PART I

Cognitive Representations of the Social World

Introduction

The selections in this Part are representative of current theory and research on cognitions *about* the social world. Within social psychology, mental representations of the self, other persons, and social groups are particularly important topics of study. Our readings include reviews of research on how such representations are constructed and stored in memory (Smith & Queller; Operario & Fiske), the structure and content of these social representations (Smith & Queller; Operario & Fiske; Collins & Allard), when they are activated and used in processing new social information (Martin, Strack, & Stapel; Operario & Fiske), and how they are changed by new learning (Rothbart).

Several interesting themes run through the various readings, reflecting current foci of interest in the social psychological literature. The first is the role of unconscious processes in our understanding and interpretation of the social world (see especially, Banaji, Lemm, & Carpenter; Oakes). The second – and related – theme is the role of affect and motivation in social cognition (Collins & Allard; Martin, Strack, & Stapel). Representations of the social world are inevitably laden with affect and emotion. Social categorization involves dividing the world into ingroups and outgroups – "us" and "them" – with associated motives to protect and enhance what is associated with the self. In addition to self-esteem motivations, human cognition is influenced by motives to achieve meaning and coherence (Oakes; Operario & Fiske) and to preserve important attachments to others (Collins & Allard). Each of these motives affects what information is attended to, how it is interpreted and remembered, and when it is accessed from memory.

Finally, a third important theme in the study of social cognition is the mutual relationship between mental representations of the world, acquired from past experience and learning, and the processing and interpretation of new information and social experiences (see especially Operario & Fiske; Rothbart). Past learning is stored in the form of social stereotypes, person memory, and attitudes and beliefs, all of which function as filters through which new persons and social events are perceived and understood. To some extent, these mental structures are resistant to change – a kind of cognitive conservatism that helps preserve meaning and consistency in our understanding of the world. But human beings are not simply passive information processors. We actively seek new information and modify our mental representations in the light of new experience. Further, we may become aware of how our prior knowledge may bias our perception of incoming information and actively correct for such biases in our judgments and decisions (Martin, Strack, & Stapel). The complex interactions among mental representations, new experience, and social motivations provide the challenge for modern theories of social cognition.

CHAPTER ONE

Mental Representations

Eliot R. Smith and Sarah Queller

Introduction

Many of the core concepts of social psychology, including attitudes, the self-concept, stereotypes, and impressions of other persons, are mental representations. Thus, most theories in social psychology, because they deal with these concepts, implicitly or explicitly make assumptions about how mental representations are constructed, stored in memory, changed, and used to make judgments or plan actions. This chapter aims to explicate and clarify the most popular general conceptions of mental representation and their respective assumptions, which in many theories remain implicit and unelaborated. We review four types of representation: associative networks, schemas, exemplars, and distributed representations. The review focuses on types of representations rather than on "theories," for a single theory often incorporates several types of representation for distinct purposes. For example, Wyer and Srull's (1989) well-known theory includes both associative networks and schemas.

For each of the four types of representation, the chapter will first review basic assumptions regarding representation formation and use. Next, a number of key empirical effects will be described and each mechanism's ability to account for these effects will be discussed. We will discuss explicit, intentional forms of memory as well as the more implicit, unintended effects of mental representations that occur when past experiences influence current perceptions or judgments. Through this discussion, cases where several mechanisms can equally account for existing data will become apparent. The chapter ends with some more general comments on the relations between the different types of representation.

Psychologists generally define a representation as an encoding of information in memory. An individual can create, retain, and access representations. Once accessed, the individual can then use the representation in various ways. For example, your impression of your neighbor is a mental representation that describes your feelings about her and your beliefs about what she is like. You might draw on your impression of your neighbor to describe her

Preparation of this chapter was facilitated by a research grant (R01 MH48640) and a Research Scientist Development Award (K02 MH01178) to the first author and an NRSA postdoctoral fellowship to the second author, from the National Institute of Mental Health.

to a friend, evaluate her as a potential dog-sitter, or decide how to behave when she says something offensive.

Effects of a representation can be explicit in that a previously stored representation is intentionally retrieved from memory, or implicit in that a previously stored representation affects current perceptions or judgments without intention, and perhaps even without conscious awareness (Schacter, 1987, 1994). We typically say "I remember" to denote explicit recall. In contrast, phrases like "I know" or "I believe" are common when implicit memory effects are at work. We rely on explicit memory when remembering a friend's phone number, for example. Explicit memory is often conceptualized using metaphors involving search and retrieval, as if memory was a warehouse filled with different objects. It is implicit memory, on the other hand, that causes us to avoid approaching a person who looks like our childhood tormentor, even if the resemblance is not consciously recognized. Implicit memory fits less well with the notion of search, and instead evokes metaphors like "resonance" to describe the way a stored representation subconsciously influences the way the individual processes new information and makes judgments.

The explicit/implicit memory distinction is one between tasks or ways in which memory has effects rather than between "memory systems" such as semantic versus episodic memory (Tulving, 1972). It is tempting to assume that semantic memory (general knowledge about the world) shows itself in implicit tasks whereas episodic memory (autobiographical memory for specific events located in time and place) affects explicit tasks. But this is misguided. A specific episode can have implicit memory effects. For example, in the phenomenon of repetition priming, reading "elephant" can improve a person's ability to complete the word fragment E – E – – A – T even when the person does not consciously remember reading the word (Tulving, Schacter, & Stark, 1982). In addition, general knowledge can influence explicit memories through reconstructive processes (Ross & Conway, 1986). Thus, the explicit/implicit distinction refers to *uses* of memory – the consciously recollective use of memory versus its use in performing some other task without conscious awareness of memory *per se* (Jacoby & Kelley, 1987).

We now turn to descriptions of the four types of representation, and then discuss how each type accounts for a number of explicit and implicit memory effects.

Associative Networks

Influential theories of associative network representation in non-social cognition can be found in Collins and Quillian (1969), Collins and Loftus (1975), and Anderson and Bower (1973). Within social psychology, the assumptions of associative networks have been described in a number of reviews (Carlston & Smith, 1996; Ostrom, Skowronski, & Nowak, 1994; Fiske & Taylor, 1991; Wyer & Carlston, 1994). The assumptions are as follows:

1 *Fundamental representational assumption*: Representations are constructed from discrete nodes connected by links.
2 *Interpretation of nodes*: Each node stands for a concept. Part of the meaning of each concept is derived from the pattern of linkages to other nodes.

3 *Formation of links through contiguity*: Links are formed between nodes when the concepts the nodes represent are experienced or thought about together.
4 *Link strength*: Existing links are strengthened to the extent the objects they link are experienced or thought about together. Strength changes only slowly with time.
5 *Activation and its spread*: Nodes have a property termed activation, which can vary rapidly over time. A node can become activated if it is perceptually present or actively thought about. An activated node spreads activation to connected nodes via the intervening links, increasing activation of the connected nodes. This process is called "spreading activation" (for a quantitative model, see Anderson, 1983).
6 *Activation in long-term memory*: Long-term memory is a single, large, interconnected associative structure. Short-term memory is the currently activated subset of this structure. Memory retrieval amounts to raising a node's activation level above some threshold.
7 *Links as pathways for retrieval in free recall*: Activating one node may result in the spread of enough activation to a neighboring node to elicit its retrieval. As a direct implication, the more links connected to a particular node, the greater its probability of retrieval.

