Representation

The form of knowledge

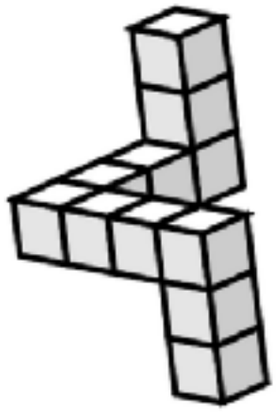

If we insist that mind and world are separate,

then we must provide a story of how it is that mind comes to know world

Most such stories involve the creation of *representations* of the world

The notion of *representation* thus needs some unpacking. It is not a simple unitary concept.

Salt Lake City

COMP 20090: Introduction to Cognitive Science

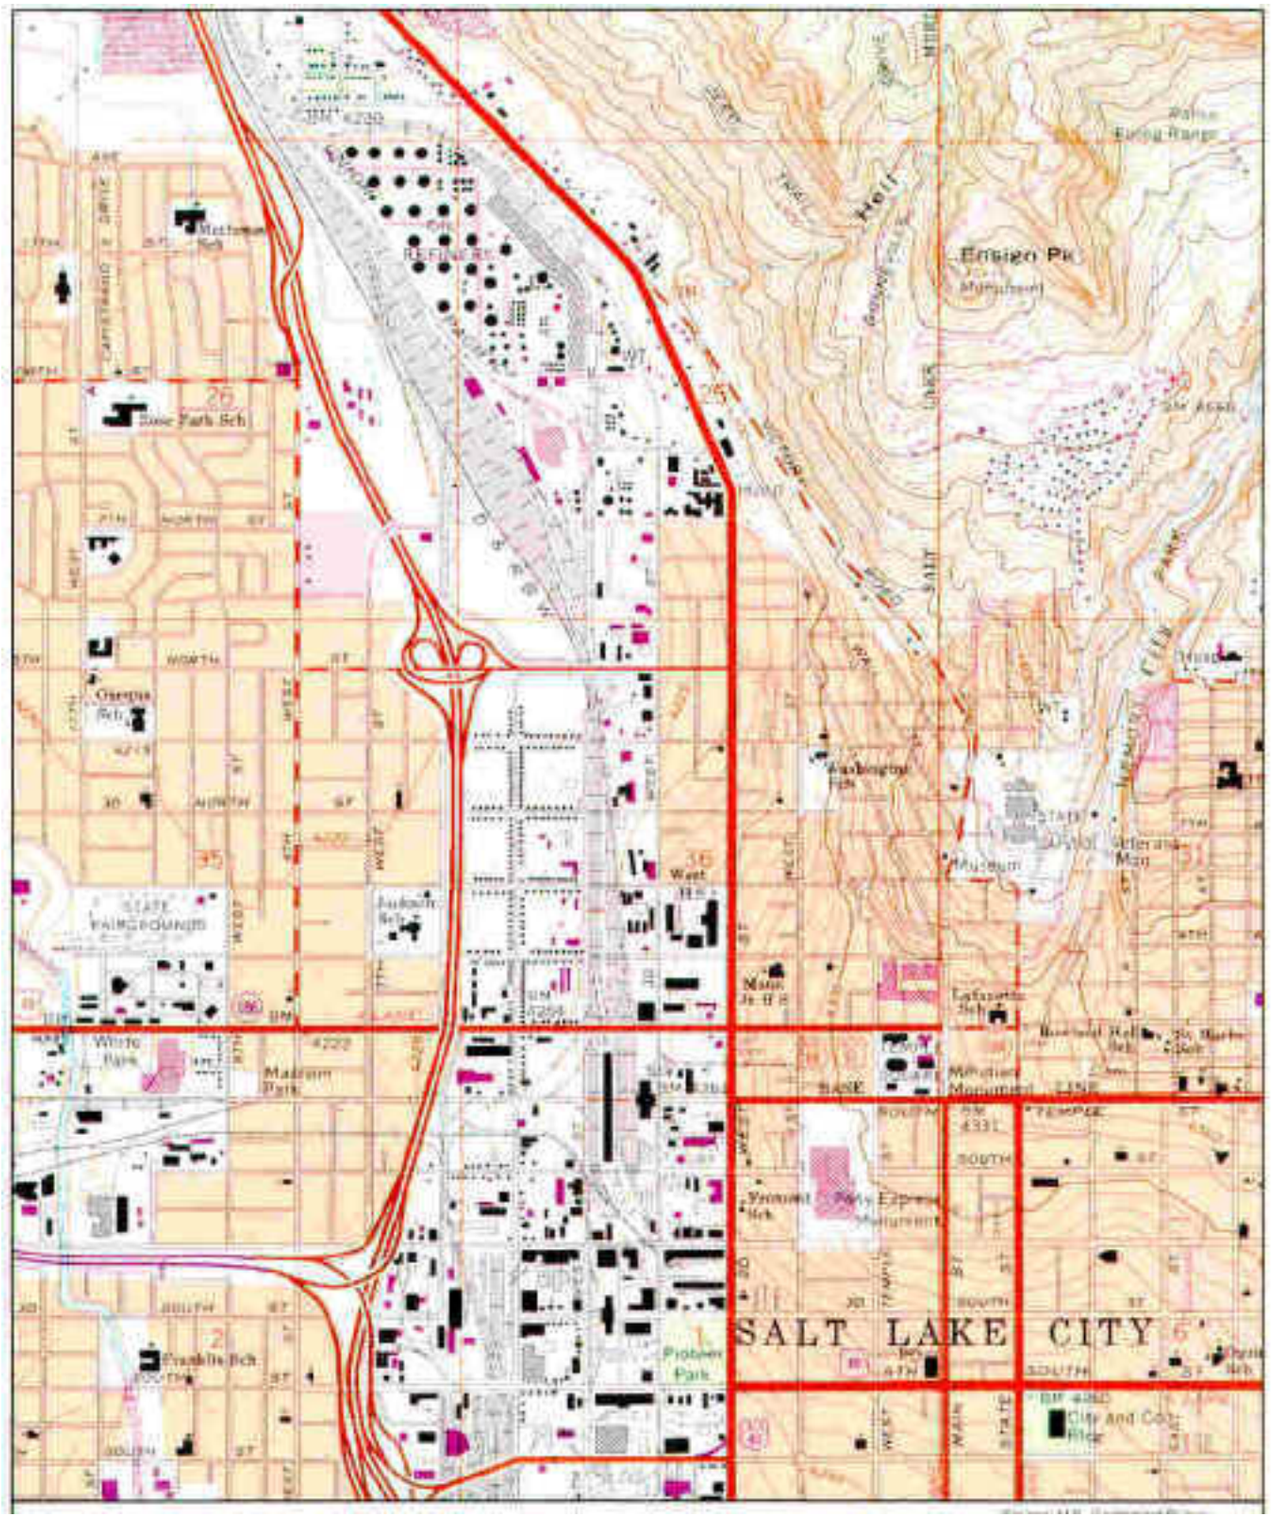

COMP 20090: Introduction to Cognitive Science

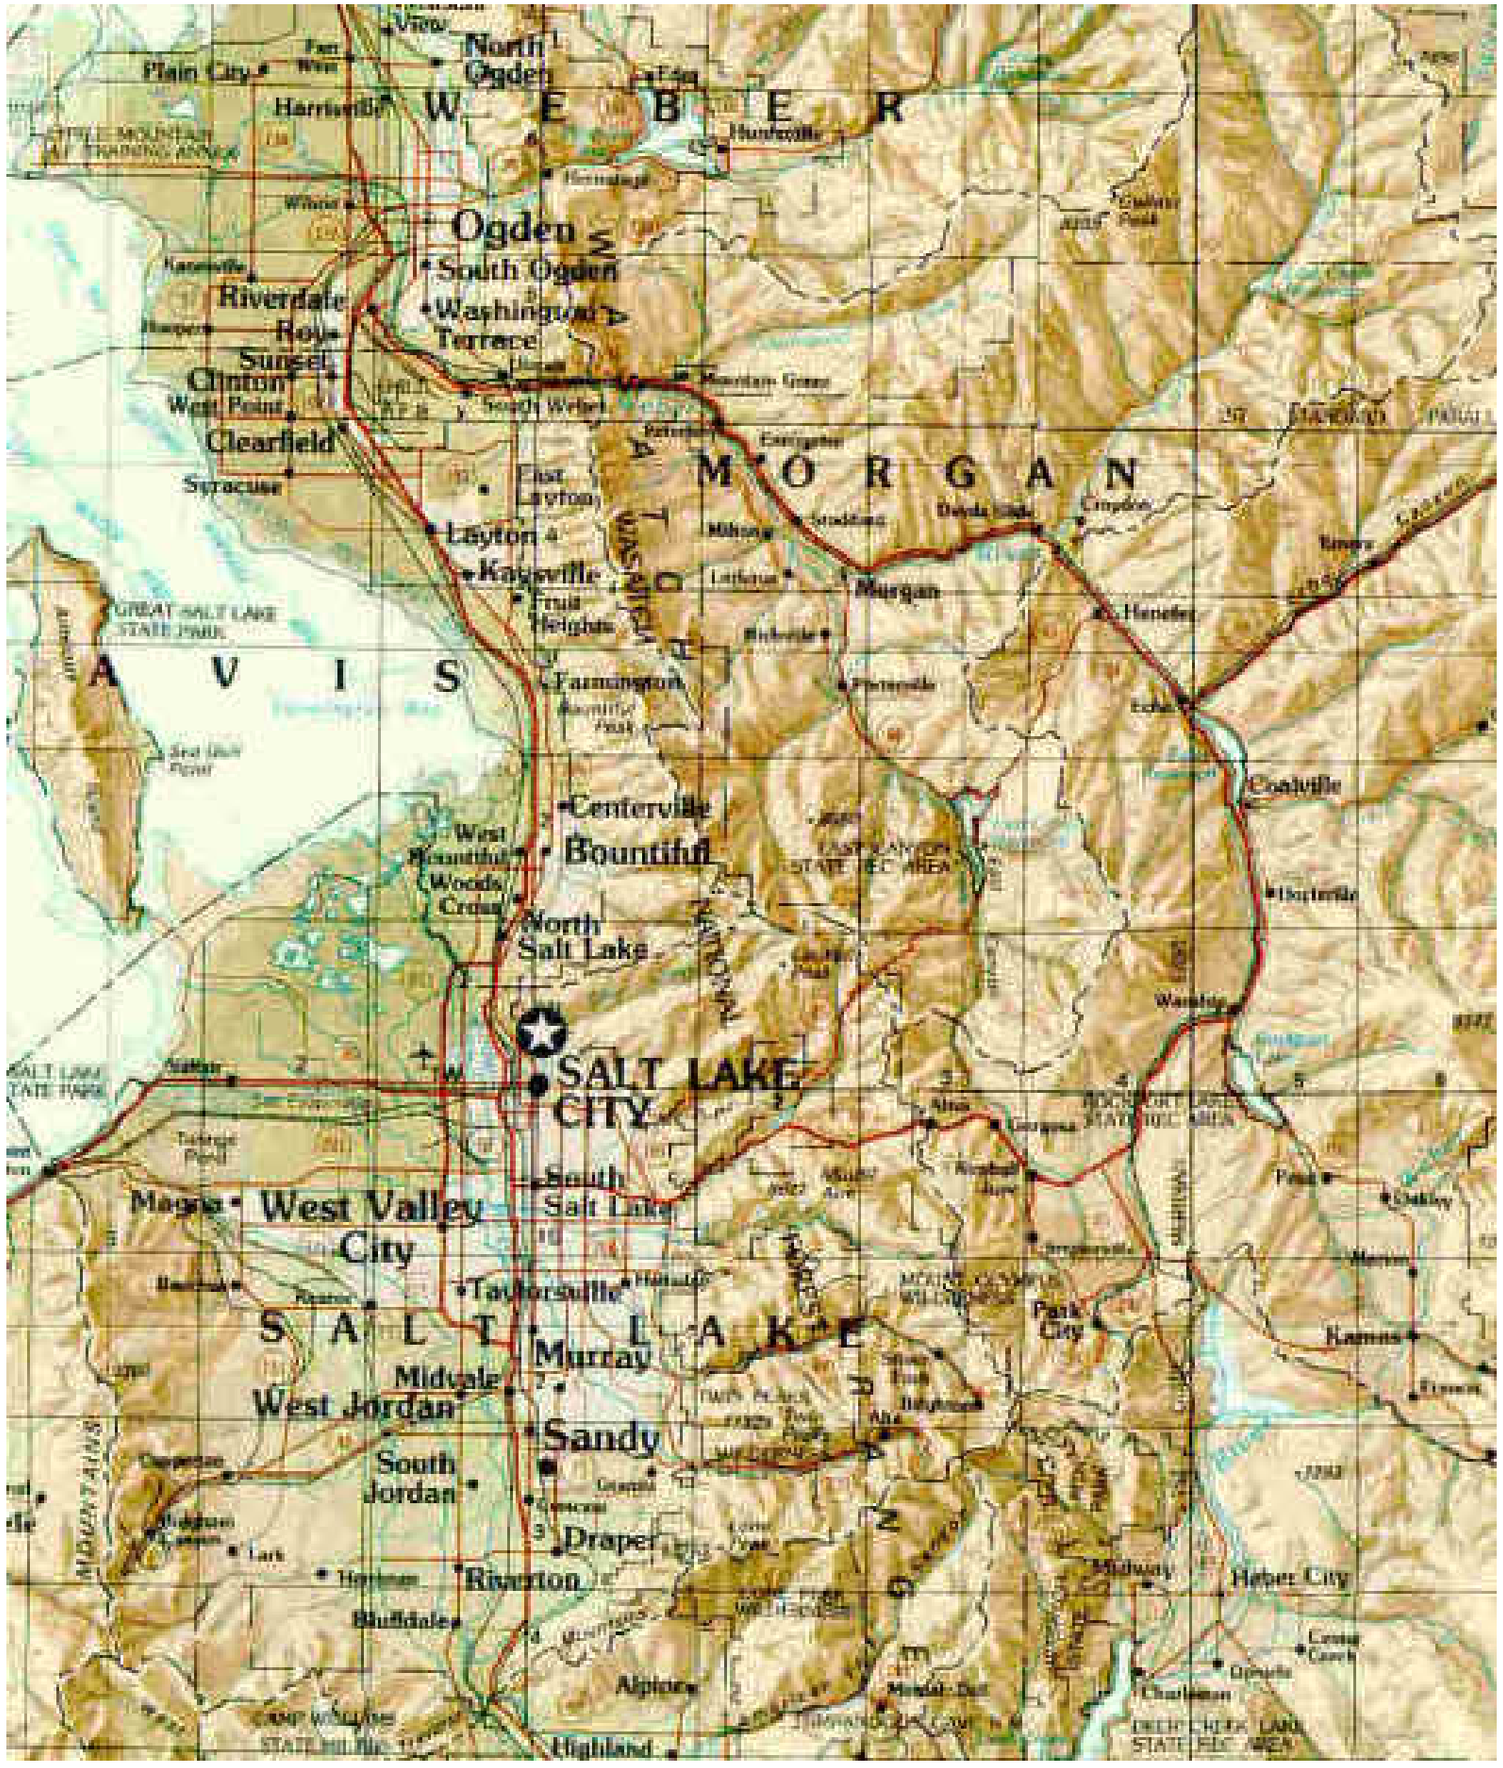
North Ogden
Plain City
Harrisville
WEBER
Ogden
South Ogden
Riverdale
Washington Terrace
Roy
Sunset
Clinton
West Point
Clearfield
Syracuse
MORGAN
East Layton
Layton
Kaysville
Fruit Heights
Morgan
Henefer
Coalville
DAVIS
Farmington
Centerville
Bountiful
West Bountiful
Woods Cross
North Salt Lake
SALT LAKE CITY
South Salt Lake
West Valley City
Taylorsville
SALT LAKE
Murray
Midvale
West Jordan
Sandy
South Jordan
Draper
Riverton
Bluffdale
Alpine
Highland
Park City
Kamas
Midway
Heber City
GREAT SALT LAKE STATE PARK

