

Automaticity: A Theoretical and Conceptual Analysis

Agnes Moors and Jan De Houwer
Ghent University

Several theoretical views of automaticity are discussed. Most of these suggest that automaticity should be diagnosed by looking at the presence of features such as unintentional, uncontrolled/uncontrollable, goal independent, autonomous, purely stimulus driven, unconscious, efficient, and fast. Contemporary views further suggest that these features should be investigated separately. The authors examine whether features of automaticity can be disentangled on a conceptual level, because only then is the separate investigation of them worth the effort. They conclude that the conceptual analysis of features is to a large extent feasible. Not all researchers agree with this position, however. The authors show that assumptions of overlap among features are determined by the other researchers' views of automaticity and by the models they endorse for information processing in general.

Keywords: automatic, unintentional, uncontrolled, autonomous, unconscious

Automaticity is a concept with a long-standing history in psychology (e.g., James, 1890; Wundt, 1903). It has been invoked in domains as diverse as perception (MacLeod, 1991), memory (Jacoby, 1991), social cognition (Wegner & Bargh, 1998), learning (Cleeremans & Jiménez, 2002), motivation (Carver & Scheier, 2002), and emotion (Scherer, 1993). Despite its central nature, there is no consensus about what automaticity means. The aim of this article is to provide an in-depth analysis of the concept and, in particular, the features that have been subsumed under the term. We also discuss implications of this analysis for future research.

In the first section, we consider several contrasting views of automaticity. This overview reveals that most scholars analyze automaticity in terms of one or more features. The implication is that automaticity can be diagnosed by looking for the presence of these features in performance or processes.¹ Contemporary approaches to automaticity, moreover, have begun to argue for the separate investigation of automaticity features. In the second section, we discuss the most essential features in greater detail. One purpose is to examine whether they can be separated on a conceptual level. We analyze the features to a point at which the overlap among them is minimized. Another purpose is to show that assumptions of overlap are pervasive in the research literature despite the recent claim that individual features should be investigated separately. These assumptions arise from authors' specific views of automaticity, as well as from the specific underlying information-processing model that they endorse.

Agnes Moors and Jan De Houwer, Department of Psychology, Ghent University, Ghent, Belgium.

The preparation of this article was made possible by Special Research Fund, Ghent University, Grant BOF-01101203 to Jan De Houwer. We thank Tom Beckers, Frank Baeyens, Dirk Hermans, Geert Crombez, and William Ickes for helpful comments on earlier versions of this article.

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

Approaches to Automaticity

James (1890), Jastrow (1906), and Wundt (1896/1897, 1903) offered some of the earliest descriptions of automaticity, and many of their ideas have reemerged in contemporary accounts. The views presented below are also based on early studies of skill development (Bryan & Harter, 1899) and early dual-task studies (Solomons & Stein, 1896; see review by Shiffrin, 1988).

Automaticity as Processing With No or Minimal Attention

Cognitive psychology is founded on the metatheoretical assumption that responses to stimuli are mediated by information processing. The processing capacity of the total system is intrinsically limited, generating the need for a mechanism of selection (Allport, 1989, 1993). Attention models describe attention as this mechanism. It selects task-relevant aspects of incoming information for further processing and storage and, in this way, avoids overloading the information system.

From a Single-Channel to a Single-Capacity View of Attention

According to early single-channel or bottleneck models, attention was a structural property of the information-processing system, situated at a fixed stage in the processing sequence. Some theorists placed the attention mechanism early (prior to perceptual analysis, e.g., Broadbent, 1958), whereas others placed it late (after semantic processing, e.g., Deutsch & Deutsch, 1963). Processes that preceded the attention mechanism (i.e., preattentive processes) were presumed to be unaffected by attentional limitations. The status of a process as dependent or independent of limitations was entirely determined by its location in the information-processing chain relative to the presumed location of the attention mechanism.

¹ Task performance is observable; the processes underlying the performance are not observable and must be inferred. We take it as a general rule that when the performance is automatic, so is the underlying process.

Single-channel models eventually became overshadowed by (single-) capacity models (e.g., Kahneman, 1973). These models no longer conceived of attention as a processing structure but instead as a source of energy that can be flexibly allocated to any stage in the information-processing chain.² Different stages in processing vary in the amount of attentional capacity they require: Early stages (sensory analysis) require no attention (preattentive processing), whereas later stages require increasing amounts. In addition to this prewired distinction in capacity requirements, these models suggested that extensive (consistent) practice resulted in diminished attentional demands. Processes were considered automatic to the extent that they operated independent of attentional resources, and the development of automaticity was described as a gradual reduction in the need for attention with practice (Hasher & Zacks, 1979; LaBerge & Samuels, 1974; Logan, 1979; Posner & Snyder, 1975a, 1975b; Shiffrin & Schneider, 1977). We call this the *capacity view* of automaticity.

Some authors have suggested that learned automatic processes may eventually behave as preattentive automatic ones (Gibson, 1969; LaBerge, 1973). Others have rejected equating learned processes with preattentive ones, because they differ with respect to the origin of their capacity independence (innate or acquired) and because they possess different functional properties (Hasher & Zacks, 1979; Treisman, Vieira, & Hayes, 1992). Logan (1992) has proposed reserving the term *automatic* exclusively for learned automatic processes.

Features of Automatic and Nonautomatic Processes

Most proponents of the capacity view regarded attentional requirements as ordered along a continuum (see Hasher & Zacks, 1979), with automatic processes at one end, depleting only minimal resources, and nonautomatic processes at the other, drawing on substantial resources.³ The difference in attentional requirements was linked to other functional properties and eventually led to the view that automatic and nonautomatic processes represent two opposite modes of processing, each characterized by a fixed set of features. Thus, it seems that the initial conception of automaticity as a continuum gave way to a dichotomous view. As an illustration, we briefly consider Shiffrin and Schneider's (1977) dual-mode model of information processing. Within this model, information processing is based on the activation of a sequence of nodes from long-term memory. The set of activated nodes is the content of short-term memory. Nodes can be activated in short-term memory in two distinct ways: either (a) by stimulus input, in which case activation spreads to remaining nodes in the sequence with little attentional demand, or (b) by a control process, through the allocation of attention. Examples of control processes are rehearsal and long-term memory search.

Shiffrin and Schneider (1977) seem to have used the term *control* in two different ways. In the first sense, a control process is one that enables a person to control the flow of information in and out of short-term memory. In the second sense, a control process is one that is under the person's control, meaning that it is chosen and established by the person (i.e., *intentional*) and that it can easily be *altered* or *suppressed*. In contrast, purely stimulus-driven activation requires no intention (i.e., is *unintentional*). A strong version of this sense is that purely stimulus-driven activation cannot be *avoided* (i.e., is *obligatory*) and, once started, is

difficult to *alter* or *suppress*. Because automatic processes use only minimal attentional capacity (i.e., are *efficient* or *effortless*), they may run in *parallel* with other concurrent automatic or control processes, with no interference. Because control processes require substantial attentional resources (i.e., are *nonefficient* or *effortful*), the simultaneous occurrence of multiple control processes is not possible, and the processes must be executed *serially*. In addition, they will be *slow*. Automatic activation travels *fast*, and individual nodes in an automatic sequence reside in short-term memory only briefly. The presence of a node in short-term memory is not sufficient for the content of that node to become conscious. Conscious experience involves only the subset of short-term memory that receives attention. Control processes are *conscious* because they use attention to bring or keep nodes in short-term memory. The nodes activated by automatic processes remain *unconscious* unless they receive attention. As soon as attention is focused on these nodes, they become part of a control process.

Not all of the enumerated features appear in various descriptions of the capacity view. The number and nature of the properties chosen, as well as the emphasis put on them, varies with details of the underlying information-processing model, with the topic under study, and with the specific research paradigm. For example, search tasks are designed to investigate the features efficient and fast, whereas Stroop and priming tasks are designed to investigate the feature unintentional (and sometimes unconscious). Participants in search tasks (and other skill development tasks) are explicitly instructed to engage in the process under study (e.g., detect a target) until, after some practice, it becomes efficient (e.g., Shiffrin & Schneider, 1977). Participants in Stroop and priming tasks, in contrast, are instructed to engage in a process (e.g., evaluation of the target) that is unrelated to the process under study (e.g., evaluation of the prime). Interference under these conditions indicates the unintentional nature of the studied process (e.g., Posner & Snyder, 1975a, 1975b). This difference in emphasis is reflected in the fact that Posner and Snyder (1975a, 1975b) chose the features unintentional, unconscious, and producing no interference as necessary features, whereas Shiffrin and Schneider (1977) considered the feature unintentional to be merely an optional feature. Bargh (1989, 1992, 1994), not a capacity theorist himself, singled out four optional features as the most frequently recurring

² It may be useful to alert the reader to a dual meaning of the term *attention*. Some formulations depict attention as a mechanism that carries out selection by allocating some of the system's limited processing resources to task-relevant input. According to other formulations, "attention itself denotes a computational resource of some kind, limited in quantity, which must be allocated selectively" (Allport, 1993, p. 184). The latter conception is apparent in the notions of attentional capacity and attentional resources. It should further be noted that we do not consider possible distinctions between the notions of capacity and resources and instead treat these constructs as interchangeable.

³ Following Logan (1988), we use the neutral term *nonautomatic* processes to refer to the counterpart of automatic processes. Other terms used are "effortful" processes (Hasher & Zacks, 1979), "control" processes (Schneider, Dumais, & Shiffrin, 1984; Shiffrin & Schneider, 1977), "controlled" processes (Bargh, 1994; Shiffrin & Schneider, 1977), "conscious" processes (Posner & Klein, 1973), "attentional" processes (Logan, 1980), "deliberate" processes (Anderson, 1982), and "strategic" processes (Posner & Snyder, 1975a, 1975b).

in the capacity literature—efficient, unintentional, uncontrollable, and unconscious—and termed them “the four horsemen of automaticity.”

The All-or-None Conception Challenged

The dual-mode model has been subject to severe criticism (see Bargh, 1992, 1994; Logan, 1985). It has been accused of propagating an all-or-none view of the distinction between automatic and nonautomatic processes. Such a view combines the idea of perfectly correlated features of each processing mode with the idea that both modes are mutually exclusive and that they exhaust the universe of possible processes. A process is either automatic or nonautomatic, and one has to investigate only one feature to know to which mode it belongs. The presence of the remaining features of that mode can be assessed via logical inference (Bargh, 1992).

The all-or-none conception has been challenged by studies showing a lack of co-occurrence among central features of automatic processes (e.g., Bargh, 1992; Kahneman & Treisman, 1984; Logan & Cowan, 1984; Logan, 1985, 1988, 1989). For example, Stroop interference, generally thought to be unintentional and uncontrollable, is not independent of attention because it can be diminished when attention is directed away from the targets (e.g., Francolini & Egeth, 1980; Kahneman & Chajczyk, 1983; Kahneman & Henik, 1981). Bargh (1992) pushed independence between features to the extreme, suggesting that for any random combination of features (some automatic, others nonautomatic), a process can be found that fits the description. He also suggested that a process matching the pure case of automaticity may be encountered only rarely (Bargh, 1994).

It became increasingly clear that an all-or-none conception of automaticity was no longer tenable. Some authors concluded, therefore, that the automaticity concept is not internally consistent and should be abandoned (e.g., Regan, 1981). They suggested that mental processes should be analyzed with respect to separate features without linking these features within a general concept of automaticity. In the following sections, we consider two arguments that call for a more cautious approach. First, although the capacity view cleared the way for an all-or-none view, early advocates of the capacity view do not seem to have actually endorsed it (see *Refining the Picture* below). Second, abandonment of the all-or-none view need not result in the abandonment of the automaticity concept (see *Gradual View* below).

Refining the Picture

A careful reading of the early writings by proponents of the capacity view reveals that its interpretation as strict all-or-none is an overstatement. For example, Shiffrin and Schneider (1977) allowed for automatic processes that are initiated under a person's control but that, once started, run to completion without further guidance. In addition, their model allowed for indirect control over the instigation of otherwise unintentional automatic processes: When a person keeps information in short-term memory that is associatively related to the nodes in an automatic sequence, the activation threshold for these processes can be lowered so that initiation is facilitated. Furthermore, Shiffrin and Schneider explicitly stated that interference between two automatic processes is inevitable when they use overlapping nodes in an incompatible

fashion. In such cases, automatic processes will not run in parallel. In addition, Shiffrin and Schneider distinguished between two types of control processes: those that are slow and consciously accessible (i.e., accessible control processes) and those that occur so quickly that they escape from consciousness and subjective control (i.e., veiled control processes). Finally, Shiffrin and Schneider made a distinction between two mechanisms of attention allocation. According to the first, a person voluntarily allocates attention to a stimulus. According to the second, attention is automatically drawn to a stimulus as a result of prior automatic relevance detection (see also Wundt's, 1896/1897, notion of active apperception). The fact that this automatic detection process calls attention is yet another element in this model that blurs the strict separation between automatic and nonautomatic processes. Moreover, as Schneider, Dumais, and Shiffrin (1984, p. 16) have pointed out, whenever an automatic call for attention takes place, other ongoing nonautomatic processing is disrupted (at least briefly). In a way, then, one could say that this type of automatic process “controls” nonautomatic processing.

Other authors, such as Hasher and Zacks (1979), have argued that even when a stimulus is attended, certain stimulus aspects are nevertheless encoded automatically. They further suggested that learned automatic processes are susceptible to disruption, reserving the feature uncontrollable exclusively for innate automatic processes. The status of “uncontrollable” as an automaticity feature has been questioned by other advocates of the capacity view as well. For example, Posner and Snyder (1975a) reviewed evidence showing that “automatic” priming effects can be blocked by effortful processes (see also Logan, 1980). Similarly, Schneider et al. (1984, p. 11) argued that automatic processes are not uncontrollable but rather that control is difficult and requires substantial effort.

Through the years, some proponents of the capacity view came close to undermining a fixed, feature-based definition of automaticity. Schneider et al. (1984, pp. 20–21) examined 12 criteria and concluded that none was necessary or sufficient for the distinction between automatic and nonautomatic processes; however, they still pointed to control and resource demands as being the least problematic. Shiffrin (1988) evaluated several criteria for automaticity, ultimately retaining not one as generally applying to all automatic processes. He eventually came to favor a gradual view. In the following section, we discuss the essence of the gradual view by considering Logan's (1985) answer to the problems posed by the lack of co-occurrence among automaticity features.

Gradual View

Logan (1985) offered an alternative to abandoning the concept of automaticity. He (re)introduced it as a continuum and (re)emphasized that the automaticity of a process is determined by the amount of training. He also suggested that “each property has its own time-course of change with practice” (Logan, 1985, p. 373). A process can switch from intentional to unintentional with very little practice, whereas it switches rather slowly from nonefficient to efficient. In most studies, different features are measured at only one point in practice; thus, the lack of co-occurrence among the features is no surprise. This lack, however, need not be interpreted as evidence against the internal consistency of the automaticity concept. In Logan's (1985) words, “we don't want to throw the

baby out with the bath water" (p. 375). Only if it can be demonstrated that different properties evolve in different directions on the automaticity continuum as practice increases does the automaticity concept lose its value. It should be noted, however, that the gradual view tends to weaken the notion of automaticity. As a gradual concept, automaticity loses its ability to distinguish one type of process (automatic) from another (nonautomatic), for any process can be labeled automatic to some degree.

Automaticity as Single-Step Memory Retrieval

Although the dichotomy assumption may have been part of the capacity view, it was not a core assumption. Therefore, it could easily be replaced by the assumption that automaticity is gradual without changing the capacity view's basic idea of automaticity as relative independence of attentional resources. However, the capacity perspective of automaticity has been challenged with respect to its basic assumptions as well (see Logan, 1991). First, many of the studies illustrating the lack of co-occurrence among automaticity features have disconfirmed the basic idea that automatic processes are unconstrained by attentional capacity (e.g., Kahneman & Chajczyk, 1983). A second weakness is that the capacity view is devoid of an explicit learning mechanism to explain the development of automaticity (Logan, 1991). The capacity view does not specify why a process needs fewer attentional resources after practice. Third, the underlying single-capacity theory of attention has been challenged by the rise of multiple-resource theories (e.g., Navon & Gopher, 1979; Wickens, 1984). The idea that processes can drain capacity from more than one resource pool undermines the conception of automaticity as processing with minimal resources. It is possible that when a process uses less of one resource pool, it taps the resources of another pool. A fourth criticism of the capacity view (single or multiple) came from scholars who rejected the notion of attention as limited capacity. Allport (1989) argued that the requirement of selective processing is not imposed by the system's limited capacity but rather by the serial nature of behavior. Attention ensures selectivity, not to protect the system from informational overload but to support coherence in behavior in the face of the multiple conflicting impulses impinging on the system.

Logan (1988) offered an alternative account of automaticity with the aim of addressing the weaknesses inherent in the capacity view. He explained the development of automaticity as a transition from algorithm computation (or multistep memory retrieval) to single-step memory retrieval. Nonautomatic processing relies on algorithm computation to produce an output. This output may be stored in memory so that future processing may bypass computation of the algorithm and rely entirely on memory retrieval. For example, when children learn to add pairs of digits, they initially count the units in both digits to produce the sum. After some practice, an association is formed in memory between a pair of digits and their sum, and children can immediately retrieve the sum from memory.

Logan (1988) postulated that single-step memory retrieval is the only (low-level) process that can account for automatic performance (or automatic high-level processes).⁴ As a next step, Logan proposed that researchers should no longer define automaticity in terms of features but instead in terms of the underlying process—single-step memory retrieval. In other words, he proposed replac-

ing the feature-based approach with a construct-based approach. Although this view still allows automatic and nonautomatic performance to be described with features, it does not accredit any diagnostic value to them. To assess the automaticity of performance, one must determine whether it relies on direct memory retrieval.

In sharp contrast with the capacity view, Logan (1992) conceived of automaticity as dependent on attention. Attending to a stimulus is sufficient to store it in memory or to retrieve from memory all information that was associated with it during a former presentation (see also Hasher & Zacks, 1979). Logan (1988) argued that both storage and retrieval are improved when attention is focused on the stimulus. In addition, when storage is improved, stronger retrieval cues are created, and automatic retrieval is more likely. Finally, Logan (1992) suggested that automatization produces a shift in attention (rather than a reduction), going from lower to higher levels of organization. When components of a skill become automatic with practice, attention is shifted from them to higher-level aspects of the skill that are concerned with integration (see also Bryan & Harter, 1899; James, 1890; Shiffrin, 1988).

Algorithm Strengthening View

Logan (1988) explained automatization as a shift from one process (algorithm computation) to another (single-step memory retrieval). In contrast, other theorists have construed automatization as a change in the same underlying process. They have emphasized the strengthening or improvement of algorithm computation itself as the main mechanism underlying automatization (e.g., Anderson, 1992; Rosenbloom & Newell, 1986). The strengthening of algorithms is one of the learning mechanisms incorporated in Anderson's (1982, 1983) adaptive control of thought (ACT*) model of skill acquisition. Although the ACT* model also contains other learning mechanisms (including a reduction of procedural steps to a single step and the strengthening of declarative facts), automatization is attributed mainly to the strengthening of algorithms (Anderson, 1992, p. 170).

Both the direct memory retrieval view and the algorithm strengthening view are construct based in that they propose one learning mechanism to account for automatization. In contrast to the direct memory retrieval view, however, the algorithm strengthening view hypothesizes that the same algorithms responsible for the initial, nonautomatic stage of performance are also responsible for the skilled, automatic stage of performance. In the latter stage, these algorithms are executed faster and more efficiently. Hence, automatic and nonautomatic algorithms differ only with regard to the features (such as speed and efficiency) they possess. In summary, the algorithm strengthening view is a construct-based view with regard to the explanation of *automatization* but a

⁴ Processes can be described at different levels of analysis (Marr, 1982; see the section (*Unconscious*)). Performance can be explained by high-level processes, and these processes can in turn be explained by low-level processes. For example, performance on the addition task can be explained by the high-level process of adding digits and further down by the low-level processes of single-step memory retrieval or algorithm computation. Logan (1988) imposes no restrictions on the types of high-level processes that can be automatic. Everything we say about performance in terms of Logan's position also applies to high-level processes.

feature-based view with regard to the definition and diagnosis of *automaticity*.

