What is cognitive science?

PSYCH-UA.46

What is this class about?

Shorter version

Experimental approaches to understanding the structure of human thought.

What is this class about?

Longer version

An introductory course on the use of various behavioral measures (accuracy, reaction time, etc...) to understand the structure of the human mind. Our goal is to use experiments to test alternative theories of cognitive function and to better understand the motivation and structure of human behavior. We will learn a basic set of skills for using computers to run experiments, collect data, analyze it, and communicate the outcome to others.

What is human cognition?

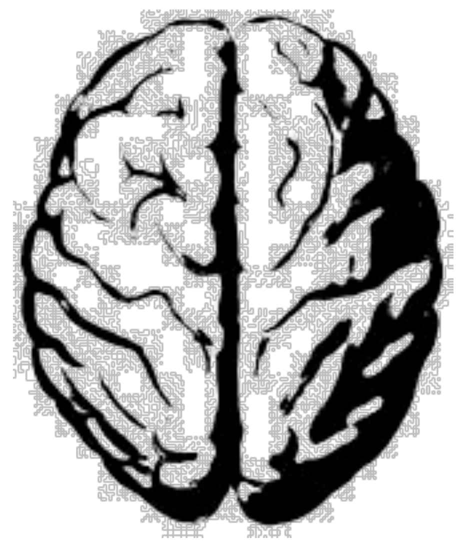

The study of how the mind (or mind/brain) works. The study of how people think, learn, and solve problems.

Examples

- How do people learn effective behavior through interaction with their environment?

- What are the stages of information processing the mind goes through to solve problems?

- How does the architecture of the mind interact with experience to determine what we know?

- How does human memory work? In what ways is human memory like computer memory? In what ways is it not?

- What is the format or "representation" of information that the mind uses?

- Can we develop theories that allow us to predict and explain human behavior? Can information we derive from these investigations enable us to build better artificial intelligence systems to solve problems?

A tale of two magnets

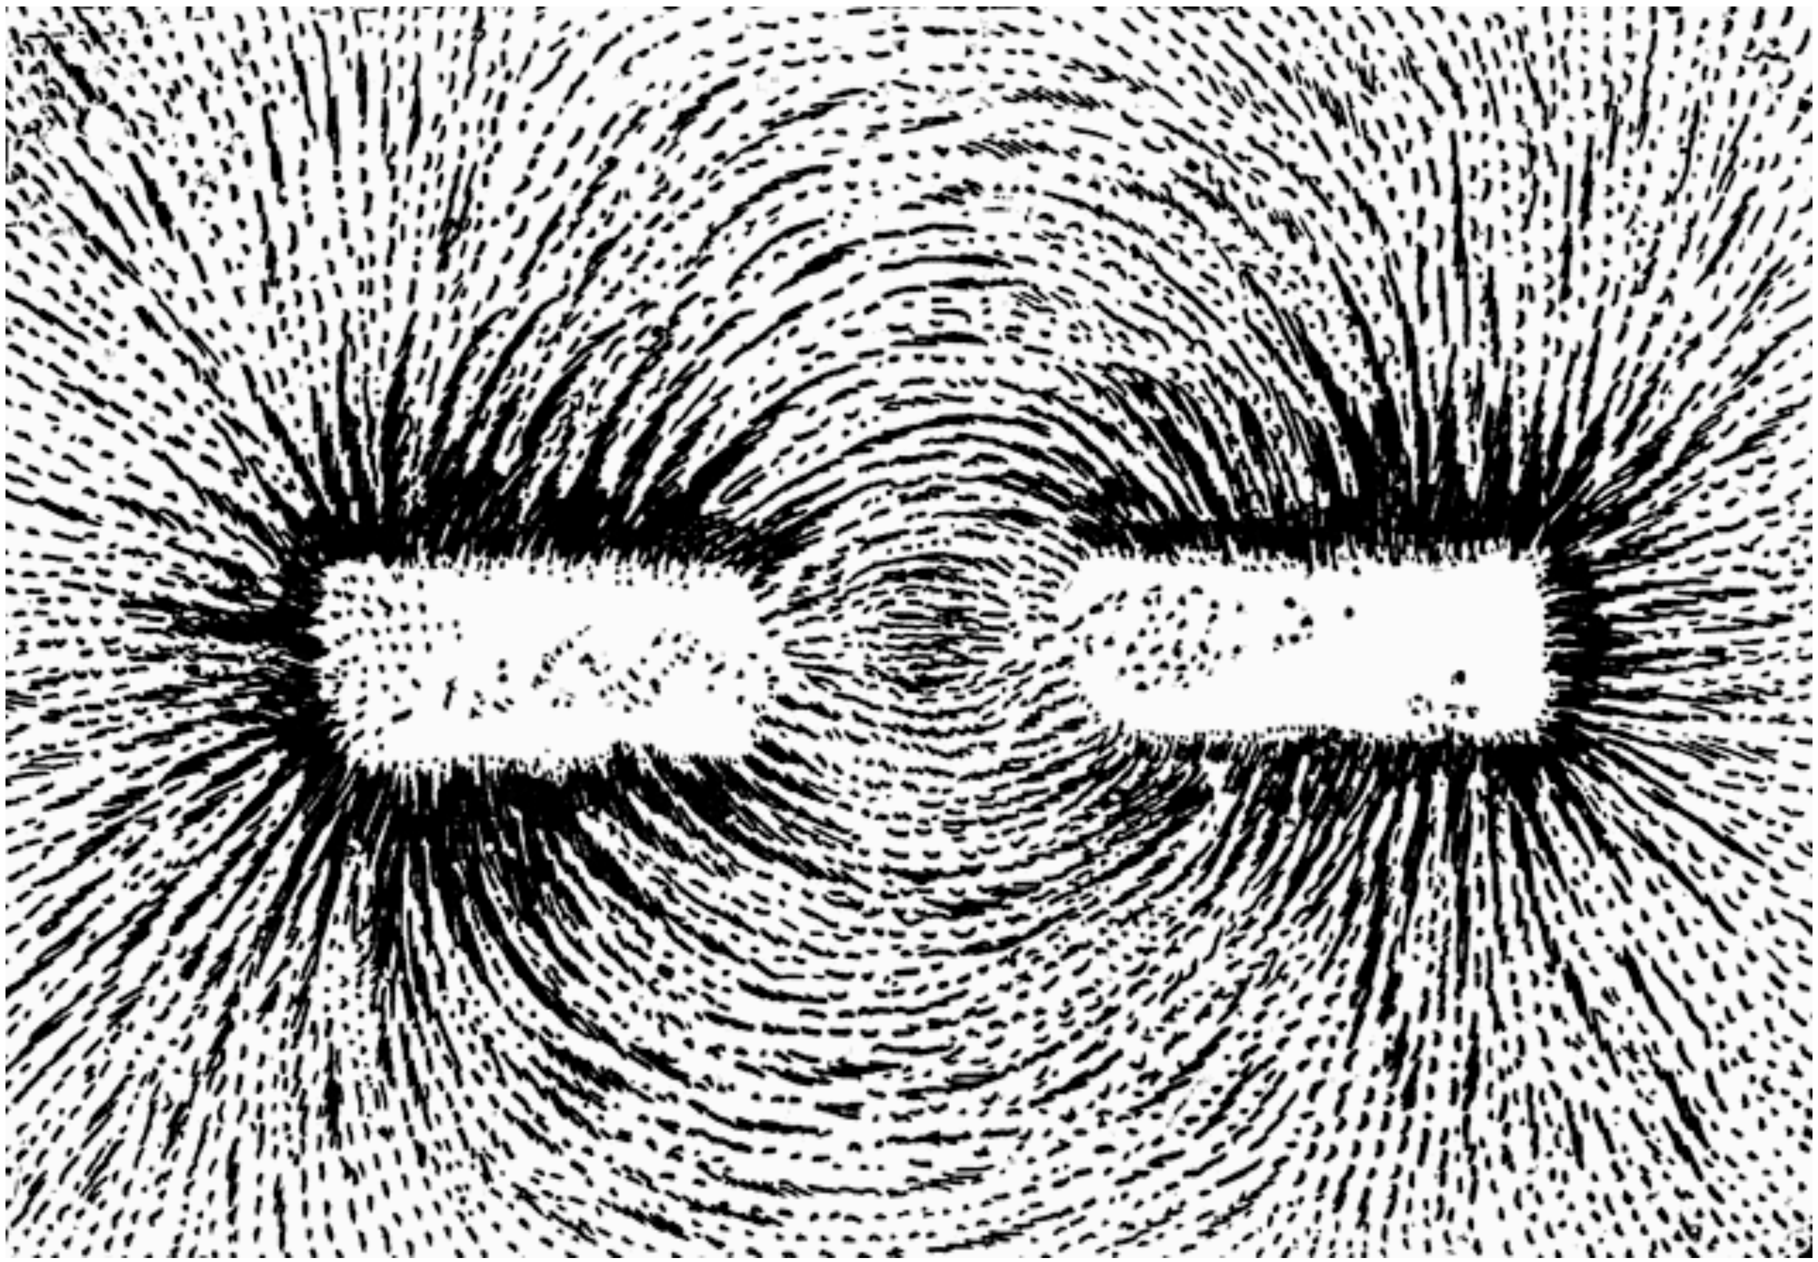

Scientific inference

Observations

Theory

Scientific inference

- Natural Behavior
- Just ask people questions
- Test/quizzes (more controlled performance measures)
- Reaction Time (RT)
- Eye tracking
- Physiological measures (Galvanic skin response, heart rate)
- Brain measures (fMRI, EEG, MEG)
- Lesion studies/surgical interventions

Scientific inference

- Theories of information processing, inspired by modern computers
- Mind-as-a-computational device
- Build computational theories of the way we think the mind works, test the implication of those theories in new experiments, refine theories when needed.

