proposition A statement relating a subject and a claim about that subject.

node In network-based models of mental representation, a "meeting place" for the various connections associated with a particular topic.

associative links In network-based models of mental representation, connections between the symbols (or nodes) in the network.

spreading activation The process through which activity in one node in a network flows outward to other nodes through associative links.

Propositions

As we've seen, mental images—and analogical representations in general—are essential for representing some types of information. Other information, in contrast, requires a symbolic representation. This type of mental representation is more flexible because symbols can represent any content we choose, thanks to the fact that it's entirely up to us what each symbol stands for. Thus, we can use the word *mole* to stand for an animal that digs in the ground, or we could use the word (as Spanish speakers do) to refer to a type of sauce used in cooking. Likewise, we can use the word *cat* to refer to your pet, Snowflake; but, if we wished, we could instead use the Romanian word *pisică* as the symbol representing your pet, or we could use the arbitrary designation X2ф. (Of course, for communicating with others, it's important that we use the same terms they do. This is not an issue, however, when we're representing thoughts in our own minds.)

Crucially, symbols can also be combined with each other to represent more complex contents—such as "San Diego is in California," or "cigarette smoking is bad for your health." There is debate about the exact nature of these combinations, but many scholars propose that symbols can be assembled into **propositions**—statements that relate a subject (the item about which the statement is being made) and a predicate (what's being asserted about the subject). For example, "Solomon loves to blow glass," "Jacob lived in Poland," and "Squirrels eat burritos" are all propositions (although the first two are true, and the last is false). But just the word *Susan* or the phrase "is squeamish" aren't propositions—the first is a subject without a predicate; the second is a predicate without a subject. (For more on how propositions are structured and the role they play in our thoughts, see J. Anderson, 1993, 1996.)

It's easy to express propositions as sentences, but this is just a convenience; many other formats are possible. In the mind, propositions are probably expressed via network structures, related to the network models we discussed for perception in Chapter 5. Individual symbols serve as **nodes** within the network—meeting places for various links—so if we were to draw a picture of the network, the nodes would look like knots in a fisherman's net, and this is the origin of the term *node* (derived from the Latin *nodus*, meaning "knot"). The individual nodes are connected to each other by **associative links** (Figure 9.5). Thus, in this system there might be a node representing *Abe Lincoln* and another node representing *President*, and the link between them represents part of our knowledge about Lincoln—namely, that he was a president. Other links have labels on them, as shown in Figure 9.6; these labels allow us to specify other relationships among nodes, and in this way we can use the network to express any proposition at all (after J. Anderson, 1993, 1996).

9.5 Associative connections Many investigators propose that our knowledge is represented through a network of associated ideas, so that the idea of "Abe Lincoln" is linked to "Civil War" and "President."

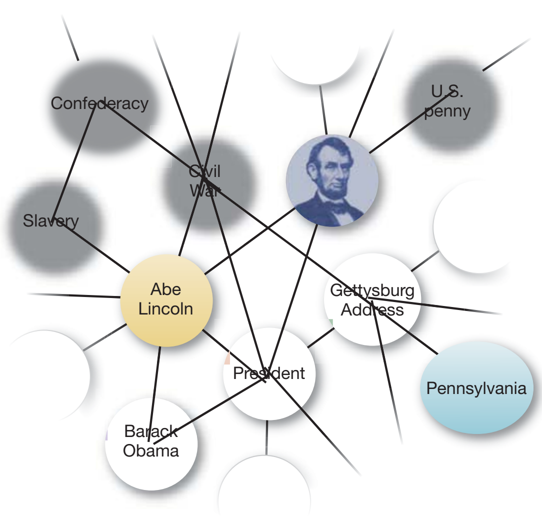

The various nodes representing a proposition are activated whenever a person is thinking about that proposition. This activation then spreads to neighboring nodes, through the associative links, much as electric current spreads through a network of wires. However, this spread of activation will be weaker (and will occur more slowly) between nodes that are only weakly associated. The **spreading activation** will also dissipate as it spreads outward, so that little or no activation will reach the nodes more distant from the activation's source.

In fact, we can follow the spread of activation directly. In a classic study, participants were presented with two strings of letters, like

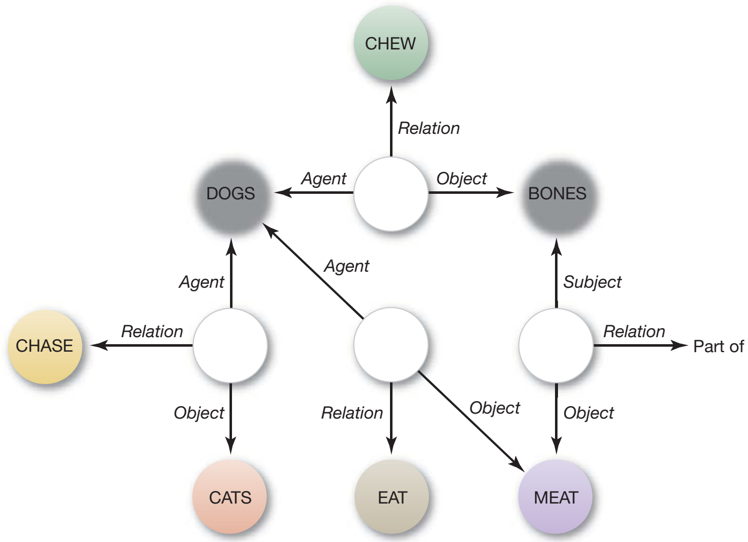

9.6 Propositions One proposal is that your understanding of dogs—what they are, what they're likely to do—is represented by an interconnected network of propositions. In this figure, each proposition is represented by a white circle, which serves as the meeting place for the elements included in the proposition. Thus, this bit of memory network contains the propositions "dogs chew bones," "dogs chase cats," and so on. A complete representation about your knowledge of dogs would include many other propositions as well.

