Markman, A. and E. Dietrich (2000). Extending the classical view of representation. *Trends in Cognitive Science* v. 4, n. 12, pp. 470-475.

Something old, Something new: Extending the classical view of representation

Arthur B. Markman
Department of Psychology
University of Texas

Eric Dietrich
Philosophy Department
Binghamton University

Summary

Representation is a central part of models in cognitive science, but recently this idea has come under attack. Researchers advocating perceptual symbol systems, situated action, embodied cognition, and dynamical systems have argued against central assumptions of the classical representational approach to mind. We review the core assumptions of the dominant view of representation and the four suggested alternatives. We argue that representation should remain a core part of cognitive science, but that the insights from these alternative approaches must be incorporated into models of cognitive processing.

Introduction

There is revolution in the air in cognitive science. Since the late 1950's, models of cognition have been dominated by representational approaches. Inspired and guided by the data structures of computer programs, these models posit some kind of internal mechanism for storing and manipulating data as well as processes that act on those representations to carry out intelligent behaviors [1, 2, 3].

While the field of cognitive science has made great strides in understanding brain, mind, and behavior, the early predictions that we would soon have autonomous robots and intelligent computers on our desktops have not yet come to pass. A number of researchers from a variety of perspectives have suggested that the standard representational assumptions made by cognitive models are to blame for the lack of progress in understanding cognitive processing. These researchers suggest remedies ranging from additional information that should be included in representations to replacement of the dominant paradigm with an alternative.

In this paper, we begin by sketching the classical view of representation that is widely employed in cognitive models in psychology and artificial intelligence. Then, we explore four new approaches to cognitive modeling: perceptual symbol systems, situated action, embodied cognition, and dynamical systems. Each of these approaches has been put forward as a successor to the classical view of representation. We argue that representation must remain at the core of Cognitive Science, but that each of the four alternative approacheshave something important to offer. We end by discussing ways to reconcile the classical view with these suggestions.

The classical view of representation

Cognitive science uses many kinds of representations, and it would be impossible to provide a complete summary of all of them [4]. On the classical view, all approaches to representation share five key assumptions: (1) representations are mediating states of an intelligent system that carry information, (2) cognitive systems require some enduring representations, (3) cognitive systems have some symbols in them, (4) some representations are tied to particular perceptual systems but others are amodal, and (5) many cognitive functions can be modeled without regard to the particular sensor and effector systems of the cognitive agent. (In this paper, we will use the term "cognitive agent" to include organisms as well as intelligent machines.)

The first assumption is that there are mediating states that are internal to the cognitive system [5]. In order for something to qualify as a mediating state, four conditions must be satisfied [3]. First, there must be some representing world. The representing world consists of the elements that serve as the representations. Second, there must be some represented world. The represented world is the information (either within the system or external to it) that is being represented. Third, there is a consistent set of representing relations that determine how elements in the representing world come to stand for elements in the represented world. Finally, there are processes that use the information in the representing world.

As an example, Tversky's [6] contrast model of similarity assumed that objects are represented by sets of features. Each feature is a symbol that stands for a particular property of the object (e.g., blue or large). Pairs of sets representing two objects are compared by finding the intersection of the sets. The features in the intersection are the commonalities of the pair, and the features that are not in the intersection are the differences of the pair. In this example, the representing world is the set of features. The represented world consists of objects external to the model. The features are assumed to bear a consistent relationship to properties in the represented world that can be used to describe the objects. Finally, the information in the representations is processed by the set operations that determine the commonalities and differences of a pair.

The second assumption is that some representations are enduring states of the system. In particular, agents must use their experience as a guide. Thus, they have internal states that endure longer than the states in the represented world that gave rise to them. To continue the example of the contrast model, an object representation can contain a particular feature (e.g., blue) regardless of whether that property is currently accessible in the environment. It is assumed that all objects that have a particular feature as part of their representation share some property in common.

