

SimVerb-3500: A Large-Scale Evaluation Set of Verb Similarity

Daniela Gerz¹, Ivan Vulić¹, Felix Hill¹, Roi Reichart², and Anna Korhonen¹

¹Language Technology Lab, DTAL, University of Cambridge

²Faculty of Industrial Engineering and Management, Technion, IIT

Abstract

Verbs play a critical role in the meaning of sentences, but these ubiquitous words have received little attention in recent distributional semantics research. We introduce SimVerb-3500, an evaluation resource that provides human ratings for the similarity of 3,500 verb pairs. SimVerb-3500 covers all normed verb types from the USF free-association database, providing at least three examples for every VerbNet class. This broad coverage facilitates detailed analyses of how syntactic and semantic phenomena together influence human understanding of verb meaning. Further, with significantly larger development and test sets than existing benchmarks, SimVerb-3500 enables more robust evaluation of representation learning architectures and promotes the development of methods tailored to verbs. We hope that SimVerb-3500 will enable a richer understanding of the diversity and complexity of verb semantics and guide the development of systems that can effectively represent and interpret this meaning.

1 Introduction

Verbs are famously both complex and variable. They express the semantics of an event as well the relational information among participants in that event, and they display a rich range of syntactic and semantic behaviour (Jackendoff, 1972; Gruber, 1976; Levin, 1993). Verbs play a key role at almost every level of linguistic analysis. Information related to their predicate argument structure can benefit many NLP tasks (e.g. parsing, semantic role labelling, information extraction) and applications (e.g. machine translation,

text mining) as well as research on human language acquisition and processing (Korhonen, 2010). Precise methods for representing and understanding verb semantics will undoubtedly be necessary for machines to interpret the meaning of sentences with similar accuracy to humans.

Numerous algorithms for acquiring word representations from text and/or more structured knowledge bases have been developed in recent years (Mikolov et al., 2013; Pennington et al., 2014; Faruqui et al., 2015). These representations (or *embeddings*) typically contain powerful features that are applicable to many language applications (Collobert and Weston, 2008; Turian et al., 2010). Nevertheless, the predominant approaches to distributed representation learning apply a single learning algorithm and representational form for all words in a vocabulary. This is despite evidence that applying different learning algorithms to word types such as nouns, adjectives and verbs can significantly increase the ultimate usefulness of representations (Schwartz et al., 2015).

One factor behind the lack of more nuanced word representation learning methods is the scarcity of satisfactory ways to evaluate or analyse representations of particular word types. Resources such as MEN (Bruni et al., 2014), Rare Words (Luong et al., 2013) and SimLex-999 (Hill et al., 2015) focus either on words from a single class or small samples of different word types, with automatic approaches already reaching or surpassing the inter-annotator agreement ceiling. Consequently, for word classes such as *verbs*, whose semantics is critical for language understanding, it is practically impossible to achieve statistically robust analyses and comparisons between different

representation learning architectures.

To overcome this barrier to verb semantics research, we introduce *SimVerb-3500* – an extensive intrinsic evaluation resource that is unprecedented in both size and coverage. *SimVerb-3500* includes 827 verb types from the University of South Florida Free Association Norms (USF) (Nelson et al., 2004), and at least 3 member verbs from each of the 101 top-level VerbNet classes (Kipper et al., 2008). This coverage enables researchers to better understand the complex diversity of syntactic-semantic verb behaviours, and provides direct links to other established semantic resources such as WordNet (Miller, 1995) and PropBank (Palmer et al., 2005). Moreover, the large standardised development and test sets in *SimVerb-3500* allow for principled tuning of hyperparameters, a critical aspect of achieving strong performance with the latest representation learning architectures.

In § 2, we discuss previous evaluation resources targeting verb similarity. We present the new *SimVerb-3500* data set along with our design choices and the pair selection process in § 3, while the annotation process is detailed in § 4. In § 5 we report the performance of a diverse range of popular representation learning architectures, together with benchmark performance on existing evaluation sets. In § 6, we show how *SimVerb-3500* enables a variety of new linguistic analyses, which were previously impossible due to the lack of coverage and scale in existing resources.

2 Related Work

A natural way to evaluate representation quality is by judging the similarity of representations assigned to similar words. The most popular evaluation sets at present consist of word pairs with similarity ratings produced by human annotators.¹ Nevertheless, we find that all available datasets of this kind are insufficient for judging verb similarity due to their small size or narrow coverage of verbs.

In particular, a number of word pair evaluation sets are prominent in the distributional semantics

¹In some existing evaluation sets pairs are scored for relatedness which has some overlap with similarity. *SimVerb-3500* focuses on similarity as this is a more focused semantic relation that seems to yield a higher agreement between human annotators. For a broader discussion see (Hill et al., 2015).

literature.

