What are representations *for?*

Rosa Cao

Representation workshop Dublin 2022

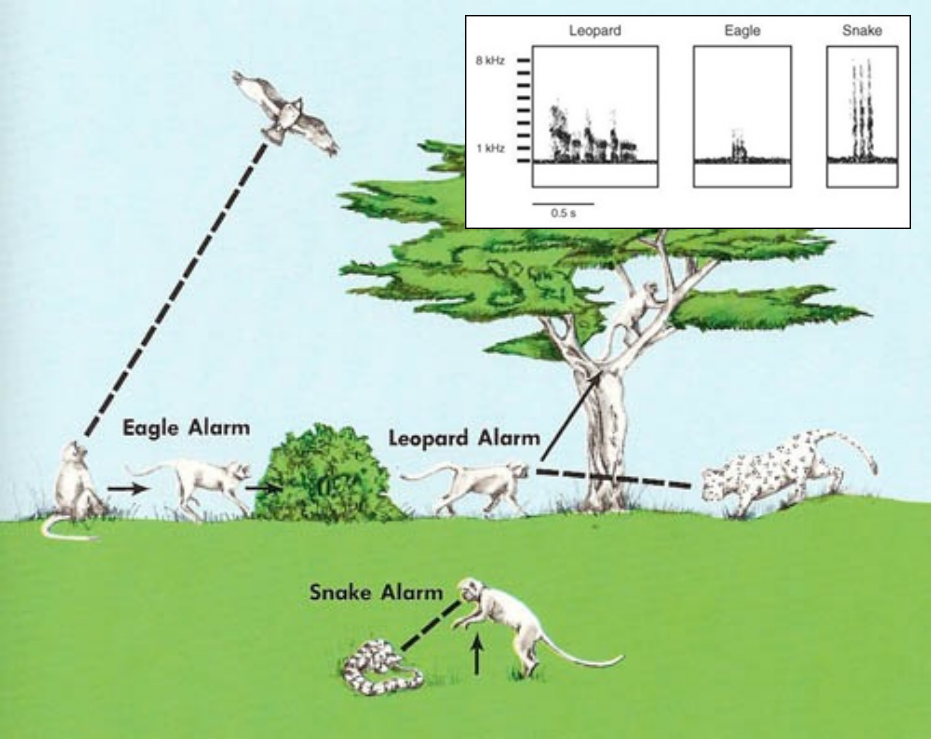

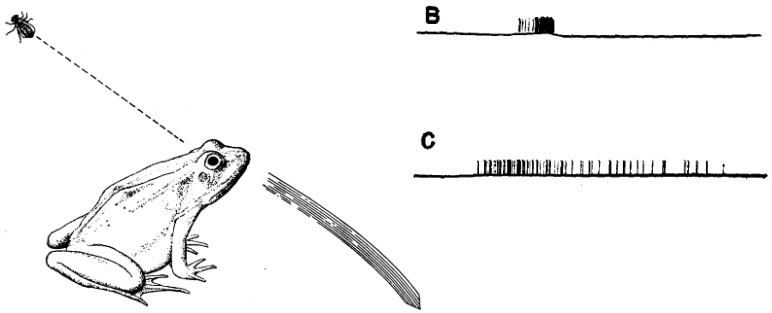

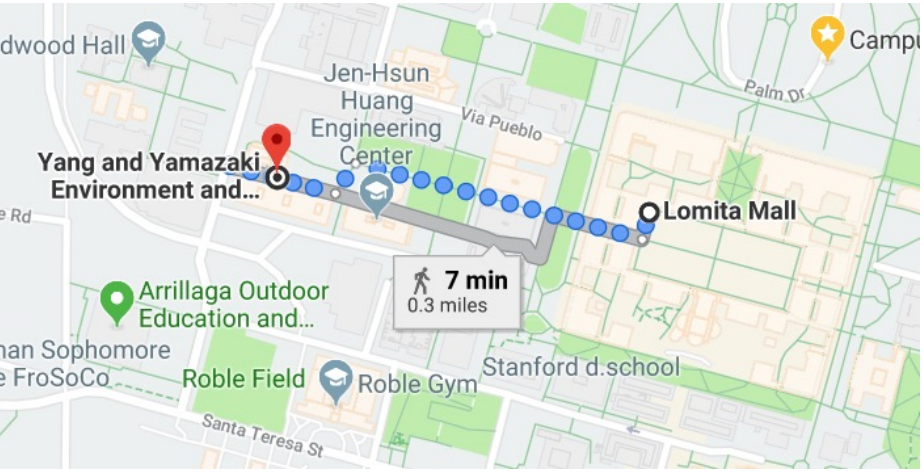

Figure 1.1
Aristotle mentally representing Graycat with a ball.