In both the direct memory retrieval view and the algorithm strengthening view, consistent practice is an important factor in the development of automaticity. However, the two views invoke different types of consistency (Carlson & Lundy, 1992). A shift toward direct memory retrieval occurs with practice on consistent data, whereas algorithm improvement requires algorithms to remain consistent with practice but not the data on which these algorithms act. Building on memory architectures that distinguish between declarative and procedural memory systems, several authors have suggested that algorithms that have reached an automatic stage may be stored in the procedural system (e.g., Anderson, 1992, 1996; Tzelgov, Yehene, & Naveh-Benjamin, 1997). The main point of debate is not whether memory is involved in automaticity but whether the system is able to store and automatically retrieve algorithms in addition to single facts. Moreover, even if retrieval is required to set the stored algorithm in place, the computation of the algorithm may fall beyond the boundaries of retrieval (at least according to those theorists who argue for a qualitative distinction between memory retrieval and algorithm computation).⁵

Support for algorithm-based automaticity has come mainly from skill development studies in which the benefits of practice with consistent algorithms transferred to data that were never encountered during practice (Carlson & Lundy, 1992; Kramer, Strayer, & Buckley, 1990; Schneider & Fisk, 1984; Smith & Lerner, 1986). These benefits were evident in increased speed and efficiency. Recent studies using Stroop-like and priming tasks also provided support for the unintentional nature of algorithm computation (Moors, De Houwer, & Eelen, 2004; Tzelgov, Yehene, Kotler, & Alon, 2000). In fact, Tzelgov and colleagues (Tzelgov, Henik, Sneg, & Baruch, 1996; Tzelgov et al., 2000) have advocated a reconciling approach, suggesting that both mechanisms of algorithm improvement and a shift toward direct memory retrieval may be equivalent mechanisms for the development of automaticity.

Separating Issues of Definition, Explanation, and Diagnosis

Logan (1988) explained automaticity as single-step memory retrieval. He proposed, moreover, to define automaticity as single-step memory retrieval and to diagnose automaticity by determining whether performance was based on single-step memory retrieval. Although we find it plausible that an important part of automatic performance is based on single-step memory retrieval, we think that defining automaticity in terms of its presumed explanation is problematic because it confuses the *explanandum* with the *explanans*, leading to circular reasoning. We prefer to define automaticity in terms that are different from those that figure in the potential explanation, as is the case when automaticity is defined in terms of features.

The definition and explanation of automatic performance should be formulated in nonoverlapping terms, but the diagnosis of automatic performance can be based (a) on the variables specified in the definition (features) or (b) on the variables specified in the explanation, provided that the explanation is exclusive and exhaustive. More precisely, one can take the presence of single-step memory retrieval as an indication that the performance was auto-

matic only if one assumes that single-step memory retrieval cannot lead to nonautomatic performance (exclusiveness). Conversely, one can take the absence of single-step memory retrieval as an indication that the performance was nonautomatic only if one assumes that no other process besides single-step memory retrieval can lead to automatic performance (exhaustiveness).

The exhaustiveness assumption has been called into question by empirical studies showing that certain algorithms may lead to automatic performance as well (cf. above). In any case, when the purpose is to investigate whether algorithms can lead to automatic performance, one cannot diagnose the automatic nature of the performance by checking whether it was based on single-step memory retrieval.

In summary, when the research purpose is to diagnose a performance as automatic or nonautomatic, one can (a) look at the features of this performance, or alternatively, (b) determine whether it was based on the process of single-step memory retrieval, given the acceptance of certain assumptions (exclusiveness and exhaustiveness). When the research purpose is to investigate which low-level process can lead to automatic task performance or which low-level process can be diagnosed as automatic, one has no choice but to look at the features of the performance and use the resulting information to make inferences about the automatic nature of the underlying process.

Autonomy and the Conditional Approach

We believe that the previous analysis provides good arguments for returning to a feature-based approach toward the definition and diagnosis of automaticity, such as the gradual approach discussed earlier. As we mentioned, one drawback of this approach is that it does not draw a clear line between automatic and nonautomatic processes. Bargh (1992) has provided a solution to this problem by acclaiming one feature as the minimal criterion for automaticity: the feature autonomous. He defined an autonomous process as one that, once started (and irrespective of whether it was started intentionally or unintentionally), runs to completion with no need for conscious guidance or monitoring. In addition, Bargh posited that all automatic processes are conditional: They are all dependent on a set of preconditions. Automatic processes may vary with regard to the set of preconditions they require (e.g., the presence of a

⁵ Some researchers define algorithm computation as multistep memory retrieval, implying that the only difference between algorithm computation and single-step memory retrieval is the number of steps involved. Other researchers hold that the distinction between algorithm computation and memory retrieval is not just a matter of quantity but also of quality. Some of these authors argue that memory retrieval is just a very simple type of algorithm (linking an input to an output on the basis of a previous correlation or contiguity relation) and that this algorithm is qualitatively different from other types of algorithms. Others argue that memory retrieval is not even a simple algorithm because algorithms capture abstract relations between data without specifying the data themselves and memory retrieval is data specific (Sloman, 1996). Indeed, the activation of memory traces (no matter how many inputs are involved) can take place only along previously established paths; hence, it can produce an output only for previously encountered stimuli or (if the notion of similarity is added, cf. Palmeri, 1997) for stimuli that resemble the previously encountered ones.

triggering stimulus, awareness of the stimulus, the intention that the process take place, a certain amount of attentional resources, the salience of the stimulus). According to this conditional perspective, the study of automaticity should be concerned with establishing the set of preconditions that must be in place for an autonomous process to occur.

Although Bargh (1989, 1992) suggested that processes can be endowed with any random combination of features (some automatic, others nonautomatic), he argued that most automatic processes fall into three broad categories depending on the conditions they require: preconscious, postconscious, and goal dependent. Roughly, preconscious processes require no conscious input and no intention. Postconscious processes require a conscious input but no intention. Goal-dependent processes require a conscious input and an intention or another processing goal to start but no goal to run to completion (i.e., autonomous).

Conclusions

The approaches to automaticity discussed above share the idea that extensive practice leads to learning, which may be accompanied by changes in attention, awareness, control, speed, and accuracy. Feature-based views focus on these changes but lack a mechanism to explain why well-practiced processes are modified in such ways. This shortcoming is addressed by the construct-based approach. We agree that it is instructive to look at the learning mechanisms underlying automatization and not just at the changes that occur in the features of processes. However, when it comes to diagnosing automaticity, turning to features seems to be the most cautious approach.

As contemporary theorists have argued, individual features should be investigated separately. This decompositional approach, as we call it, is a direct consequence of rejecting an all-or-none view and is favored by at least three of the above-discussed theoretical perspectives: (a) the view that does away with the overarching concept of automaticity, (b) the gradual view that rescues the automaticity concept but does not provide a criterion to distinguish automatic from nonautomatic processes, and (c) the approach that chooses autonomy as the minimal criterion for automaticity and allows processes to differ with regard to all features except autonomy (Bargh, 1992). Of these three views, the last has considerable conceptual appeal. It gives unity to the concept of automaticity and provides a criterion to distinguish automatic from nonautomatic processing. It is not clear, however, why the feature autonomous should be the chosen one. Should we choose it because it allows the greatest variety of processes to be included in the family of automatic processes? Given this indeterminacy about which feature(s) should be regarded as essential, we find ourselves drawn to the gradual approach. We have no choice but to investigate the presence of automaticity features one by one, without being able to draw an objective line between the automatic and the nonautomatic member in each feature pair.

Note, however, that a decompositional approach makes sense only when the features are conceptually separable. The first item on the agenda, therefore, is to consider how features of automaticity have been defined. The literature reveals that authors often do not delineate these features in a strict sense and that their definitions often display conceptual overlap. The blurring of automaticity features raises the question of how separable these

features really are and, by implication, to what extent a decompositional approach is viable.

Before we turn to a detailed discussion of automaticity features, we should point out that the early capacity view, which built on the skill development tradition, was not solely responsible for the creation of the all-or-none view. Another research tradition that contributed to this view was the new look movement in perception (Bruner, 1957), a tradition that was originally inspired by the constructivist ideas of Freud (1915/1984) and Bartlett (1932) and that generated a marked interest in the distinction between conscious and unconscious processing. Bargh (1996; Bargh & Chartrand, 2000; Wegner & Bargh, 1998) characterized the skill development tradition as concerned with goal-dependent automaticity and the new look movement as concerned with preconscious automaticity. We elaborate further on this issue in our discussion of the features conscious and unconscious.

Features of Automaticity

Below, we review a selection of the features mentioned in the research literature. The first section groups together features that can be defined in relation to goals: (un)intentional, goal directed, goal (in)dependent, (un)controlled/(un)controllable, and autonomous. We call these *goal-related features*. We also discuss the feature “purely stimulus driven” (which is not strictly goal related) in this section because it is most easily defined in contrast with the feature unintentional. The remaining sections discuss the features (un)conscious, (non)efficient, and fast (slow). In one part of each section (*Concepts*), we look at definitions of features that we have encountered in the literature. Although there may be no objective criteria to favor one set of definitions over another, in selecting our own definitions, we adopted three criteria. First, we wished to do justice to some representative philosophical writings on these features. Second, we wished to remain close to natural language. Third, to the extent that it was possible, we selected those definitions that were most successful in avoiding overlap among features. Our purpose was to consider to what extent features can be parsed to a point of irreducibility. Only if they can be parsed to this extent does it make sense to argue for the separate investigation of them. In a second part of each section (*Links*), we consider assumptions of overlap among features. Many feature descriptions in the literature contain references to other features. Some of the conceptual overlap among features may be due to a lack of precision, but some of it may result from theoretical assumptions about the cognitive apparatus. In a third part of each section (*Measurement*), we give a few illustrations of how features have been or can be operationalized. In the final part of each section (*Diagnostic Status*), we discuss different positions regarding the status that researchers confer to the feature(s) as diagnostic of automaticity. Researchers are likely to investigate those features they consider most diagnostic. A feature’s diagnostic status may refer to whether the feature is necessary or optional, as well as to whether it is a strong or weak indicator of automaticity. These dimensions do not necessarily overlap: A feature may be both necessary and weak or both optional and strong. Further, features may have diagnostic status of their own, or alternatively, they may derive their status from links to other automaticity features.

Goal-Related Features

Concepts

(Un)intentional, goal directed, and goal (in)dependent. The core meaning of the word *intentional(ity)* is directed(ness) (deriving from the Latin verb *intendere* [to be directed at]; Jacob, 2003). Several theorists have made a distinction between *Intentionality* in philosophical use and *intentionality* in ordinary use. (We capitalize the former but not the latter, following Searle's, 1983, convention.) In the philosophical use, *Intentional(ity)* refers to a property of mental phenomena by which they are directed at or about some state of affairs outside of themselves (Brentano, 1874). For example, desiring, believing, perceiving, and thinking are *Intentional* because they require an object. In ordinary language use, *intentional(ity)* derives from the mental phenomenon of intending or having an intention and is often said to be a property of actions (e.g., Wakefield & Dreyfus, 1991). The essential difference between *Intentional* and *intentional* is not that the former relates to mental phenomena and the latter to actions, for when the category of actions or acts is extended to include covert or mental acts, the distinction between both concepts does not disappear.

What we believe to be crucial is that with *Intentionality*, the mental act is directed at some state of affairs outside the mental act, whereas with *intentionality*, some person is directed to engage in the act. Formally, directedness can be considered as a type of relationship, and the source and terminus of this relationship are different with *Intentionality* and *intentionality*: With *Intentionality*, the source is the mental act and the terminus is a state of affairs outside of the mental act, whereas with *intentionality*, the source is some person and the terminus is the act. In other words, *Intentionality* speaks about a mental state or act, whereas *intentionality* speaks about the way in which a mental or other act is caused (or about the source of the act, the person). An intentional act can be defined as one that was brought about by a person who was directed at engaging in the act. The phrase *intentional act* is thus shorthand for the phrase *intentionally produced act*. The remainder of the present section elaborates on *intentionality* with a lowercase *i*. We return to *Intentionality* with a capital *I* in our discussion of the features conscious and unconscious. Most of the thinking about *intentionality* has been done with respect to overt behavior. Assuming important parallels between overt acts or actions and covert acts or processes, we use the general term *act* in the first part of our discussion. Later, we narrow our discussion to processes only (also see Figure 1).

The terminus in an *intentionality* relation has often been called a goal. One way that a system can be directed at something is by forming a representation of it (James, 1890; for an alternative account, see Wakefield & Dreyfus, 1991; but for a critique of this alternative, see Barbiero, 2004). The idea that goals can be represented in memory just as can any other type of information is now widely accepted (Bargh, 1990a; Bargh, 1996; Moskowitz, Li, & Kirk, 2004). It appears that there are two different uses of the term *goal*. In the phrase *directed toward a goal*, *goal* is used to indicate a desired state of affairs in the future. In the phrase *driven or caused by a goal*, *goal* is used to indicate a representation of that state of affairs, which, when activated, may function as an antecedent cause directing the organism toward that state. We specify which meaning we use only when there is risk of confusion.

Several authors have stated that an intention is a special kind of goal that has as its content an act and not just any state of affairs (see also James, 1890). For example, one can intend to clean the house, but one cannot intend the house to be clean. Moreover, an intention must always be carried out by the person who has it. For example, one cannot intend someone else to clean the house. A person may want or desire that the house be clean or that someone else clean it, but he or she cannot intend this state of affairs. In summary, then, an intentional act is one that is caused by the goal (in the sense of representation) to engage in the act. In contrast, an unintentional act is one that is not caused by the goal to engage in the act.

We need to delineate the concept *intentional* from the concepts *goal directed* and *goal dependent*. Because the core meaning of the word *intentional* is directed, one would expect a strong overlap between the terms *intentional* and *goal directed*. On closer consideration, however, it is evident that the source and terminus of the directedness relation are different. In an intentional act, it is a person (source) that is directed onto the act (terminus), whereas in a goal-directed act, it is the act (source) that is directed onto a further goal (terminus).

Although acts often combine the features *intentional* and *goal directed*, both features can, in principle, occur independently. Not all intentional acts are directed at a further goal, and not all acts that are directed at a goal are intentional. In the first case, a person merely has the goal to produce an act, for example, to raise his or her arm, without this act being instrumental for obtaining a further goal. The person does not pursue another goal than the goal to raise his or her arm. In the second case, a person is directed at a goal, for example, the goal to have an apple, and this goal triggers an act or a series of acts that are instrumental for the goal, such as the acts of raising one's arm and plucking the apple from a tree. It is possible, at least theoretically, that the acts of raising and plucking were not caused by the goals of engaging in these specific acts but that they were direct effects of the activation of the goal to have the apple. Something similar pertains when a person pursues a goal to engage in a complex act and the details of that act are not specifically represented in the goal. For example, the overarching goal to drive a car may give rise to several subsidiary acts, such as changing gears and turning the wheel. These subsidiary acts are goal-directed because they are instrumental for fulfilling the overarching goal of driving. At the same time, they are unintentional in that they were not caused by the subsidiary goals to engage in these specific acts.

It may be more accurate to use the term *goal dependent* instead of *goal directed* (which amounts to the same thing). The subsidiary acts in our example are dependent on the overarching goal to engage in the complex act, because without that goal, they probably would not have occurred. It is debatable, however, whether there is also something in the acts themselves that actually directs to the overarching goal. The goal directedness of these acts may be true only from a spectator's point of view. Using the term *goal dependent* also helps to clarify the relation between this concept and the concept *intentional*. A goal-dependent act is one that depends on a goal for its occurrence. This goal may be either the proximal goal to engage in the act or another, remote goal that triggers the act directly, circumventing the proximal goal. In the former case, the act is an intentional goal-dependent act; in the latter case, the act is an unintentional goal-dependent act. Thus,

