

The Interaction of Bayesian Pragmatics and Lexical Semantics in Linguistic Interpretation: Using Event-related Potentials to Investigate Hearers' Probabilistic Predictions

Markus Werning and Erica Cosentino
(`{markus.werning, erica.cosentino}@rub.de`)

Department of Philosophy, Ruhr University Bochum

Abstract

We contrast two views of how contextual influence on sentence meaning composition can be explained. The Semantic Similarity View maintains that discourse context affects sentence meaning mainly because of the semantic similarity between the words in the discourse context and the words in the sentence (as measured by Latent Semantic Analysis). The Free Pragmatic View, in contrast, defends the claim that also pragmatic aspects of the discourse context can affect sentence meaning composition. This effect can be quantitatively modelled by Bayesian Pragmatics. We introduce a Predictive Completion Task in which the hearer at every moment in a communicative situation has to generate a probabilistic prediction about how a discourse being uttered by the speaker is continued. We test the predictions of the two views in EEG using the well-established observation that the conditional probability of a word given a context is negatively correlated with the amplitude of its N400 component.

Keywords: Latent Semantic Analysis, Bayesian Pragmatics, N400, Generative Lexicon, Telicity, Affordances, Context, Predictive Coding

Introduction

It has been widely acknowledged that a preceding discourse can influence the way sentence meaning is composed from lexical meaning. In this paper we want to adjudicate between two competing views of how discourse context affects sentence meaning. A prominent view is that the contextual influence is mainly due to the semantic similarity between parts of the discourse context and the words in the target sentence (e.g., as in semantic priming; Otten & Van Berkum, 2008). It is however highly controversial whether also pragmatic aspects of the discourse context other than the mere resolution of indexicals and anaphors can immediately affect sentence meaning composition. Nieuwland and Van Berkum (2006) have argued that discourse contexts can overturn violations of animacy. For a noun denoting an object that would normally be regarded as inanimate (e.g. *peanut*) the feature of animacy can be introduced if the preceding discourse context specifies a suitable situation (e.g., a romantic story whose protagonist is a peanut). Here predicates that would normally conflict with the noun because they require animacy (*the peanut was in love*) were actually more easily predictable than canonical predicates (*the peanut was salted*), as revealed by an enhanced N400 for canonical predicates as compared to animacy-violating predicates.

To investigate the contrast between a Semantic Similarity account of contextual influence and a Free Pragmatic account that allows for free pragmatic enrichment in sentence meaning composition, we will look into the way subjects make probabilistic predictions on the completion of

a sentence given a preceding discourse. Quantitatively, the semantic similarity can be determined by Latent Semantic Analysis (Landauer & Dumais, 1997), whereas we will use the framework of Bayesian Pragmatics (Frank & Goodman, 2012) to calculate the pragmatic influence – in particular, concerning the rationality of the speaker's intentions in a narrative. As a model of lexical structure we apply the Generative Lexicon approach (Pustejovsky, 1995). In our experimental design we will use the established observation that the conditional probability of a word given a preceding context is negatively correlated with the amplitude of its N400 component measured in EEG.

Background

The general idea of Bayesian Pragmatics is to account for the rational cooperation between speaker and hearer in an act of communication by modelling the hearer's probabilistic expectations about the speaker's communicative intentions by Bayes's Theorem. Bayesian pragmatics has been successfully used e.g. to explain results in a number of behavioral experiments (e.g., Frank & Goodman, 2012). It has so far not been validated in EEG studies.

Bayesian pragmatics offers itself as a model also in what one might call the Predictive Completion Task (PCT) of communication. Predictive coding is widely acknowledged in cognitive science as a general mechanism by which the subject at every point in time generates the most probable prediction of the next event on the basis of ongoing perceptual input and learned statistical regularities (Hohwy, 2013). In a PCT the hearer at every moment in a communicative situation has to generate a probabilistic prediction about how a sentence/discourse being uttered by the speaker will be continued. To get a quantitative grasp of this task, we define $P_T(a|c)$ as the conditional probability of a sentence/discourse being continued with the word a given that the word is preceded by a context c under the assumption that the complete sentence/discourse is true. We will call P_T the truth-guided predictive probability function of the hearer. The problem of the hearer is to estimate the truth-guided predictive probability.

