

Cognitive Psychology: Overview☆

John F Kihlstrom^a and Lillian Park^b, ^a University of California, Berkeley, CA, United States; and ^b State University of New York, Old Westbury, NY, United States

© 2018 Elsevier Inc. All rights reserved.

Definition	2
A Short History of Cognition in Psychology	2
The Domain of Cognition	2
Sensation	3
Perception	3
Attention	4
Memory	5
Categorization	6
Learning	7
Language and Communication	8
Judgment, Reasoning and Problem Solving	8
Cognitive Development	9
The Ontogenetic View	10
Language, Culture, and Thought	10
The Phylogenetic View	11
Cognition in Personality, Social, and Clinical Psychology	11
Cognition Beyond Psychology	12
Beyond Cognition: Emotion and Motivation	13
References	13

Glossary

Counterfactual Emotions Counterfactual arguments involve reasoning which makes assumptions contrary to the facts in evidence (e.g., “If I were King, I’d make everyone rich”). Counterfactual emotions are feeling states, such as regret and disappointment, which require a comparison between some state of affairs and what might have been.

Gambler’s Fallacy The idea that prior outcomes, such as a string of “red” numbers in roulette, can influence the outcome of some future outcome, such as a “black” number; it is a fallacy because, in a truly random game, each outcome is independent of the others.

Dichotic Listening A technique in which different auditory messages are presented over separate earphones; the subject is instructed to repeat (shadow) one message but ignore the other.

Dissociation A statistical outcome in which one variable, either a subject characteristic (such as the presence of brain damage) or an experimental manipulation (such as the direction of attention), has different effects on two dependent measures (such as free recall or priming).

Functional Magnetic Resonance Imaging (fMRI) A brain-imaging technique using magnets to measure the changes in the ratios of deoxygenated to oxygenated hemoglobin due to neuronal activity.

Magnetoencephalography (MEG) A brain-imaging technique using Superconducting Quantum Interference Devices (SQUIDs) to measure changes in weak magnetic fields caused by the brain’s electrical activity.

Positron Emission Tomography (PET) A brain-imaging technique which uses positrons (positively charged electrons) to measure blood flow, metabolic rate, and biochemical changes in the brain.

Priming The facilitation (or, in the negative case, inhibition) of perceptual-cognitive processing of a target stimulus by prior presentation of a priming stimulus.

Schemata Organized knowledge structures representing a person’s beliefs and expectations, permitting the person to make inferences and predictions.

Sensory Thresholds In psychophysics, the minimum amount of energy required for an observer to detect the presence of a stimulus (the “absolute” threshold) or a change in a stimulus (the “relative” threshold).

Tabula Rasa From the Latin, “blank slate”; refers to the empiricist view that there are no innate ideas, and that all knowledge is gained through experience.

Definition

Cognition has to do with knowledge, and cognitive psychology seeks to understand how human beings acquire knowledge about themselves and the world, how this knowledge is represented in the mind and brain, and how they use this knowledge to guide behavior.

A Short History of Cognition in Psychology

Psychology was cognitive at its origins in the mid to late 19th century. Structuralists like Wilhelm Wundt and E.B. Titchener attempted to decompose conscious experience into its constituent sensations, images, and feelings. On the very first page of the *Principles of Psychology* (1890), the discipline's founding text, William James asserted that "the first fact for us, then, as psychologists, is that thinking of some sort goes on", and the functionalist tradition that he and John Dewey established sought to understand the role of thinking and other aspects of mental life in our adaptation to the environment. In the early 20th century, however, John B. Watson attempted to remake psychology as a science of behavior rather than, as James had defined it, a science of mental life.

For Watson, public observation was the key to making psychology a viable, progressive science. Because consciousness (not to mention "the unconscious") was essentially private, Watson argued that psychology should abandon any interest in mental life, and instead confine its interest to what could be publicly observed: behavior and the circumstances under which it occurred. In Watson's view, thoughts and other mental states did not cause behavior; rather, behavior was elicited by environmental stimuli. Thus began the behaviorist program, pursued most famously by B.F. Skinner, of tracing the relations between environmental events and the organism's response to them. Psychology, in the words of Robert S. Woodworth, lost its mind.

The behaviorist program dominated psychology between the two world wars and well into the 1950s, as manifested especially by the field's focus on learning in nonhuman animals such as rats and pigeons. Gradually, however, psychologists came to realize that they could not understand behavior solely in terms of the correlation between stimulus inputs and response outputs. E.C. Tolman discovered that rats learned in the absence of reinforcement, while Harry Harlow discovered that monkeys acquired general "sets" through learning, as well as specific responses. Noam Chomsky famously showed that Skinner's version of behaviorism could not account for language learning or performance, completely reinventing the discipline of linguistics in the process, and George Miller brought Chomsky's insights to psychology. Leon Kamin, Robert Rescorla, and others demonstrated that conditioned responses, even in rats, rabbits, and dogs, were mediated by expectations of predictability and controllability rather than associations based on spatiotemporal contiguity. These and other findings convinced psychologists that they could not understand the behavior of organisms without understanding the internal cognitive structures that mediated between stimulus and response.

The "cognitive revolution" in psychology, which was really more of a *counter*revolution against the revolution of behaviorism, was stimulated by the introduction of the high-speed computer. With input devices analogous to sensory and perceptual mechanisms, memory structures for storing information, control processes for passing information among them, transforming it along the way, and output devices analogous to behavior, the computer provided a tangible model for human thought. Perceiving, learning, remembering, and thinking were reconstrued in terms of "human information processing", performed by the software of the mind on the hardware of the brain. Artificial intelligence, simulated by the computer, became both a model and a challenge for human intelligence.

Jerome Bruner and George Miller founded the Center for Cognitive Studies at Harvard University in 1960, intending to bring the insights of information theory and the Chomskian approach to language to bear on psychology. Miller's book, *Plans and the Structure of Behavior* (1960, written with Karl Pribram and Eugene Galanter) replaced the reflex arc of behaviorism with the feedback loops of cybernetics. The cognitive (counter)revolution was consolidated by the publication of Neisser's *Cognitive Psychology* in 1967, and the founding of a scientific journal by the same name in 1970. With the availability of a comprehensive textbook on which undergraduate courses could be based, psychology regained its mind.

The Domain of Cognition

Cognition has to do with knowledge. Although some philosophers (including Plato, Descartes, and Kant) have asserted that some knowledge is innate, most also agree that at least some knowledge is acquired through experience. Accordingly, the goal of cognitive psychology is to provide an account of the sensory and perceptual processes by which the individual person forms internal, mental representations of the external world, the learning processes by which the person acquires knowledge through experience, the means by which these representations of knowledge and experience are stored more or less permanently in memory, the manner in which knowledge is used in the course of judgment, decision making, reasoning, problem solving, and other manifestations of human intelligence, and how one's thoughts and other mental states are communicated to others through language.

We cannot hope to give a comprehensive analysis of these processes in this small space. Detailed treatment is provided by other Reference Modules, and the textbooks listed in the bibliography. In this article, we seek only to orient the reader to the to the general thrust of work in this field, and to the problems and controversies which occupy its practitioners.

Sensation

Philosophers of mind have debated two views about the origins of knowledge: nativism and empiricism. Cognitive psychology, while acknowledging the possibility that some knowledge is innate, favors the empiricist view that most knowledge is acquired through the senses, including our reflections on sensory experience. Therefore, cognitive psychology begins with an analysis of the sensory mechanisms by which physical energies arising from a stimulus are transformed into neural impulses.

Research on sensation is dominated by the definition of the various sensory modalities (vision, audition, etc.), the determination of *thresholds* for sensory experience, the physical basis of various qualities of sensation (e.g., blue, C-sharp, and sour), and the search for *psychophysical laws* that would relate the physical properties of a stimulus to the psychological properties of the corresponding sensory experience. The most general psychophysical law, Stevens' Law ($S = kI^n$), holds that there is some exponent which will relate any physical property of a stimulus to the psychological property of its corresponding sensory experience. The analysis of sensation is so closely tied to the physical and biological sciences that it is often left out of cognitive psychology textbooks altogether. However, even such "lower" mental processes as sensation and perception do not escape the influence of "higher" mental processes of judgment and decision making.