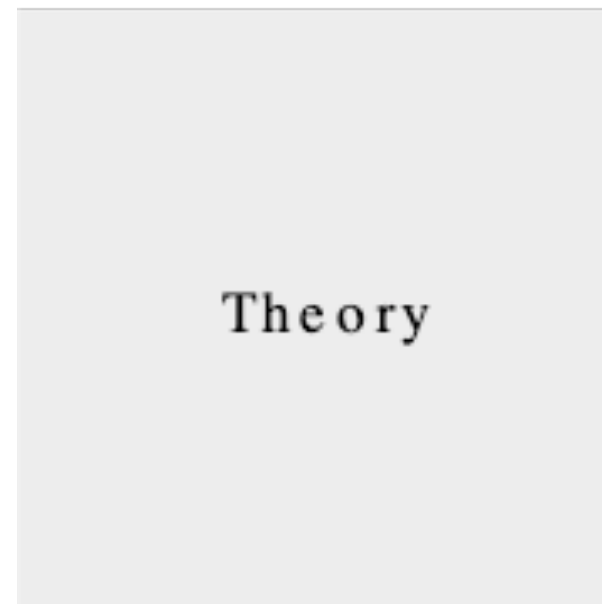

The science of the mind

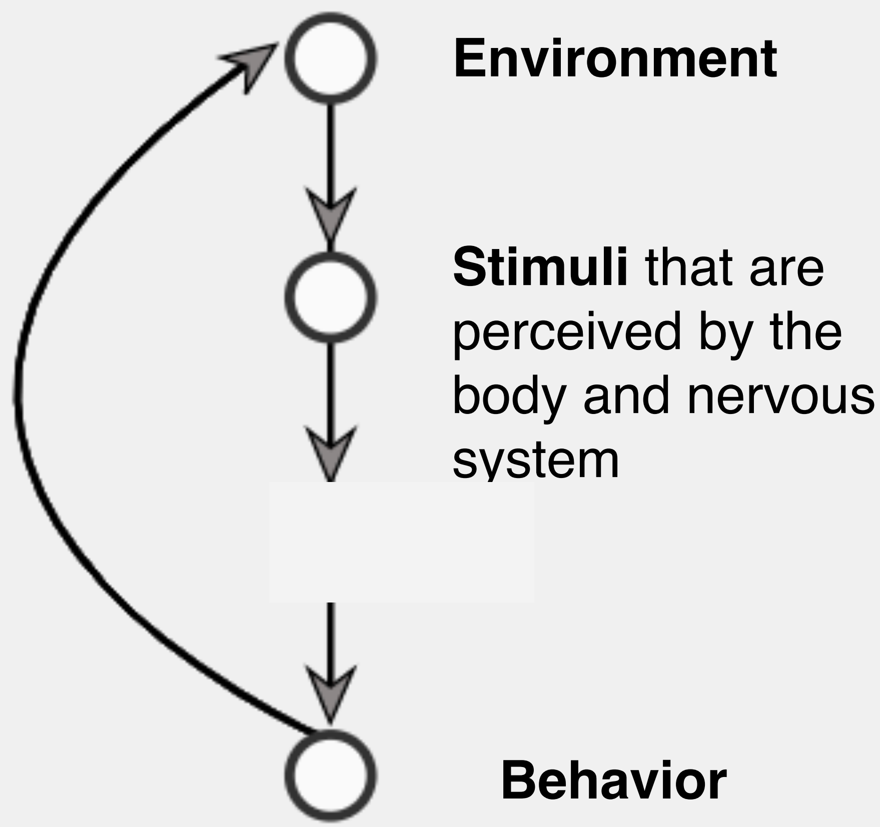

The science of the mind

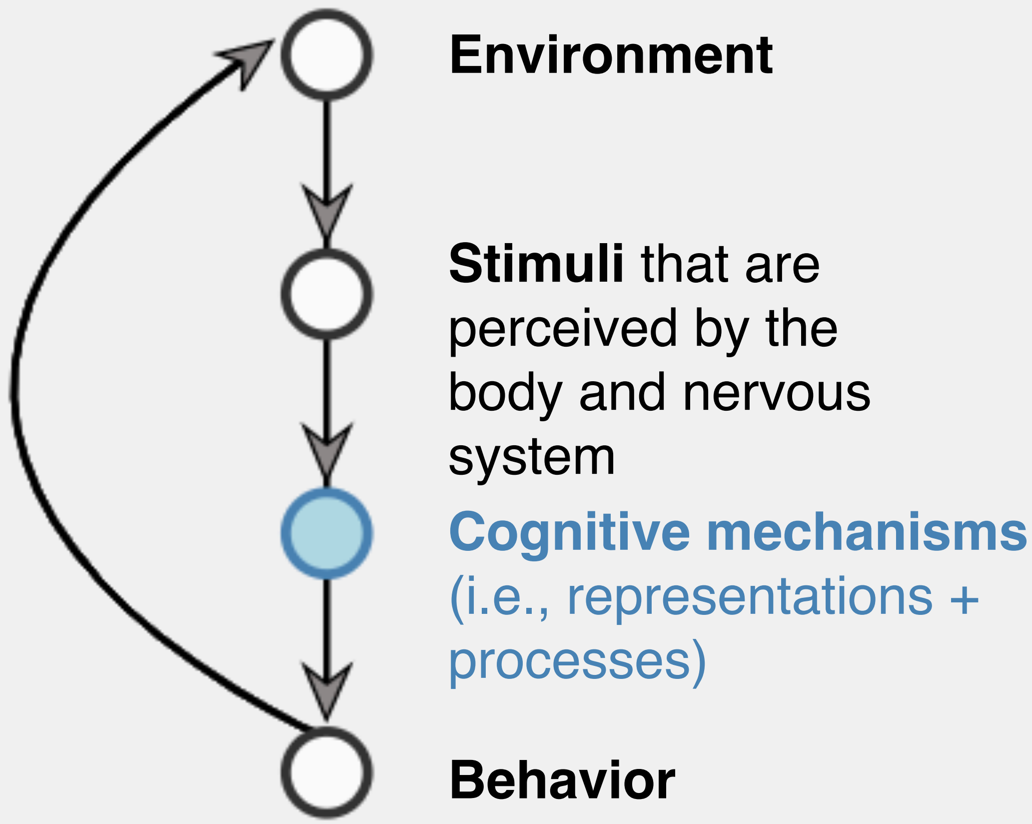

The science of the mind

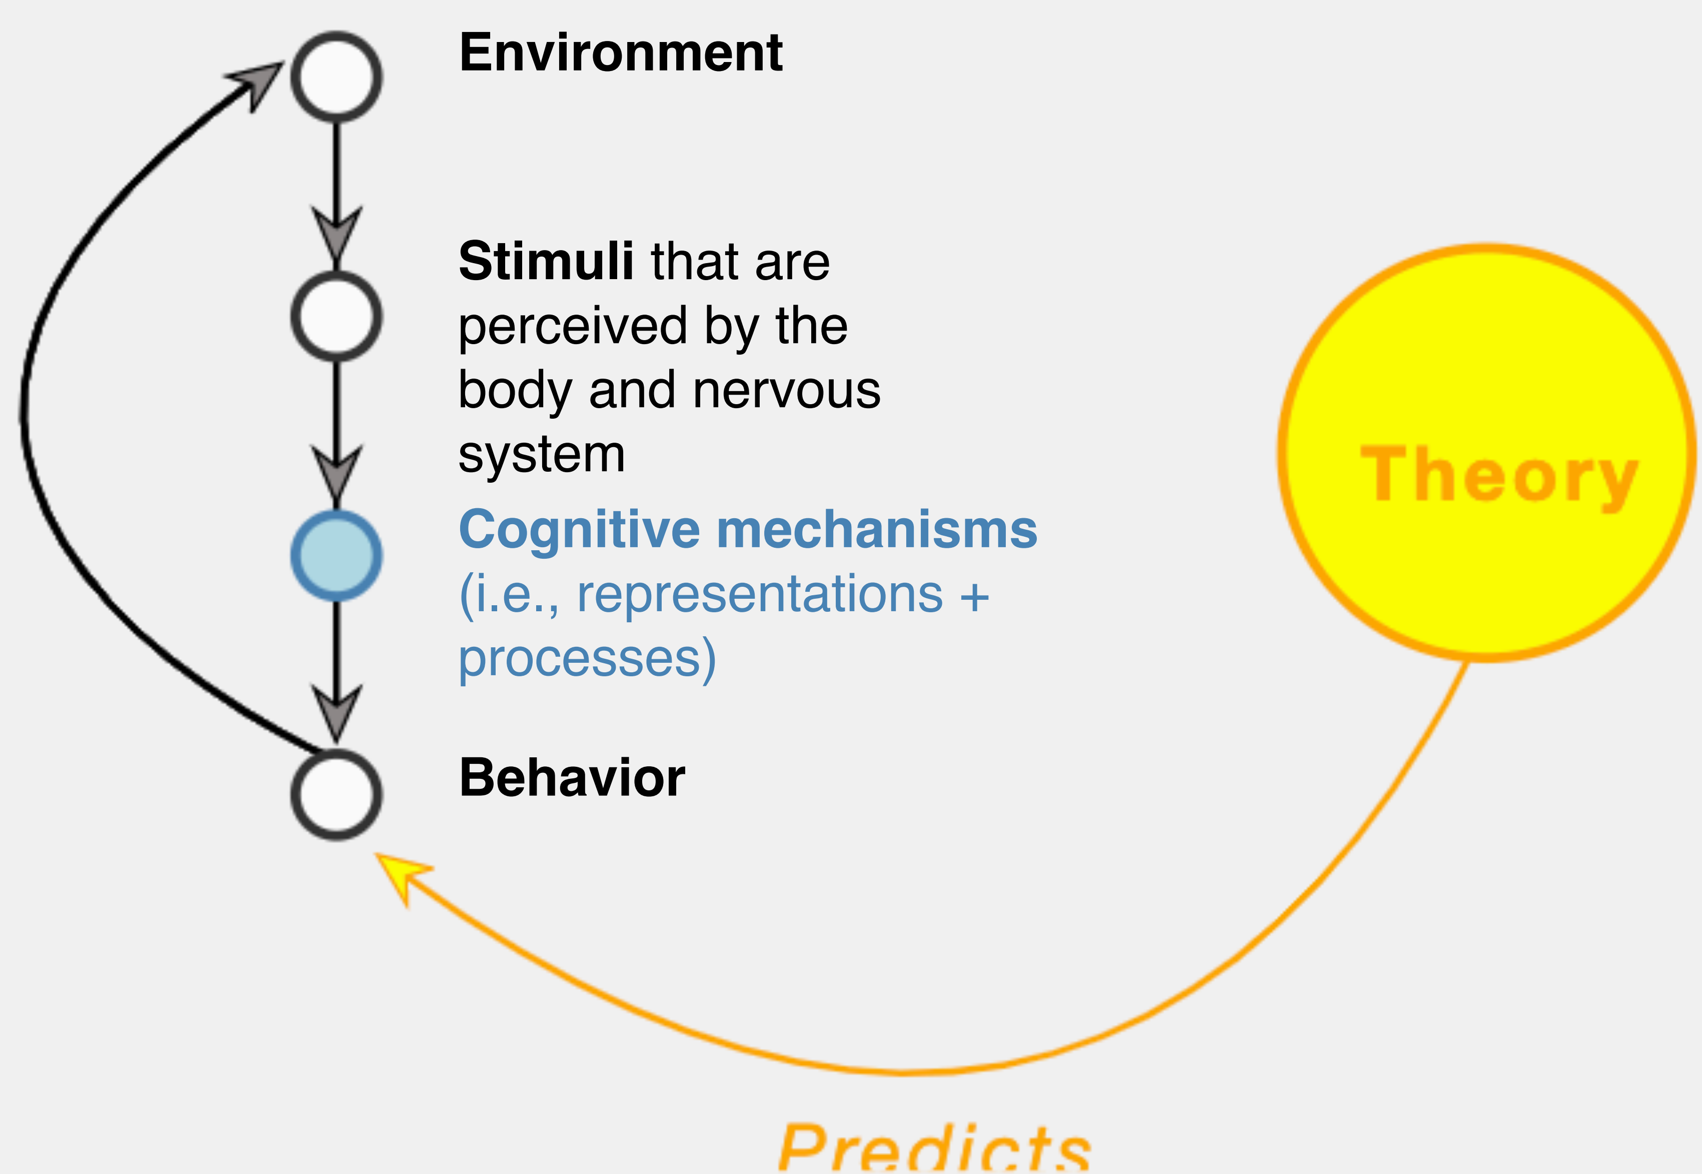

The science of the mind

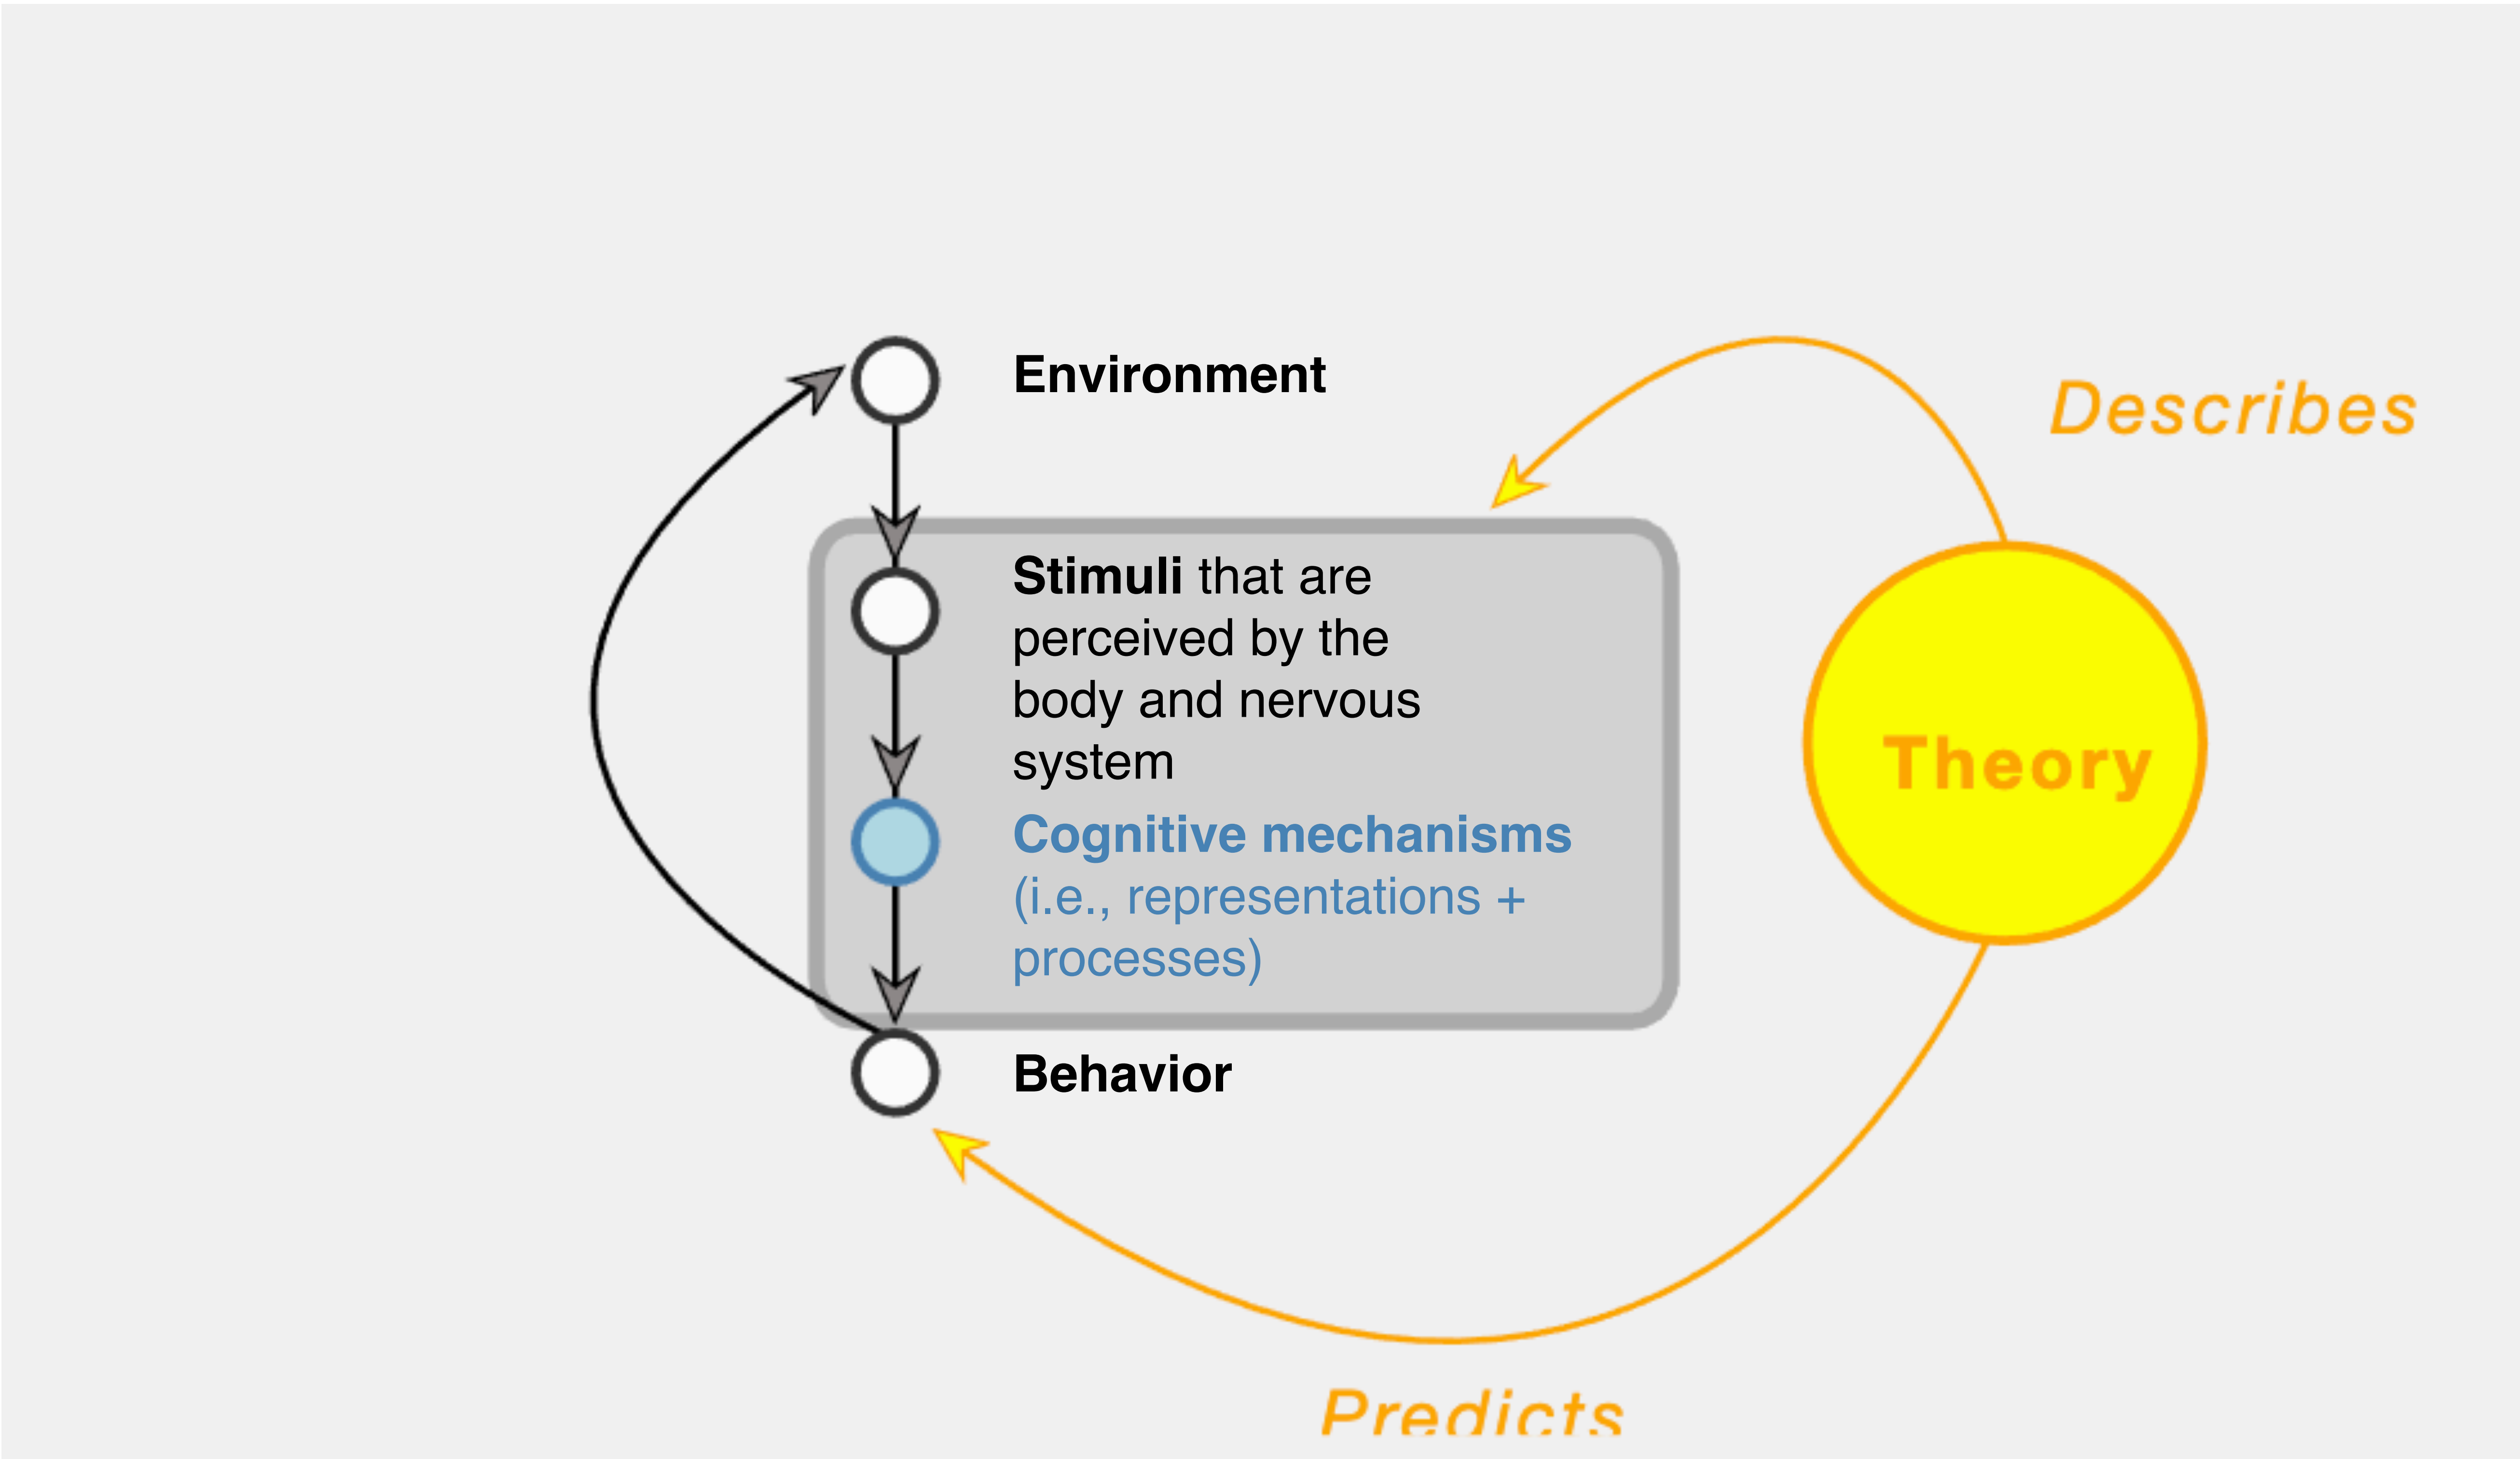

The science of the mind

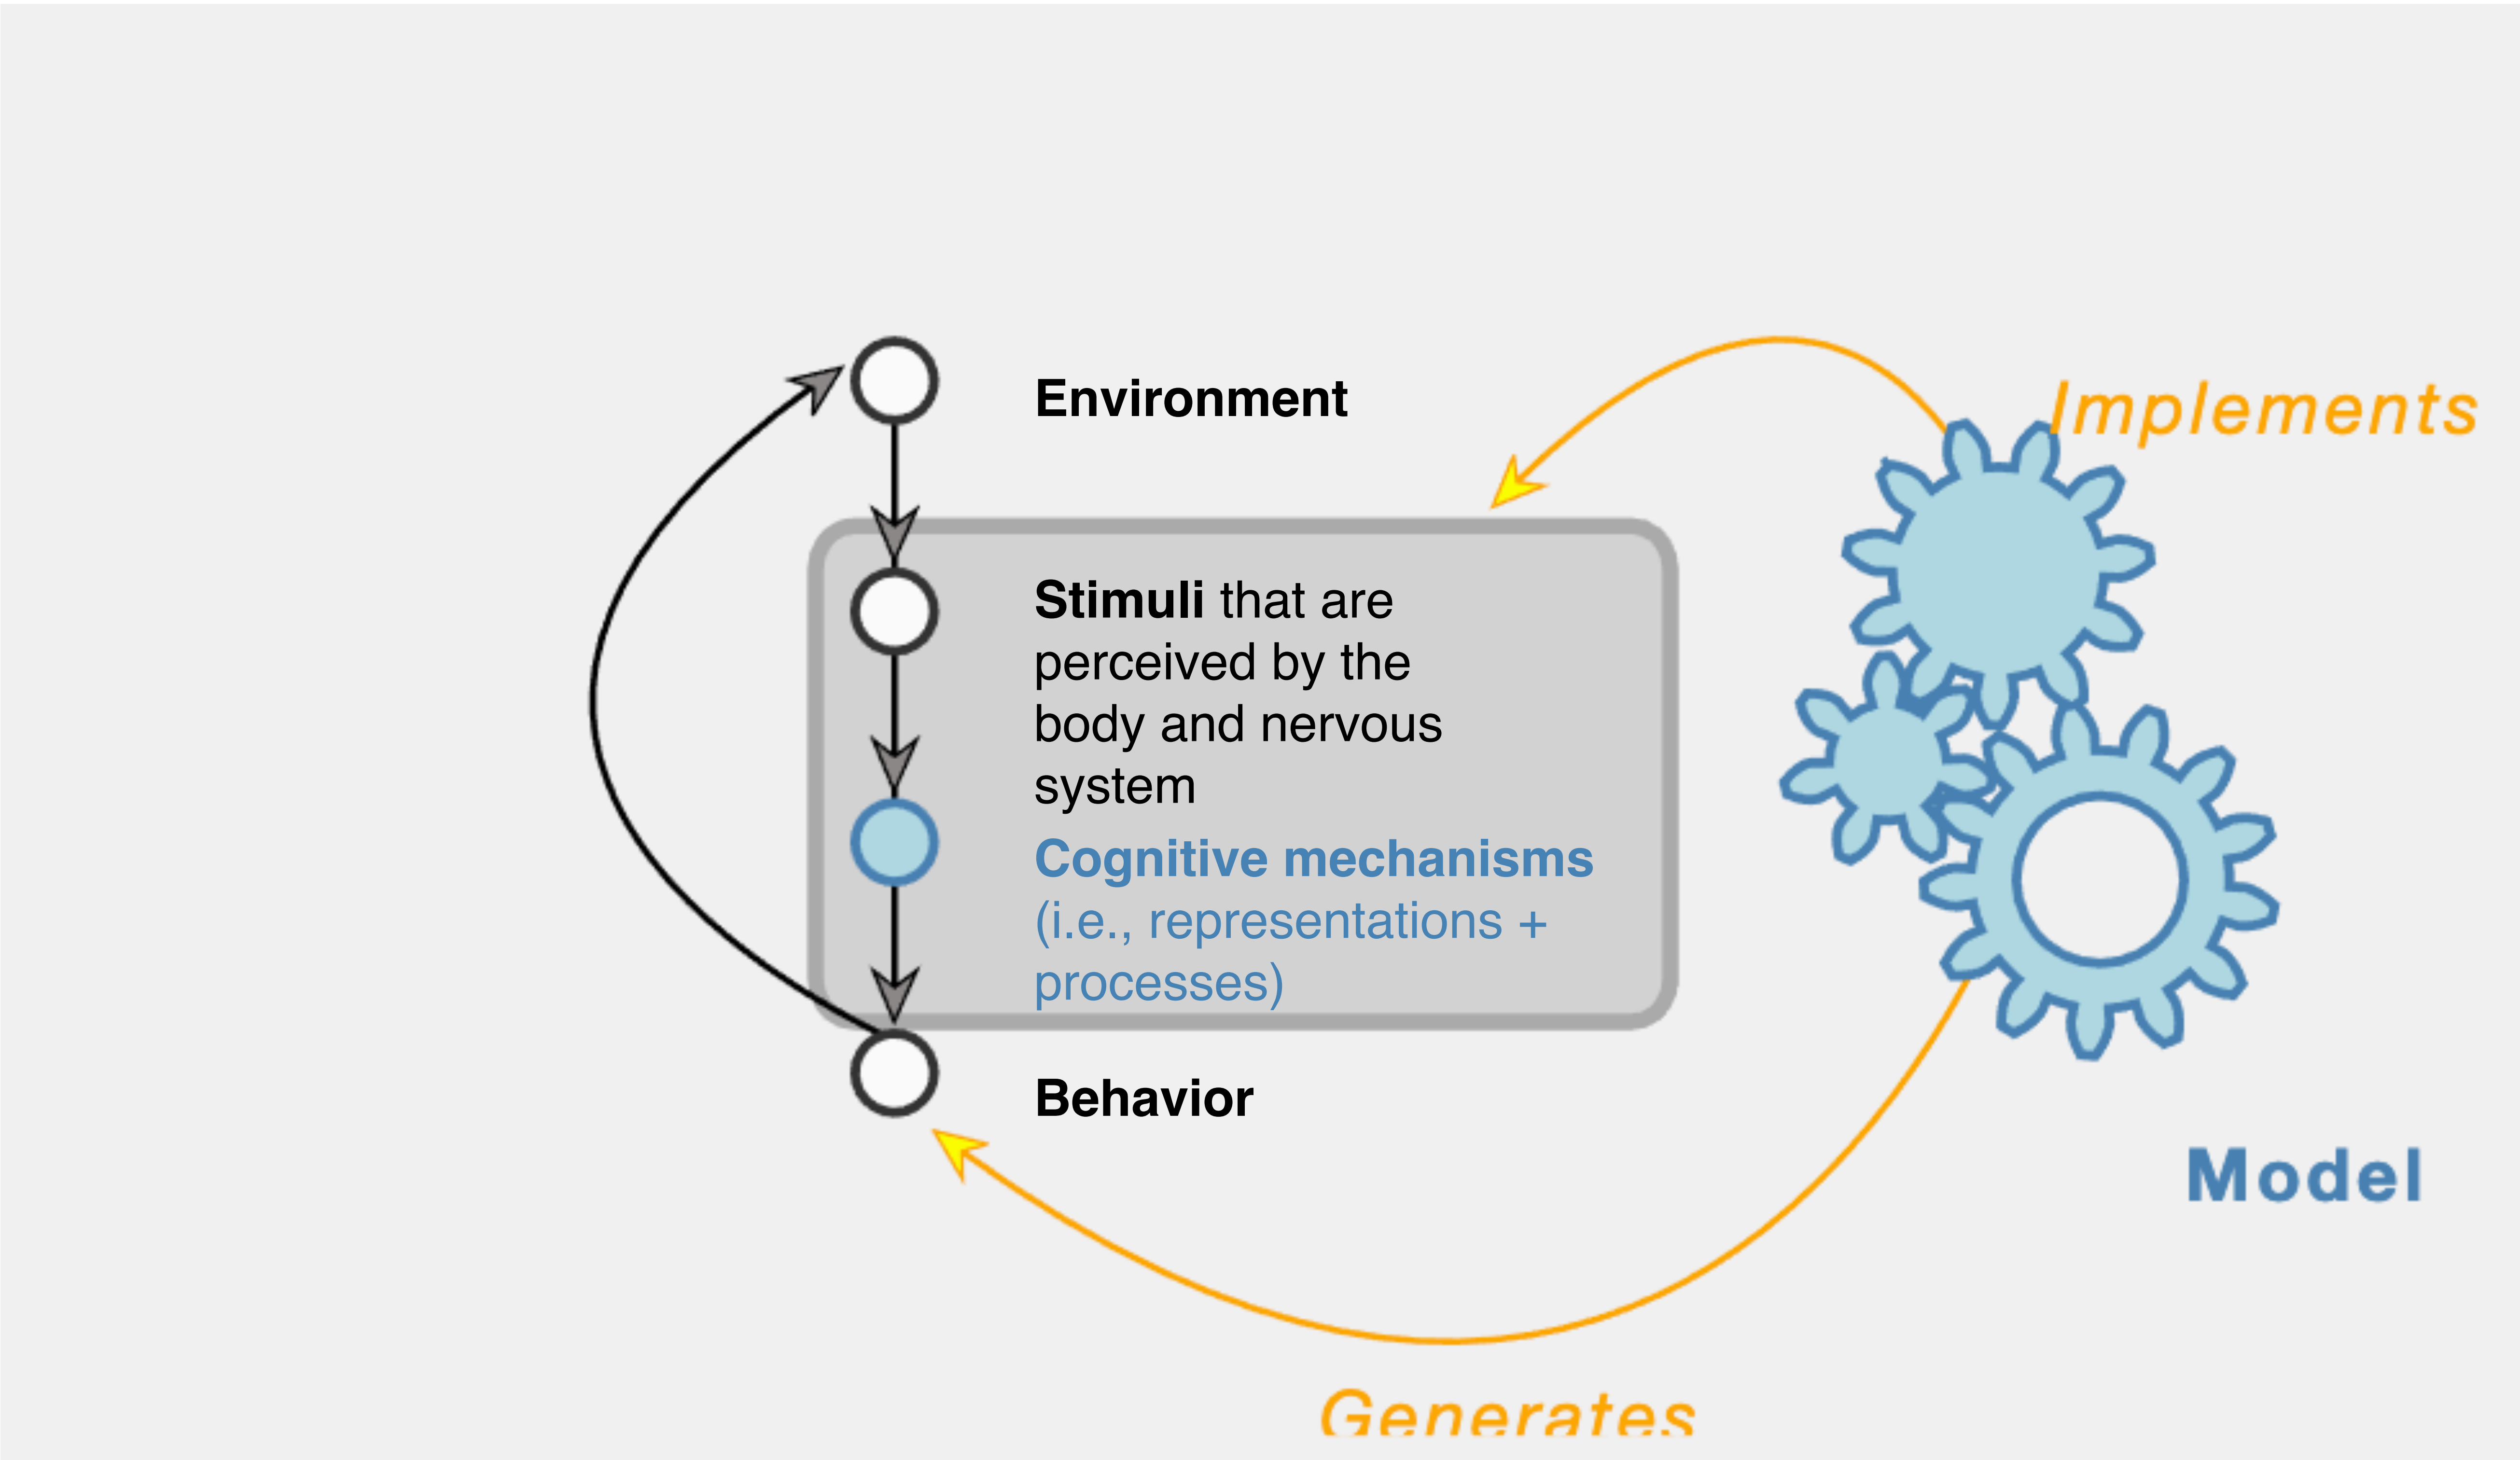

The science of the mind

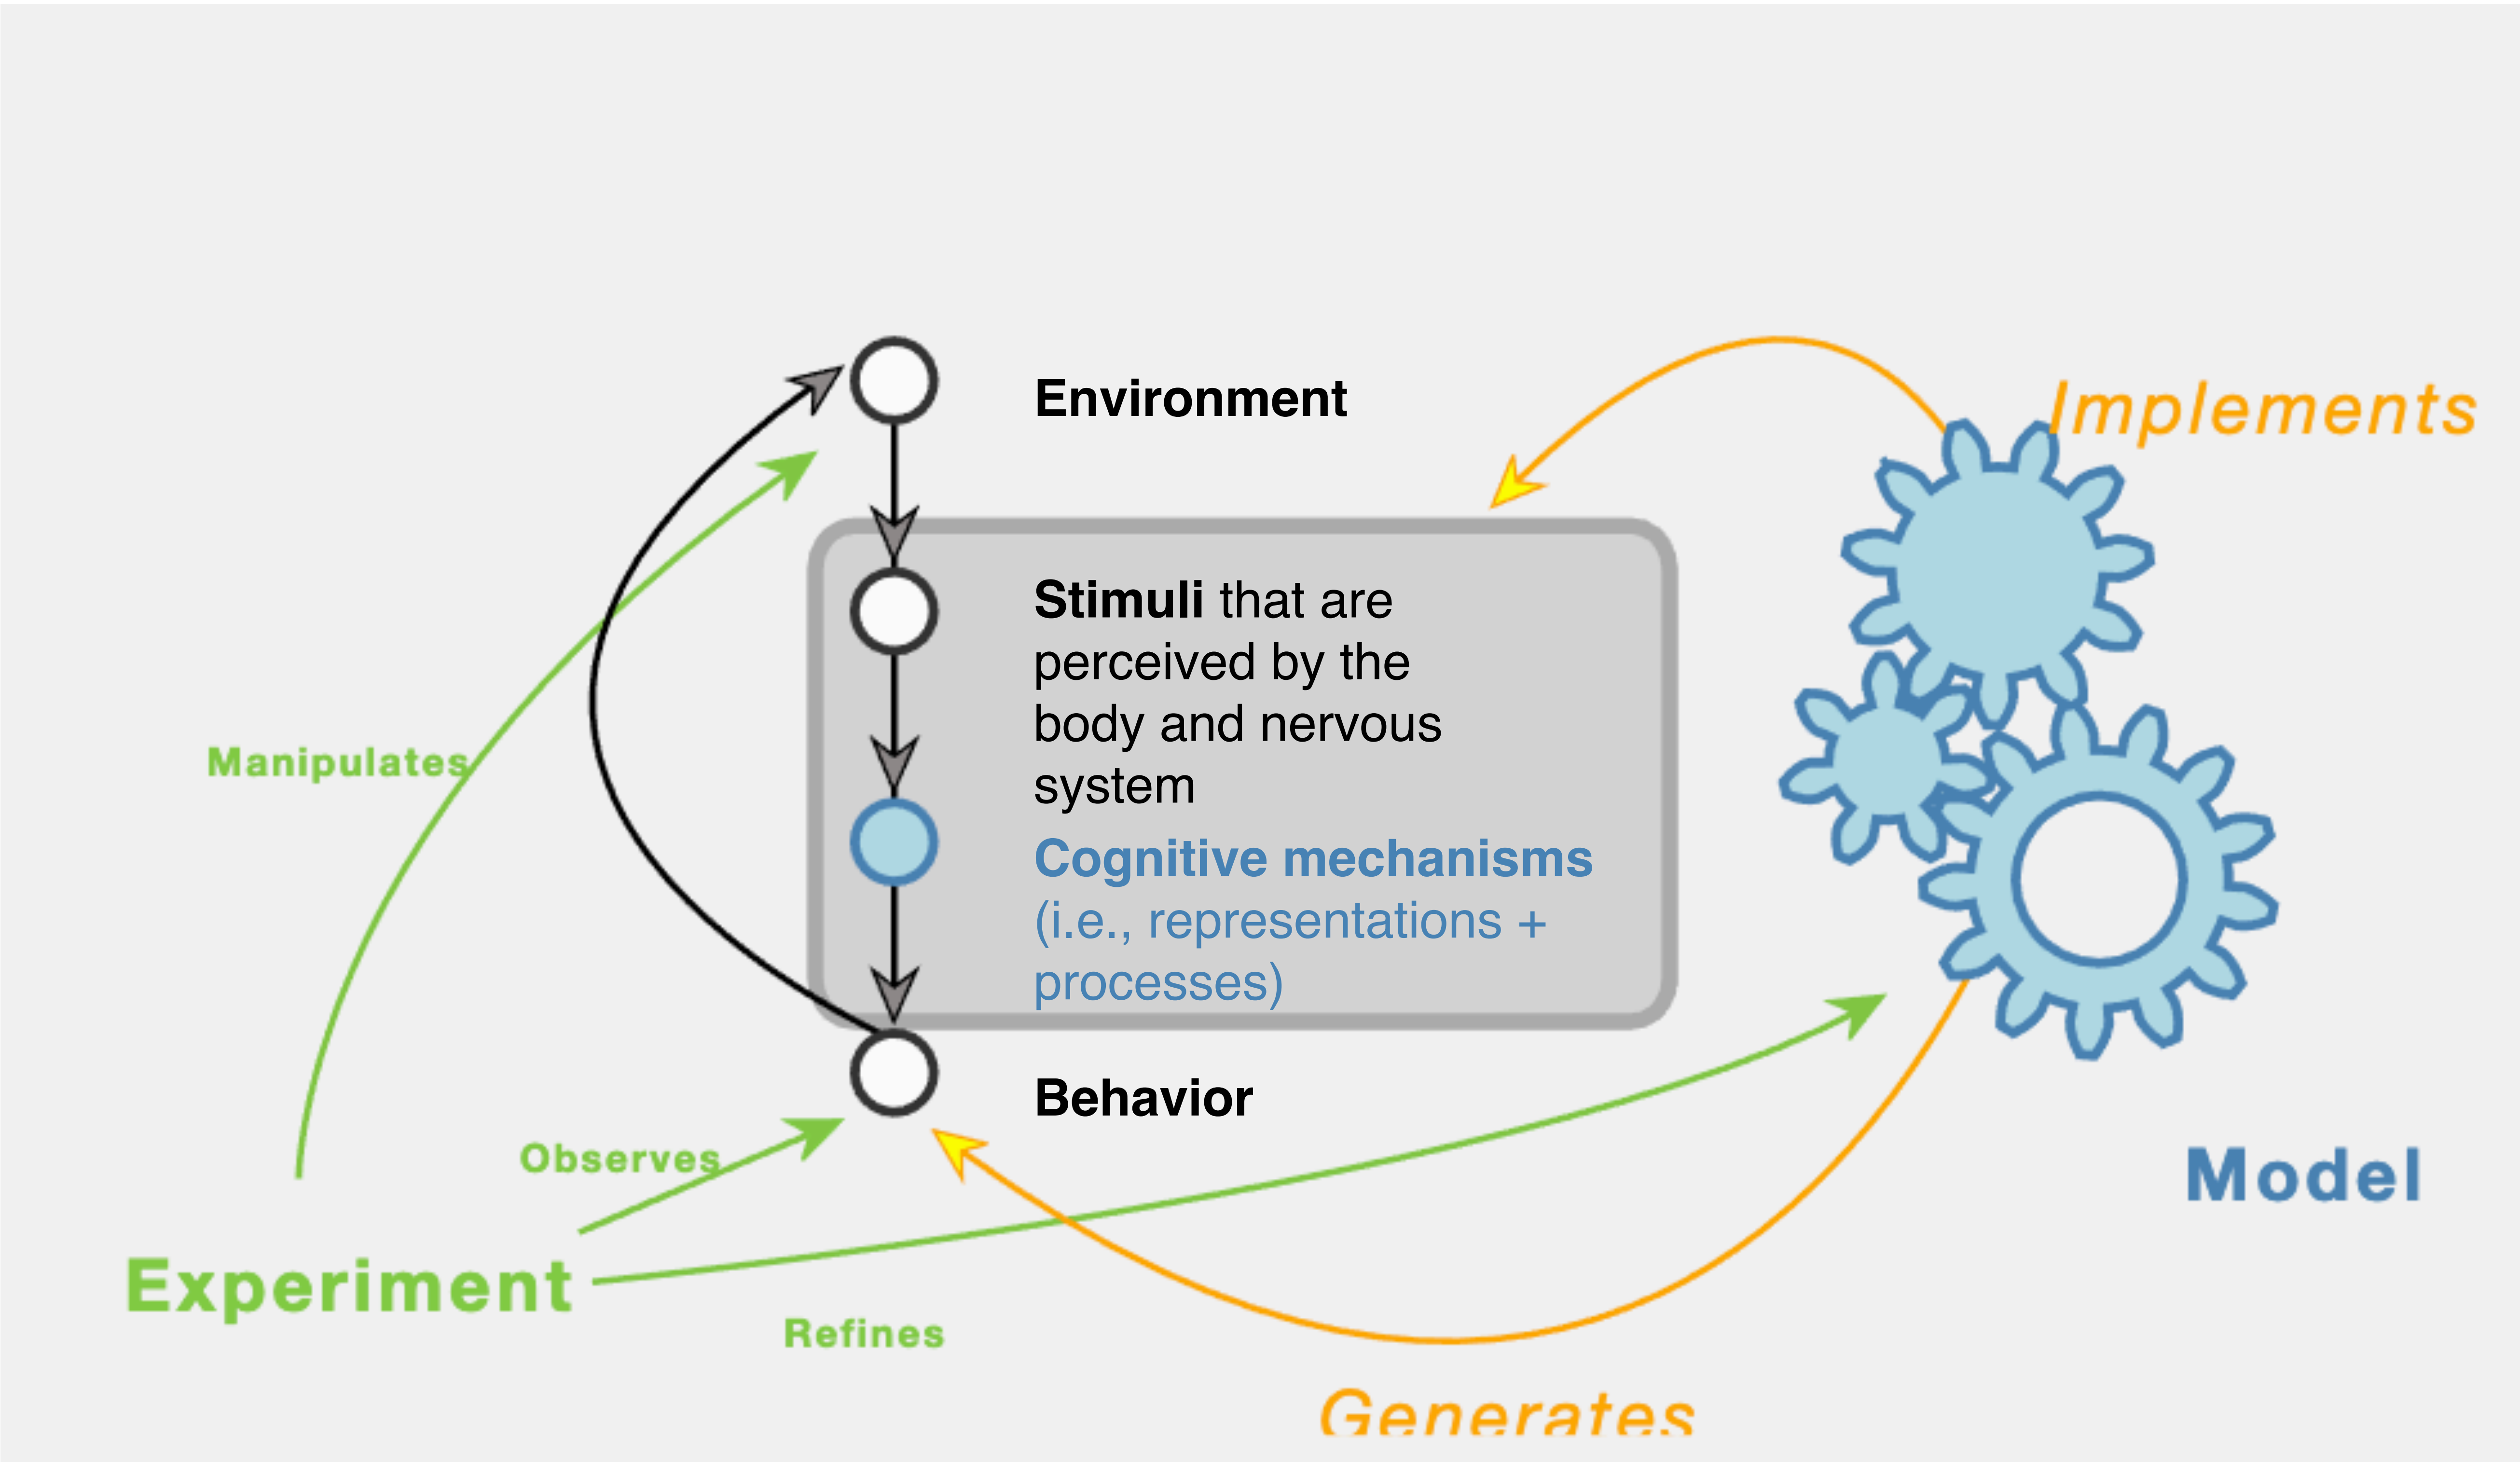

What is cognitive science?

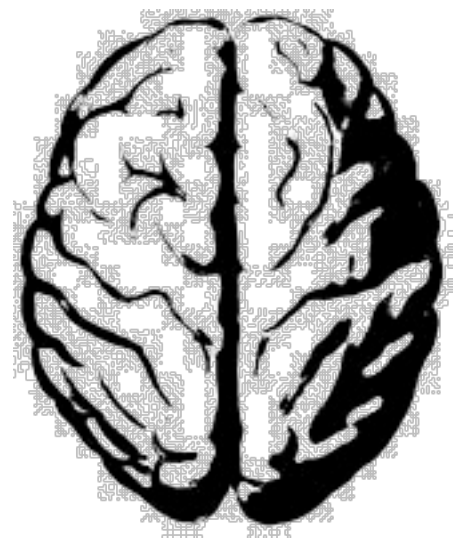

Cognitive

"Cognition, as defined by Ulrich Neisser, involves all processes by which sensory input are transformed, reduced, elaborated, stored, recovered, and used."

Science

"Science is the art of acquiring knowledge in such a manner that coherent structures of understanding can be erected on the basis of a critical evaluation of evidence."

What is a mind?

