Chapter 8

The psychology of utterance processing: context vs salience

Rachel Giora

1. Introduction

Consider the cartoon in Figure 8.1 by Adams (2009), which is an Optimal Innovation: it includes a novel stimulus intended to further activate coded, salient meanings, so that both the novel and the salient, though different, may interact and affect highest pleasurability (Giora, 2003; Giora Fein, Kronrod, Elnatan, Shuval, and Zur, 2004).

Figure 8.1

Put differently, this cartoon will be optimal innovative only to those familiar with the coded meanings it echoes (see e.g., the poster in Figure 8.2). The initiated audiences are thus invited to invoke a wide range of coded, salient meanings which might allow an insight into the ironic message of the cartoon. However, those not in the know will only have access to what is literally spelled out, superimposed on the pictorial background (in addition, of course, to the figurative title and the symbol of the crown). Indeed, if taken at face value, the cartoon will encourage people to get terrified and stop doing what they are doing on account of the nasty weather (referred to in the background). But the poster in Figure 8.2, if retrievable from memory, will bring to mind an altogether different

meaning which this cartoon must be weighed against: the British mindset on the eve of World War II – the "British restraint and stiff upper lip" (Hughes, 2009).

Figure 8.2

Additionally, one of Lance Corporal Jones' catchphrases in *Dad's Army* – the British sitcom about the Home Guard in the Second World War (Perry, 1968-1977) – might also spring to mind, echoing, via another irony, the mindset derided here:

(1) "Don't panic!"[1]

The title of the cartoon – *Spirit of the blitz 2009* – allows another ironic turn of the screw, reminding us of the spirit of blitz exercised by people of Britain, who, during the 1940 bombing of London, exhibited stoic courage and endurance. [2]

All these meanings, if coded and available, will be invoked automatically as a direct response to the stimulus in Figure 8.1, despite their apparent irrelevance to the immediate context at hand – the ferocious weather in Britain during December 2009. When weighed against contextual information – the blitz-like storm which forced many in Britain to assume the spirit of the blitz rather than the spirit of Christmas – these meanings strongly deride the panicky spirit of the Brits on Christmas Eve 2009.

The way this cartoon can be interpreted illustrates the need to take various factors into consideration when examining the end-product of utterance interpretation. Indeed, the psychology of utterance processing takes into account a number of factors that shape

---

[1] http://www.youtube.com/watch?v=ZR6wok7g7do&feature=related.

[2] http://www.roarsquadron.com/Web/FlashBacks/Ed%2011%20Page%201%20Blitz_1.htm

utterance interpretation, such as (i) salient/coded meanings, (ii) contextual information, and (iii) their unfolding interaction (or lack of it).

Debates within psycholinguistics can, thus, be viewed as divisible into two main approaches (for a review, see Giora, 2003 chapters 1–3). At one end of the spectrum are context-based models which assume that a strong context reigns supreme in that it governs early processes and facilitates contextually compatible meanings only. Consequently, the output of the interpretation processes must be seamless, involving no contextually incompatible meanings and interpretations (*the connectionist model*, e.g., Bates, 1999; Bates and MacWhinney, 1989; MacWhinney, 1987; Small, Cottrell and Tanenhaus, 1988; *the constraint-based model*, e.g. McRae, Spivey-Knowlton and Tanenhaus, 1998; Pexman, Ferretti and Katz, 2000; *the direct access view*, e.g., Gibbs 1979, 1994; Keysar, 1994; Ortony, Schallert, Reynolds and Antos, 1978).

At the other end are lexicon-based models which hold that coded meanings of stimuli are speedy responses, activated automatically, regardless of contextual information. As a result, initially accessed meanings may fail to meet context fit. Consequently, they may induce incompatible interpretations which will either feature in final outputs alongside the appropriate ones or will be subjected to revisitation or suppression processes (*the modular view*, e.g., Fodor, 1983; Swinney, 1979; *the standard pragmatic model*, e.g., Grice, 1975; Searle, 1979; *the graded salience hypothesis*, e.g., Giora, 1997, 2003; Peleg, Giora, and Fein, 2004, 2008). The context-based and the lexicon-based models, then, have different predictions, especially with regard to initial processes, which, in turn, affect later interpretation processes.

Specifically, according to the context-based view, specific and supportive contextual information penetrates lexical access and selects the contextually appropriate meaning only, the consequence of which is contextually compatible interpretations only (Section 1.1). In contrast, the lexicon-based view predicts that coded meanings cannot be blocked and, therefore, at times, end-product interpretations will also involve contextually incompatible meanings and interpretations (Section 1.2). To tease apart these two approaches, we need to look at research into utterance interpretation processes.

1.1. Context-based approaches

Context-based approaches focus on the *facilitative* effects of strong contexts which allow them to select only compatible meanings and interpretations. Thus, according to *the connectionist model*, when words (*bulb*), ambiguous between salient/dominant ("light") and less salient/subordinate ("flower") meanings, are preceded by a context strongly biased toward one of the meanings, their processing will result in selecting the contextually appropriate meaning exclusively, regardless of degree of salience. For instance, processing *The gardener dug a hole. She inserted the* ***bulb*** resulted in an exclusive activation of the less-salient, subordinate ("flower") meaning of *bulb* when probed immediately (Vu, Kellas and Paul 1998; Vu *et al.* 2000).