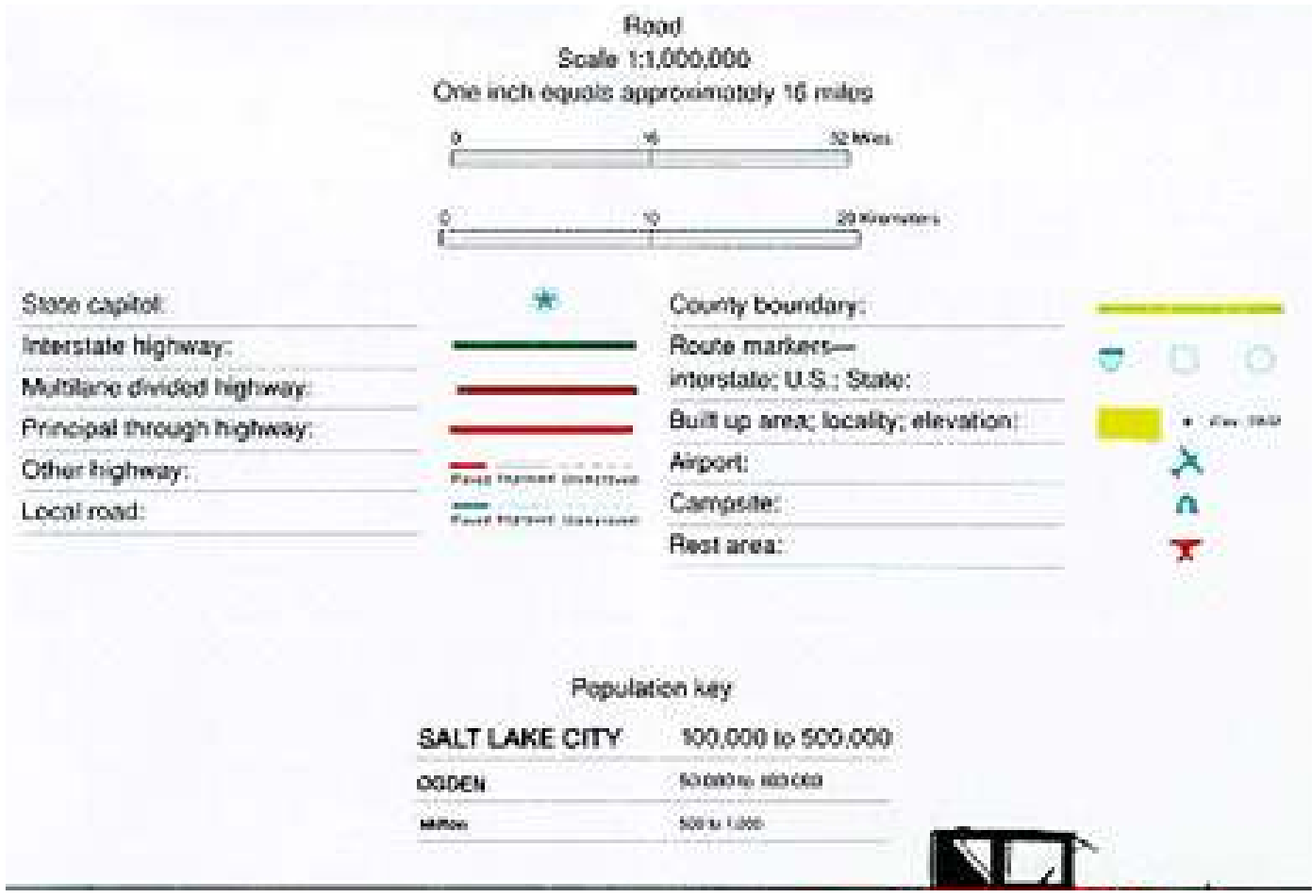

Symbols stand for things

What is explicit, and what can you figure out but is not explicit?

A symbol stands in *correspondence* to something else

Do you think a symbol has to look like the thing is corresponds with?

Which symbols you are familiar with do *not* look like the thing they refer to?


запрещено
запрещать

These are imprints. They also stand in correspondence to something.

What aspects of the lips/hand are captured by the imprints?

What aspects are not captured?

From the source event to the imprint, we can identify a *translation*, in which some properties are preserved, some are lost, and others are modified.

A SYSTEMATIC VIEW OF REPRESENTATION

On some views, we may say that something *represents* something else if:

[1] It stands in correspondence to that thing, and

[2] It plays a role within a system *because* it has that relationship of correspondence.

Note this is different from the translation/imprint view!

The relationship between a representation and its referent can be arbitrary

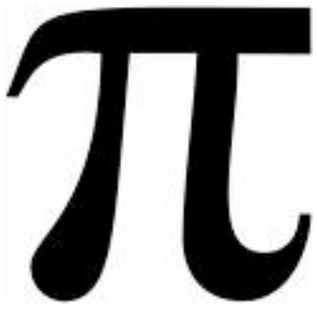

... or non-arbitrary

Which systems do these symbols function in?

Which elements on a map stand in arbitrary relationship to their referents, and which are non-arbitrary?

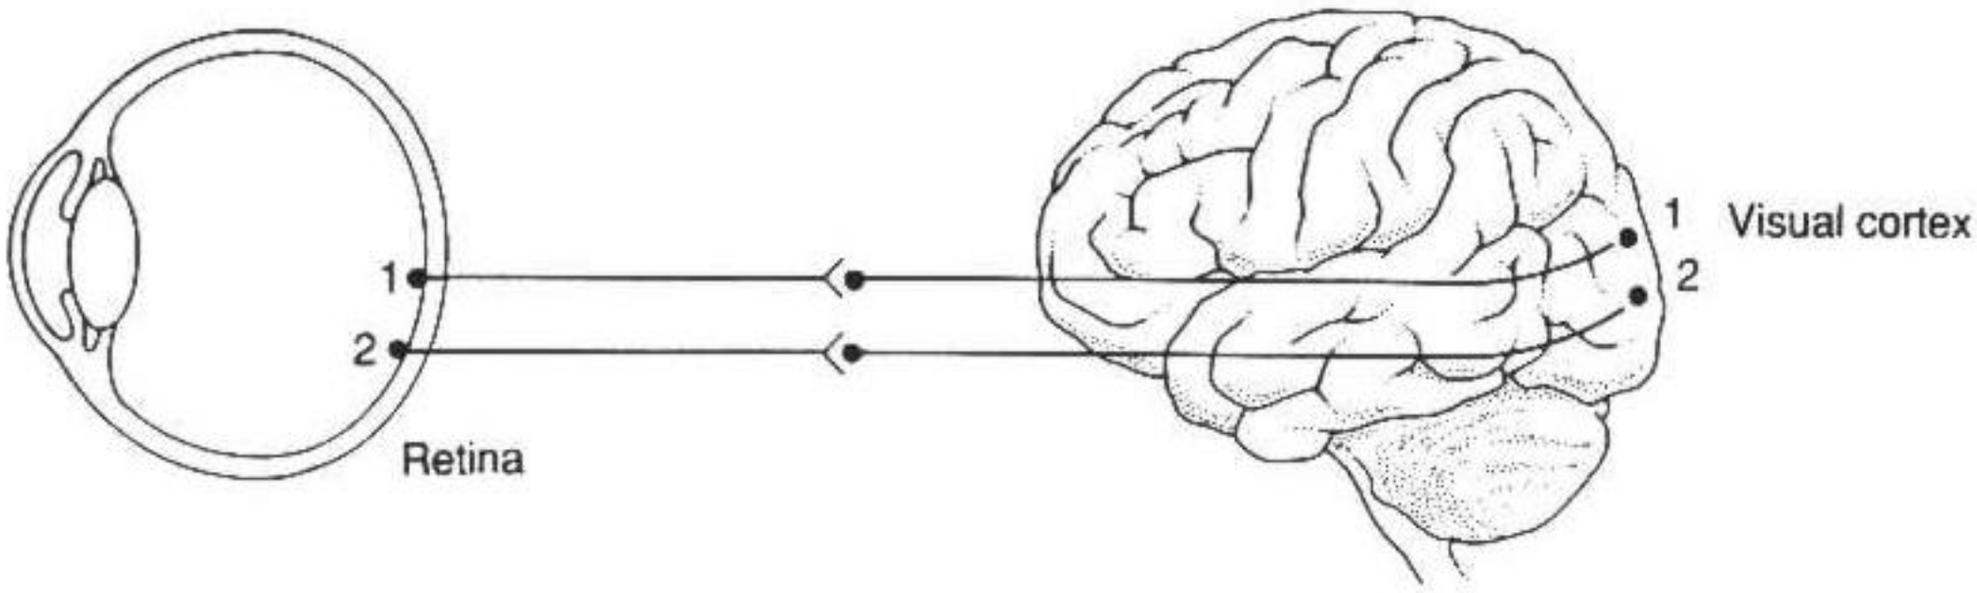

The spatial layout of the retina is retained through several transformations.

Places that are near each other stay near each other in the initial stages of processing.

Q: Is this more like a symbol or an imprint?

What is the difference between representation and re-presentation?

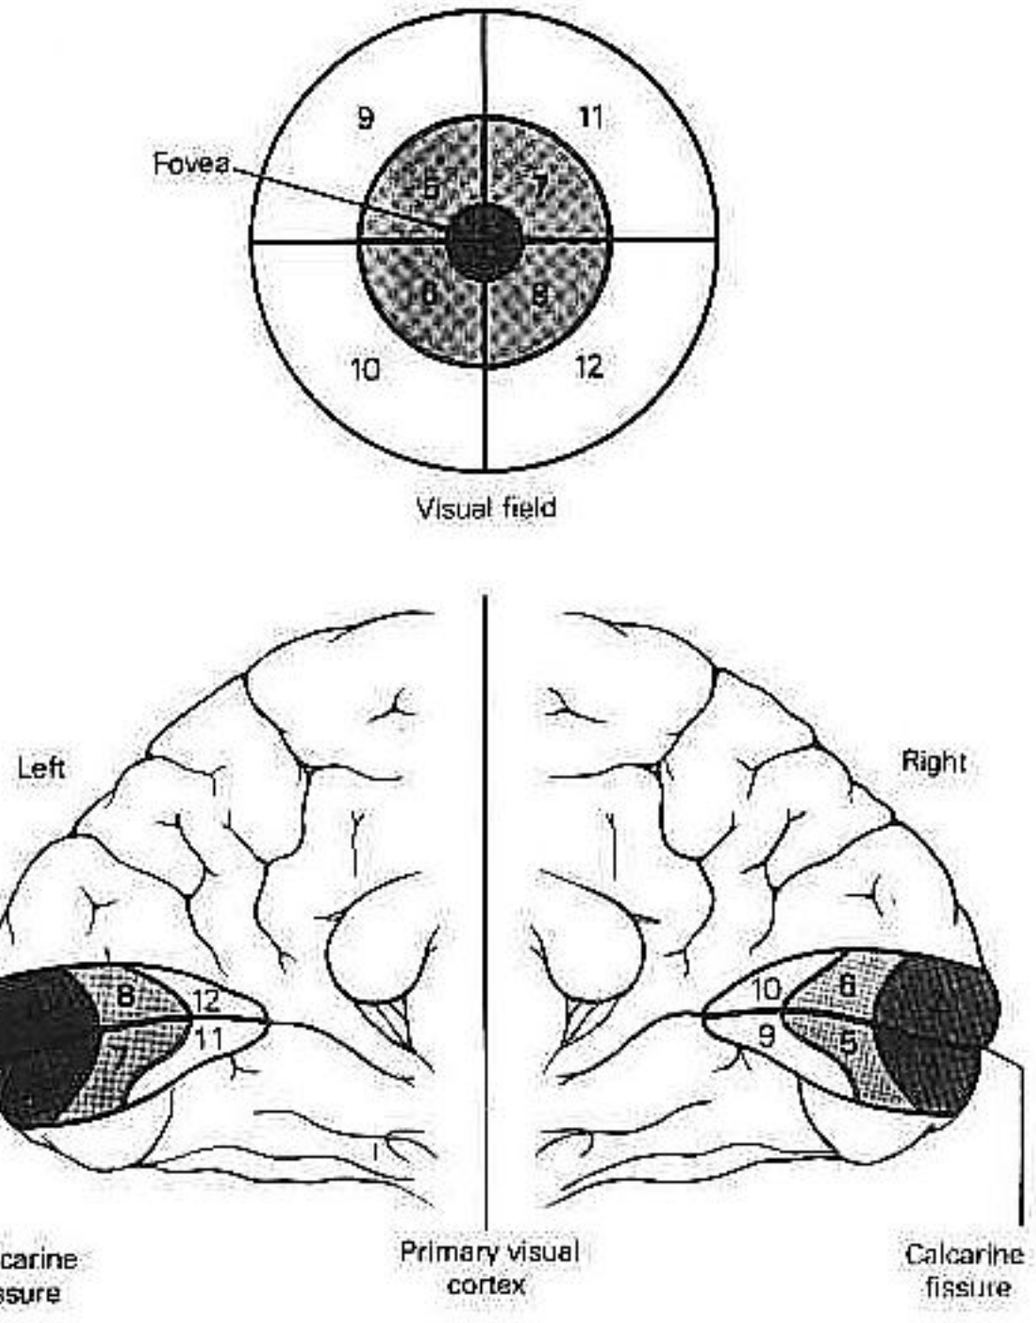

A retinotopic map in primary visual cortex of a monkey

What is preserved between retina and cortex?