Though they tend to share the above assumptions in some form, associative network models in social psychology also have some points of variation. First, activation on a node decays with time, but estimates of the rate of decay vary widely (Anderson, 1983; Higgins, 1996; Ostrom et al., 1994, p. 225). Second, some conceptualize retrieval as resulting from the parallel spread of activation across all links connected to currently activated nodes (Anderson, 1983, ch. 3), and others as sequential traversal of links where activation spreads along only one of several possible links at a time (Hastie, 1988). Third, theorists disagree regarding the conceptual level of nodes (see Wyer & Carlston, 1994, p. 7). A node could be a feature, a concept, or a whole body of knowledge ("schema"). And finally, theorists disagree as to whether the links are labeled (e.g. Fiske & Taylor, 1991, p. 297) or unlabeled (Wyer & Srull, 1989, ch. 7). Unlabeled links limit the representational power of associative structures (Carlston & Smith, 1996). For example, if links for subjects versus objects are not distinguished, the representation of the proposition "Sean killed the tiger" would be the same as that of "The tiger killed Sean."

Schemas

Influential works on schematic representation in non-social cognition include Bartlett (1932), Bruner (1957), Bransford & Franks (1971), Anderson & Pichert (1978), and Schank & Abelson (1977). Schematic mechanisms in social psychology, as reviewed by Markus & Zajonc (1985), Carlston & Smith (1996), Fiske & Taylor (1991), and Wyer & Carlston (1994), generally share the following assumptions:

1 *Fundamental representational assumption*: A schema is a structured unit of knowledge about some object or concept. Schemas represent abstract or generalized knowledge as opposed to detailed knowledge about episodes tied to a specific time or context (Fiske & Taylor, 1991; Markus & Zajonc, 1985).

2 *Activation*: A schema can be activated by explicit thought about its topic or by an encounter with relevant information. Making the schema active renders readily accessible all the structured knowledge contained therein.
3 *Level of accessibility*: A schema is likely to become activated and used to the extent that it is accessible. Accessibility is increased by recent or frequent use.
4 *Independence of units*: Schemas are independent entities. Thus, if one schema becomes active this has no necessary implications for other, related schemas.
5 *Interpretive effect of schemas*: Schemas affect the interpretation of perceptual stimuli. That is, the way ambiguous information is construed and the default values that are assumed for unavailable information are influenced by active schemas. The schema-consistent interpretation of a stimulus may be encoded in memory as if it were perceptually present in the stimulus.
6 *Attentional effect of schemas*: Activated schemas direct attention, sometimes to schema-consistent information and sometimes to unexpected or inconsistent information, depending on the circumstances.
7 *Retrieval cueing/reconstructive function of schemas*: Schemas can also influence memory retrieval and judgment. A schema can serve as a source of cues, generally facilitating retrieval of schema-consistent information. It can also serve as a guide for guessing and reconstruction when retrieval attempts fail or produce ambiguous results.

Different theorists' assumptions about schematic mechanisms differ in some respects. First, schemas are typically assumed to represent information about the typical characteristics of particular concepts, such as restaurant dining or doctors. However, in some cases, schemas are assumed to represent general rules of inference independent of any particular content domain (e.g. Heiderian balance can be viewed as a schema.) Second, theorists have modeled schema accessibility in various ways, including Storage Bins, battery, and synapse models (Wyer & Srull, 1989; Higgins, 1996).

Exemplars

Exemplar representations trace directly back to exemplar models of categorization, such as the seminal work by Medin and Schaffer (1978). These models downplay the role of abstractions (such as summaries of the average characteristics of categories) and emphasize instead the role of specific experiences. In non-social cognition, influential works include Brooks (1978), Jacoby & Brooks (1984), and Whittlesea (1987), and in social psychology see Lewicki (1985), Smith (1988, 1990), and Linville, Fischer, & Salovey (1989). Exemplar mechanisms share the following core ideas:

1 *Fundamental representational assumption*: Representations record information about specific stimuli or experiences, rather than abstracted summaries or generalizations. Such a representation may be constructed on the basis of veridical perception of a stimulus object, misperception of it, inference about it, imagination of it, second-hand communication about it, etc.

2 *Representations record feature co-occurrences*: Representations of specific stimuli record patterns of feature co-occurrences. Such representations support people's observed sensitivity to the correlation of features within categories (e.g. they know that small birds are more likely to sing than large ones; Malt & Smith, 1984). In contrast, a schema would contain only information about the typical values of features (i.e. that birds are typically small and that they typically sing), not about feature co-occurrences.
3 *Activation of exemplars by retrieval cues*: Retrieval cues (whether self-generated or external in origin) activate all stored exemplars in parallel. Each exemplar is activated to the extent that it is similar to the retrieval cue. Activation is not synonymous with retrieval, but instead makes the activated exemplars available to influence judgments or impressions (Hintzman, 1986).
4 *Parallel on-line computation*: When a new stimulus is to be evaluated, judged, or categorized, it is compared in parallel to many activated exemplar traces. Similarly, when generalizations about a type of stimulus are required they can be computed by activating all exemplars of that type and summarizing them.
5 *Effects on interpretation, attention, and judgment*: The effects of an activated mass of exemplars are assumed to be the same as those attributed to schemas. That is, the activated exemplars influence interpretation, attention, retrieval, and reconstruction at a preconscious level.

Exemplar theories differ regarding whether only exemplars are stored or, alternately, whether both abstractions and exemplars are stored.

Distributed or PDP representations

Detailed introductions to the newest category of models of mental representation, which have been termed distributed memory, connectionist, or parallel distributed processing (PDP) models, can be found in Churchland & Sejnowski (1992), Smolensky (1988), Rumelhart, McClelland, et al. (1986), as well as McClelland, Rumelhart, et al. (1986). Smith (1996) provides a brief overview oriented toward social psychologists. Distributed representation generally embodies these assumptions:

1 *Fundamental representational assumption*: A concept or object is represented by a distributed representation, where each representation is a different pattern of activation across a common set of simple nodes or units within a network. A useful analogy is a TV screen. No individual pixel has any specific meaning but different patterns of illumination over the entire array of pixels can produce a large number of different meaningful images. This assumption contrasts with associative representations, where individual nodes are semantically meaningful.
2 *Unity of representation and process*: A connectionist network is responsible for both processing and storing information. In contrast, other types of representation require additional assumptions about processes that operate on static representations.
3 *Computing with distributed representations*: Units are interconnected and send activation to each other across weighted connections. A given unit's activation level at a particular

time is a function of its previous activation level as well as the total activation flowing to it from other units across the weighted connections. Thus, the pattern of activation taken on by a given set of units is determined by the initial inputs to the network of units and the weights on the inter-unit connections.

4 *Positive or negative activation*: In most models, both the weights on inter-unit connections and the activation that flows between units can be either positive or negative. Negative activation decreases the activity level of the unit to which it flows (i.e. it has an inhibitory effect). This assumption contrasts with most associative-network models, in which "spreading activation" is always positive.

5 *Learning*: Connection weights are initially assigned random values, which are then shaped by a learning procedure that incrementally changes the weights as the network processes many stimuli.

6 *Connection weights as long-term memory*: The weights on the connections are assumed to change only slowly, in contrast to the quickly changing activation values. Thus, the connection weights are the repository of the network's long-term memory.

7 *Pattern transformation*: Networks with a feed-forward architecture (in which all connections run in one direction from inputs to outputs) can transform representations from one domain into another (Anderson, 1995). Examples are the transformation of input patterns representing behaviors into output patterns representing trait concepts, or inputs of letter sequences into output patterns representing a word's meaning or pronunciation. When the input pattern is presented to input units, activation flows over the connections and eventually produces a new pattern of activation on the network's output units.

8 *Pattern completion or memory*: Networks using a different type of architecture (recurrent connections that link units bidirectionally) can do pattern completion. After the network learns a set of patterns, when the inputs constitute a subset or an approximation of one of those patterns, flows of activation cause the network to reconstruct the entire pattern as output. Pattern completion can be viewed as a form of memory. However, the potential patterns are not explicitly "stored" anywhere. Instead, the network stores connection strengths that allow many patterns to be reproduced given the right cues.