Representative examples include RG-65 (Rubenstein and Goodenough, 1965) and WordSim-353 (Finkelstein et al., 2002; Agirre et al., 2009) which are small (65 and 353 word pairs, respectively). Larger evaluation sets such as the Rare Words evaluation set (Luong et al., 2013) (2034 word pairs) and the evaluations sets from Silberer and Lapata (2014) are dominated by noun pairs and the former also focuses on low-frequency phenomena. Therefore, these datasets do not provide a representative sample of verbs (Hill et al., 2015).

Two datasets that do focus on verb pairs to some extent are the data set of Baker et al. (2014) and Simlex-999 (Hill et al., 2015). These datasets, however, still contain a limited number of verb pairs (134 and 222, respectively), making them unrepresentative of the rich variety of verb semantic phenomena.

In this paper we provide a remedy for this problem by presenting a more comprehensive and representative verb pair evaluation resource.

3 The *SimVerb-3500* Data Set

In this section, we discuss the design principles behind *SimVerb-3500*. We first demonstrate that a new evaluation resource for verb similarity is a necessity. We then describe how the final verb pairs were selected with the goal to be representative, that is, to guarantee a wide coverage of two standard semantic resources: USF and VerbNet.

3.1 Design Motivation

Hill et al. (2015) argue that comprehensive high-quality evaluation resources have to satisfy the following three criteria: (C1) *Representative* (the resource covers the full range of concepts occurring in natural language); (C2) *Clearly defined* (it clearly defines the annotated relation, e.g., similarity); (C3) *Consistent and reliable* (untrained native speakers must be able to quantify the target relation consistently relying on simple instructions).

Building on the same annotation guidelines as Simlex-999 that explicitly targets similarity, we ensure that criteria C2 and C3 are satisfied. However, even SimLex, as the most extensive evaluation resource for verb similarity available at present, is still of limited size, spanning only 222 verb pairs and 170

distinct verb lemmas in total. Given that 39 out of the 101 top-level VerbNet classes are not represented at all in SimLex, while 20 classes have only one member verb,² one may conclude that the criterion C1 is not at all satisfied with current resources.

There is another fundamental limitation of all current verb similarity evaluation resources: automatic approaches have reached or surpassed the inter-annotator agreement ceiling. For instance, while the average pairwise correlation between annotators on SL-222 is Spearman’s ρ correlation of 0.717, the best performing automatic system reaches $\rho = 0.727$ (Mrkšić et al., 2016). SimVerb-3500 does not inherit this anomaly (see Tab. 2) and demonstrates that there still exists an evident gap between the human and system performance.

In order to satisfy C1-C3, the new SimVerb-3500 evaluation set contains similarity ratings for 3,500 *verb pairs*, containing 827 verb types in total and 3 member verbs for each top-level VerbNet class. The rating scale goes from 0 (not similar at all) to 10 (synonymous). We employed the SimLex-999 annotation guidelines. In particular, we instructed annotators to give low ratings to antonyms, and to distinguish between similarity and relatedness. Pairs that are related but not similar (e.g., *to snore / to snooze*, *to walk / to crawl*) thus have a fairly low rating. Several example pairs are provided in Tab. 1.

3.2 Choice of Verb Pairs and Coverage

To ensure a wide coverage of a variety of syntactico-semantic phenomena (C1), the choice of verb pairs is steered by two standard semantic resources available online: (1) the USF norms data set³ (Nelson et al., 2004), and (2) the VerbNet verb lexicon⁴ (Kipper et al., 2004; Kipper et al., 2008).

The USF norms data set (further USF) is the largest database of free association collected for English. It was generated by presenting human subjects with one of 5,000 cue concepts and asking them to write the first word coming to mind that is associated with that concept. Each cue concept c was normed in

²Note that verbs in VerbNet are soft clustered, and one verb type may be associated with more than one class. When computing coverage, we assume that such verbs attribute to counts of all their associated classes.

³<http://w3.usf.edu/FreeAssociation/>

⁴<http://verbs.colorado.edu/verb-index/>

Pair	Rating
to reply / to respond	9.79
to snooze / to nap	8.80
to cook / to bake	7.80
to participate / to join	5.64
to snore / to snooze	4.15
to walk / to crawl	2.32
to stay / to leave	0.17
to snooze / to happen	0.00

Table 1: Example verb pairs from SimVerb-3500.

this way by over 10 participants, resulting in a set of associates a for each cue, for a total of over 72,000 (c, a) pairs. For each such pair, the proportion of participants who produced associate a when presented with cue c can be used as a proxy for the strength of association between the two words.

