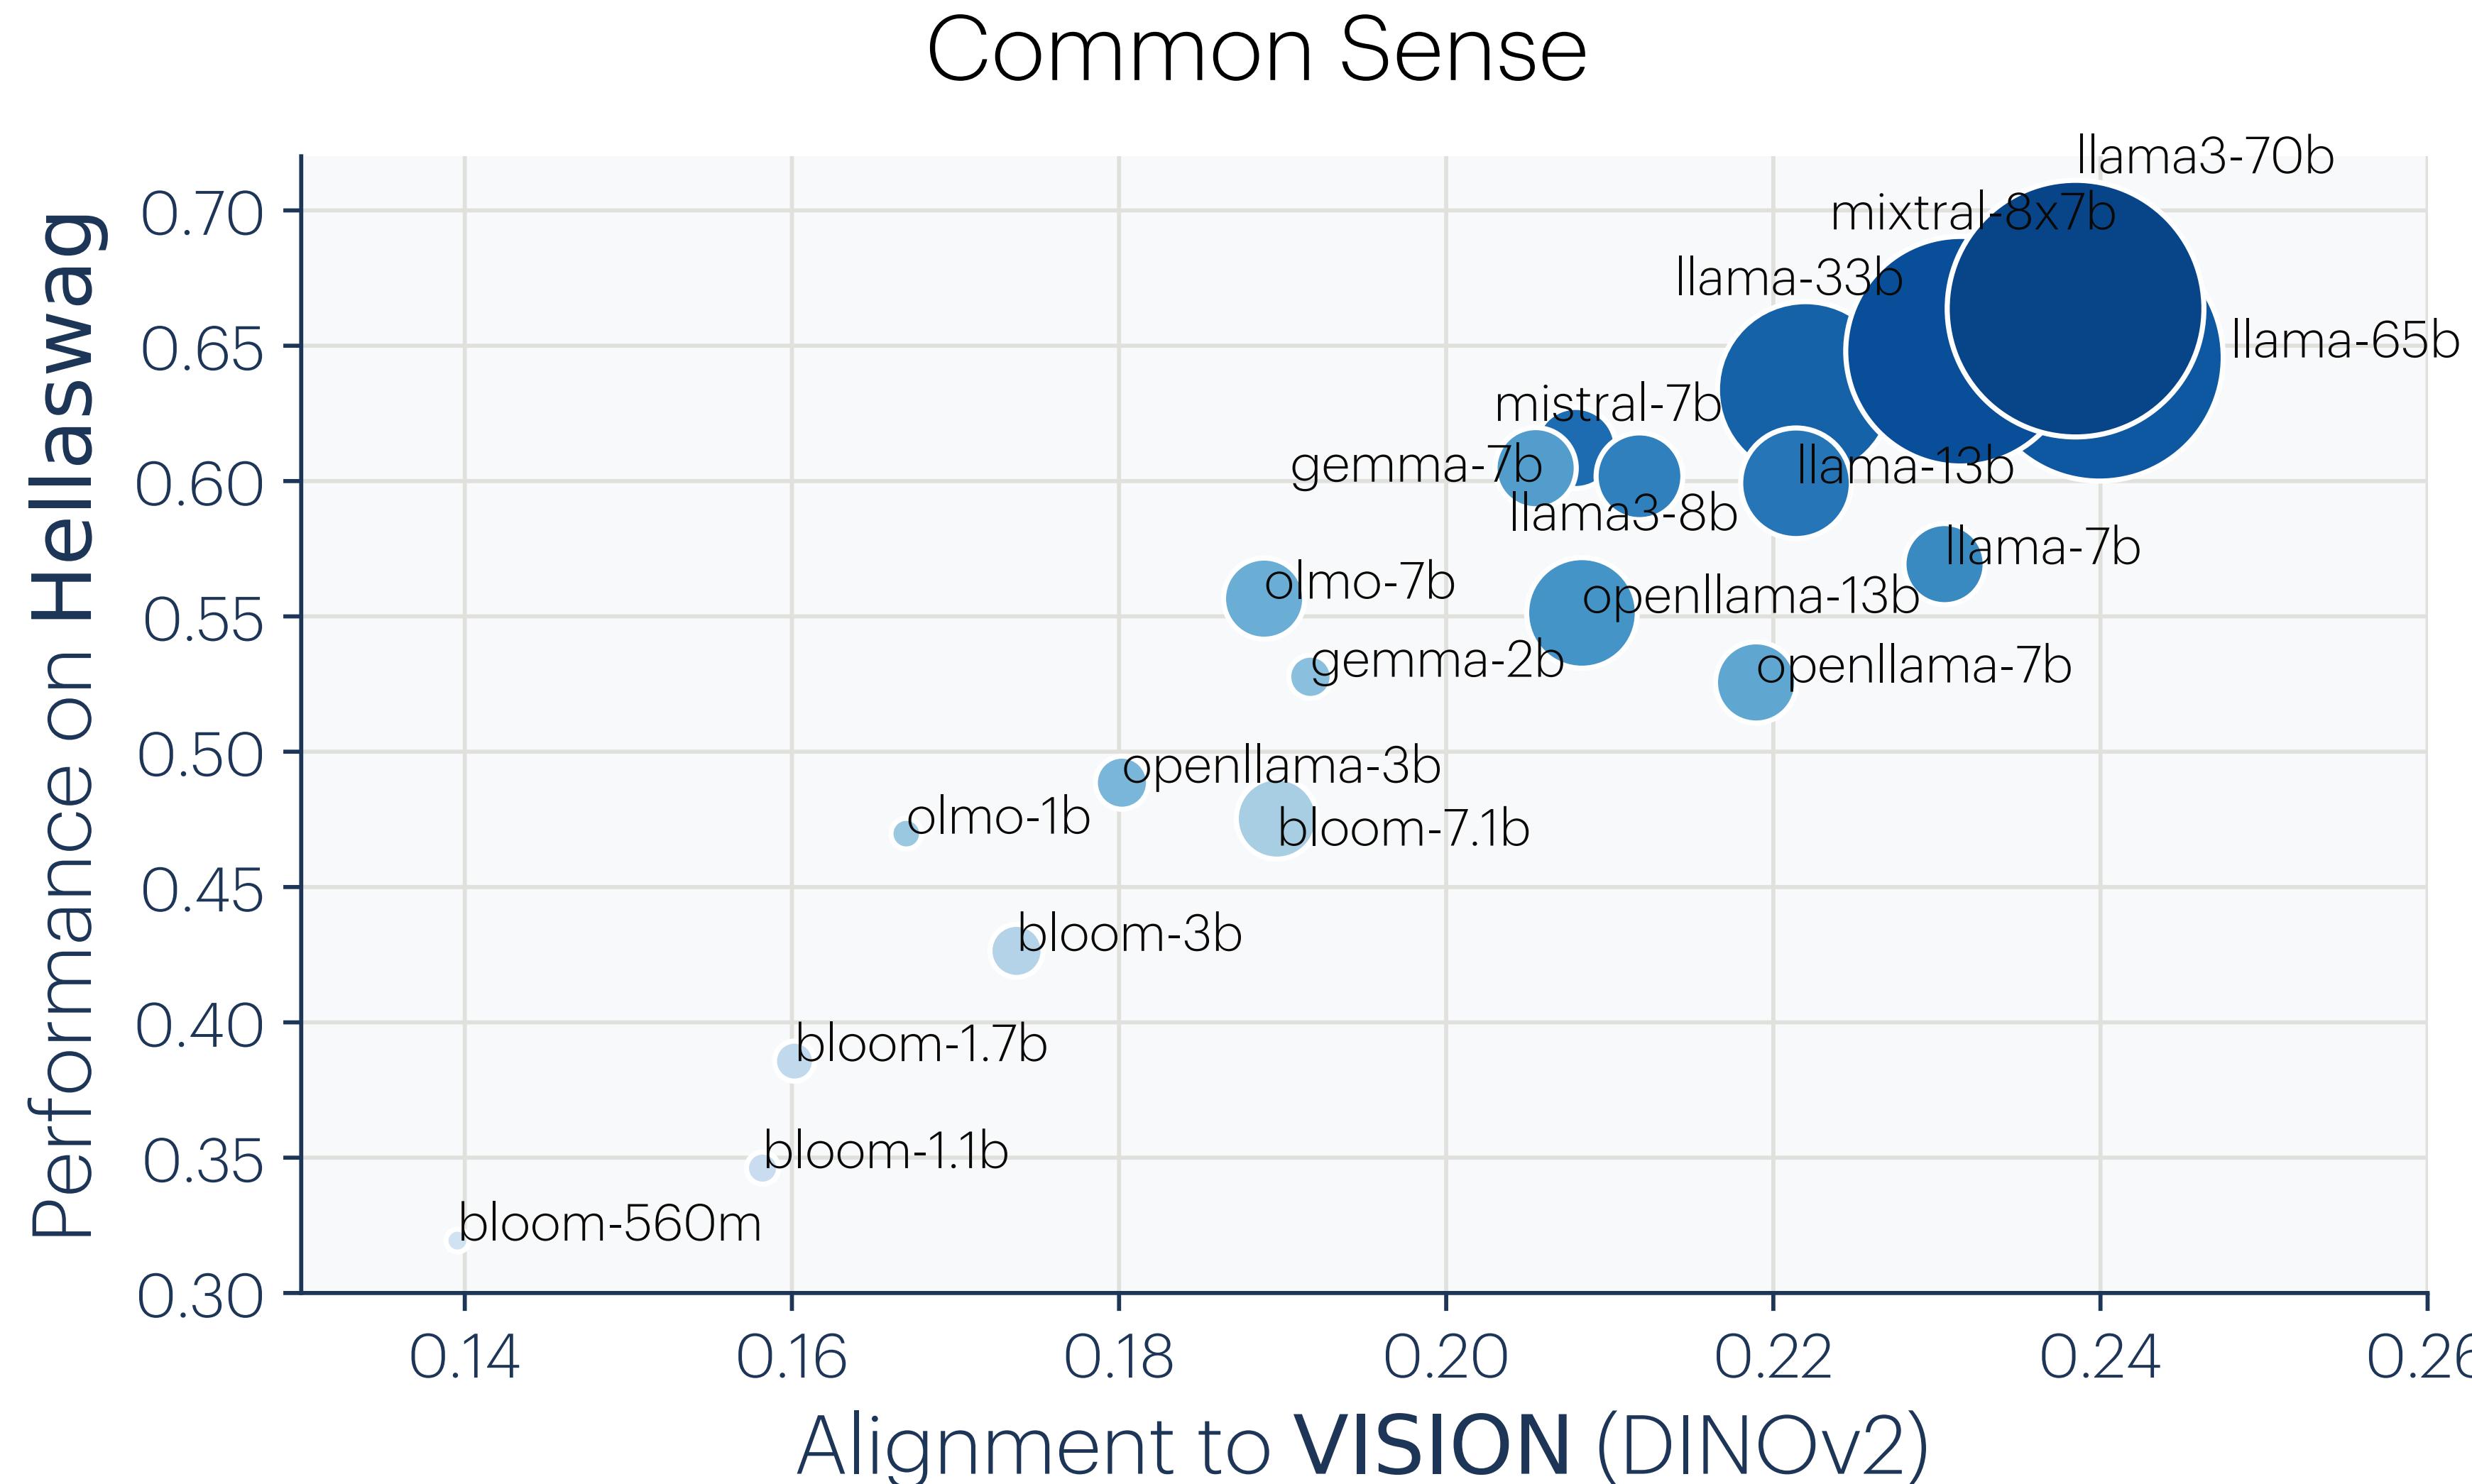
- Experiments
- Modalities
- Neural network representations

