Perceptual Symbol Systems (1999)

Paper Review by: Christian Hahm, *Temple AGI Team*

Paper by: Lawrence W. Barsalou, *Emory University*

Summary

- Abstract & Introduction
- Core Properties
- Derived Properties
- Implications
- Conclusion

Abstract & Introduction

Cognition as Perception?

- Historically, theories of cognition and knowledge were inherently **modal** (perceptual)

- With the advent of computers, cognitive science has been trying but struggling to describe coherent **amodal** (non-sensory, non-perceptual) theories of cognition that are purely symbolic, syntactic or mathematical

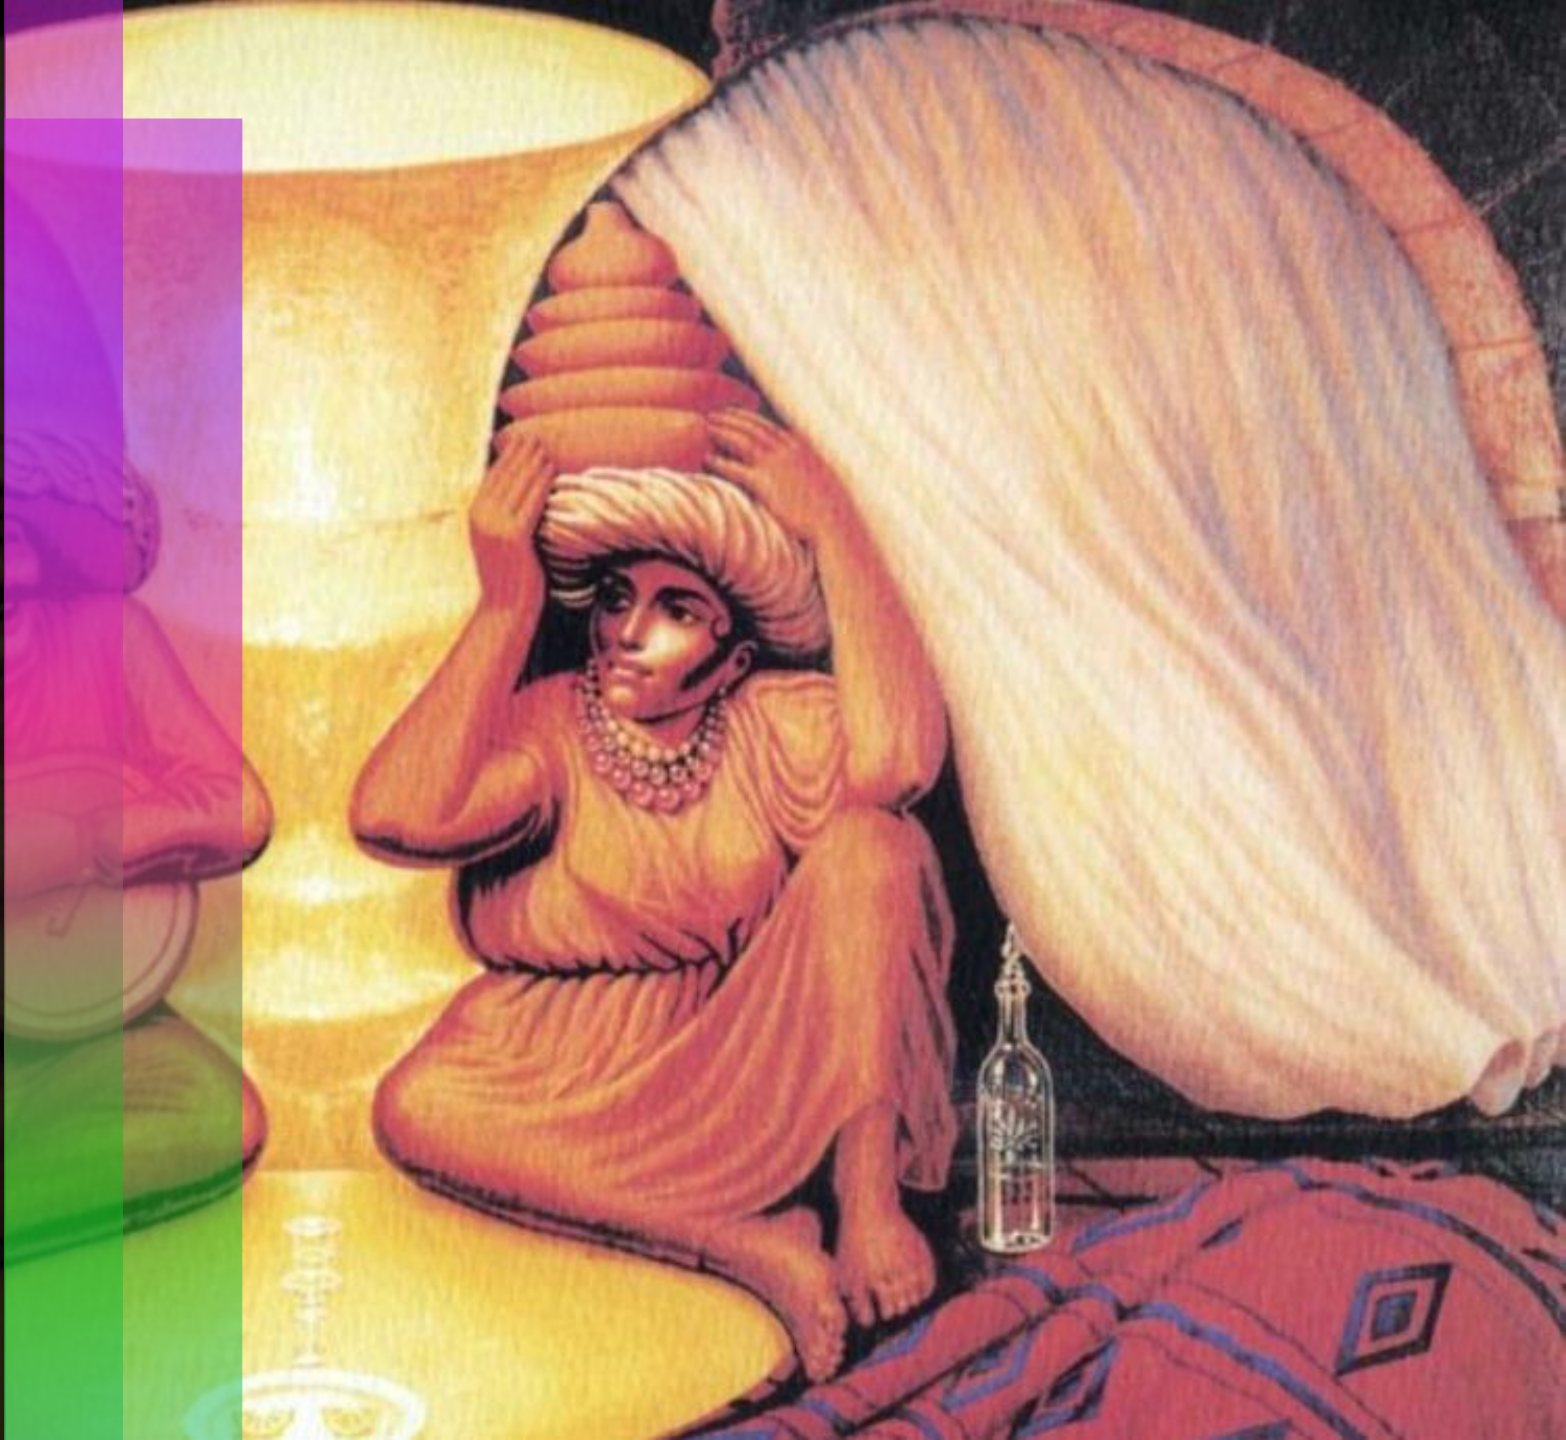

Concepts with Modality

- Perception seen as unfeasible, since it seems to require sensation “recording”, which needs massive storage and computation.
- To counter this issue, some argue that perceptual data exist solely in working memory, but are eventually discarded, whereas long-term concepts in the memory are amodal, mapped to by these perceptual data but stored separately.
- However, a cognitive system interacting with the world is already accepting **modal (sensory)** experiences, and can use components of the experience itself to form and ground its symbols and concepts.
  - Firstly, it is more direct.
  - Secondly, the sensation itself provides justification for further conceptualization.
  - Thirdly, no AI system designer can define a meaningful objective mapping between a given system’s **modal sensory experiences** (e.g. some neurons firing as photons hit the retina) and abstract **amodal symbols** (e.g. the word “dog” or “cat” in the system’s memory).

