Problems With Dividing the Realm of Processes

Agnes Moors and Jan De Houwer
Ghent University, Ghent, Belgium

Multimode models with two, three, or four modes of processing have been proposed in domains as diverse as reasoning (e.g., Sloman, 1996), categorization (e.g., Rouder & Ratcliff, 2004), learning (e.g., Shanks & St. John, 1994), social judgment (e.g., Strack & Deutsch, 2004), and emotion (e.g., Leventhal & Scherer, 1987). Multimode models have proposed several criteria (e.g., operating conditions, formal properties of the process, content, and format of the representations on which the process operates) to divide the realm of processes. As Sherman (this issue) notes, one can make as many categories as one deems useful in a certain context. Like all forms of categorization in daily life and science, making categories of cognitive processes is context dependent. However, there are at least two potential problems with dividing processes into several categories. First, most multimode models make a priori assumptions of overlap among the categories obtained by two or more criteria for division. We argue that there are many ways to cut the cake but that the different slicing methods do not necessarily result in the same slices (see also Sherman, this issue). A second problem has to do with some of the proposed criteria for division. Some criteria are not discrete but dimensional, and they do not allow for the creation of clear-cut, all-or-none categories. Other criteria are discrete but have poor explanatory value (at least according to some authors). We discuss these problems in more detail next.

Mapping Categories Obtained With Different Criteria

The realm of processes can be split up according to several criteria. To gain a better understanding of these criteria, we find it useful to start from a levels-of-analysis approach. Following Marr (1982), we distinguish three levels of process understanding. At the first level, a process is described as a *functional relation* between an input and an output. This level includes the *content* of *input* and *output*, and the *conditions* under which the process operates. The second level articulates the *formal properties* of the process (the primitive *mechanisms) involved in transforming input into output and the format of the representations* in which *input* and *output* are coded. This level addresses what is in the black box.[1] The third level is concerned with the *physical realization* of processes in the brain. The three levels are related, but only loosely. For example, the choice of a formal process is influenced by the functional process it must account for, but one functional process can be implemented by different formal processes. Other theorists (e.g., Anderson, 1987; Clark, 1990; Pylyshyn, 1980) have proposed a different number of levels and have placed the boundaries between the levels at somewhat different heights, but they share the idea that processes can be considered at qualitatively different levels of analysis and that lower levels are implementations of higher levels.

The criteria for categorization used by multimode models can be situated within this framework.

1. The characterization of a process as automatic or nonautomatic tells something about the *conditions* under which the process operates (e.g., Bargh, 1992). A process is automatic when it operates under suboptimal conditions, such as when there is minimal time, minimal attentional capacity, a subliminal stimulus input, no intention to engage in the process (or the intention is not achieved), and/or when there are attempts to stop or avoid the process. A process is nonautomatic when it operates under optimal conditions.

2. Some models distinguish processes on the basis of the *functional process* (or the *content of input* and *output*) involved in the processes. An example are the models that distinguish between the processing of heuristic information (e.g., source attractiveness, the majority's opinion, message length) and the processing of systematic information (i.e., message's persuasive arguments; e.g., Chaiken, Liberman, & Eagly, 1989).

3. The distinction between rule-based and associative processes (Sloman, 1996) refers to the *formal properties* of the process or the primitive *mechanisms*.

4. Some models use as a criterion the *format of the representations* that serve as the *input* of a process, also termed mental *codes*. For example, multimode models in the domain of emotion distinguish between processes operating on sensory codes, those operating on perceptual/analog codes, and those operating on conceptual/semantic codes (e.g., Leventhal & Scherer, 1987; Power & Dalgleish, 1997).

5. Neurophysiological models distinguish processes according to their underlying neurophysiologi-

[1]One might argue that the formal level is by definition impenetrable and that it therefore only makes sense to consider the functional level. The functional level can itself be subdivided into different sublevels. These sublevels are not qualitatively different but can be placed on a continuum ranging from the most concrete to more abstract descriptions.

cal structures or circuits (e.g., neocortical vs. subcortical pathways to the amygdala; LeDoux, 1986).

