

ARTICLE

Imparting interpretability to word embeddings while preserving semantic structure

Lütfi Kerem Şenel¹, İhsan Utlu^{2,3}, Furkan Şahinuç^{2,3}, Haldun M. Ozaktas² and Aykut Koç^{2,4*}

¹Center for Information and Language Processing (CIS), Ludwig Maximilian University (LMU), Munich, Germany, ²Electrical and Electronics Engineering Department, Bilkent University, Ankara, Turkey, ³ASELSAN Research Center, Ankara, Turkey and ⁴National Magnetic Resonance Research Center (UMRAM), Bilkent University, Ankara, Turkey

*Corresponding author. Email: aykut.koc@bilkent.edu.tr

(Received 15 May 2019; revised 3 May 2020; accepted 4 May 2020; first published online 9 June 2020)

Abstract

As a ubiquitous method in natural language processing, word embeddings are extensively employed to map semantic properties of words into a dense vector representation. They capture semantic and syntactic relations among words, but the vectors corresponding to the words are only meaningful relative to each other. Neither the vector nor its dimensions have any absolute, interpretable meaning. We introduce an additive modification to the objective function of the embedding learning algorithm that encourages the embedding vectors of words that are semantically related to a predefined concept to take larger values along a specified dimension, while leaving the original semantic learning mechanism mostly unaffected. In other words, we align words that are already determined to be related, along predefined concepts. Therefore, we impart interpretability to the word embedding by assigning meaning to its vector dimensions. The predefined concepts are derived from an external lexical resource, which in this paper is chosen as Roget's Thesaurus. We observe that alignment along the chosen concepts is not limited to words in the thesaurus and extends to other related words as well. We quantify the extent of interpretability and assignment of meaning from our experimental results. Manual human evaluation results have also been presented to further verify that the proposed method increases interpretability. We also demonstrate the preservation of semantic coherence of the resulting vector space using word-analogy/word-similarity tests and a downstream task. These tests show that the interpretability-imparted word embeddings that are obtained by the proposed framework do not sacrifice performances in common benchmark tests.

Keywords: Word embeddings; Interpretability; Computational semantics

1. Introduction

Distributed word representations, commonly referred to as *word embeddings* (Mikolov *et al.* 2013a, 2013c; Pennington, Socher, and Manning 2014; Bojanowski *et al.* 2017), serve as elementary building blocks in the course of algorithm design for an expanding range of applications in natural language processing (NLP), including named entity recognition (Turian, Ratinov, and Bengio 2010; Sienčnik 2015), parsing (Chen and Manning 2014), sentiment analysis (Socher *et al.* 2011; Yu *et al.* 2017), and word sense disambiguation (Iacobacci, Pilehvar, and Navigli 2016). Although the empirical utility of word embeddings as an unsupervised method for capturing the semantic or syntactic features of a certain word as it is used in a given lexical resource is well established (Vine *et al.* 2015; Joshi *et al.* 2016; Goldberg and Hirst 2017), an understanding of what these features mean remains an open problem (Levy and Goldberg 2014; Chen *et al.* 2016) and as such word embeddings mostly remain a black box. It is desirable to be able to develop insight into this

black box and be able to interpret what it means, while retaining the utility of word embeddings as semantically rich intermediate representations. Other than the intrinsic value of this insight, this would not only allow us to explain and understand how algorithms work (Goodman and Flaxman 2017) but also set a ground that would facilitate the design of new algorithms in a more deliberate way.

Recent approaches to generating word embeddings (Mikolov *et al.* 2013c; Pennington *et al.* 2014) are rooted linguistically in the field of *distributed semantics* (Harris 1954), where words are taken to assume meaning mainly by their degree of interaction (or lack thereof) with other words in the lexicon (Firth 1957a; 1957b). Under this paradigm, dense, continuous vector representations are learned in an unsupervised manner from a large corpus, using the word cooccurrence statistics directly or indirectly, and such an approach is shown to result in vector representations that mathematically capture various semantic and syntactic relations between words (Mikolov *et al.* 2013c; Pennington *et al.* 2014; Bojanowski *et al.* 2017). However, the dense nature of the learned embeddings obfuscates the distinct concepts encoded in the different dimensions, which renders the resulting vectors virtually uninterpretable. The learned embeddings make sense only in relation to each other, and their specific dimensions do not carry explicit information that can be interpreted. However, being able to interpret a word embedding would illuminate the semantic concepts implicitly represented along the various dimensions of the embedding and reveal its hidden semantic structures.

In the literature, researchers tackled the interpretability problem of the word embeddings using different approaches. Several researchers (Murphy, Talukdar, and Mitchell 2012; Luo *et al.* 2015; Fyshe *et al.* 2016) proposed algorithms based on nonnegative matrix factorization (NMF) applied to cooccurrence variant matrices. Other researchers suggested to obtain interpretable word vectors from existing uninterpretable word vectors by applying sparse coding (Faruqui *et al.* 2015a; Arora *et al.* 2018), by training a sparse autoencoder to transform the embedding space (Subramanian *et al.* 2018), by rotating the original embeddings (Zobnin 2017; Park, Bak, and Oh 2017), or by applying transformations based on external semantic datasets (Senel *et al.* 2018a).

Although the above-mentioned approaches provide better interpretability that is measured using a particular method such as word intrusion test, usually the improved interpretability comes with a cost of performance in the benchmark tests such as word similarity or word analogy. One possible explanation for this performance decrease is that the proposed transformations from the original embedding space distort the underlying semantic structure constructed by the original embedding algorithm. Therefore, it can be claimed that a method that learns dense and interpretable word embeddings without inflicting any damage to the underlying semantic learning mechanism is the key to achieve both high-performing and interpretable word embeddings.

Especially after the introduction of the *word2vec* algorithm by Mikolov *et al.* (2013a, 2013c), there has been a growing interest in algorithms that generate improved word representations under some performance metrics. Significant effort is spent on appropriately modifying the objective functions of the algorithms in order to incorporate knowledge from external resources, with the purpose of increasing the performance of the resulting word representations (Miller 1995; Yu and Dredze 2014; Xu *et al.* 2014; Liu *et al.* 2015; Jauhar, Dyer, and Hovy 2015; Johansson and Nieto Piña 2015; Bollegala *et al.* 2016). Significant effort is also spent on developing retrofitting objectives for the same purpose, independent of the original objectives of the embedding model, to fine-tune the embeddings without joint optimization (Faruqui *et al.* 2015b; Mrkšić *et al.* 2016, 2017). Inspired by the line of work reported in these studies, we propose to use modified objective functions for a different purpose: learning more interpretable dense word embeddings. By doing this, we aim to incorporate semantic information from an external lexical resource into the word embedding so that the embedding dimensions are *aligned* along predefined concepts. This alignment is achieved by introducing a modification to the embedding learning process. In our proposed method, which is built on top of the GloVe algorithm (Pennington *et al.* 2014), the cost function for any one of the words of concept word groups is modified by the introduction of an additive term to the cost function. Each embedding vector dimension is first associated with a

concept. For a word belonging to any one of the word groups representing these concepts, the modified cost term favors an increase in the value of this word's embedding vector dimension corresponding to the concept that the particular word belongs to. For words that do not belong to any one of the word groups, the cost term is left untouched. Specifically, *Roget's Thesaurus* (Roget 1911, 2008) is used to derive the concepts and concept word groups to be used as the external lexical resource for our proposed method. We quantitatively demonstrate the increase in interpretability using the measure given in Senel *et al.* (2018a, 2018b) as well as demonstrating qualitative results. Furthermore, manual human evaluations based on the "word intrusion" test given in Chang *et al.* (2009) have been carried out for verification. We also show that the semantic structure of the original embedding has not been harmed in the process since there is no performance loss with standard word-similarity or word-analogy tests and with a downstream sentiment analysis task.

The paper is organized as follows. In Section 2, we discuss previous studies related to our work under two main categories: interpretability of word embeddings and joint-learning frameworks where the objective function is modified. In Section 3, we present the problem framework and provide the formulation within the GloVe (Pennington *et al.* 2014) algorithm setting. In Section 4 where our approach is proposed, we motivate and develop a modification to the original objective function with the aim of increasing representation interpretability. In Section 2, experimental results are provided, and the proposed method is quantitatively and qualitatively evaluated. Additionally, in Section 2, results demonstrating the extent to which the original semantic structure of the embedding space is affected are presented using word-analogy/word-similarity tests and a downstream evaluation task. Analysis of several parameters of our proposed approach is also presented in Section 2. We conclude the paper in Section 6.

2. Related work

Methodologically, our work is related to prior studies that aim to obtain "improved" word embeddings using external lexical resources, under some performance metrics. Previous work in this area can be divided into two main categories: works that (i) *modify* the word embedding learning algorithm to incorporate lexical information and (ii) operate on pretrained embeddings with a *post-processing* step.