However, as shown by Peleg and Giora (2010) and Peleg, Giora and Fein (2001, 2004, 2008), this finding need not attest to selective access; it might just as well be the effect of a predictive context, which guesses the intended meaning without interacting with lexical processes. For instance, Peleg *et al.* (2001) show that guessing the contextually

appropriate meaning in a context biased toward the less-salient meaning occurred even before the processor encountered the relevant stimulus (*bulb*). Additionally, as shown by Peleg and Eviatar (2008, 2009), briefly following the encounter of the stimulus in question at 250 ms Stimulus Onset Asynchrony (SOA), the salient incompatible meaning as well as the less salient compatible meaning were both activated. This was further qualified by the type of homograph. In the case of homophonic homographs (*bulb*), both meanings were activated at a short (150 ms SOA) delay and remained active even 100 ms afterwards; in the case of heterophonic homographs (*tear*), the contextually appropriate, less salient meaning was activated exclusively in the left hemisphere at 150 ms SOA, but 100 ms later (at 250 ms SOA), the salient but contextually incompatible meaning also became available (Peleg and Eviatar 2008, 2009). When probed later, 1000 ms after encountering the homophonic homograph, the left hemisphere selected the contextually appropriate (less-salient) meaning, whereas both the salient and less-salient meanings were still activated in the right hemisphere. At this long delay, 1000 ms following the onset of the ambiguous word, the left hemisphere was unable to suppress the salient contextually inappropriate meaning of heterophonic homographs, while the right hemisphere was (see Peleg *et al.*, 2008). Such findings demonstrate that salient meanings cannot be blocked, not even by a strong context.

According to *the constraint-based model*, both contextual as well as lexical "constraints" may affect end-product interpretations, depending on their quantitative strength. The greater the number of the constraints favoring a specific meaning/interpretation the greater chance it stands to be selected exclusively (McRae, Spivey-Knowlton and Tanenhaus, 1998). For instance, if contextual information is biased toward an ironic interpretation of a target, and, in addition, involves other biasing factors (such as a speaker whose profession indicates s/he could be ironic), such a strong context should facilitate the appropriate (ironic) interpretation only, even though its literal meaning may be more salient.

Findings, however, show that such contexts did not facilitate non-coded, inferred or novel interpretations such irony, but instead slowed down ironic targets compared to more salient literal counterparts (Pexman *et al.*, 2000; but see Ivanko and Pexman, 2003 for similar but also for somewhat different results, argued against in Giora, Fein, Laadan, Wolfson, Zeituny, Kidron, Kaufman, and Shaham, 2007).

The direct access view (Gibbs, 1979, 1986, 1994) argues against the temporal priority of utterance literal interpretation (posited by Grice, 1975), contending instead that, in a strongly supportive context, interpretations of literal and nonliteral utterances should exhibit similar interpretive processes. Indeed, when Ortony *et al.* (1978) embedded statements, ambiguous between literal and novel (metaphoric) interpretations, in poor contexts, literal utterances were faster to read; however, when provided with rich contextual support, both literal and metaphoric statements took similarly long to read, thus testifying to context's facilitative effects on novel metaphors.

Later studies, however, failed to replicate these results. Rather, novel metaphoric items, embedded in supportive contexts, always took longer to read compared to their literal interpretation (Brisard, Frisson and Sandra, 2001**;** Giora and Fein, 1999a; Pexman *et al.*, 2000; Tartter, Gomes, Dubrovsky, Molholm and Stewart, 2002; see also Giora,

1997, 1999). Similarly, non-coded ironic utterances were always processed literally first despite a strongly supportive context (Giora *et al.*, 2007; Giora *et al.*, Giora, Fein, Kaufman, Eisenberg, and Erez, 2009).

1.2. Lexicon-based approaches

Lexicon-based approaches focus on the insensitivity of lexical processes to contextual information. According to *the modular view* (Fodor, 1983), cognitive processes are either domain-specific or domain-general. Domain-specific processes (such as lexical access) are modular: they are low level bottom-up processes, which are sensitive only to relevant stimuli (e.g., lexical items). Among other things, modular processes are informationally encapsulated, that is, impenetrable to processes occurring outside the input system. In contrast, domain-general, central systems, such as contextual information, consist in top-down, integrative, and predictive processes which are receptive to outputs of various domains.

Modular processes such as lexical access, then, are not affected by top-down feedback from higher-level representations such as contextual information or world knowledge. Rather, lexical access is autonomous and, on some views, exhaustive: *all* the meanings of a lexical stimulus are activated once this stimulus is encountered, regardless of either contextual bias or degree of salience. However, once these meanings are activated, contextual, central system processes may influence them. For instance, the central system may either integrate them with contextual information or discard them from the mental representation as contextually incompatible (for a review of other versions of modular and also hybrid models, see Giora 2003 chapters 1–3).

The standard pragmatic model (Grice, 1975; Searle 1979) may be viewed as a version of a modular view, attributing properties such as imperviousness to contextual information and, consequently, temporal priority to *literal* meanings. According to *the standard pragmatic model*, the meanings of a linguistic stimulus to be activated first are literal. On the basis of these literal meanings, utterance literal interpretations are to be constructed first. However, if literally-based (meanings and) interpretations do not meet contextual fit, suppression of these representations will take place, to be followed by their replacement with contextually appropriate alternatives.

Following the modular view (Fodor, 1983), *the graded salience hypothesis* (Giora 1997, 1999, 2003; Peleg and Giora, 2010; Peleg, Giora and Fein, 2001, 2004, 2008) assumes two kinds of mechanisms that run parallel: a bottom-up modular system, which is encapsulated and autonomous in that it is impervious to context effects (e.g., lexical access), and a top-down central system (e.g., contextual information), which is integrative but can also be strong enough so as to predict a compatible meaning or interpretation. Diverging from *the modular view*, however, *the graded salience hypothesis* posits that lexical access is ordered: salient meanings are activated faster than less-salient ones. In addition, suppression of contextually incompatible meanings is not unconditional but rather functional; it is sensitive to discourse goals and requirements, allowing for contextually incompatible meanings and interpretations to be retained if invited or if supportive or non-intrusive of the intended interpretation (see also Giora and Fein, 1999a; Giora, Fein, Aschkenazi, and Alkabets-Zlozover, 2007).

