Automatic and Controlled Processes in Pragmatics

Marco Mazzone

Abstract In utterance understanding, both personal and sub-personal aspects appear to be involved. Relevance theory (starting from Sperber and Wilson 1986/1995) and Recanati (2004) have respectively explored two alternative ways to conceive of those aspects and their interaction. Here a third account is proposed, in the light of the automatic-controlled distinction in psychology, and of recent views concerning the cooperation between these two modes of processing. Compared to Recanati (2004), the account proposed here assigns a larger role to automatic, associative processes; at the same time, it rejects the view that consciousness applies only to what Recanati calls secondary pragmatic processes. Consciousness is rather held to cooperate with associative processes in any aspect of pragmatic processing, irrespective of the pragmatic distinction between explicatures and implicatures. On the other hand, a close consideration of how associative and conscious processes plausibly interact makes it appear unnecessary the hypothesis of a specialized process for utterance understanding—such as the automatic, inferential mechanism put forth by Relevance theory.

Keywords Pragmatics · Automatic process · Controlled process · Schema · Consciousness · Working memory

I would like to thank Gina Kuperberg, whose lectures in Catania have opened new perspectives for my research on automatic and controlled processes, and Robyn Carston, whose incisive comments to a talk I gave in Leipzig helped me to acknowledge how crucial it was to address the issue at the core of this paper.

M. Mazzone (✉)
University of Catania, Catania, Italy
e-mail: mazzonem@unict.it

A. Capone et al. (eds.), *Perspectives on Linguistic Pragmatics*, Perspectives in Pragmatics, Philosophy & Psychology 2, DOI: 10.1007/978-3-319-01014-4_18,

1 Introduction

Pragmatic processing seems to be at the same time an automatic and a personal-level affair. Humans produce and understand utterances in context quite rapidly and effortlessly, just as it is expected to occur in automatic processing; nonetheless, verbal communication is thought to require an intentional involvement on the part of the speaker, and a recognition of this speaker's intention—possibly requiring rational capacities—on the part of the addressee.

However, it is far from immediately clear how pragmatic theory should account for the coexistence of the two different features considered above; this is probably one of the major challenges that current cognitive pragmatics has to cope with. There are two main explicit attempts to address that issue within pragmatic theories of language understanding: one is Relevance theory, the other is Recanati's (2004) framework.[1] The latter solve the problem by proposing a two-level model, where automatic associative processes and personal-level, inferential processes cooperate in explaining language comprehension in context. Relevance theorists assume instead that, in a sense, a single mechanism can account for both the aspects considered. Sperber and Wilson (1986/1995) in fact conceive of language comprehension as based on a process which is said to be unconscious and automatic, but nonetheless endowed with features that are normally attributed to personal-level processes: it would be meta-representational, and inferential rather than merely associative.

In this paper I intend to propose a different way to account for coexistence of automatic and personal-level features in pragmatic understanding. My proposal preserves Recanati's intuition that two different kinds of processes are involved, while accepting Relevance theorists' criticisms to Recanati with regard to his claim that these processes apply to distinct domains of pragmatic phenomena. The account I propose is based on the well-established distinction between automatic and controlled processes in psychology, and especially on the recent literature which emphasizes the constant cooperation between the two in most of our cognitive processes. Although Recanati's account is consistent with this framework to the extent that he conceives of a cooperation between automatic and conscious processes in language understanding, his claim that these processes apply to different pragmatic phenomena has no ground in linguistic and psychological evidence. In other words, the opposition between automatic and conscious processes does not seem to parallel any traditional distinction within the domain of pragmatic phenomena—such as the distinction between (the processes involved in the

[1] For sure, there are other pragmatic frameworks that could be worth discussing in this context. However, the two I have chosen are amongst the most complete and explicit attempts to analyze the overall cognitive architecture of pragmatic processing. Elsewhere I extend my analysis to other theories in the field by addressing the topic of default interpretations (Mazzone, 2013a). In particular, in that paper I address the positions of Bach, Levinson, Jaszczolt, and also Capone's (2011a, b) interesting proposal of a largely associative perspective on defaults and modularity in pragmatics.

determination of) explicit and implicit sense of utterances. On the other hand, a close consideration of the plausible division of labour between automatic and controlled processes suggests also reasons to reject Relevance theorists' proposal that a single process, conceived of as both automatic and inferential, may account for pragmatic understanding.

In practice I will proceed in the following way. First, in Sect. 2, I will survey the distinct ways in which Recanati and relevance theorists propose to combine automatic and personal-level components of utterance understanding. Section 3 will be devoted to the distinction between automatic and controlled processes in psychological and neuroscientific literature, and to a scrutiny of the collaboration between these processes, along the lines of the "distributed intentionality model" put forth by Mazzone and Campisi (2013). In that model purely associative, automatic processes play a large role, although in cooperation with conscious processes. In Sect. 4 I will analyze RT's and Recanati's positions towards associative processes in pragmatics, arguing that these processes have a key role to play in that domain too (as I argue at greater length in Mazzone 2011). Finally, in Sect. 5 I will analyze RT's and Recanati's positions towards conscious processing, and I will consider how controlled processes are to complement associative processes in order to deliver a complete account of language understanding.

2 Recanati and Relevance Theory

In the recent debate between Relevance Theory (from now on, RT) and Recanati (2004) with regard to the architecture of pragmatic processing, a key role is played by the notion of inferential process. As is well known, at the core of Grice's (1989) theoretical framework there is the distinction between two layers of utterance meaning: what is said and what is implicated by an utterance—respectively referred to in the recent literature as the explicit and implicit sense of the utterance. While the former was essentially thought to depend on the linguistic information conveyed by the utterance (except for minor appeals to context in order to obtain reference assignments and disambiguations), Grice conceived of the transition from explicit to implicit sense as a sort of rational inference requiring consideration of the current goals of the speaker. Although Grice himself insisted that the enterprise he was engaged in was a matter of rational reconstruction rather than a genuine psychological thesis about the processes involved in comprehension, in one form or another his appeal to inferential processes performed by a rational agent has framed recent cognitive accounts.

The most straightforward manifestation of this influence is Recanati's (2004) conception of what he calls secondary pragmatic processes. Recanati is between those who think that Grice had underestimated the role played by context in determining explicit meaning: in Recanati's opinion, explicit meaning is the result of pragmatic processes just as implicit meaning is. However, he thinks Grice was

right in pointing at a major difference between those two layers of meaning, or more precisely, between the processes leading to them: while the processes yielding the explicit meaning (in Recanati's terms, primary pragmatic processes) should be thought of as sub-personal, associative processes, implicit meaning would be instead the result of genuine inferential processes taking place at the personal level (secondary pragmatic processes). Let us address this proposal in some more detail.

In Recanati's (2004) account, primary pragmatic processes are conceived of as local associative processes, based on the spreading of activation within conceptual networks and the consequent degree of activation of concepts in the network. In other words, a concept would be contributed to the explicit content of the utterance insofar as that concept is the most accessible (i.e. the most activated) for the system given the situation. In practice, the literal meaning of an expression:

> is accessed first and triggers the activation of associatively related representations. That literal meaning is a natural candidate for the status of semantic value, but there are others: some of the representations activated by association contribute further candidates for the status of semantic value. All candidates, whether literal or derived, are processed in parallel and compete (Recanati 2004, 28).

Although literal meanings are said to have an initial advantage over other possible candidates, this cannot imply of course that literal meanings—or, more generally, concepts endowed with an initial advantage—always win the competition. Recanati (2004) emphasizes the importance of what he calls "accessibility shifts": in the course of processing, contextual information may change the accessibility of any concept activated previously, by adding a new train of activation to the process. According to Recanati, a key role in accessibility shift is played by abstract schemata coded in our long term memory. For a very simple example (see Carston 2007), let us consider the following utterance:

(1) I'm going to the bank now to get some cash.