The Semantic Similarity and the Free Pragmatic views provide competing theories of how the hearer accomplishes the Predictive Completion Task. To be able to discern between the two theories and given that the overall frequency of the word a in language use and the syntactic congruency of the word a relative to the preceding context

is known to have a major influence on $P_T(a|c)$, we will in the course of this paper (and in the experiment) presuppose that the frequency of a is held invariant and that the syntactic congruency between a and its context is granted.

According to the first view, $P_T(a|c)$ should be estimated solely on the basis of the semantic similarities between the lexical meaning of the word a and the semantic properties of the preceding context c . The semantic similarity can be quantified by Latent Semantic Analysis, LSA (Landauer & Dumais, 1997).

The Free Pragmatic View in contrast maintains that pragmatic aspects of the discourse directly interact with meaning components retrieved from the lexicon as well as with any further node in the sentence meaning composition tree. It thus challenges a rigorous notion of compositionality, according to which the meaning of a complex expression is determined by the meanings of its syntactic parts and the way the parts are combined (Werning, 2004, 2005a, 2005b, 2012). Pragmatic enrichment is supposed to be “free” because not only lateral modulations of a word or phrase are allowed, e.g. when the meaning of a word is modulated by the meaning of its argument – *cut the cake* (vertical cutting) vs. *cut the grass* (horizontal cutting) – but any, however remote information can in principle modulate the meaning of a linguistic expression at any stage in semantic composition. For example, *cut the grass*, given a lawn-seller situation, might be interpreted as vertical cutting. Accordingly, a situation introduced in a discourse preceding the sentence may result in the modulation of the meanings of words or phrases in the sentence and of the sentence itself (Cosentino, Adornetti, & Ferretti, 2013). These modulations will then influence the intuitive truth-conditions of the sentence. This view, as developed for example by Recanati (2012) amounts to a weakening of the rigorous notion of compositionality by introducing context-dependent semantic flexibility by means of modulation. In the Predictive Completion task Bayesian Pragmatics can be used to quantitatively model the Free Pragmatic account (see Predictions).

Design

The contrasting quantitative predictions of the Semantic Similarity and the Free Pragmatic views can be applied to a previous EEG experiment of ours (Cosentino, Baggio, Kontinen, Garwels, & Werning, 2014; Cosentino, Baggio, Kontinen, & Werning, 2017). To design the experiment we combined a particular idea of Pustejovsky’s (1995) Generative Lexicon approach with Gibson’s (1979) notion of affordances. According to the Generative Lexicon approach, the lexical entry of concrete nouns (e.g. *banana*) contain a “Qualia Structure” which, among others, specifies a so-called Telic component (e.g. *eat*) that is retrieved in sentence meaning composition. This retrieval is typically triggered by verbs like *use* and *enjoy* that take the respective noun as argument. This explains why sentences such as (a) and (c) are typically understood as having the meaning of (b) and, respectively, (d):

- (a) *The child enjoyed the banana.*
- (b) *The child enjoyed eating the banana.*
- (c) *The man used his jackknife for the cake.*
- (d) *The man used his jackknife for cutting the cake.*

In turn, Gibson proposed that many objects come with subject- and situation-dependent affordances. These are dispositional properties (e.g., sit-ability) that relate to actions to be potentially performed on that object (Werning, 2010). We distinguish between *ad-hoc affordances* and *generic affordances*. Generic affordances are affordances of a class of objects that are represented as part of the mental concept of the class (e.g., *chair – sit*). Ad-hoc affordances are affordances that a particular object has for a particular agent in a particular situation (e.g., *this chair – hide under*, for a child in a peekaboo game). In line with Pustejovsky, generic affordances are often stored as *telic components* in the lexicon of nouns and thus in semantic long-term memory.

	+TLex	-TLex
TStd-Ctx	Clare got herself a funnel to perform a little chemistry experiment at home and to this end she put some dye in water. Once she has done so, she uses the funnel to <u>pour</u> water into a container.	Clare got herself a funnel to perform a little chemistry experiment at home and to this end she put some dye in water. Being an unconventional person, she uses the funnel to <u>hang</u> her coat.
TNew-Ctx	Clare has an extra funnel and, after having decided what to do with it, she glues it to the wall leaving the narrow end facing outward. Once she has done so, she uses the funnel to <u>pour</u> water into a container.	Clare has an extra funnel and, after having decided what to do with it, she glues it to the wall leaving the narrow end facing outward. Being an unconventional person, she uses the funnel to <u>hang</u> her coat.

