Edited by

Gün R. Semin

Eliot R. Smith

Embodied Grounding

Social, Cognitive, Affective, and Neuroscientific Approaches

Cambridge

Cambridge www.cambridge.org/9780521880190

This page intentionally left blank

EMBODIED GROUNDING

How do the accomplishments that make us most deeply human – our capacity to think, to use language, to experience emotions, and to relate to one another – depend on our bodies? Exploring these issues involves laying aside an ancient Western tradition that placed body and mind in opposition, as well as more recent scientific understanding of thought as abstract, disembodied information processing. The chapters in this volume review current work on relations of the body to thought, language use, emotion, and social relationships, as presented by internationally recognized experts in these areas.

Gün R. Semin is Academy Professor, Royal Netherlands Academy of Arts and Sciences, Utrecht University, Faculty of Social Sciences.

Eliot R. Smith is Professor of Psychological and Brain Sciences, Indiana University, Bloomington, Indiana.

Embodied Grounding

Social, Cognitive, Affective, and Neuroscientific Approaches

Edited by

GÜN R. SEMIN
Utrecht University

ELIOT R. SMITH
Indiana University, Bloomington

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK
Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521880190

First published in print format 2008

ISBN-13 978-0-511-47402-6 eBook (Adobe Reader)

ISBN-13 978-0-521-88019-0 hardback

ISBN-13 978-0-521-70615-5 paperback

Contents

Contributors

Lisa Feldman Barrett Boston College, MA, USA

Lawrence W. Barsalou Emory University, GA, USA

Lera Boroditsky Stanford University, CA, USA

Pablo Briñol Universidad Autónoma de Madrid, Spain

John T. Cacioppo University of Chicago, IL, USA

Gerald L. Clore University of Virginia, VA, USA

Jens Förster International University, Bremen, Germany, and University of Amsterdam, The Netherlands

Ronald S. Friedman University of Albany, State University of New York, NY, USA

Arthur M. Glenberg University of Wisconsin at Madison, WI, USA

Kristen A. Lindquist Boston College, MA, USA

Paula M. Niedenthal Centre National de la Recherche Scientifique (CNRS) and University of Clermont-Ferrand

Lindsay Oberman University of California, San Diego, CA, USA

Richard E. Petty Ohio State University, OH, USA

Jesse Prinz University of North Carolina, Chapel Hill, NC, USA

Friedemann Pulvermüller Medical Research Council, Cognition and Brain Sciences Unit, Cambridge, UK

Simone Schnall University of Plymouth

Thomas W. Schubert Utrecht University

Beate Seibt Utrecht University

Gün R. Semin Utrecht University

Eliot R. Smith Indiana University, Bloomington, IN, USA

Sven Waldzus ISCTE, Lisbon

Piotr Winkielman University of California, San Diego, CA, USA

Introducing Embodied Grounding

Gün R. Semin and Eliot R. Smith

In just the last two decades, the embodiment perspective has inspired research and new theoretical ideas across a wide swath of the behavioral and cognitive sciences. Much of the appeal underlying this impact arises from the simple insight upon which the core ideas of embodiment rest: that nervous systems evolved for the adaptive control of action – not for abstract cogitation (chess-playing, in Brooks's memorable 1991 statement of this insight). This idea immediately has several significant implications.

First, minds co-evolved with bodies, especially sensory-motor systems. There are now many examples illustrating the importance of this fact. One is the way organisms use the physical properties of bodies to reduce the need for costly central computation (Thelen & Smith, 1994; Pfeifer & Scheier, 1999). Another is the "active vision" idea that when agents move, far from introducing problematic variation into sensory inputs, they actually create the conditions for discovering cross-modal associations, facilitating understanding of the environment and the adaptive shaping of action (Edelman, 1987; Pfeifer & Scheier, 1999).

Second, the embodiment approach suggests a renewed focus on the whole behaving organism in its natural context as the object of study. Seen from this perspective, either the isolation of specific "slices" such as central information-processing systems, or the study of organisms in environments vastly different from those in which they evolved, seem less than optimal research approaches.

Third, the context in which organisms behave is very often a social context. In the case of humans, the very capabilities that make us human – to elaborate philosophical ideas, build cities, sing, and dance – are usually collective acts or, even if performed by an individual, depend absolutely on a multitude of social products (not the least of which is language). The embodiment approach should shed light not only on the functioning of individual humans but also on the ways they cooperate, compete, and otherwise relate to each other, in groups and as individuals.

Fourth, a focus on the whole behaving organism in its context breaks down traditional disciplinary boundaries as well as distinctions between topics like cognition and motivation. The recognition of this basic fact has driven psychological theory and research to focus increasingly on the interdependences among cognition, motivation, affect, and action as they are all influenced by the body. An embodied approach is taking root not only in psychology but across a variety of disciplines ranging from the neurosciences to developmental processes to cognitive sciences and robotics with an increasingly powerful synergy between these different approaches. In this situation, fresh discoveries in one "field" lead to reformulations of the very same issues in a different "field."