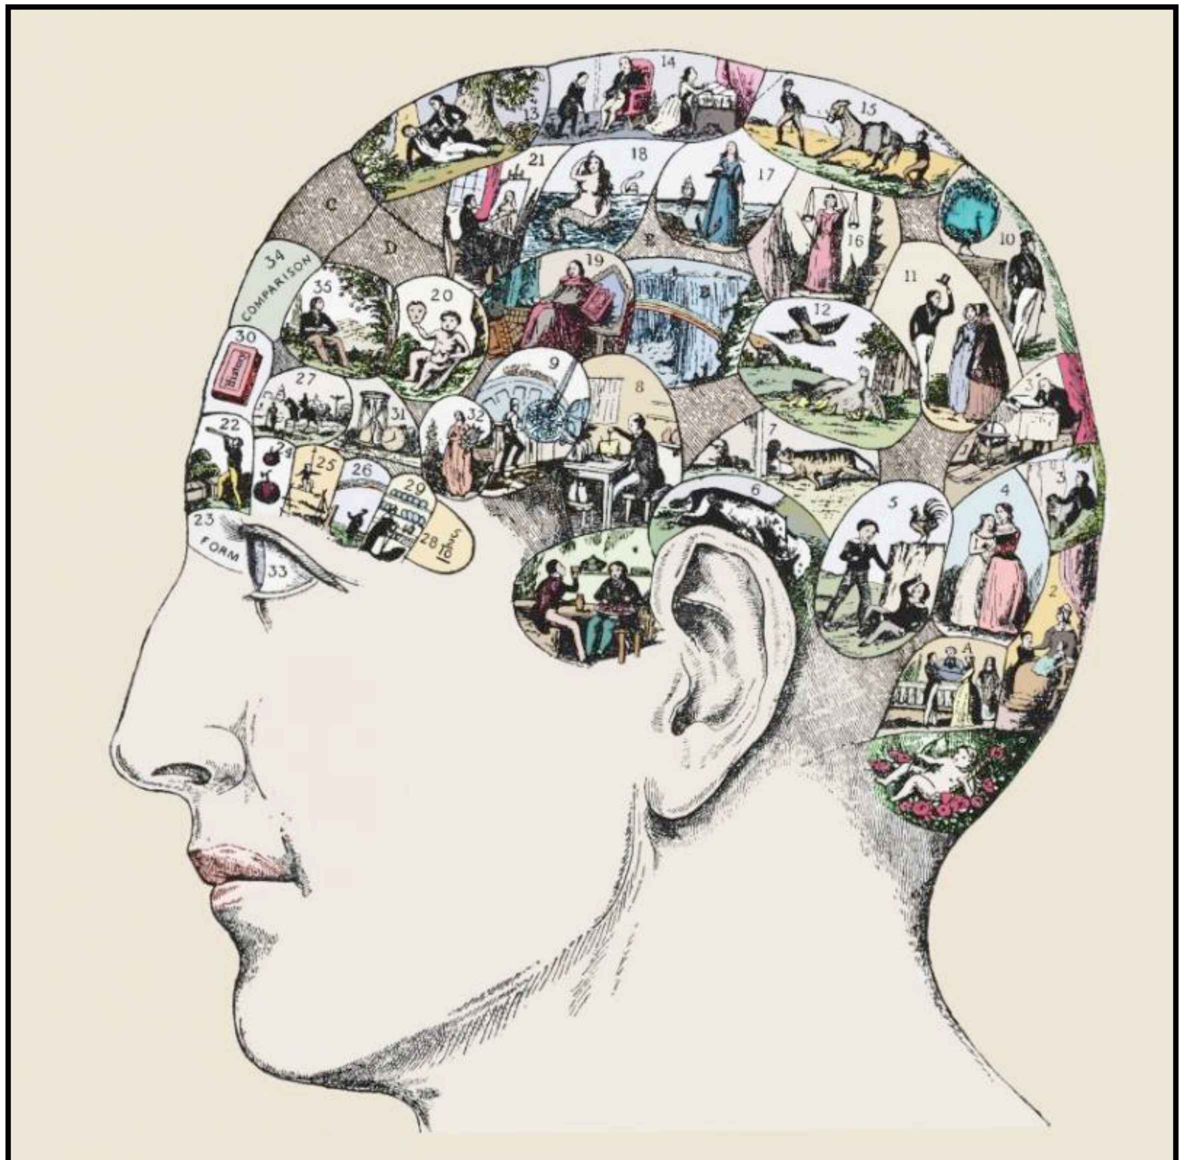

SYMBOLICAL HEAD.

This has been debated for thousands of years. If you don't have an immediate answer, don't feel bad. Various proposals have been thrown around from by Plato, Buddha, Aristotle, Zoroaster…. ancient Greek, Indian, and Islamic philosophers, and even several folks at NYU.

What is a mind?

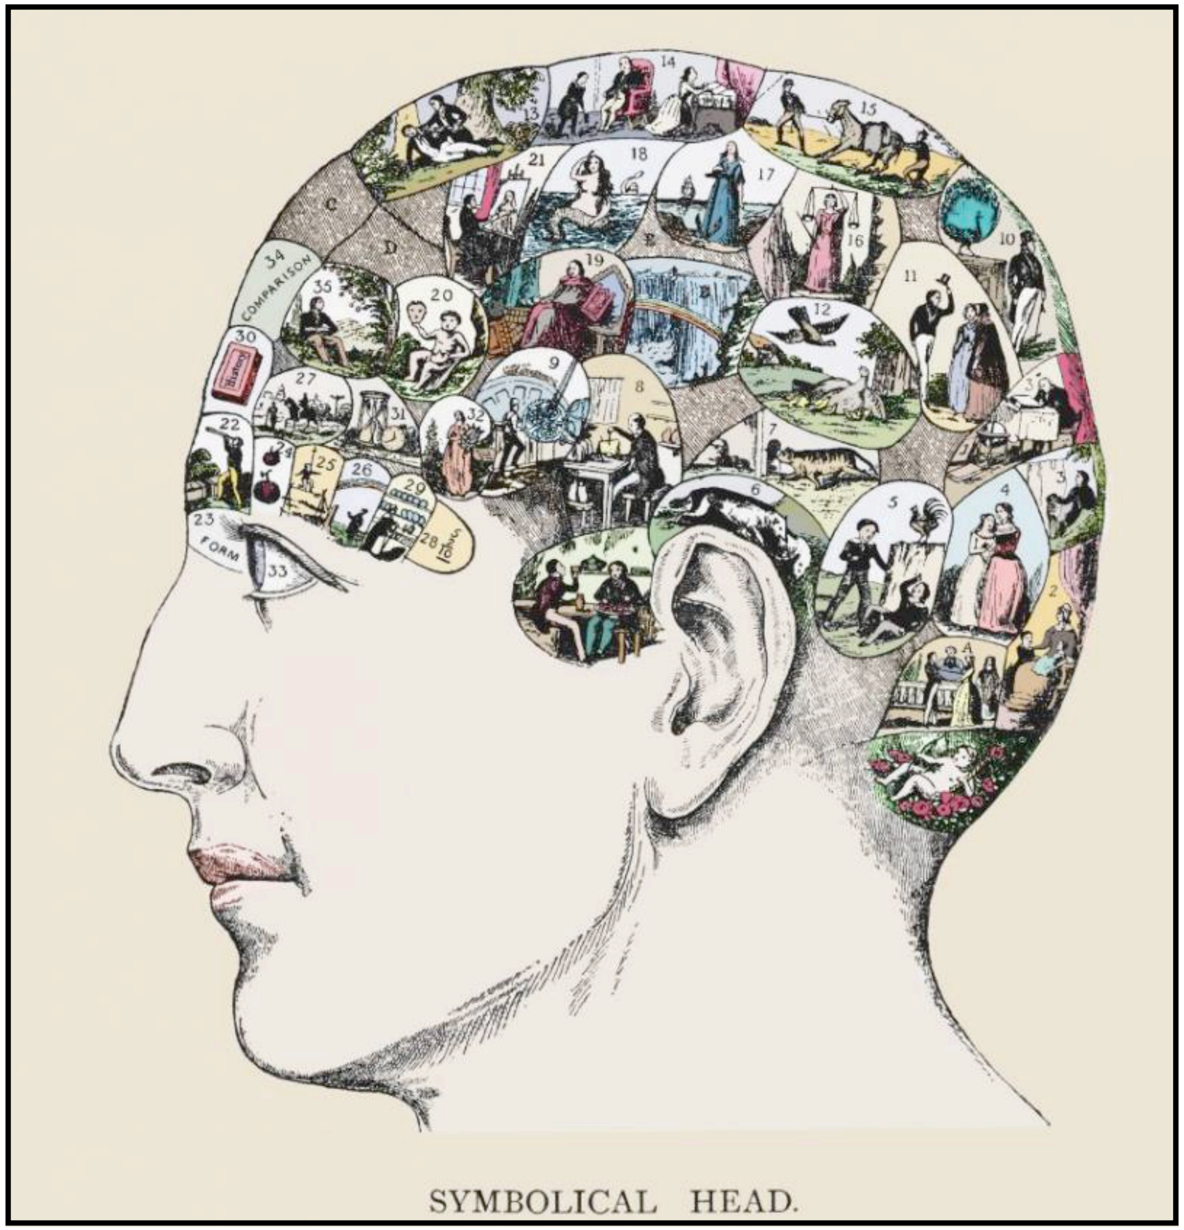

What do minds do?

Minds encompass our thoughts, which are mental processes that allow us to deal with the world. These include not only explicit wishes, desires, and intentions, but also unconscious processes.

What is a mind?

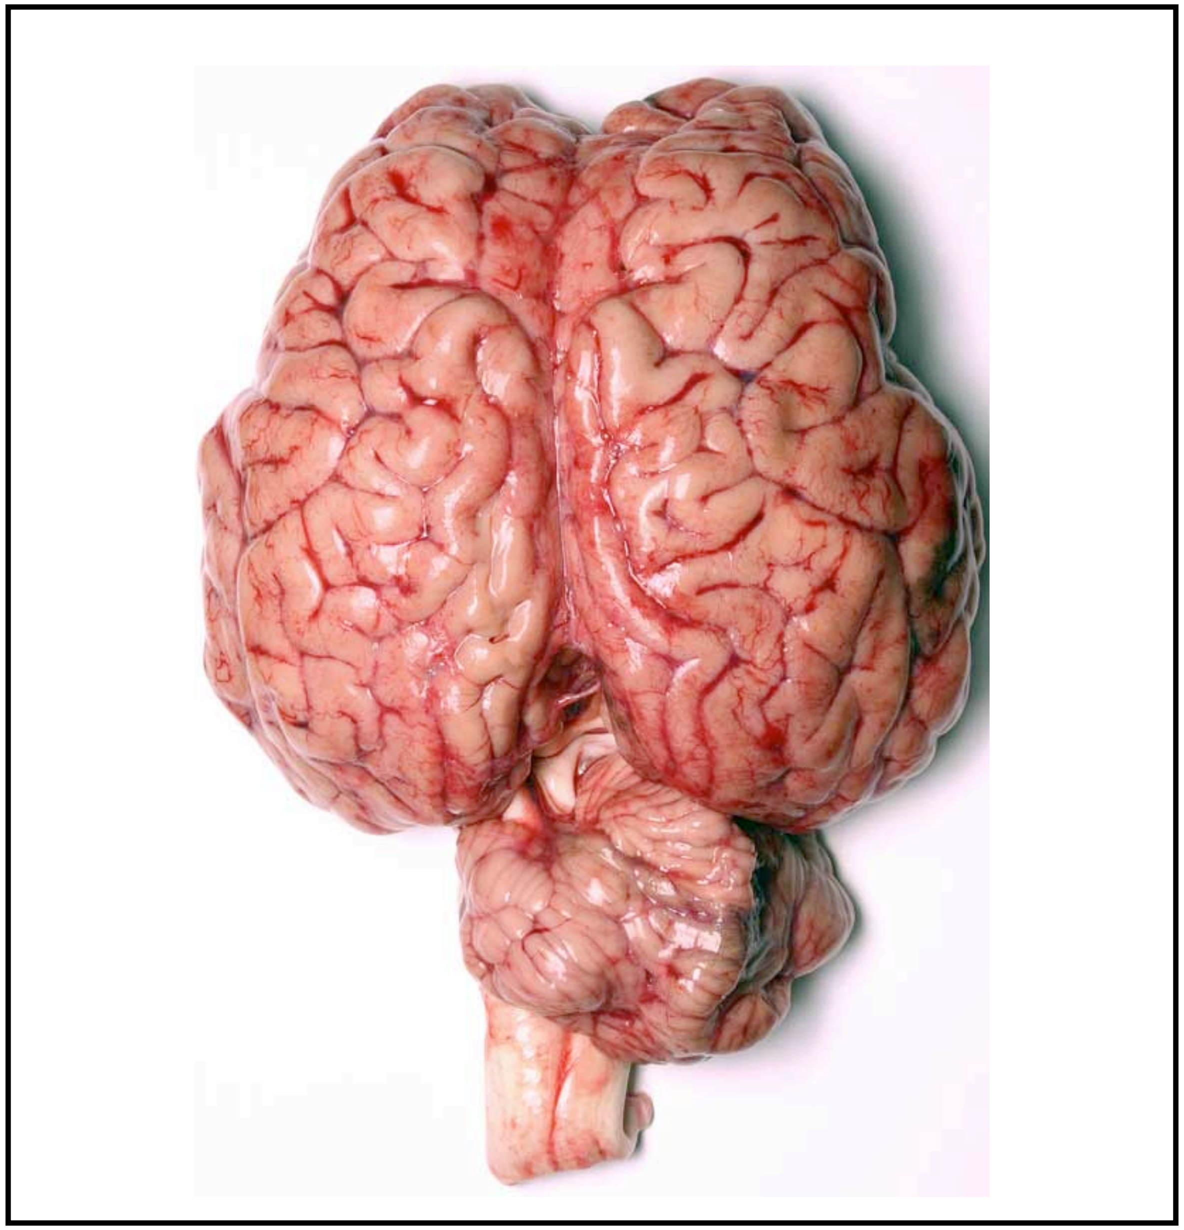

Does MIND=BRAIN?

We know that we can't have a mind or thoughts without a brain, but does that mean that minds and brain are synonymous?

What is a mind?

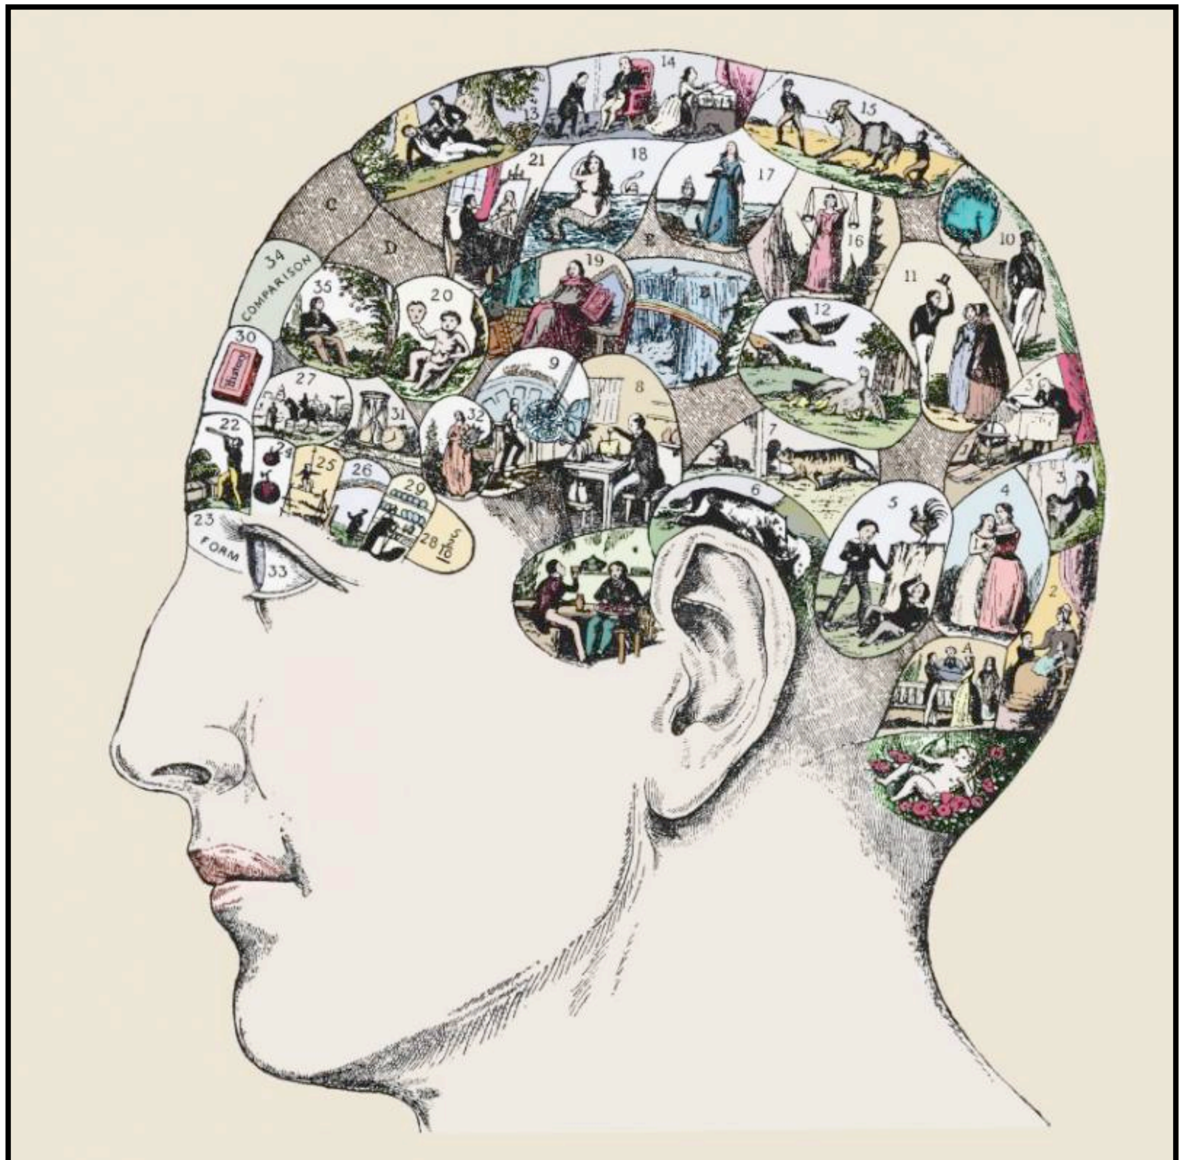

SYMBOLICAL HEAD.

A “slippery slope” argument can convince us that minds are not literally brains, but encompass anything that is organized as representational states that accurately reflect aspects of the world.

The Brain/Mind Riddle

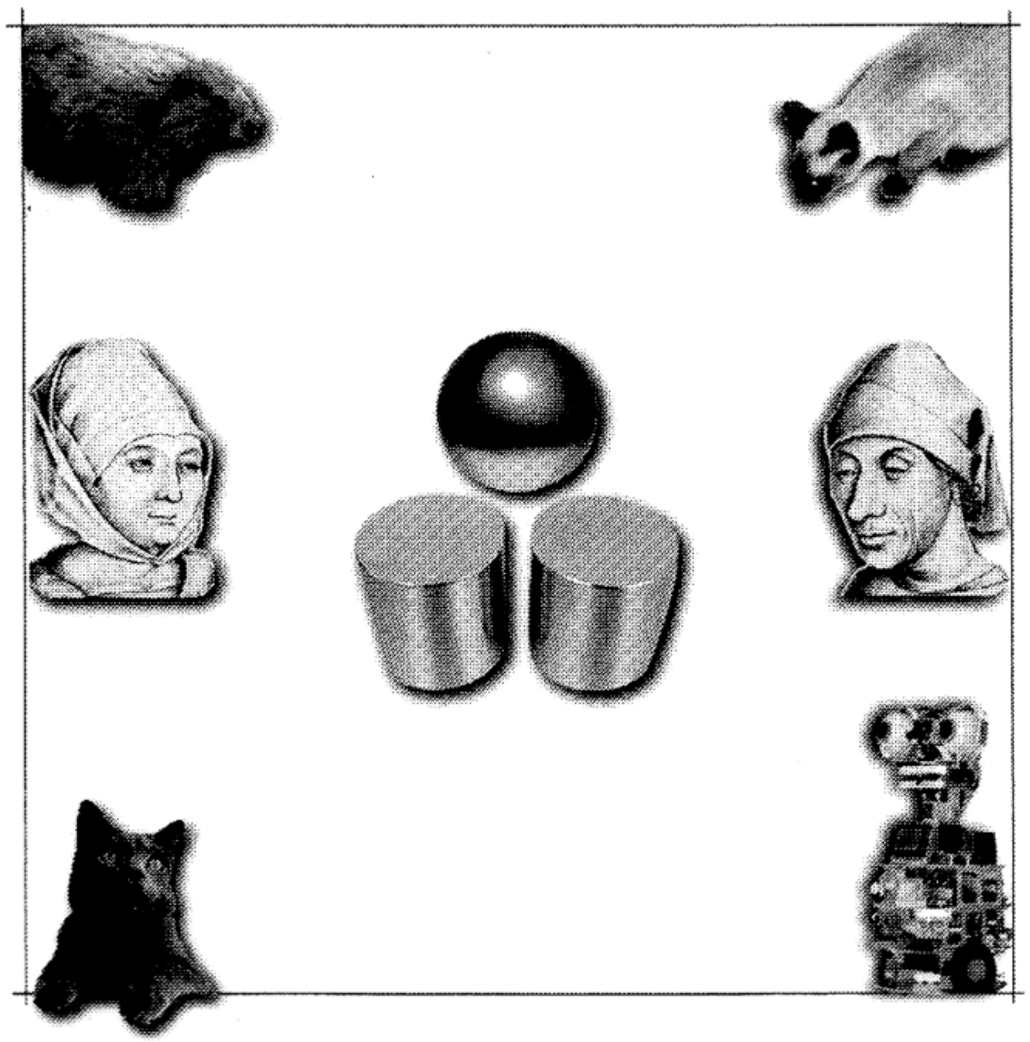

What is common to the various entities (person 1, person 2, cat 1, cat 2, robot, etc.) that look at this scene of two cylinders and a sphere and agree upon what is viewed?

Shimon Edelman's argument

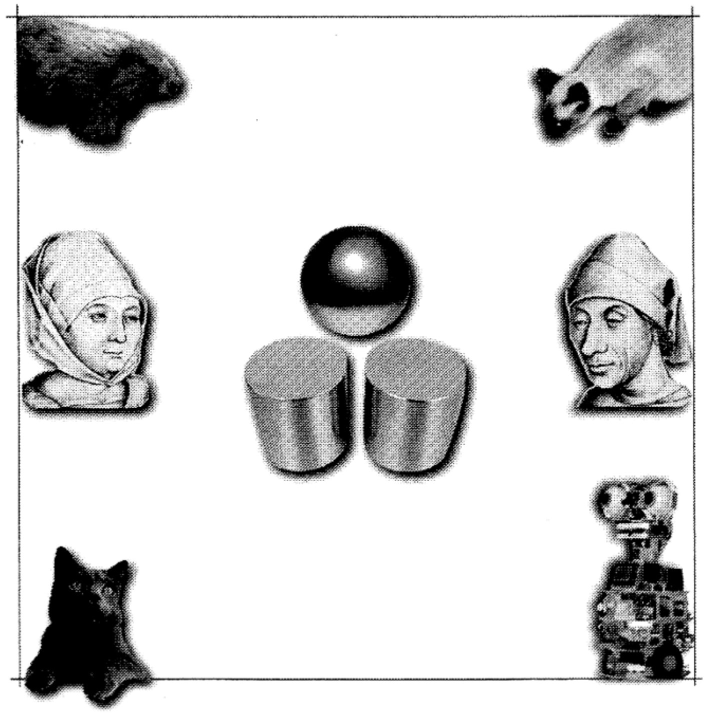

The question: What is common to observers viewing the same scene and who agree upon what is viewed?

- Can't literally be neurons. My neurons are my own, and you can't borrow them to solve your own problems.

- Is it the literal organization of the human nervous system? We know (or at least believe) that cats have a very similar visual system and view the world much like we do. Is it the mammalian visual system? What about other animals?

- What about artificial systems formed of computers and video cameras that can accurately recognize the scene as well?

- **The key to minds is not their physical substrate, but the relations that states of the system have to one another, and to the external environment.**

Minds as computers

- Minds aren’t human neurons or cat neurons or robot parts. They are dynamic, continually evolving systems that relate ongoing internal (i.e., mind) states and external (i.e., world) states
- Correspondences can be made between two systems by describing what they do, independent of their exact physical substrate.
- **We can describe these correspondences through the language of computation, simply because the THEORY OF COMPUTATION offers formal insights into how ostensibly dissimilar systems can be formally identical.**

Goals for this semester

- To explore (experimentally) how the brain **represents** and **processes** information in solving tasks

- We have to formulate **hypotheses** about how the mind might function, then design experiments to test these hypotheses

- This will involve testing various **theories** of cognitive function that can be formalized as computer programs or algorithms

- We are licensed to do this due to the fundamental idea that the mind can be understood as an **organization system** that evolves according to particular rules, steps, or procedures

A couple of good examples…

Journal of
Experimental Psychology

VOL. 77, NO. 3, PART 1 JULY 1968

ON THE GENESIS OF ABSTRACT IDEAS[1]

MICHAEL I. POSNER AND STEVEN W. KEELE
University of Oregon

Previous work has shown that Ss can learn to classify sets of patterns which are distortions of a prototype, even when they have not seen the prototype. In this paper it is shown that after learning a set of patterns, the prototype (schema) of that set is more easily classified than control patterns which are also within the learned category. As the variability among the memorized patterns increases, so does the ability of Ss to classify highly distorted new instances. These findings argue that information about the schema is abstracted from the stored instances with very high efficiency. It is unclear whether the abstraction of information involved in classifying the schema occurs while learning the original patterns or whether the abstraction process takes place at the time of the first presentation of the schema.

When a man correctly recognizes an animal he has never seen before as a dog, he has manifested an ability to generalize from previous experience. What has he learned that allows him to make the classification successfully? This question has been discussed in various forms since Aristotle. Some philosophers suggest a process of abstraction in which *S* builds up a representation of a figure (e.g., triangle)which is different from the instances he has seen. Others have denied the reality of such composite representations or abstractions. For example, Bishop Berkeley pointed out that he could search his imagination in vain for the abstraction of a triangle which was neither equilateral nor scalene but which represented both of these and all other triangles at once. The philosophical idea of abstract representations entered modern psychology from clinical neurology through the work of Barlett (1932) on schema formation (see also Oldfield & Zangwill, 1942).

In the areas of perception and pattern recognition, psychologists have studied questions related to schema formation. Attneave (1957) demon-

[1] This research was supported in part by National Science Foundation Grant GB 3939 to the University of Oregon. A preliminary

How do people represent categories?

“Bird”

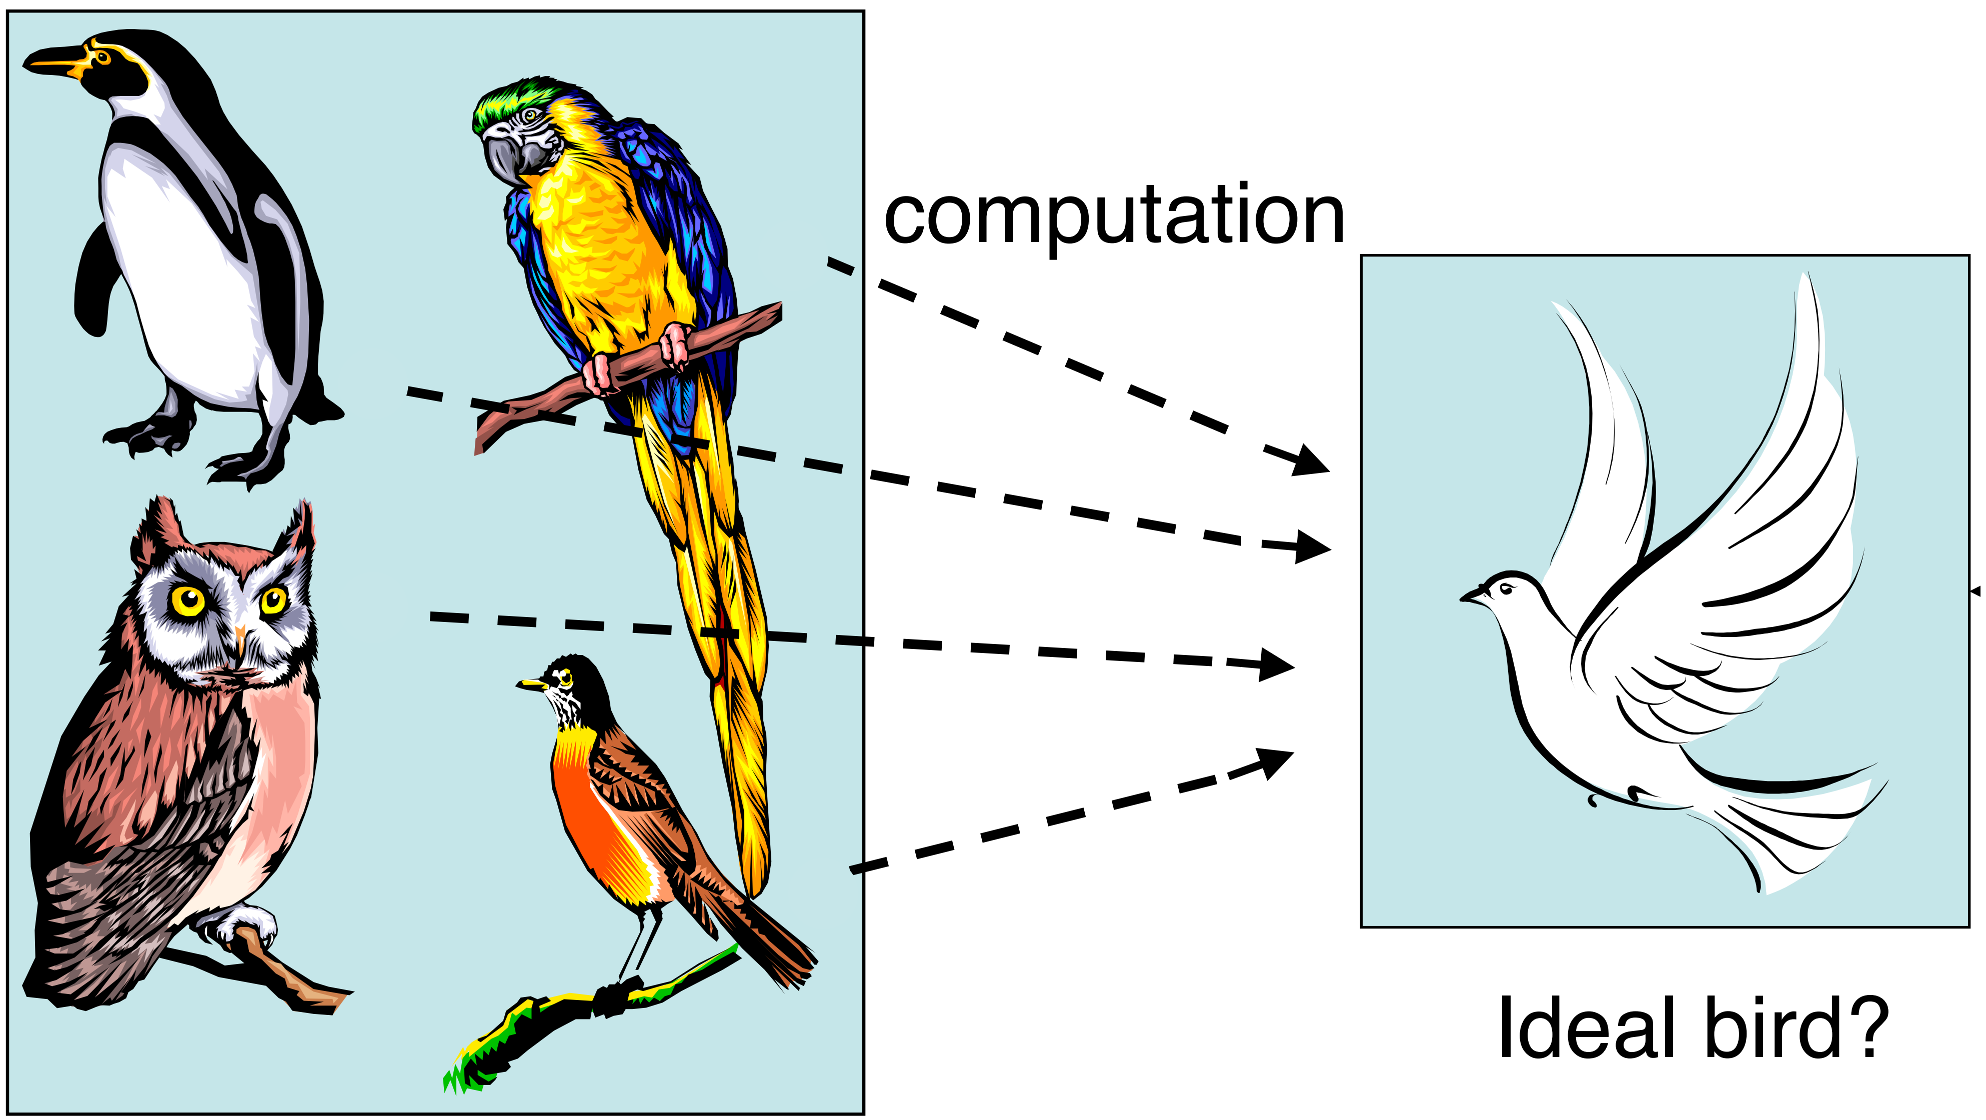

Posner and Keele's category learning task

1. **Instructions**: You will see stimuli from Category A or B. Please indicate which category you think is correct.