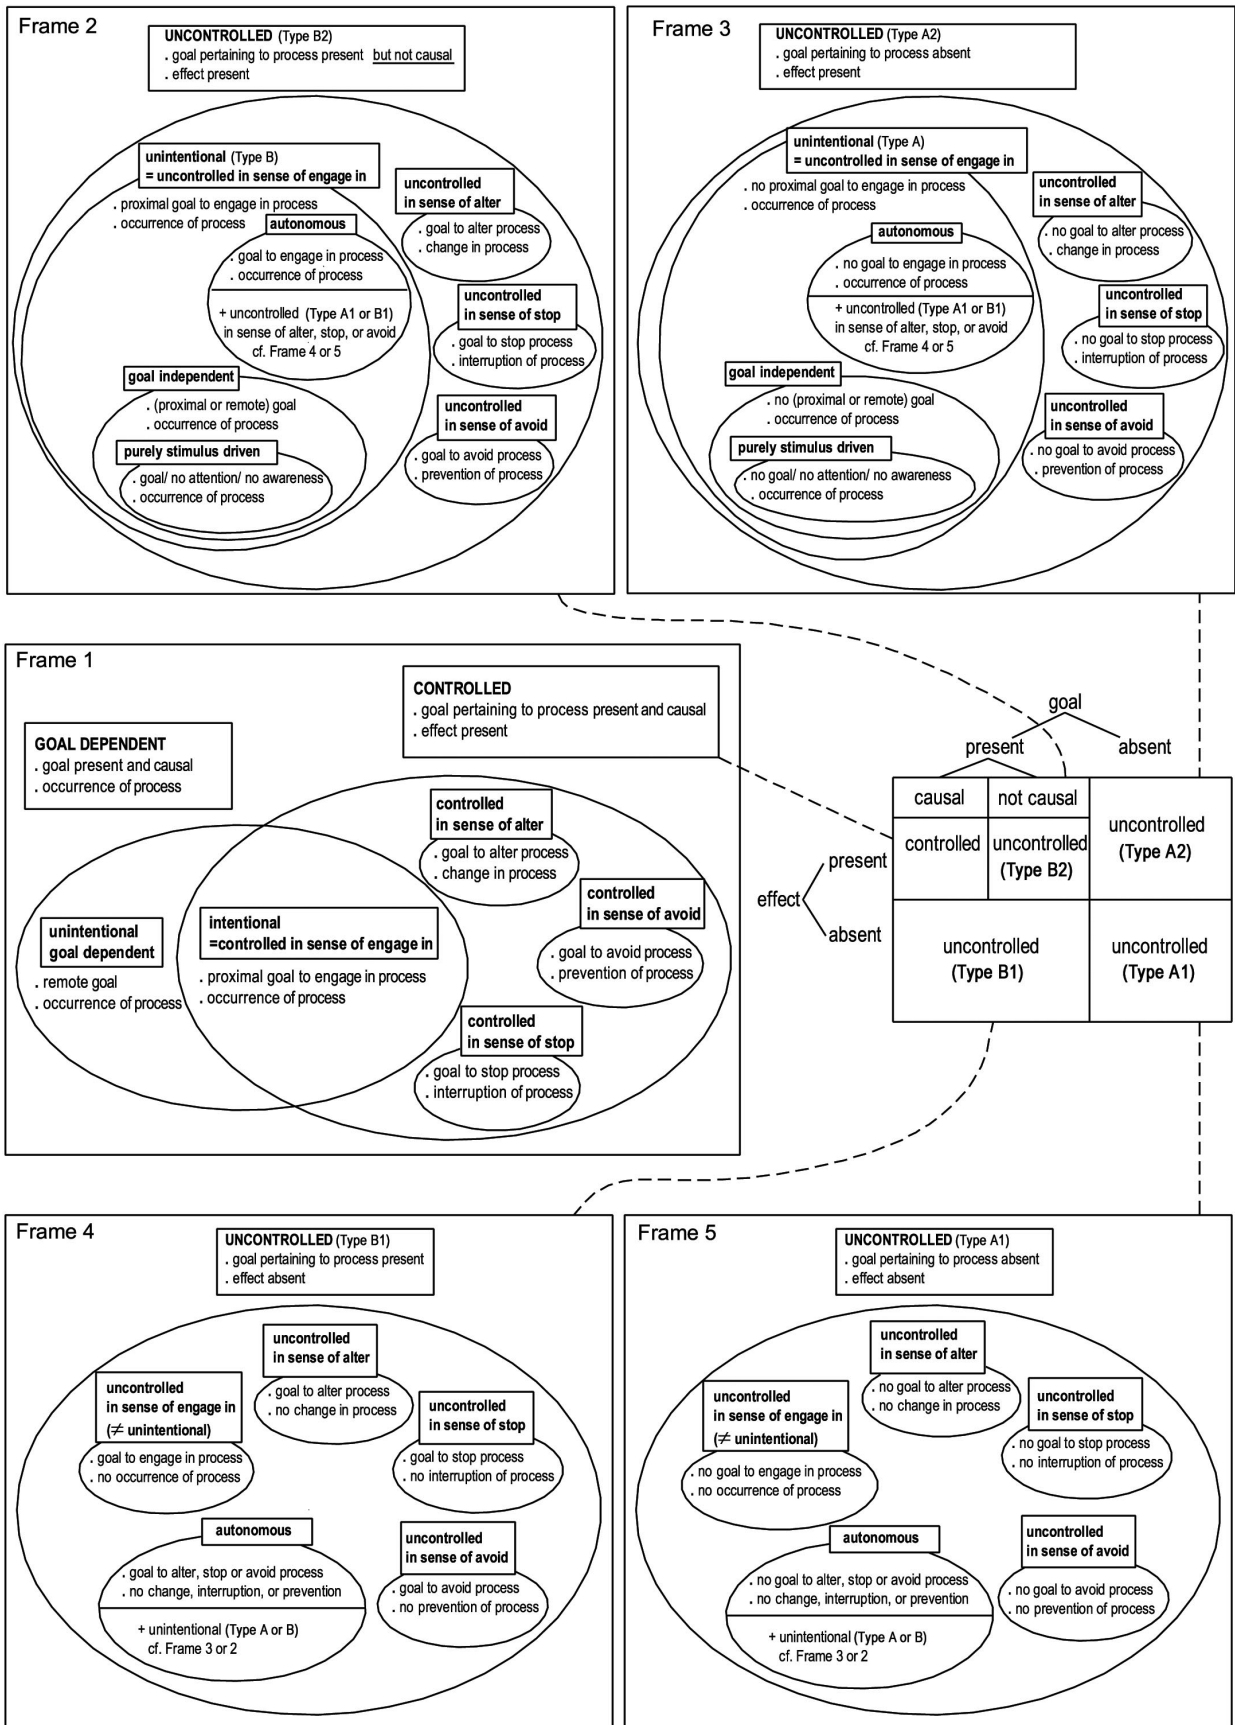


Figure 1. Conceptual overlap among the goal-related features (and the feature purely stimulus driven).

intentional acts are a subclass of goal-dependent acts (see Figure 1, Frame 1). Acts can also depend on both a remote and a proximal goal. A goal-independent act is an act that does not depend on a goal for its occurrence, neither proximal (i.e., unintentional) nor remote. By definition, then, goal-independent acts are a subclass of unintentional acts (see Figure 1, Frame 3).

The examples above illustrate why it is important to specify the unit of the act under study. Acts can be described at various levels, from low-level descriptions containing physical details to more abstract, high-level descriptions (e.g., Vallacher & Wegner, 1987). For example, “playing the piano” occupies a higher level of description than does “moving one’s fingers on the keys,” although both pertain to the same objective act sequence. Specifying the level of the act is important because it may influence one’s conclusions about intentionality. If the person’s current goal contains the high-level act but not the low-level one, then the same act sequence can be labeled intentional when it is described in high-level terms but not when it is described in low-level terms. This is akin to what philosophers mean when they say that an act can be intentional under one description but not under another. It simply means that some aspects of an act are intentional whereas others are unintentional (Searle, 1983). As another illustration of this idea, consider the example of a pedestrian who is walking in the woods and while doing so crushes a colony of ants. On the supposition that the pedestrian pursued the goal to walk but not the goal to crush ants, the same objective act sequence can be termed intentional under the description of “walking” but not under the description of “crushing ants.” Our definition of intentional acts can thus be refined as follows: An act is intentional under a certain description when it is produced by a goal representation that has as its content the act under that description.

In addition to specifying the level or aspect of an act sequence, it is also important to delineate its boundaries. Strictly speaking, an act should not be confounded with its results. For example, the act of throwing a dart should not include the arrival of the dart in the bull’s eye. A person can be directed at the state of the dart being in the bull’s eye, but this state of affairs does not strictly belong to the act of throwing the dart. Things become more complicated if we consider that high-level descriptions of acts often include their results. For example, the act of “throwing a dart” may be described at a higher level as “hitting the bull’s eye.” One may ask whether an act described in a way that includes its result can, in principle, be intentional. Some philosophers (Mele & Moser, 1994) have argued that it may be justified to call “hitting the bull’s eye” an intentional act only when a player possesses the skill to hit the bull’s eye. When an unskilled player hits the bull’s eye, it is an accident instead of an intentional act, even when the player had the intention to hit the bull’s eye (see below).

The view that every intentional act should be caused by an intention has been dubbed “the simple view” (Bratman, 1987). Opponents of this view have rejected the idea that all intentional acts must be caused by intentions. They have argued that some acts qualify as intentional even when there is no intention (in the sense of goal representation) that precedes it (Wakefield & Dreyfus, 1991). Examples of this include subsidiary acts involved in skilled behaviors such as driving a car or walking down the stairs. Opponents of the simple view believe it is unlikely that each of the subsidiary acts of a complex act are preceded by a representation of them; nonetheless, they believe that there are reasons to classify

these subsidiary acts as intentional. We wish to make a few comments. First, it seems that opponents of the simple view have confounded intentionality with goal directedness. Second, we believe that there are no a priori grounds to reject the idea that the subsidiary acts in skilled behaviors can be caused by representations of these acts. To render this idea more plausible, we should clear up two possible misunderstandings.

1. Our notion of intentionality entails the existence and activation of a goal representation. It does not, however, entail the conscious activation of a goal representation. Recent evidence shows that goal representations can be activated either consciously or unconsciously just as other types of representations can (e.g., Bargh, 1990a; Bargh, Gollwitzer, Lee-Chai, Barndollar, & Trötschel, 2001; Chartrand & Bargh, 1996; see review by Moskowitz et al., 2004). It follows that the goal representations underlying acts may be activated either consciously or unconsciously and that intentions do not equate with conscious intentions. Although it seems implausible that certain subsidiary acts would be represented in the person’s conscious intentions, they may still be represented in the person’s unconscious intentions, and these acts may still count as intentional. For example, a proficient piano player may not consciously think of every single note before playing it, but the notes may nevertheless be represented in memory. The unconscious activation of these representations may guide the musician’s performance.

2. We may have given the impression that goal representations are always verbal or conceptual in nature. If this were true, it would be unlikely that subsidiary acts described at very low levels, such as turning the wheel to a certain degree and with a certain speed, could ever be represented (even unconsciously). However, if verbal forms of representation are replaced with representations that are imagelike (Searle, 1983) or kinaesthetic (James, 1890), there is a much broader range of acts that can plausibly be produced by goal representations. According to James (1890), we have a most delicate foreshadowing of the kinaesthetic sensations that follow upon our movements. “Why else the start of surprise that runs through us if someone has filled the light-seeming box with sand before we try to lift it, or has substituted for the cannon-ball which we know a painted wooden imitation? *Surprise* can only come from getting a sensation which differs from the one we expect” (p. 502). It is of course an empirical question as to how much detail can be included in one’s intentions. Support for the existence of goal representations or motor programs preceding action comes from research showing that animals whose sensory feedback systems are impaired can still perform movements skillfully, as well as from studies showing that the time to initiate a sequence of movements can depend on characteristics of the sequence, such as its length (Rosenbaum & Saltzman, 1984).

We stated that, according to the simple view (i.e., the causal, representationalist view), an organism can be directed to an act by realizing a representation of it. Among proponents of this view, there are still disagreements about the nature of these representations. Some authors call them goal representations, imparting them with special dynamic qualities that are not shared by other kinds of representations (Bargh & Barndollar, 1996). Others argue that act representations are purely cognitive structures. According to the so-called *common-coding hypothesis*, the same representations are used both to perceive and to engage in an act (Prinz, 1990). At first glance, it seems that James (1890) subscribed to the latter view

when he stated that the mere thought of an act is sufficient to launch it (his notion of ideomotor acts). But according to James, this is because representations (which James restricted to conscious representations) are inherently impulsive (i.e., endowed with dynamic qualities).

Let us now return to our definition of an intentional act as one that was caused by the goal to engage in it. This definition contains three ingredients: (a) the intention or goal to engage in the act, (b) the act (bodily motion or mental process), and (c) the causal connection between the goal and the act (Davidson, 1980; Searle, 1983), with the goal being the cause and the act being the effect. When both the goal and the act are present but the goal is not the cause of the act, the act is not intentional. For example, suppose a person wants to kneel down to propose marriage. She kneels, but not because of her goal to do so. Rather, she falls to her knees because someone pushed her. The act of kneeling does not count as an intentional act.

A qualification of the third ingredient is that the goal should cause the act in the right way (Davidson, 1980) or proximally (Brand, 1989). When the goal to engage in an act causes it, but only indirectly by means of other intervening factors, the act is not intentional. For example, suppose a person has the goal to kneel down to propose marriage. She is overwhelmed with emotion, feels weak, and sinks to her knees. The act of kneeling cannot be classified as intentional, although it was indirectly caused by the goal to kneel (Davis, 1994).

The three ingredients—intention, act, and causal connection—can be regarded as the formal criteria needed to classify an act as intentional. In addition, the literature on goal striving (Gollwitzer, 1990; Heckhausen, 1990/1991) has identified a number of factors that determine whether an intention will translate into an act. The implementation of an intention must be understood within the context of multiple intentions competing for implementation. Within this context, it becomes clear that the strength of the critical intention must be high enough in comparison to other intentions. If this is not the case, then a fiat or decision to act on the intention is required (James, 1890; Moya, 1990; Wundt, 1902). Such a fiat can be conceived of as a second goal superimposed on the intention. According to James's (1890) notion of ideomotor action, a representation of an act discharges itself directly in the initiation of the act, unless representations of other acts prevent it. Only when there are concurrent representations is a fiat required to overcome their counteracting forces and to launch the act.

Another factor that is necessary to secure the implementation of an intention is the feasibility of the act.⁶ The factor feasibility encompasses two subfactors. First, there must be an opportunity to act. To implement the intention to pour milk into a cup, one must have milk and one must have a cup. Gollwitzer (1990) posited that in cases in which the opportunity to engage in the act is not present, the likelihood of benefiting from a future opportunity is increased when the person has specified in advance a cue that signals a suitable opportunity. Gollwitzer coined the term "implementation intention" to refer to the mental link between the representation of the cue and the representation of the act that is to follow. (This link is initially established in a deliberative manner, but once it is stored in memory, the detection of the cue triggers the act in a direct manner.) When an implementation intention is formed, the risk that opportunities to act go by unnoticed is strongly reduced. A second subfactor of feasibility is skill (Mele &

Moser, 1994). A person can be expected to intentionally produce an act only when he or she has the skill to perform it (cf. the darts example). Some acts cannot, in principle, be performed intentionally because they depend too much on luck. Winning the lottery, for example, is not an intentional act, even when the winner has the intention to win. This is because the player has no skill for winning the lottery. It is not his intention that caused the winning; it is luck that lent him a hand.

In summary, an act can be categorized as intentional when it is caused (in the right way) by an intention (i.e., a goal to engage in the act). Several factors determine whether the act will come about. Some of these factors must always be in place (e.g., opportunity, skill); other factors are necessary only when still other factors are absent (e.g., a fiat is necessary when strength is too low, and an implementation intention is necessary when there is no current opportunity). An act is unintentional when it is not caused by an intention (or when it is not in the right way). This may be because there is no intention (Type A), or there is an intention, but the intention is not the cause of the act (Type B).

(Un)controlled/(un)controllable. The word *control* has also been related to the notion of directedness and the concept of goals. Wegner and Bargh (1998, see also Bargh, 1990b) have proposed that, at the most basic level, to control something is to influence it in a certain direction and not randomly. Dennett (1984) likewise has argued that control is not mere causation but rather desired causation.⁷ A controls B if A can drive B to be in a state that A wants B to be in. These notions of control implicate a goal about a target (something or someone) and some form of influence or causation that results in the achievement of the goal. To drive or to influence a target toward a state of affairs means not only that one is directed toward that state of affairs but also that one is successful in realizing it. Thus, a person controls an act when he or she has a goal about it and when he or she achieves this goal. This definition of control contains three ingredients: (a) a goal pertaining to an act, (b) an effect (i.e., the state represented in the goal), and (c) the causal connection between the goal and the effect. The word *control* is used to characterize a person. In relation to an act, it makes sense to use the words *controlled* and *controllable*. An act is controlled when a person exerts control over it. A controllable act is one that is possible for a person to control.

As Dennett (1984) pointed out, A can drive B only into a state that belongs to B's normal range of states. Examples of normal-range goals about acts are the goal to engage in the act—which

⁶ In the motivation literature, a distinction is made between the phase in which a goal is chosen (goal setting) and the phase in which the goal is implemented, that is, translated into action or processing (goal striving; Gollwitzer, 1990; Heckhausen, 1990/1991). It should be clear that the present analysis deals only with goal striving. The feasibility factor discussed here should not be confused with the feasibility factor that has often been discussed in relation to goal setting (cf. Gollwitzer, 1990).

⁷ Whereas this use of the term *control* explicitly refers to desired causation, it has been used in the sense of mere causation or influence as well. For example, when the activation of a process is said to be under environmental control, it simply means that the environment causes activation of the process, for unless one is referring to the social environment, an environment is not supposed to have goals. But perhaps the term *control* is just used here in an anthropomorphist sense, "as if" the environment has goals.

itself can be parsed into the goals to start, continue, and complete the act—as well as the goals to avoid the act, to alter its course, or to interrupt or stop the act. Given the multitude of goals that people can pursue about acts, we recommend always specifying the goal that one envisions when using the term *controlled*. To control an act in terms of the goal to engage in the act is not the same as to control the act in terms of the goal to alter it.

Comparing the definitions of a controlled act and an intentional act, we have found that there is some overlap. We stated that an act is controlled in terms of some goal when the person has a goal about the act and this goal produces the state represented in the goal. We also stated that an act is intentional when it is produced by the goal to engage in it. An intentional act is thus also a controlled act, at least controlled in terms of the goal to engage in the act. The concept control was defined not only in relation to the goal to engage in an act but also in relation to other goals, such as the goals to alter, stop, or avoid an act. Hence, intentional acts are a subclass of controlled acts. Other subclasses are acts that are controlled in terms of the goals to alter, stop, or avoid (see Figure 1, Frame 1).

Many contemporary investigators adopt the definitions of intentional and controllable proposed by Bargh (1994). These definitions are slightly different from the ones presented so far. In Bargh's account, intentionality has to do with whether one is in control over the start of a process, whereas controllability has to do with one's ability to alter or to stop a process once started. Both features are portrayed as involving different kinds of processing goals: the feature intentional has to do with the goal to start a process, whereas the feature controllable has to do with the goals to alter or to stop a process. These goals differ with regard to the part of the time course of the process or, more generally, of the act: The goal to start the act covers the first phase, whereas the goals to alter or stop the act cover the later phases. But there is more: The goal to start an act is also a promoting (positive) goal, whereas the goals to alter or stop an act are counteracting (negative) goals.⁸ The difference between our account and Bargh's is that we apply the terms *controlled* and *controllable* in relation to any goal about the act (engage in, start, continue, complete, alter, stop, or avoid) irrespective of its sign (positive or negative) or phase (first or later). Our use of the term *intentional* or *intentionally produced*, on the other hand, is restricted to producing or promoting goals (to engage in), and in this way intentional acts form a subclass of controlled acts. Unlike Bargh (1994), we do not restrict our use of the term *intentional* to the start of an act. A person who has an intention to engage in an act often does not merely want to initiate the act but also wants to continue and complete it. For example, when a person intends to climb the stairs, he or she actually wants to reach the next floor and not be stuck at the first step. This being said, it is not unthinkable that a person would merely intend to initiate an act without also intending to continue and complete it. But in that case, we would prefer labeling the act that results from this intention an intentionally initiated act rather than an intentional or intentionally produced act. Persons may even intend to continue or complete an act that was not intentionally started (e.g., bending down to tie one's shoes, one notices a coin one has dropped and retrieves it). Such an act would be termed an intentionally continued or intentionally completed act.

In summary, an act is controlled in terms of some goal (engage in, alter, stop, or avoid the act) when (a) the goal is present, (b) the

effect (occurrence, change, interruption, or prevention of the act) is present, and (c) the goal is the cause of the effect. The term *uncontrolled* expresses the negation of the term *controlled*. An act can be classified as uncontrolled in terms of some goal when the goal is absent, in which case the effect may also be absent (Type A1) or the effect may be present because of another cause (Type A2). An act can also be classified as uncontrolled when the goal is present, but the effect is absent (Type B1), or the effect is present but the goal is not the cause of the effect (Type B2; see table in Figure 1). According to our definitions, unintentional acts are uncontrolled in terms of the goal to engage in the act and therefore count as a subclass of uncontrolled acts. More precisely, unintentional acts of Type A (i.e., the intention is absent, but the act is present) are a subclass of uncontrolled acts of Type A2, and unintentional acts of Type B (i.e., the intention and the act are present, but the intention is not causal) are a subclass of uncontrolled acts of Type B2.

The basic definition of control as the achievement of a goal that a person pursues with regard to some target corresponds to the type of control that engineers attribute to open-loop control systems (see also Wegner & Bargh, 1998). In open-loop control systems, there is a criterion (also called standard or goal) and an action that is directed at attaining that criterion. When the criterion is reached, control is a fact. This is true irrespective of whether the person or a machine has received feedback about whether the criterion is reached. Thus, a person can have control without having an accompanying experience of control (because relevant feedback is not available). Open-loop control systems are contrasted with closed-loop systems. In a closed-loop system, the result of the action is fed back into the system (hence closing the loop), and a comparison module tests whether the criterion is reached. When the criterion is reached, the action ceases; when the criterion is not reached, the action is repeated until the criterion finally is reached (Miller, Galanter, & Pribram, 1960).

Powers (1973) included yet another ingredient in his definition of control: disturbance. A controls B if A influences B in a desired direction, thereby counteracting the effect of a disturbing influence on B. For example, a sailor controls his boat when he manages to steer it in the right direction despite the disturbing influence of the wind. Keeping in mind that there may be other, more complex forms of control, we limit our discussion to the simplest form, defined as desired causation or causation by a goal (Dennett, 1984). In the remainder of the text, we speak of processes instead of acts.

Autonomous. The next feature to delineate from previous ones is the feature autonomous. The word *autonomous* (deriving from the Greek words *auto* [self] and *nomos* [law]) literally means self-governed or not controlled by outside forces. Thus, the feature autonomous overlaps with the feature uncontrolled. The only difference is that the feature uncontrolled must be specified in terms of a processing goal, whereas the feature autonomous means uncontrolled in terms of every possible processing goal. A com-

⁸ Although at first glance there seems to be no overlap between the features intentional and controllable in Bargh's (1994) definitions, a closer look reveals that both features are presented as implicating control. If one feature (intentional) is defined in terms of another feature (control), then their overlap is a fact.

pletely autonomous process is not produced, altered, stopped, or avoided by a processing goal. More precisely, the feature autonomous is a combination of the features unintentional (see Figure 1, Frames 2 and 3) and uncontrolled in the sense of alter, stop, or avoid, but only of Types A1 or B1 (see Figure 1, Frames 4 and 5).

