

On the Impact of Knowledge-based Linguistic Annotations in the Quality of Scientific Embeddings

Andres Garcia-Silva^{a,*}, Ronald Denaux^{a,*}, Jose Manuel Gomez-Perez^{a,*}

^a*Expert.ai Research Lab, Prof. Waskman 10, 28036 Madrid, Spain*

Abstract

In essence, embedding algorithms work by optimizing the distance between a word and its usual context in order to generate an embedding space that encodes the distributional representation of words. In addition to single words or word pieces, other features which result from the linguistic analysis of text, including lexical, grammatical and semantic information, can be used to improve the quality of embedding spaces. However, until now we did not have a precise understanding of the impact that such individual annotations and their possible combinations may have in the quality of the embeddings. In this paper, we conduct a comprehensive study on the use of explicit linguistic annotations to generate embeddings from a scientific corpus and quantify their impact in the resulting representations. Our results show how the effect of such annotations in the embeddings varies depending on the evaluation task. In general, we observe that learning embeddings using linguistic annotations contributes to achieve better evaluation results.

Keywords: Natural language processing, linguistic analysis, knowledge graphs, embeddings.

1. Introduction

Distributed word representations in the form of dense vectors, known as word embeddings, have become extremely popular for NLP tasks such as part-of-speech tagging, chunking, named entity recognition, semantic role labeling and synonym detection. Such interest was significantly sparked by the word2vec algorithm proposed in [1], which provided an efficient way to learn word embeddings from large corpora based on word context and negative sampling, and continues to date. Ample research has been put into the development of increasingly effective means to produce word embeddings, resulting in algorithms like GloVe [2], Swivel [3] or fastText [4]. Lately, much of this effort has been focused on neural language models that produce contextual word representations at a scale never seen before, like ELMo [5], BERT [6] and XLNet [7].

In general, word embedding algorithms are trained to minimize the distance between a token and its usual context of neighboring tokens in a document corpus. The set of tokens that occur in a text after some tokenization is

applied may include e.g. words ("making"), punctuation marks (";"), multi-word expressions ("day-care center") or combinations of words and punctuation marks ("However,"). Indeed, the information that is actually encoded in the resulting embeddings depends not only on the neighboring context of a word or its frequency in the corpus, but also on the strategies used to chunk the text into sequences of individual tokens.

In Table 1, we show several examples of tokenizations that can be produced by processing a sentence following different tokenization strategies. In the example, the original text is annotated with information on space-separated tokens (t) and surface forms (sf). Additional linguistic and semantic annotations can also be derived from the text, like the lemma (l) associated to each token, the corresponding sense or concept¹ extracted through word-sense disambiguation (c), and grammatical information about the role each token plays in the context of a particular sentence (g), like part-of-speech.

The resulting sequences of tokens used to train the embeddings differ significantly both in number of tokens and complexity. This can have a decisive impact on what information is actually encoded and consequently on the

*Corresponding author

Email addresses: agarcia@expert.ai (Andres Garcia-Silva), rdenaux@expert.ai (Ronald Denaux), jmgomez@expert.ai (Jose Manuel Gomez-Perez)

¹Note that in the example, concept information is linked to a knowledge graph through a unique identifier.

Table 1: Tokenizations for the first window of size $W=3$ for the sentence: “With regard to breathing in the locus coeruleus.”. First we show basic space-based tokenization. Next, sequences for surface forms, lemmas, syncons and grammar type.

context	t^{i-3}	t^{i-2}	t^{i-1}	t^i	t^{i+1}	t^{i+2}	t^{i+3}
t	With	regard	to	breathing	in	the	locus coeruleus,
sf	With regard to	breathing	in	the	locus coeruleus	\emptyset	
l	with regard to	breathe	\emptyset	\emptyset	locus coeruleus	\emptyset	
c	en#216081	en#76230	\emptyset	\emptyset	en#101470452	\emptyset	
g	PRE	VER	PRE	ART	NOU	PNT	

quality of the embeddings. In fact, the analysis of optimal tokenization strategies to train word embeddings working at either word (Word2Vec), character (ELMo), byte-pair encoding (BPEmb [8]) or word-piece (BERT) levels is already well covered in the literature. However, none of these approaches investigate the role that explicit linguistic and semantic annotations (henceforth, linguistic annotations) like the ones above-mentioned, as well as their combinations, can play in the generation of the embeddings.

