LANGUAGE ACQUISITION after all. However, no arguments have been given that overcome the central learnability argumentation for the innateness hypothesis. For example, neither Elman et al. nor Quartz and Senjowski explain via "learning" any of the properties of universal grammar or how they are attained. Quartz and Sejnowski attempt to critique learnability theory, but their critique does not apply to actual studies in learnability theory. For example, they characterize learnability theory as assuming that learners must enumerate every possible language in the class as part of the learning procedure. This is false of Gold (1967) and explicitly argued against in Wexler and Culicover (1980). Wexler and Culicover, who are greatly concerned with the psychological plausibility of the learning procedure, derive their results under some quite severe restrictions, making their learning procedure much more empirically adequate on psychological grounds than the procedures that are considered in so-called learning accounts. (See Gibson and Wexler 1994 for an analysis of psychologically plausible learning mechanisms in the principles and parameters framework in which there is much innate knowledge.) One simply has to say that CONNECTIONIST APPROACHES TO LANGUAGE and its acquisition are simply programmatic statements without any kind of theoretical or empirical support. To be taken seriously as a competitor to the innateness hypothesis, these approaches will have to attain real results.

See also COGNITIVE DEVELOPMENT; CONNECTIONISM, PHILOSOPHICAL ISSUES; LANGUAGE AND CULTURE; MODULARITY AND LANGUAGE; NATIVISM; NATIVISM, HISTORY OF

—Kenneth Wexler

References

Chomsky, N. (1965). *Aspects of the Theory of Syntax.* Cambridge, MA: MIT Press.

Elman, J. L., E. A. Bates, M. H. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett. (1996). *Rethinking Innateness: A Connectionist Perspective on Development.* Cambridge, MA: MIT Press.

Gibson, E., and K. Wexler. (1994). Triggers. *Linguistic Inquiry* 25(3): 407–454.

Gold, E. M. (1967). Language identification in the limit. *Information and Control* 10: 447–474.

Quartz, S. R., and T. J. Sejnowski. (1997). The neural basis of cognitive development: a constructivist manifesto. *Behavioral and Brain Sciences* 20(4): 537–596.

Wexler, K. (1991). On the argument from the poverty of the stimulus. In A. Kasher, Ed., *The Chomskyan Turn.* Cambridge: Blackwell, pp. 253–270.

Wexler, K., and P. Culicover. (1980). *Formal Principles of Language Acquisition.* Cambridge, MA: MIT Press.

Wexler, K., and H. Hamburger. (1973). On the insufficiency of surface data for the learning of transformational languages. In K. J. J. Hintikka, J. M. E. Moravcsik, and P. Suppes, Eds., *Approaches to Natural Language. Proceedings of the 1970 Standard Workshop on Grammar and Semantics.* Dordrecht: Reidel, pp. 167–179.

Inner Sense

See INTROSPECTION; SELF

Integration

See MULTISENSORY INTEGRATION

Intelligence

Intelligence may be defined as the ability to adapt to, shape, and select environments, although over the years many definitions of intelligence have been offered (e.g., see symposia in *Journal of Educational Psychology* 1921; Sternberg and Detterman 1986). Various approaches have been proposed in attempts to understand it (see Sternberg 1990). The emphasis here will be on cognitive-scientific approaches.

Historically, two major competing approaches to understanding intelligence were offered, respectively, by Sir Francis Galton in England and Alfred Binet in France. Galton (1883) sought to understand (and measure) intelligence in terms of psychophysical skills, such as an individual's just noticeable difference (JND) for discriminating weights or the distance on the skin two points needed to be separated in order for them to be felt as having occurred in distinct locations. Binet and Simon (1916), in contrast, conceptualized intelligence in terms of complex judgmental abilities. Binet believed that three cognitive abilities are key to intelligence: (1) direction (knowing what has to be done and how it should be done), (2) adaptation (selection and monitoring of one's strategies for task performance), and (3) control (the ability to criticize one's own thoughts and judgments). The "metacognitive" emphasis in this conception is apparent. Binet's views have had more impact, both because his theory seemed better to capture intuitive notions of intelligence and because Binet devised a test of intelligence that successfully predicted children's performance in school.

Charles Spearman (1923) was a forerunner of contemporary cognitive approaches to intelligence in suggesting three information processes underlying intelligence: (1) apprehension of experience, (2) eduction of relations, and (3) eduction of correlates. Spearman used the four-term ANALOGY problem (A : B :: C : D) as a basis for illustrating these processes, whereby the first process involved encoding the terms; the second, inferring the relation between A and B; and the third, applying that relation from C to D.

