

Firearms and Tigers are Dangerous, Kitchen Knives and Zebras are Not: Testing whether Word Embeddings Can Tell

Pia Sommerauer and Antske Fokkens

Computational Lexicology and Terminology Lab

Vrije Universiteit Amsterdam

De Boelelaan 1105 Amsterdam, The Netherlands

`pia.sommerauer@vu.nl`, `antske.fokkens@vu.nl`

Abstract

This paper presents an approach for investigating the nature of semantic information captured by word embeddings. We propose a method that extends an existing human-elicited semantic property dataset with gold negative examples using crowd judgments. Our experimental approach tests the ability of supervised classifiers to identify semantic features in word embedding vectors and compares this to a feature-identification method based on full vector cosine similarity. The idea behind this method is that properties identified by classifiers, but not through full vector comparison are captured by embeddings. Properties that cannot be identified by either method are not. Our results provide an initial indication that semantic properties relevant for the way entities interact (e.g. dangerous) are captured, while perceptual information (e.g. colors) is not represented. We conclude that, though preliminary, these results show that our method is suitable for identifying which properties are captured by embeddings.

1 Introduction

Word embeddings are widely used in NLP and have been shown to boost performance in a large selection of tasks ranging from morphological analysis to sentiment analysis (Lazaridou et al., 2013; Socher et al., 2013; Zhou and Xu, 2015, among many others). Despite a number of different approaches to evaluation, our understanding of what type of information is represented by the vectors remains limited. Most approaches focus on full-vector comparison which treat vectors as points in a space (Yaghoobzadeh and Schütze, 2016), which are evaluated by performance on semantic similarity or relatedness test sets and analogy questions (Mikolov et al., 2013; Turney, 2012). Previous work, however, has shown that high performance does not necessarily mean that

vectors actually contain the information required to solve the task (Rogers et al., 2017; Linzen, 2016). Better understanding of the kind of semantic information captured by word embeddings can increase our understanding of how they help improve downstream tasks. In general, understanding what information is present in (often prominent) input embeddings forms an essential component of gaining deeper understanding of the nature of information and manner in which it travels through the hidden layers of a neural network.

In this paper, we propose a method that investigates what kind of semantic information is encoded in vectors using a human-elicited dataset of semantic properties. We compare the output of supervised classifiers to an approach based on full-vector comparison that cannot access individual dimensions. The assumptions behind this approach are that (1) both full-vector comparison and the supervised classifier will perform well on identifying semantic properties that correlate highly with general similarity; (2) the classifier will outperform full-vector analysis on properties that are reflected by the context, but shared among a diverse set of entities and (3) that neither approach will perform well on properties that are not represented directly or indirectly in the text. The last two outcomes can indicate whether a semantic property is encoded in embeddings (2) or not (3).

The main contribution of this paper lies in the new method and corpus it proposes. To our knowledge, this is the first approach that aims at identifying whether specific semantic properties are captured by individual dimensions or complex patterns in the vector. In addition, we provide specific hypotheses as to which properties are captured well by which method and test them using our approach.¹ Our general hypothesis states that se-

¹The hypotheses and code for our experiments can

semantic properties that are relevant for the way entities interact with the world are well represented (e.g. functions of objects, activities entities are frequently involved in), whereas properties of relatively little consequence for the way entities interact with the world are not (e.g. perceptual properties such as shapes and colors, which either have no function or highly diverse functions). Though preliminary due to the complexity of the task, results indicate that these tendencies hold. Moreover, the overall outcome shows that the method and data are complementary to existing intrinsic evaluation methods.

The rest of this paper is structured as follows. We discuss related work in Section 2. Our method is outlined in Section 3. Section 4 presents our experiments and results. We finish with a critical discussion and overview of future work in Section 5.

2 Related work

Intrinsic evaluation of word embeddings has primarily focused on two main tasks: identifying general semantic relatedness or similarity and the so-called analogy task, where word embeddings have been shown to be able to predict missing components of analogies of the type A is to B as C is to D (Mikolov et al., 2013; Turney, 2012). Furthermore, most intrinsic evaluation methods take full vectors into consideration. The famous examples $Paris - France + Italy \approx Rome$ or $king - man + woman \approx queen$ evoke the suggestion that embeddings can capture semantic properties. The task has, however, been criticized substantially (Linzen, 2016; Gladkova and Drozd, 2016; Gladkova et al., 2016; Drozd et al., 2016, among others).

Gladkova et al. (2016) follow an observation in Levy and Goldberg (2014) on the large differences in performance on different categories in the Google analogy set (Mikolov et al., 2013). They provide a new, more challenging, analogy dataset that improves existing sets on balance (capturing more semantic categories) and size. Linzen (2016) points out more fundamental problems including the observation that the target vector in the analogy task can often be found by simply taking the vector closest to the source. Drozd et al. (2016) show that classifiers picking out the target word from a set of related terms outperform the stan-

dardly applied cosine addition or multiplication methods. Though also boosted by the aforementioned proximity bias, these results indicate that standard methods of solving analogies miss information that is captured by embeddings. Rogers et al. (2017) conclude that the analogy evaluation does not reveal if word embedding representations indeed capture specific semantic properties.