The third assumption is that some of the representations are symbols. Symbols have two central qualities. First, their relationship to the represented world is arbitrary [7]. Second, symbols are discrete packets of information. Symbols are necessary for referring to specific values or properties in the represented world, and mirror the observation that languages consist of words that permit two people to fix common reference. In the contrast model, features are symbols in the representing world.

The fourth assumption is that representational elements may exist at a variety of levels of abstraction. Some representations may correspond directly to aspects of perceptual experience. Other representations may be more interpreted, and may refer to abstract concepts like truth or justice, which are quite removed from perceptual experience. In the contrast model, there is no necessary connection between the features that describe an object and any perceptual information. Indeed, Tversky [6] explicitly avoids specifying processes that create the feature sets.

The final assumption is that some cognitive models need not be concerned with perceptual and motor representations. On this view, there are representations in the cognitive system that are sheltered from the particular body of the agent. It is assumed that such processes can be understood without considering the perceptual and effector systems of the agent. To complete our example, the contrast model does not make any assumptions about the nature of the perceptual or motor systems of cognitive agents. Instead, the model is expected to be able to operate in any agent.

Alternatives to the classical view

The four alternative approaches to representation have all taken umbrage with at least one of the core assumptions of the classical view. Each of the new approaches is motivated by some insight or example that suggests a modification of the classical view.

In each of the following sections, we describe one of the approaches, starting with its motivating insights and discussing the core assumptions it calls into question. We then argue that none of the new alternatives can replace the classical view. We conclude by discussing possible extensions to the classical view suggested by the alternatives.

1. Perceptual symbol systems

Cognitive processing is exceptionally flexible. People are able to recognize when a new situation is like one they have experienced before, but they are also quite good at handling deviations from normal situations. The classical approach to representation has assumed that flexibility requires abstraction. By abstracting away from the perceptual details of specific situations, the commonalities across situations can be preserved. Thus, the classical approach typically assumes that there are abstract *amodal* representations that play an important role in cognitive processing.

As it turns out, however, amodal representations are not as flexible as they were initially assumed to be. For example, Schank and his colleagues suggested ways to represent abstract scripts and schemas that would enable an agent to comprehend new events [8, 9]. These systems had difficulty dealing with the many potential variations of even a simple event. Indeed, in later work, Schank [9] ultimately had to posit both abstract and specific representations in order to account for human-like flexibility in dealing with the many variations on simple events like going to a restaurant.

Symbolic models have also had difficulty accounting for subtle differences in the way a property manifests itself in different items. For example, it is often noted in introductory psychology textbooks that people know that the red of a fire engine is different from the red of hair, even though the same color term is used to describe both of them. As another example, the same spatial preposition may describe a variety of subtly different situations [10, 11]. For example, the English preposition "in" normally means that one object is contained inside another, but an apple can be said to be "in" a bowl, even when it is stacked on other apples such that it rises above the top lip of the bowl. People do not notice this ambiguity in prepositions, but it can be difficult to account for it using traditional symbolic models [12].

Current research has suggested that flexibility in cognitive processing may arise from the storage and use of specific episodes in memory and their perceptual content. Barsalou [13] proposes that the perceptual system may be used to simulate objects and events. For example, to represent an apple in a bowl, a perceptual simulation of an apple on top of other apples would be formed as a part of comprehending this situation. The connection between perception and language in this case would be accomplished using principles derived from cognitive grammar [see box A].

Theoretical arguments and experimental evidence have been marshaled to suggest that cognitive science should eschew amodal representations. In the domain of categorization, Schyns, Goldstone, and Thibaut [14] point out that most amodal theories assume a fixed set of features. The goal of a classification model in this context is to use these features to predict the category to which the instance belongs. Some models calculate similarity to a prototype (i.e., an average member of the category, [15, 16]), or to various known exemplars [17, 18, 19]. Others form rules to describe the categories [20]. In each of these models, the set of features that can be used to represent objects is fixed.

