To cite this chapter:

Moors, A., & De Houwer, J. (2007). What is automaticity: An analysis of its component features and their interrelations. In J. A. Bargh (Ed.), *Social psychology and the unconscious: The automaticity of higher mental processes* (pp. 11-50). Psychology Press.

1

What is Automaticity? An Analysis of Its Component Features and Their Interrelations

AGNES MOORS and JAN DE HOUWER

The concept of automaticity is becoming increasingly important across nearly all subareas of psychology. Investigators who search for guidelines to asses the automatic nature of a certain task performance or process are faced with a multitude of views of automaticity. The variety of views is largely due to the topic studied (e.g., perception, memory, skill development, attention), the research paradigm employed (e.g., direct vs. indirect), the characteristics of the underlying information processing model (e.g., instance-based vs. rule-based), and the characteristics of the larger framework in which the information processing model fits (e.g., computational vs. connectionist). Many views define the concept of automaticity in terms of a number of features, but they differ with regard to the features they put most emphasis on, as well as with regard to the coherence they assume among the features. One contemporary account is the gradual and decompositional view, which proposes to investigate each automaticity feature separately and determine the degree to which it is present. In this chapter, we engage in a detailed analysis of the most important features in order to examine whether they can indeed be regarded as gradual, and whether they can be conceptually and logically separated.

Before embarking on a discussion of some prominent views of automaticity, we must specify what the word automatic can be a predicate of and stipulate the research questions that prevail in the study of automaticity. The word automatic can be used to describe performances or effects, which are observable, or to describe the processes underlying the performance, which are not observable and hence need to be inferred. We take it as a general rule that when the performance is classified as automatic, so can the process underlying the performance. It may be good to keep in mind that processes can be described at different levels of analysis. Marr (1982), for example, distinguished between three levels of process understanding: the computational level, the algorithmic level, 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). For example, the higher-level functional process of stimulus evaluation (i.e., is a stimulus good or bad) may be further specified at the algorithmic level as a process of direct memory retrieval (activation of a stored valence label in memory) or as a process of algorithm computation (comparison between a desired and an actual state). The hardware level is concerned with the physical implementation of processes in the brain. Other theorists (e.g., Anderson, 1987; Pylyshyn, 1984) have proposed a different number of levels and have placed the boundaries between the levels at somewhat different heights, but the important lesson is that processes described at higher levels can be explained by processes described at lower levels. Similarly, performances or effects can be explained by higher-level processes and further down by lower-level processes. It should be clear that we use the term explanation here in the sense of an explanation that specifies the underlying mechanism.

Now that we have clarified our use of the terms performance, process, and explanation, we can distinguish between two types of research purposes that automaticity researchers have been concerned with. A first purpose is to diagnose the automatic nature of a task performance or a higher-level process. For example, skill-development researchers have tried to assess whether the performance on certain tasks has reached an automatic level. Affective priming researchers have tried to find out whether the affective priming effect (i.e., the fact that responses to a target are faster when preceded by a prime with the same valence; see Fazio, 2001, for a review) occurs automatically, and, by inference, whether the higher-level process of the evaluation of the primes can take place automatically. A second purpose is to explain automaticity in general. This purpose amounts to investigating which type of lower-level processes can lead to automatic performance or automatic higher-level processes, and can be rephrased as the purpose to diagnose the automatic nature of these lower-level processes. Researchers may manipulate the lower-level process that participants will use for a certain task, and they may then assess which type of lower-level process leads to automatic performance. For example, affective priming studies may be designed to encourage the retrieval of a valence label from memory (e.g., Fazio, Sanbonmatsu, Powell, & Kardes, 1986) or, alternatively, the comparison between a desired and an actual state (e.g., Moors, De Houwer, & Eelen, 2004). It may then be assessed which of these lower-level processes produces automatic affective priming effects. Both research purposes can be rephrased as being about diagnosis: the first concerns the diagnosis of the automatic nature of a task performance or a higher-level process, and the second concerns the diagnosis of the automatic nature of a lower-level process. The diagnosis of a phenomenon is usually closely related to the way in which it is defined. We therefore start with an overview of different views (of the definition) of automaticity.

VIEWS OF AUTOMATICITY

Most theories of automaticity are feature-based, defining automaticity in terms of one or more features. Different feature-based theories vary with regard to the features they select as the crucial ones, as well as with regard to the degree of coherence they assume among the features. Another proposal is to define and diagnose automaticity, not in terms of features, but in terms of the underlying process. This view is called the mechanism-based approach.

Feature-based Approach

Features have been clustered into two modes, into three (or more) modes, or they have been considered to be entirely independent. The first view to be discussed is the dual mode view. Although this view is now largely abandoned, several researchers seem to still implicitly rely on it, as is sometimes revealed in unguarded moments.

Dual Mode View According to a dual mode view, there are two modes of processing that are each characterized by a fixed set of features. Automatic processes are characterized as unintentional, unconscious, uncontrollable, efficient, and fast. Nonautomatic processes are supposed to possess all the opposite features. The dual mode view is also an all-or-none view. Such a view combines the idea of a perfect correlation among the features of each mode with the idea that both modes are mutually exclusive and that they exhaust the universe of possible cognitive processes. In this way, any performance or process holds all of the features of one, and none of the features of the other mode. According to this all-or-none view, one can diagnose a performance or process as automatic (or nonautomatic) by assessing the presence of one feature belonging to the automatic (or the nonautomatic) mode. The presence of the remaining features of that mode can then be logically inferred (Bargh, 1992).

Two historical research traditions have been appointed as responsible for the creation of the dual mode view (Bargh, 1996; Bargh & Chartrand, 2000; Wegner & Bargh 1998). The first tradition developed from the single capacity view of attention (Shiffrin & Schneider, 1977), a view that originated from early research on skill development (Bryan & Harter, 1899) and dual tasks (Solomons & Stein, 1896; see review by Shiffrin, 1988). This tradition was also inspired by the writings of James (1890) and Jastrow (1906) on habit formation. The second research tradition grew out of the New Look program in perception research (e.g., Bruner, 1957).

Capacity view The single capacity view of attention regarded attention as a limited amount of energy that can flexibly be allocated to different stages of processing (e.g., Kahneman, 1973). It was assumed that the early stages in the processing sequence (sensory analysis) require less attention (i.e., pre-attentive) than the later stages. In virtue of extensive (consistent) practice, processes that are initially capacity-demanding can progressively develop to operate without

attention. Automatic processing was defined as processing without or with minimal attention (i.e., efficient), and automatization was defined as the gradual withdrawal of attention involvement due to practice (Hasher & Zacks, 1979; Posner & Snyder, 1975a, b; Shiffrin & Schneider, 1977). We call this view the capacity view of automaticity.

Initially, proponents of the capacity view conceived of the criterion of attentional requirements as a continuum (see Hasher & Zacks, 1979), with automatic processes depleting only a minimal amount (efficient) and nonautomatic processes drawing on a substantial amount of attentional capacity (nonefficient). Other functional feature pairs (such as unintentional vs. intentional, unconscious vs. conscious, uncontrollable vs. controllable, fast vs. slow, parallel vs. serial) were derived from the feature pair efficient–nonefficient, and eventually this led to the view that automatic and nonautomatic processes represent two opposite modes of processing, each characterized by a fixed set of features. In this way, the initial conception of automaticity as a continuum developed into a dichotomous view.

New Look The second research tradition that contributed to a dual mode view was the New Look movement in perception (Bruner, 1957). The original focus in this tradition was on the constructivist nature of perception, that is, the interaction between person variables (needs, expectancies, values, knowledge) and information available in the environment (Bartlett, 1932). Because of the hidden character of the influence of person variables on perception, the focus shifted toward unconscious perception (Erdelyi, 1992). The dual mode models that developed from this research tradition put most emphasis on the features *unconscious* and *unintentional* (e.g., Fodor, 1983).

To summarize, the first research tradition took the feature pair efficient–nonefficient as a starting point and added other feature pairs to this distinction. The second research tradition added other feature pairs to a dual mode model based on the feature pair conscious–unconscious. This different emphasis on individual features of automaticity stems for a large part from the type of research paradigms employed in both traditions. For example, Shiffrin and Schneider (1977) used search tasks, which are a special type of skill development task. In these tasks, participants are explicitly instructed to engage in the process under study (e.g., to detect a target). After extended (consistent) practice, the process becomes impervious to variations in task load, and this is taken as an indication that it has become efficient (Shiffrin & Schneider, 1977). Investigators from the New Look tradition used tasks in which participants were instructed to engage in a process that is different from the process under study or in which the process under study was concealed. For example, Bruner and Goodman (1947) asked participants to draw the physical size of coins and equally sized discs and they observed an overestimation of the size of coins compared to discs. This effect was larger for coins with a higher monetary value and for poor participants. The fact that participants processed the monetary value of coins even when they were not instructed to do so, yields support for the unintentional nature of this processing. In other studies, tachistoscopic presentations were used to establish thresholds for conscious recognition of desired and undesired words (e.g., Postman, Bruner, & McGinnies, 1948).

Despite these differences, researchers of both traditions have proposed very similar dual mode models for information processing in which perception, attention, and memory are intertwined. Both consider information processing as based on the activation of a sequence of nodes from long-term memory. Nodes can be activated in two distinct ways: by stimulus input alone, in which case activation is spread further to connected nodes with little attentional demand; or by a nonautomatic process, through the allocation of attention. The dual mode model developed by these early researchers appears to be a tenacious one. Despite criticism and recent evolutions, it is still popular in various domains of research (e.g., emotion, social cognition). One of the reasons why the dual mode view seems so difficult to shake off is that it is strongly ingrained in the classic, computational metaphor of cognition on which most dual mode models rest. The computational metaphor of cognition can thus be considered as a third factor that is responsible for the creation and persistence of the dual mode view.

Computational framework There are several elements in the computational framework that render it more susceptible to a dual mode view than the connectionist framework (see also Cleeremans, 1997). First, in computational models, knowledge is represented symbolically in long-term memory. This knowledge may consist of data (concepts, exemplars) stored in declarative memory, and programs (procedures, rules, algorithms) stored in procedural memory. Processing amounts to combining data and programs in working memory. The system takes data as its input, runs a program on it, and produces new data as its output. One commonly voiced concern is that symbolic representations, because they are abstract, need an external *interpreter* in order to inject them with meaning (i.e., symbol grounding problem; Searle, 1992). Further, the conception of processing as symbol manipulation presupposes an external *manipulator* or processor. In most classic models, the interpreter (in charge with providing conscious meaning) and the manipulator (or controller) are united in one single entity, for example, a central executive (Baddeley, 1986). This central executive is also charged with directing the attention window. Hence, it is no surprise to find that the features *conscious*, *controlled*, and *nonefficient* are often mentioned in the same breath.

Not many classic models are strictly computational. The declarative database is often depicted as an associative network, and activation can travel through the network without help from the executive. Automatic processing is equated with everything that is done without the central executive. This explains why the features *unconscious*, *uncontrolled*, and *efficient* are often glued together. In the majority of the classic models, the retrieval of single facts from a declarative database is all that can be done automatically (cf. Logan, 1988). In some classic models, algorithms can also be stored according to associative principles and they can be activated and fired without the help of an external manipulator (cf. Anderson, 1996).

