

RACHEL GIORA. *On our mind: Saliency, context, and figurative language*. Oxford University Press, 2003. 259 pp., including Bibliography and Indexes. US\$55.00.*

Reviewed by Jacob L. Mey

Rachel Giora's latest book collects the author's theoretical and experimental studies in recent years, in which she has focused mainly on the notion of 'saliency' in discourse.

Saliency can be characterized as the way a word's (or an expression's) particular meaning 'leaps to the fore', irrespective of literal meaning and sometimes even independently of context. The theory thus places itself in the center of the age-old dispute on how we understand language and, in particular, how we come from the said to the meant, from the literal to the non-literal meaning. Being an experimental psycholinguist (which the present reviewer is not), the author checks her theorizing against laboratory experiments and provides continuous evidence for the plausibility of her hypotheses.

For people like myself (and, I presume, for much of the readership of *RASK*), the main interest of Giora's work lies in the ways she combines theory and practice, hypotheses (which are formed in order to capture some observations) and experiments (in which the author's hunches are subjected to psycholinguistic testing). Especially for a pragmaticist, Giora's book provides a welcome antidote against the commonly held view that 'the context is everything' (sometimes called 'pancontextualism') – an adage that would have to be corrected, in the spirit of the book, to: 'the context is everything, provided one can properly deal with its salient elements'.

The way Giora structures her arguments is standard in psychology and psycholinguistics research work: first presenting the data informally, then forming some hypotheses, which are tested in a laboratory setting with a large number of subjects; following that, the results are tabulated and discussed, to be followed up by a general conclusion based on the confirmed findings.

In the present case, the hypothesis that salient meanings always carry precedence when we have something 'on our mind', is borne out (albeit with some reservations) by the experiments. Giora concludes

that we have to rethink our position vis-à-vis the processing of figurative language in general (including irony, sarcasm, hyperbole, jokes, and other kinds of non-literal language). This rethinking will take into account the fact that salience, as it is explained in Giora's book, can neither be pinned down to the surface form of a word or expression, nor to its syntactic function alone (e.g. salient items will often be noun phrases, but this is by no means an exclusive relationship). Neither can salience be connected to a particular semantic property or feature. One could perhaps say that salience, insofar as it depends on the user's intentions, has a contextual and pragmatic aspect. Salience can only be realized in context, but its effect is not uniquely dependent on context, let alone context-determined: rather, we should say that salience, when it is optimally realized, depends on the interplay between one user (who intentionally creates, or capitalizes on, salience) and (an)other user(s), who recognize the salient item as such, and structure their response in relation to this uptake, and thus, more generally, demonstrate 'the effect of accessible meanings on speech production and comprehension' (p. 9).

The book consists of a Preface (3 pages), followed by 9 chapters. Chapter 1, entitled 'Prologue' (pp. 3-12), introduces the reader to the problem by means of a couple of anecdotes and jokes, showing how we, for instance, in interpreting a joke, access the salient meaning of the punch word or line first, and only afterwards adjust our interpretation so as to fit the appropriate, joking context (puns and 'garden path' stories belong here, too – as when we realize that the 'baby snatchers' we hear about are not the snatchers of babies, but instead the baby counterparts of some adult snatchers; p. 7).

The main question here is about the role of the context: will it constrain a particular interpretation of an utterance or word sufficiently to rule out all other meanings (including the salient one)? Giora's answer is clearly in the negative; and exactly what (kinds of) effects 'accessible meanings [have] on speech production and comprehension', and how these effects operate, is defined as the goal of the book itself (p. 9).

In the remainder of the chapter, Giora details three models describing the role of context in speech comprehension and production: the 'interactionist', the 'modular', and the 'graded salience' model. The *interactionist* model (subscribed to by most theorists) says that we access meaning directly, depending on the context. If the context is sufficiently rich, there will be no, or only one, choice; in an ironic context, the utterance 'Thank you for your help' would only be interpreted ironically, not literally.

In contrast, the *modular* model assumes that we access all meanings of an ambiguous expression simultaneously, and then proceed to filter out the incompatible ones, by 'revisiting' the scene of the crime and adjusting our expectations (more or less as we are supposedly doing when handling conversational implicatures or indirect speech acts, in the 'standard pragmatic model' due to Searle and Grice).

Finally, under the *graded salience* hypothesis we do indeed access the lexicon directly upon encountering an item in our discourse chain; but while we tend to allow the context to guide our possible interpretations, at the same time we access meanings 'bottom-up' (pp. 11-12), as in the modular model. However, the access is not random or across the board, but 'graded': the degree of salience of the proposed occurrence determines our first choice of meaning. This access is not dependent on context, but runs in parallel with it (*ibid.*). For instance, to access the proper meaning of 'Rachel' we seemingly use the contextual information alone: we know that Rachel is a woman's name, and thus we can make the proper reference, given the context. But, as Giora points out (p. 12), this access is better explained by the fact that the non-salient meaning of Rachel ('sheep' in Hebrew) 'is not salient enough, hence of low accessibility'. However, even these hard to access meanings are available (and accessed) in case of necessity; compare (my observation, JM) the use that smart ad people have made of the word 'mouse' in connection with advertisements for 'computer literacy' ('Who is afraid of this mouse?'), where the non-salient meaning, the 'animal', is reinforced by a picture of a mouse-like tail extending from the computer device called 'mouse'.

Chapter 2, 'Salience and context' (pp. 13-38) looks into the role of context vs. salience in comprehension and production of language.¹ It turns out that in processing incoming speech, we most often prefer the

immediately accessible interpretation ('what's on our mind') to an interpretation which contextually would be more appropriate, as when we resolve anaphora by 'backtracking' (à la Halliday & Hasan) to the nearest available suitable referent (for a critique of this procedure, see Mey 2001:194).

A word of caution is in order here: if we simply take 'salient' as 'that which is most easily accessible', then we find ourselves in a circular definition: it is a tautology to say that 'salient' items are accessed first, if we have defined 'salience' as 'being accessible', or simply 'on our minds'. Therefore, the author uses the next section (2) of the chapter to refine the definition of 'salience'.

Salience is, first and foremost, a 'graded' concept: it is not 'all or nothing'. Its degrees depend on a variety of factors: frequency of occurrence, familiarity, conventionality, or proto- or stereotypical character; the more the meanings present in the mind of an individual or a community exhibit these characteristics, the more salient the information will be (p. 15). Giora then explains these features in some detail before embarking on a discussion of 'gradability'.