For example, the early psychophysicists assumed that the detection of an object in the environment was simply a matter of the physical intensity of the stimulus, and the sensitivity of the corresponding receptor organs. If a light were of sufficient intensity, given the modality and species in question, it would be detected (a "hit"); otherwise, it would be missed. However, it is also the case that observers will miss stimuli that are clearly above threshold, and commit "false alarms" by "detecting" stimuli that are not actually present. Experiments based on *signal detection theory* use the pattern of hits and false alarms to decompose performance into two parameters: *sensitivity* (often indexed by a statistic known as d'), presumably closely tied to the biology of the sensory system, and *bias* (often indexed by a statistic known as β), or the perceiver's willingness to report the presence of a stimulus under conditions of uncertainty. Interest in most signal detection experiments focuses on sensitivity; bias is generally considered to be a nuisance to be evaluated and statistically controlled. But the fact that bias occurs at all shows that the perceiver's expectations, motives, and biases influence performance even in the simplest sensory task. Thus, processes involved in reasoning, judgment, and choice, and decision percolate down even to the lowest levels of the information-processing system, and they are grist for the cognitive psychologist's mill.

Signal-detection estimates of sensitivity are independent of bias, but the reverse is not the case: shifts in criterion can lead to changes in sensitivity as well. For example, if subjects are cautioned to say "yes" only when they are *certain* that the stimulus is present, they will generally yield lower estimates of d' than when they are encouraged to guess. Findings such as this suggest that the sensory system is not unitary; instead, there may be more than one sensory system, each with a different sensitivity to stimulation, and each tapped by a different method of testing.

Perception

While sensation has to do with detecting the presence of stimuli, and changes in the stimulus field, perception has to do with forming mental representations of the objects in which give rise to sensory experiences. Much perceptual research focuses on the process by which individuals determine the size, shape, distance, and motion of objects in the environment. But as Bruner argued, every act of perception is an act of categorization, and the process of perception is not complete until the object has been recognized and classified as similar to some objects and different from others.

For most of its history, the study of perception has been dominated by the *constructivist/phenomenalist* tradition associated with Hermann von Helmholtz, Richard Gregory, Julian Hochberg, and Irvin Rock (among many others). The constructivist view assumes that the proximal stimulation impinging on sensory receptors is inherently ambiguous, and that there is an infinite array of distal configurations compatible with any momentary state of proximal stimulation. As an illustration, consider the relationship between the size and distance of an object on the one hand, and the size of the retinal image of that object on the other. Holding distance constant, the size of the retinal image is directly proportional to the size of the stimulus; but holding size constant, the size of the retinal image is inversely proportional to the distance between the object and the perceiver. Thus, given the size of a retinal image, the perceiver does not know whether s/he is viewing a large object far away or a small object close at hand.

According to the constructivist view, stimulation of this sort must be disambiguated by inference-like rules which compare the size of the object with the size of its background, or some comparison of the stimulus input with an *a priori* model of the world which tells us how large various objects are. In either case, perceiving entails thinking and problem-solving. Sometimes the solution can be wrong, as when the perceptual system overcompensates for distance cues to generate the illusion that the moon on the horizon is larger than the moon at zenith. Helmholtz famously argued that the inferential rules that guide perception are part of our tacit knowledge: they can be discovered by the scientist, but cannot be articulated by the perceiver. Because the thoughts that give rise to our percepts are unconscious, perception lacks the phenomenal quality of thought. But it remains the case that the final product of perception is a mental representation of the stimulus world, constructed by cognitive operations such as computations and symbolic transformations. We are not aware of the world itself, but only of our mental representation of it, which is projected onto the world so that the objects of perception and the objects of the world are co-referential. Even more than sensation, perception from the constructivist/phenomenalist view has cognitive underpinnings which cannot be denied.

A contrary, noncognitive view of *direct realism* was proposed by J.J. Gibson in his *ecological theory of perception* (interestingly, Neisser took a constructivist approach to perception in *Cognitive Psychology*, but later embraced a version of direct realism).

According to the ecological view, stimulation is ambiguous only at very elementary levels, but there is no ambiguity at higher levels. So, for example, in determining an object's size, the perceptual system extracts information about the ratio of the size of an object to the size of its background; it is this ratio which determines perceived size, not the size of the retinal image of the object alone. Thus, the perceived size of an object remains constant even as its distance from the viewer (and thus the size of its corresponding retinal image) varies; but this requires no computations, inferences, or *a priori* models of the world on the part of the perceiver; size is effortlessly perceived directly from information available in the environment "about the ratio of the figure to ground, without need of any mediating cognitive operations. Because perceptual systems evolved in order to support adaptive behavior, Gibson further proposed that we perceive objects in terms of their *affordances*, or the actions that we can take with respect to them. Thus, in the same way that the ecological view of perception argues that all the information required for perception is "in the light", an ecological view of semantics argues that the meanings of words are "in the world", available to be perceived directly.

The ecological theory of perception proposes that the perception of form, distance, motion, and other stimulus properties is no different from perceiving the hue of a light or the pitch of a sound. In each case, phenomenal experience occurs by virtue of the transduction of stimulation into perception, accomplished in a single step by specialized neuronal structures that have evolved to be selectively sensitive to higher-order variables of stimulation available in an organism's environmental niche. Additional stimulation is provided by the perceiver's own body – for example, movements of the eyes in the perception of motion, or accommodation of the eye in the perception of depth or distance. A famous series of experiments by Dennis Proffitt found that subjects are more accurate when they estimate the slope of a hill with their hands rather than verbally; and subjects who are tired from jogging, or who are wearing moderately heavy backpacks, give steeper estimates than those who are not.

Proponents of *embodied cognition* argue that cognitive processes are shaped by the body as well as the environment: in other words, they are affected by the sensorimotor context in which they occur. Embodiment can be viewed as an extension of Gibsonian ecological perception: perception occurs in a widely distributed system that includes the body and the environment as well as the mind.

The contrast between the constructivist/phenomenalist and direct/realist views dominates much contemporary perception research, with proponents of the ecological approach conducting clever experiments showing that percepts commonly attributed to computations, inferences, or world-models are actually given directly by higher-order variables of stimulation. Still, it is one thing to demonstrate that such information is *available*, and quite another to demonstrate that such information actually contributes to perception. The occurrence of visual illusions, ambiguous figures, and other anomalies of perception strongly suggests that we do not always see the world as it really is, and that the perceiver must, in Bruner's famous phrase, go "beyond the information given" by the environment in order to form mental representations of the world around us.

Attention

In many theories, attention is the link between perception and memory: the amount of attention devoted to an event at the time it occurs (i.e., at encoding) is a good predictor of the likelihood that it will be consciously remembered later (i.e., at retrieval).

Early cognitive theories considered attention to be a kind of *bottleneck* determining whether incoming sensory information would reach short-term memory, and thus enter into "higher-level" information processing. A major controversy in early attention research was between *early selection* theories which held that preattentive processing was limited to "low-level" analyses of physical features, and *late selection* theories which allowed preattentive processing to include at least some degree of "high-level" semantic analysis. Early selection was favored by experiments showing that subjects had poor memory for information presented over the unattended channel in *dichotic listening* experiments. Late selection was favored by evidence that such subjects were responsive to the presentation of their own names over the unattended channel.

Definitive tests of early versus late selection proved hard to come by, and beginning in the 1970s the problem of attention was reformulated by Daniel Kahneman and others in terms of *mental capacity*: According to capacity theories, individuals possess a fixed amount of processing capacity, which they can deploy rather freely in the service of various cognitive activities. Various information processing tasks, in turn, differ in terms of the amount of attentional capacity they require. Some tasks may be performed automatically, without requiring any attentional capacity at all; such tasks do not interfere with each other, or with *effortful* tasks that do make demands on cognitive resources. When the total attentional capacity required by effortful tasks exceeds the individual's capacity, they will begin to interfere with each other. Some automatic processes are innate; however, other processes, initially performed effortfully, may be automatized by extensive practice. Thus, skilled readers automatically and effortlessly decode letters and words, even while they are doing something else, while unskilled readers must expend considerable mental effort performing the same task, at great cost to other, ongoing activities.