become **interchangeable** and **alignable**.

# Alignment Loops

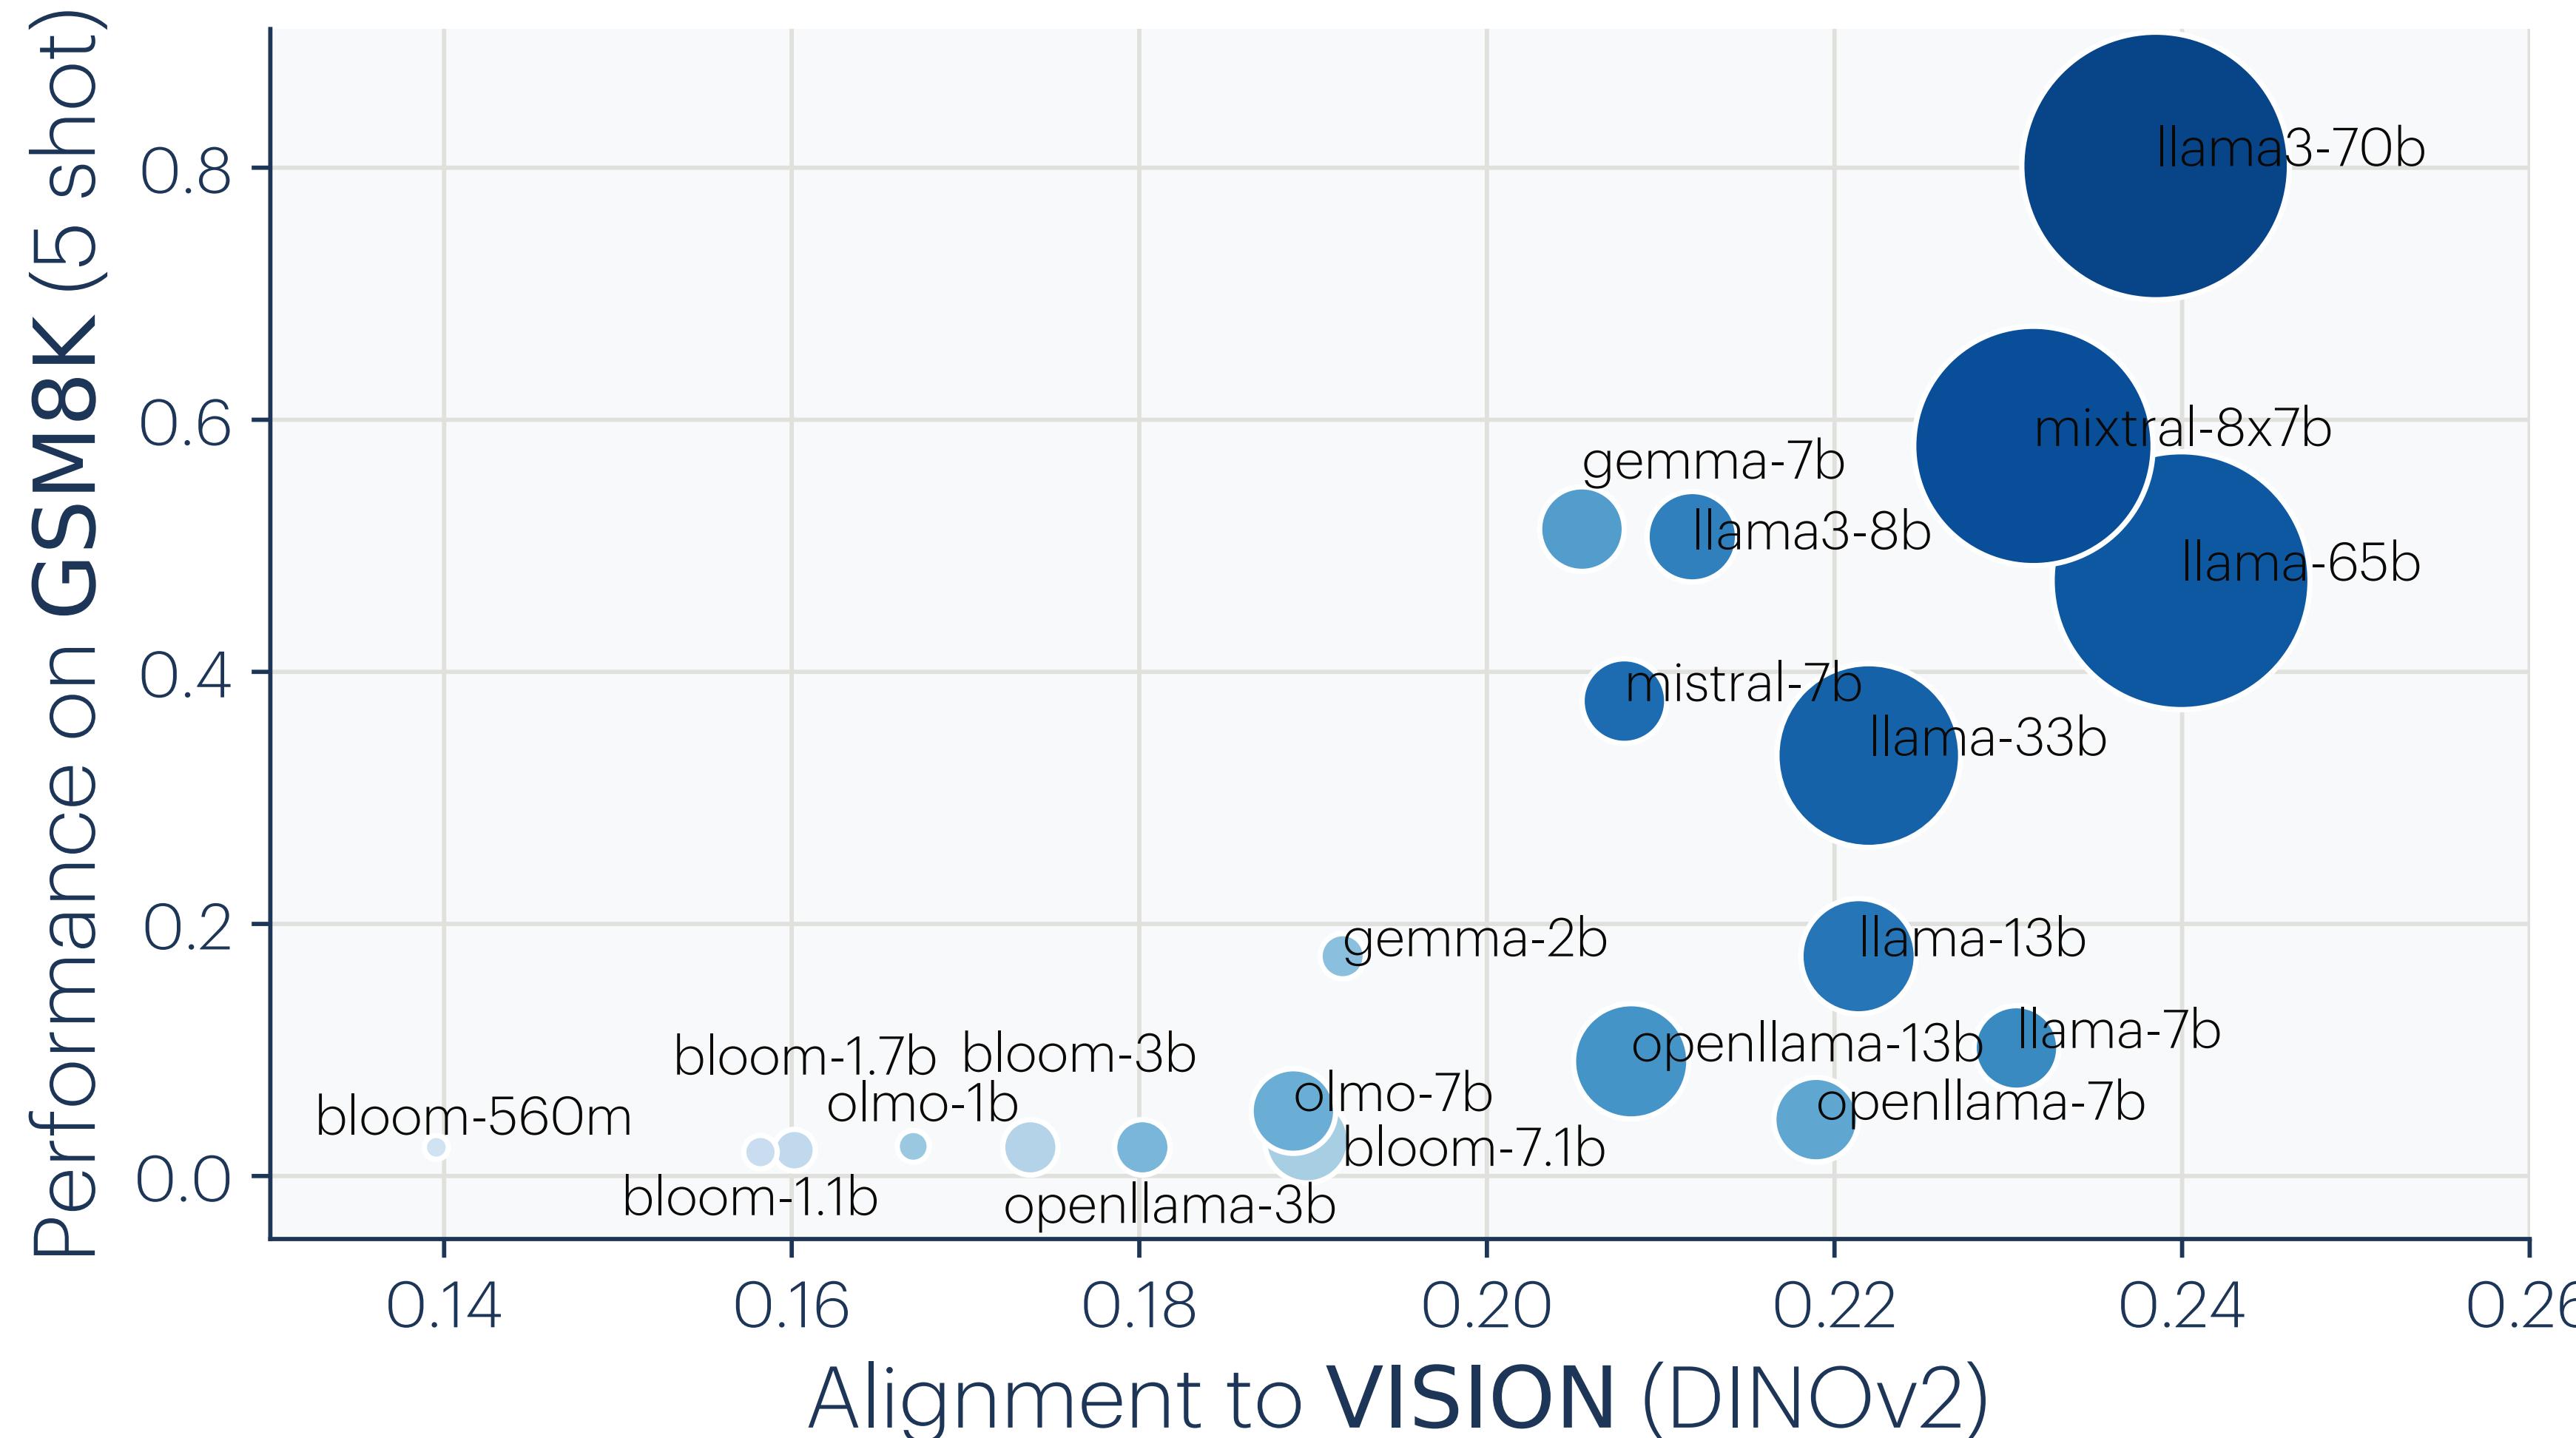


Alignment corresponds to better performance in **self consistent** domains

