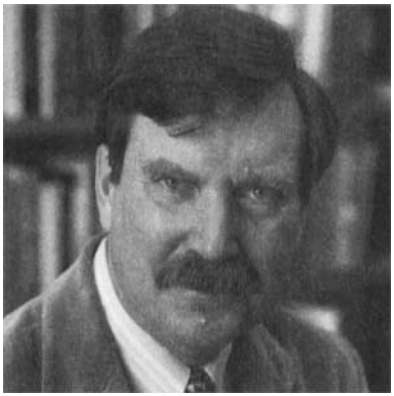

John Kihlstrom is Professor in the Department of Psychology at Yale University and the University of California, Berkeley. He received his Ph.D. from the University of Pennsylvania, and previously held positions at Harvard, Wisconsin, and Arizona. His chief research interests are in cognition in a personal and social context, and he has special interests in unconscious mental processes, autobiographical memory, and the self. He is editor of Psychological Science, *the flagship journal of the American Psychological Society, and his many awards include the APA Distinguished Scientific Award for an early career contribution to psychology.*

Stan Klein is Associate Professor of Psychology at the University of California, Santa Barbara. Klein received his Ph.D. from Harvard University and has previously held jobs at Trinity University and the University of Illinois. His research interests are in social cognition with a particular focus on the mental representation of trait and behavioral knowledge about the self. A larger goal of this research is to examine the extent to which a common set of principles can explain the way in which knowledge of self and others is represented in the mind.

Self-Knowledge and Self-Awareness[a]

JOHN F. KIHLSTROM[b]

Department of Psychology
Yale University
New Haven, Connecticut 06520

STANLEY B. KLEIN

Department of Psychology
University of California, Santa Barbara
Santa Barbara, California 93106

The self lies at the center of mental life. As William James (1890/1981, p. 221) noted in the *Principles of Psychology,*

> Every thought tends to be part of a personal consciousness. . . . It seems as if the elementary psychic fact were not *thought* or *this thought* or *that thought* but *my thought*, every thought being owned. . . . On these terms the personal self rather than the thought might be treated as the immediate datum in psychology. The universal conscious fact is not "feelings and thoughts exist" but "*I* think and *I* feel". . . .

In other words, conscious experience requires that a particular kind of connection be made between the mental representation of some current or past event, and a mental representation of the self as the agent or patient, stimulus or experiencer, of that event (Kihlstrom, 1995). It follows from this position that in order to understand the vicissitudes of consciousness, and of mental life in general, we must also understand how we represent ourselves in our own minds, how that mental representation or self gets linked up with mental representations of ongoing experience, how that link is preserved in memory, and how it is lost, broken, set aside, and restored.

The research program, described in this article (Kihlstrom & Cantor, 1984; Kihlstrom, *et al.*, 1988; Kihlstrom & Klein, 1994; Kihlstrom, Marchese & Klein, 1995; Klein & Loftus, 1993) begins with a simple question: What does the self look like? In answering this question, we assume that the self is one's mental representation of oneself—or, put another way, that the self represents our own knowledge of ourselves. Thus, the answer to the

[a]The point of view represented in this paper is based in part on research supported by Grant No. MH-35856 from the National Institute of Mental Health.

[b]Address for correspondence: Department of Psychology, Yale University, P.O. Box 208205, New Haven, Connecticut 06520-8205. Phone (203) 432-2596; fax (203) 432-7172; e-mail: Kihlstrm@Minerva.cis.Yale.edu

question of what the self looks like comes from the vast literature within cognitive psychology concerning the manner in which knowledge is represented in the mind.

THE SELF AS A CONCEPTUAL STRUCTURE

Within personality and social psychology, the self-concept is commonly taken as synonymous with self-esteem, but within the social-intelligence framework of personality (Cantor & Kihlstrom, 1987) the self-concept can be construed simply as one's concept of one's self, a concept no different, in principle, than one's concept of *bird* or *fish*. From this perspective, the analysis of the self-concept can be based on what cognitive psychology has to say about the structure of concepts in general (Smith & Medin, 1981).