Table 1. Sample stimuli for EEG experiment on context effects on Telic lexical component. The table illustrates a 2x2 design, in which two categories of noun-verb combinations, +TLex and -TLex, are combined with two categories of discourse contexts, Telic Standard Context (TStdCtx) and Telic New Context (TNewCtx). The cue verb is underlined while the corresponding noun preceding the cue verb is in bold. The original stimuli were Italian and here are translated to English.

In our 2x2 experimental design (see Tab. 1) the first variable – TelicLexicalMatch - refers back to Pustejovsky’s notion of a Telic component. The Telic component of the lexical entry specifies the function or the purpose of an object. With regard to the variable TelicLexicalMatch the two conditions, +TLex vs. -TLex varied in whether the cue verb (e.g., *pour* or, respectively, *hang*) expresses the telic component in the lexical entry of a given noun n (e.g., *funnel*). With regard to the second variable – TContext – we varied the discourse context such that in the first condition TStdCtx a standard context preceded the target sentence, whereas in the second condition TNewCtx the preceding discourse context introduced a new telic role as an ad-hoc affordance for the object denoted by the noun, facilitating an action expressible by the -TLex verb (*hang*).

Predictions

Semantic Similarity View: With word frequency held constant and syntactic congruency granted, the Semantic Similarity view now entails that the only predictor for a verb given its preceding context is the semantic similarity of the former to the latter. Given that we leave the corresponding noun n (*funnel*), which precedes the cue verb v_i (*pour/hang*), the same in all four conditions, the semantic similarity is a cumulative (i.e. in both arguments strictly monotonously increasing) function f^+ of the semantic similarity $S(v_i, n)$ between the verb and the noun and the semantic similarity $S(v_i, c_j)$ between the verb and the preceding context c_j excluding the noun. According to this view the truth-guided predictive probability $P_{T,n}(v_i|c_j)$ of the verb v_i following the noun n given the context c_j should hence be estimated by the hearer as follows (see Eqn. (4) in Table 2).

Semantic Similarity view:

$$P_{T,n}(v_i|c_j) = f^+(S(v_i, n), S(v_i, c_j)).$$

The experimental settings were chosen such that the semantic similarity – determined by LSA –between the cue verb and the context excluding the noun was invariant across all four conditions (see Eqn. (6) in Table 2). Therefore, $P_{T,n}(v_i|c_j)$ should depend solely on whether the verb expresses the telic lexical component of the noun, i.e. belongs to +TLex, and correspondingly has a high semantic similarity to the noun – determined again by LSA – as opposed to a verb belonging to –TLex with a low semantic similarity to the noun (see (5)). We can now immediately mathematically derive the four comparative predictions regarding $P_{T,n}(v_i|c_j)$ for $i \in \{+TLex, -TLex\}$ and $j \in \{TStdCtx, TNewCtx\}$. These predictions are captured by the formulae (11), (15), (16) and (17) shown in Table 2 together with their derivations.

Bayesian Pragmatics: If the Free Pragmatic as opposed to the Similarity View is true the hearer will use a different strategy to estimate $P_T(a|c)$. From a pragmatic point of view, narratives are goal-directed discourses. In our examples the speaker has the goal to attribute an action to the narrative’s protagonist which s/he performs on a given object: In the above examples, performing the action of pouring or, respectively, hanging on the funnel. In the narrative, the speaker embeds this action in a situation which he may introduce by a discourse context that precedes the description of the action. The speaker in other words has to choose a preceding context to let this action appear rational. To describe this choice situation quantitatively we can define the rationality-guided conditional probability $P_{R,n}(c_j|v_i)$ as the probability of the speaker to choose – under the assumption of narrative rationality – a context c_j given that he aims at attributing to the protagonist the action denoted by v_i to be performed on the object denoted by n . Using Bayes’s Theorem $P_{R,n}(c_j|v_i)$ can be transformed to allow the hearer, in the PCT, to estimate the truth-guided

predictive probability by equating it to the rationality-guided probability of the speaker (see Eqns. (1) and (3) in Table 2):