The norming process guarantees that two words in a pair have a degree of semantic association which correlates well with semantic relatedness and similarity. Sampling from the USF set ensures that both related but non-similar pairs (e.g., *to run / to sweat*) as well as similar pairs (e.g., *to reply / to respond*) are represented in the final list of pairs. Further, the rich annotations of the output USF data (e.g., concreteness scores, association strength) can be directly combined with the SimVerb-3500 similarity scores to yield additional analyses and insight.

VerbNet (VN) is the largest online verb lexicon currently available for English. It is hierarchical, domain-independent, and broad-coverage. VN is organised into verb classes extending the classes from Levin (1993) through further refinement to achieve syntactic and semantic coherence among class members. According to the official VerbNet guidelines,⁵ “Verb Classes are numbered according to shared semantics and syntax, and classes which share a top-level number (9-109) have corresponding semantic relationships.” For instance, all verbs from the top-level Class 9 are labelled “Verbs of Putting”, all verbs from Class 30 are labelled “Verbs of Perception”, while Class 39 contains “Verbs of Ingesting”.

Among others, three basic types of information are covered in VN: (1) verb subcategorization frames (SCFs), which describe the syntactic realization of the predicate-argument structure (e.g. *The window broke*), (2) selectional preferences (SPs), which capture the semantic preferences verbs have for their

⁵http://verbs.colorado.edu/verb-index/VerbNet_Guidelines.pdf

arguments (e.g. *a breakable physical object* broke) and (3) lexical-semantic verb classes (VCs) which provide a shared level of abstraction for verbs similar in their (morpho-)syntactic and semantic properties (e.g. *BREAK verbs*, sharing the VN class 45.1, and the top-level VN class 45).⁶ The basic overview of the VerbNet structure already suggests that measuring verb similarity is far from trivial as it revolves around a complex interplay between various semantic and syntactic properties.

The wide coverage of VN in SimVerb-3500 assures the wide coverage of distinct verb groups/classes and their related linguistic phenomena. Finally, VerbNet enables further connections of SimVerb-3500 to other important lexical resources such as FrameNet (Baker et al., 1998), WordNet (Miller, 1995), and PropBank (Palmer et al., 2005) through the sets of mappings created by the SemLink project initiative (Loper et al., 2007).⁷

Sampling Procedure We next sketch the complete sampling procedure which resulted in the final set of 3500 distinct verb pairs finally annotated in a crowdsourcing study (§ 4).

(Step 1) We extracted all possible verb pairs from USF based on the associated POS tags available as part of USF annotations. To ensure that semantic association between verbs in a pair is not accidental, we then discarded all such USF pairs that had been associated by 2 or less participants in USF.

(Step 2) We then manually cleaned and simplified the list of pairs by removing all pairs with multi-word verbs (e.g., *quit / give up*), all pairs that contained the non-infinitive form of a verb (e.g., *accomplished / finished*, *hidden / find*), removing all pairs containing at least one auxiliary verb (e.g., *must / to see*, *must / to be*). The first two steps resulted in 3,072 USF-based verb pairs.

(Step 3) After this stage, we noticed that several top-level VN classes are not part of the extracted set. For instance, 5 VN classes did not have any member verbs included, 22 VN classes had only 1 verb, and 6 VN classes had 2 verbs included in the current set.

We resolved the VerbNet coverage issue by sampling from such 'under-represented' VN classes directly. Note that this step is not related to USF at

all. For each such class we sampled additional verb types until the class was represented by 3 or 4 member verbs (chosen randomly).⁸ Following that, we sampled at least 2 verb pairs for each previously 'under-represented' VN class by pairing 2 member verbs from each such class. This procedure resulted in 81 additional pairs, now 3,153 in total.

(Step 4) Finally, to complement this set with a sample of entirely unassociated pairs, we followed the SimLex-999 setup. We paired up the verbs from the 3,153 associated pairs at random. From these random pairings, we excluded those that coincidentally occurred elsewhere in USF (and therefore had a degree of association). We sampled the remaining 347 pairs from this resulting set of unassociated pairs.

(Output) The final SimVerb-3500 data set contains 3,500 verb pairs in total, covering all associated verb pairs from USF, and (almost) all top-level VerbNet classes. All pairs were manually checked post-hoc by the authors plus 2 additional native English speakers to verify that the final data set does not contain unknown or invalid verb types.

Frequency Statistics The 3,500 pairs consist of 827 distinct verbs. 29 top-level VN classes are represented by 3 member verbs, while the three most represented classes cover 79, 85, and 93 member verbs. 40 verbs are not members of any VN class.