In this paper, we apply this observation to investigate how linguistic annotations extracted from the text can be used to produce enriched, higher-quality embeddings. In doing so, we focus on the scientific domain as a field with a significantly complex terminology, rich in multi-word expressions and other linguistic artifacts. Scientific resources, such as Springer Nature’s SciGraph [9] which contains over 3M scholarly publications, can be considerably rich in domain-specific multi-word expressions. This is particularly noticeable in Table 2 in the case of grammar types like adverbs (ADV), nouns (NOU), noun phrases (NPH), proper names (NPR) or entities (ENT); while others like conjunctions (CON) or propositions (PRO) are domain-independent.

We quantify the impact of each individual type of linguistic annotation and their possible combinations on the resulting embeddings following different approaches. First, we do this intrinsically through word similarity and relatedness, as well as analogical reasoning tasks. In this regard, we observe that current word similarity and analogy benchmarks, which have become the standard means for the intrinsic evaluation of word embeddings, are of little help in domains like science, which have highly specialized vocabularies. This observation highlights the need for domain-specific benchmarks. Second, we also evaluate extrinsically in two additional tasks: a word prediction task against scientific corpora extracted from SciGraph and Semantic Scholar [10], as well as general-purpose corpora like Wikipedia and UMBC [11], and a classification task based on the SciGraph taxonomy

Table 2: Multi-word expressions (mwe) vs. single-word expressions (swe) per grammar type (PoS) in SciGraph. Number of swe and mwe occurrences are in millions.

PoS	example	#swe	#mwe
ADJ	single+dose	104.7	0.514
ADV	in+situ	21	1.98
CON	even+though	34.6	3.46
ENT	august+2007	1.5	1.24
NOU	humid+climate	216.9	15.24
NPH	Ivor+Lewis	0.917	0.389
NPR	dorsal+ganglia	22.8	5.22
PRO	other+than	13.89	0.005
VER	take+place	69.39	0.755

of scientific disciplines.

The remainder of this paper is structured as follows: Section 2 contains an account of related work. Section 3 describes our theoretical framework, which addresses different ways to capture various types of linguistic and semantic information in a single embedding space. In this regard, we focus on Vecsgrafo [12], which allows jointly learning word, lemma and concept embeddings using different tokenization strategies, as well as linguistic annotations. Section 4 presents our experiments, where we identify significant findings. For example, embeddings learned from corpora enriched with concept information consistently improve word similarity and analogical reasoning results, while the latter additionally tends to benefit from surface form embeddings. Other results indicate e.g. that using grammatical information to generate scientific embeddings significantly improves the precision of classification but in contrast penalizes the quality of the embeddings for word prediction. Section 5 provides a more detailed discussion based on our findings. Finally, Section 6 concludes the paper.

2. Related work

The algorithms to learn word embeddings were originally categorized by Pennington et al. [2] as global matrix factorization (count-based) approaches [13, 2, 3] and local context window methods (prediction) [1, 4]. Levy et al. [13] proved that local context window methods implicitly factorize a word-context matrix. Although successfully applied in a variety of problems, word embedding algorithms usually work at a word or n-gram level. Context-free models such as word2vec or GloVe thus generate a single word embedding representation for each word in the vocabulary ("jaguar" would have the same embedding either as an animal or as a car brand). Plus, they ignore multi-word expressions, grammatical and word-sense information, which may be useful e.g. to deal with polysemy and homonymy.

Recently, considerable effort has shifted towards contextual word embeddings using neural language models based on e.g. bi-directional LSTMs [5] or transformer encoders [6]. Although these are without a doubt a powerful tool in most modern NLP architectures, we are still uncertain as to what information they actually capture. Recent studies, such as [14, 15], conclude that language models like BERT seem to learn surface features at the bottom layers, syntactic features in the middle, and semantic features at the top, following a compositional tree-like structure. In addition, other contributions have shown that BERT and ELMo output embeddings and hidden states encode parse tree information [16], while BERT can also attend at linguistic notions of syntax and coreference [17]. However, in spite of such efforts the linguistic and semantic information thus captured is still hard to interpret logically. In this paper, we use explicit linguistic and semantic features extracted from text and structured knowledge graphs and quantify their impact on the resulting embeddings.