Since there are two possible meanings for ''bank'' (FINANCIAL INSTITUTION, RIVER SIDE), one problem is how the subject may come to choose the right one. Let us suppose that, for whatever reason, at the moment when the lexical form "bank" is processed the most accessible meaning is the wrong one (RIVER SIDE). However, we can expect an accessibility shift as soon as the word "cash" is processed, since this word activates its meaning, which in turn triggers a number of concepts having to do with money, and this presumably provides further activation to the concept of bank as financial institution. In particular, the activation of CASH could recall an abstract schema—in Carston's (2007) terms, a stereotypical frame or script—for GETTING MONEY FROM A $BANK_1$ (where $BANK_1$ = FINANCIAL INSTITUTION), thus strengthening the activation of $BANK_1$. In this way, schemata drive the interpretation process by promoting the search for coherence, due to an entirely associative mechanism: on the one hand, "a schema is activated by, or accessed through, an expression whose semantic value corresponds to an aspect of the schema"; on the other hand, the "schema thus activated in turn raises

the accessibility of whatever possible semantic values for other constituents of the sentence happen to fit the schema" (Recanati 2004, 37).

Once explicit content has been recovered by means of associative processes, Recanati proposes that a quite different process leads to the determination of implicit content. In Carston's (2007, 2)[2] words, "secondary pragmatic processes are to be understood as part of a more general theory of human action and interpretation and so having the philosophically central property of being rational, personal-level (as opposed to subpersonal) processes". In practice, secondary pragmatic processes are said to be "transparently or consciously inferential" in that they satisfy the "availability condition" (Recanati 2004, 44): they are accessible to consciousness, that is, the subject is aware of what is said, of the implicature, and of the inferential process leading from the former to the latter.

For an example, let us consider the following question–answer pair:

(2) (A) Could you pay back the money you owe me?
(B) I'm going to the bank now to get some cash.

In this context, (B)—which repeats (1)–can be interpreted as a positive answer to the yes–no question (A). The explicit content of (B), in which the concept $BANK_1$ is fed thanks to associative processes, licenses a further contextual inference to this layer of implicit meaning. Recanati's assumption is that both the explicit and the implicit content are available to consciousness, in that the former conforms "to the intuitions shared by those who fully understand the utterance" (idem, 14), and the latter follows inferentially from the former insofar as the expectation for a yes–no answer raised by the previous question is taken into consideration.

It is important to emphasize that the conscious availability here appealed to is qualified by Recanati as tacit rather than fully explicit—or, to put it differently, dispositional rather than occurrent. He recognizes that conscious processes are typically effortful, slow and under voluntary control, while comprehension processes normally are not. However, although the inferences involved in comprehension cannot be conceived of as conscious in this explicit, occurrent sense, according to Recanati they are not even the sort of sub-personal inferences that are merely ascribed to a cognitive system on the grounds that the system behaves in the same way as someone who performed the relevant inferences in an explicit form (Recanati 2004, 49). An inference can also be consciously available in a tacit, dispositional sense when "the cognitive agent to which it is ascribed […] is *itself* capable of making the inference explicitly and of rationally justifying whatever methods it spontaneously uses in arriving at the 'conclusion'" (idem, 50). Therefore, in Recanati's opinion, some inferences are merely sub-personal, some are conscious in the prototypical sense—they are explicit, slow and effortful—while some others are personal-level and yet only dispositionally conscious. As it

[2] Here and below the page numbers refer to the online version of the paper: http://www.phon.ucl.ac.uk/home/robyn/Carston-Recanati-22August05%5B2%5D.pdf

should be clear, sub-personal inferences are—so to speak—inferential only in the eyes of an observer. This is the case with Recanati's primary pragmatic processes: they are thought to be merely associative processes although they may nonetheless "mimic" inferential processes (Recanati 2007). On the other hand, secondary pragmatic processes are claimed to be genuinely inferential, conscious processes at least in a dispositional sense.

This cognitive version of Grice's inferential account of comprehension differs from RT's proposal on two major points. First, relevance theorists assume that one single mechanism is sufficient to account for utterance understanding: that is, an unconscious, automatic process based on expectations of relevance. In particular, they believe that the whole distinction between primary and secondary pragmatic processes, respectively conceived of as sub-personal and personal processes, is not grounded: consciousness is not thought to play a significant role in normal episodes of comprehension. Second, the single process by which RT explains comprehension is conceived of as both automatic and yet genuinely inferential in its own right. Although this process is said to be unconscious and outside the control of the subject, relevance theorists describe it as a non-demonstrative inference that takes a set of premises as input and yields a set of conclusions as output. This means that in their account the inferences in terms of which pragmaticists reconstruct utterance understanding are literally part of the automatic process of comprehension: inferences are neither merely attributed to the subject (in particular, comprehension is not based on mere associative processes that just mimic inferences, as in Recanati's primary pragmatic processes), nor are they something that the subject is just capable of delivering explicitly if necessary (as in Recanati's dispositional account of secondary pragmatic processes). They are instead genuine, occurrent inferences, although automatic ones.

In general terms, in RT's account of communication an utterance conveys a presumption of its own relevance, and the hearer has to construct a hypothesis about the speaker's meaning which satisfies that presumption of relevance. This requires constructing appropriate hypotheses about explicit content, intended contextual assumptions, and implicated conclusions—with explicit content and contextual assumptions counting as premises from which implicated conclusions are to follow. Although explicit content provides one of the premises for the inference, this does not mean that it is wholly determined by means of a previous non-inferential process—for instance, by means of a purely associative process as in Recanati's model. In fact, the whole process is conceived of as circular rather than uni-directional: hypotheses about the implicated conclusions might be suggested straightforwardly by some contextual cues, so that those conclusions can contribute to determine the premises which are apt to draw the inference. In this sense, Relevance theorists speak of a "mutual adjustment" between explicit content, contextual assumptions and implicated conclusions. Therefore, the very same process based on the construction of inferential derivations is believed to be responsible for the determination of both explicit and implicit content.

In short, neither simple associative processes nor conscious, controlled processes are claimed to play any significant role in utterance understanding. RT

conceives of comprehension as a quite specific process: an automatic inferential process which is specialized for the purposes of communication. According to Carston (2007), reflective reasoning may well play a role in communication and comprehension but only with regard to rational reconstruction of spontaneous pragmatic processes: "this is not an exercise that people perform much off their own bat. Its most likely role is as a backup mechanism when something goes wrong with the automatic intuitive mechanisms of utterance understanding" (idem, 31).

Before we address the respective weaknesses of RT and Recanati's account, let us turn to the distinction between automatic and controlled processes: this distinction may form the basis for a different account of how sub-personal and personal features may coexist in comprehension.

3 Automatic and Controlled Processes

The view that human cognition involves two different types of processing, automatic and controlled, is a well-established theme in psychology at least since the writings of William James (1890). The issue has received renewed attention in the last decades after the seminal studies of Shiffrin and Schneider (1977) and Schneider and Shiffrin (1977)—see Schneider and Chein (2003) for a recent overview. Schneider and Shiffrin (1977) define an automatic process as the activation of a sequence of nodes in our knowledge representation that "nearly always become active in response to a particular input configuration" and that "is activated automatically without the necessity for active control or attention by the subject" (idem, 2). In contrast, controlled processes are defined as "a temporary sequence of nodes activated under control of, and through attention by, the subject" (idem, 2–3).

