Evolutionary and Neuroscience Approaches to the Study of Cognition

Warren Schmaus†‡

There is a need to bring the work of philosophers such as Millikan who take an evolutionary approach to the study of the mind more in harmony with work in the cognitive neurosciences. Studies of the brain suggest that knowledge in the cortex may be organized by sensory modality-specific properties rather than by hierarchies of substance concepts, regardless of how adaptive Millikan argues the latter would be. Ryder's work on SINBAD networks, which purports to support Millikan's conclusions, does not decide this issue. The ability to make cross-modal identifications of individuals and kinds may be only a recent evolutionary adaptation.

1. Introduction. Some philosophers of mind have carefully thought about the neuroscience literature, but have given short shrift to evolutionary considerations. Others have thought about the adaptive functions of cognition, but have not paid sufficient attention to the cognitive neurosciences. Few have given careful consideration to the methodological problem of how to integrate the evolutionary and neuroscience approaches to the study of cognition.

Ruth Millikan exemplifies philosophers who favor the evolutionary approach, endeavoring to provide natural selection accounts of mental contents. In some of her more recent work, she argues that substance concepts are adaptations that allow us to accumulate useful knowledge about individuals and real kinds in nature. She says:

> [T]he task of substance concepts is to enable us to reidentify substances through diverse media and under diverse conditions, and to enable us over time to accumulate practical skills and theoretical

†To contact the author, please write to: Lewis Department of Humanities, Illinois Institute of Technology, 3301 S. Dearborn, Chicago, IL 60616; e-mail: schmaus@iit.edu.

‡In preparing the final draft of this paper, I have been in e-mail correspondence with Ruth Millikan and Dan Ryder. I would like to thank them for their assistance. Of course, all interpretations, opinions, and errors are solely my own responsibility.

Philosophy of Science, 72 (December 2005) pp. 675–686. 0031-8248/2005/7205-0003$10.00
Copyright 2005 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/508107 Published online by Cambridge University Press

> knowledge about these substances and to use what we have learned. (Millikan 2000, 2)

For Millikan, a substance concept is a knowing how to recognize something, rather than a knowing that. To think of a substance is to "be able to represent it in a stable representational system," one that allows you to recognize that substance again under varying conditions (1998a, 61). Millikan uses the term 'substance' in the Aristotelian sense, in which it can include primary substances or individuals like Mama, as well as secondary substances or kinds of individuals like mouse and kinds of stuff like milk. These latter concepts pick out real kinds in nature for Millikan. Our terms for natural kinds function like proper names, identifying rather than classifying these kinds.

Although Millikan may be a realist about the existence of the natural kinds outside the mind, this metaphysical attitude does not seem to carry over to the kinds of things she attributes to the mind. She appears to have been content with purely functional analyses of the mind and its contents and given scant attention to the underlying neurophysiology. Her substance concepts are arranged in hierarchies that resemble the data structures of artificial intelligence and cognitive science research more than anything we know from neuroscience. Two issues arise from her neglect of neuroscience: (1) whether neuroscience supports Millikan's belief that the knowledge that allows us to identify individuals and kinds is in fact organized in the brain in hierarchies of substance concepts, or whether, for instance, it is organized in accordance with sensory modality-specific properties, and (2) whether these same substance concepts that supposedly allow us to identify individuals and kinds are also involved in the perceptual and motor skills that are employed in practical tasks. Before addressing these issues, I will begin by reviewing briefly what Millikan has to say about substance concepts. Then I shall discuss the relevant neuroscience literature. Finally, as Millikan has suggested to me that her colleague Dan Ryder's work on cortical pyramidal cells may support her interpretation, I will consider whether Ryder's work does indeed help to settle the debate about the organization of knowledge in the brain.