An interesting property of salience, which only is mentioned in passing, is that of being 'consolidated', that is to say, a salient item must be 'stored or coded in the mental lexicon'; as such, it is 'superior to unstored information such as novel information or information inferable from the context' (p. 15). But since frequency has to do with occurrence, and occurrence is always in a context, some influence of context upon salience seems to be 'built-in' into the very concept; this makes it harder to subscribe to the idea that salience is 'directly computable from the mental lexicon irrespective of inference drawn on the basis of contextual information ...' (p. 18).

The author seems to indicate something similar when she discusses the possibility of salience being reduced ('... [the "insect" meaning of "bug"] is now less salient given that it catches us by surprise at the end of the sentence ["the millennium bug bites"]; p. 18). The difficulty lies in the interplay between lexicon and context as regards salience: on the one hand, salience is a matter of degree, depending on the lexicon (which is 'hierarchically structured', *ibid.*); on the other hand, an item can be more or less salient in accordance with the context.² However, if a concept may become more salient precisely

because it is more available or accessible (cf. above, where 'frequency' of occurrence and 'familiarity' were said to be features of salience), we still aren't sure about the relationship between these two concepts. The main question is, then, if salience is an inherent trait of a word or expression, and if this is the basis for its 'resilience', even in the face of otherwise biased contexts.

This brings me to the more important matter of salient vs. less salient meanings in the case of idioms and composite expressions. Here, the salient meaning elements do not always guide us towards the meaning of the whole; on the contrary, such compositional meanings may assume a life of their own, as it is most clearly seen in the case of idioms (the salient meaning of 'bucket' in 'to kick the bucket' is not helpful for our understanding of the idiom). However, the opposite may also occur: an intended compositional meaning may be difficult to understand because the meaning relies on a less salient aspect of one or more of the components. Hence it seems safe to conclude that salience is not exclusively a matter of the lexicon and 'lexical coding'; the constraint that the author here poses on the concept may well be unfounded, and it seems also unnecessary, since compositional phrases that are not lexicalized, still may exhibit some salient characteristics, given the proper context.³

What emerges from this discussion ('What salience is not'; cf. chapter 2.4) seems to be that salience is not just (some kind of) accessibility.⁴ What we need salience for is to explain why certain meanings, in force of their salience, persist even in environments where they should 'normally' be extinguished. Salience should not be construed as belonging exclusively to lexical items (despite the prevalence of the 'lexical coding' mentioned earlier): on many occasions, the meaning of the concept is the 'salient' part of the expression, not the literal expression (the word) itself. Hence, the author is right in observing that 'the graded salience hypothesis does not subscribe to any unified view of the mental lexicon as containing just one type of entry' (p. 21); this is especially important to retain in the cases where predictive context motivates the anticipation of one entry rather than another ('meanings and concepts rather than ... specific words', p. 23). But also, one should remember that salience cannot be blocked by context: '[w]hile context may be predictive of

certain meanings, it is deemed ineffective in obstructing initial access of salient information' (p. 24) And this leads us to the (in my opinion) most important claim of the graded salience hypothesis, viz., 'that salient meanings are processed automatically (though not necessarily solely) irrespective of contextual information and strength of bias'. It is this idea that is fruitfully exploited in the following chapters, where the author explores

how salience affects ambiguity resolution vis à vis contextual information, that is, how we home in on the contextually appropriate meaning when we encounter words ... which have multiple meanings but only one that is contextually appropriate. (p. 24)

It is in this field of tension between salience and context that the book makes its most worthwhile contribution to our understanding of language comprehension and (to a lesser extent) of language production.

Chapter 3 is entitled 'Lexical Access'. In this chapter, the author takes up the question already adumbrated in the previous chapter, viz., that of context and constraint. 'Can we really constrain context to the extent that it would affect comprehension entirely so that only relevant meanings would be processed, neither more nor less?' (p. 39).

As the hedge 'really' indicates, Giora does not answer this question in the affirmative. 'Direct access' (as the single interactive model is called) fails to account for the fact that salient meanings 'escape' the constraints of contexts; in other words, '[t]hough a predictive context may avail [sic] appropriate information speedily, it cannot obstruct access of salient information when inappropriate, ...' (ibid.) The rest of the chapter is dedicated to 'an attempt to resolve this lexical versus context effects debate' (p. 40).

Giora then reviews several models that have been proposed to explain the 'rapid activation of word meanings operating upon encountering a linguistic stimulus in and out of context' (p. 40) – the phenomenon called 'lexical access'. In the *exhaustive access* model (a

variant of Fodor's modular-based model), salient meanings are accessed independently of context. This model is most outspoken in its positing of two independent agencies for comprehension: one bottom-up, the other top-down: 'lower level processes such as lexical access should not be affected by top down feedback from higher non-lexical representations such as contextual or world knowledge' (p. 41; see also p. 50). Since the independence of the lexicon also enters into the conceptual array of the graded salience hypothesis, the latter is consistent with the modular approach, while the former is unable to explain the 'graded' character of salience: it is an all-or-nothing model and thus intuitively less plausible.

The *ordered access* model similarly maintains the independence of the lexical processes, but allows for salient meanings to be accessed faster than less salient ones (p. 40).⁵ But also here, 'there is ample evidence ... that contextual information does not inhibit salient meanings' and that 'salient information cannot be bypassed' (even when access is ordered) (p. 52). By the same token, the *reordered access* model allows for both salient and non-salient meanings to be affected by context; however, 'salient meanings are not preempted' by any contextual influences (p. 53).

The *selective access* model is the one that poses the greatest challenge to the graded salience hypothesis. This model (which is sometimes called 'interactive', at other times 'interactionist'; p. 40) clearly favors an interactive conception of lexical accessing, in that both the lexicon and the context are activated simultaneously and that the two interact from the very beginning. The meanings thus selected then exclude other, less appropriate meanings:

... heavily weighted contextual information interacts with lexical processes very early on and activates contextually appropriate meanings exclusively ... contextual information blocks contextually inappropriate meanings. (p. 42)

This model is actually favored by many cognitive psychologists (for references, see p. 42), and its intuitive appeal makes it hard to refute globally.

