E
X
H
I
B
I
T

38

CHAPTER 4

Representation in Memory

Associative Networks: Organizing Memory ✦ Procedural and Declarative Memory: What Memory Does ✦ Parallel Versus Serial Processing: Coordinating Memory Processes ✦ Embodied Memory: Including Physical Representation ✦ Social Memory Structures: Why Social Memory Matters

This chapter covers social and nonsocial models of memory and closes with mental representations—the social categories and individual exemplars that we keep in mind. Several models of memory are introduced. First we outline each cognitive model, and then we describe the social cognitive models derived from it. In this area more than any other, social cognition researchers have adapted general cognitive models to develop models specific to social cognition, in particular, memory for people. We discuss associative networks, procedural memory, connectionist (parallel) models, embodied memory, and memory structures such as categories and exemplars.

ASSOCIATIVE NETWORKS: ORGANIZING MEMORY

Suppose you are standing at a busy intersection waiting for the light to change. Across the street, you see a young man knock down an elderly woman, grab her purse, and run away. By the time you can get across the street, he is long gone, so you turn your attention to the woman. Just as you have discovered that she is angry but unhurt, a police officer arrives and takes down your description of what happened. How is this event stored in your memory, and why does its mental representation matter? Several models of memory account for this phenomenon. In this section we detail the most common and well-developed type of memory model, the associative network approach. This model underlies most social cognitive studies, especially the earliest ones. Later sections address advances beyond this initial approach.

The Basic Cognitive Model of Associative Networks

The most important general principle of this approach is that the more links or associations from other concepts to any given concept in memory, the easier it is to remember that concept because many alternative routes can locate it in memory. The following sections elaborate this point in detail because social cognition research grew out of this approach. Table 4.1 summarizes some of these key features in the first column.

TABLE 4.1. Key Features of Cognitive Memory Models and Related Social Cognitive Memory Models

Associative Memory (declarative memory for content)	Procedural Memory (automatic memory for practiced cognitive processes)	Parallel Distributed Processing (connectionist models)	Embodied Memory (perceptual symbol systems)
Long-Term Memory (LTM) *Storage* (can include *intermediate*, recent memory and *remote*, longer-term memory)			Includes top-down expectations (frames)
(can include *episodic* memory for specific events and *semantic* memory for facts, word meaning, and general encyclopedic knowledge)	(includes all kinds of information)	(includes all kinds of information)	Includes bottom-up, sensory information, introspection, and proprioception
Propositions in a node-link structure	If-then (condition-action) productions	Patterns of connections	Simulators that organize related perceptual symbols
Labeled links (not used in social models)		Inhibitory and facilitative connections	
Process: spreading activation	Process: matching, selection, and execution	Process: adjusting connection strengths	Process: simulates online, in presence of object, and offline, in absence of object
Alternative retrieval routes			
Degree of activation determines retrieval	Current goals and stimuli determine selection and use of procedures	Input patterns and connection weights jointly determine retrieval	Focuses on perception for action
Implicates hippocampus immediately, cortical consolidation in longer term	Productions implicate basal ganglia loop; goals implicate dorsolateral prefrontal cortex		

cortical consolidation in longer term	basal ganglia loop; goals implicate dorsolateral prefrontal cortex

Short-Term or Working Memory	
Most activated portion of LTM but also attended stimuli, may be consolidated for storage in LTM	
Implicates ventrolateral prefrontal cortex	
Fast, easy learning requires STM capacity	Slow learning, requires practice
Automatic or controlled execution	Fast, automatic execution
Widely available use	Specific, focused use
Flexible	Durable

Person Memory Models	Social Procedural Models	Social Connectionist Models	Embodied Social Cognition
Hastie: person memory-I	E. R. Smith: proceduralization	Kunda & Thagard: parallel constraint satisfaction model	Niedenthal et al.: embodied social cognition and emotion
Srull & Wyer: person memory, and judgment	Lewicki: implicit recognition	Van Overwalle: connectionist models	Feldman Barrett: embodied experience of emotion
Marques & Hamilton: twofold retrieval and associative processing model		Kashima: tensor-product model	E. R. Smith & Semin: socially situated cognition.
		Read: Gestalt connectionism	
Carlston: associated systems theory		Vallacher & Nowak: dynamical systems	
		Smith & DeCoster: social connectionism	

EXHIBIT 38 - 3

GaState0033823

The exact format of the representation is called a memory code. A variety of possible codes are discussed later, but the best known in early cognitive psychology was a *proposition* (J. R. Anderson, 1976; Rumelhart, Lindsay, & Norman, 1972; Wickelgren, 1981). For example, "The woman stands on the corner" is one proposition; others are "The woman is elderly," "The man knocks down the woman," and so forth. Each proposition consists of nodes and links: each node being an idea (noun, verb, adjective) and each link being the relation between ideas.

One critical feature of these models of human memory is that they are *associative;* that is, most refer to associations between nodes (the woman) linked to other nodes (elderly). The associative feature of a propositional code has implications for important interpersonal events. Suppose you are called in to give eyewitness testimony on the mugging case (Figure 4.1). The organization of long-term memory into an associative node–link structure means that you will recall related facts together. That is, if you begin by thinking about the woman herself, it may be easier to recall her attributes (e.g., elderly, standing on the corner) than to recall the man's attributes (e.g., young). Moreover, the links are *labeled.* If asked who (agent) mugged whom (object), you would have the proper relationships stored in memory. (This particular labeling feature has not been adopted in social cognition.)

An important feature of propositional memory models is that the links are strengthened each time they are *activated.* That is, recall starts at one point (e.g., the woman), and activation spreads along links between the nodes (Collins & Loftus, 1975). Recalling that the woman was elderly, for example, activates both nodes in memory ("elderly" and "woman"), strengthening the link between them. The practical implication of joint activation is that frequent rehearsal (that is, mental repetition) of your testimony makes it more memorable than unrehearsed facts. The lawyer preparing you to be a witness is likely to know that frequent reviews of the testimony ahead of time will strengthen its coherence and avoid awkward surprises, such as your remembering new events on the witness stand.

In addition to strengthening links among ideas by activating them together, the more separate linkages to any given idea, the more likely it is to be recalled. More links create more *alternative retrieval routes* and enhanced memory. Thus, a smart lawyer will help you to form many alternative memory routes to any given fact to minimize last-minute forgetfulness.

Long-Term Versus Short-Term Memory

A critical feature of many network memory models is the distinction between long-term memory and short-term (or working) memory. In this view, *long-term memory* consists of the vast store of information one can potentially bring to mind. *Short-term memory* refers to the information that one is considering at any given moment, which is why it is also called "working memory." In many memory models, the activated portion of long-term memory represents short-term memory or consciousness. That is, the long-term memory nodes that are currently most active make up the contents of focal attention.

de. A vari-
arly cogni-
rt, Lindsay,
stands on
"The man
s of nodes
link being
ey are asso-
) linked to
has impli-
in to give
nization of
at you will
he woman
ing on the
e links are
the proper
s not been
e links are
point (e.g.,
Collins &
ates both
. between
ersal (that
han unre-
now that
coherence
ts on the
together,
recalled.
ry. Thus,
es to any
between
w, long-
ly bring
sidering
ory." In
repre-
nemory
tention.

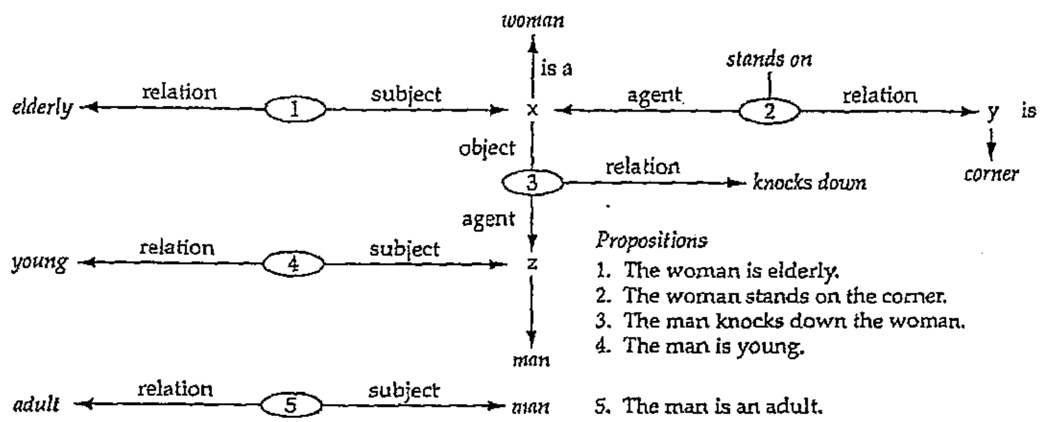

FIGURE 4.1. Propositional network model for knowledge that "The elderly woman standing on the corner is knocked down by the young man." (After J. R. Anderson, 1980a.) Each separate proposition is indicated by an ellipse. For example, the first proposition ("The woman is elderly") is represented by the nodes and links on the upper left of the figure. The numbered ellipses indicate the other propositions. Starting at ellipse 1, anything connected by arrows moving *away* from the ellipse is part of that particular proposition. For example, "elderly" and "woman" both are connected to the ellipse by arrows from it and to them. But "stands on" is not part of the first proposition; accordingly, the link from the ellipse labeled "1" to "stands on" is interrupted by an arrow pointing the wrong way. The other notation that needs explaining is the use of *x*, *y*, and *z*. They indicate that *x* is one particular woman, who is also elderly. If the proposition were "Woman are elderly" (i.e., the entire category of women is elderly), the *x* would be replaced by the word "woman" in the figure. For example, the proposition "The man is an adult" (true of all men) would be denoted differently than "This particular man is young" (see propositions 4 and 5 in the figure). For present purposes, the notation illustrates the precision with which details of meaning can be represented in a propositional network.