The Process of Perceiving

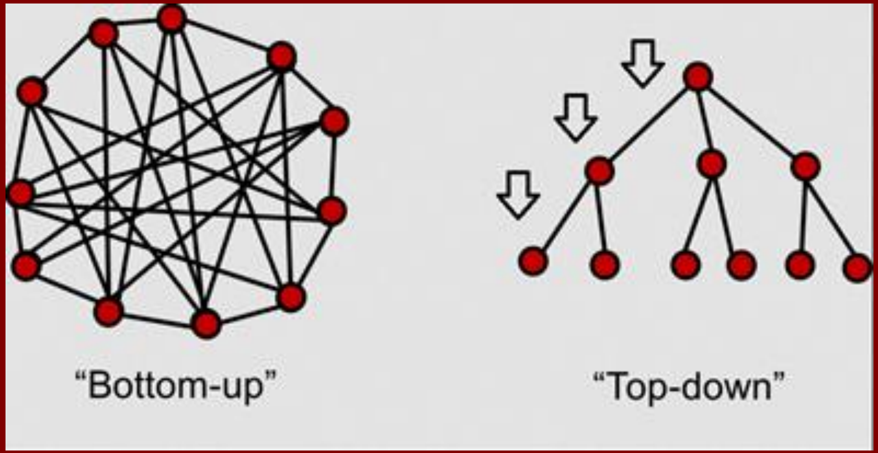

- During perceptual experience, association areas in the brain automatically capture and **associate bottom-up patterns of activation** in sensorimotor areas

- Later, **top-down** forces use **association areas** to partially **re-activate sensorimotor areas**, such as in a goal-driven process**.**

The Process of Perceiving

- The perception process is **selective**; the system perceives **specific features** or **components** of its sensory experience, **not** the entire experience.

- The system uses attentional mechanisms to extract **schematic** representations of perceptual components from its sensorimotor experience, storing and comparing them in memory.

- Over time, the system's perception + action feedback loop implements frequently extracted components as **perceptual symbols** within its memory

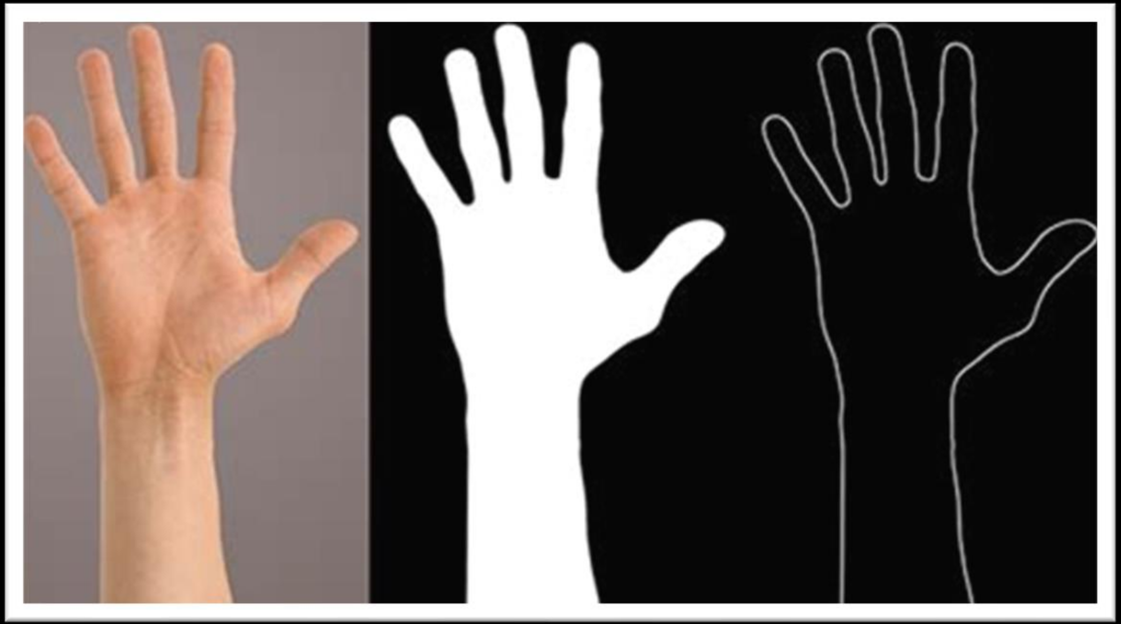

Symbol Grounding

- The same perceptual components are extracted and stored many times over many experiences and the system organizes its memories of them around a common frame.

- A well-grounded perceptual symbol acts as an internal “simulator” for that component (e.g. the color green, square shape, a “popping” sound). When the system forms many basic perceptual components, the **memory** contains a basis conceptual system that supports types, categorization, and categorical inference.

- Finally, the components are combined and used recursively to produce complex internal “simulations”. The system uses its internal simulations to formulate and ground higher-order abstract concepts.

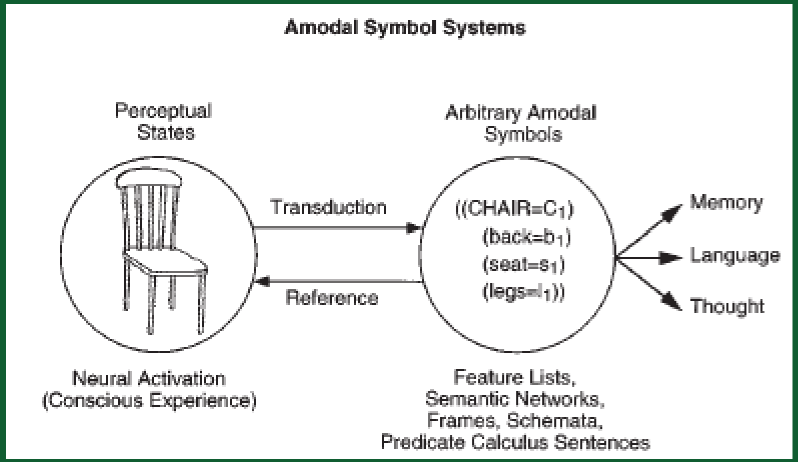

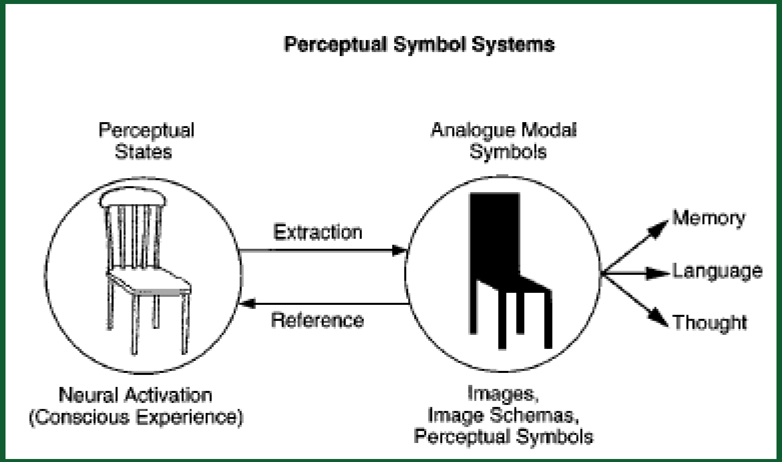

Recording Systems vs. Conceptual Systems

- Some argue that perceptual systems are infeasible, since they would be like **recording systems**: systems that store sensory inputs directly in full within the system's memory, either raw or in a compressed format, like a camera.