Schyns et al. [14] point out that people often learn new features, even new perceptual features. In order understand how categories are formed, they argue, it is necessary to further understand these processes of perceptual feature creation. In one study, Schyns and Rodet [21] taught people perceptual categories each consisting of many unfamiliar shapes. What people considered to be the basic perceptual components of the categories depended on the order in which they were exposed to the categories. For example, they might see some items that contained the complex feature XY from Figure 1 as one of its components. If this was the first category they learned, they generally treated XY as a whole unit. In contrast, if they first saw some categories with component X in them, when they later saw XY, they were likely to divide this shape up into X and Y (see Figure 1). Thus, the set of perceptual features used to construct the categories is learned. Similar demonstrations have been performed with real materials. For example, Lesgold et al. [22]suggested that the features in X-ray films used by expert radiologists to make a diagnosis are significantly different than those used by novices.

Barsalou [13] takes findings like these to suggest that all representations are closely tied to perceptual modalities. He calls this approach *perceptual symbol systems*, and suggests that concept learning and use involves the creation of perceptual simulations. As evidence for this assertion, Wu and Barsalou [23] asked people to list properties of a variety of concepts. They found that properties of these concepts were much more likely to be listed by people when the concept label made that property available in a perceptual simulation. For example, people rarely listed *roots* as a property of the concept *grass*, but were much more likely to list *roots* of the concept *rolled-up grass*. Wu and Barsalou [23] suggest that this finding reflects the perceptual simulations people construct when they process concepts[24].

To summarize, the perceptual approach calls into question the assumption of classical models that there are amodal symbols. This approach suggests that using specific representations derived from perception can allow cognitive systems greater flexibility than can be achieved with amodal symbols.

2. Situated action

The classical approach to representation often views cognition as something that can be modeled inside a computer. By taking seriously the role of perception in conceptual representations, it becomes more difficult to separate cognitive processes from the context in which they take place. In the study of situated action (sometimes called situated cognition), it is assumed that cognitive processing cannot be extracted from the environment in which it occurs [25, 26 ,27].

Two important insights follow from this focus on context. First, it may not be necessary to represent all of the information relevant to thinking about a situation, because a substantial amount of that information will be present in the environment.

Second, the problem that an agent has to solve may be eased by aspects of the environment that would be hard to foresee if the agent had to reason abstractly.

On the first point, when cognition is situated, the agent can rely on the fact that the world is enduring to avoid having to represent the world extensively. Two examples will demonstrate this point. First, studies of *change blindness* have demonstrated that people do not store much of the visual world in an enduring fashion [28, 29, 30]. In various ways, these studies find that people have difficulty detecting changes in information in visual images to which they were not attending. While it might seem inefficient to lose this information when a fixation ends, the world typically does not change drastically from moment to moment, and so there is little real cost to storing only that information that was in focal attention.

In addition, an agent may simplify its representation of the world by representing many things with respect to itself. For example, Agre and Chapman [31] developed a simulated agent that existed in a video-game world. Rather than forming a detailed map of the world and keeping track of all of the objects and their global coordinates in space, objects were represented by their relationship to the agent itself. For example, an attacking enemy in the game would be represented as something chasing the agent. The agent would use this same representation for any attacking enemy, even if it were a different one, because what was relevant was the relationship between the agent and the enemy at that moment.

The second aspect of situated action is that the problem an agent must solve depends on the environment in which it is embedded. For example, Hutchins [26] provides an extensive description of the way navigation teams aboard naval vessels keep track of a ship's position. At a general level, the problem that must be solved by a navigation team involves fixing the position of the ship in the environment and ensuring that the ship maintains a course that keeps it from running aground. However, navigation teams have many specialized tools including two-dimensional overhead perspective maps, protractors, pencils, and devices for measuring the relative location of landmarks with respect to the ship and assessing the depth of the water. These tools turn navigation into a task in which relative locations are drawn onto a map to determine the position of the ship. Navigation need not be carried out in this way, and Hutchins describes a system used by Micronesian sailors that conceptualizes navigation in terms of time rather than distance. Because they represent their task in a different way, the cognitive operations needed to solve it also differ.