What is altered or distorted?

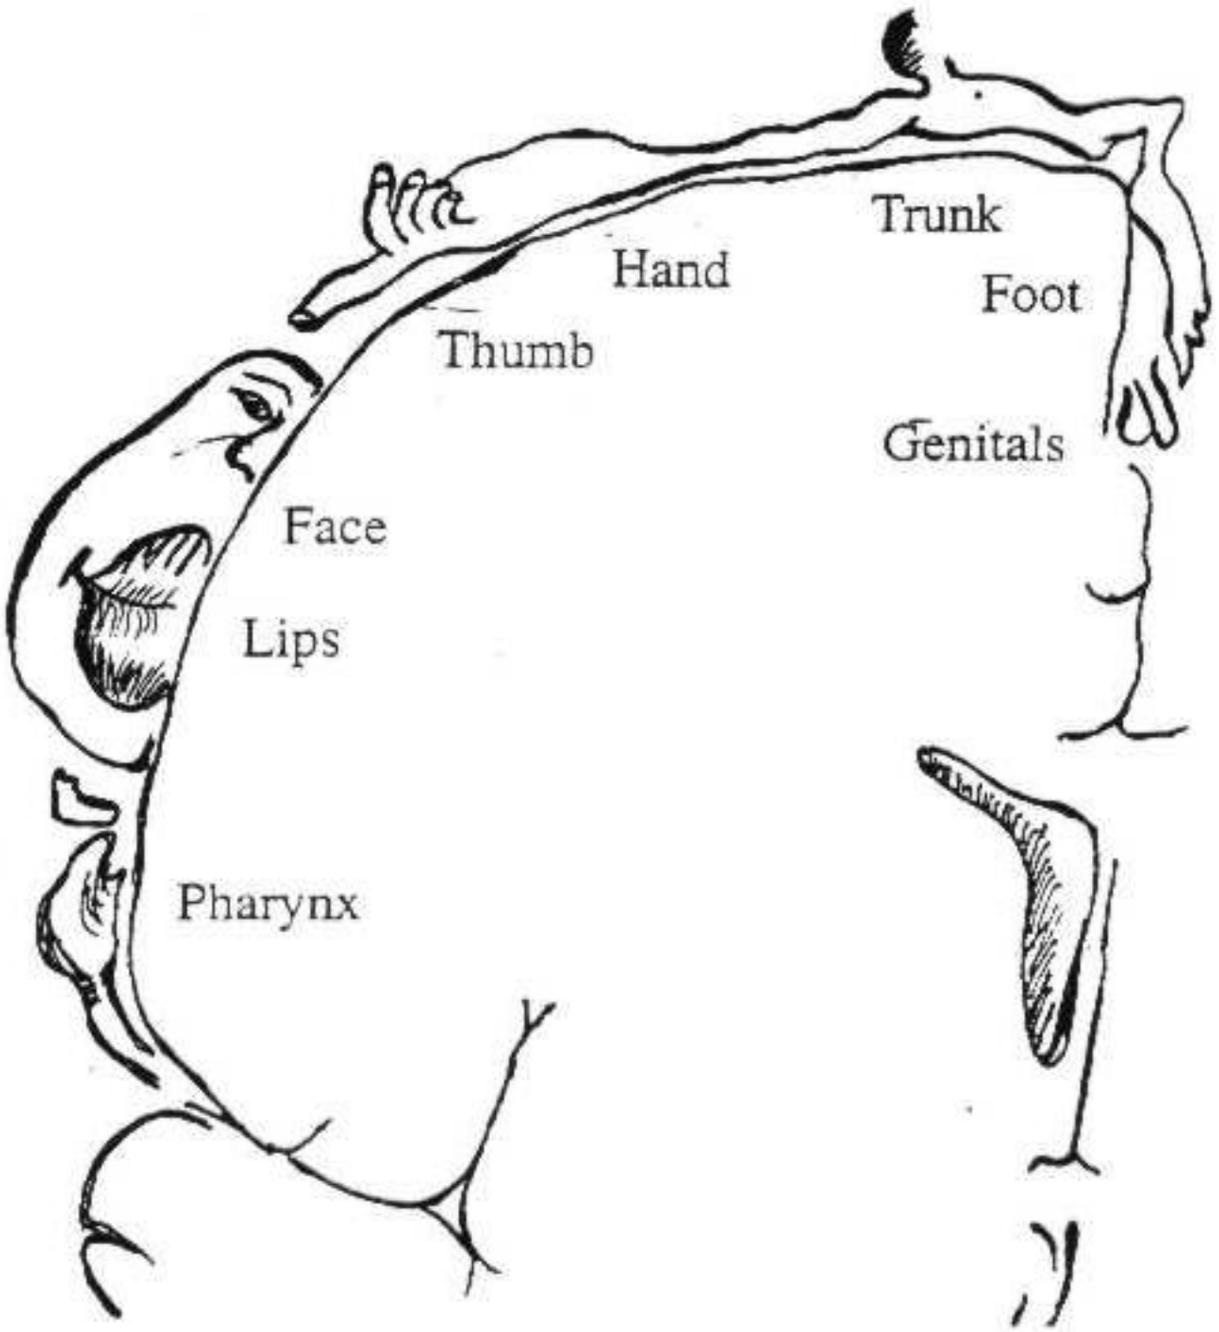

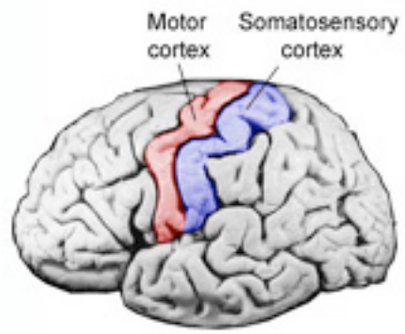

In primary somato-sensory cortex, different amounts of cortex are devoted to the various parts of the body.

This is one way the word *representation* is used.

Some would say that this shows a cortical representation of sensation in the skin

COMP 20090: Introduction to Cognitive Science

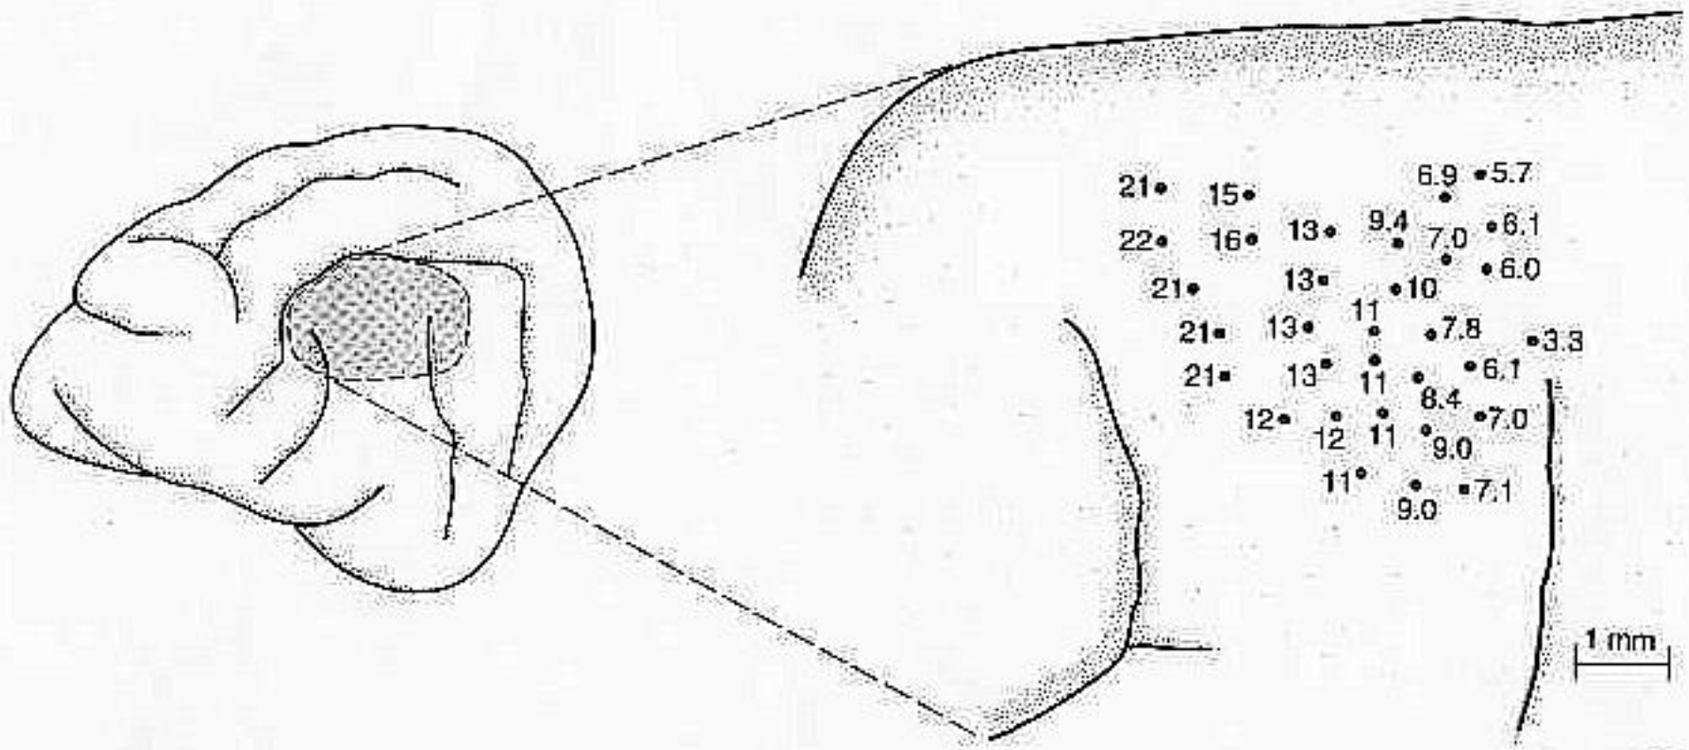

Figure 8.32

Left: The cat's brain, showing the location of the auditory receiving area (shaded). Right: A close-up of the auditory receiving area, showing the tonotopic map on the cortical surface. Each large dot represents a single neuron, and the number beside each dot represents the characteristic frequency of that neuron in thousands of Hz. (From Abeles & Goldstein, 1970, and Gidick, Gescheider, & Frisina, 1989.)

A tonotopic map in cat auditory cortex. The numbers refer to the frequencies each area is sensitive to. (All naturally occurring sounds are composed of a mixture of frequencies from low to high)

These examples are all of primary sensory cortex. We have somatotopic, retinotopic and tonotopic **maps**.

The structure of the stimulus is somewhat preserved in the pattern of neural sensitivities.

Distortions occur, because some aspects of a signal are more important than others. Compare sensitivity of lips to heels. Compare foveal vision (high focus, where you are looking directly) with peripheral vision (what you see out of the corner of your eye).

Are these maps representations?

Do they correspond to something?

Do they function within a system by virtue of that correspondence?

Can we identify the translations that lead from cause to effect. from world to neural map?

A big mystery: About-ness, or Intentionality

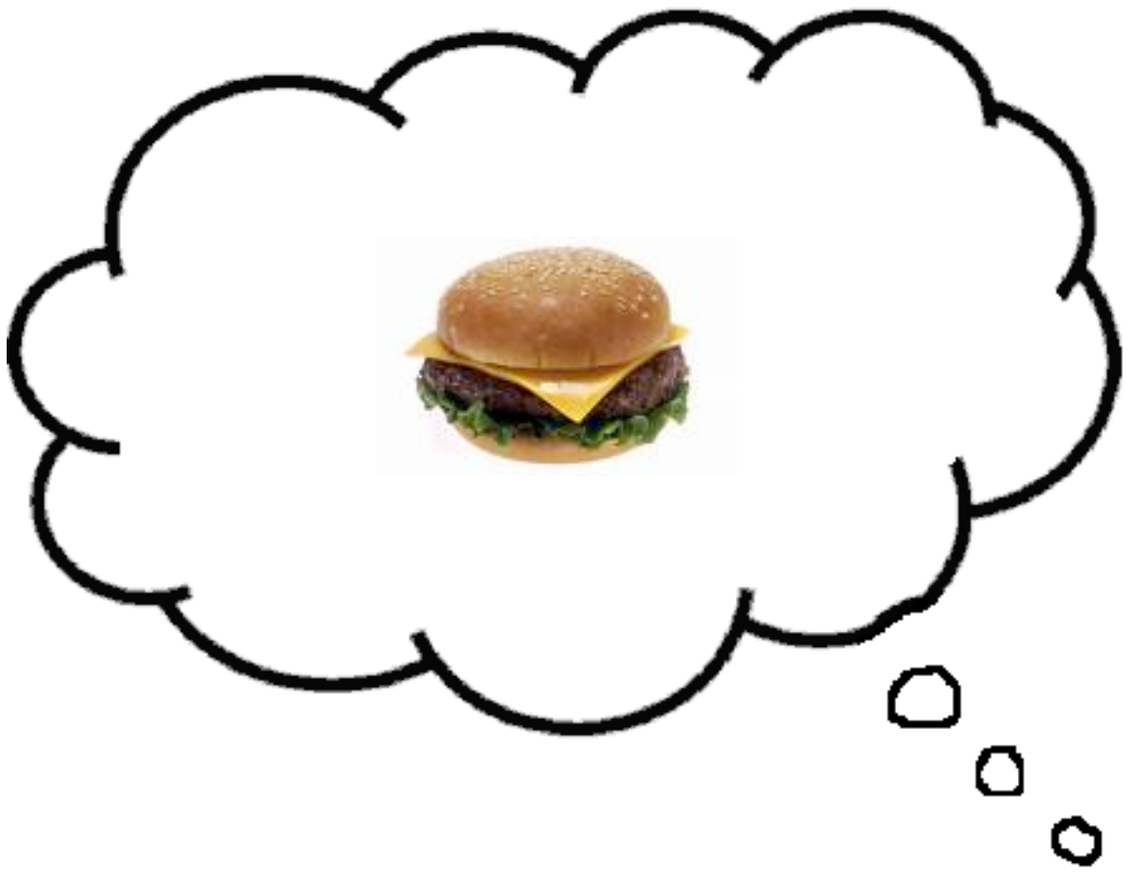

If thinking cat is thinking about a cheeseburger, and the cheeseburger is not there, there must be something in his neural activity that corresponds to the cheeseburger.

