

NOTICE: This is the author's version of a work that was accepted for publication in <Language & Communication>. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms, may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in LANGUAGE & COMMUNICATION, 33, 2013, DOI:10.1016/j.langcom.2013.01.001

Attention to the speaker. The conscious assessment of utterance interpretations in working memory

Abstract

The role of conscious attention in language processing has been scarcely considered, despite the wide-spread assumption that verbal utterances manage to attract and manipulate the addressee's attention. Here I analyse the possibility that this assumption is in fact understood not as a figure of speech but instead in terms of attentional processes proper. This hypothesis can explain a fact that has been noticed by supporters of Relevance Theory in pragmatics: the special role played by speaker-related information in utterance interpretation. I argue that attentional processes can explain this fact, in that representation of the speaker in working memory reliably enhances the activation of speaker-related information and, consequently, the role it plays in determining the content of interpretations.

Keywords: pragmatics; attention; associative processes; working memory.

INTRODUCTION

That verbal utterances are designed by the speaker so as to attract and manipulate the addressee's attention – and that, in point of fact, they normally succeed in doing it – is something that would be easily agreed upon by most scholars in language studies. To begin with, this assumption is standard

in most pragmatic research, for instance it has received great emphasis in Relevance Theory (from now on RT: Sperber and Wilson, 1986; Carston, 2002; Wilson and Sperber, 2004), whose very notion of relevance is based upon the idea that utterances deserve, and attract, attention and processing effort. Moreover, during the last twenty years or so, the notion of attention to others and specifically to their goal-directed behaviour (especially in the form of joint attention) has played a crucial role in a large body of research on social cognition, and specifically on mind reading and its contribution to language emergence in ontogeny (e.g., Tomasello, 1999; Tomasello et al., 2005).

However, in spite of the theoretical significance accorded to attention in language studies, its precise contribution to the phenomena cited above has not been addressed in any detail. This seems rather odd, since there is a huge literature on attention, and one would expect that research to be relevant for the understanding of how exactly attentional processes may contribute to mind reading and pragmatic processing.

In the present work, I intend to focus on one such possible contribution, that is, on how attention *to the speaker* may affect utterance interpretation. More precisely, I will argue that attention to the speaker may solve a major problem that has been raised by relevance theorists. They conceive it as an argument against associative accounts of comprehension, but my claim is that RT itself has no adequate solution to that problem. The point at issue is how considerations of speaker's mental states can affect utterance understanding over and beyond simple associations in the addressee's memory. It has been claimed that associative processes are a key component of pragmatic understanding. Specifically, this position has been defended by Recanati (2004) with regard to what he calls primary pragmatic processes, and generalized by Mazzone (2009, 2011, in press a, b) to pragmatic phenomena of any kind. However, relevance theorists (Carston, 2007; Mazzarella, 2011) have objected that associative accounts have to face a serious problem: they have no theoretical resources to explain the special role played by assumptions about the speaker in the interpretation

process. Just as utterances (and ostensive stimuli in general) are attention pre-empting in a way that most other stimuli are not, so – relevance theorists argue – information about the speaker gains special attention and provides crucial assumptions to the inferential processes warranting utterance interpretations. According to relevance theorists, mere spreading activation based on associative relations is too unconstrained for it to ensure that information about the speaker plays the expected role.

My line of argument in this paper is as follows. First, the objection raised by Carston (2007) and Mazzarella (2011) is substantially correct, and therefore I agree that associative accounts are in need to explain the special role of speaker-related information. In the second place, however, a close analysis of RT shows that it is subject to the very same problem: the mechanisms by which it explains utterance interpretations do not really account for how information about the speaker gains prominence. Then, a solution to the problem is needed anyway. In the third place, I argue that the required prominence of speaker-related information might be ensured by the mechanism of conscious attention.

Although relevance theorists have given great emphasis to the fact that utterances pre-empt the attention of addressees, this was essentially intended to mean that utterances trigger relevance-based *automatic* processing. RT does not accord instead any significant role to conscious attention in normal pragmatic processing. This is probably due to the reasonable assumption that conscious *reasoning* cannot be what ensures rapid and effortless language production and comprehension as they normally occur in context. However, conscious reasoning and conscious attention are quite different phenomena.

Conscious attention can be described as directedness toward information that is actively maintained in working memory (Knudsen, 2007). This active maintenance of information in working memory has important consequences. First of all, it shields some information from

distracting inputs and also prevents it from rapidly decaying in the course of processing. Second, it affects further processing by both active top-down control and automatic top-down recovery of related information. And third, within working memory information is evaluated and decisions are made (Knudsen, 2007, p. 73). Since there is large consensus that utterances and speaker's intentions are especially attention pre-empting, it can be argued that in utterance interpretation speaker-related information is likely to be mobilized in the service of self-sustained conscious processes and then possibly evaluated in working memory. This hypothesis, I claim, can solve the problem raised by relevance theorists.

In practice, next section will be devoted to state the problem and the way it is presumed to affect associative accounts; then I will argue that RT suffers of the same problem; and finally I will introduce my solution based on attention to the speaker and working memory.