Free Pragmatic view:

$$P_{T,n}(v_i|c_j) = P_{R,n}(v_i|c_j),$$

where, according to Bayes’s Theorem,

$$P_{R,n}(v_i|c_j) = K(c_j) \cdot P_{R,n}(c_j|v_i)P_{R,n}(v_i).$$

Here $K(c_j) = (\sum_{v \in V} P_{R,n}(c_j|v)P_{R,n}(v))^{-1}$ is a normalizing factor and V is the set all (syntactically congruent) verbs. We may assume that the rationality-guided prior probability $P_{R,n}(v_i)$ of an action expressed by the verb v_i being performed on the object denoted by the noun n is fully determined by the semantic similarity between the verb and the noun as computed by LSA (see Eqn. (2)):

$$P_{R,n}(v_i) = S(v_i, n).$$

For, the lexical entry of a concrete noun can be assumed to reflect the semantic memory of the learned statistical regularities between objects denoted by the noun and actions rationally performed on them. This assumption is part and parcel also of the idea that concrete nouns have telic lexical components in the sense of Pustejovsky’s (Pustejovsky, 1995) Generative Lexicon. If the verb corresponds to the telic lexical entry of the nouns, the semantic similarity between them should hence be high. In the experimental settings we can implement a comparative variation of the rationality-guided probability $P_{R,n}(c_j|v_i)$ of the speaker to choose a context c_j given that he aims at attributing to the protagonist the action denoted by v_i to be performed on the object denoted by n as captured by the inequalities (7)-(10) in Table 2. This immediately allows us to generate comparative predictions about the conditional predictive probability of the hearer as a consequence of the Bayesian interpretation of the Free Pragmatic view. Making the idealizing assumption that the normalizing factor, which is unknown not only to us, but also to the hearer, is the same for all contexts and in particular, $K(c_1) = K(c_2) = K$, we can generate predictions not only for comparisons within the same context, but also across contexts. The so attained predictions of the Free Pragmatic view are captured in formulae (11), (12), (13) and (14) in Table 2, shown together with their mathematical derivation history.

Correlation with N400 amplitude: To test the predictions, we exploited an empirically already well established relationship between the probability of a word given a preceding context and the amplitude of the N400 component of the event-related potential measured on the onset of the word in EEG. Granted that the cue word is syntactically congruent with its context (i.e., no syntactic violation) and that the frequency as well as length of the word are invariant, the truth-guided predictive probability $P_T(a|c)$ of the word a given the preceding context c is negatively correlated with the amplitude of the word’s N400. Support for the negative correlation between the truth-guided predictive probability of a word given a

preceding context and the word's N400 comes from multiple sources of evidence. Most importantly, Cloze probability is a strong predictor of the amplitude of the N400 component. Cloze probability values are obtained by asking participants in a Cloze task to complete an incomplete sentence with the word they consider to be the most likely completion. Kutas & Hillyard (1984) found that the amplitude of the N400 component measured on the target word has a nearly inverse linear relationship with its Cloze probability, that is, relative to more expected words, the N400 amplitude increases as the expectancy of a word in context decreases. DeLong, Urbach, & Kutas, 2005 confirmed that the preceding words in a sentence are used by readers to estimate relative likelihoods for upcoming words and the differences in the likelihood of the target word are reflected in differences regarding the N400 component. This effect has been observed not only in single sentences (Van Petten, Coulson, Rubin, Plante, & Parks, 1999), but also for short texts (Otten & Van Berkum, 2007; Van Berkum, Brown, Zwitserlood, Kooijman, & Hagoort, 2005). Even more subtle differences in the semantic relatedness between the words in a sentence are found to influence the conditional probability of an upcoming word and affect the amplitude of the N400 measured on that word. For instance, in the sentence *The girl was writing letters when her friend spilled coffee on the tablecloth/paper* the semantically unrelated word *tablecloth* elicits a larger N400 than the semantically related word *paper* (Baggio, van Lambalgen, & Hagoort, 2008). The integration of world knowledge during the interpretation of sentences such as *The Dutch trains are yellow/white/sour and very crowded* (the target words are underlined) also modulates the amplitude of the N400 component. This reflects the role of the (Dutch) subjects' knowledge that Dutch trains are typically yellow for establishing the conditional probability of the target word (Hagoort, Hald, Bastiaansen, & Petersson, 2004). This stresses the point that the predictive conditional probability, in fact, is guided by expectations regarding the truth of the continuing sentence. Additional evidence of a negative correlation between the predictive conditional probability of a word and the amplitude of its N400 component is provided by a study previously conducted in our laboratory. In a sentence-picture verification study on scalar implicatures, logical and pragmatic responders provide different truth-value judgements to under-informative sentences (e.g., *Some As are B*, when it is known that all As are B). Whereas logical responders evaluate these sentences as true, pragmatic responders reject them as false. These divergent responses correlate with significant differences regarding the N400 and can be explained on the basis of expected probabilities of words relative to truth presumed by the subject (Spychalska, Kontinen, & Werning, 2016). Similar findings have been reported in a study about bare numerals (Spychalska, Kontinen, Noveck, Roesch, & Werning, 2015).