The impetus to cross boundaries is also the inspiration for this volume, the foundations of which were laid during a four-day conference including the majority of the contributors in May 2006. This meeting gave rise to a unique synergy because the participants came from quite different specialties ranging from neuroscience to cognitive psychology, social psychology, affective sciences, and psycholinguistics. The discussions and exchanges resulted in the current volume, which follows a tripartite organization. The contributions that make the first part of this volume address the *embodied grounding of concepts and language*. Whereas the disciplinary breadth of the material in these contributions is considerable, ranging from neuroscience to experimental cognitive psychology, the authors of the different chapters are not only informed about each other's work but also influenced by it. These contributions examine the central issues in this field from complementary perspectives because interaction at the conference helped them build bridges across traditional boundaries. The driving theme for these four chapters is that our representations of the social world are fundamentally connected with the actions that our bodies perform, so that these actions inform our concepts, language, and thinking.

- The opening chapter by Barsalou presents an embodied account of symbolic operations, proceeding from the neural simulation of concrete concepts to how these relate to abstract concepts and symbols. He reviews the most recent research in this area, which come from a variety of different methods ranging from experimental behavioral work to neuroscientific findings. In concluding his contribution, Barsalou examines the link between language and perceptual symbols systems, arguing for links

between, for instance, syntax and the role it may play in psychological processes such as retrieval, how simulations are assembled, and the nature of recursively embedded structures.
- The next contribution by Glenberg is based on his Indexical Hypothesis, a model that is designed to interface bodily states with language and action. His chapter outlines this hypothesis and the processes underlying how meaning is understood.
- Indeed, there is considerable intellectual cross-fertilization not only between Glenberg's Indexical Hypothesis and Barsalou's perceptual symbol systems model but also with the third chapter of this section, namely, Pulvermüller's examination of the cortical mechanisms responsible for semantic grounding and embodiment concepts.
- The final chapter in this part by Boroditsky and Prinz examines the sources contributing to human knowledge and thinking. They argue for two input streams to the complex human knowledge system and thinking. The first stream they refer to is the use of stored records of sensory and motor states, inspired by Barsalou's (Chapter 1) perceptual symbols systems model. The second source that they elaborate upon is the contribution of language – treated not as an abstract inner mental "code" but as a rich store of sensorimotor regularities in the real world, whose statistical properties offer important evidence for the construction and constitution of thought.

The second part of the volume focuses on the *embodiment of social cognition and relationships.*

- Semin and Cacioppo outline a model of social cognition that breaks away from a traditional individual-centered analysis of social cognition and treats social cognition as grounded by neurophysiological processes that are distributed across brains and bodies and is manifested in the co-regulation of behaviors. The theoretical framework they introduce is an attempt to model the processes involved from joint perception to co-regulation in social interaction.
- Smith describes how social relationships that link people to other individuals or social groups are both expressed and regulated by bodily processes. Drawing on Alan Fiske's Relational Models Theory, research described in this chapter tests hypotheses that both synchronized movements and interpersonal touch operate as embodied cues to close relationships (communal sharing relationships, in Fiske's terminology).
- A related chapter by Schubert, Waldzus, and Seibt also draws upon Fiske's model, relating it to Barsalou's fundamental point that abstract concepts (such as interpersonal closeness or differences in power or authority) are

understood in terms of bodily metaphors. For example, research finds that authority differences are expressed in differences in size, height, or vertical position; the powerful literally do "lord it over" the rest of us.

- Briñol and Petty deal with a different aspect of social cognition: the effects of embodiment on processes involved in social influence, especially those leading to changes in attitudes or evaluations of particular objects. Their discussion is based on the Elaboration Likelihood Model, one of the best-supported and most far-reaching theoretical accounts of attitude change, showing how it organizes findings about the effects of bodily movements on attitudes, as well as generating new and intriguing predictions.

The third part has as its topic the *embodiment of affective processes*.

- Clore and Schnall take the viewpoint that affective reactions provide embodied evidence that people can use to validate or invalidate their evaluative beliefs about objects. In particular, affective reactions are sometimes found to have limited effects unless preexisting beliefs exist that are congruent with the affective reactions. The authors describe provocative research suggesting that holding incongruent beliefs and affective reactions has cognitive costs for the individual.
- Barrett and Lindquist address the embodiment of emotional responses. Noting that traditional theories hold that body and mind make separate and independent contributions to an emotional episode, they apply the embodiment perspective to generate the novel suggestion that the body helps constitute the mind in shaping an emotional response. That is, the conceptual knowledge that we use to categorize and understand our own (and other people's) emotions is itself represented in sensorimotor terms.
- Winkielman, Niedenthal, and Oberman similarly take an embodied approach to emotional processes in which such processes are grounded in modality-specific systems. They describe studies directly testing the hypothesis that manipulating bodily resources will influence the perception and understanding of emotional events. The chapter ranges widely, covering the role of embodiment in the formation of attitudes as well as in the representation of abstract emotion concepts.
- The chapter by Förster and Friedman presents a new conceptual model of the effects of bodily movements, facial expressions, and other embodied cues on attitudes as well as on cognitive processing styles (such as creativity and flexibility). Guided by Higgins's Regulatory Focus Theory, they show how effects of embodied cues can be accommodated and also how paradoxes in existing evidence may be resolved.

Taken together, these chapters illustrate the extraordinarily broad range of topics upon which research has been influenced by the embodiment perspective. As we observed at the beginning of this introduction, embodiment calls for a focus on the entire organism rather than on isolated slices of information-processing or behavioral systems. As a group, these chapters reflect this breadth of focus, as researchers and theorists grapple with the implications of embodiment for levels and topics ranging from neuroscience to language comprehension, social relationships to attitude change, and affect to cognition. We hope that this volume will inspire and excite still more boundary-crossing research and theory, faithful to the true underlying message of the embodiment perspective.

