Some Correspondences and Similarities of Shamanism and Cognitive Science: Interconnectedness, Extension of Meaning, and Attribution of Mental States

Timothy L. Hubbard

Texas Christian University
Department of Psychology
TCU Box 298920
Fort Worth, TX 76129
voice: 817-257-7410
email: t.hubbard@tcu.edu
fax: 817-257-7681

Abstract

Correspondences and similarities between ideas in shamanism and ideas in contemporary cognitive science are considered. The importance of interconnectedness in the web of life worldview characteristic of shamanism and in connectionist models of semantic memory in cognitive science, and the extension of meaning to elements of the natural world in shamanism and in distributed cognition, are considered. Cognitive consequences of such an extension (e.g., use of representativeness and intentional stance heuristics, magical thinking, social attribution errors, and social in-group/out-group differences) are discussed. It is suggested that attributions of mental states, beliefs, and desires to a computer on the basis of behavioral measures (e.g., the Turing test) is consistent with the extension of meaning and intentionality to nonhuman elements of the natural world in shamanism. In general, the existence of such correspondences and similarities suggests that elements of shamanism may reflect cognitive structures and processes that are also used by nonshamans and in nonshamanic settings. *Key words: shamanism, cognition, intentional stance, extension of meaning*

Western science has often considered shamanic ideas and experience to reflect psychopathology or regressive behavior (e.g., Devereaux 1961; Silverman 1967; but see Noll 1983; Walsh 2001), charlatanism or trickery (e.g., Hansen 2001; Warner 1980), or the lack of a Western education (e.g., depictions in Frazer 1963; Levi-Strauss 1966; Tylor 1871); however, of late there has been a resurgence of interest in the possible validity of shamanic ideas and phenomena (e.g., Harner 1990; Krippner 2000, 2002; Ripinsky-Naxon 1993; Winkelman 2000). The purpose here is to consider correspondences and similarities between ideas in shamanism and ideas in contemporary cognitive science. Part I focuses on similarities of the web of life worldview of shamanism and neural network models of memory in cognitive science, and argues that meaning is extended into the natural world both in shamanism and in cognitive science. Part II focuses on the social consequences of such an extension of meaning into the natural world, the attribution of intentionality to nonhuman elements of the natural world, and how such consequences and attributions are consistent with theories in social cognition. Part III suggests the attribution of cognitive abilities to so-called "artificially intelligent" computer systems in cognitive science reflects the same type of extension of meaning and intentionality as that attributed by shamans to nonhuman elements of the natural world.

Part I: Webs of Life and Neural Networks

One area in which the views of cognitive science and of shamanism exhibit correspondences and similarities is in the use of network models to describe the organization and functioning of their respective domains. Furthermore, properties of network models of semantic memory in cognitive science have drawn increasingly closer to properties of the web of life view of the natural world held by shamans.

The Web of Life

The *web of life* is an ancient idea that emphasizes the interdependence and interconnection of all living forms, and this view of the natural world characterizes pre-industrial or hunter-gatherer societies in which shamanism is more likely to be found. Different species are seen as different nodes in a larger web, and these different nodes are connected by various types of linkages that reflect how the actions, presence, or absence of members of any given species depends upon the actions, presence, or absence of members of other species. The imagery of a web of life suggests a complex tapestry in which the success and vitality of each organism is dependent upon the success and vitality of other organisms (e.g., without the oxygen given off by plants, animals would die, but without the carbon dioxide given off by animals, plants would die; similarly, if there are too many predators, prey die off, but if there are not enough predators, prey overrun their niche). The integrity and functioning of a web requires that all elements thus be in the proper place and relation to each other. If a web is sufficiently damaged, then it will cease to function as a web, and even elements that were not previously damaged may no longer be able to function without input or support from the damaged elements (for an easily accessible account, see Capra 1996).

Given the importance of interconnection in the web of life, the success and vitality of other species is in the best interest of an individual and of that individual's species, because the success and vitality of other species could directly (e.g., as a food source) or indirectly (e.g., by preying on predators) influence the vitality and success of that individual and of that individual's species.[1] Because of the relationships of different species with other elements of the web of life, a concern with the well-being of other species and with the health of the web does not have to contend with traditional difficulties of accounting for altruism and other pro-social behaviors within a Darwinian framework. Thus, a concern with the larger web of life may be consistent with evolutionary principles. Along these lines, the idea of the web of life, and the resultant notion that the functioning of different organisms may contribute to and shape the survival of the more general ecosystem, is also consistent with the Gaia hypothesis of how life on Earth regulates its environment (for discussion of Gaia, see Lovelock 1979, 1988; Schneider and Boston 1991). This larger Gaian notion is consistent with shamanic views of the whole of being as "an immense signal system" (Kalweit 1992) and of the universe as immersed in meaning (Eliade 1964).

Semantic Memory

The imagery of the web of life, and the interdependence and interconnection of all living forms, resembles recent network models of semantic memory within cognitive science. *Semantic memory* involves memory for "meaning," that is, memory for facts and descriptions, and consists of general, abstract, and conceptual knowledge (e.g., Schacter 1996; Tulving 1983). It is with recent connectionist models that the similarities of network models of semantic memory and the web of life are the strongest, but rather than jump directly to a consideration of these contemporary models, it would be useful to first present a brief digression on the history of network models of semantic memory. Such a digression will underscore how advances in theorizing about semantic memory have progressively approached the notion of connectivity that is such a critical aspect of the web of life.

Hierarchical Models. In earlier models of semantic memory, information was organized in a hierarchical manner, with subordinate categories at the bottom of the hierarchy and superordinate categories at the top of the hierarchy (e.g., Collins and Quillian 1968). Concepts were represented by nodes, nodes were connected by class inclusion linkages, and information was stored at the most general level possible (see top of Figure 1). When a node received sufficient stimulation, then that node was said to be "activated." How might a node be stimulated or activated? One way is by perceiving some object that corresponds to the information represented by that node (e.g., looking at an automobile would activate the node for "automobile," hearing the word "dog" would activate the node for "dog"). A second way is by directed memory search (e.g., searching memory for nodes that included the information "yellow" and "flies" would activate "canary"). Each node had a threshold level for activation, and more frequent, salient, or recently-activated nodes would have lower thresholds; if the stimulation level of a given node surpassed that node's threshold, the node was activated and the contents of that node were brought into conscious awareness.

Hierarchical models could account for much of the empirical data known at the time those models were developed. Such data were based on the time required to answer simple questions regarding the properties of a given category.[2] For example, the time to respond to the question "Does a canary have skin?" is slower than the time to respond to the question "Does a canary have wings?" Similarly, the time to respond to the question "Does a canary have wings?" is slower than the time to respond to the question "Does a canary sing?" Researchers accounted for these differences by positing that information was stored at the most general level possible (e.g., information regarding "skin" or "wings" was more general to "animals" and "birds," and thus stored at the "animals" and "birds" nodes, respectively, whereas information regarding "sings" was more specific to "canary," and thus stored at the "canary" node), and so more general information required more time to access because it was stored at a more distant node. In general, shorter response times suggested two pieces of information were stored more closely together, and the overall pattern of response times suggested that memory was organized in a hierarchical manner with property information (e.g., has skin, has wings, sings) stored at the most general level (e.g., canary, bird, animal) possible.