2. **Training phase**: Participants see stimuli one at a time. For each item, they respond "A" or "B". Feedback (the correct answer) is received during training.

3. **Test phase**: Participants may respond to additional stimuli. No feedback is given.

Category A or B?

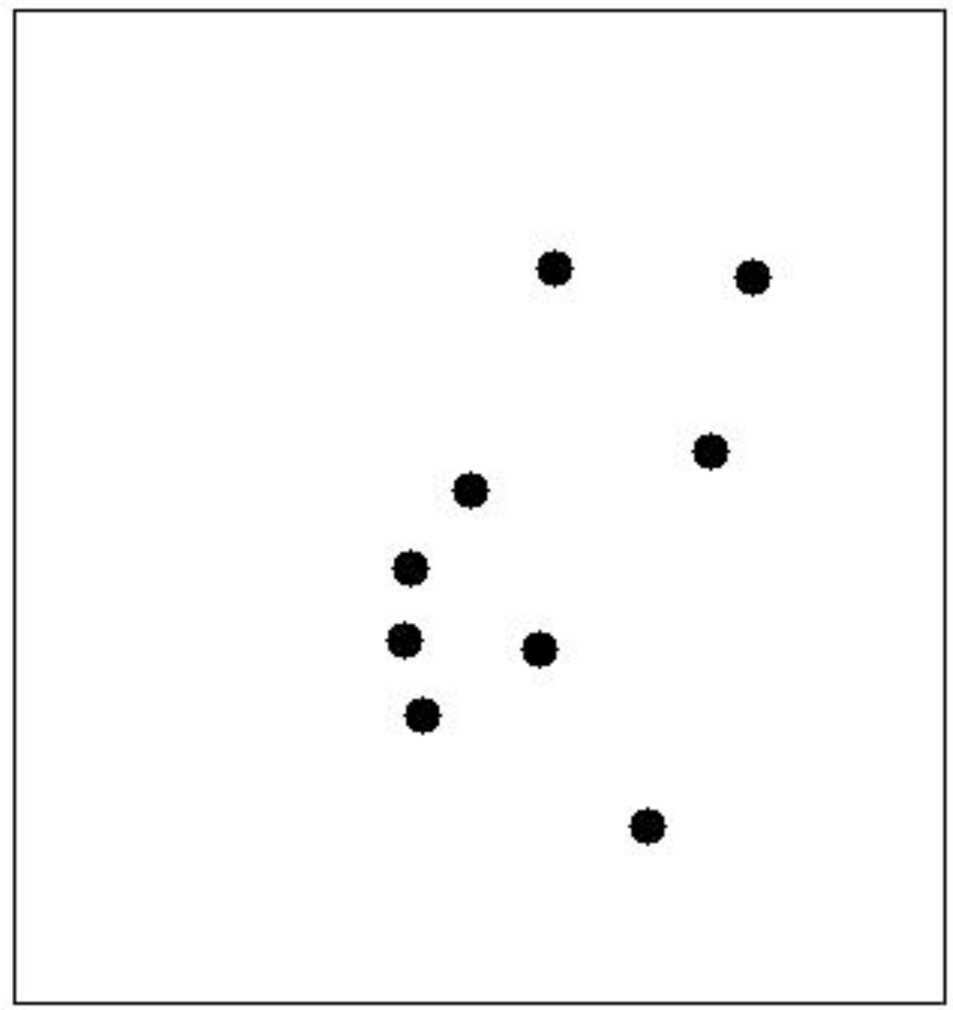

A

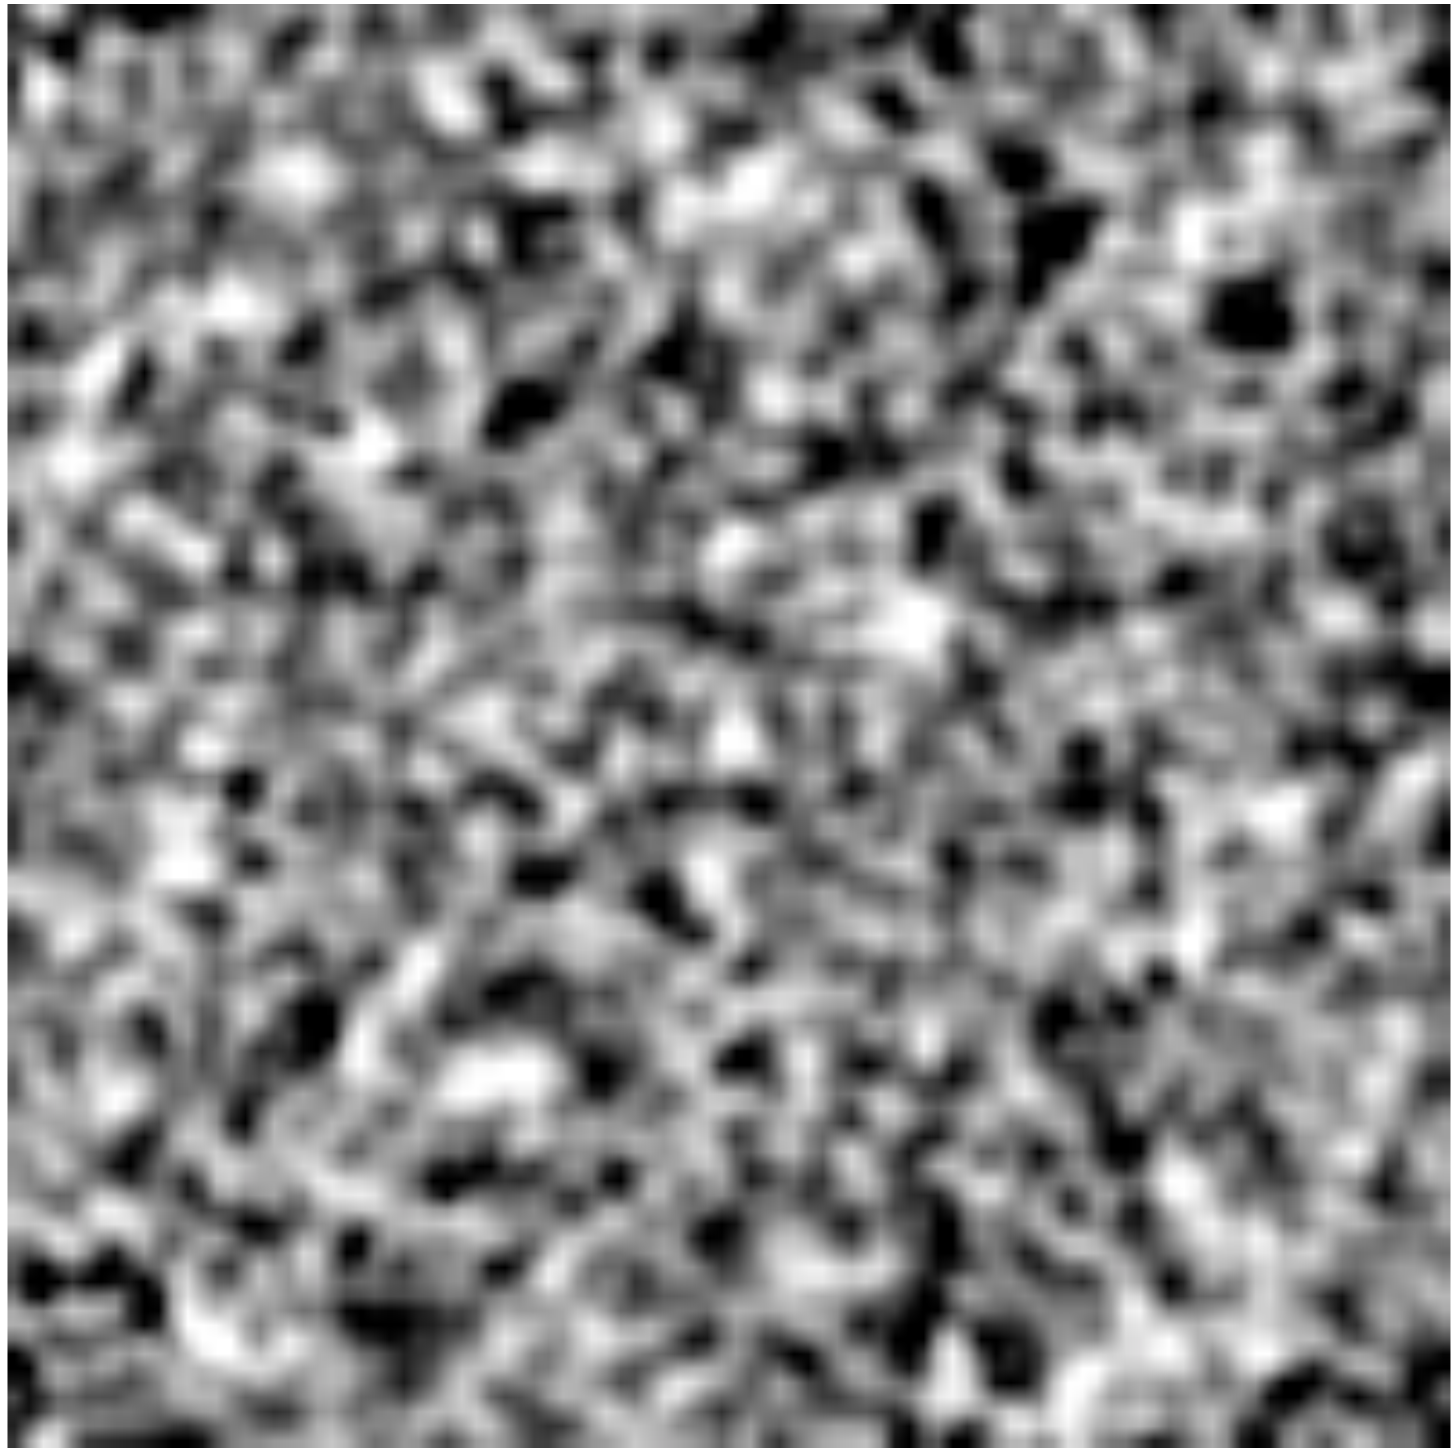

Category A or B?

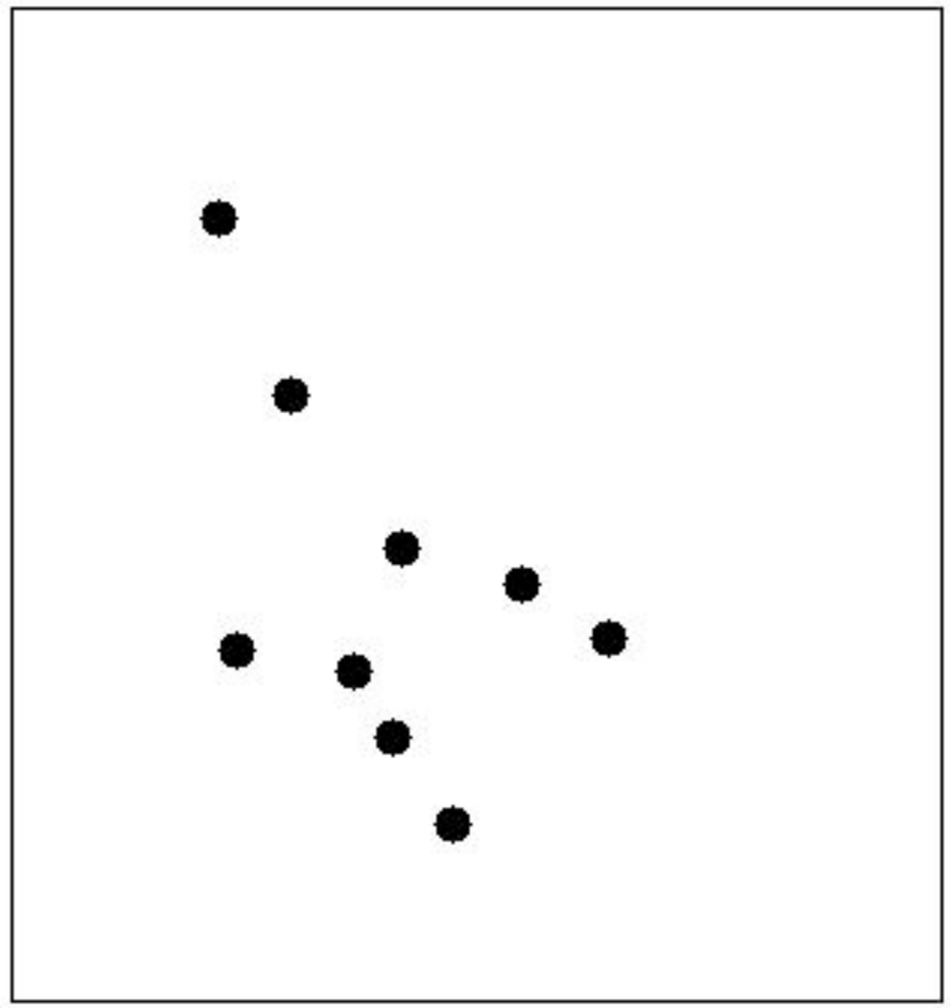

B

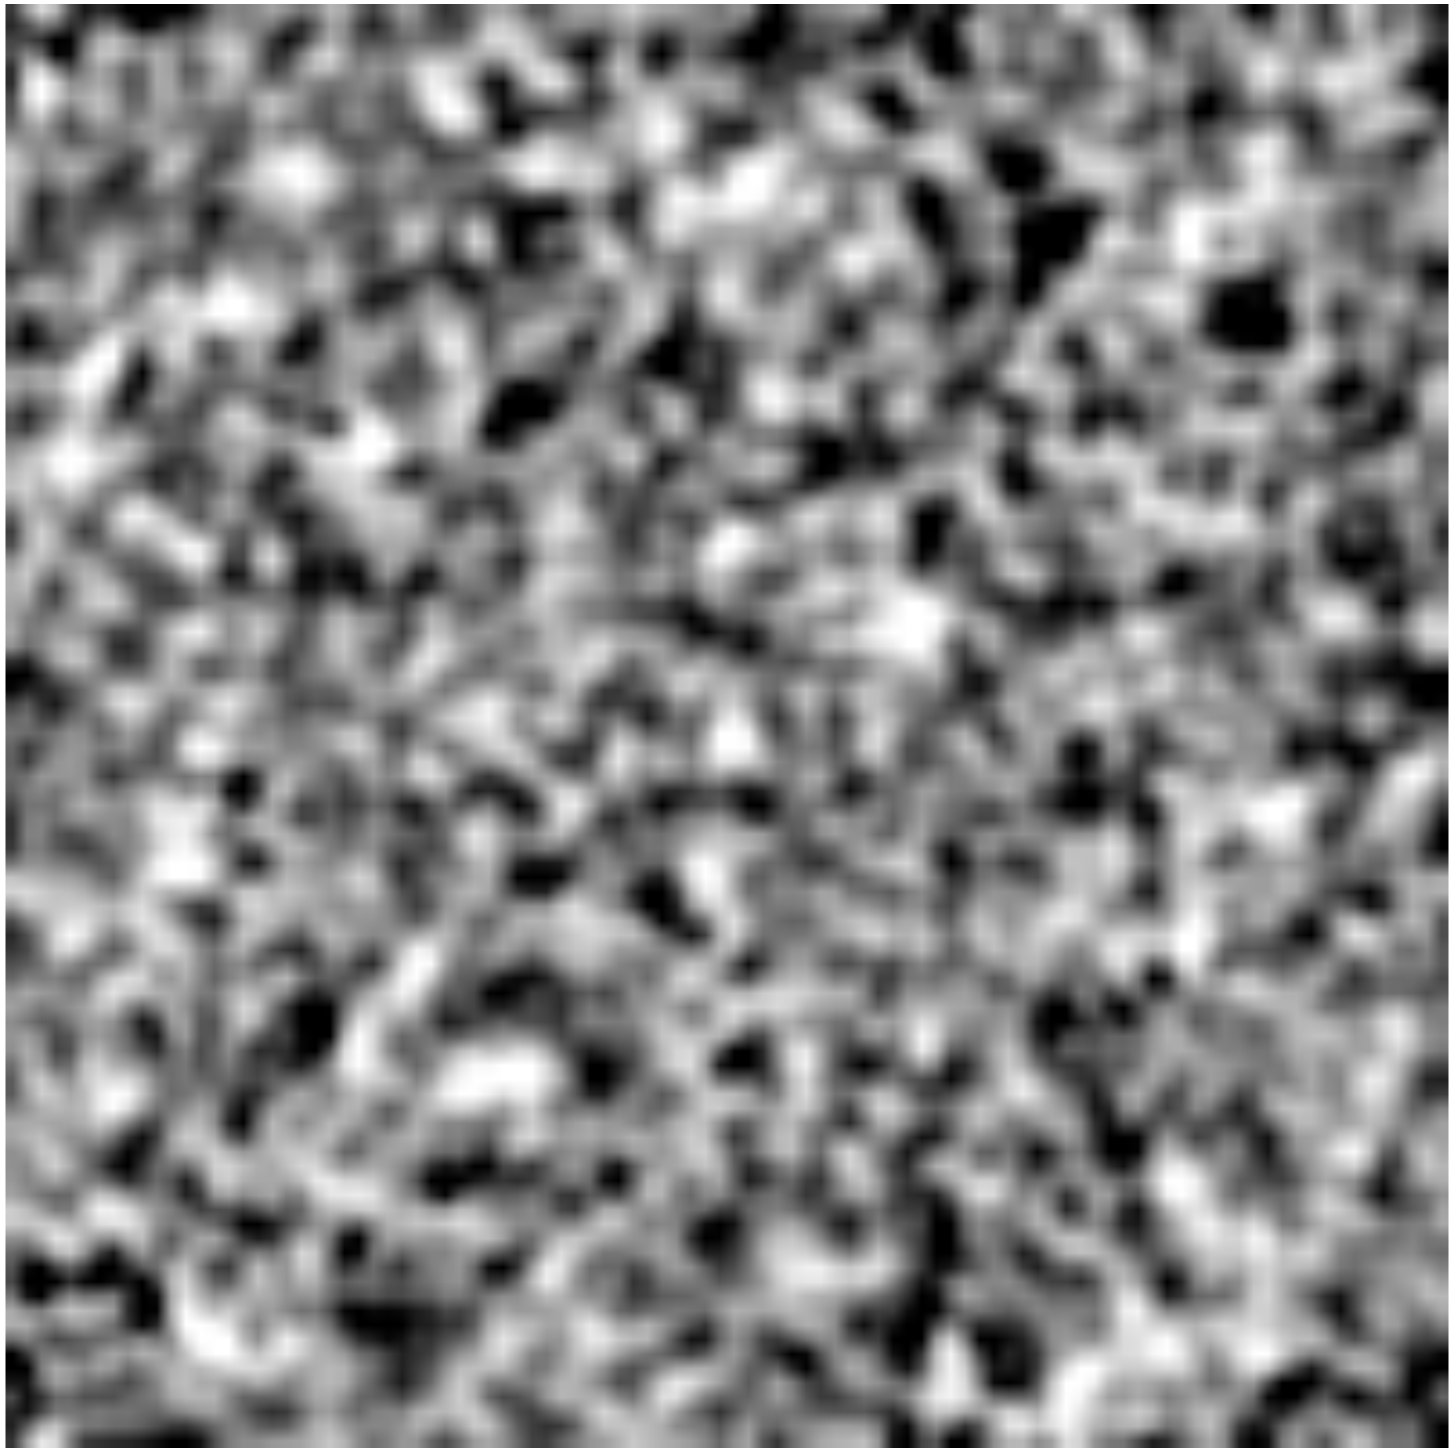

Category A or B?

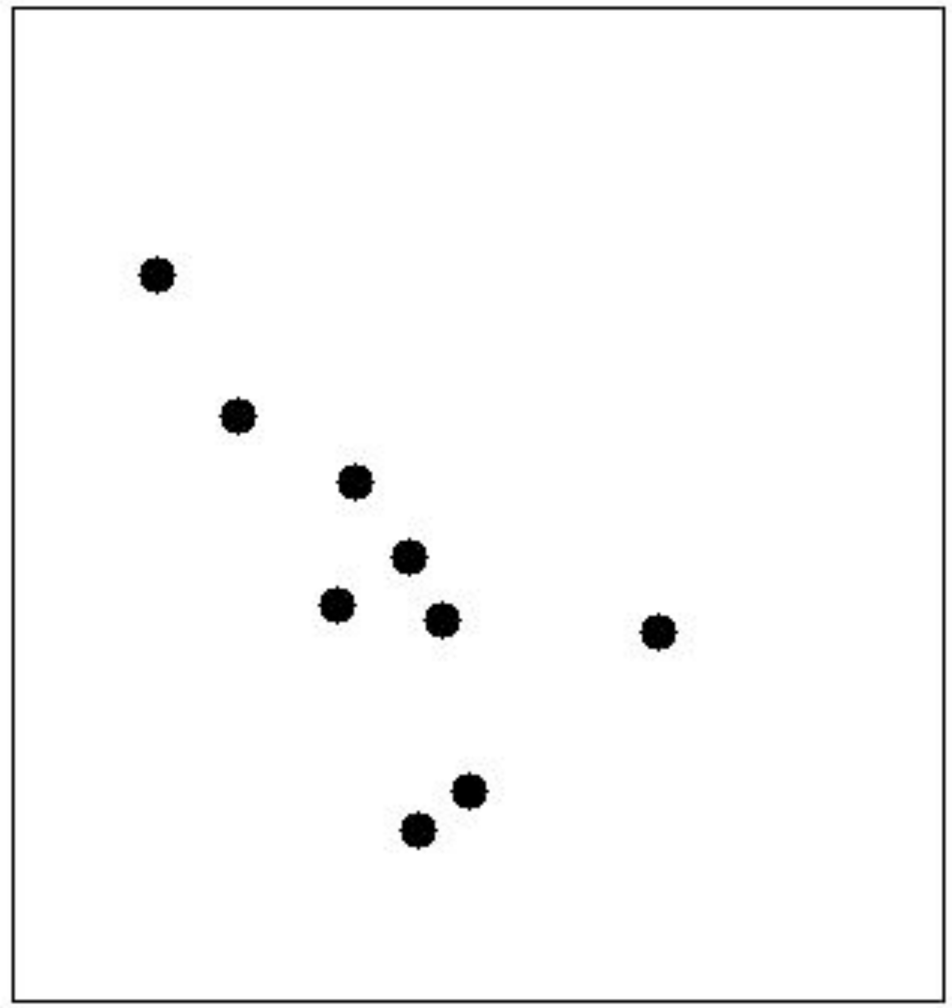

B

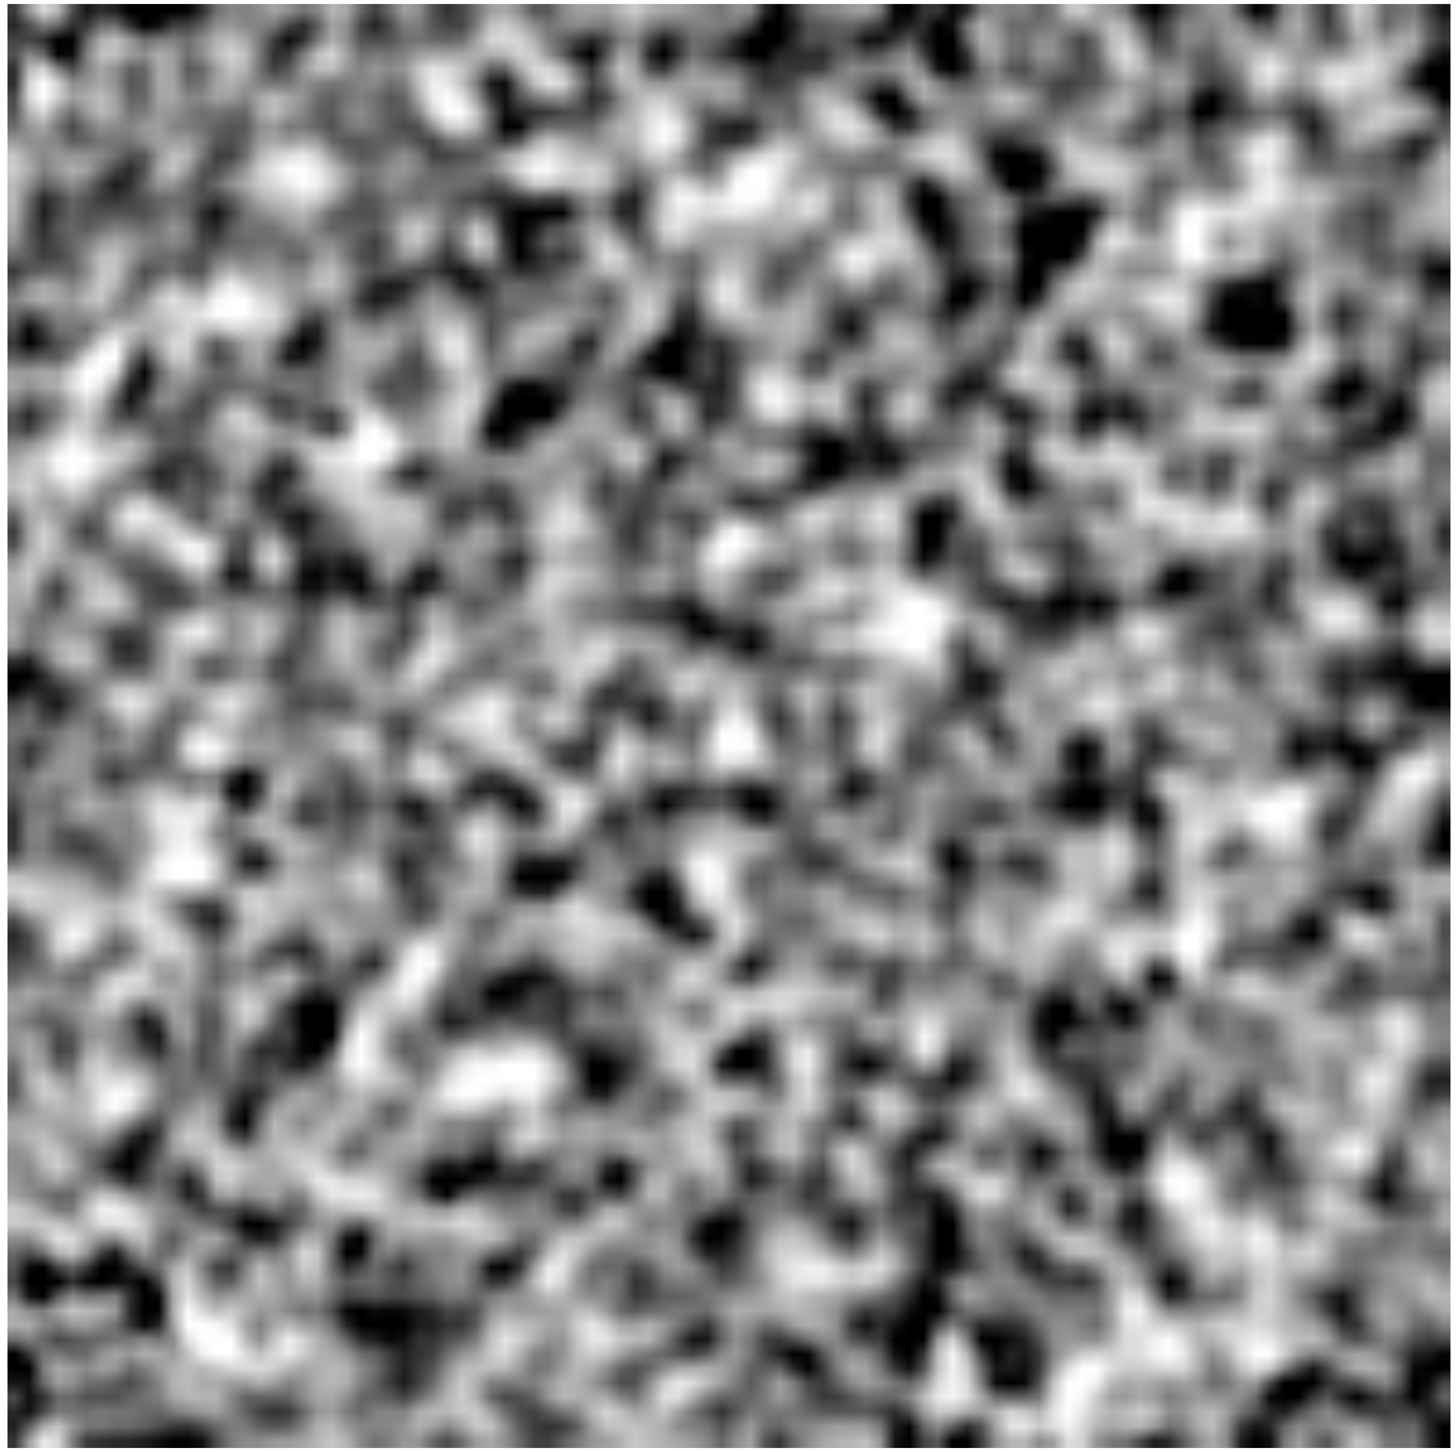

Category A or B?