The majority of multimode models map the categories obtained with two or more criteria. We discuss and evaluate five examples of such a confounding of criteria. First, some models impose a priori constraints on the *conditions* under which certain *functional processes* can take place. For example, in Sherman's (this issue) Quad Model, stimulus detection and guessing are automatic, whereas response selection and inhibition are nonautomatic.[2] For another example, the processing of heuristic information is usually thought to be automatic, whereas the processing of systematic information is thought to be nonautomatic (e.g., Chaiken, 1980). It is not difficult to find exceptions to this or alternative explanations (Kruglanski, Erb, Pierro, Mannetti, & Chun, this issue; Pierro, Mannetti, Erb, Spiegel, & Kruglanski, 2005; Sherman, this issue). Pierro et al. (2005) showed that in prior studies, observed dissociations in automaticity between processes involving different types of information (heuristic vs. systematic) were partly due to a confounding of quantitative parameters (e.g., complexity, length, presentation order) that are not meaningfully related to information type. In these studies, heuristic information was typically less complex, was shorter, and presented earlier than systematic information, thus enabling heuristic but not systematic information to exert an automatic influence on judgment.

Second, several models impose a priori constraints on the *conditions* under which certain *formal processes* can take place. The dominant view is that associative processes operate under suboptimal conditions, whereas rule-based processes can operate only under optimal conditions (cf. Logan, 1988). Opponents of this dominant view have suggested the possibility of automatic rule-based processing (e.g., in skill-development, Anderson, 1992; Tzelgov, Yehene, Kotler, & Alon, 2000; implicit grammar learning, Reber, 1989). Instead of denying automaticity to rule-based processes on an a priori basis, opponents of the dominant view argue that it should be empirically assessed which type of process can or cannot operate under suboptimal conditions. Admittedly, research aimed at establishing automatic rule-based processing is confronted with many hurdles, such as how to assess automaticity and how to assess the involvement of rule-based processes separate from, or in addition to, associative processes (see further).

Third, some models impose a priori constraints on the *format* of the *representations* or *codes* that can be acted on by each *formal process*. For example, rule-based processes are often said to operate on symbolic codes, whereas associative processes operate on perceptual or analog codes (Deutsch & Strack, this issue; Leventhal & Scherer, 1987) or on subsymbolic codes (in connectionist or hybrid models). However, other theorists have suggested that associative processes can also deal with abstract concepts (Bartlett, 1932; Hahn & Chater, 1998; James, 1890; Sloman, 1996; C. A. Smith & Kirby, 2001). Conversely, we see no principled reason to assume that rule-based processes cannot operate on perceptual codes.

Fourth, some models map different *neurophysiological routes* onto different operating *conditions*. In multimode models of emotion elicitation, for example, the subcortical pathway to the amygdala is linked to automatic and the cortical pathway to nonautomatic emotion elicitation (LeDoux, 1986). It is often recognized, however, that automaticity is not unique to the subcortical circuits of the brain and that we are only beginning to understand the subtleties of the interactions among cortical and subcortical brain structures (cf. Phelps, 2004).

Finally, so-called dual-system models (Deutsch & Strack, this issue) postulate links between the categories formed by almost all of the criteria discussed: *functional processes*, *conditions*, *formal processes*, *representations*, and *neurophysiological structures*. Each of the previous comments applies to these models.

To summarize, multimode models tend to forge links among the categories obtained by different criteria. Often these links have not been explicitly investigated (in a manner that permits falsification), and counterexamples or alternative explanations are available (see Kruglanski et al., this issue). Instead of taking as a default assumption that there is perfect overlap among the categories obtained with different criteria, we propose to take independence of categories as the starting point and to progressively investigate possible degrees of overlap. In our view, there are no compelling reasons to assume principled overlap among the categories discussed. It should be a matter of empirical research to determine whether there is some degree of actual overlap.

