

Multi-modal Representations
as the Basis of
Cognitive Architecture

B. Chandrasekaran

Laboratory for AI Research

The Ohio State University

My Goal for This Talk

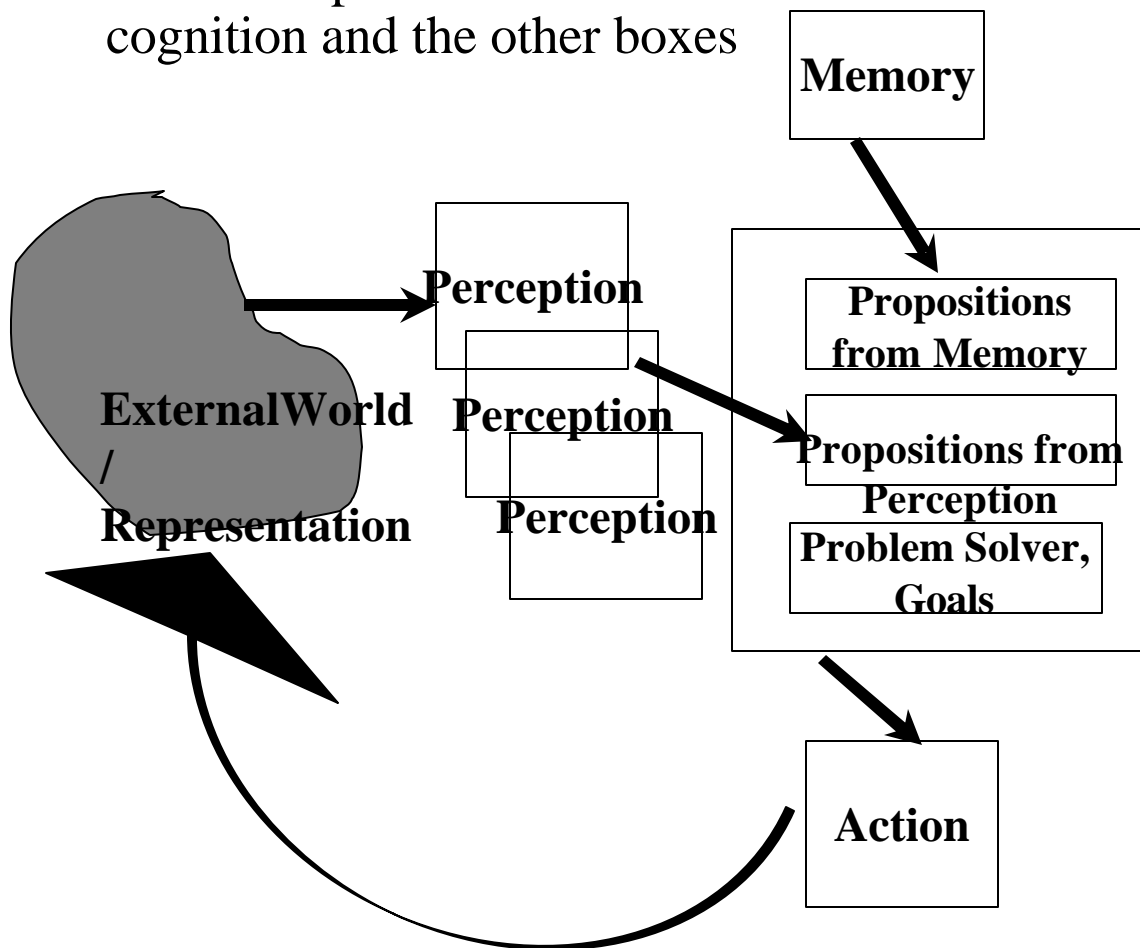
- Much of what I say won't be completely new to the ACT-R community –
 - hints of some of these ideas in ACT-R (e.g., “imaginal buffer”)
- But I think I am proposing a somewhat more radical view about the roles of perception/motor systems than standard
 - but one that is quite compatible with the ACT-R perspective.
- Want feedback from the ACT-R community
 - about how much of what is needed is already in place, and
 - how to incorporate more of the proposal in ACT-R, assuming that the proposals make sense.
- A “Big Picture” talk with all its problems
 - AI/CogSci orientation.