2. Millikan on Substance Concepts. According to Millikan, substance concepts are connected to their extensions through unconscious biological functions. The extension of a substance concept is whatever substance in the world phylogenetic and ontogenetic development have given that concept the job of tracking (2000, 50). Substance concepts do not depend on language for her. She holds that we share the ability to recognize individuals and natural kinds with the higher animals, at least to the extent that animals are able to reidentify Mama, mouse, and milk (2000, 108).

https://doi.org/10.1086/508107 Published online by Cambridge University Press

That we have this ability she regards as an empirically testable hypothesis. She says, "[E]xperiments need to be designed and interpreted with it in mind that the cognitive systems are designed by evolution and tuned by experience to find real-world substances, not random logically possible ones" (2000, 83; 1998a, 63).

Millikan explains that natural selection does not provide us with the ability to identify any particular substance, but rather the capacity to acquire this ability through interacting with the environment (2000, 51). Natural selection "selected for mechanisms that become tuned through interaction with the environment to do things of useful kinds" (2000, 63). These learning mechanisms may include "trial-and-error learning, learning by association, by imitation, by figuring something out, and so forth" (2000, 63).

For Millikan, to have the concept of a substance is to have the ability to know which of its properties can be generalized to new encounters and what sorts of questions can be asked about it. But to know what sort of questions we can ask of something, we need to know what sort of substance it is: is it a person, an animal species, a chemical element or compound, a physical object, or something else? The knowledge on which this ability depends is organized into what she calls 'substance templates' (2000, 9–10). Substance templates tell us which properties of objects we can expect to remain stable through changes in circumstances, such as color with lighting (2000, 106–108). Substance templates like animal and mammal species are also substance concepts in their own right. These substance templates can form hierarchies, such as animal, mammal, and human being. However, for Millikan these hierarchies do not belong to a single hierarchy with something like 'substance' or 'being' as the highest genus. Although the concept of a human being may belong to a nested hierarchy of kinds of animals, any one particular human individual can belong to more than one category, such as professor, parent, and neighbor, that do not all belong to the same hierarchy (2000, 29–31).

Although she believes that in general these substance templates are discovered or learned through experience, there are some, "such as physical kinds, animal kinds, plant kinds, artifact kinds, social kinds and so forth," that may be 'endogenous' (2000, 82) or 'wired in' (2000, 30). Whether such templates are built in is of course an empirical question. However, the psychological literature she cites provides only indirect evidence for the claim that these substance templates are hard-wired or endogenous by showing how this hypothesis can explain the performance of subjects on cognitive tasks.[1] The cognitive and developmental psycho-

1. Millikan (2000, 30, 82) cites Atran 1989; Boyer 1998; Carey 1985; Gallistel et al.

https://doi.org/10.1086/508107 Published online by Cambridge University Press

logical literature she cites could itself benefit from closer ties to the neurosciences.

Millikan's substance templates bear a prima facie resemblance to data structures such as Minsky's (1975) frames and Schank's scripts (1975), which were widely used in artificial intelligence in the 1970s to model such cognitive tasks as understanding natural language. Just as a template tells us what sorts of questions can or cannot be asked about something, a frame will specify certain elements as fixed, leaving open certain 'terminals' where other characteristics are to be filled in differently for different kinds of things. Also, frames and templates are arranged hierarchically. Minsky's frames include events as well as kinds of things. Similarly, Millikan (1998b, 92) is willing to include kinds of events such as siesta time or hearing Beethoven's Fifth Symphony as having definite properties that remain invariant across repeated encounters in the same way that substances do. Among the many problems with frames and scripts there is at least one that Millikan's templates appear to share. Consider the way Schank's scripts were meant to allow AI systems to provide intelligent answers to questions about stories: in order to understand the stories, they need to apply scripts. But how do they know which scripts to apply unless they already understand the story? Similarly, how does a child know whether to apply a template for an individual, a woman, or a member of the human species when it sees Mama? Millikan assures us that these are all separate concepts (1998b, 97), but what or who is keeping them separate?

One important difference between Millikan's templates and the data structures of AI is that the substance concepts to which they are linked are not just descriptions or lists of properties but are supposed to include a variety of perceptual and even motor skills. In particular, she holds that these concepts involve the skill of keeping track of something, at first perceptually and then conceptually. Citing the work of Gopnik and Meltzoff (1996), Millikan says that the perceptual mechanisms by which infants track things involve the motion, location, and trajectory, but not the properties of the tracked object (2000, 77). It is perhaps for this reason that, at least according to Xu and Carey (1996), 10-month-old infants are not surprised when one kind of object turns into another, but are surprised if one object turns into two. This changes by 12 months. That one thing does not turn into another thing appears to be something that infants have to learn. Perceptual tracking abilities make it possible for the infant to observe the object long enough to observe which sorts of properties remain the same and which do not (Millikan 2000, 77).