The contents of short-term memory can be consolidated for storage in long-term memory.

In associative models, memory retrieval consists of activating the appropriate nodes in long-term memory, which brings them to consciousness if activation is above a certain threshold. Because the most active nodes can change rapidly, the conscious part of long-term memory (that is, what you are thinking about right now) is considered short-term. Things move in and out of consciousness, or short-term memory, as they become activated and fade from being activated. The limited capacity for activation means that short-term memory is quite limited in scope. In other words, few things can be held in mind simultaneously.

The consequences of the rather severe limits on short-term memory can be illustrated by a lawyer's questioning of a witness on the stand. A witness will

GaState0033825

EXHIBIT 38 - 5

be unable to keep lots of details in mind at once, so the person may contradict earlier testimony that is out of current consciousness. Short-term memory is normally assumed to hold about seven items of information (classically, 7 ±2; G. A. Miller, 1956). An "item" of information could be as small as a single letter or digit, or it could be a "chunk" of letters (i.e., a word on the mugger's sweatshirt) or a chunk of digits (i.e., the time on your watch). In practical terms, short-term memory limits mean that people can keep only a few things in mind at once.

In contrast to the limits of short-term memory, the capacity of the overall network of long-term memory storage seems, practically speaking, to be limitless. A lawyer who urges a witness to struggle to remember crucial details may be banking on this: The information might all be there; it is only a question of finding it. For long-term memory, the issue is not capacity (or how much one knows) but retrieval (whether one can find it). Skilled performance depends in part on having efficient cues to relevant material in long-term memory (Ericsson & Kintsch, 1995). Many models of social memory are concerned primarily with how retrieval is influenced by the organization of long-term memory, by the links among items, and by accessibility.

The neat distinction between long- and short-term memory may be breaking down. Recent insights from neuroscience support distinctions among three general time periods of memory. Material may be actively attended in consciousness, in effect in extremely short-term memory, as just described. But two forms of what used to be considered long-term memory appear in distinctions between memory for recent and more remote events, sometimes called intermediate and long-term memory. These complementary intermediate and long-term memory systems allow both (a) rapid learning and recall of specific events, implicating the hippocampus, and (b) slow learning but rapid recognition of patterns, implicating the neocortex (medial temporal lobes) (e.g., McClelland, McNaughton, & O'Reilly, 1995; Norman & O'Reilly, 2003). These complementary learning systems reflect the hippocampus's role in recall of recent events and the higher brain's role in detecting regularities in the environment, slow to acquire and accept, but also slow to undo.

To review (see Table 4.1), associative models of memory share the assumptions that memory consists of nodes for ideas and associative links among the nodes. The associative links are posited to be labeled and strengthened by activation that spreads from node to node. The number of links to any given concept determines the number of alternative retrieval routes and hence the concepts' ease of retrieval. Long-term memory is the practically unlimited store of knowledge one has available, and short-term or working memory is the information one is actively using at any given time. Newer work distinguishes recent versus remote long-term memory.

Associative Network Models of Social Memory

Hastie's Person Memory Model

What do these cognitive models of memory say about social cognition? An example is one associative network model of social memory, PM-1 (for Person Memory; Hastie, 1988b). This model works as a computer simulation, which is an important test of its sufficiency (Ostrom, 1988). In a nutshell, PM-1 predicts extra attention to impression-*inconsistent* material (because it is surprising), resulting in extra associative linkages for those items, increasing their alternative retrieval paths and probability of recall. This is called the *inconsistency advantage*.

According to this model, the encoding process invokes a limited-capacity working memory, which allows one to form links among items; the length of time items stay in working memory depends on their relevance to the current impression judgment. A longer stay in working memory forges more links. Links form between items that are unexpected given the current impression because they stay in mind longer as people try to figure them out. (The model also posits that some links are formed stochastically; that is, randomly.) Subsequent retrieval from long-term memory initiates at a random point and randomly proceeds along pathways formed by interitem links. It terminates with repeated failures to retrieve an item not already retrieved.

Finally, the model also proposes a mechanism for impression formation, simultaneous with memory encoding and storage. The anchoring and adjustment process (described more fully in Chapter 7) essentially provides for an impression that is updated with each new piece of information, based on an equally weighted average of (a) the cumulative evaluations of the items so far and (b) the items currently staying in working memory, including the new item (N. H. Anderson, 1981; Lopes, 1982).

One of this model's strong points is its simultaneous modeling of the online impression formation process, plus memory storage and retrieval processes. For example, when people form impressions online—that is, as they receive information—their impressions result from this gradual process. However, when people do not form impressions online, they have to retrieve information to create an impression from memory. In this post hoc case, their memory for the information will correlate with their impression precisely because the impression is a memory-based construction. In the online, simultaneous impression case, memory is irrelevant to the spontaneous impression formed because different factors determine whether each piece of information is important to the impression versus memorable. People can remember trivia but discount them in the impression, for example.

Srull-Wyer Person Memory Model

Another model of person memory and judgment (Srull & Wyer, 1989) makes fundamentally similar assumptions. Basic processes create an impression from a target's behavior: (a) First, people interpret each behavior in terms of an applicable and accessible personality trait, which summarizes the behavior; this trait-behavior goes into its own unit, sometimes called a storage bin. (b) Next, they evaluate the target person as basically likeable or dislikeable. (c) Then, they interpret the person's behaviors in light of the evaluation, reviewing behaviors that are markedly inconsistent with it; and (d) when asked for a judgment, they prefer to use a trait already inferred in step (a), and simply retrieve that from its storage bin. If none of the already inferred traits are relevant, they will review the remembered behaviors to make the judgment.

The model provides a detailed analysis of known social cognitive processes. For example, early information most influences the evaluation. In other words, first impressions count. As another example, learning first about more positive attributes predisposes one toward a positive impression, a phenomenon called a primacy effect (cf. Asch, 1946; E. E. Jones & Goethals, 1972). In the Srull-Wyer model, an initial evaluation forms as soon as the information is clearly and consistently positive or negative, and subsequent behavior is interpreted in light of this initial information.

Like the Hastie model previously described, an important assumption of the Srull-Wyer model is that evaluatively inconsistent behaviors are considered thoughtfully in comparison to other behaviors. This consideration strengthens links between each inconsistent behavior and the remaining behaviors (compared to the links between each evaluatively consistent or neutral behavior and the remaining behaviors). The inconsistency advantage presumably occurs as people think about the inconsistencies, relating them to each other and to the consistent behaviors, creating associative links among them.

One can infer the point at which participants have formed their evaluative impression by observing impressions as they are being formed. That is, over blocks (clusters of stimulus behaviors), one can observe the start-up and increase of the inconsistency advantage in memory. (Essentially, one would plot the degree of inconsistency advantage in each successive block.) One study required at least five blocks of five behaviors each to show the inconsistency advantage (Srull, Lichtenstein, & Rothbart, 1985); that is, it took 25 behaviors to develop an evaluative impression strong enough to make people notice new information being inconsistent with an overall evaluative impression. This suggests that evaluatively mixed information does not quickly generate a strong evaluative impression, so it would not easily show an inconsistency advantage in memory. The inconsistency advantage may be limited to impressions with very few pieces of inconsistent information against an extraordinarily consistent information baseline.

Similarly, interfering with the formation of interitem links should eliminate the inconsistency advantage in memory; studies in which participants perform a competing online task (even rehearsing the stimulus items!) do not show the same inconsistency advantage (Srull, 1981; Srull et al., 1985). The same is true of participants without the capacity and time to form online links (Bargh & Thein, 1985; Barrett, Tugade, & Engle, 2004; Macrae, Bodenhausen, Schloerscheidt, & Milne, 1999). Finally, although participants show an inconsistency advantage within the main person memory paradigm, they may instead preferentially recall consistent items when they have time to think over their impression afterward, perhaps due to bolstering the overall evaluative direction of the impression (Wyer & Martin, 1986). In other words, over the long term, consistent material enjoys a memorial advantage.

One positive aspect of this model was its ability to summarize vast quantities of research by Srull, Wyer, and their colleagues. One disadvantage of the Srull-Wyer model was its assumption of multiple representations ("storage bins"), all pertaining to a single person. This cumbersome idea was not endorsed by other models of person memory, and it tends to make this model both counterintuitive (not as much a technical flaw as an aesthetic one) and also perhaps overly flexible, able to account for virtually any result or its opposite.