ASSOCIATIVE ACCOUNTS AND SPEAKER-RELATED INFORMATION

One of the best-documented facts in psychology and neuroscience is the associative dynamic in virtue of which representations are accessible – and consequently spread activation – to each other as a function of the strengthening of the synaptic connections between them, with this strengthening being in turn a function of exposition to regular co-occurrence of stimuli (in accordance with Hebbian rule). Human brain is continuously engaged in extraction of regular patterns from experience at different levels of generality (for a general survey of this view of the brain as a Hebbian machine, see Fuster, 2003). Represented patterns in the cortex code for past configurations of items and, as a consequence, license expectations about future arrangements of those items. In practice, schemata are formed which prescribe how certain objects and properties are expected to combine with each other.

In spite of the obvious generality and pervasiveness of this fact, cognitive science has too often

presumed to explain this or that cognitive ability by invoking specialized automatic processes for which no independent evidence is given, instead of considering possible explanations based on associative processes. The modest interest of early cognitive science in the mechanisms underlying functional, high-level explanations has diverted attention away from low-level processes such as associative ones. Modularism has done the rest, by encouraging the postulation of an indefinite variety of specialized, possibly innate, processes whose internal functioning was largely left unaddressed (Perner, 2010).

Language studies have been a paradigmatic case of this tendency. Not only has Chomsky's Generative Grammar provided a functional account of language processing where little attention is paid to the details of its implementation; moreover, that account explicitly denies that the extraction of linguistic information from behavioural regularities has any explanatory role to play. However, for all we know, extraction of regularities and the subsequent dynamic of associative activation are the basic affair of the cortex. Therefore, even if Chomsky's argument from poverty of the stimulus were right in concluding that coding of regularities is not *sufficient* to explain language learning, one can reasonably expect that linguistic regularities of different sorts, syntactic ones included, are coded nonetheless and play some role in automatic processing of language.

In the last decades, much evidence has been generated in support of the claim that both infants and adults are much more skilled at detecting statistical regularities within complex stimuli – and specifically, auditory/linguistic stimuli – than was previously suspected (Baldwin and Baird, 2001, p. 174). In particular (see Butterfill and Apperly, submitted), the work of Saffran and her colleagues has demonstrated that, on the mere basis of sequential probabilities, subjects are able to segment continuous auditory stimuli into phonemic chunks and then assemble these chunks into word-sized units (Saffran et al., 1996); moreover they can extract clause- and sentence-sized complexes from these units on the basis of hierarchical patterns and prosodic cues (Newport and Aslin, 2004). For

these and other reasons, in contrast with what Generative Grammar has long assumed, the ability to extract patterns from linguistic input has gained renewed importance within theory of syntactic acquisition – though possibly in conjunction with a set of further skills related to intention-reading (Tomasello, 2003). Construction grammar and usage-based grammar are examples of this renewed interest of syntactic theory in extraction of regular patterns from experience (Goldberg, 1995, 2006; Barlow and Kemmer, 2000; Tomasello, 2003). In such approaches, grammatical constructions are conceived as more or less abstract schemata prescribing combinations of classes of words on the basis of previous experience. Jackendoff (2007, p. 11) has proposed that not only constructions but also ordinary (that is, Chomskyan) phrase structure rules can be expressed by the same formalism accounting for words, regular affixes and idioms: they can all be conceived as pieces of structure stored in long-term memory. All of those pieces of structure may drive and constrain language processing by participating to an associative dynamic of parallel activation and competition – this is the key point of so-called constraint-based theories in syntax and elsewhere. In fact, linguistic representations ('pieces of structure') are taken to contain within them the information on how they can be assembled with each other, therefore we do not need to rely on specialized linguistic (namely, syntactic) processes operating in accordance with procedural rules: all we need are domain-general processes which assemble representations in accordance with that self-contained information.

Recently, this line of reasoning has been pursued in pragmatics as well. In particular Recanati (2004) has claimed that pragmatic processing occurs, at least in part, by simple spreading activation between representations as a function of their reciprocal accessibility. In his opinion, this does not apply to implicatures as conceived by Grice (1989), that is, to inferences leading from the explicit meaning of utterances to their implicit, indirect meaning. Implicatures are rather conceived by Recanati as genuine inferential processes. However, as well as the other contextualists he holds that

implicatures are far from exhausting the field of pragmatic phenomena. Contextualists assume in fact that the explicit meaning of utterances depends on context not less than the implicit one, and this is where according to Recanati associative processes are needed. In his account, then, a distinction is made between 'primary pragmatic processes' delivering explicit meaning, which are associative in nature, and 'secondary pragmatic processes' delivering implicit meaning, which are inferential instead.