Value of Individual Criteria

Now that we have discussed the problem of assumptions of overlap among various ways of categorizing processes, we turn to the second problem, whether the criteria proposed are suited to divide the realm of processes in a clear-cut manner.

Automatic Versus Nonautomatic

In the introduction, we suggested that some criteria do not allow for the creation of all-or-none categories.

[2]Although Sherman (this issue) recognizes this problem, we wonder why he goes to the trouble of explaining in great detail a model whose central assumptions he eventually comes to reject.

One such criterion is automaticity. According to a feature-based approach, automaticity is an umbrella term for a number of individual features such as fast, efficient, unintentional, uncontrolled, and unconscious. The definitions of these features can be reformulated in terms of operating conditions. For example, an efficient process is one that makes minimal use of attentional capacity, which means that it can operate when a minimal amount of attentional capacity is available. A fast process is one that can operate when there is not much time. An unintentional process is one that is not caused by an intention (i.e., the goal to engage in the process), which means that it operates when there is no causally efficacious intention. More generally, an uncontrolled process is one that is not influenced ([a] in the sense of caused, or [b] in the sense of stopped/avoided) by the goal to do so ([a] engage in the process, or [b] stop/avoid the process). This means that it operates (a) in the absence of a (causally efficacious) goal to engage in it or (b) despite the presence of a goal to stop/avoid it.[3] To say that a process is unconscious means that the process occurs when the person is not conscious of it (or of its input or output; cf. Moors & De Houwer, in 2006).

These features, it has been argued, do not hang together in an all-or-none fashion (Bargh, 1992). For example, it seems that certain processes are fast and efficient but not uncontrollable (in the sense of stop/avoid; cf. Uleman & Moskowitz, 1994). We have even raised the possibility of interdependence between some automatic and some nonautomatic features (Moors & De Houwer, in press). For example, in threshold determination studies for subliminal perception, it seems that short presentation times of the stimulus can be compensated by increased focusing of attention to the stimulus or by increased salience of the stimulus. In cases like this, there seems to be a trade-off rather than a cooccurrence among the features fast and efficient. Because of the lack of coherence among automaticity features, we favor a decompositional approach to the study of automaticity. Such an approach proposes to investigate the presence of indivual features separately (Bargh, 1992; Moors & De Houwer, 2006).

It has further been argued that each automaticity feature can be regarded as a continuum (Logan, 1985; Moors & De Houwer, 2006). For example, a process can be more or less fast, more or less efficient, more or less controlled (i.e., [a] more or less conform to one's intentions, or [b] more or less successfully stopped/avoided), and more or less unconscious. It is often not possible to conclude for the complete presence or absence of a feature. In sum, the lack of coocurrence among automaticity features as well as their gradual nature complicate the task to create separate bins of automatic and nonautomatic processes. Processes are automatic with regard to some features and to some degree, but not with regard to others. Moreover, processes may be automatic (with regard to some features and to some degree) on some occasions, but not on others, depending on conditions that are unrelated to automaticity (such as salience, complexity, length, and presentation order of the stimulus input; see Kruglanksi et al., this issue). We can thus conclude that the criterion automatic–nonautomatic is not suited for a clear-cut division of processes.

[3]Note that according to our definitions, unintentional processes are a subclass of uncontrolled processes. Unintentional processes are uncontrolled in sense (a).

Rule-Based Versus Associative

Another distinction that has been under fire is between rule-based processes and associative ones. In a rule-based process, a mental rule is applied to an input (or a representation thereof), and computation of the rule produces an output. In an associative process, an input activates stored representations of similar past inputs. This activation, in turn, spreads to associated stored representations, which determine the output. Kruglanski et al. (this issue) argue that both mechanisms have the same formal properties in that they can both be expressed in an if–then format. Given the authors' claim that an if–then format is the hallmark of rules, they argue that associations are, in fact, rules and that the activation of stored associations is a rule-based process. There are, however, reasons to challenge this definition. Rule-based processes and associative ones may have things in common (they are both processes after all), but there may still be formal distinctions left to make between them. We discuss three of these distinctions next.