We performed an initial frequency analysis of SimVerb-3500 relying on the BNC counts available online (Kilgarriff, 1997).⁹ After ranking all BNC verbs according to their frequency, we divided the list into quartiles: Q1 (most frequent verbs in BNC) - Q4 (least frequent verbs in BNC). Out of the 827 SimVerb-3500 verb types, 677 are contained in Q1, 122 in Q2, 18 in Q3, 4 in Q4 (*to enroll*, *to hitchhike*, *to implode*, *to whelp*), while 6 verbs are not covered in the BNC list. 2,818 verb pairs contain Q1 verbs, while there are 43 verb pairs with both verbs not in Q1. Further empirical analyses are provided in § 6.¹⁰

⁸The following three VN classes are exceptions: (1) Class 56, consisting of words that are dominantly tagged as nouns, but can be used as verbs exceptionally (e.g., *holiday*, *summer*, *honeymoon*); (2) Class 91, consisting of 2 verbs (*count*, *matter*); (3) Class 93, consisting of 2 single word verbs (*adopt*, *assume*).

⁹<https://www.kilgarriff.co.uk/bnc-readme.html>

¹⁰Annotations such as VerbNet class membership, relations between WordNet synsets of each verb, and frequency statistics are available as supplementary material.

⁶<https://verbs.colorado.edu/verb-index/vn/break-45.1.php>

⁷<https://verbs.colorado.edu/semink/>

4 Word Pair Scoring

We employ the Prolific Academic (PA) crowdsourcing platform,¹¹ an online marketplace very similar to Amazon Mechanical Turk and to CrowdFlower.

4.1 Survey Structure

Following the SimLex-999 annotation guidelines, we had each of the 3500 verb pairs rated by at least 10 annotators. To distribute the workload, we divided the 3500 pairs into 70 tranches, with 79 pairs each. Out of the 79 pairs, 50 are unique to one tranche, while 20 manually chosen pairs are in all tranches to ensure consistency. The remaining 9 are duplicate pairs displayed to the same participant multiple times to detect inconsistent annotations.

Participants see 7-8 pairs per page. Pairs are rated on a scale of 0-6 by moving a slider. The first page shows 7 pairs, 5 unique ones and 2 from the consistency set. The following pages are structured the same but display one extra pair from the previous page. Participants are explicitly asked to give these duplicate pairs the same rating. We use them as quality control so that we can identify and exclude participants giving several inconsistent answers.

Checkpoint Questions The survey contains three control questions in which participants are asked to select the most similar pair out of three choices. For instance, the first checkpoint is: *Which of these pairs of words is the *most* similar? 1. to run / to jog 2. to run / to walk 3. to jog / to sweat.* One checkpoint occurs right after the instructions and the other two later in the survey. The purpose is to check that annotators have understood the guidelines and to have another quality control measure for ensuring that they are paying attention throughout the survey. If just one of the checkpoint questions is answered incorrectly, the survey ends immediately and all scores from the annotator in question are discarded.

Participants 843 raters participated in the study, producing over 65,000 ratings. Unlike other crowdsourcing platforms, PA collects and stores detailed demographic information from the participants upfront. This information was used to carefully select the pool of eligible participants. We restricted the pool to native English speakers with a 90% approval

rate (maximum rate on PA), of age 18-50, born and currently residing in the US (45% out of 843 raters), UK (53%), or Ireland (2%). 54% of the raters were female and 46% male, with the average age of 30. Participants took 8 minutes on average to complete the survey containing 79 questions.

4.2 Post-Processing

We excluded ratings of annotators who (a) answered one of the checkpoint questions incorrectly (75% of exclusions); (b) did not give equal ratings to duplicate pairs; (c) showed suspicious rating patterns (e.g., randomly alternating between two ratings or using one single rating throughout). The final acceptance rate was 84%. We then calculated the average of all ratings from the accepted raters (≥ 10) for each pair. The score was finally scaled linearly from the 0-6 to the 0-10 interval as in (Hill et al., 2015).

5 Analysis

Inter-Annotator Agreement We employ two measures. **IAA-1 (pairwise)** computes the average pairwise Spearman’s ρ correlation between any two raters – a common choice in previous data collection in distributional semantics (Padó et al., 2007; Reisinger and Mooney, 2010a; Silberer and Lapata, 2014; Hill et al., 2015).

A complementary measure would smooth individual annotator effects. For this aim, our **IAA-2 (mean)** measure compares the average correlation of a human rater with the average of all the other raters. SimVerb-3500 obtains $\rho = 0.84$ (IAA-1) and $\rho = 0.86$ (IAA-2), a very good agreement compared to other benchmarks (see Tab. 2).

Vector Space Models We compare the performance of prominent representation models on SimVerb-3500. We include: (1) unsupervised models that learn from distributional information in text, including the skip-gram negative-sampling model (SGNS) with various contexts (*BOW = bag of words; DEPS = dependency contexts*) as in Levy and Goldberg (2014), the symmetric-pattern based vectors by Schwartz et al. (2015), and count-based PMI-weighted vectors (Baroni et al., 2014); (2) Models that rely on linguistic hand-crafted resources or curated knowledge bases. Here, we use sparse binary vectors built from linguistic resources (*Non-*

¹¹<https://prolific.ac/> (We chose PA for logistic reasons.)

