Distributional semantics and the conceptual foundations of verb meaning: how neural word embeddings memorize the unaccusative hypothesis

Tillmann Pross
Institute for Natural Language Processing (IMS), University of Stuttgart
prosstn@ims.uni-stuttgart.de

Abstract

In the present paper, I investigate whether and how neural word embeddings can be understood to encode not only idiosyncratic aspects of word meaning but also the kind of general and abstract concepts that are central to theoretical approaches of lexical semantics. To this end, I compute the difference between general-purpose embeddings of intransitive verbs, and task-specific embeddings of the same verbs that capture their similarity according to the unaccusative hypothesis. I approximate the resulting set of difference vectors with their nearest neighbors in the embedding space and show that the nearest neighbors correspond to prototypical linguistic realizations of unergativity and unaccusativity. My study thus suggests that word embeddings may provide a novel and empirically grounded perspective on the conceptual underpinnings of verb meaning.

1 Introduction

Distributional approaches to word meaning (Turney and Pantel, 2010), and in particular neural word embeddings like Word2Vec (Mikolov et al., 2013a) or Glove (Pennington et al., 2014) have so far found little echo in the mainstream of theoretically oriented work on lexical semantics and the syntax-semantics interface. This paper ties in with recent work (such as Asher et al. (2016), McNally and Boleda (2017), Pross et al. (2017)) that aims to show that this situation is to the detriment of both computational and theoretical approaches to lexical semantics. A main goal of lexical semantics is to decompose word meaning into idiosyncratic and recurring elements of meaning. But the type, number and determination of the recurring conceptual features relevant to verb meaning is a central, yet unsolved research problem in both theoretical and computational approaches to verb meaning: "[t]he important theoretical construct is the notion of meaning component, not the notion of verb class" (Levin, 1993, p. 18). On the one hand, the fact that distributional representations can be used to reproduce a theoretically defined gold standard of verb classification does not indicate what the concepts or semantic features are like that underlie the classification, as e.g. Lenci (2014) argues with a case study on the distributional classification of Italian verbs. On the other, popular theoretical approaches to the lexical semantics of verbs, like alternation-based lexical-conceptual semantics (Guerssel et al., 1985; Pinker, 2013; Levin, 1993; Levin and Rappaport Hovav, 1995) stipulate concepts rather than deriving them from empirical observations, as e.g. Van der Leek (1996) argues with a case study on the conative alternation.

Striving for a combination of methods from theoretical lexical semantics and distributional semantics seems a natural way to deal with the persistent problem of identifying the recurrent elements of verb meaning. But such a combination poses at first a methodological challenge. Both theoretical and computational approaches to lexical semantics have clearly defined and widely agreed methodological standards. A middle ground between theoretical and computational semantics is thus likely to fall short of the established expectations of both theoretical and computational approaches to lexical semantics. On the one hand, computational approaches to lexical semantics assess word embeddings according to their extrinsic value, which can be measured e.g. by their performance in downstream tasks. But the extrinsic assessment of word embeddings leaves open the question for whether and how word embeddings also have an intrinsic explanatory value. On the other, since the main goal of theoretical lexical semantics is to render precise and transparent the meaning of words, the way in which meaning is represented in word embeddings easily appears to be unsystematic and opaque. Keeping this methodological caveat in mind, the problem with which the present paper is concerned is that work like Levy and Goldberg (2014b) fosters relatively clear-cut intuitions about how neural word embeddings represent the meaning of particular words, and in particular nouns. But we have no similarly clear-cut intuitions about how word embeddings, and in particular verb embeddings, may be understood to characterize abstract concepts such as "agency" or "change of state". The problem can also be formulated the other way round: theoretical work on lexical semantics assigns concepts like "agency" or "change of state" a central role in the definition of thematic roles, but "[t]here is perhaps no concept in modern syntactic and semantic theory which is so often involved in so wide a range of contexts, but on which there is so little agreement as to its nature and definition, as THEMATIC ROLE" (Dowty, 1991, p. 547). The research question which the present paper aims to investigate is thus the following: what are the "cues" or "indicators" for recurrent meaning elements of verb meaning that the informed linguist should search for when inspecting interpretable verb embeddings?