According to one prominent view, automatic processes are almost reflexive in nature – although the "reflexes" in nature are mental rather than behavioral, and they are mostly acquired through extensive practice rather than innate. That is, they are inevitably engaged by the appearance of certain stimuli, and once invoked proceed inevitably to their conclusion. Because their execution consumes no attentional resources, they do not interfere with other ongoing processes, and they leave no traces of themselves in memory. This "attention-based" notion of automaticity plays a central role in many cognitive theories, and occupies a special place in social psychology. According to a revisionist "memory-based" view, however, automaticity has nothing to do with attention, but depends on the way in which skill underlying task performance is represented in memory. Automatization occurs when

performance is controlled by procedural rather than declarative knowledge representations. From either point of view, automatic processes are strictly unconscious: we have no direct introspective awareness of them, and know them only by inference from task performance.

Memory

Perceptual activity leaves traces in memory, freeing behavior from dominance by stimuli in the immediate present. Much early cognitive psychology focused on a computer-inspired “multistore” model of memory that distinguished between three different storage structures. A set of *sensory registers*, one for each modality, held a veridical representation of stimulus input which decayed over a very short period of time. By virtue of attention, some subset of this information was transferred *short-term memory*, which could be maintained in an active state by rehearsal. Given enough rehearsal, information in short-term memory could be transferred to a permanent “*long-term*” store. The capacity of short-term memory was limited to about seven items, while the capacity of “long-term” memory was essentially unlimited. Later, the structural distinction between short- and long-term stores was replaced by a unitary view in which “short-term” (or “working”) memory refers to those items in memory which are actively engaged in processing at any moment.

Terry Winograd, John Anderson, and others have distinguished between two broad forms in which knowledge is stored in long-term memory consists of two broad forms: declarative knowledge which can be either true or false; and procedural knowledge of how certain goals are to be accomplished. Procedural knowledge can be further classified into cognitive and motor skills, such as one’s knowledge of arithmetic or grammar, or of how to tie one’s shoes or drive a standard-shift car. Similarly, Endel Tulving argued that declarative knowledge can be subdivided into episodic memories of specific experiences which occurred at a particular point in space and time, such as one’s memory for eating sushi for dinner at home last Thursday, and semantic memories that are more generic in nature, such as one’s knowledge that sushi is a Japanese dish made of rice, vegetables, and fish. In theory, many semantic memories are formed by abstraction from related episodic memories, and much procedural knowledge began as declarative knowledge – essentially, as a list of instructions. Most research on memory has focused on episodic memories for specific events, and is based on an analysis of memory into three stages of encoding, storage, and retrieval.

Research on episodic memory has been dominated by the “verbal learning” paradigm invented by Ebbinghaus, in which subjects are presented with some stimulus material, usually a list of words (hence the name), strings of letters or digits, or pictures, and then attempt to remember them after some interval of time has passed. Each list, and each item on the list, thus constitutes a specific event, which is to be remembered at a later time. By varying the conditions of encoding (e.g., number of rehearsals), storage (e.g., retention interval), and retrieval (e.g., free recall vs. recognition), researchers have identified a small number of principles of human memory.

1. *Elaboration*: Memory is a function of the degree to which an event is related to pre-existing knowledge. Ebbinghaus’s emphasis on the importance of repetition has been replaced by a distinction between two types of rehearsal: *maintenance rehearsal*, or rote repetition, maintains an item in an active state, while *elaborative rehearsal* links the item to other knowledge already stored in memory.
2. *Organization*: Memory is a function of the degree to which individual events are related to each other. Whereas elaborative processing entails the “intra-item” processing of individual items, organizational processing entails “inter-item” processing of the relations between items, yielding such phenomena as category clustering in recall.
3. *Time-Dependency*: As a general rule, memory diminishes as a function of the length of the retention interval. Time-dependent forgetting from the sensory registers is mostly a function of passive decay, while forgetting from short-term memory mostly reflects the displacement of old items by new ones. In long-term memory, forgetting can reflect failure of consolidation. Short-term consolidation occurs within seconds of an event: its disruption leads to the anterograde amnesia a traumatic head injury; long-term consolidation transpires over a longer period of time, and its disruption appears to be responsible for retrograde amnesia. Under normal circumstances, however, time-dependent forgetting from long-term memory reflects retroactive or proactive interference by competing memories, as illustrated by the paradox of knowledge: the more a subject knows about an event, the harder it is to retrieve any particular bit of information.
4. *Cue-Dependency*: Memory depends on the informational value of the cues presented at the time of retrieval. Once consolidated, memories remain permanently available in storage, but they may not be accessible on any particular attempt at retrieval. Free recall, which presents only information about the spatiotemporal context of the to-be-remembered event, is generally least productive; cued recall, which provides additional information (such as category membership) more so; and recognition, in which the cue is a copy of the target, generally yields the best performance under most circumstances.
5. The relations between encoding and retrieval processes are effectively captured by a general principle of *encoding specificity* (also known as *transfer-appropriate processing*) which states that the likelihood that an event will be remembered depends on the match between the information processed at the time of encoding and the information available at the time of retrieval.
6. *Encoding Specificity*: Memory is best when the cue information processed at the time of retrieval matches the cue information processed at the time of encoding. How an item is processed at the time of encoding determines how it will be remembered at the time of retrieval. Recognition testing can compensate for poor encoding, and deep, elaborate encoding can compensate for a poverty of retrieval cues. Encoding specificity is dramatically illustrated by state-dependent memory: within limits, items encoded under the influence of a psychoactive drug are better remembered if retrieval also takes place while drugged. Similar

effects can be induced by emotional state and environmental context. State-dependent effects are strongest, however, when retrieval cues are relatively weak and impoverished.

7. *Schematic Processing*: Memory for a specific event is a function of the relationship between that event and pre-existing knowledge, expectations, and beliefs. There is a U-shaped relationship between schema-congruence and memory: schema-relevant events are remembered better than schema-irrelevant ones; among schema-relevant events, schema-incongruent events are better remembered than schema-congruent events. Schema-incongruent, surprising events receive more elaborate processing at the time of encoding – in part, reflecting attempts to explain the discrepancy between expectations and outcomes; at the time of retrieval, the schema provides additional cues which support retrieval of schema-relevant events.
8. *Reconstruction*: Memory reflects a blend of information contained in memory traces and knowledge, expectations, and beliefs derived from other sources. Most research on human memory is based on the library metaphor: encoding, storing, and retrieving a memory is analogous to purchasing, cataloging, and shelving a book, which then is looked up in a card catalog (remember them?), taken off the shelf, and read for the information it contains. But remembering and forgetting are not simply matters of encoding, storage, and retrieval. What we remember is shaped by our attitudes, expectations, and world-knowledge, which can lead to a variety of memory illusions in which we remember events inaccurately – including events that did not actually happen to us. Memories are not simply mental records of the past; instead, memories are beliefs about the past, and remembering is more like writing a book from fragmentary notes, than it is like taking a book off the shelf and reading it.

Most research on memory has employed experimental tasks requiring conscious recollection, or the ability of subjects to recall or recognize past events. However, Daniel Schacter has pointed out that episodic memory may also be expressed implicitly in tasks that do not require conscious recollection in any form. For example, a subject who has recently read the word *veneer* will be more likely to complete the stem *ven__* with that word than with the more familiar word *vendor*. A great deal of experimental research shows that such *priming effects* can occur regardless of whether the study word is consciously remembered; in fact, they can occur in amnesic patients who have forgotten the study session in its entirety. Similarly, amnesic patients can learn new concepts without remembering any of the instances they have encountered, and acquire new cognitive and motor skills while failing to remember the learning trials themselves. Along with the concept of automaticity, the dissociations observed between explicit and implicit expressions of memory have given new life to the notion of the psychological unconscious (for a further discussion of unconscious mental processes (see “The Unconscious”, this volume).