In connectionist models, a concept is represented subsymbolically, by a pattern of activation that is distributed among many nodes or units in a network,

instead of being tied to a single unit. Different representations show more or less resemblance in pattern activity as a function of the semantic resemblance of their concepts. In this way, the intrinsic properties of subsymbolic representations reflect the meaning of their concepts, and there is no need for an external interpreter.[1] Similarly, rules are represented subsymbolically, in the connection weights between units. The system is thus able to behave in a rule-like manner without actually possessing the rule as a symbolic representation (Cleeremans, 1997). Processing in the connectionist framework does not amount to combining data and rules, but to dynamic patterns of activation. Note that no distinction is made between representations (of both data and programs) and processing: both are defined in terms of patterns of activation. There is equally no distinction between a storage place and a working space, nor between encoding (or learning) and deployment (or processing). Each encounter with a stimulus both activates previously established patterns and increments the connection weights. Both the lack of a distinction between representations and processing, and the unanimous idea that activation can travel through the network without external help, render superfluous the intervention of a central executive. Because of the fact that connectionist models have spelled out the homunculus in charge of conscious interpretation, manipulation, and the orienting of attention, there is nothing that provides the glue for the features *conscious*, *controlled*, and *nonefficient*. In most connectionist models, it is the quality of the nodes or the quality of the activity between the nodes that determines their fate in terms of consciousness (Cleeremans & Jiménez, 2002). It is not always clear how control and attention must be shaped in a strictly connectionist system. Control is sometimes dispelled as an illusion, or it is taken care of by non-connectionist goal modules. Attention is sometimes translated in a biasing of connection weights. The idea that attention or selectivity is not a cause but rather a passive by-product of natural priming effects is an old idea (James, 1890, p. 447; Neisser, 1976; Johnston & Dark, 1986; Fernandez-Duque & Johnson, 2002).[2]

The Dual Mode View Challenged

The dual mode view came under attack when empirical studies showed a lack of co-occurrence among some of the central features listed for automatic processes (e.g., Bargh, 1992; Kahneman & Treisman, 1984; Logan, 1985, 1988, 1989; Logan & Cowan, 1984). For example, several studies demonstrated that the occurrence of the Stroop interference effect, which was generally thought to be unintentional and uncontrollable, was not independent of attention because it could be diluted when attention was directed away from the targets (e.g., Francolini & Egeth, 1980; Kahneman & Chajczyk, 1983; Kahneman & Henik, 1981). Uleman and Moskowitz (1994) found that trait inference processes were unintentional, unconscious, but not uncontrollable because they could be disrupted. Thus, it was concluded that the dual mode view was no longer tenable.

Several solutions can be proposed in response to the problems raised by these studies. A first solution is to distinguish between different types of automatic processes, resulting in a triple mode or a quadruple mode view. However, such views are likely to run into the same problems as the dual mode view. Ultimately, it may turn out that for any random combination of features (some automatic, others nonautomatic) a process can be found that fits the description (Bargh, 1992). A second solution is to dismiss the concept of automaticity and to investigate individual features of automaticity in their own right. A third, more cautious, solution is to investigate features separately but to preserve the term automatic and consider it as a gradual concept (e.g., Logan, 1985). A fourth alternative is to pick out one feature that is shared by all automatic processes and to allow processes to differ with regard to all features except the chosen one (e.g., Bargh, 1992). A fifth solution is to abandon the feature-based approach of automaticity altogether and to embrace a mechanism-based approach. Several authors have turned to one or a combination of the above solutions. In the next sections, we shed light on these solutions and evaluate some of their merits and problems.

Triple Mode View

Several theorists have proposed a triple mode view instead of a dual mode view: in addition to the nonautomatic mode, they have postulated two automatic modes. Bargh (1996; Bargh & Chartrand, 2000; Wegner & Bargh 1998) suggested that the New Look movement and the skill development tradition were each concerned with a different type of automatic process: the New Look with preconscious automaticity and the skill development tradition with goal-dependent automaticity. Both types of automatic processes possess a different sample of automaticity features. Preconscious processes are almost purely stimulus-driven. Goal-dependent automatic processes require a processing goal and a conscious input in order to get started, but once they are triggered they run off autonomously. Examples of goal-dependent automatic processes are those involved in skilled activities such as knitting and walking down the stairs.

Carver and Scheier (1999, 2002) and Moskowitz (2001) distinguished between two types of automaticity that arise at different stages in the evolution of a process. Initially, emergent bottom-up processes are elicited by stimulus input alone. These processes are so weak that they remain hidden to consciousness and therefore cannot be intentionally employed or controlled in other ways. With more repetition, these processes become consolidated in memory so that they can be recognized by a conscious processor, which may invoke them as a guide to behavior and adjust them. With still more top-down use, these processes become so well-established and accurate that top-down control is no longer required and that they drop out of consciousness again (although they remain potentially accesssible to consciousness). Other authors have distinguished between the same set of processing stages, but they have reserved the term *automatic* exclusively for processes in the third stage (e.g., Cleeremans & Jiménez, 2002).

Several attention researchers have distinguished between pre-attentive and learned automatic processes (Treisman, Vieira, & Hayes, 1992). Pre-attentive processes are those involved in the early stages of the processing sequence (e.g., sensory analysis). Learned automatic processes are those that become impervious to attentional capacity as a result of practice. Attention researchers who adhere to the distinction between both types of automatic processes have rejected the idea that the learned type eventually develops to behave as a pre-attentive type (Gibson, 1969; LaBerge, 1973). They argued that pre-attentive and learned automatic

processes differ not only with respect to the origin of their capacity independence (innate or acquired), but also with respect to certain functional properties. Logan (1992) recognized the distinction between the pre-attentive and the learned type, but he reserved the term *automatic* exclusively for the learned type.

Although the taxonomies proposed by these authors show certain resemblance, they are not entirely overlapping. Bargh's (1996) taxonomy is based on features; Carver and Scheier's (2002) taxonomy is based on stages of learning; and Treisman et al.'s (1992) taxonomy is based on stages in the processing chain. Preconscious processes do not equate with pre-attentive ones because they include both pre-attentive and learned automatic processes. Pre-attentive processes occur prior to any attentional involvement whereas learned preconscious processes may require that some spatial attention be directed at the triggering stimulus. Further, given that preconscious processes include learned ones, they also do not equate with emergent ones. Finally, emergent processes are also not identical to pre-attentive ones. The former occur in the first stages of learning, whereas the latter occur in the first stages of the processing sequence. The first stages in learning may comprise all stages of the processing sequence.

Gradual View Bargh (1989, 1992, 1997) has proposed three (instead of two) classes in which most automatic processes fit: preconscious, postconscious, and goal-dependent automatic processes. Roughly, preconscious processes require no conscious input and no intention. Postconscious processes require conscious input but no intention. Goal-dependent processes require both conscious input and an intention or another goal to get started, but no goal to run to completion. These categories should not be taken as all-or-none categories, but rather as heuristic devices. Bargh (1992) suggested that any random combination of features (some automatic, others nonautomatic) is, in principle, possible. Certain researchers went one step further and reasoned that due to the lack of co-ocurrence among features of automaticity, the concept of automaticity is inherently inconsistent and may just as well be abandoned (Pashler, 1998; Regan, 1981). Logan (1985) argued that the lack of momentary co-occurrence among features does not necessarily threaten the internal consistency of the automaticity concept. He took up the initial idea of capacity theorists that automaticity is a continuum and that the degree of automaticity is determined by the amount of practice. He also suggested that "each property has its own time-course of change with practice" (Logan, 1985, p. 373). For example, a process may evolve faster from conscious to unconscious than from nonefficient to efficient. Because in most studies the status of different features is measured at one point in practice, it is not surprising that a lack of co-occurrence among these features is observed. Only if it would be demonstrated that features evolve in opposite directions with practice (e.g., more efficient and more conscious) would the automaticity concept be jeopardized. The alternative that Logan proposed in 1985 is a gradual view: processes are automatic to a certain degree, possessing few or more features of automaticity. Such a view recommends the separate investigation of features. This view is also gradual with regard to individual features: features are present to a lower or higher degree. It should be noted that although a gradual view of automaticity may well prevent the rejection of the automaticity concept, it hollows it out at the same time. As a gradual concept, automaticity loses its ability to tell apart one type of process (automatic) from another type (nonautomatic), because any process can be labelled automatic to some degree.

Appointing one minimal feature of automaticity Bargh (1992) proposed a solution to the problem posed by the gradual view by choosing the feature *autonomous* as the one feature that is true of all automatic processes. He defined an autonomous process as one that can run to completion without conscious monitoring, irrespective of whether it was started intentionally or unintentionally. This option is theoretically appealing because it provides a criterion to separate automatic from nonautomatic processes. On the other hand, one may wonder how to justify that autonomy and not any other feature should be chosen as the critical feature. It is true that autonomy allows the greatest variety of processes to be included in the realm of automatic processes, but is that justification enough? Remember that other scholars proposed other critical features of automaticity. For example, attention researchers picked out the feature *efficient*. In a sense, the choice of a minimal feature of automaticity is an arbitrary matter.

Conditional View In addition to the idea that all automatic processes are autonomous, Bargh (1989, 1992) suggested that all automatic processes are dependent on a set of preconditions. Automaticity research should primarily be concerned with investigating these preconditions. For example, it may be examined whether a process requires the intention that it occurs (or another goal), whether a certain amount of capacity is needed, and whether the stimulus input must be consciously perceived. An assessment of these conditions permits the allocation of the features unintentional, efficient, and unconscious to the process under study. In addition to the conditions described above which are related to features of automaticity, the conditional approach also permits the investigation of other kinds of conditions, such as the salience of the stimulus material or the precise nature of the task instructions, independent of their relation with automaticity features. Investigators may also discover interdependencies between conditions. For example, it may be that a certain process does not rely on intentions provided that the stimulus input is salient enough (see later).

Mechanism-based Approach

All of the views that we discussed until now are feature-based in that they define and diagnose automaticity in terms of features. In 1988, Logan criticized the feature-based accounts arguing that they do not explain how these features were acquired. They lack an explicit learning mechanism that explains the development toward automaticity (i.e., automatization). Logan explained automatization as the transition from algorithm computation (which Logan defined as multi-step memory retrieval) to single-step memory retrieval. Nonautomatic processing is based on algorithm computation in order to produce an output. This output may subsequently be stored in memory so that future processing can bypass computation

of the algorithm and rely entirely on memory retrieval to produce the same outcome. For example, to solve arithmetic problems of the type "3+4+2=?", participants initially engage in an algorithm. In a first step, they retrieve the sum of 3 and 4, leading to the intermediate outcome 7. In a next step, they retrieve the sum of 7 and 2, leading to the final outcome 9. After sufficient practice with the same set of digits, a direct association is formed in memory between the set of digits and their sum, and as a result, participants are able to produce the final outcome in a single retrieval step when presented with the set of digits.

Logan (1988, 1991) explained automat*ization* as a shift toward single-step memory retrieval and he explained automat*icity* as based on single-step memory retrieval. Moreover, he postulated that single-step memory retrieval is the only lower-level process that can lead to automatic performance or automatic higher-level processes. From there, it was a small step for Logan to also *define* automaticity as single-step memory retrieval and to propose *diagnosing* automaticity by investigating whether the performance was based on single-step memory retrieval.