One of the prominent features of these person memory models (both the Hastie and Srull-Wyer versions) is that they predict a recall advantage for impression inconsistent behaviors. The inconsistency advantage is explained by increased attention, linkages, and retrieval routes for inconsistent material. The inconsistency advantage was a robust effect within the standard paradigm of instilling a trait impression, then presenting a series of consistent and inconsistent behaviors and asking participants to recall the behaviors (see Srull & Wyer, 1989). However, the inconsistency advantage is *not* obtained when the research paradigm departs from the standard one in any of several ways that complicate the perceiver's task: well-established expectancies (S. T. Fiske & Neuberg, 1990; Higgins & Bargh, 1987; Ruble & Stangor, 1986), multiple trait expectancies (D. L. Hamilton, Driscoll, & Worth, 1989), behaviors that are descriptively but not evaluatively inconsistent (Wyer & Gordon, 1982), traits and behaviors overheard in a conversation (Wyer, Budesheim, & Lambert, 1990), time to think about one's impression afterwards (Wyer & Martin, 1986), having to make a complex judgment (Bodenhausen & Lichtenstein, 1987; Bodenhausen & Wyer, 1985) under cognitive load or selective memory (Garcia-Marques, Hamilton, & Maddox, 2002), and under multitask conditions (Macrae et al., 1999). Also, when it does occur, the inconsistency advantage occurs for impressions of individuals who are expected to have coherent personalities and not so much for groups, which are viewed as containing less coherent "personalities" (Rothbart, Evans, & Fulero, 1979; Srull et al., 1985).

Marques-Hamilton Person Memory Model

Although the Hastie and Srull-Wyer models can account for many of these results, people probably use multiple processing strategies, depending on the circumstances, as in the dual-process modes described in Chapter 2. Indeed, a newer model suggests that people use distinct retrieval processes, depending on their task (Garcia-Marques & Hamilton, 1996). When they attempt to recall information, they use an exhaustive strategy, but when they want to remember instances to make a frequency judgment, for example, they use a heuristic retrieval strategy. The exhaustive strategy favors memory for inconsistency, whereas the heuristic strategy favors memory for consistency. Consistent with other dual-process models, this is called the twofold retrieval by associative pathways (TRAP) model.

Carlston's Associative Systems Theory

According to associative systems theory (AST), representations of other people develop through the use of four primary mental systems: (a) the visual system, (b) the verbal/semantic system, (c) the affective system, and (d) the action system (Carlston, 1994). These four modalities are relatively independent at lower levels related to immediate encoding but intertwine as representations become more abstract. Many of the mechanisms otherwise relate to other associative memory models. This view explicitly adds other modalities to the associative memory models' primary focus on cognitive modalities, most notably affect and action.

Conclusion

The most influential model of memory in social cognition has been the node-link structure proposed by associative memory models. In this view, each concept (trait, behavior, person's name) is represented as a node, with links formed

GaState0033829

EXHIBIT 38 - 9

by relating two items to each other. Memory retrieval proceeds along the pathways provided by the network. The recall advantage for expectancy inconsistent versus expectancy consistent information, under some circumstances, is one of the major empirical legacies of these models. Newer approaches that build on procedural memory, parallel processes, and social neuroscience go beyond these early controversies.

PROCEDURAL AND DECLARATIVE MEMORY: WHAT MEMORY DOES

The Basic Cognitive Model of Procedural and Declarative Memory

As just described, many models assume that memory includes an associative network of concepts, a long-term store of content knowledge. This *declarative memory* is sometimes contrasted with one form of automaticity described earlier, *procedural knowledge* (Squire, 1987). In Chapter 2 we explained one form of automatic processing as proceduralization—that is, the speed-up of judgments with practice. However, we did not locate procedural processes within a larger model of memory. With some additional background on memory, we can do so, building on recent memory models (e.g., J. R. Anderson et al., 2004). Respectively, they roughly cover the "what" (declarative) and the "how" (procedural) forms of memory.

Newer associative network models posit that declarative memory is activated as a joint function of its general usefulness in the past and its current relevance; together they control both probability and speed of retrieval (Anderson et al., 2004). Retrieval from declarative memory implicates the hippocampus and the temporal lobes. (Recall the taxi drivers with enlarged hippocampi for street location memory in Chapter 1.) Declarative long-term memory includes both *episodic* memory for specific events (Tulving, 1983) and *semantic* memory for facts, word meaning, and encyclopedic knowledge (Squire, 1992). Consistent with earlier ideas about short-term memory, retrieved material occupies a space-limited buffer associated with the ventrolateral prefrontal cortex (Buckner, Kelley, & Petersen, 1999; Wagner, Paré-Blagoev, Clark, & Poldrack, 2001). Related chunks of knowledge activate together, akin to short-term or working memory.

Procedural knowledge concerns skills—namely, how to do things—and is hypothesized to be represented differently from how we have been discussing it. Procedural knowledge is represented as condition-action pairs, or if-then statements, called *productions*. When an input pattern matches the "if," or condition, part of the production, the "then," or action, part immediately operates. For example, a condition might include "if the envelope is addressed to me and advertises a large cash prize." The action part of the production, for some of us, includes "then discard without opening," while for others it must include "then open, read, and respond immediately." Various common cognitive productions (such as matching, selection, and execution) make up procedural memory, which globally implicates a loop involving the basal ganglia (Anderson et al., 2004). People's current goals determine which procedures fire; the goals may be associated with the dorsolateral prefrontal cortex. The current goal may occupy a buffer that serves to track one's progress in achieving an outcome.

Procedural Models of Social Memory

As described previously, content knowledge is initially represented in the declarative associative networks long familiar to social cognition researchers. The advantages of declarative representation include easy learning, for one simply links ideas; wide availability of use, namely, in any situation that cues part of the structure; and flexibility, that is, enabling one to work in various directions among associations, depending on need (E. R. Smith, 1998). Thus declarative knowledge is general and independent of domain and is probably accessible to consciousness and verbal expression. Hence it has been amenable to the methods most commonly used by social cognition researchers.

The disadvantage of declarative knowledge is that it tends to be slower and to use up one's limited-capacity working (short-term) memory. Accordingly, as certain processes are used repeatedly, they eventually may be proceduralized to be more efficient, as noted earlier.[1]

Smith's Proceduralization

Social processes are nothing if not frequent, so it makes sense that some would be proceduralized, as we saw in the examples of automaticity. Specifically, we described the work of E. R. Smith and his colleagues, who have applied the principles of procedural memory in particular to the speed-up of social inferences with practice (E. R. Smith, 1990; E. R. Smith & DeCoster, 1998); as noted, these practice effects provide one explanation for automaticity.

In addition, Smith and his colleagues propose that procedural memory provides an alternative explanation for priming effects, discussed in light of memory models. Recall that priming demonstrates the impact of a recently or frequently activated category on the processing of category-relevant information. Priming effects are typically interpreted in terms of category accessibility, using declarative (associative network) memory. That is, models describe activation as spreading from a prime to relevant concepts along the associative network's pathways. Any linked concept can prime another, whether the priming is based on words, faces, or symbols. The process itself is general, working the same way for any content linked in the associative network.

However, studies suggest that some priming effects can instead be process specific, as a procedural account would suggest (E. R. Smith & Branscombe, 1987, 1988; E. R. Smith, 1989; cf. Higgins & Chaires, 1980). Consider the different specific processes that can prime a personality trait: For example, one might read the word and think about it, or one might generate it from an obvious behavioral instance. When a subsequent task repeats the exact same process (i.e., reading the

[1]Proceduralization is viewed as the second step of a two-step process of *compilation*. First, certain actions that are always executed together may be combined into a single mega-step via *composition*. What began as a sequence of steps ("if envelope is addressed to me and advertises prize, then open it," followed by "if envelope advertises a prize and has just been opened, then check eligibility") may end up as one more complex step ("if envelope is addressed to me and advertises a prize, then open and check eligibility"). The second way to compile certain processes is via *proceduralization*, as when one learns to sort one's mail more rapidly, applying general procedures to specific repeated experiences. When every envelope from a certain place fails to provide an unambiguous, unconditional prize, one quickly learns the procedure: "If envelope is addressed to me from Cash Grand Prizes Clearinghouse, then discard." For simplicity here and because this distinction is not universal, we discuss only proceduralization.

trait or generating it from behavior), a form of priming occurs; namely, a speed-up of subsequent access to the word. This speed-up depends on the specific process (reading or generating), not just on previous exposure to the concept. When the practice task uses *different* priming processes than the test task (e.g., generating at practice and reading at test), the trait concept is just as accessible as when the practice and test tasks use the *same* priming processes (e.g., generate at both or reading at both). Nevertheless, greater speed-up occurs when the same exact task procedure is repeated. Thus some priming effects occur from repeated procedures per se in addition to the category accessibility effects of declarative associative memory. This does not replace accessibility but essentially suggests that specific processes can be primed as well.