The standard tests employed to assess whether a process of interest is automatic or controlled are subliminal presentation and techniques based on cognitive load (see Satpute and Lieberman 2006, 91). The fact that an input is processed subliminally, that is, outside awareness, is treated as the most distinctive feature of automaticity. On the other hand, it is also expected that whenever the process of interest is automatic, it will not be influenced by load manipulations. This is because automatic processes are thought to occur in parallel, in contrast with controlled processes which operate serially, resulting in task-switching costs. Another important manifestation of the automatic-controlled distinction is that "extended consistent training is required in order to develop automatic processing, while controlled processes can be established in a few trials and under varied mapping conditions" (Schneider and Chein 2003, 528). As a consequence of extended training and parallel processing, automatic processes are fast and accurate. On the contrary, controlled processes are typically slow and inaccurate.

3.1 Old and New Approaches to the Automatic/Controlled Distinction

In the last decades, the basic distinction just outlined has been framed in largely similar ways by different scholars (Carver and Scheier 2009): intuitive versus conscious in Smolensky (1988), associative versus rule based in Sloman (1996), reflexive versus reflective in Shastri and Ajjanagadde (1993). Moreover, two-mode, or dual-process, models of functioning have emerged in personality psychology (experiential versus rational system: Epstein 1973), and in social psychology (reflexive versus reflective: Lieberman et al. 2002; impulsive versus reflective: Strack and Deutsch 2004; Strack et al. 2009).

In some of those developments of the distinction, the emphasis is on the fact that the second kind of process operates on "symbolic, or propositional structures" (Saptute and Lieberman, 2006, 88), enables symbolic logic (Lieberman et al. 2002), can be simulated by symbolic architectures allowing the binding of variables (Schneider and Chein 2003, 532)—in contrast to automatic processes which can be simulated by simple connectionist networks. This shift towards the sub-symbolic/symbolic distinction is pushed to the point that in some cases consciousness is no more considered distinctive of the controlled type of process. For instance, Lieberman et al. (2002) distinguishes between an X-system for reflexive processes and a C-system for reflective processes, where the former results itself in a state of consciousness: the X-system is said to be a parallel-processing, sub-symbolic, pattern-matching system that produces the continuous stream of consciousness that we experience as the world out there.

As useful as it can be for some theoretical purposes, this way to recast the automatic/controlled distinction is probably misleading for various reasons. First, it is hardly coherent with the most accepted way to assess automaticity, that is, by means of tests based on subliminal processing: these tests precisely aim to ascertain whether a cognitive process occurs outside consciousness. Second, at least for a crucial class of cognitive phenomena, by claiming that a process is controlled (versus automatic) scholars mostly intend to emphasize its conscious and voluntary nature, irrespective of whether it is a symbolic (propositional) process or not. This is the case with a large amount of research in neurophysiology, where the issue is at which conditions bodily movements become actions, that is, they are under conscious control (versus merely automatic: for instance, see Jeannerod 2006; Pacherie 2006). From this point of view, neurophysiological literature on intentional bodily movement is just an instance of a larger category: that of research on intentional action in general, which also includes linguistic behavior. In research on intentional action it is consciousness rather than propositionality that is held to distinguish controlled from automatic processing. Third, by downplaying the importance of the conscious/unconscious distinction one loses the crucial theoretical connection between the automatic/controlled issue on the one hand, and the notions of selective attention and executive functions on the other. Selective attention, which is tightly connected to consciousness, has

traditionally been considered a key component of executive functions. The notion of executive function is used in psychology and neuroscience to describe a loosely defined set of capacities having to do with guidance of behavior: planning, initiating appropriate actions while inhibiting inappropriate ones, cognitive flexibility etc. Selective attention is apparently a key component of this cognitive complex. On the other hand, executive functions appear as prototypical examples of controlled processes: but this is so because of the fact that executive functions involve consciousness and selective attention, not propositionality.

For all these reasons, I will rest on the most widely accepted view of the automatic/controlled distinction, that is, the view that takes consciousness as the main line of demarcation and emphasizes the connections with theories of intentional action and executive functions. Once this general framework is settled, further qualifications are suggested by evidence that has been acquired recently. One is the observation that "automatic" and "controlled" presumably are not all-or-none notions; instead, they appear to come in degrees. A second point is that, although consciousness and controlled processes have a crucial role to play in goal-directed behavior, it is possible to have goal-directed behavior outside consciousness. Third, the emphasis in recent research is less on how automatic and controlled processes may be detected and analyzed in isolation than on how they factually cooperate for most of our cognitive activities. Let us now briefly address each of these points in turn; this will prepare the ground for an updated view of the automatic/controlled issue—a view which has been explored in the "distributed intentionality model" proposed by Mazzone and Campisi (2013).

3.2 All-or-none Notions?

The traditional view according to which there is a sharp boundary between automatic and controlled processes has been challenged in the last decades, either because automaticity has been explained in terms of properties which vary gradually, or because automaticity has been analyzed in components which can, but need not, be present together (Garrod and Pickering 2007). The latter line of thought has been pursued by Bargh (1994): in his view, a process is automatic to the extent that it is unaware, mandatory, efficient, non-interruptible.[3] However, since those features do not always covary together, there may be different degrees of automaticity as a function of the number of features involved.

As for the former line of thought, it is well exemplified by Cohen et al. (1990). They propose that automaticity is a function of what they call "strength of processing", which in turn is defined in relation to processing pathways within a

[3] In Bargh's (1994) own terms, the four parameters are awareness, intentionality, efficiency, and controllability. However, Garrod and Pickering suggest that "non-interruptibility" is a more proper label for what Bargh calls "controllability". Similarly, Mazzone and Campisi (2013) observe that by the term "intentionality" Bargh properly means that a process is not mandatory.

connectionist network. A strong connection leads to fast and accurate transmission of information along the pathway; moreover, the strength of processing may determine the extent to which processes are open to interference from other processes. Considerations of this sort might be framed differently by taking into consideration the apparent coexistence of two complementary learning systems in our brain, one for rapid learning based on the hippocampus and related structures, the other for slow learning of regularly repeated sequences (McClelland et al. 1995). As it seems, the same information can be moved from the former system to the latter as a consequence of repetition and practice (Aarts and Custers 2009); this also leads to different patterns of activation in the neural circuits guiding action:

> lateral prefrontal and premotor areas are activated at the beginning of the learning of a motor sequence; with practice and repetition, however, that activation subsides, while that of subcortical structures, notably the basal ganglia, increases […]. Thus, as sequences become overlearned and automatic, their representation seems to "migrate" to lower executive stages (Fuster 2001, 321–322).

Under this hypothesis, processing can be fast, accurate and unaffected by interference to the extent that a sequence has been overlearned and therefore moved to subcortical structures—where overlearning is something that may come in degrees.

A different but possibly complementary proposal has been put forth by Dehaene et al. (2006) with regard to the neural basis of consciousness. In the model they propose, the distinction between conscious and unconscious processing is not all-or-none, for two reasons. First, they assume that besides conscious and purely subliminal processing there can be intermediate conditions. Second, in their model consciousness is also a function of the strength of activation in the interested brain areas. To be more precise, Dehaene et al. (2006) distinguish four conditions of un/consciousness, depending on the degree of activation which is found respectively in posterior sensory-motor representations and higher association cortices. Pure (i.e., unattended) subliminal processing occurs when there is a weak and rapidly decaying activation in posterior sensory-motor areas, without any significant interacting activation in anterior cortices; on the other hand, attention and task set might occasionally interact with such weak posterior activations, thus resulting in attended subliminal processing. Moreover, an intense activation which is yet confined to sensory-motor processes is thought to cause occipito-temporal loops and local synchrony, and therefore a condition of preconscious processing: processes are virtually accessible to consciousness, although attention is actually oriented away from the stimulus, so that activation is blocked from accessing higher parieto-frontal areas and establishing long-distance synchrony. Finally, conscious and controlled processes require the establishing of long-distance loops between strongly activated sensory-motor representations and higher association cortices.