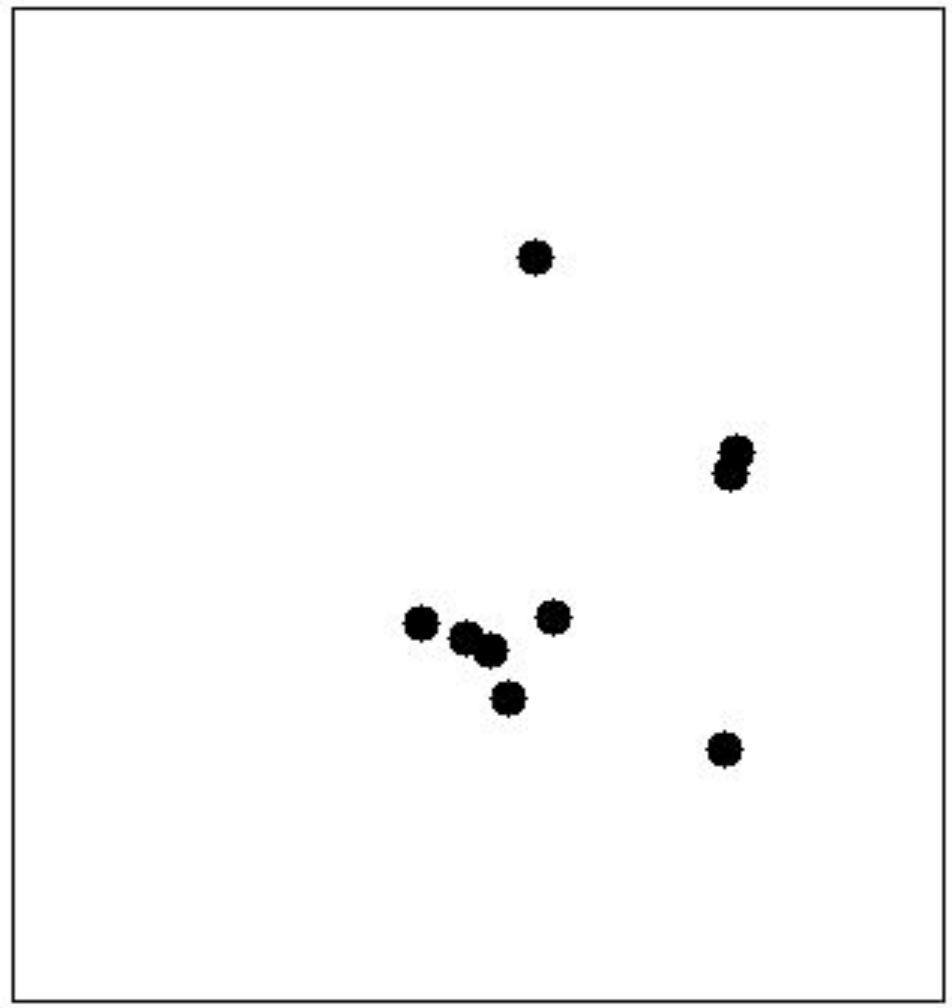

A

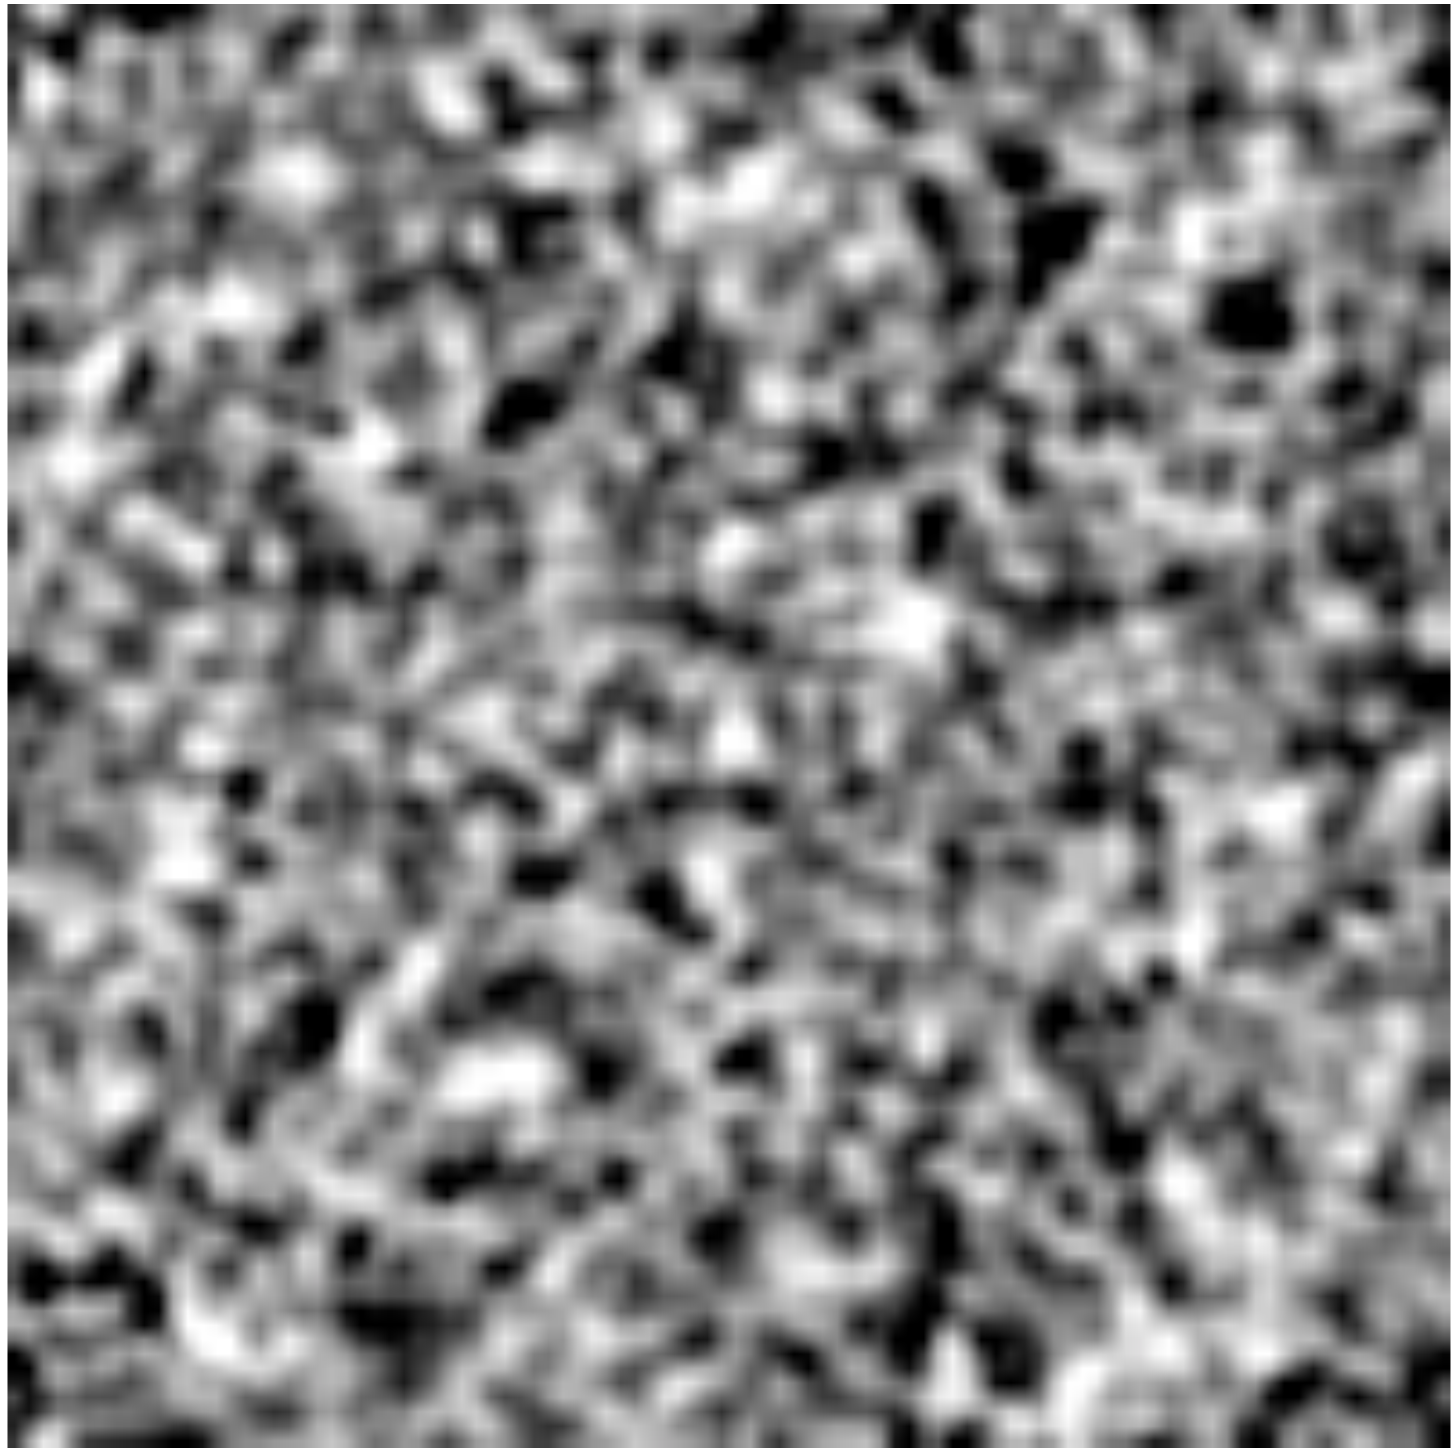

Category A or B?

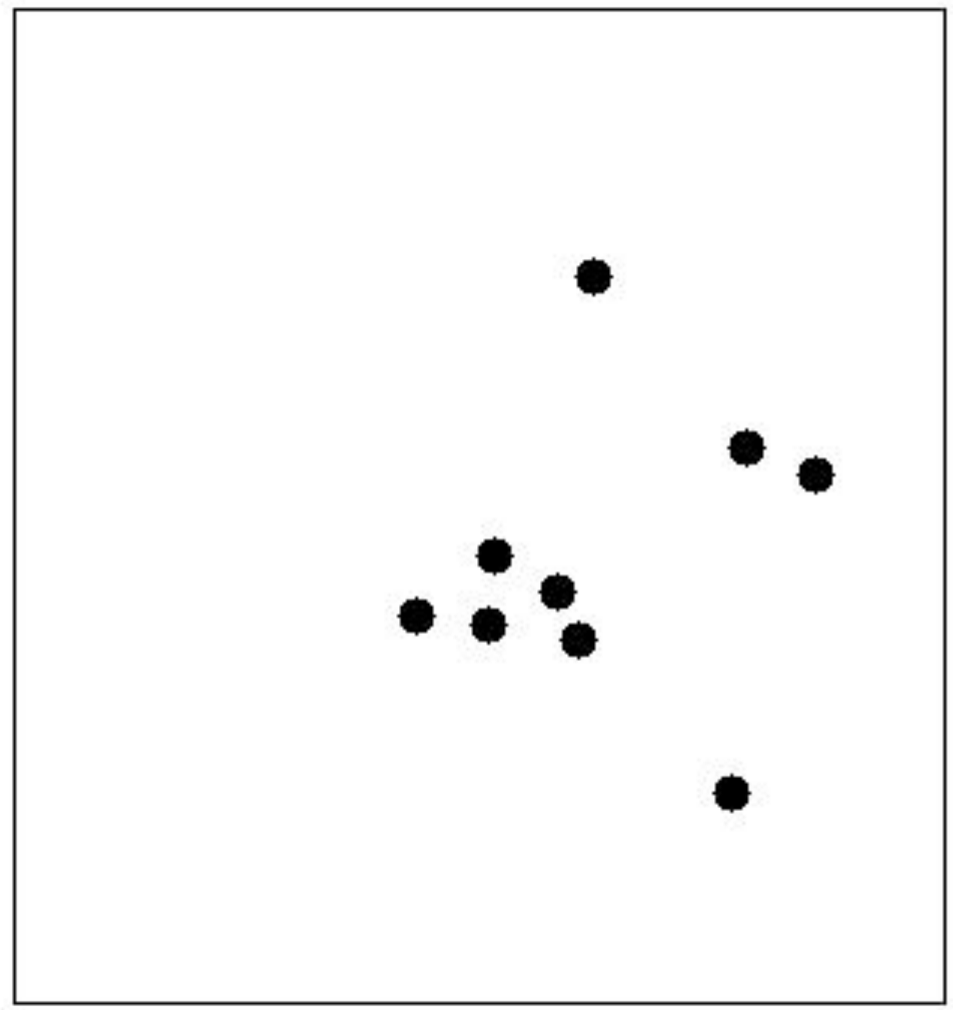

A

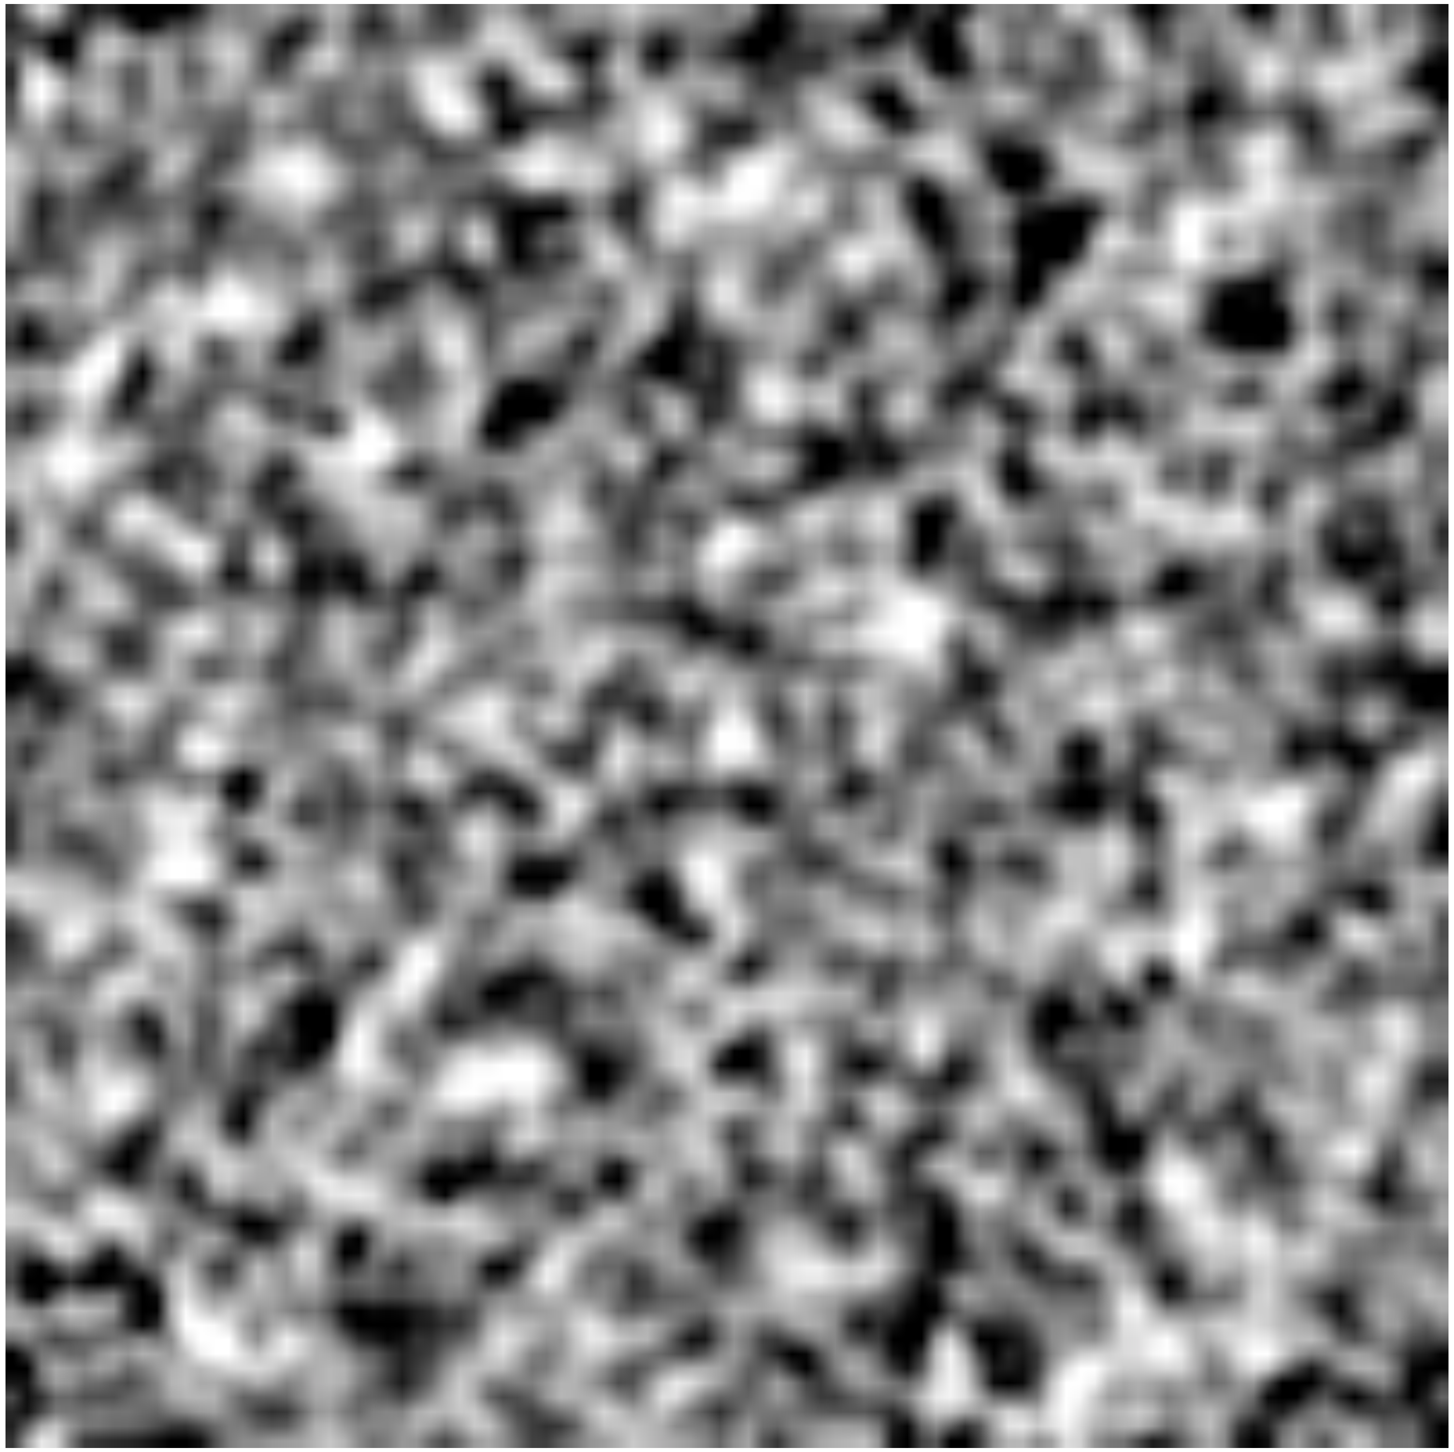

Category A or B?

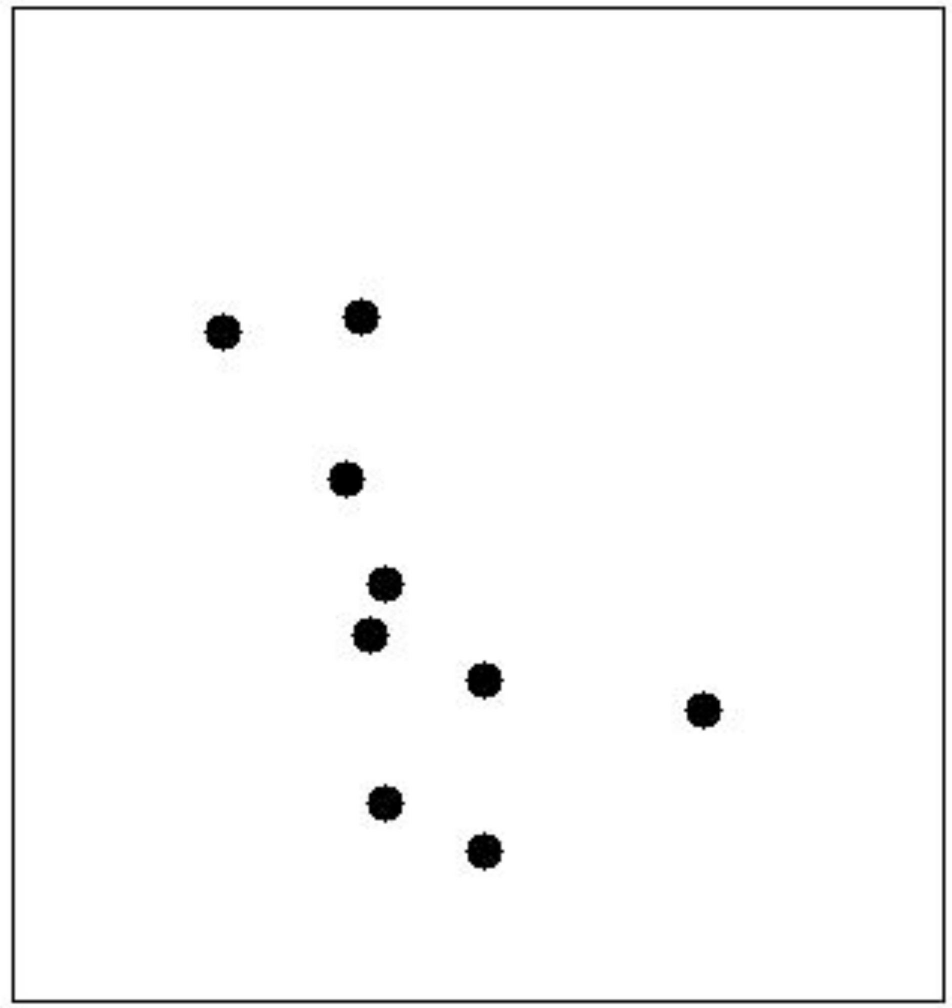

B

Training period is done. Now for testing…

Items seen during test period (after training)

(Posner & Keele, 1968)

After training, participants were tested on:
-- the prototypes (new)
-- some pattern distortions (old)
-- some pattern distortions (new)

Result:
(Accuracy for prototype = Accuracy for old distortions)

> Accuracy for new distortions

Suggests that some form of abstract representation is learned, like an "ideal image" or prototype. But also exemplars aren't lost/ forgotten

Prototype A: (not seen)

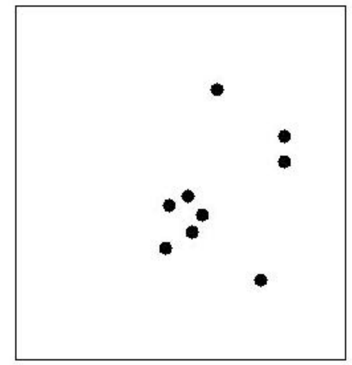

Prototype B: (not seen)

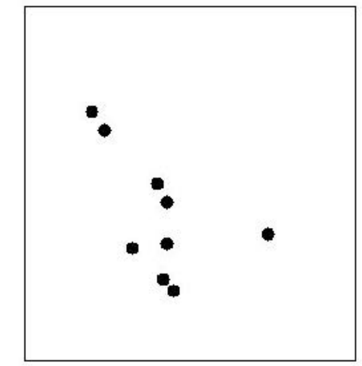

Distortions of A: Old

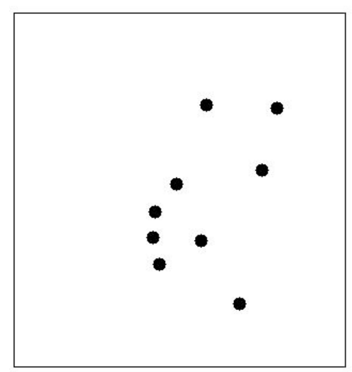

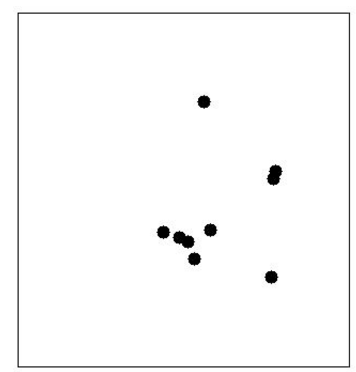

Distortions of B: Old

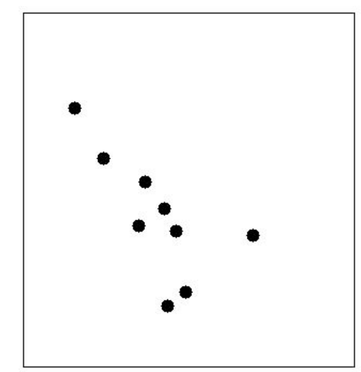

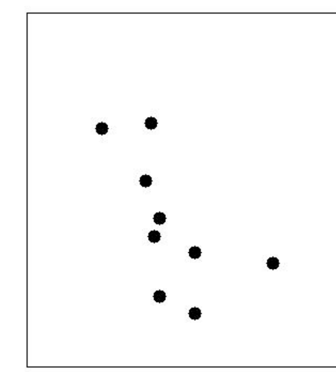

Distortions of A: New

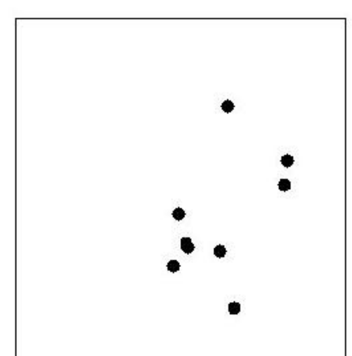

Distortions of B: New

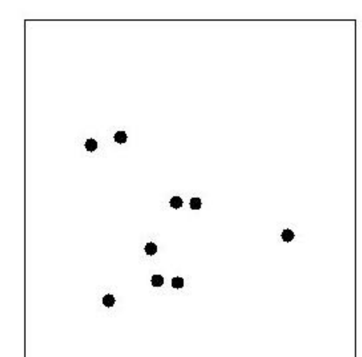

Another example…

The capacity of visual working memory for features and conjunctions

Steven J. Luck & Edward K. Vogel

Department of Psychology, University of Iowa, Iowa City, Iowa 52242-1407, USA

Short-term memory storage can be divided into separate subsystems for verbal information and visual information[1], and recent studies have begun to delineate the neural substrates of these working-memory systems[2–6]. Although the verbal storage system has been well characterized, the storage capacity of visual working memory has not yet been established for simple, suprathreshold features or for conjunctions of features. Here we demonstrate that it is possible to retain information about only four colours or orientations in visual working memory at one time. However, it is also possible to retain both the colour and the orientation of four objects, indicating that visual working memory stores integrated objects rather than individual features. Indeed, objects defined by a conjunction of four features can be retained in working memory just as well as single-feature objects, allowing sixteen individual features to be retained when distributed across four objects. Thus, the capacity of visual working memory must be understood in terms of integrated objects rather than individual features, which places significant constraints on cognitive and neurobiological models of the temporary storage of visual information[7].

To measure the capacity of working memory for simple features, we used a variant of the sequential comparison procedure developed by Phillips[8]. Subjects viewed a sample array and a test array on each trial, separated by a brief delay, and then indicated whether the two arrays were identical or differed in terms of a single feature. The accuracy of this discrimination was assessed as a function of the number of items in the stimulus array (the set size) to determine how many items could be accurately retained in working memory. In addition, control experiments were conducted to ensure that performance truly reflected the capacity of visual working memory and was not influenced by verbal working memory or by limitations in perception, memory encoding, or decision processes.

The first set of experiments examined working memory capacity for simple colours (Fig. 1a). The sample array consisted of 1–12 coloured squares and was presented for 100 ms. This was followed by a 900-ms blank delay interval and then a 2,000-ms presentation of the test array, which was either identical to the sample array or differed in the colour of one of the squares. Performance was nearly limitations in processes other than working-memory storage. To rule out limitations in perceiving the stimuli and encoding them in working memory, we varied the duration of the sample stimulus, comparing the original 100-ms duration with a 500-ms duration. This allowed substantially more time for perceiving the stimuli and encoding them in memory, which should have led to improved performance if these were limiting factors. However, performance was not significantly influenced by variations in sample duration (Fig. 1b), indicating that the errors at set sizes of 4–12 reflected limitations in storage capacity rather than limitations in perceiving or encoding the stimuli.