9 *Reconstruction, not retrieval*: When a network must encode several patterns, the connection strengths are a compromise. Hence, reproduction of any given pattern from input cues will be imperfect and will be influenced by the other patterns encoded in the network. As new patterns are learned by the network, the representation of previously learned patterns may change. Thus, distributed representations are *evoked* or *reconstructed* rather than *searched for* or *retrieved* in invariant form (McClelland, Rumelhart, & Hinton, 1986, p. 31).

10 *Parallel constraint satisfaction*: In a network in which bidirectional flows of activation between units are possible, the network can be thought of as converging on a final pattern of activation that simultaneously satisfies the constraints represented by the current inputs (representing external stimuli) and the network weights (representing learned constraints) (Barnden, 1995). Constraint satisfaction is "soft"; learned constraints and current inputs may conflict and each can only be satisfied as well as possible.

There are also points on which various models differ. Some related models accept most of these assumptions but use localist representational schemes in which single nodes have meaningful interpretations (Thorpe, 1995). A node may be interpreted as a feature, an object or concept, or a whole proposition. A connection between nodes is interpreted as encoding past experiences of covariation between the nodes (if nodes represent features or objects) or logical constraints such as consistency or inconsistency between propositions (if nodes represent propositions). Such networks perform parallel satisfaction of multiple constraints (Read & Marcus-Newhall, 1993; Kunda & Thagard, 1996). However, they lack other properties that stem from distributed representation, such as the ability to learn to represent new concepts (new concepts require the explicit addition of new nodes).

Within the class of distributed representation models, literally hundreds of competing models have been proposed with various architectures (numbers and interconnection patterns of nodes), activation equations, and learning rules (see Hertz, Krogh, & Palmer, 1991). In contrast, in each of the three types of representation considered to this point – exemplar, schematic, and associative network – there are perhaps a handful of serious, well-specified competing models. The properties of these diverse distributed models are being actively explored in ongoing theoretical and simulation studies.

Key Memory Effects in Social Psychology

With a basic understanding of the four memory mechanisms in hand, we now turn to describing how each of these mechanisms might account for a number of established effects of mental representations.

Related concepts or contextual factors as retrieval cues

One aspect of memory involves how information that has been learned in the past is retrieved at a later date. Suppose you are introduced to Arturo at a party. How can you recall Arturo's name when you meet him again? This amounts to retrieval of some information (name) from associated information (his appearance). Or how can you recall that he was among the people who attended that particular party? This is retrieval based on contextual cues (the party). These types of memory retrieval are central in most explicit memory tests including paired-associate and list-learning paradigms.

Associative representation Nodes representing concepts that are perceived together or thought about together become linked. When one of the concepts is experienced later, its node becomes active. Activation then spreads across the link to the associated node, potentially raising the activation of this node above the threshold required for retrieval. In this manner, spreading activation explains how related concepts or contextual factors can act as retrieval cues.

The counter-intuitive finding that people recall more behaviors that are inconsistent with their impression of a person than behaviors that are consistent with their impression (Hastie & Kumar, 1979) has been explained in terms of associative representations. When an impression-inconsistent behavior is encountered, the perceiver may try to resolve or explain away the inconsistency. In doing so, the perceiver thinks about the relation of the impression-inconsistent behavior to previously stored impression-consistent and inconsistent behaviors. This process establishes additional links between the inconsistent behavior and other behaviors. These additional links provide more paths along which activation can spread, thus increasing the probability of retrieval of an impression-inconsistent behavior relative to that of recalling an impression-consistent behavior (Hastie, 1980; Srull, 1981).

Recognition of expectation-inconsistent items is also enhanced relative to expectation-consistent items, but only when recognition sensitivity measures are used (measures that correct for a guessing bias; Stangor & McMillan, 1992). Perceivers will guess they have seen consistent items before even when they do not actually recall having seen them, leading to a recognition advantage for consistent items when this bias is not taken into account. While associative representations deal nicely with the finding that expectation-inconsistent items are better *recalled*, the associative representation does not offer an explanation for why *recognition* of expectation-inconsistent behaviors is also enhanced. In recall, a cue is activated and activation spreads across links until an item reaches sufficient activation for retrieval. However, in recognition, an item is presented and the perceiver is asked if he or she previously studied the item. Inter-item associative links are not required as retrieval pathways when the item is directly thought about, so the inconsistency effect for recognition does not seem to be well explained by associative memory mechanisms.

Schematic representation Schematic representations can easily account for the retrieval of items of information that are meaningfully related – that is, are part of the same schema – such as "bread" and "butter." Encountering one of these items activates the schema, which includes the other item. However, schematic models have more difficulty in accounting for a newly learned association (formed by meeting someone for the first time). One could assume that a new schema representing the person is created, but accounts of schema construction (as opposed to retrieval and use) are underdeveloped or entirely absent in most schema theories. In any case, one could argue that forming a new schema to represent a specific occurrence violates the definition of a schema as a representation of abstract, generic knowledge. The definition of schemas as abstract and generic also leads to the conclusion that schema representations do not account well for contextual cueing of retrieval.

Exemplar representation An exemplar representation may preserve information about the specific context in which a stimulus was encountered as well as information about the stimulus characteristics. Therefore an exemplar representation (e.g. incorporating a person's appearance, name, and context) could account for these types of explicit memory retrieval. However, exemplars have more often been invoked to account for implicit rather than explicit memory effects.

Distributed representation When a cue consisting of a partial pattern is presented as input to the network, retrieval occurs via reconstruction of the complete pattern that best satisfies the constraints of the cues provided as input (e.g. Chappell & Humphreys, 1994). In this explanation, a number of stimulus attributes and contextual details are all components of one large pattern, so any of these can act as retrieval cues for the entire pattern.

Accessibility

One important property of memory is that all mental constructs are not equally likely to be used. One determinant of whether a construct is applied is its fit to a current stimulus (Bruner, 1957; Higgins, 1996). Beyond that, mental representations vary in *accessibility*, affecting how readily a perceiver can apply them to the processing of an input. Thus, for example, a professor might tend to evaluate all new acquaintances in terms of intelligence. Intelligence would be an accessible construct for this person.

Associative representation In associative representations, increased accessibility in response to recent use is explained as residual activation on a recently used node. This residual activation puts the node closer to the threshold activation for conscious recall and thus facilitates retrieval of the recently used concept. Increased accessibility in response to frequent use of a concept is explained in terms of link strength. The more a concept is thought about in relation to other concepts, the stronger the links between the corresponding nodes become. Since activation flows more readily over stronger links, retrieval via spreading activation is more likely for frequently used concepts.

Applying these principles, Fazio (1986) suggests that attitudes are represented by an attitude object node linked to an evaluative node. If the attitude is expressed frequently, the link can get strong enough that simply perceiving the object can result in automatic activation of the evaluation. This in turn can lead to evaluative priming effects where a prime facilitates processing of same-valence target items relative to opposite-valence targets (Fazio, Sanbonmatsu, Powell, & Kardes, 1986).

Schema representation Higgins, Rholes, & Jones (1977) showed that recently used traits are more accessible and thus more likely to impact judgments. This finding can be interpreted in terms of increased accessibility of a trait schema due to recent use. In contrast to the short-lived effects of priming, chronic accessibility of a schema is assumed to result from frequent use of a schema over a long period of time (Higgins, King, & Mavin, 1982). For example, some people habitually interpret new information in terms of its implications for gender (Frable & Bem, 1985).