The dissociations observed between explicit and implicit expressions of episodic memory, between semantic and episodic memories, and between procedural and declarative knowledge are subject to a variety of interpretations. According to the *multiple systems* view, explicit (conscious) and implicit (unconscious) memories are served by different memory systems in the brain. The multiple systems view, in turn, is compatible with the neuroscientific view of the brain as a collection of modules, each specialized for a particular information-processing task. By contrast, researchers who prefer a *processing* view, while accepting that there is some degree of specialization in the brain, explain these same dissociations as generated by different processes that operate in the context of a single memory system. For example, according to one processing view, implicit memories are the product of automatic, attention-free processes while explicit memories are the product of effortful, attention-demanding ones. In general, processing views are compatible with computational theories of memory, which typically assume that different memory tasks require the processing of different features of memories stored in a single memory system. One of the interesting features of the debate over explicit and implicit memory is how little contact there has been between neuroscientific views of memory on the one hand, and computational views on the other.

Categorization

Memory also stores conceptual knowledge about things in general, as well as representations of specific objects and events. Bruner noted that this conceptual knowledge plays an important role in perception: in fact, every act of perception is an act of categorization. A great deal of research in cognitive psychology has sought to understand the way in which conceptual knowledge is organized in the mind.

According to the classical view handed down by Aristotle, concepts are represented by a list of features which are *singly necessary* and *jointly sufficient* to define the category in question. For example, in geometry, all triangles are closed two-dimensional figures with three sides and three angles; and a sharp boundary divides all triangles from all quadrilaterals. However, in the 1970s it became clear that however satisfying such a definition might be philosophically, it did not reflect how concepts are represented in human minds. When perceivers judge equilateral and right triangles to be “better” triangles than isosceles triangles, they are referring to something other than a list of defining features. According to the classical view, all members of a category are equally good representatives of that category. For this and other reasons, the classical view of concepts as proper sets has been replaced with a revisionist *probabilistic* view of concepts as *fuzzy sets*. According to the fuzzy set view, features are only imperfectly correlated with category membership, and concepts themselves are represented by *prototypes* (real or imagined) which possess many share many features that are characteristic of category members. The probabilistic view permits some instances (e.g., robin) of a category (bird) to be “better” than others (e.g., emu), even though all possess the same set of defining features. Moreover, it permits the boundaries between categories to be somewhat blurred (is a tomato a fruit or a vegetable?).

Both the classical and the probabilistic view regard concepts as summary descriptions of category members. However, an alternative *exemplar* view holds that concepts are represented as collections of instances rather than as summary descriptions. Thus, when

we seek to determine whether an object is a bird, we compare it to other birds we know, rather than to some abstract notion of what a bird is. Just as there is empirical evidence allowing us to firmly reject the classical view of conceptual structure as inadequate, so there are studies showing that objects are slotted into categories if they resemble particular instances of the category in question, even if they do not resemble the category prototype. Perhaps novices in a domain categorize with respect to abstract prototypes, while experts categorize with respect to specific exemplars.

Regardless of whether concepts are represented by prototypes or exemplars, categorization is a special case of similarity judgment: the perceiver assigns an object to a category by matching its features to those of his or her category representation, prototype or exemplar. There is no absolute threshold for similarity, however: categorization, like signal detection, is always a matter of judgment.

In a sense, categorization is a special case of similarity judgment, and the most recent development in theories of concepts has been stimulated by evidence of certain anomalies of similarity. For example, subjects judge gray clouds to be similar to black clouds and different from white clouds, but judge gray hair to be similar to white hair but different from black hair. The brightness of the color patches is identical; so the judgment must be based on something other than perceptual similarity, such as the perceiver's theory about how hair changes with age or how clouds change with the weather. According to the theory-based view of categorization, concepts are not represented by lists of features or instances, and categorization does not proceed by feature matching. Rather, concepts are represented by theories which make certain features and instances relevant, and which explains how features and instances are related to each other; and categorization proceeds by applying the theory to the case at hand. It remains to be seen, however, whether the theory-based view of concepts and categorization will supplant, or merely supplement, the similarity-based view.

Learning

Behaviorism was dominated by an emphasis on learning, but cognitive psychology has not abandoned the question of how knowledge is acquired. After all, while knowledge of such basic categories as time and space may be innate, most knowledge is derived from experience. Learning, then, is the process of knowledge acquisition. In fact, some of the earliest cognitive challenges to behaviorism came through alternative accounts of learning. Even such basic processes as classical and instrumental conditioning are now interpreted in terms of the organism's developing ability to predict and control environmental events. Pavlov's dogs did not salivate to the bell because it occurred in close spatiotemporal contiguity with meat powder, and Skinner's pigeons did not peck at the key because it was reinforced by the delivery of food in the presence of a certain light. Rather, and not to anthropomorphize, they did so because they *expected* food to follow the bell and the keypeck. Learning occurs in the absence of reinforcement; reinforcement controls only performance, the organism's display of what it has learned.

The importance of expectancies, and the limited role played by contingencies of reinforcement, is underscored by the development of theories of *social learning* by Julian Rotter, Albert Bandura, Walter Mischel, and others, who argued that human learning rarely involved the direct experience of rewards and punishments. Rather, most human learning is vicarious in nature: it occurs by *precept*, in the sense of sponsored teaching, or by *example*, as in observational modeling. In either case, we learn by watching and listening to other people. Bandura argued that behavior was controlled not by environmental stimuli, but by expectancies concerning the outcomes of events and behaviors, and also by *self-efficacy* expectations – that is, people's belief that they can engage in the behaviors that produce desired outcomes. Some clinical states of anxiety may be attributed to a (perceived) lack of predictability in the environment, while some instances of depression may be attributed to a perceived lack of controllability. Interestingly, a capacity for observational learning has been uncovered in nonhuman animals such as rhesus monkeys, and has been implicated in the genesis of animal "cultures".

Learning processes are obviously implicated in analyses of the encoding stage of memory processing, in the acquisition of procedural knowledge, and in concept-formation. At the same time, cognitive psychologists have generally avoided the topic of learning itself. Partly, this may reflect an overreaction to the excessive interest in learning on the part of behaviorists; partly, it may reflect the influence of Chomsky, who discounted the role of learning in the development of language. Recently, this situation has changed due to the rise of *parallel distributed processing*, *interactive activation*, *neural network*, or *connectionist* models as alternatives to traditional symbolic processing models of human information processing. In symbolic models, each individual piece of knowledge is represented by a node, and discrete nodes are connected to each other to form a network of associative links. Thus, a node representing the concept *doctor* is linked to semantically related nodes representing concepts such as *nurse* and *hospital*. Such models are very powerful, but they leave open the question of how the knowledge represented by nodes is acquired in the first place. Connectionist models assume that individual concepts are represented by a pattern of activation existing across a large network of interconnected nodes roughly analogous to the synaptic connections among individual neurons (hence the alternative label). No individual node corresponds to any concept; every concept is represented by a pattern of widely distributed nodes. Instead of one node activating another one in turn, all nodes are activated in parallel, and each passes activation to each of the others. In connectionist systems, learning occurs as the pattern of connections among nodes is adjusted (sometimes through a learning algorithm called *back propagation*) so that stimulus inputs to the system result in the appropriate response outputs. Although the link between connectionism and stimulus-response behaviorism is obvious, connectionist theories are cognitive theories because they are concerned with the internal mental structures and processes that mediate between stimulus and response. Compared to traditional symbolic models, they are extremely powerful and efficient learning devices. Unfortunately, they also display a disconcerting tendency to forget what they have learned, as soon as they are asked to learn something new – a phenomenon known as *catastrophic interference*. Moreover,

although connectionist models seem to reflect the neural substrates of learning, it has proved difficult to demonstrate the biological plausibility of specific features such as back-propagation. Accordingly, the future of connectionist models of information processing remains uncertain.