On top of that, an embedding may capture specific semantic properties in ways that are not analogous to semantic properties of related categories. Analogy methods assume that semantic properties stand in analogous relation to each other based on the information provided by the context, but there is no reason why (e.g.) things made of wood and things made of plastic result in (combinations of) embedding dimensions that are similar enough to stand in a parallel relation to each other. Our setup can determine whether the properties are represented without supposing such structures by targeting semantic properties directly rather than in relation to other concepts.

Several approaches have attempted to derive properties collected in property norm datasets from the distribution in naturally occurring texts (Kelly et al., 2014; Baroni et al., 2010; Barbu, 2008). Whereas these approaches yield indications about the potential of distributional models, they do not go beyond full-vector proximity on a low-dimensional SVD model or context words in a transparent, high-dimensional count model. Their focus lies on detecting informative contexts. We follow the idea behind this approach and make a human-elicited property dataset that is created in the same tradition, but larger. Our approach goes beyond the previous work in two ways: first, we add gold negative examples which allows us to go beyond testing for salient properties. Second, we compare full vector proximity to the outcome of a classifier which allows us to verify whether the property is captured for entities that share the property, but are not similar otherwise.

A few other studies go beyond full vector comparisons, moving towards the interpretation of word embedding dimensions. Tsvetkov et al. (2015, 2016) evaluate word embeddings by measuring the correlation between word embedding vectors and count vectors representing co-occurrences of words with WordNet supersenses. While they show that their results have a higher correlation with results obtained from

be found at: https://cltl.github.io/semantic_space_navigation

extrinsic evaluations than standardly used intrinsic evaluations, they do not provide insights into what kind of semantic information is represented well. Yaghoobzadeh and Schütze (2016) decompose distributional vectors into individual linguistic aspects by means of a supervised classification approach to test which linguistic phenomena are captured by embeddings. They test their approach on an artificially created corpus and do not provide insights into specific semantic knowledge. Faruqui et al. (2015) transform learned embedding matrices into sparse matrices to make them more interpretable, which is complementary to our approach.

Previous studies provide (indicative) support for the hypothesis that embeddings lack information people get from other modalities than language. Fagarasan et al. (2015) present a method to ground embedding models in perceptual information by mapping distributional spaces to semantic spaces consisting of feature norms. Several approaches to boosting distributional models with visual information show that the additional information improves the performance of word embedding vectors (Roller and Schulte im Walde, 2013; Lazaridou et al., 2014). Whereas this indicates that word embedding models lack visual information, it does not show to what extent different types of properties are encoded. The method proposed in this paper is, to the best of our knowledge, the first approach specifically designed to identify what semantic knowledge is captured in word embeddings. We are not aware of earlier work that provides explicit hypotheses about the kind of information we expect to learn from distributional vectors, making this the first attempt to confirm these hypotheses experimentally.

3 Method

The core of our evaluation consists of testing whether nearest neighbors and classifiers are capable of identifying which embeddings encode a given semantic property. We first describe the dataset and then present the procedure we apply. We complete this section with our hypotheses about the outcome of our evaluation.

3.1 Extended CSLB Data

We use the Centre for Speech, Language and the Brain concept property norms dataset (Devereux et al., 2014, henceforth CSLB). This dataset fol-

lows the tradition of the sets created by McRae et al. (2005); Vinson and Vigliocco (2008) and used in Kelly et al. (2014); Baroni et al. (2010); Barbu (2008) and is the largest available semantic property dataset we are aware of. In the collection process, human subjects were given concrete and mostly monosemous concepts and asked to provide a set of semantic features. Polysemous concepts were disambiguated. Properties were elicited by cues such as *has*, *is*, *does* and *made_of*. An empty slot was provided to fill in other relations. The dataset comprises 638 annotated concepts, each of which was presented to 30 participants. Properties listed by at least two participants are included in the published set.

We select features associated with at least 20 concepts. In an exploratory experiment, we count all concepts for which the target feature is listed as positive examples and all other concepts as negative examples. However, the fact that people did not list a property does not necessarily mean that a given concept is a negative example of it. For instance: *falcon* is described by *is_a_bird*, but not by *is_an_animal*.