1993; Gelman and Coley 1991; Keil 1979, 1989; Markman 1989; Marler 1993; Spelke 1989, 1993.

https://doi.org/10.1086/508107 Published online by Cambridge University Press

Perceptual tracking then makes possible what Millikan calls 'conceptual tracking', which is what comes into play when we keep track of something through its properties over discontinuities of space and time. Conceptual tracking also depends on additional perceptual mechanisms, such as those involved in perceiving color constancy, recognizing the same sizes and shapes at various angles and distances, and recognizing faces (2000, 78–79):

> The general method used for reidentification of every substance is the same. One relies on certain expected, that is, projected, continuities over times or occasions of encounter, for example, continuity in spatial projectory, in color, in shape, in odor, in general arrangement of parts, in manner of motion. Or one may rely on continuity in identity of parts or other associated features. Perhaps it is the shape of your face, or just your eyes that I recognize immediately as indicative of you, or your voice, or your signature, or your walk, or your humor, or your name . . . . (Millikan 2000, 155–156)

According to Millikan, motor skills are also involved in our tracking abilities: "My ability to track with my feet as well as my eyes and head may be involved here, hence my ability to walk over rough ground or to avoid slipping on ice" (2000, 156). For Millikan, these ways of reidentifying cannot be divided nonarbitrarily into discrete ways of identifying. Rather, identifying or keeping track of something involves a complex, integrated set of skills that reinforce each other (2000, 156).

The integration of all these skills would seem to require some sort of representational system where information from all the sensory modalities could be brought together. For instance, for a human individual, this could include information from a walk recognizer, face recognizer, voice recognizer, and a color constancy detector. Since these substance concepts are supposed to play a role in practical skills as well as theoretical knowledge (Millikan 2000, 2, quoted above), this system for representing these concepts must then be connected with our abilities to coordinate and guide bodily motions.

3. Neuroscience. Recent work in the neurosciences, however, calls into question Millikan's assumption that our substance concepts are tied to practical skills. Melvyn Goodale and his associates present pathological cases that suggest that the visual pathway involved in our ability to identify objects is not the same as that involved in reaching out and grasping them. Patients with damage to the posterior parietal cortex display optic ataxia. That is, although they have no trouble recognizing and describing objects, they are unable to reach out and grasp them when they are in the field contralateral to the brain lesion. On the other hand, patients

https://doi.org/10.1086/508107 Published online by Cambridge University Press

with damage to the occipito-temporal region display visual form agnosia. They cannot recognize or describe common objects, faces, or drawings, but have no trouble using vision to move about in the world in a way that avoids obstacles (Goodale 2000, 368–370; Goodale, Jakobson, and Servos 1996, 108*ff.*). Neuroscientists such as Goodale thus distinguish the ventral visual pathway, which is involved in our ability to identify objects, from the dorsal visual pathway, which guides action. The task of the ventral visual system, as Goodale, Jakobson, and Servos describe it, is closely related to what Millikan describes as the function of substance concepts. It is

> to identify objects and their relations, classify those objects and relations, and attach meaning and significance to them. Such operations are essential for engaging in social interactions, exchanging information with others, accumulating a knowledge base about the world, and choosing among different courses of action. (Goodale, Jakobson, and Servos 1996, 114)

Although the ventral system may be involved in selecting goals and the means to achieve them, the execution and control of action is under the direction of the dorsal system.