In summary, because cognitive agents are embedded in environments, it is not necessary for them to form complete representations of that environment at all times. Instead, they can rely on the fact that the world is relatively stable. Thus, many fewer representations in the cognitive system need to be enduring than has typically been assumed by classical approaches to representation. The situated action approach also means that agents can simplify the task they have to solve by representing information relative to themselves. Finally, the task environment determines the problem that an agent actually has to solve. Often, what appears to be a difficult problem when cast abstractly is actually much easier when embedded in an actual situation. The goal of cognitive science, in this view, is to understand how agents structure their environment in order to solve complex tasks.

3. Embodied cognition

Related to the situated action approach is embodied cognition. On this view, not only is it crucial to think about the contexts in which cognitive processing occurs, it is also necessary to *build* agents that actually interact in real environments [32, 33, 34, 35, 36]. Building real agents suggests ways that the environment can be exploited to solve difficult problems. Furthermore, while there may be many possible ways of representing information, the space of potential representations may be much narrower when the agent must achieve sensorimotor coordination. Thus, this view explicitly rejects the idea that cognitive theories can ignore perceptual and motor systems.

There are many ways that the environment can be exploited to solve difficult problems. From Gibson's classic work on perception forward, scientists have demonstrated that the visual system is sensitive to information in the environment that provides information relevant to an organism's goals. For example, many species are able to use optic flow to gauge their direction and speed of motion. In addition, the vestibular systems provide information about linear and angular acceleration that can be used to augment visual information in the construction of cognitive maps [37, 38, 39].

Sometimes building a real agent can also lead to simple solutions to potentially difficult problems. For example, Pfeifer and Scheier [36] describe a robot that is able to distinguish large cylinders from small ones. This classification is accomplished by providing the robot with simple motor routines that allow it to follow walls and therefore to circle around objects. When the robot circles a small cylinder, the ratio of the speed of the outside wheel to the inside wheel is higher than when it circles a large cylinder. By using sensors that provide information about the speed of its wheels, the robot is able to perform a classification task without an elaborate visual system.

Finally, Glenberg (1997) suggests that understanding memory requires attending to the function of memory within an organism. Many forms of memory require little effort, such as the perceptual priming observed following the presentation of a stimulus or the ability to point to the location of an object in space when the organism is navigating through that space. Glenberg [35] argues that these forms of memory are what permit organisms to carry out actions in the world. More effortful forms of memory require suppression of current input, which is what makes them more difficult to use [40]. Finally, he argues that language comprehension involves representing information as if the comprehender were going to act in the situation. In each of these cases, it is assumed that understanding cognition requires focusing on the relationship between an embodied organism and its environment.

The embodied cognition approach has had great success at building very simple machines that navigate through environments and avoid obstacles. These agents are even able to perform simple tasks like picking up cans or classifying simple objects [34, 36].

The claim these researchers make is that all of cognition, including higher cognition, can be successfully modeled using this bottom-up approach.

4. Dynamical systems

A final challenge to traditional assumptions about representation has come from proponents of *dynamical systems* as explanations of behavior [41, 42, 43]. Dynamical systems are systems of nonlinear differential equations that can be used to describe aspects of behavior (see Norton [44] for an introduction). On this view, a central problem with traditional approaches to representation is that they have discrete and enduring components. Dynamical systems do not involve discrete symbols.

In a dynamical system, there is a current state consisting of the values of some set of control variables. There is also a set of equations that combine the control variables to govern how the system changes over time. Thus, the two key aspects of dynamical systems are that they involve continuous change in the values of the control variables, and that this change occurs continuously in time. Hence, dynamical systems assume that representations are time-locked to information in the represented world. As the state of the represented world changes, the representation changes as well.