The early part of the twentieth century was dominated by psychometric approaches to intelligence, which emphasized the measurement of individual differences but had relatively less to say about the cognitive processing underlying intelligence (see Sternberg 1990 for a review). These approaches for the most part used factor analysis, a statistical technique for discovering possible structures underlying correlational data. For example, Spearman (1927) believed that a single factor, *g* (general ability), captured most of what is important about intelligence, whereas Thurstone (1938) believed in a need for seven primary factors. More recently, Carroll (1993) has proposed a three-tier hierarchical model that is psychometrically derived, but is expressed in information-processing terms, with *g* at the top and successively more narrow cognitive skills at each lower level of the hierarchy.

A change in the field occurred when Estes (1970) and Hunt, Frost, and Lunneborg (1973) proposed what has come to be called the cognitive-correlates approach to intelligence, whereby relatively simple information-processing tasks used in the laboratories of cognitive psychologists were related to scores on conventional psychometric tests of intelligence. Hunt and his colleagues found correlations of roughly –.3 between parameters of rate of information processing in tasks such as a letter-identification task (Posner and Mitchell 1967)—where participants had to say whether letter pairs like A A, A a, or A b were the same either physically or in name—and scores on psychometric tests of verbal abilities. This approach continues actively today, with investigators proposing new tasks that they believe to be key to intelligence, such as the inspection time task, whereby individuals are assessed psychophysically for the time it takes them accurately to discern which of two lines is longer than the other (e.g., Deary and Stough 1996).

An alternative, cognitive-components approach was proposed by Sternberg (1977), who suggested that intelligence could be understood in terms of the information-processing components underlying complex reasoning and problem-solving tasks such as analogies and syllogisms. Sternberg used information-processing and mathematical modeling to decompose cognitive task performance into its elementary components and strategies. Some theorists, such as Hunt (1974) and Carpenter, Just, and Shell (1990), have used computer-simulation methodology in order to identify such components and strategies in complex tasks, such as the Raven progressive matrices.

Building on his earlier work, Sternberg (1985) proposed a triarchic theory of intelligence, according to which these information-processing components are applied to experience to adapt to, shape, and select environments. Intelligence is best understood in terms of performance on either relatively novel cognitive tasks or in terms of automatization of performance on familiar tasks. Sternberg argued that intelligence comprises three major aspects: analytical, creative, and practical thinking.

Howard Gardner (1983, 1995), in contrast, has suggested that intelligence is not unitary, but rather comprises eight distinct multiple intelligences: linguistic, logical-mathematical, spatial, musical, bodily-kinesthetic, interpersonal, intrapersonal, and naturalist. Each of these intelligences is a distinct module in the brain and operates more or less independently of the others. Gardner has offered a variety of kinds of evidence to support his theory—including cognitive-scientific research—although he has not conducted research directly to test his model.

Other theorists have tried directly to link information processing to physiological processes in the brain. For example, Haier and his colleagues (Haier et al. 1988; Haier et al. 1992) have shown via POSITRON EMISSION TOMOGRAPHY (PET) scans that brains of intelligent individuals generally consume less glucose in doing complex tasks such as Raven matrices or the game of TETRIS, suggesting that the greater expertise of intelligent people enables them to expend less effort on the tasks. Vernon and Mori (1992), among others, have attempted directly to link measured speed of neural conduction to intelligence, although there is some question as to the replicability of the findings (Wickett and Vernon 1994).

The field of intelligence has many applied offshoots. For example, a number of cognitive tests have been proposed to measure intelligence (see Sternberg 1993), and a number of different programs have been developed, based on cognitive theory, to modify intelligence (see Nickerson 1994). Some investigators have also argued that there are various kinds of intelligence, such as practical intelligence (Sternberg et al. 1995) and emotional intelligence (Goleman 1995; Salovey and Mayer 1990). The field is an active one today, and it promises to change rapidly as new theories are proposed and new data collected. The goal is not to choose among alternative paradigms, but rather for them to work together ultimately to help us produce a unified understanding of intellectual phenomena.

See also CREATIVITY; MACHIAVELLIAN INTELLIGENCE HYPOTHESIS; PROBLEM SOLVING; PSYCHOPHYSICS

—*Robert J. Sternberg*

References

Binet, A., and T. Simon. (1916). *The Development of Intelligence in Children.* Baltimore: Williams and Wilkins. (Originally published in 1905.)