According to *the graded salience hypothesis*, salience is a matter of degree: a meaning is *salient* if it is coded in the mental lexicon and enjoys prominence due to cognitive priority (e.g., prototypicality, stereotypicality) or amount of exposure (e.g., experiential familiarity, frequency, or conventionality), regardless of degree of literality; a meaning is *less-salient* if it is coded but low on these variables, regardless of degree of literality; a meaning is *nonsalient* if it is non-coded – either novel or derivable (e.g., on the basis of contextual information), regardless of degree of literality.

Although salience is a property of words and fixed expressions rather than a property of utterances' compositional meaning and interpretation, utterance interpretation may often rely on the salient meanings of its components. Interpretations that are based on the salient meanings of the utterance components are salience-based interpretations and could be both literal and nonliteral (see Giora *et al.*, 2007).

Rich and specific contextual information can be predictive of an oncoming message, as well as supportive and facilitative. However, even when it is rich enough to activate meanings and interpretations on its own accord, it does not penetrate lexical processes but runs parallel. As a result, often salient but inappropriate meanings and, consequently, salience-based but inappropriate interpretations might be involved in utterance interpretation. Such inappropriate interpretations need not be suppressed; they may be retained, provided they do not interfere with the final contextually compatible interpretation. The result is the involvement of such interpretations in the final outputs of utterance interpretation (e.g., the salience-based, often literal interpretation of ironies and metaphors, see Brisard, Frisson and Sandra, 2001**;** Giora and Fein, 1999a; Pexman *et al.*, 2000; Tartter, Gomes, Dubrovsky, Molholm and Stewart, 2002; see also Giora, 1997, 1999 and Examples 1-3 above).

2. Salient meanings and salience-based interpretations are not necessarily literal

According to *the graded salience hypothesis* (Giora, 2003, 2006; Giora and Fein, 1999a), neither salient meanings nor salience-based interpretations need to be literal. Similarly, nonsalient interpretation need not be figurative.

2.1. Salient meanings are not necessarily literal

A number of studies demonstrate that salient meanings need not be literal. For instance, in Gibbs (1980), familiar (English) idioms (*spill the beans*), whose salient meaning is figurative, took less time to read in an idiomatically than in a literally biasing context, the latter inviting a salience-based literal interpretation. In Giora and Fein (1999b), familiar (Hebrew) ironies, whose salient meanings were both ironic and literal, were processed initially (at 150 ms ISI) both ironically and literally, regardless of context bias; in contrast, less familiar ironies, whose salient components were literal, were processed initially only literally, regardless of context bias.

In Colston and Gibbs (2002), familiar (English) metaphors (*This one's really* ***sharp*** said of a student) were faster to process when embedded in metaphorically than in ironically biasing contexts. Indeed, one of the salient meanings of their keyword (e.g., *sharp*) is metaphoric, which accounts for their metaphoric salience-based interpretation.

When embedded in an irony biasing context, the salience-based interpretation of the ironic use should be metaphoric and will have to be adjusted to the ironically biased contextual information.

Similarly, in Giora, Gazal, Goldstein, Fein, and Stringaris (2010), (healthy) participants were faster to respond to familiar (Hebrew) metaphors (*flower bed*), whose salient meaning is metaphoric, than to novel ones whose nonsalient interpretation is metaphoric (*golden laugh*), even though the individual words that made up the target word-pairs were similarly highly familiar.

2.2. Salience-based interpretations are not necessarily literal

Given that salient meanings need not be literal (Section 2.1), it follows that salience-based interpretations need not be sensitive to degree of literality either. After all they are derived on the basis of the salient meanings of the utterance's components, which may be either literal or nonliteral. Recall that the ironic *This one's really sharp*, whose salience-based interpretation is metaphoric (relying on the salient, metaphoric meaning of *sharp*) took longer to process in an ironically than in a metaphorically biased context (see Colston and Gibbs, 2002), since the salience-based (metaphoric) interpretation of the ironic use is contextually incompatible.

2.3. Non-salient interpretations are not necessarily nonliteral

Consider the example in Figure 8.1, which is literal: its final interpretation includes a number of salient meanings to which it adds a novel twist. Consider further *Know Hope* (in Figure 8.3)[3] which is a literal, nonsalient use, harping on the salient, literal "no hope".

Figure 8.3

[3] http://www.flickr.com:80/photos/idanska/247228762/

In Giora *et al.* (2004) we studied such (Hebrew) innovations, termed optimal innovations. To be optimally innovative, a stimulus should evoke a novel – less or nonsalient – response to a familiar stimulus, alongside a salient one from which it differs (both quantitatively and qualitatively), so that both can interact, regardless of nonliterality. Admittedly many optimal innovations may be nonliteral (novel metaphors or unfamiliar ironies). However, quite a few are literal (as can be deduced from Figures 8.1 and 8.3). Consider the following examples: *Body and sole* – the name of a shoe shop – which evokes the salient *body and soul*, both of which are literal but only the first is nonsalient; *Curl up and dye* (the name of a hair salon) which is nonsalient and literal and which evokes the salient but metaphoric *Curl up and die* (see Giora *et al.*, 2004).

Given that optimal innovation involves processing a salient meaning on top of the novel interpretation, it is no wonder that optimal innovation takes longer to process than their salient meanings, regardless of metaphoricity (as shown by Giora *et al.*, 2004). Indeed, in Giora *et al.* (2010 Experiment 2), both healthy individuals and individuals diagnosed with Asperger syndrome took longer to process and to frequently err on (Hebrew) nonsalient interpretations compared to salient ones, regardless of metaphoricity. Thus, literal optimal innovations such as *a Tverian horse* (meaning "a horse from Tiberias"), reminiscent of a salient, metaphoric collocation – *a Trojan horse*, were slower to induce correct meaningfulness judgments compared to familiar literal collocations, whose meanings are salient. Nonsalient interpretations then are not necessarily nonliteral.

