



Citation: Liuzzi, A. G., Meersmans, K., Storms, G., De Deyne, S., Dupont, P., & Vandenberghe, R. (2023). Independence of coding for affective similarities and for word co-occurrences in temporal perisylvian neocortex. *Neurobiology of Language*, 4(2), 257–279. [https://doi.org/10.1162/nol\\_a\\_00095](https://doi.org/10.1162/nol_a_00095)

DOI:  
[https://doi.org/10.1162/nol\\_a\\_00095](https://doi.org/10.1162/nol_a_00095)

Supporting Information:  
[https://doi.org/10.1162/nol\\_a\\_00095](https://doi.org/10.1162/nol_a_00095)

Received: 26 January 2022  
Accepted: 9 December 2022

Competing Interests: The authors have declared that no competing interests exist.

Corresponding Author:  
Rik Vandenberghe  
[rik.vandenberghe@kuleuven.be](mailto:rik.vandenberghe@kuleuven.be)

Handling Editor:  
Steven Small

Copyright: © 2023  
Massachusetts Institute of Technology  
Published under a Creative Commons  
Attribution 4.0 International  
(CC BY 4.0) license



# Independency of Coding for Affective Similarities and for Word Co-occurrences in Temporal Perisylvian Neocortex

Antonietta Gabriella Liuzzi<sup>1\*</sup> , Karen Meersmans<sup>1\*</sup> , Gerrit Storms<sup>2</sup> , Simon De Deyne<sup>3</sup> , Patrick Dupont<sup>1</sup> , and Rik Vandenberghe<sup>1,4</sup> 

<sup>1</sup>Laboratory for Cognitive Neurology, Department of Neurosciences, Leuven Brain Institute, KU Leuven, Leuven, Belgium

<sup>2</sup>Laboratory of Experimental Psychology, KU Leuven, Leuven, Belgium

<sup>3</sup>Computational Cognitive Science Lab, University of Melbourne, Melbourne, Australia

<sup>4</sup>Neurology Department, University Hospitals Leuven, Leuven, Belgium

\*Antonietta Gabriella Liuzzi and Karen Meersmans share first authorship.

**Keywords:** fMRI, representational similarity analysis, semantics, word embedding models

## ABSTRACT

Word valence is one of the principal dimensions in the organization of word meaning. Co-occurrence-based similarities calculated by predictive natural language processing models are relatively poor at representing affective content, but very powerful in their own way. Here, we determined how these two canonical but distinct ways of representing word meaning relate to each other in the human brain both functionally and neuroanatomically. We re-analysed an fMRI study of word valence. A co-occurrence-based model was used and the correlation with the similarity of brain activity patterns was compared to that of affective similarities. The correlation between affective and co-occurrence-based similarities was low ( $r = 0.065$ ), confirming that affect was captured poorly by co-occurrence modelling. In a whole-brain representational similarity analysis, word embedding similarities correlated significantly with the similarity between activity patterns in a region confined to the superior temporal sulcus to the left, and to a lesser degree to the right. Affective word similarities correlated with the similarity in activity patterns in this same region, confirming previous findings. The affective similarity effect extended more widely beyond the superior temporal cortex than the effect of co-occurrence-based similarities did. The effect of co-occurrence-based similarities remained unaltered after partialling out the effect of affective similarities (and vice versa). To conclude, different aspects of word meaning, derived from affective judgements or from word co-occurrences, are represented in superior temporal language cortex in a neuroanatomically overlapping but functionally independent manner.

## INTRODUCTION

Over the past decade, advances in computational modelling of word meaning have enabled novel ways of studying the representation of meaning in the human brain. Models of word meaning have relied on different types of sources. Data can be acquired from study participants performing explicit tasks, such as free word association (De Deyne et al., 2019), feature generation (De Deyne et al., 2008; McRae et al., 2005), or valence ratings (Moors et al., 2013; Van Rensbergen et al., 2016). According to such studies, word valence, or how positive/negative a word is, is one of the principal dimensions of the organisation of word meaning

(Belyk et al., 2017; De Deyne & Storms, 2008; Kotz & Paulmann, 2011; Kousta et al., 2011; Kuperman et al., 2014; Meersmans et al., 2020; Osgood et al., 1957; Pauligk et al., 2019; Troche et al., 2017; Van Rensbergen et al., 2015; Vigliocco et al., 2014). Word valence is also highly relevant for how brain patterns are organised in response to words (Belyk et al., 2017; Meersmans et al., 2020, 2022; Pauligk et al., 2019). The effect of affective similarity (modelled as the Euclidean distance between words in a three-dimensional space formed by valence, dominance, and arousal ratings) is strong in the perisylvian language network and beyond (Meersmans et al., 2020, 2022).

As another highly influential approach to the representation of word meaning, word meaning models can make use of the distributional structure (Harris, 1954) present in the vast amount of pre-existing language corpora. Words can then be represented as word embeddings, i.e., continuous, low-dimensional, real-value vectors based on co-occurrences (Devereux et al., 2010; Hollis & Westbury, 2016; Mitchell et al., 2008; Pereira et al., 2013). Natural language processing (NLP; Abnar et al., 2018; Grave et al., 2019; Honnibal, 2017) is then applied to extract and detect regularities and statistical patterns of co-occurrence and transforms real-world language into computer-friendly real-valued vectors representations. A limitation, however, is that word embedding models have a rather limited ability to model the affective content of words, such as word polarity or other aspects of text sentiment (De Deyne et al., 2021; Tang et al., 2016; Yu et al., 2018). As these models are strongly context dependent, words with opposite polarity (e.g., “bad” vs. “good”) but similar contexts may be mapped near to each other by such models (Agrawal et al., 2018). Valence, arousal, and dominance can be derived from co-occurrence-based modelling but only provided that seed words are present of which the valence has been determined beforehand based on subjective ratings or lexica (“mots germes”; Recchia & Louwerse, 2015; Vincze & Bestgen, 2011). Even then, models based on task-generated data outperform co-occurrence-based models in predicting valence, arousal and dominance (Vankrunkelsven et al., 2018).

Hence, two of the principal ways of modeling word content, based on affective judgements versus word co-occurrences, capture different aspects of word meaning. The interest of this special relationship between co-occurrence-based modelling and affective word content can be gauged based on the number of high-quality prior studies from the NLP community (Tang et al., 2016; Yu et al., 2018), from computational linguistics (Westbury et al., 2015), psycholinguistics (Bestgen & Vincze, 2012; De Deyne et al., 2021), and experimental psychology on this topic (Hollis & Westbury, 2016; Recchia & Louwerse, 2015; Vankrunkelsven et al., 2018; Vincze & Bestgen, 2011; Westbury et al., 2015).

We examined how brain activity patterns relate to these two ways of representing word meaning that have been proven to be powerful in their own right but nevertheless appear to be fundamentally distinct. Given the distinctive nature of these representations, one might predict that they would rely on separate brain regions. Affective similarities have been reported to correlate strongly with the similarity of brain activity patterns in superior temporal language cortex (Meersmans et al., 2020), whereas co-occurrence-based similarities have been associated with a neuroanatomically relatively diffuse effect (Mitchell et al., 2008).

Word co-occurrences are influenced by word associations. Word associations can also be derived from other sources, such as experimentally using cued word generation (De Deyne et al., 2019; Vankrunkelsven et al., 2018). Hence, cued-association-based similarities hold a position in between co-occurrence-based and affective similarities. To gain further insight into the mechanism behind the correlations between co-occurrence-based similarities and brain activity patterns, we evaluated whether effects were comparable when similarities were based

on task-generated word association data. This allows us to determine whether the critical factor determining the co-occurrence-based word similarity effects are the predictive computational modelling algorithms behind the word similarities or the word associations as such.

We re-analysed a functional magnetic resonance imaging (fMRI) data set from a study originally designed to investigate the representation of affective similarity between nouns. Nouns were presented in both auditory and in visual modality to further increase generalisability of our findings and analyses were pooled across modalities and tasks. Affective similarities were calculated from estimates of word valence, arousal and dominance, which were extrapolated from a data set of behavioural affective norms (Moors et al., 2013). These extrapolations were made based on a graph derived from a huge data set based on free word association (Van Rensbergen et al., 2016; for further details see Materials and Methods). Secondly, for the noun stimulus set used in this experiment, we re-calculated word similarities, using the word2vec vectors provided through the spaCy Python library, and examined how the corresponding brain activity patterns related to those obtained for the affective similarities. The primary word embeddings are based on the continuous bag of words (cBOW) approach. In this architecture, the target word is predicted based on the embeddings of its immediate context (e.g., 3–4 neighbouring words, depending on the window size) derived from a large internet-based text corpus, composed of Wikipedia entries and an extensive Webcrawl (>13 billion words; Honnibal, 2017). This study was meant to uncover commonalities and differences in coding mechanisms for similarities based on word associations/co-occurrences versus affective judgements, respectively.

## MATERIALS AND METHODS

### Subjects


Twenty-two subjects (15 women and 7 men, mean age 21.9, 19–24 years of age, all neurologically healthy and right-handed) participated. This sample size is in accordance with previous fMRI studies using the same approach and yielding replicable results (Bruffaerts et al., 2013; Carota et al., 2021; Liuzzi et al., 2015, 2019; Pauligk et al., 2019). All subjects were free of significant neurological or psychiatric history. Subjects provided informed consent before participating, and the experiments were approved by the Ethics Committee Clinical Studies UZ/KU Leuven.

### Stimuli

Stimulus words were selected from the 12,400 cue words studied as part of the Dutch Small World of Words (SWOW) data set (De Deyne et al., 2019).

A total of 66 nouns were distributed evenly over three classes: positive, neutral, and negative valence (Figure 1). Words were selected in a semiautomated manner based on the following criteria: (1) to maximise the range in semantic similarity, that is, contain word pairs that range between highly similar and highly dissimilar; and (2) to match words between the three valence classes for all relevant linguistic variables. The three groups were matched on age of acquisition, concreteness, dominance,  $\log_{10}$ (frequency), orthographic neighbourhood density, prevalence, and word length (Brysbaert et al., 2014; Keuleers et al., 2010, 2015; Marian et al., 2012; Van Rensbergen et al., 2016). The word selection script did not converge to a solution when arousal also needed to be matched between valence classes. Arousal was numerically lower for the neutral valence (mean = 3.98,  $SD$  = 0.49) class than for the positive (mean = 4.23,  $SD$  = 0.79) or negative valence class (mean = 4.54,  $SD$  = 1.05) (one-way analysis of

Small World of Words: Publicly available data set obtained by means of online responses of more than 70,000 individuals to more than 12,000 words, performed for different native languages.



**Figure 1.** Stimulus set. The two coordinates of each noun correspond to valence (x-axis) and concreteness (y-axis).

variance:  $F(2, 63) = 2.62, p = 0.08$ ). This is a consequence of the well-established U-shaped relationship between valence and arousal (Warriner et al., 2013).