Carpenter, P. A., M. A. Just, and P. Shell. (1990). What one intelligence test measures: a theoretical account of the processing in the Raven Progressive Matrices Test. *Psychological Review* 97: 404–431.

Carroll, J. B. (1993). *Human Cognitive Abilities: A Survey of Factor-Analytic Studies.* New York: Cambridge University Press.

Deary, I. J., and C. Stough. (1996). Intelligence and inspection time: achievements, prospects, and problems. *American Psychologist* 51(6): 599–608.

Estes, W. F. (1970). *Learning Theory and Mental Development.* New York: Academic Press.

Galton, F. (1883). *Inquiry into Human Faculty and its Development.* London: Macmillan.

Gardner, H. (1983). *Frames of Mind: The Theory of Multiple Intelligences.* New York: Basic.

Gardner, H. (1995). Reflections on multiple intelligences: myths and messages. *Phi Delta Kappan* 77: 200–203, 206–209.

Goleman, D. (1995). *Emotional Intelligence.* New York: Bantam.

Haier, R. J., K. H. Nuechterlein, E. Hazlett, J. C. Wu, J. Pack, H. L. Browning, and M. S. Buchsbaum. (1988). Cortical glucose metabolic rate correlates of abstract reasoning and attention studied with positron emission tomography. *Intelligence* 12: 199–217.

Haier, R. J., B. Siegel, C. Tang, L. Abel, and M. S. Buchsbaum. (1992). Intelligence and changes in regional cerebral glucose metabolic rate following learning. *Intelligence* 16: 415–426.

Hunt, E. (1974). Quote the raven? Nevermore! In L. W. Gregg, Ed., *Knowledge and Cognition.* Hillsdale, NJ: Erlbaum, pp. 129–157.

Hunt, E., N. Frost, and C. Lunneborg. (1973). Individual differences in cognition: a new approach to intelligence. In G. Bower, Ed., *The Psychology of Learning and Motivation,* vol. 7. New York: Academic Press, pp. 87–122.

"Intelligence and its measurement": A symposium. (1921). *Journal of Educational Psychology* 12: 123–147, 195–216, 271–275.

Nickerson, R. S. (1994). The teaching of thinking and problem solving. In R. J. Sternberg, Ed., *Thinking and Problem Solving.* San Diego: Academic Press, pp. 409–449.

Posner, M. I., and R. F. Mitchell. (1967). Chronometric analysis of classification. *Psychological Review* 74: 392–409.

Salovey, P., and J. D. Mayer. (1990). Emotional intelligence. *Imagination, Cognition, and Personality* 9: 185–211.

Spearman, C. (1923). *The Nature of "Intelligence" and the Principles of Cognition.* 2nd ed. London: Macmillan. (1923 edition reprinted in 1973 by Arno Press, New York.)

Spearman, C. (1927). *The Abilities of Man.* London: Macmillan.

Sternberg, R. J. (1977). *Intelligence, Information Processing, and Analogical Reasoning: The Componential Analysis of Human Abilities.* Hillsdale, NJ: Erlbaum.

Sternberg, R. J. (1985). *Beyond IQ: A Triarchic Theory of Human Intelligence.* New York: Cambridge University Press.

Sternberg, R. J. (1990). *Metaphors of Mind: Conceptions of the Nature of Intelligence.* New York: Cambridge University Press.

Sternberg, R. J. (1993). *Sternberg Triarchic Abilities Test.* Unpublished test.

Sternberg, R. J., and D. K. Detterman, Eds. (1986). *What is Intelligence? Contemporary Viewpoints on its Nature and Definition.* Norwood, NJ: Ablex.

Sternberg, R. J., R. K. Wagner, W. M. Williams, and J. A. Horvath. (1995). Testing common sense. *American Psychologist* 50(11): 912–927.

Thurstone, L. L. (1938). Primary mental abilities. *Psychometric Monographs* 1.

Vernon, P. A., and M. Mori. (1992). Intelligence, reaction times, and peripheral nerve conduction velocity. *Intelligence* 8: 273–288.

Wickett, J. C., and P. A. Vernon. (1994). Peripheral nerve conduction velocity, reaction time, and intelligence: an attempt to replicate Vernon and Mori. *Intelligence* 18: 127–132.