PART ONE

EMBODIED LANGUAGE AND CONCEPTS

1 Grounding Symbolic Operations in the Brain's Modal Systems

Lawrence W. Barsalou

A central theme of modern cognitive science is that symbolic interpretation underlies human intelligence. The human brain does not simply register images, as do cameras or other recording devices. A collection of images or recordings does not make a system intelligent. Instead, symbolic interpretation of image content is essential for intelligent activity.

What cognitive operations underlie symbolic interpretation? Across decades of analysis, a consistent set of symbolic operations has arisen repeatedly in logic and knowledge engineering: binding types to tokens; binding arguments to values; drawing inductive inferences from category knowledge; predicating properties and relations of individuals; combining symbols to form complex symbolic expressions; representing abstract concepts that interpret metacognitive states. It is difficult to imagine performing intelligent computation without these operations. For this reason, many theorists have argued that symbolic operations are central, not only to artificial intelligence but to human intelligence (e.g., Fodor, 1975; Pylyshyn, 1973).

Symbolic operations provide an intelligent system with considerable power for interpreting its experience. Using type-token binding, an intelligent system can place individual components of an image into familiar categories (e.g., categorizing components of an image as people and cars). Operations on these categories then provide rich inferential knowledge that allows the perceiver to predict how categorized individuals will behave, and to select effective actions that can be taken (e.g., a perceived person may talk, cars can be driven). Symbolic knowledge further allows a perceiver to analyze individuals in an image, predicating properties and relations that apply to them (e.g., identifying a person as an adult male, or two people as having a family resemblance). Such predications further support high-level cognitive operations, such as decision making (e.g., purchasing a gas, hybrid, or electric car), planning (e.g., finding electricity to charge an electric car while driving around town), and problem solving (e.g., how to get in if the keys are locked in the car). Symbolic operations also include a variety of operations for combining symbols, such that an intelligent system can construct complex symbolic expressions (e.g., combining word meanings during language comprehension). Finally, by establishing abstract concepts about mental states and mental operations, an intelligent system can categorize its mental life in a metacognitive manner and reason about it (e.g., evaluating one's planning and decision making strategies).

What mechanisms implement symbolic operations? Since the cognitive revolution, language-like symbols and operations have been widely assumed to be responsible. Numerous theoretical approaches have been derived from predicate calculus and propositional logic. Not only have these approaches been central in artificial intelligence (e.g., Charniak & McDermott, 1985), they have also been central throughout accounts of human cognition (e.g., Barsalou, 1992; Barsalou & Hale, 1993; Anderson, 1983; Newell, 1990).

Although classic symbolic approaches are still widely accepted as accounts of human intelligence, and also as the engine for artificial intelligence, they have come increasingly under attack for two reasons. First, classic symbolic approaches have been widely criticized for not being sufficiently statistical. As a result, neural net approaches have developed to remedy this deficiency (e.g., Rumelhart & McClelland, 1986; O'Reilly & Munakata, 2000). Second, classic symbolic approaches have been criticized for not being grounded in perception, action, and introspection. As a result, researchers have argued that higher-order cognition is grounded in the brain's modal systems.

As statistical and embodied approaches are increasingly embraced, the tendency to "throw the baby out with the bath water" has often been embraced as well. Some researchers have suggested that classic symbolic operations are irrelevant to higher cognition, especially researchers associated with neural nets and dynamical systems (e.g., van Gelder, 1990; but see Prinz & Barsalou, 2000). Notably, some neural net researchers have realized that symbolic operations are essential for implementing higher cognitive phenomena in knowledge, language, and thought. The problem in classic theories is not their inclusion of symbolic operations but *how* they implement them. For this reason, neural net researchers have developed neural net accounts of symbolic operations (e.g., Pollack, 1990; Smolensky, 1990). Interestingly, these approaches have not caught on widely, either with psychologists or with knowledge engineers. For psychologists, neural net accounts of symbolic processes have little psychological plausibility; for knowledge engineers, they are difficult and inefficient to implement. As a result, both groups continue to typically use classic approaches when symbolic operations must be implemented.

An alternative account of symbolic operations has arisen in grounded theories of cognition (e.g., Barsalou, 1999, 2003a, 2005a). Not only does this account have psychological and neural plausibility, it suggests a new approach to image analysis. Essentially, this approach develops symbols whose content is extracted from images. As a result, symbols can be bound to regions of images, thereby establishing type-token mappings. Inferences drawn from category knowledge also take the form of images, such that they can be mapped back into perception. Analysis of an individual in an image proceeds by processing its perceived regions and assessing whether perceptually grounded properties and relations can be predicated of them. Symbol combination involves the manipulation and integration of image components to construct structured images that, in effect, implement complex symbolic propositions. Abstract concepts result from situated introspection, namely, the process of perceiving internal mental and bodily states in the context of external situations and developing image-based representations of them for later use in reasoning.

The remaining sections present this framework in greater detail. The next section illustrates how symbolic operations could arise from simulation. The following section presents current empirical evidence for this account. The final section addresses the role of language in symbolic operations.