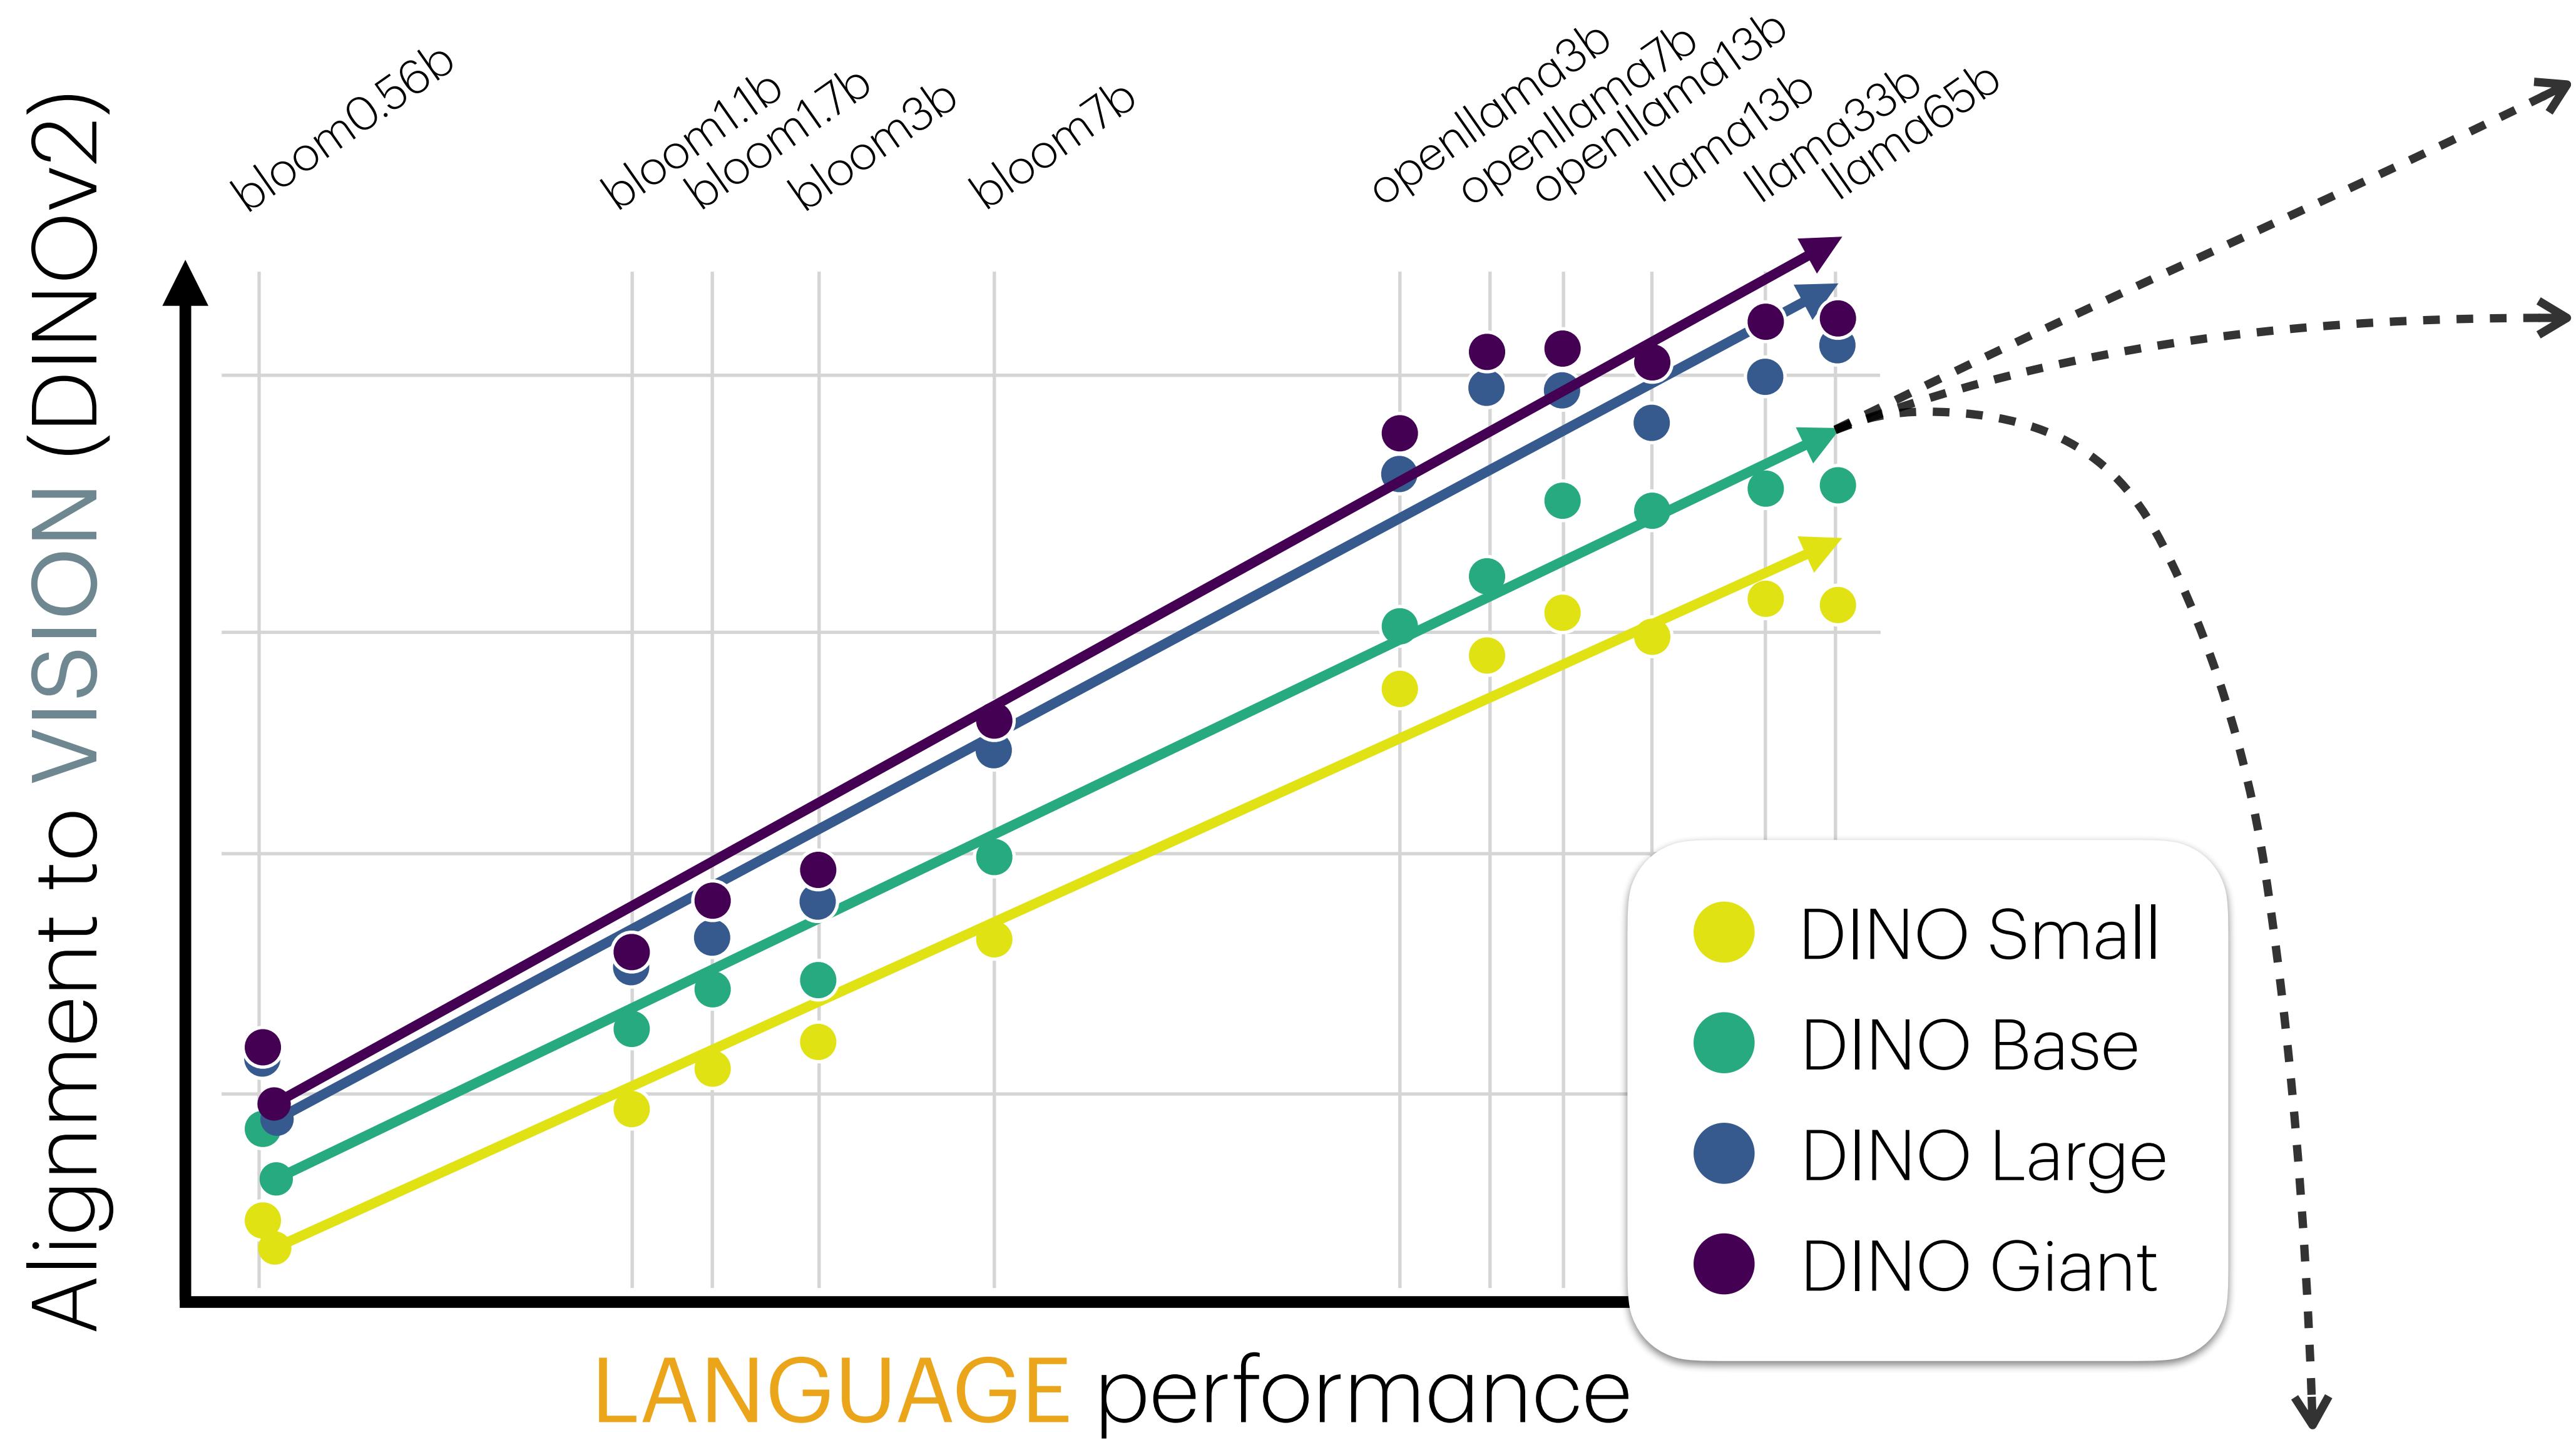


Alignment corresponds to better performance in **self consistent** domains

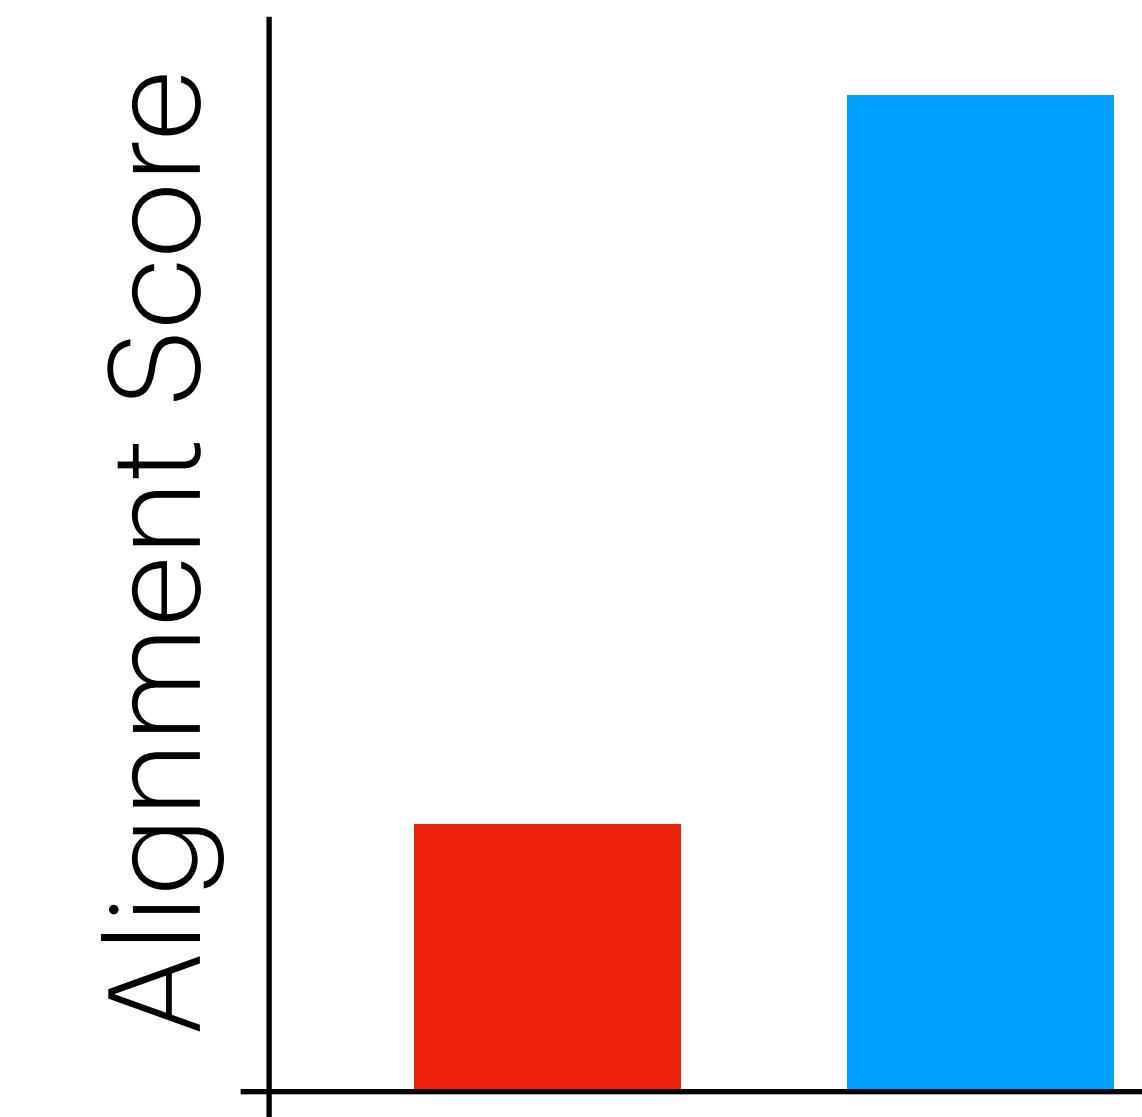