In her defense of graded salience (an hypothesis which is clearly inconsistent with the selective access model on a number of counts), Giora points to the possibility of alternative accounts. While some experiments *prima facie* seem to contradict the graded salience hypothesis in that they apparently presuppose a strong connection between context and lexicon (see Vu et al., 1998, 2000; quoted p. 47), the contextual item's facilitated access could be due to its sentential placing. When Giora and her associates (2001) replicated the experiments, but located the probe not at the end of the sentence, but *before* the test word, it turned out that subjects made the right guesses anyway, even without having seen the ambiguous word itself (p. 48). Thus, it became clear that it was not the lexical item, 'the target (ambiguous) word that primed the probes but rather the context alone' (ibid.), thus confirming the graded salience hypothesis with its two independently operating mechanisms: lexicon and context. But, as the author rightly affirms, the existence of such contextual influences makes it necessary 'to look more carefully into effects on comprehension of weighted contextual information, ... and location in the sentence context' (p. 50).

In the remainder of the chapter, the author discusses the influence of syntax on disambiguation. Earlier, the current account was that syntax steered this process by referring to tree-like structures (the 'garden path model', as Giora calls it).⁶ In Giora's view, this explanation is insufficient: one needs to count with lexical information as well, and here salience again is important. In fact, syntactic ambiguity may even depend on the availability of lexical alternatives and here, the more frequent (hence salient) interpretation prevails (p. 56). While there thus seems to be 'some evidence suggesting that syntactic ambiguity resolution is, at least in part, a function of meaning salience' (p. 58), more research is needed to tease out the different explanatory strands. One thing that is important to retain is that the automatic character of 'salience-based retrieval' (including 'suppression' and 'retention') need not imply that the lexicon acts like a robot or even that the mind is inherently 'stupid'.⁷ Automatic access of salient information is a 'rational process': the lexicon has 'a mind of its own' (p. 50).

The next chapters apply the theory of 'graded salience' to some concrete linguistic phenomena, in particular figurative language. Chapter 4, on 'Irony', first discusses various theories of irony, said to be 'consonant with either the interactionist, direct access view, or the modular view' (p. 62), discussed in the previous chapters. Consequently, the importance of the context as all-important for the processing of irony overshadows possible salience: in an ironic context, only the ironic meaning of an utterance is retrieved, while in a literal context, only the literal meaning is processed (p. 63). The processes are similar except for the context in which they occur.

This assumption leads to difficulties for some theoretical frameworks, such as the 'echoic mention' theory of irony, due to Wilson and Sperber (1992, 1993). Here, processing time for echoic irony is less than for non-echoic irony (cf. experiments due to Gibbs 1986). What happens in echoic irony is that a previous, non-ironic mention of the (sense of the) utterance to be processed reduces processing time for the ironic utterance that follows (the 'priming effect', p. 65). But this supposes that the processing still has access to the 'salient, literal meaning of irony', and thus would support a theoretical account that takes 'the linguistic meaning of what is said' (*ibid.*) seriously.

Incidentally, it is in the discussion of one of the direct access models (the 'allusion-pretense' view, p. 65) that Giora introduces the concept of 'pragmatic insincerity', which I assume may be translated as 'not behaving in accordance with the situational expectations while pretending to do so'. However, since Giora mentions 'social norms or expectation', which are typical situational features, this leads one to consider the 'situational' nature of irony as such. Unfortunately, in her discussion of irony, Giora does not refer to this concept or to the literature dealing with situational irony. If, as I have claimed in earlier work (Littman and Mey 1991), all irony basically depends on an ironic situation, such that verbal irony is only a derivative phenomenon, then of course the 'interactionist' view would be strengthened. One would like to hear Giora's views on this, especially in light of the renewed interest in the irony of the situation, as evidenced in the work of Shelley on the 'bicoherence theory of irony' (2001).

As to the modular models, they agree in positing an initial processing of literal meaning as indispensable in the activation of a subsequent ironic meaning. This is most clearly seen in the 'standard pragmatic model', due to Grice and Searle (among others), where 'understanding nonliteral [e.g. ironic, JM] language involves a sequential process [, where t]he first stage is literal and obligatory, and the second is nonliteral and optional' (p. 66). This model not only involves longer reading times for nonliteral than for literal meanings (as borne out by experimental findings), but also presupposes a 'suppression process' which 'has not gained empirical support' (p. 66; see further below on suppression and retention).

The same two-layered communicative act is at the basis of another model, called 'joint pretense' (due to Clark 1996; Clark and Gerrig 1984). An interesting aspect of this model is that it places the use of irony in an 'active' setting: 'joint pretense is conceived of as a staged communicative act' (p. 67), whereby (in my interpretation) the ironic communication may stand on its proper footing in a theory of pragmatic acts (Mey 2001:212-229). With regard to the double layer of processing, 'in which the salient literal meaning is activated and retained by both the speaker and the addressee, who reject it as the intended meaning though they pretend otherwise' (p. 67), this model comes close to Giora's own assumptions about 'retention of the contextually incompatible literal meaning' (p. 68), one of the cornerstones of the graded salience hypothesis.

In fact, this hypothesis assumes

a bipartite view of utterance comprehension including an initial phase, involving meanings accessed directly or made available by a strong prior context ... and an immediately subsequent phase of integration of activated information with contextual information. (p. 69)

In the following sections, the author discusses these phases with regard to the predictions made by the graded salience hypothesis. For instance, when processing a 'less [familiar] or unfamiliar irony' (p. 70), the salient, literal, incompatible meaning is activated in all circumstances, and mostly even before accessing the contextually compatible,

ironic meaning; in the case of an ironic context, an 'adjustment' will be necessary. In contrast, when we are dealing with a familiar irony (e.g. 'big deal' for something unimportant), both meanings (literal and ironical) are activated, since they are both salient; however, in an irony-based context, the ironical reading will prevail, and no adjustment is necessary (pp. 70-71).

It is not quite clear what Giora means by 'adjustment' in this connection. It is easily seen that salient and nonsalient interpretations of the same utterance require different processes (p. 71), and that this difference has to do with determining which of the interpretations (salient vs. nonsalient) is appropriate, given the context. But also, in light of the fact that elsewhere, Giora has talked about the dangers and 'costs' of suppressing one particular reading (chapter 3), it is not easy to see how 'adjustment' differs from 'suppression' (or contradicts its opposite number, 'retention').

After the initial phase of comprehension, the second phase, that of integration, is technically out of reach for the graded salience hypothesis, which deals explicitly with only the first phase. However, Giora furnishes an additional set of hypotheses, based on the view that irony is a kind of 'indirect negation' (p. 72): irony 'is used to implicate that a specific state of affairs is different or far from the taken for granted, expected (or more desirable) state of affairs made explicit by the expression' (*ibid.*). In this view, the utterance's literal reading should be 'retained' after initial processing, inasmuch as it is necessary for comparison between 'the explicit and the derived messages, so that the dissimilarity between them may be computed' (p. 72). While this 'retention' is plausible, and actually activates one of the graded salience hypothesis' assumptions, it is less clear what 'computing the difference' means, and how it is distinct from 'suppressing' the negated meaning. After all, the computation (whatever it stands for) will result in a preference for one reading, and by the same token a dispreference (or rejection, or suppression, or replacement) of the other one.