In sum, considerations from both psychology of learning and neurobiology of consciousness seem to suggest that the automatic/controlled distinction admits of degrees, instead of being an all-or-none affair.

3.3 Goal-Directedness without Consciousness

Traditionally, goal pursuit has been conceived of as a typical case of conscious and effortful processing. This view has been recently challenged especially by John Bargh (starting from Bargh 1989, 1990). His notion of *automatic* or *non-conscious goal pursuit* has witnessed a number of empirical demonstrations in the last decade (for recent reviews see Hassin et al. 2009; Ferguson et al. 2007). The thesis of non-conscious goal pursuit is based on the notion of habit, with habits conceived of as

> associative networks that include contexts, goals that are regularly pursued in these contexts, and means that one usually uses to attain these goals […]. These networks are shaped by one's history, and they allow for goal pursuit via spreading of activation (Hassin et al. 2009, 550–551).

Given this conception of habits as associative networks, it seems an obvious consequence that the activation of a component may spread to other components of the network, and this has been largely confirmed by experiments based on priming.[4] In particular, priming of goals appears to affect subsequent representations and behaviors in many ways. For instance, when an action is regularly selected and performed in order to obtain a goal (for instance, taking the bicycle instead of the bus to go to the university), "priming these goals automatically activates behavior representation and resultant action according to an 'if-then' rule, enabling the goal-directed behavior to occur directly and independent of conscious intentions" (Dijksterhuis et al. 2007, 105). Of particular interest is the fact that similar results have also been obtained through *unobtrusive* or *unconscious* priming. For instance, Bargh et al. (2001) unobtrusively exposed subjects to words such as "strive" and "succeed" to prime the achievement goal, and then tested their performances in an anagram puzzle task: participants primed with the achievement goal outperformed those who were not primed with the goal. Similar effects may also be obtained in more indirect ways: for instance, priming the names of significant others may lead to the automatic adoption of the goals associated with them; or for another example, thinking to a good friend may enhance the disposition to participate in a subsequent task as a possible means to help (Dijksterhuis et al. 2007, 101–102).

What these observations apparently show is that automatic processing may go deep into the guidance of behaviors which are thought of as typically intentional and controlled. That this must be the case is also shown by the fact that most of our intentional actions are nonetheless rapid and effortless. This has led to models of intentional actions where the most part of cognitive processing is thought to occur automatically. However, such models often tend to assume that automaticity does

[4] See also Gollwitzer et al. (2009, 605), where they suggest that goals may behave in accordance with simple associative (hebbian) principles: "Under the assumption that goals, too, are represented mentally and become automatically activated by the same [hebbian] principles, goal representations should also be capable of automatic activation through contact with features of the contexts in which those goals have been pursued often and consistently in the past.".

only concern the implementation of intentions: conscious representations of the goals to be pursued are held to be required for action control, while the specific behavioral means by which the goals are pursued would be activated automatically (see, for example, Levelt's 1989 model of language production). On the contrary, the evidence concerning non-conscious goal pursuit invites us to believe that goals may drive action without becoming conscious, that is, they can operate in an entirely automatic way.

3.4 How Automatic and Controlled Processes Cooperate

On the basis of our previous considerations, one could be tempted to think that consciousness does not play a significant role in human cognition. Our actions are mostly rapid and effortless, and this suggests a major role for automatic processing. To be sure, human action is essentially goal-directed, but, as we saw, goal-directedness does not imply conscious processing. Another relevant line of evidence is provided by the experiments of Libet (e.g., Libet 1992), which have shown that "[c]onsciousness of the goal of an action is not immediate, it takes time to appear" (Jeannerod 2006). More specifically,

> the first conscious awareness associated with the initiation of the movements […] occurs well after the start of the neural activity that culminates in the movement. […] This clearly suggests that whatever events one might reasonably consider to be the neural initiators of these movements, those events occur pre-consciously (Pockett 2006, 18–19).

Based on this sort of evidence, some have drawn the conclusion that consciousness is essentially a post hoc phenomenon, which has not to do with initiation and guidance of action. It would rather be (part of) a mechanism "for the cognitive rearrangement after the action is completed" (Jeannerod 2006, 37), in the service of our sense of agency and the distinction between our own and others' actions (Pockett 2006; Jeannerod 2006; Choudhury and Blakemore 2006). However, there are reasons to believe that conscious and controlled processes should be accorded instead a significant role in active online processing and guidance of action. First, conscious control appears to be occasionally required in the course of action when smooth automatic processing fails (Gollwitzer et al. 2009, 610; Bongers and Dijksterhuis 2009; Jeannerod 2006, 30). Second, sometimes we make conscious plans of action, or we are explicitly required to accomplish a task, and so on. In such cases, but possibly also in cases where initiation of action is automatic, consciousness seem to play a key role in top-down maintenance of goals and top-down inhibition: the execution of long-term plans cannot be accounted for solely in terms of automatic spreading of activation. This suggests that not only have both automatic and conscious processes a role to play in human cognition, they are also expected to cooperate in most of our cognitive performances. There is nowadays growing acceptance that "conscious and nonconscious goal pursuit are two collaborative partners taking turns in working towards goal attainment" (Gollwitzer et al. 2009, 620–621).

This cooperative view of automatic/controlled processes is entirely coherent with the neurobiological model proposed by Dehaene et al. (2006): as we saw, that model proposes that both in attended subliminal processing and in conscious processing frontal and pre-frontal activations can affect automatic processes, by amplifying the independent activation of certain representations (and presumably by causing the active inhibition of others) in posterior areas. An interesting way to frame attended subliminal processing is Neumann's (1990) theory of "direct parameter specification". According to this theory "[a] given attentional (or intentional) state might be necessary for unconscious stimuli to trigger further processes" (as Kiefer 2007, 293, puts it). More specifically:

> [Subjects] search for information in order to specify free parameters within the currently active intention/action plan. Unconsciously registered information that resembles this searched-for information is selected and processed to specify the free processing parameters. Therefore, unconsciously perceived information will translate into behavioural effects that are absent if the same information is sufficiently dissimilar from the searched-for features (Kiefer 2007, 300).

In other words, top-down intentional processing would cause stimuli to affect behavior even when they are not consciously perceived.

3.5 An Updated View of the Automatic-Controlled Issue

In the light of the sort of evidence we have reviewed so far, the distinction between automatic and controlled processes should be considered just the first step on the way to understanding their cooperation in most of our cognitive operations. Mazzone and Campisi (2013) have proposed a general approach to intentional actions—the "distributed intentionality model"—based on such a cooperative view of automatic and controlled processes. We propose that in order for actions to be intentional it is not required that action plans are consciously represented and then put into effect in a purely top-down manner. In the general case, actions are largely the result of automatic processes of activation, integration and competition between a huge number of goal-related representations. On the other hand, human behavior is intuitively intentional in essence, in that it never seems to occur without agents consciously attending this or that component of the complex goal-directed representation involved. However, conscious attention is not necessarily directed towards one specific component of that representation, be it an overall goal or whatever: conscious intentions should rather be conceived of as beams of light temporarily directed towards this or that goal-related component of a largely automatic flow of processing. In a word, intentionality is better thought of as dynamically *distributed* along the complex goal-directed representation involved in any single action, than concentrated in (the representation of) one single purpose of the action.