It is worthy to note that Recanati's primary pragmatic processes involve a dynamic of accessibility where a major role is assigned to schemata linking conceptual contents. Schemata are taken to account for the fact that spreading activation is not a wholly random process, that is, a process that merely depends on contingent, statistically based, relations of accessibility between words and concepts (Recanati, 2004; Mazzone, 2011). Instead, any word can also activate schemata which indirectly modify the initial pattern of activation for other words: an activated schema adds in turn activation to the concepts that happen to fit it and thus it introduces a principle of coherence by driving a search for the contextually relevant meanings of words. In sum, Recanati's schemata bridge contents in a principled manner, so as to mimic genuinely inferential processes (Recanati, 2007) by means of an associative device. In this sense his proposal converges with constraint-based approaches, according to which parallel activation and competition between representations substitute for procedural rules – as in Jackendoff's (2007) proposal.

In Mazzone (2009, 2011, in press a, b) I argued for an extension of Recanati's associative account in pragmatics, on the basis of two considerations. First, as noted by Carston (2007), the distinction between primary and secondary pragmatic processes is hardly tenable. In particular, the latter appear to be just as automatic as the former, and, most of all, the former can occasionally be as conscious as the latter, which makes Recanati's criteria for demarcation hardly applicable. Second, since Recanati has shown that associative processes can mimic inferential ones thanks to

schemata introducing rational motivations for selecting interpretations, one may wonder whether genuine inferential processes are needed at all. Therefore, I have explored the idea that associative explanations extend their range of application to pragmatic phenomena of any kind.

One case in point concerns how addressees take into account speaker-related information in utterance understanding. Carston (2007) has made a case against Recanati's associative account based on the following example. Let us suppose that Robyn is addressed by her student Sarah with the utterance in (1), and that Robyn knows two people called “Neil”, her young son (NEIL₁) and a colleague in the linguistics department where she works (NEIL₂).

(1) Neil has broken his leg.

Let us further suppose that Robyn is constantly worried about her son, so that her NEIL₁ concept is candidate to become more active than NEIL₂, whatever the circumstances in which (1) is uttered. According to Carston, the associative dynamics of activation cannot account for the fact that presumably, in that scenario, Robyn's preferred interpretation will be instead NEIL₂, insofar as she is aware that Sarah does not even know her son Neil, while she knows that Robyn has a colleague whose name is Neal. In other words, the challenge raised by this argument is whether associative explanations can account for the role presumably played, in understanding (1), by the addressee's (specifically, by Robyn's) knowledge of what the speaker (Sarah) knows.

In his reply to Carston (2007), Recanati (2007) has partially conceded the point by proposing that, in order to account for the influence of speaker-related information, mere associative processing must be complemented by the intrusion of a meta-representational step. In practice, in the course of associative processing, a proposition would be formed such as 'the speaker says that ...', thanks to which the identity of the speaker would be explicitly taken into consideration.

However, Mazzone (2011) has objected that in line of principle such a meta-representational step is not needed, since speaker-related information can be activated because of the mere fact that the addressee perceives the speaker as a salient component of the current communicative interaction. In other words, the hypothesis is that speaker-related information can be associatively accessed just as other information is. Mazzone (2011) made an appeal to Yeh and Barsalou's (2006) proposal that concepts have a situated nature: they are to be thought of as rich representations that may include a variety of events, entities and environments associated with their referents. Analogously, we can reasonably assume that our representation of specific people (Robyn's representation of Sarah, in the example) preserves information about related environments, events and entities — including other people they are connected to (Neil₂, in the example). Therefore, the particular world-knowledge at issue (that Sarah is acquainted with Neil₂ but not with Neil₁) could be accessed and produce its effects without any recourse to the meta-representational schema 'The speaker says that . . .'. There is in fact some evidence supporting the hypothesis that in conversation 'speaker-specific effects [on the addressee] emerge out of memory representations that incorporate episodic information, such as the identity of the speaker' (Shintel and Keysar, 2009, p. 264). Therefore, '[t]he effect of the speaker's identity may result not from a consideration of the speaker's communicative intentions, but from salience and significance of information about the speaker during encoding, allowing the speaker's identity to act as a rather potent retrieval cue' (Shintel and Keysar, 2009, p. 264).

All of this must not be taken to mean that sub-personal, associative processes are all there is to pragmatic processing. As a matter of fact, Mazzone (2011) explicitly stated that processes at a personal level may have a role to play as well. This role, however, was not further addressed, and it was suggested instead that examples such as the one proposed by Carston (2007) could not require anything more than associative access. Now, I have come to think there are reasons to doubt that

this is a proper account of the general case. The main reason is that, although associative processes are not actually unmotivated to the extent that the activation of schematic information promotes the search for rational coherence, this is not yet enough to ensure that those processes work in general as we would want them to do. It is one thing to argue that the appropriate speaker-related information can be associatively activated without any meta-representational step, it is quite another to show that this information associatively affects interpretation in the right way in most cases. To start with, as Mazzarella (2011, p. 17) has observed, one issue is what ensures that information associated to the speaker 'have the force to trigger the required accessibility shift'. If, in the example above, Robyn's NEIL1 concept has a significant advantage in accessibility over her NEIL2 concept, then the additional activation of the latter due to its association with Robyn's SARAH concept might not be sufficient to overcome the activation of the former. In other words, there seems to be no guarantee that, as a consequence of simple associative activation, speaker-related information may come to have the force to yield the desired effect. A related issue concerns the timing of processing: since associative activation is a rapidly decaying process one may ask what ensures that, as a rule, the relevant speaker-related information stays active until it can deliver the desired effect.