Thinking: The Canonical View

- Thinking is a process of manipulating a representation of the world of interest (typically in pursuit of a goal)
- The representation is something like NL sentences, LOT hypothesis
- Thoughts are propositions (or attitudes to them, such as believe P, desire P, etc.)
 - From “thoughts have propositional content,” to “thoughts are representations of propositions.”
- This general stance towards cognition is not specific to “logic” approaches in AI. Approaches based on frames, scripts, rules are all in this sense “symbolic AI.”
- For our purposes, connectionism and related ideas are alternative ways of representing the same type of information, and thus orthogonal to the issues raised here.

Interaction with the External World in the “Standard” View

- Interaction with the world takes place by taking in knowledge of the world from perception in the form of propositions and generating propositions about actions to take that are then executed by motor systems.
 - Buffers at points of contact bet. Central cognition and the other boxes



Steps in the Argument

- A view of Reasoning w. external Diagrams
 - Representations span the agent and the environment, and jointly each state is *bi-modal*.
- Make diagrams internal.
 - Problem solving with mental images of diagrams. Now internal state is bi-modal
- Generalize to multi-modal state:
 - External: [External representations in general: e.g., designer with a clay model, chemist with a 3-d molecular model, a composer with a musical instrument in hand]
 - Internal: No reason to restrict the notion of an image to the visual
- Outline a multi-modal cognitive state and elements of the associated engine.
 - Multi-modal problem-spaces
- Discuss why this kind of architecture is useful for any agent, natural or artificial.
- Issues in implementing the multi-modal state:
 - Specifically, what kind of rep. frameworks support representing images in such a way that they are both image-like and symbol-like?

Use of External Perceptual Representations in Problem Solving

- In “diagrammatic reasoning,” a problem solver uses a diagram as part of a PS episode.
 - Proposition extraction.
 - Proposition extraction is direct. Often the same proposition might take a chain of inferences if done within a propositional representation.
 - Deliberative reasoning with visually extracted & other propositions and rules of inference.
 - This is all consistent with the “Standard View,” in which perception supplies propositions.
 - Important matters which we ignore for now: Proposition projection, “simulation” of motion and changes in positions result in new external representations from which additional propositions may be extracted.

Internal Visual Representations (Images)