Eval set	IAA-1	IAA-2	ALL	TEXT
WSIM (203)	0.67	0.65	0.79	0.79
SIMLEX (999)	0.67	0.78	0.74	0.56
			Paragram+CF	SymPat+SGNS
SL-222 (222)	0.72	-	0.73	0.58
			Paragram+CF	SymPat
SIMVERB (3500)	0.84	0.86	0.63	0.36
			Paragram+CF	SGNS-DEPS

Table 2: An overview of word similarity evaluation benchmarks. ALL is the current best reported score on each data set across all models (including the models that exploit curated knowledge bases and hand-crafted lexical resources, see supplementary material). TEXT denotes the best reported score for a model that learns solely on the basis of distributional information. All scores are Spearman’s ρ correlations.

Distributional, (Faruqui and Dyer, 2015)), and vectors fine-tuned to a paraphrase database (*Paragram*, (Wieting et al., 2015)) further refined using linguistic constraints (*Paragram+CF*, (Mrkšić et al., 2016)). Descriptions of these models are in the supplementary material.

Comparison to SimLex-999 (SL-222) 170 pairs from SL-222 also appear in SimVerb-3500. The correlation between the two data sets calculated on the shared pairs is $\rho = 0.91$. This proves, as expected, that the ratings are consistent across the two data sets.

Tab. 3 shows a comparison of models’ performance on SimVerb-3500 against SL-222. Since the number of evaluation pairs may influence the results, we ideally want to compare sets of equal size for a fair comparison. Picking one random subset of 222 pairs would bias the results towards the selected pairs, and even using 10-fold cross-validation we found variations up to 0.05 depending on which subsets were used. Therefore, we employ a 2-level 10-fold cross-validation where new random subsets are picked in each iteration of each model. The numbers reported as CV-222 are averages of these ten 10-fold cross-validation runs. The reported results come very close to the correlation on the full data set for all models.

Most models perform much better on SL-222, especially those employing additional databases or linguistic resources. The performance of the best scoring Paragram+CF model is even on par with the IAA-1 of 0.72. The same model obtains the highest score on SV-3500 ($\rho = 0.628$), with a clear gap to IAA-1 of 0.84. We attribute these differences in

performance largely to SimVerb-3500 being a more extensive and diverse resource in terms of verb pairs.

Development Set A common problem in scored word pair datasets is the lack of a standard split to development and test sets. Previous works often optimise models on the entire dataset, which leads to overfitting (Faruqui et al., 2016) or use custom splits, e.g., 10-fold cross-validation (Schwartz et al., 2015), which make results incomparable with others. The lack of standard splits stems mostly from small size and poor coverage – issues which we have solved with SimVerb-3500.

Our development set contains 500 pairs, selected to ensure a broad coverage in terms of similarity ranges (i.e., non-similar and highly similar pairs, as well as pairs of medium similarity are represented) and top-level VN classes (each class is represented by at least 1 member verb). The test set includes the remaining 3,000 verb pairs. The performances of representation learning architectures on the dev and test sets are reported in Tab. 3. The ranking of models is identical on the test and the full SV-3500 set, with slight differences in ranking on the development set.

6 Evaluating Subsets

The large coverage and scale of SimVerb-3500 enables model evaluation based on selected criteria. In this section, we showcase a few example analyses.

Frequency In the first analysis, we select pairs based on their lemma frequency in the BNC corpus and form three groups, with 390-490 pairs in each group (Fig. 1). The results from Fig. 1 suggest that the performance of all models improves as the frequency of the verbs in the pair increases, with much steeper curves for the purely distributional models (e.g., SGNS and SymPat). The non-distributional non data-driven model of Faruqui and Dyer (2015) is only slightly affected by frequency.

WordNet Synsets Intuitively, representations for verbs with more diverse usage patterns are more difficult to learn with statistical models. To examine this hypothesis, we resort to WordNet (Miller, 1995), where different semantic usages of words are listed as so-called *synsets*. Fig. 2 shows a clear downward trend for all models, confirming that polysemous