Some pattern of activity is "about" the cheeseburger. This relationship of "about-ness" is also called "intentionality".

The cat, so it is argued, must have a *representation* of the cheeseburger that can stand for the cheeseburger within his thinking system.

Hard Question: Would a representation of a cheeseburger have to look like a cheeseburger? Would it have to taste like a cheeseburger? Does it need any of the properties of a cheeseburger?

Propositional Representations

If thoughts are like sentences, then a representation of a cheeseburger is a kind of symbol

I want a 

 s are tasty

Yes! We have no  s

In propositional representation, there is no need for the symbol to physically resemble its referent.

I want a π

π s are tasty

Yes! We have no π s

Moving beyond sentence-like representations...

COMP 20090: Introduction to Cognitive Science

Picture a hippopotamus.....

Are mental images like pictures?

Are the elements spatially arranged?

Can one thing block another?

Are the elements fully specified?

Can you see things you didn't think of?

Not all thought and imagination is like language!

The great mental imagery debate:

Are the representations that are used in thinking like symbols and propositions?
Do they bear any necessary relation to their referents?
Are they all symbols?
Are brains big symbol manipulating machines?

Shepard and others

- Shepard and Metzler (1974):
  - Are two stimuli identical except for rotation?

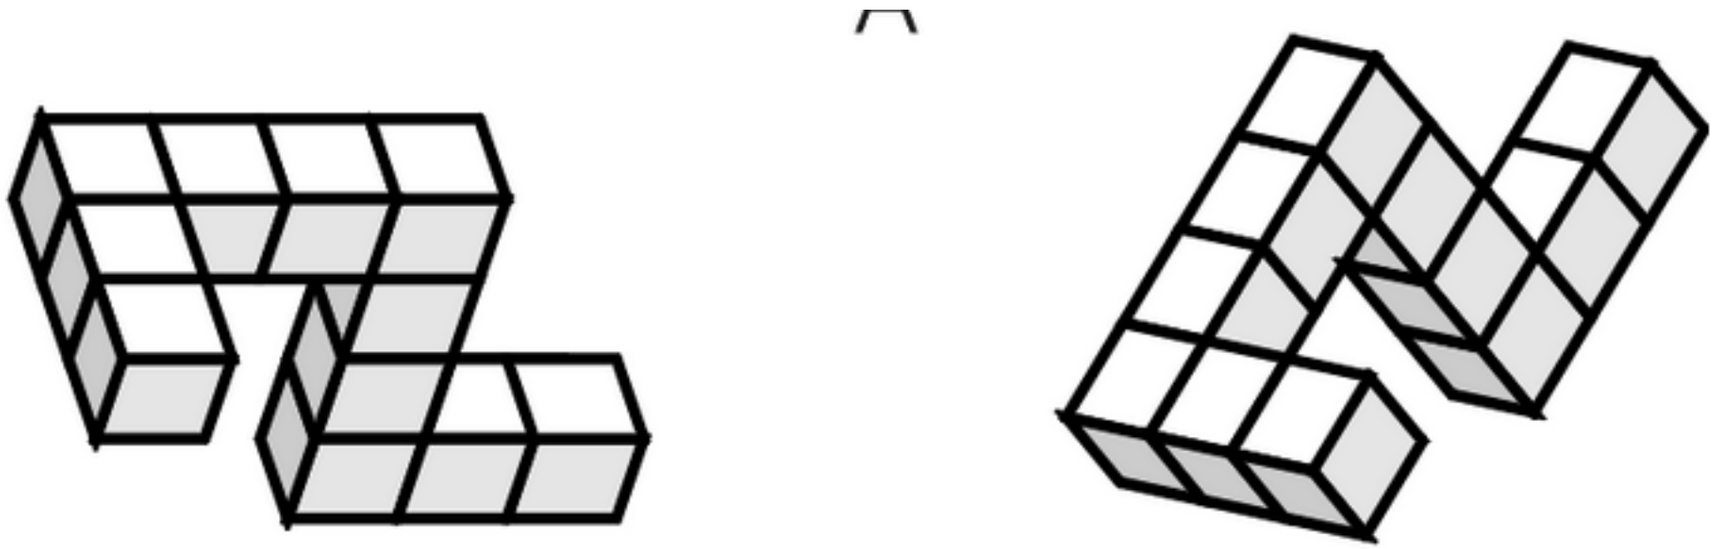

Another.....

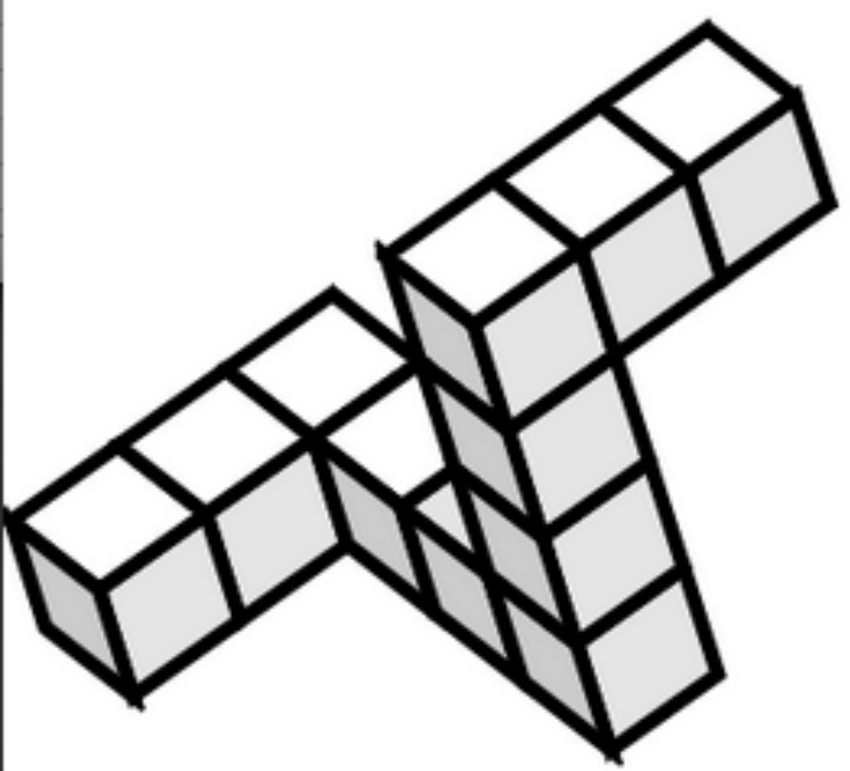

B

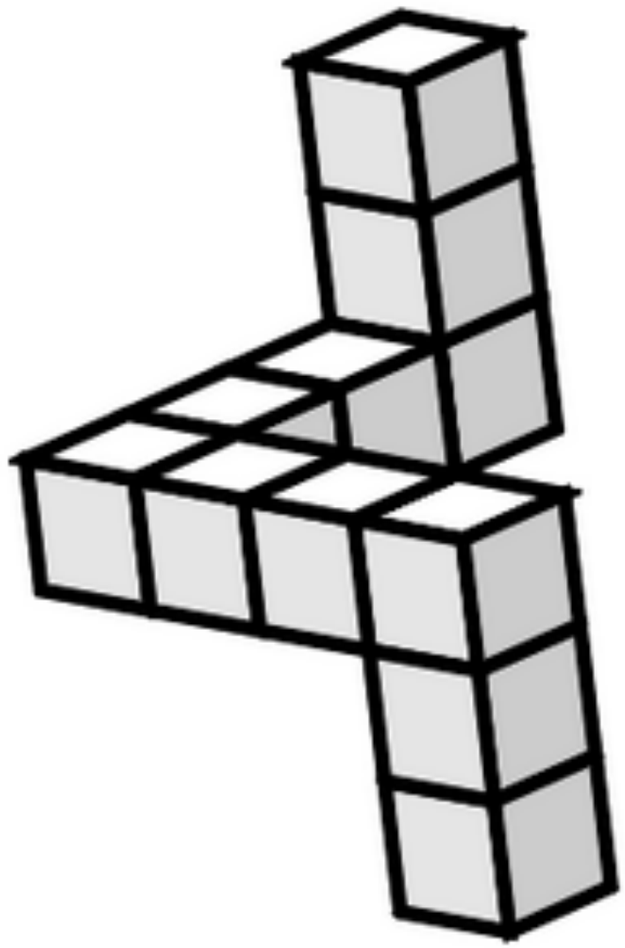

And another.....