The paper explores the question for whether and how word embeddings encode recurrent elements of verb meaning with the example of intransitive verbs. Intransitive verbs are insofar an interesting object of study, as the unaccusative hypothesis of Perlmutter (1978) predicts that intransitive verbs decompose into two distinct subclasses with a clearly delimited syntax and semantics, unaccusative and unergative verbs. To foster the search for recurrent elements of verb meaning in the word embeddings of unaccusative and unergative verbs, I first retrain general-purpose word embeddings of German and English intransitive verbs with the objective of learning to distinguish between unaccusative and unergative verbs, using the semantic role of the single argument of an intransitive verb as label (Agent for unergative, Patient for unaccusative). I then subtract the retrained task-specific embeddings from the original general-purpose embeddings and approximate the difference vectors with their nearest neighbors. These interpretable difference vectors capture unergativity and unaccusativity in the form of abstract "unergative vectors" and "unaccusative vectors" with a small set of recurring elements (i.e. nearest neighbors). I observe that these elements correspond to prototypical linguistic realizations of unergativity, like *-er* nominals for the unergative vector and descriptions of "natural" processes for the unaccusative vector. I conclude that since neural word embeddings are learned from data independent of theoretical bias and intuitions, my study suggests that neural word embeddings can be understood as providing independent empirical grounding of the semantic characterization of the distinction between unergative and unaccusative verbs proposed in the theoretical literature. Moreover, while neural word embeddings are very good at capturing specific topical similarities, our case study suggests that word embeddings it is an open question whether and how word embeddings can also capture more general and abstract concepts like agency or causation.

2 Background and previous work

2.1 Topical vs. conceptual similarity

One reason for the perceived gap between theoretical and computational approaches to word meaning is that neural word embeddings encode meaning in dense and continuous representations grounded in huge amounts of corpus data and thus do not allow for direct inspection or intuitive evaluation by humans. In contrast, lexical-conceptual representations represent word meaning with discrete feature structures grounded in human intuitions and grammaticality judgments about alternation behavior and thus allow for direct inspection and intuitive evaluation by humans. Another reason is the quite different concepts of meaning that computational and theoretical approaches embody. According to the distributional hypothesis embodied by neural word embeddings, words that occur in similar contexts tend to have similar meanings (Turney and Pantel, 2010), where similarity can be measured e.g. with the magnitude of the dot-product between two word embeddings (Levy and Goldberg, 2014a). Using the experimental setup described in section 3, (1) lists the 5 most similar words of *to eat* in the embedding space.

(1) *5 nearest neighbours of to eat*
drink.V cook.V diet.N snack.N munch.V

The nearest neighbors in (1) all have to do with the *topic* of 'nutrition', thus occur in similar contexts, and the high dot product between the embeddings of these words in the embedding space reflect this "topical" similarity. But to the best of our knowledge, there is no language which marks or distinguishes words related to the topic of nutrition. According to the distributional hypothesis, topical similarity can be read off directly from the surface distribution of words in a reasonably large corpus. But the *conceptual* similarities languages encode are more general and abstract. Such conceptual similarities often cannot be read off from the surface distribution but are encoded at a "deep" level of representation and manifest themselves through general rules of the grammar. One particularly challenging example of such a conceptual similarity is the so-called unaccusative hypothesis about intransitive verbs as in (2).

(2) a. Maria lachte.
'Maria laughed.'
b. Maria stolperte.
'Maria stumbled.'

With respect to the sentences in (2), Perlmutter (1978) argued that although verbs like *laugh* and *stumble* look the same on the surface they belong to two different classes of intransitive verbs. The single argument of unergative verbs (like *laugh* in (2-a)) behaves like the grammatical subject of a transitive verb. The single argument of unaccusative verbs (like *stumble* in (2-b)) behaves like the grammatical object of a transitive verb. The unaccusative hypothesis suggests that languages determine intransitive verbs to be similar to each other with respect to whether the single argument is a grammatical subject or object. But the way in which intransitive verbs are judged similar according to the unaccusative hypothesis is independent of the topical similarity of intransitive verbs. The verbs *sleep*, *laugh* and *work* are all unergative but not topically related. Against this background, the goal of the paper is to widen the view on both word embeddings and the lexical representation of the unaccusative hypothesis by pursuing the following research question: Are general concepts like the unaccusativity of intransitive verbs reflected in the embeddings of intransitive verbs?

2.2 Theoretical correlates of unaccusativity

Syntactically, the unaccusativity hypothesis is represented in languages like German by e.g. auxiliary selection in the present perfect (Wunderlich, 1985; Grewendorf, 1989). Unergative verbs (3-a) select HAVE in the present perfect, unaccusative verbs (3-b) select BE.