Model	SV-3500	CV-222	SL-222	DEV-500	TEST-3000
SGNS-BOW-PW (d=300)	0.274	0.279	0.328	0.333	0.265
SGNS-DEPS-PW (d=300)	0.313	0.314	0.390	0.401	0.304
SGNS-UDEPS-PW (d=300)	0.259	0.262	0.347	0.313	0.250
SGNS-BOW-8B (d=500)	0.348	0.343	0.307	0.378	0.350
SGNS-DEPS-8B (d=500)	0.356	0.347	0.385	0.389	0.351
SYMPAT-8B (d=500)	0.328	0.336	0.544	0.276	0.347
COUNT-SVD (d=500)	0.196	0.200	0.059	0.259	0.186
NON-DISTRIBUTIONAL	0.596	0.596	0.689	0.632	0.600
PARAGRAM (d=25)	0.418	0.432	0.531	0.443	0.433
PARAGRAM (d=300)	0.540	0.528	0.590	0.525	0.537
PARAGRAM+CF (d=300)	0.628	0.625	0.727	0.611	0.624

Table 3: Evaluation of state-of-the representation learning models on the full SimVerb-3500 set (SV-3500), the Simlex-999 verb subset containing 222 pairs (SL-222), cross-validated subsets of 222 pairs from SV-3500 (CV-222), and the SimVerb-3500 development (DEV-500) and test set (TEST-3000).

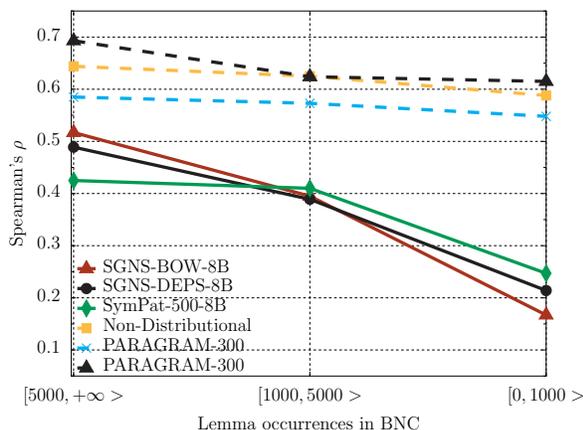


Figure 1: Subset-based evaluation, where subsets are created based on the frequency of verb lemmas in the BNC corpus. Each of the three frequency groups contains 390-490 verb pairs. To be included in each group it is required that both verbs in a pair are contained in the same frequency interval (x axis).

verbs are more difficult for current verb representation models. Nevertheless, approaches which use additional information beyond corpus co-occurrence are again more robust. Their performance only drops substantially for verbs with more than 10 synsets, while the performance of other models deteriorates already when tackling verbs with more than 5 synsets.

VerbNet Classes Another analysis enabled by SimVerb-3500 is investigating the connection between VerbNet classes and human similarity judgments. We find that verbs in the same top-level VerbNet class are often not assigned high similarity score. Out of 1378 pairs where verbs share the top-level VerbNet class, 603 have a score lower than 5. Tab. 4 reports scores per VerbNet class. When a verb be-

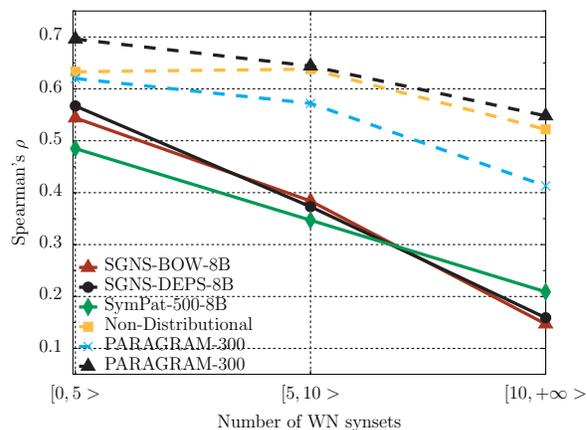


Figure 2: Subset-based evaluation, where subsets are created based on the number of synsets in WordNet (x axis). To be included in each subset it is required that both verbs in a pair have the number of synsets in the same interval.

longs to multiple classes, we count it for each class (see Footnote 2). We run the analysis on the five largest VN classes, each with more than 100 pairs with paired verbs belonging to the same class.

The results indicate clear differences between classes (e.g., Class 31 vs Class 51), and suggest that further developments in verb representation learning should also focus on constructing specialised representations at the finer-grained level of VN classes.

Lexical Relations SimVerb-3500 contains relation annotations (e.g., *antonyms*, *synonyms*, *hyper/hyponyms*, *no relation*) for all pairs extracted automatically from WordNet. Evaluating per-relation subsets, we observe that some models draw their strength from good performance across different re-

Model	#13	#31	#37	#45	#51
SGNS-BOW-8B	0.210	0.308	0.352	0.270	0.170
SGNS-DEPS-8B	0.289	0.270	0.306	0.238	0.225
<hr/>					
SYMPAT-8B (d=500)	0.171	0.320	0.143	0.195	0.113
<hr/>					
NON-DISTR	0.571	0.483	0.372	0.501	0.499
PARAGRAM (d=300)	0.571	0.504	0.567	0.531	0.387
PARAGRAM+CF	0.735	0.575	0.666	0.622	0.614

Table 4: Spearman’s ρ correlation between human judgments and model’s cosine similarity by VerbNet Class. We chose classes #13 *Verbs of Change of Possession*, #31 *Verbs of Psychological State*, #37 *Verbs of Communication*, #45 *Verbs of Change of State*, and #51 *Verbs of Motion* as examples. All are large classes with more than 100 pairs each, and the frequencies of member verbs are distributed in a similar way.