Spreading Activation Models. In hierarchical models, all members of a category were represented at the same level and were presumed to be equally good exemplars of that category. However, subsequent research on human conceptual and category learning revealed that not all members of a category are equally good exemplars of that category (for overviews, see Smith and Medin 1981; Rosch and Lloyd 1978). For example, both cats and bats are mammals, but most people consider a cat to be a more typical or prototypical mammal than is a bat. Prototypical members of a category differ from nonprototypical members of a category in a number of ways (e.g., prototypes are typically learned first and share more characteristic features), but hierarchical models in which all members of a category were at the same level côuldn't account for effects of prototypically. So, subsequent spreading activation models of semantic memory (e.g., Collins and Loftus 1975) discarded the limitation of hierarchical organization and allowed for different types

Figure 1. An illustration of models of semantic memory. The upper panel is adapted from the hierarchical model of Collins and Quillian (1969) and depicts information regarding a small subset of animal species. The middle panel is adapted from the spreading activation model of Collins and Loftus (1975) and depicts a small subset of information related to colors, fruits, flowers, vehicles, and education. The bottom panel is adapted from the connectionist model of Rumelhart and McClelland (1986) and depicts the name, gang affiliation, educational level, age, occupation, and marital status of a number of hypothetical individuals. Adapted from Hubbard (2002).

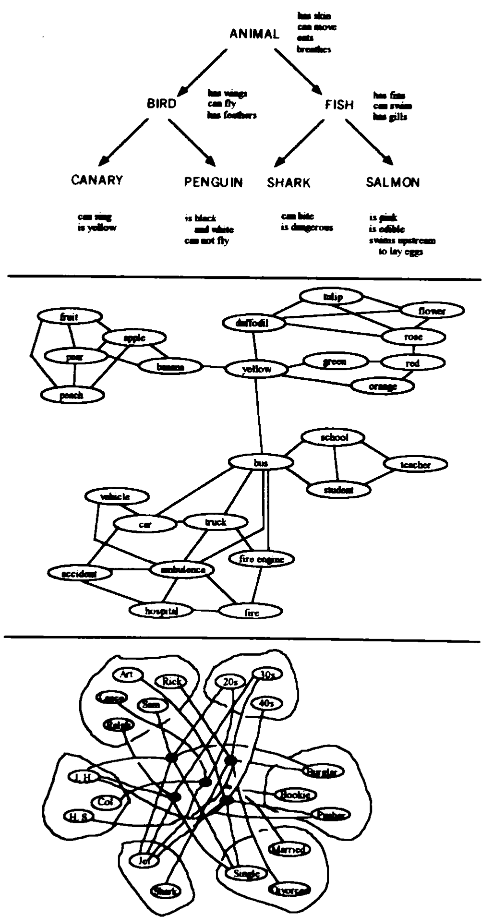

of linkages in which the length of a connection was inversely proportional to the strength of the association (see middle of Figure 1). In such models, prototypes might be more closely or strongly linked to the concept node, and effects of prototypicality could be accounted for.

Spreading activation models also provided a better accounting of data on priming effects than did hierarchical models. "Priming" refers to how the processing of an initial stimulus can facilitate the processing of a subsequent stimulus, and this is most easily illustrated in a *lexical decision task*. In this task, observers view a string of letters on each trial, and they judge whether the string of letters forms a valid word or not (e.g., the string "NURSE" forms a valid word, but the string "JEGHN" does not form a valid word). If the same string is repeated after a few trials (e.g., a person sees "NURSE" on the first trial and subsequently sees "NURSE" on a later trial), the judgment on the later trial will be faster, and this is called *repetition priming* (e.g., Brown et al. 1991; Jacoby and Dallas 1981). If a different but semantically related string is presented on a later trial (e.g., a person sees "NURSE" on the first trial and then sees "DOCTOR" on a later trial), the judgment of the related item on the later trial will be faster, and this is called *semantic priming* (e.g., Meyer and Schvaneveldt 1971; Neely 1990). Analogous types of priming have been found with judgments of shapes (Schacter, Cooper, and Delaney 1990), pictures of objects (Kroll and Potter 1984), and musical chords (Bharucha and Stoeckig 1986).

Spreading activation models account for priming effects by suggesting that as soon as a node is activated, that activation starts to decay and also starts to spread along that node's linkages to related nodes. If the same node is reactivated before the previous activation has completely decayed, then it will not require as much time for subsequent activation to reach the threshold for awareness because any residual activation from the previous stimulation will give that node a "head start." Similarly, if a related node is activated before the activation that previously spread to that node decayed, then that related node will not require as much time for subsequent activation to reach the threshold for awareness because any residual activation that spread from the original node will give the related node a "head start." The stronger the association between two nodes, the shorter the linkage between those nodes, and the stronger the semantic priming. In such spreading activation models, though, knowledge was still stored locally at a single node, and could be stored or retrieved by activation or access of just that single node. The individual node was still the primary unit of meaning, but the importance of the entire network for the functioning of each individual element was beginning to be emphasized, as even an otherwise intact node could not be accessed or would not function adequately if connections to that node were disrupted.

In more contemporary network models of semantic memory, information is not stored locally at a single node, but is instead distributed across the entire network (for overview, see Bechtel and Abrahamsen 1991; Rumelhart and McClelland 1986). For example, rather than having a single node for "canary" which contains a list of features (e.g., is yellow, can sing), there would be separate nodes for each property (e.g., a separate node for "yellow," a separate node for "sings"), and knowledge would be represented by connections between those specific nodes (e.g., the node for "canary" would connect with nodes for "yellow" and "sings," but would not connect [or would have inhibitory connections] with nodes for "blue" and "roars"). Thus, information is distributed across nodes rather than stored at a single node, and knowledge is preserved by the pattern of connectivity between nodes (see bottom of Figure 1). Storage or retrieval of a specific piece of information requires the potential activation and participation of the entire network. These distributed models have been called *connectionist* models (e.g., McClelland 1988), and because they approach more closely to the types of processing thought to be engaged in by the brain (e.g., connections between nodes have been likened to synapses between neurons, e.g., Feldman and Ballard 1982), such models have also been referred to as *neural networks* (e.g., Anderson 1995).

Storing or retrieving a specific piece of information in a connectionist network may involve the activation and participation of the entire network, and so all of the elements may be said to be connected or interrelated and to be involved in cognitive functioning. By stressing the interconnection and the importance of the entire pattern of activation across all nodes in a network rather than just the all-or-none activation of a specific node, contemporary connectionist models of semantic memory in cognitive science parallel the importance of the interconnection and the entire pattern of activity in the web of life. In the case of cognitive science, nodes may correspond to bits of information, individual neurons, or circuits and modules of cells, whereas in the case of shamanism, nodes may correspond to elements of the natural or supernatural world. Just as damage to the web of life diminishes the vitality of species, so too does damage to a connectionist network of semantic memory decrease the completeness or vividness of cognitive functioning; both in the web of life and in connectionist models of semantic memory, we see a graceful degradation of vitality or completeness rather than a complete loss of a few elements with no corresponding effect on the vitality or completeness of the remaining elements.