GROUNDING COGNITION IN THE BRAIN'S MODAL SYSTEMS

Standard architectures assume that amodal symbols are transduced from experience to represent knowledge. Figure 1.1 illustrates this general approach. On experiencing a member of a category (e.g., *dogs*), modal states arise in the visual system, auditory system, motor system, somatosensory system, and so on (i.e., the solid arrows in Figure 1.1a). These states represent sensory-motor information about the perceived category member, with some (but not all) of this information producing conscious experience. Although modal states are shown only for sensory-motor systems, we assume that modal states also arise in motivational systems, affective systems, and cognitive systems. The perception of these internal systems will be referred to as *introspection* from here forward. Once modal states arise in all relevant modal systems for a category, amodal symbols that stand for conceptual content in these states are then transduced elsewhere in the brain to represent knowledge about the category (e.g., *legs, tail, barks, pat, soft* in Panel B for

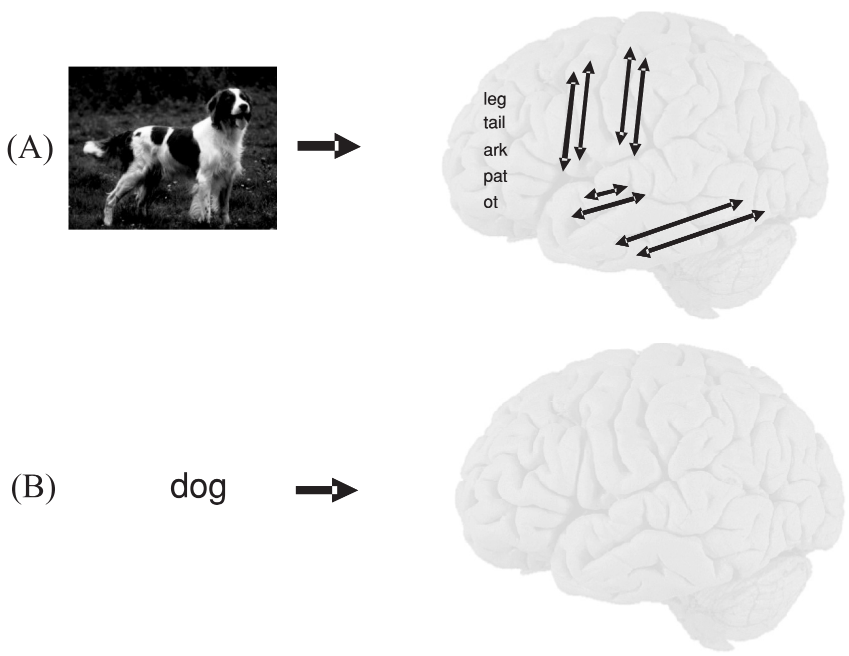

Figure 1.1. The transduction of amodal symbols from modal states in standard cognitive architectures (Panel A). Use of transduced symbols to represent the meaning of a word (Panel B). See the text for further description.

the experience of a dog). Although words often stand for transduced amodal symbols (e.g., *leg*), most theories assume that sub-linguistic symbols, often corresponding closely to words, are actually the symbols transduced (e.g., in Figure 1.1).

Once established in the brain, amodal symbols later represent knowledge about the category across a wide range of cognitive tasks (Figure 1.1, Panel B). During language comprehension, hearing the word for a category (e.g., "dog") activates amodal symbols transduced from modal states on previous occasions. Subsequent cognitive operations on category knowledge, such as inference, are assumed to operate on these symbols. Note that none of the modal states originally active when amodal symbols were transduced (Figure 1.1, Panel A) are active during knowledge representation (Figure 1.1, Panel B). Instead, amodal symbols are assumed to be sufficient, with modal states being irrelevant.

The architecture in Figure 1.1 underlies a wide variety of standard approaches to representing knowledge, such as feature lists, semantic networks, and frames. This architecture also underlies those neural net architectures where the hidden layers that represent knowledge are related arbitrarily to perceptual input layers. The architecture in Figure 1.1 does not underlie neural nets in which input units play roles in knowledge representation.

The Capture and Simulation of Modal States in Embodied Architectures

A very different approach to representing knowledge has arisen recently in theories of grounded cognition. Actually, this approach has deep roots in philosophical treatments of knowledge going back over 2000 years (e.g., Epicurus, 341–270BC/1987). Modern theories can be viewed as reinventions of these older theories in the contexts of psychology, cognitive science, and cognitive neuroscience (e.g., Barsalou, 1999, 2008; Prinz, 2002). Interestingly, the amodal architectures that currently dominate the field constitute a relatively recent and short presence in a historical context where theories grounded in the modalities have dominated.

Figure 1.2 illustrates the grounded approach to representing knowledge. On experiencing a member of a category (e.g., *dogs*), modal states are again represented as activations in the visual system, auditory system, motor system, somatosensory system, and so on (Panel A). As for Figure 1.1, modal states are only shown for sensory-motor systems, but this approach assumes that such states are also captured during the introspection of those motivational states, affective states, and cognitive states available to consciousness. Local association areas then partially capture these modality-specific states (shown in Panel A as asterisks). Higher-order cross-modal associations (also shown as asterisks) then integrate conjunctive neurons in lower-order association areas to establish a multimodal representation of the experience.