The theoretical speculations referred to above constitute the first part of the present chapter. In a second part, the author discusses the various empirical findings that have been obtained in relation to the initial processes and the subsequent phase of integration of comprehending irony. The two main parameters in this experimental

research have been *reading time* (whereby, e.g., 'as predicted by the salience-based view ..., utterances took longer to read when embedded in an irony than when they were in a literal inducing context'; p. 74), and *response time* (whereby a subject's reaction to a probe (test word) is measured as an indication of salience; p. 77).

While this second part of the chapter is naturally more oriented to the empirical findings (again divided according to the original separation between initial and integrative phases), much of the discussion of the examples repeats what already had been said in the first part. Actually, it is very difficult to separate the two; for instance, while discussing reading times, the author comes up against a crux: Gibbs' findings about processing times for sarcastic utterances seem to contradict her own hypothesis, stated previously. According to Gibbs (1986), 'irony (sarcasm) did not take longer to read than its literal "counterpart"' (p. 75); and not only that, but nonsarcastic utterances took longer to read than the ironic, sarcastic ones (e.g. 'You are not helping me' vs. 'You are a big help'; p. 77). The author explains this apparent contradiction in terms of 'discourse well-formedness': the non-sarcastic utterance is less well-formed than the literal one, hence takes longer to process (*ibid.*) I find the use of this term ('well-formedness') somewhat suspicious; as long as we are not told what is meant by 'well-formedness' (itself a relic from the glorious days of quasi-exact linguistic thinking), it is difficult to judge the value of the term as an explanatory device.

More or less the same could be said about the section on the 'retention hypothesis' (section 2.2 of this chapter, pp. 87-94). Here, too, an interesting hypothesis ('a meaning activated initially is retained for further processes if it is instrumental in constructing the intended interpretation, regardless of contextual compatibility; it is suppressed if it interferes with the process'; p. 87) is tested for response times after (a long) delay (pp. 87ff). While the experiments (at least to the present, admittedly lay, reviewer) seem convincing, the mechanism of retention vs. suppression could have deserved a stronger theoretical underpinning (especially considering the author's earlier remark on suppression as not 'having gained empirical support'; p. 66).⁸

The paragraph concluding the section serves as a good comprehensive summary of the matters discussed so far:

Interpretation of unfamiliar, innovative language, then, is a complex process. It involves activating the salient contextually incompatible meaning initially either before or alongside the contextually compatible nonsalient meaning. In the case of irony, it also involves retention of both meanings for the purpose of a contrastive comparison. (pp. 93-94)

In the chapter's penultimate section, entitled 'What does irony mean?', Giora casts her net wider. Irony is not just a matter of interpretation in accordance with a particular hypothesis – the question is what we *mean* when we are being ironic (p. 93). It is not the case, as assumed by a majority of authors, that we simply mean the opposite of what we say; for why couldn't we then just say the opposite? Merely claiming that in irony, we dissociate ourselves from someone's utterance, or that we say something that we really disapprove of, in order to distance ourselves from what is alluded to, explains part of the process called irony; but these explanations do not go far enough.

What is really at stake, and where the salient literal meaning can help us find our feet, is the fact that 'irony functions as a reference point relative to which the ironized situation is to be assessed and criticized'. In other words, it is not just a matter of criticizing an utterance; rather, what we do is to keep in touch with the literal reality via salience in order to mark 'the realization of the extent to which the state of affairs in question has fallen short of expectations usually made explicit by what is said' (p. 93).

What we have here is not only a profound and correct account of irony: it is also one that can be applied to situations (cf. my remark earlier, where I criticized the author for not taking situational irony explicitly into account). The examples from real live texts that the author adduces in the sequel of the section (harking back to the initial ironic description of a 'Jerusalem Disneyland project', complete with a roller coaster ride on the heels of the High Priest into the Holy of Holies; p. 62), show that hypotheses such as the 'echoic mention' cannot do justice to the situational character of irony. In another

example, cruelty is displayed as an ironical criticism of the cruelty inherent in patriarchal society, for which the reference point is precisely the 'literal' meaning of the cruel words and deeds attributed to the protagonist of the novel *Dolly City* (Castel-Bloom 1992), a woman physician dissecting her son – for his own best!

The author concludes that the graded salience hypothesis has indeed gained support. The following chapters expand the basis of this support by looking into related phenomena of (non-literal) language use.

Chapter 5, 'Metaphors and idioms', and the short chapter 6, on 'Jokes', follow the method applied in the preceding chapter of first explaining the impact of the graded salience hypothesis on traditional views and findings, then testing the author's own predictions experimentally. In the following, I will only highlight a few illustrative points from these two chapters, as the theory expounded is essentially the same and the experiments are conducted along the same lines.

For metaphors and idioms, as for other instances of non-literal (or figurative) language, the question is again, how 'context and privileged meanings play [a role in] shaping our linguistic behavior' (p. 103). Rather than assuming different processes with different degrees and kinds of (non-)literality, we may assign the differences to the salient vs. non- (or less) salient character of the language used. As before, the comprehension process has two stages: an initial stage with two processes running in parallel, the one 'bottom up' (lexical), the other 'top down' (contextual), and a second, 'revisitation' stage, where retention and suppression play a role. Salient meanings are never automatically excluded (as other views have it), but retained, even when contextually incompatible, for possible further processing and possible suppression or 'fading'.⁹

Giora observes that the current debates on the understanding of metaphorical expressions have focused on literality, but that 'the variable of familiarity ... has been largely overlooked' (p. 106). Familiarity is prominent in interpreting metaphors, just as it was the case for irony: due to their similar salience, familiar metaphors have identical reading times for their metaphorical as for their literal

readings. Unfamiliar metaphors, in contrast, will take longer to read as metaphors than their literal equivalents do, since 'the metaphoric interpretation does not hinge on the salient interpretation of its components, and will require extra processing time ...' (p. 107). Similar effects were observed when testing for response times and reading times, using the 'moving windows' technique (pp. 112ff). As in the case of metaphors, familiarity also affected reading times for idioms: irrespective of context, familiar idioms were always read faster. As Giora notes, it is not always easy to tease out contextual effects from other factors, such as coherence; hence, she cautions that '[m]ore research is needed to weigh context strength against meaning salience when processing figurative language' (p. 116).¹⁰