The Place of "Meaning"

Shamans traditionally placed information and connections throughout the natural world, whereas cognitive models of memory traditionally placed information and connections within the head. However, some cognitive theorists are beginning to suggest that neural networks place meaning "out in the world" as well as in the head (e.g., see discussion in Hardy 1998). By considering meaning as "out in the world" rather than as existing solely "inside the head" or as a unique construction for each individual, connectionist models make functional connections between cognitive processes of the organism and elements of the external natural world more explicit. As a result, such models suggest that individual organisms (such as people) are connected to the elements in their environment. Even so, it might be objected that the connections of the internal neural network are clearly instantiated within visible physiological structures (i.e., synapses), but it is less clear what the equivalent structures are in the case of the external web of life. Such an objection misses a fundamental point: the connections are *functional*, and as such, need not depend upon any single specific form of structural instantiation.

Given that meaning is an intrinsic part of connectionist models of semantic memory and also of the web of life, meaning could provide a common thread with which to weave together neural networks and the natural world. This is consistent with recent approaches to cultural psychology which suggest that cognition is mediated by artifacts and that what we call "mind" must include the environmental context within which cognition is embedded (e.g., see Cole 1996). For example, Bateson (1972) proposed the following thought experiment: "Suppose I am a blind man, and I use a stick. I go tap, tap, tap. Where do I start? Is my mental system bounded at the handle of the stick? Is it bounded by my skin? Does it start halfway up the stick? Does it start at the tip of the stick?" Bateson goes on to argue that such questions are nonsensical unless not just the man, but also the stick, purpose of the tapping, and environment are all taken into account; he concluded that "message pathways" exist outside the skin, and that these pathways and the messages they carry must be included as part of the mental system. Similarly, Hutchins (1995) described how the execution of one type of cognitive task (navigation) might be viewed as a single cognitive activity distributed across several individuals.

These ideas suggest a shift or even dissolution of any functional boundaries between neural structures and the external environment such that cognition is "stretched across mind, body, activity, and setting" (Lave 1993), and are consistent with the idea that an individual's neural network may simply be the microstructure of a single node in a larger functional network extending beyond the individual's body. A stretching of cognition to include "mind, body, activity, and setting" is also consistent with the shamanic view that the universe is saturated with meaning (Eliade 1964), because such a stretching places important aspects of meaning out in the natural world.[3] Patterns of meaning encoded within an individual's neural network are usually experienced as beliefs, desires, and other mental states, and so extending meaning into the natural world may be consistent with the increased likelihood of attributing human-like beliefs, desires, and other mental states to nonhuman elements of the natural world. Such an extension of meaning is consistent with the increased likelihood of animism or panpsychism observed in many preindustrial cultures. With such animism and panpsychism, not just humans, but animals, plants, rocks, and even the sky are seen as alive and as possessing beliefs, desires, and other mental states.

Part II: The Intentional Stance and Social Attribution

The attribution of human-like beliefs, desires, and mental states to nonhuman elements of the natural world has numerous implications for how humans interact with elements of the natural world and for our consideration of the correspondences and similarities of cognitive science and shamanism. Some implications of the idea that nonhuman elements of the natural world possess human-like beliefs, desires, and other mental states, and what such attributions suggest for the relationship between cognitive science and shamanism, will now be considered.

The Intentional Stance

The attribution of beliefs, desires, and other mental states onto an object has been referred to as adopting an "intentional stance" toward that object (see Dennett 1987). In the absence of other information, adoption of an intentional stance toward an object allows an individual to make predictions regarding the subsequent behavior of that object, and those predictions are based upon a conceptualization of the object as a rational agent that possesses beliefs, desires, and other mental states. An intentional stance is different from a "physical stance" (in which expectations regarding an object's behavior are based on physical properties of that object) or a "design stance" (in which expectations regarding an object's behavior are based on knowledge of how that object is designed to behave). Although in contemporary society an intentional stance may be useful in understanding or predicting the behavior of another person, a physical stance is more useful in dealing with most material objects or substances (e.g., mixing paints to match a specific hue), and a design stance is more useful in dealing with manufactured artifacts (e.g., repairing a malfunctioning automobile).

The general approach in shamanism is consistent with the adoption of an intentional stance toward nonhuman elements of the natural world. Adoption of an intentional stance toward the natural world would allow the shaman to use social interactions with other people, as well as knowledge of the self as a rational agent possessing beliefs, desires, and other mental states, as a tool or analogy to aid in attempts to predict, control, or explain the behavior of some unknown element of the natural world. For example, another person might be violent or loud when angry, and so perhaps the violence or loudness of a storm might be interpreted as the anger of nature; a person who has been angered may sometimes be appeased by sacrifices and gifts, and so perhaps an angry nature might be appeased by sacrifices and gifts. Analogously, activities that pleased nature might result in beneficent sunshine or gentle rains. Indeed, this social self might be considered as the soul or spirit (alternatively, a soul or spirit might be considered as the social self). Also, given that a person who is ill does not typically engage in social behavior, illness could be interpreted as a loss of that social agent or self; hence, illness was thought to be caused by a loss of the soul.

A New Heuristic

A shaman who attempted to cure an illness, ensure a successful hunt, or contact a supernatural realm would be faced with a great deal of uncertainty, even if knowledge of the self and of typical social interactions were used as guides. Contemporary research on problem-solving behavior in the face of uncertainty has documented a number of strategies that people use when faced with problems; these strategies are referred to as *heuristics* and are "short-cuts" or "rules-of-thumb" that offer approximate solutions for minimal cognitive effort (Kahneman, Slovic, and Tversky 1982; see also Gigerenzer and Todd 1999). Examples of well-known heuristics include *representativeness* (the likelihood that some object is a member of a set is determined by how closely that object resembles known members of that set), *availability* (the judged frequency of an item or category reflects how easily examples of that item or category may be remembered or perceived), *adjustment and anchoring* (the final estimation of the quantity of some category is influenced by the magnitude of the initial estimate), and *hindsight bias* (an overestimation of how well some outcome could have been predicted once that outcome is known to have happened).

To the list of known or suspected heuristics we might add an *intentionality heuristic*, which suggests that when individuals do not possess sufficient physical knowledge or design knowledge regarding an object, those individuals may attribute intentionality to the object and treat that object as if it were a rational agent that possessed beliefs, desires, and other mental states. Thus, if a shaman did not have sufficient physical knowledge or design knowledge regarding some nonhuman element of nature, then he or she would treat that element of nature as a rational agent that possessed beliefs, desires, and other mental states. Such a behavior is not limited to shamans or those in shamanic cultures, however. Even in our scientific and technological culture, it is a common observation that individuals without specific physical or design knowledge of technological artifacts may anthropomorphize those artifacts (e.g., "that computer just doesn't like me"). Heuristic use of an intentional stance in both shamanic and nonshamanic settings reveals a similarity between the problem-solving of contemporary individuals in Western society and the problem-solving of shamans: both groups appeal to heuristics when faced with problem-solving in the face of uncertainty.

Magical Thinking

Although nonhuman elements of the natural world are clearly related in meaningful ways (e.g., biological taxa, geological strata), the extension of intentionality to nonhuman elements of the natural world suggests that elements of the natural world may be related in additional ways. Some of these additional relationships may be revealed through different types of "magical thinking," which although common in shamanism, are considered immature or pathological by contemporary psychological science (e.g., see Zusne and Jones 1989).