- When we think of perception, we tend to imagine the entire conscious experience: a holistic perceptual image or experience. However, it may be more correct to view perception as the mind selecting out components of the sensory experiences to use and represent **productively**

- The formation and grounding of perceptual component concepts is inherently **dynamic**, since every new experience contributes to the concepts' meanings.

- Unlike recording systems, **conceptual** systems interpret sensory inputs by mapping sensory component **instances** within the data with similar "**type"** or **class** concepts in the memory. Importantly, this is different from the pre-programmed modal-to-amodal mappings discussed earlier since 1.) *the "type" concept is fluid in meaning and modally based* and 2.) *the system itself performs the extraction and mapping, and this process created / contributes to the "type" concept itself*

Core Properties

6 core properties (of a basic conceptual system)

“

1. **Perceptual symbols** are neural representations in sensory-motor areas of the brain
2. They represent **schematic components** of perceptual experience, not entire holistic experiences
3. They are **multi-modal,** arising across the sensory modalities, proprioception* and introspection
4. Related perceptual symbols become integrated into a **simulator** that produces limitless simulations of a perceptual component
5. **Frames** organize the perceptual symbols within a simulator
6. **Words** associated with simulators provide linguistic control over the construction of simulations

” – Barsalou, (p. 582)

*sense of movement and body position

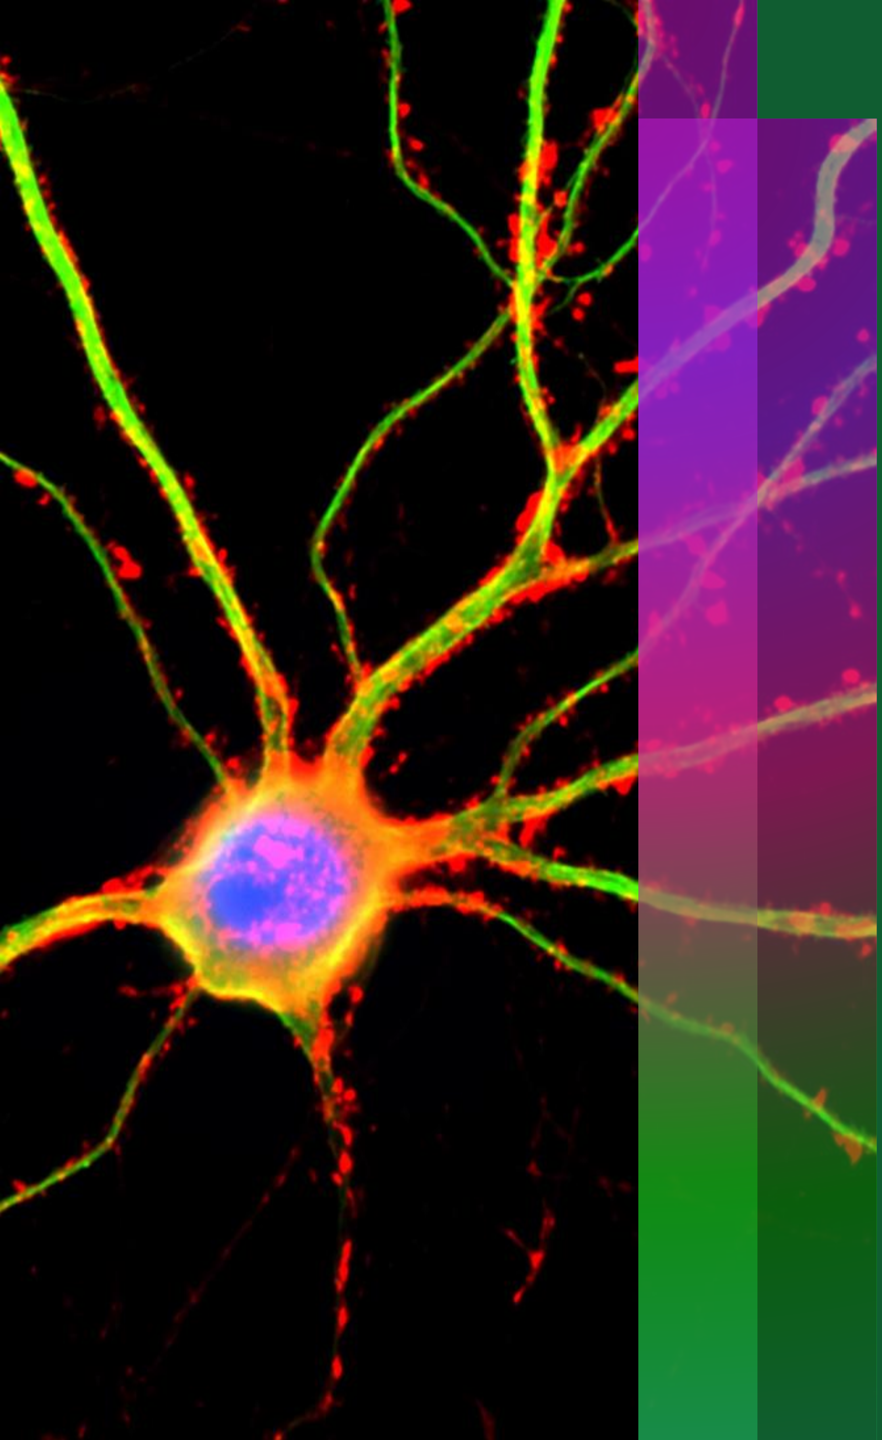

Core property #1: Neural representations in sensorimotor areas

- The main argument of this Core Property is that sensory neurons are not **pure** pre-processors; the knowledge stored here should be treated as conceptual in nature.

- Perceptual symbols are **not images or pictures**

- Instead, neurons capture **records of perceived events.** The information here is qualitative and functional (*e.g. presence of edge, shape, color, movement, etc.).*

- Certain areas of the brain associate low-level sensorimotor concepts using pre-defined (*i.e. content-agnostic*) formulas unique to the system's design, e.g. a temporal cortex temporally associates events

- Never in the perceptual process is the modal representation totally replaced by a separate amodal symbol. Nonetheless, sensorimotor neural processing may still act "pre-processor"-like in that it can be unconscious (*e.g. blindsight*), and so its results are "presented" to the conscious mind.

Core property #2: Schematic perceptual symbols

- A perceptual symbol does not represent an entire brain state during perceptual experience; only a **schematic** (i.e. subset) aspect, extracted from perceptual memories using selective attention which: 1.) *isolates specific information within the current perceptual experience*, 2.) *stores the isolated component in long-term memory.* Over time, after extracting many features, large numbers of schematics are stored in the memory.

- *Schematics* are just subsets of the neurons / concepts that were activated during a given perceptual experience. For example, if selective attention focuses on an object's shape, a record is made of the neurons that were activated during when this shape was being processed.
  - In this regard, it does not seem to matter *how* the system processes "shape"; what matters is which subset of concepts are commonly activated whenever "shape" is processed.