Language and Communication

Language is both a tool for human thought and a means of human communication. The ability of people to generate and understand sentences that have never been spoken before is the hallmark of human creative intelligence, and arguably the basis for human cultural evolution. And language permits us to convey complex information about our thoughts, feelings, and goals to other people, and provides a highly efficient mechanism for social learning.

Language was one of the first domains in which cognitive psychology broke with behaviorism. Early cognitive approaches to language were couched in terms of information theory, but a real breakthrough came in the late 1950s and early 1960s with the work of Chomsky. In his early work, Chomsky distinguished between the *surface structure* of a sentence, viewed simply as a sequence of words, its *phrase structure* in terms of noun phrases and verb phrases, and its *deep structure* or underlying meaning. He also argued that *transformational grammar* mediated between deep structure and phrase structure. Transformational grammar is universal and innate. Language acquisition consists of learning the specific rules which govern the formation of surface structures in a particular language.

The field of psycholinguistics largely arose out of attempts to test Chomsky's early views. For example, it was shown that clicks presented while sentences were being read were displaced from their actual location to the boundaries between phrases; and that the time it took to understand a sentence was determined by the number of transformations it employed. Nevertheless, over the ensuing years his theory has been substantially altered by Chomsky and his colleagues, and in some quarters it has even been discarded. In any event, it is clear that other aspects of language are important, besides grammatical syntax. Languages also have phonological rules, which indicate what sounds are permitted and how they can be combined; and morphological rules which constrain how new words can be formed. Moreover, research on the *pragmatics* of language use show how such aspects of nonverbal communication as tone of voice, gesture, facial expressions, posture, and context are employed in both the production and understanding of language.

One major controversy in the psychology of language concerns speech perception. According to the *motor theory of speech perception* proposed by Alvin Liberman and his colleagues, "speech is special" in the sense that it is processed by mechanisms which are part of a specifically human cognitive endowment; no other species has this capacity. According to a rival *auditory theory of speech perception*, understanding speech is simply a special case of auditory perception, and requires no special capacities that are unique to humans. Tests of these theories often revolve around *categorical perception*, or the ability to distinguish between related sounds such as [b] and [p]. Evidence that there are sharp boundaries between such speech categories is often attributed to innate, specifically human, mechanisms for producing speech, thus favoring the motor theory. On the other hand, evidence of categorical perception in nonhuman species which do not have a capacity for speech favors the auditory theory.

A common theme in the Chomskian approach to language is that language use is mediated by innate rules which cannot be acquired by general-purpose systems which learn solely by virtue of associations among environmental events. Chomsky claims that support for the role of innate rules comes from the errors children and others language learners make in forming the past tenses of irregular verbs: *goed* instead of *went*, or *eated* instead of *ate*. Because the children have never heard such words (the adults they listen to don't make them), they cannot have been acquired through experience; rather, they must be produced by a rule which the child has abstracted in the course of learning his or her native tongue. If computers are to have a language capacity, they must be programmed with these kinds of rules; they will never learn language without some rule-based cognitive structure. Recently, however, connectionist models of language have been developed which have no rules of syntax, and operate solely by associationistic principles, but which make precisely the errors that are traditionally attributed to the operation of grammatical rules. If so, the implication may be that human language is nothing "special" after all: it is something that can be done by any associationistic learning system which possesses sufficient computational power. Proponents of rules, in turn, have criticized these demonstrations as misleading and unrepresentative of people's actual language use. For example, the model makes mistakes that children do not make, and learns by virtue of inputs that do not resemble children's actual learning environments. Because the capacity for language is so central to our traditional conception of what it means to be human, the debate between rules and connections is likely to be vigorous and protracted.

Judgment, Reasoning and Problem Solving

Thinking played a prominent role in early psychology. The structuralist school of Wundt and Titchener was almost consumed by a fruitless debate over imageless thought, while Oswald Kulpe's *act psychology* attempted to characterize the process of *thinking* rather than the static elements of *thoughts*. Like other mentalistic topics, thinking dropped out of sight with the rise of behaviorism, but interest in the topic was preserved by Gestalt psychologists, as in Kohler's work on problem-solving in chimpanzees. Whereas behaviorists construed thinking as a matter of gradual trial-and-error learning, the Gestaltists emphasized sudden insights produced by a cognitive restructuring of the problem at hand. In England, Frederick C. Bartlett argued that perception and memory were essentially exercises in problem-solving – the problem being to construct mental representations of the present and to reconstruct mental representations of the past. After World War II, the cognitive revolution was heralded by Bruner's work on

concept learning, Jean Piaget's research on the development of thought in children, and the work of Herbert Simon and Allan Newell on computer simulations of human problem-solving.

Much early research on thinking was guided, at least tacitly, by a normative model of human rationality which held that people reason according to a logical calculus, and make rational choices based on principles of optimality and utility. According to the classical view of concept structure, for example, people categorize objects according to lists of defining features which are singly necessary and jointly sufficient to define category membership. In classical decision theory, individuals calculate the costs and benefits to themselves of various options, and then make choices that maximize their gains and minimize their losses as efficiently as possible. Much work on reasoning was dominated by the search for *algorithms*: logical, systematic rules, analogous to recipes, that specify how information should be combined to yield the correct solution to whatever problem is at hand.

Although normative rationality provided a reasonable starting point for developing theories of human reasoning and problem solving, many human judgments must be made under conditions of *uncertainty*, where there is no algorithm applicable, the information needed to apply an algorithm is unavailable. Other judgments must be made under conditions of *complexity*, where there are simply too many choices available, or too many factors entering into each choice, to permit evaluation according to some judgment algorithm. Under such conditions, people tend to rely on *judgment heuristics* – shortcut “rules of thumb” which bypass normative rules of logical inference, and thus permit judgments without recourse to algorithms. One such heuristic, associated with the work of Herbert Simon on organizational decision-making, is *satisficing*: instead of conducting an exhaustive search for the optimal choice, a judge may terminate search as soon as the first satisfactory option is encountered. Prototype or exemplar-matching constitute heuristics for categorization: instead of consulting a list of defining features that are singly necessary and jointly sufficient to assign an object to some category, people compare the object at hand to a “typical” instance, or indeed to any instance at all.

While appropriate algorithms are guaranteed to deliver the correct solution to whatever problem is at hand, use of heuristics incurs some risk of making an error in reasoning or judgment. Analysis of common judgment errors, such as the “gambler's fallacy”, by Daniel Kahneman, Amos Tversky, and others has documented a number of other commonly used judgment heuristics. *Representativeness* permits judgments of category membership, similarity, probability, and causality to be based on the degree to which an event resembles the population of events from which it has been drawn. *Availability* permits judgments of frequency and probability to be based on the ease with which relevant examples can be brought to mind, while *simulation* bases judgments on the ease with which plausible scenarios can be constructed. In *anchoring and adjustment*, initial estimates are taken as reasonable approximations to the final result of some calculation.

These and other effects show that the principles of cognitive functioning cannot simply be inferred from abstract logical considerations; rather, they must be inferred from empirical data showing how people *actually* perform. Research shows that people commonly depart from the principles of normative rationality, but a further question is what we should make of these departures. Although Aristotle defined humans as rational animals, one possible conclusion from empirical studies is that people are fundamentally irrational: that human judgment, reasoning, choice, and problem-solving is overwhelmed by a large number of fallacies, illusions, biases, and other shortcomings. At best, according to this argument, most people are “cognitive misers” who use as little information, and as little cognitive effort, as possible in their lives; at worst, people are just plain stupid – incapable, without extensive instruction (and perhaps not even then), of conforming themselves to the principles of logic and rationality.