Transmutation. One example of magical thinking involves *transmutation* in which one substance or object is transformed into another substance or object; for example, an individual who dons a false face mask in an Iroquois ceremony or a kachina mask in a Hopi ceremony is considered to literally become the entity represented by that mask. More generally, a human wearing a mask or costume depicting some other entity such as a god or an animal is not simply a human in disguise, he or she *is* that other entity. Such magical transmutation of a limited number of objects or substances is also found in cultures that are primarily nonshamanic (e.g., for a devout Western Catholic receiving communion, the wafer is transmuted into the flesh of Christ), and may involve the representativeness heuristic (e.g., a mask more closely resembles the entity represented by that mask than it resembles the human being wearing the mask). More broadly, substances or objects that would be transmuted already closely resemble the substances or objects into which they would be transmuted (e.g., red wine into blood), and perhaps this resemblance facilitates the belief in transmutation.

Law of Similarity. Another example of magical thinking is the *law of similarity* in which "like affects like;" that is, a desired effect is believed to be more likely to occur if the cause more closely resembles that effect (e.g., see Frazer 1963; Rozin, Millman, and Nemeroff 1986). For example, injuring or destroying an image of a person was believed in many cultures to bring injury or death to that person, as an image of a given individual resembles that individual more than it resembles any other individual (even if only in name). Analogously, some Northwestern Native American peoples believed that placing a painted image of a fish in the water would lead to the subsequent appearance of actual fish. Given that a painted image of a fish resembles a live fish more closely than perhaps anything else (other than another live fish) the shaman could place in the water, we again see the importance of resemblance. Similarly, in the medieval era the mandrake plant was believed to have medicinal value because its roots resemble the shape of the human body. In general, we may speculate that the law of similarity also depends (at least in part) on the representativeness heuristic, as the most effective cause is thought to be the one that most closely resembles the desired effect.

Law of Contagion. A final example of magical thinking involves the *law of contagion* which suggests that objects which have previously been in contact with each other continue to exert an influence on each other even after having been separated (e.g., see Frazer 1963). Such a notion underlies voodoo (i.e., vodoun) and other rituals in which possession of some previous part of a person (e.g., a lock of hair or a fingernail clipping) or of an object previously worn or used by a person (e.g., a locket or a ring) may be used to influence that person. Similarly, the items in the medicine bundle of a shaman may be believed to convey the power of their place of origin to the shaman. Intriguingly, the law of contagion may also be consistent with the possibility of nonlocal effects in quantum physics,[4] because what appears as a nonlocal interaction of subatomic particles if viewed within the framework of the observed particles may be a local interaction if viewed within a larger framework containing both the observer and the observed. Such a larger framework would be consistent with the integration of the neural network of the individual and the web of life discussed earlier, and so we may speculate that effects attributed to contagion in magical thinking (as well as nonlocal effects and any effects attributed to the presence or actions of the observer in quantum physics) might reflect functional connections between the individual's neural network and the natural world (see also Hardy 1998).

Locus of Control

The laws of similarity and contagion may offer (the illusion of) control over forces that would otherwise seem uncontrollable. In addition to allowing a greater perceived prediction or control, appeal to shamanic techniques may redirect the locus of control of an individual from a nonhuman element of the natural world to a human shaman. *Locus of control* refers to an individual's tendency to believe that either external or internal factors control events in his or her life (for review, see Lefcourt 1976); in general, an "internal locus" (in which the individual can predict or control events) leads to more adaptive or healthy responding than does an "external locus" (in which the individual can neither predict nor control events). Shamanic and other magical techniques may provide the perception of a more internal locus of control for the individual, family, or clan, thus helping that individual, family, or clan cope more easily with the difficulties of an uncertain situation. Consistent with this, Malinowski (1954; see also Jahoda 1969) observed that practices based on magical thinking are engaged in only during circumstances in which other knowledge is insufficient for prediction or control, and this is also consistent with the notion that an intentional stance or intentionality heuristic would be engaged in only if physical information or design information was not available.

Social Attribution

Adoption of an intentional stance toward some element of the natural world could make it seem more likely that such an element would be aware of an individual's desires or actions (and thus more likely to be influenced by appeals from that individual), whereas there might be no reason to suspect that an element of the natural world would be aware of an individual's desires or actions (or could be influenced by appeals from that individual) if an intentional stance toward that element was not adopted. Such an intentional relationship between a shaman and the natural world is strongly social, and so it might be expected that principles of social cognition would influence how the shaman interprets actions of nonhuman elements of the natural world. One such group of principles concerns the attributions individuals make regarding the causes of the behavior of other people (for review, see Ross and Fletcher 1985), and whether those other people are perceived to be within one's social group (i.e., "like us") or in a different group (i.e., "not like us") plays a key role in one's interactions with those people. Indeed, this notion could be extended to include the attributions individuals make regarding the behavior of nonhuman elements of the natural world, and also how the attributions regarding the behavior of those elements play a key role in how those individuals react to the natural world.

Fundamental Attribution Error. The intentional stance adopted by shamanism toward the natural world suggests that biases exhibited in our explanation of the behavior of other humans would also be exhibited in the explanation of the behavior of nonhuman elements of the natural world. One such bias is the *fundamental attribution error*, and this reflects a tendency to explain the actions of another human being as resulting from dispositional or internal factors (e.g., traits) rather than as resulting from situational or external factors (see Ross 1977). If another person does not acknowledge us, we are more likely to perceive that person as acting as he did because "he is just that kind of person" (e.g., he is rude or thoughtless) rather than because of any situational factors (e.g., he was in a hurry, he was preoccupied, etc.). If a similar bias is applied to the natural world, then any action of nature would be perceived to result from a trait or disposition of nature rather than from situational factors. For example, a falling rock or a violent storm would be attributed to a trait of nature (e.g., easy to anger) rather than to a situational factor (e.g., erosion of supporting soil or a clash of warm and cold air masses). Indeed, consistent with Malinowski's (1954) observation, a trait view of nature would be expected to be even more strongly favored when knowledge of the situational physical mechanisms (e.g., erosion, atmospheric dynamics) did not exist.

In-group/ Out-group Differences. The shaman's extension of intentionality into the natural world should also result in an expansion of the in-group of the shaman. Typically, the in-group of an individual consists of a small set of humans with whom that individual is in some way affiliated, and these in-group members are evaluated more positively and seen as more similar to the individual than are out-group members (for review, see Stephan 1985). Extension of intentionality into the natural world would make the elements of the natural world appear more similar to the individual, and thus more likely to be contained within the in-group of that individual; withdrawal of intentionality from the natural world would make the elements of the natural world appear less similar to the individual, and thus less likely to be contained within the in-group of that individual. Furthermore, if a portion of the natural world is considered as part of the in-group (e.g., as in totemic relationships), then that portion of the natural world would presumably be treated more positively than would other portions of the natural world or than if that portion is not considered as part of the in-group. This is consistent with observations that in shamanic societies there are prohibitions against hunting, harming, or consuming one's totem animal, and these prohibitions are analogous to the prohibitions against hunting, harming, or consuming other human beings within one's in-group.

