

Rachel Giora, Shir Givoni and Israela Becker

How defaultness affects text production: Resonating with default interpretations of negative sarcasm

Abstract: According to the Defaultness Hypothesis, interpretations of constructions, involving strong attenuation (e.g., by means of negation) of highly positive concepts, such as *S/he is not the most mesmerizing person around*; *S/he is not particularly smart*; *S/he is not extremely friendly*; and *S/he is not really the ideal teacher*, will spring to mind by default, immediately and directly. Hence, when in natural discourse, such constructions will be echoed by their environment via their default, here, sarcastic interpretation (e.g., *S/he is dull*; *S/he is stupid*; *S/he is reserved*; *S/he is the worst teacher*). Results show that, in natural discourse, default rather than nondefault interpretations prevail; indeed, the contexts of the negative constructions studied here evolve and unfold via resonating with their default interpretations.

Keywords: The Defaultness Hypothesis, default/nondefault interpretations, contextual resonance, negative sarcasm

1 The defaultness hypothesis

According to the Defaultness Hypothesis (Giora, Givoni, and Fein 2015b), **default** responses crucially affect our language comprehension and production. For a response (here, a constructed interpretation) to be considered a default, it has to be activated automatically, immediately and directly, regardless of other factors assumed to affect processing, such as degree of negation/affirmation (e.g., Clark and Clark 1977; Horn 1989), non/literalness (Grice 1975), novelty (nonsalience

Acknowledgments: We are very grateful to John Barnden and to Mihaela Popa for their illuminating comments. This research was supported by an Israel Science Foundation grant (no. 540/19) awarded to Rachel Giora. Shir Givoni was awarded a scholarship for outstanding graduate students funded by the Rector of Tel-Aviv University and matched by her supervisor – Rachel Giora. Israela Becker was awarded a scholarship for outstanding graduate students, funded by the Rector of Tel-Aviv University and by an Israel Science Foundation grant (no. 431/15) awarded to her supervisor – Mira Ariel.

Rachel Giora, Shir Givoni and Israela Becker, Tel Aviv University

<https://doi.org/10.1515/9783110652246-003>

vs. salience-based; see Giora 2003), or strength of contextual support (strong vs. weak; see e.g., Gibbs 1986, 1994, 2002).

When tested experimentally, an automatic response will be considered a default, if its stimulus, whether literal or nonliteral, is

- (i) novel, so that the response, activated automatically, will be noncoded (non-salient or salience-based), i.e., constructed, rather than accessed directly from the mental lexicon;
- (ii) free of internal cues such as semantic anomaly or internal incongruity, prompting nonliteralness (see, Beardsley 1958; Partington 2011); and
- (iii) free of explicit contextual information, including cues, so that a preference (among various responses) is allowed (for more details, see Giora et al. 2015b).

1.1 The defaultness hypothesis: Predictions

According to the Defaultness Hypothesis (Giora et al. 2015b), utterances, meeting the conditions for default interpretations (i-iii above), further involving strong attenuation (*not*) of highly positive concepts (*most candid*) as in, *S/he is not the most candid person I know* (see Giora et al. 2005; Giora et al. 2018), will be interpreted sarcastically by default.¹ Their non-attenuated (e.g., affirmative) counterparts will be interpreted compositionally (often literally) by default. Specifically,

- (a) when presented in isolation, utterances, involving strong attenuation (*not*) of highly positive concepts (e.g., *S/he is not the most candid person I know*), will be interpreted sarcastically (meaning ‘S/he is dishonest’) and rated as more sarcastic compared to their literal interpretation (meaning ‘S/he is candid but others are more candid than her/him’); their non-attenuated affirmative counterparts (*S/he is the most candid person I know*) will be interpreted compositionally (here, literally) by default (meaning ‘S/he is very honest’; see Giora et al. 2015b Exp. 1).
- (b) Hence, when in equally strong contexts, such utterances will be processed faster when embedded in contexts, biasing them towards their default nonsalient sarcastic interpretation than towards their **nondefault** (equally strongly biased) salience-based (literal) interpretation. Furthermore, when in equally strong contexts, supportive of their sarcastic interpretation, they will also be processed faster than their nondefault non-attenuated (e.g., affirmative)

¹ Sarcasm refers here to verbal irony. Verbal irony can also be classified as an understatement. For example, Gibbs (2000), who investigated five main forms of irony, found that understatements, the category to which the negative constructions studied here belong, is the least frequent one — ~0.2%.