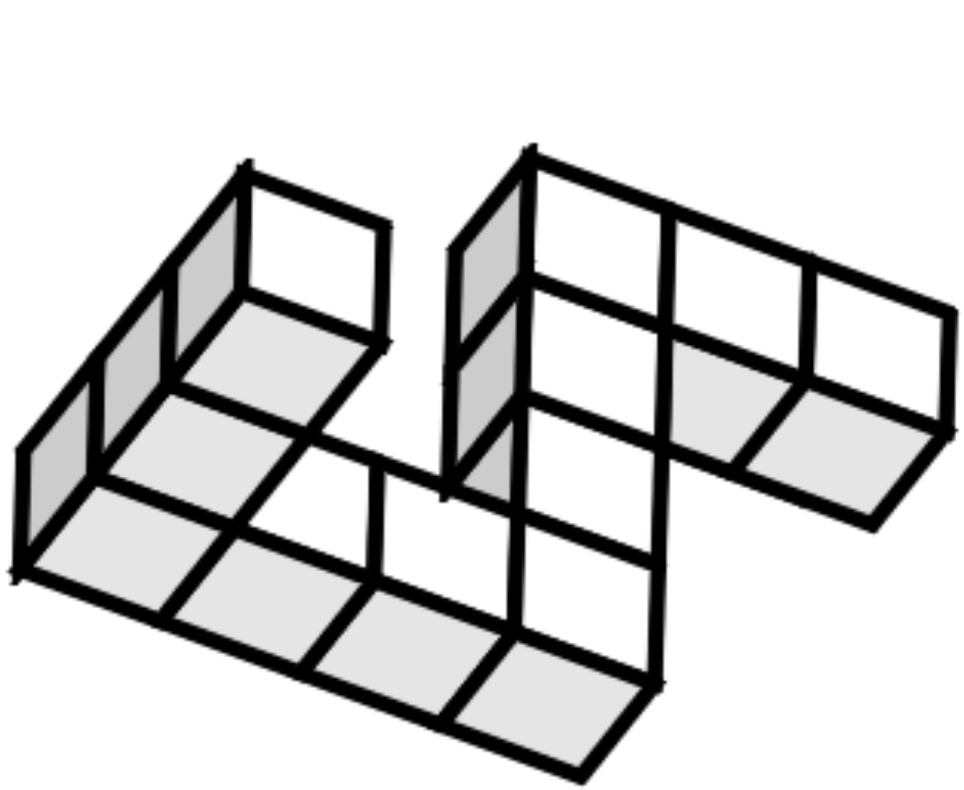

C

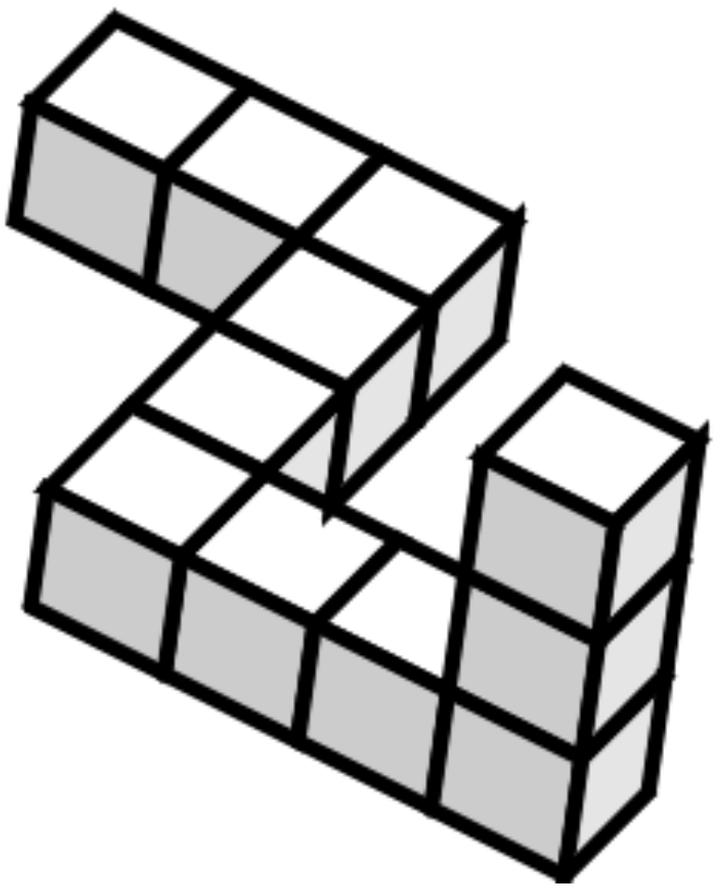

Results

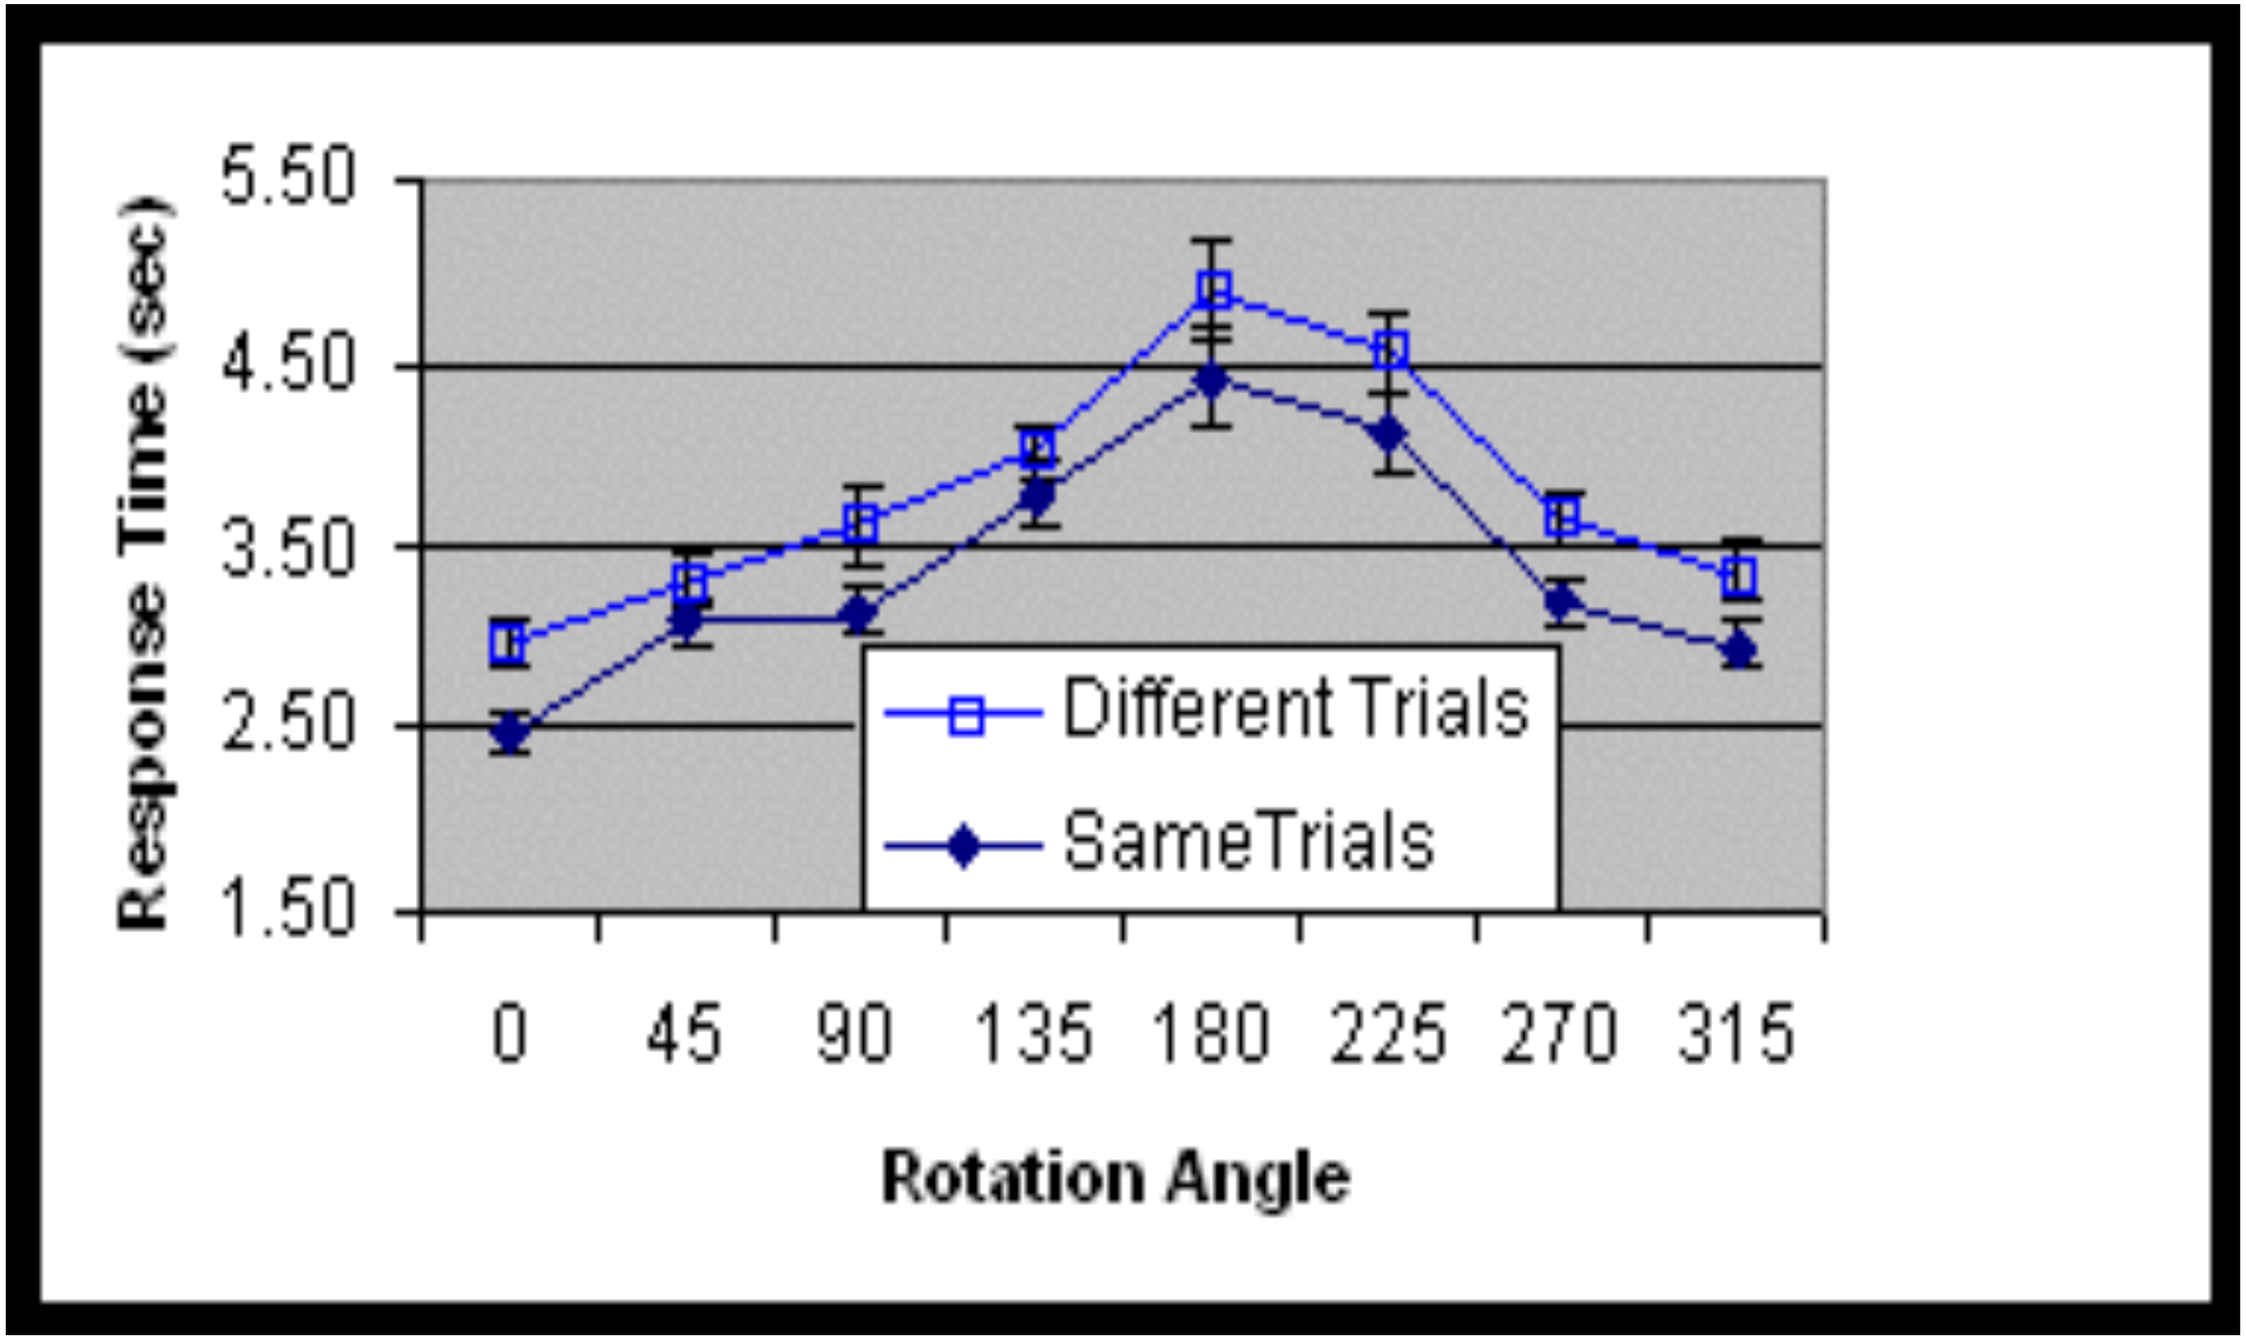
Response Time (sec)
5.50
4.50
3.50
2.50
1.50
0
45
90
135
180
225
270
315
Rotation Angle
Different Trials
SameTrials

An analogy

We can represent numerical quantities in several ways

Each form represents the same thing, but the operations that each representation supports may be different: easy with some representation and hard with others

Do the following addition using either arabic numerals or roman numerals:

3,456 + 284 or MMMCDLVI + CCLXXXIV

Grand Challenge 1:

Given a phonebook, look up the number of Mr. Thelonious Winterbotham

Grand Challenge 2:

Given a phonebook, find who owns the number 2824758

Grand Challenge 3:

Name three celebrities who are a bit like Kanye West

Grand Challenge 4:

Name three celebrities where the fourth letter in their name is ‘u’

Grand Challenge 5:

Name three words that rhyme with *flap*

Some challenges are harder than others.

For Challenges 1 & 2, you can see why this is: the information in the phonebook is organized in a fashion that makes it easy to find a number given a name, but nearly impossible to find a name given a number.

The knowledge represented in the phonebook is represented in a specific fashion. This hugely affects how that knowledge can be accessed and manipulated.

Now, what do Grand Challenges 3, 4 and 5 tell you about how your knowledge is represented or organized?

Knowledge Representation

COMP 20090: Introduction to Cognitive Science

Some things I know:

The capital of Ireland is Dublin.

I know how to tie my shoelaces.

I can find my way from my bedroom to the kitchen in the dark

Bonus question: *where* is the knowledge in each of these? Is is plausibly in your head?

The capital of Ireland is Dublin.

This is *explicit* knowledge. It is readily expressed in language.

It is disconnected from any specific *context*

I can state this nugget of wisdom, irrespective of where I currently am

There must be some sense in which my brain(+body) allow me to state this fact.

We can refer to this potential as a *representation*

We have vast amounts of this kind of explicit knowledge

It is often called *propositional* knowledge

It is something of a mistake to think that most of what you know is known in this form, or that knowledge is the accumulation of facts

But explicit propositional knowledge has been much studied

We can infer some things about how it is represented by looking at the kind of questions that are easy or hard to answer.

One area in which propositional knowledge has been studied intensively is *Artificial Intelligence*

Artificial Intelligence systems were originally conceived of as systems with a great deal of propositional knowledge

And so knowledge representation, i.e. the manner in which such knowledge might be stored, became an important research topic.

Maybe everything is unconnected to everything else

CAT
...has fur
...is an animal
...chases mice
...meows
...dislikes water
...has a tail
...likes to be petted.....

DOG
...has hair
...is an animal
...chases cats
...barks
...likes water
...has a tail
...likes to be petted.....