(At least) four notions of representation

Cummins, Meaning and Mental Representation, 1989

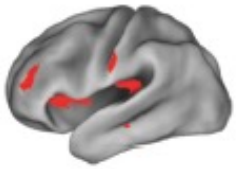

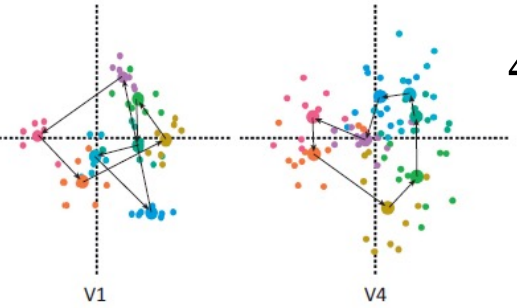

Kriegeskorte & Kievet 2013

1. Mental representation – person-level thoughts, experiences, introspectively accessible

   "Folk/intuitive" (everyday, everyone)

2. Patterns of (neural) activations identified by their co-variance with variables of interest

   "physical/neural" (neuroscientists)

3. Anything that plays a particular functional role (standing-for, carrying info for a user …)

   "theoretical posit" (philosophers, psychologists)

4. A *scheme* for mapping a class of contents (information) to a corresponding class of physical configurations

   "representation scheme" (computational)

An exercise for the reader …

"… constructing new **representations** of, or re-representing, sensory input is important because sensory receptors often deliver *information* in a form that is *unsuitable* for higher-level cognitive tasks. For example, roughly 108 photoreceptors provide a pixelated *description* of the images that appear on our retinas. A list of the membrane potentials of each of these photoreceptors provides a bulky and awkward **representation** of the visual world, from which it would be *difficult* to identify directly the face of a friend or a familiar object. Instead, *the information* provided by photoreceptor outputs is *processed* in a series of stages involving increasingly sophisticated **representations** of the visual world. In this chapter, we consider how these more complex and *useful* **representations** can be constructed.

… In **representational** learning, we seek to identify [visual] causes by analyzing the statistical structure of visual images and building a model, called the generative model, that is able to reproduce this structure."