For proper evaluation, the CSLB dataset should be extended with verified negative examples. We apply two methods to add both positive and (verified) negative properties to CSLB. First, we select properties that necessarily imply the target property (e.g. *is_a_bird* implies *is_an_animal*) or necessarily exclude the target property (e.g. *is_food* almost certainly excludes *has_wheels*). We both manually inspect the extended sets of positive and negative examples per selected property to exclude remaining noise independently, resolving disagreements after discussion.²

The resulting dataset has the disadvantage that negative examples largely consist of the same specific categories, e.g. negative examples of *has_wheels* are food, animals and plants. Based on these examples, we cannot tell whether the classifier performs well because embeddings encode the property of having wheels or because it can distinguish vehicles from food, animals and plants. We therefore need to expand the dataset so that it includes diverse negative and positive examples and preferably positive and negative examples that are closely related in semantic space.

Ultimately, we want to verify and increase the

²All annotations, guiding principles as well as notes about resolving discussions can be found at https://ciltl.github.io/semantic_space_navigation.

entire dataset and distinguish between things that always or typically have a property (e.g. *bike has_wheels*, *banana is_yellow*), things that can have a property (e.g. *bikini - is_pink*, *plate - made_of_metal*) and things that do normally not have a property (e.g. *grape - does_kill*, *beer - is_pink*). We set up a crowdsourcing task in which we ask participants whether a property applies to a word. Possible answers are *yes*, *mostly*, *possibly* and *no*.

This crowdsourcing method has currently been applied to a selection of property-concept pairs that were labeled as false-positives by at least one of our approaches in the initial setup. In addition, we extend the property-concept pairs given to crowd workers by collecting the nearest neighbors of the property centroid and a number of seed words. We aim at (1) identifying negative examples that have a high cosine similarity to positive examples in the dataset and (2) including a broader variety of words. This nearest-neighbors strategy explicitly aims at collecting words that are highly similar to positive examples of a property but are not associated with it. For instance, in order to extend the concept set for the property *has_wheels*, we used the seed words *car*, *sledge*, and *ship*.³

In the experiments reported in this paper, we only consider properties that clearly apply to a concept as positive examples (the *yes* and *mostly* cases) and properties that clearly do not apply as negative examples, leaving disputable cases and the cases that *possibly* apply for future work. We manually checked cases of disagreement in the crowd data and selected or removed data based on these criteria.⁴

3.2 Classification approaches

We use the pretrained Word2vec model based on the Google News corpus.⁵ The underlying architecture is a skip-gram with negative sampling model (Mikolov et al., 2013), which learns word vectors by predicting the context given a word.

The overall goal is to investigate whether word vectors capture specific semantic properties or not. We start from the assumption that classifiers can

³The details about our selection and full lists of seed words are provided with our code (see link in Footnote 1).

⁴Some difference in judgment are clearly the result of lack of knowledge (e.g. not knowing a something is an animal). The original outcome of the crowd and final resulting test are provided on the github repository associated with this paper.

⁵<https://code.google.com/archive/p/word2vec/>

learn properties that are represented in the embedding in a binary classification task. We apply supervised classification to see whether a logistic regression classifier or a neural network are capable of distinguishing embeddings of words that have a specific semantic property from those which do not. Specifically, we use embedding vectors corresponding to words associated or not associated with a semantic target-property (i.e. positive and negative examples) as input for a binary classifier and test whether the classifier can learn to distinguish embeddings of words that have the property from those who do not. However, word embeddings also capture semantic similarity. If a property is shared by similar entities (e.g. most animals with a beak are birds), the classifiers may perform well because of this similarity rather than identifying the actual property. We therefore compare the performance of classifiers to the performance of an approach based on full vector similarity. If only the classifiers score well, this provides an indication that the embedding captures the property. If both methods perform poorly this could mean that the property is not captured.⁶

Supervised classification

As the datasets are limited in size, we evaluate by applying a leave-one-out approach. We employ two different supervised classifiers, which we expect to differ in performance. As a ‘vanilla’ approach, we use a logistic regression classifier with default settings as implemented in SKlearn. This type of classifier is also used in Drozd et al. (2016) to detect words of similar categories in an improved analogy model.

In addition, we use a basic neural network. Meaningful properties may not always be encoded in individual patterns, but rather arise from a combination of activated dimensions. This is not captured well by a logistic regression model, as it can only react to individual dimensions. In contrast, the neural network can learn from patterns of dimensions. We use a simple multi-layer perceptron (as implemented in SKlearn⁷) with a single hidden layer. We calculate the number of nodes in the hidden layer as follows: (number of input dimensions + number of output dimensions) * 1/3. The

⁶Given the size and balance of our dataset as well as the lack of fine-tuning, we remain careful not to draw firm conclusions at this point.

⁷<http://scikit-learn.org/stable/index.html>

pretrained Google News vectors have 300 dimensions, resulting in a hidden layer of 100 nodes. We use the recommended settings for small datasets. No parameter tuning was conducted so far due to the limited size of the datasets and the use of a leave-one-out evaluation strategy. We present the runs of several models, as the neural network can react to the order in which the examples are presented as well as the randomly assigned vectors for initialization. While the performance of the model could be optimized further by experimenting with the settings, we find that the set-up presented here already outperforms the logistic regression classifier in many cases.