This seems to be inefficient

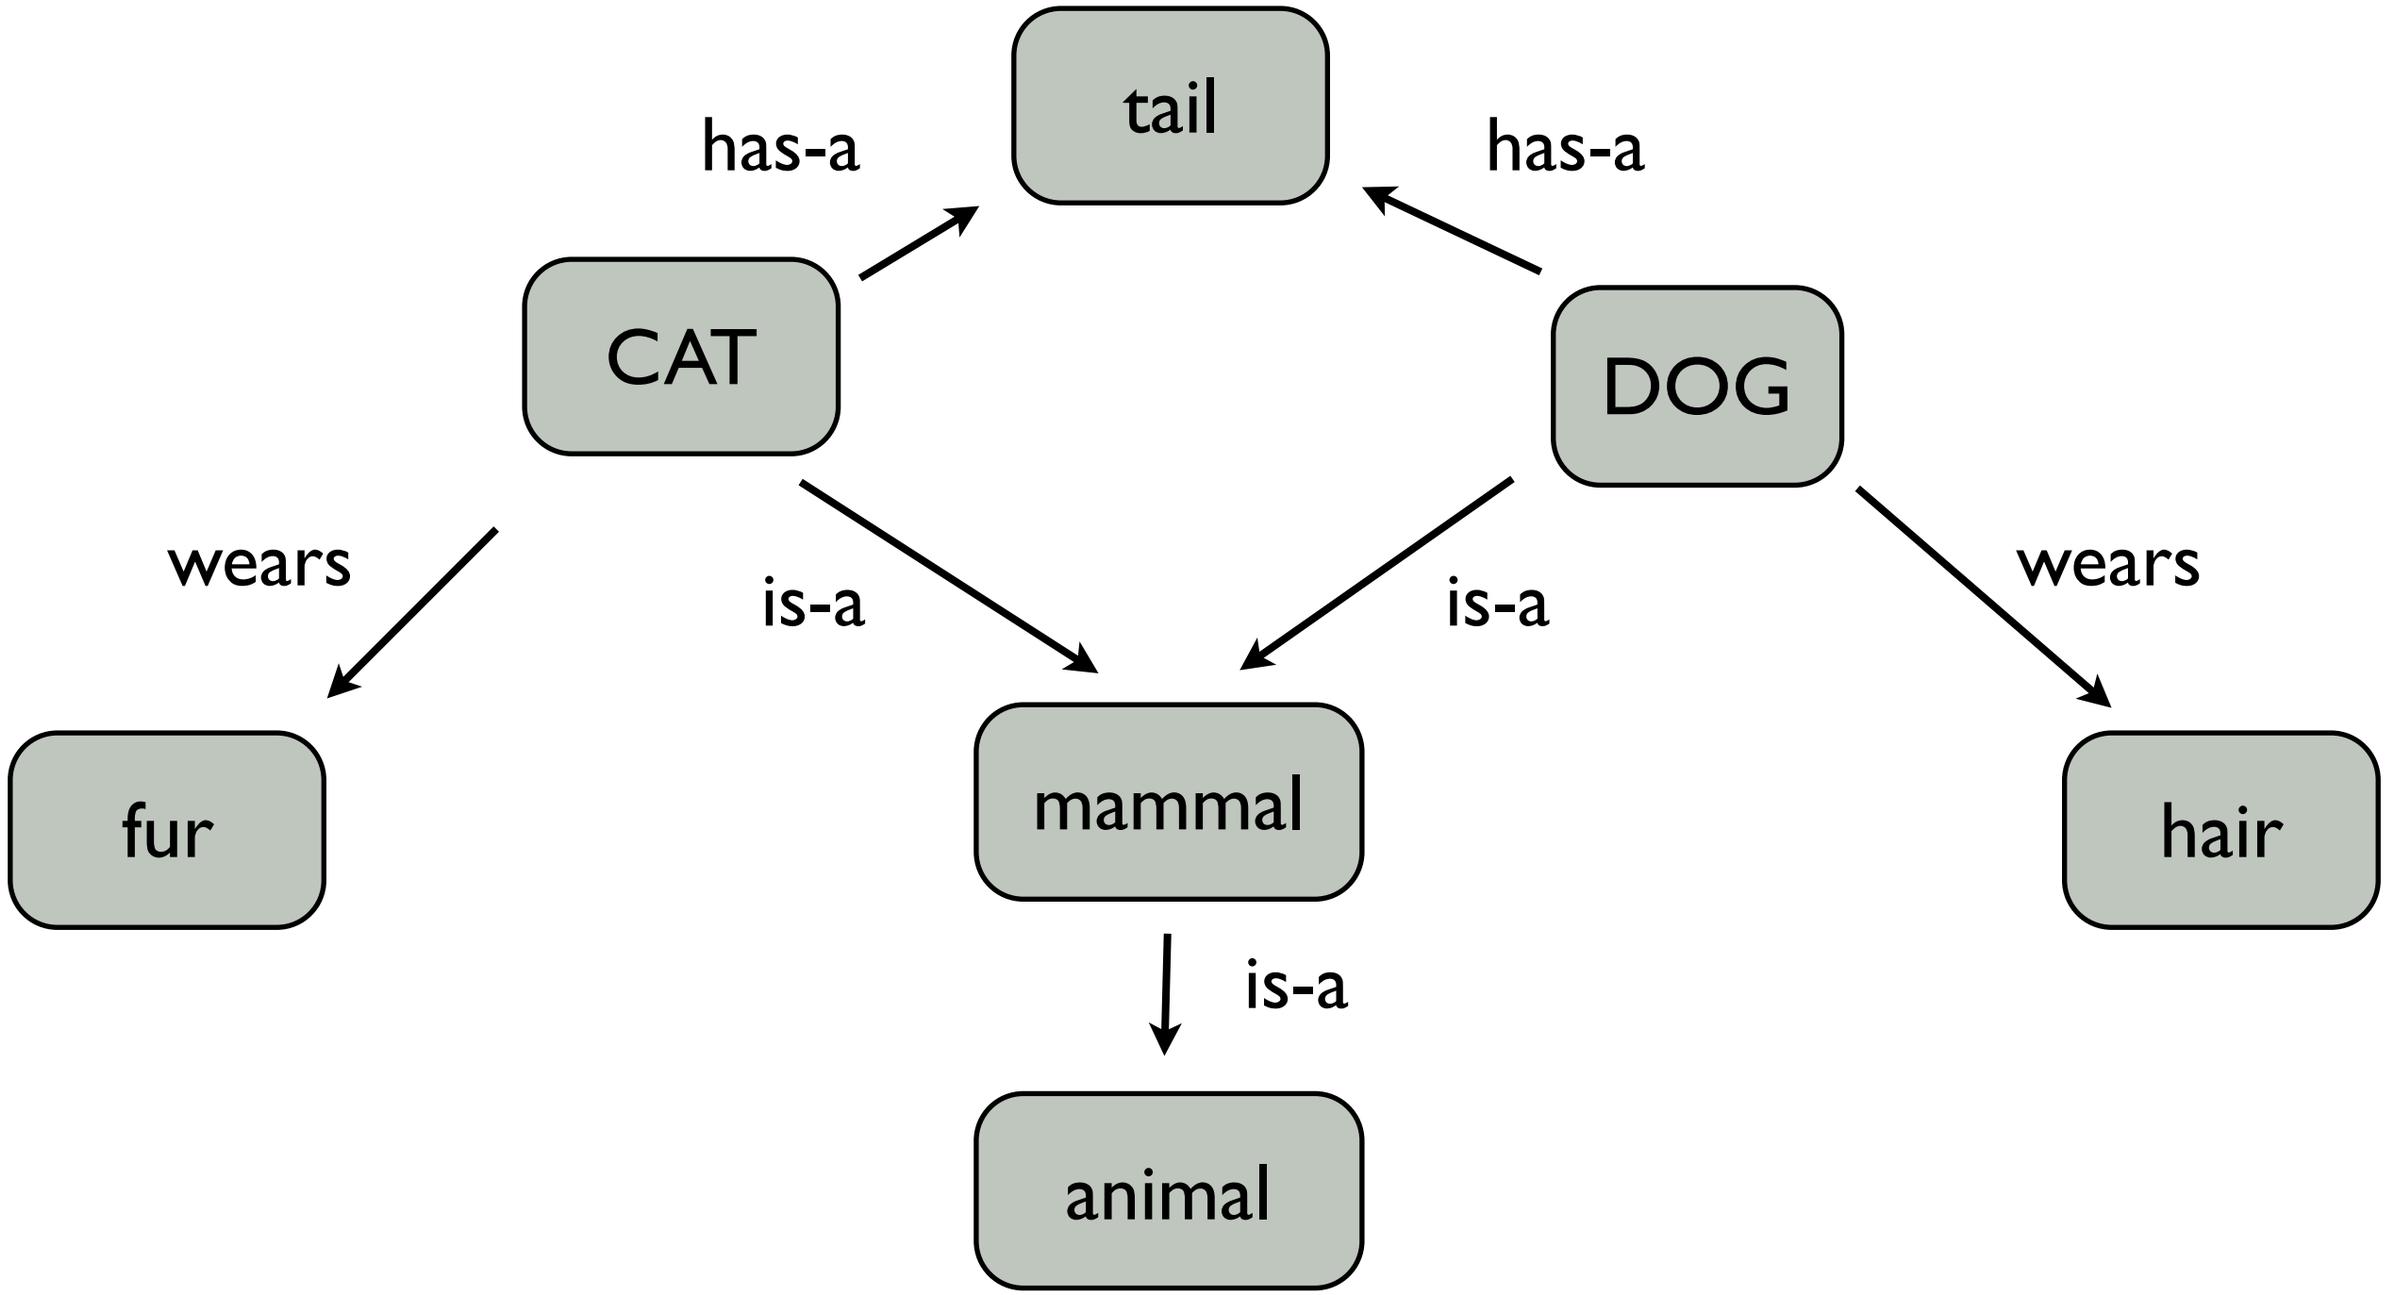

Semantic Network

A form of encoding of propositional knowledge from early AI

Storing information that will allow the expression of proposotional knowledge will have to take into account the way we organise our *categorical understanding* of the world

What, for humans, is the way explicit knowledge is organised?

Classical categories

- Some features are *necessary* for *X* to be an instance of *Y*
- Some set of features are *sufficient* for *X* to be an instance of *Y*
- What do you think the *necessary* and *sufficient* conditions for being a member of the following categories:

A bird
A pope
A planet

The idea of a category is central . . . Most symbols (i.e. words & representations) do not designate particular things or individuals in the world . . . Most of our words & concepts designate *categories*. There is nothing more basic than categorization to our thought, perception, action & speech. Every time we see something *as* a kind of thing, for example *as* a tree, we are categorizing.

George Lakoff: Women, Fire and Dangerous Things

Dyirbal language (an Australian Aboriginal language)

- Bayi
  - Men, kangaroos, possums, bats, most snakes, most fishes, some birds, most insects, moon, storms, rainbows, boomerangs, some spears
- Balan
  - Women, bandicoots, dogs, platypus, echidna, some snakes, some fishes, most birds, scorpions, crickets, hairy mary grub, water or fire, sun, stars, shields, some spears, some trees
- Balam
  - Edible fruit and plants that bear them, tubers, ferns, honey, cigarettes, wine, cake
- Bala
  - Parts of the body, meat, bees, wind, yamsticks, some spears, most trees, mud, stones, noises, language

106

(example from Lakoff...)

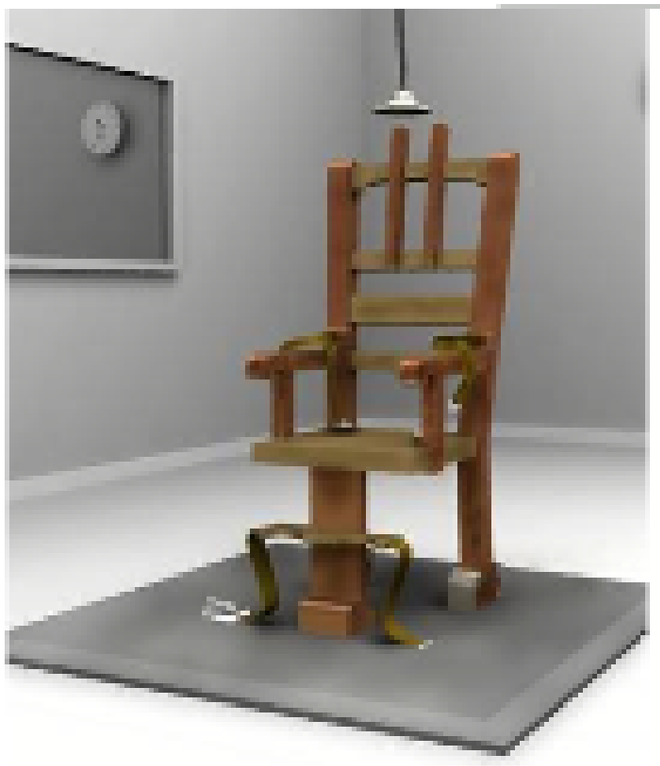


What makes a chair a chair?

Resemblance Theories: Prototype vs. Exemplar

The Standard = The acting representative of the category

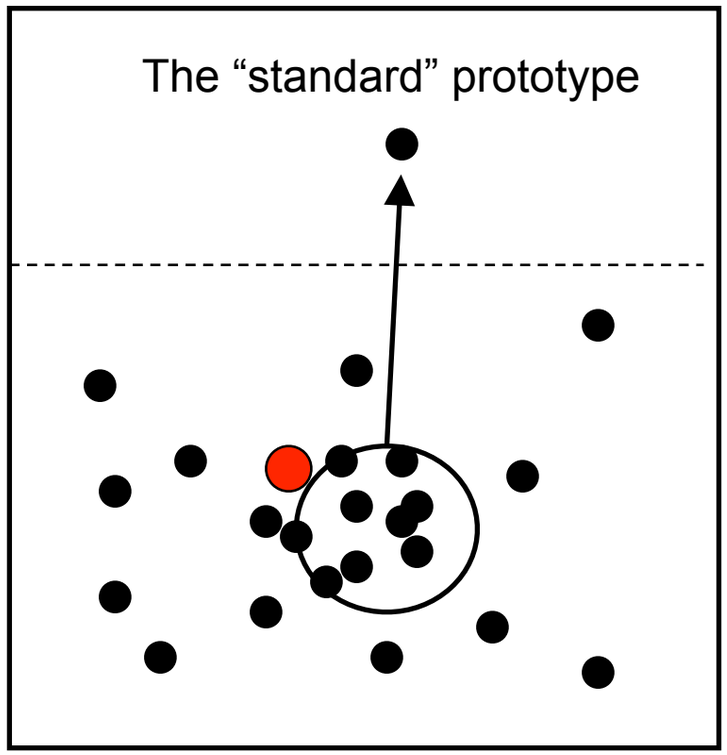

A prototype is an abstract representation that is derived from the “center of mass” of the features of all the objects in the category.

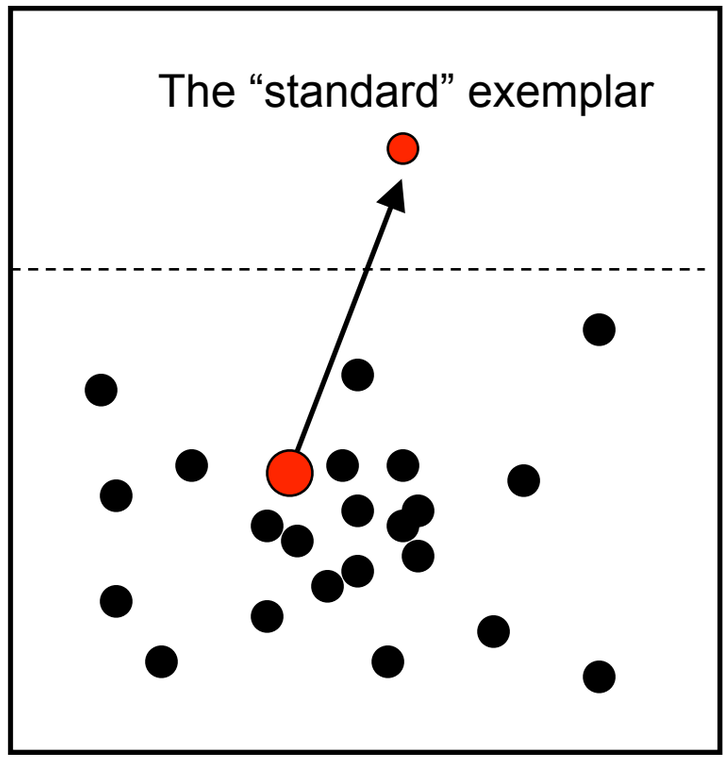

An exemplar is a concrete representation. A specific instance that happens to be the most active in memory. More typical members tend to be the most active, on average.

Prototypes	Exemplars
A prototype is an abstraction, based on previous experience	An exemplar is one specific example of a concept
Storing prototypes is cheap (there are few), but you need a lot of experience to develop a prototype	Storing exemplars is expensive, but every instance of a concept can function as an exemplar

- Eleanor Rosch
  - Categories are created by *generalization* and *specialization* of the basic level
  - At the basic level, categories are maximally distinct, i.e. They maximize perceived similarity among category members and minimize perceived similarities across contrasting categories

Categories are not merely organized in a hierarchy from the most general to the most specific, but are also organised so that the categories that are most cognitively basic are "in the middle" of a general-to-specific hierarchy. Generalisation proceeds upward from the basic level and specialization proceeds down

Which of these are basic level categories?

1. Activity, Sport, Football, Under-15s Football

2. Object, Furniture, Table, Dressing Table

3. Animal, Bird, Gull, Herring Gull

Is everything you know represented in your brain?

Talk of representation of knowledge makes most sense within the computational view of mind and brain.

This view co-exists with other views, for whom the notions of “representation” might play a different (or no) role.

There is much disagreement here.

Different approaches here will usually differ in how important they consider *context* to be. A great deal of our knowledge/memories/understanding seems to be expressible only in corresponding supporting contexts

External memory..... is it important?

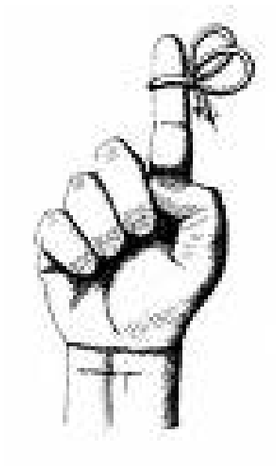

The homuncular fallacy. The Cartesian theatre.
Our essential dependence on the world around us
does not lend support to the view that knowledge is “in the head”.