Brain and behavior can be described in purely physical terms

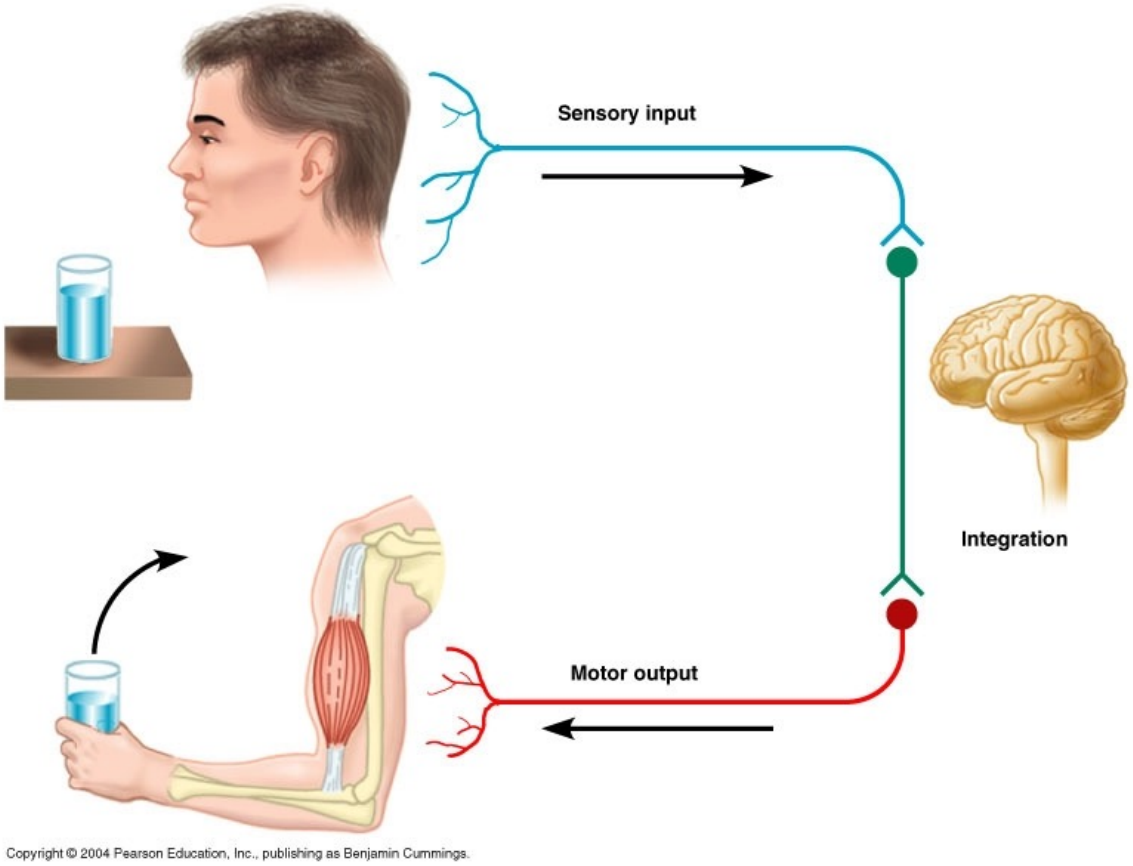

"You get spikes in … you get spikes out. This reduces our problem just from translating one set of spikes to another. Still difficult, but tractable."

- Florian Engert

So (why) do we need representations at all?

- As a bridging concept to connect the mechanistic with the psychological? The "wiring-and-connection" facts with the behavioral facts?

More generally, can we relate different levels of explanation, or different taxonomies?

e.g. Physical vehicle + Information content + Function
(cf. Decharms and Zador, Dretske, Shea, and many others differentiated primarily by their different notions of function)

- As a pragmatic shorthand compression for complex conjunctions and disjunctions of physical states?
- Or is the term now too overloaded to be helpful?

Moving forward

- What, in any given context, do we want to use the notion of representation *for* ?

Representations play a crucial role in *explanation*

As **explanandum**, or *target* of explanation:

- How is it that our brains are sophisticated enough to support our abilities to represent things?

As **explanans**, or what we appeal to in *giving* an explanation

- Our abilities depend on and are explained by the processing of certain internal content-bearing states

The interpretative exercise revisited

?

"… constructing new **representations** of, or re-representing, sensory input is important because sensory receptors often deliver *information* in a form that is *unsuitable* for higher-level cognitive tasks. For example, roughly 108 photoreceptors provide a pixelated *description* of the images that appear on our retinas. A list of the membrane potentials of each of these photoreceptors provides a bulky and awkward **representation** of the visual world, from which it would be *difficult* to identify directly the face of a friend or a familiar object. Instead, *the information* provided by photoreceptor outputs is *processed* in a series of stages involving increasingly sophisticated **representations** of the visual world. In this chapter, we consider how these more complex and *useful* **representations** can be constructed.

R2

R1 or R3?

R3

R4

… In **representational** learning, we seek to identify [visual] causes by analyzing the statistical structure of visual images and building a model, called the generative model, that is able to reproduce this structure."

Thank you

Four *different* roles in explanation?

R1. Mental representation – person-level thoughts, experiences, introspectively accessible

R2. Patterns of neural activity (often subpersonal)

R3. Informational signals subserving a particular *function*, or playing a particular computational role

R4. A *scheme* for mapping a class of contents (information, states of interest) to a corresponding class of physical configurations

R1. Explanandum in neuroscience

(perhaps explanans for behavior)

R2. Term for complex *physical* events

R3. Potential *bridging* concept

R4. An abstract mapping that exhibits useful *systematicity*

An example of R4

A neural code for three-dimensional object shape in macaque inferotemporal cortex

Yukako Yamane[1], Eric T Carlson[1,2], Katherine C Bowman[1,3], Zhihong Wang[1] & Charles E Connor[1,3]

"A primary goal in the study of object vision is to decipher **the neural code** for complex object shape. At the retinal level, object shape is **represented** isomorphically (that is, replicated point for**, stable neural code** that **efficiently captures the shape information** needed for identification and other apoint) across a two-dimensional map comprising approximately 10^6 pixels. This isomorphic representation is far too unwieldy and unstable (as a result of continual changes in object position and orientation) to be useful for object perception. The ventral pathway of visual cortex must transform the isomorphic image into a **compact**spects of object vision.

…. At the population level, these **signals** combine to **represent** complete boundary shapes as spatial configurations of constituent fragments.