This pessimistic conclusion about human nature is a little reminiscent of Sigmund Freud's argument, around the turn of the last century, that human rationality is derailed by unconscious affects and drives. On the other hand, it is possible that the case for human irrationality has been overstated. For example, the philosopher Jonathan Cohen has questioned whether the formal laws of deductive and probabilistic reasoning are properly applied to the problems which people actually encounter in the ordinary course of everyday living. Herbert Simon, for his part, has concluded that human rationality is bounded by limitations on human information-processing capacity (as captured, for example, by George Miller's famous essay on “the magical number seven, plus or minus two”). From this perspective, it is simply unreasonable to hold humans up to an impossible standard of unbounded rationality, of a sort that might characterize a computer which has the capacity to search and calculate for as long as it takes to deliver a “logical” result. Relatedly, Gerd Gigerenzer and his colleagues have argued that “fast and frugal” judgment heuristics succeed more often than they fail because they are appropriately tuned to the structure of the environments in which people actually operate. From this perspective, most “fallacies” in human reasoning emerge in performance on laboratory tasks that do not adequately reflect the real world in which judgment heuristics work. On the other hand, some evolutionary psychologists have argued that judgment heuristics are part of an “adaptive toolbox” of domain-specific cognitive devices (or modules) that evolved in the “environment of evolutionary adaptedness” – namely, the African savanna of the Pleistocene era, roughly 1.8 million years ago – to help our hominid ancestors solve fundamental problems of survival and reproduction.

Cognitive Development

Compared to other mammals, human beings are born with relatively immature brains; moreover, physical development continues even after brain development has essentially completed. These facts raise the question of how the intellectual functions characteristic of the human adult arise in the first place (and, from a life-span perspective, whether, how, and to what extent cognitive skills are lost through aging).

The Ontogenetic View

From an ontogenetic point of view, tracing the growth of cognition in the individual organism, cognitive development has recapitulated the debate between nativism and empiricism which has dominated cognitive psychology at large. From the empiricist perspective, the child is a *tabula rasa*, who acquires knowledge and skills with learning and experience. From the nativist perspective, even neonates possess at least primitive cognitive faculties, which develop further in interaction with the environment. The tension between nativism and empiricism can be seen clearly in the debate, discussed earlier, over whether language acquisition is mediated by innate grammatical rules or by a general-purpose associative learning mechanism.

A new perspective on cognitive development, combining elements of both nativism and empiricism, was offered by Jean Piaget, who proposed that children enter the world with a rudimentary set of reflex-like cognitive structures, called *sensory-motor schemata*, through which the child interacts with the world. Environmental events are interpreted through prevailing cognitive schemata but they also force these schemata to change in order to cope with an increasingly complicated stimulus environment. Through the cycle of assimilation and accommodation, the child moves through a number of qualitatively different stages, each highlighted by a particular cognitive achievement. The milestone marking the end of the sensory-motor stage, at about 18 months of age, is *object permanence* – the ability to deal with objects that are not present in the immediate physical environment. The end of the *preoperational stage*, at about age 7, is marked by *conservation*, or the ability to appreciate that quantities remain constant despite changes in physical appearance. The transition from *concrete operations* to *formal operations*, at about age 12, is marked by the child's ability to comprehend abstract concepts and formal logical relations.

For Piaget, the child proceeds through these stages of cognitive growth in a strict sequence: some tasks, requiring abilities characteristic of later stages, are simply impossible for a child who is still at an earlier stage. Such a proposal was bound to be challenged, and indeed later experiments employing extremely subtle measures often showed that, as a 1993 cover story in *Life* magazine put it, “Babies are smarter than you think”. For example, 7-month-old infants may not reach for the spot where a toy has been hidden, but they do stare at it, indicating that they have some sense of object permanence after all. Similarly, when a mouse hidden behind a screen is joined by a second mouse, 5-month-olds show surprise when the screen is revealed to reveal only one mouse, suggesting that they have some ability to conserve number.

Results such as these have suggested to some theorists that infants enter the world with a surprisingly sophisticated fund of innate knowledge about the world, which is refined and elaborated through experience. What develops, then, is *expertise*: the infant starts out as a novice with respect to objects, numbers, and the like. Development proceeds with continuous increases in motor control and information-processing capacity, and also with increased opportunities for learning through experience. Infants reach for the hidden toy not because they have acquired object permanence, but because they have acquired the ability to coordinate their actions with their thoughts. More recent experiments, however, indicate that infants' appreciation of object permanence is incomplete. Developing children acquire new knowledge and skills, not just the ability to use innate knowledge and skills more efficiently and effectively.

Development also entails the acquisition of *metacognition*: one's knowledge of what one knows, and how one's own mind works, and how this knowledge can be deployed strategically in the service of adaptive behavior. Metacognitive knowledge is sometimes characterized as a *theory of mind*, a phrase which recalls Piaget's argument that children, no less than adults, function as naïve scientists. From the first moments of life, children are constantly trying to understand themselves and the world around them, including other people, by generating hypotheses, testing them empirically, and refining their theories accordingly. Children do not just develop increasingly sophisticated theories of mind; they also develop increasingly sophisticated theories of physics, chemistry, biology, and the rest of the natural world, as well as sociological and anthropological theories about groups, organizations, societies, and cultures. The “theory” theory, as it is sometimes called, makes clear that cognitive development is not just something that happens passively to the child by virtue of maturation, learning, and the activities of adults. Rather, the child takes an active role in his or her own development, instigating the very interactions that promote cognitive growth through learning.

Language, Culture, and Thought

Cognitive development can also be viewed in cultural terms. Cognitive anthropology (see below) had its origins in the efforts of Lucien Levy-Bruhl, Franz Boas, W.H. R. Rivers, and others to determine whether there were differences in the thought patterns characteristic of members of “primitive” and “advanced” cultures. Some early Soviet psychologists, such as Lev Vygotsky and Alexander Luria, attempted to trace the effects of economic development on the way people think. In view of the central role of language in culture, considerable effort has been devoted to the question, initially raised by the American anthropologists Edward Sapir and Benjamin Whorf, whether there are cognitive differences between speakers of different languages. This is not a matter of “development”, per se, because all languages are equally complex: it is merely a matter of language and culture.

The Sapir-Whorf hypothesis takes two forms: that language determines thought or that language influences thought. The former is a much stronger view because it states that one is incapable of understanding a concept for which the language has no name (it also implies that there is no thought without language). There is no empirical evidence supporting the strong version and considerable evidence that thought can proceed without benefit of language. However, the weak version plausibly suggests that different languages can “carve up” the world into different ways – or, put another way, that conceptual thinking can be shaped and constrained by available linguistic categories. As Whorf put it, “We cut nature up, organize it into concepts, ascribe significance as

we do, largely because we are parties to an agreement to organize it in this way – an agreement that holds throughout our speech community and is codified in the patterns of our language”.

There are actually two aspects to the Sapir-Whorf hypothesis: *linguistic relativity* and *linguistic determinism*. Relativity refers to the claim that speakers are required to pay attention to different aspects of the world that are grammatically marked (e.g. shape classifiers in Japanese or verb tenses to indicate time). Determinism claims that our cognitive processes are influenced by the differences that are found in languages. The most famous example, and the most erroneous, of the Whorf hypothesis is Whorf's observation that Eskimos have many words for snow, implying that because they live in a snowy environment they needed to come up with finer distinctions for the different types of snow. But American skiers have different words for snow, too, so the example is not as remarkable as it first may appear because expertise leads to larger vocabularies for certain domains.

In a classic test of the Sapir-Whorf hypothesis, Paul Kay and his colleagues compared English speakers with Tarahumara speakers, a Uto-Aztecan language of Mexico that does not have a separate color term for blue and green. In the first experiment, the subjects were presented with a blue color chip, a green, color chip, and another color chip that was intermediate to blue and green. English speakers sharply distinguished the intermediate color chip into either blue or green by using a naming strategy, whereas the Tarahumara speakers chose randomly. In the second experiment, English speakers were first presented with two color chips and shown that one (intermediate) was greener than the other color chip (blue) and then shown that the same intermediate chip was bluer than the other color chip (green). By making the subjects call the intermediate color chip both green and blue, the bias that was demonstrated in the first experiment went away and the English speakers performed similarly to the Tarahumara speakers. More recent studies have gone beyond categorization to document language effects in motion perception, decision-making in economic and health domains, and temporal orientation.