Other theorists explained automatization not as a shift from one type of process to another, but rather as a change in the *same* underlying process. They emphasized the strengthening or improvement of algorithms as the main mechanism underlying automatization (e.g., Anderson, 1992; Rosenbloom & Newell, 1986). The algorithm strengthening view is a mechanism-based view in the sense that it emphasizes one learning mechanism as responsible for automatization. It does not, however, select one type of low-level process to explain automatic performance and another to explain nonautomatic performance. The algorithm strengthening view hypothesizes that the same algorithm responsible for the nonautomatic stage of performance is also responsible for the automatic stage of performance, but in the latter stage, the algorithm is executed faster and more efficiently. Hence, automatic and nonautomatic algorithms only differ with regard to the features (fast, efficient) they possess. In sum, the algorithm strengthening view is a mechanism-based view with regard to the *explanation* of automat*ization*, but a feature-based view with regard to the *definition* and *diagnosis* of automat*icity*.

The learning mechanisms proposed by Logan (1988) and Anderson (1992) are both dependent on consistent practice. A shift toward single-step memory retrieval benefits from practice on consistent data, whereas algorithm strengthening requires algorithms to remain consistent during practice but not the data on which these algorithms operate. Anderson (1992, 1996) distinguished between a declarative and a procedural subsystem of memory, suggesting that extended practice of algorithms results in the storage of these algorithms in the procedural subsystem. The matter in discussion is thus not whether memory is involved in automaticity, but whether the system is able to store and automatically retrieve algorithms in addition to single facts.

Although Anderson's (1992) model includes other learning mechanisms such as a reduction of procedural steps to a single step, and the strengthening of declarative facts, automaticity is mainly attributed to the strengthening of algorithms. Tzelgov and colleagues (Tzelgov, Henik, Sneg & Baruch, 1996; Tzelgov, Yehene, Kotler, & Alon, 2000) advocated a reconciling approach, maintaining that both mechanisms of algorithm improvement and a shift toward direct memory retrieval are equivalent mechanisms underlying automatization.

The debate between the direct memory retrieval view and the reconciling view is a debate about which lower-level processes can lead to automatic performance (or automatic higher-level processes) (i.e., second research purpose, see above). The direct memory retrieval view predicts that only single-step memory retrieval can lead to automatic performance, whereas the reconciling view predicts that, in addition to single-step memory retrieval, there are also certain algorithms that can lead to automatic performance.

Support for algorithm-based automaticity has come mainly from skill development tasks in which it was shown that 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 consisted in an increased speed and efficiency. Recent studies using Stroop-like and priming tasks also support the unintentional (Moors et al., 2004; Tzelgov et al., 2000) nature of algorithm computation. During a practice phase, Tzelgov et al. (2000) displayed pairs of meaningless symbols, which they presented as artificial digits, and they asked participants to decide which symbol represented the larger magnitude. Due to feedback about the correctness of their decisions, participants eventually learned to make correct magnitude comparisons of these artificial digits. During a Stroop phase, participants had to indicate as fast as possible which symbol was physically larger from a pair of artificial digits (of which one was physically large, the other small) and to ignore the magnitude that was represented by the digits. Size congruity effects were obtained, with faster response times for pairs in which the physically larger stimulus also represented the larger magnitude compared to pairs in which the physically larger stimulus represented the smallest magnitude. Importantly, this effect occurred also for pairs that were not encountered (as a pair) during practice and for which no magnitude decision could have been stored. The occurrence of a size congruity effect for these pairs suggests that participants compared their magnitudes during the Stroop phase. Because making magnitude comparisons was not part of the instructions of the Stroop task, it can be concluded that participants made these comparisons unintentionally.

Conclusion

Logan (1988) proposed to replace the feature-based approach of automaticity by a mechanism-based approach equating automaticity with single-step memory retrieval. Partisans of the reconciling approach found this proposal to be unsatisfactory because it cannot integrate findings supporting algorithm-based automaticity. The choice between the direct memory retrieval view and the reconciling view is a matter of empirical evidence, which has not been settled in favor of either view yet. Because we prefer to keep all options open at this stage, we do not believe it to be fruitful to proceed from a view that excludes certain (lower-level) processes from the realm of automatic processes on an a priori basis. As a

consequence, we have no choice but to return to an approach that uses features as diagnostic criteria for automaticity, such as a gradual, decompositional approach, which proposes to investigate each automaticity feature separately to determine the degree to which it is present.

Before we can choose for a gradual and decompositional approach to the study of automaticity, two questions need to be answered. The first question is whether features of automaticity can be seen as gradual. The second question is whether these features can be conceptually and logically separated, because only then is it useful to devote effort to their separate investigation. Before addressing these questions, we will consider the definitions that authors have given for the most important features of automaticity. Our aim is to analyse these definitions and to come up with a list of central ingredients for each feature. In a second part, we examine to what extent these features can be considered as gradual. In a third part, we explore whether our analysis in central ingredients provides ways to keep features conceptually and logically separated.

DEFINING FEATURES

In the present section, we consider definitions for the goal-related features (un)intentional, goal-(in)dependent, (un)controlled/(un)controllable, and autonomous. In addition, we consider definitions for the features "purely stimulus-driven", (un)conscious, (non)efficient, and fast(slow).

(Un)Intentional and Goal-(in)dependent

Most of the thinking about intentionality has been done with respect to overt behavior. Assuming important parallels between overt acts or actions, and covert/mental acts or processes, we will use the general term "act" throughout this section (as well as throughout the sections of the other goal-related features and the feature "purely stimulus-driven"). At least three different views can be discussed under the heading of intentionality. The first two views vie for the definition of intentionality; the third view doubts that much of our behavior and processing is intentional.

Causal View The causal view of intentionality states that an intentional act is an act that is caused by an intention. This definition contains three ingredients: an intention, an act, and the causal connection between intention and act. An intention is a special kind of goal,[3] which has as its content an act and not just any state of affairs. One can intend to paint the house blue, but one cannot intend the house to be blue. In other words, an intention is a goal to engage in an act.

When both the goal to act and the act are present, but the goal is not the cause of the act, the act does not qualify as an intentional act (Davidson, 1980; Searle, 1983). For example, when Little Red Cap wants to kneel down to pick a flower for grandmother and she does kneel down but not due to her goal to do so, but rather because the wind pushes her so that she falls on her knees, then the act of kneeling down 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 intention to act causes the act, but only indirectly, by means of other intervening factors, the act cannot be classified as intentional. For example, when Little Red Cap has the goal to kneel down to pick a flower for grandmother but the thought of this plan makes her so overcome with emotion that she suddenly feels weak and sinks to her knees, then the act of kneeling down cannot be classified as an intentional act, although it was indirectly caused by the goal to kneel down.

The same objective act sequence can be intentional under one description but unintentional under another. This means that some aspects of an act are produced intentionally, whereas other aspects are produced unintentionally. One should thus specify under which description one considers or investigates an act to be intentional or not. Consider Little Red Cap walking in the woods, thereby crushing a ladybird. Supposing that Little Red Cap had the intention to walk in the woods but not the intention to crush the ladybird, her act can be labelled intentional under the description of "walking" but not under the description of "crushing a ladybird".

Aspects that are included and aspects that are not included in one's intentions may be situated at the same level of abstraction, as in the above example, or they may be at different levels of abstraction. Vallacher and Wegner (1987) argued that acts can be described at various levels of abstraction, going from low-level descriptions containing physical details to more abstract or high-level descriptions (Vallacher & Wegner, 1987). For example, "walking" occupies a higher level of description than "bending and stretching one's legs with a certain strength". An act can be intentional under the high-level but unintentional under the low-level description (or vice versa).

In addition to specifying the aspect of an act sequence that one considers, it is also important to delineate the boundaries of the act sequence itself. Strictly speaking, an act description should not include the result or consequences of the act. For example, an act described as walking in the woods does not include the arrival at grandmother's house. It is beyond doubt that Little Red Cap can have the goal to arrive at grandmother's house, but this state of arrival does not strictly belong to the act of walking in the woods. Things become more complicated if we consider that high-level descriptions of acts often include the results of acts (Wegner & Bargh, 1998). For example, the act of "walking in the woods" may be described at a higher level as "going to grandmother's house". The latter description does include the arrival at grandmother's house.

The above paragraphs illustrate an important aspect of causation spelled out by Yablo (1992) that causes and effects must be proportional. In order for an intention and an act to stand in a causal relation, the intention should be adequate and just enough for the act. The intention to go to grandmother's house is more than enough for the act of walking in the woods and is therefore not an ideal candidate cause for the act of walking in the woods. Turning it round, the intention to walk in the woods is not adequate for the act of going to grandmother's house,

because Little Red Cap may be walking in the woods without ever arriving at grandmother's house.

Other aspects of causation can be found in Mackie's (1974, 1975) proposal that a cause C of an effect E, is an **I**nsufficient but **N**ecessary part of a combination or set of conditions that is in itself **U**nnecessary but **S**ufficient for the occurrence of the effect (short, INUS condition). For example, saying that a short circuit is the cause of a house being on fire means that the short circuit, although **I**nsufficient to start the fire by itself, is a **N**ecessary part of a set of conditions (including combustible material and oxygen), that is in itself **U**nnecessary but **S**ufficient to start the fire (another set of conditions may have started the fire as well, e.g., a set in which the short circuit is replaced by the dropping of a lit cigarette). One could say that C is conditionally necessary for E, that is, necessary within this particular sufficient set (CX). This means that E does not occur when C is removed from this set; C is not unconditionally necessary for E, that is, not necessary within every sufficient set. There may be other sets in which C is also absent (e.g., KX) that do suffice for E.

An intentional act can thus be redefined as an act for which the intention to engage in the act is an INUS condition for the act. That is, the intention is a necessary part of a sufficient set of conditions for the act. The intention is not in itself sufficient for the act. Other necessary parts of the set are the stimulus input and a host of background conditions that are usually taken for granted. For example, to intentionally process a visual stimulus, the light must be on, the stimulus must be within the visual field, and the person must not be blind. There may be other sets in which the intention is not necessary that are also sufficient for the act. Many acts are intentional in some circumstances but unintentional in others. Research can thus be aimed at identifying other conditions in the set that modulate the necessity of an intention for the act.

In addition to the three ingredients of the definition of intentional acts (intention, act, and causal connection), the literature on goal striving (Gollwitzer, 1990, 1993; Heckhausen, 1991) has presented a number of conditions that must be fulfilled in order for an intention to be implemented, that is, to actually translate into an act. Given that multiple intentions compete for implementation, 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 upon the intention is required to launch the act (James, 1890; Moya, 1990; Wundt, 1902). Such a fiat can be conceived of as a second goal (or meta-goal) superimposed on the intention.

Another condition for the implementation of an intention is the feasibility of the act.[4] This condition encompasses two subfactors. First, there must be an opportunity to act. In order to implement the intention to pick a flower, there must be a flower. Gollwitzer (1990) argued that when an opportunity to engage in the act is absent, the likelihood to benefit from a future opportunity is increased when the person has specified in advance a cue that signals a suitable opportunity. For example, "next time I see a flower, I will pick it". Gollwitzer has introduced 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.[5] When such an implementation intention is formed, the risk that opportunities to act go by unnoticed is strongly reduced. A second subfactor of feasibility is skill (Heider, 1958; Mele & Moser, 1994). A person can only be expected to intentionally produce an act when he/she has the skill to perform that act. Acts that depend too much on luck, such as winning the lottery, cannot be performed intentionally.