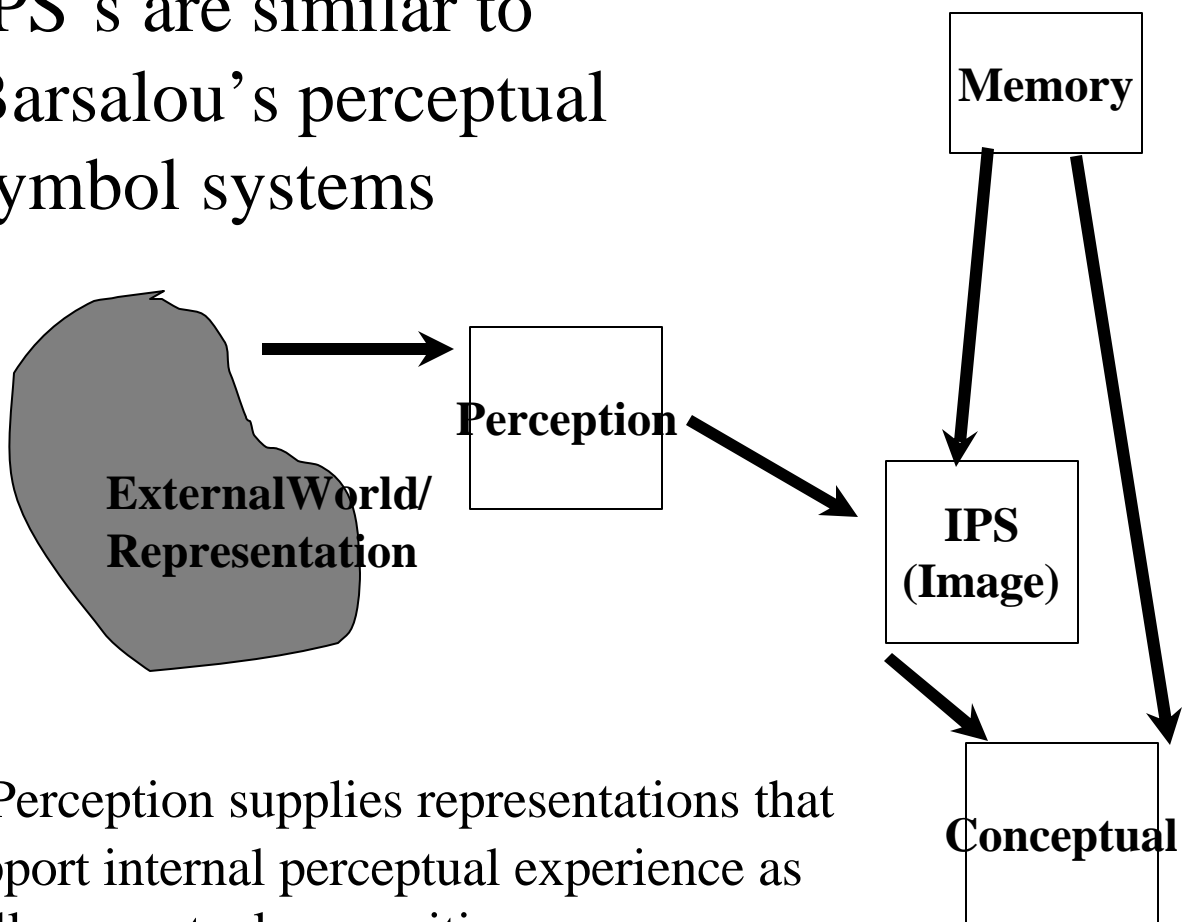
- Controversy about mental images notwithstanding, in many PS situations, mental images of diagrams are used in problem solving, playing essentially the same *functional role* as the external diagrams do.
 - Examples.
- In a real sense, the internal representation is bi-modal.
- Question: what kind of internal representation frameworks explain the functional role as diagrams without being diagrams to be perceived?
 - I have proposed a solution, which is beyond our scope for today.

What is really in a Cognitive State?

- In contrast to the “standard view”
- Phenomenologically, if we analyze the content of thought, we are aware not only of elements that have a *propositional* content, but also elements with *perceptual* content.
 - We can “see” a child swinging in the yard, we can “hear” tunes and have difficulty getting rid of them. We can decide if we can go through a passage by imagining if we can contort our bodies to the way required.
 - Even in communication, we may use, in addition to language, pictures, 3-d models, gestures, music, and so on. Thus, there is no good reason to model inner thinking purely on language and its structure.

Perception Supports Both an Inner Perceptual Experience and Conceptual Knowledge

IPS's are similar to Barsalou's perceptual symbol systems



- Perception supplies representations that support internal perceptual experience as well conceptual propositions
- Inner perceptual experience can also be created by representations from memory, and of course memory can also supply conceptual propositions as well.

Images are Not Just For the Visual Modality

- Much of the work (debates) in CogSci on images has been in the visual domain, but, as the example of tunes suggests, the phenomenology of images is not restricted to the visual.
 - “I can almost taste the food.” In addition to all the perceptual modality, one can have kinesthetic images as well
 - A sense of the contortions of the body when imagining going through a narrow passage. On looking at a design diagram: “The mouse buttons feel like they are too far apart for comfort.”