Goodale (2000, 373–374) also reports that the dorsal or visuomotor system is not subject to the same illusions as the ventral perceptual system involved in identifying objects. Studies of patients with blindsight corroborate what Goodale says about the functions of the two visual pathways. These patients have suffered damage to the primary visual cortex, but the visual pathway that projects to the motor cortex seems to remain intact. Although patients with blindsight have no conscious perceptual experience of objects, they are nevertheless able to adjust their hands to the shape, size, orientation, and location of objects when forced to grasp them. However, they never spontaneously or voluntarily reach out to grasp things, such as a glass of water when they are thirsty (Güzeldere, Flanagan, and Hardcastle 2000, 1280). This again implies that the function of planning and initiating an action is neurologically separate from that of executing it.

In her most recent work, Millikan recognizes this division in the visual system and a similar division in the auditory system (Millikan 2006, Chapter 14). But even if she were now to retract the view that substance concepts guide the execution of practical tasks, this still leaves the question whether the knowledge by which we identify things is organized in the brain by hierarchies of substance concepts or by sensory modality-specific properties. Here the evidence from the neurosciences is less clear-cut.

Alfonso Caramazza provides neuropathological evidence that would seem to support Millikan. He reports that the abilities to identify cate-

https://doi.org/10.1086/508107 Published online by Cambridge University Press

gories of animals, fruits and vegetables, artifacts, and body parts can be damaged separately (2000, 1037). As he points out, these are precisely the sorts of objects for which the ability to pick out would have adaptive significance:

> [E]volutionary pressures resulted in specialized neural mechanisms for perceptually and conceptually distinguishing specific kinds of objects, leading to the categorical organization of knowledge in the brain. (Caramazza 2000, 1043)

Curiously, he also reports on one patient who was unable to name and categorize fruits and vegetables from pictures, touch, or definitions, but who, when he heard the names, could point to the correct pictures and name their properties, and when he read their names, could categorize them correctly (Hart, Berndt, and Caramazza 1985). However, he does not consider cases like this to provide evidence that our concepts are organized by sensory modality rather than in a unified system (Caramazza 2000, 1038).

Alex Martin and his associates disagree with Caramazza, arguing that the cerebral cortex is organized in accordance with properties such as shape, color, and motion, rather than in accordance with kinds of objects such as animals or tools. Using neural imaging techniques such as positron emission tomography and functional magnetic resonance imaging that allow the researcher to see what part of the brain is active during various cognitive tasks, they have found that information about these properties is stored in different parts of the brain, near the sensory and motor parts of the brain that were active when the information was acquired (Martin, Ungeleider, and Haxby 2000, 1023). In one experiment a subject is shown a line drawing of a wagon and asked what it is (wagon), what color it usually is (red), and what action is usually associated with it (pulling). A different part of the brain is activated for each answer (2000, 1024–1025). However, they also suggest that the part of the cortex for perceiving motion may have subregions for animals and tools (2000, 1031–1032).

In reply to this experiment, Caramazza (2000, 1042) argues that the perceptual content of our thoughts may be distributed in the brain differently than the conceptual content. However, he does not make this distinction clear. Caramazza (2000, 1037) also insists that modality-specific systems cannot account for category deficits, unless we assume, say, that the visual system may be more important for recognizing animals, while the system for understanding functions is more important for recognizing artifacts. But this assumption seems entirely reasonable. Surely animals of the same kind bear a greater visual similarity to each other than do tools of the same kind. Finally, Caramazza (2000, 1038–1039) also cites evidence of patients who have trouble identifying just animals

https://doi.org/10.1086/508107 Published online by Cambridge University Press

and not fruits and vegetables. However, this may be evidence that just one part of the visual system is damaged, the part involved in recognizing animals but not fruit and vegetables.

Hence, neuroscience evidence is at least consistent with the view that the ways in which we recognize individuals and kinds may also be organized by sensory modality. That is, instead of storing all the information we have about an individual or a kind in a specific region of the brain, this information may be distributed all over the brain. Perceptual systems such as face-recognizers and voice-recognizers may each be divided into subregions for the faces or voices of the individuals that we recognize. This would explain the common experience of recognizing a face but not being able to recall the name or anything about where you have seen this person before.