Full vector similarity

To show that supervised classification can go beyond full vector comparison in terms of cosine similarity, we compare the performance of the classifiers to an n-nearest neighbors approach. We calculate the centroid vector of all positive examples in the training set. The training set consists of all positive examples in the leave-one-out split except for the one we are testing on. We then consider its n-nearest neighbors measured by their cosine distance to the centroid as positive examples. We vary n between 100 and 1,000 in steps of 100. We report the performance of the optimal number of neighbors for each property (which varies per property). In future work, we will add more fine-grained steps and investigate the performance of a classifier using the cosine similarity of words to the centroid as a sole feature.

Variety approximation

The performance of the approaches outlined above depends on the variety of words associated with a property. We approximate this variety by calculating the average cosine similarity of words associated with a property to one-another. This is done by averaging over the cosine similarities between all possible pairs of words. A high average cosine similarity means that the words associated with a concepts tend to be close to each other in the space, which should mostly apply to words associated with taxonomic categories. In contrast, a low average cosine means a high diversity, which should largely apply to general descriptions.

3.3 Specific hypotheses

We select a number of properties for closer investigation based on the clean and extended dataset

described in Section 3.1. We first formulated the hypotheses independently, before discussing and specifying them.⁸ Table 1 summarizes the agreed upon expectations. The hypotheses can be categorized in the following way:

Sparse Textual Evidence

We select properties of which we expect that textual evidence is too sparse to be represented by distributional vectors. The properties *is_black*, *is_yellow*, *is_red* and *made_of_wood* have little impact on the way most entities belonging to that class interact with the world. We expect that the only textual evidence indicating them are individual words denoting the properties themselves (e.g. *red*, *black*, *wooden*)⁹ and it is unclear how often they are mentioned explicitly. It may, however, be the case that certain subcategories in the datasets are learned regardless of this sparsity, because they happen to coincide with more relevant taxonomic categories such as *red fruits*.

Fine-grained Distinctions in Larger Categories

We expect that a supervised classifier may be able to make more fine-grained distinctions between examples of the same category when these differences are relevant for the way they interact with the world. We select two properties that introduce crucial distinctions in larger categories: *has_wheels* and *is_found_in_seas*. The former applies to a sub-group of vehicles and may be apparent in certain behaviors and contexts only applying to these vehicles (rolling, street, etc). The latter applies to animals, plants and other entities found in water, but it is unclear whether textual evidence is enough to distinguish between seawater and fresh water.

Mixed Groups

We expect that a supervised machine learning approach can find positive examples of a property that are not part of the most common class in the training set. For instance, the majority of positive examples for *is_dangerous* and *does_kill* refer to weapons or dangerous animals. We expect the classifier to (1) find positive examples from less well represented groups and (2) be able to distinguish between positive and negative examples of a well-represented category (e.g. *rhino* v.s. *hippo*

⁸Details can be found on the paper's github repository.

⁹In the case of *made_of_wood*, the evidence may be a bit broader, as it might be indicated by different types of wood occurring in the context of furniture.

Property	learnable property
is_an_animal	yes
is_food	yes
is_dangerous	yes
does_kill	yes
is_used_in_cooking	yes
has_wheels	possibly
is_found_in_seas	possibly
is_black	no
is_red	no
is_yellow	no
made_of_wood	no

Table 1: Hypotheses about whether selected semantic properties can be learned by a supervised classifier

for killing). For the property *is_used_in_cooking*, the example words refer to food items as well as utensils. We expect that classifiers can distinguish between cooking-related utensils and other tools.

Polysemy

We expect that machine learning can recognize vector dimensions indicating properties applying to different senses of a word, whereas the nearest-neighbors approach simply assigns the word to its dominant class. For instance, we expect that word vectors that can be used to describe animals as well as food (e.g. *chicken*, *rabbit* or *turkey*) record evidence of both contexts, but end up closer to one of the categories. A supervised machine learning approach should be able to find the relevant dimensions regardless of the cosine similarity to one of the groups and classify the word correctly. We test this by training on a set of monosemous words (animals and food items) and test on a set of polysemous and monosemous examples.