Intelligent Agent Architecture

Intelligent agent architecture is a model of an intelligent information-processing system defining its major subsystems, their functional roles, and the flow of information and control among them.

Many complex systems are made up of specialized subsystems that interact in circumscribed ways. In the biological world, for example, organisms have modular subsystems, such as the circulatory and digestive systems, presumably because nature can improve subsystems more easily when interactions among them are limited (see, for example, Simon 1969). These considerations apply as well to artificial systems: vehicles have fuel, electrical, and suspension subsystems; computers have central-processing, mass-storage, and input-output subsystems; and so on. When variants of a system share a common organization into subsystems, it is often useful to characterize abstractly the elements shared by all variants. For example, a family of integrated circuits might vary in clock speed or specialized data operations, while sharing a basic instruction set and memory model. In the engineering disciplines, the term *architecture* has come to refer to generic models of shared structure. Architectures serve as templates, allowing designers to develop, refine, test, and maintain complex systems in a disciplined way.

The benefits of architectures apply to the design of intelligent agents as well. An intelligent agent is a device that interacts with its environment in flexible, goal-directed ways, recognizing important states of the environment and acting to achieve desired results. Clearly, when designing a particular agent, many domain-specific features of the environment must be reflected in the detailed design of the agent. Still, the general form of the subsystems underlying intelligent interaction with the environment may carry over from domain to domain. Intelligent agent architectures attempt to capture these general forms and to enforce basic system properties such as soundness of reasoning, efficiency of response, or interruptibility. Many architectures have been proposed that emphasize one or another of these properties, and these architectures can be usefully grouped into three broad categories: the deliberative, the reactive, or the distributed.

The deliberative approach, inspired in part by FOLK PSYCHOLOGY, models agents as symbolic reasoning systems. In this approach, an agent is decomposed into data subsystems that store symbolic, propositional representations, often corresponding to commonsense beliefs, desires, and intentions, and processing subsystems responsible for perception, reasoning, planning, and execution. Some variants of this approach (Genesereth 1983; Russell 1991) emphasize formal methods and resemble approaches from formal philosophy of mind and action, especially with regard to soundness of logical reasoning, KNOWLEDGE REPRESENTATION, and RATIONAL DECISION MAKING. Others (Newell 1990) emphasize memory mechanisms, general PROBLEM SOLVING, and search. Deliberative architectures go beyond folk psychology and formal philosophy by giving concrete computational interpretations to abstract processes of representation and reasoning. Ironically, the literal-minded interpretation of mental objects has also been a source of difficulty in building practical agents: symbolic reasoning typically involves substantial search and is of high COMPUTATIONAL COMPLEXITY, and capturing extensive commonsense knowledge in machine-usable form has proved difficult as well. These problems represent significant challenges to the deliberative approach and have stimulated researchers to investigate other paradigms that might address or sidestep them.

The reactive approach to intelligent-agent design, for example, begins with the intuition that although symbolic reasoning may be a good model for certain cognitive processes, it does not characterize well the information processing involved in routine behavior such as driving, cooking, taking a walk, or manipulating everyday objects. These abilities, simple for humans, remain distant goals for robotics and seem to impose hard real-time requirements on an agent. Although these requirements are not in principle inconsistent with deliberative architectures (Georgeff and Lansky 1987), neither are they guaranteed, and in practice they have not been easily satisfied. Proponents of the reactive approach, therefore, have argued for architectures that insure real-time behavior as part of their fundamental design. Drawing on the mathematical and engineering tradition of feedback control, advocates of reactive architectures model agent and environment as coupled dynamic systems, the inputs of each being the outputs of the other. The agent contains behavioral modules that are self-contained feedback-control systems, each responsible for detecting states of the environment based on sensory data and generating appropriate output. The key is for state-estimation and output calculations to be performed fast enough to keep up with the sampling rates of the system. There is an extensive literature on how to build such behaviors (control systems) when a mathematical description of the environment is available and is of the proper form; reactive architectures advance these traditional control methods by describing how complex behaviors might be built out of simpler ones (Brooks 1986), either by switching among a fixed set of qualitatively different behaviors based on sensed conditions (see Miller, Galanter, and Pribram 1960 for precursors), by the hierarchical arrangement of behaviors (Albus 1992), or by some more intricate principle of composition. Techniques have also been proposed (Kaelbling 1988) that use off-line symbolic reasoning to derive reactive behavior modules with guaranteed real-time on-line performance.