Model	NR	SYN	HYP
SGNS-BOW-PW (d=300)	0.096	0.288	0.292
SGNS-DEPS-PW (d=300)	0.132	0.290	0.336
SGNS-BOW-8B (d=500)	0.292	0.273	0.338
SGNS-DEPS-8B (d=500)	0.157	0.323	0.378
<hr/>			
SYMPAT-8B-DENSE (d=300)	0.225	0.182	0.265
SYMPAT-8B-DENSE (d=500)	0.248	0.260	0.251
<hr/>			
NON-DISTRIBUTIONAL	0.126	0.379	0.488
PARAGRAM (d=300)	0.254	0.356	0.439
PARAGRAM+CF (d=300)	0.250	0.417	0.475

Table 5: Spearman’s ρ correlation between human judgments and model’s cosine similarity based on pair relation type. Relations are based on WordNet, and included in the dataset. The classes are of different size, 373 pairs with no relation (*NR*), 306 synonym (*SYN*) pairs, and 800 hyper/hyponym (*HYP*) pairs. Frequencies of member verbs are distributed in a similar way.

lations. Others have low performance on these pairs, but do very well on synonyms and hyper-/hyponyms. Selected results of this analysis are in Tab. 5.¹²

Human Agreement Motivated by the varying performance of computational models regarding frequency and ambiguous words with many synsets, we analyse what disagreement effects may be captured in human ratings. We therefore compute the average standard deviation of ratings per subset: $avgstd(S) = \frac{1}{n} \sum_{p \in S} \sigma(r_p)$, where S is one subset of pairs, n is the number of pairs in this subset, p is one pair, and r_p are all human ratings for this pair.

¹² Evaluation based on Spearman’s ρ may be problematic with certain categories, e.g., with antonyms. It evaluates pairs according to their ranking; for antonyms the ranking is arbitrary - every antonym pair should have a very low rating, hence they are not included in Tab. 5. A similar effect occurs with highly ranked synonyms, but to a much lesser degree than with antonyms.

While the standard deviation of ratings is diverse for individual pairs, overall the average standard deviations per subset are almost identical. For both the frequency and the WordNet synset analyses it is around ≈ 1.3 across all subsets, and with only little difference for the subsets based on VerbNet. The only subsets where we found significant variations is the grouping by relations, where ratings tend to be more similar especially on antonyms (0.86) and pairs with no relation (0.92), much less similar on synonyms (1.34) and all other relations (≈ 1.4). These findings suggest that humans are much less influenced by frequency or polysemy in their understanding of verb semantics compared to computational models.

7 Conclusions

SimVerb-3500 is a verb similarity resource for analysis and evaluation that will be of use to researchers involved in understanding how humans or machines represent the meaning of verbs, and, by extension, scenes, events and full sentences. The size and coverage of syntactico-semantic phenomena in SimVerb-3500 makes it possible to compare the strengths and weaknesses of various representation models via statistically robust analyses on specific word classes.

To demonstrate the utility of SimVerb-3500, we conducted a selection of analyses with existing representation-learning models. One clear conclusion is that distributional models trained on raw text (e.g. SGNS) perform very poorly on low frequency and highly polysemous verbs. This degradation in performance can be partially mitigated by focusing models on more principled distributional contexts, such as those defined by symmetric patterns. More generally, the finding suggests that, in order to model the diverse spectrum of verb semantics, we may require algorithms that are better suited to fast learning from few examples (Lake et al., 2011), and have some flexibility with respect to sense-level distinctions (Reisinger and Mooney, 2010b; Vilnis and McCallum, 2015). In future work we aim to apply such methods to the task of verb acquisition.

Beyond the preliminary conclusions from these initial analyses, the benefit of SimLex-3500 will become clear as researchers use it to probe the relationship between architectures, algorithms and representation quality for a wide range of verb classes. Better under-

standing of how to represent the full diversity of verbs should in turn yield improved methods for encoding and interpreting the facts, propositions, relations and events that constitute much of the important information in language.

Acknowledgments

This work is supported by the ERC Consolidator Grant LEXICAL (648909).