4 Experimental set-up and results

4.1 Concept diversity vs performance

We first investigate the relation between performance and diversity of concepts associated with a property on the full, noisy dataset using a leave-one-out approach. Table 2 shows a selection of the f1-scores achieved on properties in the CSLB dataset in relation to the average cosine similarity of the associated words. A high average cosine similarity means that the concepts overall have similar vector representations and can thus be seen as having a low diversity. The results of the Spearman Rank correlation clearly indicate that scores achieved by nearest neighbors correlate more strongly with the average cosine than the two supervised classification approaches.

feature	cos	f1-neigh	f1-lr	f1-net	type
is_heavy	0.15	0.15	0.17	0.21	op
is_strong	0.15	0.13	0.13	0.34	e
is_thin	0.16	0	0.05	0.1	vp
is_hard	0.16	0.15	0.08	0.26	op
is_expensive	0.16	0	0.28	0.37	e
...
is_black	0.2	0.29	0.23	0.24	vp
is_electric	0.21	0.48	0.5	0.69	vp
is_dangerous	0.21	0.53	0.57	0.59	e
is_colourful	0.21	0.14	0.25	0.32	vp
is_brown	0.21	0.13	0.22	0.33	vp
has_a_handle	0.22	0.44	0.57	0.58	p
_handles					
has_a_seat	0.22	0.43	0.3	0.48	p
_seats					
does_smell	0.22	0.08	0.15	0.37	op
_is_smelly					
made_of_glass	0.22	0.29	0	0.28	vp
has_a_point	0.23	0.38	0.23	0.47	p
does_protect	0.24	0.38	0.26	0.37	f
is_yellow	0.24	0.22	0	0.23	vp
is_soft	0.24	0.12	0	0.16	op
is_red	0.25	0.34	0.13	0.27	vp
is_fast	0.25	0.3	0.31	0.48	vp
is_tall	0.25	0.43	0.57	0.65	vp
is_a_tool	0.26	0.5	0.51	0.47	t
...
is_a_weapon	0.3	0.74	0.56	0.63	t
is_green	0.31	0.45	0.45	0.45	vp
has_a_	0.32	0.68	0.65	0.74	p
blade_blades					
is_worn	0.32	0.47	0.86	0.9	f
has_wheels	0.32	0.82	0.83	0.87	p
is_found	0.33	0.56	0.73	0.76	e
_in_kitchens					
does_fly	0.33	0.57	0.76	0.76	f
has_a_tail	0.33	0.53	0.68	0.69	p
is_an_animal	0.33	0.64	0.76	0.78	t
is_eaten_edible	0.33	0.37	0.88	0.85	f
has_four_legs	0.34	0.67	0.66	0.66	p
is_a_vehicle	0.34	0.76	0.69	0.79	t
does_eat	0.34	0.68	0.71	0.68	f
...
has_a_beak	0.37	0.63	0.83	0.87	p
made_of_cotton	0.37	0.68	0.56	0.64	vp
has_roots	0.37	0.3	0.65	0.72	p
is_a_mammal	0.37	0.69	0.85	0.86	t
does_grow	0.37	0.52	0.81	0.81	e
is_a_plant	0.37	0.43	0.63	0.64	t
has_leaves	0.37	0.41	0.71	0.78	p
...
has_pips_seeds	0.47	0.5	0.08	0.46	p
is_juicy	0.5	0.71	0.48	0.56	op
is_a_vegetable	0.52	0.78	0.75	0.81	t
is_played	0.53	0.9	0.98	0.98	f
_does_play					
does_make_music	0.55	0.89	0.95	0.92	f
spearman-r		0.72	0.52	0.59	

Table 2: Performance of different approaches in relation to the average cosine similarity of words associated with a property (cos). The last row shows the Spearman Rank correlation between f1-scores and average cosine similarity. Property types are listed under type (p = part, vp = visual-perceptual, op = other-perceptual, e = encyclopaedic, f = functional, t = taxonomic).

Property	pos	neg
full_does_kill	101	69
crowd_does_kill	67	49
full_has_wheels	79	349
full_is_black	42	89
full_is_dangerous	177	104
crowd_is_dangerous	131	84
full_is_found_in_seas	83	72
crowd_is_found_in_seas	47	28
full_is_red	29	80
full_is_used_in_cooking	142	61
full_is_yellow	24	68
full_made_of_wood	87	282
full_is_an_animal_test	37	20
full_is_an_animal_train	166	77
full_is_food_test	37	20
full_is_food_train	97	146

Table 3: Class distribution in dataset consisting of the clean datasets derived from the CSLB set and the additional crowd judgments (marked *full_.*). For some properties, we included the dataset consisting of crowd-judgments only, as it is more balanced across semantic categories than the full set (marked *crowd_.*). For all properties, a leave-one-out approach was applied to evaluation except for *is_animal* and *is_food*.

4.2 Outcome Specific Hypotheses

We carry out further experiments on a small extended and clean subset, consisting of carefully selected negative examples from the CSLB dataset and crowd annotations validated by the authors. The distribution of positive and negative examples per property is shown in Table 3. For some properties, the sets derived from the CSLB norms alone have an imbalanced distribution of negative examples over semantic categories, as they were selected by means of logical exclusion (e.g. concepts listed under *has_wheels* have been selected as negative examples of *is_food*). Therefore, we add the more balanced but smaller datasets created by crowd-judgments only where enough judgments have been collected. We created additional sets for words part of the food-animal polysemy to test whether supervised classifiers can successfully predict semantic properties of various senses of polysemous words. In the following sections, we will outline the most striking results. Most results confirm, but some contradict our initial hypotheses.

