Trying to make it work: Compositional effects in the processing of compound “nonwords”

Fritz Günther¹ & Marco Marelli^{1,2}

¹University of Milano–Bicocca, Milan, Italy

²NeuroMI, Milan Center for Neuroscience, Milan, Italy

Author Note

Note: This is the author’s preprint version of the article (date: 05. 12. 2019). The final article is accepted for publication in the *Quarterly Journal of Experimental Psychology*.

This work was supported by a Research Fellowship (no. 392225719) from the German Research Foundation (DFG), awarded to Fritz Günther, and by grant 2016-ATE-0516 – *Big data in psicolinguistica* from the Department of Psychology, University of Milano-Bicocca, awarded to Marco Marelli. We thank Barbara Kaup for her valuable help in the data collection. The data and analysis scripts for this study are available at <https://doi.org/10.6084/m9.figshare.7867772.v1>.

Abstract

Speakers of languages with synchronically productive compounding systems, such as English, are likely to encounter new compounds on a daily basis. These can only be useful for communication if speakers are able to rapidly compose their meanings. However, while compositional meanings can be obtained for some novel compounds such as *bridgemill*, this is far harder for others such as *radiosauce*; accordingly, processing speed should be affected by the ease of such a compositional process. To rigorously test this hypothesis, we employed a fully implemented computational model based on distributional semantics to quantitatively measure the degree of semantic compositionality of novel compounds. In two large-scale studies, we collected timed sensibility judgments and lexical decisions for hundreds of morphologically-structured nonwords in English. Response times were predicted by the constituents’ semantic contribution to the compositional process, with slower rejections for more compositional nonwords. We found no indication of a difference in these compositional effects between the tasks, suggesting that speakers automatically engage in a compositional process whenever they encounter morphologically-structured stimuli, even when it is not required by the task at hand. Such compositional effects in the processing of novel compounds have important implications for studies that employ such stimuli as filler material or “nonwords”, since response times for these items can differ greatly depending on their compositionality.

Keywords: Compound words; Distributional semantics; Compositionality; Morphological Processing; Novel Compounds

Trying to make it work: Compositional effects in the processing of compound “nonwords”

The human capacity to mentally combine existing concepts into new ones is fundamental to our creative abilities and allows us to devise new ideas and objects (Thagard, 1984). This is also mirrored in our language, where the use of compounds such as *airport* is derived from combinations of familiar words into single novel ones. Across languages, compounds are a central tool for lexical productivity and vocabulary extension (Dhar & van der Plas, 2019; Libben, 2014) – consider recent examples such as *smartwatch*, *vaporwave*, or *clickbait*. In fact, compounding has been argued to be the fundamental mechanism from which more complex forms of linguistic composition have developed (Jackendoff, 2002).

However, the successful use of such complex expressions demands certain abilities from the speakers of a language. Especially in languages with synchronically productive compounding systems (such as English and, to an even larger degree, Finnish or German), speakers can be expected to encounter novel complex expressions they have never experienced before on a daily basis (Libben, 2014). Hence, for these expressions not to disrupt communication, speakers have to be able to understand their intended meaning on the fly, without effort (Downing, 1977). This requires a fast and powerful meaning-composition system that can immediately combine the familiar constituent elements into a single novel mental representation and provide possible interpretations for it.

Empirical studies have demonstrated that speakers indeed routinely engage in such meaning-composition operations when encountering synchronically productive complex forms (e.g. El-Bialy, Gagné, & Spalding, 2013; Gagné & Spalding, 2009; Libben, 2014; Marelli & Baroni, 2015), which set in very early in processing and are initiated automatically (Amenta, Marelli, & Crepaldi, 2015). Recent studies provide strong evidence for the automaticity of such a meaning-composition process, by demonstrating that it is even applied to pseudo-compounds such as *carpet*, where it leads to incorrect results and

associated repair costs (Chamberlain, Gagné, Spalding, & Lõo, 2019; Gagné, Spalding, Nisbet, & Armstrong, 2018). In line with this evidence, Günther and Marelli (2018) observed – in a lexical decision task which does not in principle require semantic access – that existing compounds (such as *airport* or *ladybird*) are processed faster the easier a *compositional* meaning can be constructed from their constituents, and hence the more the constituent meanings contribute to the meaning of a newly-composed concept (Gagné & Spalding, 2009). These findings frame semantic effects in compound processing that are usually investigated under the notion of *semantic transparency* (Libben, Gibson, Yoon, & Sandra, 2003; Sandra, 1990) in a compositional perspective (for a detailed discussion on the relation between semantic transparency and compositionality, see Günther & Marelli, 2018).

As outlined above, the theoretical explanation for such findings is based on the argument that we need to be able to understand novel compounds rapidly when we encounter them. Therefore, it makes sense to automatically initiate meaning composition whenever a compound is encountered, irrespective of its lexical status (El-Bialy et al., 2013; Libben, 2014). Obviously, this argument directly implies that an automatic meaning-composition process should also affect the processing of novel compounds such as *radiosauce* or *bridgemill*. However, in the context of the lexical-decision experiments often employed in studies on morphological processing, such novel forms are often considered only as nonword fillers (Brysbaert et al., 2011) or morphologically-valid but non-existing items (e.g. Crepaldi, Rastle, Davis, & Lupker, 2013; Taft & Forster, 1976). This perspective however neglects the fact that these words are not a uniform class of non-existing lexical items, but can instead differ drastically in their perceived meaningfulness (Costello & Keane, 2000; Graves, Binder, & Seidenberg, 2013; Günther & Marelli, 2016; Longtin & Meunier, 2005; Meunier & Longtin, 2007): While it is intuitively easy to obtain some sensible interpretation for *bridgemill*, it is far harder for *radiosauce*. The conceptualization of these concatenations as a homogeneous class of equally

meaningless “nonwords” therefore seems counter-intuitive, especially considering that they are the very focus of the whole research field concerned with conceptual combination, that sees them as novel word candidates (e.g. Medin & Shoben, 1988; Murphy, 1988). In this line of research, it has been shown that meaning aspects of these words – such as the relational structure between the constituents – indeed affect processing times (Gagné, 2001; Gagné & Shoben, 1997; Gagné & Spalding, 2009). Thus, these words should be considered as valuable test cases for psychological theories of semantic representation and processing (Gagné & Spalding, 2006), rather than being reduced to mere nonwords.