Several approaches focus on learning embeddings that leverage structured graph representations like WordNet [18] to capture the semantics represented explicitly in such resources. Among such approaches, TransE [19], HolE [20], RDF2Vec [21], Graph Convolutions [22] and ComE+ [23] (see [24] for a recent review) directly encode the representations contained in the knowledge graph, which is typically an already condensed and filtered version of real-world data. A related research direction, with approaches like NASARI [25], SW2V [26] and Vecsgrafo [12], proposes to co-train word and concept embeddings based not only on knowledge graphs, but also on text corpora through lexical specificity or word-sense disambiguation. This has several advantages against using just the knowledge graph, including ad-

ditional coverage of the different domains of interest and the additional signal captured from the linguistic information contained in the text.

The main contribution of this paper is a comprehensive experimental study on the quality of embeddings generated from a scientific corpus using additional, explicit annotations with linguistic and semantic information extracted from the text. To this purpose, we follow the approach proposed by Vecsgrafo. However, in contrast to the experiments described in [12], which used general domain corpora, here we rely on domain-specific text from scientific publications and include the use of grammatical information (part-of-speech) in our experimentation. Also, our study covers all the possible combinations of the resulting embeddings, including those related to space-separated tokens, surface forms, lemmas, concepts, and grammatical information. Finally, we add an extrinsic evaluation task on scientific text classification to further assess the quality of the embeddings.

3. Theoretical framework

Next, we describe our approach to jointly learn embeddings using linguistic annotations on a text corpus, the strategies to mix and merge such embeddings, and the evaluation tasks used in this work.

3.1. Vecsgrafo: corpus-based word-concept embeddings

Vecsgrafo [12] is a method for jointly learning concept and word embeddings from a text corpus. Like word2vec, it uses a sliding window over the text to predict pairs of center and context words. Unlike word2vec, which relies on a simple tokenization strategy based on space-separated tokens (t), in Vecsgrafo the aim is to learn embeddings also for linguistic annotations, including surface forms ($sf \in SF$), lemmas ($l \in L$), grammar information ($g \in G$) and concept identifiers ($c \in C$) from a knowledge graph. To derive such additional information from a text corpus, in Vecsgrafo the use of a word-sense disambiguation pipeline (see Figure 1) is required to extract from the original text the corresponding surface forms, lemmas, grammar information, e.g. part-of-speech tags, and concept identifiers.

In contrast to simple tokenization strategies based on space separation, surface forms are the result of a grammatical analysis where one or more tokens can be grouped using the part of speech information e.g. in noun or verb phrases. Surface forms can include multi-word expressions that refer to concepts, e.g. "Cancer Research Center". A lemma is the base form of a word,

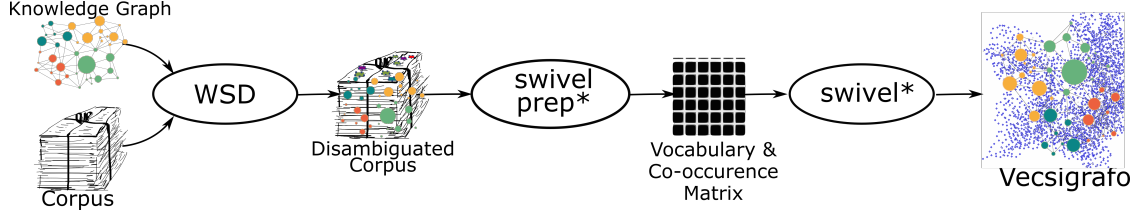


Figure 1: Vecsigrafo generation process (Figure taken from [12]).