(Also, there is nobody in your head)

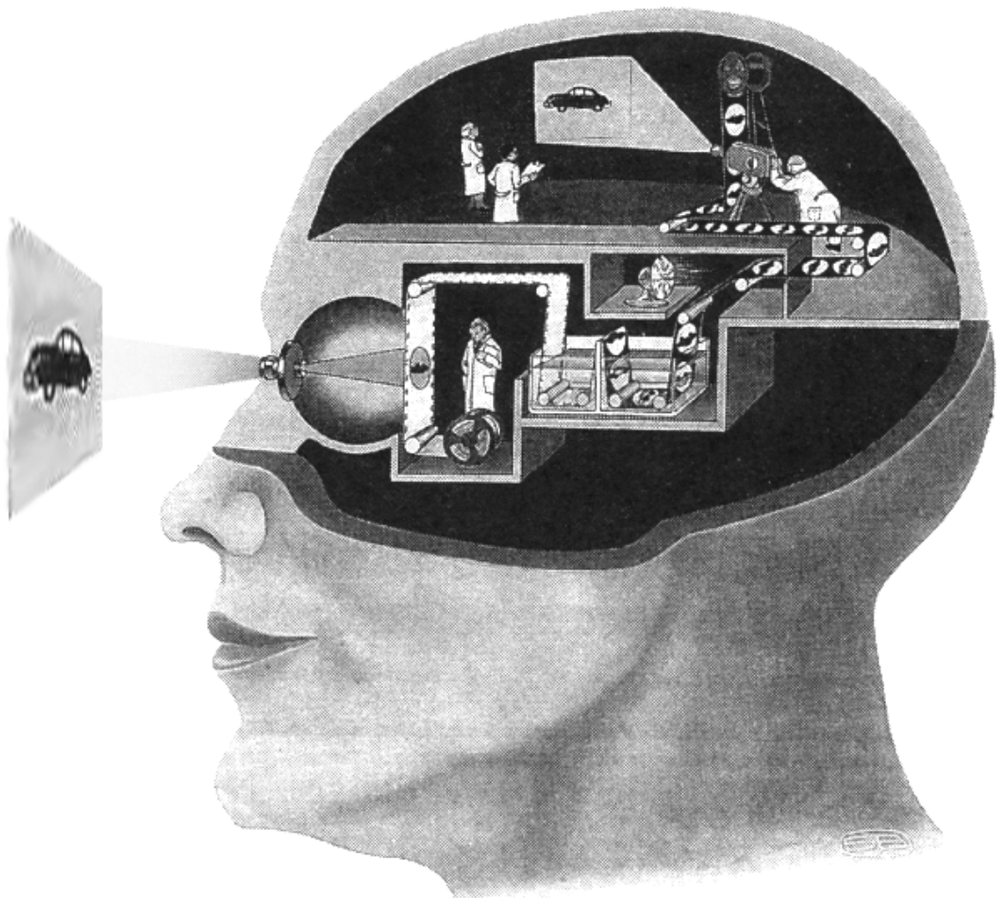

Excerpt from “Being John Malokvitch”

There is a danger of infinite regress if representations have to be presented to somebody.

In "Being John Malkovich", there are literally other people inside his head, so that what he "sees" is input to other minds/bodies. That is regress. If each of them has people in their heads, then infinite regress threatens.

The experience of building walking robots

Rodney Brooks: MIT

Brooks: When I started building walking robots, I built six legged robots and rather than have them compute ahead of time a stable way of walking I made the legs very sensitive to things that they touched in the environment and made it so that it was safe for the robot to fall down and had it learn to scramble over rough terrain, feeling its way as it went instead of sitting back looking at the whole terrain in front and computing the optimal path through, my robots got in down dirty with the environment and interacted with the environment at every step.

"The world is its own best representation"

Representationalism

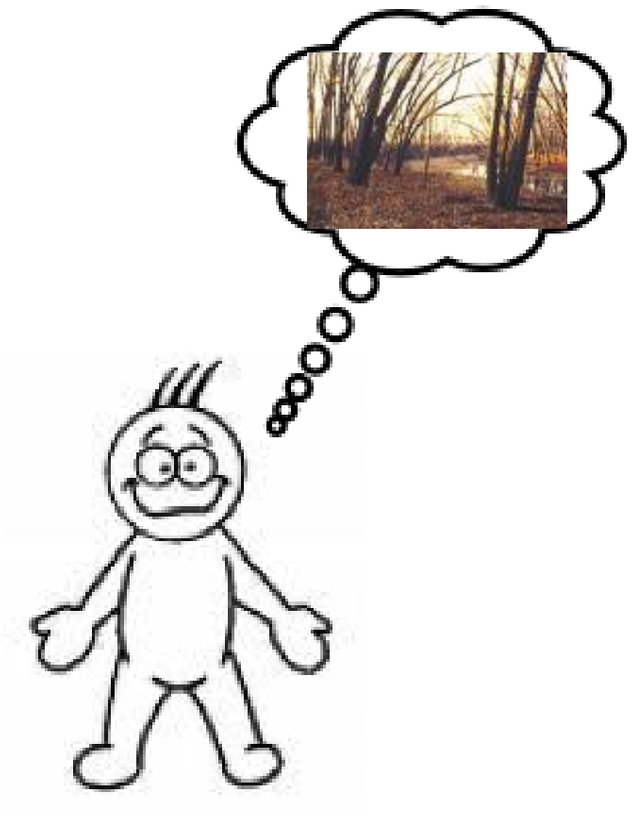

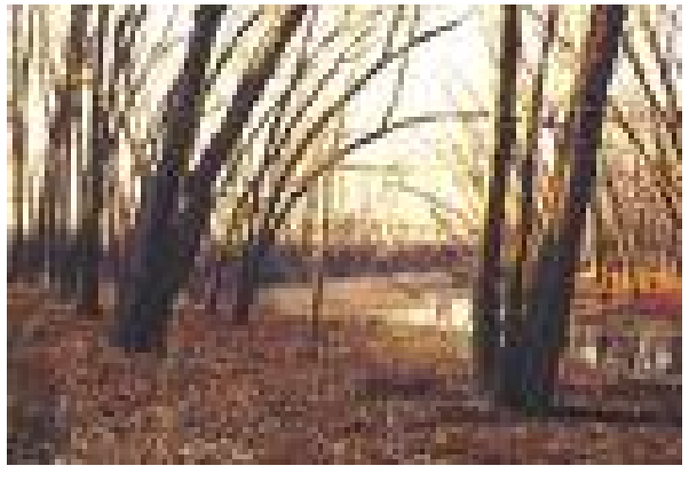

Embodied and Embedded cognition

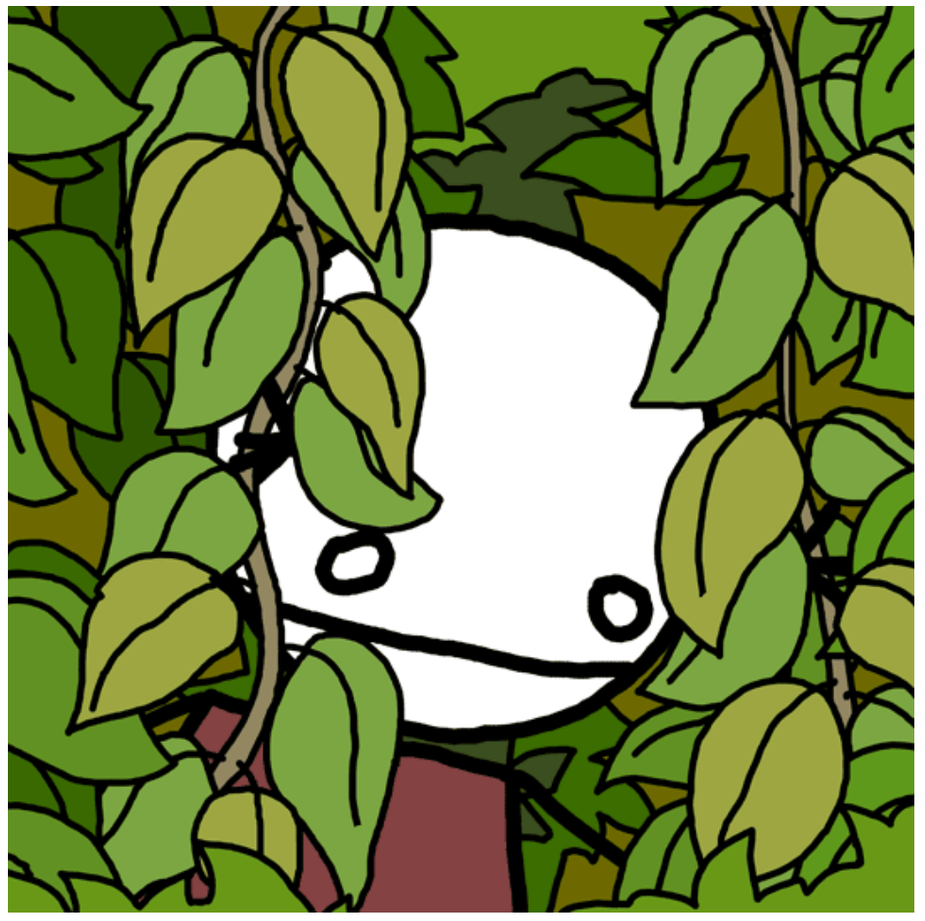

The [external] world is its own best representation

--Rodney Brooks

Problems with a ‘Brain-in-a-Vat”

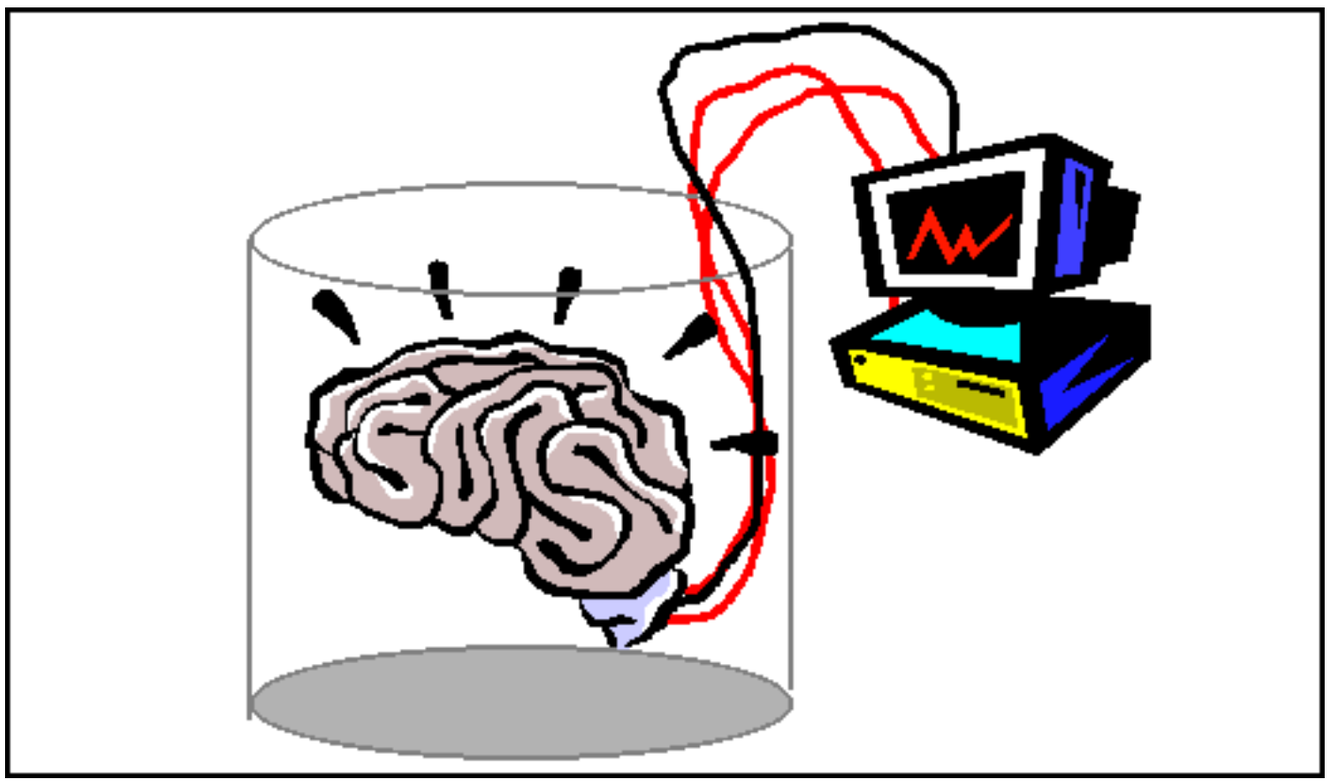

The brain-in-a-vat thought experiment is intended to make some assumptions of some approaches to brains explicit.