Among works that follow the first approach, Yu and Dredze (2014) extend the Skip-Gram model by incorporating the word-similarity relations extracted from the Paraphrase Database (PPDB) and WordNet (Miller 1995) into the Skip-Gram predictive model as an additional cost term. Xu *et al.* (2014) extend the Continuous Bag of Words model by considering two types of semantic information, termed *relational* and *categorical*, to be incorporated into the embeddings during training. For the former type of semantic information, the authors propose the learning of explicit vectors for the different relations extracted from a semantic lexicon such that the word pairs that satisfy the same relation are distributed more homogeneously. For the latter, the authors modify the learning objective such that some weighted average distance is minimized for words under the same semantic category. Liu *et al.* (2015) represent the synonymy and hypernymy-hyponymy relations in terms of *inequality constraints*, where the pairwise similarity rankings over word triplets are forced to follow an order extracted from a lexical resource. Following their extraction from WordNet, the authors impose these constraints in the form of an additive cost term to the Skip-Gram formulation. Finally, Bollegala *et al.* (2016) build on top of the GloVe algorithm by introducing a regularization term to the objective function that encourages the vector representations of similar words as dictated by WordNet to be similar as well.

Turning our attention to the post-processing approach for enriching word embeddings with external lexical knowledge, Faruqui *et al.* (2015b) have introduced the *retrofitting* algorithm that acts on pretrained embeddings such as Skip-Gram or GloVe. The authors propose an objective function that aims to balance out the semantic information captured in the pretrained embeddings with the constraints derived from lexical resources such as WordNet, PPDB, and FrameNet. One

of the models proposed in Jauhar *et al.* (2015) extends the retrofitting approach to incorporate the word sense information from WordNet. Similarly, Johansson and Nieto Piña (2015) create multisense embeddings by gathering the word sense information from a lexical resource and learning to decompose the pretrained embeddings into a convex combination of sense embeddings. Mrkšić *et al.* (2016) focus on improving word embeddings for capturing word similarity, as opposed to mere relatedness. To this end, they introduce the *counter-fitting* technique which acts on the input word vectors such that synonymous words are attracted to one another whereas antonymous words are repelled, where the synonymy–antonymy relations are extracted from a lexical resource. The *ATTRACT-REPEL* algorithm proposed by Mrkšić *et al.* (2017) improves on counter-fitting by a formulation which imparts the word vectors with external lexical information in *mini-batches*. More recently, several global specialization methods have been proposed in order to generalize the specialization to the vectors of words that are not present in external lexical resources (Glavaš and Vulić 2018; Ponti *et al.* 2018).

Most of the studies discussed above (Xu *et al.* 2014; Liu *et al.* 2015; Jauhar *et al.* 2015; Faruqui *et al.* 2015b; Bollegala *et al.* 2016; Mrkšić *et al.* 2016, 2017) report performance improvements in benchmark tests such as word similarity or word analogy, while Miller (1995) uses a different analysis method (mean reciprocal rank). In sum, the literature is rich with studies aiming to obtain word embeddings that perform better under specific performance metrics. However, less attention has been directed to the issue of interpretability of the word embeddings. In the literature, the problem of interpretability has been tackled using different approaches. In terms of methodology, these approaches can be grouped under two categories: direct approaches that do not require a pretrained embedding space and post-processing approaches that operate on a pretrained embedding space (the latter being more often deployed). Among the approaches that fall into the direct approach category, Murphy *et al.* (2012) proposed NMF for learning sparse, interpretable word vectors from cooccurrence variant matrices where the resulting vector space is called nonnegative sparse embeddings (NNSE). However, since NMF methods require maintaining a global matrix for learning, they suffer from memory and scale issue. This problem has been addressed in Luo *et al.* (2015) where an online method of learning interpretable word embeddings from corpora is proposed. The authors proposed a modified version of Skip-Gram model (Mikolov *et al.* 2013c), called Online Interpretable Word Embeddings – Improved Projected Gradient (OIWE-IPG), where the updates are forced to be nonnegative during the training of the algorithm so that the resulting embeddings are also nonnegative and more interpretable. In Fyshe *et al.* (2016), a generalized version of NNSE method, called Joint Non-Negative Sparse Embeddings (JNNSE), is proposed to incorporate constraints based on external knowledge. In their study, brain activity-based word-similarity information is taken as external knowledge and combined with text-based similarity information in order to improve interpretability.

Relatively more research effort has been directed to improve interpretability by post-processing the existing pretrained word embeddings. These approaches aim to learn a transformation to map the original embedding space to a new more interpretable one. Arora *et al.* (2018) and Faruqui *et al.* (2015a) use sparse coding on conventional dense word embeddings in order to obtain sparse, higher dimensional and more interpretable vector spaces. Motivated by the success of neural architectures, deploying a sparse autoencoder for pretrained dense word embeddings is proposed in Subramanian *et al.* (2018) in order to improve interpretability. Instead of using sparse transformations as in the above-mentioned studies, several other studies focused on learning orthogonal transformations that will preserve the internal semantic information and high performance of the original dense embedding space. In Zobnin (2017), interpretability is taken as the tightness of clustering along individual embedding dimensions, and orthogonal transformations are utilized to improve it. However, Zobnin (2017) has also shown that based on this definition for interpretability, total interpretability of an embedding is kept constant under any orthogonal transformation, and it can only be redistributed across the dimensions. Park *et al.* (2017) investigated rotation algorithms based on exploratory factor analysis in order to improve interpretability while preserving the performance. Dufter and Schutze (2019) proposed a method

to learn an orthogonal transformation matrix that will align a given linguistic signal in the form of a word group to an embedding dimension providing an interpretable subspace. In their work, they demonstrate their method for a one-dimensional subspace. However, it is not clear how well the proposed method can generalize for a larger dimensional subspace (ideally the entire embedding space). In Senel *et al.* (2018a), a transformation based on Bhattacharya distance and semantic category-based approach and category dataset (SEMCAT) categories is proposed to obtain an interpretable embedding space. In that study, also an automated metric was proposed to quantitatively measure the degree of interpretability already present in the embedding vector spaces. Taking a different approach, Herbelot and Vecchi (2015) proposed a method to map dense word embeddings to a *model-theoretic space* where the dimensions correspond to real-world features elicited from human participants.

Following a separate line of work based on the research on topic modeling domain, Panigrahi, Simhadri, and Bhattacharyya (2019) proposed a Latent Dirichlet Allocation (LDA)-based generative model to extract different senses for words from a corpus. They also proposed a method to learn sparse interpretable word embedding, called Word2Sense, based on the obtained sense distributions. Several other studies also focused on associations between word embedding models and the topic modeling methods (Liu *et al.* 2015; Das, Zaheer, and Dyer 2015; Moody 2016; Shi *et al.* 2017). They make use of LDA-based models to obtain word topics to be integrated to word embeddings. It should be noted that this literature is also related in the sense that topic modeling may be used to improve the procedures for extracting the word groups representing the concepts assigned to the embedding dimensions.

Most of the interpretability-related previous work mentioned above, except Fyshe *et al.* (2016), Senel *et al.* (2018a), and our proposed method, do not need external resources, utilization of which has both advantages and disadvantages. Methods that do not use external information require fewer resources but they also lack the aid of information extracted from these resources.

3. Problem description

For the task of unsupervised word embedding extraction, we operate on a discrete collection of lexical units (words) $u_i \in V$ that is part of an input corpus $C = \{u_i\}_{i \geq 1}$, with number of tokens $|C|$, sourced from a vocabulary $V = \{w_1, \dots, w_{|V|}\}$ of size $|V|$.^a In the setting of distributional semantics, the objective of a word embedding algorithm is to maximize some aggregate utility over the entire corpus so that some measure of “closeness” is maximized for pairs of vector representations (w_i, w_j) for words which, on the average, appear in proximity to one another. In the GloVe algorithm (Pennington *et al.* 2014), which we base our proposed method upon, the following objective function is considered:

$$J = \sum_{i,j=1}^{|V|} f(X_{ij}) \left(w_i^T \tilde{w}_j + b_i + \tilde{b}_j - \log X_{ij} \right)^2 \tag{1}$$

In Equation (1), $w_i \in \mathbb{R}^D$ and $\tilde{w}_j \in \mathbb{R}^D$ stand for word and context vector representations, respectively, for words w_i and w_j , while X_{ij} represents the (possibly weighted) cooccurrence count for the word pair (w_i, w_j) . Intuitively, Equation (1) represents the requirement that if some word w_i occurs often enough in the context (or vicinity) of another word w_j , then the corresponding word representations should have a large enough inner product in keeping with their large X_{ij} value, up to some bias terms b_i, \tilde{b}_j ; and vice versa. $f(\cdot)$ in Equation (1) is used as a discounting factor that prohibits rare cooccurrences from disproportionately influencing the resulting embeddings.

^aWe represent vectors (matrices) by bold lower (upper) case letters. For a vector \mathbf{a} (a matrix \mathbf{A}), \mathbf{a}^T (\mathbf{A}^T) is the transpose. $\|\mathbf{a}\|$ stands for the Euclidean norm. For a set S , $|S|$ denotes the cardinality. $\mathbb{1}_{x \in S}$ is the indicator variable for the inclusion $x \in S$, evaluating to 1 if satisfied, 0 otherwise.

The objective in Equation (1) is minimized using stochastic gradient descent by iterating over the matrix of co-occurrence records $[X_{ij}]$. In the GloVe algorithm, for a given word w_i , the final word representation is taken to be the average of the two intermediate vector representations obtained from Equation (1), that is, $(\mathbf{w}_i + \tilde{\mathbf{w}}_i)/2$. In the next section, we detail the enhancements made to Equation (1) for the purposes of enhanced interpretability, using the aforementioned framework as our basis.

