phrase structured grammar for English, and you already have background descriptions for noun, verb, and verb phrase, but no definition for a noun phrase. "Inventing" a definition for noun phrase would considerably simplify the overall hypothesized descriptions. However, the space of possible new predicates that could be invented is clearly large, and its topology not clearly understood.

ILP has found powerful applications in areas of scientific discovery in which the expressive power of logic programs is necessary for representing the concepts involved. Most notably, ILP techniques have been used to discover constraints on the molecular structures of certain biological molecules. These semiautomated scientific "discoveries" include a new structural alert for mutagenesis and the suggestion of a new binding site for an HIV protease inhibitor. ILP techniques have also been demonstrated capable of building large grammars automatically from example sentences.

The philosopher of science Gillies (1996) has made a careful comparison of techniques used in ILP with Bacon and Popper's conception of scientific induction. Gillies concludes that ILP techniques combine elements from Bacon's "pure" knowledge-free notion of induction and Popper's falsificationist approach. ILP has helped clarify a number of issues in the theory, implementation, and application of inductive inference within a computational logic framework.

See also DEDUCTIVE REASONING; PROBABILITY, FOUNDATIONS OF

—Stephen Muggleton

References and Further Readings

Bratko, I., and S. Muggleton. (1995). Applications of Inductive Logic Programming. *Communications of the ACM* 38(11): 65–70.

Carnap, R. (1952). *The Continuum of Inductive Methods.* Chicago: University of Chicago Press.

Gillies, D. A. (1996). *Artificial Intelligence and Scientific Method.* Oxford: Oxford University Press.

Jevons, W .S. (1874). *The Principles of Science: a Treatise on Logic and Scientific Method.* London: Macmillan. Republished in 1986 by IBIS publishing.

King R., S. Muggleton, A. Srinivasan, and M. Sternberg. (1996). Structure-activity relationships derived by machine learning: the use of atoms and their bond connectives to predict mutagenicity by inductive logic programming. *Proceedings of the National Academy of Sciences* 93: 438–442.

Lavrac, N., and S. Dzeroski. (1994). *Inductive Logic Programming.* Ellis Horwood.

Muggleton, S. (1995). Inverse entailment and Progol. *New Generation Computing Journal* 13: 245–286.

Muggleton, S., and L. de Raedt. (1994). Inductive Logic Programming: theory and methods. *Journal of Logic Programming* 19(20): 629–679.

Plotkin, G. (1971). A further note on inductive generalisation. In *Machine Intelligence 6.* Edinburgh: Edinburgh University Press.

Sternberg, M., R. King, R. Lewis, and S. Muggleton. (1994). Application of machine learning to structural molecular biology. *Philosophical Transactions of the Royal Society B* 344: 365–371.

Zelle, J. M., and R. J. Mooney. (1996). Comparative results on using inductive logic programming for corpus-based parser construction. In *Connectionist, Statistical and Symbolic Approaches to Learning for Natural Language Processing.* Berlin: Springer, pp. 355–369.

Infant Cognition

Questions about the origins and development of human knowledge have been posed for millennia. What do newborn infants know about their new surroundings, and what do they learn as they observe events, play with objects, or interact with people? Behind these questions are deeper ones: By what processes does knowledge grow, how does it change, and how variable are its developmental paths and endpoints?

Studies of cognition in infancy have long been viewed as a potential source of answers to these questions, but they face a problem: How does one find out what infants know? Before the twentieth century, most studies of human knowledge depended either on the ability to reflect on one's knowledge or on the ability to follow instructions and perform simple tasks with focused attention. Infants obviously are unfit subjects for these studies, and so questions about their knowledge were deemed unanswerable by scientists as different as HERMAN VON HELMHOLTZ and Edward Bradford Titchener.

The study of cognition in infancy nevertheless began in earnest in the 1920s, with the research of JEAN PIAGET. Piaget observed his own infants' spontaneous, naturally occurring actions under systematically varying conditions. Some of his most famous observations centered on infants' reaching for objects under different conditions of visibility and accessibility. Infants under nine months of age, who showed intense interest in visible objects, either failed to reach for objects or reached to inappropriate locations when the objects were occluded. Search failures and errors declined with age, a change that Piaget attributed to the emergence of abilities to represent objects as enduring, mechanical bodies. He proposed a domain-general, constructivist theory of cognitive development, according to which the development of object representations was just one manifestation of a more general change in cognitive functioning over the period from birth to eighteen months.

Later investigators have confirmed Piaget's central observations but questioned his conclusions. Studies of motor development suggest that developmental changes in infants' search for hidden objects stem in part from developing abilities to reach around obstacles, manipulate two objects in relation to one another, and inhibit prepotent actions. When these abilities are not required (for example, when infants are presented with an object that is obscured by darkness rather than by occlusion), successful search occurs at younger ages. The causes of developmental changes in search are still disputed, however, with different accounts emphasizing changes in action, ATTENTION, MEMORY, object representations, and physical knowledge.