References

- Eneko Agirre, Enrique Alfonseca, Keith B. Hall, Jana Kravalova, Marius Pasca, and Aitor Soroa. 2009. A study on similarity and relatedness using distributional and WordNet-based approaches. In *NAACL-HLT*, pages 19–27.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *ACL-COLING*, pages 86–90.
- Simon Baker, Roi Reichart, and Anna Korhonen. 2014. An unsupervised model for instance level subcategorization acquisition. In *EMNLP*, pages 278–289.
- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *ACL*, pages 238–247.
- Elia Bruni, Nam-Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49:1–47.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: deep neural networks with multitask learning. In *ICML*, pages 160–167.
- Manaal Faruqui and Chris Dyer. 2015. Non-distributional word vector representations. In *ACL*, pages 464–469.
- Manaal Faruqui, Jesse Dodge, Sujay Kumar Jauhar, Chris Dyer, Eduard H. Hovy, and Noah A. Smith. 2015. Retrofitting word vectors to semantic lexicons. In *NAACL-HLT*, pages 1606–1615.
- Manaal Faruqui, Yulia Tsvetkov, Pushpendre Rastogi, and Chris Dyer. 2016. Problems with evaluation of word embeddings using word similarity tasks. *CoRR*, abs/1605.02276.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppín. 2002. Placing search in context: The concept revisited. *ACM Transactions on Information Systems*, 20(1):116–131.
- Jeffrey Gruber. 1976. *Lexical structure in syntax and semantics*.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. SimLex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- Ray S. Jackendoff. 1972. *Semantic interpretation in generative grammar*.
- Adam Kilgarriff. 1997. Putting frequencies in the dictionary. *International Journal of Lexicography*, 10(2):135–155.
- Karin Kipper, Benjamin Snyder, and Martha Palmer. 2004. Extending a verb-lexicon using a semantically annotated corpus. In *LREC*, pages 1557–1560.
- Karin Kipper, Anna Korhonen, Neville Ryant, and Martha Palmer. 2008. A large-scale classification of English verbs. *Language Resource and Evaluation*, 42(1):21–40.
- Anna Korhonen. 2010. Automatic lexical classification: bridging research and practice. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 368(1924):3621–3632.
- Brenden M. Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua B. Tenenbaum. 2011. One shot learning of simple visual concepts. In *CogSci*.
- Beth Levin. 1993. *English verb classes and alternation, A preliminary investigation*. The University of Chicago Press.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *ACL*, pages 302–308.
- Edward Loper, Szu-Ting Yi, and Martha Palmer. 2007. Combining lexical resources: Mapping between PropBank and VerbNet. In *IWCS*.
- Thang Luong, Richard Socher, and Christopher Manning. 2013. Better word representations with recursive neural networks for morphology. In *CoNLL*, pages 104–113.
- Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *ICLR: Workshop Papers*.
- George A. Miller. 1995. Wordnet: A lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Blaise Thomson, Milica Gašić, Lina Maria Rojas-Barahona, Pei-Hao Su, David Vandyke, Tsung-Hsien Wen, and Steve J. Young. 2016. Counter-fitting word vectors to linguistic constraints. In *NAACL-HLT*, pages 142–148.
- Douglas L. Nelson, Cathy L. McEvoy, and Thomas A. Schreiber. 2004. The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*, 36(3):402–407.
- Sebastian Padó, Ulrike Padó, and Katrin Erk. 2007. Flexible, corpus-based modelling of human plausibility judgements. In *EMNLP-CoNLL*, pages 400–409.

- Martha Palmer, Paul Kingsbury, and Daniel Gildea. 2005. The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Joseph Reisinger and Raymond J. Mooney. 2010a. A mixture model with sharing for lexical semantics. In *EMNLP*, pages 1173–1182.
- Joseph Reisinger and Raymond J Mooney. 2010b. Multi-prototype vector-space models of word meaning. In *NAACL-HTL*, pages 109–117.
- Herbert Rubenstein and John B Goodenough. 1965. Contextual correlates of synonymy. *Communications of the ACM*, 8(10):627–633.
- Roy Schwartz, Roi Reichart, and Ari Rappoport. 2015. Symmetric pattern based word embeddings for improved word similarity prediction. In *CoNLL*, pages 258–267.
- Carina Silberer and Mirella Lapata. 2014. Learning grounded meaning representations with autoencoders. In *ACL*, pages 721–732.
- Joseph P. Turian, Lev-Arie Ratinov, and Yoshua Bengio. 2010. Word representations: A simple and general method for semi-supervised learning. In *ACL*, pages 384–394.
- Luke Vilnis and Andrew McCallum. 2015. Word representations via Gaussian embedding. *ICLR*.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. From paraphrase database to compositional paraphrase model and back. *Transactions of the ACL*, 3:345–358.