A further factor is discussed in connection with metaphor processing: the 'aptness' of metaphors. 'Apt metaphors' are 'those rated high in "goodness", that is, in getting across the figurative meaning' (p. 118). There seems to be evidence that unfamiliar, but apt metaphors take no longer to process than familiar ones, against the expectations of the graded salience hypothesis. So the question is, does aptness override salience? One possible solution is that aptness may compensate for low salience; however,

this compensation does not allow apt but less familiar metaphors to be processed as fast as or faster than (both apt and less apt) familiar metaphors, suggesting that, though aptness is a factor in metaphor comprehension, it need not override salience. (pp. 119-120)

Another solution which comes to mind is that certain apt metaphors (e.g. 'sharp' for 'intelligent'), due to their 'success', obtain near-lexical status, analogous to certain fixed expressions (compare the roles of the English modal verbs *can* and *may* in indirect speech acts and negation). Hence the apt expression would lose some of its metaphorical status. Alternatively, 'aptness' could be a feature of salience itself.

What happens to metaphors after they have been initially processed? As in the case of irony, the 'integration phase' builds on both 'retention' and 'suppression' of compatible or incompatible meanings, depending on whether they are conducive to interpreting the metaphor or interfere with the interpretation process (p. 121).

Here, the context is decisive in determining which meanings should be retained, which suppressed. The predictions were tested by the usual experimental means: response times, reading times, and word fragment completion. In all these cases, it was found that 'salient, contextually incompatible meanings [were] retained insofar as they [were] functional, ... discarded, if they [were] not' (p. 127).

One of the outstanding features of the present chapter is the attention paid to spontaneous response to metaphoric utterances in naturally occurring discourse. As I mentioned earlier (cf. footnote 8), an often heard objection to experimental psycholinguistics is that the experimental laboratory conditions hardly are conducive to a natural use of language. Giora obviates this criticism by including a long stretch of conversation in her discussion of metaphor comprehension (817 lines of recorded talk, of which 74 are used to provide an example of the analysis conducted; pp. 120-131). The question: why are salient meanings retained, and how do people go about this retention? acquires thus a new, practical perspective. It turns out that first of all (as expected), metaphors are abundant in natural conversation; but moreover, interlocutors actually respond to metaphors in ways that support the graded salience hypothesis. In particular, salient literal (e.g. spatial) meanings of familiar metaphors (such as 'coming out of things') are retained and made available for further elaboration, in this particular case with an 'irrelevant' (humorous) purpose – something which could only be warranted by appealing to (or 'resonating with', as the author calls it, creating an 'apt' metaphor in the process) the literal, salient meaning of the original metaphor (p. 131).

Thus, the salient meanings are re-activated 'and reused for all kinds of purposes', whereby

[t]he salient literal meanings of metaphors, not least familiar metaphors, are recycled, as predicted by the graded salience and retention hypotheses. Similarly, familiar metaphoric meanings are also resonated with, as predicted. (p. 132)

It turns out that of the total of 120 metaphors used in this half hour long conversation, only 20 are less familiar ones, and of these only one third are not resonated with. The conclusion is that

REVIEW

as predicted by the graded salience and retention hypotheses, metaphors, not least familiar metaphors, are processed literally as well as metaphorically. Moreover, their salient, literal meaning is retained and may be resonated with for further processing. (p. 134)

Similar hypotheses were tested and confirmed for written discourse (pp. 134-136).

I found this section on natural occurring discourse highly illuminating and of great convincing power as to the credibility of the author's hypotheses. The graded salience hypothesis, and in particular its corollary, the retention hypothesis, provide us with a better understanding of how people spontaneously create coherent discourse, both orally and in written form. This section may well be the most original and appealing of the whole book.

As to idioms, the same or similar findings were reported as for metaphors, except that in the case of highly standardized idioms, the less salient, literal meaning was more difficult to elicit, and the idiomatic meaning, because of its higher salience, 'might be difficult to suppress even when inappropriate or unintended' (example: 'having cold feet' in the *literal* meaning, rather than in the idiomatic one of 'having second thoughts'; p. 138 and cf. p. 14). While it thus seems to be the case that 'salient idioms are processed idiomatically, regardless of context' (p. 140), one can still question the impossibility of retaining a literal element that, although salient (to a lesser degree) is usually not evoked. In fact, even the most 'die-hard' idioms could provide occasional breaches in their status of being 'highly entrenched', as Giora calls it (p. 136); think of a possible variant on 'to kick the bucket' such as 'he chose to kick a different bucket' in the sense of 'a different death', said in a supportive context (my observation, JM).

The author concludes that 'the relevant factor determining differences or similarity in early processes is not the literality of nonliterality of the utterances in question but rather their degree of familiarity' (p. 140; but recall what I said earlier about 'familiarity' as a stand-in for 'salience'). Similarly, late integration processes retain the original, salient (but contextually incompatible) meaning for later processing; this holds both for irony and metaphor. An apparent counter-case (due to Gibbs, 1998), showing different processes to be

involved in irony vs. metaphor, is explained by the fact that metaphoric uses (including, and specifically so, idioms) are conventional, whereas ironic uses are not; while metaphors have salient meanings built into them, as it were, ironic utterances depend on contextual information. Hence we are again looking at a difference in salience, rather than at one of processing (p. 141).

At the end of this chapter, Giora discusses the familiar adage 'the lexicon proposes and context disposes' (p. 147). While she agrees with the first part of the sentence, it has to be remembered that the lexical 'proposing' does not just happen across the board but is salience-ordered. As to the context, it does not always 'dispose' (in the sense of 'dispose of'): rather, it selects critically, by retaining e.g. literal meaning of metaphors when they are useful for further processing (as we saw in the case of the conversation analyzed above, with its 'humorous' resonance). Furthermore, 'some highly salient meanings resist suppression even when contextually irrelevant' (p. 147-148), as we have seen in the case of idioms. While this procedure may seem to be wasteful to some, in reality it serves a good purpose: the mind has its own economy: here, so-called parsimoniousness ('economy' in the usual, 'vulgar' sense; Mey 2001:179-181) is not always the positive flip side of being wasteful.