Recent studies of cognition in infancy have tended to focus on early-developing actions such as kicking, sucking, and looking. Experiments have shown that even newborn infants learn to modify their actions so as to produce or change a perceived event: for example, babies will suck on a pacifier with increased frequency or pressure if the action is followed by a sound, and they will suck harder or longer for some sounds than for others. Studies using this method provide evidence that newborn infants recognize their parents' voices (they suck harder to hear the voice of their mother than the voice of a different woman) and their native language (they suck harder to produce speech in their own community's language). Both abilities likely depend on auditory perception and learning before birth.

A variant of this method is based on the finding that infants' sucking declines over time when followed by the same sound and then increases if the sound changes. This pattern is the basis of studies of infants' auditory discrimination, CATEGORIZATION, and memory, and it reveals remarkably acute capacities for SPEECH PERCEPTION in the first days of life. Indeed, young infants are more sensitive than adults to speech contrasts outside their native language. Studies of older infants, using similar procedures and a headturn response, reveal abilities to recognize the sounds of individual words and predictable sequences of syllables well before speech begins. The relation between these early-developing abilities and later LANGUAGE ACQUISITION is an open question guiding much current research.

Since the middle of the twentieth century, many studies of cognition in infancy have used some aspect of visual attention as a window on the development of knowledge. Even newborn infants show systematic differences in looking time to different displays, preferring patterned to homogeneous pictures, moving to stationary objects, and familiar people to strangers. Like sucking, looking time declines when a single display is repeated and increases when a new display appears. Both intrinsic preferences and preferences for novelty provide investigators with measures of detection and discrimination not unlike those used by traditional psychophysicists. They have produced quite a rich body of knowledge about early PERCEPTUAL DEVELOPMENT. Investigators now know, for example, that one-week-old infants perceive depth and the constant sizes of objects over varying distances, that two-month-old infants have begun to perceive the stability of objects over self motion and consequent image displacements in the visual field, and that three-month-old infants perceive both similarities among animals within a single species and differences across different species.

Perhaps the most intriguing, and controversial, studies using preferential looking methods have focused on more central aspects of cognitive development in infancy. Returning to knowledge of objects, experiments have shown that infants as young as three months look systematically longer at certain events that adults find unnatural or unexpected, relative to superficially similar events that adults find natural. In one series of studies, for example, three-month-old infants viewed an object that was initially fully visible on a horizontal surface, an opaque screen in front of the object rotated upward and occluded the object, and then the screen either stopped at the location of the object (expected for adults) or rotated a full half turn, passing through the space that the object had occupied (unexpected). Infants looked longer at the latter event, despite the absence of any intrinsic preference for the longer rotation. Infants' looking patterns suggested that they represented the object's continuous existence, stable location, and solidity, and that they reacted with interest or surprise when these properties were violated.

In further investigations of these abilities, the limits of early-developing object knowledge have been explored. Thus, four-month-old infants have been found to be more sensitive to the contact relations among object motions (they represent objects as initiating motion on contact with other objects) than to the inertial properties of object motions (they fail to represent objects as moving at constant or smoothly changing velocities in the absence of obstacles). Very young infants also have been shown to detect and discriminate different numbers of objects in visible and partly occluded displays when numbers are small or numerical differences are large. With large set sizes and small differences, in contrast, infants fail to respond reliably to number. Studies of cognition in infancy are most revealing where they show contrasting patterns of success and failure, as in these examples, because the patterns provide insight into the nature of the cognitive systems underlying their performance.

Where infants have shown visual preferences for events that adults judge to be unnatural, controversy has arisen concerning the interpretation of infants' looking patterns. For example, Baillargeon (1993) has proposed that the patterns provide evidence for early-developing, explicit knowledge of objects; Karmiloff-Smith (1992) has proposed that the patterns provide evidence for an initial system of object representation not unlike early-developing perceptual systems; and Haith (Haith and Bensen 1998) has proposed that preferential looking to unnatural events depends on sensory or motor systems attuned to subtle, superficial properties of the events. These contrasting possibilities animate current research.

Alongside these studies is a rich tradition of research on infants' social development, providing further insight into their cognitive capacities. Newborn infants attend to human faces, recognize familiar people, and even imitate some facial gestures and expressions in a rudimentary way. By six months, infants follow people's gaze and attend to objects on which people have acted. By nine months, infants reproduce other people's actions on objects, and they communicate about objects with gestures such as pointing. These patterns suggest that infants have considerable abilities to learn from other people, and they testify to early-developing knowledge about human action. Studies probing the nature of this knowledge, using methods parallel to the preferential looking methods just described, reveal interesting differences. Whereas infants represent inanimate object motions as initiated on contact, they represent human actions as directed to goals; whereas continuity of motion provides the strongest information for object identity, constancy of properties such as facial features provides stronger information for personal identity. Evidence for these differences has been obtained only recently and much remains to be learned, but research already suggests that distinct systems of knowledge underlie infants' reasoning about persons and inanimate objects.