In our model, the role played by automatisms is very large. This is in line with a proposal of Morsella (2009), according to which human behavior is based on a "stream of action […] driven by a continuous series of activations stemming from various sources" (idem, p. 19). In other words, our perceptions would endlessly feed automatic processes impinging on motor representations, so that plans of action are activated automatically at each moment and then compete for behavioral expression (idem, p. 16). However, this is far from implying that consciousness is either absent or purely epiphenomenal in most of our intentional actions. First, even if consciousness takes time to appear, nonetheless it may emerge in the course of action and then play a crucial role as a mechanism for goal maintenance and shielding, for reorganization of habits, or for the management of unexpected difficulties (Mazzone and Campisi, 2013). Second, it should not be forgotten that for the most part of our lives "we live in a supraliminal world" (Satpute and Lieberman 2006, 91), that is, automatic responses to perceptual inputs occur while we are engaged in conscious monitoring of the environment and our own behavior. In a sense, then, it could be true that there are conscious representations at the instigation of most of our actions: humans often respond to situations they are conscious of, and these situations set the purposes of our forthcoming actions. For instance, in dialogue we normally attend to our interlocutor's utterances. Such a conscious representation of the input we intend to respond to can be thought to drive automatic processing by constraining the kind of information which is needed to accomplish the task—as predicted by the "direct parameter specification" theory considered above.

In sum, it seems that in principle any component of the complex goal-directed representation involved in action—including goals—can be processed automatically. Nonetheless, consciousness is far from being epiphenomenal since it may focus on this or that component when needed and, as a consequence, play a role in directing automatic processing: specifically, as in Dehaene et al.'s (2006) model, by amplifying or inhibiting representations in posterior areas of our brain.

4 Pragmatics and Associative Processes

As we saw in Sect. 2, Relevance Theory and Recanati's view are not equally compatible with psychological and neurobiological accounts of the controlled/automatic distinction. Recanati's view is closer to those accounts than RT, to the extent that the former conceives of pragmatic processing in terms of a cooperation between associative and conscious processes, while the latter does not accord a role to any of these two processes within pragmatics proper. Relevance theorists propose instead a single automatic mechanism which is specialized for language comprehension. Let us now examine in more detail the positions of both RT and Recanati with regard to associative processes (this section), and conscious processes (the next section).

Within the literature we considered in Sect. 3 there is a general consensus that automatic processing occurs by way of associative mechanisms, that is, mechanisms based on associative strength in a network mainly due to "extended consistent training" (Schneider and Chein 2003, 528) and on subsequent spreading of activation in the associative network. For instance, in Sloman (1996) the automatic pole of the dichotomy is straightforwardly called "associative" (versus rule based), and in Satpute and Lieberman (2006, 88) the reflexive (versus reflective) component is claimed to be based on associations and to deliver constraint satisfaction processes. Moreover, the thesis of automatic goal pursuit depends on the notion of habits conceived of as associative networks involving representations of contexts, goals and means.

Recanati (2004) has in fact proposed that lexical items contribute their meaning to the explicit sense of utterances by way of what he calls primary pragmatic processes, conceived of as local associative processes. In Mazzone (2011) I have argued that Recanati's associative explanation may be extended beyond his intentions—in particular, beyond the domain of lexical pragmatics. But let us proceed step by step.

4.1 Associative Accounts of Primary Pragmatic Processes

A crucial notion in Recanati's account of primary pragmatic processes is that of abstract schemata driving comprehension. Not only can schemata explain shifts in accessibility of lexical meanings, they can also account for the search of coherence in associative processes: inputs activate schemata they are component of, and schemata in turn activate (or add activation to) other inputs (and their interpretations) insofar as they fit those schemata. In our previous example (Sect. 2), the schema GETTING MONEY FROM A BANK_1 (where BANK_1 = FINANCIAL INSTITUTION) may have a key role in explaining how, in the utterance "I'm going to the bank now to get some cash", the word "bank" is given an interpretation which is coherent with the context. Interestingly, the same sort of schematic information is invoked by RT in order to ensure the assumptions that behave as premises in their inferential explanations. Thus, in this respect what essentially distinguishes RT from Recanati's account is the thesis that such a schematic information is employed within genuinely inferential processes, instead of associative ones.

However, as argued in Mazzone (2011), this thesis is both highly speculative and unnecessary. As to the first point, on epistemological grounds associative activation and automatic inferential derivation are far from having the same status: the latter is not nearly as established as the former, which is in fact the only well-established explanation–both in psychology and neurobiology—of how we detect, store and exploit information by way of automatic processes. On the other hand, it is far from clear that a genuinely inferential account is needed. In particular, RT underestimates the theoretical role that schemata can play within an associative account of automatic processes.

This is clearly shown by the most extensive argument against associative accounts proposed by relevance theorists, which has been put forth by Wilson and Carston (2007) in the context of a discussion of lexical pragmatics. The key claim in their argument is that statistical associations provide no basis for drawing warranted conclusions, since the associates are not logically related to each other in any systematic way. To be sure, one could maintain that inferential relationships are also associations of some sort; for instance, the association between "shark" and "fish" could be used to derive the warranted conclusion that a shark is a fish.[5] However, although inferential relationships are associations, there are plenty of associations that are not inferential relationships. Therefore, according to Wilson and Carston associative accounts will vastly overgenerate, and so one is left without any principled method of filtering out unwanted associations (and unwarranted conclusions). This is why inferential accounts should be preferred.

The first thing to notice is that the premise of the argument is false. Far from lacking any systematic structure, associations are instead essentially schematic. In other words, associations are not stored in such a way that the relationships between their elements are in need of interpretation from the outside—so to speak. Quite on the contrary, our associative coding of contingencies yields schemata preserving information both on which content are connected with each other and *how* they are connected, be it by way of taxonomic, part-whole, temporal, causal, textual relationships or whatever. As a consequence, associative networks do not require that further mechanisms be provided in order to logically constrain their dynamics of activation. Instead, they can themselves provide—just as suggested by Recanati—the abstract schemata thanks to which the process is constrained, and unwanted associations are filtered out. For instance, although in our previous example the word "bank" may activate the meaning RIVER SIDE—not to mention all the other associations potentially activated by the utterance "I'm going to the bank now to get some cash"—this meaning either will not receive further activation from, or even will be inhibited by, other associative schemata triggered by linguistic and contextual inputs.

As it seems, the notion of schema may help to provide, after all, an associative explanation of how unwanted associations are filtered out in comprehension.

4.2 Beyond Lexical Pragmatics

It could be objected that such an explanation may only work within the limits of lexical pragmatics, where the issue is how words confer their meanings to the explicit content of utterances. Also relevance theorists grant a role to associative

[5] Within this argument, Wilson and Carston essentially identify associations with statistical relationships between lexical items in a corpus. As we are going to argue, there is no ground for that identification: there exist in fact a variety of different associative relationships, most of which concern concepts rather than words.

processes in that domain. However, according to RT not even explicit content can entirely be fixed by associative mechanisms: explicit content—no matter how it is prompted—has to become a line in an inferential derivation, and mutual adjustment between the components of the derivation is needed in order for any of those components to be fixed.

Why should we presume that such an inferential derivation is needed, and that simple associative processes will not do? Carston (2007) has an argument for this which deserves consideration.[6] Her idea is that associations suffice insofar as what is at issue is activation and deactivation of concepts (parts of concepts, schemas), while associations are not sufficient in order to understand genuinely constructive processes. Although Carston is here concerned with how concepts can be constructed online rather than simply re-activated, her argument also sheds new light on the previous claim that associations provide no basis for drawing warranted conclusions. Intuitively, in order to be justified in reaching a conclusion a cognitive system needs something more than activation merely passing from one content to another: it has to *construct* an inference that may count as a justification for the conclusion. I think Carston has a good point here, but the precise implications of the argument have to be assessed more accurately.