The stimulus set is provided in Figure 1 and in Table S1 in the Supporting Information. The stimulus selection procedure yielded both concrete and abstract words without evident taxonomic structure (as measured based on WordNet; Miller, 1995).

### Experimental Design

The experiment was originally designed for the purpose of studying the effect of valence on fMRI activity patterns. The experiment had a  $3 \times 2 \times 2$  factorial design with stimulus type (positive, neutral, negative), modality (visual or auditory presentation), and overt articulation (yes/no) as factors. After a warning dot (500 ms), stimulus words were presented in visual or auditory modality (1,500 ms), followed by the presentation of a red or green dot (2,500 ms). Presentation of a green dot indicated to the subjects that they had to repeat the word aloud after stimulus presentation had ended. The red dot indicated that subjects had to remain silent. Visual stimuli were presented in white on a black background. Auditory stimuli were presented through OptoActive II active noise cancelling headphones (Optoacoustics, 2022) to minimise interference of scanner noise. The task was chosen to monitor subject engagement in the scanner, while limiting the effect of attentional orienting and control that can occur during explicit semantic tasks. To limit head motion, subjects were explicitly instructed to

pronounce the words while restricting mouth movement to a minimum. Vocal responses were recorded using the OptoActive II FOMRI III microphone (Optoacoustics, 2022), which enables recording of responses in the fMRI environment in a sensitive manner. Head motion parameters were determined during image analysis, and runs were rejected based on a prior criterion (see below).

The initial 1,500 ms phase following word onset was identical between covert and overt trials up to response cue onset. The covert and overt trials will be pooled in the current report. Following the response phase, a white fixation dot was presented until the start of the next trial. The interstimulus interval was 8.25 s. Every noun was presented once during each of the eight runs (4× in visual and 4× in auditory modality; duration per run: 10.5 min). Visual stimuli were accompanied by an auditory control (i.e., rotated spectrogram speech), while auditory stimuli were accompanied by a visual control (i.e., consonant letter string). Audio controls were generated from the stimuli by rotating the spectrogram around 50% of the maximal frequency and applying a low-pass filter at 95% of the maximal frequency. This procedure maintains certain spectral and phonetic features and intonation, while at the same time rendering the word unintelligible (Scott et al., 2000). Visual controls were generated by replacing the vowels in the original stimulus words by consonants and shuffling the letter order. During the control trials, the visual and auditory control stimuli were presented simultaneously, and no response was required (Figure 2).

Every run consisted of 77 trials (66 word trials + 11 control trials). The control trials were included as event type in the general linear model but were not used further in the current analyses. Every subject completed eight runs over two scan sessions. The word sets were presented in a pseudorandomised order and divided in a counterbalanced manner between the two sessions. This resulted in a total of eight replications per word (4 spoken and 4 written modality).

### Word similarity matrices

The primary objective of the current analysis was to examine how similarities derived from word co-occurrence-based models relate to similarities between fMRI activity patterns and to identify the commonalities and differences of these effects with those obtained for affective

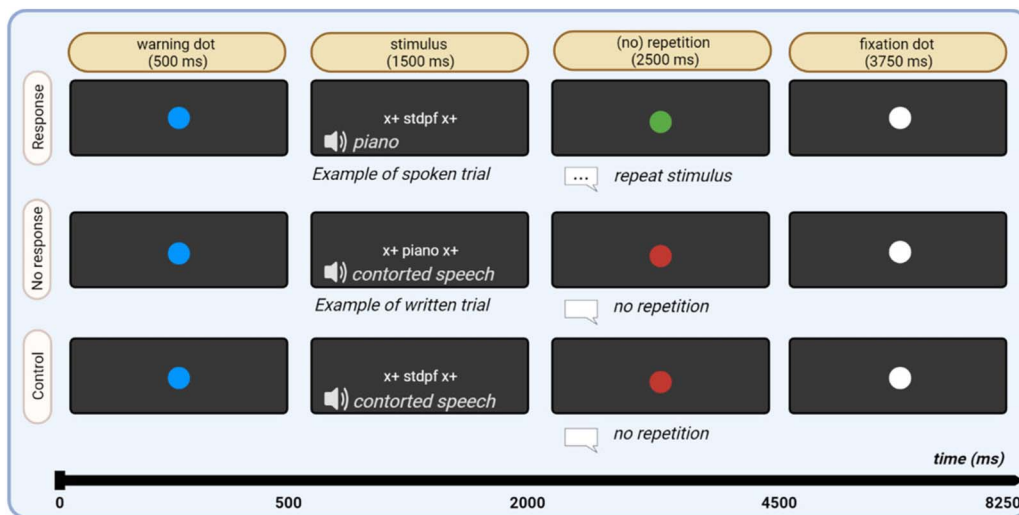


Figure 2. Experimental design.

Representational similarity analysis (RSA):

Method where the similarity between experimental stimuli (such as words) is related to the similarity between the activity patterns in response to these stimuli.

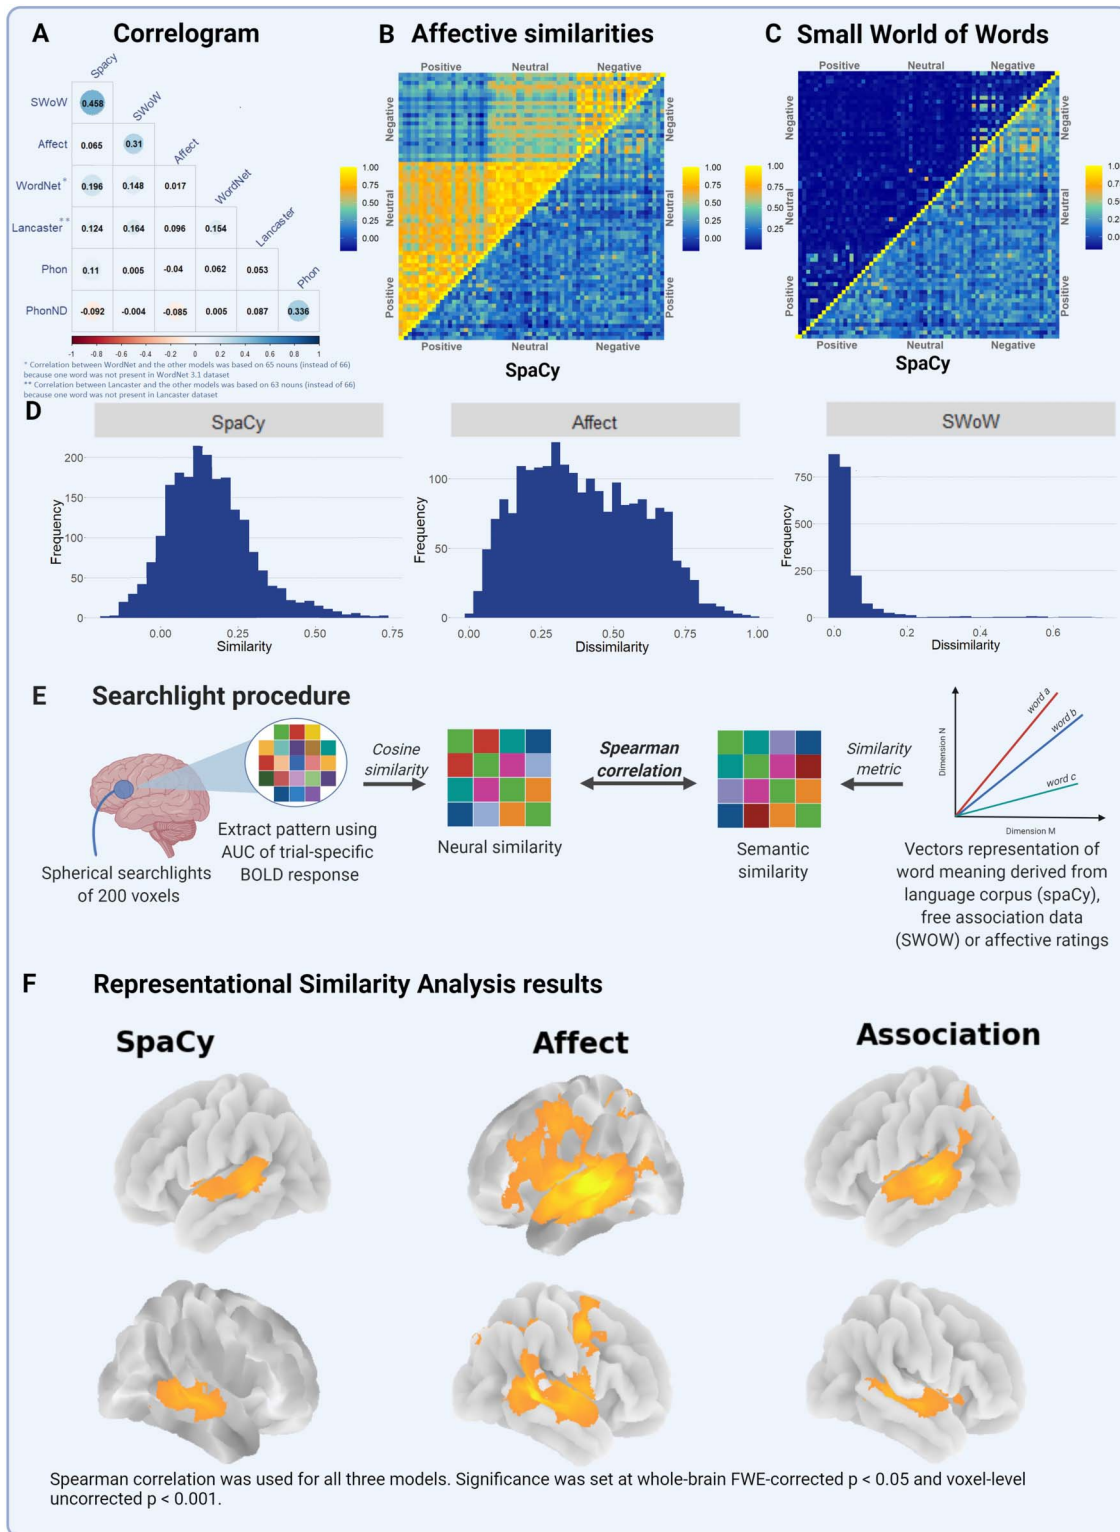
similarities. The word co-occurrence model was derived directly from spaCy (Honnibal, 2017). SpaCy is an open-source Python library for NLP. It contains built-in word embeddings (i.e., multidimensional vector representations of word meaning). Here, we use the word2vec word embeddings from spaCy's large Dutch model ([https://spacy.io/models/nl#nl\\_core\\_news\\_lg](https://spacy.io/models/nl#nl_core_news_lg)), which uses a cBOW approach on a combined corpus (web crawled text + Wikipedia). SpaCy similarities were normally distributed (Figure 3D; Figure S1 in the Supporting Information). The resulting similarity matrix was used for representational similarity analysis (RSA).