Perceived Similarity and Relatedness

The expansion of the in-group to include nonhuman elements entails that similarities between humans and those nonhuman elements now within the in-group would be more likely to be perceived. Indeed, shamanism stresses just such similarities, and the earlier consideration of the similarities of connectionist models of semantic memory and the web of life worldview also bolsters the idea of such similarities between humans and nonhuman elements of the natural world. The use of an intentional stance, as well as the attribution of dispositional or trait factors to nonhuman elements of the natural world, are also both consistent with such a perceived similarity of humans and nonhuman elements. Indeed, if nonhuman elements of the natural world were not perceived as similar to humans, then the shaman's knowledge of self and of other humans would not be used as a tool or analogy in dealing with those nonhuman elements, nor would the behavior of nonhuman elements be explained in ways similar to the explanation of human behaviors. In addition to these social and behavioral examples of perceived similarity and relatedness, linguistic examples may also be found (e.g., the Lakota greeting "Aho-mitakayeoyacin" means "hello to all my relations" and is often used in reference to all other living things and not limited to other humans or even to the immediate family or tribal group).

The perceived similarity and relatedness seen in shamanism is consistent with the idea of similarities and connections (i.e., of "relatedness") between humans and nonhuman animals highlighted by developments in Western science. Discoveries in genetics have revealed that organisms of many different species have a high overlap in their genetic codes, and thus are related through their DNA. More tellingly, neuroscientists with an ultimate goal of understanding human neural function often use giant squid, sea slugs, rats, cats, and nonhuman primates as experimental subjects; such animals are used because their overall nervous system structure is simpler (and so basic principles that govern neural functioning may be more easily studied) or because the invasive experimental procedures necessary for precise experimental control or measurement are usually considered less ethically questionable if carried out on nonhuman animals than if carried out on humans. However, use of such animal models would not be advisable unless neural functioning and elements of neural structure were identical (or at least highly similar). By endorsement of a reductionist strategy, Western science also implies that "we are all related," albeit without the extension of intentionality and the expansion of the in-group characteristic of shamanism.

Part III: Natural Shamanism and Artificial Intelligence

The extension of intentionality to elements of the natural world was not limited to bipedal animal forms similar to humans, but also included diverse elements such as the sky, wind, and rain. The extension of intentionality to a given element of the natural world typically involved a focus on the behavior of that element and the determination of attributions for why that element behaved as it did. This focus on behavior, coupled with a relative lack of concern about structural differences between humans and nonhuman elements of the natural world, suggests that ideas and practices in shamanism may assume (or otherwise be based on) a functionalist foundation. Cognitive science was also historically based on a functionalist foundation, and as a consequence emphasized connections between mental states or between mental states and either sensory inputs or motor outputs (e.g., see Dawson 1998; Johnson-Laird 1988) and did not emphasize the substrate in which such connections would be instantiated or embodied. Thus, both cognitive science and shamanism focused on explaining observable behaviors and were not concerned with the nature of the structures necessary to support those behaviors.[5]

The Legacy from Artificial Intelligence

The functionalist basis of cognitive science resulted in a computer metaphor in which "mind" was equated with "software" and in which "brain" or "body" was equated with "hardware." Software is relatively independent of hardware (e.g., the same program may run in different architectures or operating systems, the same architecture or operating system may run different programs), and so mental processes were conceived to be relatively "disembodied" and independent of specific physical mechanisms. As a consequence, developments in the understanding of human cognition and developments in the ability of humans to construct artificially intelligent devices greatly influenced each other (e.g., see Boden 1987; Crevier 1993; Kurzweil 1992), and some theorists suggested that a computer running the appropriate program could "experience" the same mental states as a human brain running the analogous program (e.g., Newell and Simon 1972; Schank and Abelson 1977). Intriguingly, such an attribution of mental states to nonhuman (and nonorganic) structures such as a computer is fully consistent with the extension of intentionality to nonhuman elements of the natural world in shamanism. In order to appreciate this point more fully, it would be useful to present a brief digression on the debate regarding computer "thinking" and "understanding."

The Turing Test. The most well-known proposal for a method to determine whether a computer could be said to think and understand (i.e., to have mental experiences) in the same way that a human could be said to think and understand is named after the mathematician Alan Turing and referred to as the *Turing test*. The Turing test involves three participants: two humans and one computer. Each participant is placed in a different room, and the only communication between rooms is via typed text messages. One human is given the task of determining which of the other two participants is the other person and which is the computer. This person engages the other participants in exchanges of typed text messages, and he or she is allowed to ask any question, make any statement, and attempt to engage the other participants in conversation on any topic. Based solely on the responses from the other participants, this person must determine which other participant is the other human and which is the computer. The logic of the Turing test is that if the first person cannot tell which other participant is the other human and which is the computer, then he or she is not justified in making any cognitive discrimination between the other two participants. If thought and understanding are attributed to the other human, and the computer's responses are indistinguishable from the other human's responses, then thought and understanding must also be attributed to the computer (for discussion, see Dreyfus 1993; Haugeland 1985; Weizenbaum 1976).

The Chinese Room. One rebuttal to the Turing test has been proposed by John Searle (1980, 1992) and become known as the *Chinese Room*. Imagine a person locked in room, and this person can read and speak only English. In the room is a stack of papers, and each sheet of paper is illustrated with a different apparently meaningless squiggle mark. Also in the room is an instruction manual written in English that specifies which squiggle marks are associated with other squiggle marks. There are two slots in one of the walls of the room, and sheets of paper illustrated with squiggle marks may enter the room through one of the slots. When a sheet containing squiggle marks enters the room, the person in the room then looks up those squiggle marks in the manual, finds the appropriate associated sheets from the stack of papers in the room, and passes those associated sheets out through the other slot. Unbeknownst to the person in the room, the squiggles are actually Chinese characters, and on the outside of the room the slot through which sheets are passed into the room is labeled "input" and the slot through which sheets are passed out of the room is labeled "output." Over time, the person in the room becomes very fast at finding and passing the appropriate sheets out of the Chinese Room. Once this occurs, and from the perspective of a person outside the room, the outputs from the Chinese Room are indistinguishable from those that might be generated by a native Chinese speaker.

The intuition Searle hoped his readers would have is that the person in the Chinese Room would not understand Chinese in the same way that a native speaker of Chinese could be said to understand Chinese. According to Searle, all that the person in the room can do is associate one set of apparently arbitrary and meaningless symbols (i.e., one set of squiggles) with another set of apparently arbitrary and meaningless symbols (i.e. another set of squiggles). By analogy, a similar associating of symbols is all that a computer running a program can do (e.g., if the symbols "2," "+," "2," and "=," are input into a properly functioning computer, then the symbol "4" will be output, but the computer has no concept or understanding of "fourness" per se). In other words, computers are syntactic (i.e., can implement rules for associating one stimulus with another stimulus) but not semantic (i.e., do not understand the meaning of the symbols that are associated). Given that meaning is critical to human forms of understanding, computers therefore could not be said to understand in the same way that humans could be said to understand. The Chinese Room scenario was meant to suggest a purely behavioral criterion such as the Turing test was inadequate as a measure of thinking or understanding; however, the majority of researchers in artificial intelligence have rejected such a conclusion (e.g., see commentary in Searle 1980; also Boden 1988).