In a previous study, Günther and Marelli (2016) demonstrated that the contributions of the constituent meanings to the combined meaning predict explicit sensibility ratings for compound words, including novel compounds (collected by Graves et al., 2013). However, it has been demonstrated that such off-line explicit ratings for complex words are affected by different semantic variables than on-line processing times, especially when only the latter ones are obtained under time pressure (Günther, Marelli, & Bölte, in press; Marelli & Baroni, 2015). The present study therefore aims at systematically investigating how these constituent contributions affect the processing times of novel compounds. Note that, in the context of tasks such as lexical decision, the correct response is still to reject these words as non-existing lexical items. We thus expect that higher constituent contributions, associated with an easier compositional process and more meaningful combined representations, should lead to larger semantic activation that in turn inhibits the rejection of the respective item (Seidenberg & McClelland, 1989).

Critically, new computational approaches implemented in the framework of distributional semantics allow us to measure the contribution of either constituent to a compositional novel-word meanings in an objective and quantitative fashion (Marelli, Gagné, & Spalding, 2017). Distributional semantic models represent word meanings as high-dimensional numerical vectors, derived from large collections of natural text (Landauer & Dumais, 1997; Lund & Burgess, 1996; Turney & Pantel, 2010; Westbury, 2016).

Importantly for our purposes, these systems have been established as powerful theories and models of human semantic memory (Günther, Rinaldi, & Marelli, 2019; M. N. Jones, Willits, & Dennis, 2015). In this distributional framework, the *Compounding as Abstract Operation in Semantic Space* model (CAOSS; Marelli et al., 2017) is designed to compute the compositional meaning of compounds from its constituents (see the Methods section). The contribution of either constituent to the compositional process is then measured as their respective similarity to this compositionally-obtained representation (Günther & Marelli, 2016, 2018). The CAOSS model has been successfully established as predicting processing times in existing compounds (with higher constituent contributions resulting in faster processing; Günther & Marelli, 2018; Günther et al., in press) and relational effects in novel compounds (Marelli et al., 2017; see Gagné, 2001; Gagné & Shoben, 1997).

The application of this computational framework enables us to investigate our hypothesis that differences in semantic compositionality (more specifically, the constituent contributions to the compositional meaning) predict processing times for novel compounds in a quantitatively precise and systematic large-scale approach. We predict that rejections of such items will be slower the more compositional, and hence closer to reaching meaningfulness, the items are. Thus, we do not expect that the compositional process will necessarily yield a result that speakers judge as meaningful or sensible (especially in the absence of any context). However, we expect that some novel compounds will be closer to the decision threshold (and hence harder to reject) than others.

We test this hypothesis in two large experiments, employing two different tasks: In Experiment 1, we first employ a timed sensibility judgment task where participants explicitly have to judge whether existing and non-existing compounds have a sensible interpretation (see, for example, Connell & Lynott, 2013; Estes, 2003; Gagné, 2001; Tagalakakis & Keane, 2006). This task necessarily requires participants to try to derive a meaning for the words presented to them, and an evaluation of the success of this process or the result obtained from it. This allows us to establish a baseline where we estimate the

influence of a meaning-composition process in an experimental paradigm where semantic access is necessary in order to perform the task. In Experiment 2, we then test the very same item material in a lexical decision task, where semantic processing is neither required nor solicited: In this task, participants just have to judge whether the presented string of letters is a word in their language (i.e., if they have encountered it before), which could in principle be done just by considering the form level, without any consideration of meaning.. If we observe compositional effects in lexical decision as well as in timed sensibility judgments, this will suggest that an automatic meaning-composition process routinely and automatically affects the processing of novel compounds, complementing the results by Günther and Marelli (2018) for existing compounds.

Experiment 1

In Experiment 1, we employed a timed sensibility judgment task (Gagné, 2001; see also Connell & Lynott, 2013; Estes, 2003), in which participants are presented with existing and non-existing compounds, and have to decide whether they have a sensible interpretation. This task encourages participants to try to build (i.e., compose) a semantic representation for the items presented to them and to examine the plausibility of this interpretation.

Method

Participants. 270 native English speaking participants were recruited in an online crowdsourcing study using the *Amazon Mechanical Turk* platform (see Buhrmester, Kwang, & Gosling, 2011). Of course, not being tested in a formal lab setting can lead participants to ignore the actual task, and answer randomly to finish the experiment as fast as possible. Therefore, we applied a false-negative and a false-positive exclusion criterion. We excluded 97 participants who judged less than 15 of 20 catch items (see below) as non-sensible, and an additional 28 participants who judged less than 80 of the 100 most frequent existing compounds as sensible. Of the remaining 145 participants ($M_{age} = 38.5$ years, $SD_{age} = 10.7$ years), 74 identified themselves as male and 70 as female (one unspecified). Participants

received \$2 for participating in the study. This study (Experiment 1 and Experiment 2) was approved by the Departmental Commission for Research Evaluation of the Department of Psychology (CRIP), University of Milano-Bicocca (project RM-2018-141).

Semantic Model. As semantic representations, we considered the best-performing word-embeddings model reported in a large-scale evaluation by Baroni, Dinu, and Kruszewski (2014)¹. The model was trained on an English ~ 2.8 billion word corpus using the *cbow* algorithm of the *word2vec* toolkit (Mikolov, Chen, Corrado, & Dean, 2013).

Using these vectors, we trained the CAOSS model (Marelli et al., 2017) on a set of 5,988 existing compounds (3,351 closed-form and 2,637 hyphenated compounds), individuated using the source corpus *TreeTagger* (Schmid, 1995) annotations. The CAOSS model aims at predicting a compound vector c from its constituent vectors u (modifier, left-hand constituent) and v (head, right-hand constituent) as

$$c = M \cdot u + H \cdot v \tag{1}$$

, and estimates a single set of weight matrices M and H so that the compound vectors in the training set are optimally approximated in a least-squares regression approach. The weight matrices M and H , once obtained via training, can be applied to compute the (compositional) compound meanings of novel compounds (see Figure 1 for an illustration).