Lewicki's Implicit Memory

Other applications of procedural memory may explain the accessibility of certain attitudes (see Chapter 10), the selection of one inference or category from among many possible ones (cf. Chapter 11), and the learning of complex patterns that cannot be articulated (Lewicki, Czyzewska, & Hoffman, 1987; Musen & Squire, 1993; for a review, see Seger, 1994). The more a particular procedure is practiced, the more likely it is to be used again instead of other equally applicable procedures. In this view, such procedural effects are a form of "implicit memory," which is a term for the influence of past judgmental processes on current judgments and reactions (E. R. Smith & Branscombe, 1988; see Jacoby & Kelley, 1987, for another discussion of unconscious memory). Procedural memory ideas have not overtaken social cognition, but their considerable influence appears in work on implicit associations (see Chapters 11–12).

Conclusion

Procedural memory introduces the idea of if-then automatic procedures built up through practice to become automatic processes. Social cognition applications of procedural memory focus mainly on priming specific operations and other instances of implicit memory (see Table 4.1).

PARALLEL VERSUS SERIAL PROCESSING: COORDINATING MEMORY PROCESSES

Cutting across issues of declarative (associative) and procedural memory is the issue of parallel versus serial processing. The traditional associative network models viewed spreading activation as a *parallel* process of activating many related pathways at once, but overall processes of encoding, memory retrieval, and response are generally viewed as a *serial* sequence of steps. Social categories are often seen as activating before the activation of individuating information (as described in Chapter 2). This could occur via either serial or parallel processing. That is, serially, first one is activated, then the other. Or if both are activated in parallel, perhaps the category information is processed faster and beats out the individuating information. The serial processing idea is more prevalent in earlier information processing models, so the newer models focus more on parallel processes, as this section explains.

The Basic Cognitive Model of Parallel Processing

Parallel distributed processing (PDP) is an approach to the structure of cognition that developed as an alternative to more traditional, mostly serial models of mental structure. For reasons that remain obscure, one of us found herself describing PDP to her 80-year-old great aunt, always an astute and intellectually challenging conversationalist, who demanded a sample of "the wave of the future" in cognition research. PDP seemed a safe bet in that the great aunt could be guaranteed to know even less about it than the author did. Precariously launched on an explanation, the metaphor for PDP that came to mind was an old-fashioned time and temperature sign board composed of lightbulbs. Such signs were made of a grid of bulbs, different combinations lighting up depending on the numerals that needed to be displayed. Each lightbulb contributes to all the times and temperatures displayed by being on or off within the overall pattern. In this oversimplified and deeply flawed PDP metaphor, born of desperation, individual memory units are lightbulbs, each unit participating in many different memory patterns, as simply one feature of the whole. The same bulb could be part of the numeral "1" or "2." Moreover, the numeral "2" could appear in different positions on the board, depending on whether the time were 2:00 or 7:32.

Consider how this approach differs from associative network models of memory. In such models, each node uniquely represents a concept, and when it is sufficiently activated, the concept is retrieved. In a PDP model, each unit helps to represent many different concepts, which are retrieved when the appropriate pattern of activation occurs across all the basic units. Thus, to return to the time and temperature sign, the specific numeral "3" could occur on the right, left, or middle of the sign, as needed, depending only on the correct configuration of lightbulbs being on. Thus no single lightbulb represents "3," but instead the pattern does, and which set of bulbs do the job is arbitrary. This differs considerably from a neon sign, for example, that has one structure dedicated to lighting up for one particular number whenever it is needed. Traditional memory models would roughly resemble a series of neon letters linked to each other. (The great aunt was skeptical but intrigued.)

PDP models essentially deal with the subatomic particles of perception and cognition. PDP models assume that memory consists of elementary units (the bulbs in our metaphor) that are connected with facilitative and inhibitory links to each other. The connections represent *constraints* about what units are associated, and the *connection strengths* represent the type and magnitude of association. Only the strengths of connections are stored, so that the pattern can be re-created by activating parts of it and waiting for the connections to reverberate throughout the system until the entire pattern is activated.

The full theory of PDP is beyond the scope of this book (for an accessible introduction, see McClelland, Rumelhart, & Hinton, 1986). For the most part, it has been applied to issues of motor control (typing, reaching) and perception at the level of individual letters in a word. One interpretation of PDP models is that they aim at a different and lower level than the network models of memory described earlier. That is, a node in the network metaphor would be not a single neuron but a pattern of activation over neurons. If one considers PDP only as a lower-level elaboration of network models—that is, operating at the level of neurons—its implications for social cognition's more macro level of analysis would be limited.

Nevertheless, PDP does have potential utility for a social level of analysis. In the more traditional associative models, knowledge is represented statically. That is, knowledge does not change its form between long-term and working memory because it is essentially just more or less activated. In PDP models, however, the patterns themselves are not stored, but the strengths of connections among basic units are stored, enabling the patterns to be re-created. From a practical perspective, this allows knowledge to be implicit in the system rather than being an explicit set of stored rules.

PDP also allows imperfect stimulus patterns to be recognized because approximations can activate part of the pattern of connections, which subsequently generate the remaining aspects of the pattern. PDP models are good at considering several sources of information simultaneously. They are parallel processors, in contrast to the more traditional serial processing models.

Recent models of memory combine serial and parallel processes (e.g., J. R. Anderson et al., 2004). For example, declarative memory retrieval might simultaneously search many related memories in parallel, but the content that comes to mind might form a serial bottleneck because one can attend to only one retrieved memory at a time. Similarly, several potential productions might be activated in parallel in procedural memory, but only a single production could fire at a time, serially.

Parallel Constraint Satisfaction in Social Cognition

One possible domain of PDP application in social cognition is to stereotypes, and in particular to how they simultaneously interact with each other. For example, combining one's knowledge about traditional Amish farmers and progressive Montessori teachers, one can imagine someone who occupies both roles by considering their shared "back to basics" perspective and shared emphasis on patience. Moreover, one can imagine the person's likely response to novel issues (e.g., cell phones in the classroom). PDP models allow for such emergent properties of previous knowledge.

Kunda-Thagard Parallel Constraint Satisfaction Model

A pertinent example in social cognition is a parallel constraint satisfaction theory (Kunda & Thagard, 1996). Introduced briefly in Chapter 2 as a single-mode alternative to the dual-mode models, it views impression formation as similar to text comprehension. The perceiver needs to interpret and integrate a variety of incoming information simultaneously with accessing the relevant knowledge base, which includes representations of stereotypes and traits. The model emphasizes the simultaneity of these processes, balancing the mutual and immediate influence of inputs ranging from the more concrete (e.g., interpretation of specific behaviors) to the more abstract (e.g., application of expectations or stereotypes). All information considered at once constrains the other currently accessed information. For example, a shove coming from a friend may come across as playful, whereas a shove coming from a stranger may come across as violent. The model posits that expectancies and new information constrain each other's interpretation, especially when information is ambiguous. Chapter 11 on stereotyping revisits these issues, but at present the point is that processing can occur in parallel, with a variety of information mutually influencing each other's interpretations. The model works as a computer simulation of impression formation processes.

Van Overwalle's Connectionist Models

The connectionist model of impression formation also applies PDP principles to social cognition (Van Overwalle & Labiouse, 2004). This model goes beyond the Kunda-Thagard parallel constraint satisfaction model because it includes a learning component as well as the perception component. This model also has an initial activation phase in which external inputs (e.g., stimuli) are balanced against internal ones (e.g., expectations). After this activation phase, the connectionist model adds *consolidation* in long-term memory. Consolidation occurs when the external inputs do not precisely match the internal linkages. The model then adjusts the long-term links based on its discrepancy from the short-term input. It's a reality check, in effect. In this model, mismatched expectations should change to fit the most typical input from the environment over time.

Two principles emerge from this computer simulation model. First, *acquisition* reflects the sheer effects of more confirming information over less of it in what is often termed the *set size effect*. People are more certain the more support they have for their perceptions. The other principle is *competition* among the links, whereby the successful (accurate) ones are strengthened at the expense of the less successful (inaccurate) ones. Thus, if the system initially believes that all cats have stubs for tails because its first cat was a Manx, subsequent encounters with fully tailed cats will strengthen the tail belief at the expense of the stub belief.

The computer simulation re-creates standard impression formation effects, such as assimilation and contrast priming effects, the inconsistency advantage in certain kinds of recall, and particular patterns of primacy and recency (greater weighting of early and recent information under specific circumstances). The same connectionist model also applies to a range of social cognition topics: causal attributions (Van Overwalle, 1998), dual-process models of attitudes (Van Overwalle & Siebler, 2005), cognitive dissonance (Van Overwalle & Jordens, 2002), group biases (Van Rooy, Van Overwalle, Vanhoomissen, Labiouse, & French, 2003), and communication (Van Overwalle & Heylighen, 2006).