The influence of language on how we think about the events that happen in our world can be demonstrated in other experiments other than those designed to confirm or disconfirm the Whorf hypothesis. Classic work by Leonard Carmichael and his colleagues demonstrated that subjects had different systematic distortions in their recall of ambiguous line drawings depending upon which verbal label they were given (e.g. dumbbells or eyeglasses). Experiments on eyewitness testimony by Elizabeth Loftus and others showed that by varying the verb (e.g. crashed or hit) one can manipulate the estimated speed of the traveling car given by the subjects. Whorf himself became interested in language when he noticed that behavior around gasoline drums changed when the drums were called “empty” though they contained dangerous vapors. Because the word *empty* connotes “lack of hazard,” careless behavior from the workers resulted in fires from the tossing of cigarette stubs or the smoking by the workers.

Beyond the influence of language on categories, the linguist George Lakoff's work on metaphors offer another way of testing the Sapir-Whorf hypothesis without depending upon the idea that language carves the world into different pieces and, as he has put it, “cultures differ only in the way they have their meat cut up”. Though some metaphors are universal (e.g. love is warmth), not all cultures share the same metaphors. By fleshing out the Sapir-Whorf hypothesis with the use of different sophisticated cognitive tasks and not relying on the differences between “exotic” and “non-exotic” languages, we can further explore the ways in which the language we speak shapes our thoughts.

The Phylogenetic View

Cognitive development can also be approached from a *phylogenetic* perspective, tracing the relations between the intellectual functions of human and children and adults and those of other animals – especially the great apes, whose genetic endowment is so similar to our own. There is a vigorous debate over whether chimpanzees and gorillas have anything like the human capacity for language, but it is clear that these and other animals do have the ability to acquire symbolic representations of objects, events, and concepts – something like semantics, if not syntax as well. Pigeons can be taught to categorize a wide variety of objects including trees, people (and their emotional expressions), fish, flowers, and automobiles. Studies of mirror-recognition indicate that chimpanzees possess a rudimentary concept of self, while the notion of a “theory of mind” initially arose out of observations that chimpanzees had the ability to attribute mental states to others of their kind. Setting aside the question of whether other species have a capacity for language, it is clear that the behavior of nonhuman animals – especially those who are closest to us in the evolutionary scheme of things – is not just a matter of innate and conditioned responses; some of them, at least, have cognitive capacities not unlike our own.

Cognition in Personality, Social, and Clinical Psychology

The cognitive revolution in psychology began with traditional experimental psychology, with a focus on learning, perception, memory, and thinking, but it quickly spread to other areas of the field. The cognitive approach to social psychology is based on the idea that interpersonal behavior is shaped by the actors' perceptions of the social situation, each other, and themselves; memories of past experiences in similar situations; and decisions about what to do in the current one. Similarly, the cognitive approach to personality assumes that individual differences in social behavior are based on individual differences in the construal (categorization) of the social situation, as well in the declarative and procedural knowledge that each participant brings to bear on the interaction. Much of experimental psychopathology is concerned with psychological deficits associated with various forms of mental illness – for example, difficulties of attention in schizophrenia. And cognitive therapy attempts to change maladaptive behavior by altering the patients' beliefs and expectations concerning themselves, other people, and the world around them.

Cognition Beyond Psychology

The cognitive revolution in psychology was paralleled by the development of the field of *cognitive science*, whose practitioners included philosophers, linguists, computer scientists, neuroscientists, behavioral biologists, sociologists, and anthropologists, as well as psychologists. In some sense, the rise of cognitive science may have been a reaction to the dominance of behaviorism within psychology: many who wished to pursue a science of mental life may have felt that they would have to go outside psychology to do so. By the same token, it seems reasonable to hope that the combined efforts of a number of different disciplines are more likely to yield a better understanding of cognitive processes than any one of them working in isolation.

Cognitive science has much to contribute to the understanding of human cognition, but its brief history goes beyond the human to include the problem of intelligent machines. While some early cognitive psychologists viewed the computer to be a model of the human mind, some early cognitive scientists believed that it offered the prospect of implementing the “mechanical mind” debated by philosophers at least since the time of Descartes.

In the formulation of the philosopher John Searle, work on *artificial intelligence* (AI) takes two broad forms. In “weak” AI, the computer provides a vehicle for writing formal theories of the mind, which can be tested by pitting the results of a computer simulation against the data of actual human performance. In terms of formal precision of its theories, weak AI is cognitive psychology at its best. By contrast, “strong” AI entails the notion that computer programs can, in principle, really *think* just as humans do. The program of strong AI has its origins in the proposal, by Alan Turing, that appropriately programmed computers are capable of performing any explicitly stated cognitive task: a machine would pass the *Turing test* if the responses a computer were indistinguishable from those of a human being. Searle believes that the program of strong AI is seriously misguided – a position that is opposed with equal vigor by other philosophers, such as Daniel Dennett.

More recently, research in artificial intelligence has shifted from an effort to make machines think the way humans do, to an effort to allow machines to “think” however they are capable of doing so, regardless of how humans might accomplish the same task. Thus, in May 1997 “Deep Blue”, a supercomputer programmed by IBM, was able to beat the world champion Gary Kasparov at chess, but nobody claimed that Deep Blue played chess the way Kasparov (or any other human) did. Furthermore in 2016, AlphaGo, an AI computer program developed by Alphabet Inc.’s Google DeepMind, beat Lee Sedol, a grandmaster in Go, a complex Chinese strategy board game, with an unconventional move never seen before in its 2500-year history. The move was described as “not a human move” and in that moment, commentators and other highly skilled players did not know what to make of that move. Indeed these non-intuitive unconventional moves in game play are called “computer moves” because these AI programs are able to formulate moves that escape human thought. Cognitive psychology remains an important component of cognitive science. But to the extent that it seeks to develop intelligent machines on their own terms, without reference to human intelligence, cognitive science departs from cognitive psychology.

Cognitive science, which once was dominated by behavioral experiments and computational models, has recently reached “down” to strengthen its connections to neuroscience. At the same time, neuroscience, which once was preoccupied with events at the molecular and cellular levels, has reached “up” to take an interest in the organismal level of experience, thought, and action. Both trends have been aided by the development of brain imaging techniques, such as positron emission tomography (PET), functional magnetic resonance imaging (fMRI), and magnetoencephalography (MEG), which open windows on the brain as it is engaged in complex cognitive activities such as perceiving, remembering, imaging, and thinking. Particularly promising is the combination of the high spatial resolution of fMRI with the high temporal resolution of MEG. Just as earlier investigators discovered specific cortical areas specialized for vision, hearing, and the like, so a new generation of brain researchers have uncovered specific areas activated in such activities as language comprehension, mathematical computation, analytical reasoning, and working memory. At the same time, it has become clear that the brain is not simply a collection of task-specific modules, as suggested by the common metaphor of a Swiss Army knife. Most cognitive activities involve large swaths of brain tissue, as information is passed among a large number of different modules.

Brain-imaging techniques such as functional magnetic resonance imaging (fMRI) promise to solve the ancient mind-body problem, but they also present the danger of lapsing into a kind of high-technology revival of phrenology. In the final analysis, brain imaging can only reveal which areas of the brain are activated when experimental subjects engage in particular tasks, such as lexical decision and mental arithmetic. Discovering what these areas *do* requires a careful analysis of the experimental tasks employed in the imaging study, and that is a matter for cognitive psychology. If researchers do not have a correct description of the components of the task at the psychological level of analysis, they will reach erroneous conclusions concerning the cognitive functions of various parts of the brain. Solving the mind-body problem is not just a matter of building bigger magnets, resolving the details of neurochemistry, and ruling out physiological artifacts. Cognitive neuroscientists sometimes engage in a “rhetoric of constraint” which asserts that psychological theories of mental, a structures and processes are constrained by the facts of brain anatomy and physiology. But cognitive neuroscientific theories are theories of function, and describing functions is the task of psychology. Genuine advances in cognitive neuroscience depend on continued progress in cognitive psychology, so that brain researchers can work with tasks that are well understood.