From the time of Aristotle until only just recently, concepts were characterized as *proper sets*: summary descriptions of entire classes of objects in terms of defining features which were singly necessary and jointly sufficient to identify an object as an instance of a category. Thus, the category *birds* includes warm-blooded vertebrates with feathers and wings, while the category *fish* included cold-blooded vertebrates with scales and fins. But both philosophical considerations and the results of experiments in cognitive psychology have persuaded us not to think about concepts in terms of proper sets and defining features, but rather in terms of family resemblance, in which category members tend to share certain attributes, but there are no defining features as such. According to this view of categories as *fuzzy sets*, category instances are summarized in terms of a *category prototype* which possesses many, but not necessarily all, of the features which are characteristic of category membership.

In the late 1970s and early 1980s the idea that the self, too, was represented as a category prototype was popular, and some very interesting experiments were done based on the assumption (Rogers, 1981). But how does one talk about family resemblance, or abstract a prototype, when there is only one member of the category—oneself? The notion of self-as-prototype, taken literally, seems to imply that the self is not unitary or monolithic. We do not have just one self: rather, each of us must have several different selves, the characteristic features of which are represented in the self-as-prototype.

A dramatic example of this may be found in *The Three Faces of Eve*, a case of multiple personality disorder, played with Academy-Award-winning skill by Joanne Woodward in the 1957 film of the same title (cf. Thigpen & Cleckley, 1954). Two personalities, Eve White and Eve Black, were separated by an asymmetrical amnesia: Eve White knew nothing about Eve Black, but Eve Black knew all about Eve White. Testing with the semantic differential technique revealed that the two Eves had markedly different self-con-

cepts. Eve White saw herself as bad, passive, and weak, while Eve Black saw herself as good, strong, and active (Osgood & Luria, 1954). During therapy a third personality, Jane, emerged, who seemed to blend the best qualities of the two Eves: she was strong and active, and recognized both her good and bad sides; in that sense, Jane might be construed as something like a self-prototype.

But one does not have to have a dissociative disorder to have a multiplicity of selves. Traditional personologists assume that behavior is broadly stable over time and consistent over space, and that this stability and consistency reflect traits which lie at the core of personality. Viewed cognitively, the self might be viewed as the mental representation of this core. But social psychologists have argued that behavior is extremely flexible, varying widely across time and place. If so, then the self-concept should represent this context-specific variability, so that each of us possesses a repertoire of context-specific self-concepts—a sense of what we are like in different classes of situations (Kihlstrom, Marchese & Klein, 1995). The self-as-prototype might be abstracted from these contextual selves. Thus, we might begin to think about a hierarchy of selves, with more or less context-specific selves at lower levels, and a very abstract prototypical self at the highest level.

This is all well and good, but maybe there is not a prototype after all. Another trend in cognitive psychology has been to abandon the notion entirely that concepts are summary descriptions of category members (Medin, 1989). Rather, according to the *exemplar view of categories*, concepts are only a collection of instances, related to each other by family resemblance perhaps, and with some instances being in some sense more typical than others, but lacking a unifying prototype at the highest level. Some very clever experiments have lent support to the exemplar view, but as yet it has not found its way into research on the self-concept. Nevertheless, the general idea of the exemplar-based self-concept is the same as that of the context-specific self, only lacking hierarchical organization or any summary prototype.

The three views of categorization presented so far—proper sets, fuzzy sets, and exemplars—all assume that the heart of categorization is the judgment of similarity. That is, instances are grouped together into categories because they are in some sense similar to each other. But similarity is not the only basis for categorization. It has been proposed that categorization is also based on one's theory of the domain in question, or, at least, that people's theories place constraints on the dimensions which enter into their similarity judgments (Medin, 1989). Application of this *theory-driven* view of categorization to the self was anticipated more than 20 years ago by Epstein (1973, p. 407), who argued that the self-concept is "a theory that the individual has unwittingly constructed about himself as an experiencing, functioning individual, and . . . part of a broader theory which he holds with respect to his entire range of significant experience."

