Luce, P. A., D. B. Pisoni, and S. D. Goldinger. (1990). Similarity neighborhoods of spoken words. In G. T. M. Altmann, Ed., *Cognitive Models of Speech Processing: Psycholinguistic and Computational Perspectives.* Cambridge, MA: MIT Press, pp. 122–147.

MacWhinney, B., and J. Leinbach. (1991). Implementations are not conceptualizations: Revising the verb learning model. *Cognition* 40: 121–157.

Marcus, M. P. (1980). *A Theory of Syntactic Recognition for Natural Language.* Cambridge, MA: MIT Press.

Marslen-Wilson, W. M., and L. K. Tyler. (Forthcoming). Dissociating types of mental computation. *Nature*.

Massaro, D. W. (1989). Testing between the TRACE model and the Fuzzy Logical Model of Perception. *Cognitive Psychology* 21: 398–421.

McClelland, J. L. (1991). Stochastic interactive processes and the effects of context on perception. *Cognitive Psychology* 23: 1–44.

McClelland, J. L., and J. L. Elman. (1986). The TRACE model of speech perception. *Cognitive Psychology* 18: 1–86.

McClelland, J. L., and D. E. Rumelhart. (1981). An interactive activation model of context effects in letter perception: 1. An account of the basic findings. *Psychological Review* 88: 375–407.

McRoy, S. W., and G. Hirst. (1990). Race-based parsing and syntactic disambiguation. *Cognitive Science* 14: 313–353.

Norris, D. (1994a). Shortlist: A connectionist model of continuous speech recognition. *Cognition* 52: 189–234.

Norris, D. (1994b). A quantitative multiple-levels model of reading aloud. *Journal of Experimental Psychology: Human Perception and Performance* 20(6): 1212–1232.

Norris, D., J. M. McQueen, and A. Cutler. (1995). Competition and segmentation in spoken word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 21: 1209–1228.

Oden, G. C., and D. W. Massaro. (1978). Integration of featural information in speech perception. *Psychological Review* 1985: 172–191.

Pinker, S., and A. Prince. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition* 28: 73–193.

Plaut, D. C., J. L. McClelland, M. Seidenberg, and K. E. Patterson. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review* 103: 56–115.

Plunkett, K., and V. Marchman. (1991). U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition* 38: 43–102.

Plunkett, K., and V. Marchman. (1993). From rote learning to system building: Acquiring verb morphology in children and connectionist nets. *Cognition* 48: 21–69.

Roelofs, A. (1992). A spreading-activation theory of lemma retrieval in speaking. *Cognition* 42: 107–142.

Roelofs, A. (1996). Serial order in planning the production of successive morphemes of a word. *Journal of Memory and Language* 35: 854–876.

Rumelhart, D. E., and J. L. McClelland. (1986). On learning the past tense of English verbs. In McClelland and Rumelhart, Eds., *Parallel Distributed Processing: Explorations in the Microstructure of Cognition.* Cambridge, MA: MIT Press.

Seidenberg, M., and J. L. McClelland. (1989). A distributed developmental model of word recognition and naming. *Psychological Review* 96: 523–568.

Sharkey, N. E. (1990). A connectionist model of text comprehension. In D.A. Balota and G. B. Flores d'Arcais, Eds., *Comprehension Processes in Reading.* Hillsdale, NJ: Erlbaum, pp. 487–514.

Computational Theory of Mind

The computational theory of mind (CTM) holds that the mind is a digital computer: a discrete-state device that stores symbolic representations and manipulates them according to syntactic rules; that thoughts are mental representations—more specifically, symbolic representations in a LANGUAGE OF THOUGHT; and that mental processes are causal sequences driven by the syntactic, but not the semantic, properties of the symbols. Putnam (1975) was perhaps the first to articulate CTM, but it has found many proponents, the most influential being Fodor (1975, 1981, 1987, 1990, 1993) and Pylyshyn (1980, 1984).