To summarize, an act is intentional (under a certain description) when it is caused (in the right way) by the goal representation of that act (under that description). Several conditions must be fulfilled in order for the act to come about. Some of these conditions must always be fulfilled (e.g., opportunity, skill); other conditions must only be fulfilled when still other conditions are not (e.g., a fiat is required when strength is too low, and an implementation intention is required when there is no current opportunity).

An unintentional act is an act that was not caused by an intention. An act can be unintentional for two reasons: there is no intention, there is only an act; or there is an intention and there is an act, but there is no causal connection between both ingredients.

We also wish to clarify the relation between the feature *intentional* and the feature *goal-dependent*. A goal-dependent act is an act for which a goal must be in place. This goal may be either the proximal goal to engage in the act, in which case the goal-dependent act is also intentional, or a more remote goal. Thus, intentional acts can be considered as a subclass of the goal-dependent acts. Acts often depend on the combination of a remote and a proximal goal. For example, Little Red Cap may have the proximal goal to walk in the woods and this act may be instrumental for her remote goal to arrive at grandmother's house. However, proximal and remote goals can sometimes occur in isolation. Little Red Cap may have the proximal goal to walk in the woods, as a goal "an sich". Another scenario is that the remote goal to be at grandmother's house directly triggers the act of walking toward the house, without the mediation of the proximal goal to walk. In this case the act is goal-dependent but unintentional.

A final note is about the nature of the representations that cause acts in the causal view. Some theorists call them goal representations, imparting them with special dynamic qualities that are not shared by other kinds of representations, such as accumulation of strength over time and persistence in the face of obstacles (Bargh & Barndollar, 1996). Others argue that these are purely cognitive structures, and that the same representations are used both to perceive and to engage in an act (i.e., common coding hypothesis; Prinz, 1990).

Emergent Intentionality

Opponents of the causal view have rejected the trivial idea that every intentional act must be caused by an intention, that is, a (goal) representation of that act. The gist of their argument is that a person may be directed at an act without possessing a representation of that act. Wakefield and Dreyfus (1991) distinguished between two types of intentionality: R-intentionality and G-intentionality. B-intentionality stands for representation-mediated intentionality; G-intentionality stands for gestalt intentionality. As examples of G-intentional acts they mention subsidiary acts involved in skilled activities such as driving a car and playing the piano. These scholars believe it to be unlikely that

each of the subsidiary acts in a routine act are preceded by a representation of the subsidiary act. Yet, they believe there are nevertheless reasons to classify such subsidiary acts as intentional. One of those reasons is that subsidiary acts *feel* different from reflexes or Penfield motions (i.e., motions caused by electrode firings in the motor cortex; Penfield, 1975). Intentional acts are accompanied by the feeling that one is the causal agent of one's acts and that these acts are appropriate to the situation.

One could, however, argue that the feeling of self-agency that one experiences with subsidiary acts stems from the fact that the overall act was intentional in a representation-mediated sense and that this produced a general feeling of self-agency. In our view then, G-intentionality corresponds to what we have previously called goal-dependence. Also, the idea that subsidiary acts can be preceded by representations should not be rejected on a priori grounds. Opponents of the causal view seem to rely on a narrow meaning of the word representation, taking representations to be always conscious and conceptual in nature. However, when one accepts that goal representations may also be unconscious (Bargh, 1990) or non-conceptual (e.g., image-like), the idea that subsidiary acts can be preceded by goal representations seems more plausible. In the end, the question of whether subsidiary acts of routine acts are intentional (in a representation-mediated sense) or unintentional but goal-dependent is an empirical question that cannot be solved by clever examples.

Emergent Behavior Another view regarding intentionality is inspired by the dynamic systems framework. According to this framework, much (if not all) of our behavior emerges from bottom-up activity that satisfies multiple constraints and there is no need for top-down intentions (e.g., Kelso, 1995; Vallacher & Nowak, 1999; see also Carver & Scheier, 2002). Many acts, even complex ones, are not produced by intentions (representations of these acts) but are the result of multiple interacting forces combined with the unspecific tendency operating within all living systems toward increased coherence. These researchers do not question so much the causal definition of intentionality, but rather that much behavior and processing "out there" actually meets the requirements of that definition.

A similar view is endorsed by proponents of the connectionist framework. Whereas dynamic systems models reject the existence (and hence the causal role) of representations altogether, connectionist models postulate the existence of subsymbolic representations. Many connectionists argue that acts are not caused by intentions, which are symbolic representations, but rather by underlying subsymbolic representations. In a similar vein, Wegner (2003; Wegner & Wheatley, 1999) argues that intentions do not have causal power but that there are underlying processes that give rise to both the felt intention and the act. According to Wegner, the folk idea that our intentions cause our acts is nothing but an erroneous attribution.

The question of which level has causal powers is a much debated issue in the philosophical literature. According to the commonsense view, acts are caused by mental states such as intentions, whereas according to the reductionist view all causal powers are situated on the lowest level of analysis. We think that the position that one adopts in this debate is highly dependent on the way in which one frames the question.

Certain authors have framed it as a problem of *causation* (e.g., Green & Vervaeke, 1996). According to the dominant view of causation, causes and effects must be situated on the same level of analysis and the cause must precede the effect. Causes must also be INUS conditions (Mackie, 1974), and they should be proportional to their effects (Yablo, 1992). Intentions and acts meet these formal criteria and can therefore be considered as potential causes and effects. The subsymbolic representations in connectionist models also precede acts, but they are situated at a lower level of analysis than most acts, and they include details that are not proportional to most acts (Green & Vervaeke, 1996). Subsymbolic representations have therefore been considered to be less ideal candidate causes of acts.

Other theorists have framed the problem as a problem of *causal explanation*. The dominant view of causal explanation is the reductionist view, which argues that the ultimate causal explanation is situated on the lowest level (for an alternative, see Menzies, 2001). Connectionists argue that the causal explanations of acts are to be found on the lower, subsymbolic level. This view can be considered as a variant of the reductionist view, although most reductionists would not stop at the subsymbolic level but go further down the ladder to neurology and finally physics. A strong reductionist view can be found in the dynamic systems framework.

(Un)Controlled/(Un)Controllable

In the previous section, we have delineated the features *intentional* and *goal-dependent*. The current section deals with the feature *controlled* and its relation to the feature *intentional*. The concept "control" has three ingredients in common with the concept "intentional": a goal and causation resulting in an effect. For instance, Wegner and Bargh (1998) posited that, at the most basic level, to control something is to influence it in a certain direction rather than in a random manner. Dennett (1984) likewise argued that control is not causation alone, but rather desired causation: A controls B if A can drive B to be in a state that A wants B to be in. Put differently, A controls B if A has a goal about B and this goal is achieved (i.e., the state represented in the goal occurs). We can thus say that a person controls an act if the person has a goal about the act and the goal is achieved. An act is controlled when a person controls it and uncontrolled when a person does not control it. An act is controllable when it is possible for a person to control and uncontrollable when it is impossible for a person to control.

As Dennett (1984) pointed out, A can only drive B into a state that belongs to B's normal range of states. Examples of normal-range goals about acts are the goals to alter, stop, avoid, and engage in the act. The goal to engage in the act can be parsed into the goals to start, continue, and complete the act. Given the multitude of goals that persons can pursue about acts, it is best to always specify which goal one envisions when using the terms control, (un)controlled, and (un)controllable.

Given the definition that an act is intentional when it is caused by the goal to

engage in the act, and the definition that an act is controlled when the person has a goal about the act and this goal is achieved, we must conclude that intentional acts are a subclass of controlled acts. An intentional act is controlled in the sense of the goal to engage in the act. Other subclasses of controlled acts are those that are controlled in the sense of the goals to change, stop, or avoid the act. Table 1.1 gives an overview of the definitions of the features *goal-dependent*, *intentional*, and *controlled*. As explained above, intentional acts are a subclass of goal-dependent and a subclass of controlled acts.

The type of control discussed here corresponds to what engineers have attributed to open-loop systems (see also Wegner & Bargh, 1998). In open-loop systems there is a criterion or goal, and an action directed at attaining that criterion. Control in such a system is a fact when the criterion is reached, irrespective of whether the system has feedback about whether the criterion is reached. This type of control is contrasted to control in closed-loop systems. In closed-loop systems, the result of the action is fed back into the system (so that the loop is closed) and it is tested in a comparison module whether the criterion is reached. The action either ceases, when the criterion is reached, or it is repeated until the criterion is reached (Carver & Scheier, 1999; Miller, Gallanter, & Pribram, 1960). Feedback (and a comparison module to match it to the criterion) is an essential ingredient of control in closed-loop systems.

According to Powers (1973), another central ingredient of control is disturbance: A controls B if A influences B in a desired direction, thereby counteracting the effect of a disturbing influence on B. For example, Little Red Cap controls her act of walking down the path when she manages to stay on the path despite attempts of the wolf to lead her astray. Keeping in mind that there are other, more complex types of control, we will focus here on the simplest type, which is control as can be found in open-loop systems and that consists of only three ingredients: a goal about an act (the cause), achievement of the goal (the effect), and the causal relation between the two elements.

An act is controlled when a goal pertaining to the act (to engage in, avoid, alter, or stop the act) is a necessary part within the sufficient set of conditions for the effect (occurrence, prevention, change, or interruption of the act). By opposition, when an act is uncontrolled, it may mean several things. First, it may mean that (a) a goal pertaining to the act is absent, in which case the effect may (a1) also be absent, or (a2) be present, but due to another cause. Second, it may mean that (b) a goal pertaining to the act is present, but (b1) the effect is absent, or (b2) the effect is present but it is not caused by the goal. In case (b1), the set of conditions including the goal is not sufficient for the effect; in case (b2), the sufficient set of conditions does not include the goal as a necessary part. In the next sections, we define the features *autonomous* and *"purely stimulus-driven"*.

TABLE 1.1 Definitions for the Features *Goal-dependent*, *Intentional*, and *Controlled*

Feature	Cause	Effect
Goal-dependent	Goal (remote=any goal/proximal = goal to engage in the act)	The act
Intentional	Goal to engage in the act	The act
Controlled	Goal about the act (stop/alter/avoid/engage in the act)	State represented in the goal (end/change/prevention/ occurrence of the act)

Autonomous The word autonomous (deriving from the Greek words *auto* = self + *nomos* = law) literally means self-governed, not controlled by outside forces. When used as a predicate of the word *person*, it means that the person determines his/her own fate, and not someone else. When used as a predicate of the word *act*, however, it means that nothing outside of the act controls the act. A completely autonomous act is an act that is uncontrolled in the sense of every possible goal. More precisely, an act is autonomous when it is (a) unintentional and (b) uncontrolled in the sense of counteracting goals. That is, (a) the set of conditions sufficient for the occurrence of the act does not include the goal to engage in the act as a necessary part, and (b) insertion of the goals to alter, stop, or avoid the act in the current set does not result in the change, interruption or avoidance of the act.

For the sake of completeness, we need to mention that the definitions of autonomous given by other authors are somewhat different from our own. For example, Bargh (1989, 1992) defined an autonomous process as a process that, once started, runs to completion without the need for conscious guidance or monitoring. This definition is different from our own in at least two respects. First, Bargh defined autonomy in terms of the redundance of *conscious* goals. This means that a process can still be termed autonomous when it is influenced by *unconscious* goals. In our definition of autonomous the term *goal* covers both conscious and unconscious goals. Second, in Bargh's definition, autonomy has to do only with the second part of the process, after the process has started. In his account, an autonomous process may begin intentionally or unintentionally, but after its start no further monitoring is required. Our definition of an autonomous process (or act) as one that is uncontrolled in the sense of every possible goal also includes the goal to start the process (or the act). The difference between both accounts is subtle. Whereas Bargh's definition suggests that a process that was intentionally started but not intentionally continued and completed is an entirely autonomous process, we suggest that it is a partially autonomous process (only the second part of the process is autonomous, i.e., not controlled by outside forces).