Effects of recency and frequency of use on schema accessibility do not follow directly from the basic assumptions of schema representation. Instead, schema theories have incorporated additional assumptions to account for accessibility. Wyer & Srull (1980) account for accessibility using a "Storage Bin" metaphor. Schemas are thought of as stacked in a Storage Bin and a search for a schema that is applicable to the current stimulus occurs in a top-down fashion. A schema that is nearer the top of the Storage Bin is more accessible. They accounted for the effects of recent use by assuming that a schema was replaced at the top of the Storage Bin after each use, increasing the schema's probability of future use. To account for effects of frequent use, a copying function was later added (Wyer & Srull, 1989), such that when a schema is used, one copy stays in the original location in the Storage Bin and another copy is placed at the top. For frequency effects to be observed, the copies below the top must contribute to accessibility so the probability of using an applicable schema in the top-down search was restricted to $p < 1$ in the revised model.

A "synapse" metaphor was suggested by Higgins, Bargh, & Lombardi (1985) that likened the activation of a schema to a charge that decays with time. Use of a schema fully charges it with activation. The activation subsequently decays. More recently used concepts have more residual activation and are thus more accessible. In order to account for the effect of frequent use on accessibility, the "synapse" model proposes that frequency of use decreases the rate of decay of activation. Note that these authors proposed two distinct mechanisms to deal with recency and frequency effects.

Exemplar representation Frequent or recent use of a concept adds additional exemplar representations to the store in memory. This means that when these exemplars are activated (when a new judgment concerning the concept is required) their summed impact on judgment or memory retrieval is greater. Smith (1988) showed that Hintzman's (1986) MINERVA exemplar model could account for accessibility effects through this mechanism.

Distributed representation Recall that distributed representations are formed through incremental changes in the common set of weights in a connectionist network. If a particular stimulus is presented frequently, the weights will be repeatedly adjusted, becoming better able to reconstruct the pattern corresponding to that stimulus. Consequently, the network will more accurately process frequently encountered patterns compared to less frequently encountered patterns. Recent exposure will also facilitate pattern recognition. In this case, the advantage in accessibility derives from a lessened opportunity for subsequent weight changes that would move the weight values away from those that best process the recent stimulus. Smith and DeCoster (1998) showed that typical effects of recency and frequency on accessibility can be modeled in a connectionist network through this mechanism.

Semantic priming

Semantic priming occurs when perceiving or thinking about one concept makes it easier to process related concepts. Thus, for example, the target word "nurse" is more quickly identified following the prime word "doctor" than following "tree" (Meyer & Schvaneveldt, 1971). The priming effect only lasts a brief period of time (Anderson, 1983; Ostrom et al., 1994, p. 225; Higgins, 1996) and it can be wiped out through the presentation of a single word intervening between prime and target (Masson, 1991; Ratcliff & McKoon, 1988). We know this is an implicit process because it occurs when the prime to target interval is too short for strategic generation of expectations about what is coming next (Neely, 1977) and because it occurs even when the prime is presented subliminally (Wittenbrink, Judd, & Park, 1997).

Associative representation If two nodes are connected with a link then activation from one node can spread across the link to the associated node. When the prime node is activated, the linked target node increases in activation but, unlike the case of explicit recall, the node does not reach the threshold necessary for retrieval. Instead, the target becomes more retrievable in a subsequent task. (Although similar, the explanation of semantic priming should not be confused with that of recency. With recency, the target concept is directly activated and becomes more accessible at a later time. With semantic priming, a concept related to the target concept is directly activated and activation is spread to produce the sub-threshold activation of the target concept.)

Associative representations typically have been relied upon in explaining priming phenomena within social psychology. For example, white subjects responded more quickly to positive trait words following the prime word "white" than following the prime word "black" (Dovidio, Evans, & Tyler, 1986). The associative explanation is that "white" is semantically linked to positive concepts to a greater extent than is "black." This finding and the related suppositions about representation have relevance for racial stereotyping. Another effect that has been explained in associative terms is evaluative priming (Fazio et al., 1986). Processing of an evaluatively laden prime ("cockroach") results in facilitated processing of evaluatively similar targets ("death") and inhibited processing of evaluatively dissimilar targets ("beautiful"). This finding is robust, although there is debate regarding whether it holds only for objects about which the perceiver holds fairly strong attitudes (Bargh, Chaiken, Raymond, & Hymes, 1996; Fazio, 1993).

Schema representation Semantic priming in schema representations occurs because a schema-relevant stimulus can activate the whole schema. Thus, for example, if the word "doctor" activates the schema for hospital or medical care, processing of the word "nurse" might be facilitated since it would be activated as part of that schema. In contrast, the word "tree" would not activate a hospital schema so "tree" would not facilitate processing of "nurse."

Exemplar representation A mass of similar exemplars can function like a schematic general knowledge structure, since the respects in which they are similar reinforce each other while contextual or nonessential differences cancel out (Hintzman, 1986). Therefore numerous exemplars of medical care or hospital situations, most of which included both doctors and nurses, might account for semantic priming in the same way as a "medical care" schema.

Distributed representation Semantically related concepts tend to share features and, thus, semantically related prime/target pairs will have overlapping patterns of activation in a distributed network. To the extent that the target's pattern of activation overlaps with the previously processed prime's pattern of activation, the network will more quickly and accurately process the target pattern. As activation is a short-lived property, the distributed mechanism accurately predicts that the priming effect should be wiped out by presentation of an unrelated stimulus between prime and target (Masson, 1991).

Repetition priming

Processing of a stimulus is facilitated when the same stimulus has been processed in the same way on a previous occasion. This phenomenon is long lasting (as long as months: Sloman, Hayman, Ohta, Law, & Tulving, 1988), in contrast with short-lived semantic priming effects. The previous exposure does not have to be consciously remembered for repetition priming to occur (Schacter, 1987; Smith, Stewart, & Buttram, 1992). For example, Smith, Stewart, & Buttram (1992) had subjects quickly decide if each of hundreds of behaviors were friendly or intelligent. Some of the behaviors were repeated and some were not. Repetition of behaviors resulted in faster judgment times, even when the delay following the initial exposure was a week. Faster judgment times occurred even when subjects could not recall the previous exposure to the behavior. Because repetition priming is so long lasting, it cannot be explained in terms of residual activation because in all four types of representation, activation is relatively short lived.

Associative representation Repetition priming can be explained by assuming that a new association is formed linking the specific stimulus being judged or processed to the results of that processing. For example, attributes of a particular behavior would become associated with a trait like "friendly." When the same behavior is encountered again, activation would spread to the trait concept, facilitating a repetition of the same judgment.

Schema representation The basic assumptions of schema representation do not readily lead to an explanation for repetition priming. Activation of a trait schema might occur when the trait judgments are made in the Smith, Stewart, & Buttram (1992) study, but the effect would be to speed *all* later judgments using that trait, rather than only judgments about specific repeated behaviors (which is what is empirically observed).

Exemplar representation When an exemplar is judged in a particular way (say a behavior is judged on a trait) that judgment becomes part of the exemplar representation that is stored in memory (e.g. the trait implications become part of the behavior representation). This illustrates the general principle that exemplar representations are always stored *as processed* or interpreted by the perceiver, not in veridical form (Whittlesea & Dorken, 1993). When the same behavior is presented again at a later time, if the exemplar can be retrieved from memory the judgment is already available. Consequently, a second judgment is performed faster than an identical first judgment (Logan, 1988; Smith, 1990).

Distributed representation Weight changes after exposure to a specific pattern will facilitate later processing of the same pattern. However, as additional patterns are presented to the network, they cause further weight changes, overwriting those that provide an advantage to the repeated stimulus. This argument is similar to that for accessibility of a pattern due to recent use. (See Wiles & Humphreys, 1993, pp. 157–163 for a quantitative analysis.)