Experiment

Method: Twenty-two right-handed native speakers of Italian (13 males; mean age = 29.2 years) were presented with a total of 160 stories in a 2x2 design (see Table 1). The ERPs recording was time-locked to the onset of the cue words, which were always verbs occurring in the midst of the sentence and matched on word length, number of syllables and mean word frequency. The preceding contexts were pair-wise matched for number of words. The experimental stimuli were translated into English and underwent Latent Semantic Analysis to check for the semantic similarity values between the cue verbs and the preceding nouns or, respectively, between the cue verbs and the preceding contexts. Whereas the difference between $S(+TLex, n)$ and $S(-TLex, n)$ was significant ($t(39)=5,449, p<.001$), there was no significant difference between the cue verbs and the preceding contexts across all experimental conditions. Using average amplitude per condition across all EEG electrodes, a $2(\text{Context: TStdCtx vs. TNewCtx}) \times 2(\text{TelicLexicalMatch: +TLex vs. -TLex})$ repeated measures analysis of variance (ANOVA) was performed in the time window between 400 and 500 ms after critical word onset. A follow-up ANOVA was performed which involved specifically a predetermined region over centro-parietal sites at which the N400 is maximal. In this case, a $2(\text{Context: TStdCtx vs. TNewCtx}) \times 2(\text{TelicLexicalMatch: +TLex vs. -TLex}) \times 7(\text{Electrodes: CP1, CP2, CPz, Pz, P1, P2, POz})$ ANOVA was conducted. Bonferroni-adjusted planned comparisons were performed to decompose the effect of trial type in this region.




Figure 1. Crossing over regarding the N400 component. Bars show the average amplitude of the N400 for the four conditions. The numbers in brackets correspond to the inequalities as predicted by the Free Pragmatic View (see Table 2, (11), (12) and (13)). Note that the fact that the difference between the +TLex and the -TLex verb in TNewCtx is not significant is also consistent with Free Pragmatic View (14).

Results: Given the standard context TStdCtx, the N400 for -TLex ($M(\text{TStdCtx}, -\text{TLex}) = -1.67 \mu\text{V}$) is significantly enhanced compared to +TLex ($M(\text{TStdCtx}, +\text{TLex}) = -.64 \mu\text{V}$, $t(20)=3.069, p=.006, \text{CI } 1.03 \pm .70$). Relative to the standard context TStdCtx, the TNewCtx significantly enhances the N400 component for +TLex ($M(\text{TStdCtx}, +\text{TLex}) = -1.15 \mu\text{V}$, $t(20)=2.276, p=.034, \text{CI } .51 \pm .47$), whereas it significantly reduces the N400 component for -TLex ($M(\text{TStdCtx}, -\text{TLex}) = -.88 \mu\text{V}$, $t(20) = -2.745$,

$p=.012$, CI $-.79 \pm .60$. Finally, given a preceding context TNewCtx, the mean amplitude of the N400 component measured on +TLex was not significantly different compared to that measured on -TLex ($t(20)=.964$, $p=.34$) See Fig. 1. The follow-up ANOVA of the predetermined N400 region showed a significant Context \times TelicLexicalMatch interaction, $F(1, 20) = 11.267$, $p<.005$. There was no interaction with electrodes in this region.