4. Imparting interpretability

Our approach falls into a joint-learning framework where the distributional information extracted from the corpus is allowed to fuse with the external lexicon-based information. An external resource in which words are primarily grouped together based on human judgments and in which the entire semantic space is represented as much as possible is needed. We have chosen to use Roget's Thesaurus not only its being one of the earliest examples of its kind but also its status of being continuously updated for modern words. *Word groups* extracted from Roget's Thesaurus are directly mapped to individual dimensions of word embeddings. Specifically, the vector representations of words that belong to a particular group are encouraged to have deliberately increased values in a particular dimension that corresponds to the word group under consideration. This can be achieved by modifying the objective function of the embedding algorithm to partially influence vector representation distributions across their dimensions over an input vocabulary. To do this, we propose the following modification to the GloVe objective given in Equation (1):

$$J = \sum_{i,j=1}^{|V|} f(X_{ij}) \left[\left(\mathbf{w}_i^T \tilde{\mathbf{w}}_j + b_i + \tilde{b}_j - \log X_{ij} \right)^2 + k \left(\sum_{l=1}^D \mathbb{1}_{i \in F_l} g(\mathbf{w}_{i,l}) + \sum_{l=1}^D \mathbb{1}_{j \in F_l} g(\tilde{\mathbf{w}}_{j,l}) \right) \right] \quad (2)$$

In Equation (2), F_l denotes the indices for the elements of the l th concept word group which we wish to assign in the vector dimension $l = 1, \dots, D$. The objective in Equation (2) is designed as a mixture of two individual cost terms: the original GloVe cost term along with a second term that encourages embedding vectors of a given concept word group to achieve deliberately increased values along an associated dimension l . The relative weight of the second term is controlled by the parameter k . The simultaneous minimization of both objectives ensures that words that are similar to, but not included in, one of these concept word groups are also “nudged” toward the associated dimension l . The trained word vectors are thus encouraged to form a distribution where the individual vector dimensions align with certain semantic concepts represented by a collection of concept word groups, one assigned to each vector dimension. To facilitate this behavior, Equation (2) introduces a monotone decreasing function $g(\cdot)$ defined as:

$$g(x) = \begin{cases} \frac{1}{2} \exp(-2x) & \text{for } x < 0.5 \\ \frac{1}{(4e)x} & \text{otherwise} \end{cases} \quad (3)$$

which serves to increase the total cost incurred if the value of the l th dimension for the two vector representations $\mathbf{w}_{i,l}$ and $\tilde{\mathbf{w}}_{j,l}$ for a concept word w_i with $i \in F_l$ fails to be large enough. Although different definitions for $g(x)$ are possible, we observed after several experiments that this piecewise definition provides decent push for the words in the word groups. $g(x)$ is graphically shown in Figure 1, and we will also present further analysis and experiments regarding the effects of different forms of $g(x)$ in Section 5.6.




Figure 1. Function g in the additional cost term.

The objective Equation (2) is minimized using stochastic gradient descent over the cooccurrence records $\{X_{ij}\}_{i,j=1}^{|V|}$. Intuitively, the terms added to Equation (2) in comparison with Equation (1) introduce the effect of selectively applying a positive step-type input to the original descent updates of Equation (1) for concept words along their respective vector dimensions, which influences the dimension value in the positive direction. The parameter k in Equation (2) allows for the adjustment of the magnitude of this influence as needed.

In the next section, we demonstrate the feasibility of this approach by experiments with an example collection of concept word groups extracted from Roget's Thesaurus. Before moving on, we would like to first comment on the discovery of the “ultimate” categories. Questions like “What are the intrinsic and fundamental building blocks of the entire semantic space?” and “What should be the corresponding categories?” are important questions in linguistics and NLP, and above all, in philosophy. Determining them without human intervention and with unsupervised means is an open problem. Roget's Thesaurus can be seen as a manual attempt for answering this question. The methodology behind it is exhaustive. Based on the premise that there should be some word in a language for any material or immaterial thing known to humans, Roget's Thesaurus conceptualizes all the words within a tree structure with hierarchical categories. So, taking it as starting point to construct the categories and concept word groups is a logical option. One can readily use any other external lexical resource as long as the groupings of words are not arbitrary (which is the case in all thesauruses) or leverage topic modeling methods to form the categories. It is also obvious that some sort of supervision is better in terms of reaching the “ultimate” categories than unsupervised approaches. On the other side, however, unsupervised methods have the advantage on not depending on external resources. This leads to the question of how supervised and unsupervised approaches compare. For that reason, in the next section, we also quantitatively compare our method against supervised and unsupervised projection-based simple baselines for concept word groups extraction.

5. Experiments and results^b

We first identified 300 concepts, one for each dimension of the 300-dimensional vector representation, by employing Roget's Thesaurus. This thesaurus follows a tree structure which starts with a *Root* node that contains all the words and phrases in the thesaurus. The root node is successively split into *Classes* and *Sections*, which are then (optionally) split into *Subsections* of various depths, finally ending in *Categories*, which constitute the smallest unit of word/phrase collections in the structure. The actual words and phrases descend from these *Categories* and make up the leaves

^bAll necessary source codes to reproduce the experiments in this section are available at <https://github.com/koc-lab/imparting-interpretability>.

Algorithm 1 Algorithm for extracting concept word groups

```

word_groups = [ ]
nodes = [ ]
nodes.add(root_node)
while length of nodes > 0 do
  large_nodes = nodes
  for node in large_nodes do
    if node.size() >  $\lambda$  then
      children_nodes = node.get_children()
      if length of children_nodes > 0 then
        | nodes.add(children_nodes)
      else
        | word_groups.add(node.truncate_to_lambda())
      end
    else
      | word_groups.add(node)
    end
  nodes.remove(node)
  end
end

```

of the tree structure. We note that a given word typically appears in multiple categories corresponding to the different senses of the word. We constructed concept word groups from Roget's Thesaurus as follows: We first filtered out the multi-word phrases and the relatively obscure terms from the thesaurus. The obscure terms were identified by checking them against a vocabulary extracted from Wikipedia. We then obtained 300 word groups as the result of a *partitioning* operation applied to the subtree that ends with *categories* as its leaves. The partition boundaries, hence the resulting word groups, can be chosen in many different ways. In our proposed approach, we have chosen to determine this partitioning by traversing this tree structure from the root node in breadth-first order, and by employing a parameter λ for the maximum *size* of a node. Here, the size of a node is defined as the number of unique words that ever-descend from that node. During the traversal, if the size of a given node is less than this threshold, we designate the words that ultimately descend from that node as a concept word group. Otherwise, if the node has children, we discard the node and queue up all its children for further consideration. If this node does not have any children, on the other hand, the node is truncated to λ elements with the highest frequency ranks, and the resulting words are designated as a concept word group. The algorithm for extracting concept word groups from Roget's Thesaurus is also given in pseudo-code form as Algorithm 1. We note that the choice of λ greatly affects the resulting collection of word groups: excessively large values result in few word groups that greatly overlap with one another, while overly small values result in numerous tiny word groups that fail to adequately represent a concept. We experimentally determined that a λ value of 452 results in the most healthy number of relatively large word groups (113 groups with size ≥ 100), while yielding a preferably small overlap among the resulting word groups (with an average overlap size not exceeding three words). A total of 566 word groups were thus obtained. Two hundred fifty-nine smallest word groups (with size < 38) were discarded to bring down the number of word groups to 307. Out of these, 7 groups with the lowest median frequency rank were further discarded, which yields the final 300 concept word groups used in the experiments. We present some of the resulting word groups in Table 1.^c

^cAll the vocabulary lists, concept word groups, and other material necessary to reproduce this procedure are available at <https://www.cambridge.org/core>. Bilkent University Library, on 29 Nov 2021 at 08:10:01, subject to the Cambridge Core terms of use, available at <https://www.cambridge.org/core/terms>. <https://doi.org/10.1017/S15324920000315>

Table 1. Sample concepts and their associated word groups from Roget’s Thesaurus

MANKIND	BUSINESS	SIMPLE QUANTITY	CONDUCT	ARRIVAL
one	living	size	government	land
population	work	way	life	home
people	line	point	game	light
world	place	force	role	airport
state	service	station	race	return
family	role	range	record	come
national	race	standard	process	complete
public	office	rate	business	port
party	act	stage	career	hit
million	case	mass	campaign	meeting
⋮	⋮	⋮	⋮	⋮

Table 2. GloVe parameters

VOCAB_MIN_COUNT	65
ALPHA	0.75
WINDOW_SIZE	15
VECTOR_SIZE	300
X_MAX	75

Using the concept word groups, we have trained the GloVe algorithm with the proposed modification given in Section 4 on a snapshot of English Wikipedia consisting of around 1.1B tokens, with the stop words filtered out. Using the parameters given in Table 2, this resulted in a vocabulary size of 287,847. For the weighting parameter in Equation (2), we used a value of $k = 0.1$ whose effect is analyzed in detail in Section 5.4. The algorithm was trained over 20 iterations. The GloVe algorithm without any modifications was also trained with the same parameters. In addition to the original GloVe algorithm, we compare our proposed method with previous studies that aim to obtain interpretable word vectors. We train the *improved projected gradient* model proposed in Luo *et al.* (2015) to obtain word vectors (called OIWE-IPG) using the same corpus we use to train GloVe and our proposed method. For the Word2Sense method (Panigrahi *et al.* 2019), we use 2250 dimensional pretrained embeddings for comparisons instead of training the algorithm on the same corpus used for the other methods due to very slow training of the model on our hardware.^d Using the methods proposed in Faruqui *et al.* (2015a), Park *et al.* (2017), and Subramanian *et al.* (2018) on our baseline GloVe embeddings, we obtain Sparse Overcomplete Word Vectors (SOV), SParse Interpretable Neural Embeddings (SPINE), and Parsimax (orthogonal) word representations, respectively. We train all the models with the proposed parameters. However, Subramanian *et al.* (2018) show results for a relatively small vocabulary of 15,000 words.