Even within the social sciences, it is clear that cognition is not just for psychologists anymore (if it ever was). Linguistics, traditionally concerned with the discovery of linguistic regularities and the origins of words, has increasingly worked to understand language as a tool for thought and means of sharing ideas. Economics, once concerned solely with the abstract description of economic systems, has more recently turned its attention to individual economic decision making (and, in the process, drawing on the insights of psychologists such as Tversky and Kahneman). For cognitive sociologists, social conventions and norms create

a framework in which individuals think their thoughts. By the same token, cognitive anthropologists are willing to entertain (and test) the hypothesis that cultural differences entail differences in modes of thought as well as differences in beliefs and behavior. Sociology and anthropology have challenged the doctrines of individualism and universalism, which have traditionally dominated psychological approaches to human thought.

Beyond Cognition: Emotion and Motivation

Cognitive psychology is about knowing, but *knowing* is not all the mind does. In his *Critique of Pure Reason* (1791), the philosopher Immanuel Kant proposed that there are three “faculties of mind”, knowledge, feeling, and desire; each of these enters into a causal relationship with behavior, and none is reducible to any other. If Kant is right, then cognitive psychology cannot be all there is to psychology: the principles of cognition must be supplemented by principles of emotion and motivation. In fact, some cognitive psychologists have argued that Kant was wrong, and that our emotional and motivational states are the byproducts of cognitive activity. For example, prominent cognitive theories of emotion hold that our emotional states are, essentially, *beliefs* about our feelings. Put another way, cognitive theories of emotion hold that our emotional states depend on our interpretation of environmental events and our own behaviors. As William James famously put it: we do not run from the bear because we are afraid; we are afraid because we run from the bear. In response, some theorists have argued that emotions are not dependent on cognitive processing, but rather are governed by their own, independent systems. To some degree, such proposals reflect a reaction to the hegemony of the cognitive point of view within psychology. At the same time, the question of the independence of emotion and motivation from cognition is a legitimate one, and has given rise to a new interdisciplinary field, *affective neuroscience*, proceeding in parallel with cognitive neuroscience.

Regardless of how the independence issue is resolved, it is clear that cognitive processes can influence emotions and motives. Emotions can be induced by remembering past events, and they can be altered by construing events differently. Certain “counterfactual” emotions, such as disappointment and regret, require that the person construct a mental representation of what might have been. While some emotional reactions may be innate and reflex-like, others are acquired through conditioning and social learning. As noted earlier, there is evidence that some emotional states, such as anxiety and depression, result from the perception that environmental events are unpredictable or uncontrollable; they may disappear when such beliefs are corrected. Surgical patients’ fears can be allayed (and the outcome of treatment improved) if their doctors carefully explain what is going to happen to them, and why it is necessary. The ability to use cognitive processes to regulate one’s own feelings and desires is an important component of *emotional intelligence*.

Turning to the other side of the coin, it is clear that emotional and motivational states can have an impact on cognition. In an important sense, the “affective revolution” in psychology was initiated by studies of the effects of mood on memory; these led psychologists to become more interested in the nature of the moods themselves. Five such effects have been well documented: the *affective intensity effect* (better memory for positive or negative events, compared to neutral events); the *affective valence effect* (better memory for positive than for negative events); *mood-congruent memory* (better memory for material whose affective valence matches the mood in which it is encoded or retrieved); *resource allocation effects* (depression impairs performance on effortful, but not automatic, aspects of memory function); and *mood-dependent memory* (memory is better when there is congruence between the emotional state present at the time of encoding, and the state present at the time of retrieval). Although clinical lore holds that emotional trauma can render people amnesic, the overwhelming finding in both the clinical and experimental literature is that traumatic experiences are remembered all too well.

There is also a growing literature on the emotional effects of other cognitive processes, such as perception and judgment. Signal-detection theory has already demonstrated that goals and motives can percolate “down” to affect the most elementary psychological functions. Common metaphors speak of happy people viewing the world through rose-colored glasses, and that things look dark when we’re unhappy, and in fact mood and emotion do seem to serve as filters on perception, just as they do on memory. Similarly, emotions have a considerable effect on judgment and decision-making. *Prospect theory*, proposed by Kahneman and Tversky as an alternative to rational choice, holds that decisions are affected by the way that choices are framed, and emotions and motives form an important element in these frames. Happy people are more likely to take risks than unhappy ones. In his Nobel Prize acceptance speech, Daniel Kahneman noted that because our experiences are finite, we tend to base our decisions on what we think will lead to happy memories. Even if feelings and desires prove to be largely independent of knowledge and belief, the interest of cognitive psychologists in our emotional and motivational lives gives eloquent testimony to the breadth of the field as it approaches its second half-century.

References

- Anderson, J.R., 2014. *Cognitive Psychology and Its Implications*, eighth ed. Worth, New York.
Baars, B.J., 1986. *The Cognitive Revolution in Psychology*. Guilford, New York.

☆ *Change History*: June 2017, Author John F. Kihlstrom updated the sections on Memory and Reasoning, and made smaller changes elsewhere, updated the Reference list. Author Lillian Park updated the section on Language and made smaller changes elsewhere.
July 2017: 1. Author John F. Kihlstrom added Keywords and Abstract and sections.

- Boroditsky, L., 2018. *7,000 Universes: How the Language We Speak Shapes the Way We Think*. Doubleday, New York.
- Carlston, D.E. (Ed.), 2013. *The Oxford Handbook of Social Cognition*. Oxford University Press, New York.
- Chipman, S.E.F. (Ed.), 2017. *The Oxford Handbook of Cognitive Science*. Oxford University Press, New York.
- D'Andrade, R., 1995. *The Development of Cognitive Anthropology*. Cambridge University Press, Cambridge, U.K.
- Gardner, H., 1985. *The Mind's New Science: A History of the Cognitive Revolution*. Basic Books, New York.
- Gazzaniga, M.S., Ivry, R.B., Mangun, G.R., 2014. *Cognitive Neuroscience: The Biology of the Mind*, fourth ed. Norton, New York.
- Gopnik, A., Meltzoff, A.N., Kuhl, P.K., 1999. *The Scientist in the Crib*. Morrow, New York.
- Gould, J.L., Gould, C.G., 1994. *The Animal Mind*. Scientific American Library, New York.
- Holyoak, K.J., Morrison, R.G. (Eds.), 2012. *The Oxford Handbook of Thinking and Reasoning*. Oxford University Press, New York.
- Kronenfeld, D.G., Bennardo, G., de Munck, V.C., Fischer, M.D., 2011. *A Companion to Cognitive Anthropology*. Wiley-Blackwell, New York.
- Medin, D.L., Ross, B.H., Markman, A.B., 2004. *Cognitive Psychology*, fourth ed. Wiley, New York.
- Ochsner, K.M., Kosslyn, S.M. (Eds.), 2013. *The Oxford Handbook of Cognitive Neuroscience*. Oxford University Press, New York.
- Park, D.C., Schwarz, N. (Eds.), 2000. *Cognitive aging: A primer*. Psychology Press, Philadelphia.
- Reisberg, D.N., 2015. *Cognition: Exploring the Science of the Mind*, sixth ed. Norton, New York.
- Reisberg, D.N. (Ed.), 2013. *The Oxford Handbook of Cognitive Psychology*. Oxford University Press, New York.
- Schwartz, B., Wasserman, E.A., Robbins, S.J., 2002. *Psychology of Learning and Behavior*, fifth ed. W.W. Norton, New York.
- Shettleworth, S.J., 1998. *Cognition, Evolution, and Behavior*. Oxford University Press, New York.
- Siegler, R.S., 1996. *Emerging Minds: The Process of Change in Children's Thinking*. Oxford University Press, New York.
- Sternberg, R.J., Sternberg, K., 2016. In: *Cognitive Psychology*, seventh ed. Cengage.
- Tulving, E., Craik, F.I.M. (Eds.), 2000. *The Oxford Handbook of Memory*. Oxford University Press, New York, NY, US.
- Wilson, R.A., Keil, F.C., 1999. *The MIT Encyclopedia of the Cognitive Sciences*. MIT Press, Cambridge, MA.
- Zerubavel, E., 1997. *Social Mindscapes: An Invitation to Cognitive Sociology*. Harvard University Press, Cambridge, MA.