Discussion

In order to test the predictions made by the Free Pragmatic View and the Semantic Similarity View, we rely on the empirically well-founded observation that the truth-guided conditional probability $P_T(a|c)$ of the word a given the preceding context c (granted that no syntactic violation is involved and that features such as the frequency and length of the word are constant) is negatively correlated with the amplitude of the N400 component measured on that word succeeding the context.

In our experimental settings, we determined the semantic similarity $S_n(v_i)$ between the meaning of the verb v_i and the preceding noun n using LSA. This value gives us the prior probability of $P_{R,n}(+TLex)$ and, respectively, $P_{R,n}(-TLex)$ with the former (*pour* corresponds to the telic lexical component of *funnel*) being higher than the latter (*hang* does not correspond to the lexical component of *funnel*), as reported in equation (5). We also determined the semantic similarity values between the verb v_i and the discourse context c_j excluding the noun and kept these values constant across all experimental conditions (see (6)). With regard to what is relevant for the Semantic Similarity View, the experimental conditions differ only in the semantic similarity values between the verb v_i and the preceding noun n . The Semantic Similarity view entails that these values are the only predictor of differences in the truth-guided conditional probability of the verb given the preceding context and the noun (see (11), (15), (16) and (17)) and, hence, they are the only predictor of differences regarding the amplitude of the N400 component.

The Free Pragmatic view focuses instead on the differences in the rationality-guided probability of the speaker choosing a certain context given that he attributes to the narrative subject the aim of performing a certain action with an object. In the Free Pragmatic framework, the rationality-guided probability that the speaker chooses a standard context TStdCtx (e.g., *funnel* in a chemistry experiment) given that he attributes to the narrative subject the aim of performing the action denoted by the +TLex verb (*pour*) with the object denoted by the noun n (*funnel*) is higher than that of choosing this context given the attributed aim of performing with that object the action denoted by the -TLex verb (*hang*) (see inequality (7)). Furthermore, it is more rational for the speaker to choose a context TNewCtx, which introduces a new ad-hoc affordance for the object (*funnel* glued to the wall), compared to choose the standard context TStdCtx, given that he attributes to the narrative subject the aim of performing the action denoted by the -

TLex verb (*hang*) with the object denoted by n (see (8)). As captured by inequality (9), the rationality-guided probability of the speaker choosing the standard context TStdCtx compared to TNewCtx is higher given that he attributes to the narrative subject the aim of performing the action denoted by the +TLex verb with the object denoted by n . Finally, as expressed by (10), the rationality-guided probability that the speaker chooses TNewCtx given that he attributes to the narrative subject the aim of performing the action denoted by the -TLex verb with the object denoted by n is higher than the rationality-guided probability of choosing this context given the attribution to the narrative subject of the aim of performing the action denoted by the +TLex verb with the object denoted by n .

Given that the Free Pragmatic view estimates the truth-guided predictive probability of a word by equating it with its rationality-guided probability (see (1)), the Free Pragmatic view not only predicts (11), in line with the Semantic Similarity View, but, in contrast to the Semantic Similarity View, predicts a crossing-over regarding the N400 component, as expressed by the inequalities (12), (13). With regard to the comparison expressed in (14) the Free Pragmatic View does not make an unambiguous prediction. For, a greater/smaller comparison of the values of the product $P_{R,n}(v_i|c_j) = K(c_j) \cdot P_{R,n}(c_j|v_i)P_{R,n}(v_i)$ depends not only on the numerical value of the prior probability $P_{R,n}(v_i)$, which is given through the equation $P_{R,n}(v_i) = S(v_i, n)$, but also on the unknown numerical value of and not just the inequalities between the likelihoods $P_{R,n}(c_j|v_i)$.