3. Opting for the literal interpretation is not necessarily a default strategy

To further test the claim that opting for a nonliteral rather than a literal interpretation may be a default strategy, independent of explicit contextual information (including information about the speaker and the addressee), one needs to neutralize factors affecting nonliterality such as degree of salience (recall that familiar items may have a lexicalized nonliteral meaning, see Section 2.1), semantic anomaly (known to trigger metaphoricity, see e.g., Beardsley, 1958), and contextual information (since breach of pragmatic maxims or contextual misfit may invite a nonliteral interpretation, see e.g., Grice, 1975). It thus follows that for this claim to be experimentally substantiated, testing it should involve novel items susceptible to a literal interpretation and presented outside a specific context.

Indeed, in Giora, Fein, Metuki and Stern (2010), materials were affirmative statements of the form X is Y (such as *This is Memorial Day; I am your doctor*) and their negative versions (*This is not Memorial Day; I am not your doctor*), all of which could, potentially, be assigned a literal interpretation in that they were all equally novel, semantically intact, and presented in isolation (see Figures 8.4 and 8.5).

Participants were asked to rate the interpretation of these targets on a 7 point scale ranging between two specific (either literal or nonliteral) interpretations presented randomly at the scale's ends.

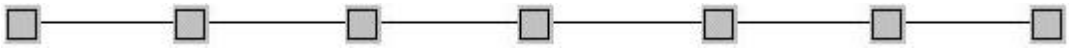

Everybody is sad today

We are celebrating Memorial Day today

Figure 8.4. This is Memorial Day

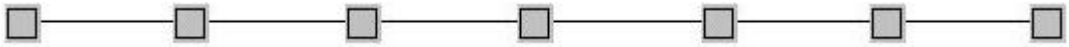

No need to be so sad

We are not celebrating Memorial Day today

Figure 8.5. This is not Memorial Day

Results show that the negative statements were rated as significantly more metaphoric than their affirmative counterparts, supporting the claim that opting for the literal interpretation need not be a default strategy.

Indeed, when the negative statements were embedded in contexts equally supportive of either their metaphorical or their literal interpretation, reading times were faster for items embedded in metaphorically than in literally biasing contexts (Giora, Fein and Livnat, in prep.), further supporting the claim that a nonliteral interpretation may be a default interpretation.

Additional evidence supportive of this claim comes from research into irony interpretation. Giora, Fein, Ganzi, Alkeslassy Levi, and Sabah (2005) show that negating affirmative overstatements results in assigning these statements an ironic interpretation, even though these statements are amenable to literal interpretation as well. Specifically, findings demonstrate that negative overstatements (*He is not exceptionally bright*) come across as ironic even outside a specific context. When presented in isolation, they are rated as more ironic than other alternatives such as affirmative overstatements (*He is exceptionally bright*) or negated non-overstatements (*He is not bright*).

Taken together, findings from various negative statements indicate that it is not the case that, when available, literal interpretation is a default. Instead, our findings demonstrate that even nonconventional utterances, susceptible to literal interpretation, are often perceived as nonliteral, even when no supportive information of that interpretation is made manifest.

5. Context effects – later interpretation processes

Although the various approaches outlined above differ in their predictions with regard to the effects of a "strong context" on early lexical processes (Section 1), there seems to be an agreement that contextual information should affect later interpretation processes. But what these effects should look like is still a matter of debate. To test the various

predictions of the approaches with regard to later processes, I will focus here on late effects of a "strong context" (such that anticipates an ironic utterance) on irony interpretation and on the effects of coherence on later processes of negated information.

5.1. Irony interpretation

According to context--based approaches, if context is strongly biased in favor of the appropriate (ironic) interpretation, only that interpretation should be activated immediately and feature exclusively in the final product. According to some lexicon-based approaches, inappropriate meanings as well as inappropriate (literal) interpretations (of irony) should be activated immediately even in the presence of a strong context. However, later, they should be discarded from the mental representation so that the final product features only contextually appropriate (ironic) interpretations (Grice, 1975; Fodor, 1983).

Whereas these two approaches have similar predictions with regard to how contextual information should affect the final (ironic) representation, *the graded salience hypothesis* has different predictions. According to *the graded salience hypothesis*, salience-based yet inappropriate interpretations (the salience-based interpretation of novel metaphors and unfamiliar ironies), which are activated immediately, need not be discarded. They may be retained since they contribute to (or at least do not disrupt) the final interpretation processes (Giora, 2003).

But what makes up a strong context? According to Gibbs (1986, 2002), a context may be strong enough to facilitate the ironic interpretation of an utterance exclusively if it sets up an "ironic situation" through contrast between what is expected and the reality that frustrates it. Inducing an expectation for an ironic utterance should allow ironic interpretation to be tapped directly, with no recourse to contextually inappropriate utterance level interpretations. According to Gibbs, then, a strong context is one that allows addressees to anticipate an ironic utterance. This expectation should, in turn, render irony interpretation frictionless (*The expectation hypothesis*).

But a close look at "ironic situations" reveals that they need not promote an expectation for an ironic utterance nor do they facilitate irony interpretation (Giora, Fein, Kaufman, Eisenberg, and Erez, 2009; see also Ivanko and Pexman, 2003). Rather, such contexts encouraged readers to select literal utterances (*This demonstration is a remarkable failure*), which were by far the most preferred option, over ironic ones (*This demonstration is a remarkable success*), whether following a context featuring a frustrated Section 5.1.1 or a fulfilled Section 5.1.2 expectation (see Giora *et al.* (2009).

5.1.1. Frustrated expectation

Shirley is a feminist activist. Two weeks ago, she organized a demonstration against the closure of a shelter for victimized women, and invited the press. She hoped that due to her immense efforts many people would show up at the demonstration, and that the media would cover it widely. On the day of the demonstration, 20 activists arrived, and no journalist showed up. In response to the poor turn out, Shirley muttered:

a. This demonstration is a remarkable success. (Ironic)

b. This demonstration is a remarkable failure. (Literal)

5.1.2. Realized expectation

Shirley is a feminist activist. Two weeks ago, she organized a demonstration against the closure of a shelter for victimized women, and invited the press. As always, she prepared herself for the idea that despite the hard work, only a few people will show up at the demonstration and the media will ignore it entirely. On the day of the demonstration, 20 activists arrived, and no journalist showed up. In response to the poor turn out, Shirley muttered:

a. This demonstration is a remarkable success. (Ironic)

b. This demonstration is a remarkable failure. (Literal)

In addition, "ironic situations" did not facilitate irony. Rather, ironic statements (*Ski vacation is recommended for your health*) took as long to read following a context featuring a contrast between what is expected and the reality that frustrates it (Section 5.1.3) as following a context in which this expectation is met (Section 5.1.4).