- Since a given perceptual symbol is schematic in nature, it can be used flexibly during perception depending on context, analogy, etc. Perceptual symbols can be generic or "indeterminate", like when they detect the presence of certain features in general, e.g. *lines, stripes, bumps, etc..* as opposed to specific features

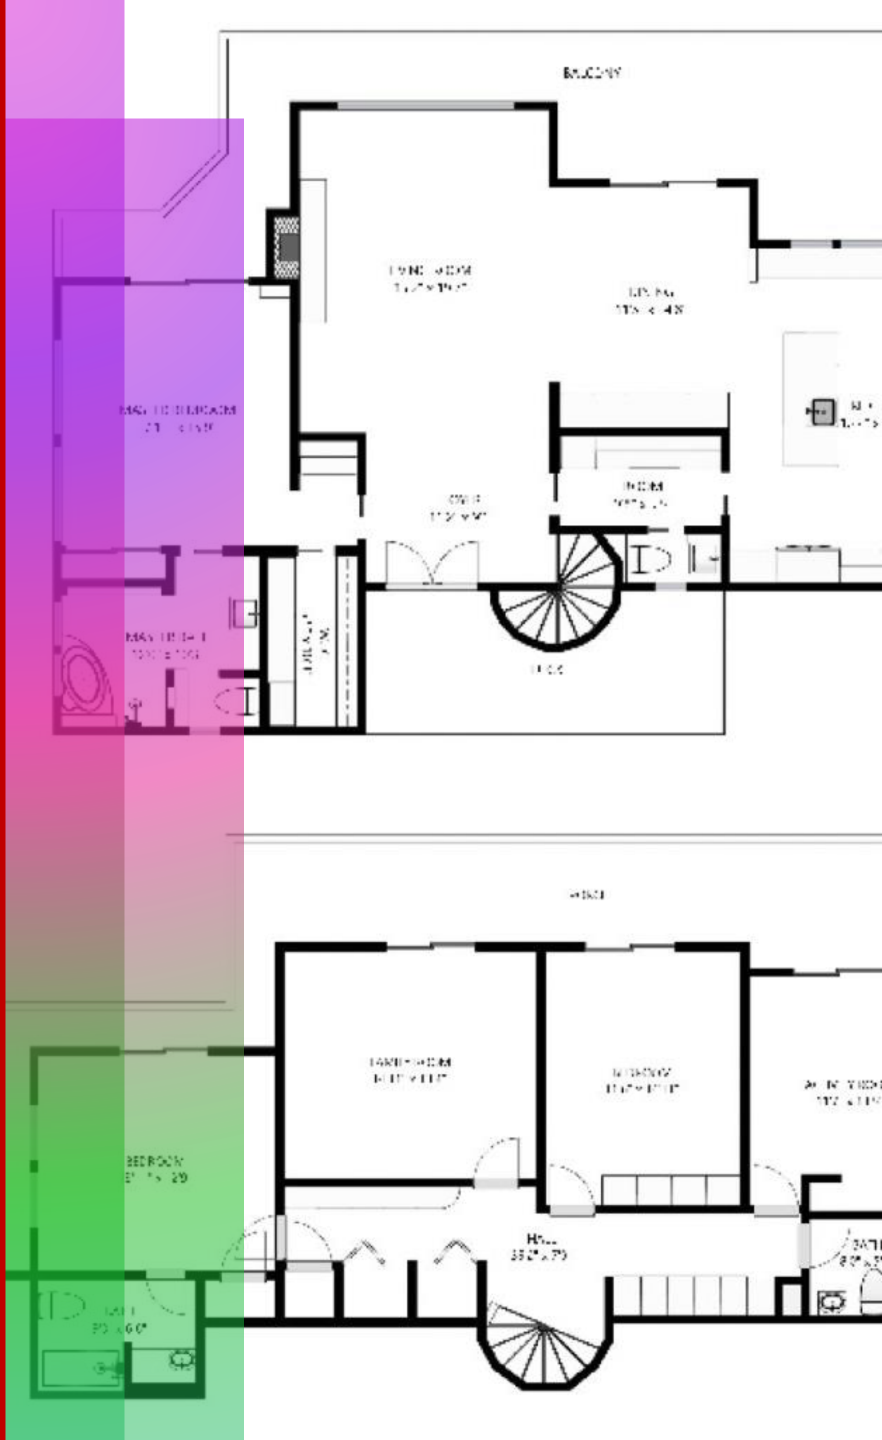

Core property #3: Multimodal Perceptual Symbols

- The process of schematic symbol formation works the same for **any** modality, and the symbols stored in the corresponding area of the brain (e.g. visual areas, motor areas, etc.):
  - *Auditory*: symbols for commonly heard sounds, such as speech
  - *Touch*: symbols for temperature, texture
  - *Proprioception*: symbols for limb movements and positions
  - Studies show when people conceptualize animals, the visual part of the brain is active. When people conceptualize using tools, the motor part of the brain is active, etc.

- **Introspection**: sensorimotor perception is clearly important, but so is introspective perception. Three types of important introspective experience can also form perceptual symbols by abstracting away specific events:
  - *Representational states*: symbols representing an object in its sensory absence, or determining a perceived instance belongs to a certain category
  - *Cognitive operations*: symbols for the comparison between 2 symbols
  - *Emotional states*: symbols representing which schematic aspects of a sensory experience associate with certain emotions

- Just like animals are evolutionarily hardwired to notice certain dimensions of vision (e.g. color, shape, etc.), they may also be hardwired for introspection abilities (e.g. visualize, reason, guess, compare)

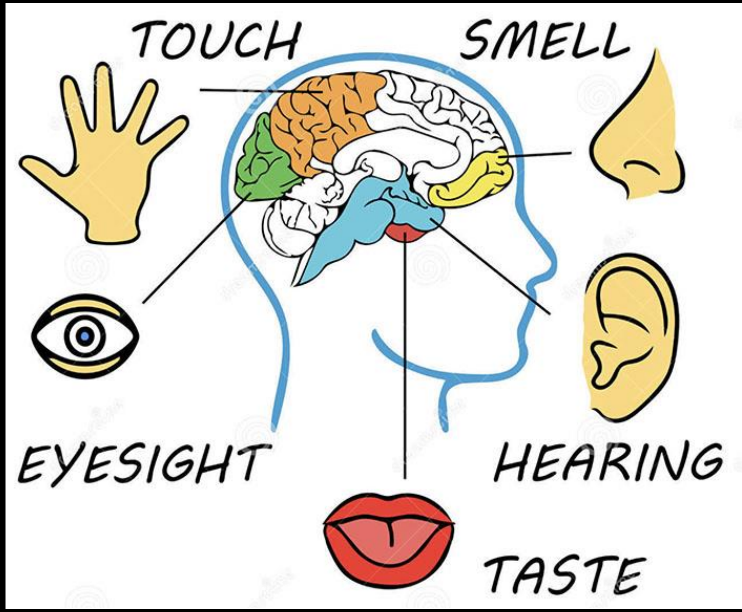

Core property #4: Simulators and Simulations

- Perceptual symbols do not exist independently of each other; they are to be organized into internal “simulators” / “frames” which allow the system to construct specific simulations of objects when perceiving the object, and in the object’s absence.
  - Event sequences can be simulated to make predictions before events occur
  - Feature simulators can be used to perceive features in their absence
  - ***NARS**:* Component **simulators** are equivalent to **concepts** in NARS, especially concepts of conjunctions and implications. **Simulating** is equivalent to reasoning. **Simulations** are specific concept activations (i.e. inference results, events) that occur when using complex simulators together with current perceptual events, concept selection, and reasoning.