- counterparts, biased towards the same, equally strongly supported sarcastic interpretation (Giora et al. 2015b Exp. 2).
- (c) Therefore, when in natural discourse, (i) these negative constructions will be rated as sarcastic rather than literal; their non-attenuated affirmative counterparts will be rated as (salience-based, often) literal (as shown by Becker and Giora 2018). Hence (ii) the neighboring utterances of these attenuated (e.g., negative) items will echo or resonate with their highly accessible default interpretations more often than with their nondefault literal interpretations, low on accessibility. Along the same line, their nondefault, non-attenuated, contextually compatible affirmative sarcastic counterparts will also be echoed via their default, contextually incompatible, salience-based (often literal) interpretation (as shown for similar constructions by Giora, Drucker, and Fein 2014a; Giora et al. 2014b).
- (d) Furthermore, when in natural discourse, it is nondefault interpretations (not high on speakers' and addressees' mind) that rely on cues, rejecting utterances default automatic interpretations (e.g., the sarcastic interpretation of specific negative constructions, such as those studied here, and the compositional often literal interpretation of their affirmative counterparts). Indeed, prompting contextually appropriate yet *nondefault interpretations* by cues, rejecting default counterparts, while inviting nondefault alternatives, has been attested to by Becker and Giora (2018) for both nondefault negative literalness and nondefault affirmative sarcasm, and by Givoni, Giora, and Bergerbest (2013) for nondefault *meanings*.

Here we test predictions (1.1ci-cii) with regard to strongly attenuated highly positive concepts, such as *S/he is **not** the **most** candid person I know*, which, according to Giora et al. (2005, 2015b, 2018), are interpreted sarcastically by default (1.1ci). Hence, as predicted by the Defaultness Hypothesis, they will affect both the prevalence of their sarcastic interpretations (1.1ci) and their effect on discourse production (1.1cii).

Note that the defaultness of such constructions has been established earlier, both experimentally and via corpus-based studies. When tested experimentally, constructions, negatively attenuating highly positive concepts (e.g., *Candidness is not her/his forte/best attribute*; *Candid s/he is not*; *S/he is not the most candid person I know*) were shown to be interpreted sarcastically when in isolation ('S/he is dishonest'). They were therefore processed faster in contexts biasing them towards their default sarcastic interpretation ('S/he is dishonest') than towards their equally strongly biased nondefault literal counterpart ('S/he has other strong attributes'/'S/he is honest but others are more honest than her/him'; see Giora et al. 2015a; Giora et al. 2013). They were also shown to be processed faster than nondefault affirmative sarcasm, as in *S/he is the most candid person I know* (Giora et al. 2015b).

Usage-wise, stimuli such as *Candidness is not her/his forte/best attribute; Candid s/he is not*, extracted from Hebrew corpora, were further shown to be interpreted by 3 external expert judges as sarcastic and to be echoed by their natural context via their default sarcastic interpretation (Giora et al. 2013, 2014a). Additionally, stimuli, such as the negative constructions studied here (*S/he is not the most candid person I know*,) were also shown to be rated as sarcastic when in natural use. For instance, Becker and Giora (2018) conducted a corpus-based study by compiling a list of 171 instances of such negative constructions, preceded by a metalinguistic comment *to put it mildly*, which acts as an additional attenuator, mitigating a harsh message. These 171 items accommodated adjectival phrases, noun phrases, or verb phrases. The vast majority of these items (167/171=97%) were rated as sarcastic by at least 2 out of the 3 expert judges, versed in the field of non/literal language. In addition, the judges were asked to indicate the polarity of the concept in the scope of these constructions, whether it is positive, neutral, or negative. Results show that 146 items of these 167 sarcastic items were rated as positive (by 2 or more judges), whereas only 21 items were rated as non-positive (by 2 or more judges). Namely, there are significantly more positively-oriented concepts (146/167=87%) in the scope of these negative constructions than non-positive concepts (21/167=13%), binomial test, $p=2.68 \times 10^{-24}$ (highly more significant than $p < 0.001$).²

These recent results provide converging corpus-based evidence to our online findings, attesting to the close association between the sarcastic interpretation of attenuated constructions and the polarity of the concepts which they host. Put differently, the default sarcastic interpretation of these negative constructions is induced by the strong attenuation of the highly positive polarity of the concept it scopes over.