^dPretrained Word2Sense model has advantage over our proposed method and the other alternatives due to being trained on a nearly three times larger corpus (around 3B tokens).




Figure 2. Most frequent 1000 words sorted according to their values in the 32nd dimension of the original GloVe embedding are shown with “●” markers. “○” and “+” markers show the values of the same words for the 32nd dimension of the embedding obtained with the proposed method where the dimension is *aligned* with the concept JUDGMENT. Words with “○” markers are contained in the concept JUDGMENT while words with “+” markers are not contained.

When we trained their model on our baseline GloVe embeddings with a large vocabulary of size 287,847, the resulting vectors performed significantly poor on word-similarity tasks compared to the results presented in their paper. In addition to these alternatives, we also compare our method against two simple projection-based baselines. Specifically, we construct two new embedding spaces by projecting the original GloVe embeddings onto (i) randomly sampled 300 different tokens and (ii) average vectors of the words for the 300 word groups extracted from Roget’s Thesaurus. We evaluate the interpretability of the resulting embeddings qualitatively and quantitatively. We also test the performance of the embeddings on word-similarity and word-analogy tests as well as on a downstream classification task.

In our experiments, vocabulary size is close to 300,000, while only 16,242 unique words of the vocabulary are present in the concept groups. Furthermore, only dimensions that correspond to the concept group of the word will be updated due to the additional cost term. Given that these concept words can belong to multiple concept groups (two on average), only 33,319 parameters are updated. There are 90 million individual parameters present for the 300,000 word vectors of size 300. Of these parameters, only approximately 33,000 are updated by the additional cost term. For the interpretability evaluations, we restrict the vocabulary to the most frequent 50,000 words^e except Figure 2 where we only use most frequent 1000 words for clarity of the plot.

5.1 Qualitative evaluation for interpretability

In Figure 2, we demonstrate the particular way in which the proposed algorithm Equation (2) influences the vector representation distributions. Specifically, we consider, for illustration, the 32nd dimension values for the original GloVe algorithm and our modified version, restricting the plots to the top 1000 words with respect to their frequency ranks for clarity of presentation. In Figure 2, the words in the horizontal axis are sorted in descending order with respect to the

^eSince Word2Sense embeddings have a different vocabulary, we first filter and sort it based on our vocabulary.

values at the 32nd dimension of their word embedding vectors coming from the original GloVe algorithm. The dimension values are denoted with “●” and “○”/“+” markers for the original and the proposed algorithms, respectively. Additionally, the top 50 words that achieve the greatest 32nd dimension values among the considered 1000 words are emphasized with enlarged markers, along with text annotations. In the presented simulation of the proposed algorithm, the 32nd dimension values are encoded with the concept JUDGMENT, which is reflected as an increase in the dimension values for words such as *committee*, *academy*, and *article*. We note that these words (denoted by +) are *not* part of the predetermined word group for the concept JUDGMENT, in contrast to words such as *award*, *review*, and *account* (denoted by ○) which are. This implies that the increase in the corresponding dimension values seen for these words is attributable to the joint effect of the first term in Equation (2) which is inherited from the original GloVe algorithm, in conjunction with the remaining terms in the proposed objective expression Equation (2). This experiment illustrates that the proposed algorithm is able to impart the concept of JUDGMENT on its designated vector dimension above and beyond the supplied list of words belonging to the concept word group for that dimension. It should also be noted that the majority of the words in Figure 2 are denoted by “+,” which means that they are *not* part of the predetermined wordgroups and are semi-supervisedly imparted.

We also present the list of words with the greatest dimension value for the dimensions 11, 13, 16, 31, 36, 39, 41, 43, and 79 in Table 3. These dimensions are aligned/imparted with the concepts that are given in the column headers. In Table 3, the words that are given with regular font denote the words that exist in the corresponding word group obtained from Roget’s Thesaurus (and are thus explicitly forced to achieve increased dimension values), while bolded words denote the words that achieve increased dimension values by virtue of their cooccurrence statistics with the thesaurus-based words (indirectly, without being explicitly forced). This again illustrates that a semantic concept can indeed be coded to a vector dimension provided that a sensible lexical resource is used to guide semantically related words to the desired vector dimension via the proposed objective function in Equation (2). Even the words that do not appear in, but are semantically related to, the word groups that we formed using Roget’s Thesaurus are indirectly affected by the proposed algorithm. They also reflect the associated concepts at their respective dimensions even though the objective functions for their particular vectors are not modified. This point cannot be overemphasized. Although the word groups extracted from Roget’s Thesaurus impose a degree of supervision to the process, the fact that the remaining words in the entire vocabulary are also indirectly affected makes the proposed method a semi-supervised approach that can handle words that are not in these chosen word groups. A qualitative example of this result can be seen in the last column of Table 3. It is interesting to note the appearance of words such as *guerilla*, *insurgency*, *mujahideen*, *Wehrmacht*, and *Luftwaffe* in addition to the more obvious and straightforward *army*, *soldiers*, and *troops*, all of which are not present in the associated word group WARFARE.

Most of the dimensions we investigated exhibit similar behavior to the ones presented in Table 3. Thus generally speaking, we can say that the entries in Table 3 are representative of the great majority. However, we have also specifically looked for dimensions that make less sense and determined a few such dimensions which are relatively less satisfactory. These less satisfactory examples are given in Table 4, where the regular and bolded fonts carry the same meanings as in Table 3. These examples are also interesting in that they shed insight into the limitations posed by polysemy and existence of very rare outlier words.

5.2 Quantitative evaluation for interpretability

One of the main goals of this study is to improve the interpretability of dense word embeddings by aligning the dimensions with predefined concepts from a suitable lexicon. A quantitative measure is required to reliably evaluate the achieved improvement. One of the methods to measure the

Table 3. Words with largest dimension values for the proposed algorithm

GOVERNMENT	CHOICE	BOOK	NEWS	PROPERTY IN GENERAL
republic	poll	editor	radio	lands
province	shortlist	publisher	news	land
provinces	vote	magazine	tv	ownership
government	selection	writer	broadcasting	possession
administration	televoting	author	broadcast	assets
prefecture	preference	hardcover	broadcasts	acquired
governor	choosing	paperback	simulcast	property
county	choose	books	channel	acres
monarchy	choice	page	television	estate
region	chosen	press	cnn	lease
territory	elect	publishing	jazeera	inheritance
autonomous	list	edited	fm	manor
administrative	election	volume	programming	holdings
minister	select	encyclopedia	bbc	plows
senate	preferential	published	newscast	estates
districts	option	publications	simulcasts	owner
democratic	voters	bibliography	syndicated	feudal
legislature	ballots	periodical	media	heirs
abolished	votes	publication	reporter	freehold
presidency	sssis	essayist	cbs	holding
⋮	⋮	⋮	⋮	⋮
	NUMBERS IN			
TEACHING	THE ABSTRACT	PATERNITY	WARFARE	FOOD
curriculum	integers	family	battle	meal
exam	Polynomial	paternal	war	dishes
training	integer	maternal	battles	bread
school	Polynomials	father	combat	eaten
students	logarithm	grandfather	military	dessert
toefl	modulo	grandmother	warfare	cooked
exams	formula	mother	fighting	foods
teaching	coefficients	ancestry	battlefield	dish
schools	multiplication	son	guerrilla	food

Table 3. Continued

GOVERNMENT	CHOICE	BOOK	NEWS	PROPERTY IN GENERAL
education	finite	hemings	fought	meat
teach	logarithms	ancestor	campaign	eating
karate	algebra	patrilineal	fight	cuisine
taught	integrals	daughter	insurgency	beverage
courses	primes	grandson	armed	soup
civics	divisor	descent	tactics	snack
instruction	compute	house	operations	pork
syllabus	arithmetic	parents	army	eat
test	algorithm	descendant	mujahideen	wine
examinations	theorem	grandparents	armies	beef
instructor	quadratic	line	soldiers	fried
⋮	⋮	⋮	⋮	⋮

interpretability is the word intrusion test (Chang *et al.* 2009), where manual evaluations from multiple human evaluators for each embedding dimension are used. Deeming this manual method expensive to apply, Senel *et al.* (2018a) introduced a SEMCAT to automatically quantify interpretability. We use both of these approaches to quantitatively verify our proposed method in the following two subsections.