Epstein's views have not been translated into programmatic experimental research on the self, but we can perhaps see examples of theory-based construals of the self in the variety of "recovery" movements in American society today (Kaminer, 1992). Whether we are healing our wounded inner child, freeing our inner hairy man, dealing with codependency issues, or coping with our status as an adult child of alcoholics or a survivor of child abuse, what links us to others, and literally constitutes our definition of ourselves, is not so much a set of attributes as a theory of how we got the way we are. And what makes us similar to other people of our kind is not so much that they resemble us but that they went through the same kind of formative process, and hold the same theory about *themselves* as we do of *ourselves.* Dysfunctional or not, it may well be that we all have theories—we might call them origin myths—about how we became what we are, and these theories are important parts of our self-concept. Such self-theories could give unity to our context-specific multiplicity of selves, explaining why we are one kind of person in one situation, and another kind of person in another (Kihlstrom *et al.*, 1995).

THE SELF AS A STORY

Something of the flavor of the self-as-theory is provided by a second form of mental representation, knowledge as stories. Recently, Schank and Abelson (1995, p. 80) have asserted that "from the point of view of the social functions of knowledge, what people know consists almost exclusively of stories and the cognitive machinery necessary to understand, remember, and tell them." As they expand on the idea (p. 1),

> (1) Virtually all human knowledge is based on stories constructed around past experiences; (2) New experiences are interpreted in terms of old stories; (3) The content of story memories depends on whether and how they are told to others, and these reconstituted memories form the basis of the individual's *remembered self.* Further, shared story memories within social groups define particular social selves, which may bolster or compete with individual remembered selves.

Schank and Abelson concede that knowledge also can be represented as facts, beliefs, lexical items like words and numbers, and rule systems (like grammar), but they also argue that, when properly analyzed, non-story items of knowledge actually turn out not to be knowledge (for example, they may constitute indexes used to organize and locate stories). Their important point is that from a functional standpoint, which considers how knowledge is used and communicated, knowledge tends to be represented as stories.

The idea of knowledge as stories, in turn, is congruent with Pennington and Hastie's story model (1993) of juror decision-making. According to Pennington and Hastie, jurors routinely organize the evidence presented to them into a story structure with initiating events, goals, actions, and consequences. According to Schank and Abelson (1995), each of us does the same sort of thing with the evidence of our own lives. From this point of view, the self consists of the stories we tell about ourselves—stories which relate how we got where we are, and why, and what we have done, and what happened next. We rehearse these stories to ourselves to remind ourselves of who we are; we tell them to other people to encourage them to form particular impressions of ourselves; and we change the stories as our self-understanding, or our strategic self-presentation, changes. When stories aren't told, they tend to be forgotten—a fact dramatically illustrated by Nelson's studies (1993) of the development of autobiographical memory in young children. Furthermore, when something new happens to us, the story we tell about it, to ourselves and to other people, reflects not only our understanding of the event, but our understanding of ourselves as participants in that event. Sometimes, stories are assimilated to our self-understanding; on other occasions, our self-understanding must accommodate itself to the new story. When this happens, the whole story of our life changes.

THE SELF AS AN IMAGE

Our discussion of the self as concept and as story illustrates a strategy that we have found particularly useful in our work: beginning with some fairly informal, folk-psychological notion of the self-concept, we see what happens when we apply our technical understanding of what that form of self-knowledge looks like. Much the same attitude (or, if you will, heuristic) can be applied to another piece of ordinary language: the self-image. Schilder (1938), p. 11) defined the self-image as "the picture of our own body which we form in our mind, that is to say, the way in which the body appears to ourselves." What follows from this?

First, there is the question whether, in talking about our mental images of ourselves, we should be talking about mental images *at all.* Beginning in the 1970s, a fairly intense debate raged about the way in which knowledge is stored in the mind. At this point, most cognitive psychologists are comfortable distinguishing between two forms of mental representation: meaning-based and perception-based (Anderson, 1983). Meaning-based representations store knowledge about the semantic relations among objects, features, and events, and take the form of propositions—primitive sentence-like units of meaning which omit concrete perceptual details. We will discuss these later. For now, we wish to focus on perception-based representations, which store knowledge about perceptual structure.