Bargh's (1992) definition of autonomous bears close resemblance to Logan's (Logan & Cowan, 1984, p. 296; Zbrodoff & Logan, 1986) definition of ballistic. Logan defined a ballistic process as one that cannot be inhibited once it begins and so must run to completion. Note, however, that Bargh's definition of autonomous speaks about the goals to *continue* and *complete*, whereas Logan's definition of ballistic speaks about the goal to *stop*. A process that does not require a goal for its continuation and completion is not always impossible to stop.

Purely Stimulus-driven An act is purely stimulus-driven when it is caused by the mere presence of the triggering stimulus. By assigning the stimulus as the

cause (i.e., INUS condition) of an act, we are not saying that there are no other necessary conditions in the sufficient set for the act. We take the usual background conditions (e.g., that the light is on and the person's eyes are open in case of the processing of a visual stimulus) for granted. What we are saying is that awareness of the input, attention to the input, goals, and the outcomes of other acts are not necessary conditions in the sufficient set of conditions for the act.

According to our definitions, there is no complete overlap between the features *purely stimulus-driven* and *unintentional*. Purely stimulus-driven acts count as a subclass of unintentional acts. Other subclasses of unintentional acts do not require an intention to occur but they do require other conditions such as other goals or awareness of the input, and hence they are not purely stimulus-driven.

(Un)Conscious Several authors have emphasized that the features *conscious* (or aware) and *unconscious* (or unaware) are not unitary concepts. Two often-mentioned ingredients of consciousness are Intentional and phenomenal aspects.

The word Intentional is used here in the philosophical sense, which is different from intentionality in the ordinary sense (cf. former sections we capitalize the former but not the latter, following Searle's convention, 1983). The common core in both terms is "directedness" (deriving from the Latin verb *intendere*: to be directed at). In philosophy, Intentionality 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, remembering, and intending are Intentional because they require an object. In ordinary language, intentionality derives from the mental phenomenon of intending or having an intention and is a property of actions (e.g., Wakefield & Dreyfus, 1991). In our view the essential difference between Intentional and intentional is not that the former relates to mental phenomena and the latter to actions or acts, for when the category of acts is extended to include mental acts or processes the distinction between both concepts does not vanish. What we do believe to be crucial is that with Intentionality it is the mental act itself that is directed at some state of affairs outside the mental act, whereas with intentionality it is a person who is directed at engaging in the act. Formally, directedness can be considered as a 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 the mental act, whereas with intentionality, the source is a person and the terminus is the act. Otherwise put, Intentionality speaks about a mental state or act, whereas intentionality speaks about the way in which an act is produced (or about the source of the act, i.e., the person).

The phenomenal aspect is also known under the names of immediate qualia (Tolman, 1932), the qualitative aspect (Natsoulas, 1981), or the "what it is like" aspect of awareness (Nagel, 1994). The difference between the Intentional and the phenomenal aspect of awareness can be illustrated by the difference between knowing that one is hurt versus feeling pain.

Block (1995) issued the distinction between two types of consciousness: phenomenal consciousness (P-consciousness) and access consciousness

(A-consciousness). This distinction is partially founded on the distinction between the two aspects of consciousness discussed above: P-conscious states have a phenomenal content; and A-conscious states have a content that is accessible for use in reasoning and in rational control of behavior and/or verbal report. Contents are accessible by virtue of their representational or Intentional format, but they are actually accessed when attention is directed to 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. This means that contents that are outside of the focus of attention and not A-conscious may nevertheless be P-conscious. For example, the noise of a pneumatic drill in the street or the ticking of a clock may be P-conscious all along, but not A-conscious unless attention is directed to it. Turning it round, contents may be A-conscious without being P-conscious. Block cited the hypothetical case of superblindsight to illustrate this. Patients suffering from "regular" blindsight report a lack of P-consciousness but are nevertheless able to make successful guesses about presented stimuli. Because these guesses are not considered rational, they are not classified as demonstrations of A-consciousness. In the hypothetical case of superblindsight, however, it is conjectured that patients know (that is, rationally) what stimuli have been presented, not via visual P-conscious experience, but in much the same way as a thought that suddenly pops to mind. A-consciousness without P-consciousness is also imputed to the virtual case of a philosophical zombie. A zombie is a creature that is functionally equivalent to a human being but with one difference: all is dark inside (Chalmers, 1996), which means that there is no P-consciousness. Although both illustrations for A-consciousness without P-consciousness are non-empirical, Block (1995) took them as support for the idea that the two types of consciousness can be separated on a conceptual level. He did acknowledge, however, that they go together under normal circumstances, and he even conceded that they may interact. Attention can alter the content or flavor of experience (see also Hardcastle, 1997). For example, there is a phenomenal difference between contents that are in and those that are outside the focus of attention (as in figure and ground).

Much of the conceptual debate is centred on Block's (1995) assumption that the two types of consciousness are independent. A first alternative states that P-consciousness is a precondition for A-consciousness (e.g., Schacter, 1989). According to this position, information must be represented in phenomenal consciousness before it can be accessed for use in voluntary behavior. A second alternative, termed epiphenomenalism, holds the reverse: only those contents that enter the executive system, and hence are A-conscious, can become P-conscious. Such assimilation of P-consciousness to A-consciousness probably stems from the desire to measure P-consciousness, together with the fact that measurement is based on verbal report or behavior. These methods are more suitable to measure A-consciousness than P-consciousness. A third alternative, called the collapse hypothesis, states that the distinction between the two types of consciousness is artificial and that P- and A-consciousness are one and the same thing (e.g. Baars, 1995; Kobes, 1995; Morton, 1995).

Two additional types of consciousness that Block described are monitoring

consciousness (M-consciousness) and self-consciousness (S-consciousness). To be M-conscious is to be conscious of the fact that one is P- or A-conscious, and to be S-conscious is to be conscious about oneself. Several authors have made a distinction between only two types of consciousness, called, for example, minimal and full consciousness (Wegner & Bargh, 1998) or Type 1 and Type 2 experience (O'Shaughnessy, 1991). The first type corresponds to some combination of A- and P-consciousness, and the second type corresponds to M-consciousness.

Apart from the fact that the terms conscious and unconscious have different ingredients or come in different varieties, both terms may also be used as a predicate of different things. They may pertain to (a) the stimulus input that evokes a process, (b) the output of a process, (c) the process itself, and (d) the consequences of the process, such as its influence on subsequent processes (see also Bargh, 1994). It is best to specify which of these elements one has in mind when using the terms conscious and unconscious. This is important because there are many otherwise unconscious processes that require a conscious input in order to operate (cf. class of postconscious and goal-dependent processes, Bargh, 1992; see also Bargh & Barndollar, 1996). Note that the consequences of a process do not strictly belong to the critical process, but rather to a subsequent processing unit.

Keeping in mind that processes can be described at different levels of analysis (cf. Marr, 1982), we should point out that the only level that seems in principle to be accessible to consciousness is the functional level. Most theorists would probably agree that everything below that level is structurally inaccessible. All that the functional level of process understanding requires is that the person is able to link an output to an input. For example, a person may be aware that a certain advertisement influenced his/her decision to buy a product, without being aware of the formal characteristics of this process or the neural circuits involved. Further, we need to be more precise 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, not the raw stimulus input. It is important to distinguish between potentially accessible and structurally inaccessible parts because it only makes sense to invest effort in the investigation of the conscious or unconscious nature of parts that are potentially accessible to consciousness.

(Non)Efficient The concept efficient has been defined in relation to limited attentional capacity or processing resources. An efficient process is a process that uses only a minimal amount of this capacity. The amount of available capacity is depicted as a continuum and hence the feature *efficient* is easily regarded as a gradual feature. Because efficient processes demand only minimal attentional capacity, they are likely to remain unaffected when resources are scarce due to the high demand of other concurrent processes. Degree of capacity consumption is therefore 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, which is assumed to consume a certain amount of resources. The processes involved in the primary task are efficient when varying (and especially, augmenting) the amount of capacity used by the secondary task has no impact on performance of the primary task (e.g., Hermans, Crombez, & Eelen, 2000). Because it is impossible to determine whether a particular secondary task exhausts all processing resources, the processes involved in the primary task can never be shown to be totally independent of resources.

The use of dual task paradigms is founded on the single resource view (or single-capacity view) of attention and especially on a number of assumptions inherent to that view. One of these assumptions is that the resource requirements for the simultaneous execution of two processes equal the sum of the individual resource requirements for each process. Another assumption is that the total amount of available resources remains constant under different conditions of secondary task load (see Navon & Gopher, 1979, for a complete list of assumptions). To the extent that these assumptions can be challenged, however, the usefulness of the dual task method can be called into question. From the multiple resource perspective, which postulates that the cognitive system has more than one resource and that different tasks may tap from different resources, the dual task paradigm is only useful when both tasks call on the same processing resource (Navon & Gopher, 1979; Wickens, 1984).

In the wake of the capacity view of attention, attention has been conceptualized as a limited source of energy, much like a barrel filled with liquid or a tank of gaz. Efficiency has been defined as the ability of a process to do without (or to do with very little of) this limited energy. However, besides quantity, attention may also be said to have a direction, as is embodied in another metaphor, that of attention as a spotlight (Baars, 1998; LaBerge, 1983; Posner, Snyder, & Davidson, 1980; Wundt, 1897). Both quantity and direction are aspects of attention that go together. However, there may be processes that depend on the specific direction of attention but at the same time use only a small amount of attention. For example, Stroop interference effects have been shown to diminish when relevant and irrelevant dimensions are spatially separated (e.g., Kahneman & Henik, 1981). The amount of capacity used by the processes tapped in these Stroop experiments may have been so small that these processes still deserve the label efficient.

Attention may be directed at the process (the relation between input and output) or, alternatively, at the 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.

Fast(slow) The concept *fast* is also a gradual feature. This feature has sometimes been ignored in lists of automaticity features, but the speedup of task performance is taken as an important indicator of automaticity in search studies (Shiffrin & Schneider, 1977) and other skill development studies (Smith & Lerner, 1986). Although the feature *fast* can only be used as a predicate of a process, it is useful to distinguish between the duration of the process and the duration of the stimulus input on which the process operates (see later).

GRADUAL CONCEPTUALIZATION OF FEATURES

A first question that we set out to examine is whether features are gradual. Only when they are gradual does it make sense to choose a gradual approach. A gradual feature is a feature that can be present to a certain degree. We already pointed out that the features *efficient* and *fast* are gradual features. Although less immediately obvious, the feature *(un)conscious* and the goal-related features can also be conceptualized as gradual. A gradual conception of consciousness is a natural position for connectionist models that consider the transition from unconscious to conscious cognition to be a matter of strength (cf. Cleeremans & Jiménez, 2002). It is more problematic for computational models that depict conscious and unconscious cognition as two separate systems or routes. Empirical evidence is mixed: certain studies support a gradual transition from unconscious to conscious perception (Bar et al., 2001), whereas others support a sharp non-linear transition (Sergent & Dehaene, 2004).