Having shown that constructions, strongly attenuating highly positive concepts, such as *S/he is not the most candid person I know*, are rated as sarcastic, both in natural context and out of context (see predictions (1.1a) & (1.1ci)), and are processed faster than nondefault counterparts (see prediction (1.1b)), here we will look for converging evidence showing that these items' default interpretation further affects discourse production via resonating with default interpretations. Based on items from the Hebrew *HeTenTen* web-corpus (see Appendix A below), we will test here prediction (1.1cii) of the Defaultness Hypothesis, related to resonating with default sarcastic interpretations.

² Note that in Becker and Giora (2018), the authors also provided evidence supportive of prediction (d), regarding cueing.

2 Resonating with default interpretations

In this section, we test predictions (1.1cii) of the Defaultness Hypothesis, related to discourse resonance with default interpretations.

2.1 What is resonance?

Resonance is defined as “the catalytic activation of affinities across utterances” (Du Bois 2014: 359; Du Bois and Giora 2014: 351), involving given and new information, uttered within and between speakers, in both prior and subsequent context (Du Bois, 2007, 2014; Du Bois and Giora 2014; Giora 2007). Resonance is, therefore, a property of the relation between elements in discourse. The affinities activated may be based on similarity (e.g., ‘smart’ and ‘bright’) or difference (e.g., ‘smart’ and ‘stupid’), and resonance can be perceived whenever a suitable structural parallelism, supporting the affinity, occurs. Resonance can arise from parallelism in pairs across meanings (e.g., ‘not the smartest’ and ‘stupid’) as well as structures or constructions (e.g., ‘not the brightest’ and ‘not exactly bright’). Whenever language users reproduce some aspect of a prior or subsequent utterance, they create parallelisms and resonances. These, in turn, result in an environment that aligns with utterances’ interpretations (Du Bois 2004, 2014). As such, resonance can reflect the interpretation of a given ambiguous (i.e., sarcastic or literal) utterance.

Previous research, attesting experimentally to the defaultness of sarcastic interpretations of various negative constructions (e.g., *Candidness is not her/his forte/best attribute*; see Giora, et al. 2014a; or *Candid s/he is not*; see Giora et al. 2013), further attests to contextual resonance with these default sarcastic interpretations, rather than their nondefault literal interpretation. For instance, in Giora et al. (2014a), findings from 2 corpus-based studies of (Hebrew and English) negative constructions lend usage-based support to the Defaultness Hypothesis (see also Giora et al. 2010, 2013, 2014b). They show that, when in natural discourse, such utterances are interpreted sarcastically and rated as more sarcastic than their affirmative alternatives. Hence, their neighboring utterances further reflect their default nonsalient sarcastic interpretation rather than their nondefault salience-based often literal counterpart. In contrast, affirmative sarcasm, whose default interpretation is a salience-based often literal interpretation (see Giora et al. 2015a,b), is echoed by its contextual environment via that literal interpretation, even if contextually inappropriate (Giora et al. 2014b).

2.2 A corpus-based experiment

In order to test predictions (1.1ci-cii), we conducted an experiment in which 3 judges, versed in the field of sarcasm, were presented with 151 negative constructions of the following 4 variants – *s/he is not the most X*; *s/he is not particularly X / s/he is not X in particular*³; *s/he is not extremely X*; and *s/he is not really X* – embedded in their natural context. The judges' tasks were to decide (1) whether these constructions are sarcastically or literally interpreted and (2) whether the context, in which they are embedded, resonates with their sarcastic (here, default) or literal (here, nondefault) interpretation.⁴