We need to mention that the definitions of autonomous given by other authors are somewhat different from our own. For example, Bargh's (1992) definition of an autonomous process is one that can run to completion without conscious monitoring after it was started (irrespective of whether it was started intentionally or unintentionally). This differs from our definition in two ways. First, the goal that is invoked in Bargh's definition is a conscious goal. In our definition, the term *goal* refers to both conscious and unconscious goals. Second, Bargh's definition of autonomous pertains only to the later phase of the process (after its initiation) and corresponds to what we would call *partial autonomy*. Bargh's definition of autonomous corresponds most closely to Logan and Cowan's (1984) definition of ballistic. They defined a ballistic process as one that cannot be inhibited once it begins and so must run to completion. It may be noted, however, that Bargh's definition of autonomy is formulated in terms of the goals to continue and complete the process, whereas Logan and Cowan's definition of ballistic is formulated in terms of the goal to stop the process. It may not be impossible to stop a process that does not require a goal for its continuation and completion. It may also be noted that Zbrodoff and Logan (1986) did not identify the feature ballistic with autonomous. They defined a process as completely autonomous (a) if it begins without intention, (b) if it runs to completion without intention, and (c) if it runs to completion despite one's best efforts to stop it (i.e., ballistic). This definition of autonomous seems closest to our own.

Issues of causation. The causal connection between goals and effects (the occurrence, interruption, or change of a process) occupies a key position in our definitions of the features intentional and controlled or controllable. We therefore consider what it means to speak of causation. Whereas attempts to analyze causation in noncausal terms are necessarily reductionist and perhaps not entirely satisfactory (Tooley, 1990), attempts to analyze causation in causal terms tend to be circular (Sosa, 1993). One proposal within the reductionist approach is to define causation in terms of necessary and sufficient conditions. The best-known example is Mackie's (1965, 1974) characterization of a cause *C* of an effect *E* as an *insufficient* but *necessary* part of a set of conditions that is itself *unnecessary* but *sufficient* for the occurrence of the effect (in short, an INUS condition, from the initial letters of the words italicized above). For instance, when experts declare that a short circuit caused a fire in a house, they are saying that the short circuit, although insufficient to start the fire by itself, was a necessary part of a set of conditions (including requisite oxygen and combustible material). Although this set of conditions was unnecessary to start the fire (the fire may have been started by a different set of conditions such as one that includes dropping a lit cigarette instead of a short circuit), it was sufficient to start it.⁹ We think that *C* (the short circuit) can be said to be only conditionally necessary for *E*, that is, only within this particular set of conditions. This means that the omission of *C* from this set does not lead to *E*. *C* cannot be unconditionally necessary, however, because there may be other sets in which *C* is also absent that still lead to *E* (e.g., a set in which the short circuit is replaced by dropping a lit

cigarette). More formally, let *X* stand for the remaining conditions in the set so that *CX* is a sufficient set for *E*. Saying that *C* is necessary within set *CX* for *E* (or *C* is necessary for *E* given that the remaining conditions are *X*) means that *X* alone is not sufficient for *E*. However, *C* is not unconditionally necessary for *E* because there may be other sets that also lack *C* (e.g., *KX*, *LY*) but that do suffice for *E*.

We may now reconsider our definitions of the features (un)intentional, (un)controlled, and (un)controllable in light of this analysis of cause as an INUS condition. A process (i.e., the effect) is intentional when the goal to engage in it is a necessary part of the sufficient set of conditions for the process. Another necessary part of this set may be the stimulus input. In contrast, a process is unintentional when the goal to engage in the process is not a necessary part of the sufficient set of conditions for the process. This may be because the goal is absent (Type A) or the goal is present but it is not necessary (i.e., redundant) within the set (Type B).

More generally, a process is controlled in some sense (engage in, alter, stop, or avoid the process) when the goal (to engage in, alter, stop, or avoid the process) is a necessary part of the sufficient set of conditions for the effect (occurrence, change, interruption, or prevention of the process). In contrast, a process is uncontrolled in some sense when the goal is not a necessary part of the sufficient set of conditions for the effect. This may be because the goal is absent, in which case the effect is also absent (Type A1) or the effect is present but because of another cause (Type A2). This may also be because the goal is present, but the effect is absent (Type B1), or the effect is present but not because of the goal (Type B2). In Type B1, the set of conditions including the goal is not sufficient for the effect; in Type B2, the set of conditions for the effect does not include the goal as a necessary part. The features controlled and uncontrolled say something about the current state of a process. The feature controllable expresses the potential for control, without specifying when and where or, more precisely, in which particular set insertion of a goal will turn the set into a sufficient set for the effect. The feature uncontrollable, on the other hand, specifies that no such set exists.

Purely stimulus driven. At this point, we can delineate the feature unintentional from the feature "purely stimulus driven." When theorists say that a process is purely stimulus driven or produced by the mere presence of the stimulus, they usually do not mean that the stimulus is a uniquely sufficient condition for the

⁹ This concept of causation is complemented by a recommendation for how to select "the" cause among the set of conditions that are all equally necessary to produce the effect. When comparing the occasion in which the effect occurred with another occasion in which the effect did not occur, there is often one condition that makes the difference or fills in the blank. For example, when the time in which the fire occurred is compared with a time in which the fire did not occur, the short circuit may be identified as the condition that was different and therefore be designated as the cause. An important implication is that selection of the cause depends on the occasion one chooses as a standard for comparison, and this choice, in turn, is for the most part driven by interest (e.g., Mill, 1843; for other factors, see McGill & Tenbrunsel, 2000). Given the arbitrary nature of this choice, our focus is on the set of necessary conditions that are jointly sufficient to produce an effect rather than on discriminating the triggering cause from the background conditions.

process. Indeed, there are a host of other necessary conditions that are always presumed to be in place (e.g., when the stimulus is visually presented, the light must be switched on, the person's eyes must be open, and he or she must be awake). What the theorists do mean is that the usual suspects, such as processing goals, other processes or their outputs, consciousness, and attention are not necessary parts of the sufficient set of conditions for the process. We argue that purely stimulus-driven processes are a subclass of unintentional processes. All purely stimulus-driven processes are unintentional, but not all unintentional processes are purely stimulus driven. An unintentional process may require conditions other than the presence of a stimulus to be in place, such as more remote goals (i.e., unintentional goal dependent), other processes or their output, consciousness, and attention.

Links Among Goal-Related Features

The foregoing analysis reveals that goal-related features can, to a certain extent, be conceptualized as nonoverlapping. Our definitions have partial overlap among several sets of goal-related features. First, intentional processes are a subclass of goal-dependent processes and a subclass of controlled processes. That is, intentional processes are dependent on the goal to engage in the process, and they are controlled in terms of the goal to engage in the process (see Figure 1, Frame 1). Second, unintentional processes and autonomous processes are both subclasses of uncontrolled processes. That is, unintentional processes are uncontrolled in terms of the goal to engage in the process (see Figure 1, Frames 2 and 3), and autonomous processes are uncontrolled in terms of every possible goal (see Figure 1, Frames 2–5). Autonomous processes are thus also a subclass of unintentional processes (see Figure 1, Frames 2 and 3). Third, both goal-independent and purely stimulus-driven processes are a subclass of unintentional processes. Unintentional processes are independent of the proximal goal to engage in the process, whereas goal-independent processes are independent of any goal (proximal and remote). Purely stimulus-driven processes are independent not only of (proximal and remote) goals but also of other conditions such as awareness and attention. Purely stimulus-driven processes are also a subclass of goal-independent processes (see Figure 1, Frames 2 and 3). Despite these assumptions of partial overlap, we avoid overlap among different types of control. For example, it is possible that a person controls the occurrence of a process but not the interruption of it. Precisely because we separate the different goals that are involved in types of (un)controlled processes, we believe that our set of definitions is consistent with a decompositional approach.

Measurement

Besides being helpful for the delineation of features, the conceptualization of a cause as an INUS condition suggests directions for the investigation of these features. First, it must be noted that investigators seek to formulate conclusions that go beyond individual instances of processes: They want to discover general patterns. For example, they try to examine whether a certain *type* of process is intentional or unintentional. Second, the purpose of an investigation is usually to demonstrate that a process *can* be intentional or unintentional, not that it always is. This purpose

squares with the fact that an INUS condition is presented as conditionally necessary, but not as unconditionally necessary, for the effect. When an intention is necessary within one sufficient set of conditions for a process, it does not follow that an intention is also necessary within other sufficient sets. In other words, a process may be intentional in some circumstances but unintentional in others.

When the purpose is to show that a process can be intentional, the investigator should try to find a set of conditions that is sufficient for the occurrence of the process and show that the goal of engaging in it is a necessary element in this set. In contrast, when the purpose is to show that a process can be unintentional, an investigator should find one set of conditions that is sufficient for the occurrence of the process and show that the goal of engaging in it is absent from the set or, if the goal is present, that it is not necessary.

Demonstrating the sufficiency of a set of conditions is easy: When the process has occurred, the set must have been sufficient. Investigating whether a goal is necessary within a set amounts to comparing two sets of conditions, one in which the goal is present (CX) and another in which the goal is absent (X), keeping other conditions as much alike as possible. The goal is unnecessary when the process occurs in both sets (CX and X), but necessary when the process obtains only from the set in which the goal is included (CX).

If intentionality involves not just the mere presence of an intention and a process, but also a causal connection between these elements, then the evidence in favor of intentionality should comprise something over and above mere presence. Treating a cause as an INUS condition is one way to deliver something over and above the mere presence of the relata: Comparing the data of two sets (X and CX) is more than simply considering the data of only one set (CX or X).

When a goal is necessary in one set, it does not follow that it is also necessary in another. Formally, C may be necessary in set CX but unnecessary in set CKX, so that KX is sufficient for E as well. In cases such as this, an important next step in research would be to pinpoint some of the remaining conditions (K) that modulate the necessity of C for E. It is a step forward to demonstrate that a process can be intentional or unintentional; it is a further step forward to identify the circumstances under which a process is intentional and those under which it is unintentional.

A similar procedure must be followed when the purpose of an investigation is to demonstrate that a process can be controlled or is controllable in terms of the goals to alter, stop, or avoid the process. For example, to demonstrate that a process is controllable in terms of the goal to stop the process, the investigator should try to find a set of conditions that is sufficient to stop the process and then show that the goal of stopping it is a necessary element within this set. To demonstrate the reverse, that a process can be uncontrolled in terms of the goal to stop it, requires the investigator to find a set of conditions in which the goal of stopping the process is included and then show that the set is not sufficient for stopping it (Type B1) or, if the set is sufficient to stop the process, that the goal is not a necessary element in the set (Type B2). Finally, to demonstrate that a process cannot be controlled or is uncontrollable in terms of some processing goal, investigators must show that no set will ever be found in which inclusion of the goal makes the set sufficient for the effect. Demonstrating this is not a realistic

purpose of an investigation. It seems that one can only provisionally maintain that a process is uncontrollable until other investigators come up with evidence that rebuts the claim.

Studies demonstrating the simple presence or absence of a goal are undertakings that present their own difficulties. We first discuss problems related to demonstrating the presence of a goal and then problems related to demonstrating the absence of a goal. Establishing the presence of a goal is often operationalized by giving explicit instructions to engage in or to suppress some process. Suppose that the instructions actually lead to the occurrence of a process. One might then wonder whether there is any reason to postulate the intermediate step of goal activation, because the instructions might instead have triggered the process via nonmotivational representations. Some useful guidelines for testing the mediating influence of goals may be found in articles by Bargh and Barndollar (1996) and Bargh et al. (2001). These investigators proposed two dynamic qualities by which goal representations are different from nonmotivational representations. First, when a goal is the source of some activation, this activation should persist even in the face of obstacles. Second, goal activation may accrue over time, whereas activation of purely cognitive representations should decrease over time. Bargh et al. (2001) primed one group of participants with achievement-related words and another group with neutral words in an initial task. They did this to activate an achievement goal in the former group but not in the latter one. Bargh et al. found that achievement-primed participants were more determined to work on a second, ostensibly unrelated task in the face of obstacles than were the neutral-primed participants (i.e., evident of persistence). They also found that performance of the achievement-primed participants was better after a delay than when there was no delay (i.e., evident of accumulation). This led the authors to conclude that they had successfully primed a goal instead of some nonmotivational schema. Their studies offer useful suggestions for how to test for the presence of processing goals: The presence of goals to initiate, continue, or alter some process can be investigated by testing for persistence and accumulation on a subsequent task.

Establishing the absence of a goal to engage in a process is often operationalized by the use of indirect tasks, in which participants are instructed to engage in a process other than the critical one. In addition, investigators carefully omit from the instructions any element that might induce the goal of engaging in the critical process. Even if their attempts in this regard are successful, one cannot exclude the possibility that other characteristics of the experimental setting induce the goal. For example, in recent affective priming studies, investigators have replaced the original task of evaluating the targets (Fazio, Sanbonmatsu, Powell, & Kardes, 1986) with a nonevaluative task such as a lexical decision or a pronunciation task (e.g., Bargh, Chaiken, Raymond, & Hymes, 1996; Hermans, De Houwer, & Eelen, 1994). When a nonevaluative target task is used, participants are not encouraged to focus on the valence of stimuli. A priming effect under these circumstances has led investigators to conclude that the process of evaluating the primes can be unintentional. Note, however, that the activation of an evaluation goal was not prevented in these studies; it was only discouraged. It is still possible that the polarized valence of the stimuli induced participants to pursue an evaluation goal and that this goal was necessary to produce the priming effect in this set of conditions. This objection is weakened

by a pronunciation study by Bargh et al. (1996) in which affective priming effects were obtained despite the fact that only targets with moderate valence were presented. Still, it is possible that the participants pursued an evaluation goal by default. Removing processing goals from a set of conditions is even more complicated if one considers that goals are not always consciously activated but may be unconscious. Eliminating the conscious goal to engage in some process is one thing; eliminating an unconscious goal is yet another.

In the section *Issues of Causation*, we stated that treating a cause as an INUS condition is a reduction: Showing that a goal is an INUS condition may be indicative of the causal role of the goal, but it is far from being proof. Nevertheless, we accepted this reduction because it delivers something over and above the mere presence or absence of goals and processes (or other effects). Showing that a goal is not the cause of an effect does not suffer from this reduction because when a goal is shown to be absent, a fortiori it cannot have caused the effect. In summary, it may be far more difficult to demonstrate that processes are intentional than to show the reverse, that they are unintentional. We make a further note about the factors that we have previously cited as involved in the implementation of an intention (opportunity, skill). Knowledge about the status of these factors may guide researchers in setting up their experiments and may strengthen the arguments that they adduce in favor of their conclusions about goal-related features.

The dominant view seems to regard nonautomatic processing as the default and puts the burden of proof entirely on models that invoke automatic mechanisms (Bargh, 1997, p. 4). Although several authors have expressed dissatisfaction with this view, most automaticity research is guided by the assumption that processes must be produced by intentions until it is demonstrated that they can be produced otherwise. The emphasis is on showing that a processing goal is not a necessary part of the set of conditions that suffices for the process rather than on showing the reverse, that it is a necessary part. But the tide may turn, as some theorists have come to advocate a view that is exactly opposite to the dominant view sketched above. These theorists, inspired by a dynamic systems framework, argue that much behavior (even complex behavior) emerges from bottom-up activity that satisfies multiple constraints without the need for top-down intentions (e.g., Kelso, 1995). Similarly, advocates of a connectionist framework argue that subsymbolic patterns of activation give rise to both acts and intentions. According to this view, humans' phenomenal experience of intentions being the causal motor of acts is nothing but an illusion based on post hoc causal attributions (see also Wegner, 2002, 2003; Wegner & Wheatley, 1999; but see Green & Vervaeke, 1996; Menzies, 2004). Finally, it seems that different purposes of investigation can arise depending on the nature of the goal under study. With regard to positive processing goals, automaticity researchers have, for the most part, been trying to show that processes can be produced without them (cf. affective priming studies). With regard to counteracting goals, researchers have been trying to show that these goals can be unsuccessful (cf. thought suppression studies; Wegner, 1994) and that they can be successful in certain cases (cf. stereotype priming studies; Blair, 2002; Moskowitz, Gollwitzer, Wasel, & Schaal, 1999).

Diagnostic Status

According to the conditional view of automaticity (Bargh, 1992), the fewer conditions that are needed, the stronger the automatic nature of the process. In such a view, the features purely stimulus driven and completely autonomous are likely to be considered strong features of automaticity. The feature partially autonomous (what Bargh called simply “autonomous”) is the only necessary feature of all automatic processes. At the same time, however, the feature partially autonomous is only a minimal feature of automaticity. Processes that possess no other features than partial autonomy are considered to be automatic, but only weakly so.

(Un)conscious

The next features to be discussed in our list are the features conscious (or aware) and unconscious (or unaware). According to some theorists, the conscious–unconscious distinction reflects the presence or absence of an isolated feature (e.g., Neisser, 1976). Other theorists have attached two separate modes of information processing to this distinction. We argue that such information-processing models continue to be an important breeding ground for the dual-mode view of automaticity. Before we can elucidate this argument, we need to give a conceptual account of what it is that these information-processing models try to explain. As described in the next sections, the terms *conscious* and *unconscious* are not unitary concepts (Block, 1995; Cleeremans & Jiménez, 2002).

Concepts

Conscious. Traditionally, theorists of consciousness have emphasized the distinction between two aspects of consciousness: a phenomenal aspect and an intentional one. The phenomenal aspect refers to the qualitative (Natsoulas, 1981), the experiential (Block, 1995), or the nonrepresentational property (Eilan, 1998) of awareness, also known as “immediate qualia” (Tolman, 1922). This aspect of consciousness captures, for example, the difference between seeing red and green or what it is like to experience colors at all (Nagel, 1994). The phenomenal aspect has primarily been ascribed to perceptual experiences and bodily sensations (e.g., pain) but not exclusively so (e.g., internal images may also have a phenomenal quality). It is the phenomenal aspect that has often been mentioned as the ultimate difference between the human brain on the one hand and an entirely functionally equivalent computer on the other (at least by theorists who contest a purely functional account of the phenomenal quality of experience; cf. Chalmers, 1995; for a functionalist account, see Dennett, 1991). Explaining how a physical system can produce phenomenal states has been dubbed the hard problem of consciousness (Chalmers, 1996), the explanatory gap (Huxley, 1866, as cited in Block, 1995), or the mystery of consciousness (Dennett, 1991).

The Intentional aspect of conscious states refers to their directedness at some state of affairs outside of themselves, also called the *aboutness* aspect (see above). An organism can be directed at something by forming a representation of it. Therefore, the Intentional aspect has also been termed the *representational* property (Eilan, 1998) of awareness. (Note that some theorists have claimed that Intentionality is what separates humans from functionally equivalent computers, e.g., Searle, 1992.)

Different positions have been defended with respect to the relation between the phenomenal and the Intentional aspects of consciousness. Certain authors hold that phenomenality requires Intentionality to be classified as consciousness, as in Kant’s (1787/1929) dictum that sensations without concepts are blind. Intentionality without phenomenality, on the other hand, corresponds to abstract thoughts (Natsoulas, 1981). Opponents of this view argue that Intentionality must be injected with phenomenality to make sense. For example, the abstract thought of a circle remains meaningless without the sensory experience of a circle. Other authors argue that the phenomenal and the Intentional aspects are interwoven in a mutually dependent manner (Eilan, 1998).