Material. We first generated a list of novel closed-form compounds by re-combining all constituents in our training set with a frequency of at least 1,000 in our source corpus (1,282 modifiers and 1,540 heads), and computed the (compositional) compound meaning vectors for all resulting items. All existing compounds observed in the source corpus (including those found in open form, such as *tree house*) were then removed

¹ context window size of 5 words, 400-dimensional vectors, negative sampling with $k = 10$, subsampling with $t = 1e^{-5}$

Figure 1. Illustration of the compositional system: A single set of weight matrices M and H is trained to predict existing compound vector representations from their constituent vector representations (*Training*). These weight matrices are then applied to obtain compositional meanings for novel compounds on the basis of their constituent vectors (*Composition*).

from the list. For all novel compounds, we computed *modifier composition* and *head composition* as the cosine similarity between the (compositional) compound vector and the modifier vector or head vector, respectively. In cognitive terms, these composition values measure the constituents’ individual contributions in the meaning-composition process (Günther & Marelli, 2016, 2018). Thus, they provide a graded measure of the extent to which the original constituent meaning, when used as a single word, is maintained in the combined concept. Item examples for different grades of modifier and head composition are provided in Table 1. Note that these contributions can be high or low for both constituents independently: Neither constituent contributes much to *nickgiving*, both contribute highly to *sailship*, only the modifier really contributes to *consulstuff*, and only the head to *meancotton*. We then selected as item material a subset of 1,499 novel compounds whose modifier composition and head composition values were (a) both uniformly distributed over the whole range of possible values and (b) not correlated ($r < .001$), thus systematically and orthogonally manipulating modifier and head composition as independent variables. In this subset, no modifier or head appeared in more than 10 different compounds.

Table 1

Item examples for novel compounds with varying degrees of modifier and head composition.

		head composition		
		low	medium	high
mod. comp.	low	<i>nickgiving, wormgrade</i>	<i>factcounty, tompea</i>	<i>standquake, meancotton</i>
	medium	<i>heirdrop, knifegrade</i>	<i>deckwave, prayertie</i>	<i>factmedal, weathersail</i>
	high	<i>waistset, consulstuff</i>	<i>farmroom, boattank</i>	<i>sailship, godlore</i>

Additionally, a subset of 1,499 existing compounds of the dataset by Günther and Marelli (2018) was included in the item material, to ensure that participants also would be presented with sensible compounds. These items were collected by identifying all compounds from the English Lexicon Project (Balota et al., 2007) using the annotations of the CELEX database (Baayen, Piepenbrock, & Gulikers, 1995), and selecting those with a frequency higher than 50 in our source corpus. The resulting 2,998 items were randomly divided into six sets (five containing 250 novel – *boatday, armflour* – and 250 existing – *rulebook, pineapple* – compounds, one containing one less), and each participant was presented with one of these sets. As “catch compounds” (items aimed at checking whether participants actually performed sensibility judgments as instructed), each set additionally included 20 novel compounds (such as *circleshoe* and *guylawn*) that received extremely low sensibility ratings in a study by Graves et al. (2013).

Word frequencies for modifiers and heads (when used as free words) were collected from the source corpus, as well as modifier and head family sizes (the number of different compound types sharing the constituent in the respective position) and as-modifier and as-head frequencies (the number of compound tokens containing the constituent in the respective position). All frequency-based measures were log-transformed.

Procedure. Participants were instructed to judge, as fast as possible, whether the items presented to them had a sensible interpretation (by pressing the C key on the

Table 2

Descriptive measures for Experiment 1 and 2.

	Experiment 1	Experiment 2
data points total (N)	36228	36475
mean data points per item	21.3	22.7
mean RT in ms	1027.5	990.4
item-wise RT correlation		.32

keyboard) or not (by pressing the N key). Lowercase items in black font were presented in a full-screen window on white background, and remained on the screen until a response was provided. Each item was preceded by a fixation cross at the center of the screen (250 ms) and followed by a blank screen (250 ms). The presentation order was randomized for each participant. For most participants, the experiment took between 5 and 20 minutes to complete (median = 12.5 minutes). The online experiment was implemented using the *jsPsych* library (de Leeuw, 2015).

Results and Discussion

The analysis focused on rejection times in novel compound trials. The initial exclusion of participants left us with 21 to 27 data points per item. Further, outlier trials with response times faster than 100 ms and slower than 5,000 ms were removed from the analysis (1.0 % of the data), as well as trials with *yes*-answers (11.0 % of the data). Descriptive measures for Experiment 1 are displayed in Table 2.

Logarithmic response times were analyzed using Linear Mixed Effect Models as implemented in the R packages *lme4* (Bates, Mächler, Bolker, & Walker, 2015) and *lmerTest* (Kuznetsova, Brockhoff, & Christensen, 2017). We first estimated a baseline model including lexical variables (log-transformed constituent and as-constituent frequencies, family sizes, compound length) as well as random intercepts for participants,

Table 3

Parameters for the final model predicting log-transformed rejection times for novel compounds in the timed sensibility judgments task (Experiment 1), after model criticism.

parameter	<i>b</i>	<i>t</i>	<i>p</i>
intercept	6.44	203.54	< .001
modifier frequency	0.01	4.18	< .001
as-modifier frequency	0.01	6.68	< .001
head family size	0.02	7.74	< .001
compound length	0.02	13.74	< .001
modifier composition	0.05	3.09	.002
head composition	0.04	2.24	.026

modifier constituent and head constituent. We then optimized the baseline model by eliminating non-predictive variables (see Table 3).

In a next step, we tested whether including semantic parameters significantly improved this baseline model in a likelihood-ratio test. This was the case when including a parameter for modifier composition ($\chi^2(1) = 6.21$, $p = .013$), as well as for head composition ($\chi^2(1) = 4.28$, $p = .039$). Furthermore, a model that contained both parameters outperformed the modifier composition-only model ($\chi^2(1) = 4.67$, $p = .031$) as well as the head composition-only model ($\chi^2(1) = 6.61$, $p = .010$). The final model, including both semantic parameters, is displayed in Table 3 (after model criticism, i.e. the removal of overly influential data points on the basis of a threshold of 2.5 SD standardized residual errors; Baayen, 2008). All parameter values are positive, so higher values predict longer rejection times.