In general terms, the issue can be framed as follows: in contrast with the intuition that what the speaker knows (and intends) has a special importance in communication, associative accounts appear to put speaker-related information on a par with information of any other sort, so that, due to contingencies, the former can either prevail over the latter or not in the competition for activation. More basically, Carston (2007) observes that utterances are special in the way they pre-empt attention. As a consequence, a general, nonspecific mechanism such as spreading activation seems inadequate to account for the case of communication: '[m]ere accessibility, even coherence-based accessibility, doesn't seem sufficient to motivate the automatic investment of attention and effort

typical of our cognitive response to utterances' (Carston, 2007, p. 18). In sum, utterances pre-empt attention in a special way, and, in particular, they seem to focus the addressee's attention on the speaker.

If these intuitions are right – and, along with many others, I think they are – we had better provide an explanation of pragmatic phenomena that preserve them. And associative accounts seem lacking in this regard. In the next section, however, I will argue that Relevance Theory too lacks an explanation of how speaker-related information can play the expected role in comprehension. Then I will propose a solution for this problem which is compatible with an associative framework.

RELEVANCE THEORY AND SPEAKER-RELATED INFORMATION

RT has defended the thesis that pragmatic processes are genuinely inferential versus associative. It has also given great emphasis to the speaker's mental states, to the point that they are included as a crucial component in the definition of optimal relevance. According to RT, an utterance conveys to the addressee a presumption of its optimal relevance, where an ostensive stimulus is said to be optimally relevant to an audience if and only if (a) it is at least relevant enough to be worth processing and (b) it is the most relevant one compatible with the *speaker's abilities and preferences*. The expectation of optimal relevance is then taken to drive the comprehension process: the addressee is said to follow a path of least effort in accessing interpretative hypotheses, and stop when that expectation is satisfied.

The comprehension process is further analyzed in terms of the construction of hypotheses – concerning explicit content, contextual assumptions and contextual implications – that act as premises and conclusions in non-demonstrative inferences. This hypothesis construction is to be seen as taking place 'not in sequence but in parallel, with tentative hypotheses about context, explicit content and cognitive effects being mutually adjusted or elaborated as online

comprehension proceeds' (Wilson and Carston, 2006, p. 409). In the end, a successful overall interpretation has to meet two criteria: it is one that 'yields enough implications, at a low enough cost, to satisfy the hearer's expectations of relevance', and it 'is internally consistent in the sense that these implications are properly warranted by the context, the presumption of relevance and the enriched explicit content' (Wilson and Carston, 2006, p. 409).

As it seems, expectations of (optimal) relevance are key to the comprehension process, insofar as they both determine when that process has to stop and contribute to the assessment of consistency between contextual premises and conclusions. Since the speaker's abilities and preferences are in turn crucial in the definition of what optimal relevance is, then speaker-related information apparently plays a major role in RT's account of comprehension. One can ask, however, through which cognitive processes speaker-related information is presumed to play that role. It is important here to distinguish between two explanatory levels within cognitive pragmatics. One is the level of functional description, where one aims to provide an adequate conceptual analysis of the phenomena at issue; the other is the level of the actual cognitive mechanisms which are needed in order to account for those phenomena. On the one hand, functional descriptions are thus explanations in their own right, while, on the other hand, they set the requirements that a proper cognitive explanation has to meet. Now, as far as I can tell, the role assigned by RT to expectations of optimal relevance, and specifically to expectations concerning how the speaker's abilities and preferences affect communication, lie at the level of the functional description of communication. Therefore, one can without contradiction acknowledge that RT assigns a crucial role to speaker-related information at the level of functional description, *and* ask whether the actual cognitive mechanism described by relevance theorists does really account for that role.

The cognitive mechanism invoked by RT is essentially one of mutual adjustment between premises and conclusions of non-demonstrative inferences (in short, MAIS: Mutual Adjustment

between Inferential Steps). In the light of my previous distinction between functional description and proper cognitive explanation, I insist that one cannot simply take for granted that RT is in a better position than associative theories in order to account for examples such as (1) above, simply because RT has at its core the notion of optimal relevance (as is suggested in Mazzarella, 2011, p. 13). One should rather show that MAIS hypothesis has the resources to account for the role that, at a functional level, the notion of optimal relevance assigns to speaker-related information.

There is another point that requires careful consideration. When relevance theorists analyze a specific utterance interpretation, and specifically when they reconstruct the derivation that could justify that interpretation, speaker-related information can always be injected as contextual assumptions into the derivation. However, this injection of speaker-related information is indeed ensured by relevance theorists in the course of what can be legitimately called a rational reconstruction of the processes at play: it seems reasonable to any of us that the audience takes into consideration such and such information, and draw such and such inferences, and so on and so forth. But this leaves entirely open the question of how speaker-related information is injected into the derivation in the course of actual pragmatic processing – both because pragmatic processing is almost universally thought to be different from post-hoc rational reconstruction, and because, in any case, the fact that theorists are able to inject the right assumptions is far from being an explanation of how they do it. In sum, although relevance theorists can easily introduce speaker-related information in their account of specific utterance interpretations, and although this is coherent with their functional description of communication, what is of interest to us is the entirely separate issue of whether, and in case how, MAIS explains how speaker-related information gets to be processed and play the expected role.