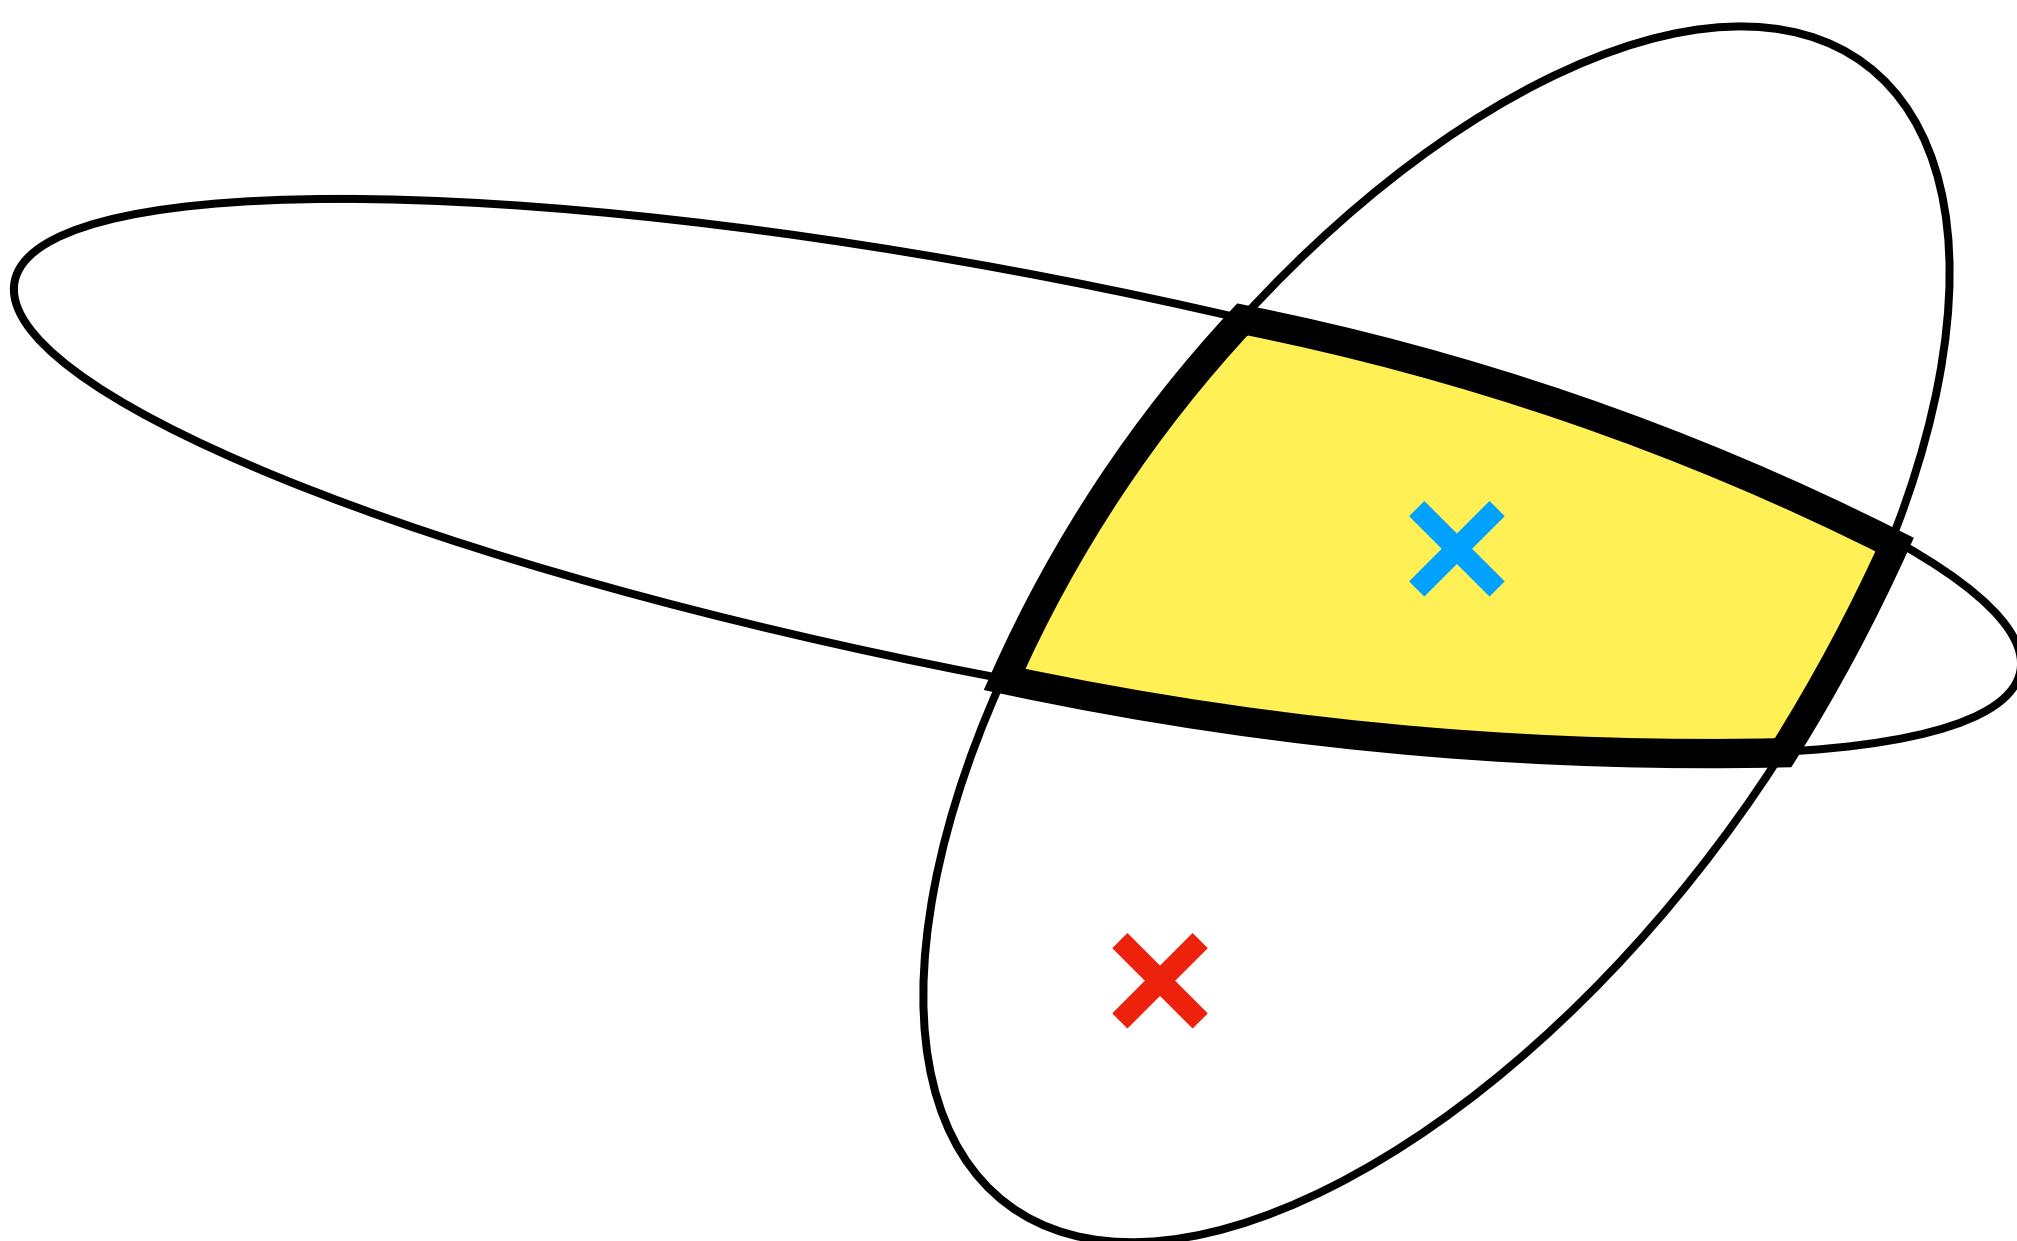
Math



# We can choose where this goes



# Alignment measures when hypotheses are **not consistent**



# Math reasoning is an alignment loop

Please add a single pair of parentheses to the incorrect equation:  
 $2 * 1 - 5 + 4 + 4 + 3 + 2 + 2 = 22$   
to make the equation true.