Let us first notice that in current linguistics there is a family of theories assuming that associative relations can explain cognitive phenomena which had previously been thought to require rule-based, specialized processes instead. This is the case with what are known as constraint-based models, that is, models in which parallel activation of, and competition between, representations substitute for procedural rules, in syntax and elsewhere (e.g. Trueswell et al. 1994; Ferreira et al. 2002; Jackendoff 2007; Breheny et al. 2006). Constraint-based processes and associative processes can be seen as two sides of the same coin: as a consequence of activation within an associative network, each activated representation may act as a constraint on the overall process, insofar as it contributes to selecting the outcomes which are compatible with it.

One insightful example of constraint-based model in linguistics has been put forth by Jackendoff (2007). Although his theory has its roots in Generative Grammar, Jackendoff maintains that linguistic phenomena—syntax included—may be explained by a general-domain, constraint-based mechanism. Crucially, while in the mainstream view of Generative Grammar phrase structure has been represented in terms of procedural rules, Jackendoff proposes that any linguistic information[7] including phrase structure is instead captured by regular patterns of representation essentially abstracted away from experience: words, regular affixes, idioms, constructions, and ordinary phrase structure rules are conceived of as nothing but "pieces of structure stored in long-term memory" (Jackendoff 2007, 11). As a consequence, Jackendoff's explanation does not rely on specialized

[6] She also proposes another interesting argument we will consider in the next section, since it concerns the role of consciousness.

[7] With the possible exception of a very restricted number of innate constraints.

linguistic (namely, syntactic) processes operating in accordance with procedural rules. Rather, linguistic representations (pieces of structure) are thought to contain within them the information on how they can be assembled with each other, and all we would need is a general-domain process which mechanically assembles representations in accordance with that information: this process is called *unification*.

Since what is at stake in unification is the building of occurrent linguistic structures, Jackendoff rightly points out that this process necessarily requires something like a "workbench", or a "blackboard" where structures are constructed online. Typically, such a workbench is what working memory is thought to provide. But Jackendoff also emphasizes that in order to accomplish the task, working memory cannot be conceived of as just the part of long-term memory that is currently activated—as it is in some connectionist architecture. In his opinion, working memory should rather be thought of as physically separate from long-term memory. For our purposes, though, the point is that building conceptual structures requires more than simple spreading of activation in a network: it requires that the cognitive system is able to keep certain pieces of structure active until the whole process of activation, competition and unification is accomplished.

What these considerations suggest is that the construction of complex conceptual structures can be accounted for within an associative framework, provided that working memory is added to simple spreading of activation. Carston's objection has the merit of calling attention to this important qualification, but it does not speak in favor of inferential processes, since we have an associative account of how complex structures can be constructed in the course of online processing. One may speculate that those complex structures possibly include exemplifications of inferential schemata. This could explain how associative processes may mimic inferential processes, as suggested by Recanati (2007). There could occur indeed processes of mutual adjustment between assumptions counting as premises and utterance interpretations counting as conclusions, insofar as those assumptions and interpretations are unified in working memory by means of inferential schemata: such schemata would activate, or strengthen the activation of, the components (premises and conclusions) which fit them.

Would that count as an inferential account of the sort recommended by RT? Not at all. A couple of things should be emphasized. First, the mechanisms Jackendoff's model makes use of (spreading activation, working memory) are domain-general. On the contrary, inferential processes hypothesized by RT are specialized for utterance understanding. Second, in constraint-based accounts a crucial role is played by structures of representation: those structures (plus simple general-domain mechanisms) substitute for rules and derivations. Analogously, in pragmatic processing a variety of schemata (together with spreading activation and working memory) could explain how warranted conclusions could be granted without genuine inferential processes. Third, there is a clear sense in which working memory is just part of the general associative dynamics of our brain: from a neurobiological perspective, working memory is generally taken to consist in self-sustained loops occurring in cortical circuits. In other words, working memory is essentially a specific modality of activation within associative networks.

These considerations support the view that associative processes can explain linguistic and pragmatic phenomena well beyond the limits of lexical pragmatics.

Before we conclude this section, one qualification is in order. The fact that we have introduced working memory in the picture does not necessarily imply that consciousness is at play as well. For sure, the idea of a global workspace—ensured by self-sustained loops in the cortex, and accounting for the active maintenance and integration of information—has traditionally been tightly associated with consciousness (e.g., Baars 1997). Nonetheless, there are reasons to think that working memory is independent from consciousness. For instance, Hassin et al. (2009) have noticed that there is an apparent contradiction between the fact that we can only engage in a very limited number of high-order cognitive processes (and specifically, conscious processes) at any given point in time, and the simple intuition that there are points in time in which we seem to be advancing multiple goals, decisions and plans. Hassin and colleagues propose to solve this apparent contradiction by adopting the thesis of an implicit working memory. In other words, they argue that working memory can operate outside of conscious awareness and therefore it may ensure parallel processing. Incidentally, they also observe that none of the major views on this issue suggests that people have conscious access to everything that goes on within working memory.

The obvious implication of the "implicit working memory" hypothesis is that conceptual integration may also occur automatically—i.e., outside consciousness. Automatic integration of spatial information has been in fact argued for by Hommel (1996, 1998, 2002). In ERP studies of language comprehension, it could be argued that a similar notion has been invoked. Many have proposed to interpret the difference between the best known event-related potentials in that domain—N400 and P600—in terms of a difference between two modes of processing. The N400 (a negative deflection having its peak 400 ms after the stimulus that elicits it) is thought to reflect a process of semantic integration (van Berkum et al. 1999; Vissers et al. 2006; Chwilla et al. 2007), while the P600 (a positive deflection with its peak at 600 ms from the stimulus) would reflect instead a process of monitoring and "continued algorithmic analysis" presumably under executive control (Kuperberg 2007, 42; Vissers et al. 2007). The sort of integration that is proposed as an explanation of the N400 is thus conceived of as different from both simple spreading of activation (Chwilla et al. 1998) and processes involving conscious monitoring and executive control.

From a neurobiological perspective, the hypothesis of implicit working memory is compatible with the model proposed by Dehaene et al. (2006) we considered above (Sect. 3). Under the assumption that working memory is ensured by self-sustained loops in the cortex, the model distinguishes between local loops located in occipito-temporal areas and long-distance loops also involving anterior association cortices. While the latter are thought to grant conscious and controlled processes, the former are claimed to cause preconscious processing. Therefore, local loops could be the neurobiological basis for implicit working memory and preconscious integration of representations.

5 Pragmatics and Consciousness

As we saw, Recanati's account assigns a significant role to consciousness within pragmatics. More precisely, the "availability condition" (Recanati 2004, 44) posits that subjects have conscious awareness of the explicit content, the implicit content, and the inferential process leading from the former to the latter. In a sense, Recanati conceives of the transition from explicit to implicit sense in terms of conscious verbal reasoning, with the important qualification that consciousness may be only dispositional: subjects are capable of making the relevant inferences explicitly, but in the normal case they reach the implicit content without any actual involvement of conscious verbal reasoning.

Carston (2007) makes two objections that are easily agreed upon. First, moving from occurrent to (merely) dispositional reflective inferences leaves us with no idea of how the real process of implicature derivation works: what we do know is just that conscious verbal reasoning is *not* the occurrent process by which implicit content is normally obtained. As we saw, Carston's own view is that conscious reasoning should better be assigned a role in post hoc rational reconstruction, which is something that people are actually engaged in only as a backup mechanism when something goes wrong with automatic processing.

Second, Carston argues that there is no ground for the distinction between primary and secondary pragmatic processes in terms of conscious availability. She makes various examples of cases in which people seems to be aware of how explicit content may depart from linguistic meaning as a function of contextual factors. Let us consider the following example:

(3) Mother to young child just before bedtime: Have you brushed your teeth?
Child (grinning): Yes I have—[pause]—last night.