It is easy to see that this explanation lies outside the scope of MAIS. Such hypothesis predicts that interpretation is ensured by a process of mutual adjustment between the premises and

conclusions of pragmatic derivations, but it does not account for how the information is selected in order to be injected into those premises and conclusions. Therefore, MAIS cannot explain how speaker-related information is accorded prominence over information of other sorts. In this sense, relevance theorists are subject to the very same objection they raise against associative accounts. To be sure, RT has not been entirely silent on the way in which information is fed into pragmatic derivations. What relevance theorists say on this issue, however, is clearly outside the core of their cognitive explanation (i.e., outside MAIS) and it is so generic as to be compatible with associative accounts. As we have seen before, according to RT the addressee follows a path of least effort in accessing interpretive hypotheses, and stops when her expectation of optimal relevance is satisfied. At times (e.g., in Sperber and Wilson, 2004) the first part of this claim is also phrased by saying that interpretive hypothesis are tested *in order of accessibility*. This could plausibly be intended to apply to each of the premises and conclusions in a pragmatic derivation: each of them is fed into the derivation because of the fact that it is the first to be accessed compared with other alternatives. As I said, this formulation is quite generic and entirely compatible with an associative framework. And in fact relevance theorists seem to be ready to assign such a role – peripheral as it can be – to associative relations, as is shown by some of their considerations in the domain of lexical pragmatics (for further details, see Mazzone, 2011). For instance, Wilson and Carston (2007, p. 243) state that, when the addressee listens to any given word, certain associated concepts 'are likely to be highly *accessible* as a result of frequent use' (emphasis mine), and therefore 'they are likely to be strongly activated'. Here accessibility is clearly thought of in terms of the associative strength resulting from frequency of use and affecting in turn the strength of activation. As it seems, for the relevant aspects RT's account of how information is selected appears to be strictly similar to associative accounts: it is based on a general notion of accessibility that in itself does not explain the special salience of speaker-related information. As a consequence, RT can hardly provide a

better explanation of the phenomenon than associative accounts do.

ATTENTION IN PRAGMATIC PROCESSING

In whatever way the initial salience of speaker-related information is to be explained, there is a further ingredient in MAIS which is worth considering. According to MAIS, both mutual adjustment between premises and conclusions and the final decision about the intended interpretation depend on a process of evaluation of the rational consistency between those premises and conclusions. Associative accessibility can possibly explain how information is fed into the derivation, but then there must be some further mechanism assessing whether the resulting derivation is consistent. Let us repeat Wilson and Carston's (2006, p. 409) words: a successful overall interpretation is one that 'is internally consistent in the sense that [its] implications are properly warranted by the context, the presumption of relevance and the enriched explicit content'. As the arguments offered by Carston (2007) and Mazzarella (2011) make clear, what the audience knows about the speaker's abilities and preferences is taken to have a crucial role in this process. This corresponds, as we saw, to a wide-spread intuition: in pragmatic processing, not only is speaker-related information preferentially activated, it also contributes to some sort of rational evaluation which seems to go beyond the contingencies of mere associative relations.

Does this assumption put RT in a better position than associative accounts? There are two complementary points that need to be addressed here.

On the one hand, RT does not provide a specific account of how such an evaluation process could be accomplished. Once again, just as for the selection of information, this idea lies outside the scope of MAIS proper. What MAIS predicts is that contextual assumptions, explicit content and contextual implications are fed into pragmatic derivations in order to be rationally evaluated, and then a mutual adjustment occurs if needed. But the mechanism by which the evaluation would be

accomplished is not addressed at all. In this respect, then, RT has no advantage over associative accounts: both need to be complemented by some plausible evaluation process if they want to be regarded as viable explanations of utterance interpretation.

On the other hand, there is a cognitive process that has strong independent evidence and that has the resources to explain the sort of rational evaluation that relevance theorists have in mind: it is the attention process. However, as I am going to argue, while attention can naturally play the envisaged role in associative accounts, it is less clear that the same applies to RT. Let me start from this last point.

Relevance theorists frequently refer to the fact that inputs, and especially ostensive stimuli, may deserve and in fact attract attention. A cornerstone of their approach is that relevance is a basic feature of human cognition, since cognitive systems need to pick out from the mass of competing stimuli what is worth being processed. As it should be clear, then, what is meant in the first place by saying that some inputs may attract attention is that they may be selected to be processed by the cognitive system, irrespective of whether this is actually accomplished by attention proper. In fact, in RT attention (to ostensive stimuli and other inputs) is not an explanans, it is rather what the theory aims to explain by postulating automatic heuristics oriented to the maximisation of relevance. In Wilson and Carston's (2006, p. 407) words: 'According to the Cognitive Principle of Relevance, human cognition tends to be geared to the maximisation of relevance, so that perceptual, memory retrieval and inferential processes are likely to include automatic heuristics for selecting potentially relevant inputs and processing them in the most relevance-enhancing way'. The inferential comprehension heuristic by which relevance theorists explain utterance understanding would be but one of those automatic, domain-specific mechanisms that they presume to account for the selective processing of certain inputs.