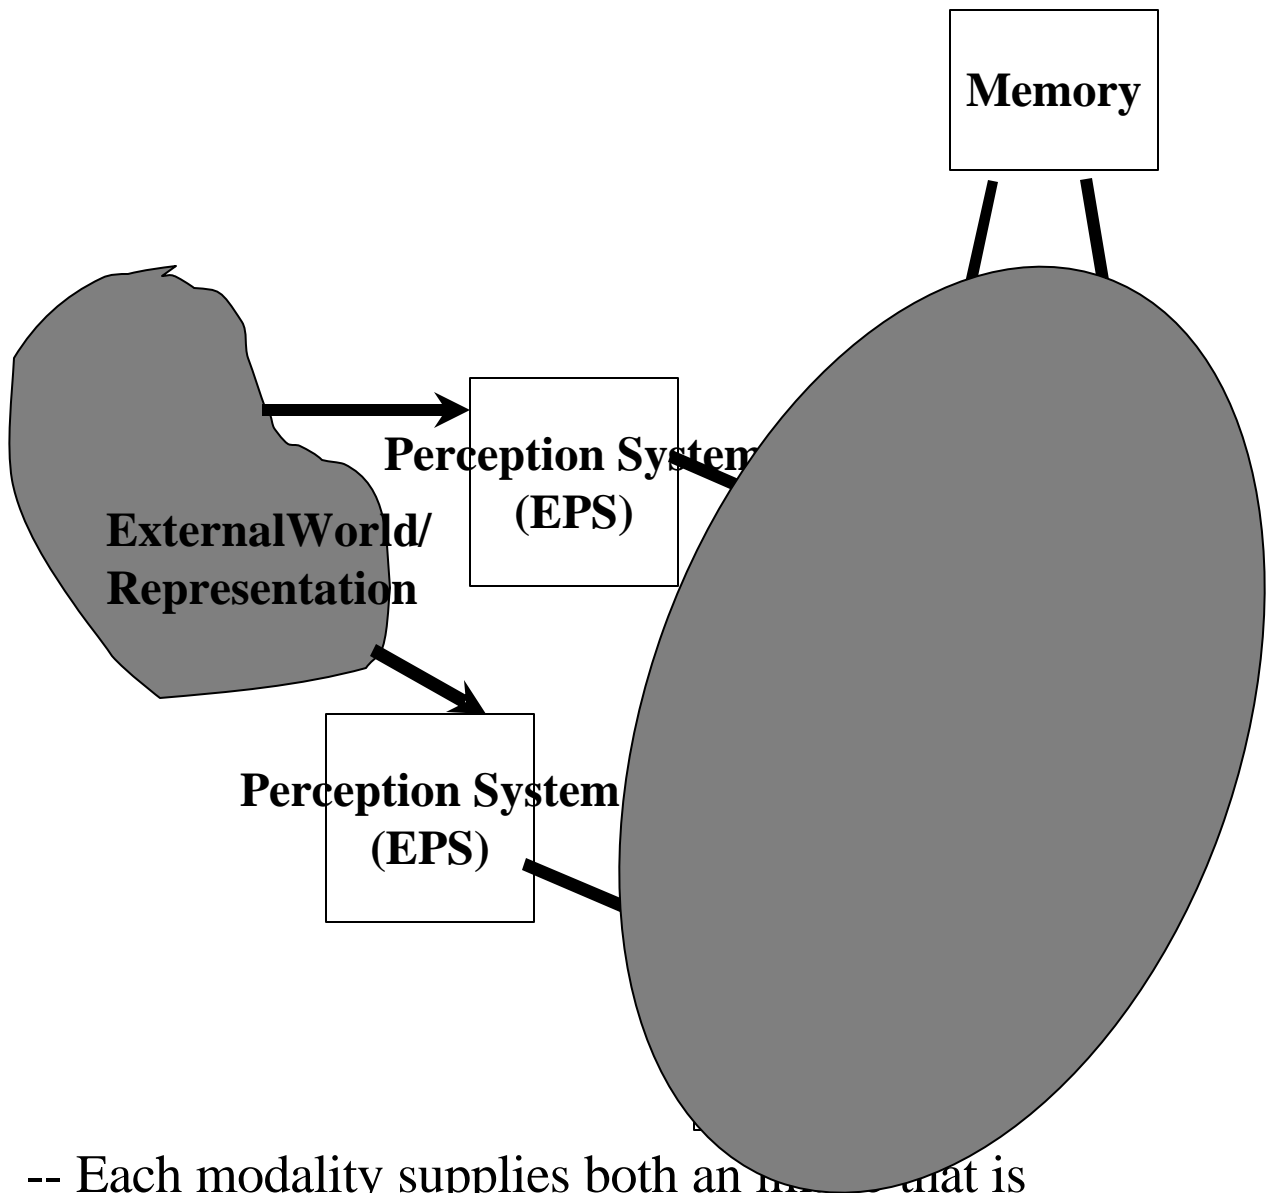
Images are not identical to their propositional content