We next examined the possibility that performance was limited by decision factors. At larger set sizes, more decisions must be made, and this can lead to an increase in errors even in the absence of any capacity limitations[10,11]. To rule out this explanation, we conducted an experiment in which the memory requirements were the same as in the original experiment but only a single decision was necessary, regardless of the set size. Specifically, we used a partial report procedure in which we cued the observers to make a decision about only one of the items in the test array by presenting an outline box around the one item that might have been different from the sample array. This required them to retain information from all of the items in the sample array, but allowed them to restrict decision processes to a single item in the test array. As shown in Fig. 1b, this manipulation did not significantly alter performance, indicating that accuracy was not limited by decision factors (or, alternatively, that the subjects were unable to use the cue box effectively, which seems unlikely given that previous studies have found similar cues to be very effective in improving performance in decision-limited tasks[12,13]).

To determine whether capacity is different for different feature dimensions, memory for orientation was compared with memory for colour using 4, 8 or 12 bars that varied both in colour and in orientation. The observers were instructed to detect either colour changes or orientation changes (in different trial blocks), and a verbal load was used in both cases. The effects of set size on accuracy were nearly identical for colour and orientation, with a capacity of about four items for both feature types.

We then assessed whether visual information is stored in working memory as individual features or as integrated objects. This was tested by comparing memory for simple features with memory for objects defined by a conjunction of features. Observers performed the same sequential comparison task used above (while performing a concurrent verbal load task) with arrays of 2, 4 or 6 coloured bars of varying orientations. Relatively small set sizes were used so that the objects could be widely spaced, which was necessary to avoid 'illusory conjunctions' in the perception of the bars[14]. In one condition, only colour could vary between the sample array and the test array, and the observers were instructed to look for a colour change. In a second condition, only orientation could vary, and the

Get ready…

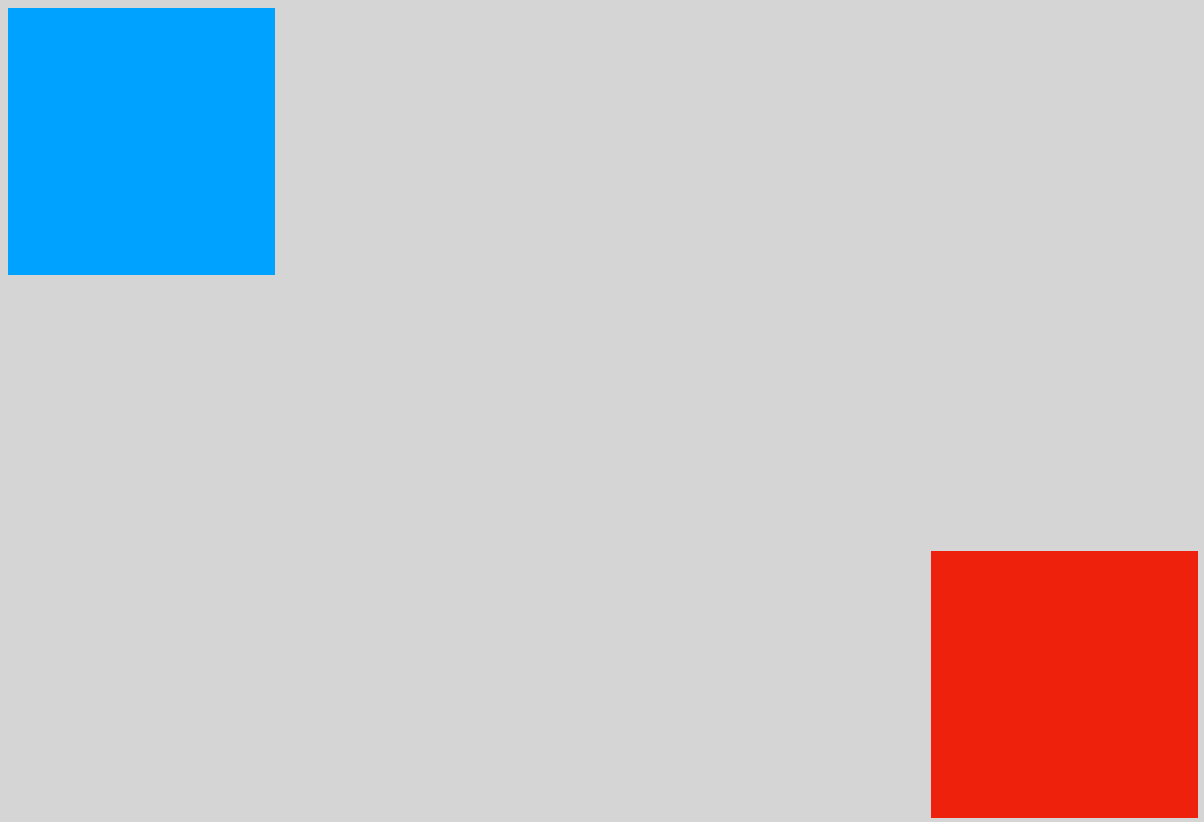

Did anything change?

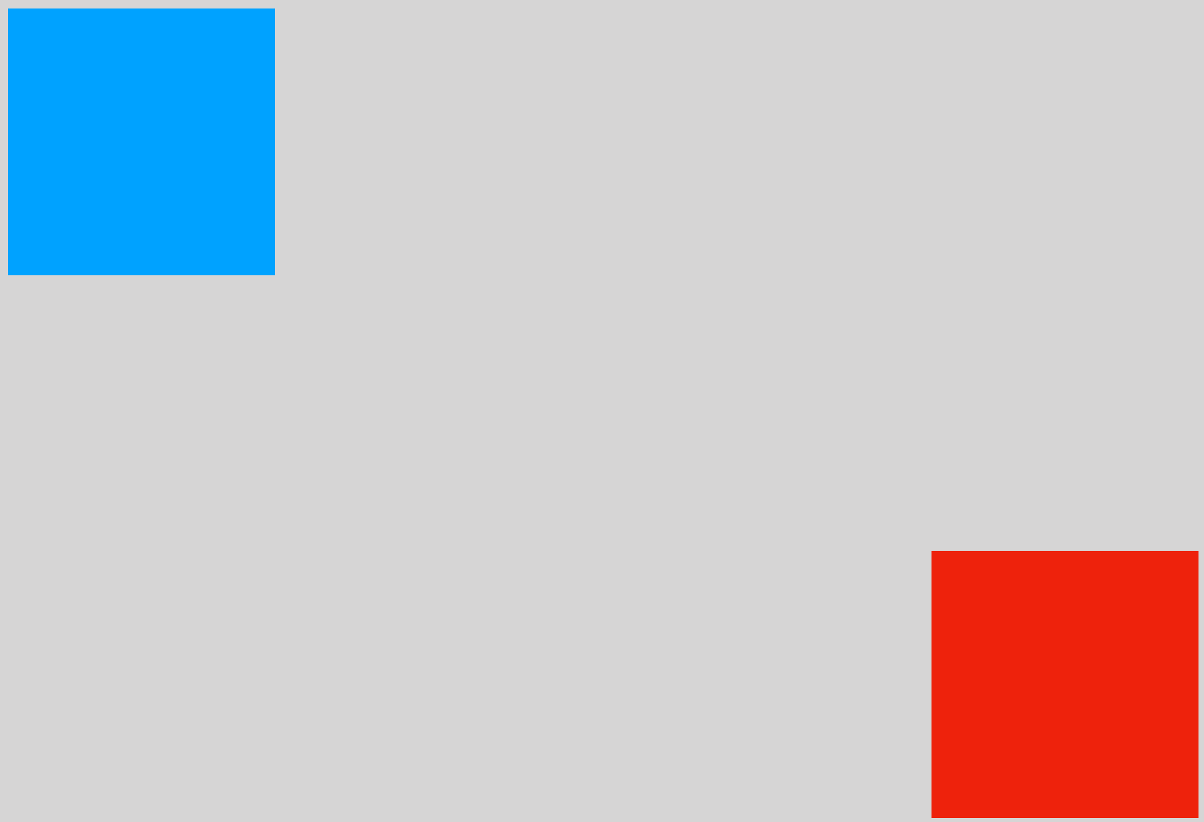

Get ready…

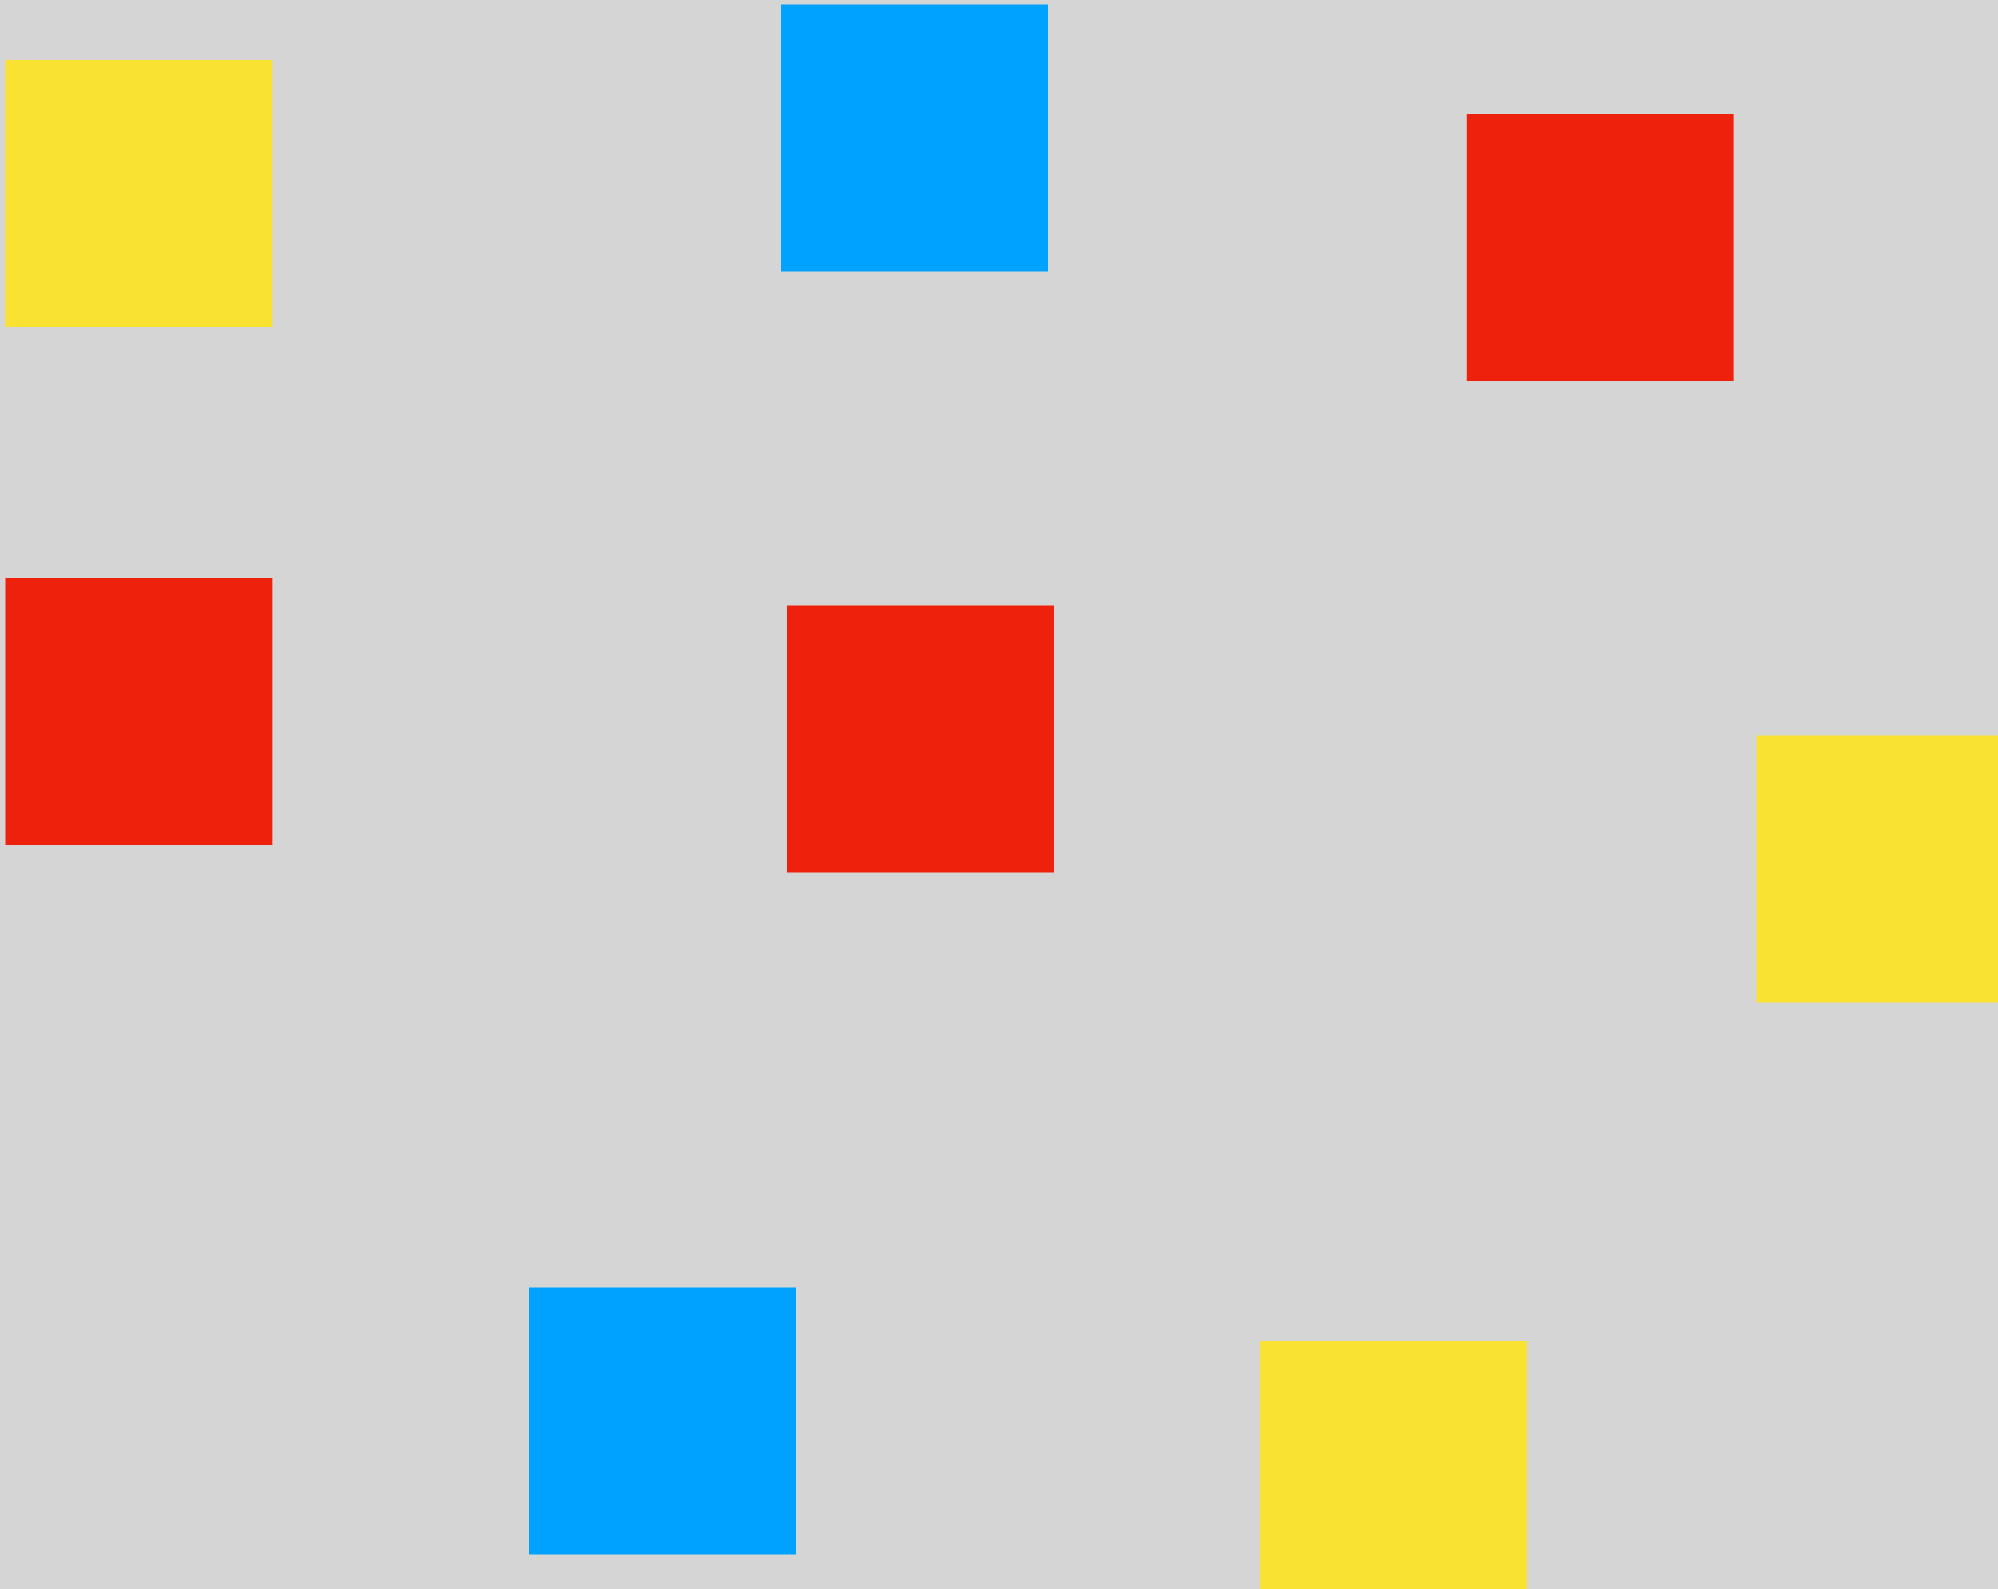

Did anything change?