CTM's proponents view the theory as an extension of the much older idea that thought is MENTAL REPRESENTATION—an extension that shows us how a commitment to mental states can be compatible with a causal account of mental processes and with a commitment to materialism and the generality of physics. Older breeds of representationalism were unable to explain how mental processes could be semantically coherent—how thoughts could follow one another in a fashion appropriate to their meanings, while also being bona fide causal processes that did not depend on an inner homunculus who understood the meanings of the representations. Using formalization and digital computers, however, we can explain how this occurs. Formalization shows us how to link semantics to syntax. For any formalizable symbol system, it is possible to develop a set of formal derivation rules, based wholly on syntactic properties, that license all and only the inferences permissible on semantic grounds. Computers show us how to link syntax to causation. For any finite formal system, it is possible to construct a digital computer that automates the derivations of that system. Thus, together, formalization and computation show us how to link semantics to causation in a material system like a digital computer: design a set of syntactic rules that "track" the semantic properties of the symbols (i.e., formalize the system), and then implement those rules in a computer. Because digital computers are purely physical systems, this shows us it is possible for a purely physical system to carry out symbolic inferences that respect the semantics of the symbols without recourse to a homunculus or to any other nonphysical agency. Syntactic properties are the causal determinants of reasoning, syntax tracks semantics, and syntactic properties can be implemented in a physical system.

CTM has been touted both for its connections to successful empirical research in cognitive science and for its promise in resolving philosophical problems. The main argument in favor of the language of thought hypothesis and CTM has been the "only game in town" argument: cognitive theories of language, learning, and other psychological phenomena are the only viable theories we possess, and these theories presuppose an inner representational system. Therefore we have a prima facie commitment to the existence of such a representational system (Fodor 1975). Some have claimed that CTM also explains the INTENTIONALITY of mental states and that it reconciles mentalism with materialism. The meanings and intentionality of mental states are "inherited from" the meanings and intentionality of the "mentalese" symbols (Fodor 1981). And because symbols, the ultimate bearers of semantic properties and intentionality, can both have meaning and be physical objects, there is not even a prima facie conflict between a commitment to semantics and intentionality and a commitment to materialism. Finally, CTM has been held to explain the generative and creative powers of thought that result from the COMPOSITIONALITY of the language of thought. Chomskian linguistics shows us how an infinite number of possible sentences can be generated out of a finite number of atomic lexical units, syntactic structures, and transformation rules. If the basis of thought is a symbolic language, these same resources can be applied directly to explain the compositionality of thought.

Although CTM gained a great deal of currency in the late 1970s and 1980s, it has since been criticized on a number of fronts. First, with philosophers' rediscovery in the late 1980s of alternative approaches to psychological modeling, represented in NEURAL NETWORKS and dynamic adaptive systems, the empirical premise of the "only game in town" argument has been brought into question. Indeed, the main thrust of philosophical debate about neural networks and connectionism has been over whether their models of psychological phenomena are viable alternatives to rule-and-representation models.

Second, writers such as Dreyfus (1972, 1992) and Winograd and Flores (1986) have claimed that much human thought and behavior cannot be reduced to explicit rules, and hence cannot be formalized or reduced to a computer program. Thus, even if CTM does say something significant about the parts of human cognition that can be formalized, there are large portions of human mental life about which it can say nothing. Dreyfus and others have attempted to argue that this includes all expert knowledge and such simple skills as knowing how to drive a car or order in a restaurant.

A third line of criticism has been directed at CTM's use of symbolic meaning to explain the semantics of thought, on the grounds that symbolic meaning is derivative from the intentionality of thought, either causally (Searle 1980; Haugeland 1978; Sayre 1986) or conceptually (Horst 1996). Thus the attempt to explain intentionality by appeal to symbols is circular and regressive. Searle (1990) and Horst (1996) have taken this line of argument even further, claiming that the "representations" in computers are not even symbolic or syntactic in their own right, but possess these properties by virtue of the intentions and conventions of computer users: a digital machine not connected to our interpretive practices has a "syntax" only in a metaphorical sense of that word. Horst's version of these criticisms also yields an argument against the claim to reconcile mentalism with materialism: what digital computers show us how to do is to link convention-laden symbolic meaning with CAUSATION by way of convention-laden syntax, not to link the sense of "meaning" attributed to mental states with causation.