- Two dimensions to images:
 - The first is the experiential dimension
 - Listening to music or looking at a sunset is not identical to being given propositions about them.
 - The second is the informational dimension
 - An image corresponds potentially to an infinite number of propositions.
 - Inferential or information extraction operations are different for images.

Internal Multimodal Representations: A Proposed Functional Architecture

- Each state of thinking is *potentially* multimodal.
 - Potentially, because not all instances have all modalities.
- Perceptual modalities (PM) includes kinesthetic modality for this discussion

More than One Perceptual Modality



-- Each modality supplies both an image that is experienced as well conceptual predicates

-- Remember that kinesthetic modality is one of the perceptual modalities in the above.


Functional Architecture: Multimodality of Cognitive State

- The effects of acts of perception and of imagining are very similar, except that EPS can maintain the richness of the images in IPS without strain on memory.
- Awareness (cognitive state) is multimodal, as a rule.
 - Its components are the various IPSs and the *conceptual* modality.
 - For the sake of uniformity, we refer to representations in IPS's as well as the conceptual modality as images. And when we say IPS, we will include the conceptual mode as well.

Functional Architecture: Memory is multimodal too

- Agent's memory is also multimodal, paralleling the organization of the cognitive state
 - Views, postures, tunes, concepts, episodes that have all these.
 - Elements in one memory mode are associated in various ways with elements in other modes
 - Concept of apple in memory may be associated with the memory of its shape and color in the visual modality, the act of biting into it in the kinesthetic memory and so on