Once established in the brain, these multimodal associative structures represent knowledge about the category across a wide range of cognitive tasks (Figure 1.2, Panel B). During language comprehension, for example, hearing the word for a category (e.g., "dog") activates conjunctive neurons in higher-order cross-model association areas that have previously encoded experiences of the respective category. In turn, these conjunctive neurons activate lower-order conjunctive neurons that partially reactivate modal states experienced previously for the category. These neural reenactments attempt to simulate the modal states likely to occur when actually encountering category members. As the dashed arrows in Panel B illustrate, the modal states are only partially reenacted, not fully reenacted. For the remainder of this chapter, these reenactments will be referred to as *simulations*, given that they result from the brain attempting to simulate previous experience.

(A)

(B)

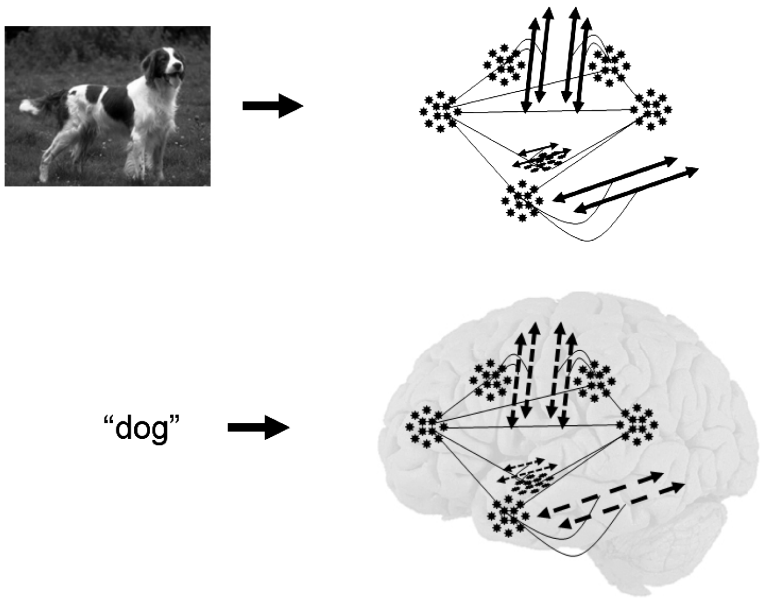

Figure 1.2. Conjunctive units in hierarchically organized association areas capture modal states across modalities in grounded theories of knowledge (Panel A). Captured multimodal states are simulated to represent the meaning of a word (Panel B). See the text for further description.

The capture-and-simulate architecture in Figure 1.2 underlies a wide variety of traditional and modern approaches to representing knowledge. Whereas some of these approaches focus on mental images, others focus on neural reenactments of modal states in the brain, which may often be unconscious. All share the common assumption that the representation of knowledge is grounded in modal states rather than in amodal symbols transduced from them. All share the common assumption that simulations of previous experience are central to cognition.

Additional Assumptions

Several misunderstandings often arise about the simulation architecture in Figure 1.2. First, this architecture is often viewed as an instantiation of classic empiricist theories. Second, this architecture is often viewed as a recording system incapable of interpretation. Third, this architecture is often viewed as capturing knowledge only from perception of the external world. Each of these misunderstandings will be addressed in turn.

First, the architecture in Figure 1.2 need not be a classic empiricist theory. In principle, specific simulations could be genetically determined rather than being acquired from experience. If particular simulations had evolutionary significance, they could be genetically encoded into the brain such that they are ready to become active in relevant situations. A more likely possibility is that strong genetic constraints determine the feature systems and association areas that exist in the brain. Because certain categories of objects, events, and introspective states were important over evolution, a neural architecture evolved that is optimally prepared to represent and process these particular categories (cf. Caramazza & Shelton, 1998). Although experience with categories is necessary to establish representations of them in memory, the feature systems relevant for representing these categories originate genetically, as do the associative systems for integrating them. Indeed, nativists have often assumed that images are central to cognition (e.g., Kant, 1787/1965; Reid, 1785/1969). There is no a priori reason why a system that uses simulation to represent knowledge cannot have strong genetic constraints as well.

Second, the architecture in Figure 1.2 need not be a recording system. Unlike a camera or video recorder, this architecture is not limited to capturing holistic bit-mapped records of experience. Instead, the interpretation of experience lies at the heart of this account, where interpretation arises from the application of symbolic operations (Barsalou, 1999, 2003a, 2005a). Because this account is the focus of this chapter, it is not described further. As will become clear, this architecture is highly capable of implementing interpretation.

Third, the architecture in Figure 1.2 need not only acquire knowledge from perception of the external world. Instead, it can capture extensive amounts of knowledge from introspection, namely, from the internal perception of affect, motivation, proprioception, pain, pleasure, cognitive states, cognitive operations, and so on. Because the human brain is capable of perceiving internal states (to some extent), it captures these states and simulates them on subsequent occasions, just as it captures and simulates the perception of external states. As described later, the capture and simulation of internal states appears central to the representation of abstract concepts, specifically, and to human cognition, more generally. Thus, the architecture in Figure 1.2 can be applied to the capture and simulation of both external and internal states.

The Status of Empirical Evidence for Grounded Knowledge

Accumulating empirical evidence supports the simulation architecture in Figure 1.2. Many findings indicate that the brain's modal systems for perception, action, and introspection are active during the higher cognitive activities of memory, knowledge, language, and thought. For reviews of the evidence from cognitive psychology, see Barsalou (2003b, 2008) and Barsalou, Simmons, Barbey, and Wilson (2003). For reviews of evidence from social psychology, see Barsalou, Niedenthal, Barbey, and Ruppert (2003) and Niedenthal, Barsalou, Winkielman, Krauth-Gruber, and Ric (2005). For reviews of evidence from cognitive neuroscience, see Martin (2001, 2007), Pulvermüller (1999), and Thompson-Schill (2003). For reviews of developmental evidence, see Smith (2005), Smith and Gasser (2005), and Thelen (2000). The rapidly accumulating findings across these diverse literatures indicate that the higher cognitive processes engage modal systems frequently and robustly.