It may seem irrational or suboptimal for representations of the various properties of the same individual to be distributed all over the cortex rather than all held in one place. However, there is evidence that even different visual properties are handled by different areas in the brain. Jon Kaas (1989, 467) has argued that it actually makes more sense for the brain to have multiple visual areas:

> More specifically, the local connections that would be needed to sharpen the tuning of individual neurons for one of the many attributes of a complex stimulus would be long if neurons for all the different attributes were in the same visual area. Having many retinotopic maps allows most of the computing to be done with short local connections, although the receiving and sending information would requires long connections between areas and subcortical structures. (Kaas 1989, 467)

Having multiple visual areas not only simplifies the problem of interconnecting functionally related groups of neurons, but also makes it easier for the genetic programming to guide brain development (Kaas 1989, 467). These advantages for having multiple visual areas should apply with even greater force when it comes to the question of having separate cortical areas for different sensory modalities. Alternatively, a brain organized by substance concepts would require having neurons with unrelated functions in the same area. There would have to be either long connections to functionally related neurons elsewhere in the brain, or massive functional redundancy, with assemblies of neurons performing similar functions at each area dedicated to a particular substance.

A problem for the view that the cortex is organized by sensory properties rather than substance concepts is to explain our cross-modal ability to identify individuals and kinds of things. Our ability to make cross-modal identifications raises an issue for the cognitive neuroscience account ac-

https://doi.org/10.1086/508107 Published online by Cambridge University Press

cording to which the cortex is organized by sensory modality: exactly how, for instance, are the face-recognizer and a voice-recognizer connected, so as to make it possible for us to know that the people we see are the same as those we hear?

4. The SINBAD Account. Dan Ryder thinks he has found neuroscience evidence that supports Millikan's philosophy. On his account, pyramidal cells in the cortex may be 'tuned' or 'correspond' to 'sources of correlation' among properties in the environment. According to Ryder,

> Millikan's (1998a, 1998b, 2000) 'substances' are very similar to what I am calling sources of correlation. . . . Her account of substance concepts is, to my knowledge, the closest cousin to my neurosemantics in the current literature. In fact, I am greatly encouraged by the partial resemblance, especially since the convergence arises from opposite directions: she comes from abstract psychological considerations, while I come from the biology. This sort of concordance between top-down and bottom-up is surely a sign of truth. (Ryder 2004, 235)

The tuning of the pyramidal cells takes place through the back-propagation of output to the principal dendrites, which then correct their input to cell bodies, such that if there are N dendrites, each contributes $1/N$ of the activity of the cell's output. It is from this tuning process that Ryder takes the name 'SINBAD', which stands for "Set of INteracting BAck-propagating Dendrites" (2004, 212). In one of his examples, he says that a pyramidal cell for 'Banana' may have one dendrite picking up an input from a neuron that spikes for 'green-or-yellow', another for 'banana-shaped', another for 'peelable', and yet another for 'coming in bunches' (2004, 219). He also suggests that there may be a pyramidal cell tuned to the natural kind 'tiger'.

Ryder's pyramidal cells sound suspiciously like the elusive grandmother cell. In his correspondence with me, he says that single cells may be highly equivocal, tuning to more than one source of correlation, and that normally it will be populations of cells that explain how we pick out individuals and kinds. He also tells me that there seems to be circumstantial evidence for cells that are sensitive to information that is fairly remote from sensory input, such as Bill Clinton cells, at least in the hippocampus.

But whether it is one cell, a population of cells, or a network of interacting cells that are tuned to sources of correlation is not the issue that divides us. The issue is more methodological. Ryder is in effect arguing that pyramidal cells have the right sort of structure and function to be tuned to kinds and individuals. However, even if this were true, this would not show that in fact the pyramidal cells are tuned to substances. Even

https://doi.org/10.1086/508107 Published online by Cambridge University Press

the computer model he has constructed with Favorov of an artificial network of 15 pyramidal cells that is tuned to kitchen sinks (Favorov and Ryder 2004) shows only that such a network *could* do this. We need evidence other than from computer modeling that there are pyramidal cells tuned to individuals and kinds. Such evidence might consist of someone who had lost all memories of a particular person, although his face-recognizer, voice-recognizer, and other sensory modality-specific recognition mechanisms still seemed to be otherwise intact. Or it might consist of someone who remembers that animals that appear to be dogs bark, but has forgotten that those that appear to be cats meow.