Perception-based representations are relatively unstudied in social cognition, but it is quite clear that we have them. For example, we have visual and auditory images of the faces and voices of people we know, which permit us to recognize these people as familiar. That we have perception-based representations of others makes it more likely that we have perception-based representations of *ourselves* as well. In fact, Head (1926) coined the term *body schema* to refer to our postural models of our own bodies—models which allow us to maintain stability and adjust to our environment, and which are distorted in the classical experiments on prism adaptation. The fact that we can adjust our movements when our vision is distorted, and that these adjustments persist when objective stimulus conditions change, indicates that we have internal representations of our bodies, and their parts, which are independent of sensory stimulation.

As with the self-concept, the self-image can be illustrated with clinical data. Acute schizophrenics often complain of distortions in their perception of their own bodies. In *autopagnosia*, a neurological syndrome associated with focal lesions in the left parietal lobe, the patient can name body parts touched by the examiner, but cannot localize body parts on demand. In *phantom limb pain*, amputees perceive their lost arms and legs as if they were still there. In *body dysmorphic disorder*, the patient complains of bodily defects where there really aren't any. In *eating disorders* such as anorexia and bulimia, the sufferer sees fat where the body is objectively normal, lean, or even gaunt.

But little of this clinical folklore has been studied experimentally. In the laboratory, studies of the self-image qua image are very rare. One exception is a fascinating study by Mita, Derner, and Knight (1977) on the mere exposure effect (Zajonc, 1968), in which subjects view a series of unfamiliar objects (for example, nonsense polygons or Turkish words), and later make preference ratings of these same objects and others which had not been previously presented. On average, old objects tend to be preferred to new ones, and the extent of preference is correlated with the number of prior exposures. In the experiment of Mita *et al.* (1977), subjects were presented with pairs of head-and-shoulder photographs of themselves and their friends, and asked which one they preferred. In each pair, one photo was the original, and the other was a left–right reversal. The result was that when viewing photos of their friends, subjects preferred the original view (that is, the view as seen through the lens of the camera), but when viewing photos of themselves, the same subjects preferred the left-right reversal (that is, the view as would be seem in a mirror). Thus, our preferences for pictures match the way we typically view ourselves and others. Mita took this as evidence for the mere exposure effect, which it is; but it is also evidence for a highly differentiated self-image which preserves information about both visual details and the spatial relations among them.

THE SELF AS ASSOCIATIVE NETWORK

For most of its history the study of memory has been the study of verbal learning. Accordingly, many psychologists have come to think of memory as a set of words (or phrases or sentences), each representing a concept, joined to each other by associative links representing the relations between them, the whole system forming an associative network of meaning-based knowledge (Anderson, 1983)—Schank and Abelson's (1995) theory of knowledge as stories is explicitly opposed to this conventional thinking. It is also commonplace to distinguish between two broad types of verbal knowledge stored in memory (Tulving, 1983). *Episodic* memory is autobiographical memory for a person's concrete behaviors and experiences: each episode is associated with a unique location in space and time. *Semantic* memory is abstract, generic, context-free knowledge about the world. Almost by definition, episodic memory is part of the self-concept, because episodic memory is about *the self*: it is the record of the individual person's past experiences, thoughts, and actions. But semantic memory can also be about the self, recording information about physical and psychological traits of the sort that might be associated with the self-concept.

Within the verbal-learning tradition, knowledge about other people has been studied extensively in a line of research known as *person memory* (Hastie, Ostrom, Ebbesen, Wyer, Hamilton, & Carlston, 1980). Several different models of person memory have been proposed (Kihlstrom & Hastie, 1993), and some of these have been appropriated for the study of memory for one's self (Kihlstrom & Klein, 1994; Klein & Loftus, 1993). The simplest person-memory model is an associative network with labeled links. Each person (or perhaps only his or her name) is represented as a single node in the network, and knowledge about that person is represented as fanning out from that central node. The person-nodes are also connected to each other, to represent relationships among them, but that is another matter. The point is that in this sort of model the various nodes are densely interconnected, so that each item of knowledge is associatively linked to a great number of other items. In theory, the interconnections among nodes form the basis for associative priming effects, in which the presentation of one item facilitates the processing of an associatively related one.