5.1.3. Frustrated expectation

Sagee went on a ski vacation abroad. He really likes vacations that include sport activities. A relaxed vacation in a quiet ski-resort place looked like the right thing for him. Before leaving, he made sure he had all the equipment and even took training classes on a ski simulator. But already at the beginning of the second day he lost balance, fell, and broke his shoulder. He spent the rest of the time in a local hospital ward feeling bored and missing home. When he got back home, his shoulder still in cast, he said to his fellow workers:

"A ski vacation is recommended for your health". (Ironic)

Everyone smiled.

5.1.4 Realized expectation

Sagee went on a ski vacation abroad. He doesn't even like skiing. It looks dangerous to him and staying in such a cold place doesn't feel like a vacation at all. But his girlfriend wanted to go and asked him to join her. Already at the beginning of the second day he lost balance, fell, and broke his shoulder. He spent the rest of the time in a local hospital ward feeling bored and missing home. When he got back home, his shoulder still in cast, he said to his fellow workers:

"A ski vacation is recommended for your health". (Ironic)

Everyone smiled.

Importantly, however, both ironic targets took longer to read than a salience-based (literal) interpretation which followed a context featuring no expectation (Section 5.1.5).

5.1.5. No-expectation

Sagee went on a ski vacation abroad. He has never practiced ski so it was his first time. He wasn't sure whether he would be able to learn to ski and whether he will handle the weather. The minute he got there he understood it was a great thing for him. He learned how to ski in no time and enjoyed it a lot. Besides, the weather was nice and the atmosphere relaxed. When he got back home, he said to his fellow workers:

"A ski vacation is recommended for your health". (Salience-based, literal)

Everyone smiled.

5.1.6. Will expecting an ironic utterance facilitate it initially?

What, then, can make a strong context, such that would induce an expectation for an ironic utterance? In Giora, Fein, Laadan, Wolfson, Zeituny, Kidron, Kaufman, and Shaham (2007), we showed that the involvement of an ironic speaker in vivo (in context mid-position, in bold for convenience) induced an expectation for another such utterance on the part of that speaker when these contexts were presented without the final utterances and had to be completed by participants. We therefore used these contexts, completed by an utterance which was biased either toward the ironic (2) or toward its salience-based (literal) interpretation (3) in a reading experiment.

(2) Barak: I finish work early today.
Sagit: So, do you want to go to the movies?
Barak: I don't really feel like seeing a movie.
Sagit: So maybe we could go dancing?
Barak: No, at the end of the night my feet will hurt and I'll be tired.
Sagit: **You're a really active guy** …
Barak: Sorry, but I had a rough week.
Sagit: So what are you going to do tonight?
Barak: I think I'll stay home, read a magazine, and go to bed early.
Sagit: Sounds like you are going to have a really interesting evening.
Barak: So we'll talk sometime this week.

(3) Barak: I was invited to a film and a lecture by Amos Gitai.
Sagit: That's fun. He is my favorite director.
Barak: I know, I thought we'll go together.
Sagit: Great. When is it on?
Barak: Tomorrow. We will have to be in Metulla in the afternoon.[4]
Sagit: **I see they found a place that is really close to the center**.
Barak: I want to leave early in the morning. Do you want to come?
Sagit: I can't, I'm studying in the morning.
Barak: Well, I'm going anyway.
Sagit: Sounds like you are going to have a really interesting evening.
Barak: So we'll talk sometime this week.

[4] Metulla is a town in in the Northern District of Israel.

Although both contexts raised an expectation for an ironic utterance, identical targets (*Sounds like you are going to have a really interesting evening*) took longer to read following ironically biasing contexts (2) than following salience-based, literally biasing contexts (3). Strong contexts, then, inducing an expectation for an ironic utterance, did not facilitate ironic interpretations nor did they slow down salience-based, literal interpretations.

To further test *the expectation hypothesis* we attempted to induce an expectation for an ironic utterance by presenting participants only with contexts that ended in an ironic utterance (4) so that they are trained to anticipate an ironic utterance. This (+Expectation) condition was compared to a weaker (–Expectation) condition in which only half the contexts ended in an ironic utterance; the other half ended in a nonironic utterance (5). Results from lexical decisions to probes related to ironic ("harmful") and salience-based ("healthy) utterance-level interpretations showed facilitation of the salience-based interpretation only, regardless of context bias. This was true when short (250 ms) as well as long (750–1000 ms) processing time was allowed. This pattern of results was not different from the one obtained in the weaker condition. Such results suggest that, contra *the direct access view*, but in keeping with *the graded salience hypothesis*, inducing an expectation for an ironic utterance does not facilitate ironic interpretation immediately and does not affect a seamless interpretation process:

(4) Yuval and Omry went out for their lunch break after a morning of work. They went to the cafeteria in their office building and each filled a platter with food. They stood in line for a long while and were eager to start the meal. When they had sat down, Yuval saw that his colleague chose fried sausage, chips, a glass of coke for a drink, and a sugar-glazed doughnut for desert. Then Yuval said: "I see that you picked the *ideal* meal today!"

(5) Yuval and Omry went out for their lunch break after a morning of work. They went to the cafeteria in their office building and each filled a platter with food. They stood in line for a long while and were eager to start the meal. When they had sat down to eat, Yuval saw that his colleague filled his platter with salad, tofu, and sprouts and chose natural carrot juice for a drink. Then Yuval said: "I see that you picked the *ideal* meal today!"