Table 4 shows the f1-scores on the full clean datasets. As hypothesized, the color properties *is_yellow* and *is_red* perform low in all approaches, with slightly better results yielded by supervised learning.

The properties involved in functions and activi-

ties or with high impact on the interaction of entities with the world all perform highly in the classification approaches. For *does_kill*, *is_dangerous* and *is_used_in_cooking*, there is a large difference between the best nearest neighbors approach and the best classification approach (between 60 and 19 points), indicating that the classification approaches are able to infer more information from individual dimensions than is provided by full vector similarity. The property *is_dangerous* has, as can be expected, a particularly high diversity of associated words (comparable to the colors). *Has_wheels* and *is_found_in_seas* can be expected to have high correlations with other taxonomic categories (fish and water animals, vehicles), which is reflected in the lower diversity and comparatively high nearest neighbor performance.

Cases contradicting our expectations are the visual properties *is_black* and *made_of_wood*. Both have comparatively high classification performance with a big difference to the nearest neighbor results. Most likely, this is due to a category bias in the negative examples. For instance, a large portion of the negative examples for *is_made_of_wood* consist of animals and food. In the dataset for *is_black*, a large proportion of the positive examples consists of animals. A classifier can perform highly by simply learning to distinguish these two categories from the rest.

The biases in semantic classes mentioned above partially result from the way we generated the negative examples from the original CSLB dataset. This means that a classifier may learn to distinguish two semantic categories rather than being able to find vector dimensions indicative of the target property. We therefore also present selected results on crowd-only datasets shown in Table 4, which do not have this bias. It can be observed that for all three properties,¹⁰ the performance of the classification approaches drops marginally, whereas it rises for nearest neighbors.

We investigate the outcome on a number of individual examples to gain more insights into whether the subtle differences hypothesized in Section 3 hold. Since we only formulate a general hypothesis for Sparse Textual Evidence, we do not dive deeper into the results for that category here.

¹⁰We only included properties for which we had enough positive and negative examples in our set

Fine-Grained Category Distinctions

The full clean *has_wheels* dataset includes a number of instances for which the classifiers can make more fine-grained distinctions than nearest neighbors. As hypothesized, classifiers, in contrast to nearest-neighbors, can recognize that neither *sled* nor a *skidoo* have wheels, but a *unicycle*, a *limousine*, a *train*, *carriage*, an *ambulance*, a *porsche* do. Another fine-grained distinction can be identified in the *is_found_in_seas* crowd-only set: *Sculpin* is correctly identified as a seawater fish by all classifiers but not by nearest-neighbors.

Mixed Groups

Whereas nearest neighbors predominantly identify weapons as *is_dangerous* in the crowd-only set, the classifiers go beyond this category. The neural network approach correctly identifies that *imitation pistol*, *imitation handgun*, and *screwdriver* are negative examples of *is_dangerous*. Furthermore, no animals are labeled as dangerous based on proximity to the centroid, but the classifiers are able to distinguish between some dangerous and non-dangerous animals (e.g. *rhinoceros* is labeled positive, while *giraffe* and *zebra* are labeled as negative). All three classifiers recognize that *meth*, *cocaine* and *oxycodone* are considered dangerous substances, despite the fact that they are far away from the centroid of dangerous things. Of the only two disease-like concepts, *Hepatitis C* and *allergy*, the former is recognized by all classifiers and the latter only by logistic regression. The performance on the smaller, but also weapon-dominated *does_kill* crowd-only set is comparable, but the variety of atypical cases is lower. Among the only two disease-related items, *dengue* is identified by all classifiers and *dengue virus* only by the neural network.

In the crowd-only *is_found_in_seas* set, *seabird* and *gannet* are correctly labeled as positive, even though positive examples almost exclusively consist of fish or underwater-animals, whereas the negative examples encompass a vast variety of animals, including *bird* and some freshwater fish.