5.2.1 Automated evaluation for interpretability

Specifically, we apply a modified version of the approach presented in Senel *et al.* (2018b) in order to consider possible sub-groupings within the categories in SEMCAT.^f Interpretability scores are calculated using Interpretability Score (*IS*) as given below:

$$\begin{aligned}
 IS_{i,j}^+ &= \max_{n_{min} \leq n \leq n_j} \frac{|S_j \cap V_i^+(\lambda \times n)|}{n} \times 100 \\
 IS_{i,j}^- &= \max_{n_{min} \leq n \leq n_j} \frac{|S_j \cap V_i^-(\lambda \times n)|}{n} \times 100 \\
 IS_{i,j} &= \max (IS_{i,j}^+, IS_{i,j}^-) \\
 IS_i &= \max_j IS_{i,j}, \quad IS = \frac{1}{D} \sum_{i=1}^D IS_i
 \end{aligned} \tag{4}$$

^fPlease note that the usage of “category” here in the setting of SEMCAT should not be confused with the “categories” of Roget’s Thesaurus.

Table 4. Words with largest dimension values for the proposed algorithm—less satisfactory examples

MOTION	TASTE	REDUNDANCY	FEAR
nektonic	polish	eusebian	horror
rate	classical	margin	fear
mobile	taste	drug	dread
movement	culture	arra	trembling
motion	corinthian	overflow	scare
evolution	przeworsk	overdose	terror
gait	artistic	extra	panic
velocity	judge	excess	anxiety
novokubansk	aesthetic	bonus	$\phi\delta\beta\sigma\zeta$
brownian	amateur	synaxarion	phobia
port	critic	load	fright
flow	kraków	padding	terrible
gang	elegance	crowd	frighten
roll	esthetics	redundancy	pale
stride	plaquemine	overrun	vacui
run	judgment	boilerplate	haunt
kinematics	connoisseur	excessive	afraid
stream	katarzyna	$\tau\iota\tau\lambda\omicron\iota$	fearful
walk	cucuteni	lavish	frightened
drift	warsaw	gorge	shaky
⋮	⋮	⋮	⋮

In Equation (4), $IS_{i,j}^+$ and $IS_{i,j}^-$ represent the interpretability scores in the positive and negative directions of the i th dimension ($i \in \{1, 2, \dots, D\}$, D is the number of dimensions in the embedding space) of word embedding space for the j th category ($j \in \{1, 2, \dots, K\}$, K is the number of categories in SEMCAT, $K = 110$) in SEMCAT, respectively. S_j is the set of words in the j th category in SEMCAT, and n_j is the number of words in S_j . n_{min} corresponds to the minimum number of words required to construct a semantic category (i.e., represent a concept). $V_i(\lambda \times n_j)$ represents the set of $\lambda \times n_j$ words that have the highest (V_i^+) and lowest (V_i^-) values in i th dimension of the embedding space. \cap is the intersection operator and $|\cdot|$ is the cardinality operator (number of elements) for the intersecting set. In Equation (4), IS_i gives the interpretability score for the i th dimension and IS gives the average interpretability score of the embedding space.

Figure 3 presents the measured average interpretability scores across dimensions for original GloVe embeddings, for the proposed method and for the other five methods we compare, along with a randomly generated embedding. Results are calculated for the parameters $\lambda = 5$ and $n_{min} \in \{5, 6, \dots, 20\}$. Our proposed method significantly improves the interpretability for all n_{min}




Figure 3. Interpretability scores averaged over 300 dimensions for the original GloVe method, the proposed method, and five alternative methods along with a randomly generated baseline embedding for $\lambda = 5$.

compared to the original GloVe approach. and it outperforms all the alternative approaches by a large margin especially for lower n_{min} .

The proposed method and interpretability measurements are both based on utilizing concepts represented by word groups. Therefore, it is expected that there will be higher interpretability scores for some of the dimensions for which the imparted concepts are also contained in SEMCAT. However, by design, word groups that they use are formed using different sources and are independent. Interpretability measurements use SEMCAT, while our proposed method utilizes Roget’s Thesaurus.

5.2.2 Human evaluation for interpretability: word intrusion test

Although measuring interpretability of imparted word embeddings with SEMCAT gives successful results, making another test involving human judgment surely enhances reliability of imparted word embeddings. One of the tests that includes human judgment for interpretability is the word intrusion test (Chang *et al.* 2009). Word intrusion test is a multiple choice test where each choice is a separate word. Four of these words are chosen among the words whose vector values at a specific dimension are high and one is chosen from the words whose vector values are low at that specific dimension. This word is called an intruder word. If the participant can distinguish the intruder word from others, it can be said that this dimension is interpretable. If the underlying word embeddings are interpretable across the dimension, the intruder words can be easily found by human evaluators.

In order to increase the reliability of the test, we both used imparted and original GloVe embeddings for comparison. For each dimension of both embeddings, we prepare a question. We shuffled the questions in a random order so that participant cannot know which question comes from which embedding. In total, there are 600 questions (300 GloVe + 300 imparted GloVe) with 5 choices for each.⁸ We apply the test on five participants. Results tabulated in Table 5 show that our proposed method significantly improves the interpretability by increasing the average correct answer percentage approximately from 28% for baseline to 71% for our method.

⁸Questions in this word intrusion test are available at <https://github.com/koc-lab/imparting-interpretability>.

Table 5. Word intrusion test results: correct answers out of 300 questions

	GloVe	Imparted
Participants 1–5	80/88/82/78/97	212/170/207/229/242
Mean/std	85/6.9	212/24.4

Table 6. Correlations for word-similarity tests

Dataset (EN-)	GloVe	Word2Vec	OIWE-IPG	SOV	SPINE	Word2Sense	Proposed
WS-353-ALL	0.612	0.7156	0.634	0.622	0.173	0.690	0.657
SIMLEX-999	0.359	0.3939	0.295	0.355	0.090	0.380	0.381
VERB-143	0.326	0.4430	0.255	0.271	0.293	0.271	0.348
SimVerb-3500	0.193	0.2856	0.184	0.197	0.035	0.234	0.245
WS-353-REL	0.578	0.6457	0.595	0.578	0.134	0.695	0.619
RW-STANF.	0.378	0.4858	0.316	0.373	0.122	0.390	0.382
YP-130	0.524	0.5211	0.353	0.482	0.169	0.420	0.589
MEN-TR-3k	0.710	0.7528	0.684	0.696	0.298	0.769	0.725
RG-65	0.768	0.8051	0.736	0.732	0.338	0.761	0.774
MTurk-771	0.650	0.6712	0.593	0.623	0.199	0.665	0.671
WS-353-SIM	0.682	0.7883	0.713	0.702	0.220	0.720	0.720
MC-30	0.749	0.8112	0.799	0.726	0.330	0.735	0.776
MTurk-287	0.649	0.6645	0.591	0.631	0.295	0.674	0.634
Average	0.552	0.6141	0.519	0.538	0.207	0.570	0.579

5.3 Performance evaluation of the embeddings

It is necessary to show that the semantic structure of the original embedding has not been damaged or distorted as a result of aligning the dimensions with given concepts, and that there is no substantial sacrifice involved from the performance that can be obtained with the original GloVe. To check this, we evaluate performances of the proposed embeddings on word-similarity (Faruqui and Dyer 2014) and word-analogy (Mikolov *et al.* 2013c) tests. We also measure the performance on a downstream sentiment classification task. We compare the results with the original embeddings and the four alternatives excluding Parsimax (Park *et al.* 2017) since orthogonal transformations will not affect the performance of the original embeddings on these tests.

Word-similarity test measures the correlation between word-similarity scores obtained from human evaluation (i.e., true similarities) and from word embeddings (usually using cosine similarity). In other words, this test quantifies how well the embedding space reflects human judgments in terms of similarities between different words. The correlation scores for 13 different similarity test sets and their averages are reported in Table 6. We observe that, let alone a reduction in performance, the obtained scores indicate an almost uniform improvement in the correlation values for the proposed algorithm, outperforming all the alternatives except word2ve baseline on average. Although Word2Sense performed slightly better on some of the test sets, it should be noted that it is trained on a significantly larger corpus. Categories from Roget's Thesaurus are groupings

Table 7. Precision scores for the analogy test

Methods	# dims	Analg. (sem)	Analg. (syn)	Total
GloVe	300	78.94	64.12	70.99
Word2Vec	300	81.03	66.11	73.03
OIWE-IPG	300	19.99	23.44	21.84
SOV	3000	64.09	46.26	54.53
SPINE	1000	17.07	8.68	12.57
Word2Sense	2250	12.94	19.44	16.51
Proposed	300	79.96	63.52	71.15

Table 8. Precision scores for the semantic analogy test

Questions subset	No. of questions seen	GloVe	Word2Vec	Proposed
All	8783	78.94	81.03	79.96
At least one concept word	1635	67.58	70.89	67.89
All concept words	110	77.27	89.09	83.64

of words that are similar in some sense which the original embedding algorithm may fail to capture. These test results signify that the semantic information injected into the algorithm by the additional cost term is significant enough to result in a measurable improvement. It should also be noted that scores obtained by SPINE are unacceptably low on almost all tests indicating that it has achieved its interpretability performance at the cost of losing its semantic functions.