Still other authors argue that the two aspects are, in principle and sometimes even de facto, separable. For example, Block (1995) proposed a distinction between two types of consciousness that are themselves founded on the distinction between the phenomenal and the Intentional aspect: phenomenal consciousness (*P-consciousness*) and access consciousness (*A-consciousness*). P-conscious states have a phenomenal content; A-conscious states have a content that is accessible for use in reasoning and in the rational control of behavior and/or verbal report. Contents are accessible by virtue of their representational or Intentional format, but they are accessed when attention is directed at them as well. Therefore, A-conscious content coincides with those aspects that are in the focus of attention. In Block’s view, contents may be P-conscious without being A-conscious. For example, the noise of a pneumatic drill outside in the street may be P-conscious but become A-conscious only when attention is directed to it. Turning it around, contents may be A-conscious without being P-conscious. Block describes the nonempirical cases of philosophical zombies and persons with superblindsight to illustrate this. Both are able to use and report contents without having an accompanying feeling of these contents. Opponents of this view have argued (a) that P-consciousness is a prerequisite for A-consciousness (e.g., Schacter, 1989), (b) that A-consciousness is a prerequisite for P-consciousness (i.e., epiphenomenalism), or (c) that the distinction between the two is artificial (i.e., the collapse hypothesis; cf. Baars, 1995).

Block (1995) proposed two additional types of consciousness: monitoring consciousness (*M-consciousness*) and self-consciousness (*S-consciousness*). One has M-consciousness when one is conscious that one has P-consciousness or A-consciousness. S-consciousness refers to awareness of or thoughts about oneself. Most authors have distinguished only between a first- and second-order consciousness (e.g., O’Shaughnessy, 1991; Wegner & Bargh, 1998), of which the first corresponds to some combination of P-consciousness and A-consciousness and the second corresponds to M-consciousness.

Unconscious. The term *unconscious* is not unitary either. Several theorists have distinguished between unconscious mental entities that are structurally inaccessible and those that are generally unconscious but potentially accessible to consciousness. According to some scholars, the accessible unconscious requires only the allocation of attention to become conscious (e.g., Baars, 1998). Although the distinction between levels of unconscious is generally acknowledged, some authors use the term *unconscious* exclusively for the accessible variant (e.g., Searle, 1992, posits that the inaccessible unconscious is not mental), whereas others reserve it for the inaccessible variant (e.g., Power & Brewin, 1991).

There are several ways to divide the concept unconscious into an accessible and an inaccessible part. Some scholars have distinguished between processes that are so weak that they remain hidden to consciousness and those that are so well practiced that they have dropped out of consciousness (Carver & Scheier, 2002; Cleeremans & Jiménez, 2002; James, 1890; Moskowitz, 2001). The evolution of the former (inaccessible) type into the latter (accessible) type has been described as follows. When emergent bottom-up processes are newly organizing, they are so weak that they are inaccessible to consciousness. With repetition, they become consolidated in memory such that they can be recognized by a conscious processor, which may invoke them as a guide to behavior. With increasing top-down use of these processes, they become so reliable that top-down control of them is no longer required, and as a result, they drop out of consciousness again.

Other scholars have equated the accessible–inaccessible distinction with the contents–processes distinction (e.g., Gillett, 1996). Many have argued that only the contents (inputs and outputs) of processes may become consciously accessible and not the processes themselves. Not all authors subscribe to this position, however (e.g., Bargh, 1994). One way to solve the tension between the two groups might be to distinguish between different levels of process understanding as posited by Marr (1982): the computational level, the algorithmic level,¹⁰ and the hardware level. The computational level articulates the functional relation between input and output, whereas the algorithmic level contains information about the formal properties of the processes involved in transforming input into output (i.e., what is actually in the black box). The hardware level is concerned with the physical implementation of processes in the brain. Whereas both the algorithmic and the hardware levels can be confidently classified as structurally inaccessible, the computational level may be accessible to consciousness. All that is required for a person to be conscious of the computational level is that he or she is able to link an output to a given input. People may, for instance, be aware that they evaluated a stimulus (i.e., input) as good or bad (i.e., output) without being aware of the formal characteristics of this process or of the neural circuits that were involved.

The foregoing discussion hints that the term *unconscious* can be used as a predicate to different things. The feature unconscious may apply (a) to the stimulus input that evokes a process, (b) to the output of a process, (c) to the process or its constituent steps, or (d) to the consequences of a process such as its influence on subsequent processing (see also Bargh, 1994). It is advisable to make explicit what one considers to be unconscious when one uses the term. Whereas the consequences of a process do not, strictly speaking, belong to the critical process but rather to another, subsequent processing unit, the elements of input, process, and output all relate to the critical process, and special care must be given to disentangling them. Of course, one may also specify the elements of input, output, or process for the feature conscious.

Distinction between conscious and unconscious. Now that we have analyzed the concepts of conscious and unconscious, we turn to theories that concern the distinction between them. These theories differ with regard to a number of issues. We highlight five of these below.

1. Theories embrace different metatheoretical assumptions about the proportion of mental life that is occupied by conscious and unconscious cognition. Cleeremans and Jiménez (2002) have

discussed two extreme positions. Zombie models (e.g., Lewicki, 1986; Reber, 1993) assume that cognition is largely unconscious (as in zombies) and that only a part of it, the metaphorical tip of the iceberg, is optionally available to consciousness. At the other extreme, Commander Data theories (e.g., Shanks & St. John, 1994) claim that all information processing is (at least potentially) available for introspection. Commander Data is a Star Trek android that has conscious access to all of its internal bodily and cognitive processes. Commander Data theories deny the unconscious and dismiss empirical evidence in favor of unconscious processing as artifactual. Commander Data theories usually accept the existence of unconscious neural processes, but they do not consider these to be mental (e.g., Dulany, 1997; Perruchet & Vinter, 2002).

2. A related issue is whether the unconscious is smart or dumb. Loftus and Klinger (1992) identified two factors that indicate a smart mental process: complexity and flexibility. Complexity is the sophistication of the stimulus material that can be handled (line patterns, letters, words, sentences) and the depth of analysis that can be performed (sensory, perceptual, semantic, propositional). Greenwald (1992), argued that the unconscious is not very smart on the basis of the observation that subliminally presented input yields only partial meanings of single words (Abrams & Greenwald, 2000). Bargh and Barndollar (1996) criticized the fact that Greenwald (1992) equated unconscious processes with unconscious processes that operate on unconscious input and noted that many other unconscious processes require conscious input to operate: “To assess the intelligence of the unconscious by seeing how it does in the absence of informational input (i.e., by withholding from it any attentionally supplied information) is like taking a fish out of water and concluding from the fact that it just lies there that it is pretty stupid.” (Bargh & Barndollar, 1996, p. 461; see also Erdelyi, 1992). Support for complex unconscious processes operating on conscious input is provided by implicit learning studies showing that complex relations among consciously perceived stimuli or attributes can be discovered unconsciously (e.g., Lewicki, Czyzewska, & Hoffman, 1987; but see Hendrickx, De Houwer, Baeyens, Eelen, & Van Avermaet, 1997). In contrast to Greenwald (1992), Lewicki (1992) proposed that the unconscious system might be far smarter than the conscious one. It is important to understand, however, that the two investigators envisioned a different type of unconscious process (one that operates on unconscious input vs. one that operates on conscious input).

With regard to flexibility, a process is flexible when it can be applied to novel stimulus material (the same algorithm operating on new data), but it is even more flexible when it can adapt its own form to the demands of the situation (the algorithm changes itself). The currently dominant view is that unconscious processes cannot handle new material and that they are applied in essentially the same way time after time (e.g., Loftus & Klinger, 1992; Neumann, 1984).

¹⁰ Marr (1982) used the term *algorithmic* in a different sense than we did in the previous section. Both the process of single-step memory retrieval and the process of algorithm computation can be situated at Marr’s algorithmic level.

The cleverness of the unconscious (Issue 2) and the proportion of cognition occupied by it (Issue 1) are related issues. Zombie theorists, who consider most mental life to be unconscious, must assume the unconscious to be fairly smart to account for the richness of human cognition. By contrast, Commander Data theorists do not need a smart unconscious, because they assume that most processing is conscious.

3. A third issue is how theories construe the transition from unconscious to conscious cognition. A. P. Atkinson, Thomas, and Cleeremans (2000) have proposed a taxonomy for categorizing theories based on the kinds of factors they hold responsible for consciousness. Theories can be classified according to two independent dimensions: a specialized versus nonspecialized dimension and a representation versus process dimension. Specialized theories postulate a special system that is dedicated to consciousness, such as a short-term memory (e.g., R. C. Atkinson & Shiffrin, 1971) or a central executive (Norman & Shallice, 1986). Nonspecialized theories lack a special consciousness system and explain consciousness exclusively in terms of particular representations or processes. Representation or vehicle theories characterize consciousness in terms of intrinsic properties of representations or nodes, such as strength of activation, stability, and distinctiveness. For example, some models posit that representations become conscious when they receive activation above a certain threshold (e.g., R. C. Atkinson & Shiffrin, 1971; Shallice, 1978; Wundt, 1896/1897). Process theories characterize consciousness in terms of the specific processes or computations operating over representations (e.g., Schacter, 1989). When both dimensions are combined, four broad categories result (for a detailed discussion, see A. P. Atkinson et al., 2000).

4. A fourth issue is whether conscious and unconscious cognition are alike or are qualitatively different. According to Cleeremans and Jiménez (2002), zombie models assume that the unconscious system mirrors the conscious one in all respects except for the feature conscious (unconscious system = conscious system minus consciousness). Marcel (1983a, 1983b) likewise pointed out that much of the extant theorizing rests on the identity assumption. According to this assumption, unconscious representations are the same ones that constitute conscious experience. Not all theorists agree that the conscious system is a blueprint of the unconscious system. For example, Marcel regarded conscious perception as the result of a constructive process in which unconscious hypotheses are matched against unconscious records of sensory input. Conscious perception imposes a particular structure on otherwise unstructured unconscious percepts. A qualitative distinction between conscious and unconscious cognition is also present in Wundt's (1896/1897) conception of consciousness as an interconnection of mental processes that reaches beyond the individual elements, as well as in Dennett's (1991) metaphor of consciousness as a virtual serial computer running on a parallel computer. Some empirical data suggest a qualitative distinction between conscious and unconscious processing (e.g., Marcel, 1980; review by Merikle & Daneman, 1998), whereas other data reflect merely quantitative differences, suggesting that unconscious cognition is a weak form of conscious cognition. Reingold (1992) argued that the conscious-unconscious distinction is worth agonizing about only when conscious and unconscious cognition have qualitatively different consequences.

We believe that this fourth issue is related to the third one. Theories assuming that the transition from unconscious to conscious cognition is a mere function of the activation strength of representations (i.e., vehicle models) are likely to see the distinction between conscious and unconscious cognition as a matter of quantity. On the other hand, Cleeremans (2005) has argued that gradual changes in quantity can sometimes produce qualitative differences, as when water turns into ice as a result of a continuous change in temperature. However, such qualitative differences can be considered "accidents" rather than active constructions. Models that include a specialized consciousness system are more likely to ascribe a different quality to conscious and unconscious cognition, except when the consciousness system is merely conceived of as the director of an attention spotlight.

5. The final issue concerns differences arising from the underlying metaphor of cognition (computational vs. connectionist) that is adopted. Classic, computational models conceive of information processing as symbol manipulation performed by some central executive or external processor (Newell & Simon, 1972). The processor takes symbolic representations of data (concepts or exemplars) from one database as its input, runs a symbolic program or process (algorithm or rule) from another database on this input, and produces new symbols as its output. In most models, the processor is also charged with interpreting the output. Eliminating the processor in such a model amounts to halting processing and leaving symbols meaningless. Because unconscious cognition is defined by the absence of external access, it follows that computational models cannot account for unconscious processing and unconscious meaning. Most zombie models and some Commander Data models rely on the computational framework. Zombie models offer no solution to the problem that the computational framework is incompatible with unconscious meaning and processing. Commander Data models elude the incompatibility problem because they deny the unconscious (Cleeremans, 1997).

The connectionist framework possesses several features by which it is thought to provide a more natural account of unconscious meaning and processing than does the computational framework (Cleeremans, 1997). In connectionist models, a concept is represented subsymbolically, as a pattern of activation distributed among many nodes or units in a network. Different representations show more or less resemblance in activity patterns as a function of the semantic resemblance of their concepts. Through these resemblances, the intrinsic properties of subsymbolic representations reflect the meaning of their concepts (Clark, 1993; but see Christiansen & Chater, 1992). Symbols, in contrast, have no intrinsic relation to what they mean because they are abstract; they require an external interpreter to inject them with meaning (i.e., the symbol grounding problem; Searle, 1992). In connectionist models, rules are also represented subsymbolically, specifically, in the connection weights among units. The system behaves in a rulelike manner without possessing the rule in some symbolic form (Cleeremans, 1997). This is illustrated by the fact that people can conform to grammatical rules without being able to report them. Processing in the connectionist framework does not amount to combining rules and data but to dynamic patterns of activation. This property, together with the assumption that activation can travel through the network without external help, renders superfluous the intervention of an external processor and fits well with the notion of unconscious processing, just as the subsymbolic

nature of the representations fits well with the notion of unconscious meaning.¹¹

Another aspect that distinguishes classic from connectionist models is modularity. Classic models tend to view conscious and unconscious cognition as two separate systems, typically represented as separate boxes in schematic figures of information-processing models (e.g., Greenwald, 1992). This aspect is related to the fact that many classic models can be categorized as specialized theories of consciousness, proposing a specialized consciousness system (e.g., R. C. Atkinson & Shiffrin, 1971; Norman & Shallice, 1986; Schacter, 1989). In contrast, connectionist models typically hold a decentralized, nonspecialized view of consciousness. That is, there is no specific location where consciousness is produced. It is primarily the quality of the nodes (in vehicle theories) or the quality of the activity among the nodes (in process theories) that determines their fate in terms of consciousness.

In summary, we have discussed theories that are concerned with the distinction between conscious and unconscious cognition and have contrasted them with regard to the properties they assign to each type, both in terms of territory (zombie vs. Commander Data theories) and competence (dumb vs. smart). In addition, we have classified them according to the factors they consider to be responsible for the transition from unconscious to conscious cognition (a special top-down system and/or bottom-up properties of representations or processes). A related issue we considered was whether theories assume the conscious–unconscious distinction to be a matter of quantity or quality. We then illustrated the idea that several divergences may arise from the metaphor of cognition (computational vs. connectionist) that underlies a given theory. Our main objective was to show that the theoretical choices one makes with regard to the distinction between conscious and unconscious cognition and the metaphor of cognition one adopts have some bearing on the controversy between the dual-mode and the decompositional approach to automaticity. Because specialized theories of consciousness and the computational model of cognition in general tend to view conscious and unconscious cognition as separate systems, authors have been encouraged to ascribe additional contrasting features (intentional vs. unintentional, controllable vs. uncontrollable, slow vs. fast, nonefficient vs. efficient) to the systems so that the conscious–unconscious distinction has developed into an automatic–nonautomatic distinction (e.g., Fodor, 1983). Other computational theorists have worked in the opposite direction, adding the feature pair conscious–unconscious to a dual-mode model based on other automaticity features, such as efficiency and control (e.g., Posner & Snyder, 1975b; cf. Neumann, 1984). In summary, the modularity assumption in specialized and computational models has led to clustering of features and, in this way, contributed to the dual-mode view of automaticity. Nonspecialized models that do not embrace the notion of architectural modularity, such as most connectionist models, are less likely to cluster the feature unconscious with other automaticity features (e.g., Cleeremans & Jiménez, 2002). The decompositional view thus seems to be a more natural position for these models to adopt.

Links Between Conscious and Goal-Related Features

Previously, we defined the features intentional and controlled in terms of processing goals that are achieved. So defined, these

features bear no overlap with the feature conscious, at least not if one accepts that goals may be either consciously or unconsciously activated (see Bargh, 1990a; Fiske, 1989; Freud, 1915/1984). Bargh (1990a; Bargh & Gollwitzer, 1994) introduced an automatic model that is specifically concerned with the unconscious activation of goals by relevant situational cues. He asserted that an unconscious intention is just as “intentional” as a momentary conscious goal (Bargh & Barndollar, 1996, p. 465). Several authors agree that intentionality does not require a conscious subject (e.g., Natsoulas, 1981; Searle, 1983; but see Libet, 1985). The same has been suggested for control. Logan and Cowan (1984) applied the feature controlled to all goal-directed activity, irrespective of whether it is under the command of an executive system (and hence conscious) or exerted by unconscious, subordinate systems that interpret these commands and carry them out. Moskowitz (2001; Moskowitz et al., 2004) has reviewed numerous studies that can be subsumed under the heading of unconscious control.

Overlap between the feature conscious and certain goal-related features is nevertheless pervasive in the literature on automaticity (see also Natsoulas, 1981; Uleman, 1989). When a process is called intentional or controlled, what is often meant is “consciously intentional” or “consciously controlled.” As we discuss below, this overlap may stem from the implicit notion that there is a fundamental relation between consciousness and these features. Consciousness has often been conceived of as a defining ingredient and/or a precondition for the features intentional and controlled.

In lay conceptions, consciousness has been identified as an ingredient of intentional actions (Malle & Knobe, 1997). Wegner and Bargh (1998) noted that “most people think of control in humans as conscious control, so much that the term unconscious control doesn’t seem right at all” (p. 453). Scientific writings also suffer from the overinclusive use of the words *intentional* and *controlled*. Oatley (1993) made a distinction between conscious and unconscious goals, but he classified actions as fully intentional only when there is consciousness of the goal and it has been explicitly chosen by the agent. Wegner and Bargh (1998) acknowledged a form of unconscious control, but they singled out the conscious variant as the one that matters to automaticity. Bargh (1989, 1992, 1994) explicitly referred to conscious goals in his definitions of automaticity features. He identified unintentional effects with effects that occur in the absence of explicit goals (Bargh, 1989, p. 5). His definition of controllable made reference to an act of will that is able to stop or alter the process once it has begun (Bargh, 1994, p. 16). Similarly, he defined an autonomous process as one that runs to completion without the need for conscious guidance (Bargh, 1992, p. 186).

In addition to its *defining* role for the features intentional and controlled, consciousness has also been identified as a *condition* for the feature intentional. It is often suggested that intentional processes require conscious stimulus input. Note that in the pre-

¹¹ Not many classic models are strictly computational. Most classic models incorporate associative principles for encoding and deployment in addition to symbol manipulation. This allows data and/or algorithms to be activated and fired without the help of an external manipulator (cf. Anderson, 1996) and symbols to be grounded in experience (Harnad, 1990).

vious paragraph, the word *conscious* was used as a predicate of the word *goal*, whereas in the present paragraph it is used as a predicate of the phrase *stimulus input*. Several theorists have stated that the reason for the existence of consciousness is to motivate behavior (Johnson-Laird, 1988; Shallice, 1988). This function has primarily been ascribed to A-conscious content, which was defined earlier as content that is accessible for use in reasoning and in the rational control of behavior and/or verbal report (Block, 1995). Yet, some authors also ascribe this function to P-conscious content, claiming that phenomenal experience motivates behavior in a different way than would accessible information without phenomenal experience. For example, “a being that feels an acute pain must behave differently from a being that just cognises the existence of the pain source” (Navon, 1995, p. 265; see also Marcel, 1988).

In addition to the goal to engage in a process, consciousness may set the stage for other goals, such as the goals to alter, stop, or avoid a process. Consciousness may constitute a condition for the feature controlled in terms of these goals (Cleeremans & Jiménez, 2002; Dienes & Perner, 1999; Umiltà, 1988). According to Uleman (1989), “when subjects are unaware of some process, they should have no reason to halt or change it prior to completion” (p. 427). Similarly, Bargh (1992, 1994, 1999) noted that a lack of awareness of some process precludes controlling that process. The idea that awareness of a process allows a person to control it is also inherent in the parsimony principle of consciousness expressed in James’s (1890) much-cited phrase that “consciousness deserts all processes when it can no longer be of use” (p. 496). As processes become routinized because of practice, they become so well established that they are in no danger of going astray from their path; therefore, control is no longer required. The parsimony principle says that when control of the process becomes superfluous, so does awareness of it. This principle rests on the assumption that awareness of the process enables the person to control it. Note that in this paragraph, the word *conscious* is used as a predicate of the word *process*. The present discussion applies to an ongoing process that, once it enters consciousness, enables the setting and implementation of the goals to stop, alter, continue, and complete the process. It may also apply to a process that has not yet started but that one is aware may (or expects will) happen. This awareness may be a precondition for the activation of the goal to avoid the process.