The answer clearly shows that the child is well aware of the normal pragmatic enrichment by which the relevant time is assigned to the temporal parameter of the question: in fact, the child openly violates the expectations raised by that normal enrichment. But in Recanati's terms that sort of enrichment is a case of primary pragmatic process for the determination of explicit content. Therefore, conscious availability seems not to be an exclusive property of secondary pragmatic processes.

The claim that any stage of pragmatic processing may be consciously attended accords well with our previous considerations on conscious processes (Sect. 3), with particular regard to Mazzone and Campisi's (2013) "distributed intentionality model": speaking is a prototypical case of intentional action and, as I argued above, intentional action involves complex goal-directed representations across which consciousness is dynamically distributed. In other words, there are no specific components of goal-directed representations such that consciousness is necessarily directed towards them; consciousness may be directed instead towards different aspects in different occasions, and also in the course of the same action. But Carston also claims that consciousness have a role to play in utterance

understanding only in very special circumstances. However, it should be emphasized, Carston essentially refers to the role of conscious *reasoning*, which is a quite specific sort of conscious process. Although Carston is presumably right in pointing out that we rarely resort to conscious reasoning in utterance understanding, this in itself does not speak against consciousness having a role in pragmatics. To put it differently, while verbal reasoning proper is a prototypical instance of (largely) controlled processes, consciousness may also cooperate with processes which are mainly automatic: it is this latter kind of process, not the former, that is apparently involved in normal cases of utterance understanding.

Carston (2007) has made an argument against associative processes that in my perspective can be seen as an involuntary step in this direction. She points out that mere accessibility—even coherence-based accessibility—cannot account for the fact that utterances virtually inevitably trigger attentional focus and the expenditure of some processing effort. The conclusion this argument is aimed at is RT's thesis according to which utterance understanding cannot be explained by a general-domain associative process; it would require instead a specialized automatic mechanism based on relevance. However, one may speak of attention (attentional focus, and the like) in two quite different ways. First, one may refer to the mere fact that a cognitive system has to select somehow the direction of processing. Second, one may specifically refer to conscious attention. Carston cannot presumably be interested in the latter sense, since she argues in favor of unconscious processing of utterances. Nonetheless, it is difficult to deny that utterances do normally trigger *conscious* attention in humans. And this may contribute to explain how cognitive resources are allocated in utterance understanding, beyond mere spreading of activation: automatic processing, as we saw above, can be driven by consciously attended representations, which have a role in amplifying or inhibiting other representations in posterior cortical areas, in maintaining certain representations activated, and in creating expectations about the inputs to be automatically processed.[8]

In a word, Carston calls our attention to a fact that, again, can be easily described in terms of a general and well-established mechanism—conscious attention—although she argues in favor of a highly speculative explanation—the hypothesis of an automatic, inferential, relevance-based mechanism. While one may agree with her that there is no ground for Recanati's distinction between primary and secondary pragmatic processes, that conscious verbal reasoning is not involved in normal cases of comprehension, and that comprehension is instead a largely automatic process, it seems reasonable to acknowledge nonetheless that conscious attention may play a role in utterance understanding, in cooperation with automatic processes.

[8] The role of conscious attention in pragmatics is further explored in Mazzone, 2013b.

6 Conclusions

We live in a supraliminal (Satpute and Lieberman 2006, 91), personal-level world. Language perception, in particular, does not normally occur outside consciousness. According to Grice, comprehension involves personal-level, rational abilities on the part of the hearer. At the same time, however, pragmaticists have not neglected that utterance understanding appears to be a spontaneous, rapid and effortless process. For that reason, Recanati has hypothesized an automatic, associative stage in utterance understanding, and has assigned only a dispositional (versus occurrent) role to conscious verbal reasoning. On the other hand, Relevance theorists have proposed that comprehension is a wholly automatic, though inferential, process, with conscious verbal reasoning being assigned only a peripheral role as a backup mechanism.

I have proposed here a different account, where consciousness plays a significant role in cooperation with automatic, associative processes. This account is based on the automatic/controlled distinction in psychology, and on recent views about the cooperation between these two kinds of process. In that perspective, not only do automatic and controlled processes cooperate, they are also closer to each other than it was previously thought. Specifically, I have argued that automatic processes are based on schemata which may also be recruited in reflective reasoning, while the main difference between reflexive and reflective processing concerns just the dynamics of activation within cortical networks. This is why automatic processes are apt to mimic inferential ones.

This is not to say that spreading activation is sufficient to account for utterance understanding. I have claimed instead that working memory is also needed, and argued that it may come in two different varieties: implicit and conscious. In the framework I propose, spreading activation, implicit working memory, and conscious attention are all present in normal episodes of utterance understanding, with the first two components doing the greatest part of the work, although conscious attention has also a key role in maintenance, amplification/inhibition and anticipation of representations.

The present account has the advantage of explaining both automaticity and personal-level, rational features of comprehension by an appeal to mechanisms that are general and well-supported in psychology and neurobiology.

References

Aarts, H., Custers, R. (2009). Habit, action, and consciousness. In W. P. Banks (Ed.), *Encyclopedia of Consciousness* (Vol. 1), 315–328. Oxford: Elsevier.

Baars, B. J. (1997). In the theatre of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies, 4*(4), 292–309.

Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence on social perception and cognition. In J. Uleman, and J. Bargh (Eds.), *Unintended thought*, 3–51. New York: Guilford Press.

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), 93–130. New York: Guilford Press.

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, efficiency, intention, and control in social cognition. In R. S. Wyer Jr., T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed.), 1–40. Hillsdale: Erlbaum.

Bargh, J. A., Gollwitzer, P. M., Lee–Chai, A., Barndollar, K., & Troetschel, R. (2001). The automated will: Nonconscious activation and pursuit of behavioral goals. *Journal of Personality and Social Psychology, 81*, 1014–1027.

Bongers, K. C. A., & Dijksterhuis, A. (2009). Consciousness as a trouble shooting device? The role of consciousness in goal–pursuit. In E. Morsella, J. A. Bargh, & P. Gollwitzer (Eds.), *Oxford handbook of human action* (pp. 589–604). New York: Oxford University Press.

Breheny, R., Katsos, N., & Williams, J. (2006). Are scalar implicatures generated by default? *Cognition, 100*(3), 434–463.

Capone, A. (2011a). Default Semantics and the architecture of the mind. *Journal of Pragmatics, 4*(3), 1741–1754.

Capone, A. (2011b). The Attributive/Referential distinction, pragmatics, modularity of mind and modularization. *Australian Journal of Linguistics, 31*(2), 153–186.

Carston, R. (2007). How many pragmatic systems are there? In M. J. Frapolli (Ed.), *Saying, Meaning, Referring. Essays on the Philosophy of François Recanati*, 1–17. New York: Palgrave.

Carver, C. S., & Scheier, M. F. (2009). Action, affect, and two–mode models of functioning. In E. Morsella, J. A. Bargh, & P. Gollwitzer (Eds.), *Oxford handbook of human action* (pp. 298–327). New York: Oxford University Press.

Choudhury, S., & Blakemore, S. J. (2006). Intentions, actions, and the self. In S. Pockett, W. P. Banks, & S. Gallagher (Eds.), *Does consciousness cause behavior?* (pp. 39–51). Cambridge: MIT Press.

Chwilla, D. J., Hagoort, P., & Brown, C. M. (1998). The mechanism underlying backward priming in a lexical decision task: Spreading activation versus semantic matching. *The Quarterly Journal of Experimental Psychology, 51*, 531–560.