As an example, Kelso [41] describes a number of studies involving the coordination among limbs. For example, put your hands in a fist and place them in front of you. Then, extend the index fingers on both hands. Now flex and extend these index fingers in synchrony, increasing the speed of movement. Most people are able to maintain this coordination, even at high speeds. In contrast, try this same task, except flex one index finger as you extend the other. As people increase the speed of this movement, it often becomes difficult to maintain, until finally, they end up performing the first movement (flexing and extending both fingers at the same time). Kelso is able to describe this movement, as well as many more complex kinds of motor coordination using dynamical systems. Further, he makes a convincing case that this type of explanation is superior to an explanation of these behaviors involving other types of representations. In this model, the state of the system changes through time as the positions of the fingers change. Thus, this model contains no enduring representations.

Some researchers have argued that this success in describing motor behavior can be extended to all of cognitive processing [42, 43]. They suggest that dynamical systems have two advantages over other approaches to cognition. First, by focusing on processes that evolve continuously, they are able to account for the plasticity of cognition. Second, it is assumed that continuous processes allow dynamical systems to account for the fine details of processing, which in turn allows them to account for individual differences. This focus on individual differences contrasts with much research in cognitive science, which focuses on commonalities in behavior across individuals.

Semantics and representation

The four alternative approaches to representation have focused primarily on low-level perceptual and motor processes. They have not had success at explaining higher-level cognition. There is a good reason for the problems these models have with complex cognitive processes. To some degree, each of the alternative approaches ties representations to perceptual and motor pathways. On the positive side, this coupling of representation with perceptual and effector systems provides a basis for the semantics of the representation. In particular, one important way that representations come to have meaning is for them to correspond to something external to the agent.

On the negative side, using correspondence as the primary basis for semantics is more likely to be successful for perceptual and motor processes than for high-level cognition. People's ability to represent abstract concepts involves a second aspect of semantics: functional role. That is, the meaning of a representational element is also determined by its relationship to other representational elements. If a theory of representation focuses primarily on correspondence, then processes that require functional role information will be difficult to explain.

So, how should the classical view be extended?

The classical view of representation has served as the basis for research in cognitive science since the late 1950s. This is a long time for a single framework to hold sway in a young science. Nonetheless, there is no reason to abandon the classical view yet. None of the problems identified by advocates of the four alternative approaches are fatal to the classical approach to representation. Instead, they are simply signs of growing pains.

All of the approaches to representation discussed here agree on the fundamental assumption that cognitive processing involves internal mediating states that carry information. Thus, the exploration of representation can be fruitfully described as an examination of the types of properties that must be added to the basic concept of a mediating state in order to capture cognitive processing.

Each of the alternative approaches essentially highlights particular properties that must be added to mediating states in order to account for cognitive processing [5]. Thus, the remaining assumptions of the classical view all require some change in light of the issues raised by alternative approaches, but it is always a change in scope. Not all representations are enduring, not all are symbols, not all are amodal, and not all are independent of the sensory and effector systems of the agent.

The assumption that some representations are amodal is the one that requires the most future scrutiny. The studies described in the section on perceptual symbol systems suggest that tying representations to specific modalities may provide the basis for considerable flexibility in cognitive processing, and may even account for the use of abstract concepts. While it is too early to argue that cognitive science can dispense with amodal representations, it may be able to go a long way without them.

The other three assumptions of the classical view are likely to survive intact for most aspects of higher cognitive processing. The assumption that cognitive systems have enduring states was challenged both by situated action and dynamical systems approaches. The situated action approach captures the important insight that many aspects of the world remain stable and thus do not need to be incorporated into enduring representations. Classical models will have to focus on ways that agents can use the world as a representation. The dynamical systems view further asserts that representations undergo continuous change in relation to changes in the external environment. This criticism seems less problematic for classical models, as there are many cases where an agent must be able to represent the past in order to be able to reason.