Of course, knowledge about a person can build up rather fast: consider how much we know about even our casual acquaintances. According to the spreading activation theory that underlies most associative network models of memory (Anderson, 1983), this creates a liability known as the *fan effect*: the more information you know about someone or something, the longer it takes to retrieve any particular item of information.

Is there any way around the fan effect in person memory? One possibility which has been suggested is that our knowledge about ourselves and others

(especially those whom we know well) is organized in some way—perhaps according to its trait implications. There is some evidence that organization does abolish the fan effect (Smith, Adams & Schorr, 1978), but this evidence is rather controversial, and some have concluded that memory isn't really organized in this manner after all (Reder & Anderson, 1980). Nevertheless, the notion of a hierarchically organized memory structure is so sensible that many person-memory theorists, such as Hamilton and Ostrom, have adopted it anyway (see, for example, their contributions to Hastie *et al.*, 1980; see also Kihlstrom & Hastie, 1993).

How can we generalize from person memory to the structure of memory for the self? Actually, there are three possibilities here, distinguished particularly with respect to the relations between behavioral and trait information, or episodic and semantic knowledge. FIGURE 1A depicts the conventional wisdom: that memory for specific behavioral episodes is organized by their trait implications. In this model, nodes representing traits fan off the node representing the self, and nodes representing specific episodes which exemplify these traits fan off the trait-nodes. This *hierarchical* model implies that retrieval has to pass through traits to effect access to information about behaviors. Thus, traits will be activated in the course of gaining access to information about behaviors. FIGURE 1B depicts an alternative model, which holds that the self contains only episodic knowledge about experiences and behaviors, and that semantic knowledge about traits is known only indirectly, by inference. One such inferential process would involve sampling autobiographical memory, and integrating the trait implications of the memories so retrieved. In this *self-perception* model, retrieval must pass through behaviors in order to reach traits. Put another way, nodes representing behaviors will be activated in the course of recovering—to be precise, in the course of *constructing*—information about traits. FIGURE 1C depicts the original default model, in which episodic self-knowledge is encoded independently of semantic self-knowledge—or, put another way, in which knowledge of behaviors is represented separately from knowledge of traits.

An extensive series of studies by Klein and Loftus (1993) has produced a compelling comparative test of these models. These studies adapted for the study of the self the *priming* paradigm familiar in studies of language and memory, in which the presentation of one item facilitates the processing of another associatively related item. Subjects were presented with trait adjectives as probes, and performed one of three tasks. In the *define* task, they simply defined the word; in the *describe* task, they rated the degree to which the term described themselves; in the *recall* task they remembered an incident in which they displayed behavior relevant to the trait. For each probe, two of these tasks were performed in sequence—for example, *describe* might be followed by *recall*, or *define* by *recall*, or *recall* by *describe*. There were nine possible sequences, and the important datum was the subject's response latency when asked the second question of each pair.