Importantly, in Giora and Fein (2010a), we strengthened the ironically biasing condition used in Giora, Fein, Laadan *et al.* (2007) by introducing an additional constraint, informing participants that the aim of the experiment was to test irony interpretation. The control group, whose experimental design was mixed, raising no expectation, were not informed about this specific aim of the experiment; their contextual information was therefore weaker compared to that in which expectation for ironic utterances was made more pronounced, both implicitly and explicitly. Still, although contextual information was now more strongly biased in favor of the ironic interpretation, this did not affect the pattern of results which replicated those obtained earlier (in Giora, Fein, Laadan *et al.*, 2007). Even this multiple constraints condition did not facilitate irony interpretation; only salience-based albeit incompatible interpretations were made available in both the strongly and weakly biasing conditions.

Would allowing extra processing time make a difference? In Giora and Fein (2010b), we allowed participants longer (1500 ms) processing time. We predicted that even if, at this stage, irony is understood, salience-based but incompatible interpretations would still be available. Indeed, as predicted, even at such a long delay, pattern of results did not change: the salience-based though incompatible interpretation was never less accessible than the ironic yet compatible interpretation.

Such results demonstrate that understanding utterances in a strong context supportive of and anticipating their nonsalient (ironic) interpretation via inducing an expectation for such an interpretation does not unconditionally involve dispensing with the salience-based but incompatible interpretation. Do salience-based interpretations have a role in shaping up utterance final products?

We have seen that processing utterances in context may involve entertaining meanings and interpretations on account of their salience rather than because of their contextual compatibility. According to *the suppression/retention hypothesis* (Giora, 2003), such meanings and interpretations, when incompatible, will either be retained or discarded from the mental representation depending on the role they might play in shaping the contextually appropriate interpretation. When they might contribute to the final representation they will be retained. For instance, on *the indirect negation view*, the involvement of salience-based interpretation in irony processing allows computing the gap between what is said and the reality referred to (Giora, 1995 and Giora *et al.*, 2009); on the *tinge hypothesis* (Dews, Kaplan and Winner, 1995; Dews and Winner, 1995, 1997, 1999), the involvement of salience-based interpretations in irony processing is functional in mitigating the negativity of ironic criticism and the positivity of ironic praises. Although salience-based interpretations might not be the intended interpretation, they are instrumental in shaping it up and are therefore retained.

The involvement of salience-based but incompatible interpretations in the late stages of irony interpretation supports *the graded salience hypothesis* and *the suppression/retention hypothesis* (Giora, 1997, 1999, 2003). However, it argues against both *the standard pragmatic model* (Grice, 1975)[5] and *the direct access view* (Gibbs, 1986, 2002) which assume that the products of irony processing will not involve incompatible interpretations.

5.2. Negation interpretation

Consistent with *the graded salience hypothesis*, results obtained from studies of irony interpretation support the view that salient meanings and salience-based interpretations are activated initially, regardless of a strong contextual bias toward the nonsalient interpretation. They further demonstrate that, as predicted by the *retention/suppression hypothesis* (Giora, 2003; Giora and Fein, 1999a), suppression of contextually incompatible meanings and interpretations is not unconditional. Instead, suppression is attuned to contextual goals and requirements and would not operate if apparently incompatible information may be conducive to the final interpretation.

---

[5] See also Carston, this volume.

A great number of studies probing the effect of negation (*no*, *not*) on patterns of activation of negated concepts show that, outside a specific context, negated information ("fast" in *The train to Boston was* ***no rocket***) is activated initially (between 100-500 ms) as does nonnegated information (Hasson and Glucksberg, 2006; Giora, Balaban, Fein, and Alkabets, 2005; Kaup, Yaxley, Madden, Zwaan, Lüdtke, 2007; MacDonald and Just, 1989 Experiments 1-2 reading phase). However, later on (between 500-1000 ms) initial levels of activation of negated concepts drop to base-line levels (Hasson and Glucksberg, 2006). When given extra processing time (1500 ms), negated information is suppressed and replaced by an alternative opposite (Kaup, Lüdtke, and Zwaan, 2006; for conflicting results, however, see Lüdtke, Friedrich, De Filippis, and Kaup, 2008). Outside a specific context, then, the suppressive effect of negation is a default strategy.

However, when provided with late coherent vs. incoherent context (*The train to Boston was* ***no rocket****. The trip to the city was *fast*, though*), such negated items did not dispense with the negated, albeit incompatible concept ("fast") even as long as 1000 ms following its mention. Retention of apparently irrelevant, salience-based interpretations, however, allow for late context to resonate with earlier context, despite indication to the contrary invited by the negation marker (Giora, Fein, Aschkenazi *et al.*, 2007; on discourse resonance, see Du Bois, 1998, 2001; on discourse resonance following negation, see Giora, 2007).

Contextual effects on the retention of negated information were also found in Kaup (2001) and Kaup and Zwaan (2003). Results of these studies show that a concept's accessibility may be affected by its presence in the situation model rather than by negation. In the following (6), the fact that a referent (*photographs*) is not removed from the situation allows its retention despite it being within the scope of negation (Kaup, 2001):

(6) Elizabeth tidied up her drawers. She burned the old letters but not the photographs. Afterwards she cleaned up.

Similarly, in Kaup and Zwaan (2003), only concepts absent from the situation described lost accessibility, regardless of negation. In contrast, negated concepts present in the situation described retained their accessibility even after being entertained for as long as 1500 ms.

Suppression and retention are functional, then, and conform to global rather than to local coherence considerations (Giora 2006).