Goal-related features can also be conceptualized as gradual. Indeed, the correspondence between a processing goal and the outcome that is actually obtained can be less than perfect. We should point out that this is only true in cases where a goal is present. The presence of a goal was assumed in the definition of the features *intentional* and *controlled*, as well as in one type of unintentional and one type of uncontrolled process, in which only the causal connection between the goal and the effect is missing. In cases where there is no goal (there is only an effect), the features *unintentional* and *uncontrolled* are not gradual.

In addition to a gradual approach to the *conceptualization* of features, we also favor a gradual approach to the *measurement* of features. We accept support in favor of automaticity features even if this support is not infallible evidence. However, we do argue that empirical support for features of automaticity should be related to the ingredients of the features. For example, arguments for the intentional nature of a process may be stronger or weaker, but they have to be related to the ingredients of this feature, that is, the concepts *goal* and *causation*.

DECOMPOSITIONAL APPROACH

A second question that we set out to examine is to what extent it is possible to draw a conceptual distinction between features. This is an important issue because the separate investigation of features is only worthwhile when features can be conceptually separated. In the previous parts, we analyzed the most important features of automaticity and isolated their ingredients. Overlap is inevitable when features are defined in such a way that they share central ingredients. Apart from conceptual overlap, there are other types of relations between features that have implications for research. A certain feature can be a necessary and/or sufficient condition for another feature. We will refer to this type of relation as a modal relation. Modal relations have led authors to logically infer the presence or absence of one feature on the basis of the presence or absence of another feature.

We first discuss the conceptual overlap between features. After that, we look at assumptions of modal relations.

Conceptual Overlap

Overlap among Goal-related Features (and Purely Stimulus-driven) In the section about the definition of features, we already pointed out that there is partial overlap among several sets of goal-related features. Intentional processes are a subclass of the goal-dependent processes and a subclass of the controlled processes. Unintentional processes and autonomous processes are both subclasses of uncontrolled processes. Purely stimulus-driven processes are a subclass of unintentional processes. (For a more extended description of the conceptual overlap among goal-related features, see Moors & De Houwer, 2006). On the other hand, we do not assume overlap among different types of controlled processes that are based on different types of goals.

The definitions of certain other researchers, however, seem to lead to somewhat different conclusions. The features *(un)intentional* and *(un)controlled* are often presented as non-overlapping, whereas the features *unintentional* and *purely stimulus-driven* are often regarded as completely overlapping. For example, in Bargh's (1994) account, the feature *intentional* has to do with being in control of the start of a process, whereas the feature *controllable* has to do with the ability to alter or stop a process once it is 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 stop a process. These goals differ with regard to the part in the time course of the process (first vs. second part) as well as with regard to their content (promoting vs. counteracting goals). At first glance there seems to be no overlap between the features *intentional* and *controllable* in Bargh's definitions. However, a closer look reveals that his definition of *intentional* makes reference to the word *control* ("the feature intentional has to do with being in control of the start of the process"; see above). When one feature is defined in terms of another, then overlap is a fact. Another point of difference is that in our definitions, the word "intentional", which is shorthand for "intentionally produced", is not restricted to the start of the process but covers the whole process from start to completion. We stated that the goal to engage in a process can be parsed into the goals to start, continue, and complete a process. It is not unthinkable that a person would have the goal to merely start the process, or merely complete it, but in these cases we would call the process an intentionally started or an intentionally completed process rather than an entirely intentional or intentionally produced one.

Some researchers treat the features *unintentional* and *purely stimulus-driven* as completely overlapping. This is because they assume that processes generate from a two-factor universe of possible causes inhabited by stimuli and intentions. We have argued to take into account other possible causes such as remote goals, outputs of other processes, consciousness, and attention. The next paragraph deals with conceptual overlap between the goal-related features and consciousness.

Overlap among Goal-related Features and the Feature Conscious

Goal-related features involve a causal relation between a goal (to engage in, alter, or stop a process) and an effect (a process, a change, or an interruption of a process). No mention was made of the requirement that this goal be conscious. If one accepts that goals can also be activated in an unconscious manner, there is no need to suppose overlap between goal-related features and consciousness (cf. Bargh, 1990; Bargh, Gollwitzer, Lee-Chai, Barndollar, & Trötschel, 2001; Chartrand & Bargh, 1996; see review by Moskowitz, Li, & Kirk, 2004). Despite this, the terms intentional and controlled are often used as shorthand for the notions "consciously intentional" and "consciously controlled". Some theorists even argue that consciousness is an ingredient of the features intentional and controlled (e.g., Oatley, 1993). These theorists are likely to infer a lack of control on the basis of a lack of consciousness.

Modal Relations

In addition to conceptual overlap, investigators have also assumed modal relations among features. When feature A is a necessary condition for feature B, then evidence that A is absent allows one to logically infer that B is absent as well. When feature A is a sufficient condition for B, then evidence that A is present allows for the logical inference that B is present as well. These relations have contributed to the clustering of features and hence to the persistence of the all-or-none view of automaticity.

Many investigators have formally abandoned the all-or-none view, and grant that the relations between goals, attention, and consciousness are quite complex. However, in discussions of empirical results regarding the diagnosis of automaticity, they often seem to rely on a very simple set of one-to-one relations between these elements. First, we discuss what has almost been considered the holy trinity of goals, attention, and consciousness (see Figure 1.1 for a schematic overview). After that, we discuss why it is problematic to use this simple set of relations to draw logical inferences about features.

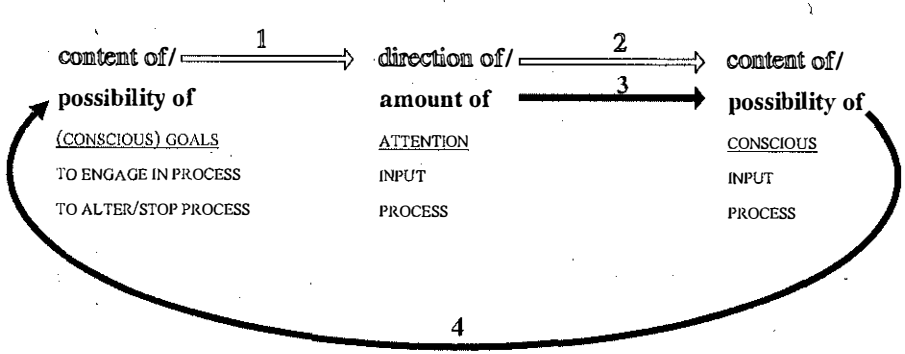


FIGURE 1.1 Simple set of interrelations between goals, attention, and consciousness.

Simple Set of Modal Relations

From goals to attention The relation between goals and attention is widely accepted. One scenario that connects goals to attention goes as follows. When a person has set the goal to engage in some process, this goal must still be implemented. To implement a goal, a suitable occasion is required, for example, in the form of an appropriate stimulus input. Attention may be conceived of as a tool that helps the organism to incorporate suitable input to launch the process. For instance, to implement the goal to evaluate a stimulus, there must at least be a stimulus in sight. If the stimulus is not in sight or not optimally so, the person may direct attention or zoom in on the stimulus (or on some aspect of the stimulus). In sum, goals determine where attention will be directed. To implement the goal to engage in a process, attention must be directed to an appropriate stimulus input; to implement the goal to alter/stop the process, attention must be directed to the process itself (Figure 1.1, Step 1).

From attention to consciousness Another close team is attention and consciousness. The common view is that the direction of attention determines the content of consciousness (Figure 1.1, Step 2). Whatever is in the focus of attention is conscious, everything else is unconscious (e.g., Baars, 1998; Merikle & Joordens, 1997; Posner, 1982, 1994; Shiffrin & Schneider, 1977; but see Baars, 1997).

Thus far, we stated that goals determine where attention will be directed (Step 1), and that the focus of attention, in turn, determines the content of consciousness (Step 2). By way of transitivity then, goals also determine the content of consciousness (Step 1 + Step 2).

Further, the *amount* of attention spent (to input or process) determines whether consciousness (of this input or process) will occur, or (for gradual models of consciousness) the amount or degree of consciousness that will occur (Step 3).

From consciousness to goals Many authors have suggested that consciousness is a necessary condition for the features *intentional* and *controlled* in the sense of alter/stop (Figure 1.1, Step 4). When stimulus input is conscious, it permits the setting and implementation of the goal to engage in the process. When the process itself is conscious, it permits the setting and implementation of the goals to alter or stop the process. This idea is truly pervasive (e.g., Bargh, 1992, 1994, 1999; Cleeremans & Jiménez, 2002; Dehaene & Naccache, 2001; Dienes & Perner, 1999, Uleman, 1989; Shallice, 1988). Let us reconsider Carver and Scheier's (2002) description of the evolution of a process. When an initially weak, emergent automatic process grows stronger due to repetition, it can be recognized by a conscious processor, who can use and manipulate it. With more top-down use, the process becomes so accurate that top-down control eventually becomes superfluous, and as a result the process drops out of consciousness (conform to James', 1890, parsimony principle that "consciousness deserts all processes when it can no longer be of use"). Both phases in the evolution of a process rest heavily on the assumption that it is awareness of the process that allows a person to control it. It is duly noted that the terms "intentional" and "controlled" are probably used here in the sense

of "consciously intentional" and "consciously controlled" (see section on conceptual overlap; see also below).

Until now, we stated that the amount of attention (to input or process) determines the possibility of consciousness (of input or process) (Step 3), and that consciousness, in turn, determines the possibility of the setting and implementation of goals (to engage in or to alter/stop a process) (Step 4). By way of transitivity then, the amount of attention also determines the possibility of the setting and implementation of goals (Step 3 + Step 4). We can thus say that goals not only determine where attention will be *directed* (Step 1), but that the *amount* of attention available also determines whether information (input or process) will be detected (Step 3) and hence whether it can be used for the implementation of goals (to engage in or to alter/stop the process) (Step 4).

Problems with the Simple Set of Modal Relations There are two problems that complicate the customary practice to generalize support from one feature to the next. The first problem is that the picture of interrelations is more complex than what is just described. The second problem is related to the gradual nature of the features. We discuss these two problems in turn.

Empirical research and theoretical analysis have extended the simple set of interrelations presented above and we describe a number of these additional relations below (see Figure 1.2 for a schematic overview).