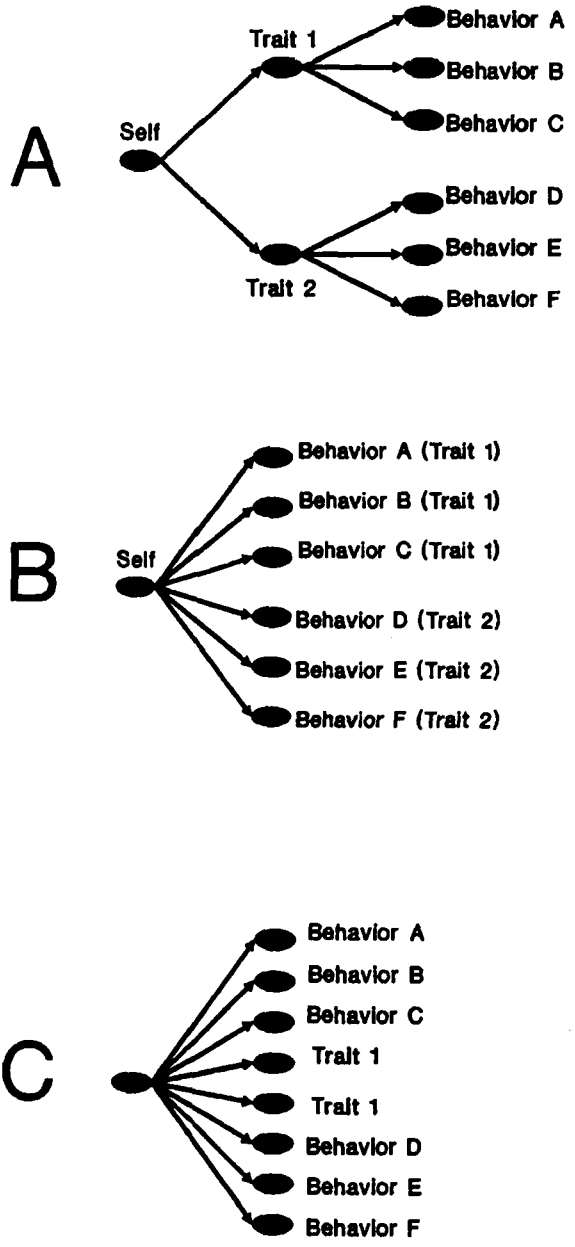

FIGURE 1. Three formats for representing episodic (behavioral) and semantic (trait) information about the self. **(A)** behaviors organized by traits; **(B)** traits inferred from behaviors, **(C)** behaviors and traits represented independently.

Because priming occurs as a function of overlap between the requirements of the initial task and the final task, systematic differences in response latencies will tell us whether activation passes through trait information on the way to behaviors, or vice versa, or by neither way. When the two processing tasks were identical, there was a substantial repetition priming effect of the first one on the second. But when Klein and Loftus (1993) examined the effect of *recall* on *describe*, they saw no evidence of semantic priming compared to the effects of the neutral *define* task. Nor was there semantic priming when they examined the effect of *describe* on *recall* (again, compared to the effects of the neutral *define* task). Contrary to the hierarchical model, the retrieval of autobiographical memory does not automatically invoke trait information. And contrary to the self-perception model, retrieval of trait information does not automatically invoke memory for behavioral episodes. Because self-description and autobiographical retrieval do not prime each other, Klein and Loftus concluded that items of episodic and semantic self-knowledge must be represented independently of each other.

This conclusion has received some additional support from an unusual source: Tulving's study of a neurological patient with an exceptionally dense amnesia (Tulving, 1993). The patient, known as K. C., was involved in a serious motorcycle accident at age 30. As a consequence, he suffered a dense anterograde amnesia affecting his memory for events which occurred since the time of the accident, but he also suffered a profound retrograde amnesia covering his entire personal history. To this day, K. C. remembers absolutely nothing of any of his experiences since birth. Interestingly, K. C. also underwent a marked change in personality. K. C.'s premorbid personality was quite extroverted, but his postmorbid personality was quite introverted. Tulving learned of the work of Klein and Loftus (1993) and wondered whether K.C. had any appreciation of his current personality.

With Klein's cooperation, Tulving (1993) assembled a list of personality trait adjectives and administered them to K. C. and his mother in a two-alternative forced-choice format. K. C. was asked to choose the adjective from each pair which better described himself as he is now. His mother was asked to complete the task twice, once for his premorbid personality, once for his postmorbid personality. K. C.'s choices were reliable, and agreed with his mother's judgments of his postmorbid personality, but disagreed with her judgments of his premorbid personality. This means that he has generic, semantic knowledge of what he is like *now* even though he cannot consciously recollect anything he has ever said or done. So he has acquired semantic knowledge of himself without retaining any episodic knowledge of the specific actions and experiences on which that self-knowledge is based.