Polysemy

For polysemy between food and animals (Table 4), we observe that when trained on pure animal and food words and tested on polysemous animal and food words, the classifiers perform highly with a large difference to nearest neighbors. For food versus pure animal words, the classifier perfor-

property	av-cos	neigh	lr	net1	net2
full_is_yellow	0.23	0.19	0.47	0.64	0.64
full_is_used_in_cooking	0.37	0.29	0.98	0.98	0.98
full_is_black	0.19	0.35	0.75	0.77	0.77
full_is_red	0.23	0.36	0.51	0.54	0.52
full_is_dangerous	0.24	0.58	0.88	0.88	0.87
crowd_is_dangerous	0.26	0.61	0.86	0.86	0.86
full_has_wheels	0.38	0.90	0.96	0.96	0.95
full_is_found_in_seas	0.44	0.87	0.97	0.98	0.98
crowd_is_found_in_seas	0.50	0.87	0.94	0.96	0.96
full_does_kill	0.27	0.67	0.83	0.86	0.82
crowd_does_kill	0.30	0.70	0.82	0.84	0.80
full_made_of_wood	0.17	0.14	0.84	0.85	0.85
full_is_food_test	0.37	0.00	0.36	0.36	0.36
full_is_animal_test	0.37	0.52	0.88	0.88	0.88

Table 4: F1 scores achieved by logistic regression (lr) two runs of a neural net classifier (net1 and net2 and the n-best nearest neighbors evaluated with leave-one-out on the full datasets (marked as *full_* and the crowd-only sets (marked as *crowd_*).

mance is much lower. We expect the extremely low nearest neighbor performance to be due to the fact that the centroid is calculated over pure food items (without a single animal-related item, not even culinary meat terms such as *pork* or *beef*) which is far away from the animal-region in the space. Despite the classifiers outperforming nearest neighbors, the outcome does not confirm our original hypotheses. We expected that the classifiers could identify that edible animals have both animal properties and food properties, but upon inspection of the results, the classifiers only identified entities with a predominant animal sense correctly as animals and those with a predominant food sense correctly as food.

5 Discussion & Future Work

The experiments presented in this approach have several limitations. First, our semantic datasets are still limited in size. Second, the implication method we applied to generate negative examples led to biases for some properties where most negative examples belong to a small set of (taxonomic) classes. Third, no parameter tuning has been carried out so far. Careful parameter tuning would ensure that the best possible classification approaches are chosen and that the obtained results truly exploit the informative power of the embeddings. Due to the limited size of the dataset and the leave-one-out approach to evaluation, this

has not been possible in this preliminary study. Fourth, the experiments presented here only concern a small subsection of semantic properties too limited to draw general conclusions.

Despite these limitations, our results provide preliminary insights that lead us to conclude that the overall idea behind our methods works and opens up promising directions for future work. We first aim to address the limitations of the current dataset. We intend to incorporate other sets designed for similar insights, such as the analogies presented in [Drozd et al. \(2016\)](#) and the SemEval 2018 discriminative property set ([Krebs et al., 2018](#)). In addition, we plan to extend and refine the sets with crowd annotations asking for graded judgments (e.g. a property can *mostly* or *possibly* apply) and exploit these judgments in future experiments.

Once we created a bigger and more balanced dataset, we can carry out experiments on different train and test splits in order to overcome the limitations of the leave-one-out evaluation. Furthermore, we will apply careful parameter tuning on a development set in order to ensure our results are representative of the information captured by the embeddings. The increased size of the set will allow us to conduct more experiments that take the distributions of semantic categories in the splits into account as was done for the polysemy set. This way, we ensure that we do not train on examples that belong to the same semantic category as the ones in the test set.

Going beyond the method introduced in this paper, we plan on investigating the type of information encoded in linguistic context by testing which properties can be learned from textual context directly. In addition, applying the method presented by [Faruqui et al. \(2015\)](#) may provide stronger indications about the information represented by word embedding dimensions. Adding these to experiments allows us to trace which information is provided by the context and what ends up being present in word embeddings.

6 Conclusion

The main contribution of this paper is that it introduces a new method aimed at investigating the kind of semantic information captured by word embedding vectors. We have taken the first steps towards constructing a dataset suitable for this investigation on the basis of an existing dataset

of human-elicited semantic properties. We introduced a set of hypotheses concerning which semantic properties are captured by embeddings and presented exploratory experiments verifying them.

The current results are limited by the size and balance of our dataset, as discussed in detail in the previous section. Nevertheless, we can report preliminary insights based on our experiments. We show that classifiers, in particular neural networks, can identify which entities have a specific property in cases where this does not follow from general similarity or the overall semantic class the entity belongs to. This can be seen as a first indication that (some) semantic properties are encoded in individual (patterns of) vector dimensions, which can be identified.

The results on the extended datasets partly confirm that visual properties are not well represented by embeddings, while properties relating to function (e.g. cooking, having wheels) and interactions with other entities (e.g. being dangerous or killing) tend to be represented well. Some of these indications could be the result of the bias in our current dataset, but others have been confirmed on the smaller crowd-only sets for properties with enough available data (*is_dangerous* and *does_kill*). Further evidence is provided by the full dataset for *has_wheels* which encompasses a large group of vehicles to which the property does not apply. In addition, we support these indications by qualitative insights through examples of the kinds of distinctions made by the classifiers, but not the nearest neighbor approach. Results achieved for polysemous words and two visual properties currently do not confirm our hypotheses.