NARDE–DOCTOR, or GARDEN–DOCTOR, or NURSE–DOCTOR (Meyer & Schvaneveldt, 1971). The participants' job was to press a "yes" button if both sequences were real words (as in the second and third examples here), and a "no" button if either was not a word (the first example). Our interest here is only in the two pairs that required a yes response. (In these tasks, the no items serve only as *catch trials,* ensuring that participants really are doing the task as they were instructed.)

Let's consider a trial in which participants see a related pair, like NURSE–DOCTOR. In choosing a response, they first need to confirm that, yes, NURSE is a real word in English. To do this, they presumably need to locate the word NURSE in their mental dictionary; once they find it, they can be sure that these letters do form a legitimate word. What this means, though, is that they will have searched for, and activated, the node in memory that represents this word—and this, we have hypothesized, will trigger a spread of activation outward from the node, bringing activation to other, nearby nodes. These nearby nodes will surely include the node for DOCTOR, since there's a strong association between "nurse" and "doctor." Therefore, once the node for NURSE is activated, some activation should also spread to the node for DOCTOR.

Once they've dealt with NURSE, the participants can turn their attention to the second word in the pair. To make a decision about DOCTOR (is this string a word or not?), the participants must locate the node for this word in memory. If they find the relevant node, then they know that this string, too, is a word and can hit the "yes" button. But of course the process of activating the node for DOCTOR has already begun, thanks to the activation this node just received from the node for NURSE. This should accelerate the process of bringing the DOCTOR node to threshold (since it's already partway there), and so it will take less time to activate. Hence, we expect quicker responses to DOCTOR in this context, compared to a context in which it was preceded by some unrelated word and therefore not primed. This prediction is correct. Participants' lexical decision responses are faster by almost 100 milliseconds if the stimulus words are related, so that the first word can prime the second in the way we just described.

We've described this sequence of events within a relatively uninteresting task—participants merely deciding whether letter strings are words in English or not. But the same dynamic—with one node priming other, nearby nodes—plays a role in, and can shape, the flow of our thoughts. For example, we mentioned in Chapter 6 that the sequence of ideas in a dream is shaped by which nodes are primed. Likewise, in problem solving, we sometimes have to hunt through memory, looking for ideas about how to tackle the problem we're confronting. In this process, we're plainly guided by the pattern of which nodes are activated (and so more available) and which nodes aren't. This pattern of activation in turn depends on how the nodes are connected to each other—and so the arrangement of our knowledge within long-term memory can have a powerful impact on whether we'll locate a problem's solution.

JUDGMENT: DRAWING CONCLUSIONS FROM EXPERIENCE

So far we've been discussing the content of thought, with an emphasis on how thoughts are represented in the mind. Just as important, though, are the processes of thought—what psychologists call **directed thinking**—the ways people draw conclusions or make decisions. What's more, these two broad topics—the contents of thought and the processes—are linked in important ways. As we've discussed, representing ideas with images will highlight visual appearance in our thoughts and thus may call to mind objects with similar appearance. Likewise, representing ideas as propositions will cause activation to spread to other, associated, nodes; and this too can guide our thoughts in one direction rather than another.

directed thinking Thinking aimed at a particular goal.

judgment The process of extrapolating from evidence to draw conclusions.

heuristics A strategy for making judgments quickly, at the price of occasional mistakes.

But, of course, the flow of our thoughts also depends on what we're trying to accomplish in our thinking. So it will be useful to divide our discussion of thought processes into four sections, each corresponding to a type of goal in our thinking: We will, therefore, consider *judgment, reasoning, decision making,* and *problem solving.* Let's begin with judgment.

The term **judgment** refers to the various steps we use when trying to reach beyond the evidence we've encountered so far, and to draw conclusions from that evidence. Judgment, by its nature, involves some degree of extrapolation because we're going beyond the evidence; and as such, this always involves some risk that the extrapolation will be mistaken. If, for example, we know that Jane has enjoyed many trips to the beach, we might draw the conclusion that she will always enjoy such trips. But there's no guarantee here, and it's surely possible that her view of the beach might change. Likewise, if you have, in the past, preferred spending time with quiet people, you might draw a conclusion about how much you'd enjoy an evening with Sid, who's quite loud. But here, too, there's no guarantee—and perhaps you'll have a great time with Sid.

9.7 Daniel Kahneman and Amos Tversky Much of what we know about judgment and decision making comes from pioneering work by Daniel Kahneman (A) and Amos Tversky (B); their work led to Kahneman receiving the Nobel Prize in 2002.

Even with these risks, we routinely rely on judgment to reach beyond the evidence we've gathered so far—and so we do make forecasts about the next beach trip, whether the evening with Sid would be fun, and more. But how do we proceed in making these judgments? Research suggests that we often rely on a small set of shortcuts called *judgment heuristics.* The word **heuristics,** borrowed from computer science, refers to a strategy that's relatively efficient but occasionally leads to error. Heuristics, in other words, offer a trade-off between efficiency and accuracy, helping us to make judgments more quickly—but at the price of occasional mistakes.

Let's start our discussion with two of these shortcuts—the *availability* and *representativeness heuristics,* first described by Amos Tversky and Daniel Kahneman (Figure 9.7); their research in this domain is