In the affective similarity model, we used valence, dominance, and arousal ratings (Van Rensbergen et al., 2016) to create a three-dimensional space. The pairwise Euclidean distance  $d$  between stimulus words was calculated, normalised, and converted to a similarity metric by  $1 - d$ . Valence, dominance, and arousal ratings were taken from a large data set of affective estimates (Van Rensbergen et al., 2016). In this study, word valence, dominance, and arousal were extrapolated for 14,000 words from a smaller data set of behavioural norms ( $n = 4,300$  words; Moors et al., 2013) using a  $k$ -nearest neighbour approach. The nearest neighbours were determined based on word embedding similarities derived from a free association graph (De Deyne et al., 2019). This graph is created from a large data set of continued associations (>12,000 words; see below for more details) with the words as nodes in the graph and the edges weighted by the association strength  $p(\text{response}|\text{cue})$ . In the affective ratings data set by Moors et al. (2013), word valence, dominance, and arousal were rated on a 1–7 scale. Only words for which these extrapolations were available were eligible for this study.

In a secondary analysis, for comparison with the spaCy based model, we included a behavioural similarity model derived from the Small World of Words word association data set (De Deyne et al., 2019). In this similarity model, word embeddings were derived from a large-scale continued association task, during which subjects were instructed to provide three associates for every cue (De Deyne et al., 2019). The SWOW data set comprises responses for over 12,000 cue words from 70,000+ subjects. From this data set, a graph is constructed using associative strength (i.e., probability of the response given a cue) as edge weight. Word embeddings were extracted from the graph using a random walk algorithm, and semantic similarity was calculated using cosine similarity (De Deyne et al., 2019; Liuzzi et al., 2019; Meersmans et al., 2020). Valence (pleasantness; positive [*holiday*] – negative [*murder*]), dominance (power; strong [*avalanche*] – weak [*sleep*]), and arousal (activity; active [*explosion*] – passive [*silence*]) are important dimensions in this graph (Van Rensbergen et al., 2015). Furthermore, a direct comparison between word embeddings and word association indicated that the latter is better suited to capture affective information (De Deyne et al., 2021) The SWOW similarity matrix was not normally distributed (Figure 3D and Figure S1 in the Supporting Information).

To evaluate and control for the possible contribution of phonological similarities to the effects obtained, we calculated a phonological similarity matrix as  $1 -$  the Levenshtein distance between phonological transcriptions of the stimuli. As a separate phonological model, we also calculated phonological neighbourhood density similarities as the difference in phonological neighbourhood density between stimuli (i.e., the number of words that can be constructed by substitution, addition, or deletion of one phoneme). Phonological transcriptions and neighbourhood density estimates were retrieved from the CLEARPOND database for Dutch (Marian et al., 2012; <https://clearpond.northwestern.edu/dutchpond.php>). These models were included to evaluate whether the observed semantic effects are in part driven by phonology.

Since prior studies have emphasised the role of taxonomic structure and of experiential similarities in the organization of brain activity pattern (Fernandino et al., 2022), we also calculated word similarity based on taxonomy and word similarity derived from experiential



**Figure 3.** Overview of the word embedding models, the whole-brain representational similarity analysis (RSA) procedure, and the results. (A) Correlations between the different similarity matrices. (B) Affective similarity matrix (upper half of the matrix) next to the distributional similarity matrix derived from spaCy (lower half of the matrix). (C) Compares the Small World of Words and spaCy matrices in the same manner. (D) Distribution of the similarities from spaCy, affective ratings and Small World of Words (SWoW). (E) Provides a schematic overview of the whole-brain RSA using a 200 voxel searchlight. (F) Results of the RSA for all three models.

judgements for our set of nouns and tested the correlation with the dimensions of primary interest, the affective similarities, and the co-occurrence-based similarities. Taxonomic structure was modelled based on WordNet (<https://wordnet.princeton.edu/>). WordNet similarity matrix was calculated based on Wu and Palmer (1994) similarity (WPsim). WPsim calculates the relatedness between two words by considering the depths of the two synonyms (i.e., synsets) in the WordNet taxonomies, along with the depth of the least common subsumer (i.e., the most specific common hypernym). Experiential strength was based on the Lancaster sensorimotor norms (Lynott et al., 2020). Lancaster similarity matrix was calculated by computing the cosine similarity for all pairs of concept vectors (see Table 1 for an overview of the models). The taxonomic and experiential similarities were calculated from English translations of the Dutch nouns. Three nouns were absent from the Lancaster norms and one noun was absent from WordNet, hence these were left out (Table S1, for Dutch nouns and English translations).

### **Correlations between similarities from the different models**

Correlations between the different word embedding models show the difference in overlap between models (Figure 3A). The correlations between the co-occurrence model and the association-based model was 0.458, showing the partial collinearity between the models. This is in line with the view that co-occurrence-based models capture some of the meaning encoded in word associations.

In contrast, the affective similarity model correlated poorly with the co-occurrence-based model ( $r = 0.065$ ). Affective similarity was better captured by the association-based model ( $r = 0.310$ ) than by the co-occurrence-based models. This difference was formally evaluated using the cocor package for R (Diedenhofen & Musch, 2015), indicating that the correlation between the affective and association-based model was significantly stronger than the correlation between the affective and co-occurrence-based model (Pearson and Filon's  $z = -11.4$ ;  $p < 0.0001$ ).

The correlations between the phonological (neighbourhood density) matrix and the semantic matrices were overall very low (Figure 3A). The Spearman correlation of the co-occurrence-based similarities with the taxonomy-based similarities was 0.196, and with the experiential strength-based similarities 0.124. The Spearman correlation of the affective similarities with the taxonomy-based similarities was 0.017, and with the experiential-strength-based similarities 0.096. Despite the low or negative correlations with phonological distance and neighbourhood density, taxonomy-based or experiential-strength-based similarity matrices, we also calculated the RSA effect of co-occurrence-based and affective similarities after partialling out each of these factors.

### **Image Acquisition**

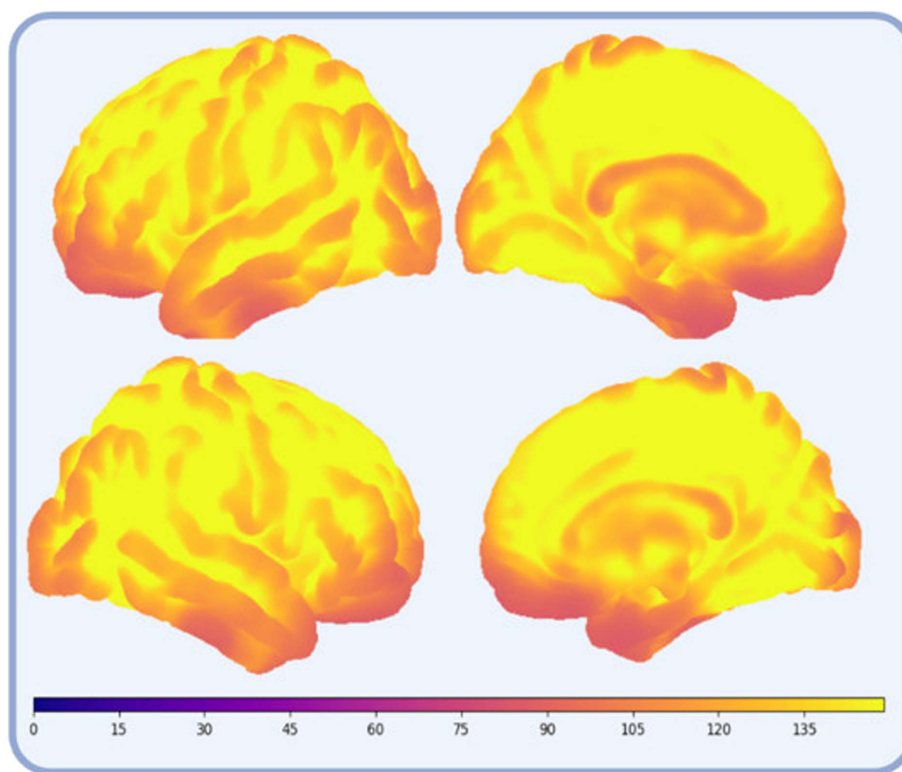
Functional and structural fMRI images were acquired on a 3T Philips Achieva with a 32-channel head coil and equipped with the OptoActive II ANC headphones and FOMRI III microphone. Functional images were acquired using a T2\* sequence with 36 slices (multiband acceleration factor = 2; repetition time = 1 s; echo time = 30 ms, field of view =  $220 \times 135 \times 220$  mm<sup>3</sup>, voxel size =  $2.75 \times 2.75 \times 3.75$  mm<sup>3</sup>). The temporal signal-to-noise ratio was calculated by dividing the mean of the residual time series by its standard deviation (Murphy et al., 2007; Figure 4). Structural images were obtained using a T1-weighted 3D turbo-field-echo (repetition time = 9.6 ms, echo time = 4.6 ms, in-plane resolution = 0.97 mm, slice thickness = 1.2 mm) in every subject.



**Table 1.** Overview of the models.

	Type	Data	N	Procedure	Similarity
SpaCy	Co-occurrence	Wikipedia and webcrawl	300	cBOW	Cosine
Affective similarities	Task-based	Valence, dominance, and arousal ratings	3	Human ratings extrapolated based on SWOW similarities	Euclidean
Small World of Words	Task-based	Semantic associations	NA	Random walk	Cosine
Phonological similarities	Phonological	Transcriptions from CLEARPOND	NA	Number of insertions, deletions or substitutions needed to change one string into the other	Levenshtein
Phonological neighbourhood density	Phonological	Total number of words that differed in the addition, deletion, or substitution of a single phoneme	1	NA	Absolute value of the difference in phonological neighbourhood density
Taxonomic structure	Taxonomy	WordNet v31	NA	Knowledge graph where words are grouped into synsets	Wu & Palmer similarity
Experiential strength	Task-based	Lancaster sensorimotor norms	12	Human ratings for 39,707 concepts across six perceptual modalities and five action effectors	Cosine similarity

*Note.* N: number of dimensions in the model. cBOW: continuous bag of words. SWOW: Small World of Words.



**Figure 4.** Temporal signal-to-noise ratio, calculated by dividing the mean of the residual time series by its standard deviation.

### Image Preprocessing

Images were submitted to a preprocessing pipeline in SPM12 (Ashburner et al., 2021). Functional images were realigned to correct for head motion, followed by slice timing correction. Next, functional and structural images were coregistered and the anatomical images were segmented into white matter, grey matter, and cerebrospinal fluid. The grey matter probability maps are used to restrict the analysis to cortical voxels. The anatomical and functional images were normalised to Montreal Neurological Institute (MNI) space. No smoothing was applied to suit the multivariate nature of our analyses. During realignment, motion parameters were generated. Runs where the framewise displacement (the sum of the absolute values of the differentiated realignment parameters) exceeded 1 mm were excluded from further analysis.