Chwilla, D. J., Kolk, H. H. J., & Vissers, C. T. V. M. (2007). Immediate integration of novel meanings: N400 support for an embodied view of language comprehension. *Brain Research, 1183*, 109–123.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review, 97*, 332–361.

Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences, 10*(5), 204–211.

Dijksterhuis, A., Chartrand, T. L., & Aarts, H. (2007). Effects of priming and perception on social behavior and goal pursuit. In J. A. Bargh (Ed.), *Social psychology and the unconscious: The automaticity of higher mental processes* (pp. 51–132). Philadelphia: Psychology Press.

Epstein, S. (1973). The self–concept revisited or a theory of a theory. *American Psychologist, 28*, 404–416.

Ferreira, F., Bailey, K. G. D., & Ferraro, V. (2002). Good–enough representations in language comprehension. *Current Directions in Psychological Science, 11*, 11–15.

Ferguson, M. J., Hassin, R. R., & Bargh, J. A. (2007). Implicit motivation. In J. Shah & W. Wood (Eds.), *Handbook of motivation* (pp. 150–168). New York: Guilford Press.

Fuster, J. (2001). The prefrontal cortex-an update: Time is of the essence. *Neuron, 2*, 319–333.

Garrod, S., & Pickering, M. J. (2007). Automaticity in language production in monologue and dialogue. In A. S. Meyer, L. R. Wheeldon, & A. Krott (Eds.), *Automaticity and control in language processing* (pp. 1–21). Hove: Psychology Press.

Gollwitzer, P. M., Parks-Stamm, E. J., & Oettingen, G. (2009). Living on the edge: Shifting between nonconscious and conscious goal pursuit. In E. Morsella, J. A. Bargh, & P.

Gollwitzer (Eds.), *Oxford handbook of human action* (pp. 603–624). New York: Oxford University Press.

Grice, H. P. (1989). *Studies in the way of words*. Cambridge MA: Harvard University Press.

Hassin, R. R., Aarts, H., Eitam, B., Custers, R., & Kleiman, T. (2009). Non–conscious goal pursuit and the effortful control of behavior. In E. Morsella, J. A. Bargh, & P. Gollwitzer (Eds.), *Oxford handbook of human action* (pp. 442–457). New York: Oxford University Press.

Hommel, B. (1996). The cognitive representation of action: Automatic integration of perceived action effects. *Psychological Research, 59*, 176–186.

Hommel, B. (1998). Event files: Evidence for automatic integration of stimulus–response episodes. *Visual Cognition, 5*, 183–216.

Hommel, B. (2002). Responding to object files: Automatic integration of spatial information revealed by stimulus–response compatibility effects. *Quarterly Journal of Experimental Psychology, 55A*, 567–580.

Jackendoff, R. (2007). A parallel architecture perspective on language processing. *Brain Research, 1146*, 2–22.

James, W. (1890). *The principles of psychology* (Vol. 1). New York: Holt.

Jeannerod, M. (2006). Consciousness of action as an embodied consciousness. In S. Pockett, W. P. Banks, & S. Gallagher (Eds.), *Does consciousness cause behavior?* (pp. 25–38). Cambridge MA: MIT Press.

Kiefer, M. (2007). Top–down modulation of unconscious 'automatic' processes: A gating framework. *Advances in Cognitive Psychology, 3*(1–2), 289–306.

Kuperberg, G. R. (2007). Neural mechanisms of language comprehension: Challenges to syntax. *Brain Research, 1146*, 23–49.

Lieberman, M. D., Gaunt, R., Gilbert, D. T., & Trope, Y. (2002). Reflection and reflexion: A social cognitive neuroscience approach to attributional inference. *Advances in Experimental Social Psychology, 34*, 199–249.

Libet, B. (1992). The neural time–factor in perception, volition and free will. *Revue de métaphysique et de morale, 97*(2), 255–271.

Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge MA: MIT Press.

Mazzone, M. (2011). Schemata and associative processes in pragmatics. *Journal of Pragmatics, 43*, 2148–2159.

Mazzone, M. (2013a). A pragmatic Pandora's box. Regularities and defaults in pragmatics. In F. Liedtke and C. Schulze (Eds.), *Beyond Words, Content, context, and inference* (pp. 307–329). Berlin, Boston: Mouton.

Mazzone, M. (2013b). Attention to the speaker. The conscious assessment of utterance interpretations in working memory. *Language & Communication, 33*, 106–114.

Mazzone, M., Campisi, E. (2013). Distributed intentionality. A model of intentional behavior in humans. *Philosophical Psychology*, *26*, 267–290.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102*, 419–457.

Morsella, E. (2009). The mechanisms of human action: Introduction and background. In E. Morsella, J. A. Bargh, & P. Gollwitzer (Eds.), *Oxford handbook of human action* (pp. 1–32). New York: Oxford University Press.

Neumann, O. (1990). Direct parameter specification and the concept of perception. *Psychological Research, 52*, 207–215.

Pacherie, E. (2006). Towards a dynamic theory of intentions. In S. Pockett, W. P. Banks, & S. Gallagher (Eds.), *Does consciousness cause behavior?* (pp. 145–167). Cambridge MA: MIT Press.

Pockett, S. (2006). The Neuroscience of movement. In S. Pockett, W. P. Banks, & S. Gallagher (Eds.), *Does consciousness cause behavior?* (pp. 9–24). Cambridge MA: MIT Press.

Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.

Recanati, F. (2007). Reply to Carston. In M. J. Frapolli (Ed.), *Saying, Meaning, Referring. Essays on the Philosophy of François Recanati*, 49–54. New York: Palgrave.

Satpute, A. B., & Lieberman, M. D. (2006). Integrating automatic and controlled processing into neurocognitive models of social cognition. *Brain Research, 1079*, 86–97.

Schneider, W., & Chein, J. M. (2003). Controlled and automatic processing: Behavior, theory and biological mechanism. *Cognitive Science, 27*, 525–559.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing. I. Detection, search, and attention. *Psychological Review, 84*(1), 1–66.

Shastri, L., & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic binding using temporal synchrony. *Behavioral and Brain Sciences, 16*, 417–451.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing. II. Perceptual learning, automatic attending and a general theory. *Psychological Review, 84*(2), 127–190.

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

Smolensky, P. (1988). On the proper treatment of connectionism. *The Behavioral and Brain Sciences, 11*, 1–23.

Sperber, D., Wilson, D. (1986/1995). *Relevance: Communication and Cognition*. Oxford: Blackwell.

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

Strack, F., Deutsch, R., & Krieglmeyer, R. (2009). The two horses of behavior: Reflection and impulse. In E. Morsella, J. A. Bargh, & P. Gollwitzer (Eds.), *Oxford handbook of human action* (pp. 104–117). New York: Oxford University Press.

Trueswell, J. C., Tanenhaus, M. K., & Kello, C. (1994). Semantic influences on parsing: Use of thematic role information in syntactic disambiguation. *Journal of Memory and Language, 33*, 285–318.

Van Berkum, J. J. A., Hagoort, P., & Brown, C. M. (1999). Semantic integration in sentences and discourse: Evidence from the N400. *Journal of Cognitive Neuroscience, 11*(6), 657–671.

Vissers, C. T. W. M., Chwilla, D. J., & Kolk, H. H. J. (2006). Monitoring in language perception: The effect misspellings of words in highly constrained sentences. *Brain Research, 1106*, 150–163.

Vissers, C. T. W. M., Chwilla, D. J., & Kolk, H. H. J. (2007). The interplay of heuristics and parsing routines in sentence comprehension: Evidence from ERPs and reaction times. *Biological Psychology, 75*(1), 8–18.

Wilson, D., & Carston, R. (2007). A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics* (pp. 230–260). New York: Palgrave.