Kashima's Tensor-Product Model

A related model is the tensor-product model (Kashima, Woolcock, & Kashima, 2000), which uses a Hebbian learning approach instead of the competition approach used in the Van Overwalle connectionist model. A Hebbian approach describes some kinds of associative learning by changes in the strength of links between nerve cells; simultaneous activation strengthens the links, but there is no provision for inhibition of unactivated links. Also, it is not viewed as a literal representation of neural networks. Nevertheless, this model nicely describes several phenomena in forming impressions of groups.

Conclusion

PDP connectionist models have proved popular for computer simulation models of social cognition in various forms. Another connectionist model addresses basic Gestalt principles of causal reasoning, cognitive consistency, and goal-directed behavior (see Chapters 1, 6, 9, and 15; Read, Vanman, & Miller, 1997). Still others address the self-concept (Nowak, Vallacher, Tesser, & Borkowski, 2000), attitude learning (Eiser, Fazio, Stafford, & Prescott, 2003), and perceptions of outgroups (Read & Urada, 2003). This approach also goes by the name dynamical perspective (Vallacher, Read, & Nowak, 2002).

EMBODIED MEMORY: INCLUDING PHYSICAL REPRESENTATION

A Basic Cognitive Model of Perceptual Symbol Systems

PDP approaches originally developed ways to understand perceptual recognition of familiar patterns, for example, a blurred, chipped, or degraded letter. Perceptual approaches in general were not typically viewed as tenable conceptual systems, just as recording systems. However, a newer perceptual theory of knowledge aims to fill that gap (Barsalou, 1999b).

Perceptual symbols encode experience, both external and internal. This form of memory representation incorporates bottom-up perceptual processes, that is, processes that work from direct perceptual experiences and resulting associations in the brain that activate sensory motor areas. Perceptual experiences can include all the senses, plus introspection and proprioception (sense of one's own bodily position from internal feedback). Selective attention picks out features of the environment, for example, not every available stimulus but specific components. This perceptual side captures useful information about edges, colors, movement, temperature, and the like. It is *embodied* in the sense that the information includes both external stimuli (temperature) and bodily experiences (pain). It is also embodied in the sense that it prepares the perceiver for appropriate action (avoidance). In this way, it is a more sophisticated 21st century version of Gibsonian ecological perception described in Chapter 3, as in "perceiving is for doing."

In addition to the perceptual side, this form of memory representation also captures top-down expectations, after perceptual experiences. The perceptual symbol systems (PSS) record the neural activation during stimulus input, but they also reactivate later for conceptual processing. PSS can represent objects in their presence (perception) or in their absence (imagery or conception). *Imagery* differs from conception in being more conscious and specific about the sensory motor representation, as when you close your eyes and visualize your childhood home. *Conception* entails knowing about it without consciously retrieving visual (or other sensory motor) details.

The key components of PSS include the *simulator* that first registers and later re-creates a perceptual experience. Essentially, it is the pattern of brain activation created by selective attention at the perceptual stage. Related perceptual symbols organize into simulators that allow the system to represent specific entities (e.g., what you see the first time to arrive at a new home). The representation is dynamic, changing as you experience more information (walking in, living there over time). The simulator contains two kinds of structures: the underlying *frame* that integrates across experience with that category (that home, homes in general), and the *simulations* that can be created from the frame; simulations create the experience of a particular example on a particular occasion.

In this view, at one extreme cognitive processes involve bottom-up sensory motor perceptions. At the other extreme are top-down sensory motor representations that include conception and imagery. Intermediate are processes such as priming, filling gaps where information is missing, anticipating future events, and interpreting ambiguous information. These all entail complementary processes that are part perceptual and part top-down. The PSS view differs from the associative network models in a particularly important way: The perceptual symbols represent sensory modality, whereas the earlier models were *amodal* in the sense that memory represented abstract structures regardless of original type of sensory or internal input. PSS represent abstract concepts by the integrative frames across specific experiences, selecting the core features, and by incorporating both physical and introspective experiences. From the PSS perspective, working memory runs the perceptual simulations (e.g., as a buffer containing information just experienced). Long-term memory records the experience and relevant simulations. We have described this model at some length because it may prove useful to social cognition researchers in the future, as some current work suggests.

Social Cognitive Models of Embodiment

The PSS view appeals to social cognition settings because of its focus not just on archiving information but on preparation for situated action (Barsalou, 1999b). Social interaction is nothing if not situated action. The main intellectual message of social psychology is that social situations greatly influence thoughts, feelings, and behaviors. Social psychologists thus are adopting the embodied cognition viewpoint, as anticipated by Zajonc and Markus's (1984) description of the "hard interface."

Embodied cognition places the actor squarely in interpersonal context, in socially situated cognition (E. R. Smith & Semin, 2004). Embodiment plausibly underlies social information processing both in direct perception—*online cognition* during interaction—and in the absence of the considered social object—*offline cognition* (Niedenthal, Barasiou, Winkielman, Krauth-Gruber, & Ric, 2005). PSS applies to several recent findings involving effects of facial or bodily positions on social cognition. For example, when people are induced to nod their head vertically or shake their head horizontally while receiving a persuasive communication, they are, respectively, more or less likely to agree with it (Wells & Petty, 1980). When people hold a pencil horizontally in their teeth—making them contract their zygomaticus smile muscles—they find cartoons funnier than when they hold the pencil with their lips (Strack, Martin, & Stepper, 1988). People like unfamiliar Chinese ideographs better when they flex their arms, mimicking an approach motion, than when they extend their arm, making an avoidance motion (Cacioppo, Priester, & Berntson, 1993). When people adopt different postures characteristic of emotions such as sadness, they report feeling those emotions more (Duclos et al., 1989), and making a fist relates to people's conceptions of power (Schubert, 2004). All these examples suggest an effect of embodiment on emotions and evaluations.

In Chapter 3 we saw the reverse: Conceptual priming through words or visual priming through faces makes people enact relevant stereotypic behaviors (Bargh, Chen, & Burrows, 1996). But the effects of bottom-up sensory motor representation and top-down generalized concepts and images are reciprocal between embodiment and affective experience: People feel particularly understood when they imitate each other's nonverbal behavior during interaction (Chartrand & Bargh, 1999). Empathy operates partly through bodily channels (Decéty & Chaminade, 2003). The representation of an emotional experience is especially likely to occur through embodied knowledge about the emotion (Barrett, 2006).

Conclusion

Memory models often build on the idea that associative networks organize concepts in long-term memory and that the currently most activated associations represent short-term or working memory. This framework underlies most work on social memory. When researchers turn to memory in action—what it does—they distinguish declarative memory (the associative networks of what is recalled) and procedural memory (the if-then pairs that determine how recall triggers other operations). Memory models increasingly focus on simultaneous parallel procedures to coordinate memory processes, not just serially sequenced processes. And models of embodied cognition rely on sensory motor perceptual systems for both bottom-up perception and top-down conceptual representation.

SOCIAL MEMORY STRUCTURES: WHY SOCIAL MEMORY MATTERS

"I love you, Jane . . . I mean, Sally." What happens when you call someone by another person's name? People typically are annoyed at being treated as equivalent to the other person (or as less cognitively accessible than the other person). And well they should be annoyed, although we know of worse cases in which people confuse their children with the dog (A. P. Fiske et al., 1991). What happens when you forget someone's name or their confidences? Feelings are hurt, and personal tragedies can ensue. People use broad expectations and categories to organize their memory for other people. This is often functional, but it has its drawbacks. Social cognition research has long focused on this tension between the general and the particular. In this final section we discuss mental representations of social categories and specific social exemplars because these structures cut across different memory theories, and all are crucial to social life.

Social Categories

Categories describe our expectations about, for example, people, entities, or social groups. Like it or not, we all make assumptions about other people, ourselves, and the situations we encounter. Sometimes we are dramatically misled by our expectations. However, much of the time our expectations are functional, and indeed, we would be unable to operate without them. Such expectations, assumptions, and generic prior knowledge allow us some sense of prediction and control, which is essential to our well-being (Chapter 2).

Could we do without categories? Consider the seemingly objective alternative of operating within situations and with people about whom we have virtually no expectations or prior knowledge. Arriving on a new campus the first day, coming into an unfamiliar culture for the first time, or meeting a stranger whose gender, age, and role are mysterious—these all are disorienting encounters that challenge our ability to function without the normal level of prediction and control provided by expectations. Prior knowledge about the campus (a map, for instance), guides to the culture (from travel books), or an introduction to the stranger (by a mutual friend) would facilitate each encounter. Of course, our inevitable reliance on such prior knowledge is not perfectly adaptive. We may rely on the wrong assumptions, or our assumptions may be overly rigid. But on the whole, such expectations are useful.

Category-Driven Versus Data-Driven Processes

Categories represent knowledge about a concept; sometimes termed a *schema*, such an abstract representation includes the concept's attributes and the relations among them (S. T. Fiske & Linville, 1980; Macrae & Bodenhausen, 2001; Rumelhart & Ortony, 1977; S. E. Taylor & Crocker, 1981). Categorical person perception facilitates what is variously termed top-down, conceptually driven, or theory-driven processes, which simply means processes heavily influenced by one's organized prior knowledge as opposed to processes that are more bottom-up, stimulus driven, or data driven (Abelson, 1981; Bobrow & Norman, 1975; Rumelhart & Ortony, 1977). As people's theories and concepts about the world, categorical perceptions are concerned with the general case, abstract generic knowledge that holds across many particular instances. The basic message of this research has been that people simplify reality by storing knowledge at a molar, inclusive level rather than squirreling away, one-by-one, all the original individual experiences in their raw forms, which would be pure data-driven processing.