Shamanic Elements of Artificial Intelligence

On one reading, the Chinese Room scenario may be interpreted as suggesting meaning is not present in (or understood by) a computer running a program. As noted earlier, the shamanic universe is saturated with meaning, and so the relative lack of meaning in a purely syntactic computer program might seem to suggest the framework of artificial intelligence is incompatible with the framework of shamanism. However, to the extent that semantics may result from syntax (e.g., Chomsky 1957, 1965), an artificial intelligence, as well as a neural network or a web of life, might be able to bootstrap meaning from the initial syntactic pattern of connectivity (e.g., the importance or role of an individual organism or species is determined by the relationship of that organism or species to other elements in the web of life). Given this, the apparent incompatibility between artificial intelligence and shamanism regarding the presence of meaning may be lessened or even eliminated, and cognition in an artificially intelligent system would be distributed or stretched across the (artificial) mind, activity, and setting (cf. Lave 1993) and functionally connected with "message pathways" in the environment (cf. Bateson 1972). Thus, the relationship between artificial intelligence and shamanism would exhibit the same correspondences and similarities as would the relationship between human cognition and shamanism.

The emphasis on functional connections between mental states in artificial intelligence (and in cognitive science more generally) is consistent with several aspects of shamanism. For example, the importance of functional connections between the neural network of an individual and the larger web of life, as well as the possibility of transmutation, are both consistent with the idea of recreating similar or even identical mental states in a variety of different physical substrates (e.g., in silicon chips and in neural tissue). Along these lines, the extension of intentionality to elements of the natural world, as well as the expansion of the social in-group of the shaman to include nonhuman elements of the natural world, is based on the idea that elements of the natural world possess mental states, beliefs, and desires similar to those of humans. This similarity is a functional one that occurs regardless of differences in the physical structure or composition of humans or of other elements in the natural world. In a sense, behaviors of nonhuman elements of the natural world are explained using the same criteria that are used to explain behaviors of other humans.

The functionalist foundation of theory and research in artificial intelligence suggests that behaviors and mental states, rather than physical structure or composition per se, is related to intelligence and understanding. Acceptance by mainstream researchers and theorists in artificial intelligence of a behavioral criterion such as the Turing test for determining whether "thinking" or "understanding" occurs underscores a surprising convergence of cognitive science and shamanism, as the Turing test explicitly suggests mental states and understanding similar to or even identical with those experienced by humans may be experienced by nonhuman (and nonorganic) structures. In the Turing test and in shamanism, mental states, beliefs and desires are attributed to some object on the basis of a perceived similarity of the behaviors of people and the behaviors of that object. Indeed, if we consider intentionality and the social aspects of shamanism discussed earlier, then a strong form of artificial intelligence (in which human-like mental states and understanding are attributed to a computer running the appropriate program) might even appear as a type of shamanism in which the recipient of the extension of meaning and intentionality is an artificially constructed device rather than an element of the natural world.[6]

Given that computer models and simulations of cognitive processes are usually based on verbal protocols (i.e., verbalizations of introspections of human subjects as they completed various tasks, for review, see Ericsson and Simon 1993) or other behavioral measures (e.g., priming), many such models and simulations possess at least a superficial similarity to human cognitive performance. The apparent extension of intentionality to artificially constructed devices, or more specifically, the attribution of thought and understanding to artificially intelligent systems, may thus also reflect a type of magical thinking similar to the law of similarity. In this case, the similarity of the "outputs" of a computer and a human would lead to an assumption of similar causes, that is, similar mental states. However, and as has been noted by numerous critics of strong forms of artificial intelligence, computer simulations of hurricanes do not experience high winds and flooding, and so why should computer simulations of cognitive processes experience mental states? Indeed, the attribution of mental states to computers running programs seems very similar to the attribution of mental states to nonhuman elements of the natural world involved in some action; in both examples, an intentional stance toward a complex and not-completely-understood system is adopted.

Summary and Conclusions

There are numerous correspondences and similarities between ideas in contemporary cognitive science and ideas in shamanism. Connectionist models of semantic memory in cognitive science suggest cognition is distributed across elements in an individual neural network, and recent accounts in cultural psychology and distributed cognition extend the range of this distribution and suggest cognition is distributed across individuals, artifacts, and elements of the natural world. Therefore, the cognitive processes of an individual may be intimately connected to and distributed across elements of the natural world, and both the individual's neural network and the natural world may be saturated with meaning. The extension of meaning into the natural world results in an individual's neural network being functionally integrated into the larger web of life in the natural world, and this is consistent with the emphasis on interconnection in the shamanic worldview. Cognitive models of distributed cognition and shamanic views of the web of life thus converge on similar notions regarding the importance of meaningfulness in the natural world and the interconnectedness of people with elements of the natural world.

The extension of meaning into the natural world is consistent with the idea that nonhuman elements of the natural world experience human-like beliefs, desires, and other mental states. Consistent with this, shamans may heuristically adopt an intentional stance toward elements of the natural world when faced with problem-solving in the face of uncertainty. Certain types of magical thinking are more prevalent in shamanic settings than in nonshamanic settings, but this may reflect a difference in degree rather than in kind, as some types of magical thinking may simply reflect an extreme application of heuristic thinking (e.g., transmutation and the law of similarity might reflect an extreme application of the representativeness heuristic). Shamanic explanations for the behavior of nonhuman elements of nature may exhibit biases such as the fundamental attribution error and in-group/out-group differences that are similar to nonshamanic explanations for the behavior of other humans. These biases, coupled with the general use of an intentional stance, suggest that nonhuman elements of the natural world are perceived as similar or related to the human elements of the natural world. Overall, the use of similar heuristics and similar patterns of attribution in shamanic thought and in nonshamanic thought highlight general patterns in human cognition.

Both shamanism and cognitive science were traditionally based on functionalist foundations in which attributions regarding beliefs, desires, and mental states are based on behavior and in which differences in physical structures are irrelevant to the attribution (e.g., both a storm cloud and a person may act as if angry, even though the physical structures of storm clouds and of persons are very different). Thus, the extension of intentionality to nonhuman elements of the natural world that occurs in shamanism is consistent with the use of behavioral criteria regarding whether artificially intelligent computer systems may be said to think and understand in the same way that humans may be said to think and understand. Similarly, attributions of beliefs, desires, and mental states to an artificially intelligent computer system on the basis of behavioral measures (e.g., the

Turing test) are consistent with attributions of mental states, beliefs, and desires to some element of the natural world on the basis of an extension of intentionality. Furthermore, attributions of beliefs, desires, and mental states to artificially intelligent computer systems on the basis of a behavior (output) that is similar to or even identical to that of a human may simply reflect an extreme application of the representativeness heuristic, and if so, would be consistent with the suggested role of the representativeness heuristic in magical thinking within shamanism.

In sum, there are a surprising number of correspondences and similarities between ideas in shamanism and ideas in cognitive science. In shamanism and in cognitive science, meaning is extended into the world, and the extension of that meaning has clear social consequences for the attributions that a person may make concerning the behaviors or actions of humans and of nonhuman elements of the natural world. Intriguingly, extensions of intentionality and attributions of mental states to nonhuman entities are not limited to shamanism, but also appear to occur in contemporary research and theory in the development of artificial intelligence. Furthermore, even when ideas in shamanism differ from those in cognitive science (e.g., magical thinking), such differences do not reflect pathological or regressive thinking by shamans, but may rather simply reflect a more extreme use of heuristics that are also used by nonshamans and by shamans in nonshamanic settings. Overall, the correspondences and similarities between ideas in shamanism and ideas in cognitive science do not support the hypothesis that shamanism reflects a pathological or regressive type of thought; rather, the correspondences and similarities between ideas in shamanism and ideas in cognitive science suggest that ideas in shamanism may reflect the application or use of structures and processes used in nonshamanic (and normative) thought.