2.2.1 Materials

We exhaustively extracted from *HeTenTen* web-corpus (see Appendix A) all instances of *not really X* (*what we learnt at school [...] is not really accurate*) (14159 instances); *not extremely X* (*Their food is not extremely tasty*) (8126 instances); *not the most X* (*on the face of it, not the most respectable profession*) (4000 instances); *not particularly X/not X in particular* (*It was probably not particularly successful; the leadership was not creative in particular*) (761 instances), hosting an adjective in the *X* slot. From each of the 4 lists, we then pseudo-randomly sampled 300 instances. Based on the authors' judgements, we further narrowed down the resulting list of 1200 results into a preliminary list of 1140 instances, hosting potentially positively-oriented (rather than non-positively-oriented) adjectives in the *X* slot of the negative constructions listed above. Then, following a strict selection procedure, we further filtered out instances hosting adjectives that were (a) semantically ambiguous (e.g., *gay*) (b) or positively-oriented only when in context (such as *ambitious, relevant, or quiet*). We ended up with 516 instances, following this filtering procedure.

Striving to produce as diverse as possible a list of highly-positive adjectives (embedded in the 4 negative constructions listed above), while minimizing repetition of adjectives but keeping the list as balanced as possible construction-wise, we compiled a set of 151 items – *not the most X* (47); *not particularly X/not X in particular* (40); *not extremely X* (34); *not really X* (30), comprising 103 diverse positive adjectives – comprising 73 unique adjectives, 13 adjectives repeated twice, 16 adjectives repeated 3 times, and 1 adjective repeated 4 times. Note that

³ This ordering of the construction (*s/he is not X in particular*) is equivalent to its Hebrew ordering.

⁴ The 3 external expert judges were, at the time, MA students. They rated the items individually without consulting each other or any of the authors.

adjectives repeated more than once were not repeated within the same constructions and were different from each other with regard to the grammatical gender and number, since in Hebrew the grammatical gender and number are explicitly marked on adjectives.

Instances of resonance are shown in (1–2), translated from our Hebrew items. Target sentences are underlined and in bold, as presented to the judges. For resonance with *default* sarcastic interpretation of sarcastic targets (in bold here, for convenience), see example (1); for resonance with *nondefault* literal interpretation of literal targets (in italics here, for convenience) see example (2):

- (1) A year ago Ha'aretz⁵ disclosed that some of Kafka's manuscripts were held in a small flat [...] in Tel Aviv. Max Brod, Kafka's friend, is the person who held the originals [...] some of which he later shared with the appropriate authorities but still kept many of the manuscripts. He left his property, including the manuscripts, to his secretary, Ilse Esther Hoffe. Hoffe dealt with the inheritance the way she felt like, which probably **wasn't particularly successful**. Apart from this hidden treasure being inaccessible to the public, she probably **wasn't particular about its preservation**. In 2006, Ha'aretz reports that Hoffe's neighbors complained about **bad smell coming up from her flat**. **The municipality's inspectors found out that the flat was populated by tens of cats and some dogs...**⁶
- (2) I saw in your response to the question on the origin of the Ashkenazim,⁷ that what we learnt in school about the expulsion of the Jewish people by the Romans **isn't really accurate**. As we get nearer to Tisha Be-Av⁸ (even though the massive deportation was after the Bar Kokhba revolt, according to what I have learnt), perhaps tell us *what is accurate*. *When did the exile begin? Why? Where were the Jewish people exiled to in the beginning? What happened with the Jewish people who did stay in the land of Israel?*⁹

2.2.2 Procedure

Our 3 expert judges were presented with the above list of 151 pseudo-randomly ordered items. Each item comprised one of the 4 negative constructions, embedded in their natural context, namely, 2–5 sentences, both preceding and following

⁵ *Ha'arets* is an Israeli daily newspaper.

⁶ http://mitzidlav.blogspot.co.il/2009/09/blog-post_08.html

⁷ *Ashkenazim* refers to Jewish people descending from (Eastern) European countries.

⁸ *Tisha B'Av* refers to an annual fast day in Judaism.

⁹ <http://news.nana10.co.il/Article/?ArticleID=502067>

the target construction, as in examples (1) and (2) above.¹⁰ Note that the target construction was marked in bold and was further underlined for the judges' convenience.