A fourth line of criticism has come from advocates of externalist theories of meaning. For many years, advocates of CTM tended also to be advocates of a "methodological solipsism" (Fodor 1980) or INDIVIDUALISM who held that the typing of mental states needed to be insensitive to features outside of the cognizer because the computational processes that determined thought have access only to mental representations. At the same time, CTM required that the typing of mental states reflect their semantic properties. These two commitments together seemed to be incompatible with externalist theories of content, which hold that the meanings of many terms are at least partially determined by factors that lie outside of the cognizer, such as its physical (Putnam 1975) and linguistic (Burge 1979, 1986) environment. This was used by some externalists (e.g., Baker 1987) as an argument against computationalism, and was used at least at one time by Fodor (1980) as a reason to reject externalism. Nevertheless, at least some computationalists, including Fodor (1993), have now embraced strategies for reconciling computational theories of mental processes with externalist theories of meaning for mental representations.

See also CHINESE ROOM ARGUMENT; COMPUTATION AND THE BRAIN; CONNECTIONISM, PHILOSOPHICAL ISSUES; FUNCTIONALISM; NARROW CONTENT; RULES AND REPRESENTATIONS

—Steven Horst

References

Baker, L. R. (1987). *Saving Belief: A Critique of Physicalism.* Princeton: Princeton University Press.

Burge, T. (1979). Individualism and the mental. In P. French, T. Euhling, and H. Wettstein, Eds., *Studies in Epistemology,* Midwest Studies in Philosophy, vol. 4. Minneapolis: University of Minnesota Press.

Burge, T. (1986). Individualism and psychology. *Philosophical Review* 95(1): 3–45.

Dreyfus, H. (1972). *What Computers Can't Do.* New York: Harper and Row.

Dreyfus, H. (1992). *What Computers Still Can't Do.* Cambridge, MA: MIT Press.

Fodor, J. (1975). *The Language of Thought.* New York: Crowell.

Fodor, J. (1980). Methodological solipsism considered as a research strategy in cognitive science. *Behavioral and Brain Sciences* 3: 63–73.

Fodor, J. (1981). *Representations.* Cambridge, MA: MIT Press.

Fodor, J. (1987). *Psychosemantics.* Cambridge, MA: MIT Press.

Fodor, J. (1990). *A Theory of Content and Other Essays.* Cambridge, MA: MIT Press.

Fodor, J. (1993). *The Elm and the Expert.* Cambridge, MA: MIT Press.

Haugeland, J. (1978). The nature and plausibility of cognitivism. *Behavioral and Brain Sciences* 2: 215–226.

Haugeland, J., Ed. (1981). *Mind Design.* Cambridge, MA: MIT Press.

Horst, S. (1996). *Symbols, Computation and Intentionality: A Critique of the Computational Theory of Mind.* Berkeley and Los Angeles: University of California Press.

Putnam, H. (1975). The Meaning of "Meaning." In K. Gunderson, Ed., *Language, Mind and Knowledge.* Minnesota Studies in the Philosophy of Science, vol. 7. Minneapolis: University of Minnesota Press.

Pylyshyn, Z. (1980). Computation and cognition: Issues in the foundation of cognitive science. *Behavioral and Brain Sciences* 3: 111–132.

Pylyshyn, Z. (1984). *Computation and Cognition: Toward a Foundation for Cognitive Science.* Cambridge, MA: MIT Press.

Sayre, K. (1986). Intentionality and information processing: An alternative model for cognitive science. *Behavioral and Brain Sciences* 9(1): 121–138.

Searle, J. (1980). Minds, brains and programs. *Behavioral and Brain Sciences* 3: 417–424.

Searle, J. (1984). Minds, brains and science. Cambridge, MA: Harvard University Press.

Searle, J. (1990). Presidential address. *Proceedings of the American Philosophical Association.*

Searle, J. (1992). *The Rediscovery of the Mind.* Cambridge, MA: MIT Press.

Winograd, T., and F. Flores. (1986). *Understanding Computers and Cognition.* Norwood, NJ: Ablex.

Further Readings

Cummins, R. (1989). *Meaning and Mental Representation.* Cambridge, MA: MIT Press.

Garfield, J. (1988). *Belief in Psychology: A Study in the Ontology of Mind.* Cambridge, MA: MIT Press.

Newell, A., and H. Simon. (1975). Computer science as empirical inquiry. (1975 Turing Lecture.) Reprinted in J. Haugeland, Ed., *Mind Design.* Cambridge, MA: MIT Press, 1981, pp. 35–66.