Author Notes

The author thanks Stanley Krippner for comments on a previous draft of the manuscript.

Notes

[1]The interdependence and interconnection of the web of life was beautifully summed up in a quote attributed to Chief Seattle in 1854: "This we know. All things are connected like the blood which unites one family. Whatever befalls the Earth, befalls the sons and daughters of the Earth. Man did not weave the web of life; he is merely a strand in it. Whatever he does to the Earth, he does to himself." Although whether or not Chief Seattle actually spoke those words has been disputed, that dispute does not diminish the intensity or the validity of the description.

[2]Researchers focused on differences in the time required to answer such questions because people were almost always very accurate, and because there weren't any differences in accuracy, a consideration of accuracy rates wasn't informative regarding the organization of memory. However, differences in response time were consistently found, and researchers used these differences to guide and constrain the development of cognitive models.

[3]Such a stretching of cognition is also consistent with (1) the possibility that many natural objects and phenomena may be fractal, as the organization and functioning of networks would be the same at smaller scales (e.g., the nervous system of a single individual) and at larger scales (e.g., a given ecosystem or biome); (2) ecological theories of perception which suggest a dynamic and informationally-rich world presents stimuli that do not require additional processing and that the observer may simply pick up or sample meaning (e.g., Gibson 1979); (3) suggestions that cognition reflects the embodiment of the individual within the natural world (e.g., Varela, Thompson, and Rosch 1992); and (4) the application of dynamic theory to connectionist models referred to as *semantic field theory* and which suggests that all physical systems (including humans and nonhuman elements of the natural environment) involve nested semantic fields (e.g., Hardy 1998). Indeed, such a stretching of cognition might also provide a mechanism for the Native American adage that "we are all related," albeit in computational terms not originally envisioned by Native Americans. Thus, not only is a stretching of cognition consistent with ideas from shamanism, it is also consistent with ideas in several other domains.

[4]At the quantum level, particles that are separated may behave as if connected, and this is referred to as a *nonlocal effect* or as *nonlocal causality*. Discussion of nonlocal effects often involves a paradigm originally proposed by Einstein, Podolsky and Rosen (1935) and updated in popular accounts: A particle decays and emits two photons that travel in

opposite directions. If we change the polarization of one photon, the other photon will spontaneously change to "match" the changed photon more often than expected by chance. For this to happen, the photon not initially changed needs some way of "knowing" the polarization of the changed photon. This "knowing" occurs even though there was no known connection between the photons; indeed, the photon that was not initially changed seemed to respond to the change in the other photon before any signal could travel between the photons.

[5]Although this notion is true of cognitive approaches historically (e.g., Lachman, Lachman, and Butterfield 1979; Neisser 1967), it is less true of more recent approaches (e.g., Gazzaniga, Ivry, and Mangun 2002; Posner and Raichle 1994). The advent of brain imaging technologies such as PET and fMRI has led to a resurgence of interest in the physiological structures underlying cognitive processing. For our purposes, though, this recent interest is more technologically driven, and does not change the fact that historically cognitive science has not concerned itself with issues of the underlying substrate.

[6]If the argument suggested by Chinese Room scenario is correct and a computer may pass the Turing test and yet not have understanding or possess mental states, beliefs and desires similar to a human, many people might nonetheless treat that computer as if it did have understanding and possess mental states, beliefs, and desires (e.g., see responses to the program ELIZA which emulates a therapist, Weizenbaum 1976). This would be consistent with the attributions of intentionality that are typical of shamanism, and so even if the Turing test is inappropriate as a test of understanding, the validity of objections raised by the Chinese Room scenario would not necessarily challenge the apparent similarity of attributions in artificial intelligence and attributions in shamanism.

References Cited

Anderson, J. A.
1995 An Introduction to Neural Networks. Cambridge, MA: MIT Press.

Bateson, G.
1972 Steps to an Ecology of Mind. New York: Ballantine.

Bechtel, W. and A. Abrahamsen
1991 Connectionism and the Mind: An Introduction to Parallel Processing in Networks. Cambridge, MA: Basil Blackwell Ltd.

Bharucha, J. J. and K. Stoeckig
1986 Reaction Time and Musical Expectancy: Priming of Chords. Journal of Experimental Psychology: Human Perception and Performance 12: 403-410.

Boden, M. A.
1987 Artificial Intelligence and Natural Man. 2nd edition. New York: Basic Books, Inc.
1988 Computer Models of Mind. New York: Cambridge University Press.

Brown, A. S., D. R. Neblett, T.C. Jones, and D.B. Mitchell
1991 Transfer of Processing in Repetition Priming: Some Inappropriate Findings. Journal of Experimental Psychology: Learning, Memory, and Cognition 17: 514-525.

Capra, F.
1996 The Web of Life. London: HarperCollins Publishers.

Chomsky, N.
1957 Syntactic Structures. The Hague: Mouton.
1965 Aspects of a Theory of Syntax. Cambridge, MA: MIT Press.

Cole, M.
1996 Cultural Psychology. Cambridge, MA: Belknap Press/Harvard University Press.

Collins, A. M. and E.F. Loftus
1975 A Spreading Activation Theory of Semantic Processing. Psychological Review 82: 407-428.

Collins, A. M. and M.R. Quillian
1969 Retrieval Time from Semantic Memory. Journal of Verbal Learning and Verbal Behavior 8: 240-247.

Crevier, D.
1993 AI: The Tumultuous History of the Search for Artificial Intelligence. New York: Basic Books, Inc.

Dawson, M. R. W.
1998 Understanding Cognitive Science. Malden, MA: Blackwell Publishers.

Dennett, D. C.
1987 The Intentional Stance. Cambridge, MA: MIT Press/Bradford Books.

Devereaux, G.
1961 Shamans as Neurotics. American Anthropologist 63: 1080-1090.