$f_{text}$



$2 * 1 - 5 + (4 + 4 + 3 + 2 + 2) = 12$

$f_{text}$



$2 * (1 - 5 + 4 + 4 + 3 + 2 + 2) = 22$

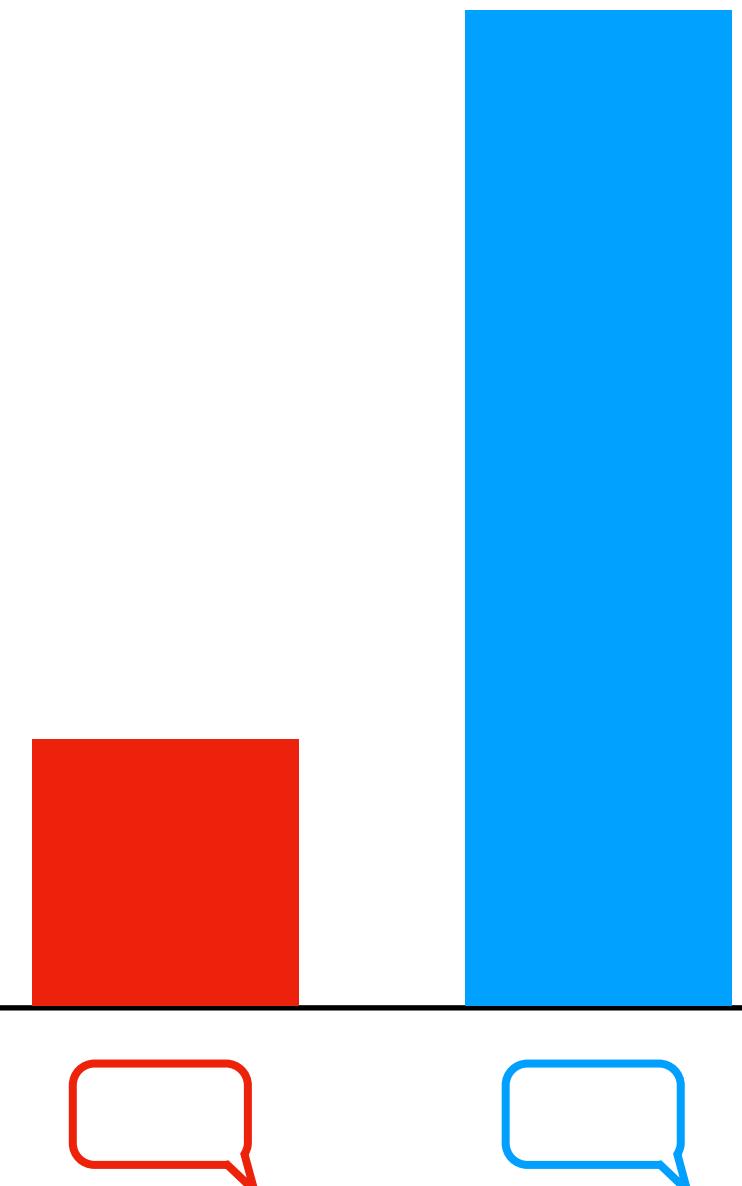
$f_{text}$



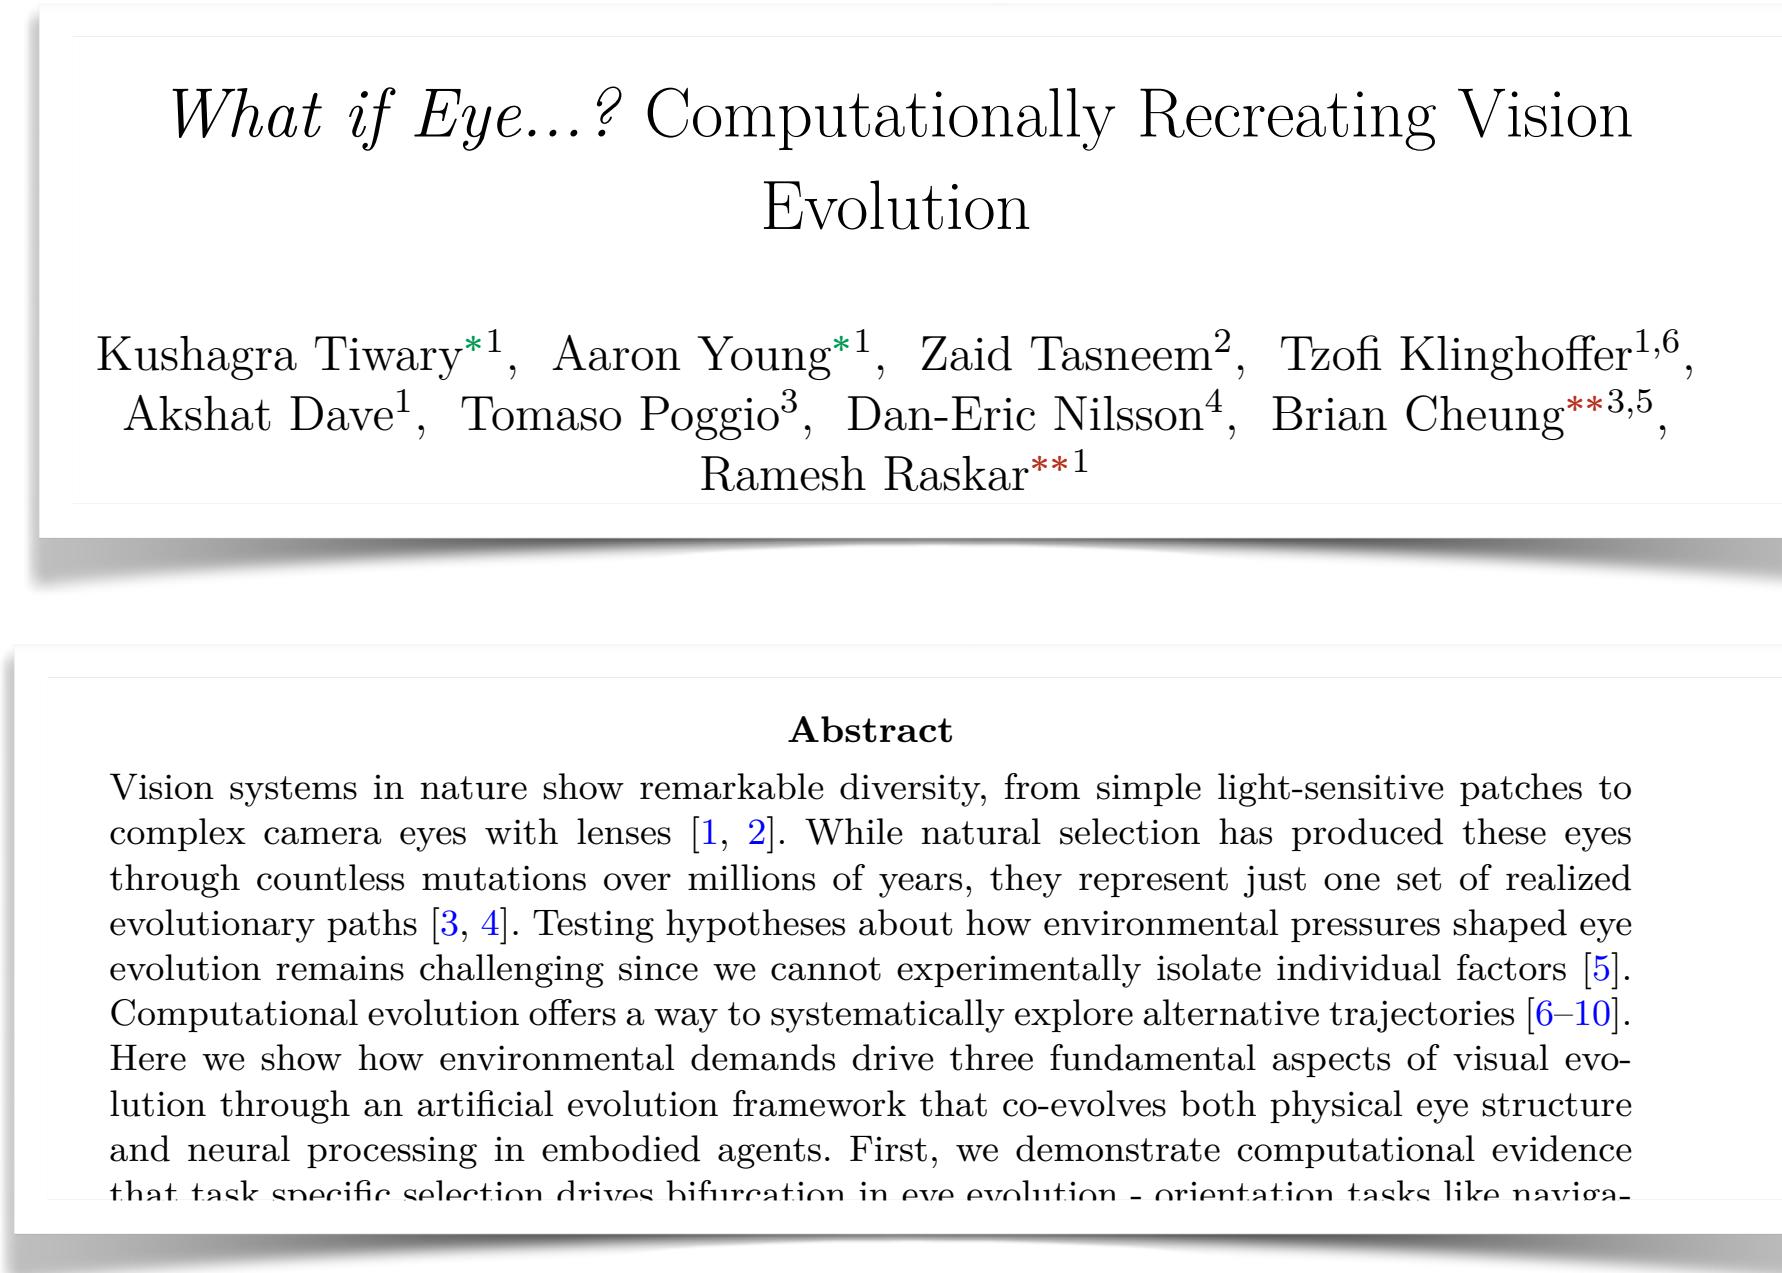
Hypothesis