First, some authors have argued that rule-based processes are governed by *abstract* rules (e.g., Sloman, 1996). Abstract rules not only fit the if–then format, they also require that the premise contains variables. Variables are abstract representations that can be instantiated in more than one way (i.e., with more than one constant). Consider the abstract rule that could underlie the elicitation of positive emotions such as happiness: "if X = Y then q" in which X stands for an actual situation, Y stands for a desired situation, and q = happiness. The rule applies to an infinite set of actual and desired situations (e.g., if you desire chocolate cake [y1] and you are offered chocolate cake [x1], happiness occurs; if you desire success at work [y2] and success is what you achieve [x2], happiness occurs). Associations, on the other hand, can be said to fit the format of nonabstract rules in which the premise merely consists of constants. Constants are representations of concrete or even unique instances. For example, "if p then q," with p = chocolate cake and q = happiness. For a more complex example, "if (p and r) then q," with p = chocolate cake, r = desire for chocolate

cake, and q = happiness. Because p and r are constants, the rule cannot be applied to new input (e.g., s = new car or t = success at work) unless there is some resemblance with p (e.g., u = strawberry cake). Nonabstract rules can thus be applied to new input but only by virtue of similarity among the input and the constants specified in the premise.

This brings us to the second distinction between rule-based processes and associative ones. In the case of a rule-based process, the premise of the rule must be strictly matched, whereas in the case of an associative process, the premise may be partially matched (Hahn & Chater, 1998[4]). Rule-based processes and associative ones can both account for generalization, due to the complementary forces of abstraction and partial matching. In the case of rule-based processes, generalization is obtained by virtue of abstract variables. In the case of associative processes, generalization is obtained by virtue of partial matching (partial matching compensates for the lack of variables). Also note that abstraction is a relative notion (Hahn & Chater, 1998). Abstraction has to do with a loss of information: abstract representations contain less unique features than concrete ones. The variables figuring in abstract rules and the constants figuring in nonabstract rules thus occupy two points on a continuum. Variables can be substituted by a larger class of instances than constants can, but the variables that figure in abstract rules can often not be substituted by just any constant (e.g., in the previous example of an abstract rule, X must be an actual state and Y must be a desired state). At the very extreme are logical rules in which the variables can be sustituted by any constant (e.g., if [X and Y] then X). Conversely, concrete representations often contain some level of abstraction (e.g., in the previous example of a nonabstract rule, the representation of chocolate cake can itself be instantiated by more than one unique chocolate cake). At the very extreme are constants that represent a unique instance.

A third, often mentioned distinction between rule-based and associative processes is that rule-based processes must *follow* rules, whereas associations—at most—*conform* to rules. Rule *following* requires that a mental rule sits between the input and the output of a process and *causally affects* the output; rule *conforming* merely requires that the relation between input and output can be described or summarized according to a rule (Hahn & Chater, 1998; Pylyshyn, 1980; Searle, 1980; Sloman, 1996; E. E. Smith, Langston, & Nisbett, 1992). If this constraint of rule-based processing would be relaxed, anything that can be described by rules, such as the swimming pattern of a school of fish or planetary motions, would have to be categorized as rule-based. Although associations (or patterns of associations) also mediate between input and output, some do not consider them to be mental rules in that they are not symbolic representations of a rule; they do not have a rule as their content (e.g., Hahn & Chater, 1998). According to others (Clark, 1990; Fodor & Pylyshyn, 1988) the rules in rule-based processes must not be symbolically represented. They may also be wired in the system from birth or through learning. Associations can be seen as rules on the latter view.