In the short chapter 6, on 'Jokes', the theory of graded salience is applied to a related domain, that of the joking utterance. The usual interpretation (and one that at first blush appears most plausible) is that 'jokes involve entertaining two incompatible interpretations' (p. 167). While this certainly is true, it does not explain the mechanisms underlying comprehension (or non-comprehension) of jokes. One possible interpretation of joke understanding is that of 'frame-shifting'; this assumption would at least in part explain why jokes take longer to process than do 'normal' utterances. But a more comprehensive explanation would have to take into account that it is not enough to observe a clash in meanings: we have to be more specific and ask what kind of meanings are retained, what kind are suppressed; and the answer is, not surprisingly: we have to look at their *salience* (p. 168).

In respect to salience, jokes do not differ from other non-literal language (such as metaphor) in the initial phase of processing; in fact, '[m]ost jokes make up a discourse that best exposes our tendency to opt for the salient interpretation first'. What happens at the punch line (the point called the 'disjuncter' by some) is that 'a sudden incongruity forces reinterpretation' (pp. 168-169), a 'revisitation', as it was called earlier. As to context, this may either promote or weaken the salient interpretation which is activated initially and subsequently revisited; such a revisitation should be more costly (because of the suppression of incompatible meanings involved; see above, and also p. 171) than merely registering a surprise ending following a 'frame shift'. The graded salience hypothesis is consistent also with the latter approach (see especially p. 174), but in addition, furnishes a better explanation of 'cost'. This 'suppression hypothesis' has been tested initially in a pilot study, but (as the author remarks) more research may be needed here (p. 170).

In general, in accordance with the graded salience hypothesis, 'the more salient the information[,] the more difficult it is to suppress' (p. 172); this finding is supported independently by other studies, such as on 'cloze probability' (p. 173).¹¹ Overall, the graded salience hypothesis seems to offer an account of jokes that is better suited to explain what happens at the 'disjuncter' point; differently from metaphor interpretation, where the integration phase does not involve suppression of earlier, 'salient, though contextually incompatible information' (p. 175), in jokes, we not only 'reshuffle salience', but actually suppress a non-suitable meaning. However, to quote Gernsbacher (1990), suppression is more costly, and the extra cost is reflected in extra processing difficulty for jokes, as opposed to metaphors and the like (p. 173).

The chapter ends with a joke, that has to be read twice to be understood (extra processing time required!). In general, this chapter reads very well, not least due to the many jokes used as example material (a pleasant light-heartedness is an overall positive feature of Giora's book, by the way).

The thoughts expressed in the likewise brief chapter 7, called 'Innovation' (pp. 176-184), rank among the most original and valuable presented in the book. Giora proposes an 'optimal innovation

hypothesis' (p. 176), dealing with the observation that innovative use of language, or creativity, is most pleasurable when it allows us to recognize the familiar (in accordance with Freud; p. 176). It turns out that salience and creativity are, in a way, complementary notions: just like salience, creativity is a graded concept. While pure novelty can be said to be the contrary opposite of salience, creativity not only requires a change in salience in order to be optimal, but in addition, this change has to be such that it 'allow[s] for the recovery of the salient meaning from which that meaning stems, in order that the similarity and difference between them may be assessable' (ibid.). Thus, while pure innovation may create novelty, it is not always pleasurable, and even may be less pleasurable because it does not allow us to retrieve the original meaning.¹²

Creative innovation is not restricted to the domain of tropes, even though most of the examples in the book refer to the latter: the hypothesis of optimal creativity holds for literal innovations, too. In addition, there is the added benefit of considering salience as a factor in creativity. Salient metaphors may be used innovatively: in fact, 'most novel metaphors are not pure innovations, but rooted in salient metaphoric concepts', as the author remarks (p. 179). But metaphor in itself is not sufficient to explain creativity: '[the mind] is constantly in search of novelty, regardless of whether it is figurative or literal' (ibid.). Indeed, the more salient meanings are often the ones that gets most easily 'de-automatized', re-novated, because they happen to be 'on our minds' (with an allusion to the metaphorical innovation of the book's title; p. 179). This is also the reason that the technique of 'bestrangement' (*ostranenie*, originally due to Shklovsky; see Mey 2000:254) is such a successful innovative device: it de-familiarizes the familiar, the salient.

The author then subjects her hypothesis to a number of experimental tests, all showing that just like suppression, optimal pleasurability comes with a cost: 'the optimally creative interpretation is more effortful, because it involves deriving the salient meaning and more' (p. 182). Needless to say, removing the element of familiarity would result in even higher processing costs, as the intended effects could only be achieved (or barely) at the price of great effort. To borrow the author's pithy formulation: 'It is not the most familiar,

then, that is least enjoyable, but rather the most novel that is least pleasing. Pleasure, however, resides half way between high salience and high novelty' (p. 182).

In addition, there is a political aspect to innovation: it may be the expression of a 'bestrangement' in yet another sense, namely, the dissatisfaction with the current state of affairs in society. In this way, innovation has always been the realm of the revolutionary (as it is beautifully expressed by the Brazilian poet Vinícius de Moraes in his 'Triptych to Sergey Eisenstein'; see Mey 2000:380-382). The feminists have been particularly active in creating novel, 'politically correct' terminology, with varying success and pleasurability. The danger of a dialectic development lurks right around the corner, though: innovative expressions (such as 'his/her') may themselves become familiar, and open to innovative attacks by the 'opposition', and thus in the end do more harm than good, as I have shown elsewhere.¹³

The moral of the chapter cannot be better captured than in the the author's own words:

The morale [*sic*] of this chapter is highly predictable. Highly novel language and thought will be less attractive and catchy, and easier to shirk off. ... Effective novelty (attractive, affecting change), on the other hand, is such that induces change but is rooted in salience to the extent that it allows for the recoverability of the familiar. (p. 184)

The two final chapters are entitled 'An overview' (chapter 8, pp. 185-195) and 'Coda: Unaddressed questions. Food for future thought' (chapter 9, pp. 196-199). The former of the two merely restates the findings obtained in the preceding chapters, and reconfirms the author's conclusions regarding the graded salience hypothesis: anything the others can do, it can do too, and better. The chapter recapitulates the author's thoughts and experimental findings; as such, it is useful but somewhat repetitive, with the exception of section 3, 'On the superiority of the salient meaning: The case of literal language', which discusses important work by Richard Gerrig (1989) on *degrees of salience*. It would perhaps have been more appropriate to include this section in

the body of the book, as it fits nicely in after the chapter on 'Jokes' (chapter 6). Also Gerrig, when dealing with literal language, contrasts other models of sense creation (or 'error recovery'; p. 192) with a parallel-process model. It turns out that '[i]n spite of some different assumptions, both processing models [Gerrig's and the graded salience model, JM] assume that the salient (conventional) meaning should be activated on encounter, regardless of context'. Thus, Gerrig's studies 'support the view that salience plays a major role in language comprehension [as such; hence not only of metaphorical or other figurative language, JM]' (p. 194) – a finding that well could have been placed earlier in the book. (Gerrig *is* mentioned briefly on p. 29, under the heading 'embodied meaning').