Thus, we observe effects of semantic compositionality – as measured by the CAOSS model (Marelli et al., 2017) – on processing times in a timed sensibility judgment task. As expected and required by the task, participants engaged in a meaning-composition process

to build a semantic representation for the novel compound items presented to them, in line with the literature on conceptual combination (e.g. Gagné, 2001). Notably, this experiment shows that differences in semantic compositionality predict processing times even for strings that speakers are rejecting as non-sensible. Thus, the meaning-composition process affects linguistic processing, even when it clearly does not result in a representation the participants themselves can intuitively interpret in a meaningful way.

Notably, acceptance rates for novel compounds in Experiment 1 were relatively low. This is most likely caused by the selection criteria – none of these compounds are attested even once in a huge corpus of natural language, and there might be semantic as well as pragmatic reasons for these words never being used by speakers. In addition, half of the items in this experiment were existing compounds, which could have served as an implicit reference point for high sensibility. The inclusion of these existing compounds was necessary in order to ensure comparability between the results of the timed sensibility task in Experiment 1 and the lexical decision task in Experiment 2 (where existing words need to be included in the material). Still, an interesting venue for future research on compositional effects could be to include only novel compounds in a timed sensibility judgment task. In this case, we would expect higher overall acceptance rates, and faster acceptances as well as slower rejections for more compositional compounds.

Since semantic judgments as requested in Experiment 1 encourage a compositional process (Connell & Keane, 2006; Gagné, 2001; Gagné & Spalding, 2009), the present result could be specific to such a semantics-focused task, and not allow us to conclude that such a process *automatically* takes place in compound processing. Rather, they establish a baseline for which semantic effects can be expected in a task where a novel-compound meaning is constructed, to which semantic-composition effects in a lexical decision task can then be compared. In Experiment 2, we now turn to this lexical decision task while employing the same experimental material as in Experiment 1, to investigate whether this meaning composition is indeed an automatic process.

Experiment 2

Method

Participants. In Experiment 2, 218 native English speaking participants were recruited in an online crowdsourcing study, again using the *Amazon Mechanical Turk* platform. Applying the same exclusion criteria as in Experiment 1 (correct answers in less than 15 of the 20 catch trials, correct answers for less than 80 of the 100 most frequent existing compounds), we omitted data from 72 participants from the analysis. Of the remaining 146 participants ($M_{age} = 38.4$ years, $SD_{age} = 12.1$ years), 86 identified themselves as male and 59 as female (one unspecified). As in Experiment 1, participants received \$2 for participating in the study.

Material and Procedure. The semantic model, the material (real words as well as nonwords), and the procedure of Experiment 2 were completely identical to Experiment 1, except for a change in the instruction. In Experiment 2, participants were now instructed to perform a lexical decision task instead of a timed sensibility judgment task. More specifically, they were instructed to judge as fast as possible whether or not the items presented to them were existing English words. Most participants completed the experiment in between 5 and 25 minutes (median = 12.7 minutes).

Results

In Experiment 2, the initial exclusion of participants left us with 21 to 27 data points per item. Further, outlier trials with response times faster than 100 ms and slower than 5,000 ms were removed from the analysis (1.4 % of the data), as well as trials with *yes*-answers (7.2 % of the data). Descriptive measures for Experiment 2 are displayed in Table 2. As can be seen there, the item-wise correlation between response times in Experiment 1 and Experiment 2 is not very high ($r = .32$), indicating some behavioral divergence between the timed sensibility judgment task and the lexical decision task.

As in Experiment 1, we first identified a baseline Linear Mixed Effect Model by

Table 4

Parameters for the final model predicting log-transformed rejection times for novel compounds in the lexical decision task (Experiment 2), after model criticism.

parameter	<i>b</i>	<i>t</i>	<i>p</i>
intercept	6.46	216.52	< .001
as-modifier frequency	0.01	3.49	.001
as-head frequency	0.01	4.98	< .001
modifier family size	0.01	3.28	.001
head family size	0.01	3.44	.001
compound length	0.02	11.31	< .001
modifier composition	0.06	3.62	< .001

eliminating non-predictive lexical variables (see Table 4). In the likelihood-ratio tests assessing the effects of semantic parameters, the inclusion of a modifier composition parameter significantly improved the baseline model ($\chi^2(1) = 17.52$, $p < .001$), but a similar improvement was not observed with the inclusion of a head composition parameter ($\chi^2(1) = 0.34$, $p = .560$). Additionally including a head composition parameter did also not improve the model already containing the modifier composition parameter ($\chi^2(1) = 0.44$, $p = .507$). Therefore, the final model contained only a parameter for modifier composition (see Table 4 for the parameters of the final model after model criticism). Again, the parameter is positive, and therefore higher values predict longer rejection times.

Task-Interaction Analysis

The previous single-experiment analyses seem to suggest qualitative differences in the pattern of results (see Table 3 and Table 4). Most relevant to our study – which is focused on semantic effects – we find an effect of head composition in Experiment 1 (timed sensibility judgments), but not in Experiment 2 (lexical decision). However, this does not

Table 5

Parameters for the final model predicting log-transformed rejection times for novel compounds in both tasks (Experiment 1: timed sensibility judgments, Experiment 2: lexical decision judgments), after model criticism. The timed sensibility task serves as the reference level for the factor Experiment.

parameter	<i>b</i>	<i>t</i>	<i>p</i>
intercept	6.45	199.27	< .001
modifier frequency	0.004	3.19	.001
as-modifier frequency	0.01	3.93	< .001
as-head frequency	0.01	3.95	< .001
modifier family size	0.01	2.69	.007
length	0.02	14.53	< .001
Exp. : as-head freq.	0.003	3.67	< .001
Exp. : length	-0.01	-3.97	< .001
modifier composition	0.06	4.14	< .001

necessarily indicate that semantic effects differ between the two tasks, as the difference between a significant and a non-significant effect does not have to be significant itself (Gelman & Stern, 2006; Nieuwenhuis, Forstmann, & Wagenmakers, 2011). We therefore performed an additional combined analysis, in which we pooled the data from both experiments, and introduced a new dummy variable coding for the Experiment (Experiment 1 vs. Experiment 2). We then tested for interactions between this factor and the semantic parameters (modifier and head composition), which would reflect differential semantic effects depending on the task at hand.