(3) a. Maria hat gelacht.
Maria HAVE laugh
'Maria has laughed.'
b. Maria ist gestolpert.
Maria BE stumble
'Maria has stumbled.'

Semantically, unaccusativity is determined by an intuition that Dowty (1991) characterizes as follows: "intransitive predicates argued to be unaccusative on syntactic grounds usually turned out to entail relatively patient-like meanings for their arguments [...], while those argued to be syntactically unergative were usually agentive in meaning." Levin and Rappaport Hovav (1995) define this semantic intuition more precisely and argue that unergative verbs describe internally caused events in which "inherent properties of the single argument like will, volition, emotion or physical characteristics are 'responsible' for bringing about the eventuality" (Levin and Rappaport Hovav, 1995, p. 91) that the verb describes. Unaccusative verbs describe externally caused events for which an agent, an instrument, a natural force or a circumstance has "immediate control over bringing about the eventuality described by the verb" (Levin and Rappaport Hovav, 1995, p. 92).

2.3 Previous Work

Our work is inspired by Pross et al. (2017), who argue that the embeddings of a certain class of *over*-prefixed verbs do not represent the expected topical meaning. For example, when the embedding of the verb *to overrun* (as in "the horse overran the girl") is rendered interpretable by its nearest neighbors in the embedding space as in (4), the nearest neighbors do not reflect any meaning aspect of the base verb *to run* nor do they reflect one of the literal meanings of *over*.

(4) *5 nearest neighbours of to overrun*
invade.V pillage.V horde.N incursion.N destroy.V

Pross et al. (2017) hypothesize that one way to think about the nearest neighbor characterization in (4) is that the embedding does not reflect topical similarities but rather prototypical properties of the patient argument that is only licensed by the *over*-prefix, given that (following Levin and Rappaport Hovav (1995)) one way to think of a patient is as being subjected to some external power. Basically, the goal of the present paper is to identify similar cues of agency and patiency in word embeddings, but with the example of unergative and unaccusative verbs.

3 Experimental Setup

The basis for the results reported in the present paper are German and English word embeddings learned with the skip-gram algorithm and negative sampling (Word2Vec, 300 dimensions, parameter settings as suggested in Mikolov et al. (2013b)). The German embeddings were learned from SdeWac, a 0.88 billion word corpus of parsable German web data (Faaß and Eckart, 2013). Sentences were filtered to consist only of content words (i.e. verbs, adjectives and nouns). SdeWac was parsed with the syntactic and semantic dependency parser described in Björkelund et al. (2010) and I extracted verbs that the parser saw more than 90 percent in an intransitive construction together with the semantic role label of the single argument (grammatical subject or grammatical object). I manually corrected the semantic role labels, using auxiliary selection in the present perfect (see (3)) as a diagnostics. In a further step of cleaning, I removed two classes of intransitive verbs that have been argued to involve an unaccusativity mismatch (Zaenen, 1988) and thus are not unambiguously unergative or unaccusative, so-called verbs of emission and particle verbs of directed movement. In total, I ended up with a vocabulary of 972 unergative and 840 unaccusative German verb embeddings. The English word embeddings were learned from ukWac, a 1.3 billion word corpus of English web data (Ferraresi et al., 2008) (same algorithm and parameter settings as for German). Sentences were filtered to consist only of content words. Since English doesn't have reliable markers of unaccusativity (such as auxiliary selection) determining whether an English intransitive verb is unergative or unaccusative is more involved. I thus relied on existing lexical resources and used a subset of the unambiguously internally and externally caused verbs listed in the appendix of Levin and Rappaport Hovav (1995). As examples of externally caused verbs, I chose the classes of "alternating change of state" verbs and "cooking verbs" (251 verbs, class labels according to Levin (1993)). For internally caused verbs, I chose "run"-verbs, verbs that partake in the unspecified object alternation and verbs that alternate with a cognate object construction (275 verbs). I used the classification of Levin and Rappaport Hovav (1995) to label intransitive verbs as subcategorizing either a grammatical subject or a grammatical object. In the following, I refer to the German and English verb embeddings described in this section as "baseline embeddings". I retrained the baseline embeddings for the intransitive verbs using the semantic role of the single argument of the intransitive verb as labels. Retraining took place in a simple neural network architecture consisting of an embedding layer (of size 1812×300 for German and 526×300 for English) fully connected to a single output neuron with a sigmoid activation function. I chose the sigmoid activation to yield a continuous probability distribution over the binary training labels and binary crossentropy as loss function. To make sure the embeddings memorize the distinction between unergative and unaccusative verbs I intentionally overfit the embeddings to 1.0 accuracy on the training data, which I achieved after 20 epochs of training for German and 25 epochs for English.