In another more recent paper, Ryder (2006) says merely that kinds belong to the more general class of 'sources of correlation', thus leaving it open that pyramidal cells could be tuned to things other than substances, such as just faces or voices. It would seem that for most practical purposes, correlations among properties are all that is needed. Substance concepts may be unnecessary.

In sum, Ryder's work with SINBAD networks does not decide the question as to whether the cortex is organized by sensory modality or by concepts of objects. It remains plausible that representations of our concepts of individuals or kinds may be distributed over the entire cortex. In correspondence, Ryder agrees that he has found no evidence to decide the issue between Caramazza and Martin. Rather, his account of SINBAD networks should be regarded only as an hypothesis that, if true, would support Caramazza's view that the cortex is organized by concepts of objects.

5. Discussion. One might wish to defend Millikan by arguing that she is characterizing the mind at an abstract functional level and that it is not really relevant to her concerns whether substance concepts are localized or distributed over the cortex. Knowledge in the cortex could be organized by sensory modality-specific properties, as long as there are connections that allow one to recognize the same individuals or kinds across sensory modalities. For Millikan, these would be simply engineering details.

This defense raises the methodological issue of how to link evolutionary theory with the neurosciences. It assumes that we can reason from what would be adaptive to how the mind is functionally organized. How the brain actually carries out these functions is then left to the neuroscientists. The danger with this approach is that natural selection typically does not result in what from our point of view would seem to be the optimal design. Natural selection jury-rigs from the parts ready to hand whatever it can to get the job done. The notion that the brain is optimally organized to adapt us to our environment belongs more to natural religion than the theory of natural selection.

https://doi.org/10.1086/508107 Published online by Cambridge University Press

An alternative method would be to look at how the brain is organized, and then inquire into the evolutionary pathways that may have led to this sort of organization, and how they may have been selected for at each step. We may find, for instance, that organization by sensory modality is evolutionarily older, and that connections across sensory modalities that allow cross-modal identification of individuals and kinds evolved later. In sum, neuroscience should guide evolutionary theorizing as well as provide ways of testing these theories, and not just seek the mechanisms that could carry out the functions that evolutionary theory decides are necessary.

REFERENCES

Atran, Scott (1989), "Basic Conceptual Domains", *Mind & Language* 4: 7–16.

Boyer, Pascal (1998), "If 'Tracking' Is Category-Specific a 'Common Structure' May Be Redundant", *Behavioral and Brain Sciences* 21: 67–68.

Caramazza, Alfonso (2000), "The Organization of Conceptual Knowledge in the Brain", in Michael S. Gazzaniga (ed.), *The New Cognitive Neurosciences*, 2nd ed. Cambridge, MA: MIT Press, 1037–1046.

Carey, Susan (1985), *Conceptual Change in Childhood*. Cambridge, MA: MIT Press.

Favorov, Oleg, and Dan Ryder (2004), "SINBAD: A Neocortical Mechanism for Discovering Environmental Variables and Regularities Hidden in Sensory Input", *Biological Cybernetics* 90: 191–202.

Gallistel, Charles, A. Brown, Susan Carey, R. Gelman, and Frank Keil (1993), "Lessons from Animal Learning for the Study of Cognitive Development", in Susan Carey and Rochel Gelman (eds.), *The Epigenesis of Mind: Essays on Biology and Cognition*. Hillsdale, NJ: Erlbaum, 3–36.

Gelman, Susan A., and J. D. Coley (1991), "Language and Categorization: The Acquisition of Natural Kind Terms", in Susan A. Gelman and James P. Byrnes (eds.), *Perspectives on Language and Thought*. Cambridge: Cambridge University Press, 146–196.

Goodale, Melvyn A. (2000), "Perception and Action in the Human Visual System", in Michael Gazzaniga (ed.), *The New Cognitive Neurosciences*, 2nd ed. Cambridge, MA: MIT Press, 365–377.