Word-analogy test is introduced in Mikolov *et al.* (2013a) and looks for the answers of the questions that are in the form of “X is to Y, what Z is to ?” by applying simple arithmetic operations to vectors of words X, Y, and Z. We present precision scores for the word-analogy tests in Table 7. It can be seen that the alternative approaches that aim to improve interpretability have poor performance on the word-analogy tests. However, our proposed method has comparable performance with the original GloVe embeddings. Our method outperforms GloVe in semantic analogy test set and in overall results, while GloVe performs slightly better in syntactic test set. This comparable performance is mainly due to the cost function of our proposed method that includes the original objective of the GloVe.

To investigate the effect of the additional cost term on the performance improvement in the semantic analogy test, we present Table 8. In particular, we present results for the cases where (i) all questions in the dataset are considered, (ii) only the questions that contain at least one concept word are considered, and (iii) only the questions that consist entirely of concept words are considered. We note specifically that for the last case, only a subset of the questions under the semantic category family.txt ended up being included. We observe that for all three scenarios, our proposed algorithm results in an improvement in the precision scores. However, the greatest performance increase is seen for the last scenario, which underscores the extent to which the semantic features captured by embeddings can be improved with a reasonable selection of the lexical resource from which the concept word groups were derived.

Lastly, we compare the model performances on a sentence-level binary classification task based on the Stanford Sentiment Treebank which consists of thousands of movie reviews (Socher *et al.*

Table 9. Accuracies (%) for sentiment classification task

GloVe	Word2Vec	OIWE-IPG	SOV	SPINE	Word2Sense	Proposed
72.62	77.91	73.47	77.45	74.07	81.32	78.31

Table 10. Comparisons against Random Token Projection/Roget Center Projection baselines

Task	GloVe	Random Token Projections	Roget Center Projections	Proposed
Word Analogy Test Semantic Results (%)	78.94	2.84	12.35	79.96
Word Analogy Test Syntactic Results (%)	64.12	4.42	19.70	63.52
Word Similarity Test Average Results	0.552	0.115	0.252	0.579
Sentiment Classification Test Results (%)	72.62	51.76	77.28	78.31
Interp _{-nim=5}	29.23	47.62	64.95	65.19
Interp _{-nim=10}	19.22	33.83	54.02	45.87

2013) and their corresponding sentiment scores. We omit the reviews with scores between 0.4 and 0.6, resulting in 6558 training, 824 development, and 1743 test samples. We represent each review as the average of the vectors of its words. We train an Support Vector Machine classifier on the training set, whose hyperparameters were tuned on the validation set. Classification accuracies on the test set are presented in Table 9. The proposed method outperforms the original embeddings and performs on par with the SOV. Pretrained Word2Sense embeddings outperform our method; however, it has the advantage of training on a larger corpus. This result, along with the intrinsic evaluations, shows that the proposed imparting method can significantly improve interpretability without a drop in performance.

In addition to the comparisons above, we also compare our method against two projection-based simple baselines. First, we project the GloVe vectors onto randomly selected 300 tokens which results in a new 300-dimensional embedding space (Random Token Projections (RTPs)). We repeat this process for 10 times independently and report the average results. Second, we calculate the average of the vectors for the words in each of the 300 word groups extracted from Roget's Thesaurus. Then, we project the original embeddings onto these average vectors to obtain Roget Center Projections (RCPs). Table 10 presents the results for the task performance and interpretability evaluations for these two baselines along with the original GloVe embeddings and imparted embeddings. Although, these simple projection-based methods are able to improve interpretability, they distort the inner structure of the embedding space and reduce its performance significantly.

5.4 Effect of weighting parameter k

The results presented in the above subsections are obtained by setting the model weighting parameter k to 0.1. However, we have also experimented with different k values to find the optimal value for the evaluation tests and to determine the effects of our model parameter k to the performance. Figure 4 presents the results of these tests for $k \in [0.02 - 0.4]$ range. Since parameter k adjusts the magnitude of the influence for the concept words (i.e., our additional term), average interpretability of the embeddings increases when k is increased. However, the increase in the interpretability saturates and we almost hit the diminishing returns beyond $k = 0.1$. It can also be observed that




Figure 4. Effect of the weighting parameter k is tested using interpretability (top left, $n_{min} = 5$ and $\lambda = 5$), word-analogy (top right) and word-similarity (bottom) tests for $k \in [0.02 - 0.4]$.

by further increasing k beyond 0.3, no additional increase in the interpretability can be obtained. This is because interpretability measurements are based on the ranking of words in the embedding dimensions. With increasing k , concept words (from Roget’s Thesaurus) are strongly forced to have larger values in their corresponding dimensions. However, their ranks will not further increase significantly after they all reach to the top. In other words, a value of k between 0.1 and 0.3 is sufficient in terms of interpretability.

We now move to test whether high k values will harm the underlying semantic structure or not. To test this, our standard analogy and word-similarity tests that are given in the previous subsections are deployed for a range of k values. Analogy test results show that larger k values reduce the performance of the resulting embeddings on syntactic analogy tests, while semantic analogy performance is not significantly affected. For the word-similarity evaluations, we have used 13 different datasets; however, in Figure 4, we present four of them as representatives along with the average of all 13 test sets to simplify the plot. Word-similarity performance slightly increases for most of the datasets with increasing k , while performance slightly decreases or does not change for the others. On average, word-similarity performance increases slowly with increasing k and it is less sensitive to the variation of k value than the interpretability and analogy tests.

Combining all these experiments and observations, empirically setting k to 0.1 is a reasonable choice to compromise the trade-off since it significantly improves interpretability without sacrificing analogy/similarity performances.

5.5 Effect of number of dimensions

All the results presented above for our imparting method are for 300-dimensional vectors which is a common choice to train word embeddings. We trained the imparting method with the 300

Table 11. Effect of embedding dimension to the imparting performance

Task	200 Dimensions		300 Dimensions		400 Dimensions	
	Original	Proposed	Original	Proposed	Original	Proposed
Word Analogy Test Semantic Results (%)	77.13	77.64	78.94	79.96	80.46	80.95
Word Analogy Test Syntactic Results (%)	61.91	61.29	64.12	63.52	64.70	64.06
Word Similarity Test Average Results	0.541	0.557	0.552	0.579	0.554	0.586
Sentiment Classification Test Results (%)	76.13	76.82	72.62	78.31	77.57	77.85
Interp. $n_{nim}=5$	29.64	69.91	29.23	65.19	26.85	61.81
Interp. $n_{nim}=10$	19.73	51.09	19.22	45.87	17.29	40.42

word groups extracted from Roget's Thesaurus for all experiments in order to make full use of embedding dimensions for interpretability. To investigate the effect of number of dimensions on interpretability and performance of the imparted word embeddings, we also trained the proposed method using 200 and 400 dimensions. In both cases, we make full use of the embedding dimensions. To achieve this, we extracted 200 and 400 word groups from Roget's Thesaurus by discarding categories that have less than 76 and 36 words, respectively.

Table 11 presents the results for interpretability and performance evaluations of the GloVe and imparted embeddings for 200, 300, and 400 dimensions. For the word-similarity evaluations, results are averaged across the 13 different datasets. On the performance evaluations, imparted embeddings perform on par with the original embeddings regardless of the dimensionality. It can be seen that performance on the intrinsic tests slightly improve with increasing dimensionality for both embeddings, while the performance on the classification task does not change significantly. For the interpretability evaluations, the trend is the opposite. In general, interpretability decreases with increasing dimensionality since it is more difficult to consistently achieve good interpretability in more dimensions. However, imparted embeddings are significantly more interpretable than the original embeddings in all cases. Based on these results, we argue that 300 is a decent selection for dimensionality in terms of performance, interpretability, and computational efficiency.

5.6 Design of function $g(x)$

As presented in Section 4, our proposed method encourages the trained word vectors to have larger values if the underlying word is semantically close to a collection of concept word groups, one assigned to each vector dimension. However, a mechanism is needed to control amounts of these inflated values. To facilitate and control this behavior, a function $g(x)$ should be used. This function serves to increase the total cost incurred if the value of the l th dimension for the two vector representations $w_{i,l}$ and $\hat{w}_{j,l}$ for a concept word w_i with $i \in F_l$ fails to be large enough. In this subsection, we will elaborate more on the design and selection methodology of this function $g(x)$.

First of all, it is very obvious to see that $g(x)$ should be a positive monotone decreasing function because the concept words taking small values in the dimension that corresponds to word groups that they belong to should be penalized more harshly so that they are forced to take larger values. On the other hand, if their corresponding dimensions are large enough, contribution to the overall cost term coming from this objective should not increase much further. In other words, value of $g(x)$ should go to positive infinity as x decreases and go to 0 as x increases. Outstanding natural candidates for $g(x)$ with these features are then exponential




Figure 5. Interpretability scores averaged over 300 dimensions for the proposed method for different forms of function $g(x)$.

decay and reciprocal of polynomials with odd degrees, such as $1/x$ or $1/(x^3 + 1)$. Such polynomial functions satisfy the decaying requirement. However, for negative values of x , function can take negative values. This situation is undesirable because it leads to decrease in the cost function. Therefore, polynomials can only be used for the positive x values. What is left is exponential decays. In exponential case, there is no concern for $g(x)$ taking negative values. All we need to do is just to adjust the decaying rate. Too fast decays can make the additional objective lose its meaning, while too slow decays can break the structure of the general objective function by disproportionately putting more emphasis on our proposed modification term that is added to the original cost term of the embedding. (Here, it should also be noted that the adjustment of this parameter needs to be done in conjunction with parameter k that controls the blending of the two sub-objectives, which has also been analyzed in detail in Section 5.4.)