Although these three criteria provide formal distinctions between rule-based and associative processes, one can argue that these distinctions remain meaningless at the functional level because they do not seem to lead to different testable predictions. The two mechanisms seem able to account for much the same functional observations. First, as previously explained, both mechanisms are able to account for generalization toward new stimuli (abstract rules by virtue of variables, associations by virtue of partial matching). Second, given the relative nature of abstraction, no objective line can be drawn between variables and constants. This is reflected in the idea that activation of stored stimuli can be based on concrete as well as abstract similarities (e.g., similar function). Similarity may even pertain to abstract relations among variables (cf. Redington & Chater, 1996). Thus, evidence for generalization toward stimuli that share abstract (but not concrete) features with previously acquired ones (Marcus, Vijayan, Bandi Rao, & Vishton, 1999; Reber, 1989) is equally compatible with rule-based as with activation-based accounts (Redington & Chater, 1996; but see Sloman & Rips, 1998; E. E. Smith et al., 1992). Third, even if an arbitrary line would be drawn between abstract and concrete features in a way that everybody would agree with, there still is the problem that every abstract rule can be translated in a set of nonabstract rules (one for each combination of values that can be entered in the variable slots of the abstract rule) and vice versa. Both can thus account for the same input–output relations. Fourth, to empirically assess whether the output of a process is causally determined by a symbolically represented rule is not an easy task. Some authors have proposed to use verbal protocols to investigate the content of representations. If people's performance is mediated by mental rules, it is possible that they can verbally report these rules. There are two problems with this proposal. For one thing, the inability to report a rule cannot be taken as proof for its absence, because a rule may affect performance without being consciously accessible. The criterion of ver-

[4]Actually, Hahn and Chater (1998) distinguished between rule-based processes and similarity-based ones, linking strict matching to rule-based processes and partial matching to similarity-based ones. They distinguished both types of processes from purely associative ones such as those described in connectionist models, in which no actual computation of a match takes place but which are nevertheless sensitive to similarity. We use the term *associative* in the broad sense, including all processes that are in some sense dependent on similarity between features of the input and features of a stored (symbolic or subsymbolic) representation.

bal reportability is thus useless for research concerned with automatic (in the sense of unconscious) rule-based processing. In addition, the ability to report a rule cannot always be taken as proof for its causal role, because people may report rules that they did not actually use (Nisbett & Wilson, 1977). Recall, moreover, that some authors do not think that the rules governing rule-based processes should be symbolically represented but may also be built-in (Clark, 1990). Built-in rules are not likely to be available for introspection. The literature contains several other proposals for how to empirically assess whether performance is rule based or associative (cf. reviews by Hahn & Chater, 1998; Sloman, 1996; E. E. Smith et al., 1992). However, there does not seem to be general consensus about what constitutes evidence for each process. A further problem noted by many authors (Jacoby, Toth, & Yonelinas, 1993; Reingold & Merikle, 1993; Sherman, this issue) is that no task is process pure. Task performance may be determined simultaneously by rule-based and associative processes. Evidence for one type of process thereby does not exclude presence of the other.

We have thus seen that, in contrast to Kruglanksi et al. (this issue), one can make formal distinctions between rule-based processes and associative ones. These distinctions do not seem to lead to decisive distinctions at the functional level. Despite the controversial status of many empirical results, we prefer to leave open the issue of whether rule-based processes and associative ones can be distinguished empirically. It is up to proponents of dual-mode models that rely on this distinction to investigate the issue further. If no functional difference can be found between the two processes, or if no agreement can be obtained about what this difference should be, rule-based, associative, and hybrid models remain empirically indistinguishable theories (Marcus, Berent, Seidenberg, MacDonald, & Saffran, 2003).

Is There a Future for Multimode Models?

We have seen that several of the criteria for the categorization of processes are problematic and that there is a lack of overlap among the categories obtained with different criteria. Because of these problems, certain proponents of multimode models have toned down some of their initial claims. Some models now grant exceptions, or they abandon some of their initial criteria. For example, some dual-mode models that were originally based on the criterion of type of information processed (e.g., heuristic vs. systematic) have now abandoned that criterion (cf. Sherman, this issue). Others grant that the automaticity criterion is not able to create separate bins of processes (Deutsch & Strack, this issue). We agree with Sherman (this issue) that when the original criteria for categorization are given up, the models risk losing their ground for being a dual-mode model. Multimode models should be able to keep at least one criterion to which they attach their categories.