Recently, however, social and cognitive psychologists have taken a closer look at data-driven processes by focusing on their interplay with category-driven processes. Data-driven processes demonstrate ordinary people's sensitivity to the specific qualities of another individual or situation. Purely categorical theories have, in the extreme case, portrayed people as blithely glossing over important details, as stubbornly refusing to see the information in front of them, and as maintaining their schemas at any cost. In contrast, data-driven approaches show that people do indeed care about the information given (Higgins & Bargh, 1987). We consider both types of phenomena in this section, beginning with categorical processes.

As noted, our perceptions of the world reflect an interplay between what's out there and what we bring to it. We are, paradoxically, more aware of the contributions of the world out there than of our own contributions to our cognitive processes. That is, we know we are encoding information, but we underestimate the roles of selective attention, interpretation, and gap-filling. Expectations are structured knowledge that we bring to everyday perceptions, so expectations emphasize our active construction of reality. This is not to say that we are unconstrained by the stimuli themselves; contrary to Gertrude Stein, "there is a there there."

Categorical expectations emphasize the part of the perceptual interplay—our own contribution—that is mostly preconscious. We experience the world as if we have added nothing to it, so common sense tells us that we perceive an unchanged or literal copy of the environment. We experience perception as instantaneous and direct, as if our brains were simply videotaping the surroundings. Both ordinary people and some philosophers have held this commonsense view that perceptions are unfiltered and veridical (Aristotle, 1931; J. Mill, 1869). An emphasis on the importance of less-filtered experience continues in the present-day study of exemplars (covered later in this section).

In contrast, Gestalt psychology encouraged a different view of perception (Brunswik, 1956; Koffka, 1935). As we noted in Chapter 1, Gestalt psychologists argued that perception is constructive and that perceptions are mediated by the interpretive faculties of nerves and brain. What we "see" in any given stimulus depends on context; for example, the "l"s in "1952" and in "life" objectively are quite similar, but we interpret them differently because of their respective contexts. Context provides a different Gestalt, or configuration, that alters the meaning of the individual elements. Hence, the whole is more than the simple combination of its parts. The Gestalt emphasis on people's perception of configurations in context anticipates social categories and expectations as configurations-in-context actively contributed by the perceiver. This organized prior knowledge shapes what is perceived and remembered in much the same way that context-based Gestalt configurations do, but generally as more complex types of configurations involving people and situations. Gestalt stimulus configurations guided two theoretical developments that directly precede current categorical theories: Solomon Asch's (1946) configural model of forming impressions of others (Chapter 1) and Fritz Heider's (1958) theory of social configurations that produce psychological balance (Chapter 9).

Categorical Perception

The commonsense view and the classical view have been that one can precisely define the boundaries of everyday categories (E. E. Smith & Medin, 1981), just as one imagines being able to do in science or mathematics (although even there the classical assumption is questioned). On closer examination, this proves not to be possible. Building on principles first noted by Wittgenstein (1953), several principles derived in cognitive psychology (Mervis & Rosch, 1981; Rosch, 1978, 1987) and social-personality psychology (Cantor & Mischel, 1979) describe how people categorize things, situations, and other people.

One core notion is that natural categories do not have necessary and sufficient attributes. Instead, the category members fall within fuzzy boundaries, so it is not always clear which instances belong in the category. For example, Monopoly, baseball, and charades are good examples of the category games, but what about playing house, torturing ants, and betting on the Superbowl? The perception that some instances are more typical than others led to the idea that instances range from being quite typical to atypical, with a most typical or prototypical instance best representing the category. The *prototype* is the "central tendency" or average of the category members.[2]

People may never actually encounter their prototypes in real life because they are abstracted from experiences with examples. Even though none of the instances may itself be a perfect prototype, people abstract out the most typical features (Hayes-Roth & Hayes-Roth, 1977; Posner & Keele, 1968, 1970; Reed, 1972). People then decide if a new instance fits the category by assessing its similarity to the prototype.

In this view, category members are related by the criterion of family resemblance. Any given pair of category members will share some features with each other and other features with other category members. For example, Twenty Questions and baseball both include a certain number of turns (the 20 questions or the nine innings), whereas tag and Monopoly do not. On the other hand, Twenty Questions and tag are both played without specific equipment, whereas

[2]This central tendency could come from the arithmetic mean or mode, depending on whether the particular dimension is continuous (e.g., height) or discrete (color).

baseball and Monopoly both require specialized equipment. The more features an instance shares with other category members, the more consistently, consensually, and quickly it is identified as a category member (McCloskey & Glucksberg, 1978; Rosch, 1978; E. E. Smith, Shoben, & Rips, 1974). Thus any given feature is not necessarily present in all category members, only more or less probably so. The point is that the internal structure of categories is more fluid than the classical view would have it, and that it can be well described as a fuzzy set centering around a prototype.

Moving from within-category to between-category structure, categories are often thought to be organized hierarchically at varying levels of inclusion. That is, under the broad category "entertainment," one might include (at least) games, parties, television, books, and movies. Under each of these subcategories, one would have several more subordinate categories, such as car games, board games, outdoor games. Different levels of categorization are useful for different purposes (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). For example, people propose "let's play a game" or "let's go to a movie" (intermediate level) more often than the generic "let's have some entertainment" and more often than the specific "let's play a car game." In this view, such basic-level or intermediate categories for objects are rich in the attributes people associate with them, are easily distinguished from related categories, and involve well-practiced everyday behaviors.

Categorical Person Perception

Just as we categorize different kinds of things and activities, so we also categorize different kinds of people, often according to their personalities. Suppose you see a list of another person's prototypically extraverted attributes, such as being energetic, entertaining, and friendly. Later, you may be unsure whether you also saw other prototypical attributes, such as outgoing and lively (Cantor & Mischel, 1979; Tsujimoto, 1978). Activating some of the attributes activated related ones, so it is hard to remember which you saw and which you inferred from the prototype. This fits the view that people seem to extract a trait prototype from exposure to category-consistent information. Social categories thus may be viewed as fuzzy sets (without rigid boundaries) centering on a prototype.

One implication of this view is that category-based thinking can generate false memories. People store the general prototype and fit the gist of new information to it. Hence, they may remember category-consistent information that never occurred. In one study (Macrae, Schloerscheidt, Bodenhausen, & Milne, 2002), participants saw 60 first names, half male and half female, paired equally often with the gender-typed occupations of mechanic and hairdresser; that is, half were gender consistent and half gender inconsistent. Later, they had to recognize which names they had previously seen, out of a batch of 120; if they said the name was previously seen, they also had to state its occupation. As is typical with simple materials, and as explained by several person memory models described earlier, people more easily recalled expectancy–inconsistent names (male hairdressers and female mechanics).

More interesting from the categorical perspective were the false memories. False alarms mistakenly identify distracter items on the test as being part of the original set. These false alarms were twice as likely to be expectancy consistent as expectancy inconsistent. That is, when they did falsely identify a name as familiar, they most often identified those names as having the gender-consistent occupation. What's more, these false memories were accompanied by a feeling of knowing, so when people falsely recognized a name and attributed a gender-consistent occupation, they were more likely to have a feeling of familiarity for the gender-consistent than gender-inconsistent cases. All this especially occurred when people were either distracted or elderly.

The implication is that people can experience a feeling of knowing when a memory is entirely false, and that category-based consistency underlies this phenomenon. Recall the opening example about witnessing a mugging. Do you remember the older woman, the busy intersection, the parked cars, and the traffic light? What about the mugger's sweatshirt, mentioned later? And his cap? What did he steal? Her purse, you answer. Right. What color was his cap? Where were the parked cars? Well, actually, we never mentioned either the cap or the parked cars, but in dozens of controlled experiments, people are led to experience clear and persistent false memories for events that never happened (Loftus, 2004). Lawyers can plant false memories by asking leading questions ("Did you see the parked cars?" when there weren't any). This throws into question, for example, allegedly repressed and recovered memories that may result from other people's suggestions (Loftus & Davis, 2006). Importantly, the suggestions have to fit with known categories of experience that are plausible in retrospect. Various forms of categorical processing could account for these effects.[3]

Critiques of Category and Prototype Views

Social categories have well-established effects, but the models are changing as the field progresses. First, one modification of the basic approach suggests that social categories are more often represented by ideals or extremes (Barsalou, 1985; Chaplin, John, & Goldberg, 1988). That is, the best example of a nun may in fact be the ideal nun rather than an average nun. In this view, ideals are especially likely to be used for goal-oriented categories with a particular purpose intrinsic to them, in which the best example is one that most completely fulfills the category's goal. It may also be that people have separate representations of the best category member and the average category member; if you ask "Could you suggest the best example of a local restaurant?" my interpretation of the question may lead me to recommend either the best restaurant or the most typical (not at all the same!).