Filling in default values and resolving ambiguity

Theorists have long understood that perceivers do not process new information in a strictly unbiased manner but, instead, rely on prior knowledge to help make sense of new information (Arbib, 1995; Bartlett, 1932; Bruner, 1957; Markus & Zajonc, 1985; Minsky, 1975;

Neisser, 1976). Prior knowledge is helpful in resolving ambiguities in incoming information or filling in default values for unobserved characteristics. These effects of prior experience are implicit: we don't "remember," for example, that the bird we saw standing in the tall grass had feet, we just "know" it. As an example, imagine observing a car that does not slow down as a pedestrian crosses in front of it. The pedestrian raises his hand and a moment later, the driver extends his arm out the window. They might be waving at each other or they might be exchanging rude gestures.

Associative representation Spreading activation across links makes linked nodes available for use in resolving ambiguity. In an associative framework, the way you interpret a situation and make assumptions about unknown aspects would depend on the concepts you have strongly linked together based on past experience (Anderson, 1983). If you thought the driver had been careless, links from that concept to the concept of anger might lead you to interpret the hand gestures as insults. You might also think that the driver shouted a curse, even though you could not hear the driver.

Schema representation Schemas are activated in an all or none fashion and the content of an activated schema is applied to incoming information in an implicit manner. The ready explanation of the resolution of ambiguous information and the use of default values to fill in missing information have led to the popularity of schema models within social psychology. Stereotypes are often conceived of as schemas that allow us to generate expectations about types of people. For example, a girl who turned in a mixed performance on a test was rated as more academically talented by subjects who believed she came from an upper-middle-class background than by subjects who thought she was working class (Darley & Gross, 1983). Similarly, scripts are schemas that store generalized knowledge about a type of event, such as going to a birthday party or dining at a restaurant (Schank & Abelson, 1977). If a script for "road rage" is activated by the driver-pedestrian encounter, it may lead to interpretations and inferences that are consistent with an angry interchange.

Exemplar representation The common characteristics of the mass of exemplars called to mind when making a judgment may be applied to the novel stimulus (Hintzman, 1986; Smith & Zarate, 1992). The effect is similar to that of a schema and the difference is only that a set of exemplars serve as the prior knowledge instead of a schema that contains generalized knowledge.

Distributed representation The filling in of default values and resolution of ambiguities occurs in distributed representations through the flow of activation in a connectionist network whose weights have been tuned by past experiences (Rumelhart, Smolensky, McClelland, & Hinton, 1986). The input to the network may be a partial pattern with only a few characteristics filled in. But as activation flows through the connection weights, the network outputs a complete pattern that best satisfies the constraints of the input and the knowledge currently stored in the connection weights (Smith & DeCoster, 1998). Thus, default values are automatically generated and ambiguities are resolved through constraint satisfaction in distributed networks.

Flexibility and context sensitivity

Recent thinking in social psychology has emphasized flexibility and context sensitivity in the areas of the self-concept (Markus & Wurf, 1987; Linville & Carlston, 1994; Higgins, Van Hook, & Dorfman, 1988; Turner, Oakes, Haslam, & McGarty, 1994), attitudes (Wilson & Hodges, 1992; Tourangeau & Rasinski, 1988; Strack & Martin, 1987; Schwarz & Clore, 1983; Wilson, 1990), and even stereotypes (Bodenhausen, Schwarz, Bless, & Wanke, 1995). Not only the accessibility of a representation, but the content of the representation can be altered by contextual information. That is, particular features are emphasized or de-emphasized in different contexts. In addition, Barsalou (1987) has shown that ad hoc concepts such as "things that might fall on your head" are structured in the same way as concepts with which one has had a great deal of experience and prior learning. Extending this finding to the social domain, people may easily generate stereotypes (summaries of and feelings about a group's typical characteristics) for previously unconsidered groups such as "people who fly only at night" or "adopted Latinos."

Associative representation Associative representations are flexible to the extent that distinct sets of cues may activate different sets of associates. Thus, for example, the cue "bird" plus the cue "barnyard" might activate the concept "chicken," whereas the cue "bird" plus the cue "suburban backyard" may activate the concept "robin" (see Barsalou, 1987). This is consistent with research showing that although people rate "robin" as a better exemplar of the bird category than "chicken" without context, they rate "chicken" higher in a barnyard context. Interestingly, however, if the cue "bird" alone activates the concept "robin," in associative terms it seems that the compound cue "bird" plus "barnyard" should activate "robin" in addition to "chicken." (The same activation should spread from the "bird" node regardless of whether the "barnyard" node is also active.) Only if assumptions about spreading inhibition as well as activation are added might the "chicken" concept be retrieved without also retrieving the "robin" concept. Classic models using associative representations and spreading activation did not invoke inhibition (e.g. Anderson, 1983). However, recent associative formulations within social psychology have assumed that associative links can be inhibitory as well as facilitative (e.g. Carlston, 1994).

Schema representation Schematic representations have difficulty accounting for context sensitivity. One would have to postulate, for instance, that distinct "bird-in-suburbs" and "bird-in-barnyard" schemas exist instead of a single "bird" schema. However, this leads to an explosion of the number of concepts that must be represented. Alternatively one could assume that general "bird" and "barnyard" schemas are combined on-line in some fashion to yield the new context-specific concept representation, but models of schema combination have received little attention (Wisniewski, 1997), and none in social psychology.

Exemplar representation Exemplar representations can accommodate flexible use of prior knowledge (Smith, 1990). Different sets of exemplars may be activated when making different judgments, depending on context or other details of the specific set of cues provided. Judgments are then based on the activated set of exemplars.

Distributed representation Distributed representations also allow fluid use of prior knowledge. The set of cues provided can include incidentally activated goals (e.g. enhancing self-esteem), context, mood, perceptually present objects, and/or objects of current thought. All of these cues are represented in the common set of weight values so retrieval is influenced by all of them (Rumelhart, Smolensky, McClelland, & Hinton, 1986; Clark, 1993).

Dissociations between recall, recognition, and judgment

If different memory tasks such as recall, recognition, and judgment all access the same underlying representational structure, one would expect dependence between the different tasks. For example, if a specific representation can be demonstrated to influence judgment, such as through a priming effect, the representation would be expected also to be available to explicit retrieval. However, this is not always the case. For example, priming manipulations often have similar or even greater effects on judgment when they cannot be consciously remembered as when they are explicitly retrievable (Lombardi, Higgins, & Bargh, 1987; Martin, Seta, & Crelia, 1990). And although a person might be judged to be honest, recall of the person's behaviors might include a large number of dishonest behaviors (Hastie & Kumar, 1979).

These and other dissociations have been shown to be a function of how the information was processed at learning (Carlston & Smith, 1996; Hastie & Park, 1986; Jacoby, 1983). For example, whether a target is processed perceptually (read the word "honest") or conceptually (read several honest behaviors and generate the relevant trait) will affect performance on different memory tasks. Fragment completion (complete the word "H – N – – T") is more strongly enhanced following a perceptual task, whereas recall (recall the trait words from the previous task) and category accessibility (read an ambiguous behavior and pick a trait that fits) show greater effects of a previous conceptual task (Smith & Branscombe, 1988). For a review of dissociations among implicit and explicit memory measures in nonsocial cognition, see Richardson-Klavehn & Bjork (1988) and Hintzman (1990).