Coherently, relevance theorists never mention conscious attention as part of their explanation of

utterance interpretation. Their inferential comprehension heuristic is indeed presumed to be a modular mechanism whose functioning is not affected by domain-general processes (e.g., Sperber and Wilson, 2002). As it seems, in their view conscious attention could rather play a role as an *alternative* to the comprehension heuristic. This is suggested, for instance, by Carston (2007) in her critical assessment of Recanati's notion of conscious availability: in her words, the most likely role of reflective reasoning in communication/comprehension is 'as a backup mechanism when something goes wrong with the automatic intuitive mechanism of utterance understanding'. In sum, on the one hand RT does not specify a mechanism which is apt to ensure the rational evaluation of coherence between premises and conclusions, on the other hand it assumes that utterance interpretation can be entirely explained in terms of an automatic modular process, thus precluding a role for domain-general processes such as attention.

Quite on the contrary, simple associative processes are widely thought to interact with conscious attention in more than one way.

First of all, a given input is consciously attended as a consequence of competitive selection between a number of inputs, and this selection is thought to result from a simple associative dynamic of activation and competition. In his synthesis of the fundamental components of attention – where the assumption is made that attention involves the representation of inputs in conscious working memory – Knudsen (2007) summarizes this associative dynamic for competitive selection in the following way:

The selection of information for entry into working memory is a highly competitive process [...]. Information about the external world, from memory stores, and about the animal's internal state is processed extensively and automatically in parallel hierarchies of networks in the central nervous system. [...C]ompetitive selection reflects a computation that is intrinsic to a network, a competition for representation that is based on *the relative strength of activity (saliency) across*

the entire network. (Knudsen, 2007, pp. 69-70; emphasis mine).

This bottom-up, stimulus-driven mechanism is not all there is to the selection of relevant information. On the contrary, it is the case that what is already inside working memory may affect its future content in a top-down way. In Knudsen's words:

In the context of attention, not only does working memory accept, store, and manipulate information, but it also generates signals that improve the quality of the information that it processes (Knudsen, 2007, p. 62).

This can be done in two quite different ways: by directly orienting movements (gazes, etc.) toward targets, and by modulating the sensitivity of neural circuits that represent information (Knudsen, 2007, p. 62). In both cases, working memory apparently tends to ensure a certain consistency between its current focus of attention and future contents. The latter case – that is, the modulation of sensitivity of neural circuits – means in practice that the accessibility of certain information is enhanced due to its connection with the current content of working memory. This enhancement in accessibility can result either in new information being fed itself into working memory, or else in unconscious processing of that information. A well-known model of this sort of unconscious 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, p. 293, puts it). More precisely, consciously attended action plans are thought to contain free processing parameters that need to be filled by information which may then be searched for and processed in a wholly unconscious way. Thus, conscious attention can drive unconscious processing thanks to associative priming of the stimuli to be

searched for.

As it seems, associative and attentional processes are expected to interact both in bottom-up and top-down direction: as for the former, the object of attention is at least in part determined by an associative process of competitive selection; as for the latter, what is in the focus of attention associatively modulates the sensitivity of neural circuits and thus the kind of information that is either consciously or unconsciously processed. For these reasons, it seems that associative accounts are naturally compatible with cognitive explanations based on attentional processes.

In fact, what I suggest is that attentional processes must be added to simple associative accounts in order to explain how speaker-related information come to gain prominence in pragmatic processing and, especially, in the evaluation of consistency that either leads to mutual adjustment of premises and conclusions or causes the process to stop. Based on the above considerations, my argument is the following. There is large evidence that, in communication, the speaker, her utterance, and her goals are largely attended to by the addressee. In other words, representation of the speaker is apparently favored in the process of associative selection between competing inputs, and thus it normally gains access to working memory. As a consequence, such a representation can in turn modulate the associative sensitivity of neural circuits representing information and, more specifically, it can activate speaker-related information in the service of the conscious evaluation of information that is a characteristic feature of working memory.