It is *not* just there to induce you to say “Gosh, what if it’s all a simulation”

- Treats the brain as an input-output device
- Assumes that experience arises based on inputs
- Assumes that electrical signals going into the brain generate experience
- Separates perception from action

Much of the discipline of Artificial Intelligence has been the attempt to find the most useful form of knowledge representation


KTSC

The Uncanny Valley

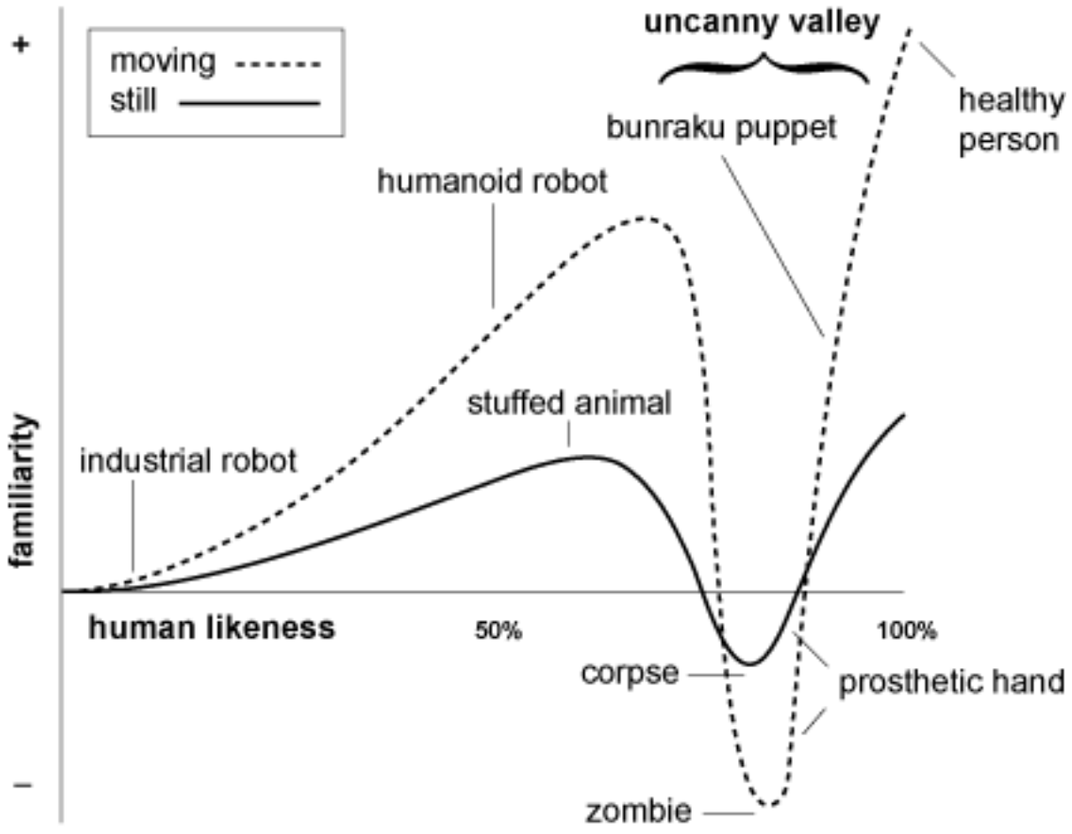

http://www.youtube.com/watch?v=CNdAlPoh8a4

How do we build a system that incorporates human-like knowledge of the world?

Approach 1: Add More Facts

Terry Winograd's SHRDLU program

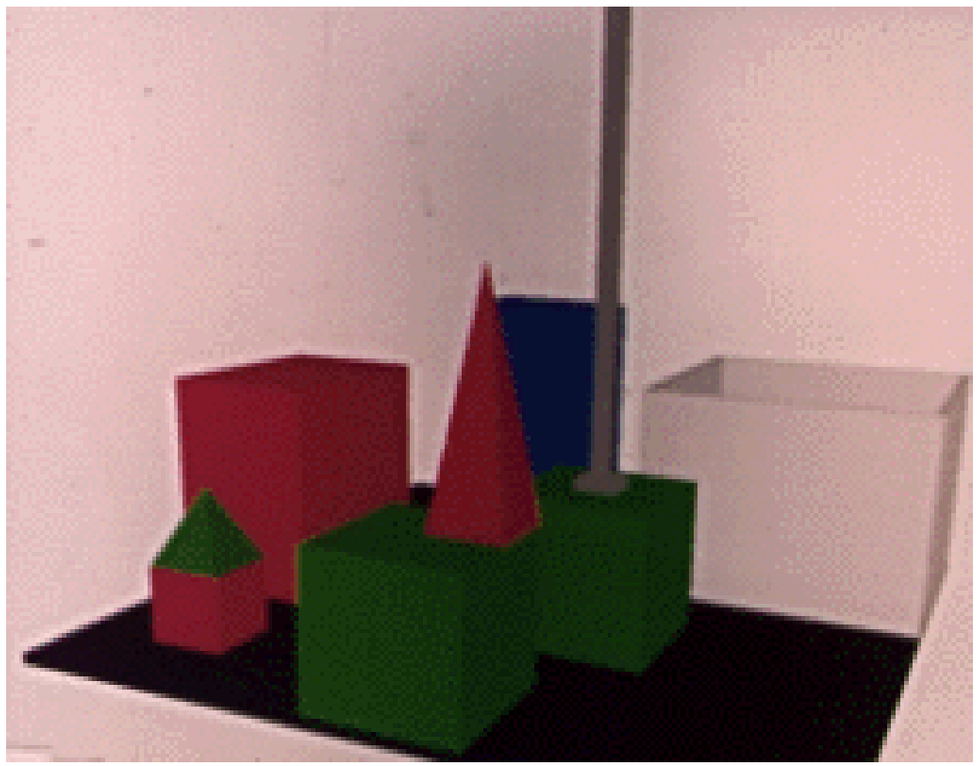

Person: PICK UP A BIG RED BLOCK.
Computer: OK. (does it)
Person: GRASP THE PYRAMID.
Computer: I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN.
Person: FIND A BLOCK WHICH IS TALLER THAN THE ONE YOU ARE HOLDING AND PUT IT INTO THE BOX.
Computer: BY "IT", I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER THAN THE ONE I AM HOLDING.
Computer: OK. (does it)
Person: WHAT DOES THE BOX CONTAIN?
Computer: THE BLUE PYRAMID AND THE BLUE BLOCK.
Person: WHAT IS THE PYRAMID SUPPORTED BY?
Computer: THE BOX.
... etc

Sample dialogue available under 'Reading and Resources'

Some great things about SHRDLU:

It could learn (some things)

It "knew" what it "knew" and what it didn't "know"

It interacted using natural language

Problems: No common sense

Successes in various 'blocks worlds' did not scale to systems dealing with the real world

Sample problem:

Nora is in the house

Where is Nora's nose?

Cyc Foundation:

"Imagine a world in which every single person is given free access to ***programs that reason with*** the sum of all human knowledge. That's what we're doing."

The CYC project took this approach to extremes:

Add more facts.....

Cyc Reasoning System

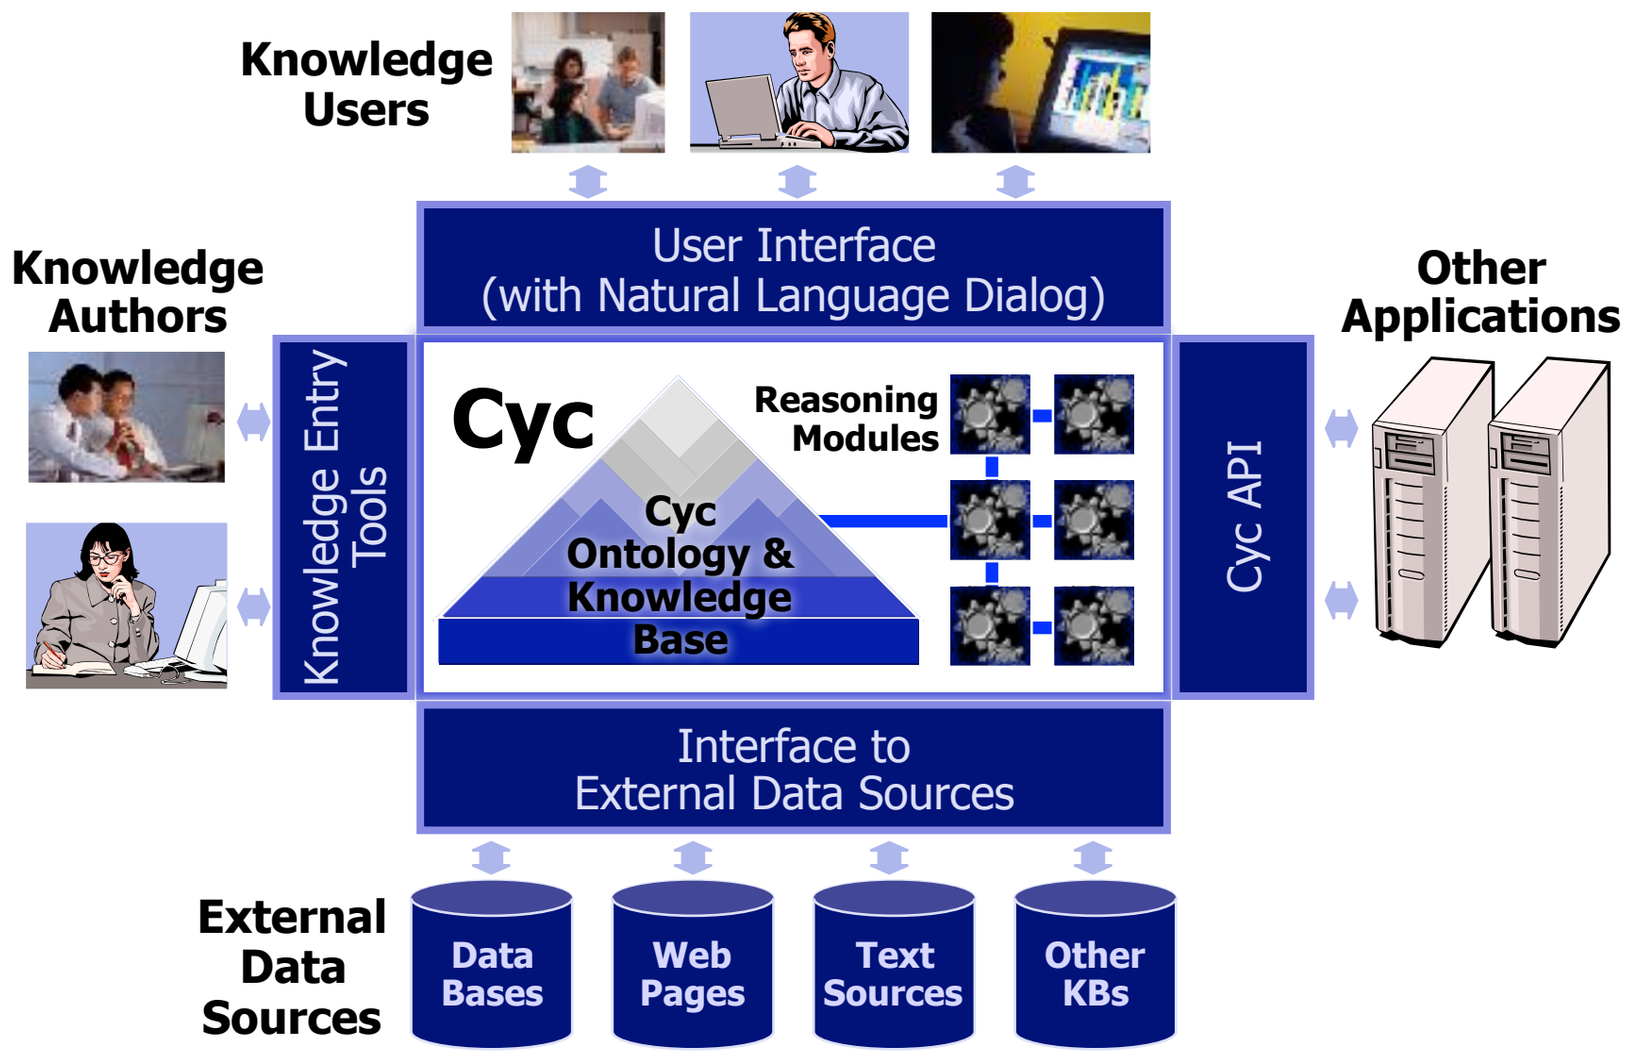

Common sense reasoning is hard (for computers)

The set of facts which might be relevant to solving any problem is essentially infinite. Small toy world solutions don't scale up to real world situations.

This is the *frame problem*.

Nora goes into the house. Where is Nora's nose?

One recent product of this line of reasoning: IBM's Watson, the Jeopardy-busting computer.

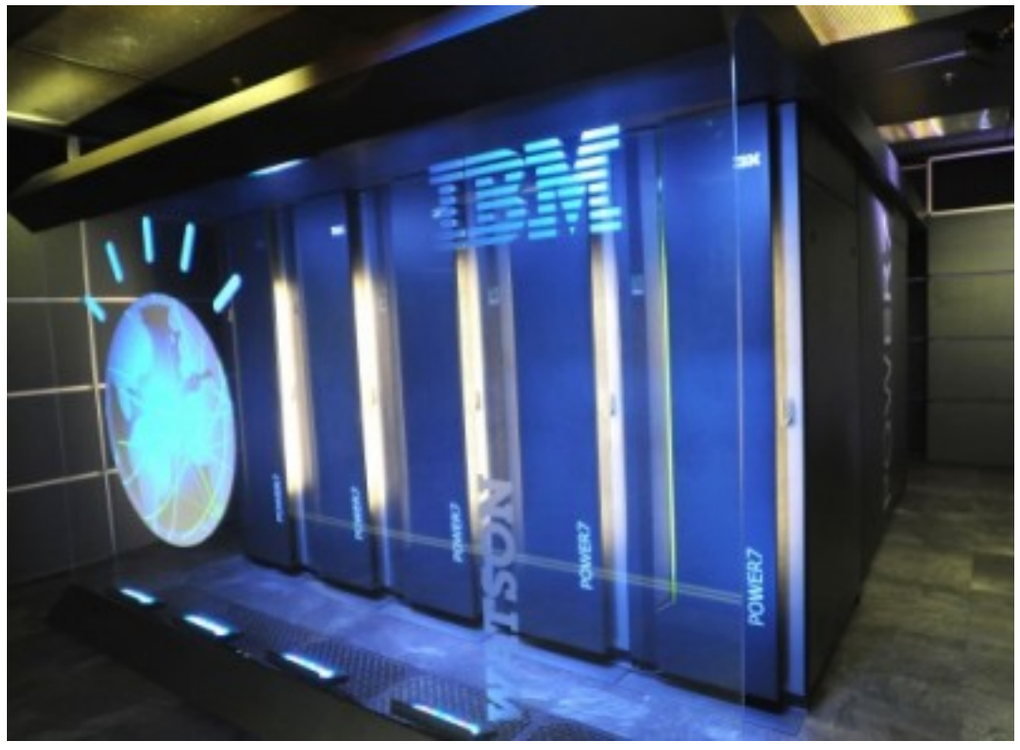

The optional reading material for this week includes a good account of just what Watson does and does not do.

Moravec's paradox

contrary to traditional assumptions, the uniquely human faculty of *reason* (conscious, intelligent, rational thought) requires very little computation, but that the unconscious *sensorimotor* skills and instincts that we share with the animals require enormous computational resources. (source: wikipedia)

Steven Pinker

The main lesson of thirty-five years of AI research is that the hard problems are easy and the easy problems are hard. The mental abilities of a four-year-old that we take for granted – recognizing a face, lifting a pencil, walking across a room, answering a question – in fact solve some of the hardest engineering problems ever conceived.... As the new generation of intelligent devices appears, it will be the stock analysts and petrochemical engineers and parole board members who are in danger of being replaced by machines. The gardeners, receptionists, and cooks are secure in their jobs for decades to come.

(In Defense of Dangerous Ideas, 2007)

Perceptually guided exploratory action, such as searching, may not require a complex cognitive architecture supporting abstract representation: