

How direct is the link between words and images?

Hassan Shahmohammadi, Maria Heitmeier,
Elnaz Shafaei-Bajestan, Hendrik P.A. Lensch, and
R. Harald Baayen
University of Tübingen

Günther et al. (2022) investigated the relationship between words and images in which they concluded the possibility of a direct link between words and embodied experience. In their study, participants were presented with a target noun and a pair of images, one chosen by their model and another chosen randomly. Participants were asked to select the image that best matched the target noun. Building upon their work, we addressed the following questions. 1. Apart from utilizing visually embodied simulation, what other strategies subjects might have used? How much does this setup rely on visual information? Can it be solved using textual representations? 2. Do current visually-grounded embeddings explain subjects' selection behavior better than textual embeddings? 3. Does visual grounding improve the representations of both concrete and abstract words? For this aim, we designed novel experiments based on pre-trained word embeddings. Our experiments reveal that subjects' selection behavior is explained to a large extent on text-based embeddings and word-based similarities. Visually grounded embeddings offered modest advantages over textual embeddings in certain cases. These findings indicate that the experiment by Günther et al. (2022) may not be well suited for tapping into the perceptual experience of participants, and the extent to which it measures visually grounded knowledge is unclear.

Keywords: visual grounding, word embeddings, grounded cognition

1. Introduction

How do knowledge and specifically the knowledge of word meanings emerge? For more than 2000 years, it was thought that mental representations of meanings at least resemble the perceptual states producing them (Barsalou, 1999; Zwaan

and Madden, 2005). Though heavily debated in the early to mid 20th century (overview in Barsalou, 1999), grounded cognition has regained acceptance over the last decades (Barsalou, 2010) and is especially prominent in cognitive linguistics (e.g. Lakoff, 1987; Langacker, 1999). It is supported by a range of empirical studies reporting evidence in favour of an embodied view of cognition (e.g. Simmons et al., 2005; Martin, 2007; Goldstone, 1995; Solomon and Barsalou, 2001, 2004; Barsalou, 2008).

Nevertheless, over the last 30 years, computational models of meaning called Distributed Semantic Models (DSM) have emerged, which derive meaning representations entirely from text corpora, i.e., from language alone. They are based on the hypothesis that words occurring in the same context are semantically related (Harris, 1954). Prominent examples are Latent Semantic Analysis (LSA; Landauer and Dumais, 1997), word2vec (Mikolov et al., 2013) and GloVe (Pennington et al., 2014). DSM have been applied widely and successfully not only in Natural Language Processing (NLP; e.g. Wang et al., 2019), but also in cognitive science (overview in Günther et al., 2019). For example, they have been shown to be able to predict aspects of brain imaging data such as fMRI and EEG (Bulat et al., 2017; Hollenstein et al., 2019; Anderson et al., 2015), explain variance in a range of behavioural data (Mandera et al., 2017; Westbury, 2014; Westbury and Hollis, 2019; Marelli and Amenta, 2018), and even reflect relations between colours and geographical locations respectively (Abdou et al., 2021; Louwerse and Zwaan, 2009).

All these purely text-based methods, while providing a practical way for dealing with largescale data, suffer from an obvious limitation often referred to as the symbol grounding problem (Harnad, 1990). That is, the meanings of words are solely based on other words without links to the outside world. Moreover, they also have practical problems; for example, because of the fundamental assumption of DSM, antonyms belonging to the same topical class (e.g., *small* and *big*) typically end up very close together in purely text-based vector spaces (see, e.g., Shahmohammadi et al., 2021, and references cited there). As a consequence, applications for, e.g., sentiment analysis, cannot well distinguish between “it was a good movie” and “it was a bad movie”¹.

These issues, as well as the empirical findings regarding embodied views of cognition, have sparked interest in grounding word meaning in experience. One approach to gain insight into speakers’ conceptual knowledge and to obtain such representations is to simply ask participants what they associate with a number of concepts. McRae et al. (2005) and Buchanan et al. (2019), among others, asked participants to provide perceptual, functional, and encyclopedic features for a

1. For recent progress for sentential negation, see Anschütz et al. (2023).

range of target words, resulting in so-called feature norms. Lynott et al. (2020) targeted the sensorimotor experience with words more specifically by asking participants to rate words on 11 sensorimotor dimensions, including auditory, visual, and interoceptive experience, but also the extent to which they involve body parts such as arms, legs, or heads. Both of these norms have been useful in cognitive science and used to, e.g., capture the semantic richness of concepts in memory (e.g. Grondin et al., 2009; Buchanan et al., 2019). Nevertheless, such norms do not necessarily mirror mental representations exactly, as only especially salient features may be recalled or only the features of a particular memory might be elicited (e.g. Cree and McRae, 2003; Barsalou, 2003; Buchanan et al., 2019).

A second approach has been to utilise image data to integrate visual information into word embeddings. Visual information can be extracted from images by means of image classification models. From these models, features can be extracted representing a condensed, distributed representation of an image. Baroni (2016) differentiates between two methods of grounding textual information in vision: *multimodal fusion* and *cross-modal mapping*.

Multimodal fusion methods range from simple concatenation (Kiela and Bottou, 2014; Rotaru and Vigliocco, 2020a; Utsumi, 2022), to projecting both into a shared space (Silberer and Lapata, 2014; Hasegawa et al., 2017; Kiela et al., 2018; Chrupała et al., 2015), or aligning textual embeddings with image information (Shahmohammadi et al., 2021; Bordes et al., 2019). Multimodal fusion approaches are successful in NLP (Bruni et al., 2014; Shahmohammadi et al., 2021; Bordes et al., 2019), as well as showing some promise when it comes to predicting cognitive phenomena (e.g. Bulat et al., 2017; Lazaridou et al., 2016, 2017; De Deyne et al., 2021; Anderson et al., 2015; Rotaru and Vigliocco, 2020a; Utsumi, 2022). However, there remain a number of open questions, such as how much image information should be allowed to be fused into textual embeddings, whether image information should be linked to individual words or words in context, and how much grounding helps when a lot of textual training data is available. These questions have been explored extensively in Shahmohammadi et al. (2023, henceforth Shahmohammadi et al.). They propose a new method for grounding textual embeddings in vision. This method, a multimodal fusion model (Baroni, 2016), is trained on full image descriptions instead of single words. Hence, it leverages both the textual context and all visual information. After the training phase, the algorithm can be used to obtain a visually grounded representation of any single target word even if it has not been seen during training (so-called “zero-shot”). Accordingly, we will refer to this model as “ZSG”, Zero-shot Grounded, for the remainder of this paper. In Shahmohammadi et al., ZSG was evaluated on a range of different tasks. This allowed the authors to investigate several questions related to how much visual information should be

allowed to fuse into textual embeddings, how much training on full sentences instead of single words benefits the grounding process, and finally, how the amount of textual training data available for a particular task influences the benefit of visual grounding.

Cross-modal mapping, on the other hand, involves mapping between visual and textual information, for instance, to model how visual concepts can be translated into language. One example of such a model is proposed in a study by Günther, Petilli, Vergallito, and Marelli (2020), henceforth GPVM, who employ a simple linear mapping to map from textual to visual information. Their model is designed to account for the grounding of both concrete and abstract nouns. They trained a linear mapping to map from language-based textual embeddings to visual embeddings of concrete nouns. Then, they mapped the text embeddings of a set of target nouns, including unseen concrete and abstract nouns, to the vision domain. They selected the image that was most similar to the generated image vector as the image that corresponds to the target noun. Thus they could predict images also and critically for unseen abstract nouns. In the following experiments, the images provided by this setup were compared with random control images by asking participants to select the image which best represented the target noun. They concluded that their model was able to account for the visual grounding of abstract words but that concreteness was nevertheless a significant predictor of model performance, i.e. the model performed much better on concrete than on abstract words.

GPVM's model and subsequent behavioural experiments raise a number of questions and issues of how textual representations should be grounded in vision, which deserve further investigation. The model shows how knowledge based on language can be translated into knowledge based on vision. However, previous work suggests that rather than mapping from one modality to the other, combining the two sources of information into a multimodal representation might be preferable. Notably, Shahmohammadi et al. conducted various experiments comparing cross-modal mapping accounts such as the one by GPVM to more sophisticated ones of multimodal fusion and found that a simple linear mapping from language to vision (which they call a "Word-Level" model) results in representations which are much worse at predicting human relatedness and similarity judgments compared to a number of different multimodal fusion architectures (compared to ZSG, their best model, the Word-Level model was on average less accurate by 30 percentage points). In other words, the vision representations predicted from textual embeddings in a model similar to the one in GPVM do not perform well in tasks assessing words' semantic similarity and relatedness (as perceived by human raters, such as the MEN task, Bruni et al., 2014).

Nevertheless, the behavioural data collected by GPVM provides insights into speakers' semantic representations of words, specifically highlighting the role of visual information available in their experiments. In the present study, we therefore take the behavioural experiment by GPVM as point of departure and conduct novel computational experiments to evaluate the multimodal embeddings by Shahmohammadi et al. compared to a baseline of simple textual embeddings. This is especially worthwhile since both the model underlying the behavioural experiment of GPVM and the ZSG model claim to be able to ground unseen abstract words into vision. Using current multi-modal embeddings to model participants' behavior provides useful insights for analyzing the interplay between vision and language in human cognition.

In contrast to GPVM, who used their behavioural experiment as a simple verification of their model, we want to model participants' responses. As a first step, we need a task analysis, working out what participants might actually be doing in the experiment. GPVM explicitly state that their model is not a processing model, i.e., they do not describe how and whether participants in their experiment actually access visual information. For instance, it is unclear whether GPVM infer from their experimental results that abstract words evoke images in the mind, which subsequently might influence lexical processing. However, their mapping model appears to suggest that participants translate language representations to visual representations in order to solve the task of which presented image better matches the presented target word.

A number of both behavioural and neuropsychological studies has explored to what extent humans actually generate mental images in knowledge retrieval. It was found, for example, that perceptual characteristics (e.g., size) best accounted for verification times and errors. This has been taken to indicate that participants simulated concepts visually to verify their properties (Solomon and Barsalou, 2004). Further evidence comes from lesioning studies: it was found that damage to certain areas in the brain increased the probability of losing categories relying on that area. For example, if visual areas are damaged, patients are more likely to lose the animal category since it relies on visual processing (overview in Barsalou, 2008). However, this is rarely interpreted as providing evidence for concrete images being created "inside the brain". "[T]he claim is not that there are pictures in the mind. Rather, the claim is that traces of visual and other experiences are (partly) reactivated and recombined in novel ways by associated words" (Zwaan and Madden, 2005, p.242). Or, in the words of Barsalou (1999, p.582), "[perceptual symbols] are records of the neural states that underlie perception". And indeed, Solomon and Barsalou (2004) found that participants only used perceptual simulation in a feature verification task if the task could not be solved based on association alone (but see Barsalou et al. (2008, p.267): in a

task where participants were presented with words for 5 seconds, and only then asked whether the word applied to an image that was next presented, involvement of the simulation system was found).

In the present paper, therefore, we aim to answer three questions:

1. How can we explain the data of GPVM without assuming that participants generate mental images?

Specifically, we propose that rather than actually translating textual representations into visual representations, as suggested by the model of GPVM, participants might identify the objects in the images and associate the object names with the target word. The image with closer associations to the target word would then be selected as the image that better fits the target word. In GPVM, similarity judgments appear to take place at the level of visual similarity only. We on the other hand propose that rather than at the visual level, similarity judgments occur within a semantic space influenced by our understanding of words through their co-occurrence and potential visual attributes. For example, in one of the examples by GPVM (see Table 1), the model prediction shows an image with three children, while the other random image depicts a green plant. The target word is “childhood”. Based on the objects “children” and “plant” it can be predicted that participants will choose the image depicting “children”, as it is clearly more associated with “childhood” than “plant”. Solving the task at hand can therefore in principle be accomplished without ever generating an internal image of “childhood”.

2. As discussed above, previous research nevertheless indicates that word embeddings grounded in vision are better at predicting behavioural data than purely textual embeddings (e.g. Bruni et al., 2014; Shahmohammadi et al., 2023). Our understanding of the task still warrants participants to make comparisons between words (albeit in semantic rather than in visual space), and here, it will be interesting to see whether their responses are best predicted by purely textual (GloVe/Word2Vec; Pennington et al., 2014; Mikolov et al., 2013) or by the multimodal embeddings of Shahmohammadi et al. (2023).

Thus, assuming that the alternative explanation proposed above finds support, a central question remains: Is participants’ behaviour best accounted for by purely textual or multimodal word embeddings?

3. The last question relates to the treatment of abstract words in the process of visual grounding. Abstract words are differentiated from concrete words by the (un)availability of their denotations to the human senses. Concepts become gradually more abstract as they are separated further from sensible physical entities and become more associated with mental states (Barsalou,

2003). Evidently, obtaining images for an abstract word such as *malice* is much more difficult, if at all possible, than for a straightforwardly concrete word such as *apple*. Accordingly, image corpora usually provide images exclusively for common concrete objects. Nonetheless, some abstract concepts can be elicited from the contextual situations in which these objects occur.

In the grounding approach by Shahmohammadi et al. (2023), initial grounding takes place for concrete words, after which other words, including abstract words, can also be grounded using zero-shot learning. They show that zero-shot learning of visual grounding is highly effective, not only for concrete words but also for abstract words. In this study, we ask: (1) Does this finding replicate using the data of GPVM? If so, (2) Does the indirect grounding of abstract words afford a better understanding of the experimental results reported by GPVM?

These results are particularly interesting given the extensive evidence from case reports and behavioral and neural studies suggesting that abstract and concrete words are processed differently, involving overlapping but distinct brain regions (see Montefinese, 2019; Mkrtychian et al., 2019, for reviews).

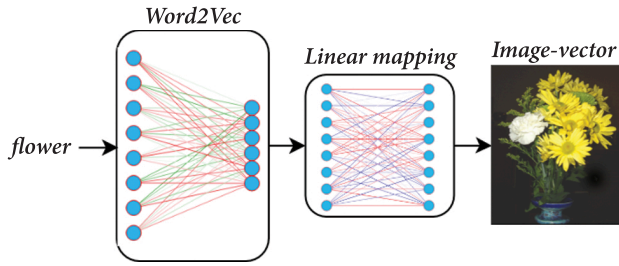
The rest of the paper is structured as follows: in Section 2 we first introduce in detail the model and experiment presented by GPVM, followed by a description of the grounding model of Shahmohammadi et al. Subsequently, in Section 3 we model the experiment utilizing both textual and grounded word embeddings, aiming to answer the three questions formulated above. In Section 4 we discuss our results.

2. Methodology

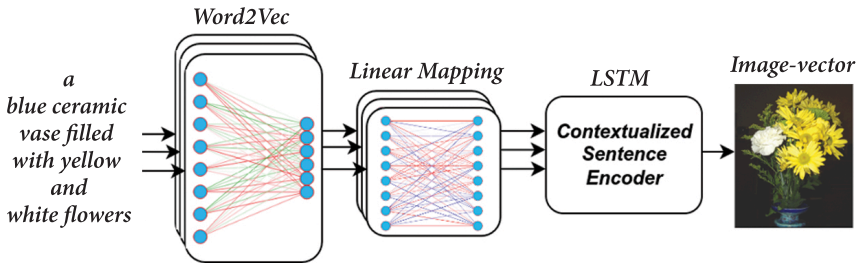
This section details the model and the behavioral experiment by GPVM and briefly explains the ZSG grounding approach proposed by Shahmohammadi et al.

2.1 Materials from Günther et al. (2022)

GPVM proposed a grounding model that combines vision and language information. It maps textual representations (obtained from a pre-trained *Word2Vec* model, Mikolov et al., 2013) onto image vectors (obtained from VGG-F, a pre-trained image classification model, Chatfield et al., 2014) via a single linear mapping (see Figure 1a). It is first trained on a set of isolated words for which images are available in ImageNet (Deng et al., 2009) and later tested on both concrete and abstract words which did not occur in the training set. For instance,



a. Cross-modal mapping model by Günther et al. (2022)



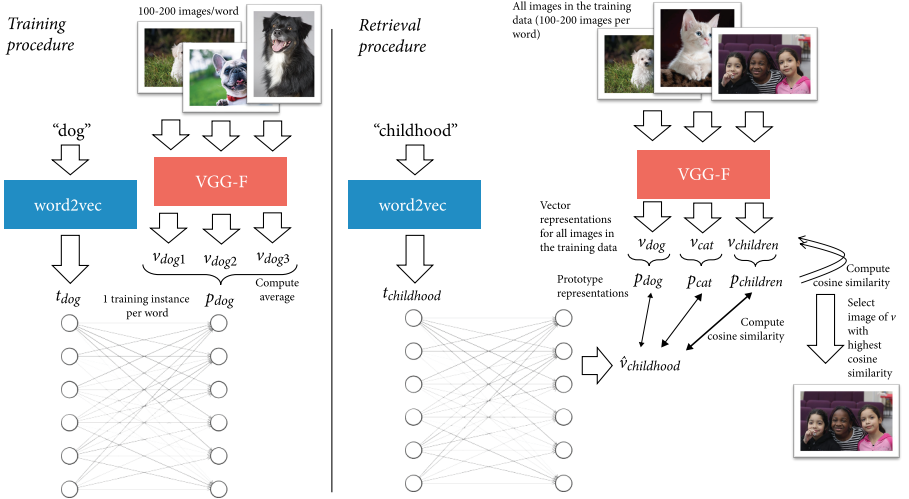
b. Multimodal fusion model by Shahmohammadi et al. (2023)

Figure 1. Comparison of the cross-modal mapping model by Günther et al. (2022) and the multimodal fusion model proposed by Shahmohammadi et al. (2023). The latter model takes the context into account and first applies the mapping matrix M to the textual embeddings of all words in the context. Then, it applies a contextualized sentence encoder to predict the image vector

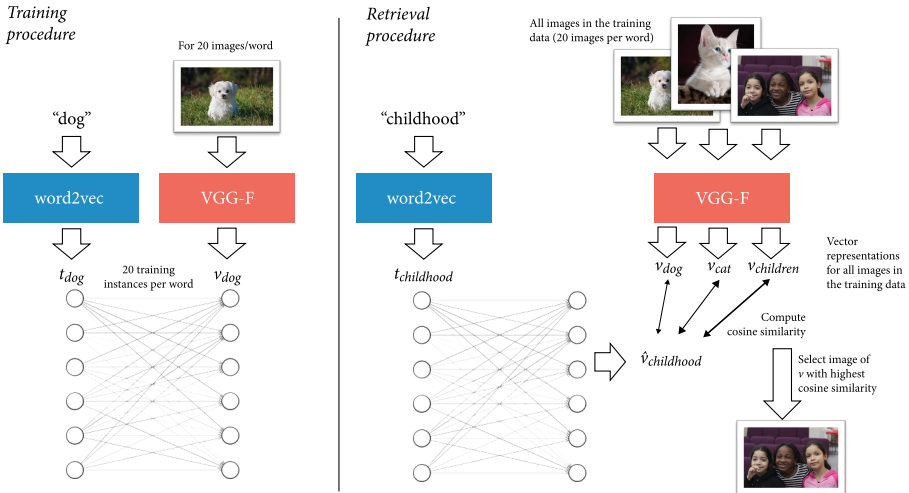
it is first trained to predict an image vector of a *dog*, given a word vector of *dog*. The trained model is then used to generate visual representations for unseen words including abstract words such as *jealousy* or *childhood*.

GPVM trained two versions of their model: a prototype model, where for each word, the image representations of 100 to 200 images (depending on how many were available in ImageNet, Deng et al., 2009) were averaged to obtain a “prototype” representation, and an exemplar model, for which the model was trained on 20 different image representations per word.

GPVM tested their model by predicting image representations for a range of both abstract and concrete nouns. Since the model is not able to actually generate images, they simply selected existing images whose representations were as close as possible to the predicted image representation. For the exemplar model, they straightforwardly selected the image from the set of training images closest to the predicted image representation. For the prototype model, on the other hand, they first selected the prototype image vector closest to the predicted representation



a. Prototype model



b. Exemplar model

Figure 2. Training and retrieval procedures in the prototype and exemplar models in Günther et al. (2022). t indicates textual embeddings, v image representations as generated by VGG-F (Chatfield et al., 2014) and p prototype image representations. Images are for illustration purposes only

and then selected the training image closest to this prototype image vector. The two training and retrieval strategies are visualised in Figure 2. Examples of generated (selected) images can be seen in the center column of Table 1.

In their first two experiments, GPVM asked participants to select one of two images for each of the abstract and concrete nouns: either the image predicted by the model, or a random control image (see Table 1 for examples of predicted and random control images). Two variables were controlled for: the *concreteness* (*concrete* vs. *abstract*) of the nouns (Brysbaert et al., 2014), as well as the number of semantic neighbours which had an associated image in ImageNet (*visual neighbours*). *Visual neighbours* are divided into no (*far* condition), few (*near*) and many visual neighbours (*maximum*).





The participants selected the predicted images above chance level for concrete nouns and abstract nouns with many visual neighbours, and for the exemplar model also for abstract nouns with no visual neighbours. Both *concreteness* and *visual neighbours* are correlated, and they were both found to be predictive of the participant's performance only in the prototype model: the more concrete the noun, and the more *visual neighbours* it had, the more likely the participants were to pick the image predicted by the model. These two variables failed to be predictive for the exemplar model. GPVM propose that a possible reason for this is that the exemplar model, being presented with many images for each word instead of just a single average representation as in the prototype model, is able to pick up more "idiosyncratic visual information" than the prototype model in the abstract and far conditions, thereby removing any effect of concreteness and visual neighbours. By further disentangling the variables of *visual neighbours* and *concreteness* in a third experiment with the prototype model where both were represented as continuous variables, the authors found that only *concreteness* was significant at predicting participants' accuracy, but not the number of *visual neighbours*. They, therefore, concluded from their experiment that (a) it is possible to predict image representations for unseen words from a purely language-based representation even for abstract words, and (b) that *concreteness* is a significant, graded predictor of model performance. This result is interesting insofar as the model relies on visual neighbours for predicting images for unseen words but its performance is nevertheless not significantly influenced by the number of visual neighbours.

In the present study, we restricted ourselves to data from experiments 1 and 2. They included 53 and 57 participants, respectively. In both experiments, participants were presented with 115 items (5 conditions with 23 items each). Additionally, they included ten catch trials where the target and random control images were selected manually. Target images were generated by their prototype and exemplar models as described above, while the control images were picked

Table 1. Examples of predicted (by the prototype model) and random control images for target words from various conditions of *concreteness* and *visual neighbours*. Table adapted from Günther et al. (2022), all images replaced by visually similar public domain images

Condition	Target word	predicted image	Random control image
concrete/ maximum	stallion		
concrete/ near	scout		
concrete/far	aspirin		

Table 1. (continued)

Condition	Target word	predicted image	Random control image
abstract/ near	childhood		
abstract/far	jealousy		

randomly from a set of images not included in the target images. Further details on the experimental materials can be found in Günther et al. (2022).

2.2 Model from Shahmohammadi et al. (2023)

Contrary to GPVM, the model from Shahmohammadi et al. takes the textual context of words into account in the grounding process. As illustrated in Figure 1b, they use image-caption pairs from an image captioning dataset (Lin et al., 2014). This dataset contains images paired with multiple human-annotated descriptions in the form of sentences. For example, instead of having the word *dog* associated with different pictures of dogs, they have access to multiple descriptions of scenes depicting dogs in various situations. Example sentences from their training data are *a dog leaping into the air to catch a frisbee* and *two dogs poking their heads through curtains at windows*. In the process of visually grounding, Shahmohammadi et al. use all the words in such sentences. As a consequence, words are still ‘aware’ of their textual co-occurrence patterns, allowing the model to predict the image vectors with higher fidelity. In what follows, we describe their proposed model in more detail.

For any textual word vector t_i (e.g., from GloVe, Pennington et al., 2014), the goal is to train a linear alignment/mapping \mathbf{M} to visually ground t_i as $g_i = t_i \cdot \mathbf{M}$, where g_i denotes the visually grounded version of t_i . The alignment \mathbf{M} ideally should: (a) preserve the essence of the statistical semantics captured from

textual corpora, and (b) align the textual word vectors with their corresponding visual features in images (Shahmohammadi et al., 2021, 2023). Given the properties of linear transformations, the grounded word vector g_i still respects the textual vector space while being informed about its corresponding perceptual properties in images.

The Microsoft_COCO_2017 dataset (Lin et al., 2014) was used to train such an alignment. This dataset is split into 118k train and 5k validation samples, with each sample consisting of a single image along with five different captions describing the image. Unlike many previous approaches (Collell Talleda et al., 2017; Günther et al., 2022), according to which a single textual word vector is mapped into its associated image features, Shahmohammadi et al. argue that textual context should be taken into account and not discarded, as it contains valuable information about the specific properties of images. For instance, instead of learning a model to map from the word *dog* to an image vector of a dog, their model maps a whole sentence such as *a dog is playing in the grass* into its associated image features. By taking the context *the grass* in the image caption into account, the model is given the opportunity to learn about the grass and to dissociate the dog from the grass. Each word vector is, therefore, aware (connected) of (to) its textual context while being aligned with its visual features. Here, we note that also the images are complex and seldom contain only a single isolated object. We also note that when presented with photographs, human viewers typically fixate at many different locations (Castelhano and Rayner, 2008). For complex pictures presented for 12 ms, the mean number of fixations can be as large as 32 (Cronin et al., 2020). Instead of taking human image interpretation as wholistic and undifferentiated, we take image understanding to be a rich and diversified process in which multiple objects in an image are scanned and interpreted.

For training the linear alignment/mapping \mathbf{M} , first, all the textual word vectors are piped through the alignment \mathbf{M} . Then, a contextualized Encoder (e.g., LSTM, Hochreiter and Schmidhuber, 1997) produces a single representation vector for the whole sentence while taking all the words and their relative order into account. The contextualized encoder is trained to map the representation of the whole sentence to the image features. In this setup, the encoder provides feedback on how words co-occur and how they should be rearranged to be aligned with the image features. After training the model, the trained alignment \mathbf{M} is used to ground all the textual word vectors, including the word vectors of abstract words. That is, given any textual word vector t_p , its visually grounded version g_i is easily computed as $g_i = t_p \cdot \mathbf{M}$. In this way, they created a visually grounded version of a given textual word embedding model.

Shahmohammadi et al. evaluated their grounded embeddings on word similarity/relatedness benchmarks, which have been commonly used for evaluation

of multi-modal embeddings (Park and Myaeng, 2017a; Kiros et al., 2018; Kiela et al., 2018; Collell Talleda et al., 2017). These benchmarks measure the relatedness/similarity of word pairs based on human-annotated scores. They used the following datasets for evaluation: WordSim353 (Finkelstein et al., 2001), MEN (Bruni et al., 2014), SimLex999 (Hill et al., 2015), Rare-Words (Luong et al., 2013), SimVerb3500 (Gerz et al., 2016), and MTurk771 (Halawi et al., 2012). Even though each of these benchmarks evaluates words' meaning from a particular perspective, their grounded embeddings boosted the mean score of GloVe (Pennington et al., 2014) textual embeddings across all benchmarks by almost 7 percentage points (56.7% to 63.6%). They further explored the benefit of grounding in further detail, leveraging from the SimLex999 dataset. This dataset categorizes words into different categories, such as 'highly abstract' and 'highly concrete'. Their grounding approach is beneficial for both concrete and abstract words and boosts results in other categories (e.g., *verbs* and *adjective*) as well (see Shahmohammadi et al., 2023, for further details).

While both Shahmohammadi et al. and GPVM apply a linear mapping to the textual embeddings, the former argues that bridging the gap between language and vision solely using a linear transformation is not ideal. They carried out multiple experiments with increasing complexity in terms of technique and network architecture. They concluded that the right balance is necessary to obtain high-quality embeddings that perform well on word similarity tasks. Their simplest approach, which they refer to as 'Word-Level', is similar to the model by GPVM in which textual word vectors of words in isolation (e.g., *dog*) are mapped to corresponding image vectors through a linear transformation. The grounded embeddings are then constructed by mapping the textual word vectors through the trained linear mapping. Evaluation on various word similarity benchmarks (e.g., MEN, Bruni et al., 2014), revealed that embeddings obtained in this way underperform severely (by > 30 percentage points) compared to purely textual embeddings. Applying very deep and complex models (e.g., applying multiple layers of LSTMs or transformer-based models (Kenton and Toutanova, 2019) for the mapping M) also resulted in a drop in performance. The right balance in terms of the depth and complexity of the network is therefore necessary. For further implementation details about their model and analysis, see Shahmohammadi et al. (2023).

2.3 Procedure

In order to clarify whether participants associated the target noun with object names depicted in the images rather than generating mental images and basing their decisions on comparisons of these images, we first extracted the names of

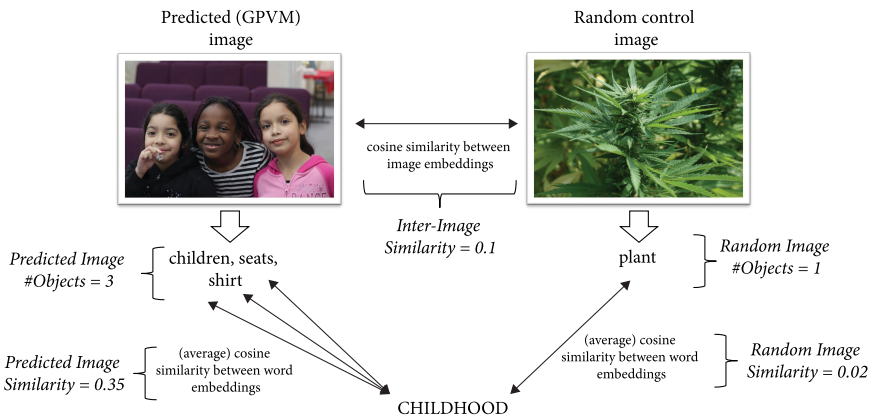


Figure 3. In Günther et al. (2022), participants are presented with two images (predicted and random control image) and a target word (here: childhood). They then have to select the image which better fits the target word. We calculate five measures: First, we use a CNN to automatically extract the object names visible in the predicted and random images and count them (*Predicted Image #Objects* and *Random Image #Objects*). Then we compute the average cosine similarity between the embeddings of the target word and the object names respectively (*Predicted Image Similarity* and *Random Image Similarity*). We also calculate the cosine similarity between the image embeddings of the two images (*Inter-image similarity*)

the objects (or labels) in both the predicted image and the random control image, using a pre-trained CNN model (Tan and Le, 2019). We then used the object names to obtain the corresponding word embeddings. In other words, our goal here is to retrieve the semantics of the objects in the images, and *not* to extract their orthographic written forms. The cognitive process that we are approximating with an engineering solution is the process of understanding what the objects in an image are. Note that empirical studies using eye-tracking to trace image interpretation show that images are typically scanned with many fixations at many different image locations (Castelhamo and Rayner, 2008; Cronin et al., 2020).

More specifically, for each image, we extracted the names of the top 10 classes² predicted by the CNN model. Examples of predicted classes for a particular image are ‘*bagel*’, ‘*plate*’, ‘*pretzel*’, ‘*dough*’, ‘*bakery*’, ‘*butternut_squash*’, ‘*cheeseburger*’, ‘*spaghetti_squash*’, ‘*chocolate_sauce*’, and ‘*acorn_squash*’. Here, the underscore represents a space character. For most of these very specific subcategories, no

2. The CNN classes include the true classes for all the images used by GPVM as they both utilized the same image database.

embeddings are available. As a result, the number of objects detected by our algorithm closely follows the number of objects in the images.

Then we modelled participants' behaviour using two approaches. In our first, very simple baseline approach (called "Max" in the remainder of this paper), we used the average cosine similarity between the embedding of the query (i.e., target) word and the embeddings of all the object embeddings detected in the predicted image, and in the control image, presented to a participant at a given trial. The average cosine similarity for a given image provides a measure of the likelihood of selecting that image for the given query word. This resulted in two measures:

- *Predicted Image Similarity*: The mean cosine similarity of the target word's embedding with the embeddings of the objects in the image predicted by the model of GPVM (henceforth the GPVM image);
- *Random Image Similarity*: The mean cosine similarity of the target word's embedding with the embeddings of the objects detected in the random control image.

To model human selection behavior, the image with the higher image similarity was selected as our model's choice. For example, if in trial 1, *Predicted Image Similarity* was higher than *Random Image Similarity*, we selected the GPVM image. We did this for the prototype and exemplar models separately. This cognitively rich model of how participants solve the experimental task contrasts with the lean, vision-only model of GPVM, who assume that a target query word makes contact with its corresponding embedding. That, in turn, generates an internal image (using a pre-trained mapping from word embeddings to images) that is subsequently compared with the two images presented to a participant, without any further involvement of higher cognitive processes evaluating what the objects present in images actually are.

In our second approach (in the following called GAM), we used the following additional predictors:

- *Inter-Image Similarity*: The mean cosine similarity between the GPVM image and the random control image vectors;
- *Predicted Image #Objects*: The number of object labels (e.g., dog, tree, ...) in the GPVM image for which word embeddings were available in the set of embeddings;
- *Random Image #Objects*: The number of objects labels in the random control image for which word embeddings were available in the set of embeddings.

An overview of the extracted measures can be found in Figure 3. Additionally, we used the two predictors provided by GPVM, which capture the number of

visual neighbours (Distance) and *concreteness* (WordType). All of these metrics were used as predictors in a Generalised Additive Model (GAM) with a logistic link function, as implemented in the *mgcv* package (Wood, 2011) in R. GAMs can model non-linear relationships between independent and dependent variables. We tested a range of different GAMs with these predictors, which we compared using the Akaike Information Criterion (AIC). While models with interactions between predictors (using tensor product smooths) gave a substantially better fit to the data, they also proved to be much harder to interpret. Therefore, we selected GAMs with main effects only for both the prototype and the exemplar models. We compared two sets of GAMs inspecting both AIC and predictions, one based on grounded vectors and one based on purely textual embeddings.³

3. Results

3.1 Q1: Can we model participant behaviour without assuming participants generate mental images?

3.1.1 *Max models*

Our hypothesis is that participants compare the meanings of the objects in the two images with the meaning of the target noun (using the respective embeddings), and select the image with the higher similarity. This idea is operationalized in our Max approach. We test Max using both textual and grounded GloVe and a version of (textual and grounded) Word2Vec (Mikolov et al., 2013) that was used by GPVM. The first minimal evaluation criterion for our approach is that it is able to differentiate the GPVM image from the random control image with a higher-than-chance probability. We found that this is indeed the case for the Max approach based on both textual and grounded, GloVe and Word2Vec embeddings, and prototype and exemplar setups (proportions test; $p < 0.0001$). This shows that our approach provides at least a theoretical possibility of how to solve the task.

Secondly, to investigate how well the Max approach approximates human behaviour, we measure the proportion where Max selected the GPVM image. We can thus view our four embedding types as *virtual participants*. If our hypothesis for predicting participants' selection behaviour holds true, we expect the virtual participants to show a similar preference (compared to participants' preference) for the GPVM image.

3. Generated measures and analysis notebooks can be found in the Supplementary Materials at <https://osf.io/7rxde/>.

The results of our experiments are reported in Table 2 for both the exemplar and prototype setups. We observe that the mean scores of the virtual participants are quite close to the mean scores of real participants (i.e., “Participants” in Table 2). The absolute difference between the mean score of the virtual participants and that of the participants is reported and labeled as Δ in the Table. Lower Δ values indicate a better fit for modeling the participants’ preferences.

Viewing our four embedding types as “virtual participants” begs the question of whether their performance fits into the distribution of human performance in GPVM. Figure 4 shows that the performance of most of the embedding types falls well within plausible participant performance across all categories of WordType and Distance. The clearest outlier is the model with textual GloVe embeddings in the abstract far category for the prototype setup. We will return to the question of grounded vs. textual embeddings and differences across concreteness/visual neighbour conditions below.

Thirdly, we expect our virtual participants to show the same effects of Word-Type and Distance as human participants. Unfortunately, the low number of data points per embedding type (114) made running individual, embedding-specific logistic regression models akin to the one by GPVM for human participants impossible. However, we do note that the (all nonsignificant) effects pointed in the same directions as for human participants for all embedding types for the prototype setup: both higher concreteness and more Distance lead to a higher probability of selecting the GPVM image. When combining the data of all four embedding types into one logistic regression model (without by-embedding random effects, since they prevented the model from converging), the effects still pointed in the same directions and were significant. In the exemplar setup, GPVM did not find a significant effect of any of the conditions other than a significant intercept (indicating that human participants generally performed above chance). The combined logistic regression model for all four embedding types based on the exemplar setup showed no effect of concreteness which is in line with findings by GPVM but did find a positive effect of Distance which was not found in GPVM (models in Supplementary Materials).

Fourthly, we can use the predictions of the Max model to directly predict participants’ behaviour. The results for the prototype experiment are reported in Table 3a. The bottom rows of both tables report the proportion of trials in which participants select the GPVM image over the random control image across categories in the two experiments by GPVM. Participants tended to select the GPVM image, the more concrete the target words were and the more visual neighbors they had (see also Figure 4). Since here we are interested in predicting participant behaviour rather than trying to differentiate the GPVM image and random control image, the numbers in the upper and lower parts of the two tables cannot

Table 2. Modeling participants' preference for the predicted images. The numbers in the top section of each table represent the percentage of trials in which each virtual participant chooses the GPVM image over the random image. The numbers in the last row of each table show the mean percentage of trials in which the GPVM image was selected, averaged over all human participants. The best results are marked in bold in each category. The absolute difference between the mean score of our model's prediction and that of the participants is reported and labeled as Δ in the Table. Lower Δ values indicate a better fit for modeling the participants' preferences

a. Prototype model

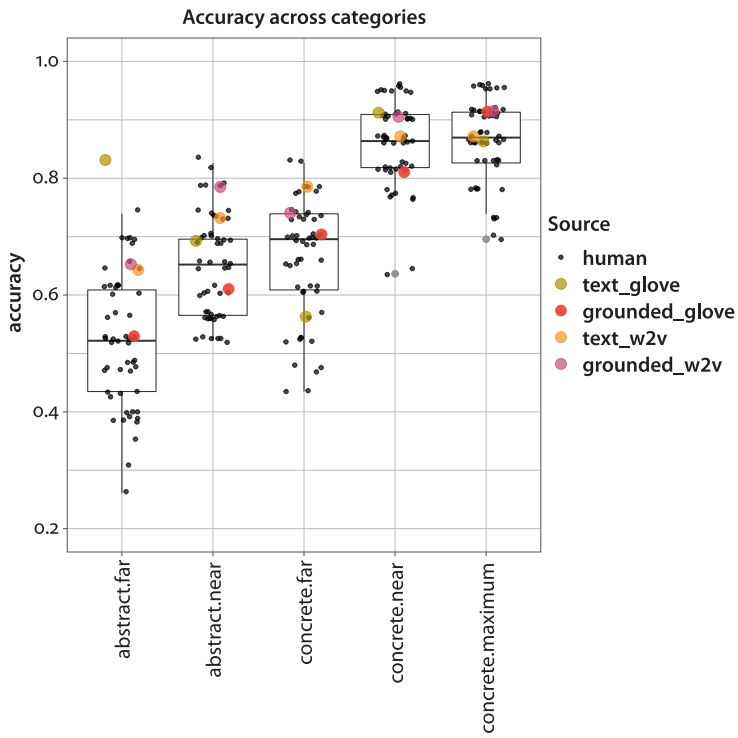
Embeddings	A. far	A. Near	C. Far	C. Near	C. Max	Mean (Δ)
Max: GloVe	82.61	69.57	56.52	90.91	86.96	77.31 (07.06)
Max: ZSG-GloVe	52.17	60.87	69.57	81.82	91.30	71.15 (00.90)
Max: W2V	65.22	73.91	78.26	86.36	86.96	78.14 (07.89)
Max: ZSG-W2V	65.22	78.26	73.91	90.91	91.30	79.92 (09.67)
Participants	52.00	64.00	66.00	84.25	85.00	70.25

b. Exemplar model

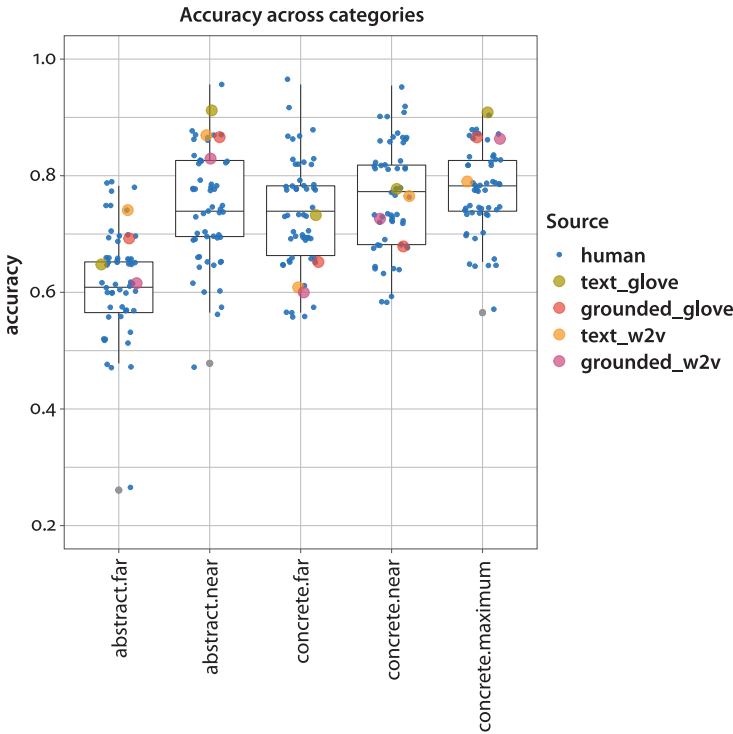
Embeddings	A. far	A. Near	C. Far	C. Near	C. Max	Mean (Δ)
Max: GloVe	65.22	91.30	73.91	77.27	91.30	79.80 (07.40)
Max: ZSG-GloVe	69.57	86.96	65.22	68.18	86.96	75.38 (02.98)
Max: W2V	73.91	86.96	60.87	77.27	78.26	75.45 (03.05)
Max: ZSG-W2V	60.87	82.61	60.87	72.73	86.96	72.81 (00.41)
Participants	62.00	74.00	73.00	76.00	77.00	72.40

be compared directly. Rather, we want our models to score as closely to 100% accuracy as possible, indicating a complete matching to participants' selections. Focusing on the mean performance across all concreteness/distance categories, Table 3a shows that participants' preferences are predicted fairly well (Mean accuracies for Max models are 68%, 70%, 69%, and 71% for textual and grounded GloVe and textual and grounded Word2Vec). Table 3b presents the results for the exemplar experiment. Mean accuracy for Max models tends to be somewhat lower compared to the prototype setup (70%, 66%, 67%, and 66% for textual and grounded GloVe and textual and grounded Word2Vec).

We can therefore conclude that it is indeed possible to model the behavioural experiment by GPVM without assuming that participants generate mental images



a. Prototype model



b. Exemplar model

Figure 4. Performance of the four embedding types compared to human participants in Günther et al. for the prototype and exemplar setup. Boxplots are based on human data points only. The performance of most embedding types is well within the range of human participants

and that even meaning representations based on textual information only are able to model participants' behaviour fairly well.

3.1.2 GAM models

Thus far, our predictions for participants' selection preferences have been based on average similarity scores for the images. Accuracy can be improved by also taking into account the similarity of the GPVM images and their controls, the number of objects in these images, as well as the two factorial predictors considered by GPVM: Distance and WordType. As mentioned above, we use logistic GAMs to obtain predictions for participants selection decisions. Using GAMs also enables us to investigate what effects *Predicted Image Similarity* and *Random Image Similarity* have on participant behaviour. In the prototype setup, accuracy

improved for all pairs of comparisons. For instance, the accuracy of Max: ZSG-GloVe, 70.20, improved for GAM: ZSG-GloVe to an accuracy of 72.19.

The GAM (see Table 4a) indicated that prototype images for concrete words were more often selected compared to abstract words, and that images for words with more image neighbors were also selected more often. These effects mirror those observed by GPVM.

The GAM also indicated (see Figure 5a) that a greater *Predicted Image Similarity* comes with a higher probability of selection. This effect aligns with our hypothesis that in this task, participants are scanning images for the visible objects, basing their decision on the match between these objects and the printed word stimulus.

A greater *Random Image Similarity* goes hand in hand with a lower probability of selection, but this effect is present only for higher similarity values. Although *Random Image Similarity* is in general lower than *Predicted Image Similarity* (ranges $(-0.06, 0.19)$ and $(-0.03, 0.42)$ respectively), *Predicted Image Similarity* leads to higher selection rates in the interval $(0.00, 0.10)$ whereas *Random Image Similarity* does not. This suggests that the objects in the GPVM images are more tightly and consistently interconnected, so that even for low similarity values they provide consistent evidence for selection.

A non-linear effect emerged for *Inter-image Similarity*. Setting aside the most extreme values of the predictor, this non-linear effect reduces to a U-shaped effect for where there is good data support. This U-shaped effect suggests that atypical similarities (values away from the mean) induced higher ratings. Apparently, the selection task induced an image scanning strategy that is based on whether the degree of similarity of a pair of images is remarkable and surprising. Both highly similar and very dissimilar images attract attention, resulting in more careful selection in favor of the GPVM image.

Results are subtly but informatively different for the exemplar setup. We first note that of the two factorial predictors, WordType was again supported, but Distance was not. This is in line with participant behaviour in the two setups: while they were more accurate for concrete target words than for abstract ones in both setups, the difference is much stronger in the prototype setup. Changing the kind of image – from prototype to exemplar – resulted in changed selection behavior. Apparently, training the mapping on image exemplars instead of on an averaged image results in more informative images in categories with less concrete words with fewer visual neighbours.

This has further consequences for participants' selection behavior. Although in the prototype setup, the number of objects in the generated image (i.e. *Predicted Image #Objects*) was predictive, the number of objects in the random image (*Random Image #Objects*) was not. However, in the exemplar experiment,

Table 3. Evaluation of two textual embeddings and their grounded versions on the behavioural experiment by Günther et al. (2022). The numbers in the top section of each table represent the percentage of trials in which the models correctly predict the participants' choices. The numbers in the last row of each table (labelled 'Participants') show the mean percentage of trials in which the GPVM image was selected, averaged over all human participants. A. and C. indicate abstract and concrete words respectively. Far, near and max refer to the distance of visual neighbours

a. Prototype model

Embeddings	A. Far	A. Near	C. Far	C. Near	C. Max	Mean
Max: GloVe	52.56	62.66	56.44	84.74	83.46	67.97
Max: ZSG-GloVe	61.41	64.52	60.79	76.79	87.5	70.20
GAM: GloVe	45.42	71.35	61.41	87.99	85.02	70.24
GAM: ZSG-GloVe	53.96	69.49	62.03	87.99	87.5	72.19
Max: W2V	54.27	66.54	67.16	81.33	77.72	69.40
Max: ZSG-W2V	55.82	67.93	58.93	87.34	87.5	71.05
GAM: W2V	49.77	66.54	60.95	87.99	85.79	70.21
GAM: ZSG-W2V	55.98	65.45	62.19	87.99	87.5	71.82
Participants	52.00	64.00	66.00	84.25	85.00	70.25

b. Exemplar model

Embeddings	A. Far	A. Near	C. Far	C. Near	C. Max	Mean
Max: GloVe	57.71	72.73	68.30	70.33	81.5	70.11
Max: ZSG-GloVe	49.09	66.17	65.45	73.00	78.58	66.45
GAM: GloVe	58.02	72.73	73.28	74.63	81.50	72.03
GAM: ZSG-GloVe	62.85	68.93	74.31	75.87	73.7	70.94
Max: W2V	55.81	64.58	63.00	74.00	75.42	66.56
Max: ZSG-W2V	52.73	63.40	63.48	73.22	76.36	65.83
GAM: W2V	55.77	69.09	69.25	75.54	81.11	70.95
GAM: ZSG-W2V	57.23	67.27	70.91	75.21	79.13	70.00
Participants	62.00	74.00	73.00	76.00	77.00	72.40

the number of objects in not only the generated but also in the random image were both significant predictors of selection behavior. For both, more objects in the image corresponded to higher selection probabilities.

A possible explanation for the high variation in participants' selection preferences across highly concrete and abstract queries, between the exemplar and prototype models, may be attributed to the distinct training schemes utilized in each model. In the prototype setup, each noun is associated with the mean image vectors of a specific class (e.g., *horse*), resulting in a feature vector conveying a typical characteristic of a given class and thereby reducing the number of feasible model outputs. Consequently, more distinct boundaries are established for concrete words, which leads to greater discriminability between GPVM images for concrete target words and random control images. In the Exemplar model, on the other hand, each noun is linked with various image vectors that contain the target class in different contexts (other classes). For example, the word *horse* may be associated with multiple images of horses in distinct settings per training sample. This results in the establishment of an association between the target noun and a diverse yet related set of classes. Although the boundaries for concrete words are not as distinct as those for the prototype model, as evidenced by a lower rating for highly concrete words, participants are more inclined to associate the target words with the different yet related set of words encountered by the model during training. Hence, in the Exemplar setup, participants are more likely to prefer the GPVM image for abstract words and concrete words in the far category. Overall, the Exemplar model might retrieve images which contain useful hints indicating its selection and encourage participants to deeply analyse the semantics of the given images with the target word, whereas the prototype model seems to be less flexible in this regard.

Further evidence for a deeper processing of the random control images is provided by the *Random Image Similarity* measure, which was predictive for a larger range of values in the exemplar setup (0, 0.20), compared to the range (0.10, 0.20) observed for the prototype setup. For where *Random Image Similarity* has dense data support, it predicted mostly a decrease in selection probability.

Finally, a greater *Inter-image Similarity* corresponded to lower selection probabilities, although for large values its effect leveled off. Whereas in the prototype setup, *Inter-image Similarity* revealed that unexpected similarity boosted the selection of the generated image, in the exemplar setup, a greater similarity led to the more frequent selection of the random image.

3.2 Q2: Is participants' behaviour best accounted for by purely textual or multimodal word embeddings?

The previous subsection discussed results for both purely textual and grounded embeddings without further discussing any differences in their performance.

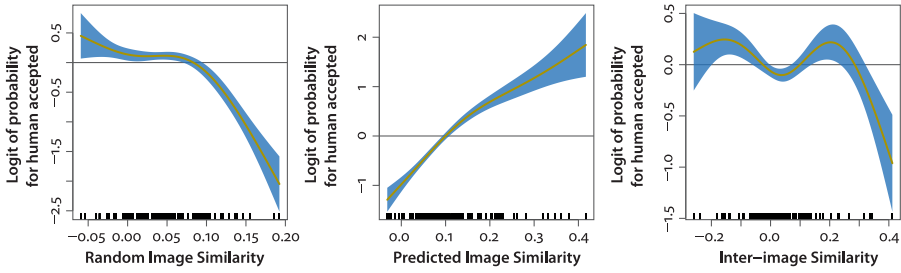
However, our second question is whether using visually grounded embeddings instead of purely textual embeddings will improve prediction accuracy.

Similar to the previous subsection, we first focus on the Max approach selecting the GPVM image or the random control image for each target word. If participants generated visual images for the target word, then the grounded vectors developed by Shahmohammadi et al. are expected to provide enhanced prediction accuracy for participants' behavior. Table 2 shows that using grounded embeddings results in lower Δ values for both grounded GloVe and grounded Word2Vec embeddings in the exemplar model, and for grounded GloVe embeddings in the prototype model, suggesting that the grounded embeddings model participants' preference somewhat better compared to textual embeddings. We conducted a sign test for both setups, comparing the average accuracy in each WordType/Distance category between the textual and grounded version of each embedding type, thus resulting in 10 comparisons overall. We counted a comparison as a "success" if the average accuracy of the grounded embeddings was closer to participants' performance than the textual one. Interestingly, the sign test clarified that actually the grounded embeddings were not significantly better than their purely textual counterparts, neither in the prototype nor in the exemplar setup.

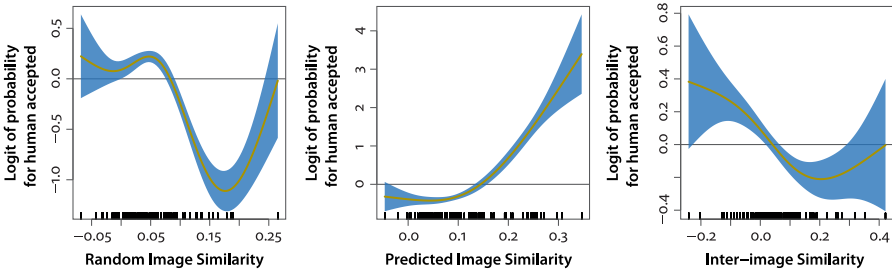
Moving on to predicting participants' behaviour directly, we find that for the prototype setup (see Table 3a), using visually grounded vectors improves mean accuracy by 2% for GloVe and 1% for Word2Vec (compare Max: GloVe with Max: ZSG-GloVe, and Max: W2V with Max: ZSG-W2V). However, for the exemplar setup (see Table 3b), accuracy decreased by 4% for ZSG-GloVe and by 1% for ZSG-W2V.

Turning to the results of the GAMs (rows denoted "GAM", column "Mean" in Table 3), we observe that the GAMs for the prototype models show better accuracy (by 1–2 percentage points) when based on grounded rather than on purely textual embeddings. In the exemplar setup, they do not show better results in predicting human responses. In terms of AIC, the GAMs for both prototype and exemplar models show a better model fit when based on grounded GloVe embeddings than on textual ones (by 80.4 and 86.74 AIC points respectively) and on grounded Word2Vec embeddings only in the prototype setup (by 123.1 AIC points; in the exemplar setup the difference was -3.1 points), again compared to textual ones.

In summary, in terms of numerical differences, the evidence of multimodal embeddings improving prediction for participants' behaviour is mixed. For GPVM images in the prototype setup, both the Max and GAM evaluation methods show a slight advantage for grounded embeddings. For exemplar images, no such advantage is visible in neither the Max nor the GAM evaluation. We again



a. Prototype model



b. Exemplar model

Figure 5. Partial effects (using thin plate regression splines) of the predictors in GAMs for prototype and exemplar models based on grounded GloVe vectors. Plots for textual GloVe vectors as well as the Word2Vec vectors used by Günther et al. (2022) can be found in the supplementary materials

ran a sign test for the prototype and exemplar setups respectively in the same way as in the previous section (20 comparisons per setup, this time a comparison was a “success” if the grounded version of the embeddings showed a larger accuracy). The difference was again not significant in both setups.

The reason that the grounded embeddings by Shahmohammadi et al. are on average, numerically, somewhat less effective for the exemplar model is likely to be that grounded embeddings cluster by semantic similarity rather than by semantic relatedness (Shahmohammadi et al., 2021, 2023). Since in the exemplar model, due to the way in which images are processed, relatedness plays a much stronger role than in the prototype approach (see Section 3.1), the grounded embeddings are less effective for the experimental data obtained from the exemplar-based set-up.

Considered jointly, these results lead us to conclude that participants’ behaviour appears to be equally well accounted for by purely textual vectors and multi-modal vectors. This result raises doubts about participants actually generating visual images of the target words.

Table 4. GAM summary tables for prototype and exemplar models for grounded GloVe. Summary tables for textual/grounded W2V models can be found in the supplementary materials

a. Prototype model

A. parametric coefficients	Estimate	Std. Error	t-value	p-value
(Intercept)	-0.0351	0.1869	-0.1878	0.8510
WordType=concrete	0.5216	0.0851	6.1271	< 0.0001
Distance=near	0.1866	0.0878	2.1247	0.0336
Distance=max	0.4936	0.1195	4.1296	< 0.0001
Predicted Image #Objects	0.0928	0.0188	4.9400	< 0.0001
Random Image #Objects	-0.0078	0.0187	-0.4157	0.6776
WordType=concrete:Distance=near	0.4303	0.1392	3.0913	0.0020
B. smooth terms	edf	Ref.df	F-value	p-value
s(Random Image Similarity)	3.5038	3.8671	111.9450	< 0.0001
s(Predicted Image Similarity)	2.7066	3.2410	227.4596	< 0.0001
s(Inter-Image Similarity)	3.7710	3.9649	24.0678	0.0001

b. Exemplar model

A. parametric coefficients	Estimate	Std. Error	t-value	p-value
(Intercept)	0.1015	0.1689	0.6009	0.5479
WordType=concrete	0.3432	0.0906	3.7898	0.0002
Distance=near	0.1636	0.0931	1.7574	0.0789
Distance=max	-0.1399	0.1139	-1.2286	0.2192
Predicted Image #Objects	0.0544	0.0202	2.6973	0.0070
Random Image #Objects	0.0704	0.0177	3.9813	0.0001
WordTypeconcrete:distance_near1	-0.1100	0.1399	-0.7859	0.4319
B. smooth terms	edf	Ref.df	F-value	p-value
s(Random Image Similarity)	3.8844	3.9923	132.6433	< 0.0001
s(Predicted Image Similarity)	3.2650	3.6982	184.3976	< 0.0001
s(Inter-Image Similarity)	2.7887	3.3389	25.7459	< 0.0001

3.3 Q3: Does the indirect grounding of abstract words afford a better understanding of the experimental results reported by GPVM?

Is visual grounding beneficial not only for concrete words but also for abstract words? On the basis of a series of human-annotated semantic similarity datasets, Shahmohammadi et al. argued that indeed abstract words do benefit from indirect visual grounding. Can the same conclusion be drawn for the data of GPVM?

Analogously to the previous two subsections we again first consider the proportions of the Max approach selecting the GPVM image over the random control image, this time broken down for each of the combinations of WordType and Distance. Considering Table 2a, in the prototype experiment, participants' scores are close to random for the abstract far condition, they are somewhat higher for the abstract near and concrete far conditions, and they are highest for the concrete near and concrete max conditions. Focusing on the best model, Max: ZSG-GloVe, we find a very similar pattern. Comparing this model to the predictions of Max: GloVe, we find that grounding moves the predictions close to human performance in all conditions except the concrete max. The most notable difference here can be found in the abstract far category, where Max: GloVe selects the GPVM image far more often than human participants (see also Figure 4a).

Next, consider Table 2b, which concerns participants' selection preferences for the GPVM images in the exemplar setup. Compared to their performance in the prototype setup, participants' accuracy scores are down considerably for Concrete Near and Concrete Max, and up considerably for Abstract Far and Abstract Near. Performance for Abstract Far words clearly lags behind performance for the other four subsets of words. The GPVM images in the exemplar setup elicited flatter scores, consistent with our conclusion in the preceding sections that in this setup participants scan the control images more carefully. Both the Max:ZSG-GloVe and Max:ZSG-W2V models perform reasonably similar to the participants' preferences, but there are conditions where textual embeddings capture their preferences better. However, it should be noted that all models are well within the range of participants' performance (Figure 4b).

Next, we turn to predict participants' selection behaviour directly. First consider Table 3a, which concerns the prototype setup. Focusing on the best model, GAM:ZSG-GloVe, we find that prediction accuracy is clearly higher for Concrete Near and Concrete Max compared to the Abstract and Concrete Far conditions. Comparing this model with GAM:GloVe, we see that visual grounding improves accuracies for 3 of the five subsets: Concrete Far, Concrete Max, and Abstract Far. There is one subset where grounding leads to lower scores, Abstract Near, and one where grounding does not change performance, Concrete

Near. Averaging over both subsets of abstract words, it seems that there is a modest advantage overall for the visual grounding of abstract words.

Table 3b concerns the exemplar setup. The GAM:ZSG-GloVe model performs reasonably similar to the participants. Compared to GAM:GloVe, for the abstract words, the model shows an improvement of 4% in the Abstract Far category and a reduction in accuracy of 4% for the Abstract Near category.

In summary, it appears that visual grounding aligns more closely with participants' selection behavior in the prototype setup, but its effect is somewhat mixed for abstract words in the exemplar setup. A potential explanation for this finding, same as in the previous section, is that visual grounding tends to create clusters of similar words rather than clusters of related words (Shahmohammadi et al., 2021, 2023). Given that the exemplar model establishes an association between the target nouns and diverse yet related concepts, its behavior for abstract nouns may be better explained by textual embeddings. Therefore, shifting the focus toward similarity appears to benefit highly concrete words but has a negative impact on modeling abstract words.

4. Discussion and conclusion

We started this investigation with three questions: first, can we predict the behaviour of participants in the experiments reported by Günther et al. (2022) without assuming that they generated mental images? Second, is participant behaviour predicted better by visually grounded or purely textual word embeddings? And third, how does the visual grounding process affect performance on abstract words?

Regarding the first question, we found that an approach taking into account the objects present in the presented images is able to predict participants' behaviour quite well. The covariates that we derived from the embeddings for the objects in the images, random image similarity, predicted image similarity, and inter-image similarity, all helped improve the logistic GAMs that we fitted to predict participants' choice behavior. This finding dovetails well with the eye-tracking literature on image scanning: typically, images illicit multiple fixations, reflecting attention being directed to different parts of images and different objects in images (Cronin et al., 2020). From these findings, we infer that participants probably based their decisions on comparisons in semantic space, and not only in visual space. Our experiments suggest that it is unlikely that participants really generated mental images for the words presented to them. Although eye-tracking experiments suggest that participants can get the gist of an image within a time span of 40 ms (see Castelhana and Rayner, 2008, for a review),

understanding images usually requires a series of fixations. This finding does not fit well with the assumption made by GPVM that the images presented to the participants in the experiments of GPVM were processed holistically and were compared with an equally holistic image projected from the target word. Furthermore, it is well-established that our perception of the world is shaped by the limitations of our sensory organs and the constraints imposed by the cultures we live in (see, e.g., Kant et al., 1999; Hoffman, 2019). The way in which Shahmohammadi et al. implement visual grounding – constraining the extent to which vision can change embeddings from human texts – does justice, however crude, to this insight.

Our conclusions are also in line with previous findings on mental imagery: for example, Louwse and Connell (2011) argue that modality-information (such as visual information) is to some extent already included in linguistic information and that only for more precise information embodied simulation is required (thus arguing that both linguistic and embodied processes contribute to conceptual processing). According to their study, linguistic processes account for early processing (short reaction times) and embodied ones for later processing (longer reaction times). Our results also dovetail well with the views on grounding proposed by Zwaan and Madden (2005) and Barsalou (1999) mentioned in the introduction.

As to question 2, we found that our models for predicting participants' performance are slightly, though not statistically significant, improved by using grounded embeddings compared to purely textual embeddings for GPVM's prototype setup.

While there was a slight numerical improvement in the prototype setup, Shahmohammadi et al.'s grounded embeddings were not able to improve on the textual baseline in the exemplar model. Our interpretation of this result is that the images predicted by GPVM in the exemplar setup are driven more by semantic relatedness than by semantic similarity (by virtue of how the model is trained), and as the visual grounding method of Shahmohammadi et al. enhances semantic similarity rather than semantic relatedness, it is less effective for the experimental data of the exemplar model.

GPVM argued that their exemplar model picked up more “idiosyncratic information”, leading to a loss of predictivity of concreteness and number of visual neighbors. The above comparison of the performance of textual and grounded embeddings suggests that the exemplar model is not picking up just noise (idiosyncratic information), but rather that it is more influenced by semantic relatedness, mediated by the objects that co-occur with the objects that are actually targeted in the images used to train the models (e.g., a doctor co-occurring in an image selected to depict a nurse). The employed visually grounded vectors,

by their design zoom in on semantic similarity, whereas standard textual vectors are somewhat more sensitive to semantic relatedness. These considerations lead us to conclude that in the experiments of GPVM, subjects' decisions were guided by both semantic similarity and semantic relatedness, and that the way in which images were selected (prototype vs. exemplar) influenced the relative importance of similarity and relatedness in participants' decision making.

Here, it is important to stress the differences between the grounding model proposed by GPVM and the one by Shahmohammadi et al. GPVM posit a simple linear mapping from textual to visual embeddings. This model works as a proof of concept that a connection between language and vision can be drawn. However, it is not necessarily the best way to combine the two modalities. Shahmohammadi et al. showed that the embeddings generated by such a simple linear mapping perform much worse at predicting human ratings in a number of semantic similarity and relatedness datasets. This indicates that it is not enough to show that language and vision *can* be linked: the real challenge is to understand *how* humans *combine* information from both modalities in order to form meaning representations and make similarity judgments.

Regarding question 3, we observed that grounding tends to yield better predictions of human judgments than textual embeddings for abstract words at least in the prototype setup. This finding, together with previous results suggesting that grounding improves performance on similarity/relatedness judgment tasks even for abstract words (Shahmohammadi et al., 2023), begs the question of why a grounding model trained only on concrete words for which images were available in COCO improves performance also for abstract words. Our interpretation of these results is that by improving the relative position of concrete words in semantic space, the representations of abstract words are also improved. In other words, the visual alignment trained on concrete words also benefits abstract words by transferring them into a more precise semantic space. This view is in line with conceptual metaphor theory which posits that abstract words are understood in terms of concrete words (e.g. Lakoff and Johnson, 1980).

With respect to how abstract and concrete words are learned, Vigliocco et al. (2018) conclude that only by the age of 10 children have sufficient experience with their language to be able to start making use of distributional semantics for learning abstract words. They argue that children are more likely to be using the strong association between abstractness and emotional valence for learning abstract words. Images come with emotional values that are visible in the EEG even under scrambling (Rozenkrants et al., 2008). By visually grounding abstract words, it is possible that the embeddings of abstract words are not only more precisely profiled with respect to concrete words, but also that abstract words are better grounded with respect to their emotional loadings. Empirical studies

have provided extensive support for the significant role of emotion in human cognition, especially in abstract concepts (see Dolan, 2002, for a review). It has been shown that the addition of emotional representations to textual embeddings improves classification tasks on datasets that contain mostly abstract words (Rotaru and Vigliocco, 2020b). It is also possible that once visually grounded embeddings are used, instead of purely textual vectors, children may turn out to be sensitive to distributional aspects of their language at an earlier age than reported by Vigliocco et al. (2018).

We finish with three comments: Firstly, with regard to other visually grounded models: Despite the tremendous amount of work on the visual grounding of textual embeddings, the common belief holds that grounding words visually is beneficial for words with concrete meaning and has an adverse effect on abstract words (Pezzelle et al., 2021; Kiros et al., 2018; Kiela et al., 2018; Park and Myaeng, 2017b). Some new studies further suggest that visual grounding of current sentence-level contextualized textual models does not add extra knowledge for downstream NLP tasks (Yun et al., 2021; Tan and Bansal, 2020). However, the common core idea among many previous grounded approaches is fusing vision and language into a single modality. That is, image vectors and word/sentence vectors are either 1. mapped to a common semantic space where similar visual and textual concepts are forced to have similar representations, most often using non-linear transformation, or 2. word vectors of concrete words are replaced by image vectors during the training process.

Shahmohammadi et al. (2023) on the other hand showed that allowing the complete fusion of both modalities, while benefiting the concrete words, is detrimental to modeling abstract words. They argue that language benefits from vision the most once it is guided by perceptual knowledge as opposed to being merged with it. Using this idea, Shahmohammadi et al. showed that visual grounding is highly beneficial for modeling abstract words and further boosts the performance on downstream NLP tasks when limited training is available. Even though their model learns a linear alignment based on a limited number of captions describing concrete scenes, it was used to indirectly generate grounded representations for unseen abstract words. This is in line with the indirect grounding perspective that implies the direct grounding of concrete words and indirect grounding of abstract words via language (Howell et al., 2005; Louwerse, 2011). The indirect grounding theory of abstract words has been recently shown effective at predicting abstract concepts using distributional semantic models (Utsumi, 2022). Indirect grounding, therefore, seems to be a plausible cognitive mechanism for grounding abstract words.

Secondly, as a final conclusion regarding GPVM, we note that while their experiment clearly highlights that a simple linear mapping is able to predict

images that are chosen by participants above chance level, our work highlights that the conclusions which can be drawn from this study are far from clear. Firstly, we know from previous research (Shahmohammadi et al., 2023) that a simple linear mapping as used by GPVM is not a good grounding model: vectors that are grounded in such a way perform much worse than purely textual embeddings on word similarity and relatedness datasets. Secondly, we are able to demonstrate that the task which GPVM used to evaluate their embeddings can be solved to a large extent without taking into account grounded meaning representations. Thus, it is unclear how much the experiment actually taps into aspects of grounded meaning representations. In this study, we used grounded embeddings which have been shown to improve upon purely textual embeddings in previous work (Shahmohammadi et al., 2023). We found a modest numerical, though not statistically significant, improvement of these embeddings over purely textual embeddings in the prototype setup and no major improvement in the exemplar setup. This suggests to us that if at all, the effects of grounding are more pronounced in the prototype setup. Thus, the method for grounding used in GPVM may not provide the optimal method for generating grounded word representations that can be utilized in psycholinguistics tasks, and the extent to which their experiment is well-suited for detecting effects of grounded meaning representations in human cognition remains somewhat uncertain.

Thirdly, what can we ultimately conclude from this about how humans represent meanings? On the one hand, the present paper together with the results from Shahmohammadi et al. support the conclusion that meaning representations are not based on language alone, but also include information based on vision.⁴ Moreover, both Shahmohammadi et al., as well as the present study, show that this applies not only to concrete words, where such an effect may be expected but also to abstract ones. On the other hand, however, Shahmohammadi et al. found that if textual information is overwhelmed by image information, resulting embeddings as predictors of human similarity ratings suffer. This dovetails well with many previous studies which reported that representations based on textual information alone can predict behavioural data very successfully (e.g. Mandera et al., 2017; Westbury, 2014; Westbury and Hollis, 2019). These results underline that human meaning representations are largely based on the experience humans have with the world through language.

4. Note that we restricted ourselves to visual information here. It is very likely that other multi-modal information such as auditory and olfactory information also plays a role, see Kiela et al. (2015); Kiela and Clark (2015)

Funding

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC number 2064/1 – Project number 390727645 as well as by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A, and the European research council, project WIDE-742545.

Acknowledgements

The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Hassan Shahmohammadi.

References

1 On: Sat, 03 Feb 2024 10:16:51
IP: 134.2.72.21
(Guest)



- [doi](#) Abdou, M., Kulmizev, A., Hershovich, D., Frank, S., Pavlick, E., and Søgaard, A. (2021). Can Language Models Encode Perceptual Structure Without Grounding? A Case Study in Color. In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 109–132, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [doi](#) Anderson, A. J., Bruni, E., Lopopolo, A., Poesio, M., and Baroni, M. (2015). Reading visually embodied meaning from the brain: Visually grounded computational models decode visual-object mental imagery induced by written text. *NeuroImage*, 120:309–322.
- [doi](#) Anschütz, M., Lozano, D. M., and Groh, G. (2023). This is not correct! negation-aware evaluation of language generation systems.
- [doi](#) Baroni, M. (2016). Grounding distributional semantics in the visual world. *Language and Linguistics Compass*, 10(1):3–13.
- [doi](#) Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22(4).
- [doi](#) Barsalou, L. W. (2003). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 358(1435).
- [doi](#) Barsalou, L. W. (2008). Grounded Cognition. *Annual Review of Psychology*, 59(1).
- [doi](#) Barsalou, L. W. (2010). Grounded cognition: Past, present, and future. *Topics in cognitive science*, 2(4):716–724.
- [doi](#) Barsalou, L. W., Santos, A., Simmons, W. K., and Wilson, C. D. (2008). Language and simulation in conceptual processing. In *Symbols and Embodiment: Debates on meaning and cognition*. Oxford University Press.
- [doi](#) Bordes, P., Zablocki, E., Soulier, L., Piwowski, B., and Gallinari, P. (2019). Incorporating visual semantics into sentence representations within a grounded space. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 696–707, Hong Kong, China. Association for Computational Linguistics.
- [doi](#) Bruni, E., Tran, N.-K., and Baroni, M. (2014). Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49:1–47.

- [doi](#) Brysbaert, M., Warriner, A. B., and Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3):904–911.
- [doi](#) Buchanan, E. M., Valentine, K. D., and Maxwell, N. P. (2019). English semantic feature production norms: An extended database of 4436 concepts. *Behavior Research Methods*, 51(4).
- [doi](#) Bulat, L., Clark, S., and Shutova, E. (2017). Speaking, Seeing, Understanding: Correlating semantic models with conceptual representation in the brain. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Castelhano, M.S. and Rayner, K. (2008). Eye movements during reading, visual search, and scene perception: An overview. *Cognitive and cultural influences on eye movements*, 2175:3–33.
- [doi](#) Chatfield, K., Simonyan, K., Vedaldi, A., and Zisserman, A. (2014). Return of the Devil in the Details: Delving Deep into Convolutional Nets. *arXiv preprint arXiv:1405.3531*.
- [doi](#) Chrupała, G., Kádár, Á., and Alishahi, A. (2015). Learning language through pictures. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 112–118, Beijing, China. Association for Computational Linguistics.
- [doi](#) Collell Talleda, G., Zhang, T., and Moens, M.-F. (2017). Imagined visual representations as multimodal embeddings. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*, pages 4378–4384. AAAI.
- Cree, G.S. and McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, 132(2).
- [doi](#) Cronin, D.A., Hall, E.H., Goold, J.E., Hayes, T.R., and Henderson, J.M. (2020). Eye movements in real-world scene photographs: General characteristics and effects of viewing task. *Frontiers in Psychology*, 10:2915.
- [doi](#) De Deyne, S., Navarro, D.J., Collell, G., and Perfors, A. (2021). Visual and Affective Multimodal Models of Word Meaning in Language and Mind. *Cognitive Science*, 45(1).
- [doi](#) Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A largescale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee.
- [doi](#) Dolan, R. J. (2002). Emotion, cognition, and behavior. *Science*, 298(5596):1191–1194.
- [doi](#) Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., and Ruppín, E. (2001). Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*, pages 406–414.
- [doi](#) Gerz, D., Vulić, I., Hill, F., Reichart, R., and Korhonen, A. (2016). SimVerb-3500: A large-scale evaluation set of verb similarity. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2173–2182, Austin, Texas. Association for Computational Linguistics.
- [doi](#) Goldstone, R. L. (1995). Effects of Categorization on Color Perception. *Psychological Science*, 6(5).
- [doi](#) Grondin, R., Lupker, S. J., and McRae, K. (2009). Shared features dominate semantic richness effects for concrete concepts. *Journal of Memory and Language*, 60(1):1–19.

- [doi](#) Günther, F., Petilli, M.A., Vergallito, A., and Marelli, M. (2022). Images of the unseen: extrapolating visual representations for abstract and concrete words in a data-driven computational model. *Psychological Research*.
- [doi](#) Günther, F., Rinaldi, L., and Marelli, M. (2019). Vector-Space Models of Semantic Representation From a Cognitive Perspective: A Discussion of Common Misconceptions. *Perspectives on Psychological Science*, 14(6):1006–1033.
- [doi](#) Halawi, G., Dror, G., Gabrilovich, E., and Koren, Y. (2012). Large-scale learning of word relatedness with constraints. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1406–1414.
- [doi](#) Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1–3):335–346.
- [doi](#) Harris, Z.S. (1954). Distributional Structure. *WORD*, 10(2–3).
- [doi](#) Hasegawa, M., Kobayashi, T., and Hayashi, Y. (2017). Incorporating visual features into word embeddings: A bimodal autoencoder-based approach. In *IWCS 2017 – 12th International Conference on Computational Semantics – Short papers*.
- [doi](#) Hill, F., Reichart, R., and Korhonen, A. (2015). Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- [doi](#) Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Hoffman, D. (2019). *The case against reality: Why evolution hid the truth from our eyes*. WW Norton & Company.
- [doi](#) Hollenstein, N., de la Torre, A., Langer, N., and Zhang, C. (2019). CogniVal: A Framework for Cognitive Word Embedding Evaluation. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [doi](#) Howell, S.R., Jankowicz, D., and Becker, S. (2005). A model of grounded language acquisition: Sensorimotor features improve lexical and grammatical learning. *Journal of Memory and Language*, 53(2):258–276.
- Kant, I., Guyer, P., and Wood, A.W. (1781/1999). *Critique of pure reason*. Cambridge University Press.
- Kenton, J.D.M.-W.C. and Toutanova, L.K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- [doi](#) Kiela, D. and Bottou, L. (2014). Learning image embeddings using convolutional neural networks for improved multi-modal semantics. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 36–45, Doha, Qatar. Association for Computational Linguistics.
- [doi](#) Kiela, D., Bulat, L., and Clark, S. (2015). Grounding semantics in olfactory perception. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 231–236.
- [doi](#) Kiela, D. and Clark, S. (2015). Multi- and Cross-Modal Semantics Beyond Vision: Grounding in Auditory Perception. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Stroudsburg, PA, USA. Association for Computational Linguistics.

- doi Kiela, D., Conneau, A., Jabri, A., and Nickel, M. (2018). Learning visually grounded sentence representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 408–418, New Orleans, Louisiana. Association for Computational Linguistics.
- doi Kiros, J., Chan, W., and Hinton, G. (2018). Illustrative language understanding: Largescale visual grounding with image search. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 922–933, Melbourne, Australia. Association for Computational Linguistics.
- doi Lakoff, G. (1987). *Women, Fire, and Dangerous Things*. University of Chicago Press.
- doi Lakoff, G. and Johnson, M. (1980). The metaphorical structure of the human conceptual system. *Cognitive science*, 4(2):195–208.
- doi Landauer, T.K. and 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(2).
- doi Langacker, R.W. (1999). A view from cognitive linguistics. *Behavioral and Brain Sciences*, 22(4).
- doi Lazaridou, A., Chrupała, G., Fernández, R., and Baroni, M. (2016). Multimodal Semantic Learning from Child-Directed Input. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Stroudsburg, PA, USA. Association for Computational Linguistics.
- doi Lazaridou, A., Marelli, M., and Baroni, M. (2017). Multimodal Word Meaning Induction From Minimal Exposure to Natural Text. *Cognitive Science*, 41.
- doi Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C.L. (2014). Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- doi Louwse, M. and Connell, L. (2011). A Taste of Words: Linguistic Context and Perceptual Simulation Predict the Modality of Words. *Cognitive Science*, 35(2):381–398.
- doi Louwse, M.M. (2011). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, 3(2):273–302.
- doi Louwse, M.M. and Zwaan, R.A. (2009). Language Encodes Geographical Information. *Cognitive Science*, 33(1):51–73.
- Luong, T., Socher, R., and Manning, C. (2013). Better word representations with recursive neural networks for morphology. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 104–113, Sofia, Bulgaria. Association for Computational Linguistics.
- doi Lynott, D., Connell, L., Brysbaert, M., Brand, J., and Carney, J. (2020). The Lancaster Sensorimotor Norms: multidimensional measures of perceptual and action strength for 40,000 English words. *Behavior Research Methods*, 52(3).
- doi Mander, P., Keuleers, E., and Brysbaert, M. (2017). Explaining human performance in psycholinguistic tasks with models of semantic similarity based on prediction and counting: A review and empirical validation. *Journal of Memory and Language*, 92.
- doi Marelli, M. and Amenta, S. (2018). A database of orthography-semantics consistency (osc) estimates for 15,017 english words. *Behavior research methods*, 50:1482–1495.


- [doi](#) Martin, A. (2007). The Representation of Object Concepts in the Brain. *Annual Review of Psychology*, 58(1):25–45.
- [doi](#) McRae, K., Cree, G. S., Seidenberg, M. S., and Mcnorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37(4).
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- [doi](#) Mkrtychian, N., Blagovechtchenski, E., Kurmakaeva, D., Gnedykh, D., Kostromina, S., and Shtyrov, Y. (2019). Concrete vs. Abstract Semantics: From Mental Representations to Functional Brain Mapping. *Frontiers in Human Neuroscience*, 13(August):267.
- [doi](#) Montefinese, M. (2019). Semantic representation of abstract and concrete words: A minireview of neural evidence. *Journal of Neurophysiology*, 121(5):1585–1587.
- Park, J. and Myaeng, S.-h. (2017a). A computational study on word meanings and their distributed representations via polymodal embedding. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 214–223, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Park, J. and Myaeng, S.-h. (2017b). A computational study on word meanings and their distributed representations via polymodal embedding. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 214–223.
- [doi](#) Pennington, J., Socher, R., and Manning, C. (2014). Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Stroudsburg, PA, USA. Association for Computational Linguistics.
- [doi](#) Pezzelle, S., Takmaz, E., and Fernández, R. (2021). Word representation learning in multimodal pre-trained transformers: An intrinsic evaluation. *Transactions of the Association for Computational Linguistics*, 9:1563–1579.
- [doi](#) Rotaru, A. S. and Vigliocco, G. (2020a). Constructing semantic models from words, images, and emojis. *Cognitive science*, 44(4):e12830.
- [doi](#) Rotaru, A. S. and Vigliocco, G. (2020b). Constructing Semantic Models From Words, Images, and Emojis. *Cognitive Science*, 44(4):e12830.
- Rozenkrants, B., Olofsson, J. K., and Polich, J. (2008). Affective visual event-related potentials: arousal, valence, and repetition effects for normal and distorted pictures. *International Journal of Psychophysiology*, 67(2):114–123.
- [doi](#) Shahmohammadi, H., Heitmeier, M., Shafaei-Bajestan, E., Lensch, H., and Baayen, H. (2023). Language with vision: a study on grounded word and sentence embeddings. *Behavior Research Methods*, accepted for publication.
- [doi](#) Shahmohammadi, H., Lensch, H. P. A., and Baayen, R. H. (2021). Learning zero-shot multifaceted visually grounded word embeddings via multi-task training. In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 158–170, Online. Association for Computational Linguistics.
- [doi](#) Silberer, C. and Lapata, M. (2014). Learning grounded meaning representations with autoencoders. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 721–732, Baltimore, Maryland. Association for Computational Linguistics.

-  Simmons, W.K., Martin, A., and Barsalou, L. W. (2005). Pictures of Appetizing Foods Activate Gustatory Cortices for Taste and Reward. *Cerebral Cortex*, 15(10):1602–1608.
-  Solomon, K. O. and Barsalou, L. W. (2001). Representing Properties Locally. *Cognitive Psychology*, 43(2):129–169.
-  Solomon, K. O. and Barsalou, L. W. (2004). Perceptual simulation in property verification. *Memory & Cognition*, 32(2):244–259.
-  Tan, H. and Bansal, M. (2020). Vokenization: Improving language understanding with contextualized, visual-grounded supervision. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2066–2080, Online. Association for Computational Linguistics.
- Tan, M. and Le, Q. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR.
-  Utsumi, A. (2022). A test of indirect grounding of abstract concepts using multimodal distributional semantics. *Frontiers in psychology*, 13.
-  Vigliocco, G., Ponari, M., and Norbury, C. (2018). Learning and processing abstract words and concepts: Insights from typical and atypical development. *Topics in cognitive science*, 10(3):533–549.
-  Wang, B., Wang, A., Chen, F., Wang, Y., and Kuo, C.-C.J. (2019). Evaluating word embedding models: Methods and experimental results. *APSIPA transactions on signal and information processing*, 8.
-  Westbury, C. (2014). You Can't Drink a Word: Lexical and Individual Emotionality Affect Subjective Familiarity Judgments. *Journal of Psycholinguistic Research*, 43(5).
- Westbury, C. and Hollis, G. (2019). Wiggly, squiffy, lummoX, and boobs: What makes some words funny? *Journal of Experimental Psychology: General*, 148(1).
-  Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society (B)*, 73(1):3–36.
-  Yun, T., Sun, C., and Pavlick, E. (2021). Does vision-and-language pretraining improve lexical grounding? In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4357–4366, Punta Cana, Dominican Republic. Association for Computational Linguistics.
-  Zwaan, R. A. and Madden, C. J. (2005). Embodied Sentence Comprehension. In *Grounding Cognition*. Cambridge University Press.

Address for correspondence

Hassan Shahmohammadi
 Department of Computer Science
 University of Tübingen
 Room 20–31 A6, Maria-von-linden strasse 6
 72076, Tuebingen
 Germany

hassan.shahmohammadi@uni-tuebingen.de

 <https://orcid.org/0000-0002-3112-3252>

Co-author information

Maria Heitmeier
Department of Linguistics
University of Tübingen
maria.heitmeier@uni-tuebingen.de

Elnaz Shafaei-Bajestan
Department of Linguistics
University of Tübingen
elnaz.shafaei@gmail.com

Hendrik P.A. Lensch
Department of Computer Science
University of Tübingen
hendrik.lensch@uni-tuebingen.de

R. Harald Baayen
Department of Linguistics
University of Tübingen
Harald.Baayen@protonmail.com

Publication history

Date received: 4 July 2022
Date accepted: 14 November 2023
Published online: 11 January 2024