6. Coda

Is our mind "efficient" enough to engage in processing only contextually appropriate interpretations given that contextual information is strongly supportive of that interpretation, as argued by the conceptualist school (e.g., Gibbs, 1994; Vu *et al.*, 2000; see Section 1.1)? Luckily, it is not. In fact, there is enough evidence now to allow the conclusion that processing utterances, even inside highly biasing contexts, may involve entertaining meanings and interpretations solely on account of their meaning salience and consequently their salience-based interpretation, regardless of contextual fit, as argued by the graded salience hypothesis (see Section 1.2). Additionally, the activation of such

incompatible meanings and interpretations in utterance processing is not unconditionally aborted by suppression processes (as assumed by Fodor, 1983 or Grice, 1975). Rather, such meanings and interpretations are retained because they are deemed functional in shaping up final interpretations (*the suppression/retention hypothesis*, Giora, 2003). Such a functional view of suppression and retention allows for the poetics of linguistic and nonlinguistic stimuli such as optimal innovations (whether literal or nonliteral; Section 1), for discourse resonance, and humor (Section 5).

It has also become clear that, contra to Grice (1975), nonliteral interpretations may be a default interpretation even when innovative, free of semantic anomaly, and context-less (Section 2.3). Indeed, when affirmative statements of the form X is Y are negated (*This is not Memorial Day*), they are assigned a nonliteral interpretation even when presented in isolation. Similarly, negative overstatements such as *He is not exceptionally bright* are assigned an ironic interpretation even outside a specific context.

The review of the literature introduced in this chapter reveals that the psychology of utterance processing is a multi-faceted phenomenon; its products may, at times, be surprisingly creative and even amusing.

Acknowledgments

This paper was supported by a grant to the second author by THE ISRAEL SCIENCE FOUNDATION (grant No. 652/07). Thanks also go to Ran Abramson for the cartoons in Figures 8.1, 8.2 and for example (1) and to Keith Allan and Kasia Jaszczolt for their very valuable comments.

References

Adams**,** December 22, 2009. Spirit of the Blitz 2009. http://www.telegraph.co.uk:80/comment/cartoon/?cartoon=6863333&cc=6707029

Bates, E. A. (1999). On the nature and nurture of language. In B. E., C. P. and V. V. (Eds.), *Frontiere della biologia [Frontiers of biology]. Il cervello di Homo sapiens [The brain of homo sapiens]* (pp. 241-265). Rome: Istituto della Enciclopedia Italiana fondata da Giovanni Trecanni S.p.A.

Bates, E. A., and MacWhinney, B. J. (1989). Functionalism and the Competition Model. In B. J. MacWhinney and E. A. Bates (Eds.), *The crosslinguistic study of sentence processing*. New York: Cambridge University Press.

Beardsley, M. C. 1958. *Aesthetics*. New York: Harcourt, Brace and World.

Brisard, F., Frisson, S., and Sandra, D. (2001). Processing unfamiliar metaphors in a self-paced reading task. *Metaphor and Symbol, 16*(1and2), 87-108.

Colston, H. L., and Gibbs, R. W. (2002). Are irony and metaphor understood differently? *Metaphor and Symbol, 17*, 57-60.

Dews, S., Kaplan, J., and Winner, E. (1995). Why not say it directly? The social functions of irony. *Discourse Processes, 19*, 347-367.

Dews, S., and Winner, E. (1995). Muting the meaning: A social function of irony. *Metaphor and symbolic activity, 10*(1), 3-19.

Dews, S., and Winner, E. (1997). Attributing Meaning to Deliberately False Utterances: The Case of Irony. In C. Mandell and A. McCabe (Eds.), *The Problem of Meaning: Behavioral and Cognitive Perspectives* (pp. 377-414). Amsterdam, Holand: Elsevier.

Dews, S., and Winner, E. (1999). Obligatory Processing of the Literal and the Nonliteral Meanings of Ironic Utterances. *Journal of Pragmatics, 31*, 1579-1599.

Du Bois, John W. 1998 Dialogic syntax. Paper presented at the Cognitive Theories of Intertextuality Meeting. Tel Aviv University.

Du Bois, John W. 2001 Towards a dialogic syntax. Unpublished ms., University of California Santa Barbara.

Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.

Frazier, Lyn, and Rayner, Keith (1990). Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language, 29*, 181-200.

Gibbs, R. W. (1979). Contextual effects in understanding indirect requests. *Discourse Processes, 2*, 1-10.

Gibbs, R. W. (1986). On the psycholinguistics of sarcasm. *Journal of Experimental Psychology, General*(115), 3-15.

Gibbs, R. W. (1994). *The poetics of mind: Figurative thought, language, and understanding*. New York: Cambridge University Press.

Gibbs, R. W. (2002). A new look at literal meaning in understanding what is said and implicated. *Journal of Pragmatics, 34*, 457-486.

Giora, R. (1997). Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics, 8*(3), 183-206.

Giora, R. (1999). On the priority of salient meanings: Studies of literal and figurative language. *Journal of Pragmatics, 31*, 919-929.

Giora, R. (2003). *On our mind: Salience, context, and figurative language*. New York: Oxford University Press.

Giora, R. (2007). A good Arab is not a dead Arab - a racist incitement: On the accessibility of negated concepts. In L. R. Horn and I. Kecskes (Eds.), *Explorations in Pragmatics: Linguistic, Cognitive and Intercultural Aspects* (pp. 129-162). Berlin: Mouton de Gruyter.

Giora, R., Federman, S., Kehat, A., Fein, O., and Sabah, H. (2005). Irony aptness. *Humor, 18*, 23-39.

Giora, R., and Fein, O. (1999a). On understanding familiar and less-familiar figurative language. *Journal of Pragmatics, 31*, 1601-1618.

Giora, R., and Fein, O. (1999b). Irony: Context and salience. *Metaphor and Symbol*, *14*, 241-257.

Giora, R., and Fein, O. (2010a). On the priority of salience-based interpretation: The case of irony *Paper submitted for publication*.

Giora, R., and Fein, O. (2010b). On the involvement of salience-based but inappropriate interpretation in the products of irony. *Paper Submitted for Publication*.