A more drastic critique of prototypes rejects altogether the idea of a summary representation (ideal *or* typical), arguing that categories are represented as a collection of exemplars previously encountered. As it is rather a strong critique, we will discuss this perspective later. For now, the point is that prototypes may not be the only way that categories are stored.

Second, social categories also had been viewed as being automatically activated upon perceiving an instance. The original views held that categorization was essentially automatic and inevitable, but later views have interpreted category activation as conditionally automatic, depending on various factors including one's goals (Macrae & Bodenhausen, 2000; also see Chapter 2).

[3]A failure of *source-monitoring* (remembering the origin of a memory; Johnson, 2006) also contributes in another way. The idea suggested by the lawyer fits expectations about the event, and also one forgets where the idea came from. Failure of source-monitoring accounts for some forms of inadvertent plagiarism (Macrae, Bodenhausen, & Guglielmo, 1999).

Third, recall that the basic or intermediate levels of nonsocial categories apparently dominate people's everyday usage; the same may not be so true for social categories (Holyoak & Gordon, 1984). In using social categories, people's specific goals and expertise most likely determine which levels they choose to use (Cantor & Kihlstrom, 1987; Hampson, John, & Goldberg, 1986). This raises the further possibility that goals and expertise also determine people's choice of levels in nonsocial categories. In both social and nonsocial categories, however, some levels may function as default options, all else being equal.

A fourth issue is the extent to which social categories form a clear hierarchy, with higher levels subsuming lower levels. Social categories had been viewed as being hierarchically organized, with categories becoming more inclusive as they become broader (Cantor & Mischel, 1979); that is, people can list more instances and more attributes that go with broader categories (Goldberg, 1986; Hampson, Goldberg, & John, 1987). This overall approach proved initially useful but provoked several critiques.

Many social categories are not so neat, representing a fuzzy hierarchy in which class inclusions do not work strictly. For example, some data indicate that refusing to share (behavioral level) is a subset of being stingy (lower level), which is clearly one specific way of being unkind (upper level); all this would fit a hierarchy, but not all traits are structured so neatly. People report, for example, that being passive (more concrete, lower level) is not necessarily one specific way of being introverted (more abstract, upper level). Accordingly, the vertical connections between upper and lower levels are not always clear-cut (Hampson et al., 1986). Moreover, for some categories such as gender, the upper levels (male) do not provide richer associations than do the lower levels (the stereotypic businessman) (Deaux, Winton, Crowley, & Lewis, 1985), contrary to hierarchical predictions.

Perhaps the hierarchy notion needs to be discarded altogether; people may not think in hierarchies at all except when psychologists constrain them to do so. People may actually make associations in complex networks that resemble a tangled web rather than a hierarchy (Cantor & Kihlstrom, 1987). Imagine a neat hierarchy, from extravert at the top level, to politician, comedian, and bully at the intermediate level, with specific attributes (socially skilled, self-confident) at the lowest level. In a strict hierarchy, of course, some attributes are true of the superordinate level, and they are inherited by each of the lower levels. However, people's associations are not always so neat (Andersen & Klatzky, 1987). Attributes perceived to hold for the top-level category (e.g., extraverts are self-confident) do not always hold for all the intermediate-level categories (politicians are self-confident, but comedians and bullies are not). Moreover, some intermediate-level categories (bullies) are associated with other intermediate-level categories (politicians) but not with the overarching category (extravert). The overlapping nature of social categories differentiates them from object categories (Lingle, Altom, & Medin, 1984), at least as nonsocial categories have been studied so far.

To summarize, the prototype approach has introduced several ideas: namely, that social categories do not have rigid boundaries but rather operate as fuzzy sets, organized around a prototypical or average category member, against which other members are judged for typicality, so that category members are related by family resemblance rather than by necessary and sufficient rules for inclusion. However, some social categories may be organized around ideal or extreme cases rather than prototypes. The prototype approach originally proposed that social categories are arranged hierarchically, with a most useful intermediate level functioning as basic. But organization in terms of hierarchies and basic levels has been questioned, suggesting that social categories are related in more flexible and complex ways.*

Usage of Social Categories

Current research views social categories as conditionally activated and applied, though everyday life often meets the conditions that activate categories, as we indicated in Chapter 2. Categories can be activated, applied, and even inhibited, depending on social conditions (Macrae & Bodenhausen, 2000).

Category activation depends on attentional resources (Gilbert & Hixon, 1991). That is, under some rare circumstances, people may not notice another person's race, gender, or age. Under extreme cognitive overload, perceivers may identify the person's category, but they may not activate the associated stereotypes. *If the stereotypes are indeed activated, however, cognitive load increases* the odds of applying them.

People belong to multiple categories, and when one is activated, others are inhibited (Bodenhausen & Macrae, 1998). Category salience, chronic accessibility (Chapter 3), and processing goals (Chapter 2) determine which categories are activated and which are inhibited. People sometimes can inhibit category-congruent information intentionally, although this is costly to attentional resources (Macrae, Bodenhausen, Milne, & Ford, 1997). Relevant information can also be inhibited inadvertently. When people repeatedly retrieve one kind of information, related but not retrieved information is inhibited, a kind of adaptive forgetting that enables the mind to focus on what is currently relevant (MacLeod & Macrae, 2001).

Even if activated, categories may be applied in different ways, or not applied at all. We will come back to this point, but the point here is that category activation and application depend on a variety of conditions.

Exemplars

Cognitive Models of Exemplars

Just as the prototype view of categories was developed in reaction to the shortcomings of the classical view, so the exemplar view originally developed in reaction to the shortcomings of the prototype view (see E. E. Smith & Medin, 1981, for a review), and it may be best understood in that light. As a counterpoint to the prototype perspective, the exemplar approach suggests that one remembers separate instances or exemplars one has actually encountered rather than some average prototype one has abstracted from experience. One has several

*Besides the criticisms specific to social cognition just discussed, the prototype approach has been criticized on three basic points: (1) people know more about a concept than simply its features, for they also know which features tend to go together; for example, birds that sing tend to be small and to fly, but large birds tend not to sing or fly; (2) people also know the constraints, that there are limits on what features *may be posited; for example, one does not posit how many pistons a bird has; and* (3) context matters; for example, "bird" in the woods suggests a different sort than "bird" in a nursery school classroom (Lingle et al., 1984; Murphy & Medin, 1985; E. E. Smith & Medin, 1981).

exemplars {
ing whether
Although tl
by many re
understand
The exe
for people'
the categor
numbers o.
specific me
asserts that
ample to re
of course, c
dispute tha
not the foc
More i
category n
sameness
resent info
only the c
to model
knowledg
be possibl
Anoth
ment for i
For exam
tend to gc
ing throu
convenie
1985). A s
knowledg
especially
about wh
are both a
retrieve e
Final
new insta
another e
(Lingle e
but we w

Social C
The soci
preted i
that peo
example
player n
diagnos
advocat

GaState0033844

EXHIBIT 38 - 24

exemplars for a category and, in this view, people categorize a thing by seeing whether it resembles a lot of remembered exemplars from a single category. Although the prototype and exemplar view have subsequently been integrated by many researchers, especially those in social cognition, it is useful first to understand the pure exemplar view.

The exemplar view has several advantages. First, it most directly accounts for people's knowledge of specific examples that guide their understanding of the category. To illustrate, if asked whether restaurants typically contain equal numbers of tables for two or tables for more, one may have to consult several specific mental examples of restaurants to answer the question. Or if someone asserts that all restaurants have cashiers, one may retrieve a specific counterexample to refute the statement. This reliance on concrete instances suggests but, of course, does not require the idea of exemplars. (The prototype view does not dispute that people can remember some specific instances, but the instances are not the focus of the prototype view.)

More important, people know a lot about the possible variability between category members; consider the big variety of Chinese restaurants versus the sameness of a particular fast food chain. A true prototype theory cannot represent information about variability; a pure prototype model would represent only the category average, not its variability. Exemplars provide a simple way to model variability across category members. It is easy to describe people's knowledge of such variation by positing exemplars, although equivalents might be possible within the prototype approach.

Another major advantage of the exemplar view, and probably the best argument for it, is its ability to account for correlations of attributes within a category. For example, people know that within the category restaurant, formica tables tend to go with paying the cashier directly, and tablecloths tend to go with paying through the server. They know this through their theories about inexpensive, convenience-oriented businesses as compared to other kinds (Murphy & Medin, 1985). A single summary prototype does not easily handle this knowledge. The knowledge of which attributes tend to be correlated among category members is especially important in social perception because people have implicit theories about which traits go together (e.g., based on remembering several friends who are both ambitious and hard-working) and which do not (e.g., not being able to retrieve examples of people who are both ambitious and kind; Schneider, 1973).