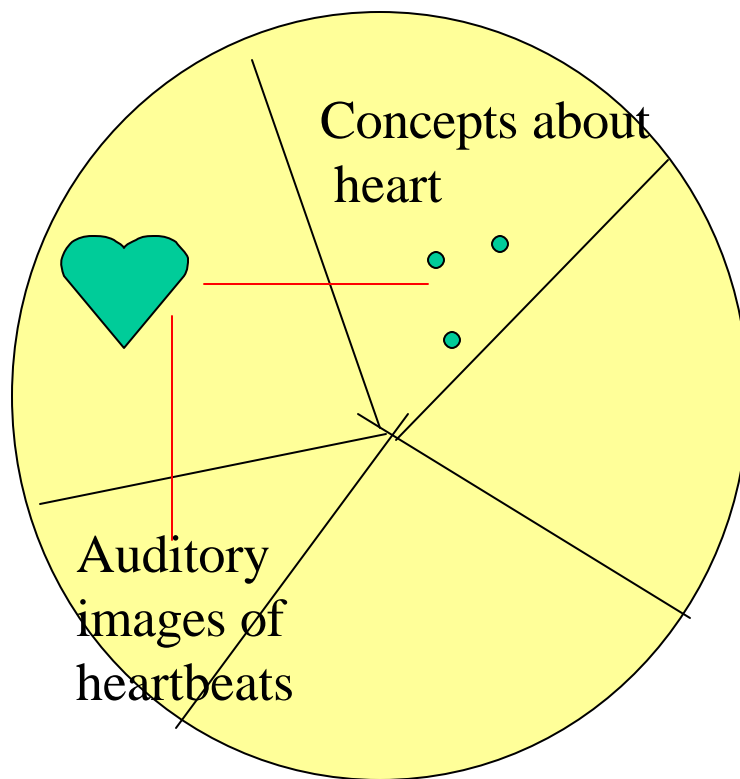
Memory is Multimodal Too



Representations May Evoke Associated Representations


- A representation in any IPS (I.E, in the cognitive state) may evoke images in other IPSs
 - The evoked images are those associated in memory
 - E.G., The visual image of a heart may evoke conceptual information about its role in life and health issues, the auditory image of heartbeats,....,....
 - These multimodal evocations occur whether or not the image in an IPS came from EPS or from memory