- A whole frame is not activated during perception, only certain subsets. The activated simulators produce specific simulations in working memory*. e.g. an olfactory subset of a “car” frame is activated when you smell the inside of a car, or a visual subset of “car” frame is activated when you see a car etc.*

- Over time as the system explores its environment, massive amounts of multi-modal experiences are incorporated into these simulators until the system has a reference for “what it is like” to have certain perceptual experiences *e.g. walking with or petting a dog, etc.*

- Simulators should not be arbitrary, but should be geared to describe meaningful units. What “meaningful” means depends on the system’s goals and intrinsic mechanisms.

Core property #4 continued - Categorization!

- One of the most important functions of perception is assigning a known **category** to a perceived **object**. Many theories assume that an object can be characterized using its amodal structure, by some definition or universal classification algorithm of features, etc.
- However, the idea of simulators contends a different approach: if the simulator for a category can produce a coherent simulation of the perceived object, then the object belongs in that category
- One benefit to categorization using perceptual symbols is that the category symbol's representation / structure is analogous to the perceived object to be categorized and is easily matched. Every successful categorization stores another simulation of the categorized object.
- Categorization is important because it lets the system make **categorical inferences**. Knowledge about the category suggest predictions, history, ways of interaction, etc. Rather than learning about every new entity from scratch, inferences let the perceiver simulate objects mentally.
- Categorization is helped when the system associates the perceived object with ***affordances***, i.e. **actions** associated with the object category
- Because all symbols, perceptual and abstract, reside in the same overall memory structure of the system, the process of **perception** is deeply unified with **cognition** in the system. The difference between perceptual symbols and "cognitive" symbols is not explicit, they are a gradient of generalization and complexity.
- Reasoning activates simulators (concepts) in a top-down manner, but illusions show that bottom-up inference results tend to dominate conflicting top-down inference results. But, when they don't conflict, they tend to complement each other (e.g. by filling in, anticipation, etc.)

Core property #5: Constructing Frames

- A **Frame** is a collection of perceptual symbols used to construct specific simulations of a category (*e.g. a frame for "car" in the image is used to construct simulations of cars*)
- A frame represents spatial information and frame content separately. Volume is implicitly represented by spatial layout (i.e. frame skeleton); the specific content of a subregion is represented as a specialization of the subregion (i.e. fleshed out frame).
- **Image A:** 1st time a car is perceived, the system extracts its overall shape and some components, integrating the subregions into an object-centered reference. (*e.g. if the car window and door handle are given attention, the system extracts symbols for the subregions and the content that specializes the subregions*)
  Image B: 2nd time a car is perceived, the system is reminded of the 1st time. The retrieved symbols / subregions from memory guide the top-down process for extracting *additional* symbols in the same subregions. New subregions might also be extracted.
- After these 2 "car" perceptual experiences, the system organizes the extracted perceptual information into a knowledge structure that resembles a "car" frame. Extracted components from the same perceptual instance are interconnected, to show they are related and instantiate them later together.
- Connections gradually weaken over time, unless they are reinforced by later perceptions, in which case they become stronger. Inhibitory connections form between components competing for the same subregion.
- Over time, the frame contains lots of information; for a given subregion, many specializations exist. Better or "default" specializations develop, either by averaging specializations or choosing the strongest associated specialization with the subregion.