Finally, the exemplar view makes it easier to modify existing categories with new instances. In the exemplar approach, the new instance may be added as yet another exemplar, which will then contribute to subsequent category judgments (Lingle et al., 1984). In comparison, it is less clear how prototypes are modified, but we will return to this issue.

Social Cognitive Models of Exemplars

The social evidence for exemplars is growing. For example, although not interpreted in terms of exemplars, one series of studies (Gilovich, 1981) showed that people's judgments are affected by irrelevant similarities to specific past examples; a new football player from the same hometown as a famous football player may be considered a good bet even though hometown is not especially diagnostic of athletic ability. People's political judgments are also shaped by advocates drawing lessons from history, even when the prior example is similar to the current one only in nondiagnostic ways. As another example, people's judgments about a member of a foreign culture are more heavily influenced by the foreigner's irrelevant similarity to a previously encountered exemplar rather than by general rules they had learned about the culture (Read, 1983, 1984, 1987). If people have causal theories (e.g., that the exemplar somehow influences the new instance), they are even more likely to use the exemplar. And in more everyday settings, people deciding which of two strangers is more approachable will choose the one most superficially similar to another stranger who was recently kind to them (or, conversely, least similar to a recently mean person), without apparently being aware of the reasons for their choice (Lewicki, 1985).

Familiarity, similarity to a previously known individual, whether consciously perceived or not, may create a sense of shared attitudes, attraction, predictability, and safety (S. T. Fiske, 1982, pp. 62–66; Genaro & Cantor, 1987; White & Shapiro, 1987). Familiarity may be a required mechanism for the operation of exemplars. When people encounter unfamiliar compound categories (e.g., male elementary school teachers), they rely on memory for specific instances, whereas when they encounter familiar compound categories (e.g., female elementary school teachers), they rely on abstract stereotypes. It is the familiarity per se that changes the judgment strategy (Groom, Sherman, Lu, Conrey, & Keijzer, 2005).

Although some of these studies were explicitly based on exemplar models from cognitive psychology and some were not, they all demonstrate the effect of single, concrete, prior experiences on subsequent judgments and behavior, and as such they provide a counterpoint to the influence of more global, abstract categories.

Some more direct discussion of the evidence for exemplars comes from work on the perceived variability of social categories. For example, people perceive increased variability in groups as their impressions become more differentiated (Linville, Salovey, & Fischer, 1986). As you learn more about a particular group of foreigners through an exchange program, you perceive them to come in more different varieties than you did beforehand (and they view citizens of your country as more variable as well). Linville, Fischer, and Salovey (1989) have explicitly argued that this effect can be best described by exemplar models. Some effects of abstract, category-level information can be explained in terms of pure exemplar models to show that the abstract representations are not necessary, although they certainly may be involved (E. R. Smith, 1988).[5] People do seem to use both abstract and exemplar information when provided; that is, they consider both their generalizations about the other nationality as well as specific citizens they have known. Moreover, when people are first given abstract information, followed by information about specific instances, they perceive less variability and make their judgments based on such prototypes more than when the order is reversed or when they are given no abstract information (Medin, Altom, & Murphy, 1984; Park & Hastie, 1987; E. R. Smith & Zarate, 1992).

Problems

Evidence for exemplars is not at all clear-cut. That is, people do understand that some groups are more variable than others, and they use this information, first, in deciding whether to generalize from an individual to the group and,

[5]One may also reinterpret earlier work on the retrieval of behavioral instances versus global traits as compatible with this view (e.g., Lingle, Geva, Ostrom, Lieppe, & Baumgardner, 1979; see Chapter 4).

second, in classifying new individuals. This might seem to argue for exemplars, but knowledge about variability does not seem to be based on memory for exemplars (Park & Hastie, 1987). Similarly, people often perceive minimal variability in outgroups (i.e., groups of which they are not a member), but differences in exemplar frequencies may (Linville et al., 1986, 1989; Ostrom & Sedikides, 1992) or may not (Judd & Park, 1988) be responsible. Researchers attempt to clarify the role of exemplars.

One intriguing possibility is that people may be most likely to use exemplars when they are trying to account for something out of the ordinary. Sometimes we need to know whether something that has just happened is normal or not; that is, we need an immediate check on how surprising it was (assuming we do not already have a relevant schema). For example, when involved in an accident, people think of similar accidents in their past and of the events leading up to the incident to judge how surprising or avoidable it was, and even how upset to be. An elegant and explicit model of exemplars was developed to describe just this process of people's post hoc normality judgments (Kahneman & Miller, 1986). Unlike category models, which focus on anticipation and prediction based on what seems typical or probable in the future given one's abstracted prior experience, norm theory focuses on post hoc interpretation based on an encounter with a particular stimulus in a particular context, with the aim of judging whether the stimulus was normal or surprising. Category and schema theories describe reasoning forward, norm theory describes reasoning backward. According to norm theory, people consider a particular stimulus in light of exemplars it brings to mind. These exemplars allow people to compare the instance to the sum of the previous experiences to see the degree to which the instance is normal or surprising. People compute this sum on the fly, so it is ad hoc rather than explicit prior knowledge. We return to this model in Chapter 7, but for present purposes norm theory illustrates the uses of exemplars in active judgment processes.

Prototypes or Exemplars: A Resolution

Exemplar models are not sufficient by themselves any more than prototype models were sufficient by themselves. People rely on a mixture of representations (cf. M. B. Brewer, 1988; Cantor & Kihlstrom, 1987; J. B. Cohen & Basu, 1987; Groom et al., 2005; Lingle et al., 1984; Linville et al., 1989; Messick & Mackie, 1989; E. E. Smith & Medin, 1981). People clearly do recall specific instances and use them to classify new instances, but specific instances also give rise to category generalizations that in turn facilitate classification of new instances, so people are using both (Elio & Anderson, 1981). People can rely on direct experience with exemplars or on previously provided prototypes to classify new instances, depending on the task and the information available (Medin et al., 1984).

Moreover, because use or development of abstract representations depends on the demands of the task, abstraction of a prototype is not automatic (Whittlesea, 1987). Indeed, exemplars may be more basic (and therefore more likely to be automatic) because they are used (a) when people's cognitive capacity is strained, (b) for more complex concepts, and (c) especially by younger children (Kossan, 1981; cf. Kemler Nelson, 1984).

Exemplars may be more basic building blocks for abstract generalizations such as categories, but once the category is established, exceptions require unpacking the category to return to the more concrete individual exemplar level. That is, cognitive generalizations may start with exemplars and return to them for troubleshooting. As we saw earlier, whereas familiar groups encourage abstractions, unfamiliar groups encourage exemplar usage (Groom et al., 2005). A summary-plus-exception model captures this idea in individual impression formation (Babey, Queller, & Klein, 1998). People generalize across individual instances of behavior and retain both the summary and the instances. Trait judgments rely on both the summary and the specific exceptional episodes (Klein, Cosmides, Tooby, & Chance, 2001).

Exemplars might be useful when the summary abstraction is insufficient for reasons besides exceptions to the rule. When people are motivated or simply have more information, exemplars provide elaborated processing. People apparently use both exemplars and prototypes to represent groups to which they belong, but only prototypes to represent groups to which they do not belong and about which they therefore know less (Judd & Park, 1988). People could also use exemplars to represent both their own and other groups, but they have more exemplars available for their own group (Linville et al., 1989).

Clearly, people can use either abstract category-level information, such as prototypes, or instances and memory for exemplars to make categorical judgments. Which people do so and when is the interesting question (cf. M. B. Brewer, 1988; Park & Hastie, 1987; S. J. Sherman & Corty, 1984, pp. 237–245; E. R. Smith & Zárate, 1992). When people do each doubtless depends on task demands and individual differences. For example, the capacity and the motivation to be accurate or to focus on individuals would probably encourage exemplar-based processes over prototype-based processes (S. T. Fiske & Neuberg, 1990; Kruglanski, 1990; Messick & Mackie, 1989).

When all is said and done, what ultimate use are fuzzy concepts and concrete exemplars in loose hierarchies in tangled webs? Cantor and Kihlstrom (1987) argue that this framework (a) captures the social perceiver's need to represent both the gist of a category and its variability, allowing an economical, functional core representation as well as acknowledging the variability of instances within the category, and (b) describes the multiple paths people use in responding flexibly to the fluidity of social interaction.

Summary

This chapter on mental representation focused on memory; that is, what stays in our heads. We began with associative networks that organize memory in the basic cognitive model, with the basic distinction of long-term versus short-term memory. Associative network models of social memory build on this work. In addition, procedural memory informs some models of social memory. Finally, cognitive models of parallel versus serial processing for coordinating memory processes inform ideas about parallel constraint satisfaction in social cognition.

Social memory structures matter to social cognition, as we saw in discussing the unique features of categorical person perception and uses of social categories. In response to critiques of classic category and prototype views, the exemplar view, with its own advantages and disadvantages, offers an alternative and a combined resolution.

We have focused here on general principles of mental representation that will be useful as we encounter representations of self, causality, attitudes, and stereotypes. Just as in Chapter 2 on dual-mode models, these ideas about mental representation are converging on a consensus useful to the wider field of social cognition.