Giora, R., Fein, O., Aschkenazi, K., and Alkabets-Zlozover, I. (2007). Negation in context: A functional approach to suppression. *Discourse Processes, 43*, 153 - 172.

Giora, R., Fein, O., Kaufman, R., Eisenberg, D., and Erez, S. (2009). Does an "ironic situation" favor an ironic interpretation? . In B. G. and V. J. (Eds.), *Cognitive poetics. Goals, gains and gaps (Applications of Cognitive Linguistics series).* (pp. 383-399). Berlin/New York: Mouton de Gruyter.

Giora, R., Fein, O., Kronrod, A., Elnatan, I., Shuval, N., and Zur, A. (2004). Weapons of mass distraction: Optimal Innovation and Pleasure Ratings. *Metaphor and Symbol, 19*, 115-141.

Giora, R., Fein, O., Laadan, D., Wolfson, J., Zeituny, M., Kidron, R., *et al.* (2007). Expecting irony: Context vs. salience-based effects. *Metaphor and Symbol 22*(2), 119 - 146.

Giora, R., Fein, O., Metuki, N., and Stern, P. (2010). Negation as a metaphor-inducing operator. In L. R. Horn (Ed.), *The Expression of Negation.* (pp. 225-256). Berlin/New York: Mouton de Gruyter.

Giora, R., Gazal, O., Goldstein, I., Fein, O., and Stringaris, A. (2010). Salience and context: Interpreting metaphors and literals by young adults diagnosed with Asperger's syndrome. *Paper submitted for publication.*

Grice, H. P. (1975). Logic and Conversation

In D. Davidson and g. Harman (Eds.), *The Logic of Grammar* (pp. 64-75). Encino, CA: Dickenson.

Hasson, U., and Glucksberg, S. (2006). Does negation entail affirmation? The case of negated metaphors. *Journal of Pragmatics, 38*, 1015-1032.

Hughes, Stuart. February 4, 2009. The greatest motivational poster ever? http:// news.bbc.co.uk/2/hi/uk_news/magazine/7869458.stm

Ivanko, S. L., and Pexman, P. M. (2003). Context incongruity and irony processing. *Discourse Processes, 35*, 241-279.

Kaup, Barbara (2001). Negation and its impact on the accessibility of text information, *Memory and Cognition* 29, 960–967.

Kaup, Barbara, Lüdtke, Jana Zwaan, Rolf A 2006 Processing negated sentences with contradictory predicates: Is a door that is not open mentally closed? *Journal of Pragmatics* 38: 1033–1050.

Kaup, B., Yaxley, R. H., Madden, C. J., Zwaan, R. A., and Lüdtke, J. (2007). Experiential simulations of negated text information. *Quarterly journal of experimental psychology 60*(7), 976-990.

Kaup, B., and Zwaan, R. A. (2003). Effects of negation and situational presence on the accessibility of text information. *Journal of experimental psychology. Learning, memory, and cognition, 29*(3), 436-439.

Keysar, B. (1994). Discourse context effects: Metaphorical and literal interpretations. *Discourse Processes, 18*, 247-269.

Lüdtke, J., Friedrich, C. K., De Filippis, M., and Kaup, B. (2008). Event-related Potential Correlates of Negation in a Sentence-Picture Verification Paradigm. *Journal of Cognitive Neuroscience, 20*(8), 1355-1370.

MacDonald, M. C., and Just, M. A. (1989). Changes in activation levels with negation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 633-642.

MacWhinney, B. J. (1987). The competition model. In B. J. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 249-308). Hillsdale, NJ: Erlbaum.

McRae, K., Spivey-Knowlton, M. J., and Tanenhaus, M. K. (1998). Modeling thematic fit (and other constraints) within an integration competition framework. *Journal of Memory and Language, 38*, 283-312.

Ortony, A., Schallert, D. L., Reynolds, R. E., and Antos, S. J. (1978). Interpreting metaphors and idioms: Some effects of context on comprehension. *Journal of Verbal Learning and Verbal Behavior, 17*(465-477).

Peleg, O., and Eviatar, z. (2008). Hemispheric sensitivities to lexical and contextual constraints: Evidence from ambiguity resolution. *Brain and Language, 105*(2), 71-82.

Peleg, O., and Eviatar, z. (2009). The disambiguation of homophonic versus heterophonic homographs in the two cerebral hemispheres. *Brain and Cognition, 70*, 154-162.

Peleg, O., and Giora, R. (2010). Salient meanings: The whens and wheres. In K. M. J. K. Allan (Ed.), *Salience and Defaults in Utterance Processing*. Berlin/New York: Mouton de Gruyter.

Perry, Jimmy and David Croft. 1968-1977. *Dad's Army*. British Broadcasting Corporation (BBC). UK.

Pexman, P. M., Ferretti, T. R., and Katz, A. N. (2000). Discourse factors that influence on-line reading of metaphor and irony. *Discourse Processes, 29*(201-222).

Searle, J. R. (1979). *Expression and meaning: Studies in the theory of speech acts*. Cambridge, London, New York, Melbourne: Cambridge University Press.

Small, S. L., Cottrell, G. W., and Tanenhaus, M. K. (Eds.). (1998). *Lexical ambiguity resolution and language comprehension: Computational, psycholinguistic and neurolinguistic perspectives*. San Diego, CA: Morgan-Kaufman.

Swinney, D. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18*(645-660).

Tartter, Vivien C., Hilary Gomes, Boris Dubrovsky, Sophie Molholm and Rosemarie Vala Stewart (2002). Novel metaphors appear anomalous at least momentarily: Evidence from N400. Brain and Language 80/3, 488-509.

Vu, H., Kellas, G., Metcalf, K., and Herman, R. (2000). The influence of global discourse on lexical ambiguity resolution. *Memory and Cognition, 28*, 236-252.

Vu, H., Kellas, G., and Paul, S. T. (1998). Sources of Sentence Constraint on Lexical Ambiguity Resolution. *Memory and Cognition, 26*(5), 979-1001.