Associative representation With associative representations, dissociations occur when different cognitive structures are drawn on for different tasks. As Hastie and Park (1986) suggest, perceivers create different types of representations in memory depending on whether they initially process the incoming information in a "memory-based" fashion or in an "on-line" fashion. "Memory-based" processing initially stores representations of the input stimuli. When a judgment is called for, the stimulus details are recalled and summarized at the time of retrieval. In this case, recall of the details and the judgment should correspond, since the judgment is based directly on what can be recalled. With "on-line" processing, abstraction or summarization of the stimuli occurs as the stimuli are encountered, resulting in a summary representation as well as stimulus-specific details being stored in memory. When a judgment is required, the summary representation is accessed. When recall is performed, the detailed representations are accessed. Dissociation is explained by the use of these two distinct representations that may not contain exactly the same content. Similarly, Wyer and Gordon (1984) suggested that people store a target's behaviors in both trait-based associative clusters and an evaluative summary, stored independently in the "Storage Bin." An evaluative judgment task accesses the evaluative summary, whereas explicit recall of behaviors accesses the trait based clusters. Because the two tasks access different cognitive representations, there may be little relationship between the judgment and the behaviors recalled.

Schema representation As noted earlier, models of schema formation are typically weak or underdeveloped, making schema models poor candidates for explaining phenomena that involve the creation of new representations. Leaving this issue aside, one could explain dissociations by suggesting that people form multiple schemas organized along different lines (as in the Wyer & Gordon model just described) as they process stimulus inputs.

Exemplar representation Dissociations are assumed to be due to the distinct subsets of exemplars that are activated by the different cues provided by different tasks (such as implicit versus explicit memory tasks; Roediger & McDermott, 1993). Thus the exemplar representations that can be explicitly retrieved (e.g. the members of a category that one can recall or recognize) may not be the same as those that implicitly influence categorization judgments.

Distributed representation Dissociations among different memory tasks may be due to differences in tasks that cause the perceiver to draw on distinct network representations. For example, Wiles and Humphreys (1993) suggest that explicit recall and semantic priming draw on pattern completion networks that reconstruct the features of past stimuli when partial cues are presented. In contrast, they suggest that repetition priming is due to changes in weights in pattern transformation networks that translate from perceptual (e.g. visual) to internal (e.g. semantic) types of representations. This proposal explains why repetition priming is specific to a given perceptual modality such as visual or auditory, while semantic priming is not modality specific. Another explanation for dissociations in distributed representations is that different cues might be presented in different types of memory tasks and this may give rise to independence between tasks (Humphreys, Bain, & Pike, 1989).

Summary and Conclusions

The two types of representation most frequently invoked in social psychology are associative and schema representations. Comparison of these two is informative. First, although associative representations focus on the acquisition of new knowledge as well as its use, schema models emphasize the use of existing general knowledge. Loosely, we might say that the construction of new associative representations accounts for episodic memory for particular events or stimuli, whereas schematic representations seem better aligned with semantic or generic memory. In other words, these different types of representation may be complementary rather than competing. Indeed, a number of theorists (McClelland, McNaughton, & O'Reilly, 1995; Humphreys, Bain, & Pike, 1989; Hirst, 1989; Masson, 1989; Macleod & Bassili, 1989; Squire, 1994; Schacter, 1994; Moscovitch, 1994) have recently posited two functionally independent memory systems. One system handles one-shot learning by constructing new representational structures (akin to social psychological assumptions regarding associative representations). The other system learns slowly, gradually building representations of the general characteristics of objects or events. These dual memory theories are supported by a number of psychological and neuropsychological studies of different types. For example, evidence suggests that rapid, episodic learning is mediated by the hippocampus and related structures whereas slow, semantic learning relies on cortical structures (Squire, 1994).

Comparison between associative and schematic representations also suggests that associative representations can mimic schema representations. A schema can be conceptualized in associative terms as a set of units that are so strongly interlinked that activating any one of them necessarily activates them all (Ostrom et al., 1994, p. 221; Anderson, 1983). As discussed above, exemplar and distributed representations can also mimic the effects of schemas that are typically emphasized in social psychology, such as using prior knowledge to interpret new inputs. Thus, we suggest that in social psychology, a schema is more a description of a *function* that can be performed by a learned knowledge representation, rather than a description of an actual entity inside our heads.

In fact, associative and exemplar representations as well as schemas may be more descriptions of memory function than they are accounts of actual underlying representations and processes. This consideration raises a distinction between associative, schema, and exemplar representations on the one hand and distributed representations on the other. The former three have all been formulated based on specific empirical phenomena. That is, a function was observed (e.g. accessibility) and a mechanism of memory was proposed to account for that function (e.g. position in a Storage Bin). This leads to a one-to-one correspondence between empirical evidence and theoretical mechanisms that makes theories relatively clear and understandable. Distributed representations tend to proceed in the opposite direction. Instead of starting from psychological findings, theorists who use distributed representations begin with a theoretical vocabulary (computationally simple units and interconnections, modeled very loosely on the properties of biological neurons). They then use this vocabulary to build specific models and see whether they can replicate known memory findings. It is possible that distributed representations may turn out to more often predict novel empirical effects, precisely because they are not originally formulated to provide a one-to-one correspondence with known phenomena.

This distinction between empirical observations and theoretical vocabulary as starting points for modeling may be regarded as a distinction between levels of theory (Smith, 1998; Smolensky, 1988). If a higher level of theory is sufficient to explain social psychological phenomena of interest, it might be argued that we need not consider the lower-level details of mental representations. However, there are several counterarguments. First, the details of how mental representations are formed and used necessarily constrain higher level theories. As an analogy, a theory about the chemical reactions between molecules will not stand if it is inconsistent with known properties of atoms. Second, new mechanisms display new properties. Distributed representations are certainly new to social psychology but have already been shown to generate novel predictions regarding social psychological issues such as accessibility (DeCoster & Smith, 1998) and stereotype learning and change (Queller & Smith, 1998). Third, distributed representations provide a dynamic approach that emphasizes learning. Many of the most interesting social psychological questions – about attitude change and stereotype change, for example – involve *changes* in representation. Finally, modeling at a lower level can lead to greater integration and parsimony. Traditionally within social psychology, theories have been developed and refined to account for the specifics of a relatively small domain. Taken to extremes, this approach can lead to a profusion of fragmentary minitheories that have unclear relations with one another and no common basis of assumptions. Connectionist models offer the possibility of a broad integration not only within social psychology, but beyond, including areas of cognition, perception, development, and neuroscience (Elman, Bates, Johnson, Karmiloff-Smith, Parisi, & Plunkett, 1995; McClelland, McNaughton, & O'Reilly, 1995).

REFERENCES

Anderson, J. A. (1995). Associative networks. In M. A. Arbib (Ed.), *Handbook of brain theory and neural networks* (pp. 102–107). Cambridge, MA: MIT Press.

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington, DC: Winston & Sons.

Anderson, R. C., & Pichert, J. W. (1978). Recall of previously unrecallable information following a shift in perspective. *Journal of Verbal Learning and Verbal Behavior*, *17*, 1–12.

Arbib, M. A. (1995). Schema theory. In M. A. Arbib (Ed.), *Handbook of brain theory and neural networks* (pp. 830–834). Cambridge, MA: MIT Press.

Bargh, J. A., & Thein, R. D. (1985). Individual construct accessibility, person memory, and the recall–judgment link: The case of information overload. *Journal of Personality and Social Psychology*, *49*, 1129–1146.

Bargh, J. A., Lombardi, W. J., & Higgins, E. T. (1988). Automaticity of chronically accessible constructs in person X situation effects on person perception: It's just a matter of time. *Journal of Personality and Social Psychology*, *55*, 599–605.

Bargh, J. A., Lombardi, W. J., & Tota, M. E. (1986). The additive nature of chronic and temporary sources of construct accessibility. *Journal of Personality and Social Psychology*, *50*, 869–879.