In a first step, we estimated a baseline model containing interactions between Experiment and all lexical baseline variables, and then removed all non-predictive

parameters from the model (see Table 5).² The only semantic parameter whose inclusion significantly improved this baseline model is modifier composition ($\chi^2(1) = 18.43$, $p < .001$). Adding an interaction parameter between Experiment and modifier composition does not improve this model ($\chi^2(1) = 1.38$, $p = .240$), and neither does a parameter for head composition ($\chi^2(1) = 3.16$, $p = .076$) or an interaction parameter between Experiment and head composition ($\chi^2(2) = 3.71$, $p = .157$). Thus, the final model contained only a task-independent main effect for modifier composition (see Table 5 for the parameters of the final model after model criticism).

General Discussion

In the present study, we investigated semantic effects in the processing of compound “nonwords” (possible word combinations that are never attested in a very large text corpus) in two large-scale experiments, employing a timed sensibility judgment and a lexical decision task, respectively. In both experiments, we systematically manipulated the constituent contributions to the compositional meaning of such nonwords, as measured by a fully implemented computational model (the CAOSS model; Marelli et al., 2017). As hypothesized, participants are slower in rejecting as non-sensible novel compounds for which a compositional meaning is easier to compute, as indicated by the inhibitory effect of modifier composition. In other words, the easier the modifier meaning can be integrated into a concept combined from both constituents, the more difficult it is to reject the compound “nonword”. In line with previous findings (e.g., Meunier & Longtin, 2007; Gagné, 2001; Gagné & Shoben, 1997), this demonstrates that semantic factors, such as semantic interpretability, influence the processing of novel compounds overtly presented to participants.

² We observe the same pattern of results as reported here when we estimate the baseline model from combining the baseline models of the individual analyses for Experiment 1 and 2, and including interactions for those lexical parameters that are only significant in one baseline model (or also, alternatively, for all lexical parameters in both baseline models).

We observe this pattern of results in two different experimental paradigms entailing different task demands: A timed sensibility judgment task, explicitly requiring semantic processing and meaning composition, and a lexical decision task, which does not require semantic processing. And indeed, on the basis of the combined analysis of Experiment 1 and Experiment 2, there is no statistical evidence for the influence of such a compositional process differing between the tasks: We find no interaction between the task and the measures of semantic compositionality (modifier and head composition). However, this is not to say that the two tasks involve the same processes in general: The item-wise response time correlation is relatively low ($r = .32$), there are task interactions for as-head frequency and word length, and there certainly are influential factors that we did not consider here, including semantic ones. Our study focused only on one type of semantic effects – compositional effects – and we find no task interaction here. The most parsimonious explanation for the present results is that the processing of morphologically-structured words entails an automatic meaning-composition process in both tasks, even when it is not necessarily required (El-Bialy et al., 2013; Libben, 2006, 2014). In other words, speakers automatically engage in a meaning-composition process whenever they encounter a form that is, at a surface level, morphologically complex (Günther & Marelli, 2018; Marelli & Baroni, 2015). This emphasizes the remarkable creative abilities of humans to rapidly combine known elements into new concepts, even if such combinations are encountered for the very first time, and underline the cognitive system’s pursuit of extracting meaning whenever a stimulus constellation offers an opportunity to do so – even if semantic access is not required by the task at hand.

The combined analysis further reveals a reliable effect of modifier composition, while results concerning head composition are less clear. The latter is only present in the isolated analysis of the timed sensibility experiment, but does not reach significance in the combined analysis. Given this pattern, we remain agnostic regarding the presence of a head composition effect. However, irrespective of this issue, the strong modifier composition

effect clearly demonstrates semantic effects related to a meaning-composition process. Furthermore, the absence of a head composition effect does not imply that the head is irrelevant in processing: Modifier composition is measured as the similarity between modifier and the compositional compound meaning, and this compositional meaning is constructed as a combination of the modifier and the head meaning. The results rather indicate that the critical factor affecting processing speed is the ease of integrating the modifier into this combined concept. Evidence for the central role of the modifier in compound processing was reported in several previous studies (Gagné, 2000; Gagné & Shoben, 1997; Spalding, Gagné, Mullaly, & Ji, 2010), including the study by Günther and Marelli (2018) employing the same CAOSS-derived measures for estimating constituent contributions. This is in line with predictions from theoretical accounts of conceptual combination, ascribing a more prominent and active role to the modifier in the meaning-composition process (e.g. Gagné, 2000; L. L. Jones, Estes, & Marsh, 2008; Spalding et al., 2010): For example, according to the RICE model by Spalding et al. (2010), the modifier meaning guides the compositional process, as it makes available possible semantic relations between modifier and head, which are then evaluated with respect to both constituents.

In this perspective, the application of the CAOSS model does not only demonstrate a semantic component in the processing of novel compounds, but also sheds light on how this process would unfold, highlighting its relation with prominent theoretical proposals from conceptual combination (such as the RICE model, Spalding et al., 2010). Hence, the present study also serves as an important validation of the CAOSS model and its ability to adequately represent the meanings of novel compounds. This is especially relevant since, by definition, these words are not attested in the source corpus on which the distributional vectors are trained, and therefore also not part of the training set for the CAOSS algorithm. This is an important quality feature of computational models, as stated by Glenberg and Robertson (2000): “A computational model should be able to account for

material beyond its training set. It is especially important that a theory of language and meaning be flexible and productive beyond its training set because humans are flexible and productive” (p. 397). In this sense, it is important also from a validation perspective that the results by Günther and Marelli (2018) on existing compounds (which are part of the training set) are mirrored in the present study for novel compounds.