Nature Neuroscience, 2008

Not all useful biological processes involve representations

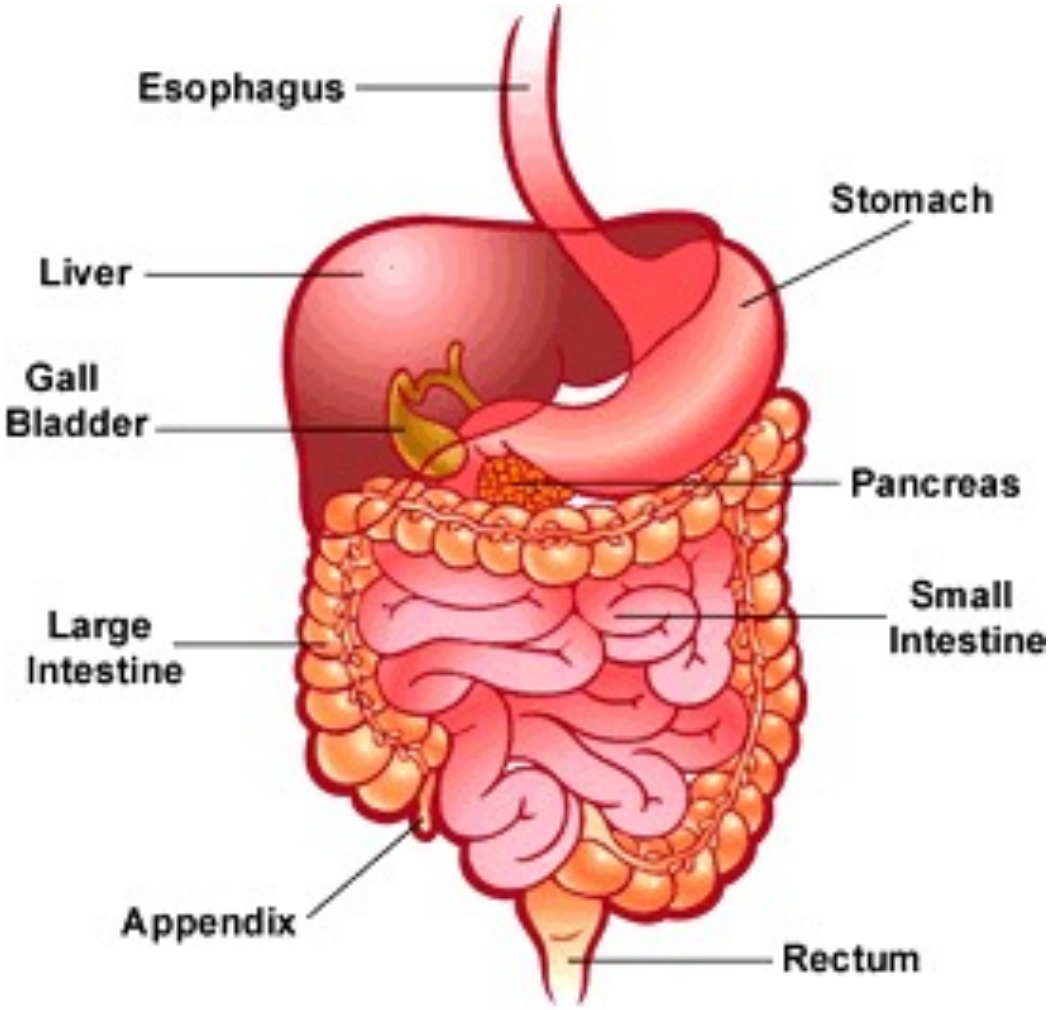

Sterelny's decoupled intermediaries

"…internal states that track aspects of our world, but which do not have the function of controlling particular behaviors"

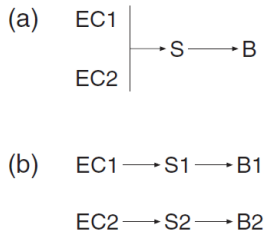

Figure 1. Godfrey-Smith's two stages of behavioral complexification. (**a**) Stage 1: indifference to environmental variation. (**b**) Stage 2: flexible response to environmental variation. *B* behavior, *EC* environmental condition, *S* sensory

Detection system (a) EC → S → B

Robust tracking (b)

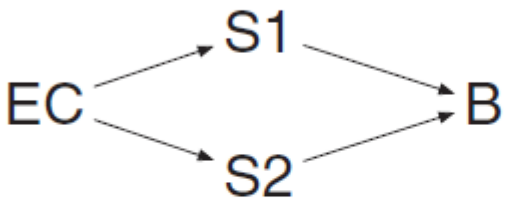

Decoupled representation (c)

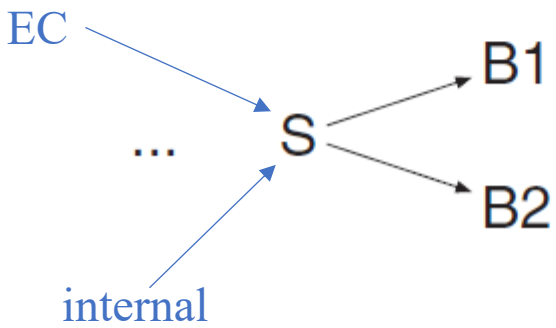

Transparency

Translucency

Christensen 2010