The dynamical systems view also challenges the importance of discrete symbols. Dynamical systems has successfully demonstrated that continuous representational states are important for capturing low-level perceptual and motor processes. However, there is good reason to believe that many cognitive processes do require discrete symbols. For example, many of the aspects of cognition that make perceptual symbol systems attractive argue against dynamical systems as the sole mode of cognitive representation. For example, people's ability to represent spatial relations in language, and to freely substitute arguments into those relations suggests that there must be discrete components that endure beyond particular sensory stimulation.

Finally, the degree to which perceptual and motor systems must be considered when modeling cognitive processing is an open question. The embodied cognition approach suggests that building actual agents is necessary for constructing cognitive models. The perceptual symbol system view requires that representational assumptions must be compatible with what is known about perception. Furthermore, the situated action approach to cognitive processing suggests that the environment is an important source of information that is used by cognitive agents to solve problems. Thus, models must be able to take advantage of information in the environment.

Despite the clear importance of perception in cognitive processing, cognitive science must continue to develop models of higher cognitive processes. Perception is not a purely bottom-up process. As discussed above, expertise in a domain changes the way people perceive the basic features of that domain. Thus, without models of how complex reasoning and expertise develops, we will not be able to understand how perceptual representations are constructed. While cognitive science would ultimately like to have explanations that span from sensation to higher level cognition, these models cannot be developed in a purely bottom-up fashion.

In summary, the classical approach to representation must be extended, but not replaced. The fundamental assumptions that there are internal mediating states and that many of those states are symbolic, enduring, and amodal form the core of the computational view of mind. Because these assumptions can be retained, the basic approach to cognitive science remains intact. The core insights of the alternative approaches to representation, however, do require significant changes to the base view. In particular, cognitive models must be more sensitive to perceptual representation. In doing this, we must now seriously address the problem of how high-level concepts are formed from low-level percepts.

Outstanding questions

1. Are there any amodal symbols in cognition?

2. How do abstraction representations of lower-level mediating states, especially those of continuously sensing dynamical systems, get made in such a way that they are useful for high-level cognition?

3. To what degree do higher-level cognitive processes vs. lower-level perceptual processes place constraints on the form of cognitive representations?

4. Can effective models of cognitive processes be developed without first modeling sensory and effector systems?

Acknowledgements

This work was supported in part by NSF grant SBR-9905013. The authors thank Dedre Gentner and Ryan Gossen.

References

[1] Anderson, J. R. (1978). Arguments concerning representations for mental imagery. Psychological Review, 85(4), 249-277.

[2] Marr, D. (1982). Vision. New York: W.H. Freeman and Company.

[3] Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), Cognition and Categorization (pp. 259-302). Hillsdale, NJ: Lawrence Erlbaum Associates.

[4] Markman, A. B. (1999). Knowledge representation. Mahwah, NJ: Lawrence Erlbaum Associates.

[5] Markman, A. B., & Dietrich, E. (2000). In defense of representation. Cognitive Psychology, 40(2), 138-171.

[6] Tversky, A. (1977). Features of similarity. Psychological Review, 84(4), 327-352.

[7] Peirce, C. S. (1897/1955). Logic as semiotic: The theory of signs. In J. Buchler, ed., *The philodophical writings of Pierce* (1955). New York: Dover, 98-119.

[8]Schank, R.C., & Abelson, R. (1977) Scripts, plans, goals, and understanding. Hillsdale, NJ: Lawrence Erlbaum Associates.

[9]Schank, R.C. (1982). Dynamic memory. New York: Cambridge University Press.

[10]Herskovits, A. (1986). Language and spatial cognition: An interdisciplinary study of the prepositions in English. New York: Cambridge University Press.

[11]Landau, B., & Jackendoff, R. (1993). "What" and "where" in spatial language and spatial cognition. Behavioral and Brain Sciences, 16, 217-266.

[12]Miller, G.A., & Johnson-Laird, P.N. (1976). Language and perception. Cambridge, MA: Harvard University Press.

[13] Barsalou, L. W. (1999). Perceptual symbol systems. Behavioral and Brain Sciences, 22(4), 577-660.