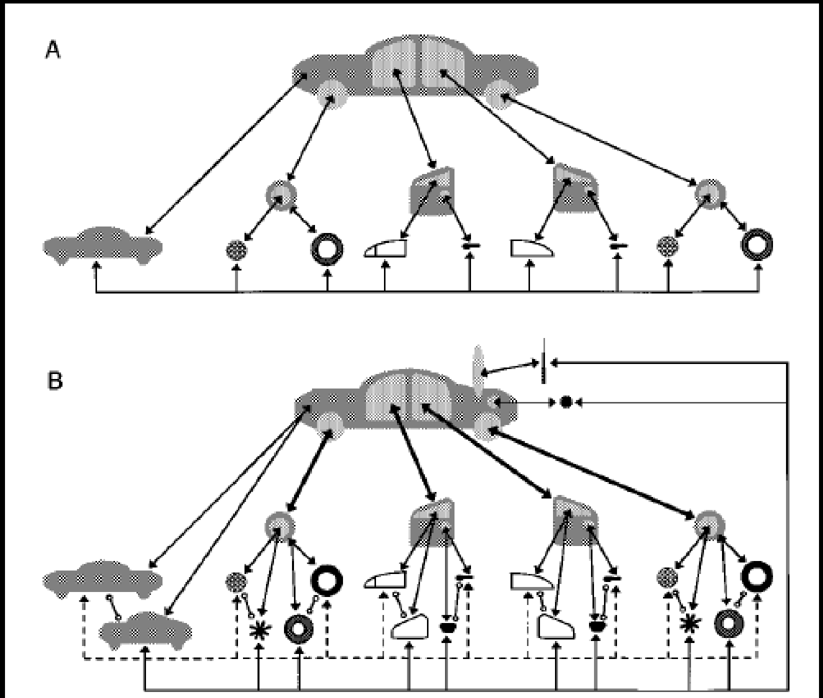

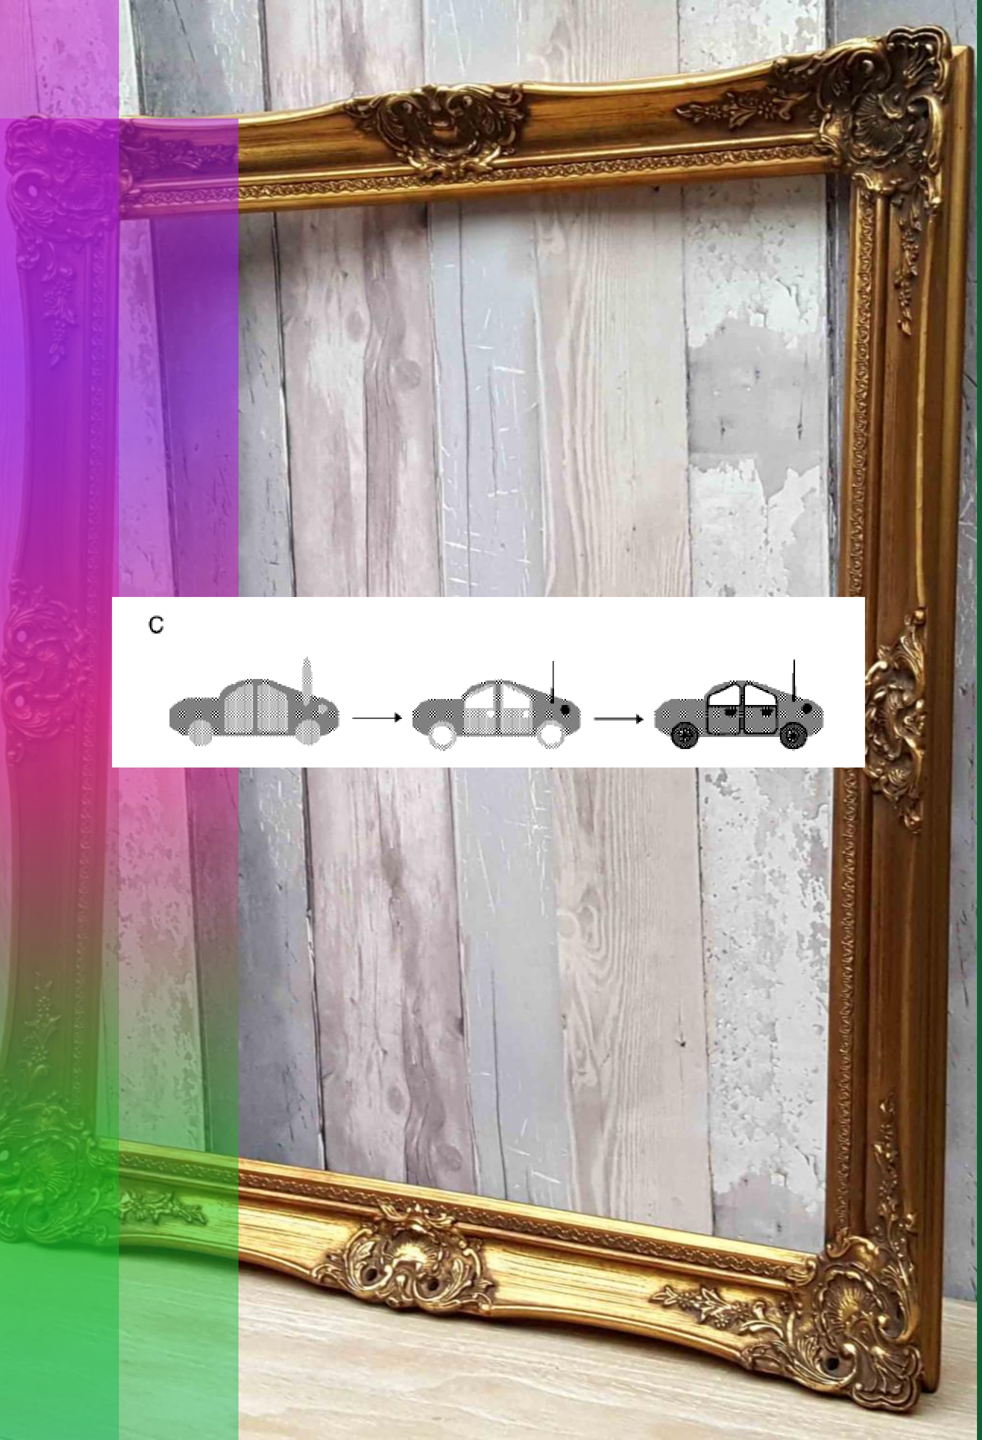

Core property #5 continued: Using Frames

- 4 basic properties of frames:
  1. Predicates (the unspecialized frame / skeleton) *e.g. CAR(Door=x, Window=y..).* Frames hold a schematic structure, which is more generic.
  2. Attribute-value bindings (specializations [value] are binded to subregions [attribute]). New stored specializations can be used as values when simulating subregions, filling in the frame.
  3. Constraints (associative connections formed when constructing the frame, constrains activation spreading). When a specialization is activated, its connections are activated.
  4. Recursion (simulators form within an existing simulator; *e.g. “tire” and “hubcap” simulators may be develop within “wheel” simulator*). The recursion can be arbitrarily deep, bounded by the perceiver’s motivation and perceptual resolution
- The system uses a category frame to create specific simulations. First, the spatial representation is activated. Then, some of the subregions are activated (recursively). Then, the specialization most strongly associated with each subregion is activated.
- The simulation that emerges is the “strongest attractor” (evolved state) in the frame’s state space, unless the simulation activations are restrained by context, in which case the simulation strength is context-dependent.
- While simulating, the system not only retrieves stored frame information, but can also transform it, *e.g. size scaling, deformations, temporal/spatial translation, spatial rotation, deconstructed, etc.* These operations seem intrinsic to cognitive systems and related to perceptual experience

Core property #6: Linguistic Indexing and Control

- Linguistic symbols (in humans) develop alongside their associated perceptual symbols. (e.g. “car” and *car*)
- The linguistic symbol itself is a perceptual symbol, except the perceived event is a spoken / written word. As the system focuses its attention on these events, over time the schematic representations create a simulator for the linguistic word.
- The linguistic symbol also becomes associated in memory with the entities and events to which they refer. They can be linked to entire complex simulators, or simulator’s subregions and specializations. (e.g. “car” is linked to car frame; “trunk” is linked to subregion of car frame.). Over time, they associate with many properties of simulations.
- Once linguistic “word” simulators become associated with “pure” perceptual simulators, they can control simulations. Once the system recognizes a word, it activates associated concept simulators as potential referents.
- As people hear or read language, they construct a simulation that represents their semantic interpretation of the words.
- Conversely, simulations within the cognitive system activate linguistic symbols which become candidates for the system’s own speech / written text. In social communication, these symbols are meant to activate a similar simulation in the listener.
- Using this linkage, linguistic symbols are useful *indexers* for “pure” perceptual symbols.

Derived Properties

4 derived properties (of a fully functional conceptual system)

“

1. Simulators can be combined and used **recursively** to implement **productivity**
2. They can become **bound to perceived individuals** to implement **propositions**
3. Because perceptual symbols are in sensory-motor areas, they implement **embodiment not functionalism**
4. Using **complex simulations** of combined physical and introspective events, perceptual symbol systems represent **abstract concepts**

” – Barsalou, (p. 582)

Derived property 1: Productivity

- **Productivity** is the ability to construct an unlimited number of complex representations from a finite number of symbols using combination and recursion

- Productivity processing:
  A. Select **object simulators** (*e.g. "balloon" simulator, leftmost*)
  B. Select **spatial relation simulators** (*e.g. "above" simulator [same as "below"], leftmost*)
  C. Simulator **combination** (*e.g. "above" simulator produces a specific simulation for "above", "balloon" and "cloud" produce specific simulations to specialize "above"/ "below" simulation. The result is a balloon simulated above a cloud. Such simulations can take many forms, including different balloons, clouds, or above relations.*)
  D. Simulator **recursion**; recursion occurs when specializing the specialization of a schematic region (*e.g. the lower region of "above" is specialized with a simulation of "left-of", which are further specialized with "plane" and "cloud"*).

- Productivity is the reverse process of symbol formation. During symbol formation, the system filters out information to create a schematic; during productivity, the system adds information. (*e.g. during formation, the system creates a schematic for "sphere", filtering out specific color and texture. During productivity, the system uses "sphere" schematic, but fills in color and texture using other simulators*.)

- Productivity is not limited to filling in schematic regions, but can also do transformation, replacements, deletions, etc. (*e.g. imagine a lamp with a white shade; replace lamp shade with brown box in your mind*)

- Productivity can surpass sensory experience, as a form of **imagination**. You can imagine objects with various properties (*e.g. shapes, textures*) and can violate properties of the real-world since there are no physical restrictions required in the simulations. Conceive of imaginary entities and events.

- Although productivity is flexible, there are still constraints; you cannot always effectively use any 2 perceptual symbols together. There are also emergent features (*e.g. imagine a running watermelon; did it have legs?*)

- **Language**: words are combined productively into sentences, just like conceptual structures. A productively constructed sentence corresponds to a productively constructed complex simulation. This correspondence allows humans to convey simulations to each other without using visible objects.