In the assumptions mentioned above in which consciousness (of input and processes) allows for the setting and implementation of goals (to engage in, alter, stop, or avoid the process), the term *goals* is again most likely used in the sense of “conscious goals.” It seems plausible that conscious knowledge (about input and process) is a precondition for the setting and implementation of certain conscious goals (to engage in and to alter, stop, or avoid the process), but it may not be a precondition for the setting and implementation of unconscious goals. We prefer not to implant the feature conscious inside the definitions of the goal-related features intentional and controlled; thus, we do not take the position that consciousness is a precondition for these goal-related features. We do believe it is plausible to assume that consciousness is a precondition for the features “consciously intentional” and “consciously controlled.”

Measurement

Before describing guidelines for the measurement of unconscious cognition, we address a few issues. We mentioned that the notion of unconscious can be divided into an accessible and an inaccessible part. This distinction is important because it makes sense to investigate the unconscious nature only of parts that are potentially accessible to consciousness. With regard to the notion of process, we argued that it is only meaningful to investigate whether the computational level of a process (Marr, 1982) can remain unconscious (that the algorithmic and the hardware level of a process can remain unconscious seems obvious). Awareness of the computational level of a process involves awareness of the relation between the input and the output of the process and hence of three elements: the input, the output, and the relation between the two. In addition, we need to be more specific about the notion of input. This notion conceals some kind of initial perceptual process operating on the raw stimulus input, and it is the output of this process that is potentially accessible to consciousness; the raw stimulus input is never accessible. Hence, saying of a process (e.g., evaluation) that it operates on unconscious input actually means that it operates on the unconscious output of an initial perceptual process that has already operated on the raw stimulus input. The term *perceptual* process remains open to interpretation until it is specified which level or aspect of the stimulus input is processed. Examples of such aspects are the mere presence of the stimulus, its physical properties, and its semantic properties.

To investigate whether a process can operate on the basis of unconscious input, one should find a case in which (a) the process is present but (b) awareness of some aspect (e.g., presence, meaning) of the input is absent. To investigate whether a process can operate and result in an output that remains unconscious requires one to find a case in which (a) the process is present but (b) awareness of the output is absent. And finally, to investigate whether the process can remain unconscious requires one to find a case in which (a) the process is present but (b) awareness of the relation between input and output is absent. To demonstrate unawareness of the relation between input and output, it is enough to demonstrate the absence of awareness of either the input or the output. This explains why the combination of a conscious process operating on unconscious input is problematic. It is unlikely that one will engage in the conscious processing of a stimulus of which one is not aware.

Problems related to the measurement of unconscious cognition are well-documented in the study of subliminal semantic activation (SSA), in which the central question is whether the process of semantic activation (activation of the meaning of stimuli) can take place when awareness of the output of this process (meaning) is absent. The presence of a process can be assessed with an indirect measure; the absence of awareness must be assessed with a direct one. Direct and indirect tasks differ only with regard to the instructions that accompany them. In direct tasks, instructions explicitly refer to the manipulated dimension, whereas in indirect tasks, no such reference is made. Typically, the indirect measure is a semantic priming task in which briefly presented primes are followed by visual masks and participants are instructed to respond to the targets. The direct measure is identical to the indirect measure except that participants are instructed to make forced-choice discriminations regarding the meaning (i.e., recognition

task) or the presence (i.e., detection task) of the primes. (The detection task is sometimes used because it is thought to be a more stringent test of awareness than is the recognition task. This idea is based on the assumption that without detection there can be no recognition.) The so-called *classic dissociation logic* states that to demonstrate SSA, the indirect measure must yield a greater-than-zero effect, whereas the direct measure must indicate no effect (i.e., an *indirect-without-direct-effect* data pattern; Holender, 1986).

Other examples of direct measures are free verbal reports and confidence ratings. Whereas forced-choice tasks for recognition and detection are assumed to yield an objective criterion of awareness, free verbal reports reflect a subjective criterion. Cheesman and Merikle (1984) have recommended replacing the objective criterion with the subjective one. On the one hand, the subjective criterion is stringent enough to be consistent with a definition of unconscious as the lack of subjective awareness (i.e., M-consciousness; Block, 1995). On the other hand, the subjective criterion is more relaxed, and therefore it is easier to demonstrate SSA in terms of this criterion.

The validity of an indirect-without-direct-effect data pattern as evidence for SSA is dependent on two conditions. The direct measure must detect conscious information (a) exhaustively and (b) exclusively (Holender, 1986). That is, it must detect (a) all consciously available information and (b) nothing but conscious information. These conditions imply that any effect produced by the direct measure must reflect all conscious information that is available but none of the unconscious information that is activated. To interpret any indirect-without-direct-effect data pattern as evidence in favor of perception without awareness, one's direct measure must be exhaustive. Otherwise, one cannot exclude the possibility that there is some conscious information that is not captured by the direct measure but that is nevertheless responsible for the greater-than-zero effect on the indirect measure. To interpret an indirect-with-direct pattern as evidence against perception without awareness, one's direct measure must be exclusive. If the direct measure also captured unconscious information, the direct measure would yield an overestimate of awareness.

Unfortunately, one cannot determine on an a priori basis whether a particular direct measure is an exhaustive and/or exclusive index of conscious knowledge. One cannot prove that the conditions of exhaustiveness and exclusiveness are satisfied; thus, one has to assume their fulfillment on the basis of plausibility arguments. Several scholars have cast doubt on the plausibility of these assumptions. With respect to exhaustiveness, Holender (1986) has argued that the (objective) direct measures used in classic subliminal priming studies (e.g., Balota, 1983; Fowler, Wolford, Slade, & Tassinari, 1981; Marcel, 1980, 1983b) are not sensitive enough to capture all conscious knowledge. Therefore, the null effect found on the direct measure in these studies does not reflect the absence of awareness but rather the inadequacy of the direct measure to detect conscious knowledge. Exhaustiveness is particularly problematic for subjective measures of awareness. With free verbal reports, there is no guarantee that participants report all of their conscious knowledge because they may use very stringent confidence criteria before reporting anything at all. Objective measures such as forced-choice discrimination tasks suffer less from this drawback because responding is obligatory (Reingold & Merikle, 1993). Even so, it remains logically and statisti-

cally impossible to prove that an observed null effect reflects a true absence of awareness, for null effects are always open to the alternative interpretation that the direct measure was not completely exhaustive (Reingold & Merikle, 1988). No matter how sensitive the direct measure, there is no way to determine whether the absence of a direct effect is due to the stimulus being unavailable to consciousness at the time of presentation or due to forgetting during the interval between presentation and response (Erdelyi, 1992; Shiffrin, 1988).

With respect to exclusiveness, Reingold and Merikle (1988, 1993) have argued that investigators have no way to ascertain whether a direct task is process pure. Free verbal report is probably a more exclusive measure of awareness than are discrimination tasks because discrimination tasks "tend to involve some contextual cueing and hence are not free from the influence of potential unconscious determinants" (Jiménez, Méndez, & Cleeremans, 1996, p. 948). Similarly, Reingold and Merikle (1993) have suggested that discrimination tasks probably rely on both conscious and unconscious knowledge.

Reingold and Merikle (1988) argued that the exhaustiveness and exclusiveness assumptions are unverifiable and implausible and that the conclusion that unconscious cognition exists dissociated from conscious cognition remains beyond empirical reach. They therefore abandoned the search for an indirect-without-direct-effect data pattern and proposed a new criterion for SSA based on an *indirect-greater-than-direct-effect* data pattern. Instead of demanding that the direct measure be an absolute index of conscious knowledge, they suggested comparing the effects produced by comparable direct and indirect measures of some discrimination. SSA is evidenced when the effect of an indirect measure is greater than the effect of a comparable direct measure, and this is true even when the latter effect is different from zero. The only condition that must be fulfilled for this *relative sensitivity approach* to warrant the successful detection of SSA is that sensitivity to conscious, task-relevant information should not be smaller (i.e., it should be greater or equal) for the direct measure than for a comparable indirect measure. It seems plausible that this condition is met when comparable direct and indirect measures are used, because participants should make more use of conscious information when they are instructed to do so (in direct tasks) than when they are not (in indirect tasks). Suppose that this condition is fulfilled and that a larger effect is found with the indirect than the direct task; then one could not attribute the surplus effect on the indirect task to conscious information, precisely because the direct task is said to make more use of conscious information than does the indirect task. Although comparable direct and indirect tasks are allowed to differ only with respect to instructions, the notion "comparable" still leaves room for interpretation and debate (see Holender, 1986).

Despite Reingold and Merikle's (1988) critical analysis, Greenwald, Klinger, and Schuh (1995) were not prepared to abandon the search for an indirect-without-direct-effect data pattern. To overcome the statistical impossibility of accepting a null finding, Greenwald et al. (1995; Draine & Greenwald, 1998; Klauer, Greenwald, & Draine, 1998) developed a regression method, in which the indirect measure was regressed onto the direct measure. A positive significant intercept indicates above-zero performance on the indirect measure (i.e., priming) at a point of zero performance on the direct measure (i.e., no awareness of the primes). In

response, Merikle and Reingold (1998) argued that the use of a regression method does not change the fact that a dissociation conclusion requires the exhaustiveness assumption, which remains as problematic as before. Greenwald and Draine (1998) admitted that the exhaustiveness assumption cannot be tested empirically but defended the plausibility of it for at least a portion of their data. Recognizing that evidence for unconscious cognition dissociated from conscious cognition may never be conclusive, these authors still found it worthwhile to collect data that add credence to such a dissociation.

In addition to explaining how unconscious cognition can be studied (and the difficulties inherent to it), this section has shown that disagreement about what constitutes the right measure for a certain feature can sometimes be traced to a disagreement in the underlying metatheoretical view about what measures are supposed to deliver: an approximation or a conclusive argument. Apart from our suspicion that conclusive arguments remain practically unattainable, it seems that more conclusive arguments are possible only when the initial questions are attenuated.

Diagnostic Status

Although many scholars have bracketed together the words *unconscious* and *automatic* (e.g., Jacoby, 1991; Kihlstrom, 1987; Posner & Snyder, 1975a), others have argued against a complete overlap between the concepts. In fact, several authors do not think the feature unconscious is a necessary feature of automatic processes (e.g., Tzelgov, 1997). In Bargh's (1992) conditional approach, the feature unconscious is treated as an optional feature of automaticity. At the same time, however, the feature unconscious is regarded as a strong indication of automaticity. We think that this is due in part to the links that have been assumed between the feature unconscious and other features of automaticity. Recall in this respect the ubiquitous view that unconscious input prohibits intentional processing and that unconscious processes cannot be controlled in the sense of being altered, stopped, or avoided. We have argued that these ideas make sense only given the assumption that the terms *intentional* and *controlled* include the feature conscious as a defining ingredient, an assumption to which we do not subscribe. In our view, the status of the feature unconscious may have been overestimated.

In addition to the position that not all automatic processes are unconscious, there is a position that holds that not all unconscious processes are automatic. Certain unconscious processes are non-automatic, such as the veiled control processes in the Shiffrin and Schneider (1977) model. Other unconscious processes, such as preattentive processes or emergent processes, have been excluded from the realm of automatic processes by certain authors (Cleeremans & Jiménez, 2002; Logan, 1992).

(Non)efficient

Concept, Measurement, and Diagnostic Status

A process is defined as efficient when it consumes little or no processing resources or attentional capacity. Efficiency leads to the subjective experience that processing is effortless (Bargh, 1989), and for that reason, the terms *efficient* and *effortless* have been used interchangeably, as have the terms *nonefficient* and *effortful*.

The criterion of attentional requirement was originally depicted as a continuum and efficiency was regarded as a gradual concept (Hasher & Zacks, 1979).

The degree of capacity consumption is often operationalized as dual-task interference (e.g., Baddeley, 1986; Logan, 1978, 1979; but see Shiffrin & Schneider, 1977). In dual-task studies, participants perform a primary task simultaneously with a secondary task that is assumed to consume resources. The efficiency of the processes involved in the primary task is inferred when performance is unaffected by varying (and especially augmenting) the resources used by the secondary task. For example, Hermans, Crombez, and Eelen (2000) had participants perform an affective priming task while simultaneously reciting a series of digits. Manipulation of the mental load imposed by the digit series did not influence the magnitude of the priming effect, suggesting the efficient nature of the evaluation process underlying priming. However, one can never demonstrate that the processes involved in the primary task are totally independent of resources because it is impossible to know whether the secondary task exhausts all processing resources.

The dual-task method is founded on the single-resource view and on a number of assumptions inherent to it. To the extent that these assumptions can be challenged, the usefulness of the method is called into question. From the multiple-resource perspective, the dual-task method is useful only when both tasks use the same processing resource (Navon & Gopher, 1979). On the other hand, multiple-resource theories are difficult to falsify because they often do not specify how many resource pools there are and which tasks draw on which pools.

Efficiency occupies a central place in the capacity view (e.g., Shiffrin & Schneider, 1977). The notion of efficiency rests heavily on the metaphor of attention as a barrel filled with liquid (e.g., Kahneman, 1973). This metaphor captures the energy or quantity aspect of attention. Another aspect of attention is its direction, which is captured by another metaphor, that of attention as a spotlight (e.g., Posner, Snyder, & Davidson, 1980). The notion of efficiency has primarily been linked to the quantity aspect. Although it is fair to say that both aspects, quantity and direction, cannot occur independent of each other, there may be processes that depend on the specific direction of attention but use only a minimal amount of attention and hence are efficient. Recall, in this respect, processes underlying Stroop interference, which is diminished when attention is directed away from the location or the critical feature of the stimulus (e.g., Kahneman & Henik, 1981).

A final note is that attention can be directed either at the process or at the stimulus input on which the process operates. The feature efficient can thus be used as a predicate of the critical process or, alternatively, as a predicate of the initial perceptual process operating on the raw stimulus input. A clear distinction between the uses should be kept in mind.

Links Between Efficient and Other Features

Link between efficient and goal-related features. The idea that goals determine where attention should be directed is widely accepted (Allport, 1989; Shallice, 1978). This may also hold true for processing goals. To implement the goal to engage in some process, attention must be directed at an appropriate stimulus input. To implement the goal to alter or stop a process, attention

must be directed at the process. These goals need not be conscious. This argument is consistent with the distinction in several theories between an attention system guided by conscious goals and one guided by unconscious goals (e.g., Baars, 1998; Luria, 1973; Posner & Rothbart, 1989; Shiffrin & Schneider, 1977).

One way to separate goal-related features and efficiency is by considering the distinction between the quantity and the direction of attention. It may be that goals relate more to direction than to quantity. Although implementation of the goal to engage in (or alter or stop) some process requires attention to be directed to suitable input (or to the process itself), the detection of this input (or process) may be accomplished with minimal attention. In this way, an intentional (or otherwise controlled) process may still qualify as efficient.

Another way to separate goal-related features and efficiency is by asking whether goals are necessary for attention orienting. Several authors have argued that goal relevance and, in particular, the discrepancy between a stimulus and a goal or desired state are unnecessary for attention orienting. Attention orienting may also be elicited by novelty or the discrepancy between a stimulus and an expected state (Gati & Ben-Shakar, 1990; Sokolov, 1963; but see Bernstein, 1969).

When attention orienting is driven by unconscious goals or by the novelty of the stimulus, there is no accompanying conscious experience of orienting or directing. It feels as if the stimulus captures one's attention. This phenomenon is reflected in use of the terms *attentional capture*, *orienting reflex* (Sokolov, 1963), and *automatic attention response* (Shiffrin & Schneider, 1977). Also note that attention need not always be directed to be located in a particular place. The attention spotlight may randomly travel around the visual field until it crosses goal-relevant or novel stimuli, and it may then dwell on these stimuli (Fox, Russo, Bowles, & Dutton, 2001; Posner, 1980).

Link between efficient and conscious. Consciousness is another feature that has close ties to attention. Most models assume that the focus or direction of attention determines the content of consciousness: Whatever is in the focus of attention is conscious; everything else is unconscious. In the Shiffrin and Schneider (1977) model, automatically activated nodes that enter short-term memory remain unconscious unless attention is directed or attracted to them. Some theorists assume that consciousness and attention amount to the same thing (e.g., Baars, 1988; Posner, 1982, 1994; but see Baars, 1997). Other theorists argue that attended information does not exhaust humans' conscious experience (e.g., Block, 1995; Hardcastle, 1997; Nelkin, 1993). Recall that Block (1995) coupled A-consciousness with attention but reserved P-consciousness as a form of consciousness that is not strictly reliant on attention (cf. his example of the pneumatic drill). The idea of an incomplete overlap between consciousness and attention was described very early by Wundt (1896/1897, p. 209). In his view, elements that reach a certain threshold enter the field of consciousness, where they are interconnected in compounds. These compounds remain obscure (they are merely apprehended) until attention is directed to them. They then enter the so-called "fixation point of consciousness" (and are apperceived, i.e., clearly grasped). If one is willing to accept a form of consciousness that escapes attention (P-consciousness), then that form must be undisturbed by attentional constraints and be efficient. Further, even if A-consciousness is determined by the focus or direction of

attention, it does not force use a large amount of attentional capacity. This leaves the door open for A-conscious but efficient processes.

Along with the idea that the direction of attention determines the content of consciousness, it has been assumed that the amount of attention determines whether consciousness will occur (or, for gradual models, the degree of consciousness). Note that several theorists have argued that the amount of attention is not the only determinant of consciousness. For example, Cleeremans and Jiménez (2002) argued that the top-down element of attention allocation must be combined with the bottom-up element of the quality of the nodes involved. When a stimulus impinging on the senses is so weak or occurs so briefly that activation of the corresponding nodes remains below a certain threshold, the stimulus may not be accessible to consciousness even if attention is adequately focused.

Apart from its essential role in consciousness, attention may be of critical importance for unconscious encoding as well. This is the central thesis of Naccache, Blandin, and Dehaene (2002). They reported that subliminal priming effects vanished when temporal attention was focused away from the time window in which the prime-target pair was presented. Similar conclusions were drawn in older studies in which subliminal priming effects disappeared when primes were no longer presented in the same position as the targets but rather in a "flanker" position (e.g., Inhoff, 1982). Such data are at odds with the common view that subliminal priming reflects attentionless processing (e.g., Greenwald, 1992), and they provide more evidence in the case against conflating attention with consciousness.

In summary, several arguments were presented for separating the features conscious and nonefficient. First, if one is willing to accept a form of consciousness that escapes attention (P-consciousness), then that form must be undisturbed by attentional constraints and hence be efficient. Second, even if A-consciousness is determined by the focus or direction of attention, it does not necessarily use a large amount of attentional capacity. This means that A-conscious processes can be efficient. Third, given the evidence that attention is neither sufficient nor exclusive for awareness, nonefficient processes are not always conscious.

Fast (Slow)

Concept and Measurement

The concept fast does not require much analysis apart from noting that it is gradual. A fast process is a process with a short duration. It is useful to distinguish between the duration of a process and the duration of the stimulus input on which the process operates. For example, in affective priming tasks, the short duration of the primes does not demonstrate the short duration of the evaluation process operating on the primes. Participants may use almost the entire time between the onset of the prime and the target response to evaluate the prime. Therefore, the short duration of the primes is combined with a short stimulus onset asynchrony (the time between the onset of the prime and the onset of the target) and with quick target responses (because of speeded response instructions). The fact that *fast* is a gradual notion has consequences for measurement. Because there are no objective thresholds available, investigators must rely on common sense arguments for calling some interval short or long and for deciding whether the process under study was fast or slow.