Apart from being asked to decide (1) whether these constructions are sarcastically or literally interpreted and (2) whether the embedding context resonates with any of the two optional interpretations (i.e., sarcastic or literal), the judges were also asked (3) to mark the resonance (if found) either in red, when sarcastic, or in blue, when literal; in addition, (4) they were further asked to state explicitly whether the resonance is content-related or structure-related, that is whether one of the 4 kinds of negative constructions is also structurally replicated. Note that our focus here is on the results of the first two assignments (1–2). However, the other two assignments (3–4) aimed at making the judges reflect upon their answers to their second (2) assignment, the one in which they were asked to check for any potential resonance with content.

2.2.3 Results

(a) Rating degree of non/literalness

Rating degree of non/literalness reveals that, as predicted (see 1.1ci), the list of 151 items comprises significantly more sarcastically-oriented constructions than literally-oriented constructions. Specifically, 2 or more judges rated 118 items of the 151 cases as sarcastic ($118/151=78\%$), whereas 2 or more judges rated the remaining 33 items as literal ($33/151=22\%$), binomial test, $p=1.96\times 10^{-12}$ (highly more significant than $p<0.001$). Such results provide support for the prediction that the negative constructions of the kind studied here convey their default sarcastic interpretation more often than their nondefault literal interpretations, when hosting positive concepts.

(b) Resonance with non/literalness

Deciding whether the neighboring utterances resonate with any of the two optional interpretations (i.e., sarcastic or literal), reveals that, as predicted (see 1.1cii), 2 or more judges indicated that the environment of 109 out of 118 sarcastic cases, resonated with any of their interpretations. (In the remaining 9 items, no

¹⁰ Our predictions don't distinguish between resonance with targets by an early context as opposed to resonance with targets by a late context; still it makes more sense to expand on targets by late context rather than by early context. Indeed, generally speaking, less resonating segments were found before as opposed to after targets, but that doesn't seem to make a difference.

resonance was detected). Specifically, there were significantly more items whose environment echoed their default sarcastic interpretation (77/109=71%) than their nondefault literal interpretation (32/109=29%), binomial test, $p=1.94 \times 10^{-5}$ (more significant than $p < 0.001$).

In all, these results support the Defaultness Hypothesis. They show that, as predicted, these negative constructions (a) convey their default sarcastic interpretation more often than their nondefault literal interpretation; namely, significantly more such negative items are interpreted sarcastically than literally. They further show that, as predicted, when interpreted sarcastically, (b) their environment resonates with their default sarcastic interpretation rather than their nondefault literal interpretation, when hosting positively-oriented adjectives.

3 Discussion and conclusion

According to the Defaultness Hypothesis (Giora, et al. 2015b, 2018), default interpretations of some negative constructions will spring to mind unconditionally, initially and directly, irrespective of negation, novelty, nonliteralness, or contextual support. A case in point is the sarcastic interpretation of e.g., *not particularly successful* in *Hoffe dealt with the inheritance the way she felt like, which probably **wasn't particularly successful*** (rated here as sarcastic, meaning ‘unsuccessful’, see example (1) above). Indeed, rating degree of non/literalness of such naturally occurring instances of the form *not the most X*; *not particularly X*; *not X in particular*; *not extremely X*; and *not really X*, involving strong attenuation (e.g., by negation) of highly positive concepts (e.g., *particularly successful*) reveals that their vast majority are judged as sarcastic, as predicted by the Defaultness Hypothesis (see ci, section 1.1 above).

Prediction (1.1b), when first tested experimentally in Giora et al. (2015b Exp. 2), was supported by results attesting to the speed superiority of default interpretations over nondefault counterparts. Here, as per predictions (1.1ci-cii), it has gained further corpus-based support. Findings adduced here indicate superior prevalence of default negative sarcasm (see example (1) above) over nondefault negative literalness (see example (2) above, and as shown in subsection (a) of the Results, in subsection 2.2.3).