The Importance of Going beyond Demonstration Experiments

Problematically, these findings do not indicate what roles the modalities play in higher cognition. When these findings were acquired, it was not accepted widely that modal systems participated in higher cognition at all. Researchers holding this hypothesis therefore attempted to evaluate it primarily using demonstration experiments. Typically, these researchers did not attempt to establish the roles that modal processing played in the experimental phenomena studied. Instead, the primary goal was simply to demonstrate that the brain's modal systems become active during cognitive processes. Now that the presence of modal processing is becoming well established in higher cognition, however, demonstration experiments are likely to have diminishing returns. Instead, it is becoming increasingly important to establish the specific roles that modal processing plays.

One possibility is that the brain's modal systems play relatively peripheral, or even epiphenomenal, roles in higher cognition. Although these systems become active, other systems that operate on amodal symbols implement basic cognitive operations (Figure 1.1). Alternatively, the brain's modal systems may provide the core computational engine in higher cognition (Figure 1.2). In particular, modal systems may implement fundamental symbolic operations, such as binding types to tokens, binding arguments to values, drawing inductive inferences from category knowledge, predicating properties and relations of individuals, combining symbols to form complex symbolic expressions, and representing abstract concepts that interpret introspective states metacognitively.

An Example

The following experiment illustrates both the utility and limitations of demonstration experiments. Simmons, Martin, and Barsalou (2005) assessed the neural representation of food categories. Participants lay passively in a functional magnetic resonance imaging (fMRI) scanner while briefly viewing food pictures and evaluating whether the picture currently present physically matched the previous picture. Participants were not asked to categorize the foods, to think about how they taste, or to conceptualize them in any manner. Nevertheless, the pictures activated brain areas that represent how foods taste and how rewarding they are to consume. Although participants did not actually taste anything, category knowledge about foods became active and produced activations in the brain's gustatory and reward systems.

This experiment offers a demonstration that the brain's modal systems become active during cognitive processing. When participants looked at visual pictures, their gustatory systems became active, even though they were not tasting anything. Nevertheless, it is not clear what role these gustatory activations play in higher cognition. Are they epiphenomenal? Or are they implementing symbolic operations associated with drawing taste and reward inferences about the perceived visual objects? Although this experiment, and many others like it, demonstrate that modal systems become active during higher cognition – a possibility that seemed unlikely until recently – this experiment does not shed any light on the specific role that these activations play.

Symbolic Operations

One possibility is that activations in modal systems underlie symbolic operations. Rather than amodal symbols implementing these operations, simulations implement them. The following subsections describe how simulations could, in principle, implement the symbolic operations of predication, conceptual combination, and abstract concepts. Later subsections summarize preliminary evidence for these accounts and their limitations to date.

Representing Concepts and Instances: Simulators and Simulations

To implement symbolic operations, it is essential for an intelligent system to have some means of learning and representing concepts. The lack of concepts is what prevents many recording devices, such as cameras and video recorders, from implementing the symbolic operations that would allow them to interpret the images they capture. The central innovation of Perceptual Symbol Systems theory (PSS) is its ability to implement concepts and their interpretive functions using image content as basic building blocks (Barsalou, 1999, 2003a, 2005a).

According to PSS, concepts develop in the brain as follows. Much research has shown that categories have statistically correlated features (e.g., *wheels,* *steering wheel*, and *engine* for *cars*; McRae, de Sa, & Seidenberg, 1997). Thus, encountering different instances of the same category should activate similar neural patterns in feature systems (cf. Farah & McClelland, 1991; Cree & McRae, 2003). Furthermore, similar populations of conjunctive neurons in the brain's association areas – tuned to these particular conjunctions of features – should tend to capture these similar patterns (Damasio, 1989; Simmons & Barsalou, 2003). Across experiences of a category's instances, this population of conjunctive neurons integrates the modal features of a category, establishing a distributed multimodal representation of it (Figure 1.2).

PSS refers to these distributed multimodal representations as *simulators*. Conceptually, a simulator functions as a type and is roughly equivalent to a concept in standard theories. Specifically, a simulator integrates the multimodal content of a category across instances and provides the ability to interpret later individuals as tokens of the type. Consider the simulator for the category of *cars*. Across learning, visual information about how cars look becomes integrated in the simulator along with auditory information about how they sound, somatosensory information about how they feel, motor programs for interacting with them, emotional responses to experiencing them, and so on. The result is a distributed system throughout the brain's feature and association areas that accumulates modal representations of the category.

In principle, an indefinitely large number of simulators can develop in memory for all forms of knowledge, including objects, properties, settings, events, actions, introspections, and so forth. Specifically, a simulator develops for any component of experience that attention selects repeatedly. When attention focuses repeatedly on a type of object in experience, such as *cars*, a simulator develops for it. Analogously, if attention focuses on a type of action (*driving*) or on a type of introspection (*fear*), simulators develop to represent it as well. Such flexibility is consistent with Schyns, Goldstone, and Thibaut's (1998) proposal that the cognitive system acquires new properties as they become relevant for categorization. Because selective attention is flexible and open-ended, a simulator develops for any component of experience that attention selects repeatedly.