Links Between Fast and Other Features

Fast and efficient. According to the capacity view of automaticity, nonefficient processes are more hampered than efficient processes when resources are scarce. This obstruction can be manifested in slowing of the process or in increased latency in starting the process because other nonefficient processes must be completed first. Of interest, some authors have defined the term *efficient* in terms of low demand on processing capacity and increased speed and accuracy of processing (e.g., Smith & Lerner, 1986). Consequently, these authors have considered faster performance and diminished errors as evidence for greater efficiency.

Fast and unconscious. In the conceptual sections that addressed the features fast and unconscious, we found it useful to specify whether these features describe the process or, alternatively, describe the stimulus input. We believe there is a link between the features fast and unconscious. First, awareness of the stimulus input is related to its duration. Some stimulus presentations are too short to permit the conscious perception of the stimulus. The fact that, in subliminal perception research, stimuli are presented very briefly (in combination with masks) to prevent their conscious perception illustrates just how deep-rooted this link is. Second, awareness of the process is related to its duration. Some processes occur too quickly to enter consciousness. In the Shiffrin and Schneider (1977) model, the constituent nodes of both automatic and veiled control processes are presumed to remain unconscious because they reside in short-term memory only briefly. Similarly, Logan (1988) contended that single-step memory retrieval occurs so fast that it is usually impervious to consciousness: "Automatic processes may not be available to the mind's eye long enough to provide conscious evidence of their inner workings" (p. 514).

Fast and goal-related features. Another set of features that has close ties with the feature fast are the goal-related features unintentional and uncontrolled (in the sense of alter or stop). First, implementing the goal to engage in a process depends on the duration of the stimulus input. When a stimulus occurs too briefly, there may not be enough time to intentionally engage in a process that uses the stimulus as its input. Second, implementing the goals to alter or stop a process depends on the duration of the process. When a process occurs too rapidly, there may not be enough time to alter or stop it before its completion (Logan, 1989).

Gollwitzer (e.g., 1993, 1999) showed that when a goal is chosen in advance and an implementation intention is formed (i.e., a decision has been made about the circumstances in which the goal is to be implemented), the actual implementation may happen instantly upon detection of the predetermined circumstances. It thus seems that implementation of a processing goal may be less affected by the short duration of the stimulus input (in the case of the goal to engage in) or the short duration of the process (in the case of the goals to alter or stop) when the goal is chosen in advance and when an implementation intention is established. This complicates the strategy of drawing inferences about goal-related features on the basis of the short duration time of input or processes. Another complication is the possibility of interdependence among the factors time, attention, and intensity (of either input or process). For example, higher levels of attention and/or stimulus strength may compensate for short stimulus durations.

Diagnostic Status

The feature fast has sometimes been omitted from lists of automaticity features. Yet, it is present in the writings of many authors and has close links with other features. The feature fast derives an important part of its diagnostic status for automaticity from these links. It is a consequence of the feature efficient and a condition for certain goal-related features and the feature unconscious. Some processes are too short-lived for some forms of control to be successful, and some stimulus presentation times are too brief to permit the (conscious) perception of these stimuli and the occurrence of processes that rely on their (conscious) input. A major difficulty is determining where exactly to draw the line between those time intervals that do and those that do not allow processing goals or consciousness to take effect. This determination is especially difficult given that implementation intentions can help people take advantage of very short time windows to implement their processing goals. Moreover, we question whether the search for fixed time intervals independent of other conditions is useful. The appraisal of a process as fast does not allow one to draw significant inferences about the presence of other features.

Regardless of whether the feature fast has diagnostic value for other features of automaticity, it may be a fully fledged automaticity feature in its own right. According to the algorithm-based account of automaticity, when a process becomes automatic because of practice, all initial steps are preserved, but they are carried out faster and with fewer mistakes (Anderson, 1992; Carlson & Lundy, 1992). From an adaptational point of view, when complex operations can be executed within a few hundred milliseconds, they can be extremely useful for survival, even if they require effort, goals, and/or consciousness.

Conclusions

We analyzed eight concepts that have often been associated with (non)automaticity: (un)intentional, goal (in)dependent, (un)controlled/(un)controllable, autonomous, purely stimulus driven, (un)conscious, (non)efficient, and fast (slow) (see Table 1). For each feature, we specified the minimal ingredients of their definitions. For certain features, we specified different types, on the basis of different sets of minimal ingredients. For example, we distinguished between different types of uncontrolled processes (Types A1, A2, B1, and B2) and different types of consciousness (A-consciousness vs. P-consciousness). In addition, we specified whether features can be used as a predicate of the critical process, its input, its output, or the initial perceptual process operating on the raw stimulus input.

Our definitions show partial conceptual overlap among certain goal-related features but not among goal-related features and the features (un)conscious, (non)efficient, and fast (slow) or among any of these other features. In addition to examining the conceptual overlap between features, we also noted connections of a different kind than identity relations, as when one feature is considered to be a necessary or sufficient condition for another. We showed that the relations between factors such as goals, attention, time, salience, and consciousness are not one-to-one relations. Goals may direct attention, but so may novelty. The direction of attention determines the content of A-consciousness but not P-consciousness. The amount of attention available determines the possibility of

Table 1
Types and Ingredients of (Non)automaticity Features

Feature	Type	Ingredients		
		Condition	Effect	Causal connection
Controlled		Goal pertaining to process present To engage in process To alter process To stop process To avoid process	Effect present Occurrence of process Change in process Interruption of process Prevention of process	Present
Uncontrolled	A1	Goal pertaining to process absent	Effect absent	—
	A2	Goal pertaining to process absent	Effect present	—
	B1	Goal pertaining to process present	Effect absent	—
	B2	Goal pertaining to process present	Effect present	Absent
Intentional ^a		Intention ^b present	Process present	Present
Unintentional	A	Intention absent	Process present	—
	B	Intention present	Process present	Absent
Goal dependent		Goal present	Process present	Present
Goal independent	A	Goal absent	Process present	—
	B	Goal present	Process present	Absent
Purely stimulus driven		Stimulus ^c	Process present	Present
Autonomous ^d		Goal to engage in process absent or present	Process present	— or absent
		And goal to alter, stop, or avoid process absent or present	Effect absent	—
Conscious	A-conscious	Intentional aspect		
	P-conscious	Phenomenal aspect		
	M-conscious	Conscious of A-consciousness or P-consciousness		
	S-conscious	Conscious of oneself		
Unconscious	Accessible			
	Inaccessible			
Efficient		Minimal quantity of attention		
Nonefficient		Substantial quantity of attention		
Fast				
Slow				

Note. A = access; P = phenomenal; M = monitoring; S = self.

^a Controlled in terms of goal to engage in process. ^b Proximal goal to engage in process. ^c Plus usual background conditions. ^d Uncontrolled in terms of every possible goal.

consciousness, but it probably does so in interaction with the factors stimulus strength (salience) and time. Moreover, the factors amount of attention, stimulus strength, and time are probably also important for unconscious processing. Finally, consciousness may determine whether conscious goals can be implemented but not whether unconscious goals can be implemented. Given this complex picture of interrelations, it is dangerous to draw inferences about the presence or absence of one feature on the basis of the presence or absence of another.

General Conclusions

Many researchers are concerned with the problem of how to diagnose the automatic nature of performance and processes. After considering a few of the leading views of automaticity, our first conclusion was that it can be diagnosed by looking at the presence of features. Our second conclusion was that these features should be investigated separately. To see whether such a decompositional approach to the study of automaticity is worthwhile, we examined to what extent features of automaticity can be conceptually and logically separated, and we concluded that this is, to a large extent, feasible.

Few authors have pushed the distinctions among features to the limits as we have, and assumptions of overlap are frequent in the extant literature. This overlap is often attributable to authors' theoretical choices about automaticity (e.g., all-or-none view) and about the cognitive system in general (e.g., specialized theories and computational framework). Authors' assumptions of overlap shape the diagnostic status and definitions of features. In turn, these elements influence which feature(s) one will examine and what evidence one considers to be valuable in demonstrating the presence of this (these) feature(s).

Given our conclusion that the conceptual and logical separation of features is feasible, it makes sense to adopt a decompositional approach toward automaticity. This being said, however, investigating features in isolation is not the only useful approach. There may be reasons to investigate the presence of some particular combination of features. For example, with regard to the issue of whether people can be held responsible for their prejudices, it may be legitimate to investigate the role of conscious intentions in prejudice rather than investigate intentionality irrespective of consciousness. Another consideration is that it may not always be easy to investigate features separately. For example, isolating the role of intentions that are not conscious may be extremely difficult. De-

spite these complications, we argue that researchers should try to investigate features separately whenever possible. Although we acknowledge that the separate investigation of features is not always practical, we invite researchers to seek maximal separation when it comes to definitional issues.

In addition to a decompositional approach to automaticity, we also favor a gradual approach. Our approach is gradual in two senses. First, several of the features we discussed are inherently gradual. A feature is gradual when it can be present to some degree. This is most obvious for the features fast and efficient. Although less obvious, the feature unconscious and the goal-related features can also be construed as gradual. For those theorists who see the transition between unconscious and conscious cognition as a matter of strength of activation, the notion of gradual awareness (or unawareness) is self-evident. We would argue that the presence of the goal-related features intentional and controlled can also be seen as a matter of degree. A person can be said to have only a degree of control over some process when there is a less than perfect match between the goal the person had regarding the process and the outcome that was obtained (Uleman, 1999). The same goes for the feature intentional. The process one produces may match what one intended to produce in varying degrees. But this is not the whole story. There is more than one way in which a process can be classified as uncontrolled or unintentional: Either the person did not pursue a goal with regard to the process, or the person did pursue a goal, but the goal was not, or only partially, achieved. Only in the case in which the person did pursue a goal may there be something gradual to the words *uncontrolled* and *unintentional*; when there is a goal, it may be achieved to varying degrees. Of course nothing gradual occurs in the case in which a person simply did not pursue a processing goal.

For each continuum that represents more or less of something, there is one pole that represents no quantity of that thing. Thus, it is reasonable to ask whether a process can occur in the complete absence of time, resources, and consciousness and despite the fact that the person's goals regarding the process were completely unsuccessful. Presenting features as continua does not overrule questions concerning the extreme poles of the continuum. This brings us to the second sense in which our approach can be called gradual. Although investigating the extreme poles of the feature continua is legitimate, it is often not practically attainable, because available research methods seldom allow for conclusive arguments. Of the various methods that we reviewed for testing the presence of automaticity features, not a single test gave airtight guarantees. In the end, every test is based on one or more plausibility assumptions. Recall, for example, the debates revolving around SSA or the fact that measures of efficiency do not allow one to conclude that a process is completely independent of resources. The message to be drawn from such debates is not to give up investigating automaticity features altogether but rather to gather empirical support one step at a time. Our approach can thus be called gradual in the sense that we accept gradual support for the diagnosticity of various automaticity features.

As noted above, an important weakness of the gradual approach is that it does not provide objective criteria for distinguishing between the automatic and the nonautomatic member of each feature pair: Every process is uncontrolled, efficient, unconscious, and fast, to some degree. One solution may be to provide a

subjective criterion. In fact, we argue for a relative approach. The term *relative* covers more than the term *gradual*; it involves a standard of comparison. For example, saying that a process is relatively controlled implies that there is some degree of control and that it is controlled relative to some standard, that is, more controlled than something else. Although it is not currently a common practice, it would be helpful if investigators would specify which standard for comparison they have in mind when calling a process relatively controlled. For example, a process may be controlled relative to other processes, relative to what is generally thought about the process, or relative to the way it was before practice. Although the choice of a standard for comparison is arbitrary, once an investigator has revealed his or her choice, it becomes possible to draw a line between controlled versus uncontrolled processes (and between the members of other feature pairs). The term *automaticity* may be kept as an umbrella term, for reasons discussed earlier (cf. Logan, 1985), as long as investigators specify in what sense they believe a process to be automatic or, in other words, which features they think apply to it.

The present analysis is far from complete, but we believe that it provides a useful framework for thinking about the features that have been associated with automaticity. By mapping out several theoretical ideas about automaticity and its constituent features, we hope that investigators will be able to choose their theoretical standpoints aware of the available options. But most of all, we hope that our analysis will help investigators formulate their empirical arguments for the presence or absence of automaticity features in a more transparent way so as to improve communication about this important issue. Finally, our framework may help identify gaps in the existing automaticity research and encourage future research that is intended to address these unresolved issues.

References

- Abrams, R. L., & Greenwald, A. G. (2000). Parts outweigh the whole (word) in unconscious analysis of meaning. *Psychological Science, 11*, 118–124.
- Allport, A. (1989). Visual attention. In M. I. Posner (Ed.), *Foundation of cognitive science* (pp. 631–682). Cambridge, MA: MIT Press.
- Allport, A. (1993). Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 183–218). Cambridge, MA: MIT Press.
- Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review, 89*, 369–406.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1992). Automaticity and the ACT* theory. *American Journal of Psychology, 105*, 165–180.
- Anderson, J. R. (1996). ACT: A simple theory of complex cognition. *American Psychologist, 51*, 355–365.
- Atkinson, A. P., Thomas, M. S. C., & Cleeremans, A. (2000). Consciousness: Mapping the theoretical landscape. *Trends in Cognitive Sciences, 4*, 372–382.
- Atkinson, R. C., & Shiffrin, R. M. (1971). The control of short-term memory. *Scientific American, 224*, 82–90.
- Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge, England: Cambridge University Press.
- Baars, B. J. (1995). Evidence that phenomenal consciousness is the same as access consciousness. *Behavioral and Brain Sciences, 18*, 249.

- Baars, B. J. (1997). Some essential differences between consciousness and attention, perception, and working memory. *Consciousness and Cognition*, 6, 363–371.
- Baars, B. J. (1998). Metaphors of consciousness and attention in the brain. *Trends in Neurosciences*, 21, 58–62.
- Baddeley, A. (1986). *Working memory*. Oxford, England: Clarendon Press.
- Balota, D. A. (1983). Automatic semantic activation and episodic memory. *Journal of Verbal Learning and Verbal Behavior*, 22, 88–104.
- Barbiero, D. (2004, May). Phenomenological critique of representationalism. In C. Eliasmith (Ed.), *Dictionary of philosophy of mind*. Retrieved August 5, 2004, from <http://philosophy.uwaterloo.ca/MindDict/phenomcritique.html>
- Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 3–51). New York: Guilford Press.
- Bargh, J. A. (1990a). Auto-motives: Preconscious determinants of social interaction. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition* (Vol. 2, pp. 93–130). New York: Guilford Press.
- Bargh, J. A. (1990b). Goal ≠ intent: Goal-directed thought and behavior are often unintentional. *Psychological Inquiry*, 1, 248–251.
- 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.
- Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 1, pp. 1–40). Hillsdale, NJ: Erlbaum.
- Bargh, J. A. (1996). Automaticity in social psychology. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 169–183). New York: Guilford Press.
- Bargh, J. A. (1997). The automaticity of everyday life. In R. S. Wyer (Ed.), *Advances in social cognition* (Vol. 10, pp. 1–49). Mahwah, NJ: Erlbaum.
- Bargh, J. A. (1999). The cognitive monster: The case against the controllability of automatic stereotype effects. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 361–382). New York: Guilford Press.
- Bargh, J. A., & Barndollar, K. (1996). Automaticity in action: The unconscious as repository of chronic goals and motives. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 457–481). New York: Guilford Press.
- Bargh, J. A., Chaiken, S., Raymond, P., & Hymes, C. (1996). The automatic evaluation effect: Unconditional automatic activation with a pronunciation task. *Journal of Experimental Social Psychology*, 32, 104–128.
- Bargh, J. A., & Chartrand, T. L. (2000). The mind in the middle: A practical guide to priming and automaticity research. In H. Reis & C. Judd (Eds.), *Handbook of research methods in social psychology* (pp. 253–285). New York: Cambridge University Press.
- Bargh, J. A., & Gollwitzer, P. M. (1994). Environmental control of goal-directed action: Automatic and strategic contingencies between situations and behavior. In W. D. Spaulding (Ed.), *Nebraska Symposium* (Vol. 41, pp. 71–124). Lincoln: University of Nebraska Press.
- Bargh, J. A., Gollwitzer, P. M., Lee-Chai, A., Barndollar, K., & Trötschel, R. (2001). The automated will: Nonconscious activation and pursuit of behavioral goals. *Journal of Personality and Social Psychology*, 81, 1014–1027.
- Bartlett, F. C. (1932). *Remembering*. Cambridge, England: Cambridge University Press.
- Bernstein, A. (1969). To what does the orienting response respond? *Psychophysiology*, 6, 338–350.
- Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, 6, 242–261.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18, 227–287.
- Brand, M. (1989). Proximate causation of action. *Philosophical Perspectives*, 3, 423–442.
- Bratman, M. E. (1987). *Intention, plans, and practical reason*. Cambridge, MA: Harvard University Press.
- Brentano, F. C. (1874). *Psychologie vom empirischen Standpunkt* [Psychology from an empirical standpoint]. Leipzig, Germany: Duncke & Humblot.
- Broadbent, D. (1958). *Perception and communication*. London: Pergamon Press.
- Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, 64, 123–152.
- Bryan, W. L., & Harter, L. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review*, 6, 345–378.
- Carlson, R. A., & Lundy, D. H. (1992). Consistency and restructuring in cognitive procedural sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 127–141.
- Carver, C. S., & Scheier, M. F. (2002). Control processes and self-organization as complementary principles underlying behaviour. *Personality and Social Psychology Review*, 6, 304–315.
- Chalmers, D. J. (1995). Facing up the problem of consciousness. *Journal of Consciousness Studies*, 2, 200–219.
- Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory*. New York: Oxford University Press.
- Chartrand, T. L., & Bargh, J. A. (1996). Automatic activation of impression formation and memorization goals: Nonconscious goal priming reproduces effects of explicit task instructions. *Journal of Personality and Social Psychology*, 39, 752–766.
- Cheesman, J., & Merikle, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics*, 36, 387–395.
- Christiansen, M. H., & Chater, N. (1992). Connectionism, learning, and meaning. *Connection Science*, 4, 227–252.
- Clark, A. (1993). *Associative engines*. Cambridge, MA: MIT Press.
- Cleeremans, A. (1997). Principles for implicit learning. In D. Berry (Ed.), *How implicit is implicit learning?* (pp. 195–234). Oxford, England: Oxford University Press.
- Cleeremans, A. (2005). Computational correlates of consciousness. *Progress in Brain Research*, 150, 81–98.
- Cleeremans, A., & Jiménez, L. (2002). Implicit learning and consciousness: A graded, dynamic perspective. In R. M. French & A. Cleeremans (Eds.), *Implicit learning and consciousness* (pp. 1–40). Hove, England: Psychology Press.
- Davidson, D. (1980). *Essays on actions and events*. Oxford, England: Oxford University Press.
- Davis, L. H. (1994). Action. In S. Guttenplan (Ed.), *A companion to the philosophy of mind* (pp. 111–117). Oxford, England: Blackwell.
- Dennett, D. C. (1984). *Elbow room: The varieties of free will worth wanting*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1991). *Consciousness explained*. Boston: Little, Brown.
- Deutsch, J. A., & Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review*, 70, 80–90.
- Dienes, Z., & Perner, J. (1999). A theory of implicit knowledge. *Behavioral and Brain Sciences*, 22, 735–808.
- Draine, S. C., & Greenwald, A. G. (1998). Replicable unconscious semantic priming. *Journal of Experimental Psychology: General*, 127, 286–303.
- Dulany, D. E. (1997). Consciousness in the explicit (deliberative) and implicit (evocative). In J. D. Cohen & J. W. Schooler (Eds.), *Scientific approaches to the study of consciousness* (pp. 179–212). Hillsdale, NJ: Erlbaum.
- Eilan, N. (1998). Perceptual intentionality. Attention and consciousness. In