Acknowledgements

This research is funded by the PhD in the Humanities Grant provided by the Netherlands Organization of Scientific Research (Nederlandse Organisatie voor Wetenschappelijk Onderzoek, NWO) PGW.17.041 awarded to Pia Sommerauer and NWO VENI grant 275-89-029 awarded to Antske Fokkens. We would like to thank anonymous reviewers for feedback and Piek Vossen for feedback and discussion that helped improve this paper. All remaining errors are our own.

References

Eduard Barbu. 2008. Combining methods to learn feature-norm-like concept descriptions. In *Proceed-*

- ings of the *ESSLLI Workshop on Distributional Lexical Semantics*, pages 9–16.
- Marco Baroni, Brian Murphy, Eduard Barbu, and Massimo Poesio. 2010. Strudel: A corpus-based semantic model based on properties and types. *Cognitive science*, 34(2):222–254.
- Barry J. Devereux, Lorraine K. Tyler, Jeroen Geertzen, and Billi Randall. 2014. The centre for speech, language and the brain (cslb) concept property norms. *Behavior research methods*, 46(4):1119–1127.
- Aleksandr Drozd, Anna Gladkova, and Satoshi Matsuo. 2016. Word embeddings, analogies, and machine learning: Beyond king-man+ woman= queen. In *COLING*, pages 3519–3530.
- Luana Fagarasan, Eva Maria Vecchi, and Stephen Clark. 2015. From distributional semantics to feature norms: grounding semantic models in human perceptual data. In *Proceedings of the 11th International Conference on Computational Semantics*, pages 52–57.
- Manaal Faruqui, Yulia Tsvetkov, Dani Yogatama, Chris Dyer, and Noah A. Smith. 2015. Sparse overcomplete word vector representations. In *Proceedings of ACL*.
- Anna Gladkova and Aleksandr Drozd. 2016. Intrinsic evaluations of word embeddings: What can we do better? In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 36–42.
- Anna Gladkova, Aleksandr Drozd, and Satoshi Matsuo. 2016. Analogy-based detection of morphological and semantic relations with word embeddings: what works and what doesn't. In *Proceedings of the NAACL Student Research Workshop*, pages 8–15.
- Colin Kelly, Barry Devereux, and Anna Korhonen. 2014. Automatic extraction of property norm-like data from large text corpora. *Cognitive Science*, 38(4):638–682.
- Alicia Krebs, Alessandro Lenci, and Denis Paperno. 2018. Semeval-2018 task 10: Capturing discriminative attributes. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 732–740.
- Angeliki Lazaridou, Elia Bruni, and Marco Baroni. 2014. Is this a wampimuk? cross-modal mapping between distributional semantics and the visual world. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1403–1414.
- Angeliki Lazaridou, Eva Maria Vecchi, and Marco Baroni. 2013. Fish transporters and miracle homes: How compositional distributional semantics can help np parsing. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1908–1913.
- Omer Levy and Yoav Goldberg. 2014. Linguistic regularities in sparse and explicit word representations. In *Proceedings of the eighteenth conference on computational natural language learning*, pages 171–180.
- Tal Linzen. 2016. Issues in evaluating semantic spaces using word analogies. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 13–18.
- Ken McRae, George S Cree, Mark S Seidenberg, and Chris McNorgan. 2005. Semantic feature production norms for a large set of living and nonliving things. *Behavior research methods*, 37(4):547–559.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013. Linguistic regularities in continuous space word representations. In *HLT-NAACL*, volume 13, pages 746–751.
- Anna Rogers, Aleksandr Drozd, and Bofang Li. 2017. The (too many) problems of analogical reasoning with word vectors. In *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (*SEM 2017)*, pages 135–148.
- Stephen Roller and Sabine Schulte im Walde. 2013. A multimodal lda model integrating textual, cognitive and visual modalities. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1146–1157.
- Richard Socher, Alex Perelygin, Jean Y. Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Yulia Tsvetkov, Manaal Faruqui, and Chris Dyer. 2016. Correlation-based intrinsic evaluation of word vector representations. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 111–115.
- Yulia Tsvetkov, Manaal Faruqui, Wang Ling, Guillaume Lample, and Chris Dyer. 2015. Evaluation of word vector representations by subspace alignment. In *Proc. of EMNLP*.
- Peter D Turney. 2012. Domain and function: A dual-space model of semantic relations and compositions. *Journal of Artificial Intelligence Research*, 44:533–585.
- David P. Vinson and Gabriella Vigliocco. 2008. Semantic feature production norms for a large set of objects and events. *Behavior Research Methods*, 40(1):183–190.

- Yadollah Yaghoobzadeh and Hinrich Schütze. 2016. Intrinsic subspace evaluation of word embedding representations. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 236–246.
- Jie Zhou and Wei Xu. 2015. End-to-end learning of semantic role labeling using recurrent neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1127–1137.