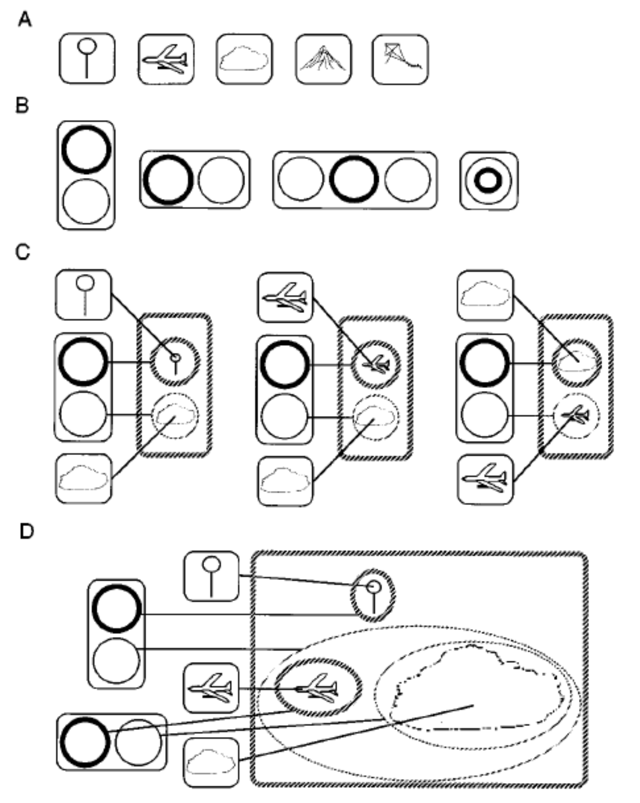

"Productivity" with simulations

Derived property 2: Propositions

- Creating propositions or "true" statements in perception require mapping **perceived instances** to **concepts in memory**. Perceptual symbols implement type-token mappings:
    A. Bind perceived entity with a simulator. As extracted individual visual information produces good simulations from a specific simulator, the simulator becomes more highly active.(*e.g. individual visual information extracted from the left part of the image, the plane, partially activates the "plane" simulator in memory.*). The perception is a fusion of bottom-up (visual information) and top-down (conceptual knowledge). The simulator is the "type", the simulation is the "instance".

- When a perceived instance is bound to a simulator, the simulator can be used to make categorical inferences (i.e. reasoning).

- **False propositions**: When a binding fails, it produces a false proposition.

- **Negative propositions**: To explicitly note a failed binding, use a negation.

- **Many interpretations:** depending on the system's attention to specific details, different propositions may be constructed.

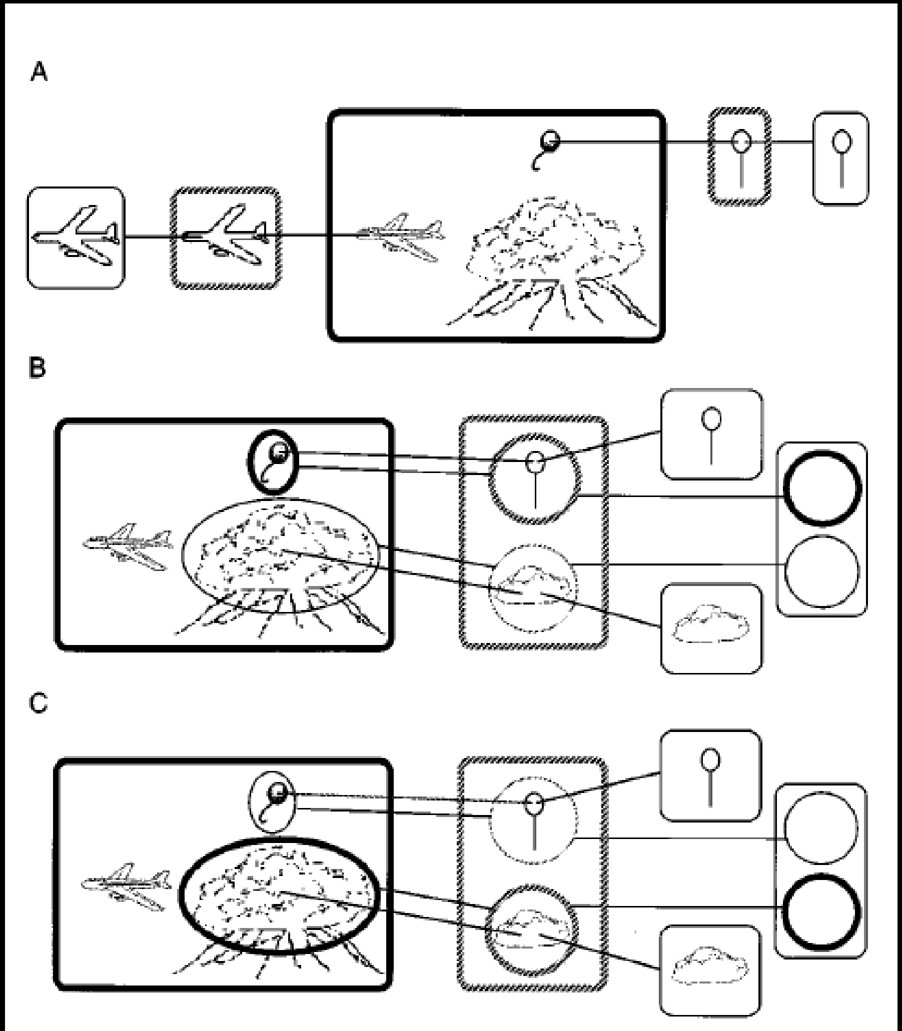

Derived property 3: Variable Embodiment

- A symbol's meaning cannot be disembodied from its cognitive system. A symbol only reflects the system in which it is represented, due to every system's unique physical structure and experience.

- Perceptual symbols are adaptive, since the system forms concepts specific to its environment

- Modal symbols are more expressive, they can be transformed to represent different meaning (rotated, scaled, etc.); amodal symbols cannot

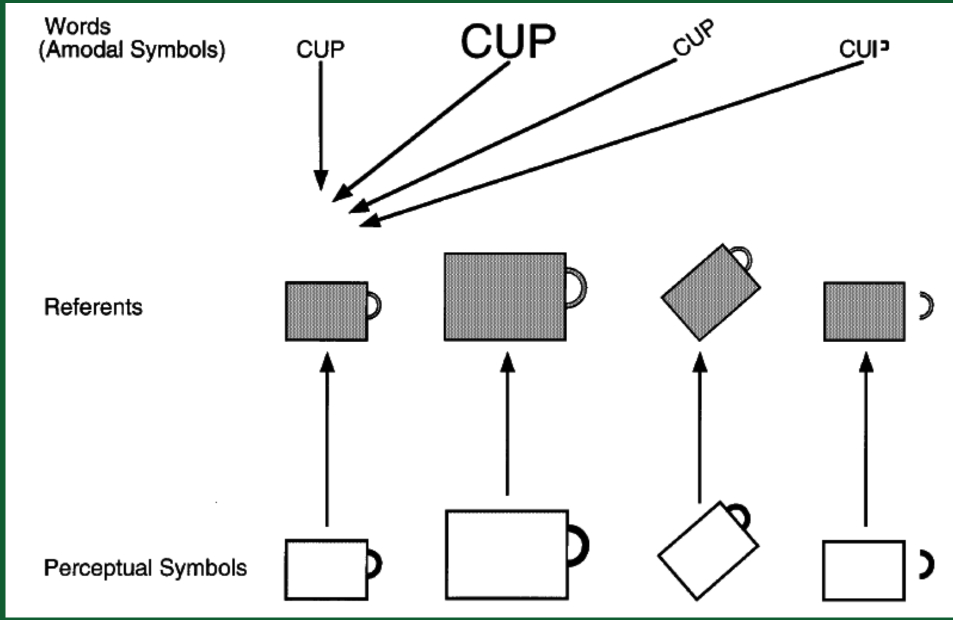

Derived property 4: Abstract Concepts

- Possibly the system can develop abstract concepts by:
  1. **Framing:** An abstract concept is framed against an entire simulated event sequence.
  2. **Selectivity:** The system's attention selects core parts of the abstract concept represented within the event context and associates them.
  3. **Instropection:** Abstract concepts are represented using both sensory and introspective event symbols.

- "Truth" and "Disjunction" (OR) are 2 useful abstract concepts

- Subevent 1: System constructs a perceptual simulation
  Subevent 2: System tries to map the perceptual simulation into the current situation, to decide whether they match

- **Truth:** Truth is established when a simulation and situation are successfully binded, *e.g. the captions in the image 1.* If the system perceives introspective events, over time the system can learn "truth" , what it feels like in general or means to successfully bind reality to a concept. Falsity occurs from a failed binding. Negation is the explicit noticing of a failed binding.