- A. O'Hear (Ed.), *Current issues in philosophy of mind* (pp. 181–202). Cambridge, England: Cambridge University Press.
- Erdelyi, M. H. (1992). Psychodynamics and the unconscious. *American Psychologist*, *47*, 784–787.
- Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, *50*, 229–238.
- Fiske, S. 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 Press.
- Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Fowler, C. A., Wolford, G., Slade, R., & Tassinari, L. (1981). Lexical access with and without awareness: A methodological survey and evaluation. *Journal of Experimental Psychology: General*, *110*, 341–362.
- Fox, E., Russo, R., Bowles, R., & Dutton, K. (2001). Do threatening stimuli draw or hold visual attention in subclinical anxiety? *Journal of Experimental Psychology: General*, *130*, 681–700.
- Francolini, C. M., & Egeth, H. (1980). On the automaticity of “automatic” activation: Evidence of selective seeing. *Perception & Psychophysics*, *27*, 331–342.
- Freud, S. (1984). The unconscious. In A. Richards (Ed.), *The Pelican Freud library* (Vol. 11, pp. 159–222). Harmondsworth, England: Penguin. (Original work published 1915)
- Gati, I. & Ben-Shakar, G. (1990). Novelty and significance in orientation and habituation: A feature-matching approach. *Journal of Experimental Psychology: General*, *119*, 251–263.
- Gibson, E. J. (1969). *Principles of perceptual learning and development*. New York: Appleton-Century-Crofts.
- Gillett, E. (1996). Searle and the “deep unconscious”. *Philosophy, Psychiatry, and Psychology*, *3*, 191–200.
- Gollwitzer, P. M. (1990). Action phases and mind-sets. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, pp. 53–92). New York: Guilford Press.
- Gollwitzer, P. M. (1993). Goal achievement: The role of intentions. *European Review of Social Psychology*, *4*, 141–185.
- Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist*, *54*, 493–503.
- Green, C. D., & Vervaeke, J. (1996). What kind of explanation, if any, is a connectionist net? In C. W. Tolman, F. Cherry, R. van Hezewijk, & I. Lukbek (Eds.), *Problems of theoretical psychology* (pp. 201–208). North York, Ontario, Canada: Captus.
- Greenwald, A. G. (1992). New look 3: Unconscious cognition reclaimed. *American Psychologist*, *47*, 766–779.
- Greenwald, A. G., & Draine, S. C. (1998). Distinguishing unconscious from conscious cognition—Reasonable assumptions and replicable findings: Reply to Merikle and Reingold (1998) and Doshier (1998). *Journal of Experimental Psychology: General*, *127*, 320–324.
- Greenwald, A. G., Klinger, M. R., & Schuh, E. S. (1995). Activation by marginally perceptible (“subliminal”) stimuli: Dissociation of unconscious from conscious cognition. *Journal of Experimental Psychology: General*, *124*, 22–42.
- Hardcastle, V. G. (1997). Attention versus consciousness: A distinction with a difference. *Cognitive Studies: Bulletin of the Japanese Cognitive Science Society*, *4*, 56–66.
- Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, *42*, 335–346.
- Hasher, L., & Zacks, R. T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General*, *108*, 356–388.
- Heckhausen, H. (1991). *Motivation and action* (P. K. Leppman, Trans.). Berlin, Germany: Springer-Verlag. (Original work published 1990)
- Hendrickx, H., De Houwer, J., Baeyens, F., Eelen, P., & Van Avermaet, E. (1997). Hidden covariation detection might be very hidden indeed. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 201–220.
- Hermans, D., Crombez, G., & Eelen, P. (2000). Automatic attitude activation and efficiency: The fourth horseman of automaticity. *Psychologica Belgica*, *40*, 3–22.
- Hermans, D., De Houwer, J., & Eelen, P. (1994). The affective priming effect: Automatic activation of evaluative information in memory. *Cognition and Emotion*, *8*, 515–533.
- Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *Behavioral and Brain Sciences*, *9*, 1–23.
- Inhoff, A. W. (1982). Parafoveal word perception: A further case against semantic preprocessing. *Journal of Experimental Psychology: Human Perception and Performance*, *8*, 137–145.
- Jacob, P. (2003, Fall) Intentionality. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Retrieved August 5, 2004, from <http://plato.stanford.edu/entries/intentionality>
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic and intentional uses of memory. *Journal of Memory and Language*, *30*, 513–541.
- James, W. (1890). *The principles of psychology*. New York: Holt, Rinehart & Winston.
- Jastrow, J. (1906). *The subconscious*. Boston: Houghton Mifflin.
- Jiménez, L., Méndez, C., & Cleeremans, A. (1996). Comparing direct and indirect measures of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 948–969.
- Johnson-Laird, P. N. (1988). A computational analysis of consciousness. In A. J. Marcel & E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 357–368). Oxford, England: Oxford University Press.
- Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice Hall.
- Kahneman, D., & Chajczyk, D. (1983). Test of the automaticity of reading: Dilution of Stroop effects by color-irrelevant stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 497–509.
- Kahneman, D., & Henik, A. (1981). Perceptual organization and attention. In M. Kubovy & J. Pomerantz (Eds.), *Perceptual organization* (pp. 181–212). Hillsdale, NJ: Erlbaum.
- Kahneman, D., & Treisman, A. (1984). Changing views of attention and automaticity. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 29–61). Orlando, FL: Academic Press.
- Kant, I. (1929). *Critique of pure reason* (N. K. Smith, Trans.). New York: Macmillan. (Original work published 1787)
- Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: MIT Press.
- Kihlstrom, J. F. (1987, September 18). The cognitive unconscious. *Science*, *137*, 1445–1452.
- Klauer, K. C., Greenwald, A. G., & Draine, S. C. (1998). Correcting for measurement error in detecting unconscious cognition. *Journal of Experimental Psychology: General*, *127*, 318–319.
- Kramer, A. F., Strayer, D. L., & Buckley, J. (1990). Development and transfer of automatic processing. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 505–522.
- Kruglanski, A. W. (1996). Goals as knowledge structures. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action* (pp. 599–618). New York: Guilford Press.
- LaBerge, D. (1973). Attention and measurement of perceptual learning. *Memory & Cognition*, *1*, 268–276.
- LaBerge, D., & Samuels, S. J. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology*, *6*, 293–323.
- Lewicki, P. (1986). *Nonconscious social information processing*. New York: Academic Press.
- Lewicki, P. (1992). Nonconscious acquisition of information. *American Psychologist*, *47*, 796–801.
- Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acqui-

- sition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 523–530.
- Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8, 529–566.
- Loftus, E., & Klinger, M. R. (1992). Is the unconscious smart or dumb? *American Psychologist*, 47, 761–765.
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, 91, 295–327.
- Logan, G. D. (1978). Attention in character classification tasks: Evidence for the automaticity of component stages. *Journal of Experimental Psychology: General*, 107, 32–63.
- Logan, G. D. (1979). On the use of concurrent memory load to measure attention and automaticity. *Journal of Experimental Psychology: Human Perception and Performance*, 5, 189–207.
- Logan, G. D. (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *Cognitive Psychology*, 12, 523–553.
- 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.
- Logan, G. D. (1989). Automaticity and cognitive control. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 52–74). New York: Guilford Press.
- Logan, G. D. (1991). Automaticity and memory. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 347–366). Hillsdale, NJ: Erlbaum.
- Logan, G. D. (1992). Attention and preattention in theories of automaticity. *American Journal of Psychology*, 105, 317–339.
- Luria, A. R. (1973). *The working brain: An introduction to neuropsychology*. New York: Basic Books.
- Mackie, J. L. (1965). Causes and conditionals. *American Philosophical Quarterly*, 2, 245–264.
- Mackie, J. L. (1974). *The cement of the universe: A study of causation*. Oxford, England: Clarendon Press.
- MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, 109, 163–203.
- Malle, B. F., & Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology*, 33, 101–121.
- Marcel, A. J. (1980). Conscious and preconscious recognition of polysemous words: Locating the selective effect of prior verbal context. In R. S. Nickerson (Ed.), *Attention and Performance VIII* (pp. 435–457). Hillsdale, NJ: Erlbaum.
- Marcel, A. J. (1983a). Conscious and unconscious perception: An approach to the relation between phenomenal experience and perceptual processes. *Cognitive Psychology*, 15, 238–300.
- Marcel, A. J. (1983b). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology*, 15, 197–237.
- Marcel, A. J. (1988). Phenomenal experience and functionalism. In A. J. Marcel & E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 121–158). Oxford, England: Oxford University Press.
- Marr, D. (Ed.). (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: Freeman.
- McGill, A. L., & Tenbrunsel, A. E. (2000). Mutability and propensity in causal selection. *Journal of Personality and Social Psychology*, 79, 677–689.
- Mele, A. R., & Moser, P. K. (1994). Intentional action. *Noûs*, 28, 39–68.
- Menzies, P. (2004). The causal efficacy of mental states. In J. M. Monnoyer (Ed.), *The structure of the world: The renewal of metaphysics in the Australian School*. Paris: Vrin Publishers.
- Merikle, P. M., & Daneman, M. (1998). Psychological investigations of unconscious perception. *Journal of Consciousness Studies*, 5, 5–18.
- Merikle, P. M., & Reingold, E. M. (1998). On demonstrating unconscious perception: Comment on Draine and Greenwald (1998). *Journal of Experimental Psychology: General*, 127, 304–310.
- Mill, J. S. (1843). *A system of logic, ratiocinative and inductive*. London: Parker.
- Miller, G. A., Galanter, E., & Pribram, K. (1960). *Plans and the structure of behavior*. New York: Holt.
- Moors, A., De Houwer, J., & Eelen, P. (2004). Automatic stimulus–goal comparisons: Support from motivational affective priming studies. *Cognition and Emotion*, 18, 29–54.
- Moskowitz, G. B. (2001). Preconscious control and compensatory cognition. In G. B. Moskowitz (Ed.), *Cognitive social psychology: The Princeton symposium on the legacy and future of social cognition* (pp. 333–358). Hillsdale, NJ: Erlbaum.
- Moskowitz, G. B., Gollwitzer, P. M., Wasel, W., & Schaal, B. (1999). Preconscious control of stereotype activation through chronic egalitarian goals. *Journal of Personality and Social Psychology*, 77, 167–184.
- Moskowitz, G. B., Li, P., & Kirk, E. R. (2004). The implicit volition model: On the preconscious regulation of temporarily adopted goals. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 34, pp. 317–414). San Diego, CA: Academic Press.
- Moya, C. J. (1990). *The philosophy of action: An introduction*. Cambridge, England: Polity Press.
- Naccache, L., Blandin, E., & Dehaene, S. (2002). Unconscious masked priming depends on temporal attention. *Psychological Science*, 13, 416–424.
- Nagel, T. (1994). What is it like to be a bat? *Philosophical Review*, 83, 434–450.
- Natsoulas, T. (1981). Basic problems of consciousness. *Journal of Personality and Social Psychology*, 41, 132–178.
- Navon, D. (1995). A-consciousness: The local newspaper of the mind? *Behavioral and Brain Sciences*, 18, 265.
- Navon, D., & Gopher, D. (1979). On the economy of the human processing system. *Psychological Review*, 86, 214–255.
- Neisser, U. (1976). *Cognition and reality*. San Francisco: Freeman.
- Nelkin, N. (1993). The connection between intentionality and consciousness. In M. Davies & G. W. Humphreys (Eds.), *Consciousness: Psychological and philosophical essays* (pp. 224–239). Oxford, England: Blackwell.
- Neumann, O. (1984). Automatic processing: A review of recent findings and a plea for an old theory. In W. Prinz & F. Sanders (Eds.), *Cognition and motor processes* (pp. 255–293). Berlin, Germany: Springer-Verlag.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice Hall.
- Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behaviour. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation* (Vol. 4, pp. 1–18). New York: Plenum Press.
- Oatley, K. (1993). Freud's cognitive psychology of intention: The case of Dora. In M. Davies & G. W. Humphreys (Eds.), *Consciousness: Psychological and philosophical essays* (pp. 90–104). Oxford, England: Blackwell.
- O'Shaughnessy, B. (1991). Searle's theory of action. In E. Lepore & R. Van Gulick (Eds.), *John Searle and his critics* (pp. 271–280). Cambridge, MA: Blackwell.
- Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology*, 23, 324–354.
- Perruchet, P., & Vinter, A. (2002). The self-organizing consciousness. *Behavioral and Brain Sciences*, 25, 297–330.
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3–25.

- Posner, M. I. (1982). Cumulative development of attentional theory. *American Psychologist*, 37, 168–179.
- Posner, M. I. (1994). Attention: The mechanism of consciousness. *Proceedings of the National Academy of Science, USA*, 91, 7398–7403.
- Posner, M. I., & Klein, R. (1973). On the functions of consciousness. In S. Kornblum (Ed.), *Attention and performance IV* (pp. 21–35). New York: Academic Press.
- Posner, M. I., & Rothbart, M. K. (1989). Intentional chapters on unintended thoughts. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 450–469). New York: Guilford Press.
- Posner, M. I., & Snyder, C. R. R. (1975a). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 153–175). Hillsdale, NJ: Erlbaum.
- Posner, M. I., & Snyder, C. R. R. (1975b). Facilitation and inhibition in the processing of signals. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V* (pp. 669–682). New York: Academic Press.
- Posner, M. I., Snyder, C. R. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, 109, 160–174.
- Power, M., & Brewin, C. R. (1991). From Freud to cognitive science: A contemporary account of the unconscious. *British Journal of Clinical Psychology*, 30, 289–310.
- Powers, W. T. (1973). *Behavior: The control of perception*. Hawthorne, NY: Aldine de Gruyter.
- Prinz, W. (1990). A common coding approach to perception and action. In O. Neumann & W. Prinz (Eds.), *Relationships between perception and action* (pp. 167–201). Heidelberg, Germany: Springer-Verlag.
- Reber, A. S. (1993). *Implicit learning and tacit knowledge*. Oxford, England: Oxford University Press.
- Regan, J. E. (1981). Automaticity and learning: Effects of familiarity on naming letters. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 180–195.
- Reingold, E. M. (1992). Conscious versus unconscious processes: Are they qualitatively different? *Behavioral and Brain Sciences*, 15, 218–219.
- Reingold, E. M., & Merikle, P. M. (1988). Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics*, 44, 563–575.
- 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.
- Rosenbaum, D. A., & Saltzman, E. (1984). A motor-program editor. In W. Prinz & F. Sanders (Eds.), *Cognition and motor processes* (pp. 51–61). Berlin, Germany: Springer-Verlag.
- Rosenbloom, P. S., & Newell, A. (1986). The chunking of goal hierarchies: A generalized model of practice. In R. S. Michaliski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 247–288). Los Altos, CA: Morgan Kaufmann.
- Schacter, D. L. (1989). On the relation between memory and consciousness: Dissociable interactions and conscious experience. In H. L. Roediger & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 355–389). Mahwah, NJ: Erlbaum.
- Scherer, K. R. (1993). Neuroscience projections to current debates in emotion psychology. *Cognition and Emotion*, 7, 1–41.
- Schneider, W., Dumais, S. T., & Shiffrin, R. (1984). Automatic and control processing and attention. In R. S. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 1–27). New York: Academic Press.
- Schneider, W., & Fisk, A. D. (1984). Automatic category search and its transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 1–15.
- Searle, J. R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge, England: Cambridge University Press.
- Searle, J. R. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.
- Shallice, T. (1978). The dominant action system: An information-processing approach to consciousness. In K. S. Pope & J. L. Singer (Eds.), *The stream of consciousness* (pp. 117–157). New York: Plenum Press.
- Shallice, T. (1988). Information-processing models of consciousness: Possibilities and problems. In A. J. Marcel & E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 305–333). Oxford, England: Oxford University Press.
- Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17, 367–447.
- Shiffrin, R. M. (1988). Attention. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology* (Vol. 2, pp. 739–811). New York: Wiley.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84, 127–190.
- Slooman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22.
- Smith, E. R., & Lerner, M. (1986). Development of automatism of social judgements. *Journal of Personality and Social Psychology*, 50, 246–259.
- Sokolov, E. N. (1963). *Perception and the conditioned reflex*. New York: Macmillan.
- Solomons, L., & Stein, G. (1896). Normal motor automatism. *Psychological Review*, 3, 492–512.
- Sosa, E. (1993). Introduction. In E. Sosa & M. Tooley (Eds.), *Causation*. Oxford, England: Oxford University Press.
- Tolman, E. C. (1922). Concerning the sensation quality: A behavioristic account. *Psychological Review*, 29, 140–145.
- Tooley, M. (1990). Causation: Reductionism versus realism. *Philosophy and Phenomenological Research*, 50, 215–236.
- Treisman, A., Vieira, A., & Hayes, A. (1992). Automaticity and preattentive processing. *American Journal of Psychology*, 105, 341–362.
- Tzelgov, J. (1997). Automatic but conscious: That is how we act most of the time. In R. S. Wyer, Jr. (Ed.), *Advances in social cognition* (Vol. 10, pp. 217–230). Mahwah, NJ: Erlbaum.
- Tzelgov, J., Henik, A., Sneg, R., & Baruch, B. (1996). Unintentional word reading via the phonological route: The Stroop effect with cross-script homophones. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 336–349.
- 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.
- Tzelgov, J., Yehene, V., & Naveh-Benjamin, M. (1997). From automaticity to memory and vice versa: On the relations between automaticity and memory. In J. Brzezinski, B. Krause, & T. Maryszewski (Eds.), *Idealization in psychology: Poznan studies of sciences and the humanities* (Vol. 55, pp. 239–261). Amsterdam: Rodopi.
- Uleman, J. S. (1989). A framework for thinking intentionally about unintended thoughts. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 425–449). New York: Guilford Press.
- Uleman, J. S. (1999). Spontaneous versus intentional inferences in impression formation. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 141–160). New York: Guilford Press.
- Umiltà, C. (1988). The control operations of consciousness. In A. J. Marcel & E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 334–356). Oxford, England: Oxford University Press.
- Vallacher, R. R., & Wegner, D. M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review*, 94, 3–15.
- Wakefield, J., & Dreyfus, H. (1991). Intentionality and the phenomenology

- of action. In E. Lepore & R. Van Gulick (Eds.), *John Searle and his critics* (pp. 259–270). Cambridge, MA, Blackwell.
- Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review*, *101*, 34–52.
- Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
- Wegner, D. M. (2003). The mind's best trick: How we experience conscious will. *Trends in Cognitive Sciences*, *7*, 65–69.
- Wegner, D. M., & Bargh, J. A. (1998). Control and automaticity in social life. In D. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology* (4th ed., Vol. 1, pp. 446–496). New York: McGraw-Hill.
- Wegner, D. M., & Wheatley, T. (1999). Apparent mental causation: Sources of the experience of will. *American Psychologist*, *54*, 480–492.
- Wickens, C. D. (1984). Processing resources in attention. In R. Parasuraman & R. Davies (Eds.), *Varieties of attention* (pp. 63–102). New York: Academic Press.
- Wundt, W. (1897). *Outlines of psychology* (C. H. Judd, Trans.). Leipzig, Germany: Wilhelm Engelmann. (Original work published 1896)
- Wundt, W. M. (1902). *Grundzüge der physiologischen Psychologie* [Principles of physiological psychology] (Vol. 1). Leipzig, Germany: Wilhelm Engelmann.
- Wundt, W. M. (1903). *Grundzüge der physiologischen Psychologie* [Principles of physiological psychology] (Vol. 3). Leipzig, Germany: Wilhelm Engelmann.
- Zbrodoff, N. J., & Logan, G. D. (1986). On the autonomy of mental processes: A case study of arithmetic. *Journal of Experimental Psychology: General*, *115*, 118–130.

Received April 12, 2005

Revision received July 12, 2005

Accepted August 25, 2005 ■