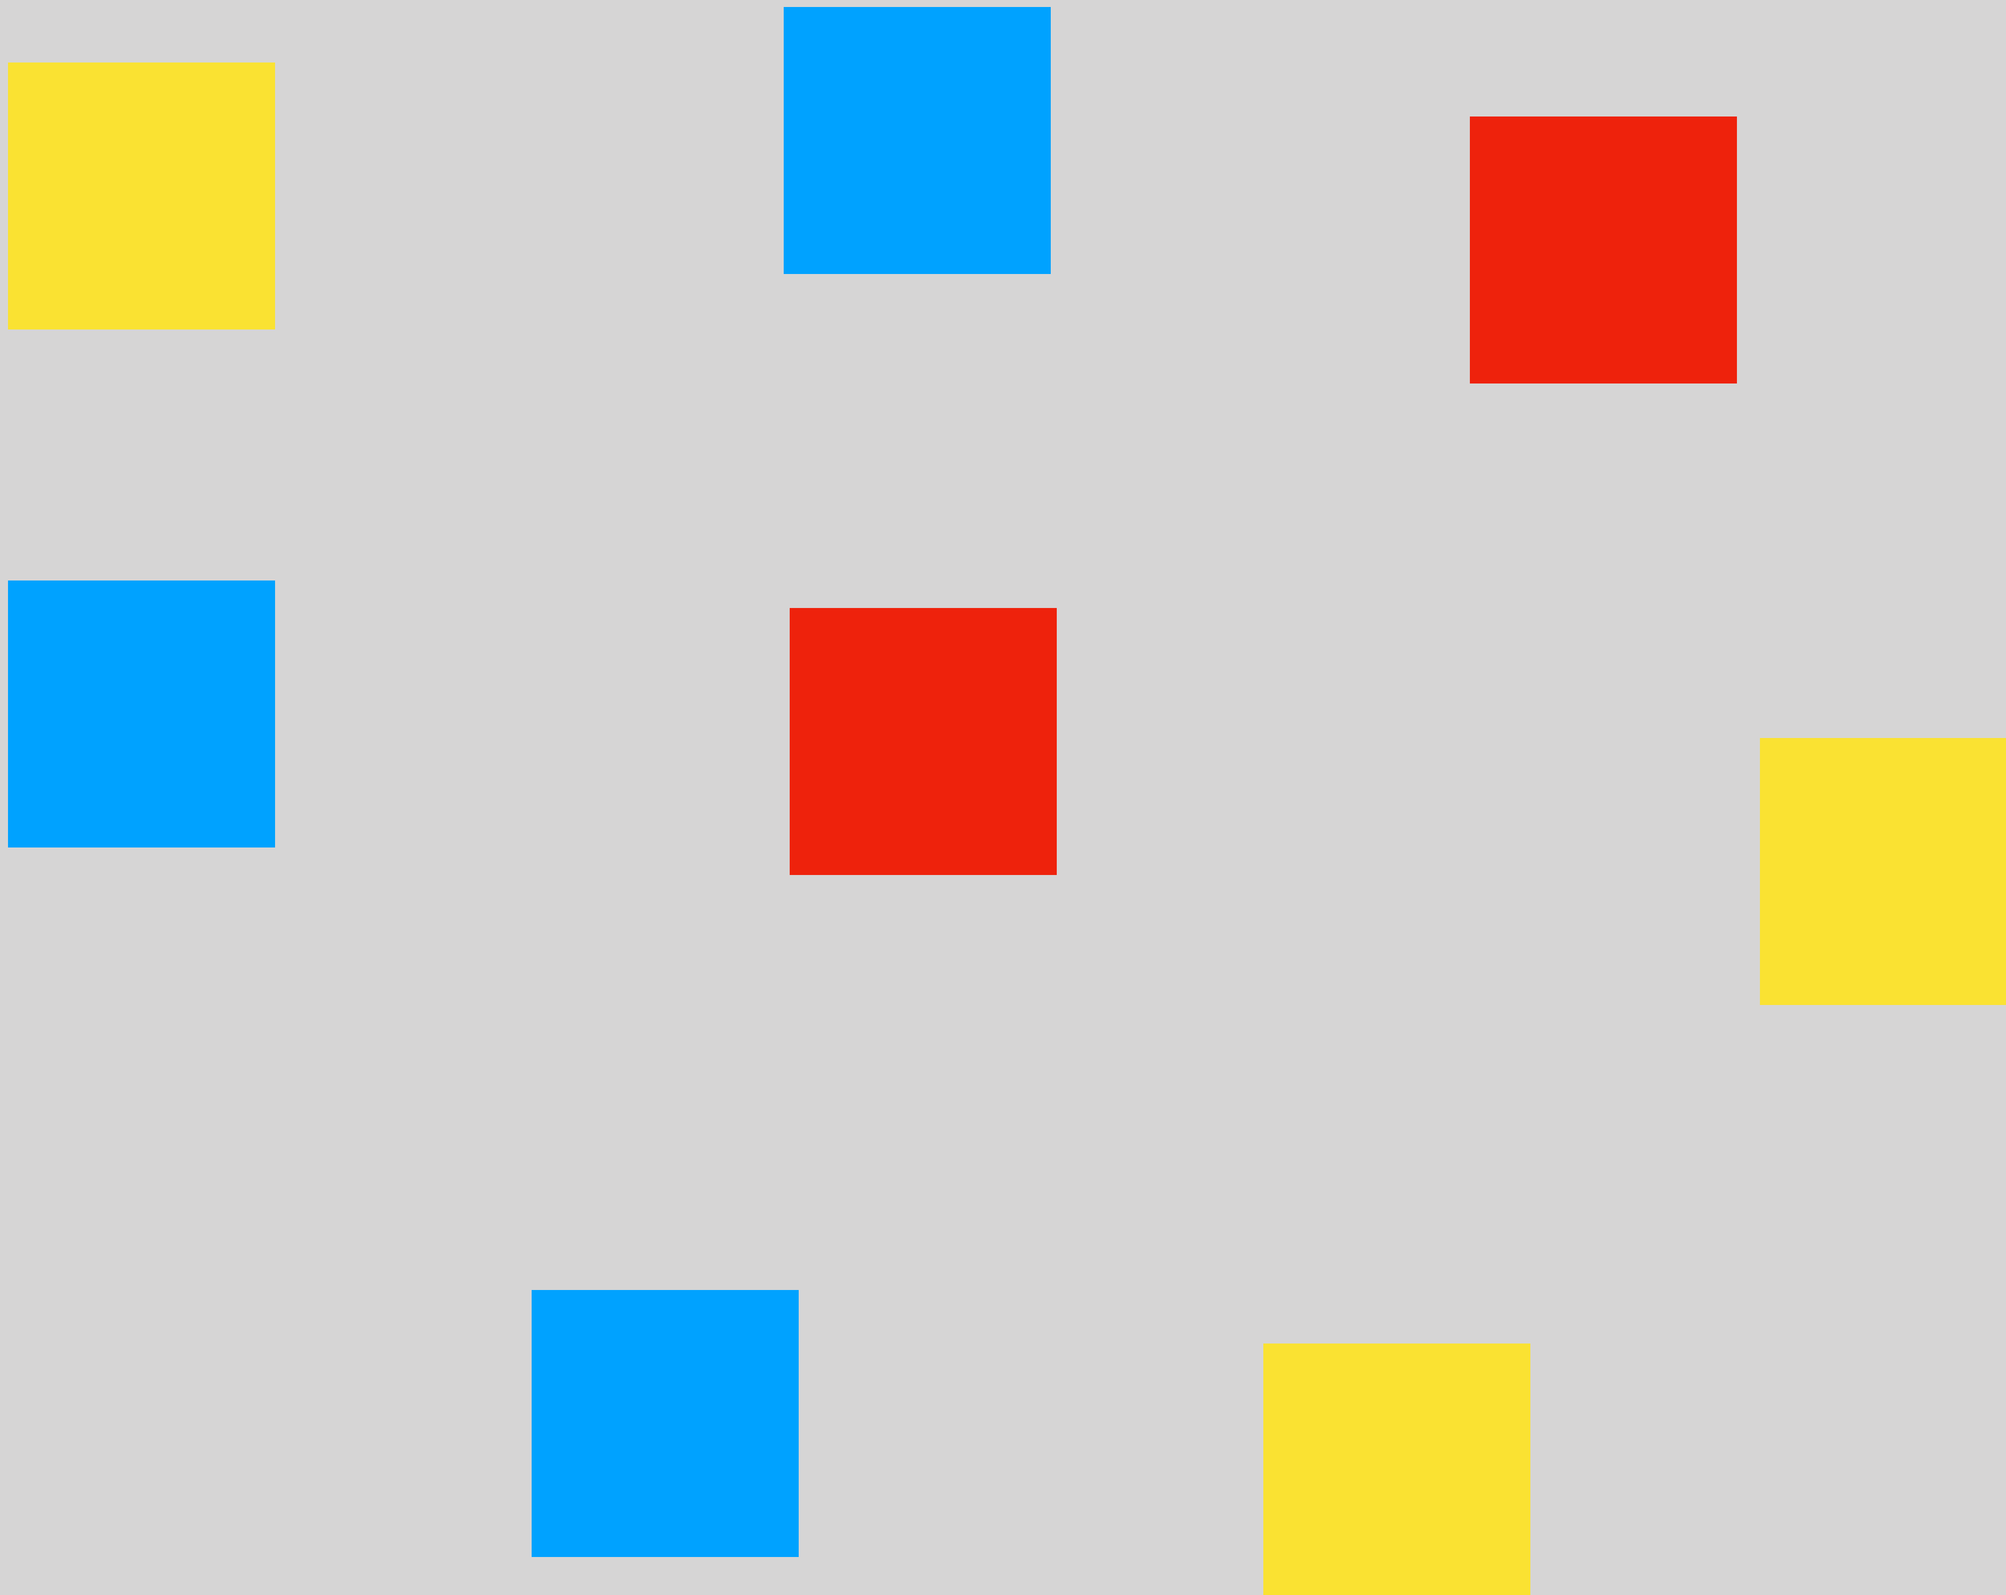

Results and analysis

Performance nearly perfect for 3 items, starts to drop at 4

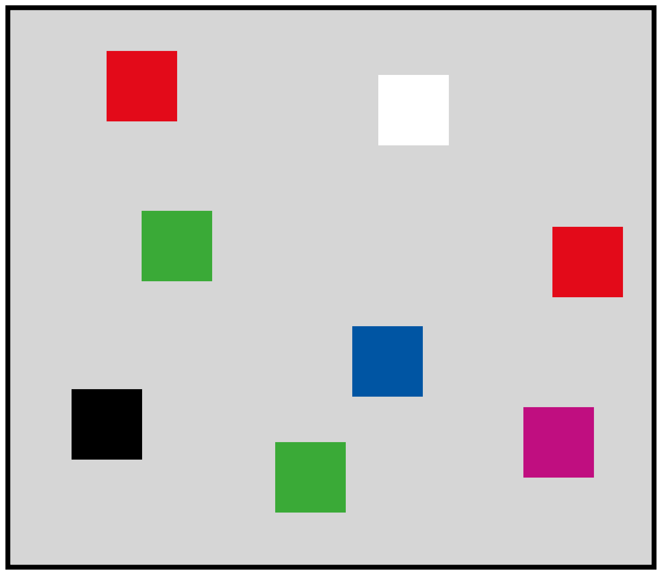

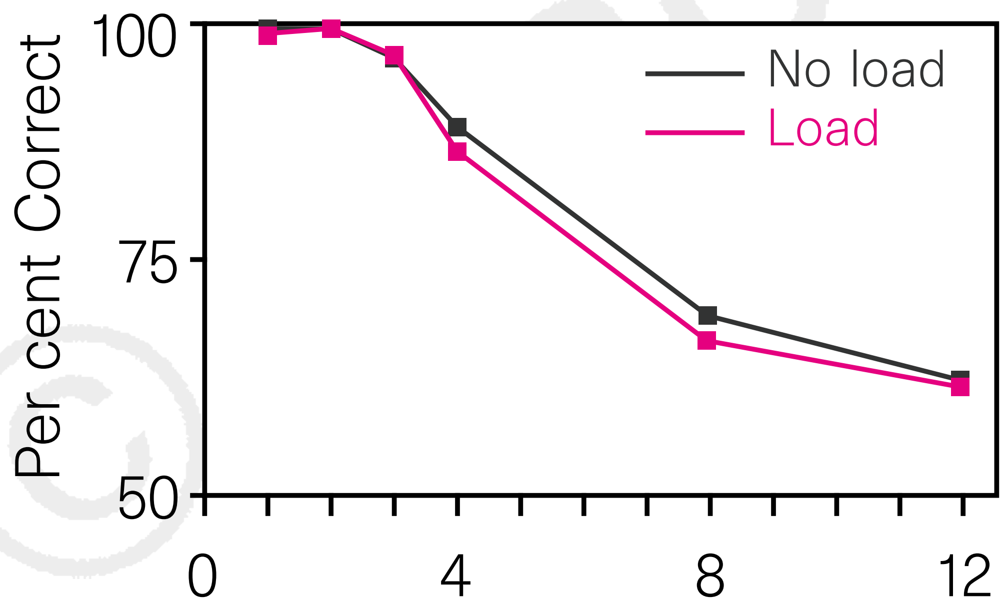

Remembering two features per item is no harder than one feature

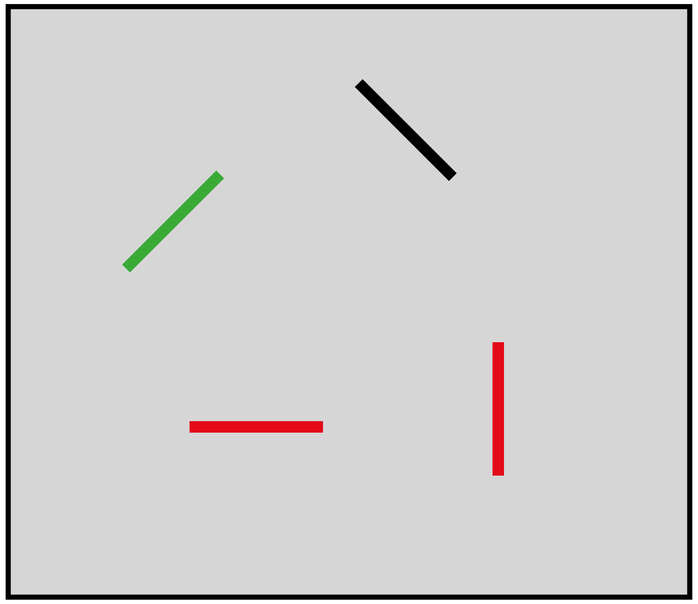

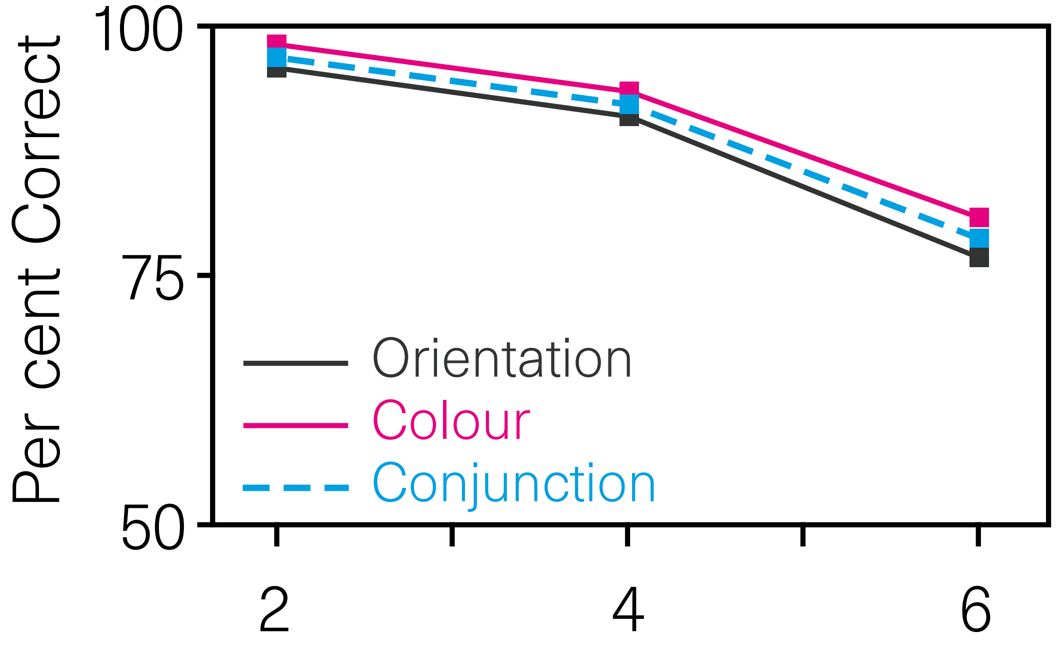

four slot theory

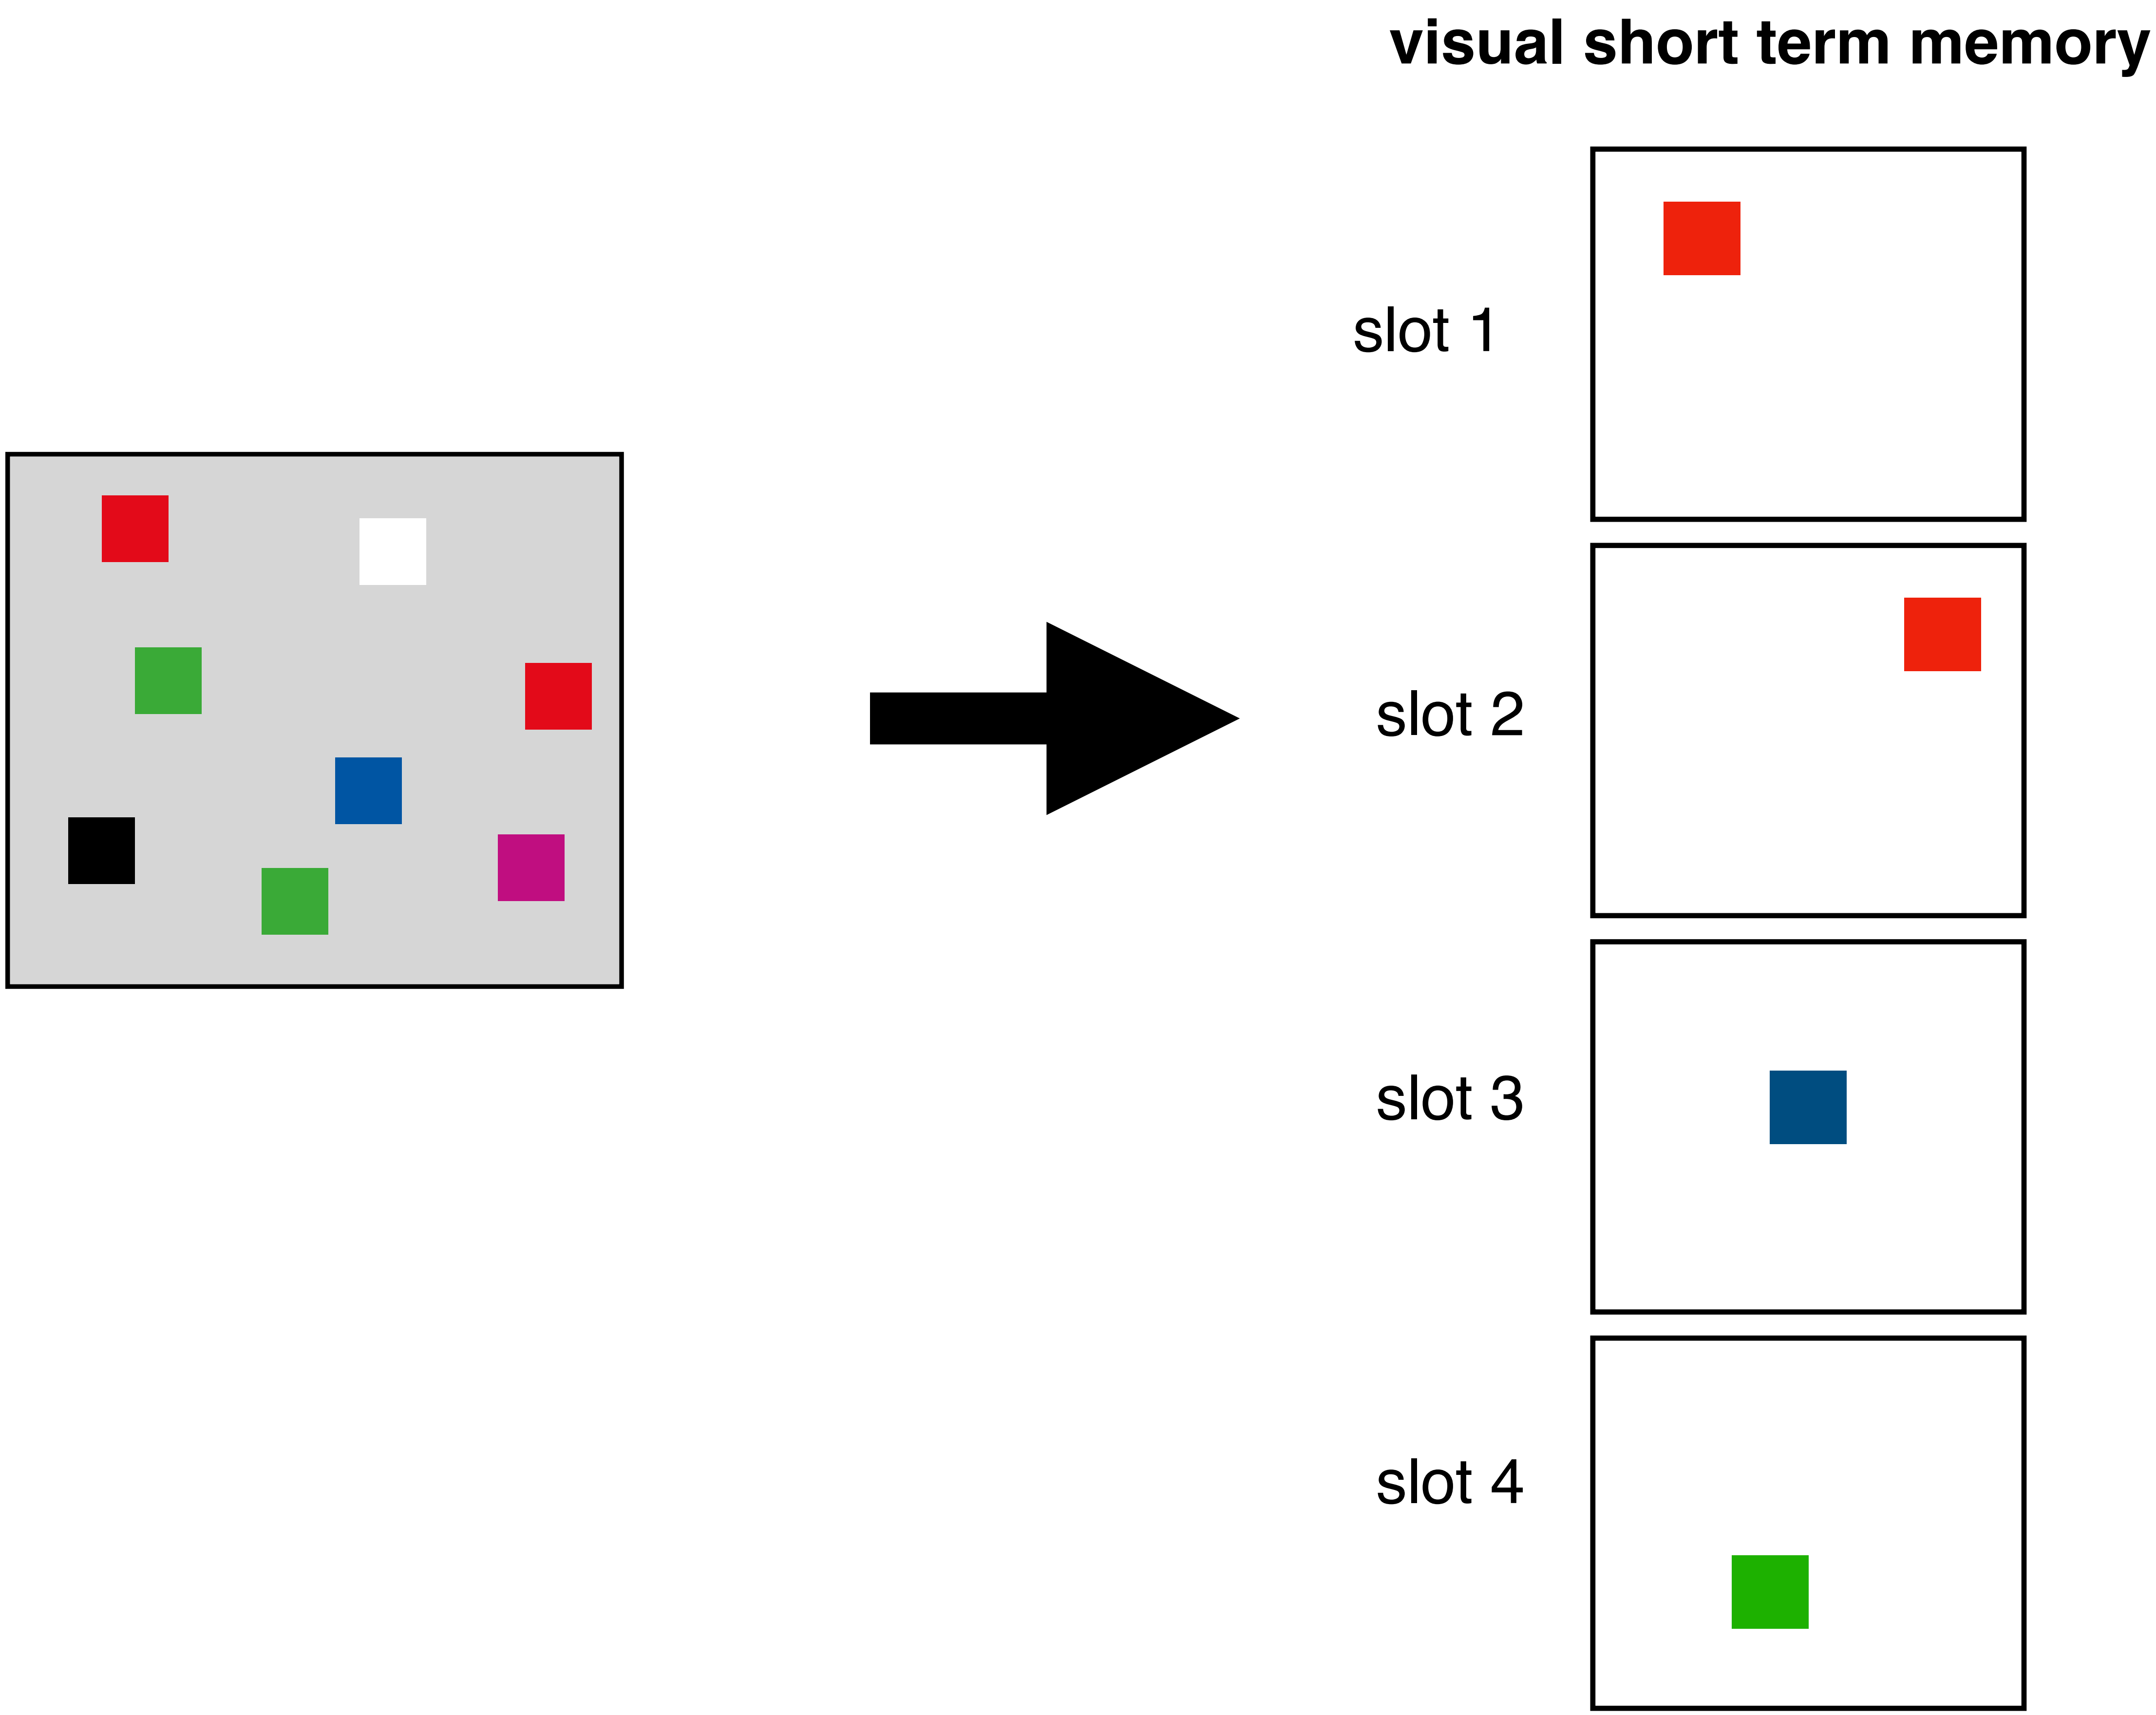

A detection theory account of change detection

Patrick Wilken

Division of Biology, California Institute of Technology, Pasadena, CA, USA

Wei Ji Ma

Division of Biology, California Institute of Technology, Pasadena, CA, USA

Previous studies have suggested that visual short-term memory (VSTM) has a storage limit of approximately four items. However, the type of high-threshold (HT) model used to derive this estimate is based on a number of assumptions that have been criticized in other experimental paradigms (e.g., visual search). Here we report findings from nine experiments in which VSTM for color, spatial frequency, and orientation was modeled using a signal detection theory (SDT) approach. In Experiments 1-6, two arrays composed of multiple stimulus elements were presented for 100 ms with a 1500 ms ISI. Observers were asked to report in a yes/no fashion whether there was any difference between the first and second arrays, and to rate their confidence in their response on a 1-4 scale. In Experiments 1-3, only one stimulus element difference could occur (T = 1) while set size was varied. In Experiments 4-6, set size was fixed while the number of stimuli that might change was varied (T = 1, 2, 3, and 4). Three general models were tested against the receiver operating characteristics generated by the six experiments. In addition to the HT model, two SDT models were tried: one assuming summation of signals prior to a decision, the other using a max rule. In Experiments 7-9, observers were asked to directly report the relevant feature attribute of a stimulus presented 1500 ms previously, from an array of varying set size. Overall, the results suggest that observers encode stimuli independently and in parallel, and that performance is limited by internal noise, which is a function of set size.

Keywords: feature judgment, visual short-term memory (VSTM), signal detection theory, change blindness, high-threshold theory, capacity limitations

Introduction

A critical aspect of any creature's ability to function effectively within a changing environment is the facility to efficiently utilize information from a variety of sensory sources in both its present and its immediate past. The high evolutionary value of such information is implied by the ability of human observers to store various perceptual dimensions, such as spatial frequency, orientation, and hue, with a high degree of fidelity and stability over extended periods of time (Magnussen & Greenlee, 1992; Magnussen, Greenlee, Asplund, & Dyrnes, 1991; Magnussen, Greenlee, & Thomas, 1996; Regan, 1985). It has been shown, for instance, that observers are readily able to detect spatial frequency changes for time periods of upwards of 60 s that are smaller than the Nyquist frequency associated with the spacing between adjacent cones on the fovea (Magnussen, Greenlee, Asplund, & Dyrnes, 1990; Regan, 1985).

In a typical *visual short-term memory* (VSTM) experiment,

ments present leads to a monotonic decrease in the sensitivity of observers to differences between the two displays; although for experiments employing suprathreshold stimuli, this decrease is typically only observed after set size has reached around three to four elements (Luck & Vogel, 1997; Pashler, 1988; Vogel, Woodman, & Luck, 2001).

A prominent class of VSTM model proposes that the performance decline associated with increasing set size is caused by a fundamental limit of the number of items that can be encoded, either because the capacity of VSTM itself is limited (Cowan, 2001; Luck & Vogel, 1997; Pashler, 1988; Vogel et al., 2001), or because of a bottleneck in the number of items that can be attended to during the encoding process (Rensink, 2000).

This type of model assumes that VSTM is restricted in storage capacity to only a few items, C (often estimated to lie in the range of 4 to 5), within a set size N (Pashler, 1988). The probability that a suprathreshold change will be reported (H) is then