Experiments

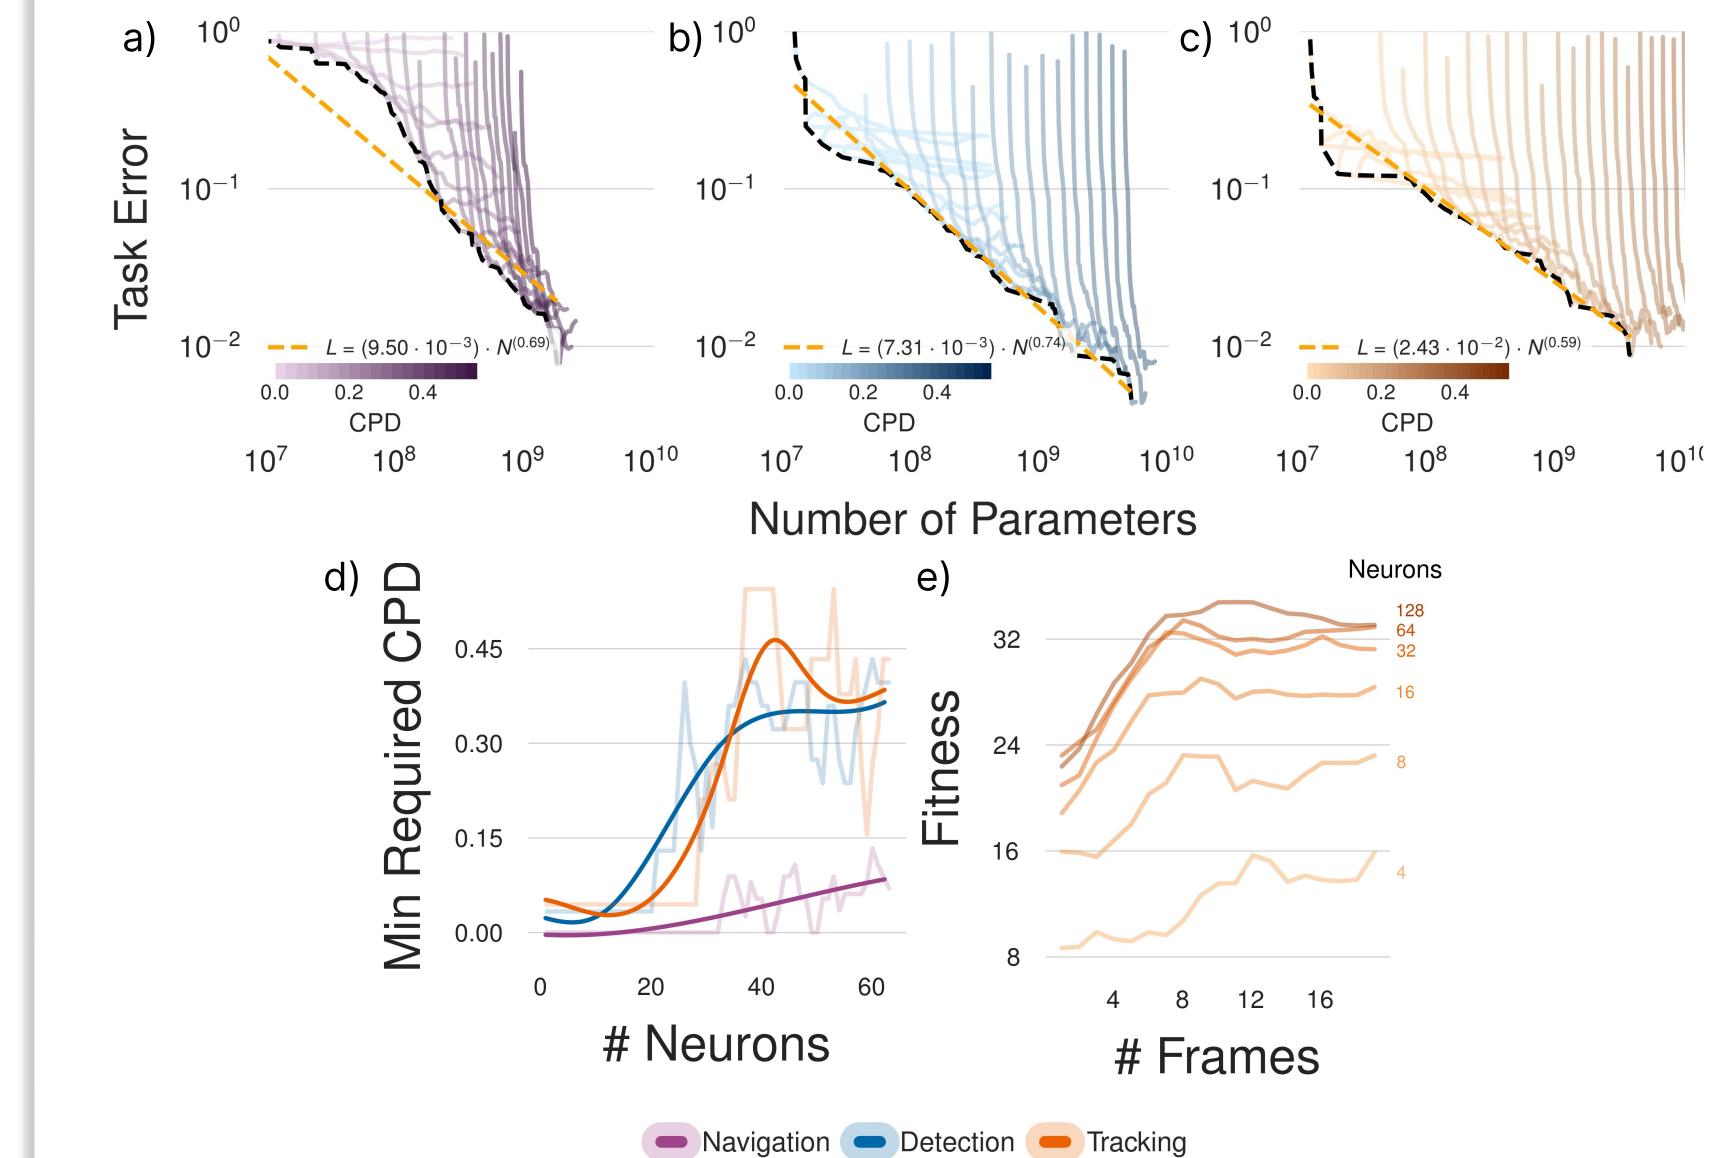
Alignment Score



# Representation Alignment for 'soft' verification



## 2.6 What if animal brains stayed small throughout evolution?



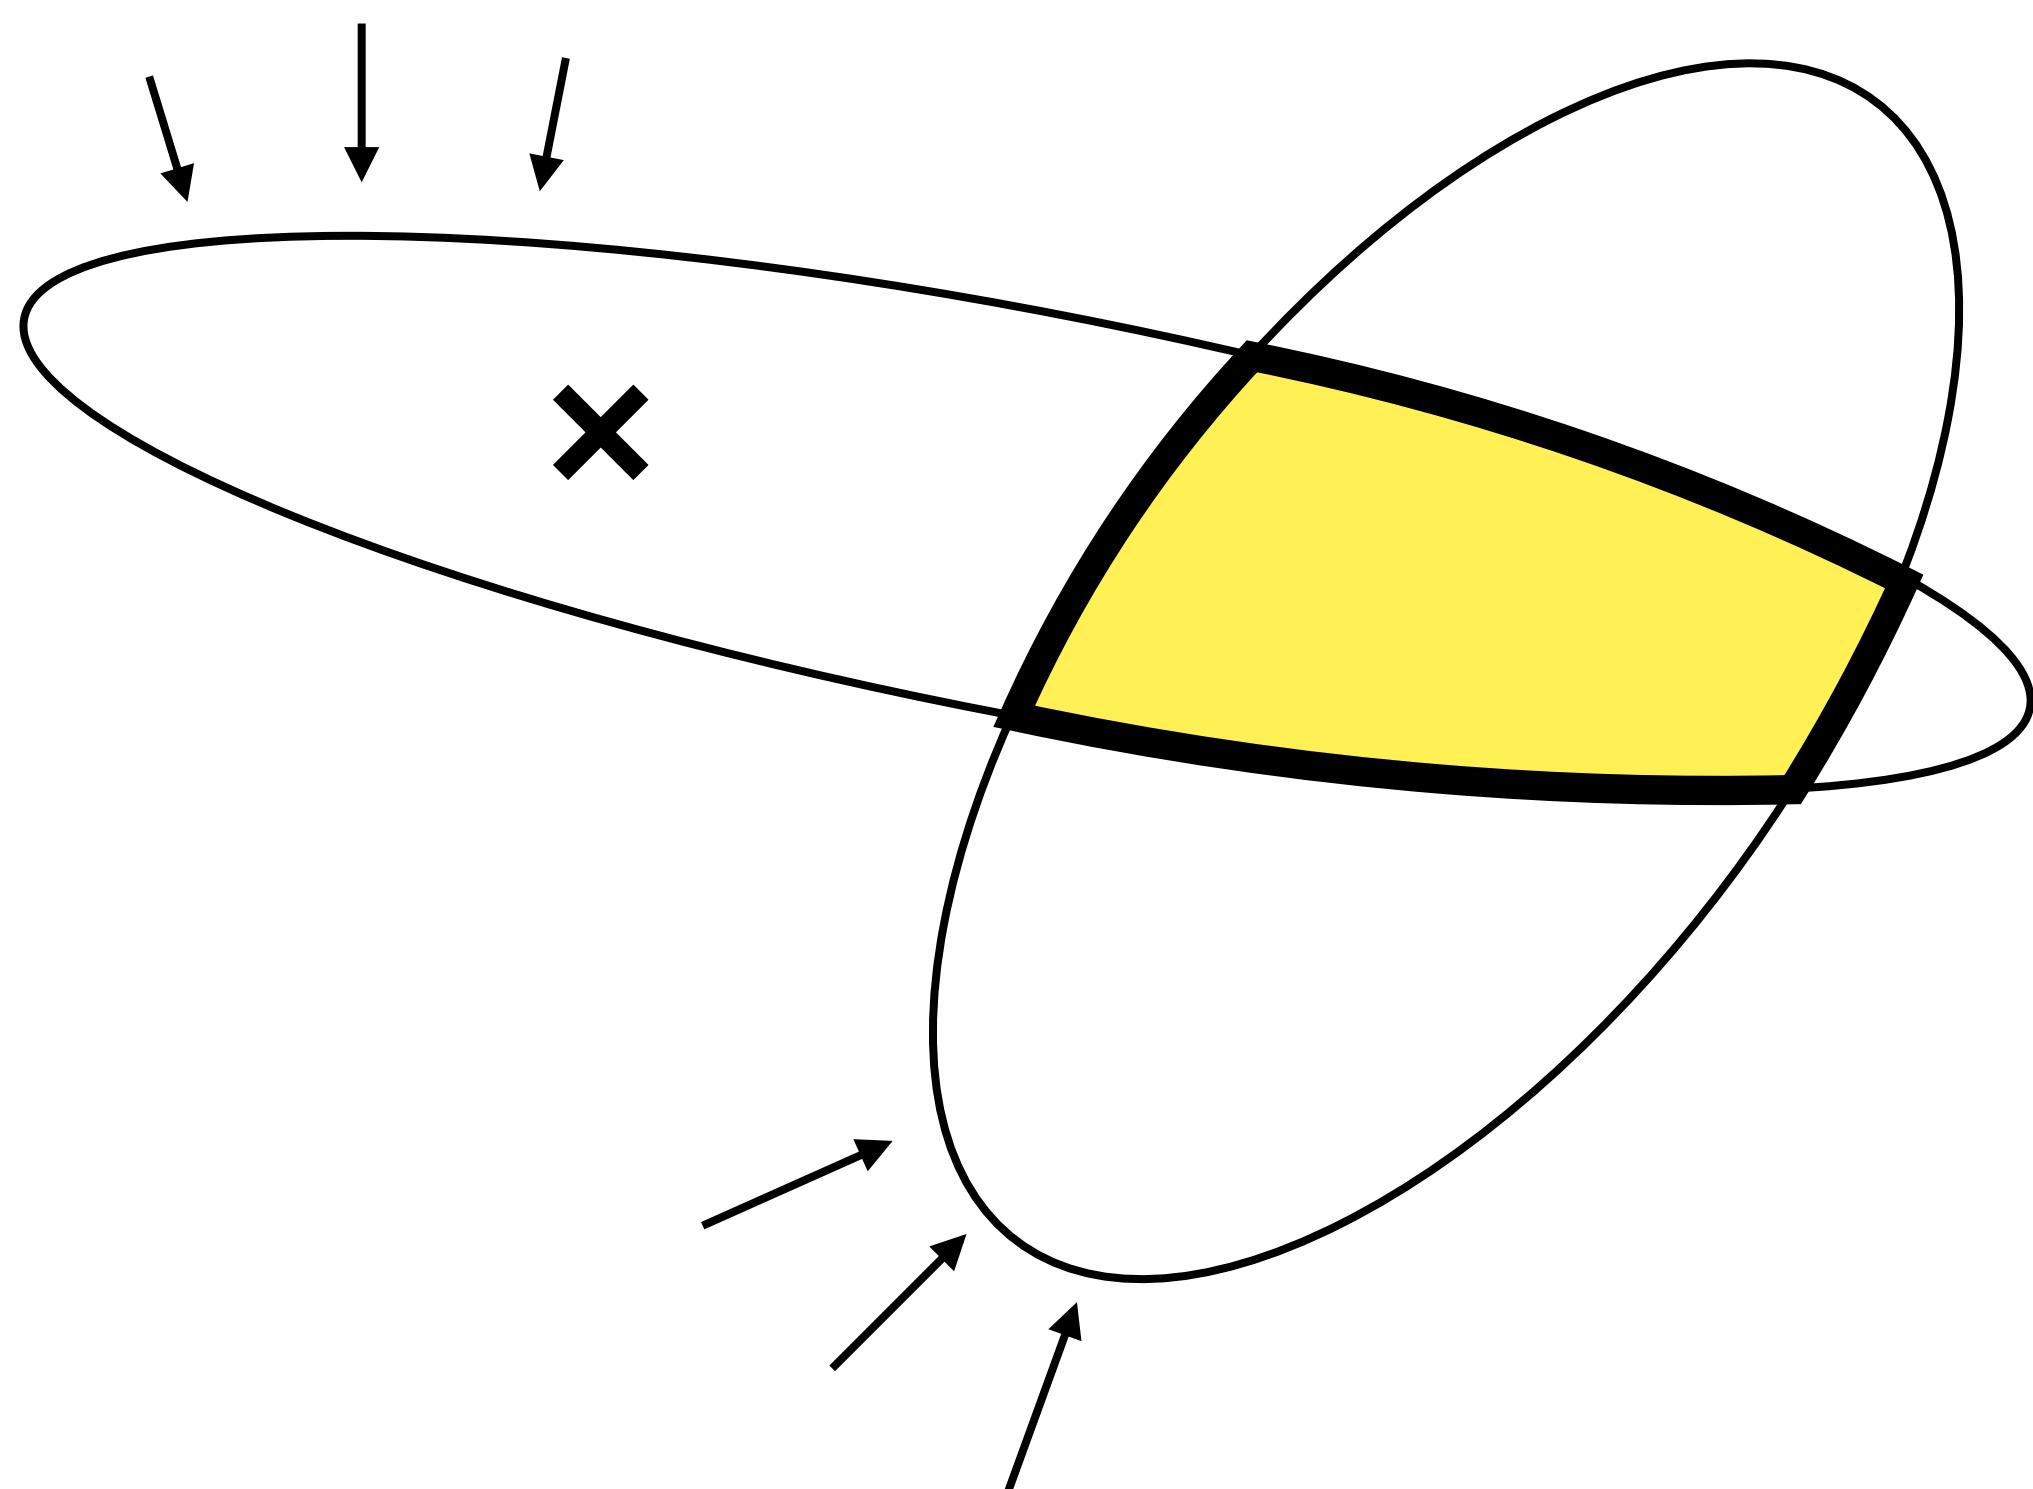
$f_{text}$

Similarity is **easier** to verify than exact matching

Alignment Enables:

AI for Science

Science for AI





You can ask questions

cheungb@mit.edu

**UCSF** in 2026

Ask me about PhD and Post-Docs!