It seems that the safest criterion to choose is the type of functions that can be performed by each system (e.g., evaluation, counting, detection, guessing, metacognition). So rather than dividing processes at the formal level (mechanisms), one could divide processes at the functional level. Sherman (this issue) seems to have reached the same conclusion near the end of his target article. This criterion also seems to be the key criterion in the Deutsch and Strack (this issue) model: The reflective system can perform functions that the impulsive system cannot, such as metacognition and generating new action plans. Even if metacognition would turn out to rely on associative mechanisms (Theofilou & Cleeremans, 2005) or be able to operate under suboptimal conditions (Reder & Schunn, 1996), the functional distinction between cognition and metacognition remains valid and useful.

Conclusion

Categorization has been marked as a normal aspect of information processing. It serves to reduce information in a way that enables people to focus on the relevant aspects in some context and to facilitate communication. There is nothing fallacious to dividing a group of people according to age, gender, skin color, or shoe size. The fallacy is in attributing to the resulting categories features that have not been verified but that are consistent with our implicit theories. Similarly, talking about processes using contrasting categories permits focusing on key features in some context and may facilitate communication. We should be cautious, however, not to let our implicit theories dictate overlap among categories obtained with different criteria. The tenacious link between the categories of associative and automatic perhaps stems from an implicit metaphor of associative processing as a kind of electrical signal that effortlessly and uncontrollably flows through a copper wire. Conversely, the link between the categories of rule based and nonautomatic perhaps originates from an implicit metaphor of rule application as an active manipulation performed by some homunculus. Such metaphors have considerable intuitive appeal, but so do our implicit theories and prejudices about people in daily life. This should not be an excuse for not making "the hard choice" (Fiske, 1989), which consists in postponing conclusions until sufficient converging evidence supports them. Until that happens, it is best to make as few presuppositions as possible to leave open the debate and the opportunity for careful empirical research.

Notes

Agnes Moors is a postdoctoral researcher for the Fund of Scientific Research, Flanders.

Correspondence should be sent to Agnes Moors, Department of Psychology, Ghent University, Henri Dunantlaan 2, B-9000 Ghent, Belgium. E-mail: agnes.moors@ugent.be

References

Anderson, J. R. (1987). Methodologies for studying human knowledge. *Behavioral and Brain Sciences, 10*, 467–505.

Anderson, J. R. (1992). Automaticity and the ACT* theory. *American Journal of Psychology, 105*, 165–180.

Bargh, J. A. (1992). The ecology of automaticity. Toward establishing the conditions needed to produce automatic processing effects. *American Journal of Psychology, 105*, 181–199.

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

Chaiken, S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology, 39*, 752–766.

Chaiken, S., Liberman, A., & Eagly, A. H. (1989). Heuristic and systematic information processing within and beyond the persuasion context. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 212–252). New York: Guilford.

Clark, A. (1990). Connectionism, competence, and explanation. *British Journal of Philosophical Science, 47*, 195-222.

Fiske, S. T. (1989). Examining the role of intent: Toward understanding its role in stereotyping and prejudice. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 253–283). New York: Guilford.

Fodor, J. A., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28*, 3–71.

Hahn, U., & Chater, N. (1998). Similarity and rules: Distinct? Exhaustive? Empirically distinguishable? *Cognition, 65*, 197–230.

Jacoby, L. L., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology: General, 122*, 139–154.

James, W. (1890). *The principles of psychology*. New York: Holt, Rinehart & Winston.

LeDoux, J. E. (1986). The neurobiology of emotion. In J. E. LeDoux & W. Hirst (Eds.), *Mind and brain: Dialogues in cognitive neuroscience* (pp. 301–354). Cambridge, England: Cambridge University Press.

Leventhal, H., & Scherer, K. (1987). The relationship of emotion to cognition: A functional approach to a semantic controversy. *Cognition and Emotion, 1*, 3–28.