Color judgement experiment

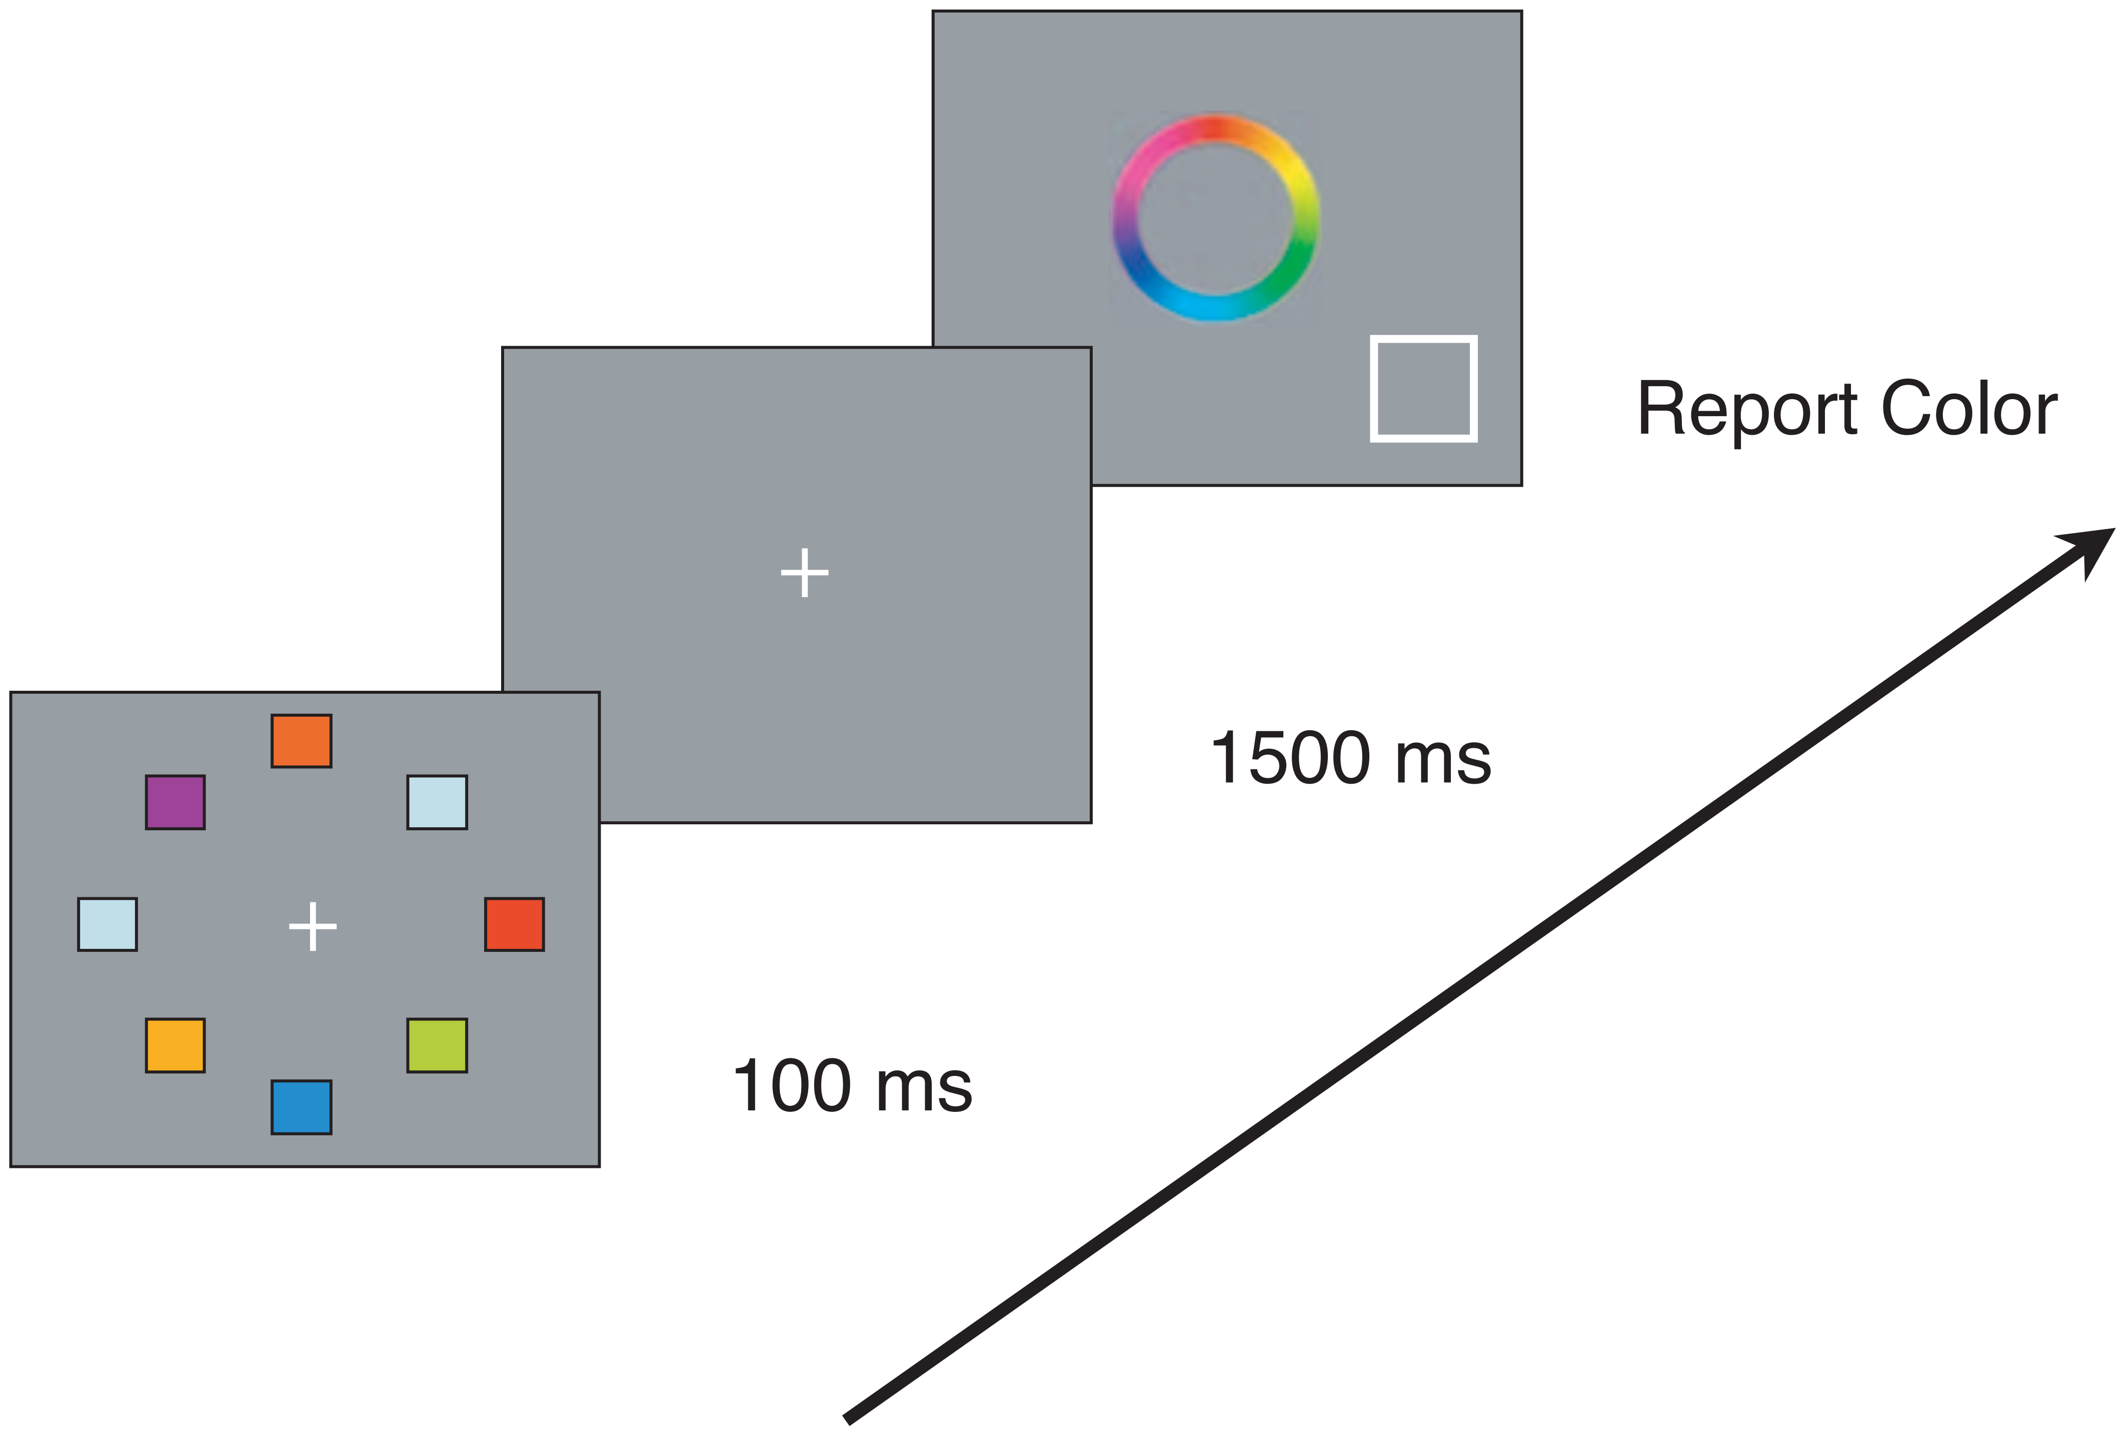

Get ready…

+

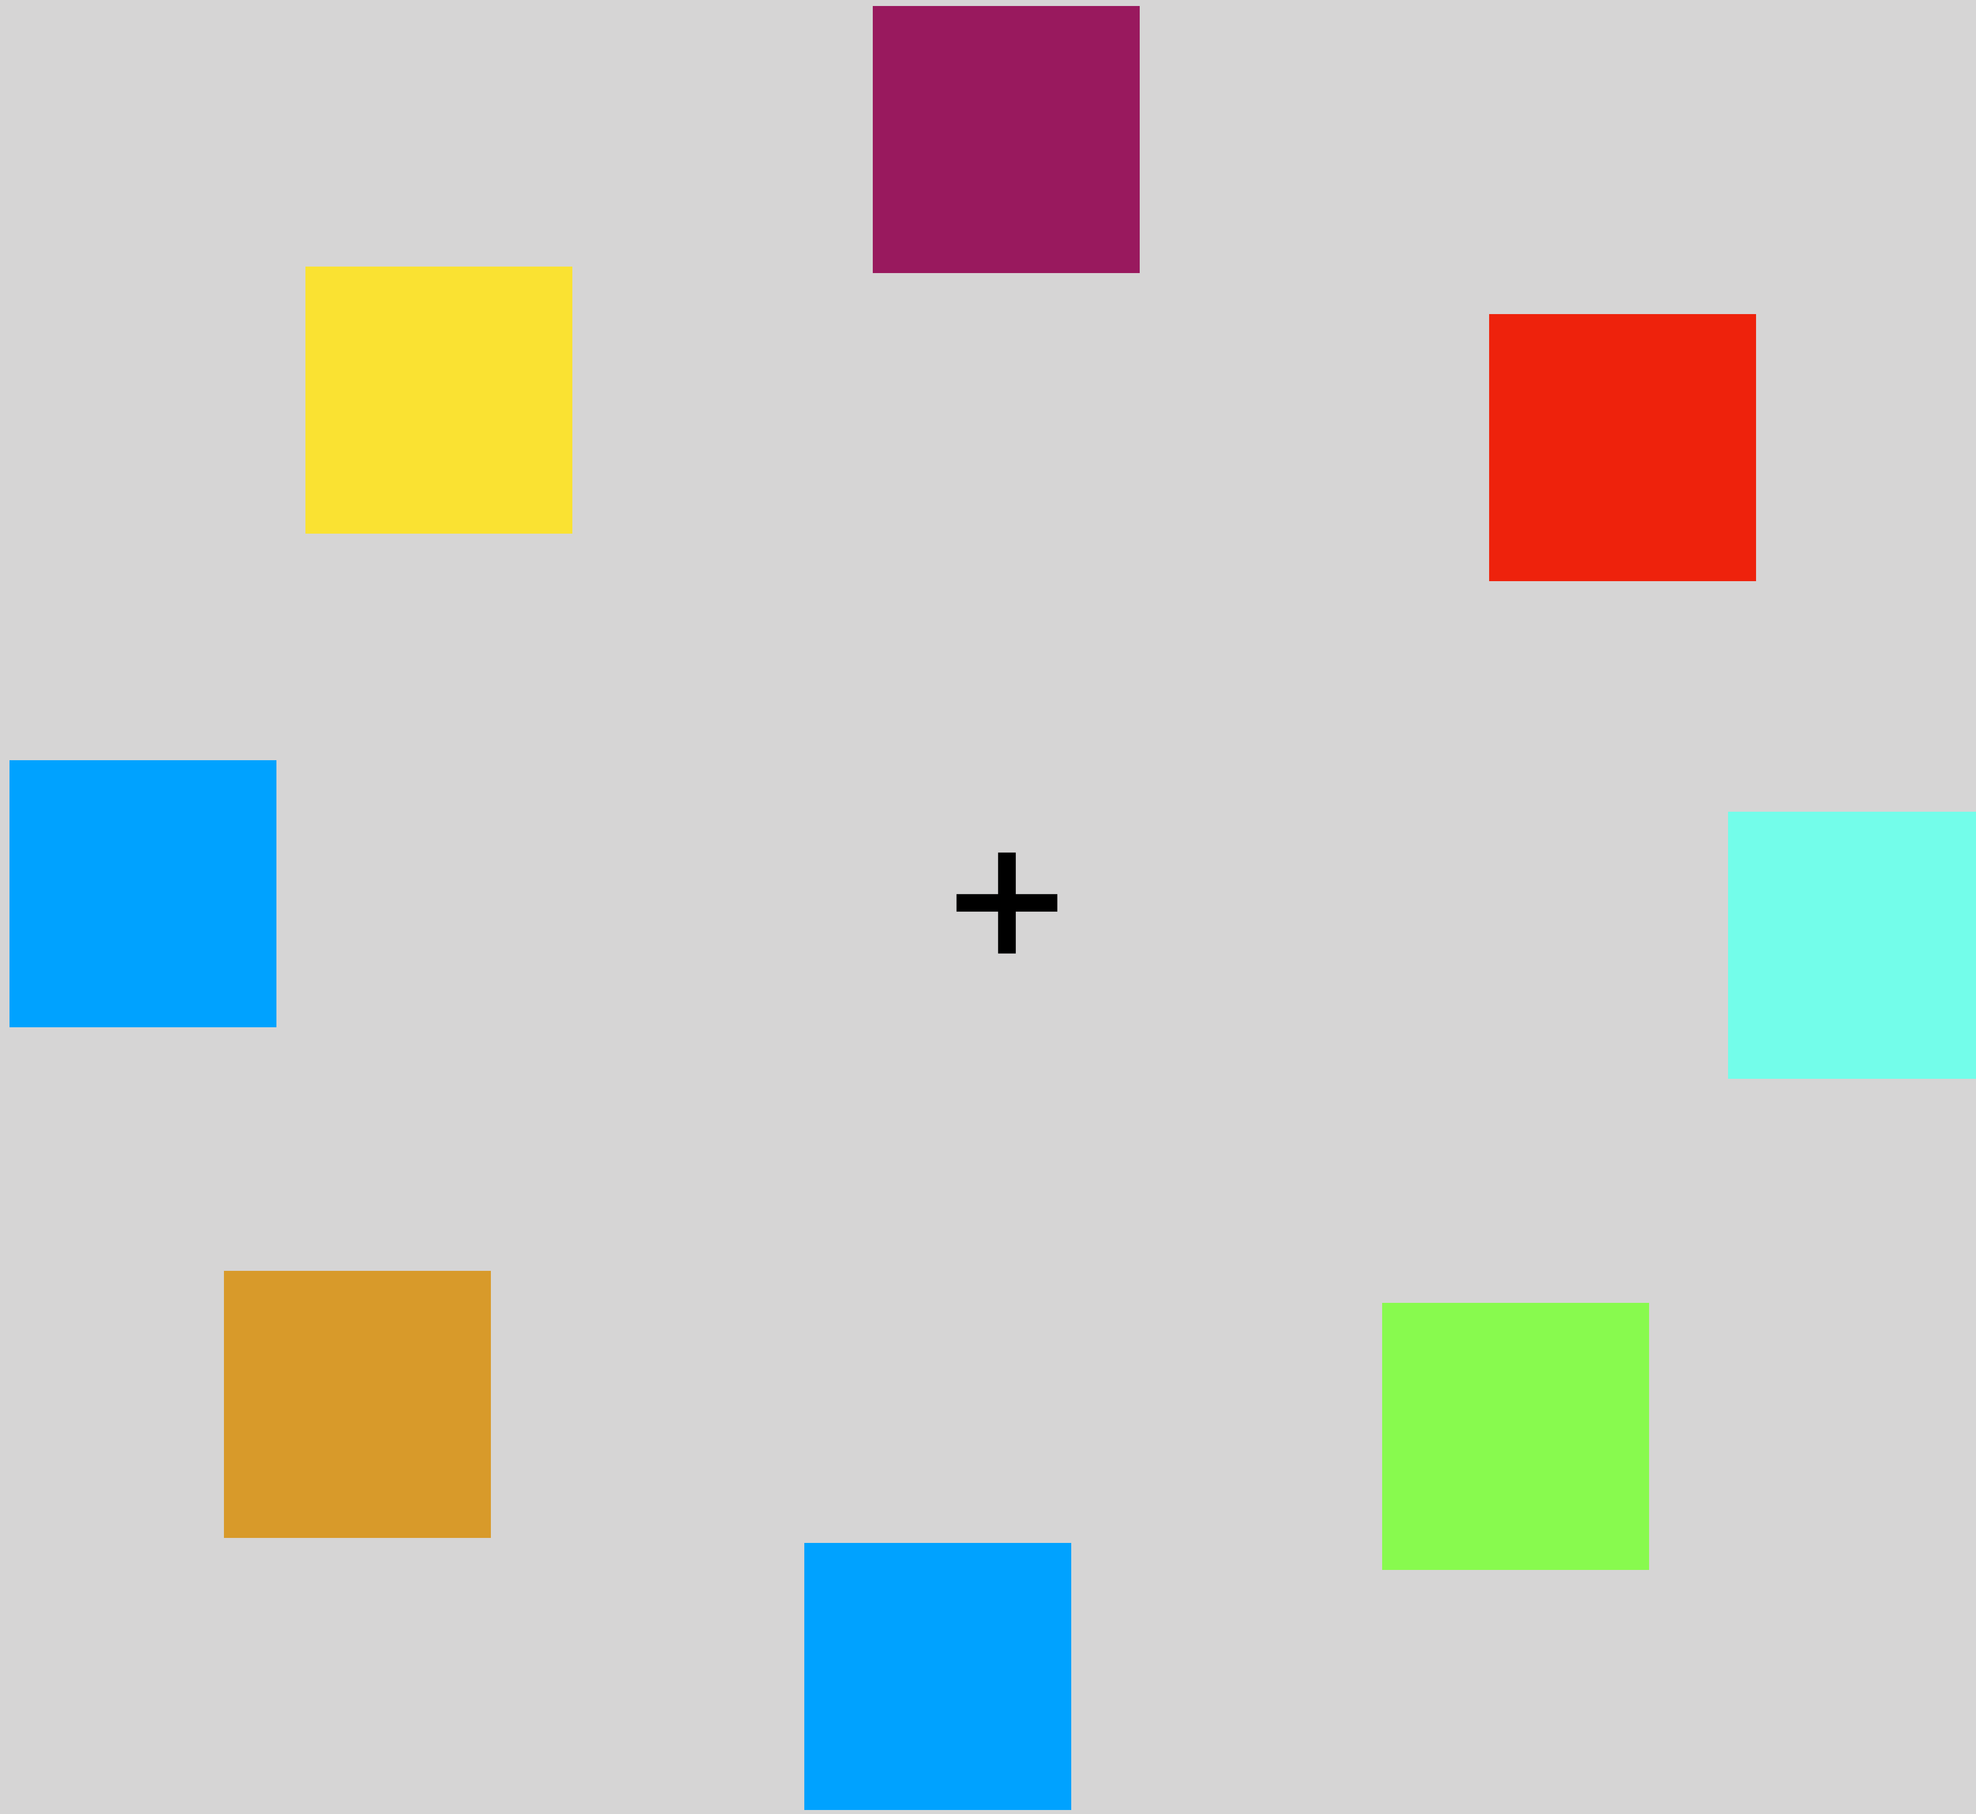

What color was it?

+

Results

As set size increases, approximately linear increase in error

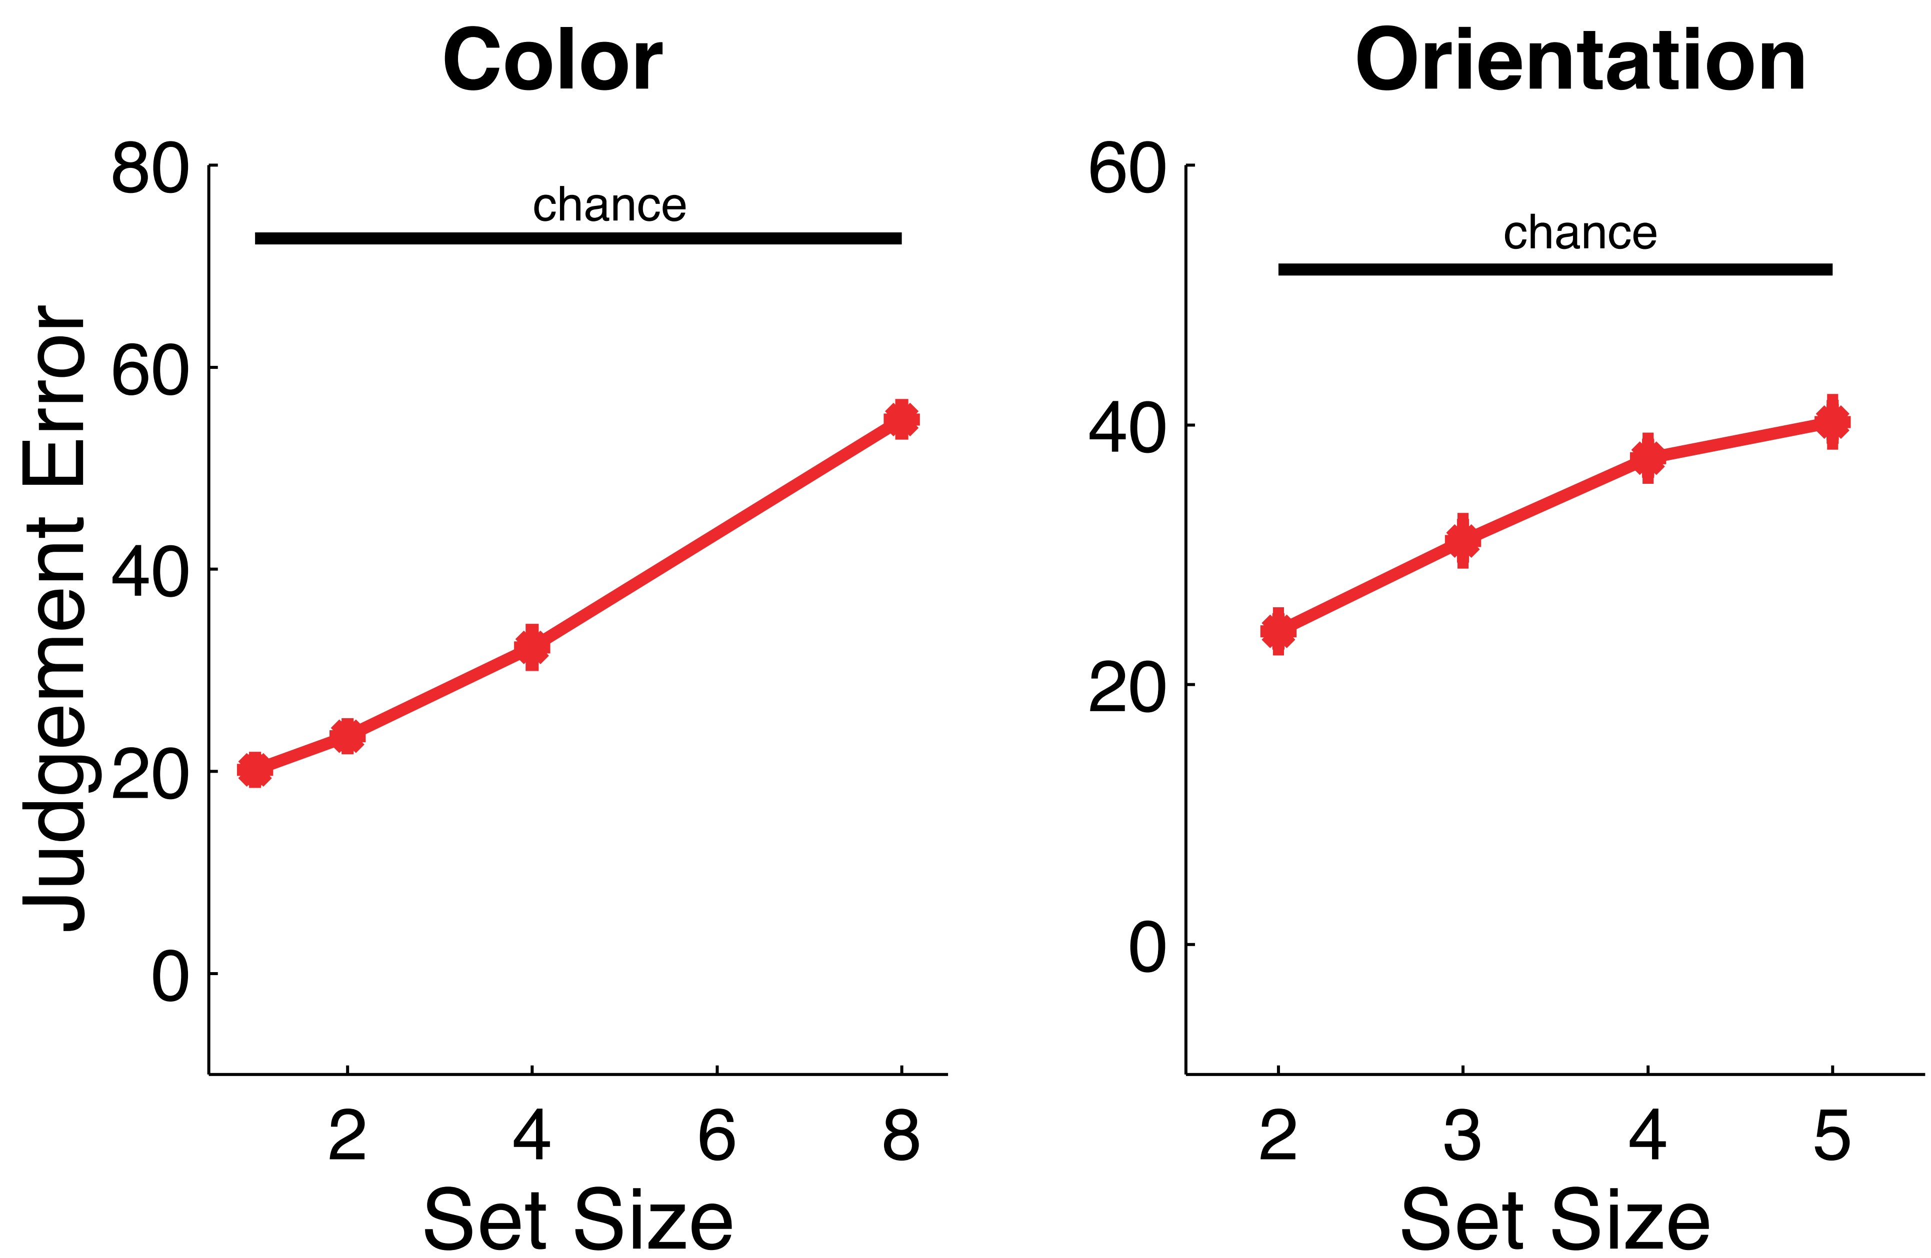

more continuous resource than slots..

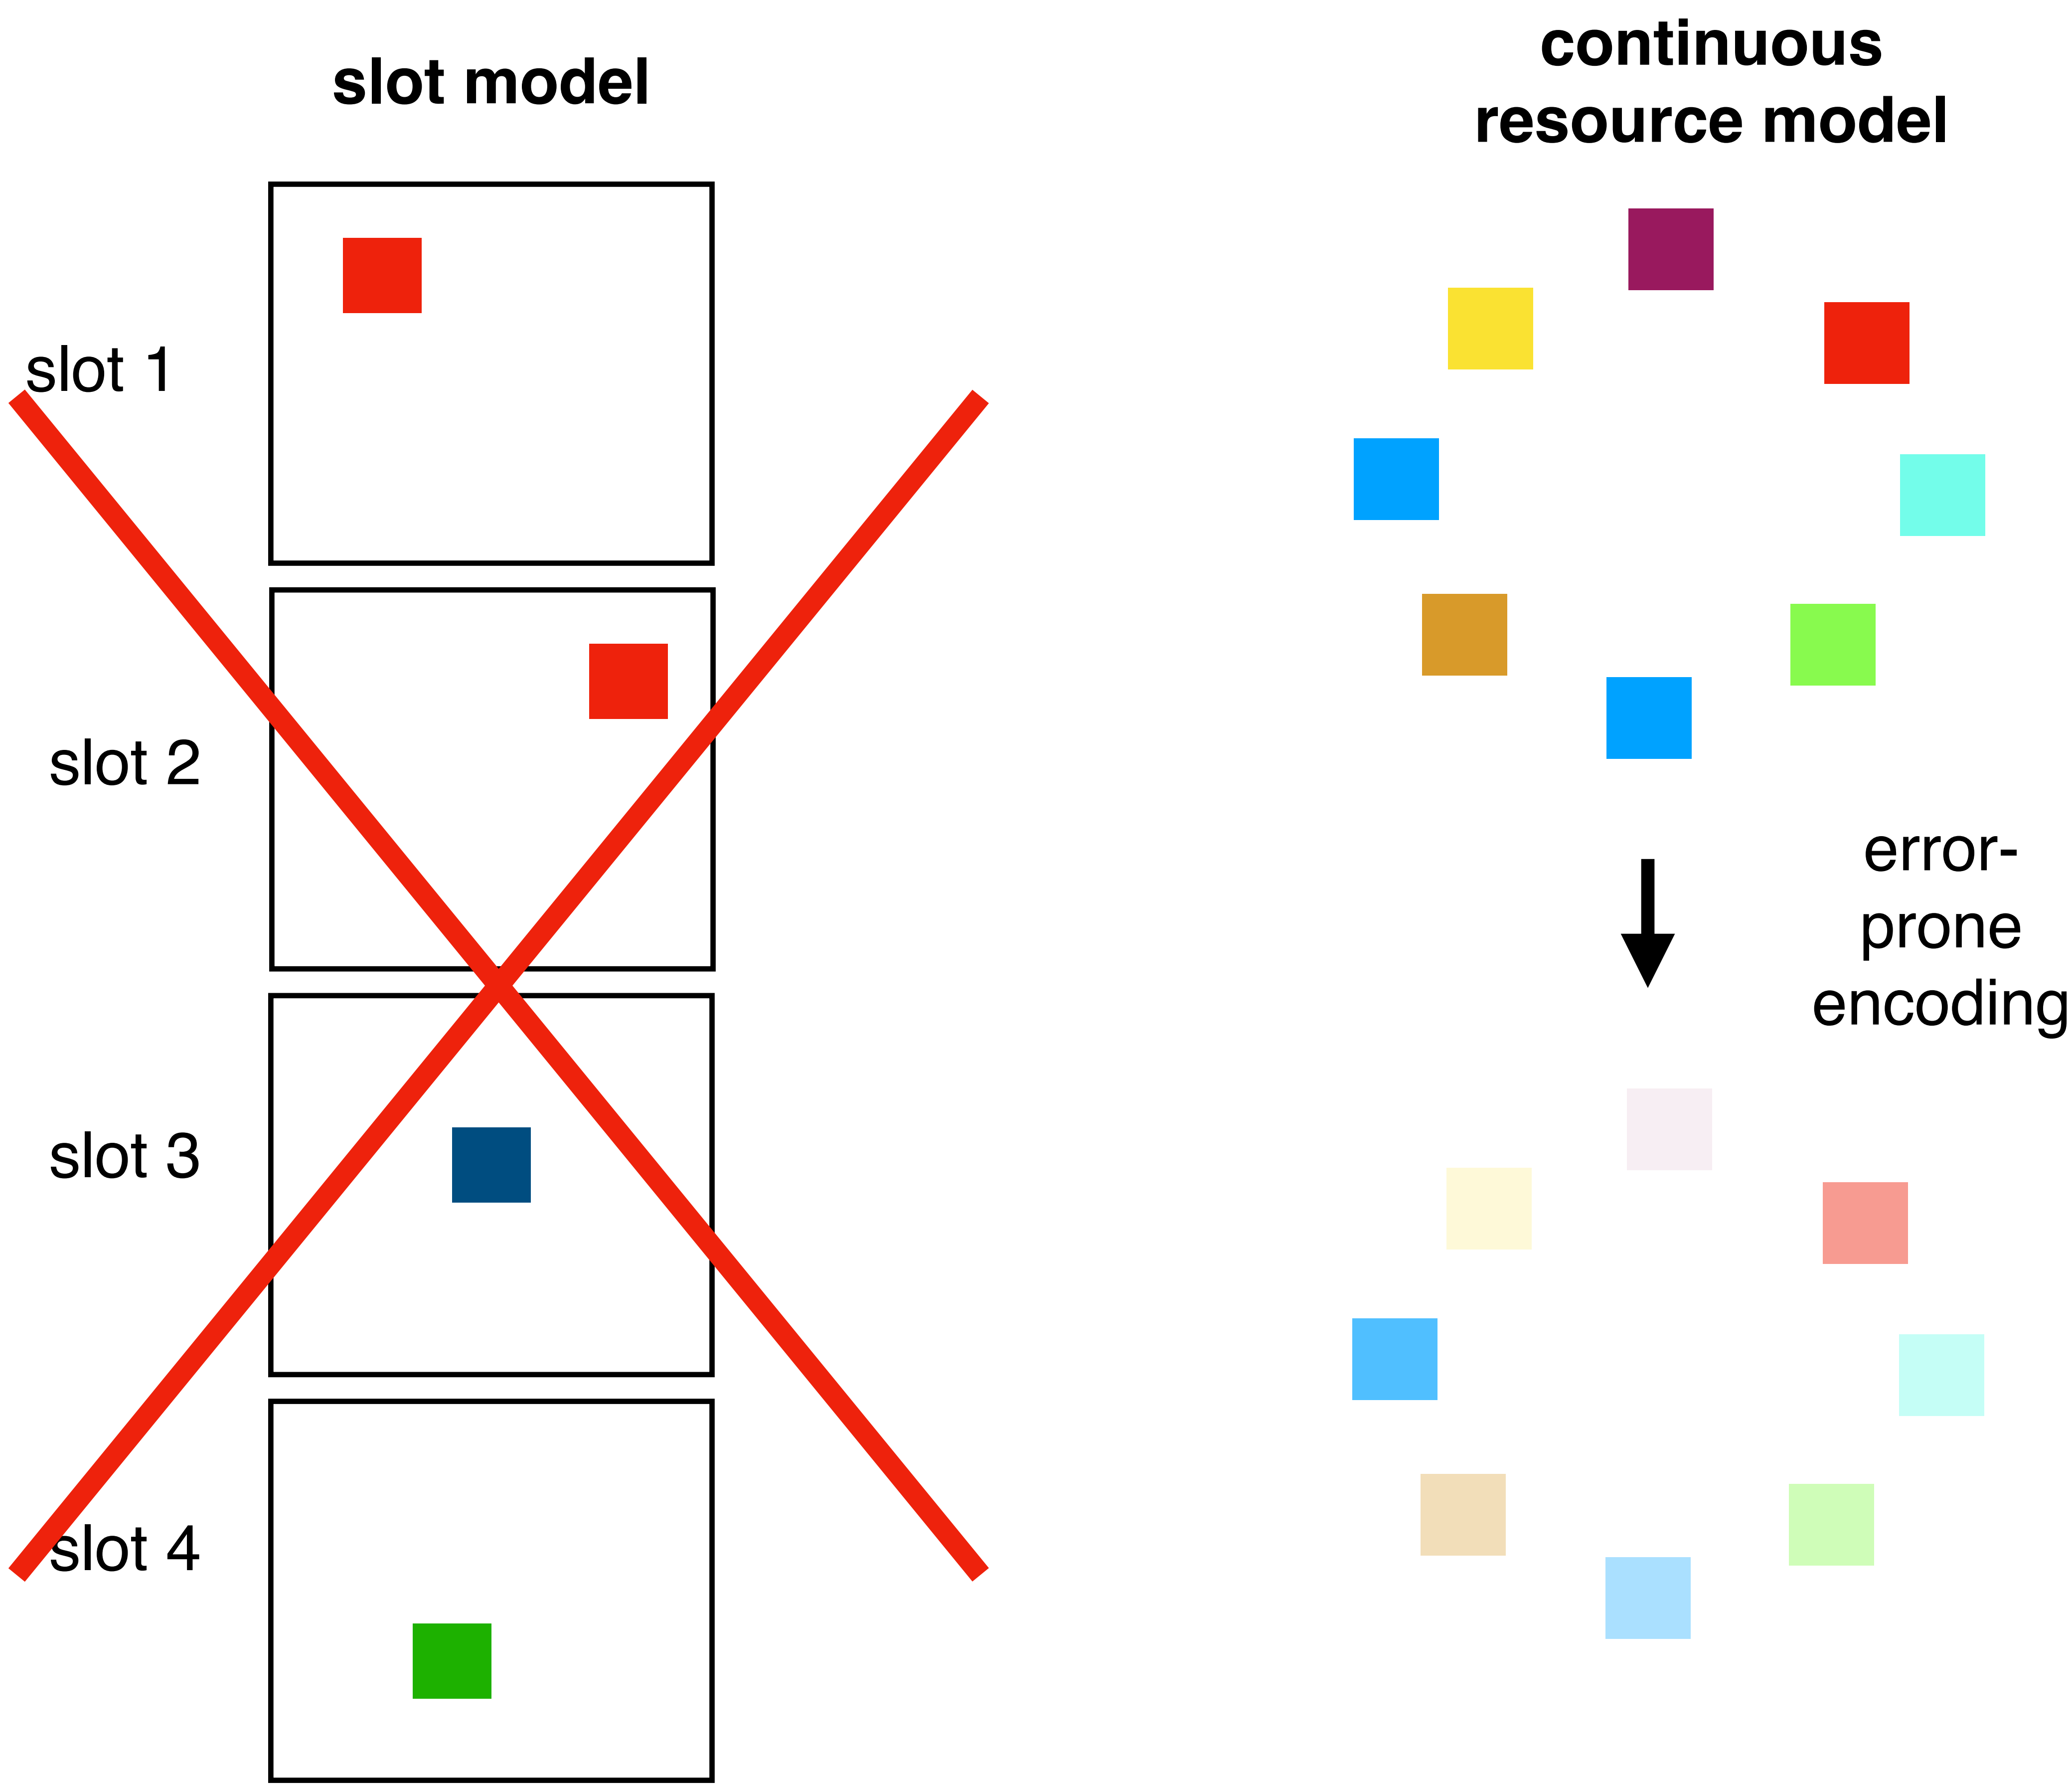

Well-designed studies that get at underlying representations/ computations

Now, time for the in-class activity

Acknowledgements

- Initial version of many of these slides are from Todd Gurecks