In sum, the descriptive enterprise of characterizing infants' developing knowledge is well under way, both in the preceding domains and in others not mentioned. In contrast, the deeper and more important enterprise of explaining early cognitive development has hardly begun. Most investigators agree that knowledge is organized into domain-specific systems at a very early age, but they differ in their characterizations of those systems and their explanations for each system's emergence and growth. Elman et al. (1996) suggest that infants are endowed with a collection of connectionist learning systems whose differing architectures and processing characteristics predispose them to treat information from different domains. Spelke and others suggest that infants are endowed with systems of core knowledge that remain central to humans as adults. Carey (1991) proposes that infants are endowed with modular systems for processing perceptual information, but CONCEPTUAL CHANGE occurs as these systems are partly superceded over development by more central systems of representation. As research on infants' learning, knowledge, and perception progresses, these views and others will become more amenable to empirical test.

References to the experiments discussed earlier can be found in Bertenthal (1996), Haith and Bensen (1998), Mandler (1998), and Spelke and Newport (1998). Discussions of infant cognition from diverse theoretical perspectives are listed in the references.

See also COGNITIVE DEVELOPMENT; INTERSUBJECTIVITY; NAIVE PHYSICS; NATIVISM; PHONOLOGY, ACQUISITION OF

—*Elizabeth S. Spelke*

References

Baillargeon, R. (1993). The object concept revisited: New directions in the study of infants' physical knowledge. In C. Granrud, Ed., *Perception and Cognition in Infancy.* Hillsdale, NJ: Erlbaum.

Bertenthal, B. I. (1996). Origins and early development of perception, action, and representation. *Annual Review of Psychology* 47: 431–459.

Carey, S. (1991). Knowledge acquisition: enrichment or conceptual change? In S. Carey and R. Gelman, Eds., *Epigenesis of Mind: Essays on Biology and Cognition.* Hillsdale, NJ: Erlbaum, pp. 1257–1291.

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

Haith, M. M., and J. B. Bensen. (1998). Infant cognition. In D. Kuhn and R. Siegler, Eds., *Handbook of Child Psychology,* vol. 3: *Cognition, Perception, and Language.* 5th ed. New York: Wiley.

Jusczyk, P. (1997). *The Discovery of Spoken Language.* Cambridge, MA: MIT Press.

Karmiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Science.* Cambridge, MA: MIT Press.

Kellman, P. J., and M. Arterberry. (Forthcoming). *The Cradle of Knowledge: Development of Perception in Infancy.* Cambridge, MA: MIT Press.

Leslie, A. M. (1988). The necessity of illusion: perception and thought in infancy. In L. Weiskrantz, Ed., *Thought Without Language.* Oxford: Clarendon Press.

Mandler, J. M. (1998). Representation. In D. Kuhn and R. Siegler, Eds., *Handbook of Child Psychology,* vol. 3: *Cognition, Perception, and Language.* 5th ed. New York: Wiley.

Munakata, Y., J. L. McClelland, M. H. Johnson, and R. S. Siegler. (1997). Rethinking infant knowledge: toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review* 104: 686–713.

Piaget, J. (1954). *The Construction of Reality in the Child.* New York: Basic Books.

Spelke, E. S., K. Breinlinger, J. Macomber, and K. Jacobson. (1992). Origins of knowledge. *Psychological Review* 99: 605–632.

Spelke, E. S., and E. L. Newport. (Forthcoming). Nativism, empiricism, and the development of knowledge. In R. Lerner and W. Damon, Eds., *Handbook of Child Psychology,* vol. 1: *Theoretical Models of Human Development.* 5th ed. New York: Wiley.

Thelen, E., and L. B. Smith. (1994). *A Dynamical Systems Approach to the Development of Cognition and Action.* Cambridge, MA: Bradford Books/MIT Press.

Wynn, K. (1995). Infants possess a system of numerical knowledge. *Current Directions in Psychological Science* 4: 172–176.

Information Theory

Information theory is a branch of mathematics that deals with measures of information and their application to the study of communication, statistics, and complexity. It originally arose out of communication theory and is sometimes used to mean the mathematical theory that underlies communication systems. Based on the pioneering work of Claude Shannon (1948), information theory establishes the limits to the shortest description of an information source and the limits to the rate at which information can be sent over a communication channel. The results of information theory are in terms of fundamental quantities like entropy, relative entropy, and mutual information, which are defined using a probabilistic model for a communication system. These quantities have also found application to a number of other areas, including statistics, computer science, complexity and economics. In this article, we will describe these basic quantities and some of their applications. Terms like *information* and *entropy* are richly evocative with multiple meanings in everyday usage; information theory captures only some of the many facets of the notion of information. Strictly speaking, information theory is a branch of mathematics, and care should be taken in applying its concepts and tools to other areas.

Information theory relies on the theory of PROBABILITY to model information sources and communication channels. A source of information produces a message out of a set of possible messages. The difficulty of communication or storage of the message depends only on length of the representation of the message and can be isolated from the meaning of the message. If there is only one possible message, then no information is transmitted by sending that