### Statistical Analysis

We applied RSA in a whole-brain searchlight procedure using the CosMoMVPA toolbox (Oosterhof et al., 2016). Trials were pooled over tasks (overt articulation and covert trials), over word classes (positive, neutral, negative) and over modalities (auditory and visual). Control trials where no words were presented were not included. Figure 3E provides an overview of the searchlight procedure.

Trial-specific activation maps were used as input for this analysis. These maps were generated by extracting the voxelwise blood oxygen level dependence (BOLD) response from the normalised, unsmoothed images and calculating the area under the curve of the BOLD response in every voxel from 2 to 8 s. This procedure is identical to that applied in previous

multivariate pattern analysis (MVPA) studies (Bruffaerts et al., 2013; Liuzzi et al., 2015, 2017; Meersmans et al., 2020, 2022). The number of voxels per searchlight was set a priori to 200 (corresponding to a sphere radius = 11 mm; voxel size equal to  $2 \times 2 \times 2 \text{ mm}^3$ ) and the recommended data centring was applied before the analysis. The searchlight analysis was limited to subject-specific grey matter masks, created by thresholding the grey matter probability maps at 0.3. In every searchlight, the activity pattern was extracted from the activation maps and pairwise similarities were estimated using cosine similarity. The search sphere in CosMoMVPA may contain voxels from opposite walls of a sulcus. During the analysis, subject-specific voxelwise RSA maps were created containing the correlations between the activity similarities of the searchlight centred on a specific voxel and the word similarity model. The distribution of the task-based similarities (affective matrix and free association-based matrix) deviated from normality (Figure 3D; Figure S1). Therefore, Spearman correlation was used in the searchlight RSA for all models to facilitate between-model comparison. These maps were *r* to *z* transformed and smoothed before being submitted to a group-level *t* test using SPM12. At every voxel, the *t* test evaluates whether the correlation is significantly different from zero across subjects. These *t* tests were performed for every searchlight RSA separately. To find significant differences between pairs of word similarity models, the same maps were submitted to a paired *t* test. Significance at the group level was set to a whole-brain familywise error (FWE)-corrected threshold of  $p < 0.05$  (with uncorrected voxelwise  $p < 0.001$ ).

Finally, in order to address the unique relation between a specific model (e.g., co-occurrence) and the neural representation, partial correlation searchlight RSAs were conducted. The procedure was identical to the one described above, with the only difference that the correlation between a model representational dissimilarity matrix (RDM; e.g., co-occurrence) and the neural RDM was controlled for another model (e.g., affective).

## RESULTS

On average, subjects made 3.7 (range 0–14) repetition errors out of 264 overt articulation trials (i.e., wrong word or no repetition) over the total of eight runs. Over all subjects, a total of eight runs (4.5%) had to be excluded due to excessive head (framewise displacement > 1 mm).

### Effect of Co-Occurrence-Based Similarities

Using searchlight RSA, similarities between words estimated using word2vec in spaCy correlated significantly with the similarity of activity patterns in the lateral temporal neocortex surrounding the superior temporal sulcus (STS) from the posterior end to the middle third, to the left and also to the right (Figure 3F; Table 2). Table 3 provides the Spearman correlation and 95% confidence interval (CI) of the local maxima yielded by the group-level (SPM12) analysis of the searchlight RSA for co-occurrence similarity. Since the analysis is based on a random-effects model, the Spearman correlations provided are the average of the Spearman correlations obtained in each of the individuals. No other regions showed a significant similarity effect.

We verified the similarity between the word2vec based matrix (spaCy) and the fMRI activity patterns after partialling out the effect of affective similarity and association-based similarities. The RSA results remained essentially the same when controlling for affective similarities (L pSTS peak coordinates:  $-57, -34, 2$ ;  $t(21): 4.80$ ; cluster size: 377 voxels; cluster level FWE-corr  $p = 0.002$ ; R pSTS peak coordinates:  $48, -28, 2$ ;  $t(21): 4.83$ ; cluster size: 702 voxels; cluster level FWE-corr  $p < 0.001$ ; Figure 5A). This is logical given the very low correlation between the co-occurrence-based model and the affective similarity model. When controlling for the association-based similarities, no significant effects of the co-occurrence similarities were observed, in line

**Table 2.** Results of the whole-brain searchlight representational similarity analysis (RSA) with word2vec similarities (spaCy) and affective similarities.

<b>RSA results for co-occurrence-based similarities</b>						
Label	Size	FWE <i>p</i>	Peak coordinates			<i>t</i> (21)
L posterior STS	782	<0.001	-57	-34	2	5.48
			-54	-13	-4	4.75
			-69	-1	-4	4.56
R posterior STS	1,023	<0.001	54	-37	5	5.37
			48	-28	5	5.33
			57	-22	-1	5.05
<b>RSA results for affective similarities</b>						
Label	Size	FWE <i>p</i>	Peak coordinates			<i>t</i> (21)
Left and right STS, left MTG, precentral gyrus, precuneus. Left middle and inferior FC, left FG	13,319	<0.001	-69	-25	5	8.60
			-51	-34	5	7.54
			-36	-64	-1	7.27
<b>RSA results for free association-based similarities</b>						
Label	Size	FWE <i>p</i>	Peak coordinates			<i>t</i> (21)
Left STS, left rostral AG, left FG	2,423	<0.001	-57	-34	-1	6.78
			-42	-43	-7	5.65
			-66	-1	-4	5.05
Right STS	1,210	<0.001	51	-37	-1	5.32
			54	-16	-1	5.17
			48	-28	5	5.12

*Note.* Spearman correlation was used for all analyses. Significance was set at a whole-brain FWE-corrected threshold of  $p < 0.05$  at the cluster level (after applying an uncorrected voxelwise threshold  $p < 0.001$ ). Size: cluster size in number of voxels ( $2 \times 2 \times 2 \text{ mm}^3$ ), in MNI space. FWE *p*: FWE-corrected *p*-value at the cluster-level. Abbreviations: STS: superior temporal sulcus; MTG: middle temporal gyrus; FC: frontal cortex; FG: fusiform gyrus; AG: angular gyrus; FWE: familywise error.

with the moderate degree of collinearity between these two similarity models. As expected based on the correlogram between the different matrices (Figure 3A), the results remained unchanged when controlling for phonological similarity or similarity in phonological neighbourhood density using partial correlations (Table S2A and S2B). It also remained unchanged when controlling for taxonomy-based or for experiential similarities (Table S2C and S2D).

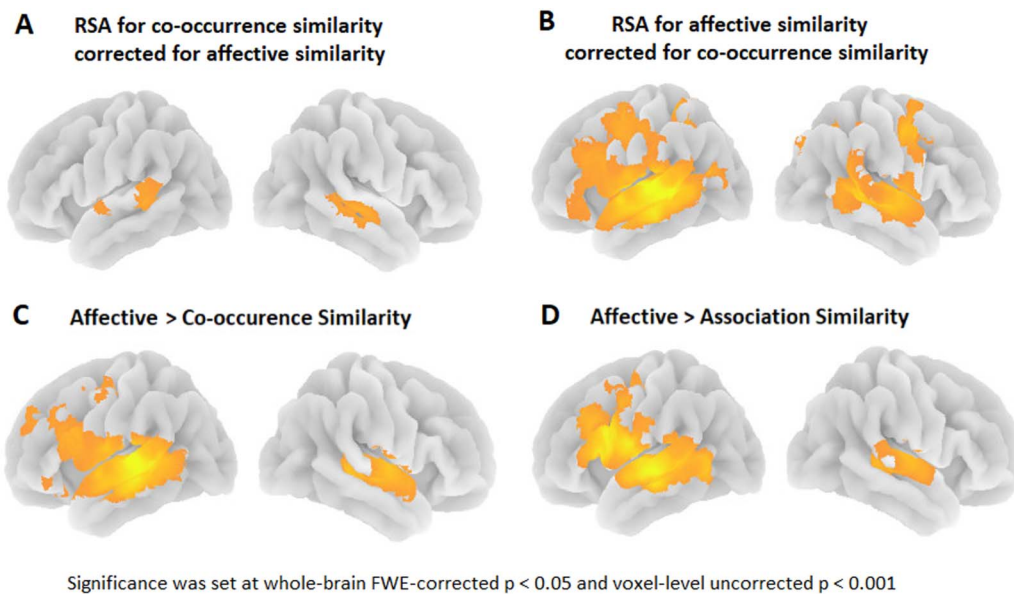
**Comparison to the Effect of Affective Similarities**

For the affective similarities, strong effects were obtained in the lateral temporal perisylvian neocortex, and more distributed effects outside of lateral temporal cortex were also observed, namely in left middle and inferior frontal regions and fusiform gyrus (Figure 3F; Table 3F–), which provides the Spearman correlation and 95% CI of the local maxima yielded by the group-level (SPM12) analysis of the searchlight RSA for affective similarity. Using a paired

**Table 3.** Correlational values and 95% confidence interval (CI) based on average correlation value across volunteers.

(A) RSA for co-occurrence similarity						(F) RSA for affective similarity					
Local maxima			Rho	Lower CI	Upper CI	Local maxima			Rho	Lower CI	Upper CI
-57	-34	5	0.026	-0.018	0.066	-69	-25	5	0.034	-0.007	0.076
(B) RSA for co-occurrence similarity corrected for affective similarity						(G) RSA for affective similarity corrected for co-occurrence similarity					
Local maxima			Rho	Lower CI	Upper CI	Local maxima			Rho	Lower CI	Upper CI
-57	-34	2	0.020	-0.021	0.063	-69	-25	5	0.034	-0.008	0.076
(C) RSA for co-occurrence similarity corrected for association-based similarity						(H) RSA for affective similarity corrected for association-based similarity					
Local maxima			Rho	Lower CI	Upper CI	Local maxima			Rho	Lower CI	Upper CI
Not significant						-69	-25	5	0.030	-0.012	0.072
(D) RSA for co-occurrence similarity corrected for taxonomic similarity (WordNet)						(I) RSA for affective similarity corrected for taxonomic similarity (WordNet)					
Local maxima			Rho	Lower CI	Upper CI	Local maxima			Rho	Lower CI	Upper CI
-57	-34	2	0.026	-0.016	0.068	-69	-25	5	0.033	-0.009	0.075
(E) RSA for co-occurrence similarity corrected for experiential similarity (Lancaster)						(J) RSA for affective similarity corrected for experiential similarity (Lancaster)					
Local maxima			Rho	Lower CI	Upper CI	Local maxima			Rho	Lower CI	Upper CI
-57	-34	2	0.021	-0.021	0.064	-69	-25	5	0.033	-0.008	0.076

*Note.* Values are reported for the left posterior temporal peak local maxima based on the group-level (SPM12) of the (A) RSA for co-occurrence similarity, RSA for co-occurrence similarity (B) corrected for affective similarity, (C) corrected for association-based similarity (D) corrected for taxonomic similarity, (E) corrected for experiential similarity. (F) RSA for affective similarity, RSA for affective similarity (G) corrected for co-occurrence similarity, (H) corrected for association-based similarity, (I) corrected for taxonomic similarity, (J) corrected for experiential similarity.



**Figure 5.** 3D rendering of the (A) semantic similarity effect for co-occurrence similarity corrected for affective similarity; (B) semantic similarity effect for affective similarity corrected for co-occurrence similarity; (C) paired  $t$  test showing the correlation strength of affective similarity stronger than co-occurrence similarity; (D) paired  $t$  test showing the correlation strength of affective similarity stronger than association-based similarity. RSA: representational similarity analysis; FWE: familywise error.

$t$ -test across subjects, we evaluated significant differences between the models. Correlations were significantly higher for the affective similarities compared to spaCy- or association-derived similarities in bilateral superior and middle temporal gyrus and left middle and inferior frontal gyrus (paired  $t$  test; FWE-corrected cluster-level  $p < 0.05$  combined with uncorrected voxel-level  $p < 0.001$ ; Figure 5C and D). No regions were identified where spaCy- or association-derived similarities had a stronger effect than affective similarities. The effects of affective similarities remained unaltered when partialling out the effects of co-occurrence-based (peak-coordinates:  $-69, -25, 5$ ;  $t(21) = 8.68$ ; cluster size = 13,143 voxels; cluster level FWE-corr  $p < 0.001$ ; Figure 5B) or association-based similarities. Similarly, results did not change after partialling out phonological distance, phonological neighbourhood density, taxonomy-based and experiential-based similarity (Table S3).

#### Effect of Free Word Association-Based Similarities

Co-occurrence-based models are influenced by word associations. We also evaluated the effect of association-based similarities derived from SWOW. In a secondary analysis, we repeated the RSA using a behavioural similarity matrix derived from free association data (SWOW data set). The results of the SWOW similarity matrix overlapped with those for the word2vec matrix in STS but extended more posteriorly into the rostral angular gyrus and inferiorly into the fusiform gyrus (LH: 2,423 voxels, peak coordinate  $-57, -34, -1$ ; RH: 1,210 voxels, peak coordinate  $51, -37, -1$ ; Table 2; Figure 3F).

#### DISCUSSION

There was a striking correspondence between word co-occurrence-based similarities (word2vec) and similarities in activity patterns in the neocortex surrounding the left STS, and this was also true for affective similarities. Despite this co-localisation in superior temporal

neocortex, the co-occurrence and affective similarity effects in lateral temporal neocortex were independent, in line with the low correlation between the two similarity matrices. In contrast, the results obtained with the co-occurrence-based and SWOW word association model showed a strong resemblance to each other, in line with the correlation between the word similarities derived from these two types of word association models.

### **Comparison to Previous fMRI Studies Using Word2vec Modelling**

The co-occurrence-based similarities were localised to a region confined to the lateral superior temporal cortex. This contrasts with some of the earlier studies of co-occurrence-based similarities showing a neuroanatomically distributed correlation between co-occurrence-based word similarities and the similarity of brain activity patterns. In a landmark study, Mitchell et al. (2008) described a predictive relation between co-occurrence data and neural activity. Each object stimulus word was represented as a 25-dimensional vector, with each value corresponding to the normalised sentence-wide co-occurrence of that word with 25 sensorimotor verbs (e.g., see, hear, eat; see also Anderson et al., 2013). While dimensionality reduction is an integral part of the word2vec approach applied here, the dimensions retained are entirely data driven based on the corpora analysed (300 dimensions in the spaCy model). The low-dimensional sensorimotor approach could potentially explain the high congruence between predicted and observed data in visual, motor, or gustatory cortex in the Mitchell et al. (2008) study. Representations in lateral temporal cortex seem to be less dependent on sensorimotor information, but instead rely on language-internal sources. These representations are also better modelled in a data-driven space with more dimensions than when dimensions are defined by the investigator. When the Mitchell et al. (2008) data set was analysed using word2vec, good concordance was achieved between the model and the overall brain activity pattern (Abnar et al., 2018). However, in that study the neuroanatomical network underlying this result was not the main topic of interest and relatively poorly defined.

Another study using co-occurrence-based vectors also reported a more distributed pattern than the focalised lateral temporal region we found (Pereira et al., 2018). Between-study differences in the similarity modelling can explain the difference in outcome. In Pereira et al. (2018) the co-occurrence-based vectors form the basis for spectral clustering that divides the semantic space in approximately 180 clusters that are then each represented by a representative word selected by the examiners. Next the authors derive dimensions from the set of representative words. During the fMRI experiment the words were embedded in sentences that also contained other words from the same cluster. In contrast to Pereira et al. (2018) the current paper directly imports the word embeddings as determined by spaCy, without further processing steps. It is conceivable that the sequence of meaning-oriented procedures in Pereira et al. (2018) may lead to a semantic representation that is richer in content than the representations based purely on corpus-derived word embeddings. This may then explain the wider neuroanatomical distribution of the effects in Pereira et al. (2018) compared to the current findings.

In an approach similar to the current one, Wang et al. (2018) performed a RSA with three similarity matrices: based on co-occurrences, based on semantic features, and based on subjective ratings. All words were abstract. Co-occurrence-based similarities were correlated with similarities of activity patterns in perisylvian language regions in temporal and frontal cortex (based on a volume of interest encompassing the perisylvian language network in its entirety), whereas the more semantically based similarities resulted in a more distributed pattern. In their globality, the findings by Wang et al. (2018) indicate that the perisylvian language network codes for co-occurrence-based similarities. The current findings confirm this observation in

a whole-brain RSA in a larger group of participants (22 vs. 6), a wider range of words and input modalities, and without prior word selection step based on stability of responses. Another finding in common with Wang et al. (2018) as well as other studies (e.g., Meersmans et al., 2020) is the strength of the effect of affective similarities. As a major difference with Wang et al., in a whole-brain search analysis, the co-occurrence-based-similarities were limited to the cortex surrounding the STS in our study while they were found in ventral and middle premotor cortex in Wang et al. (2018) (mostly in the ventral and middle premotor cortex according to Figure 4B). Furthermore, the correlation with the co-occurrence-based model in the cortex surrounding the STS was bilateral in the current study and left lateralised in the study by Wang et al. (2018). The difference in localisation may be due to a variety of reasons, for instance, the use of a broader range of words in our study compared to abstract words only in Wang et al., and the use of auditory along with visual modality.

A fourth study (Fernandino et al., 2022) revealed that representational similarity for co-occurrence-based similarities was weaker than for taxonomy-based similarities or, even more so, similarities based on componential experiential semantics (Binder et al., 2016). For that reason, we repeated the RSA after partialling out taxonomic and experiential similarities. Results remained essentially the same. The difference between studies can be accounted for by the difference in composition of the stimulus set. In the current study, words were equally distributed between three valence classes and selected semiautomatically to maximise the range of the pairwise semantic similarities controlling for a list of relevant variables. The stimulus set did not contain an evident taxonomic structure and many of the experiential features of Binder et al. (2016) were not applicable to the words we used. The difference between studies demonstrates that RSA reflects the structure present in the word set. Comparisons between the effects of different dimensions within a word set are valid but should not be interpreted in absolute terms, and the outcome will depend on the structure present in the experimental set of words. Had we composed a word set where words were evenly distributed over the traditional semantic categories, it is likely that the effect of taxonomic structure would have been much more prominent (Fernandino et al., 2022).

#### **Comparison Between the Results for Co-Occurrence and Affective Similarity-Based Models**

The lateral temporal region that showed co-occurrence-based similarities in our study also showed strong affective similarity effects. When the affective similarity effect was partialled out, the co-occurrence-based effect remained unchanged, as expected given the low correlation between the co-occurrence-based similarities and the affective similarities (Figure 3). Hence, the cortex surrounding the STS codes for co-occurrences and for affective similarities likely through relatively separate operational mechanisms.

Affective information is particularly difficult to capture from co-occurrence-based word embeddings (De Deyne et al., 2021; Tang et al., 2016; Yu et al., 2018): The correlation between affective and co-occurrences models equalled 0.065, while the correlation between affective and free word association-based models was 0.310. In the current study, the affective model resulted in stronger and notably more extensive effects than the co-occurrence model in the RSA. The effect of affective similarities was present in the STS, replicating previous findings (Meersmans et al., 2020), but was more widespread in line with the distributed representation of word meaning (Binder & Desai, 2011; Huth et al., 2016). The effect of co-occurrence-based similarities cannot be accounted for by the affective similarity effect, as it was unaltered when partialling out the affective similarity effect. Models that are derived from evaluative judgements by human subjects capture word meaning to a fuller extent than those that are derived



from co-occurrences in text corpora, and this is a plausible explanation of the stronger and more extensive effects. However, the effect of co-occurrence-based similarities remain after partialling out the affective similarities, and this indicates that the two models represent different information about words and not that one is simply a weaker version of the other.

The comparison with SWOW-based similarities is of interest because they occupy a position somewhat in between co-occurrence-based similarities and affective similarities. SWOW is based on a cued free association task. It shares with affective similarities the origin in subjective judgements by human participants and with co-occurrence-based similarities its strong link with word associations. This intermediary position between co-occurrence-based and affective-judgement-based modelling is also evident from the correlogram (Figure 3A): SWOW-based similarities correlate both with co-occurrence-based similarities and with affective similarities. When SWOW word association similarities were partialled out, the co-occurrence-based effect was no longer significant. Co-occurrence and free word association-based similarities share a substantial amount of information ( $r = 0.458$ ; Figure 3A), but association-based models are known to better capture affective information. With the current correlational approach, it is impossible to disentangle whether the co-occurrence-based or the word-association-based similarities are driving the effect more strongly given the collinearity between the two matrices (Figure 3A).

#### **Hypothetical Functional-Anatomical Model of Superior Temporal Cortex**

Taken together, the findings from the co-occurrence-based and the affective model may be understood in the context of the theory of a hybrid conceptual system with symbolic and embodied representations (Andrews et al., 2009; Louwerse & Jeuniaux, 2010; Paivio, 1991). The co-occurrence-based model yielded strong effects specifically in the left STS, the affective similarity model yielded even stronger, but independent, effects in this region, and beyond. A hybrid model could explain why both affective similarity and word2vec similarities correlate significantly with the superior temporal neocortex, despite a relatively low correlation between the two similarity matrices (Figure 3A). The two processes, one co-occurrence/association-based and the other driven by valence as one of the principal dimensions of word meaning, may operate in a same region in a relatively independent manner.

The lateral temporal region to the left corresponds to what is also known as Wernicke's area, although its exact anatomical boundaries are debated (Binder, 2015). Wernicke already hypothesised about the role of co-occurrences in the formation of associations between cortical responses (Gage & Hickok, 2005), a mechanism that has since been shown to be important in processing at distinct linguistic levels (Blank & Davis, 2016; Forseth et al., 2020; Schuster et al., 2021). The activated zone surrounding the STS has been previously implicated in human word intelligibility in a very consistent manner (Davis & Johnsruide, 2003; Evans et al., 2014; Obleser & Kotz, 2010; Scott et al., 2000). Word intelligibility requires integration of multiple, hierarchical levels, from the acoustic over the phonological to the semantic level (Obleser, 2014) and can be manipulated bottom up (e.g., by filtering the acoustic signal) or top down (e.g., predictability based on sentence context; Obleser & Kotz, 2010). Obviously, in connected speech, statistical regularities are extremely useful for predicting and recognizing words efficiently despite the high word rate and any degradation of the acoustic signal that may exist. Even for artificially isolated words, word identification may be guided by a scaffold of learnt co-occurrences. Here, the terms context and prediction do not refer to actual experimental context or predictions, but to probabilities learnt from language exposure. In connected speech, these co-occurrence-based regularities may be very informative for predictive coding (Arnal et al., 2011; Blank & Davis, 2016).

The role of superior temporal cortex in coding word associations/co-occurrences was relatively independent from its role in coding affective similarities. Estimating affective similarities requires subjective ratings and can be based on human ratings or on dictionaries and lexica of data, such as emoticon labels (Tang et al., 2016), labelled by experts. The current data confirm that these affective dimensions, known to be important for meaning representation for decades (Osgood et al., 1957), are also expressed in the perisylvian lateral temporal cortex (Meersmans et al., 2020). Intuitively one could have expected a more prominent role of so-called emotion processing regions rather than language regions for coding affective word similarities. However, the current findings confirm earlier and consistent observations of the coding of affective similarities by superior lateral temporal regions (Meersmans et al., 2020). Words with positive or negative valence are generally more abstract than neutral valence words. Abstract words more heavily rely on the linguistic system for their meaning than concrete words (Borghi & Binkofski, 2014; Kousta et al., 2011; Louwerse, 2011; Vigliocco et al., 2014). We have previously postulated that this may relate to the strong effect of affective similarities in superior temporal perisylvian language cortex (Meersmans et al., 2020). Neurobiologically, this is a correlate of the entwinement between affect and language (Wilce, 2009).

Our results do not imply that superior temporal cortex activity patterns are exclusively explained by either association-based similarities or affective similarities for any set of nouns. As mentioned before, results of RSA mirror the structure within the stimulus set employed. Activity patterns in superior temporal language cortex certainly code for other dimensions too, such as taxonomic structure or componential experiential semantic structure (Fairhall & Caramazza, 2013; Fernandino et al., 2022). This by no means contradicts our findings. Which model yields the best correlations between word similarity and similarity in activity patterns depends on the type of stimulus set and maybe also the task performed (Meersmans et al., 2022). For a word set with a high proportion of concrete words and hence a clear taxonomic structure and strong sensory embodiment, an experiential model may be superior to a co-occurrence-based model (Fernandino et al., 2022). However such an observation should not be generalised to all types of word sets. The current stimulus set has been selected semi-automatically to maximise the range in pairwise similarities and to control for the proportion of positive, neutral, and negative valence stimuli. It was not generated in order to have a representative and consistent sample of traditional semantic categories distributed over the three valence classes. As a consequence, the current stimulus set did not contain a clear taxonomic structure and several of the 64 features of a prevailing componential experiential semantic model (Binder et al., 2016) were not applicable to the current stimulus set. It is important to put fMRI representational similarity results in the context of the type of word stimuli examined. As a consequence, exclusionary claims towards the representation of dimensions of word meaning in the lateral superior temporal language cortex are unwarranted. When the stimulus set contains a clear categorical structure, categorical effects can be found in superior temporal cortex (Carota et al., 2021; Devereux et al., 2013), and likewise for taxonomy-based structure (Fernandino et al., 2022) or experiential sensorimotor strengths (Carota et al., 2021; Fernandino et al., 2022). The structure within the experimental set of nouns will determine which dimensions are most strongly represented by superior temporal cortex in response to these nouns in a dynamic, context-dependent manner.

### **Limitations**

The use of overt articulation during fMRI can potentially introduce motion artefacts. The state-of-the-art voice recording system with noise cancellation enables the recording of responses in a sensitive manner even when speech volume is low. Subjects were specifically instructed to

pronounce the words without moving their mouths excessively and without head motion. Head motion was evaluated and runs where the a priori threshold for framewise displacement (>1 mm) was exceeded were excluded from further analysis.

Trials with overt articulation and trials without overt responses were pooled for RSA. This increases the generalisability of our findings. The subjects received the task instruction only after the word stimulus had been shown so that the initial phase (1,500 ms duration) of conscious word processing is matched regardless of response or judgement. For repetition semantic processing is not necessary (Jefferies et al., 2006) although in the intact brain conscious word perception will inevitably activate a certain level of semantic processing. This is also evidenced by the empirical effects of affective and word association-based similarities upon the brain activity pattern.

The word sets were selected to be evenly distributed across different word valence classes. They may therefore not be representative for the total noun thesaurus which is not distributed evenly across negative, neutral, and positive valence classes. This could introduce a bias in favour of affective similarity modelling versus word co-occurrence-based modelling. However, the main finding is the independency of coding between the two types of similarities rather than the strength of the effect of one versus the other approach.

The affective ratings used to create the three-dimensional affective similarity space were extrapolated from a smaller set of behavioural ratings based on similarities derived from the SWOW association data set. As such, the stronger overlap between the association-based and affective models can be in part explained by this procedure. On the other hand, this underlines our central conclusion of independent coding of co-occurrence-based and affective similarities: Using extrapolations could potentially have introduced a similarity-based confound into the affective estimates.

Behaviourally, there was no correlation between the affective similarity matrix and the co-occurrence-based similarity. This observation fits with the initial rationale of the current experiment, namely that co-occurrence-based similarities capture affective similarities poorly and that for affective similarities to be calculated requires subjective input. The word set we chose was not specifically selected for an absence of correlation between affective and co-occurrence-based similarities, nevertheless for other word sets such a correlation may be found, although we would expect it in general to be relatively weak. The conclusion of functional independence of the superior temporal representation of affective and co-occurrence-based similarities is based on the absence of a behavioural correlation, the co-localisation as well as the absence of any effect of partialling out the effects of co-occurrence or affective similarities, respectively. The absence of an effect of partialling out the effects of co-occurrence or affective similarities follows mostly from the absence of a correlation between the two matrices behaviourally. In the future, for word sets where affective and co-occurrence-based similarities show a weak but significant correlation, it will be worthwhile to examine what remains after partialling out the effect of the respective matrices.

The visual control stimuli consisted of consonant letter strings and the auditory control stimuli of rotated spectrograms. Hence, in the visual modality, the stimuli still contain linguistic symbols but not in the auditory control stimuli. This is unlikely to affect the results, which are based on pooled analysis of auditory and visual word trials.

The sensitivity of the RSA is difficult to determine accurately as there is no groundtruth. This is the case for most functional imaging studies of cognition. The sensitivity depends on a number of factors and not only on the number of words examined. Other factors that determine the sensitivity and specificity are the number of within-subject replications per word (8 in our

study), the range of similarities between the words in the word set (maximised in our study using as semiautomated procedure), the number of individuals scanned, the intertrial interval duration and the hemodynamic modelling procedure. These multiple factors together determine sensitivity and specificity.

When association-based similarities were partialled out, the representational similarity effect for co-occurrence-based similarities disappeared. This emphasises the shared information in both word association models. Co-occurrence-based models and free word association models both tap into word meaning as derived from word context. With a correlational approach the collinearity between the different measures requires caution for exclusionary explanations. Hence the findings should not be interpreted in terms of a general superiority of one type of models versus another in modelling word associations. Neither is it possible to empirically ascertain whether the correlation between the co-occurrence-based matrix and the fMRI-activity-based similarity matrix reflects any correspondence in underlying computational mechanisms. The correlation between the co-occurrence-based matrix and the fMRI-activity-based similarity matrix may also be indirect via the associations between word meanings. Hence the proposal that co-occurrence-based mechanisms may play a role in coding of word meaning in lateral temporal cortex remains hypothetical at the current stage.

## CONCLUSIONS

The results of this study lead us to hypothesise that the brain uses several behavioural and distributional mechanisms to construct word meaning, in accordance with a hybrid framework of embodied and symbolic representations. While experiential information is represented in a more distributed fashion, corpus-based information correlates with patterns in a focalised region surrounding the STS. The latter relies on statistical regularities that allow for predictive coding.

## FUNDING INFORMATION

Rik Vandenberghe, Onderzoeksraad, KU Leuven (<https://dx.doi.org/10.13039/501100004497>), Award ID: C14/17/108. Rik Vandenberghe, Onderzoeksraad, KU Leuven (<https://dx.doi.org/10.13039/501100004497>), Award ID: C14/21/109. Rik Vandenberghe, Fonds Wetenschappelijk Onderzoek (<https://dx.doi.org/10.13039/501100003130>), Award ID: G094418N. Antonietta Gabriella Liuzzi, Fonds Wetenschappelijk Onderzoek (<https://dx.doi.org/10.13039/501100003130>), Award ID: 1247821N.

## AUTHOR CONTRIBUTIONS

**Antonietta Gabriella Liuzzi:** Formal analysis: Supporting; Methodology: Supporting; Writing – review & editing: Supporting. **Karen Meersmans:** Conceptualization: Lead; Data curation: Lead; Formal analysis: Lead; Investigation: Lead; Methodology: Lead; Writing – original draft: Lead; Writing – review & editing: Lead. **Gerrit Storms:** Conceptualization: Supporting; Writing – review & editing: Supporting. **Simon De Deyne:** Formal analysis: Equal; Methodology: Equal; Software: Supporting; Writing – review & editing: Supporting. **Patrick Dupont:** Methodology: Equal; Software: Equal; Writing – review & editing: Equal. **Rik Vandenberghe:** Conceptualization: Lead; Funding acquisition: Lead; Methodology: Equal; Supervision: Lead; Writing – original draft: Equal; Writing – review & editing: Lead.

## DATA AVAILABILITY STATEMENT

Data are available on <https://github.com/lcn-kul/affective-similarities>.