A third architectural paradigm, explored by researchers in distributed artificial intelligence, is motivated by the following observation. A local subsystem integrating sensory data or generating potential actions may have incomplete, uncertain, or erroneous information about what is happening in the environment or what should be done. But if there are many such local nodes, the information may in fact be present, in the aggregate, to assess a situation correctly or select an appropriate global action policy. The distributed approach attempts to exploit this observation by decomposing an intelligent agent into a network of cooperating, communicating subagents, each with the ability to process inputs, produce appropriate outputs, and store intermediate states. The intelligence of the system as a whole arises from the interactions of all the system's subagents. This approach gains plausibility from the success of groups of natural intelligent agents, for example, communities of humans, who decompose problems and then reassemble the solutions, and from the parallel, distributed nature of neural computation in biological organisms. Although it may be stretching the agent metaphor to view an individual neuron as an intelligent agent, the idea that a collection of units might solve one subproblem while other collections solve others has been an attractive and persistent theme in agent design.

Intelligent-agent research is a dynamic activity and is much influenced by new trends in cognitive science and computing; developments can be anticipated across a broad front. Theoretical work continues on the formal semantics of MENTAL REPRESENTATION, models of behavior composition, and distributed problem solving. Practical advances can be expected in programming tools for building agents, as well as in applications (spurred largely by developments in computer and communications technology) involving intelligent agents in robotics and software.

See also BEHAVIOR-BASED ROBOTICS; COGNITIVE ARCHITECTURE; FUNCTIONAL DECOMPOSITION; MODULARITY OF MIND; MULTIAGENT SYSTEMS

—*Stanley J. Rosenschein*

References

Albus, J. S. (1992). RCS: A reference model architecture for intelligent control. *IEEE Comput*. 25(5): 56–59.

Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Trans. Rob. Autom*. 2: 14–23.

Genesereth, M. R. (1983). An overview of metalevel architecture. *Proceedings AAAI* 83: 119–123.

Georgeff, M., and A. Lansky. (1987). Reactive reasoning and planning. *Proceedings AAAI* 87.

Kaelbling, L. (1988). Goals as parallel program specification. *Proceedings AAAI* 88.

Miller, G., E. Galanter, and K. H. Pribram. (1960). *Plans and the Structure of Behavior.* New York: Henry Holt and Company.

Newell, A. (1990). *Unified Theories of Cognition.* Cambridge, MA: Harvard University Press.

Russell, S., and E. Wefald. (1991). *Do the Right Thing.* Cambridge, MA: MIT Press.

Simon, H. A. (1969). *The Sciences of the Artificial.* Cambridge, MA: MIT Press.

Intentional Stance

The *intentional stance* is the strategy of interpreting the behavior of an entity (person, animal, artifact, or the like) by treating it as if it were a rational agent that governed its "choice" of "action" by a "consideration" of its "beliefs" and "desires." The distinctive features of the intentional stance can best be seen by contrasting it with two more basic stances or strategies of prediction, the physical stance and the design stance. The physical stance is simply the standard laborious method of the physical sciences, in which we use whatever we know about the laws of physics and the physical constitution of the things in question to devise our prediction. When I predict that a stone released from my hand will fall to the ground, I am using the physical stance. For things that are neither alive nor artifacts, the physical stance is the only available strategy. Every physical thing, whether designed or alive or not, is subject to the laws of physics and hence behaves in ways that can be explained and predicted from the physical stance. If the thing I release from my hand is an alarm clock or a goldfish, I make the same prediction about its downward trajectory, on the same basis.

Alarm clocks, being designed objects (unlike the rock), are also amenable to a fancier style of prediction—prediction from the design stance. Suppose I categorize a novel object as an alarm clock: I can quickly reason that if I depress a few buttons just so, then some hours later the alarm clock will make a loud noise. I do not need to work out the specific physical laws that explain this marvelous regularity; I simply assume that it has a particular design—the design we call an alarm clock—and that it will function properly, as designed. Design-stance predictions are riskier than physical-stance predictions, because of the extra assumptions I have to take on board: that an entity is designed as I suppose it to be, and that it will operate according to that design—that is, it will not malfunction. Designed things are occasionally misdesigned, and sometimes they break. But this moderate price I pay in riskiness is more than compensated for by the tremendous ease of prediction.

An even riskier and swifter stance is the intentional stance, a subspecies of the design stance, in which the designed thing is an agent of sorts. An alarm clock is so simple that this fanciful anthropomorphism is, strictly