To further study several alternatives, we have also considered piecewise functions composed of both decaying exponentials and polynomials to facilitate and study properties of both. After several experiments for hyperparameter estimation, we conclude that the functions like $(1/2)\exp(-2x)$ or $(1/3)\exp(-3x)$ give best results, and we have proposed the function in Equation (3) which switches continuously from a decaying exponential to a reciprocal of a polynomial when x becomes greater than 0.5 (transition from $(1/2)\exp(-2x)$ to $1/(4ex)$ is adjusted such that $g(x)$ is continuous). A piecewise function is formed such that its polynomial part decays more slowly than its exponential part does, so that the objective function can keep pushing the words with small x values a little bit longer. Furthermore, $\exp(-x)$ has also been tried but since its decay is too slow and it takes large values, GloVe did not converge properly. On the other hand, faster decays also did not work since they quickly neutralize the additional interpretability cost term.

Experimental results for interpretability scores averaged over 300 dimensions are presented in Figure 5 for $g(x)$ options $(1/2)\exp(-2x)$, $(1/3)\exp(-3x)$, and the piecewise function given in Equation (3) (k is kept at 0.10). Experiments show that there is not much difference between single exponential function and piecewise function, but the single exponential decay and piecewise $g(x)$ with decay rates of -2 seem to be slightly better. To sum up, one can choose the piecewise option since it includes both natural options.

6 Conclusion

We presented a novel approach to impart interpretability into word embeddings. We achieved this by encouraging different dimensions of the vector representation to align with predefined concepts, through the addition of an additional cost term in the optimization objective of the GloVe algorithm that favors a selective increase for a prespecified input of concept words along each dimension.

We demonstrated the efficacy of this approach by applying qualitative and quantitative evaluations based on both automated metrics and on manual human annotations for interpretability. We also showed via standard word-analogy and word-similarity tests that the semantic coherence of the original vector space is preserved, even slightly improved. We have also performed and reported quantitative comparisons with several other methods for both interpretability increase and preservation of semantic coherence. Upon inspection of Figure 3 and Tables 6–8 altogether, it should be noted that our proposed method achieves both of the objectives simultaneously, increased interpretability and preservation of the intrinsic semantic structure.

An important point was that, while it is expected for words that are already included in the concept word groups to be aligned together since their dimensions are directly updated with the proposed cost term, it was also observed that words not in these groups also aligned in a meaningful manner without any direct modification to their cost function. This indicates that the cost term we added works productively with the original cost function of GloVe to handle words that are not included in the original concept word groups but are semantically related to those word groups. The underlying mechanism can be explained as follows. While the outside lexical resource we introduce contains a relatively small number of words compared to the total number of words, these words and the categories they represent have been carefully chosen and in a sense, “densely span” all the words in the language. By saying “span,” we mean they cover most of the concepts and ideas in the language without leaving too many uncovered areas. With “densely,” we mean all areas are covered with sufficient strength. In other words, this subset of words is able to constitute a sufficiently strong skeleton, or scaffold. Now remember that GloVe works to align or bring closer related groups of words, which will include words from the lexical source. So the joint action of aligning the words with the predefined categories (introduced by us) and aligning related words (handled by GloVe) allows words not in the lexical groups to also be aligned meaningfully. We may say that the non-included words are “pulled along” with the included words by virtue of the “strings” or “glue” that is provided by GloVe. In numbers, the desired effect is achieved by manipulating less than only 0.05% of parameters of the entire word vectors. Thus, while there is a degree of supervision coming from the external lexical resource, the rest of the vocabulary is also aligned indirectly in an unsupervised way. This may be the reason why, unlike earlier proposed approaches, our method is able to achieve increasing interpretability without destroying underlying semantic structure, and consequently without sacrificing performance in benchmark tests.

Upon inspecting the second column of Table 4, where qualitative results for concept TASTE are presented, another insight regarding the learning mechanism of our proposed approach can be made. Here, it seems understandable that our proposed approach, along with GloVe, brought together the words *taste* and *polish*, and then the words *Polish* and, for instance, *Warsaw* are brought together by GloVe. These examples are interesting in that they shed insight into how GloVe works and the limitations posed by polysemy. It should be underlined that the present approach is not totally incapable of handling polysemy but cannot do so perfectly. Since related words are being clustered, sufficiently well-connected words that do not meaningfully belong along with others will be appropriately “pulled away” from that group by several words, against the less effective, inappropriate pull of a particular word. Even though *polish* with lowercase “p” belongs where it is, it is attracting *Warsaw* to itself through polysemy and this is not meaningful. Perhaps because *Warsaw* is not a sufficiently well-connected word, it ends being dragged

along, although words with greater connectedness to a concept group might have better resisted such inappropriate attractions. Addressing polysemy and meaning conflation deficiency is beyond the scope and intention of our proposed model that represents all senses of a word with a single vector. However, our model may open up new directions of research in the intermingled studies on representing word senses and contextualized word embeddings (Jang, Myaeng, and Kim 2018; Camacho-Collados and Pilehvar 2018). For example, the associated concepts of our model assigned to the word embedding dimensions may be chosen such that different senses are assigned to different dimensions.

In this study, we used the GloVe algorithm as the underlying dense word embedding scheme to demonstrate our approach. However, we stress that it is possible for our approach to be extended to other word embedding algorithms which have a learning routine consisting of iterations over cooccurrence records, by making suitable adjustments in the objective function. Since word2vec model is also based on the cocurrences of words in a sliding window through a large corpus, we expect that our approach can also be applied to word2vec after making suitable adjustments, which can be considered as an immediate future work for our approach. Although the semantic concepts are encoded in only one direction (positive) within the embedding dimensions, it might be beneficial to pursue future work that also encodes opposite concepts, such as good and bad, in two opposite directions of the same dimension.

The proposed methodology can also be helpful in computational cross-lingual studies, where the similarities are explored across the vector spaces of different languages (Mikolov, Le, and Sutskever 2013b; Senel *et al.* 2017).

Acknowledgments. We thank Dr. Tolga Cukur (Bilkent University) for fruitful discussions. We would also like to thank the anonymous reviewers for their many comments which significantly improved the quality of our manuscript. We also thank the anonymous human subjects for their help in undertaking the word intrusion tests.

Financial support. H. M. Ozaktas acknowledges partial support of the Turkish Academy of Sciences.

References

- Arora S., Li Y., Liang Y., Ma T. and Risteski A. (2018). Linear algebraic structure of word senses, with applications to polysemy. *Transactions of the Association of Computational Linguistics* 6, 483–495.
- Bojanowski P., Grave E., Joulin A. and Mikolov T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* 5, 135–146.
- Bollagela D., Alsuhaibani M., Maehara T. and Kawarabayashi K. (2016). Joint word representation learning using a corpus and a semantic lexicon. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. Phoenix, AZ, USA: Association for the Advancement of Artificial Intelligence (AAAI), pp. 2690–2696.
- Camacho-Collados J. and Pilehvar M.T. (2018). From word to sense embeddings: a survey on vector representations of meaning. *Journal of Artificial Intelligence Research* 63(1), 743–788.
- Chang J., Gerrish S., Wang C., Boyd-Graber, J.L. and Blei D.M. (2009). Reading tea leaves: how humans interpret topic models. In Bengio Y., Schuurmans D., Lafferty J. D., Williams C. K. I. and Culotta A. (eds.), *Advances in Neural Information Processing Systems*, pp. 288–296. Curran Associates, Inc.
- Chen D. and Manning C. (2014). A fast and accurate dependency parser using neural networks. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar: Association for Computational Linguistics, pp. 740–750
- Chen X., Duan Y., Houthoofd R., Schulman J., Sutskever I. and Abbeel P. (2016). Infogan: interpretable representation learning by information maximizing generative adversarial nets. In Lee D. D., Sugiyama M., Luxburg U. V., Guyon I. and Garnett R. (eds.), *Advances in Neural Information Processing Systems*, pp. 2172–2180. Curran Associates, Inc.
- Das R., Zaheer M. and Dyer C. (2015). Gaussian LDA for topic models with word embeddings. 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)*. Beijing, China: Association for Computational Linguistics, pp. 795–804
- De Vine L., Kholgi M., Zuccon G., Sitbon L. and Nguyen A. (2015). Analysis of word embeddings and sequence features for clinical information extraction. In *Proceedings of the Australasian Language Technology Association Workshop 2015*. Parramatta, Australia, pp. 21–30.