## REFERENCES

- Abnar, S., Ahmed, R., Mijnheer, M., & Zuidema, W. (2018). Experiential, distributional and dependency-based word embeddings have complementary roles in decoding brain activity. In *Proceedings of the 8th Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2018)* (pp. 57–66). Association for Computational Linguistics. <https://doi.org/10.18653/v1/W18-0107>
- Agrawal, A., An, A., & Papagelis, M. (2018). Learning emotion-enriched word representations. In *Proceedings of the 27th International Conference on Computational Linguistics* (pp. 950–961). Association for Computational Linguistics.
- Anderson, A. J., Bruni, E., Bordignon, U., Poesio, M., & Baroni, M. (2013). Of words, eyes and brains: Correlating image-based distributional semantic models with neural representations of concepts. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing* (pp. 1960–1970). Association for Computational Linguistics.
- Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, *116*(3), 463–498. <https://doi.org/10.1037/a0016261>, PubMed: 19618982
- Arnal, L. H., Wyart, V., & Giraud, A. L. (2011). Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. *Nature Neuroscience*, *14*(6), 797–801. <https://doi.org/10.1038/nn.2810>, PubMed: 21552273
- Ashburner, J., Barnes, G., Chen, C.-C., Daunizeau, J., Flandin, G., Friston, K., Gitelman, D., Glauche, V., Henson, R., Hutton, C., Jafarian, A., Kiebel, S., Kilner, J., Litvak, V., Mattout, J., Moran, R., Penny, W., Phillips, C., Razi, A., ... Zeidman, P. (2021). *SPM12 manual*. Wellcome Centre for Human Neuroimaging.
- Belyk, M., Brown, S., Lim, J., & Kotz, S. A. (2017). Convergence of semantics and emotional expression within the IFG pars orbitalis. *NeuroImage*, *156*, 240–248. <https://doi.org/10.1016/j.neuroimage.2017.04.020>, PubMed: 28400265
- Bestgen, Y., & Vincze, N. (2012). Checking and bootstrapping lexical norms by means of word similarity indexes. *Behavior Research Methods*, *44*, 998–1006. <https://doi.org/10.3758/s13428-012-0195-z>, PubMed: 22396137
- Binder, J. R. (2015). The Wernicke area: Modern evidence and a reinterpretation. *Neurology*, *85*(24), 2170–2175. <https://doi.org/10.1212/WNL.0000000000002219>, PubMed: 26567270
- Binder, J. R., Conant, L. L., Humphries, C. J., Fernandez, L., Simons, S. B., Aguilar, M., & Desai, R. H. (2016). Toward a brain-based componential semantic representation. *Cognitive Neuropsychology*, *33*(3–4), 130–174. <https://doi.org/10.1080/02643294.2016.1147426>, PubMed: 27310469
- Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends in Cognitive Sciences*, *15*(11), 527–536. <https://doi.org/10.1016/j.tics.2011.10.001>, PubMed: 22001867
- Blank, H., & Davis, M. H. (2016). Prediction errors but not sharpened signals simulate multivoxel fMRI patterns during speech perception. *PLOS Biology*, *14*(11), Article e1002577. <https://doi.org/10.1371/journal.pbio.1002577>, PubMed: 27846209
- Borghi, A. M., & Binkofski, F. (2014). *Words as social tools: An embodied view on abstract concepts*. Springer. <https://doi.org/10.1007/978-1-4614-9539-0>
- Bruffaerts, R., Dupont, P., Peeters, R., De Deyne, S., Storms, G., & Vandenberghe, R. (2013). Similarity of fMRI activity patterns in left perirhinal cortex reflects semantic similarity between words. *Journal of Neuroscience*, *33*(47), 18597–18607. <https://doi.org/10.1523/JNEUROSCI.1548-13.2013>, PubMed: 24259581
- Brysbaert, M., Stevens, M., De Deyne, S., Voorspoels, W., & Storms, G. (2014). Norms of age of acquisition and concreteness for 30,000 Dutch words. *Acta Psychologica*, *150*, 80–84. <https://doi.org/10.1016/j.actpsy.2014.04.010>, PubMed: 24831463
- Carota, F., Nili, H., Pulvermüller, F., & Kriegeskorte, N. (2021). Distinct fronto-temporal substrates of distributional and taxonomic similarity among words: Evidence from RSA of bold signals. *NeuroImage*, *224*, Article 117408. <https://doi.org/10.1016/j.neuroimage.2020.117408>, PubMed: 33049407
- Davis, M. H., & Johnsrude, I. S. (2003). Hierarchical processing in spoken language comprehension. *Journal of Neuroscience*, *23*(8), 3423–3431. <https://doi.org/10.1523/JNEUROSCI.23-08-03423.2003>, PubMed: 12716950
- De Deyne, S., Navarro, D. J., Collell, G., & Perfors, A. (2021). Visual and affective multimodal models of word meaning in language and mind. *Cognitive Science*, *45*(1), Article e12922. <https://doi.org/10.1111/cogs.12922>, PubMed: 33432630
- De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2019). The “Small World of Words”: English word association norms for over 12,000 cue words. *Behavior Research Methods*, *51*(3), 987–1006. <https://doi.org/10.3758/s13428-018-1115-7>, PubMed: 30298265
- De Deyne, S., & Storms, G. (2008). Word associations: Network and semantic properties. *Behavior Research Methods*, *40*(1), 213–231. <https://doi.org/10.3758/BRM.40.1.213>, PubMed: 18411545
- De Deyne, S., Verheyen, S., Ameel, E., Vanpaemel, W., Dry, M. J., Voorspoels, W., & Storms, G. (2008). Exemplar by feature applicability matrices and other Dutch normative data for semantic concepts. *Behavior Research Methods*, *40*(4), 1030–1048. <https://doi.org/10.3758/BRM.40.4.1030>, PubMed: 19001394
- Devereux, B. J., Clarke, A., Marouchos, A., & Tyler, L. K. (2013). Representational similarity analysis reveals commonalities and differences in the semantic processing of words and objects. *Journal of Neuroscience*, *33*(48), 18906–18916. <https://doi.org/10.1523/JNEUROSCI.3809-13.2013>, PubMed: 24285896
- Devereux, B. J., Kelly, C., & Korhonen, A. (2010). Using fMRI activation to conceptual stimuli to evaluate methods for extracting conceptual representations from corpora. In *Proceedings of the NAACL HLT 2010 First Workshop on Computational Neurolinguistics* (pp. 70–78). Association for Computational Linguistics.
- Diedenhofen, B., & Musch, J. (2015). Cocor: A comprehensive solution for the statistical comparison of correlations. *PLOS ONE*, *10*(3), Article e0121945. <https://doi.org/10.1371/journal.pone.0121945>, PubMed: 25835001
- Evans, S., Kyong, J. S., Rosen, S., Golestani, N., Warren, J. E., McGettigan, C., Mourao-Miranda, J., Wise, R. J. S., & Scott, S. K. (2014). The pathways for intelligible speech: Multivariate and univariate perspectives. *Cerebral Cortex* *24*(9), 2350–2361. <https://doi.org/10.1093/cercor/bht083>, PubMed: 23585519
- Fairhall, S. L., & Caramazza, A. (2013). Brain regions that represent amodal conceptual knowledge. *Journal of Neuroscience*, *33*(25), 10552–10558. <https://doi.org/10.1523/JNEUROSCI.0051-13.2013>, PubMed: 23785167
- Fernandino, L., Tong, J. Q., Conant, L. L., Humphries, C. J., & Binder, J. R. (2022). Decoding the information structure underlying the neural representation of concepts. *Proceedings of the National Academy of Sciences*, *119*(6), Article e2108091119. <https://doi.org/10.1073/pnas.2108091119>, PubMed: 35115397
- Forseth, K. J., Hickok, G., Rollo, P. S., & Tandon, N. (2020). Language prediction mechanisms in human auditory cortex. *Nature*