[14] Schyns, P. G., Goldstone, R. L., & Thibaut, J. P. (1998). The development of features in object concepts. Behavioral and Brain Sciences, 21(1), 1-54.

[15] Hampton, J. A. (1995). Testing the prototype theory of concepts. Journal of Memory and Language, 34, 686-708.

[16] Reed, S. K. (1972). Pattern recognition and categorization. Cognitive Psychology, 3, 382-407.

[17] Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification. Psychological Review, 85(3), 207-238.

[18] Nosofsky, R. M. (1986). Attention, similarity and the identification-categorization relationship. Journal of Experimental Psychology: General, 115(1), 39-57.

[19] Porter, B. W., Bareiss, R., & Holte, R. C. (1990). Concept learning and heuristic classification in weak theory domains. Artificial Intelligence, 45, 229-263.

[20] Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. Psychological Review, 101(1), 53-97.

[21] Schyns, P. G., & Rodet, L. (1997). Categorization creates functional features. Journal of Experimental Psychology: Learning, Memory, and Cognition, 23(3), 681-696.

[22]Lesgold, A., Rubinson, H., Feltovich, P., Glaser, R., Klopfer, D., & Wang, Y. (1988). Expertise in a complex skill: Diagnosing x-ray pictures. In M.T.H. Chi, R.

Glaser, & M.J. Farr (Eds.) The nature of expertise. Hillsdale, NJ: Lawrence Erlbaum Associates.

[23] Wu, L. L., & Barsalou, L. W. (in preparation). Grounding concepts in perceptual simulations: I. Evidence from property generation.

[24] Glenberg, A. M., Kruley, P., & Langston, W. E. (1994). Analogical processes in comprehension: Simulation of a mental model. In M. A. Gernsbacher (Ed.), Handbook of Psycholinguistics . New York: Academic Press.

[25] Clancey, W. J. (1997). Situated cognition: On human knowledge and computer representations. New York: Cambridge University Press.

[26] Hutchins, E. (1995). Cognition in the wild. Cambridge, MA: The MIT Press.

[27] Suchman, L. A. (1987). Plans and situated actions: The problem of human-machine communication. New York: Cambridge University Press.

[28] Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.), Perception (pp. 89-110). New York: Oxford University Press.

[29] Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. Psychological Science, 8(5), 368-373.

[30] Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people during a real-world interaction. Psychonomic Bulletin and Review, 5(4), 644-649.

[31] Agre, P. E., & Chapman, D. (1987). Pengi: An implementation of a theory of activity. Paper presented at the Proceedings of AAAI-87, Seattle, WA.

[32] Agre, P. E. (1995). Computational research on interaction and agency. Artificial Intelligence, 72, 1-52.

[33] Brooks, R. A. (1991). Intelligence without representation. Artificial Intelligence, 47, 139-159.

[34] Brooks, R. A. (1999). Cambrian intelligence. Cambridge, MA: The MIT Press.

[35] Glenberg, A. M. (1997). What memory is for. Behavioral and Brain Sciences, 20(1), 1-55.

[36] Pfeifer, R., & Scheier, C. (1999). Understanding intelligence. Cambridge, MA: The MIT Press.

[37] Golledge, R. G. (1999). Human wayfinding and cognitive maps. In R. G. Golledge (Ed.), Wayfinding Behavior: Cognitive mapping and other spatial processes (pp. 5-45). Baltimore, MD: The Johns Hopkins University Press.

[38] Golledge, R. G., & Stimson, R. J. (1997). Spatial behavior: A geographic perspective. New York: The Guilford Press.

[39] Richardson, A. E., Montello, D. R., & Hegarty, M. (1999). Spatial knowledge acquisition from maps and from navigation in real and virtual environments. Memory and Cognition, 27(4), 741-750.

[40] Glenberg, A. M., Schroeder, J. L., & Robertson, D. A. (1998). Averting the gaze disengages the environment and facilitates remembering. Memory and Cognition, 26(4), 651-658.