Dreyfus, H. L.
1993 What Computers Still Can't Do. Cambridge, MA: MIT Press.
Eliade, M.
1964 Shamanism: Archaic Techniques of Ecstasy. Princeton, NJ: Princeton University Press/Bollingen Series LXXVI.
Einstein, A., B. Podolsky and N. Rosen
1935 Can Quantum-Mechanical Description of Reality be Considered Complete? Physical Review 47: 777ff.
Ericsson, K. A. and H.A. Simon
1993 Protocol Analysis: Verbal Reports as Data. Revised edition. Cambridge, MA: MIT Press.
Feldman, J. A. and D.H. Ballard
1982 Connectionist Models and Their Properties. Cognitive Science 6: 205-254.
Frazer, J. G.
1963[1922] The Golden Bough: A Study in Magic and Religion (Abridged). New York: Macmillian Publishing Company.
Gazzaniga, M. S., R.B. Ivry, and G.R. Mangun
2002 Cognitive Neuroscience: The Biology of the Mind. 2nd Edition. W. W. Norton & Co.
Gibson, J. J.
1979 The Ecological Approach to Visual Perception. Boston: Houghton Mifflin.
Gigerenzer, G. and P.M. Todd, eds.
1999 Simple Heuristics That Make Us Smart. New York: Oxford University Press.
Hansen, G. P.
2001 The Trickster and the Paranormal. New York: Xlibris.
Hardy, C.
1998 Networks of Meaning. Westport, CT: Praeger Publishers.
Harner, M.
1990 The Way of the Shaman. 3rd edition. San Francisco: HarperCollins Publishers.
Haugeland, J.
1985 Artificial Intelligence: The Very Idea. Cambridge, MA: MIT Press.
Hubbard, T. L.
2002 Cognitive Science and Shamanism I: Webs of Life and Neural Nets. Shamanism 15: 4-10.
Hutchins, E.
1995 Cognition in the Wild. Cambridge, MA: MIT Press.
Jacoby, L. L. and M. Dallas
1981 On the Relationship between Autobiographical Memory and Perceptual Learning. Journal of Experimental Psychology: General 3: 306-340.
Jahoda, G.
1969 The Psychology of Superstition. London: The Penguin Press.
Johnson-Laird, P. N.
1988 The Computer and the Mind: An Introduction to Cognitive Science. Cambridge, MA: Harvard University Press.
Kahneman, D., P. Slovic, and A. Tversky, eds.
1982 Judgment Under Uncertainty: Heuristics and Biases. New York: Cambridge University Press.
Kalweit, H.
1992 Shamans, Healers, and Medicine Men. Boston, MA: Shambhala Press.
Krippner, S.
2000 The Epistemology and Technologies of Shamanic States of Consciousness. Journal of Consciousness Studies 7: 93-118.
2002 Conflicting Perspectives on Shamans and Shamanism: Points and Counterpoints. American Psychologist 57: 962-977
Kroll, J. F. and M.C. Potter
1984 Recognizing Words, Pictures, and Concepts: A Comparison of Lexical, Object, and Reality Decision. Journal of Verbal Learning and Verbal Behavior 23: 39-66.
Kurzweil, R.
1992 The Age of Intelligent Machines. Cambridge, MA: MIT Press.
Lachman, R., J.L. Lachman, and E.C. Butterfield
1979 Cognitive Psychology and Information Processing: An Introduction. Hillsdale, NJ: Erlbaum.

Lave, J.

1993 The Practice of Learning. *In* Understanding Practice: Perspectives on Activity and Context. S. Chaiklin and J. Lave, eds. New York: Cambridge University Press.

Lefcourt, H. M.

1976 Locus of Control: Current Trends in Theory and Research. Hillsdale, NJ: Erlbaum.

Levi-Strauss, C.

1966 The Savage Mind. London: Weidenfeld & Nicolson.

Lovelock, J.

1979 Gaia: A New Look at Life on Earth. New York: Oxford University Press.

1988 The Ages of Gaia. Norton.

Malinowski, B.

1954 Magic, Science and Religion, and Other Essays. Garden City, NY: Doubleday.

McClelland, J. L.

1988 Connectionist Models and Psychological Evidence. Journal of Memory and Language 27: 107-123.

Meyer, D. E. and R.W. Schvaneveldt

1971 Facilitation in Recognizing Pairs of Words: Evidence of a Dependence between Retrieval Operations. Journal of Experimental Psychology 90: 227-234.

Neely, J. H.

1990 Semantic Priming Effects in Visual Word Recognition: A Selective Review of Current Findings and Theories. *In* Basic Processes in Reading: Visual Word Comphrension. D. Besner and G. Humphreys, eds. Hillsdale, NJ: Erlbaum.

Neisser, U.

1967 Cognitive Psychology. New York: Appleton-Century-Crofts.

Newell, A. and H.A. Simon

1972 Human Problem Solving. Englewood Cliffs, NJ: Prentice-Hall.

Noll, R.

1983 Shamanism and Schizophrenia: A State-specific Approach to the "Schizophrenia Metaphor" of Shamanic States. American Ethologist 10: 433-459.

Posner, M. I. and M.E. Raichle

1994 Images of Mind. New York: W.H. Freeman.

Ripinsky-Naxon, M.

1993 The Nature of Shamanism: Substance and Function of a Religious Metaphor. Albany, NY: State University of New York Press.

Rosch, E. and B.B. Lloyd, eds.

1978 Cognition and Categorization. Hillsdale, NJ: Erlbaum.

Ross. L.

1977 The Intuitive Psychologist and his Shortcomings: Distortions in the Attribution Process. *In* Advances in Experimental Social Psychology (Vol. 10).L. Berkowitz, ed. New York: Academic Press.

Ross, M. and G.J.O. Fletcher

1985 Attribution and Social Perception. *In* The Handbook of Social Psychology, Vol. II. 3rd edition. Lindzey and E. Aronson, New York: Random House.

Rozin, P., L. Millman, and C. Nemeroff

1986 Operation of the Laws of Sympathetic Magic in Disgust and Other Domains. Journal of Personality and Social Psychology 50, 703-712.

Rumelhart, D. E. and J. L. McClelland

1986 Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations. Cambridge, MA: MIT Press.

Schacter, D. L.

1996 Searching for Memory: The Brain, the Mind, and the Past. New York: Basic Books.

Schacter, D. L., L. Cooper, and S. Delaney

1990 Implicit Memory for Unfamiliar Objects Depends on Access to Structural Descriptions. Journal of Experimental Psychology: General 119: 5-24.

Schank, R. C. and R. Abelson

1977 Scripts, Plans, Goals, and Understanding. Hillsdale, NJ: Erlbaum.

Schneider, S. and P. Boston, eds.

1991 Scientists on Gaia. Cambridge, MA: MIT Press.

Searle, J. R.

1980 Minds, Brains, and Programs. Behavioral and Brain Science 3: 417-456.

1992 The Rediscovery of the Mind. Cambridge, MA: MIT Press.

Silverman, J.
1967 Shamans and Acute Schizophrenia. American Anthropologist 69: 21-31.
Smith, E. E. and D.L. Medin
1981 Categories and Concepts. Cambridge, MA: Harvard University Press.
Stephan, W. G.
1985 Intergroup Relations. *In* The Handbook of Social Psychology, Vol. II, 3rd edition. Lindzey and E. Aronson, eds. New York: Random House.
Tulving, E.
1983 Elements of Episodic Memory. New York: Oxford University Press.
Tulving, E. and D.L. Schacter
1980 Priming and Human Memory Systems. Science 247: 301-306.
Tylor, E. B.
1871 Primitive Culture: Researches into the Development of Mythology, Philosophy, Religion, Art, and Custom. London: John Murray.
Varela, F. J., E. Thompson, and E. Rosch
1992 The Embodied Mind: Cognitive Science and Human Experience. Cambridge, MA: MIT Press.
Walsh, R.
2001 Shamanic Experiences: A Developmental Analysis. Journal of Humanistic Psychology 41: 31-52.
Warner, R.
1980 Deception and Self-deception in Shamanism and Psychiatry. International Journal of Social Psychiatry 26: 41-52.
Weizenbaum, J.
1976 Computer Power and Human Reason. New York: Freeman.
Winkelman, M.
2000 Shamanism: The Neural Ecology of Consciousness and Healing. Westport, CT: Bergin & Garvey.
Zusne, L. and W.H. Jones
1989 Anomalistic Psychology: A Study of Magical Thinking. 2nd edition. Hillsdale, NJ: Erlbaum.