As noted earlier, our study focussed on compositional effects as one type of semantic effect in compound processing, and our item material was designed so that compositionality was experimentally manipulated as an independent variable (see Günther, Dudschig, & Kaup, 2016). However, there are semantic variables beyond those considered in our present study that potentially further influence the processing of novel compounds: For example, the similarity between modifier and head has been discussed as an index of conceptual combination (Lynott & Ramscar, 2001; Schmidtke & Kuperman, 2019). In addition, it has been observed that semantic factors such as valence (of the constituents as well as the compound, Schmidtke & Kuperman, 2019) and concreteness (at the compound level, Feldman, Basnight-Brown, & Pastizzo, 2006) affect processing times (see also Juhasz, Lai, & Woodcock, 2015; Kuperman, 2013). Notably, this line of research has so far mostly focused on existing compounds, and extending it to novel compounds (where any semantic aspect at the compound level needs to be induced from the constituents) is a promising direction for future research. Following this reasoning, also measures of semantic richness (Pexman, Hargreaves, Siakaluk, Bodner, & Pope, 2008), such as the number of semantic neighbors or the density of the semantic neighborhood (Buchanan, Westbury, & Burgess, 2001; Danguécan & Buchanan, 2016; Macdonald, 2013), are interesting candidates to be investigated in novel word combinations. Post-hoc analyses of constituent similarity and semantic neighborhood density, which we can compute using our distributional vectors but which were not controlled as independent variables in the construction of our dataset, are provided in Supplementary Material A.

The present results have important implications for theories of morphological

processing, and specifically the *morpheme interference effect* (Crepaldi, Rastle, & Davis, 2010; Crepaldi et al., 2013; Taft & Forster, 1975, 1976) highlighting how morphologically-valid but non-existing words (such as *radiosauce*) are rejected slower than combinations of morphemes and non-lexical elements (such as *ridiosauce*). These findings are typically interpreted as clear evidence that surface morphological structure, i.e., the parsability of an orthographic sequence into familiar (sub)lexical units, can hinder the rejection of unfamiliar strings. That is, the very fact that such representational units are available, and that they are activated during processing, is assumed to interfere with the response. However, as argued in the present article, morphologically-valid but non-existing words are in fact possible novel complex words for which a compositional meaning can, more or less easily and successfully, be derived. And the present results demonstrate that, depending on the specific compound at hand, these differences in semantic compositionality lead to considerable differences in processing times: All other variables being set to their mean value, the model prediction for the lowest modifier composition value (.05) is 898 ms, while for the highest value (.69) it is already at 933 ms, amounting to a difference of 35 ms. Thus, the size of a morpheme interference effect will vary considerably based on the nonwords selected – not only because of its lexical properties such as length and frequencies, but also due to its degree of semantic compositionality. Due to this potential for substantial variability, morpheme interference can only be reliably examined with careful consideration for “nonword” semantics and should not be reduced to a purely form-level effect.

In summary, the present study highlights the vital role of semantics in the processing of morphologically-structured linguistic input, irrespective of whether we formally consider them as words or nonwords. It is the very purpose of a (productive) morphological system to efficiently express new meanings by systematically combining familiar, meaning-bearing elements, and semantically plausible novel compounds are likely candidates to eventually enter the lexicon (Dhar & van der Plas, 2019). In such a context, novel compounds cannot

simply be treated as collocations of two constituents not interacting in any meaningful way, as if they were *busneighbors*, two strangers silently sitting next to each other during a bus ride. Rather, their constituents should be seen as two persons on a first date, who – more or less successfully – try to build a new, meaningful unit together: a *datecouple*.

Supplementary Material A: Post-Hoc Analyses – Constituent Similarity and Semantic Neighborhood Density

As noted in the General Discussion of the main article, there is a range of other possible semantic variables that can potentially affect the processing compounds, besides the compositional ones that are the focus of the article. In the post-hoc analyses presented here, we investigate the influence of two other variables that can directly be derived from our distributional vector representations: constituent similarity (Lynott & Ramscar, 2001; Schmidtke & Kuperman, 2019) and semantic neighborhood density (Buchanan et al., 2001; Danguécan & Buchanan, 2016; Macdonald, 2013).

We computed constituent similarity for each novel compound item as the cosine similarity between the distributional vectors of its modifier and head. Furthermore, we computed semantic neighborhood density (SND) as the average similarity between the compositional compound vector and all its neighbors (Buchanan et al., 2001; Danguécan & Buchanan, 2016). In this context, these neighbors are determined by first computing the cosine similarity between a distributional vector and *all* other word vectors in the semantic space, and then selecting only those whose cosine similarity lies at least 3.5 SD over the mean value of all these cosines (Danguécan & Buchanan, 2016; Macdonald, 2013).

Note however that, unlike modifier and head composition, both constituent similarity and SND were not explicitly manipulated as independent variables in the construction of our item material. This has two direct consequences: Firstly, unlike for modifier and head composition, these two variables don’t necessarily cover the whole range of their possible values in a uniform distribution. This is illustrated in Figure 2. Secondly, again unlike modifier and head composition, these variables were not controlled to be non-correlated. While the correlations are moderate for SND ($r = .13$ with modifier composition, $r = .34$ for head composition), they are rather substantial for constituent similarity ($r = .52$ with modifier composition, $r = .50$ for head composition). As a consequence, simultaneously including constituent similarity and the composition values in

Figure 2. The distributions of constituent similarity (left) and semantic neighborhood density (right) in the 1,499 items in our datasets

a model can lead to suppression effects due to collinearity (in this case, constituent similarity would have, as the only predictor variable in this model, a variance inflation factor (VIF) larger than 2). To counteract this, we de-correlated constituent similarity from modifier and head composition for the analyses presented here, by taking the residuals from a standard linear regression model predicting constituent similarity from modifier and head composition.

Given these considerations, the results presented here should be considered exploratory, and one should be cautious in drawing theoretical conclusions from them. In order to more thoroughly investigate the effect of constituent similarity and SND however, future studies should construct new item material that explicitly considers them as independent variables and experimentally manipulates them.

Analyses

For Experiment 1, additionally including the de-correlated constituent similarity to the model reported in the main article does not improve model predictions ($\chi^2(1) = 0.65$,

Table 6

Parameters for the extended model predicting log-transformed rejection times for novel compounds in the timed sensibility judgments task (Experiment 1), after model criticism.

parameter	<i>b</i>	<i>t</i>	<i>p</i>
intercept	6.49	185.58	< .001
modifier frequency	0.01	4.52	< .001
as-modifier frequency	0.01	7.01	< .001
head family size	0.02	8.11	< .001
compound length	0.02	13.96	< .001

modifier composition	0.05	3.30	.001
head composition	0.04	2.74	.006
SND	-0.12	-2.84	.004

$p = .421$). On the other hand, additionally including SND significantly improves the model ($\chi^2(1) = 5.57$, $p = .018$). The resulting model (after model criticism) is displayed in Table 6.