[41] Kelso, J. A. S. (1995). Dynamic patterns: The self-organization of brain and behavior. Cambridge, MA: The MIT Press.

[42] Port, R. F., & Van Gelder, T. (Eds.). (1995). Mind as Motion. Cambridge, MA: The MIT Press.

[43] Thelen, E., & Smith, L. B. (1994). A Dynamic Systems Approach to the Development of Cognition and Action. Cambridge, MA: The MIT Press.

[44] Norton, A. (1995). Dynamics: An introduction. In R. F. Port & T. v. Gelder (Eds.), Mind as motion (pp. 45-68). Cambridge, MA: The MIT Press.

(Peirce, 1897/1955 ----)

Figure Captions

Figure 1. Example of a perceptual feature like the ones used by Schyns and Rodet, (1997).

XY

X

Y

Box A: Cognitive grammar

Cognitive grammar is a prominent part of the cognitive linguistics movement, which attempts to account for grammatical phenomena using representations and processes that are continuous with those used by other cognitive processes [A1]. On this view, grammar facilitates the construction of representations and uses both perceptual and attentional processes. For example, the representation of the prepositions "above" and "below" shown in Figure A1a involve setting up locations in a semantic space, and then focusing attention on one of the objects (the top on in the case of "above" and the bottom one in the case of "below").

In this example, the productivity of grammar is accomplished by allowing representations of the arguments of the prepositions to be freely bound to the circles in this representation. Thus, representing the phrase "the lamp is above the table" would involve binding a symbol for the lamp to the top argument to "above" and a symbol for the table to the bottom argument.

A variety of different kinds of grammatical structures can be represented using these principles. For example, temporal events can be represented by extending the representations in time. For example, the concept "arrive" can be represented by a situation in which one argument gradually gets closer to a second fixed argument over time until they eventually meet. The moving object is the focus of attention in this representation.

In sum, cognitive grammar uses principles of perceptual representation and attention to account for structural aspects of language. These perceptual properties form the basis of Barsalou's perceptual symbol systems.

[A1]Langacker, R.W. (1987). Foundations of Cognitive Grammar. Stanford, CA: Stanford University Press.

Figure A.1

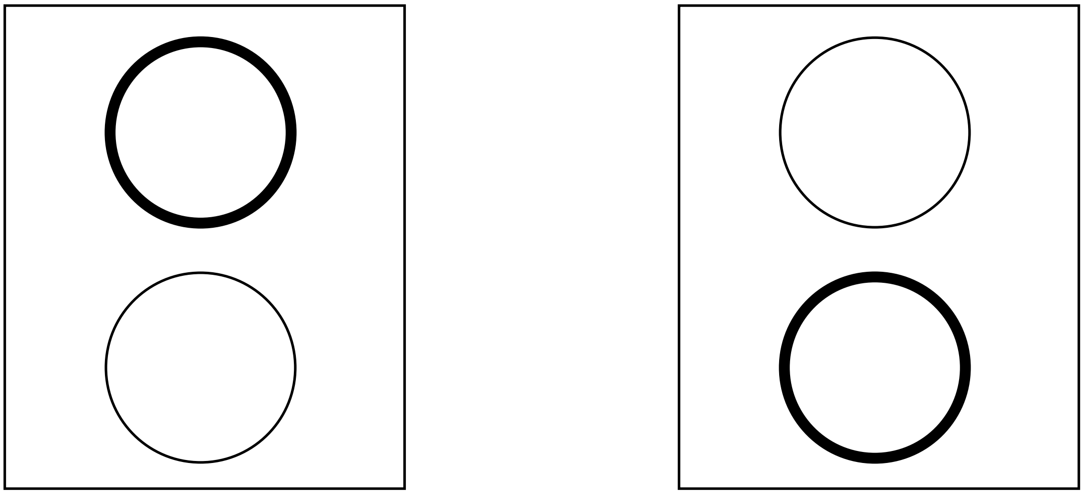

above

below