- **Disjunction:** Disjunction can be established when the system tries to remember a specific simulation, but during reconstructive mechanisms finds multiple matches for a subregion, which can be swapped in the simulation. The alternating simulations ground "or" (disjunction)

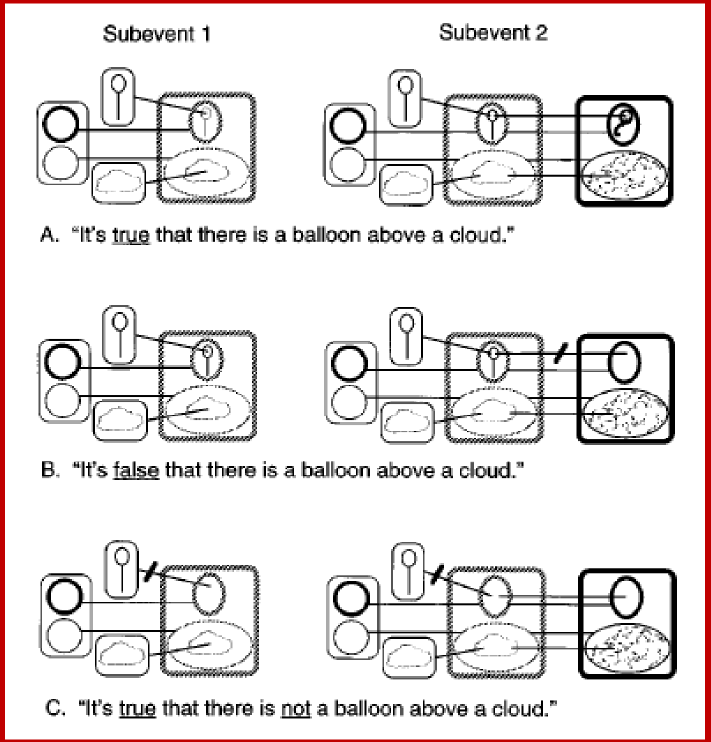

Image 1: truth

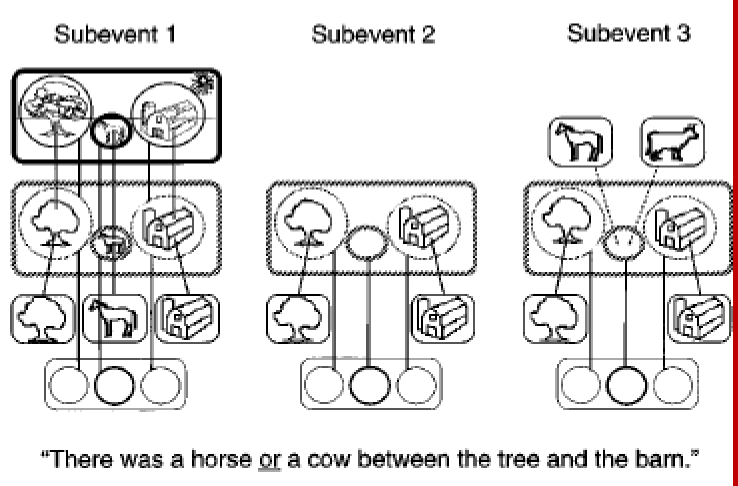

Image 2: disjunction

Disjunction for ad hoc categories

- The system forms an ad hoc category when a schematic region is specialized using disjunction.
- For example, in the image, if the system is standing in a room trying to construct a successful simulation for changing a light bulb, it might fill in the region of “thing to stand on” using a disjunction of objects in the room (chair, stool, or table, etc). It can use this same ad hoc disjunctive category if it wants to simulate other tasks about the room such as propping open the door.

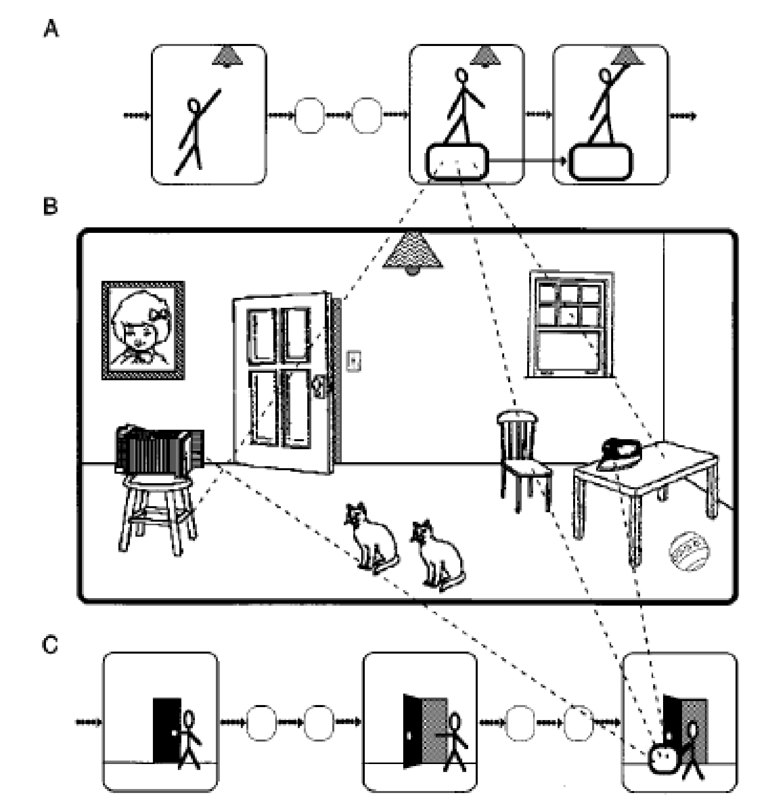

What can I stand on to change a light bulb? To prop open the door?

Implications and Conclusion

Implications

- Perceptual symbols are not merely recording systems, but use attention mechanisms to extract components of perceptual experience to establish concepts
- Fairly generic, well-grounded perceptual concepts represent “types” that lets the system produce categorical inferences
- Viewing cognition as grounded in perceptual simulations provides a different way of thinking about human memory, language, and thought.
- From the perspective of perceptual symbol systems, there is no dissociation between perception and overall cognition. They share many of the same mechanisms.
- Categorization is integral to perception, matching an individual to the best-fitting simulator / type. The process runs schematic simulations to find the best fit for the individual’s features, adding inferred features.
- Since concepts are not amodal, spatiotemporal relations and transformations are central to concepts. This is not a problem since the concepts are stored modally so their modal relations and transformation follow naturally.
- People don’t use background theories or knowledge to frame their perceptual concepts, it’s the other way around, they use simulations of perceptual concepts to ground and frame their theories.
- **Attention** is integral to perceptual symbol formation, since it determines what is extracted and transferred to long-term memory. Not everything can be stored like in a recording system.
- Working memory is conceived as having a central processor and several modality-specific buffers. It is the system that runs perceptual simulations. The visual short-term buffer runs the current visual simulation; the motor short-term buffer simulates movements just performed or to be performed. Long-term memory contains simulators; working memory contains specific events or simulations
- Long term simulators should compile more and more information from many experienced simulations; ideally it should improve they stability and usefulness over time.

Conclusion

- Cognition and perception use many of the same mechanisms and are performed with the same concepts.
- Concepts are inherently perceptual and stored in the same modality (same signal type) that produced them
- NARS fits naturally as a perceptual symbol system, since it:
  - uses symbols
  - can perform inferences
  - inherently represents taxonomic relationships
  - concepts and implications as simulators in long-term memory
  - events as specific simulations in modality-specific working memories
  - attentional control

The end! ☺