Intermodal Evocations

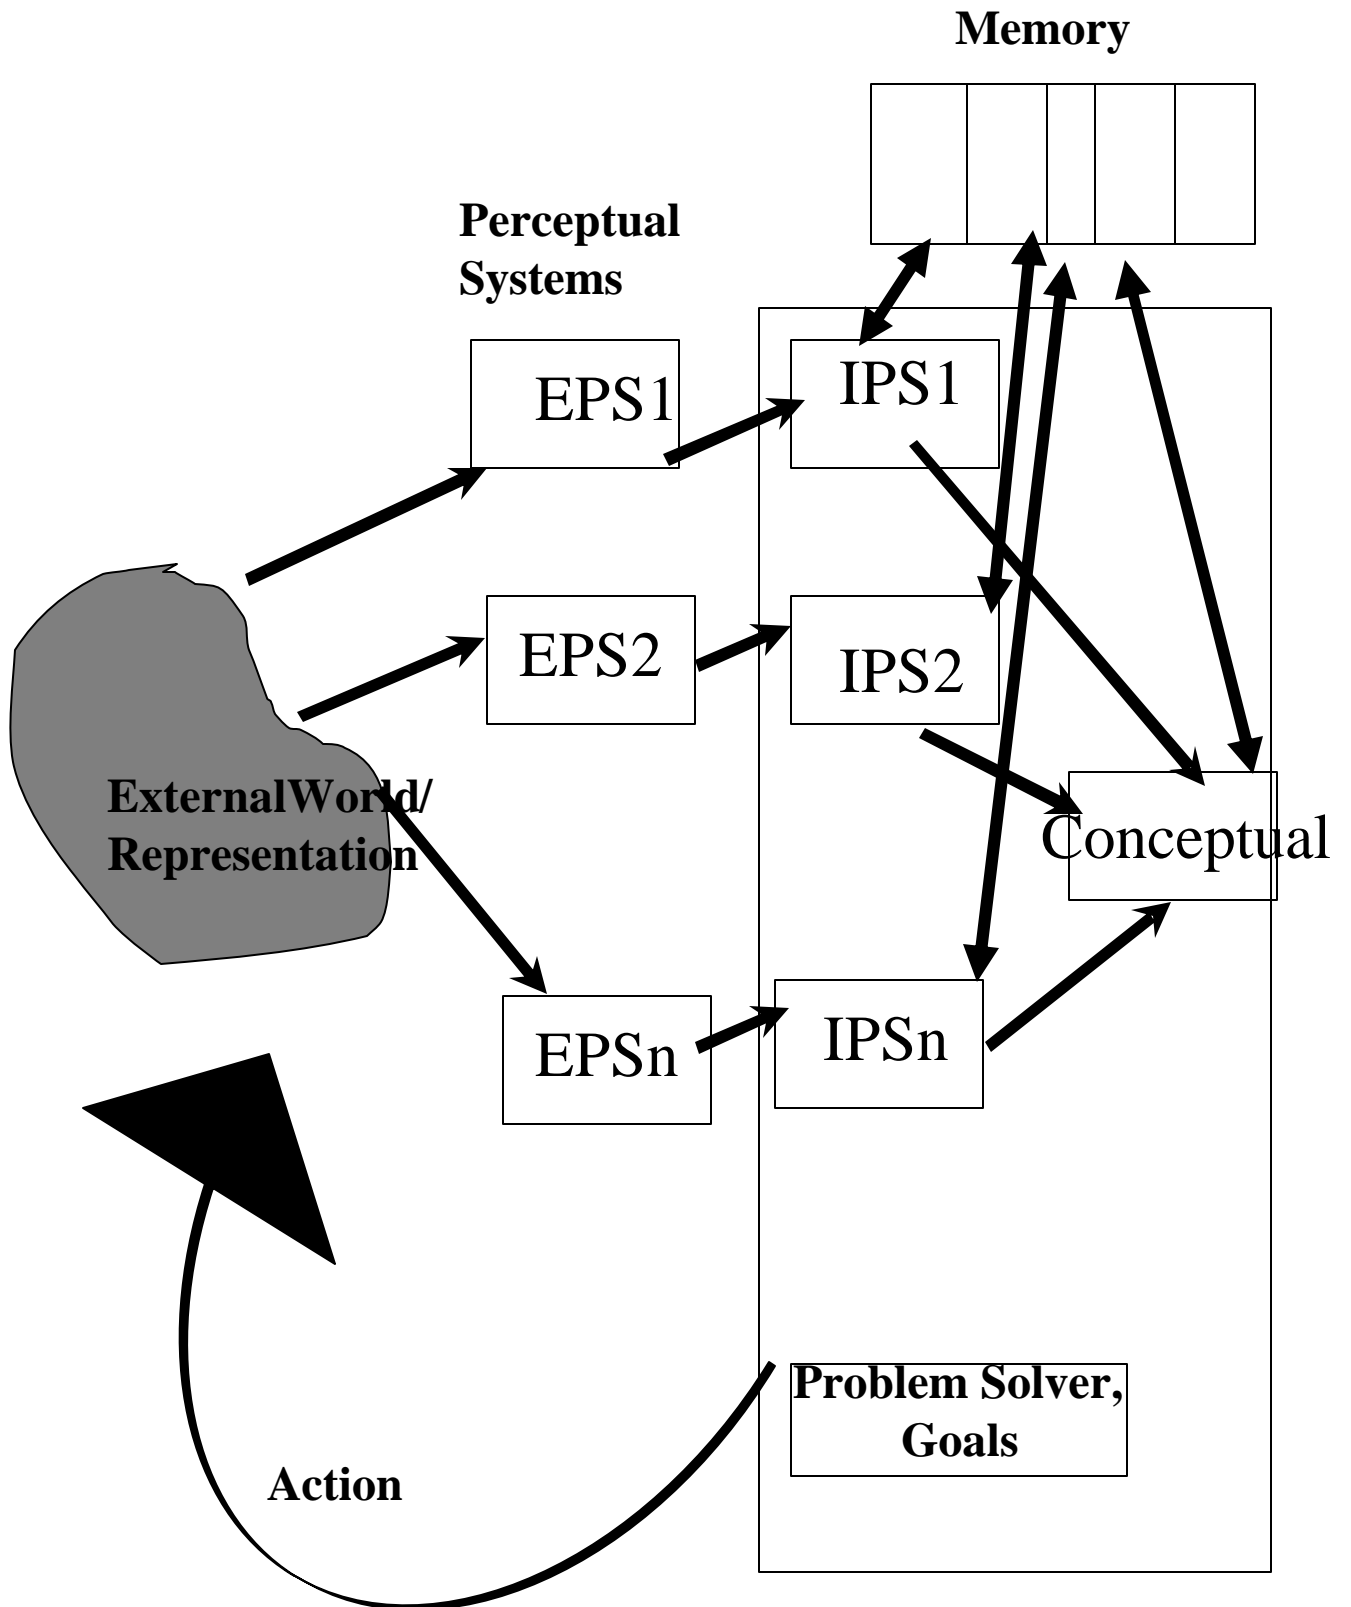


Intermodal Evocations

A representation in modality may evoke associated representations in memory of other modalities, that in turn would evoke images in the corresponding IPSs.