4 Quantitative interpretation of results

Retraining made the baseline embeddings linearly separable with respect to unaccusativity, an effect which is clearly depicted in figures 1 and 2, where the baseline and retrained German verb embeddings are projected onto two dimensions with Principal Component Analysis (red=unergative, green=unaccusative). The linear separability achieved through retraining is reflected in an increase of the F1 score of simple baseline classifiers reported in table 1. Table 1 shows a constant increase of F1 score across different classifiers and languages, with the increase being downright spectacular for German. Interestingly, the adjustments to embeddings weights through retraining were so tiny that the embeddings retained their original position in the baseline embedding space, i.e. apart from some slight reordering the nearest neighbors of the retrained embeddings stayed the same. In turn, this suggests that injection of linguistic knowledge into word embeddings with an architecture as ours is a cheap but effective way to achieve an improvement of performance in downstream tasks like semantic role labeling, an observation which I leave to further research.

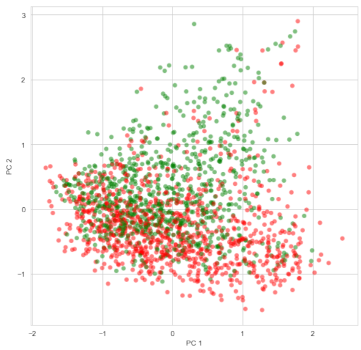

Fig. 1. PCA baseline embeddings

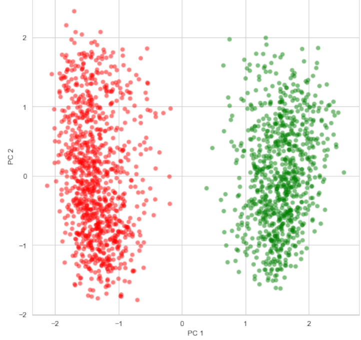

Fig. 2. PCA retrained embeddings

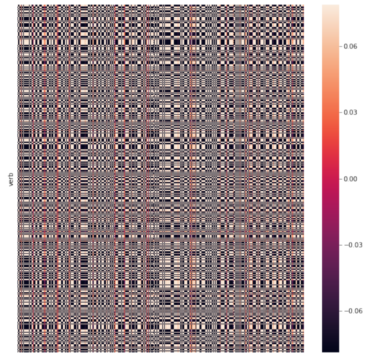

Fig. 3. Heatmap of the difference between baseline and retrained embeddings. x=dimensions, y=vocabulary

	LR	linear SVM	SGD
baseline-de	0.69	0.68	0.70
retrained-de	1.0	1.0	1.0
baseline-en	0.90	0.89	0.87
retrained-en	0.98	0.97	0.95

Table 1: F1-score (10-fold cross-validation) of binary classification with logistic regression (LR), linear support vector machine (SVM), and stochastic gradient descent (SGD) for German (de) and English (en). Embeddings and labels as described in section 3.

5 Qualitative interpretation of results

Given the considerable quantitative effect of retraining on the ability of intransitive verb embeddings to distinguish between unergativity and unaccusativity, I was wondering whether this effect has an explanation that is related to the characterization of the unaccusative hypothesis in theoretical-conceptual semantics. To investigate this question, I isolated the "surplus" that the retraining made to the embeddings by subtracting for each of the embeddings of our German and English vocabulary the weights of the baseline embedding from the weights of the retrained embedding. In a manner of speaking, I computed through subtraction "unergative vectors" and "unaccusative vectors" for German and English which represent the relevant semantic information that is responsible for the linear separability reported in table 1. To foster intuitions about the regularities memorized through retraining of the baseline embeddings, I visualize the difference vectors (which I assume to represent unergativity and unaccusativity in the embedding space) with the heatmap in figure 3. Visual inspection of figure 3 clearly shows that there are systematic patterns in the adjustment of weights through retraining and thus that the effect reflected in the increase of F1 score reported in table 1 does not come about just by chance. To make the pattern visualized in figure 3 interpretable, I approximated the difference vectors resulting from subtraction of baseline from retrained embeddings with their nearest neighbors in the embedding space (similar to the approximation of the representation of *to eat* in the introductory example (1)). Interestingly, the pattern appearing in figure 3 crystallizes in a quite restricted set of shared nearest neighbors of the individual difference vectors that characterize the general unergative and unaccusative difference vectors, respectively. The unergative difference vectors are characterized by a total of 16/14 German/English shared neighbors and the unaccusative difference vectors by 15/24 shared neighbors. This "uniformity" of the interpretation of the difference vectors of the two classes of intransitive verbs allows to abstract away from specific verbs and to consider the sets of shared nearest neighbors as general characterizations of how the retrained embeddings memorize unergativity and unaccusativity. The punchline of our experiment, then appears when (5)/(6) and (7)/(8) are considered, where I list the five most informative shared nearest neighbors of the unergative and unaccusative vectors for German and English.