The final chapter 9 contains, as the title indicates, 'food for further thought' (p. 196). Here, Giora raises several interesting questions, such as whether literality is linguistic or conceptual, and whether metaphor is necessarily linguistic. While the latter question seems resolved in the negative sense, the former is still being debated: is literality based on the 'resemblance obtaining between a propositional form of an utterance and the *thought* it represents?' (p. 197). Similarly, the jury is still out on questions having to do with nonverbal irony and its relation to literal or figurative language (p. 198).

Giora further remarks that there seems to be no need to assume different strategies for the understanding of the different kinds of (non-)figurative language: for non-salient language, the process involves both an 'inferential' and a 'predictive' component, whereas the understanding of salient language always involves a lexical look-up, followed by a contextual readjustment (p. 198).

An interesting development is signaled on pp. 198-199, where it is mentioned that learners of a second language may have trouble processing the salient meanings of what is called 'situation-bound utterances'; instead, they fall back on the expressions' literal meanings, thus missing out on the pragmatic functions of the expressions in question (Kecskés and Papp 2000). To those speakers, the literal meanings appear to have been more salient.

REVIEW

A further interesting suggestion is that salient meanings, because of their availability ('they spring to mind'), may be difficult to ignore. Giora suggests that this is the reason why we are so comfortable in our prejudices, and insists, with Reinhart (2000), that 'deautomatizing salient meanings, concepts and ideas is ... one of the most important roles of art and science' (p. 199). On the other hand, contextual information is not automatic, either, and cannot 'dominate our thinking entirely'. Between being trapped in salience and being totally contextualized, we must make a reasoned choice, allowing us to 'keep an open mind in face of biasing or manipulative information' (*ibid.*).

Before ending this review, I want to point to a few less fortunate aspects of Giora's work (in addition to the critique I have leveled above on some particular points). First off, there is the problem of 'overkill' (mentioned in connection with chapter 3). While I find the discussion there overall relevant and enlightening, not being an experimental psycholinguist, I must withhold an evaluation of the experiments and their design (for a detailed account of these matters, the reader is directed to the literature referenced in the chapter). But even for one who generally is sympathetic to the author's claims, it is not necessarily the case that more of the same also is more convincing (compare Giora et al.'s (2001) rebuttal of the experiments by Vu et al., 1998 and 2000, discussed on pp. 47-49).

Also, as a point of method, it is not a good idea to introduce concepts in a certain order, and then to discuss them in another order, to finally explain them in detail in yet another order. This is what happens on p. 40, paragraphs 3 and 4, where Giora talks about the five different models for lexical access, and in the subsequent sections (pp. 41-54), where she discusses the individual models in more detail.

A minor negative point may be made in connection with the somewhat incomplete 'legend' pertaining to the long conversational stretch quoted on pp. 148-166 (cf. also the extracts on pp. 128-130 and 131). Many of the abbreviations are not, or only scantily, explained, and why are some of them capitalized (to name just one example: what does 'VOX' stand for?). In particular, the way numerals and other

(presumably) intonation signaling devices are used should have been spelled out in some detail.

Occasional quaint uses of English should have been weeded out (e.g. 'to avail' with a direct object on p. 40 and elsewhere). There are a few typos (e.g. author Hogabaom (correctly: Hogaboam) is misspelled on p. 43 and the entry is consequently missing in the Index.

A more serious, and also persistent problem concerns what I would call the 'context-lessness' of much psycholinguistic research. The experiments are strictly designed to show what they should show, and this is how it should be; but sometimes one wonders if a less discipline-oriented or -limited approach would not have resulted in alternative, more plausible explanations. Take for example the case of the married couple conversing on p. 40, where the husband interprets the word 'jeans' as 'genes', even though the context for the latter is wholly inappropriate. It belongs to the story (as Giora points out) that the husband is interested in genetics, and (presumably) not in repairing his daughter's jeans. To me, this is a case of what a cognitive linguist would name the 'grounding' or 'embodiment' of a concept. Thus, for a geneticist, the most salient meaning of [ji:nz] is not a piece of apparel, but something to do with genetics; this belongs to his or her mental make-up. I don't think that the reasoning behind Giora's rejection of grounding in connection with salience is altogether compelling. Salience *could* be based on grounding, and salient meanings *could* be embodied without the graded salience hypothesis losing any of its ascendancy and scope; in addition, our view of the process of lexical access would perhaps become more balanced.

Concluding this review, it needs to be said that this is a most important contribution to the study of the relationships between mind and language, meaning and context, human and environment. Giora's theory of 'graded salience' provides us with a new angle on the old question of how language is processed. Her approach is not one of mere theorizing, but of experimentally testing her hypotheses, mostly with positive results (when the results are not entirely satisfactory, we are told this right away). An additional advantage of the book is that there is room for human aspects, including the use of poetry and

REVIEW

humor (motivated by the subject matter, but at the same time making for refreshing breaks from the purely psycholinguistic and linguistic reasoning).

Giora endeavors (in my opinion, successfully) to modify the way we think, and do experiments, about language comprehension (and to a lesser degree, also production). She has single-handedly lifted a current paradigm out of its hinges and placed it on a new footing, by drawing attention to the overriding role played by salience in language comprehension, in particular with regard to figurative (non-literal) language. Despite the many technical explanations, the book is easy to read and the experiments are (even for one who is not an experimental psychologist) easy to follow. The well-argued plausibility of Giora's model is at least an eye-opener, if not a refreshing alternative for people like the present reviewer, who subscribe to a different paradigm.

*1100 West 29th Street
Austin, TX 78703-1915
U.S.A.*

Notes

* An abridged version of this article has appeared in *Lingua* (2004).

1. I must confess that I have a little difficulty with restricting myself to 'speech' here, as Giora does; actually, her experiments, to be discussed below, are mostly based on *written* tests, not on actually occurring speech – with the exception of the long oral discourse analyzed in chapter 5 (see below). But of course this does not invalidate the results.
2. This is supposed to include the extra-linguistic context: depending on developments in society, once salient meanings may lose their salience, as in the 'bug' example quoted above; another example is that of 'mail', now in many countries, e.g. Italy and apparently also Israel, synonymous with the salient meaning of 'electronic mail'.