Cognitive state of agent is multimodal



Caveat on the Narrow View of Modularity

- For simplicity of exposition, I am adopting an oversimplified view of modularity.
 - Goal-Perception-Action schemas provide a somewhat different sets of modules than in my schematic diagrams

The Architecture is a Functional One

- I am only appealing to the *capabilities* or *functionalities* of imagining and perceiving and reasoning.
 - Alternative implementations of these capabilities in different mechanisms (Turing symbolic, connectionist, dynamical systems, whatever) possible.

Why This is Not Just Another Instance of a “Standard” Propositional Representation

- The symbols and relations in IPS do of course refer to objects in some world, but relations between them is not abstracted into a relational symbol.
 - Each image potentially corresponds to an infinity of propositions.
 - The set of operators for each modality is an “analog” of the corresponding modality.

What do agents do with IPS representations?

- Each of the IPS components supports its own characteristic form of inference
 - Perceptual modalities support perceptual proposition extractions, while the conceptual modality supports “reasoning”
- Intermodal evocations provide a powerful way to bring to bear distal knowledge in all modalities to bear on the current situation
 - Shape & color of apple (visual) --> taste of apple (taste) --> appropriateness for the pie recipe (conceptual) --> decision to buy apple (conceptual)

IPS Representations Drive Problem Solving

- Problem solving is a process in which the agent's cognitive state changes as a function of the contents of the current cognitive state, the state of the external world, and the PS goals
- Changes in the external world (or attention) cause EPS to deliver new percepts and relations to IPS and the conceptual component
 - Proposition projection is one cause of the change in the external world
- Changes in one IPS representation may evoke associated images in other IPSs
- New propositions may be extracted in those IPSs, including in the conceptual mode
- Conceptual inferences may include action items that change the external world
- Goals determine control of which inference options are pursued

Advantages to the Agent

- A wide-variety of modality-specific information extraction operators are directly available -- obviating the need for complex inferences from propositional abstractions.
- Each image corresponds literally to an infinity of propositions. Thus, is experience is stored closer to perception -- how it was experienced -- propositions can be extracted as appropriate for the task at hand.
- Builds on top of perceptual machinery already needed for other purposes.
- Continuity with animal intelligence in general.

What does this mean for AI

- Especially important in the emerging integrated systems approach to AI -- robotic-based intelligence systems, with perception, action and reasoning rolled into one.
- As the robot experiences the world, its memory and reasoning exploit the structure of perception.
- Knowledge retrieval vs knowledge construction

Relation to Barsalou's Perceptual Symbol Systems

- The approach being proposed has many points of contact with Barsalou's proposal, though my motivating problem is robot-building, while he motivates his discussion quite a bit by neural representation issues.
- My proposal – mentioned earlier, but not presented here – for IPS is about what a composable perceptual symbol system might look like.

Concluding Remarks

- Experiencing, problem solving, reasoning, takes place in the context of an external world that we perceive, act on, imagine and reason about in multiple modalities
 - Conceptual modality is always present
 - Different tasks emphasize different perceptual modalities
 - Music composition vs mechanical design
- In our framework, the internal representational life of the agent is multimodal, with representations in one mode evoking allied representations in other modes, and each mode making inferential contributions for which it is best suited. Mental images come in all modalities.

Concluding Remarks (cont.)

- Treats conceptual component as just another component with equal status with inner perceptual and kinesthetic components. In one way of thinking, having concepts is imaging the world in the conceptual world, just as having images is imaging the world in the perceptual mode.
- Logical rules of inference are just a very small part of the information extraction operators in the conceptual part.

ACT-R and the Multimodal proposal

- ?