More recently, Klein, Loftus, and Kihlstrom (1995) obtained similar findings from another patient, W. J., who, as a result of a severe head injury, suffered a temporary loss of her episodic memory. Both during her amnesia and following its remission, W. J. was asked to make trait judgments about

herself. Because responses given when she could access episodic self-knowledge were reliable compared with responses given when she could not do so, Klein and his colleagues concluded that the loss of her trait-relevant behavioral memories did not greatly affect the availability of her trait self-knowledge.

These two cases of amnesia, taken together, offer strong support for the general proposition that knowledge of one's traits is represented and accessed independently from knowledge of one's behavior. Tulving (1993) interprets this outcome as support for his notion that episodic and semantic memory are represented in different memory systems of the brain. In the present context, it is also generally inconsistent with the self-perception theory of self-knowledge, and is consistent with the proposition that trait self-knowledge is encoded independently of behavioral self-knowledge.

WHAT DOES THE SELF LOOK LIKE?

To this question, cognitive psychology offers four answers:

1. *Viewed as a concept*, the self is a fuzzy set of context-specific selves, perhaps united by a prototypical self, or by an overarching theory of why we are one person in some situations, and another person in others.
2. *Viewed as a story*, the self is a narrative, or a set of narratives, which we have constructed, rehearsed to ourselves, and related to others, which answers Gaugin's three questions: Where do we come from? What are we? Where are we going? And perhaps a fourth: What does it all mean?
3. *Viewed as an image*, the self is a percept-based representation which stores knowledge about both spatial relationships and visual details about our faces, bodies and gestures.
4. *Viewed as an associative network*, the self is a bundle of propositions about our abstract traits, and our specific experiences, thoughts and actions, in which semantic self-knowledge is represented independently of episodic self-knowledge.

One feels a little like someone watching four blind people describe an elephant. But this is all right: why shouldn't self-knowledge be rich and multifaceted? In the end it's all neural connections anyway, but at the psychological level of analysis, which is where we as psychologists *should* be operating, there's no reason why we shouldn't consider different representational formats for different kinds of self-knowledge.

One thing is sure: the self, which plays such an important role in social interaction, is also a knowledge structure represented in the mind of the actor. Far from being a mystical entity, it appears that we can study and understand the self using the conceptual and methodological tools of modern

cognitive psychology—bringing what we know about category structure, story understanding, image processing, and priming effects to bear on the self. And we can gather data for research on the self from a wide variety of sources: from conventional personality and social psychology, from conventional cognitive psychology, from cognitive neuropsychology, and from clinical psychology. So, both conceptually and empirically, the study of the self seems to serve a unifying function in psychology, bringing cognitive, social, personality, developmental, and clinical psychology together in common cause—just as William James thought it would.

ACKNOWLEDGEMENTS

We thank Robert Abelson, Mahzarin Banaji, Marilyn Dabady, David DiSteno, William Hirst, and Judith Loftus for their comments.

REFERENCES

ANDERSON, J. (1983). *The architecture of cognition.* Hillsdale, NJ: Erlbaum.

CANTOR, N. & KIHLSTROM, J. F. (1987). *Personality and social intelligence.* Hillsdale, NJ: Erlbaum.

EPSTEIN, S. (1973). The self-concept revisited: Or a theory of a theory. *American Psychologist, 28,* 404–416.

HASTIE, R., OSTROM, T. M., EBBESEN, E. B., WYER, R. S., HAMILTON, D. L. & CARLSTON, D. E. (EDS.). (1980). *Person memory: The cognitive basis of social perception.* Hillsdale, NJ: Erlbaum.

HEAD, H. (1926). *Aphasia and kindred disorders of speech.* Cambridge, MA: Cambridge University Press.

JAMES, W. (1981). *Principles of psychology,* (p. 221). Cambridge, MA: Harvard University Press. Original work published 1890.

KAMINER, W. (1992). *I'm dysfunctional, you're dysfunctional: The recovery movement and other self-help fashions.* Reading, MA: Addison-Wesley.

KIHLSTROM, J. F. (1995). Consciousness and me-ness. In J. Cohen & J. Schooler (Eds.), *Scientific approaches to the question of consciousness* (in press). Hillsdale, NJ: Erlbaum.