- Dufter P. and Schutze H.** (2019). Analytical methods for interpretable ultradense word embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, pp. 1185–1191.
- Faruqui M. and Dyer C.** (2014). Community evaluation and exchange of word vectors at wordvectors.org. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Baltimore, MD, USA: Association for Computational Linguistics, pp. 19–24.
- Faruqui M., Tsvetkov Y., Yogatama D., Dyer C. and Smith, N.A.** (2015a). Sparse overcomplete word vector representations. 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)*. Beijing, China: Association for Computational Linguistics, pp. 1491–1500.
- Faruqui M., Dodge J., Juahar S.K., Dyer C., Hovy E. and Smith N.A.** (2015b). Retrofitting word vectors to semantic lexicons. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*. Denver, CO, USA: Association for Computational Linguistics, pp. 1606–1615.
- Firth J.R.** (1957a). *Papers in Linguistics, 1934–1951*. London: Oxford University Press.
- Firth J.R.** (1957b). A synopsis of linguistic theory. In Philological Society (Great Britain) (ed.), *Studies in Linguistic Analysis*, Oxford: Blackwell, pp. 1930–1955.
- Fyshe A., Talukdar P.P., Murphy B. and Mitchell T.M.** (2014). Interpretable semantic vectors from a joint model of brain- and text-based meaning. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Baltimore, MD, USA: Association for Computational Linguistics, pp. 489–499.
- Glavaš G. and Vulić I.** (2018). Explicit retrofitting of distributional word vectors. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia: Association for Computational Linguistics, pp. 34–45.
- Goldberg Y. and Hirst G.** (2017). Neural network methods in natural language processing. In *Synthesis Lectures on Human Language Technologies*. Morgan & Claypool Publishers.
- Goodman B. and Flaxman S.** (2017). European Union regulations on algorithmic decision-making and a “right to explanation”. *AI Magazine* 38(3), 50–57
- Harris Z.S.** (1954). Distributional structure. *Word* 10(2–3), 146–162.
- Herbelot A. and Vecchi E.M.** (2015). Building a shared world: mapping distributional to model-theoretic semantic spaces. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Lisbon, Portugal: Association for Computational Linguistics, pp. 22–32.
- Jang K.-R., Myaeng S.-H. and Kim S.-B.** (2018). Interpretable word embedding contextualization. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*. Brussels, Belgium: Association for Computational Linguistics, pp. 341–343.
- Juahar S.K., Dyer C. and Hovy E.** (2015). Ontologically grounded multi-sense representation learning for semantic vector space models. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*. Denver, CO, USA: Association for Computational Linguistics, pp. 683–693.
- Johansson R. and Nieto P.L.** (2015). Embedding a semantic network in a word space. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*. Denver, CO, USA: Association for Computational Linguistics, pp. 1428–1433.
- Joshi A., Tripathi V., Patel K., Bhattacharyya P. and Carman M.** (2016). Are word embedding-based features useful for sarcasm detection?. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Austin, TX, USA: Association for Computational Linguistics, pp. 1006–1011.
- Iacobacci I., Pilehvar M.T. and Navigli, R.** (2016). Embeddings for word sense disambiguation: an evaluation study. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany: Association for Computational Linguistics, pp. 897–907.
- Levy O. and Goldberg Y.** (2014). Dependency-based word embeddings. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Baltimore, MD, USA: Association for Computational Linguistics, pp. 302–308.
- Liu Y., Liu Z., Chua, T.-S. and Sun M.** (2015). Topical word embeddings. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. Austin, TX, USA: Association for the Advancement of Artificial Intelligence (AAAI), pp. 2418–2424.
- Liu Q., Jiang H., Wei S., Ling Z.-H. and Hu Y.** (2015). Learning semantic word embeddings based on ordinal knowledge constraints. 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)*. Beijing, China: Association for Computational Linguistics, pp. 1501–1511.

- Luo H., Liu Z., Luan H.-B. and Sun M. (2015). Online learning of interpretable word embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Lisbon, Portugal: Association for Computational Linguistics, pp. 1687–1692
- Mikolov T., Chen K., Corrado G. and Dean J. (2013a). Efficient estimation of word representations in vector space. arXiv preprint [arXiv:1301.3781](https://arxiv.org/abs/1301.3781).
- Mikolov T., Le Q.V. and Sutskever I. (2013b). Exploiting similarities among languages for machine translation. arXiv preprint [arXiv:1309.4168](https://arxiv.org/abs/1309.4168).
- Mikolov T., Sutskever I., Chen K., Corrado G.S. and Dean J. (2013c). Distributed representations of words and phrases and their compositionality. In Burges C. J. C., Bottou L., Welling M., Ghahramani Z. and Weinberger K. Q. (eds.), *Advances in Neural Information Processing Systems*, pp. 3111–3119. Curran Associates, Inc.
- Miller G.A. (1995). WordNet: a lexical database for English. *Communications of the ACM* 38(11), 39–41.
- Moody C.E. (2016). Mixing dirichlet topic models and word embeddings to make lda2vec. arXiv preprint [arXiv:1605.02019](https://arxiv.org/abs/1605.02019).
- Mrkšić N., Ó Séaghdha D., Thomson B., Gašić M., Rojas-Barahona L.M., Su P.-H., David V., Wen T.-H. and Young S. (2016). Counter-fitting word vectors to linguistic constraints. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*. San Diego, CA, USA: Association for Computational Linguistics, pp. 142–148.
- Mrkšić N., Vulić I., Ó Séaghdha D., Leviant I., Reichart R., Gašić M., Korhonen A. and Young S. (2017). Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints. *Transactions of the Association for Computational Linguistics* 5, 309–324.
- Murphy B., Talukdar P.P. and Mitchell T.M. (2012). Learning effective and interpretable semantic models using non-negative sparse embedding. In *Proceedings of COLING 2012*. Mumbai, India: The COLING 2012 Organizing Committee, pp. 1933–1950.
- Panigrahi A., Simhadri H.V. and Bhattacharyya C. (2019). Word2Sense: sparse interpretable word embeddings. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, pp. 5692–5705.
- Park S., Bak J. and Oh A. (2017). Rotated word vector representations and their interpretability. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Copenhagen, Denmark: Association for Computational Linguistics, pp. 401–411.
- Pennington J., Socher R. and Manning C. (2014). GloVe: global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar: Association for Computational Linguistics, pp. 1532–1543.
- Ponti E.M., Vulić I., Glavač G., Mrkšić N. and Korhonen A. (2018). Adversarial propagation and zero-shot cross-lingual transfer of word vector specialization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Brussels, Belgium: Association for Computational Linguistics, pp. 282–293.
- Roget P.M. (1911). *Roget's Thesaurus of English Words and Phrases*. New York: T.Y. Crowell Company.
- Roget P.M. (2008). *Roget's International Thesaurus, 3/E*. New Delhi: Oxford & IBH Publishing Company Pvt. Limited.
- Senel L.K., Utlu I., Yucesoy V., Koç A. and Cukur T. (2018). Semantic structure and interpretability of word embeddings. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 26(10), 1769–1779.
- Senel L.K., Yucesoy V., Koç A. and Cukur T. (2018). Interpretability analysis for Turkish word embeddings. In *26th Signal Processing and Communications Applications Conference (SIU)*. Izmir, Turkey: IEEE, pp. 1–4.
- Senel L.K., Yucesoy V., Koç A. and Cukur T. (2017). Measuring cross-lingual semantic similarity across European languages. In *40th International Conference on Telecommunications and Signal Processing (TSP)*. Barcelona, Spain: IEEE, pp. 359–363.
- Sienčnik S.K. (2015). Adapting word2vec to named entity recognition. In *Proceedings of the 20th Nordic Conference of Computational Linguistics (NODALIDA 2015)*. Vilnius, Lithuania: Linköping University Electronic Press, Sweden, pp. 239–243.
- Shi B., Lam W., Jameel S., Schockaert S. and Lai K.P. (2017). Jointly learning word embeddings and latent topics. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. Shinjuku, Tokyo, Japan: ACM, pp. 375–384.
- Socher R., Pennington J., Huang E.H., Ng, A.Y. and Manning C.D. (2011). Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Edinburgh, Scotland, UK: Association for Computational Linguistics, pp. 151–161.
- Socher R., Perelygin A., Wu J., Chuang J., Manning C.D., Ng A.Y. and Potts C. (2013). Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Seattle, WA, USA: Association for Computational Linguistics, pp. 1631–1642.
- Subramanian A., Pruthi D., Jhamtani H., Berg-Kirkpatrick T. and Hovy E. (2018). SPINE: sparse interpretable neural embeddings. In *Proceedings of the Thirty Second AAAI Conference on Artificial Intelligence* New Orleans, LA, USA: Association for the Advancement of Artificial Intelligence (AAAI), pp. 4921–4928.

- Turian J., Ratinov L.-A. and Bengio Y.** (2010). Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Uppsala, Sweden: Association for Computational Linguistics, pp. 384–394.
- Xu C., Bai Y., Bian J., Gao B., Wang G., Liu X. and Liu T.-Y.** (2014). RC-NET: a general framework for incorporating knowledge into word representations. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. Shanghai, China: ACM, pp. 1219–1228.
- Yu L.-C., Wang J., Lai K.R. and Zhang X.** (2017). Refining word embeddings for sentiment analysis. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Copenhagen, Denmark: Association for Computational Linguistics, pp. 545–550.
- Yu M. and Dredze M.** (2014). Improving lexical embeddings with semantic knowledge. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Baltimore, MD, USA: Association for Computational Linguistics, pp. 545–550.
- Zobnin A.** (2017). Rotations and interpretability of word embeddings: the case of the Russian language. In *International Conference on Analysis of Images, Social Networks and Texts*. Moscow, Russia: Springer International Publishing, pp. 116–128.