Logan, G. D. (1985). Skill and automaticity: Relations, implications, and future directions. *Canadian Journal of Psychology, 39*, 367–386.

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

Marcus, G. F., Berent, I., Seidenberg, M. S., MacDonald, M. C., & Saffran, J. R. (2003, April 4). Are there limits to statistical learning? *Science, 300*, 53–55.

Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999, January 1). Rule learning by seven-month-old infants. *Science, 283*, 77–80.

Marr, D. (1982). Vision: The philosophy and the approach. In D. Marr (Ed.), *Vision: A computational investigation into the human representation and processing of visual information* (pp. 103–126). New York: Freeman.

Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin, 32*, 297–326.

Moors, A., & De Houwer, J. (in press). What is automaticity: An analysis of its component features and their interrelations. In J. A. Bargh (Ed.), *Social psychology and the unconscious: The automaticity of higher mental processes*. New York: Psychology Press.

Nisbett, R., & Wilson, T. (1977) Telling more than we know: Verbal reports on mental processes. *Psychological Review, 8*, 231–259.

Phelps, E. A. (2004). Human emotion and memory: Interactions of the amygdala and hippocampal complex. *Current opinion in Neurobiology, 14*, 198-202.

Pierro, A., Mannetti, L., Erb, H.-P., Spiegel, S., & Kruglanski, A.W. (2005). Informational length and order of presentation as determinants of persuasion. *Journal of Experimental Social Psychology, 41*, 458–469.

Power, M., & Dalgleish, T. (1997). *Cognition and emotion: from order to disorder*. Hove, England: Psychology Press.

Pylyshyn, Z. W. (1980). Computation and cognition: Issues in the foundations of cognitive science. *Behavioral and Brain Sciences, 3*, 111–132.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General, 118*, 219–235.

Reder, L. M., & Schunn, C. D. (1996). Metacognition does not imply awareness: Strategy choice is governed by implicit learning and memory. In L. M. Reder (Ed.), *Implicit memory and metacognition* (pp. 45–78). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Redington, M., & Chater, N. (1996). Transfer in artificial grammar learning: A re-evaluation. *Journal of Experimental Psychology: General, 125*, 123–138.

Reingold, E. M., & Merikle, P. M. (1993). Theory and measurement in the study of unconscious processes. In M. Davies & G. W. Humphreys (Eds.), *Consciousness: Psychological and philosophical essays* (pp. 40–57). Oxford, England: Blackwell.

Rouder, J. N., & Ratcliff, R. (2004). Comparing categorization models. *Journal of Experimental Psychology: General, 133*, 63–82.

Searle, J. (1980). Rules and causation. *Behavioral and Brain Sciences, 3*, 37.

Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17, 367–447.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*, 3–22.

Sloman, S. A., & Rips, L. J. (1998). Similarity as an explanatory construct. *Cognition, 65*, 87–101.

Smith, C. A., & Kirby, L. D. (2001). Toward delivering on the promise of appraisal theory. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.), *Appraisal processes in emotion* (pp. 121–138). New York: Oxford University Press.

Smith, E. E., Langston, C., & Nisbett, R. (1992). The case for rules in reasoning. *Cognitive Science, 16*, 1–40.

Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of human behavior. *Personality and Social Psychology Review, 8*, 220–247.

Theofilou, D., & Cleeremans, A. (2005). Metarepresentation in a dual connectionist model. (abstract). In B. Hommel, G. Band, W. La Heij, & G. Wolters (Eds.), *Proceedings of the 14th meeting of the European Society for Cognitive Psychology*, Leiden, The Netherlands.

Tzelgov, J., Yehene, V., Kotler, L., & Alon, A. (2000). Automatic comparisons of artificial digits never compared: Learning linear ordering relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 103–120.

Uleman, J. S., & Moskowitz, G. B. (1994). Unintended effects of goals on unintended inferences. *Journal of Personality and Social Psychology, 66*, 490–501.