The premise of this argument, as I have already said, is a wide-spread assumption in language studies. That ostensive stimuli pre-empt attention, and that this is related to a concern with the speaker and her possible mental states, are also key claims of RT – although, as we saw, these claims are not intended to mean that any explanatory role is to be acknowledged to attention proper. However, the very same assumption is made in the literature on joint attention and its role in the emergence of communication. What is at issue in this domain is essentially overt attention – that is,

attention as manifested by behavior – which is considered a visible cue of the genuine process of attention (Hommel, 2010, p. 128). In joint attentional frames a subject is mainly concerned with other agents and the focus of their attention. The reason why this is thought to be crucial for the emergence of communication is that a subject can exploit the others' attention to the focus of her attention in order to produce ostensive stimuli. In this case, a subject *A* makes manifest her intention to attract the attention of a subject *B* towards an object *C*, so that *B* not only directs his attention to *C* but also becomes aware of *A*'s intention that this occurs. Both the attention of *A* and *B* is then focussed on each other together with the aspects of entities in the shared situation that are made “mutually manifest” and so potentially “relevant” for acts of interpersonal communication (Sperber and Wilson 1986)' (Tomasello et al., 2005, p. 683).

In practice, in communication each subject directs her attention to the other and its communicative goal(s). Insofar as the subjects involved in communication tend to make manifest to each other their reciprocal attention, this is a case of overt attention and thus, it seems, a visible cue of genuine attentional processes. That is, the attention to others that is at play in communication plausibly amounts to conscious representation of others in working memory. Thus, the premise that the speaker and her utterances pre-empt the attention of addressees is to be understood in terms of the cognitive process of attention proper. Consequently, for all we know about this process, we can presume that the conscious representation of the speaker in working memory activates speaker-related information in a way that is relatively stable: a key function of working memory is in fact the active maintenance of information and its shielding from distracting inputs, in order to ensure some consistency in processing and in action (Hommel, 2009). The stable representation of the speaker, then, can be expected to cause the activation of speaker-related information in the service of both conscious and unconscious processing.

A reasonable assumption that is implicit in this argument is that concepts have a situated nature

in the sense defended by Yeh and Barsalou (2006): they include, or are connected to, a variety of events, entities and environments associated with their referents. As a particular case, our representation of specific people can be expected to preserve information about related environments, events and entities, so that this information can be activated by the representation of a given person in working memory.

Working memory, as we saw, is described as a cognitive space where information is evaluated and 'analyzed in detail' and where 'decisions about that information can be made' (Knudsen, 2007, p. 58). Processes of this sort could account for the intuition that pragmatic processing involves a final evaluation of consistency of the attempted interpretation, in which information about the speaker may play a crucial role over and beyond simple associations in the addressee's memory. Once this and other information enter working memory, they can be attentively confronted with each other. As Knudsen (2007, p. 64) puts it, the importance of information that enters in working memory 'can be evaluated and compared with the importance of other information already being processed in working memory'. Considerations of this sort suggest that working memory enables forms of conscious monitoring of information which parallel analogous associative operations. After all, both evaluation of relative importance (salience) and evaluation of consistency (thanks to the search for coherence driven by schemata) are something that occurs within spreading activation processes as well. What is different in attentional processes is essentially the active maintenance and conscious monitoring of information. Given the similar nature of the processes, it could be of no much importance whether speaker-related information gains access to working memory; what really matters is that representation of the speaker does. This is enough to ensure that speaker-related information gains prominence in processing, and plays a special role in evaluation and decision making occurring in working memory. For instance, in Carston's (2007) example, it is not necessary that Robyn consciously represents in working memory the fact that her student Sarah knows Neil₂

but not Neil₁. The simple fact that Robyn has a conscious representation of Sarah might explain why this information strongly affects the evaluation of the utterance interpretation. In other words, when Robyn makes the conscious decision that

(2) NEIL₂ HAS BROKEN HIS LEG

is the right interpretation of Sarah's utterance 'Neil has broken his leg.', this decision does not require that Robyn represents in working memory the relation between Sarah and Neil₂. This is a possibility, but in order to ensure that this information properly affects the decision it is sufficient that Robyn's NEIL₂ concept is reliably prompted by the representation of Sarah in working memory. In sum, by interpreting attention to the speaker in terms of conscious representation of the speaker in working memory, we can explain how speaker-related information may associatively gain prominence in the service of the final evaluation of consistency, whether or not this information itself gains access to working memory.

CONCLUSIONS

The role of conscious attention in language processing has been scarcely considered, despite the wide-spread assumption that verbal utterances are designed by the speaker so as to attract and manipulate the addressee's attention, and that they normally succeed in doing it. In the present work I have analysed the possibility that this assumption is in fact understood in terms of attentional processes proper. This hypothesis can explain a fact that has been noticed by relevance theorists in pragmatics: the special role played by speaker-related information in utterance interpretation. RT has considered this fact as an objection to associative accounts in pragmatics. However, I have argued, RT is affected by this problem not less than associative accounts are. Moreover, relevance theorists tend to exclude that attentional processes have a role to play in normal utterance

interpretation, therefore they cannot rely on these processes in order to provide an explanation of the prominence of speaker-related information.

On the contrary, associative accounts can naturally be complemented by attentional processes. More generally, analysing the cooperation between associative and controlled processes can be expected to give us a better comprehension of the cognitive role of the former: associative processes in humans do not work in isolation. The same can be said, on the other hand, for conscious processes. They are less independent from associative processes than it is ordinarily thought (Mazzone and Campisi, in press).