Once a simulator becomes established for a category, it reenacts small subsets of its content as specific *simulations*. All the content in a simulator never becomes active at once. Instead, only a small subset becomes active to represent the category on a particular occasion (cf. Barsalou, 1987, 1989, 1993, 2003b). For example, the *car* simulator might simulate a sedan on one occasion but simulate a sports car or off-road vehicle on others. Because all the experienced content for cars resides implicitly in the *car* simulator, many different subsets can be reenacted as simulations.

The presence of simulators in the brain makes the implementation of symbolic operations possible. Indeed, symbolic operations follow naturally from the presence of simulators. Because simulators are roughly equivalent to concepts, the symbolic functions made possible by concepts are also made possible by simulators. The next three subsections illustrate how simulators enable three classic symbolic functions: predication, conceptual combination, and the representation of abstract concepts. For further details, see Barsalou (1999, 2003a, 2005a).

Implementing the Symbolic Function of Predication in PSS

To implement predication, an intelligent system must first distinguish types from tokens. In PSS, simulators implement types because they aggregate multimodal information across category members (e.g., cars). Conversely, simulations implement tokens because they represent category members (e.g., individual cars). Thus, the simulator-simulation distinction in PSS naturally implements the type-token distinction essential for predication.

This distinction further allows PSS to explain a wide variety of phenomena related to predication, including type-token predication, true versus false propositions, and interpretive spin. Type-token predication results from binding simulators to simulations (or perceptions). For example, binding the *car* simulator to a simulated (or perceived) car produces the predication that the individual is an instance of the *car* category. These type-token bindings essentially implement propositions, where binding the simulator to the individual represents a claim about the individual, namely, that the individual is a *car*. Notably, such propositions can be false, as when the *car* simulator is applied mistakenly to a small truck. Furthermore, the potential predications of an individual are infinite, thereby producing interpretative spin. Because indefinitely many simulators (and combinations of simulators) can be used to interpret an individual, indefinitely many true and false interpretations are possible. For example an individual car could be interpreted as a *car*, *vehicle*, *artifact*, *sedan, junked sedan, truck* (false), and so on. Thus, the simulator-simulation distinction allows PSS to implement classic symbolic functions related to predication.

Implementing Conceptual Combination in PSS

To see how PSS implements conceptual combination, first consider a simulator for the spatial relation, *above*. An *above* simulator could result from having pairs of objects pointed out in perception where the focal object always has a higher vertical position than the other object (e.g., a helicopter above a hospital). As each *above* relation is pointed out, selective attention focuses on the spatial regions containing the two objects, filters out the objects, and captures modal information about the shapes and sizes of the regions, the vertical distance between them, their horizontal offset, and so on (the parietal lobe and parahippocampal place areas are plausible regions where the *above* simulator might be represented in the brain). Over time, the *above* simulator captures many such pairs of regions and the spatial relations between them. On later occasions, this simulator can produce a wide variety of *above* simulations, each containing a pair of spatial regions not containing objects. An *above* simulation could represent two round regions of equal size, nearly touching vertically, with no horizontal offset; it could represent two rectangular regions of different size, distant vertically, with a large horizontal offset; and so on.

The *above* simulator lends itself to producing conceptual combinations. Imagine that this simulator produces one particular *above* simulation. Next, imagine that the *helicopter* simulator runs a simulation in the upper region of this *above* simulation, and that the *hospital* simulator runs a simulation in the lower region. The resulting simulation represents a helicopter above a hospital, thereby implementing a conceptual combination that expresses the proposition *ABOVE (helicopter, hospital)*. The individual simulators for *above, helicopter*, and *hospital* remain bound to their respective regions of the complex simulation, thereby continuing to interpret them symbolically in embedded type-token propositions.

Infinitely many other conceptual combinations can be implemented by simulating different types of objects or events in the regions of the *above* simulation, thereby expressing related propositions, such as *ABOVE (jet, cloud)*, *ABOVE (lamp, table)*, and so on. In general, this is how PSS implements conceptual combination. Because simulators represent components of situations and relations between components, their simulations can be combined into complex, multicomponent simulations. Much like an object-oriented drawing program, PSS extracts components of situations so that it can later combine them to represent either previously experienced situations or novel ones.

Representing Abstract Concepts in PSS

Relatively little is known about abstract concepts (e.g., *truth*, *thought*), given that most research has addressed concrete concepts (e.g., *car*, *bird*). Abstract concepts however, are likely to provide deep insights into the nature of human cognition, and to help produce increasingly sophisticated forms of intelligent computation, given the metacognitive capabilities that they enable.

Recent theory suggests that one central function of abstract concepts is to represent introspective states (e.g., Barsalou, 1999). In an exploratory study, more content about introspective states was observed in abstract concepts than in concrete concepts (Barsalou & Wiemer-Hastings, 2005). In another exploratory study, the abstractness of a concept increased with the amount of introspective content that it contained (Wiemer-Hastings, Krug, & Xu, 2001). These studies further found that abstract concepts typically relate introspective states to situations and events. For example, *intend* relates introspective states about goals to events in the world that follow from these states causally (intending to seek information, which then leads to asking someone a question and receiving an answer).