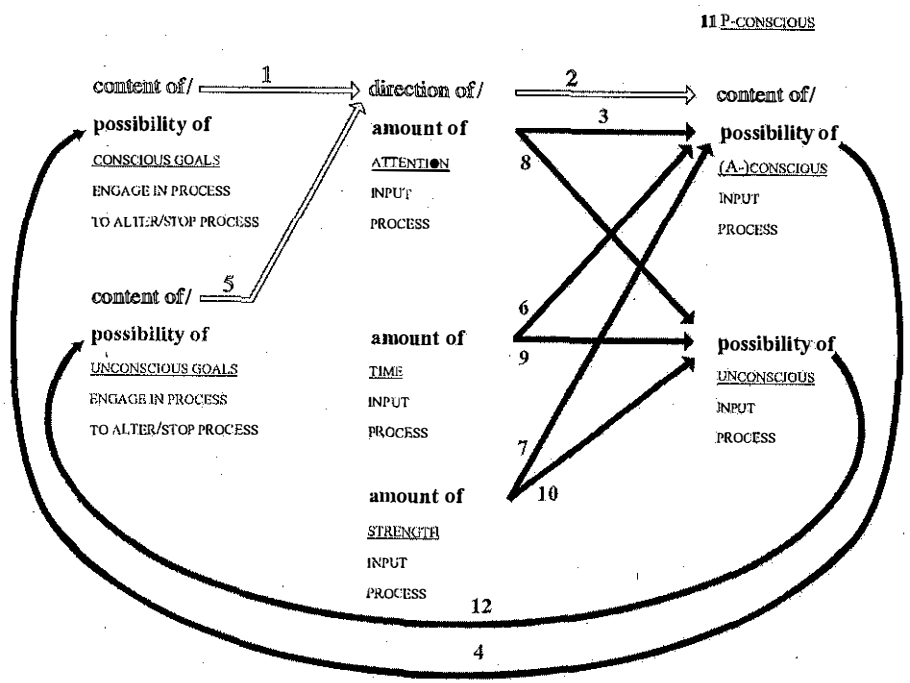


FIGURE 1.2 Extended set of interrelations between goals, attention, time, strength, and consciousness.

From goals to attention Attention may not only be directed by conscious goals, but also by unconscious goals (Figure 1.2, Step 5). The distinction between an attention system guided by conscious goals and one guided by unconscious goals turns up in several theories (e.g., Baars, 1998; Luria, 1973; Posner & Rothbart, 1989; Shiffrin & Scheider, 1977; Yantis & Jonides, 1996).

There is considerable controversy about whether goals are necessary for attention orienting. Some authors have argued that goal relevance, particularly the discrepancy of a stimulus with a goal or *desired* state, is a necessary condition for attention orienting. Other authors have argued that attention orienting may also be elicited by novelty, or the discrepancy between the 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 mere novelty of the stimulus, there is no accompanying conscious experience of orienting or directing. It feels as if the stimulus grabs one's attention. This is reflected in use of the terms attention attraction, attentional grasp, orienting reflex (Sokolov, 1963), and "automatic" attention response (Shiffrin & Schneider, 1977). Of course, unless they are animate, stimuli cannot "do" something; they are simply there.

We should add that attention must not always be *directed* in order to be in a certain place. The attention spotlight may randomly travel around the visual field until it crosses goal-relevant or novel stimuli and it may then remain focused on these stimuli. Technically speaking, there is no *directing* of attention in this scenario, but there is nevertheless *allocation* of attention to a stimulus (Fox, Russo, Bowles, & Dutton, 2001).

From attention to consciousness Several researchers argue that attention is not a sufficient condition for consciousness. The possibility or the degree of consciousness (of input or process) is not only determined by the amount of attention (to input or process) (Figure 1.2, Step 3), but also by other variables such as the amount of time or duration (of input or process) (Step 6), and the amount of strength (of input or process) (Step 7). When a stimulus is too weak or occurs too briefly so that activation of the corresponding nodes remains below a certain threshold, the stimulus may not be accessible to consciousness, despite adequate focusing of attention. This is illustrated by the fact that in subliminal perception research, stimuli are presented very briefly (in combination with masks) or in degraded form.

Certain authors argue that attention is not exclusive for consciousness. Some amount of attention, time, and strength may not only be required for conscious access (to input or process), but also for unconscious "access" (of input or process) (Steps 8–10). For example, Naccache, Blandin, and Dehaene (2002) observed that subliminal priming effects disappeared when attention was not focused on the time window in which the prime–target pair was presented. These observations are consistent with older studies in which subliminal priming effects disappeared when primes were not presented in the same position as the targets but rather in a "flanker" position (e.g., Inhoff, 1982). There also appear to be lower limits to the strength and presentation duration of stimuli that can lead to subliminal perception. We should note that attention, strength, and time are only a few of the many

parameters that have been studied in subliminal perception research (for other parameters, see Verleger, Jaskowski, Aydemir, van der Lubbe, & Groen, 2004).

Most theorists assume that the *direction* of attention determines the content of consciousness: whatever is in the focus of attention is conscious; everything else is unconscious (e.g., Baars, 1988; Posner, 1982, 1994; Shiffrin & Schneider, 1977). Other theorists argue that attended information does not exhaust our conscious experience (e.g., Block, 1995; Hardcastle, 1997; Nelkin, 1993). Recall that Block (1995) coupled A-consciousness with attention, but saved P-consciousness as a form of consciousness that escapes attention (cf. his example of the pneumatic drill) (Step 11). Note that if P-consciousness is independent of the *direction* of attention, it should also be independent of the *amount* of attention. One could therefore say that attention is not even a necessary condition for consciousness.

From consciousness to goals We think that the necessity of consciousness (of input and process) for the setting and implementation of the goals to engage in or to alter/stop a process (Figure 1.2, Step 4) may be true only when these goals are understood in the sense of *conscious* goals. Although it seems reasonable to assume that consciously intentional processes require a conscious input, this may not be true for unconsciously intentional processes. It is quite possible that unconsciously intentional processes can operate on unconscious input as well (Step 12). Similarly, in order to alter or change a process in a conscious way, it seems plausible that the person must be at least aware that the process is taking place. However, unconscious altering or stopping of a process may also be possible when processes are operating below the surface (e.g., Glaser & Banaji, 1999). In sum, the setting and implementation of unconscious goals (to engage in or alter/stop a process) may not only be possible on the basis of conscious, but also on the basis of unconscious material (input or process).

The picture of interrelations is undoubtedly more complex than what we just sketched, but the main point is to show that the relations between features are not one-to-one relations and hence that arguments for the presence or absence of automaticity features based on logical inferences are likely to be flawed.

A second problem for the customary practice to generalize support from one feature to the next is posed by the gradual nature of the features. Let us consider the relations between the amount of attentional capacity, time, and strength on the one hand, and the possibility of consciousness on the other (Steps 3, 6, & 7). One should be able to conclude that awareness of a stimulus is prevented when attentional capacity, stimulus duration, and stimulus strength are too low. A major difficulty, however, is to determine exactly how much of these factors is needed to ensure awareness of the stimulus. This is further complicated by the fact that awareness of a stimulus can be realized with only a small amount of capacity and time (especially when an implementation intention is formed; cf. Gollwitzer, 1999). Thus, the process that uses conscious input may still count as efficient and fast. Another complication is the possibility of interdependence between these factors. For example, lack of time to incorporate a stimulus input may be compensated by increased attention to this stimulus or increased strength (salience) of the stimulus. Thus, it is dangerous to conclude on the basis of short presentation duration alone that some input was not consciously accessed (and therefore not processed in a consciously intentional way).

CONCLUSION

In the first part, we justified our choice for a feature-based approach to the definition and diagnosis of automaticity. In particular, we advocated a gradual and decompositional approach. Such an approach argues to examine each automaticity feature separately to determine the degree to which it is present. To see whether this gradual, decompositional approach can be carried through, we examined whether features are gradual and whether they can be conceptually and logically separated. Before we were able to answer these questions, we engaged in a conceptual analysis of the most important features: (un)intentional, goal-(in)dependent, (un)controlled, autonomous, purely stimulus-driven, (un)conscious, (non)efficient, and fast(slow). For each feature, we specified the minimal ingredients. In this way, different types could be distinguished for certain features. For example, we distinguished between two types of consciousness (A-consciousness vs. P-consciousness), each of which possess a different ingredient (Intentional vs. phenomenal aspect). We also distinguished between different types of control (open-loop vs. closed-loop; with vs. without disturbance; conscious vs. unconscious), each of which possess a different set of ingredients. We further specified what the features (or the ingredients of the features) can be a predicate of: the input, the output, the critical process (relation between input and output), or the initial perceptual process (which operates on the raw stimulus input and produces an output that is used by the critical process as an input).

This analysis led to the conclusion that features can to a large extent be defined in a nonoverlapping way. We could not avoid partial conceptual overlap between several goal-related features. On the other hand, because we argued for a separation of the different goals (engage in, alter, stop, avoid) that are involved in different types of controlled processes, we believe that our set of definitions is truly consistent with a decompositional approach. Other theorists have been less concerned with a rigorous conceptual separation of features, and assumptions of overlap seem to be the rule rather than the exception. In addition to conceptual overlap, many investigators assume modal relations among features, treating one feature as a necessary and/or sufficient condition for another feature. By drawing a more complex picture of interrelations, we showed that logically inferring the presence or absence of one feature on the basis of the presence or absence of another feature is not warranted. A decompositional approach is thus not only possible, but also desirable.

We also examined whether the features in our list can be conceptualized as gradual, and we concluded that this is indeed possible. The good news is that a gradual approach is justified. The bad news is that we have no objective criterion to draw a line between the automatic and the nonautomatic member in each feature pair: every process is efficient, uncontrolled, unconscious, and fast to some extent. The best solution that we can think of at present is to choose a subjective

criterion to draw this line. In fact, most investigators are already doing this. When investigators say that a certain process is efficient, they implicitly rely on a subjective criterion to call the process efficient. We invite investigators to make explicit the subjective criteria they are using. This would greatly improve the communication about features. Examples of subjective criteria to call a process efficient are: "process A is more efficient than process B", "process A is more efficient than what is usually thought about the process", or "process A is more efficient than before practice". In fact then, our approach is not only a gradual but also a relative approach. The term "relative" involves more than the term "gradual": it involves a standard or criterion for comparison. Once a standard for comparison is specified, it becomes possible to draw a line between nonefficient and efficient processes (and between processes characterized by other feature pairs). The general term automatic may be kept as an umbrella term, as long as it is specified in what sense a process is automatic, that is, which features apply to the process and to what degree.

NOTES

1. Some authors have argued that connectionist models are no more grounded than their computational counterparts, and that they are prone to problems proper to any kind of correlational semantics (Christiansen & Chater, 1992).
2. Different meta-theoretical views exist about the relation between the computational and connectionist framework. According to a first view, both frameworks propose different metaphors of cognition that are incompatible. Eliminativists argue to replace the metaphor of symbol manipulation with the metaphor of pattern activation (e.g., Ramsey, Stich, & Garon, 1991; Smolensky, 1988). According to a second view, both frameworks pertain to a different level of analysis (as described by Anderson, 1987; Marr, 1982). Implementationalists have argued that connectionist models deal with a lower or more fine-grained level than computational models. In principle, processes described at different levels are not expected to make different predictions, and hence connectionist and computational models are not expected to compete with each other. However, there is disagreement about the level that has causal powers (Smolensky, 1995; see Green & Vervaeke, 1996), or about the level that should be of concern to psychologists (Anderson, 1987). Such disagreements rest on the assumption that the different levels that both frameworks occupy are more than just different descriptions of the same thing. Computational implementationalists argue that the higher level adds something extra which is missing at the lower level. For example, Fodor and Pylyshyn (1988) argue that the processes proposed by connectionist models are not in themselves sufficient to explain certain characteristics of human cognition. Connectionist implementationalists argue that the causal powers are situated on the lowest level. Still others propose genuinely hybrid models (e.g., Harnad, 1990).
3. The word goal can be understood in two senses. In the first sense, the word goal indicates a desired state in the future. This meaning is apparent in the phrase "to be directed toward a goal". One way in which a person can be directed at something is by forming a representation of it. In the second sense then, the word goal indicates a representation of the desired state, which, when activated, acts as an antecedent cause that directs the organism toward that state. This meaning is apparent in the phrase "to be driven or caused by a goal".
4. In this context, we can point at the distinction commonly made in the motivation literature between goal setting (the phase in which a goal is chosen) and goal striving (the phase in which the goal is implemented) (Gollwitzer, 1990; Heckhausen, 1991). It should be clear that the present analysis only deals with goal striving. The feasibility factor that we discuss should thus not be confused with the feasibility factor mentioned in the literature on goal setting (cf. Gollwitzer, 1990).
5. This link is initially established in a deliberative manner, but once it is stored in memory, the detection of the cue can trigger the act directly.