Given the negative correlation between the truth-guided conditional probability of a word given a preceding context and the amplitude of its N400 component, the results of our EEG study can be used to directly evaluate the different predictions of the two views. In our experiment, we found that, if preceded by a standard discourse context TStdCtx, a -TLex verb incongruent with the noun's telic component (*funnel-hang*) elicited an enhanced N400 compared to a +TLex verb congruent with the telic component (*funnel-pour*) (confirming (11)). However, given a discourse context TNewCtx, in which a new function for the object is introduced as an ad-hoc affordance, we observed a crossing-over regarding the direction of the N400 effect: Comparing TNewCtx with TStdCtx, first, the N400 for the -TLex verb was significantly smaller in TNewCtx than in TStdCtx (disconfirming (15) and confirming (12)). Second, the N400 for the +TLex verb was significantly greater in TNewCtx than in TStdCtx (disconfirming (16) and confirming (13)). Finally, given a preceding context TNewCtx, the N400 measured on the +TLex verb was not significantly different compared to that measured on the -TLex verb (see Fig. 1). This result is not decisive between the two views (neither confirming nor disconfirming (17) and being consistent with (14)).

The reported differences regarding the N400 component are best explained by the assumption that hearer accomplishes the Predictive Completion Task as envisaged

by the Free Pragmatic View, namely by estimating $P_{T,n}(v_i|c_j)$ through equating it with $P_{R,n}(v_i|c_j)$ and applying Bayes's Theorem to it. Indeed, the crossing-over regarding the N400 cannot be explained solely in terms of

the differences in the semantic similarity values between the target verb and the preceding noun, as assumed by the Semantic Similarity account.

	Free Pragmatic View	Semantic Similarity View
Th. Asspts.	(1) $P_{T,n}(v_i c_j) = P_{R,n}(v_i c_j)$ (2) $P_{R,n}(v_i) = S(v_i, n)$ (3) $P_{R,n}(v_i c_j) = K \cdot P_{R,n}(c_j v_i)P_{R,n}(v_i)$	(4) $P_{T,n}(v_i c_j) = f^+(S(v_i, n), S(v_i, c_j))$
Exp. Settings	(5) $S(+T\text{Lex}, n) > S(-T\text{Lex}, n)$ (6) $S(+T\text{Lex}, T\text{StdCtx}) = S(-T\text{Lex}, T\text{StdCtx}) = S(+T\text{Lex}, T\text{NewCtx}) = S(-T\text{Lex}, T\text{NewCtx})$ (7) $P_{R,n}(T\text{StdCtx} +T\text{Lex}) > P_{R,n}(T\text{StdCtx} -T\text{Lex})$ (8) $P_{R,n}(T\text{NewCtx} -T\text{Lex}) > P_{R,n}(T\text{StdCtx} -T\text{Lex})$ (9) $P_{R,n}(T\text{StdCtx} +T\text{Lex}) > P_{R,n}(T\text{NewCtx} +T\text{Lex})$ (10) $P_{R,n}(T\text{NewCtx} -T\text{Lex}) > P_{R,n}(T\text{NewCtx} +T\text{Lex})$	
Predictions	(11) $P_{T,n}(+T\text{Lex} T\text{StdCtx}) > P_{T,n}(-T\text{Lex} T\text{StdCtx})$ (from(1),(2), (3), (5), (7)) (12) $P_{T,n}(-T\text{Lex} T\text{NewCtx}) > P_{T,n}(-T\text{Lex} T\text{StdCtx})$ (from(1),(2), (3), (8), (10)) (13) $P_{T,n}(+T\text{Lex} T\text{StdCtx}) > P_{T,n}(+T\text{Lex} T\text{NewCtx})$ (from(1),(2), (3), (9)) (14) $P_{T,n}(-T\text{Lex} T\text{NewCtx}) \geq P_{T,n}(+T\text{Lex} T\text{NewCtx})$ (from(1),(2), (3), (5), (10))	(11) $P_{T,n}(+T\text{Lex} T\text{StdCtx}) > P_{T,n}(-T\text{Lex} T\text{StdCtx})$ (from (4), (5), (6)) (15) $P_{T,n}(-T\text{Lex} T\text{StdCtx}) = P_{T,n}(-T\text{Lex} T\text{NewCtx})$ (from (4), (6)) (16) $P_{T,n}(+T\text{Lex} T\text{StdCtx}) = P_{T,n}(+T\text{Lex} T\text{NewCtx})$ (from (4), (6)) (17) $P_{T,n}(-T\text{Lex} T\text{NewCtx}) < P_{T,n}(+T\text{Lex} T\text{NewCtx})$ (from (4), (5), (6))

Table 2. Overview of the different theoretical assumptions and predictions of the Free Pragmatic View and the Semantic Similarity View given our experimental settings.