Putnam, H. (1960). Minds and machines. In S. Hook, Ed., *Dimensions of Mind.* New York: New York University Press, pp. 138–164.

Putnam, H. (1961). Brains and behavior. Reprinted in Ned Block, Ed., *Readings in Philosophy of Psychology.* Cambridge, MA: Harvard University Press, 1980, pp. 24–36.

Putnam, H. (1967). The nature of mental states. In W. H. Capitan and D. D. Merrill, Eds., *Art, Mind and Religion.* Pittsburgh: University of Pittsburgh Press. Reprinted in Ned Block, Ed., *Readings in Philosophy of Psychology.* Cambridge, MA: Harvard University Press, 1980, pp. 223–231.

Rumelhart, D. E., J. McClelland, and the PDP Research Group. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition.* Cambridge, MA: MIT Press.

Sayre, K. (1987). Cognitive science and the problem of semantic content. *Synthèse* 70: 247–269.

Smolensky, P. (1988). The proper treatment of Connectionism. *Behavioral and Brain Sciences* 11(1): 1–74.

Computational Vision

The analysis of a visual image yields a rich understanding of what is in the world, where objects are located, and how they are changing with time, allowing a biological or machine system to recognize and manipulate objects and to interact physically with its environment. The computational approach to the study of vision explores the information-processing mechanisms needed to extract this important information. The integration of a computational perspective with experimental studies of biological vision systems from psychology and neuroscience can ultimately yield a more complete functional understanding of the neural mechanisms underlying visual processing.

Vision begins with a large array of measurements of the light reflected from object surfaces onto the eye. Analysis then proceeds in multiple stages, with each producing increasingly more useful representations of information in the scene. Computational studies suggest three primary representational stages. *Early representations* may capture information such as the location, contrast, and sharpness of significant intensity changes or edges in the image. Such changes correspond to physical features such as object boundaries, texture contours, and markings on object surfaces, shadow boundaries, and highlights. In the case of a dynamically changing scene, the early representations may also describe the direction and speed of movement of image intensity changes. *Intermediate representations* describe information about the three-dimensional (3-D) shape of object surfaces from the perspective of the viewer, such as the orientation of small surface regions or the distance to surface points from the eye. Such representations may also describe the motion of surface features in three dimensions. Visual processing may then proceed to higher-level representations of objects that describe their 3-D shape, form, and orientation relative to a coordinate frame based on the objects or on a fixed location in the world. Tasks such as object recognition, object manipulation, and navigation may operate from the intermediate or higher-level representations of the 3-D layout of objects in the world. (See also MACHINE VISION for a discussion of representations for visual processing.)

Models for computing the early representations of intensity edges typically begin by filtering the image with filters that smooth and differentiate the image intensities. Smoothing at multiple spatial scales allows the simultaneous representation of the gross structure of image contours, while preserving the fine detail of surface markings and TEXTURE. The differentiation operation transforms the image into a representation that facilitates the localization of edge contours and computation of properties such as their sharpness and contrast. Significant intensity changes may correspond to maxima, or peaks, in the first derivative, or to zero-crossings in the second derivative, of the image intensities. Subsequent image analysis may operate on a representation of image contours. Alternative models suggest that later processes operate directly on the result of the filtering stage.

Several sources of information are used to compute the 3-D shape of object surfaces. Binocular stereo uses the relative location of corresponding features in the images seen by the left and right eyes to infer the distance to object surfaces. Abrupt changes in motion between adjacent image regions indicate object boundaries, while smooth variations in the direction and speed of motion within image regions can be used to recover surface shape. Other cues include systematic variations in the geometric structure of image textures, such as changes in the orientation, size, or density of texture elements; image shading, which refers to smooth variations of intensity that occur as surfaces bend toward or away from a light source; and perspective, which refers to the distortion of object contours that results from the perspective projection of the 3-D scene onto the two-dimensional (2-D) image. (See STRUCTURE FROM VISUAL INFORMATION SOURCES and STEREO AND MOTION PERCEPTION for further discussion of visual cues to structure and form.)

The computation of 3-D structure cannot proceed unambiguously from the 2-D image alone. Models also incorpo-