(5) *5 nearest neighbours of the German unergative vector*
Kulturmanager.N (cultural manager) Diätassistent.N (dietitian) Informatikkaufmann.N (information technology officer) Prüferinnen.N (examiners) Beköstigung.N (feeding)

(6) *5 nearest neighbours of the German unac-*

cusative vector
aushärten.V (to harden) ionisieren.V (to ionize) Bremsvorgang.N (braking process) Ladungstrennung.N (charge separation) Spaltprodukt.N (fission product)

(7) *5 nearest neighbors of the English unergative vector*
beginner.N ceilidh.N sewing.N lug.V crafty.A

(8) *5 nearest neighbors of the English unaccusative vector*
refract.V purify.V eruption.N redness.N irritant.A

The intriguing observation about the approximations in (5)-(8) is that they represent the unergative/unaccusative distinction with those morphological and semantic cues that have been argued to be relevant from a theoretical point of view. The most prominent feature of unergativity are *-er* nominals. Although the correlation of unergativity and *-er* nominalization is no longer considered a reliable diagnostics, there is a general tendency that unergative (9-a) but not unaccusative (9-b) verbs license *-er* nominalizations (Wunderlich (1985) for German, Levin and Rappaport (1988) for English).

(9) a. Tänzer, Arbeiter, Träumer
dancer, worker, dreamer

b. *Faller, *Einschläfer, *Ankommer
*faller, *asleeper, *arriver

Another prominent cue for unergativity are internally controlled activity descriptions like *sewing* or *ceilidh* or *Beköstigung* ('feeding') and adjectives that describe properties of intentional Agents like *crafty*. In contrast, the nearest neighbors of the unaccusative vector characterize "natural" processes that are non-agentive, uncontrolled and externally caused like *aushärten* ('to harden'), *to refract* or *eruption*, results of such processes like *redness* or *Spaltprodukt* ('fission product'), or dispositional properties of objects like *irritant* that require an external stimulation to manifest themselves.

In sum, an informed linguist is able to interpret the nearest neighbors of the unergative and unaccusative vectors as characterizations of the distinction between agent- and patient-like meanings by detecting word formation patterns connected to unergativity, such as *-er* nominals, and shared lexical entailments of proto-agent and proto-patient properties in the sense of (Dowty, 1991). Since these observations can be obtained independently for both English and German intransitive verbs, this suggests that the representations of the semantic correlates of the unaccusative hypothesis by the approximated embeddings of intransitive verbs are not random outliers but rather point towards a systematic effect of our retraining of intransitive word embeddings. One explanation for this systematic effect may be that retraining of the embeddings is a method for strengthening those latent dimensions of the embedding space that involve the same recurrent meanings relative to the retraining objective.

6 Conclusion and Outlook

I showed in a proof-of-concept manner how neural word embeddings can be understood to encode the general conceptual similarity underlying the two classes of intransitive verbs predicted by the unaccusative hypothesis. Word embeddings may thus provide a fresh perspective on the conceptual foundations of verb meanings. The proof-of-concept nature of the present paper comes with a number of limitations that should be addressed by future research. First, I investigated word embeddings learned with a relatively simple neural network architecture but did not consider more recent advances in machine learning, where word embeddings are learned with complex deep bidirectional neural network architectures (Peters et al., 2018; Devlin et al., 2018). I believe an understanding of how more complex neural network architectures learn and represent abstract word meanings is only possible on the basis of an understanding of the simple SGNS algorithm I consider in the present paper. Second, I chose intransitive verbs as my subject of study because the relatively clear-cut semantic dichotomy of unergative and unaccusative verbs provides an accessible starting point with respect to the search for "cues" of recurrent meaning elements in approximated word embeddings. Whether and how the methodology employed in the present paper can also be applied to cases where no similarly clear-cut intuitions about possible cues are available (e.g. mass and count nouns) is a question I leave to future research.