- Communications*, 11(1), Article 5240. <https://doi.org/10.1038/s41467-020-19010-6>, PubMed: 33067457
- Gage, N. M., & Hickok, G. (2005). Multiregional cell assemblies, temporal binding and the representation of conceptual knowledge in cortex: A modern theory by a “classical” neurologist, Carl Wernicke. *Cortex* 41(6), 823–832. [https://doi.org/10.1016/S0010-9452\(08\)70301-0](https://doi.org/10.1016/S0010-9452(08)70301-0), PubMed: 16353368
- Grave, E., Bojanowski, P., Gupta, P., Joulin, A., & Mikolov, T. (2019). Learning word vectors for 157 languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. European Language Resources Association (ELRA).
- Harris, Z. S. (1954). Distributional structure. *WORD*, 10(2–3), 146–162. <https://doi.org/10.1080/00437956.1954.11659520>
- Hollis, G., & Westbury, C. (2016). The principals of meaning: Extracting semantic dimensions from co-occurrence models of semantics. *Psychonomic Bulletin & Review*, 23(6), 1744–1756. <https://doi.org/10.3758/s13423-016-1053-2>, PubMed: 27138012
- Honnibal, M. (2017). *spaCy2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing*. Sentometrics Research. <https://spacy.io/>
- Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E., & Gallant, J. L. (2016). Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532(7600), 453–458. <https://doi.org/10.1038/nature17637>, PubMed: 27121839
- Jefferies, E., Crisp, J., & Lambon Ralph, M. A. (2006). The impact of phonological or semantic impairment on delayed auditory repetition: Evidence from stroke aphasia and semantic dementia. *Aphasiology*, 20(9), 963–992. <https://doi.org/10.1080/02687030600739398>
- Keuleers, E., Brysbaert, M., & New, B. (2010). SUBTLEX-NL: A new measure for Dutch word frequency based on film subtitles. *Behavior Research Methods*, 42(3), 643–650. <https://doi.org/10.3758/BRM.42.3.643>, PubMed: 20805586
- Keuleers, E., Stevens, M., Mandra, P., & Brysbaert, M. (2015). Word knowledge in the crowd: Measuring vocabulary size and word prevalence in a massive online experiment. *Quarterly Journal of Experimental Psychology*, 68(8), 1665–1692. <https://doi.org/10.1080/17470218.2015.1022560>, PubMed: 25715025
- Kotz, S. A., & Paulmann, S. (2011). Emotion, language, and the brain. *Language and Linguistics Compass*, 5(3), 108–125. <https://doi.org/10.1111/j.1749-818X.2010.00267.x>
- Kousta S.-T., Vigliocco, G., Vinson, D. P., Andrews, M., & Del Campo, E. (2011). The representation of abstract words: Why emotion matters. *Journal of Experimental Psychology: General*, 140(1), 14–34. <https://doi.org/10.1037/a0021446>, PubMed: 21171803
- Kuperman, V., Estes, Z., Brysbaert, M., & Warriner, A. B. (2014). Emotion and language: Valence and arousal affect word recognition. *Journal of Experimental Psychology: General*, 143(3), 1065–1081. <https://doi.org/10.1037/a0035669>, PubMed: 24490848
- Liuzzi, A. G., Bruffaerts, R., Dupont, P., Adamczuk, K., Peeters, R., De Deyne, S., Storms, G., & Vandenberghe, R. (2015). Left perirhinal cortex codes for similarity in meaning between written words: Comparison with auditory word input. *Neuropsychologia*, 76, 4–16. <https://doi.org/10.1016/j.neuropsychologia.2015.03.016>, PubMed: 25795039
- Liuzzi, A. G., Bruffaerts, R., Peeters, R., Adamczuk, K., Keuleers, E., De Deyne, S., Storms, G., Dupont, P., & Vandenberghe, R. (2017). Cross-modal representation of spoken and written word meaning in left pars triangularis. *NeuroImage*, 150, 292–307. <https://doi.org/10.1016/j.neuroimage.2017.02.032>, PubMed: 28213115
- Liuzzi, A. G., Dupont, P., Peeters, R., Bruffaerts, R., De Deyne, S., Storms, G., & Vandenberghe, R. (2019). Left perirhinal cortex codes for semantic similarity between written words defined from cued word association. *NeuroImage*, 191, 127–139. <https://doi.org/10.1016/j.neuroimage.2019.02.011>, PubMed: 30753925
- Louwerse, M. M. (2011). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, 3(2), 273–302. <https://doi.org/10.1111/j.1756-8765.2010.01106.x>, PubMed: 25164297
- Louwerse, M. M., & Jeuniaux, P. (2010). The linguistic and embodied nature of conceptual processing. *Cognition*, 114(1), 96–104. <https://doi.org/10.1016/j.cognition.2009.09.002>, PubMed: 19818435
- Lynott, D., Connell, L., Brysbaert, M., Brand, J., & Carney, J. (2020). The Lancaster Sensorimotor Norms: Multidimensional measures of perceptual and action strength for 40,000 English words. *Behavior Research Methods*, 52(3), 1271–1291. <https://doi.org/10.3758/s13428-019-01316-z>, PubMed: 31832879
- Marian, V., Bartolotti, J., Chabal, S., & Shook, A. (2012). Clearpond: Cross-linguistic easy-access resource for phonological and orthographic neighborhood densities. *PLOS ONE*, 7(8), Article e43230. <https://doi.org/10.1371/journal.pone.0043230>, PubMed: 22916227
- McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37(4), 547–559. <https://doi.org/10.3758/BF03192726>, PubMed: 16629288
- Meersmans, K., Bruffaerts, R., Jamouille, T., Liuzzi, A. G., De Deyne, S., Storms, G., Dupont, P., & Vandenberghe, R. (2020). Representation of associative and affective semantic similarity of abstract words in the lateral temporal perisylvian language regions. *NeuroImage*, 217, Article 116892. <https://doi.org/10.1016/j.neuroimage.2020.116892>, PubMed: 32371118
- Meersmans, K., Storms, G., De Deyne, S., Bruffaerts, R., Dupont, P., & Vandenberghe, R. (2022). Orienting to different dimensions of word meaning alters the representation of word meaning in early processing regions. *Cerebral Cortex*, 32(15), 3302–3317. <https://doi.org/10.1093/cercor/bhab416>, PubMed: 34963135
- Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM*, 38(11), 39–41. <https://doi.org/10.1145/219717.219748>
- Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K.-M., Malave, V. L., Mason, R. A., & Just, M. A. (2008). Predicting human brain activity associated with the meanings of nouns. *Science*, 320(5880), 1191–1195. <https://doi.org/10.1126/science.1152876>, PubMed: 18511683
- Moors, A., De Houwer, J., Hermans, D., Wanmaker, S., van Schie, K., Van Harmelen A.-L., De Schryver, M., De Winne, J., & Brysbaert, M. (2013). Norms of valence, arousal, dominance, and age of acquisition for 4,300 Dutch words. *Behavior Research Methods*, 45(1), 169–177. <https://doi.org/10.3758/s13428-012-0243-8>, PubMed: 22956359
- Murphy, K., Bodurka, J., & Bandettini, P. A. (2007). How long to scan? The relationship between fMRI temporal signal to noise ratio and necessary scan duration. *NeuroImage*, 34(2), 565–574. <https://doi.org/10.1016/j.neuroimage.2006.09.032>, PubMed: 17126038
- Obleser, J. (2014). Putting the listening brain in context. *Language and Linguistics Compass*, 8(12), 646–658. <https://doi.org/10.1111/Inc3.12098>
- Obleser, J., & Kotz S. A. (2010). Expectancy constraints in degraded speech modulate the language comprehension network. *Cerebral Cortex*, 20(3), 633–640. <https://doi.org/10.1093/cercor/bhp128>, PubMed: 19561061

- Oosterhof, N. N., Connolly, A. C., & Haxby, J. V. (2016). CoSMoMVA: Multi-modal multivariate pattern analysis of neuroimaging data in Matlab/GNU octave. *Frontiers in Neuroinformatics, 10*, Article 27. <https://doi.org/10.3389/fninf.2016.00027>, PubMed: 27499741
- Optoacoustics. (2022). *OptoActive II headphones and microphones* [Apparatus]. <https://www.optoacoustics.com/medical/>
- Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The measurement of meaning*. University of Illinois Press Urbana.
- Paivio, A. (1991). Dual coding theory: Retrospect and current status. *Canadian Journal of Psychology/Revue Canadienne de Psychologie, 45*(3), 255–287. <https://doi.org/10.1037/h0084295>
- Pauligk, S., Kotz, S. A., & Kanske, P. (2019). Differential impact of emotion on semantic processing of abstract and concrete words: ERP and fMRI evidence. *Scientific Reports, 9*(1), Article 14439. <https://doi.org/10.1038/s41598-019-50755-3>, PubMed: 31594966
- Pereira, F., Botvinick, M., & Detre, G. (2013). Using Wikipedia to learn semantic feature representations of concrete concepts in neuroimaging experiments. *Artificial Intelligence, 194*, 240–252. <https://doi.org/10.1016/j.artint.2012.06.005>, PubMed: 23243317
- Pereira, F., Lou, B., Pritchett, B., Ritter, S., Gershman, S. J., Kanwisher, N., Botvinick, M., & Fedorenko, E. (2018). Toward a universal decoder of linguistic meaning from brain activation. *Nature Communications, 9*(1), Article 963. <https://doi.org/10.1038/s41467-018-03068-4>, PubMed: 29511192
- Recchia, G., & Louwerse, M. M. (2015). Reproducing affective norms with lexical co-occurrence statistics: Predicting valence, arousal, and dominance. *Quarterly Journal of Experimental Psychology, 68*(8), 1584–1598. <https://doi.org/10.1080/17470218.2014.941296>, PubMed: 24998307
- Schuster, S., Himmelstoss, N. A., Hutzler, F., Richlan, F., Kronbichler, M., & Hawelka, S. (2021). Cloze enough? Hemodynamic effects of predictive processing during natural reading. *NeuroImage, 228*, Article 117687. <https://doi.org/10.1016/j.neuroimage.2020.117687>, PubMed: 33385553
- Scott, S. K., Blank, C. C., Rosen, S., & Wise, R. J. (2000). Identification of a pathway for intelligible speech in the left temporal lobe. *Brain, 123*(12), 2400–2406. <https://doi.org/10.1093/brain/123.12.2400>, PubMed: 11099443
- Tang, D., Wei, F., Qin, B., Yang, N., Liu, T., & Zhou, M. (2016). Sentiment embeddings with applications to sentiment analysis. *IEEE Transactions on Knowledge and Data Engineering, 28*(2), 496–509. <https://doi.org/10.1109/TKDE.2015.2489653>
- Troche, J., Crutch, S. J., & Reilly, J. (2017). Defining a conceptual topography of word concreteness: Clustering properties of emotion, sensation, and magnitude among 750 English words. *Frontiers in Psychology, 8*, Article 1787. <https://doi.org/10.3389/fpsyg.2017.01787>, PubMed: 29075224
- Vankrunkelsven, H., Verheyen, S., Storms, G., & De Deyne, S. (2018). Predicting lexical norms: A comparison between a word association model and text-based word co-occurrence models. *Journal of Cognition, 1*(1), Article 45. <https://doi.org/10.5334/joc.50>, PubMed: 31517218
- Van Rensbergen, B., De Deyne, S., & Storms, G. (2016). Estimating affective word covariates using word association data. *Behavior Research Methods, 48*(4), 1644–1652. <https://doi.org/10.3758/s13428-015-0680-2>, PubMed: 26511372
- Van Rensbergen, B., Storms, G., & De Deyne, S. (2015). Examining assortativity in the mental lexicon: Evidence from word associations. *Psychonomic Bulletin & Review, 22*(6), 1717–1724. <https://doi.org/10.3758/s13423-015-0832-5>, PubMed: 25893712
- Vigliocco, G., Kousta S.-T., Della Rosa, P. A., Vinson, D. P., Tettamanti, M., Devlin, J. T., & Cappa, S. F. (2014). The neural representation of abstract words: The role of emotion. *Cerebral Cortex, 24*(7), 1767–1777. <https://doi.org/10.1093/cercor/bht025>, PubMed: 23408565
- Vincze, N., & Bestgen, Y. (2011). Une procédure automatique pour étendre des normes lexicales par l'analyse des cooccurrences dans des textes. *Traitement Automatique des Langues, 52*(3), 191–216.
- Wang, X., Wu, W., Ling, Z., Xu, Y., Fang, Y., Wang, X., Binder, J. R., Men, W., Gao J.-H., & Bi, Y. (2018). Organizational principles of abstract words in the human brain. *Cerebral Cortex, 28*(12), 4305–4318. <https://doi.org/10.1093/cercor/bhx283>, PubMed: 29186345
- Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods, 45*(4), 1191–1207. <https://doi.org/10.3758/s13428-012-0314-x>, PubMed: 23404613
- Westbury, C., Keith, J., Briesemeister, B. B., Hofmann, M. J., & Jacobs, A. M. (2015). Avoid violence, rioting, and outrage; approach celebration, delight, and strength: Using large text corpora to compute valence, arousal, and the basic emotions. *Quarterly Journal of Experimental Psychology, 68*(8), 1599–1622. <https://doi.org/10.1080/17470218.2014.970204>, PubMed: 26147614
- Wilce, J. (2009). *Language and emotion*. Cambridge University Press.
- Wu, Z., & Palmer, M. (1994). Verb semantics and lexical selection. In *Proceedings of the 32nd annual meeting of the Association for Computational Linguistics* (pp. 133–138). Association for Computing Machinery. <https://doi.org/10.3115/981732.981751>
- Yu, L.-C., Wang, J., Lai, K. R., & Zhang, X. (2018). Refining word embeddings using intensity scores for sentiment analysis. *IEEE/ACM Transactions on Audio, Speech, Language Processing, 26*(3), 671–681. <https://doi.org/10.1109/TASLP.2017.2788182>