Giora's statement on p. 18 should be qualified in this sense; cf. the author's observation in footnote 10, p. 202: 'In my previous writings (Giora, 1997b) I

REVIEW

said that context can contribute to the salience of meanings. I should have said that context may contribute to the availability of accessibility of a concept'.

3. Note that the author's example on p. 14 – Brigitte Nerlich's 'cold feet' – could perhaps be explained by interlingual interference: the speaker is German, but addresses an Englishman, who may well have been tempted to exploit this difference to construct a joking reply; such a possible humorous effect is intimated by the author herself on p. 27.

Note also that Giora, following her discussion of salience in relation to other, related concepts, makes the point that 'the mental lexicon is amenable to learning and change', hence that 'the salience-ordered, hierarchical structure of the lexicon may be unstable and in a constant state of flux' (pp. 34-35). This is, one could say, the diachronic equivalent of the 'synchronic' coding of the lexicon that I had some difficulties with above.

4. With regard to this discussion, it should be remarked that Giora distinguishes two kinds of accessibility: that of sense and that of reference. Salience has to do with the former, not the latter (p. 31; cf. also chapter 2.4).
5. Notice Giora's wording here (and compare my earlier discussion of 'salience' vs. 'frequency'): '... the more salient (frequent, 'dominant') meaning is accessed faster ...' (p. 40).

Here, the author seems to assume (a certain degree of) synonymy between these concepts.

6. Usually, the expression 'garden path' is used in a somewhat different fashion, denoting the willful weighting of one interpretation in favor of another in cases of ambiguity. This 'leading' effect can be due to either syntactic or semantic features, or both. (For a discussion and an example, see Mey 1992).
7. Fodor's (1983) expression, quoted p. 60.
8. Another criticism that could be vented here, viz., that the experiments were conducted in a laboratory context rather than in the context of naturally occurring conversations, is partly obviated by the evidence from text production that the author adduces in section 2.2.3 of this chapter. Her findings confirm the hypothesis formulated earlier: that 'literal meaning was not suppressed as irrelevant, but was retained for further processes' (p. 93). (See further the data provided on pp. 131ff. of the book and discussed later in this review).
9. This term appears on p. 106, but has not previously been used or defined, as far as I can see; presumably it denotes suppression viewed as a kind of 'graceful degradation'.

REVIEW

10. I have a terminological quibble here: the author seems to conflate the two concepts: familiarity and salience, such that they could seem to be used interchangeably. Whereas familiarity earlier was defined as a feature of *salience* (familiar meanings of familiar stimuli/expressions; cf. p. 16), here, it concerns familiar *stimuli*.
11. On p. 174, Giora remarks that cloze and salience are not just inversely proportionate, even though her formulation on the preceding page, '... N400 brain wave amplitude is largest for items with low cloze probability', could lead one to believe just that. In addition, the earlier made reservation as to cloze procedures, viz., that they rest upon word rather than on conceptual 'fit', could have been repeated here as well.
12. As to the example that Giora provides here, *spandy wear* as female apparel, it should be noted that this particular item is easily recoverable, at least for persons over a certain age, who remember the female fashion in foundation garments called *spandex*. Apparently, the author does not fall into this category.
13. When the Norwegian Parliament (the *Storting*) decided to stop referring to its members as *stortingsmenn*, by dropping the offensive and sexist plural *menn* ('men') and replacing it by the neutral *representanter* ('representatives'), the newspapers immediately caught on and started talking about 'representatives and their wives!' (See Mey 1985)

References

- Castel-Bloom, Orly. 1992. Dolly City. Tel Aviv: Zmora Bitan. (In Hebrew)
- Clark, Herbert H. 1996. Using language. Cambridge: Cambridge University Press.
- Clark, Herbert H. & Richard J. Gerrig 1984. 'On the pretense theory of irony'. *Journal of Experimental Psychology: General*: 113.121-126.
- Fodor, Jerry A. 1983. The modularity of mind. Cambridge, Mass.: MIT Press.
- Gernsbacher, Morton Ann. 1990. Language comprehension as structure building. Hillsdale, N.J.: Erlbaum.
- Gerrig, Richard J. 1984. 'The time course of sense creation'. *Memory & Cognition* 17.194-207.
- Gibbs Jr., Raymond W. 1986. 'On the psycholinguistics of sarcasm'. *Journal of Experimental Psychology: General* 115.3-15.
- Gibbs Jr., Raymond W. 1998. 'The fight over metaphor in thought and language'. In: A. Katz, C. Cacciari, C. Gibbs & M. Turner (eds.), *Figurative language and thought*. New York: Oxford University Press, pp. 88-118.

REVIEW

- Giora, Rachel, Orna Peleg & Ofer Fein. 2001. 'Resisting contextual information: You can't put a salient meaning down'. International Workshop 'Towards an Experimental Pragmatics', Lyon, France, May 2001.
- Kecskés, István & Tünde Papp. 2000. Foreign language and mother tongue. Mahwah, N.J.: Erlbaum.
- Littman, David & Jacob L. Mey. 1991. 'The nature of irony: Towards a computational model of irony'. *Journal of Pragmatics* 15.131-151.
- Mey, Jacob L. 1985. *Whose language? A study in linguistic pragmatics*. Amsterdam & Philadelphia: John Benjamins.
- Mey, Jacob L. 1992. 'Pragmatic gardens and their magic'. *Poetics* 20.233-245.
- Mey, Jacob L. 2000. *When voices clash: A study in literary pragmatics*. Berlin & New York: Mouton de Gruyter.
- Mey, Jacob L. 2001. *Pragmatics: An introduction*. Oxford & Boston: Blackwell. (Second, revised edition).
- Reinhart, Tanya. 2000. *Self-inspection in the art of the 20th century*. Tel Aviv: Ha-kibbutz ha-meuchad. (In Hebrew).
- Shelley, Cameron. 2001. 'The bicoherence theory of situational irony'. *Cognitive Science* 25.775-818.
- Vu, H., G. Kellas & S. T. Paul. 1998. 'Sources of sentence constraint in lexical ambiguity resolution'. *Memory & Cognition* 26.979-1001.
- Vu, H., G. Kellas, K. Metcalf & R. Herman. 2000. 'The influence of global discourse on lexical ambiguity resolution'. *Memory & Cognition* 28.236-252.
- Wilson, Deirdre & Dan Sperber. 1992. 'On verbal irony'. *Lingua* 87.53-76.
- Wilson, Deirdre & Dan Sperber. 1993. 'Linguistic form and relevance'. *Lingua* 90.1-25.