REFERENCES

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

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

Anderson, J. R. (1996). ACT: A simple theory of complex cognition. *American Psychologist, 51*, 355–365.

Baars, B. J. (1988). *A cognitive theory of consciousness*. Cambridge, UK: 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, UK: Clarendon Press.

Bar, M., Tootell, R. B., Schacter, D. L., Greve, D. N., Fischl, B., Mendola, J. D., Rosen, B. R., & Dale, A. M. (2001). Cortical mechanisms specific to explicit visual recognition. *Neuron, 29*, 529–535.

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.

Bargh, J. A. (1990). 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.

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.

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.

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.

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., 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, UK: Cambridge University Press.

Bernstein, A. (1969). To what does the orienting response respond? *Psychophysiology, 6*, 338–350.

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.

Brentano, F. C. (1874). *Psychologie vom empirischen Standpunkt*, Leipzig: Duncke & Humblot.

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review, 64*, 123–152.

Bruner, J. S., & Goodman, C. C. (1947). Value and need as organizing factors in perception. *Journal of Abnormal Social Psychology, 42*, 205–215.

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. (1999). Themes and issues in the self-regulation of behaviour. In R. S. Wyer (Ed.), *Advances in social cognition* (Vol. 12, pp. 1–105). Mahwah, NJ: Erlbaum.

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. (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.

Christiansen, M. H., & Chater, N. (1992). Connectionism, meaning and learning. *Connection Science, 4*, 227–252.

Cleeremans, A. (1997). Principles for implicit learning. In D. Berry (Ed.) *How implicit is implicit learning?* (pp. 195–234). Oxford, UK: Oxford University Press.

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, UK: Psychology Press.

Davidson, D. (1980): *Essays on actions and events*. Oxford, UK: Oxford University Press.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition, 79*, 1–37.

Dennett, D. C. (1984). *Elbow room: The varieties of free will worth wanting*. Cambridge, MA: MIT Press.

Dienes, Z., & Perner, J. (1999). A theory of implicit knowledge. *Behavioral and Brain Sciences, 22*, 735–808.

Erdelyi, M. H. (1992). Psychodynamics and the unconscious. *American Psychologist, 47*, 784–787.

Fazio, R. H. (2001). On the automatic activation of associated evaluations: An overview. *Cognition and Emotion, 15*, 115–141.

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.

Fernandez-Duque D., & Johnson, M. L. (2002). Cause and effect theories of attention: The role of conceptual metaphors. *General Psychology, 6*, 153–165.

Fodor, J. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

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

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 automaticty of "automatic" activation: Evidence of selective seeing. *Perception and Psychophysics, 27*, 331–342.

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.

Glaser, J., & Banaji, M. R. (1999). When fair is foul and foul is fair: Reverse priming in automatic evaluation. *Journal of Personality and Social Psychology, 77*, 669–687.

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.

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. Lubek (Eds.), *Problems of theoretical psychology* (pp. 201–210). North York, ON: Captus University Publications.

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, 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: Springer-Verlag. (Original work published 1990.)

Hermans, D., Crombez, G., & Eelen, P. (2000). Automatic attitude activation

and efficiency: The fourth horseman of automaticity. *Psychologica Belgica, 40*, 3–22.

Inhoff, A. W. (1982). Parafoveal word perception: A further case against semantic preprocessing. *Journal of Experimental Psychology: Human Perception and Performance, 8*, 137–145.

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

Jastrow, J. (1906). *The subconscious*. Boston, MA: Houghton-Mifflin.

Johnston, W. A., & Dark, V. J. (1986). Selective attention. *Annual Review of Psychology, 37*, 43–75.

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*. 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.

Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: MIT Press.

Kobes, B. W. (1995). Access and what it is like. *Behavioral and Brain Sciences, 18*, 260.

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.

LaBerge, D. (1973). Attention and measurement of perceptual learning. *Memory & Cognition, 1*, 268–276.

LaBerge, D. (1983). Spatial extent of attention to letters and words. *Journal of Experimental Psychology: Human Perception and Performance, 9*, 371–379.

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. (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.

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.

Logan, G., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review, 91*, 295–327.

Luria, A. R. (1973). *The working Brain: An introduction to neuropsychology*. New York: Basic Books.

Mackie, J. L. (1974). *The cement of the universe: A study of causation*. Oxford, UK: Clarendon Press.

Mackie, J. L. (1975). Causes and conditionals. In E. Sosa (Ed.), *Causation and conditionals*. Oxford, UK: Oxford University Press.

Marr, D. (Ed.). (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: Freeman.

Merikle, P. M., & Joordens, S. (1997). Parallels between perception without attention and perception without awareness. *Consciousness and Cognition, 6*, 219–236.

Mele, A. R. & Moser, P. K. (1994). Intentional action. *Noûs, 28*, 39–68.

Menzies, P. (2001). 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.

Miller, G. A., Galanter, E., & Pribram, K. (1960). *Plans and the structure of behavior*. New York: Holt.

Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin, 132*, 297–326.

Moors, A., De Houwer, J., & Eelen, P. (2004). Automatic stimulus–goal comparisons: Support from motivational affective priming studies. *Cognition and Emotion, 18*, 29–54.

Morton, A. (1995). Phenomenal and attentional consciousness may be inextricable. *Behavioral and Brain Sciences, 18*, 263–264.

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*. Hillsdale, NJ: Erlbaum.

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, UK: 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., & 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, UK: Blackwell.

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, UK: Blackwell.

O'Shaughnessy, B. (1991). Searle's theory of action. In E. Lepore & R. Van Gulick (Eds.), *John Searle and his critics*. Cambridge, MA: Blackwell.

Pashler, H. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.

Penfield, W. (1975). *The mystery of the mind*. Princeton: Princeton University Press.

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., & 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.

Posner, M. I., & Snyder, C. R. R., (1975a). 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., (1975b). 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., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General, 109*, 160–174.

Postman, L., Bruner, J. S., & McGinnies, E. (1948). Personal values as selective factors in perception. *Journal of Abnormal Social Psychology, 43*, 142–154.

Powers, W. T. (1973). *Behavior: The control of perception.* Hawthorne, NY: Aldine DeGruyter.

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: Springer-Verlag.

Pylyshyn, Z. W. (1984). *Computation and cognition: Toward a foundation for cognitive science.* Cambridge, MA: MIT Press.

Ramsey, W., Stich, S. P., & Garon, J. (1991). Connectionism, eliminativism, and the future of folk psychology. In W. Ramsey, S. P. Stich, & D. E. Rumelhart (Eds.), *Philosophy and connectionist theory* (pp. 199–228). Hillsdale, NJ: Lawrence Erlbaum.

Regan, J. E. (1981). Automaticity and learning: Effects of familiarity on naming letters. *Journal of Experimental Psychology: Human Perception and Performance, 7*, 180–195.

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.

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: Cambridge University Press.

Searle, J. R. (1992). *The Rediscovery of the Mind.* Cambridge, MA: MIT Press.

Sergent, C. & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science, 15*, 720–728.

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, UK: Oxford University Press.

Shiffrin, R. M. (1988). Attention. In R. C. Atkinson, R. J. Hernstein, 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.

Smith, E. R., & Lerner, M. (1986). Development of automatism of social judgements. *Journal of Personality and Social Psychology, 50*, 246–259.

Smolensky, P. (1995). Constituent structure and explanation in an integrated connectionist/symbolic architecture. In C. MacDonald & G. MacDonald (Eds.), *Connectionism: Debates on psychological explanation* (pp. 223–290). Oxford: Blackwell.

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.

Tolman, E. C. (1932). Concerning the sensation quality: A behavioristic account. *Psychological Review, 29*, 140–145.

Treisman, A., Vieira, A., & Hayes, A. (1992). Automaticity and preattentive processing. *American Journal of Psychology, 105*, 341–362.

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.

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.

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

Vallacher, R. R., & Nowak, A. (1999). The dynamics of self-regulation. In R. S. Wyer (Ed.), *Advances in social cognition* (Vol. 12, pp. 241–259). Mahwah, NJ: Erlbaum.

Vallacher, R. R., & Wegner, D. M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review, 94*, 3–15.

Verleger, R., Jaskowski, P., Aydemir, A., van der Lubbe, R. H. J., & Groen, M. (2004). *Journal of Experimental Psychology: General, 133*, 494–515.

Wakefield, J. and Dreyfus, H. (1991). Intentionality and the Phenomenology of Action. In E. Lepore & R. Van Gulick (Eds.), *John Searle and his critics.* Cambridge, MA: Blackwell.

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: Wilhelm Engelmann. (Original work published 1896.)

Wundt, W. M. (1902). *Grundzüge der physiologischen Psychologie* (Vol. 1). Leipzig: Wilhelm Engelmann.

Yablo, S. (1992). Cause and essence. *Synthese, 93*, 403–449.

Yantis, S., & Jonides, J. (1996). Attentional capture by abrupt onsets: new perceptual objects or visual masking? *Journal of Experimental Psychology: Human Perception and Performance*, 22, 1505–1513.

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.

2

Effects of Priming and Perception on Social Behavior and Goal Pursuit

AP DIJKSTERHUIS, TANYA L. CHARTRAND, and HENK AARTS

Let's say you pick up a rock and you throw it. And in midflight you give that rock consciousness and a rational mind. That little rock will think it has free will and will give you a highly rational account of why it has decided to take the route it's taking . . . (From *I Am Charlotte Simmons* (p. 283) by Tom Wolfe (2004), published by Jonathan Cape. Reprinted by permission of the Random House Group Ltd.)

If the editor would have asked us to write about automaticity in social behavior 25 years ago, he would have met a blank stare. And we do not mean that we would have been puzzled because we were kids at the time, more interested in playing with Lego than in writing chapters. The whole concept of automatic or unconscious behavior would have struck anyone as odd at that time. Sure, we do not need consciousness to breathe or perhaps not even to brush our teeth, but automaticity in behavior that social psychologists are traditionally interested in? Automaticity in social interaction? Automatic goal pursuit? Of course not! We are conscious agents and automaticity was reserved for relatively unimportant routine actions only. It would have become a very short chapter indeed.

Needless to say, things have changed. The idea of the sovereign conscious agent determining our behavior received many blows over the years from social cognition (and other) researchers. We learned that impressions of people are influenced by contextual factors of which we are unaware (Bargh & Pietromonaco, 1982; Higgins, Rholes, & Jones, 1977; Srull & Wyer, 1979). Subsequently, we were confronted with the fact that social stereotypes are activated automatically (Devine, 1989; Macrae, Milne, & Bodenhausen, 1994). Only just recovered from that major setback, we had to accept that changes in social behavior and intellectual performance could be completely automatic (Bargh, Chen, & Burrows, 1996; Dijksterhuis & van Knippenberg, 1998). And finally, we had to accept the fact that even goal-directed behavior was largely under "unconscious