KIHLSTROM, J.F. & CANTOR, N. (1984). Mental representations of the self. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 17, pp. 1–47). New York: Academic Press.

KIHLSTROM, J. F., CANTOR, N., ALBRIGHT, J. S., CHEW, B. R., KLEIN, S. B. & NIEDENTHAL, P.M. (1988). Information processing and the study of the self. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 21, p. 145–177). San Diego, CA: Academic Press.

KIHLSTROM, J. F. & HASTIE, R. (1993). Mental representations of persons and personality. In R. Hogan, J.H. Johnson, & S. R. Briggs (Eds.), *Handbook of personality psychology.* Orlando, FL: Academic Press.

KIHLSTROM, J. F. & KLEIN, S. B. (1994). The self as a knowledge structure. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 1, pp. 154–208). Hillsdale, NJ: Erlbaum.

KIHLSTROM, J. F., MARCHESE, L. & KLEIN, S. B. (1995). Situating the self in interpersonal space. In D. Jopling & U. Neisser (Eds.), *The conceptual self in context.* Cambridge, UK: Cambridge University Press.

KLEIN, S. B. & LOFTUS, J. (1993). The mental representation of trait and autobiographical knowledge about the self. In R.S. Wyer & T.K. Srull (Eds.), *Advances in social cognition* (Vol. 5, pp. 1–49). Hillsdale, NJ: Erlbaum.

KLEIN, S. B., LOFTUS, J. & KIHLSTROM, J. F. (1995). Self-knowledge of an amnesic patient: Toward a neuropsychology of social cognition. *Journal of Experimental Psychology: General,* in press.

MEDIN, D. L. (1989). Concepts and conceptual structure. *American Psychologist, 44,* 1469–1481.

MITA, T. H., DERNER, M. & KNIGHT, J. (1977). Reversed facial images and the mere-exposure hypothesis. *Journal of Personality & Social Psychology, 35,* 597–601.

NELSON, K. (1993). The psychological and social origins of autobiographical memory. *Psychological Science, 4,* 7–14.

OSGOOD, C. E. & LURIA, Z. (1954). A blind analysis of a case of multiple personality using the semantic differential. *Journal of Abnormal Psychology, 49,* 579–591.

PENNINGTON, N. & HASTIE, R. (1993). The story model for juror decision making. In R. Hastie (Ed.), *Inside the juror: The psychology of juror decision making* (pp. 192–221). Hillsdale, NJ: Erlbaum.

REDER, L. M. & ANDERSON, J. R. (1980). A partial resolution of the paradox of interference: The role of integrating knowledge. *Cognitive Psychology, 12,* 447–472.

ROGERS, T. B. (1981). A model of the self as an aspect of human information processing. In N. Cantor & J. F. Kihlstrom (Eds.), *Personality, cognition, and social interaction* (pp. 193–214). Hillsdale, NJ: Erlbaum.

SCHANK, R. C. & ABELSON, R. P. (1995). Knowledge and memory: The real story. In R.S. Wyer (Ed.), *Advances in social cognition* (Vol. 8, pp.1–85). Hillsdale, NJ: Erlbaum.

SCHILDER, P. (1938). *Image and appearance of the human body.* London: Kegan Paul, Trench, and Tubner.

SMITH, E. E., ADAMS, N. & SCHORR, D. (1978). Fact retrieval and the paradox of interference. *Cognitive Psychology, 10,* 438–464.

SMITH, E. E. & MEDIN, D. L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

THIGPEN, C. H. & CLECKLEY, H. (1954). A case of multiple personality. *Journal of Abnormal Psychology, 49,* 135–151.

TULVING, E. (1983). *Elements of episodic memory.* Oxford: Oxford University Press.

TULVING, E. (1993). Self-knowledge of an amnesic individual is represented abstractly. In T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition* (Vol. 5, pp.147–156). Hillsdale, NJ: Erlbaum.

ZAJONC, R.B. (1968). Attitudinal effects of mere exposure. *Journal of Personality & Social Psychology Monographs, 9* (2, Pt. 2), 1–27.