Because abstract concepts focus on introspective states to a large extent, they provide a window on metacognition. Similar to how people perceive the external world through vision and audition, people perceive their internal worlds through introspection. During introspection, people perceive motivations, affective states, cognitive states, and cognitive operations. Clearly, only a small subset of brain activity is perceived introspectively, but this subset supports impressive understanding and control of internal mechanisms.

According to PSS, simulators develop to represent the internal world, just as they develop to represent the external world. As people perceive the internal world, they focus attention on salient regions of it repeatedly, such that simulators develop for these regions. Thus, simulators develop for cognitive states, such as *image* and *belief*, cognitive operations such as *retrieve* and *compare*, affective states such as *happiness* and *fear*, and motivational states such as *hunger* and *ambition*. Once these simulators exist, they support symbolic operations in the introspective domain, thereby establishing metacognition. Type-token propositions result when simulators for introspective content become bound to regions of introspective activity. These categorizations then license associated inferences and support symbolic analyses of introspective activity. Predications result from mapping relevant simulators into regions of introspective activity. Novel conceptualizations of how to organize introspective processing to achieve goals result from combining relevant simulators.

EMPIRICAL EVIDENCE FOR THE PSS ACCOUNT OF SYMBOLIC OPERATIONS

The next three subsections describe preliminary evidence for the PSS accounts of predication, conceptual combination, and abstract concepts, respectively. The final subsection discusses limitations of the current evidence.

Predication

The Property Verification Task

One way to study predication is with the property verification task. In one version of this task, the word for a concept appears (e.g., "pony"), followed by the word for a property (e.g., "mane"), and the participant indicates as quickly as possible whether the property is true or false of the concept (e.g., "pony–mane" versus "pony–horns"). This task assesses predication because participants are asked whether a property predication is true of a concept (i.e., can the property *mane* be predicated of the concept *pony*?). By manipulating and measuring variables potentially related to simulation, this task can be used to assess whether simulation underlies predication.

The experiments to be reviewed next assess three specific predictions about predication that follow from the PSS account. First, if simulations represent the properties predicated of objects during property verification, then different types of properties should activate different modal systems in the brain. For example, verifying visual properties should activate visual areas, whereas verifying auditory properties should activate auditory areas. Second, when the modality of the property being verified changes, verification should take longer than when the modality remains the same. Third, perceptual aspects of the properties being verified, such as their size and shape, should affect the time to verify them.

Evidence from Brain Activation

Earlier we saw that viewing pictures of food activated the brain's gustatory system. Because gustatory processing was not relevant for the task that participants had to perform (one-back perceptual matching), it is not clear what role the gustatory activations were playing, although a plausible interpretation is that they were representing taste inferences.

Similar activations for properties were observed in the next two studies described. In these studies however, participants were asked to predicate properties as the modality of the properties varied. If modal activations represent the properties being predicated in these experiments, then as different types of properties are predicated (e.g., visual versus auditory), different corresponding brain areas should become active (e.g., visual versus auditory). If the brain represents these properties with simulations, activations across modal systems would be expected to vary in this manner.

In a positron emission tomography (PET) experiment by Kellenbach, Brett, and Patterson (2001), participants received the names of animate and inanimate objects in blocks of 24 names each. At the start of each block, participants received a property and had to decide whether it could be predicated of each object in the block or not (e.g., was each object in the block *colorful*). Across blocks, three different properties were tested. One property was whether the object is *colorful* (e.g., *banana*), not *monochrome* (e.g., *skunk*). The second property was whether the object is *loud* (e.g., *drill*), not *silent* (e.g., *pillow*). The third property was whether the object is *small* (e.g., *coin*), not *large* (e.g., *bus*). In a control condition, participants determined whether a letter string had an "X" in it or not. In half the trials, the object possessed the property; in the other half, it did not. The objects assessed for each property were controlled carefully.

Of interest was how brain activations across the block for a particular property (e.g., *colorful*) differed from activations for blocks in the control condition. Kellenbach et al. found that relevant modal activations became active for each property relative to the control. Color judgments activated color processing areas in the fusiform gyrus. Sound judgments activated auditory processing areas in the superior temporal gyrus. Size judgments activated parietal areas associated with processing space. Other brain areas were also active in each property condition, but the areas just mentioned were most relevant to the simulation hypothesis.

These findings suggest that participants simulated the property being assessed on each block (e.g., *loud*). As each name of an object in the subsequent block was read (e.g., "drill"), participants simulated the object and assessed whether it contained the simulated property (e.g., is a drill loud). If the property simulation matched a modal region of the simulated object, participants responded "true;" if not, they responded "false."

If this account is correct, then it follows that the symbolic operation of predication is grounded in simulation. To assess whether a property can be predicated of an object, participants simulate the property and then assess whether it can be found in a simulation of the object. Kellenbach et al.'s results do not show definitively that this account is correct but their results are highly consistent with it. As the property being predicated changes, so do the modal areas active for it.

Goldberg, Perfetti, and Schneider (2006) found similar results in an fMRI experiment. As in Kellenbach et al., participants received a property name prior to a block of object names and had to assess whether the property could be predicated of each object. Goldberg et al. assessed four properties: *green* (color), *loud* (sound), *soft* (touch), and *sweet* (taste). Analogous to Kellenbach et al.'s results, Goldberg et al. observed relevant modal activation for each property. Predicating the color property activated object processing regions in the temporal lobe. Predicating the sound property activated the