References

Baggio, G., van Lambalgen, M., & Hagoort, P. (2008). Computing and recomputing discourse models: An ERP study. *Journal of Memory and Language*, 59, 36–53.

Cosentino, E., Adornetti, I., & Ferretti, F. (2013). Processing Narrative Coherence: Towards a Top-Down Model of Discourse. *Proceedings of the 4th Workshop on Computational Models of Narrative*, 32, 61–75.

Cosentino, E., Baggio, G., Kontinen, J., Garwels, T., & Werning, M. (2014). Lexicon in action: N400 contextual effect on affordances and telicity. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 2079–2084). Austin, TX: Cognitive Science Society.

Cosentino, E., Baggio, G., Kontinen, J., & Werning, M. (2017). The time-course of sentence meaning composition. N400 effects of the interaction between context-induced and lexically stored affordances. *Frontiers in Psychology*. doi:10.3389/fpsyg.2017.00813

DeLong, K. A., Urbach, T. P., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience*, 8, 1117–1121.

Frank, M. C., & Goodman, N. D. (2012). Predicting Pragmatic Reasoning in Language Games. *Science*, 336, 998–998.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston, MA: Houghton Mifflin.

Hagoort, P., Hald, L., Bastiaansen, M., & Petersson, K. M. (2004). Integration of word meaning and world knowledge in language comprehension. *Science (New York, N.Y.)*, 304, 438–41.

Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press.

Kutas, M., & Hillyard, S. (1984). Brain potentials during reading reflect word expectancy and semantic association. *Nature*, 307, 161–163.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211–240.

Nieuwland, M. S., & Van Berkum, J. J. A. (2006). When peanuts fall in love: N400 evidence for the power of discourse. *Journal of Cognitive Neuroscience*, 18, 1098–1111.

Otten, M., & Van Berkum, J. J. a. (2008). Discourse-Based Word Anticipation During Language Processing: Prediction or Priming? *Discourse Processes*, 45, 464–496.

Otten, M., & Van Berkum, J. J. A. (2007). What makes a discourse constraining? Comparing the effects of discourse message and scenario fit on the discourse-dependent N400 effect. *Brain Research*, 1153, 166–177.

Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, MA: MIT Press.

Recanati, F. (2012). Compositionality, Semantic Flexibility, and Context-Dependence. In M. Werning, W. Hinzen, & E. Machery (Eds.), *Oxford Handbook of Compositionality* (pp. 175–191). Oxford: Oxford University Press.

Spychalska, M., Kontinen, J., Noveck, I., Roesch, L., & Werning, M. (2015). Exploring the processing costs of the exactly and at least readings of bare numerals with event-related brain potentials. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 2260–2265). Austin, TX: Cognitive Science Society.

Spychalska, M., Kontinen, J., & Werning, M. (2016). Investigating scalar implicatures in a truth-value judgement task: evidence from event-related brain potentials. *Language, Cognition and Neuroscience*, 31, 817–840.

Van Berkum, J. J. A., Brown, C. M., Zwitserlood, P., Kooijman, V., & Hagoort, P. (2005). Anticipating Upcoming Words in Discourse: Evidence From ERPs and Reading Times. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 443–467.

Van Petten, C., Coulson, S., Rubin, S., Plante, E., & Parks, M. (1999). Time course of word identification and semantic integration in spoken language. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 394–417.

Werning, M. (2004). Compositionality, Context, Categories and the Indeterminacy of Translation. *Erkenntnis*, 60, 145–178.

Werning, M. (2005a). Right and wrong reasons for compositionality. In M. Werning, E. Machery, & G. Schurz (Eds.), *The Compositionality of Meaning and Content* (Vol. I, pp. 285–309). Frankfurt: Ontos Verlag.

Werning, M. (2005b). The Temporal Dimension of Thought: Cortical Foundations of Predicative Representation. *Synthese*, 146.

Werning, M. (2010). Complex First? On the Evolutionary and Developmental Priority of Semantically Thick Words. *Philosophy of Science*.

Werning, M. (2012). Non-symbolic compositional representation and its neuronal foundation: Towards an emulative semantics. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 633–654). Oxford University Press.