Bargh, J. A., Chaiken, S., Raymond, P., & Hymes, C. (1996). The automatic evaluation effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology*, *22*, 104–128.

Barnden, J. A. (1995). Artificial intelligence and neural networks. In M. A. Arbib (Ed.), *Handbook of brain theory and neural networks* (pp. 98–102). Cambridge, MA: MIT Press.

Barsalou, L. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 101–140). Cambridge, UK: Cambridge University Press.

Bartlett, F. C. (1932). *Remembering*. Cambridge, UK: Cambridge University Press.

Bodenhausen, G. V., Schwarz, N., Bless, H., & Wanke, M. (1995). Effects of atypical exemplars on racial beliefs: Enlightened racism or generalized appraisals? *Journal of Experimental Social Psychology*, *31*, 48–63.

Bransford, J. D., & Franks, J. J. (1971). The abstraction of linguistic ideas. *Cognitive Psychology*, *2*, 331–350.

Brooks, L. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169–211). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bruner, J. S. (1957). Going beyond the information given. In H. Gruber, G. Terrell, & M. Wertheimer (Eds.), *Contemporary approaches to cognition*. Cambridge, MA: Harvard University Press.

Carlston, D. E. (1994). Associated systems theory: A systematic approach to cognitive representations of persons. In T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition: A dual process model of impression formation* Vol. 7 (pp. 1–78). Hillsdale, NJ: Lawrence Erlbaum Associates.

Carlston, D. E., & Smith, E. R. (1996). Principles of mental representation. In E. T. Higgins & A. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 184–210). New York: Guilford Press.

Chappell, M., & Humphreys, M. S. (1994). An auto-associative neural network for sparse representations: Analysis and application to models of recognition and cued recall. *Psychological Review*, *101*, 103–128.

Churchland, P. S., & Sejnowski, T. J. (1992). *The computational brain*. Cambridge, MA: MIT Press.

Clark, A. (1993). *Associative engines: Connectionism, concepts, and representational change*. Cambridge, MA: MIT Press.

Collins, A., & Loftus, E. F. (1975). A spreading activation theory of semantic memory. *Journal of Verbal Learning and Verbal Behavior*, *8*, 240–247.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, *8*, 240–247.

Darley, J. M., & Gross, P. H. (1983). A hypothesis-confirming bias in labelling effects. *Journal of Personality and Social Psychology*, *44*, 20–33.

DeCoster, J., & Smith, E. R. (1998). Rapid recovery of accessibility: Empirical tests of a novel prediction of a connectionist memory. Unpublished paper, Purdue University.

Dovidio, J. F., Evans, N., & Tyler, R. B. (1986). Racial stereotypes: The content of their cognitive representations. *Journal of Experimental Social Psychology*, *22*, 22–37.

Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1995). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.

Fazio, R. H. (1986). How do attitudes guide behavior? In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition* (pp. 204–243). New York: Guilford Press.

Fazio, R. H. (1993). Variability in the likelihood of automatic attitude activation: Data reanalysis and commentary on Bargh, Chaiken, Govender, and Pratto (1992). *Journal of Personality and Social Psychology, 64*, 753–758.

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50*, 229–238.

Fiske, S., & Taylor, S. (1994). *Social cognition*. 2nd. edn. Reading, MA: Addison-Wesley.

Frable, D. E. S., & Bem, S. L. (1985). If you are gender schematic, all members of the opposite sex look alike. *Journal of Personality and Social Psychology*, *49*, 459–468.

Hastie, R. (1980). Memory for information which confirms or contradicts a general impression. In R. Hastie, T. M. Ostrom, E. B. Ebbesen, R. S. Wyer, D. Hamilton, & D. E. Carlston (Eds.), *Person memory* (pp. 155–177). Hillsdale, NJ: Lawrence Erlbaum Associates.

Hastie, R. (1988). A computer simulation model of person memory. *Journal of Experimental Social Psychology*, *24*, 423–447.

Hastie, R., & Kumar, P. A. (1979). Person memory: Personality traits as organizing principles in memory for behaviors. *Journal of Personality and Social Psychology*, *37*, 25–38.

Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review*, *93*, 258–268.

Hertz, J., Krogh, A., & Palmer, R. G. (1991). *Introduction to the theory of neural computation*. Redwood City, CA: Addison-Wesley.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133–168). New York: Guilford Press.

Higgins, E. T., Bargh, J. A., & Lombardi, W. (1985). Nature of priming effect on categorization. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *11*, 59–69.

Higgins, E. T., King, G. A., & Mavin, G. H. (1982). Individual construct accessibility and subjective impressions and recall. *Journal of Personality and Social Psychology*, *43*, 35–47.

Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology*, *13*, 141–154.

Higgins, E. T., Van Hook, E., & Dorfman, D. (1988). Do self-attributes form a cognitive structure? *Social Cognition*, *6*, 177–207.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*, 411–428.

Hintzman, D. L. (1990). Human learning and memory: Connections and dissociations. *Annual Review of Psychology*, *41*, 109–140.

Hirst, W. (1989). On consciousness, recall, recognition, and the architecture of memory. In S. Lewandowsky, J. C. Dunn, & K. Kirsner (Eds.), *Implicit memory: Theoretical issues* (pp. 33–46). Hillsdale, NJ: Lawrence Erlbaum Associates.

Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks. *Psychological Review*, *96*, 208–233.

Jacoby, L. L. (1983). Perceptual enhancement: Persistent effects of an experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 21–38.

Jacoby, L. L., & Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception and concept learning. In G. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* Vol. 18. New York: Academic Press.

Jacoby, L. L., & Kelley, C. M. (1987). Unconscious influences of memory for a prior event. *Personality and Social Psychology Bulletin*, *13*, 314–336.

Kunda, Z., & Thagard, P. (1996). Integrating stereotypes with individuating information: A parallel constraint satisfaction model of impression formation. *Psychological Review*, *103*, 284–308.

Lewicki, P. (1985). Nonconscious biasing effects of single instances of subsequent judgments. *Journal of Personality and Social Psychology*, *48*, 563–574.

Linville, P. W., & Carlston, D. E. (1994). Social cognition of the self. In P. G. Devine, T. M. Ostrom, & D. L. Hamilton (Eds.), *Social cognition: Impact on social psychology* (pp. 143–193). Orlando, FL: Academic Press.

Linville, P. W., Fischer, G. W., & Salovey, P. (1989). Perceived distributions of the characteristics of in-group and out-group members: Empirical evidence and a computer simulation. *Journal of Personality and Social Psychology*, *57*, 165–188.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.

Lombardi, W. J., Higgins, E. T., & Bargh, J. A. (1987). The role of consciousness in priming effects on categorization: Assimilation versus contrast as a function of awareness of the priming task. *Personality and Social Psychology Bulletin*, *13*, 411–429.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–457.

McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986). The appeal of parallel distributed processing. In D. E. Rumelhart, J. L. McClelland, et al. (Eds.), *Parallel distributed processing* Vol. 1 (pp. 3–44). Cambridge, MA: MIT Press.

McClelland, J. L., Rumelhart, D. E., et al. (Eds.) (1986). *Parallel distributed processing* Vol. 2. Cambridge, MA: MIT Press.

MacLeod, C. M., & Bassili, J. N. (1989). Are implicit and explicit tests differentially sensitive to item-specific vs. relational information? In S. Lewandowsky, J. C. Dunn, & K. Kirsner (Eds.), *Implicit memory: Theoretical issues* (pp. 159–172). Hillsdale, NJ: Lawrence Erlbaum Associates.