The Defaultness Hypothesis further predicts (see cii, section 1.1) that, as a result of their speed superiority (attested to by Giora et al. 2015b), the negative constructions’ neighboring utterances will resonate with their default sarcastic interpretation (in bold, in example (1) above) rather than with nondefault literal alternatives (in italics, in example (2) above). Being so prominent on our mind,

these interpretations will affect discourse production, which will evolve while reflecting these utterances' default interpretations. As predicted by the Defaultness Hypothesis, resonating with default interpretations will supersede resonating with nondefault counterparts. Investigating these constructions' effect on prior and ongoing natural discourse, we show here that default rather than nondefault interpretations prevail. Specifically, the default sarcastic interpretations of the negative constructions studied here are reflected by their neighboring utterances significantly more often than their nondefault literal counterparts (see subsection (b) of Results, in subsection 2.2.3).

In sum, results in this corpus-based study provide converging usage-based evidence supportive of the Defaultness Hypothesis, both in terms of the prevalence of their sarcastic interpretations and their effect on the discourse production.

References

- Adler, Meni. 2007. *Hebrew morphological disambiguation: An unsupervised stochastic word-based approach*. Beer-Sheva, Israel: Ben-Gurion University of the Negev dissertation.
- Baroni, Marco, Silvia Bernardini, Adriano Ferraresi & Eros Zanchetta. 2009. The WaCky wide web: A collection of very large linguistically processed web-crawled corpora. *Language Resources and Evaluation* 43 (3). 209–226.
- Beardsley, Monroe C. 1958. *Aesthetics*. New York, NY: Harcourt, Brace and World.
- Becker, Israella, & Rachel Giora. 2018. The Defaultness Hypothesis: A quantitative corpus-based study of non/default sarcasm and literalness production. *Journal of pragmatics* 138. 149–164.
- Clark, Herbert H. & Eve V. Clark. 1977. *Psychology and Language*. New York, NY: Harcourt Brace Janovich, Inc.
- Du Bois, John W. 2004. Searching for intersubjectivity: 'Too' and 'either' in stance alignment. Paper presented at the the 25th Conference of the International Computer Archive of Modern and Medieval English (ICAME). University of Verona, 19–23 May.
- Du Bois, John W. 2007. The stance triangle. In Robert Englebretson (ed.), *Stancetaking in discourse: Subjectivity, evaluation, interaction*, 139–182. Amsterdam: John Benjamins.
- Du Bois, John W. 2014. Towards a dialogic syntax. *Cognitive Linguistics* 25 (3). 359–410.
- Du Bois, John W. & Rachel Giora. 2014. From cognitive-functional linguistics to dialogic syntax. *Cognitive Linguistics* 25 (3). 351–357.
- Gibbs, Raymond W. 1986. On the psycholinguistics of sarcasm. *Journal of Experimental Psychology: General* 115 (1). 3–15.
- Gibbs, Raymond W. 1994. *The Poetics of Mind: Figurative Thought, Language, and Understanding*. New York, NY: Cambridge University Press.
- Gibbs, Raymond W. 2000. Irony in talk among friends. *Metaphor and Symbol* 15 (1–2). 5–27.
- Gibbs, Raymond W. 2002. A new look at literal meaning in understanding what is said and implicated. *Journal of Pragmatics* 34 (4). 457–486.