In my proposal, attentional processes can explain the prominence of speaker-related information in utterance interpretation thanks to the fact that representing the speaker in working memory reliably enhances the activation of speaker-related information and, consequently, the role played by this information in determining the content of the interpretations then evaluated in working memory.

References

- Baldwin, D., & Baird, J. A. (2001). Discerning intentions in dynamic human action. *Trends in Cognitive Science*, 5, 171–178. doi: 10.1016/S1364-6613(00)01615-6
- Barlow, M., & Kemmer, S. (eds.) (2000). *Usage-Based Models of Language*. Stanford, CA.: CSLI Publications.
- Butterfill, S. A., & Apperly, I. A. (in press). How to construct a minimal theory of mind. *Mind and Language*.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Carston, R. (2007). How many pragmatic systems are there?. In M. J. Frapolli (Ed.), *Saying, Meaning, Referring. Essays on the Philosophy of François Recanati* (pp. 1–17). New York: Palgrave.
- Fuster, J. (2003). *Cortex and Mind*. Oxford: Oxford University Press.
- Goldberg, A. E. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*. University of Chicago Press: Chicago.
- Goldberg, A. E. (2006). *Constructions at Work. The Nature of Generalization in Language*. Oxford University Press: Oxford.
- Grice, P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hommel, B. (2010). Grounding attention in action control: The intentional control of selection. In B.J. Bruya (Ed.), *Effortless attention: A new perspective in the cognitive science of attention and action* (pp. 121-140). Cambridge, MA: MIT Press.
- Jackendoff, R. (2007). A parallel architecture perspective on language processing. *Brain Research*, 1146, 2–22. doi:10.1016/j.brainres.2006.08.111
- Kiefer, M. (2007). Top-down modulation of unconscious ‘automatic’ processes: A gating

- framework. *Advances in Cognitive Psychology*, 3(1-2), 289–306. doi: 10.2478/v10053-008-0032-2
- Knudsen, E. I. (2007). Fundamental components of attention. *Annual Review of Neuroscience*, 30, 57–78. doi: 10.1146/annurev.neuro.30.051606.094256
- Mazzarella, D. (2011). Accessibility and relevance. A fork in the road. *UCL Working Papers in Linguistics*, 1, 11–20.
- Mazzone, M. (2009). Pragmatics and cognition: intentions and pattern recognition in context. *International Review of Pragmatics*, 1 (2), 321–347. doi: 10.1163/187730909X12535267111615
- Mazzone, M. (2011). Schemata and associative processes in pragmatics. *Journal of Pragmatics*, 43, 2148–2159. doi: 10.1016/j.pragma.2011.01.009
- Mazzone, M. (in press a). A pragmatic Pandora's box. Regularities and defaults in pragmatics. In F. Liedtke, & C. Schulze (eds.), *Beyond the Words*. Berlin: Mouton de Gruyter.
- Mazzone, M. (in press b). Controlled processes in pragmatics?. In A. Capone, F. Lo Piparo, & M. Carapezza (Eds.), *Perspectives on Pragmatics and Philosophy*. Springer.
- Neumann, O. (1990). Direct parameter specification and the concept of perception. *Psychological Research*, 52, 207-215. doi: 10.1007/BF00877529
- Newport, E. L., & Aslin, R. N. (2004). Learning at a distance I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, 48(2), 127–162. doi:10.1016/S0010-0285(03)00128-2
- Perner, J. (2010). Who took the cog out of cognitive science? Mentalism in an era of anticognitivism. In P.A. Frensch, & R. Schwarzer (Eds.), *Cognition and neuropsychology: International perspectives on psychological science* (Vol. 1, pp. 241–261). Psychology 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* (pp. 49–54). New York: Palgrave.

- Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Statistical Learning by 8-Month-Old Infants. *Science*, *274*(5294), 1926–1928. doi: 10.1126/science.274.5294.1926
- Shintel, H., & Keysar, B. (2009). Less is more: A minimalist account of joint action in communication. *Topics in Cognitive Science*, *1*, 260–273. doi: 10.1111/j.1756-8765.2009.01018.x
- Sperber, D., & Wilson, D. (1986/1995). *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Sperber, D., & Wilson, D. (2002). Pragmatics, modularity and mind-reading. *Mind and Language*, *17*, 3–23. doi: 10.1111/1468-0017.00186
- Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Cambridge: Harvard University Press.
- Tomasello, M. (2003). *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Cambridge (Mass.) and London: Harvard University Press.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, *28*, 675-691. doi: 10.1017/S0140525X05000129
- Wilson, D., & Carston, R. (2006). Metaphor, relevance and the emergent property issue. *Mind and Language*, *21* (3), 404–433. doi: 10.1111/j.1468-0017.2006.00284.x
- 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.
- Wilson, D., & Sperber, D. (2004). Relevance theory. In L. Horn, & G. Ward (Eds.), *Handbook of Pragmatics* (pp. 607–632). Oxford: Blackwell.
- Yeh, W., & Barsalou, L. W. (2006). The situated nature of concepts. *American Journal of Psychology*, *119*, 349–384.