For Experiment 2 however, additionally including the de-correlated constituent similarity to the model reported in the main article did improve the model ($\chi^2(1) = 14.91$, $p = < .001$), whereas this was not the case for SND ($\chi^2(1) = 3.00$, $p = .084$). The resulting model (after model criticism) is displayed in Table 7.

Table 7

Parameters for the extended model predicting log-transformed rejection times for novel compounds in the lexical decision task (Experiment 2), after model criticism. Constituent similarity is de-correlated from modifier and head composition.

parameter	b	t	p
intercept	6.46	217.67	< .001
as-modifier frequency	0.01	3.51	.001
as-head frequency	0.01	5.02	< .001
modifier family size	0.01	3.80	< .001
head family size	0.01	3.88	< .001
compound length	0.02	11.31	< .001

modifier composition	0.06	2.71	.007
constituent similarity	0.15	6.06	< .001

References

- Amenta, S., Marelli, M., & Crepaldi, D. (2015). The fruitless effort of growing a fruitless tree: Early morpho-orthographic and morpho-semantic effects in sentence reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*, 1587–1596.
- Baayen, R. H. (2008). *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge, UK: Cambridge University Press.
- Baayen, R. H., Piepenbrock, R., & Gulikers, L. (1995). *The CELEX lexical data base (CD-ROM)*. University of Pennsylvania, Philadelphia, PA: Linguistic Data Consortium.
- Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., . . . Treiman, R. (2007). The English Lexicon Project. *Behavior Research Methods*, *39*, 445–459.
- Baroni, M., Dinu, G., & Kruszewski, G. (2014). Don’t count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of ACL 2014* (pp. 238–247). East Stroudsburg, PA: ACL.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48.
- Brysbaert, M., Buchmeier, M., Conrad, M., Jacobs, A. M., Bölte, J., & Böhl, A. (2011). The word frequency effect. *Experimental Psychology*, *58*, 412–424.
- Buchanan, L., Westbury, C., & Burgess, C. (2001). Characterizing semantic space: Neighborhood effects in word recognition. *Psychonomic Bulletin & Review*, *8*, 531–544.
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon’s Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, *6*, 3–5.
- Chamberlain, J. M., Gagné, C. L., Spalding, T. L., & Lõo, K. (2019). Detecting spelling

- errors in compound and pseudocompound words. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Advance online publication. doi: doi.org/10.1037/xlm0000748
- Connell, L., & Keane, M. T. (2006). A Model of Plausibility. *Cognitive Science*, 30, 95–120.
- Connell, L., & Lynott, D. (2013). Flexible and fast: Linguistic shortcut affects both shallow and deep conceptual processing. *Psychonomic Bulletin & Review*, 20, 542–550.
- Costello, F. J., & Keane, M. T. (2000). Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*, 24, 299–349.
- Crepaldi, D., Rastle, K., & Davis, C. J. (2010). Morphemes in their place: Evidence for position-specific identification of suffixes. *Memory & Cognition*, 38, 312–321.
- Crepaldi, D., Rastle, K., Davis, C. J., & Lupker, S. J. (2013). Seeing stems everywhere: Position-independent identification of stem morphemes. *Journal of Experimental Psychology: Human Perception and Performance*, 39, 510–525.
- Danguécan, A. N., & Buchanan, L. (2016). Semantic neighborhood effects for abstract versus concrete words. *Frontiers in psychology*, 7, 1034.
- de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47, 1–12.
- Dhar, P., & van der Plas, L. (2019). Learning to predict novel noun-noun combinations. In *Proceedings of the Joint Workshop on Multiword Expressions and WordNet (MWE-WN 2019)* (pp. 30–39). Florence, Italy: ACL.
- Downing, P. (1977). On the creation and use of English compound nouns. *Language*, 53, 810–842.
- El-Bialy, R., Gagné, C. L., & Spalding, T. L. (2013). Processing of English compounds is sensitive to the constituents’ semantic transparency. *The Mental Lexicon*, 8, 75–95.
- Estes, Z. (2003). Attributive and relational processes in nominal combination. *Journal of Memory and Language*, 48, 304–319.

- Feldman, L. B., Basnight-Brown, D. M., & Pastizzo, M. J. (2006). Semantic influences on morphological facilitation: Concreteness and family size. *The Mental Lexicon*, 1, 59–84.
- Gagné, C. L. (2000). Relation-based versus property based combinations: A test of the CARIN theory and dual-process theory of conceptual combination. *Journal of Memory and Language*, 42, 365–389.
- Gagné, C. L. (2001). Relation and lexical priming during the interpretation of noun–noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 236–254.
- Gagné, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier–noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 71–87.
- Gagné, C. L., & Spalding, T. L. (2006). Using conceptual combination research to better understand novel compound words. *SKASE Journal of Theoretical Linguistics*, 3, 9–16.
- Gagné, C. L., & Spalding, T. L. (2009). Constituent integration during the processing of compound words: Does it involve the use of relational structures? *Journal of Memory and Language*, 60, 20–35.
- Gagné, C. L., Spalding, T. L., Nisbet, K. A., & Armstrong, C. (2018). Pseudo-morphemic structure inhibits, but morphemic structure facilitates, processing of a repeated free morpheme. *Language, Cognition and Neuroscience*, 33, 1252–1274.
- Gelman, A., & Stern, H. (2006). The difference between "significant" and "not significant" is not itself statistically significant. *The American Statistician*, 60, 328–331.
- Glenberg, A. M., & Robertson, D. A. (2000). Symbol Grounding and Meaning: A Comparison of High-Dimensional and Embodied Theories of Meaning. *Journal of Memory and Language*, 43, 379–401.
- Graves, W. W., Binder, J. R., & Seidenberg, M. S. (2013). Noun–noun combination:

- Meaningfulness ratings and lexical statistics for 2,160 word pairs. *Behavior Research Methods*, 45, 463–469.
- Günther, F., Dudschig, C., & Kaup, B. (2016). Latent semantic analysis cosines as a cognitive similarity measure: Evidence from priming studies. *The Quarterly Journal of Experimental Psychology*, 69, 626–653.
- Günther, F., & Marelli, M. (2016). Understanding Karma Police: The perceived plausibility of noun compounds as predicted by distributional models of semantic representation. *PLOS ONE*, 11(10). doi: 10.1371/journal.pone.0163200
- Günther, F., & Marelli, M. (2018). Enter sandman: Compound processing and semantic transparency in a compositional perspective. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Advance online publication. doi: 10.1037/xlm0000677
- Günther, F., Marelli, M., & Bölte, J. (in press). Semantic transparency effects in German compounds: A large dataset and multiple-task investigation. *Behavior Research Methods*.
- Günther, F., Rinaldi, L., & Marelli, M. (2019). Vector-space models of semantic representation from a cognitive perspective: A discussion of common misconceptions. *Perspectives on Psychological Science*, 14, 1006–1033.
- Jackendoff, R. (2002). *Foundations of knowledge*. Oxford, UK: Oxford University Press.
- Jones, L. L., Estes, Z., & Marsh, R. L. (2008). An asymmetric effect of relational integration on recognition memory. *Quarterly Journal of Experimental Psychology*, 61, 1169–1176.
- Jones, M. N., Willits, J., & Dennis, S. (2015). Models of semantic memory. In J. Busemeyer, Z. Wang, J. Townsend, & A. Eidels (Eds.), *Oxford Handbook of Mathematical and Computational Psychology* (pp. 232–254). New York, NY: Oxford University Press.
- Juhasz, B. J., Lai, Y.-H., & Woodcock, M. L. (2015). A database of 629 English compound

- words: ratings of familiarity, lexeme meaning dominance, semantic transparency, age of acquisition, imageability, and sensory experience. *Behavior Research Methods*, 47, 1004–1019.
- Kuperman, V. (2013). Accentuate the positive: Semantic access in English compounds. *Frontiers in psychology*, 4, 203.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13), 1–26.
- 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.
- Libben, G. (2006). Why study compounds? An overview of the issues. In G. Libben & G. Jarema (Eds.), *The representation and processing of compound words* (pp. 1–21). Oxford, UK: Oxford University Press.
- Libben, G. (2014). The nature of compounds: A psychocentric perspective. *Cognitive Neuropsychology*, 31, 8–25.
- Libben, G., Gibson, M., Yoon, Y. B., & Sandra, D. (2003). Compound fracture: The role of semantic transparency and morphological headedness. *Brain and Language*, 84, 50–64.
- Longtin, C.-M., & Meunier, F. (2005). Morphological decomposition in early visual word processing. *Journal of Memory and Language*, 53, 26–41.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instrumentation, and Computers*, 28, 201–208.
- Lynott, D., & Ramscar, M. J. A. (2001). Can we model conceptual combination using distributional information? In *Proceedings of the Twelfth Irish Conference on Artificial Intelligence and Cognitive Science* (pp. 1–10). Maynooth, Ireland.
- Macdonald, G. (2013). *Aging and semantic processing*. (Doctoral dissertation, University

of Windsor, Windsor, ON)

- Marelli, M., & Baroni, M. (2015). Affixation in semantic space: Modeling morpheme meanings with compositional distributional semantics. *Psychological Review*, *122*, 485–515.
- Marelli, M., Gagné, C. L., & Spalding, T. L. (2017). Compounding as Abstract Operation in Semantic Space: A data-driven, large-scale model for relational effects in the processing of novel compounds. *Cognition*, *166*, 207–224.
- Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, *20*, 158–190.
- Meunier, F., & Longtin, C.-M. (2007). Morphological decomposition and semantic integration in word processing. *Journal of Memory and Language*, *56*, 457–471.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv:1301.3781v3*.
- Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, *12*, 529–562.
- Nieuwenhuis, S., Forstmann, B. U., & Wagenmakers, E.-J. (2011). Erroneous analyses of interactions in neuroscience: a problem of significance. *Nature Neuroscience*, *14*, 1105–1107.
- Pexman, P. M., Hargreaves, I. S., Siakaluk, P. D., Bodner, G. E., & Pope, J. (2008). There are many ways to be rich: Effects of three measures of semantic richness on visual word recognition. *Psychonomic Bulletin & Review*, *15*, 161–167.
- Sandra, D. (1990). On the representation and processing of compound words: Automatic access to constituent morphemes does not occur. *The Quarterly Journal of Experimental Psychology Section A*, *42*, 529–567.
- Schmid, H. (1995). Improvements in part-of-speech tagging with an application to German. In *Proceedings of the ACL SIGDAT-Workshop*. Dublin, Ireland.
- Schmidtke, D., & Kuperman, V. (2019). A paradox of apparent brainless behavior: The time-course of compound word recognition. *Cortex*, *116*, 250–267.

- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*, 523–568.
- Spalding, T. L., Gagné, C. L., Mullaly, A. C., & Ji, H. (2010). Relation-based interpretation of noun-noun phrases: A new theoretical approach. *Linguistische Berichte Sonderheft*, *17*, 283–315.
- Taft, M., & Forster, K. I. (1975). Lexical storage and retrieval of prefixed words. *Journal of Verbal Learning and Verbal Behavior*, *14*, 638–647.
- Taft, M., & Forster, K. I. (1976). Lexical storage and retrieval of polymorphemic and polysyllabic words. *Journal of Verbal Learning and Verbal Behavior*, *15*, 607–620.
- Tagalakakis, G., & Keane, M. T. (2006). Familiarity and relational preference in the understanding of noun-noun compounds. *Memory & Cognition*, *34*, 1285–1297.
- Thagard, P. (1984). Conceptual combination and scientific discovery. In P. Asquith & P. Kitcher (Eds.), *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* (Vol. 1, pp. 3–12). East Lansing, MI: Philosophy of Science Association.
- Turney, P. D., & Pantel, P. (2010). From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, *37*, 141–188.
- Westbury, C. (2016). Pay no attention to that man behind the curtain: Explaining semantics without semantics. *The Mental Lexicon*, *11*, 350–374.