Goodale, Melvyn A., Lorna S. Jakobson, and Philip Servos (1996), "The Visual Pathways Mediating Perception and Prehension", in Michael Gazzaniga (ed.), *Cognitive Neuroscience: A Reader.* Malden, MA: Blackwell Publishers, 106–123.

Gopnik, Alison, and Andrew Meltzoff (1996), *Words, Thoughts, and Theories*. Cambridge, MA: MIT Press.

Güzeldere, Güven, Owen Flanagan, and Valerie Gray Hardcastle (2000), "The Nature and Function of Consciousness: Lessons from Blindsight", in Michael Gazzaniga (ed.), *The New Cognitive Neurosciences*, 2nd ed. Cambridge, MA: MIT Press, 1277–1284.

Hart, John, Jr., Rita Sloan Berndt, and Alfonso Caramazza (1985), "Category-Specific Naming Deficit Following Cerebral Infarction", in Michael Gazzaniga (ed.), *Cognitive Neuroscience: A Reader*. Malden, MA: Blackwell Publishers, 406–410.

Kaas, Jon H. (1989), "Why Does the Brain Have So Many Visual Areas?", in Michael Gazzaniga (ed.), *Cognitive Neuroscience: A Reader*. Malden, MA: Blackwell Publishers, 448–472.

Keil, Frank C. (1979), *Semantic and Conceptual Development: An Ontological Perspective*. Cambridge, MA: Harvard University Press.

——— (1989), *Concepts, Kinds, and Cognitive Development*. Cambridge, MA: MIT Press.

Markman, Ellen M. (1989), *Categorization and Naming in Children*. Cambridge, MA: MIT Press.

Marler, Peter (1993), "The Instinct to Learn", in Susan Carey and Rochel Gelman (eds.),

https://doi.org/10.1086/508107 Published online by Cambridge University Press

The Epigenesis of Mind: Essays on Biology and Cognition. Hillsdale, NJ: Erlbaum, 37–66.

Martin, Alex, Leslie G. Ungeleider, and James V. Haxby (2000), "Category Specificity and the Brain: The Sensory/Motor Model of Semantic Representation of Objects", in Michael Gazzaniga (ed.), *The New Cognitive Neurosciences*, 2nd ed. Cambridge, MA: MIT Press, 1023–1036.

Millikan, Ruth Garrett (1998a), "A Common Structure for Concepts of Individuals, Stuffs, and Real Kinds: More Mama, More Milk, and More Mouse", *Behavioral and Brain Sciences* 21: 55–65.

——— (1998b), "Words, Concepts, and Entities: With Enemies Like These, I Don't Need Friends", *Behavioral and Brain Sciences* 21: 89–100.

——— (2000), *On Clear and Confused Ideas*. Cambridge: Cambridge University Press.

——— (2006), *Varieties of Meaning*. Cambridge, MA: MIT Press.

Minsky, Marvin (1975), "A Framework for Representing Knowledge", in Patrick Henry Winston (ed.), *The Psychology of Computer Vision*. New York: McGraw-Hill, 211–277.

Ryder, Dan (2004), "SINBAD Neurosemantics: A Theory of Mental Representation", *Mind & Language* 19: 211–240.

——— (2006), "On Thinking of Kinds: A Neuroscientific Perspective", in Graham Macdonald and David Papineau (eds.), *Teleosemantics*. Oxford: Oxford University Press.

Schank, Roger C. (1975), "The Primitive Acts of Conceptual Dependency", Theoretical Issues in Natural Language Processing, workshop in Cambridge, MA, June 10–13.

Spelke, Elizabeth S. (1989), "The Origins of Physical Knowledge", in Lawrence Weiskrantz (ed.), *Thought without Language*. Oxford: Oxford University Press, 185–210.

——— (1993), "Physical Knowledge in Infancy", in Susan Carey and Rochel Gelman (eds.), *The Epigenesis of Mind: Essays on Biology and Cognition*. Hillsdale, NJ: Erlbaum, 133–170.

Xu, F., and Susan Carey (1996), "Infant Metaphysics: The Case of Numerical Identity", *Cognitive Psychology* 30: 111–153.

https://doi.org/10.1086/508107 Published online by Cambridge University Press