Acknowledgements

DFG grant PR1860/1-1 is gratefully acknowledged. I would like to thank Sebastian Padó for discussion and three anonymous reviewers for helpful comments on an earlier version of the paper.

References

Nicholas Asher, Tim Van de Cruys, Antoine Bride, and Márta Abrusán. 2016. Integrating type theory and distributional semantics: A case study on adjective—noun compositions. *Computational Linguistics*, 42(4):703 – 725.

Anders Björkelund, Bernd Bohnet, Love Hafdell, and Pierre Nugues. 2010. A high-performance syntactic and semantic dependency parser. In *Coling 2010: Demonstrations*, pages 33–36. Coling 2010 Organizing Committee.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.

David R. Dowty. 1991. Thematic proto-roles and argument selection. *Language*, 67(3):547 – 619.

Gertrud Faaß and Kerstin Eckart. 2013. SdeWaC - a corpus of parsable sentences from the web. In I. Gurevych, C. Biemann, and T. Zesch, editors, *Proceedings of the International Conference of the German Society for Computational Linguistics and Language Technology*.

Adriano Ferraresi, Eros Zanchetta, Marco Baroni, and Silvia Bernardini. 2008. Introducing and evaluating ukwac, a very large web-derived corpus of english.

Günther Grewendorf. 1989. *Ergativity in German*. Foris, Dordrecht.

Mohamed Guerssel, Kenneth Hale, Margaret Laughren, Beth Levin, and Josie White Eagle. 1985. A cross-linguistic study of transitivity alternations. In *CLS 21: Papers from the Parasession on Causatives and Agentivity.*, volume 2, pages 48–63. Chicago Linguistic Society.

Alessandro Lenci. 2014. Carving verb classes from corpora. In Raffaele Simone and Francesca Masini, editors, *Word Classes: Nature, typology and representations*, pages 17 – 36. John Benjamins.

Beth Levin and Malka Rappaport Hovav. 1995. *Unaccusativity at the syntax-semantics interface*. MIT Press.

Beth Levin and Malka Rappaport. 1988. Non-event -er nominals: a probe into argument structure. *Linguistics*, 26:1067–1083.

Beth Levin. 1993. *English verb classes and alternations: a preliminary investigation*. University of Chicago Press, Chicago.

Omer Levy and Yoav Goldberg. 2014a. Linguistic regularities in sparse and explicit word representations. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 171–180. Association for Computational Linguistics.

Omer Levy and Yoav Goldberg. 2014b. Neural word embedding as implicit matrix factorization. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2177–2185. Curran Associates, Inc.

Louise McNally and Gemma Boleda, 2017. *Conceptual vs. Referential Affordance in Concept Composition*, pages 245–267. Springer International Publishing, Cham.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Proceedings of NIPS*, pages 3111–3119.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

David M. Perlmutter. 1978. Impersonal passives and the unaccusative hypothesis. In *Proceedings of the 4th Annual Meeting of the Berkeley Linguistics Society*, pages 157–190.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *CoRR*, abs/1802.05365.

Steven Pinker. 2013. *Learnability and Cognition: The Acquisition of Argument Structure. New edition.* MIT Press.

Tillmann Pross, Antje Roßdeutscher, Sebastian Padó, Gabriella Lapesa, and Max Kisselew. 2017. Integrating lexical-conceptual and distributional semantics: a case report. In *Proceedings of the Amsterdam Colloquium 2017*, pages 75–85.

Peter D. Turney and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, 37:141–188.

Frederike Van der Leek. 1996. The english conative construction: A compositional account. In *CLS 32: Papers from the Main Session*, volume 32, pages 363–378. Chicago Linguistic Society.

Dieter Wunderlich. 1985. Über die Argumente des Verbs. *Linguistische Berichte*, 97:183–227.

Annie Zaenen. 1988. Unaccusatives in dutch and the syntax-semantics interface. CSLI Report 123, Center for the Study of Language and Information.