Malt, B. C., & Smith, E. E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior*, *23*, 250–269.

Markus, H., & Wurf, E. (1987). The dynamic self-concept: A social psychological perspective. *Annual Review of Psychology*, *38*, 299–337.

Markus, H., & Zajonc, R. B. (1985). The cognitive perspective in social psychology. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology.* 3rd. edn., Vol. 1 (pp. 137–230). New York: Random House.

Martin, L. L., Seta, J. J., & Crelia, R. A. (1990). Assimilation and contrast as a function of people's willingness and ability to expend effort in forming an impression. *Journal of Personality and Social Psychology*, *59*, 38–49.

Masson, M. E. J. (1989). Fluent reprocessing as an implicit expression of memory for experience. In S. Lewandowsky, J. C. Dunn, & K. Kirsner (Eds.), *Implicit memory: Theoretical issues* (pp. 123–138). Hillsdale, NJ: Lawrence Erlbaum Associates.

Masson, M. E. J. (1991). A distributed memory model of context effects in word identification. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 233–263). Hillsdale, NJ: Lawrence Erlbaum Associates.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*, 227–234.

Minsky, M. (1975). A framework for representing knowledge. In P. H. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill.

Moscovitch, M. (1994). Memory and working with memory: Evaluation of a component process model and comparisons with other models. In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 269–310). Cambridge, MA: MIT Press.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, *1*, 226–254.

Neisser, U. (1976). *Cognition and reality*. San Francisco: Freeman.

Ostrom, T. M., Skowronski, J. J., & Nowak, A. (1994). The cognitive foundation of attitudes: It's a wonderful construct. In P. G. Devine, T. M. Ostrom, & D. L. Hamilton (Eds.), *Social cognition: Impact on social psychology* (pp. 195–257). Orlando, FL: Academic Press.

Queller, S., & Smith, E. R. (1998). Stereotype change from dispersed and concentrated disconfirming information: New insights from connectionist models. Unpublished paper, Purdue University.

Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. *Psychological Review*, *95*, 385–408.

Read, S. J., & Marcus-Newhall, A. (1993). Explanatory coherence in social explanations: A parallel distributed processing account. *Journal of Personality and Social Psychology*, *65*, 429–447.

Richardson-Klavehn, A., & Bjork, R. A. (1988). Measures of memory. *Annual Review of Psychology*, *39*, 475–544.

Roediger, H. L., & McDermott, K. B. (1993). Implicit memory in normal human subjects. In F. Boller & J. Erafman (Eds.), *Handbook of neuropsychology* Vol. 8 (pp. 63–131). Amsterdam: Elsevier.

Ross, M., & Conway, M. (1986). Remembering one's own past: The construction of personal histories. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition* (pp. 122–144). New York: Guilford Press.

Rumelhart, D. E., McClelland, J. L., et al. (Eds.). (1986). *Parallel distributed processing* Vol. 1. Cambridge, MA: MIT Press.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* Vol. 2 (pp. 7–57). Cambridge, MA: MIT Press.

Schacter, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 501–518.

Schacter, D. L. (1994). Priming and multiple memory systems: Perceptual mechanisms of implicit memory. In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 233–268). Cambridge, MA: MIT Press.

Schank, R., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology*, *45*, 513–523.

Sloman, S. A., Hayman, C. A. G., Ohta, N., Law, J., & Tulving, E. (1988). Forgetting in primed fragment completion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 223–239.

Smith, E. R. (1988). Category accessibility effects in a simulated exemplar-only memory. *Journal of Experimental Social Psychology*, *24*, 448–463.

Smith, E. R. (1990). Content and process specificity in the effects of prior experiences. Target Article in T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition* Vol. 3 (pp. 1–60). Hillsdale, NJ: Lawrence Erlbaum Associates.

Smith, E. R. (1996). What do connectionism and social psychology offer each other? *Journal of Personality and Social Psychology*, *70*, 893–912.

Smith, E. R. (1998). Mental representation and memory. In D. Gilbert, S. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology.* 4th. edn., Vol. 1 (pp. 391–445). New York: McGraw-Hill.

Smith, E. R., & Branscombe, N. R. (1988). Category accessibility as implicit memory. *Journal of Experimental Social Psychology*, *24*, 490–504.

Smith, E. R., & DeCoster, J. (1998). Knowledge acquisition, accessibility, and use in person perception and stereotyping: Simulation with a recurrent connectionist network. *Journal of Personality and Social Psychology*, *74*, 21–35.

Smith, E. R., & Zárate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review*, *99*, 3–21.

Smith, E. R., Stewart, T. L., & Buttram, R. T. (1992). Inferring a trait from a behavior has long-term, highly specific effects. *Journal of Personality and Social Psychology*, *62*, 753–759.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, *11*, 1–74.

Squire, L. R. (1994). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 203–232). Cambridge, MA: MIT Press.

Srull, T. K. (1981). Person memory: Some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 440–463.

Stangor, C., & McMillan, D. (1992). Memory for expectancy-congruent and expectancy-incongruent information: A review of the social and social developmental literatures. *Psychological Bulletin*, *111*, 42–61.

Strack, F., & Martin, L. L. (1987). Thinking, judging, and communicating: A process account of context effects in attitude surveys. In H.-J. Hippler, N. Schwarz, & S. Sudman (Eds.), *Social information processing and survey methodology* (pp. 123–148). New York: Springer-Verlag.

Thorpe, S. (1995). Localized and distributed representations. In M. A. Arbib (Ed.), *Handbook of brain theory and neural networks* (pp. 549–552). Cambridge, MA: MIT Press.

Tourangeau, R., & Rasinski, K. A. (1988). Cognitive processes underlying context effects in attitude measurement. *Psychological Bulletin*, *103*, 299–314.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory*. New York: Academic.

Tulving, E., Schacter, D. L., & Stark, H. (1982). Priming effects in word-fragment completion are independent of recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *8*, 336–342.

Turner, J. C., Oakes, P. J., Haslam, S. A., & McGarty, C. (1994). Self and collective: Cognition and social context. *Personality and Social Psychology Bulletin, 20*, 454–463.

Whittlesea, B. W. A. (1987) Preservation of specific experiences in the representation of general knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 3–17.

Whittlesea, B. W. A., & Dorken, M. D. (1993). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General*, *122*, 227–248.

Wiles, J., & Humphreys, M. S. (1993). Using artificial neural nets to model implicit and explicit memory test performance. In P. Graf & M. E. J. Masson (Eds.), *Implicit memory: New directions in cognition, development, and neuropsychology* (pp. 141–165). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wilson, T. D. (1990). Self-persuasion via self-reflection. In J. M. Olson & M. P. Zanna (Eds.), *Self-inference processes: The Ontario Symposium* Vol. 6 (pp. 43–68). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wilson, T. D., & Hodges, S. D. (1992). Attitudes as temporary constructions. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 37–65). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin and Review*, *4*, 167–183.

Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, *72*, 262–274.

Wyer, R. S., & Carlston, D. E. (1994). The cognitive representation of persons and events. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of Social Cognition*. 2nd. edn. (pp. 41–98). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wyer, R. S., & Gordon, S. E. (1984). The cognitive representation of social information. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* Vol. 2 (pp. 73–150). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wyer, R. S., & Srull, T. K. (1980). The processing of social stimulus information: A conceptual integration. In R. Hastie, T. M. Ostrom, E. B. Ebbesen, R. S. Wyer, D. Hamilton, & D. E. Carlston (Eds.), *Person memory* (pp. 227–300). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wyer, R. S., & Srull, T. K. (1989). *Memory and cognition in its social context*. Hillsdale, NJ: Lawrence Erlbaum Associates.