- Giora, Rachel. 2003. *On our Mind: Saliency, Context, and Figurative Language*. New York, NY: Oxford University Press.
- Giora, Rachel. 2007. "A good Arab is not a dead Arab – a racist incitement": On the accessibility of negated concepts. In Istvan Kecskes & Laurence R. Horn (eds.), *Explorations in Pragmatics: Linguistic, Cognitive and Intercultural Aspects*, 129–162. Berlin: Mouton de Gruyter.
- Giora, Rachel, Ari Drucker & Ofer Fein. 2014a. Resonating with default nonsalient interpretations: A corpus-based study of negative sarcasm. *Belgian Journal of Linguistics* 28. 3–18.
- Giora, Rachel, Ari Drucker, Ofer Fein & Itamar Mendelson. 2015a. Default sarcastic interpretations: On the priority of nonsalient interpretations. *Discourse Processes* 52 (3). 173–200.
- Giora, Rachel, Ofer Fein, Jonathan Ganzi, Natalie Alkeslassy Levi & Hadas Sabah. 2005. On negation as mitigation: The case of negative irony. *Discourse Processes* 39 (1). 81–100.
- Giora, Rachel, Ofer Fein, Nili Metuki & Pnina Stern. 2010. Negation as a metaphor-inducing operator. In Laurence R. Horn (ed.), *The Expression of Negation*, 225–256. Berlin: Mouton de Gruyter.
- Giora, Rachel, Shir Givoni & Ofer Fein. 2015b. Defaultness reigns: The case of sarcasm. *Metaphor and Symbol* 30 (4). 290–313.
- Giora, Rachel, Inbal Jaffe, Israela Becker & Ofer Fein. 2018. Strongly mitigating a highly positive concept: The case of default sarcastic interpretations. *Review of Cognitive Linguistics* 6 (1). 19–47.
- Giora, Rachel, Elad Livnat, Ofer Fein, Anat Barnea, Rakefet Zeiman & Iddo Berger. 2013. Negation generates nonliteral interpretations by default. *Metaphor and Symbol* 28 (2). 89–115.
- Giora, Rachel, Moshe Raphaely, Ofer Fein & Elad Livnat. 2014b. Resonating with contextually inappropriate interpretations in production: The case of irony. *Cognitive Linguistics* 25 (3). 443–455.
- Givoni, Shir, Rachel Giora & Dafna Bergerbest. 2013. How speakers alert addressees to multiple meanings. *Journal of Pragmatics* 48 (1). 29–40.
- Grice, Herbert Paul. 1975. Logic and conversation. In Peter Cole & Jerry L. Morgan (eds.), *Syntax and Semantics: Speech Acts* (Vol. 3), pp. 41–58. New York, NY: Academic Press.
- Horn, Laurence R. 1989. *A Natural History of Negation*. Chicago, IL: University of Chicago Press.
- Kilgarriff, Adam, Vít Baisa, Jan Bušta, Miloš Jakubíček, Vojtěch Kovář, Jan Michelfeit, Pavel Rychlý & Vít Suchomel. 2014. The Sketch Engine: Ten years on. *Lexicography* 1 (1). 7–36.
- Partington, Alan. 2011. Phrasal irony: Its form, function and exploitation. *Journal of Pragmatics* 43 (6). 1786–1800.

Appendix A

The corpus

As a consequence of the strict criteria we followed compiling the list of negative constructions, which are practically rare,¹¹ and given the need to compile a fair number of items for the sake of reliable statistics, we used the largest corpus of Modern Hebrew available, *HeTenTen*,¹² a web-corpus which comprises about 1×10^9 tokens.

HeTenTen was compiled using a web-crawler. It contains approximately 1.2×10^6 web documents which were mined, filtered, and processed using a generic algorithm suggested by Baroni et al. (2009). Crucially, the corpus was crawled in a way that would not allow it to be biased by topic while still covering a wide range of language varieties as represented over the web, including semi-spoken languages.

Each word of *HeTenTen* (i.e., ‘surface form’) was Part-of-Speech-tagged and morphologically annotated for additional morphological features (see Adler 2007), such as gender, number, affixes, etc.¹³ The corpus can be queried by using an extended version of a Corpus Querying Language (CQL)¹⁴ (Kilgarriff et al. 2014), which enables users to retrieve lines whose patterns are defined by specifying sequences of token and sub-token-level features. Due to the morphological annotation (Adler 2007), the queries can be defined over many features which address the rich morphology of Hebrew. Additionally, upon request, a wider context can be suggested beyond the line/sentence level, and there is always a pointer to the URL from which the web page was crawled.

¹¹ Only few instances of the 4 negative constructions (e.g., *not the most X*) are found in the Corpus of Spoken Israeli Hebrew (a free-access popular corpus of spoken Israeli Hebrew), which includes ~40,000 tokens, in contrast to several tens of the affirmative counterparts (e.g., *the most X*). For instance, only a single case of *not the most X* is found, whereas 34 cases of *the most X* are detected, binomial test, $p=1.2 \times 10^{-14}$ (highly more significant than $p < 0.001$). <http://cosih.com/>

¹² <http://www.sketchengine.co.uk>

¹³ <https://www.sketchengine.co.uk/hetenten-corpus/#Tokenattributes>

¹⁴ <http://cwb.sourceforge.net/temp/CQPTutorial.pdf>