Table 3: Example of research paper title and linguistic annotations added to jointly learn vecsigrafo embeddings for surface forms, lemmas and concepts. Vecsigrafo learns embeddings for each instance of the annotation types. Each box is a set of annotations with the following pattern: <surface form>|<lemma>|<concept>. Lemmas use the prefix "lem_" and concepts "en%".

t	a novel wasp gene mutation in a chinese boy with wiskott-aldrich syndrome			
sf_l_c	a	novel lem_novel en%23100008261	wasp lem_Wiskott-Aldrich+syndrome+protein	
		gene+mutation lem_gene+mutation en%23101415380	in a	chinese lem_Chinese en%2398003
		boy lem_boy en%2346011	with	wiskott-aldrich+syndrome lem_Wiskott-Aldrich+syndrome

e.g. surface forms "is", "are", "were" all have lemma "be". Due to the linguistic annotations that are included in addition to the original text, the word2vec approach is no longer directly applicable. The annotated text is a sequence of words and annotations. Thus, sliding a context window from a central word needs to take into consideration both words and annotations. For this reason, Vecsigrafo extends the Swivel [3] algorithm to first use the sliding window to count co-occurrences between pairs of linguistic annotations in the corpus. These are then used to estimate the point-wise mutual information (PMI) between each pair. The algorithm then tries to predict all of the pairs as accurately as possible.

To illustrate how Vecsigrafo works let us analyse the text, "*a novel wasp gene mutation in a chinese boy with Wiskott-Aldrich syndrome*". Swivel splits the text into 12 space-separated tokens, then builds the co-occurrence matrix with these tokens as features in rows and columns. In contrast, Vecsigrafo first enriches the text with the linguistic annotations that are shown in Table 3 and then each annotation (surface forms, lemmas, and concepts) is used as a feature in the swivel co-occurrence matrix. Note that Swivel understands "gene" and "mutation" as distinct tokens, while Vecsigrafo identifies them as a single surface form, a lemma and a concept, and then generates embeddings for each of them. More importantly, while for the ambiguous term "wasp" Swivel generates an embedding, Vecsigrafo realizes that in this context the surface form "wasp" refers to the lemma "Wiskott-Aldrich syndrome protein", and encodes this contextual information in the corresponding embeddings.

In the end, Vecsigrafo produces an embedding space:

$\Phi = \{(x, e) : x \in SF \cup L \cup G \cup C, e \in \mathbf{R}^n\}$. That is, for each entry x , it learns an embedding of dimension n , such that the dot product of two vectors predicts the PMI of the two entries. For the interested reader, a complete description of Vecsigrafo algorithm is available in [12].

3.2. Joint embedding space

To use Vecsigrafo embeddings, we need to annotate a corpus D with the linguistic annotations used to learn the embeddings. An annotated document in D is a sequence of annotation tuples $D_j = ((sf, l, g, c)_i)_{i=0}^{|D_j|-1}$. Then we need to generate out of D a vocabulary for each linguistic annotation $V(D) = \{v : v \in SF_D \cup L_D \cup G_D \cup C_D\}$, and represent any annotated document as a finite sequence of items in this vocabulary $S(D_j) = (sf_i, l_i, g_i, c_i)_{i=0}^{|D_j|-1}$. Following this approach, there is a direct mapping between any sequence element in $S(D_j)$ and Vecsigrafo embeddings in Φ .

In addition, it is possible to merge Vecsigrafo embeddings, such as those corresponding to surface forms and concepts. This is in line with common practice in the literature involving e.g. the combination of conventional and contextualized embeddings [5, 6] or the creation of multimodal embedding spaces [27]. Different operations can be applied on the embeddings, including concatenation, average, addition, and dimension reduction techniques such as principal component analysis PCA, singular vector decomposition SVD or neural auto-encoders.

3.3. Embeddings evaluation

Embeddings evaluation can be classified in intrinsic and extrinsic methods [28]. The former assesses whether the embedding space actually encodes the distributional context of words, while the latter makes an assessment according to the performance of a downstream task. The evaluation tasks described in this section are intended to shed light on the quality of the scientific embeddings generated using linguistic annotations and to compare them with regular, space-separated token embeddings. Since the embeddings are learned from a scientific corpus, it is not expected that they outperform the state of the art in intrinsic tasks, such as word similarity and relatedness [29] or analogical reasoning, where datasets comprise mostly terms and relations from general lexicon with little overlap with ours.

Word similarity and relatedness [30] is an intrinsic evaluation method that assesses whether the embedding space captures the semantic relatedness between words. To evaluate the embeddings in this task, we compute the correlation between an embedding similarity measure, usually the cosine function, and human similarity or relatedness scores. In the word embeddings literature, Spearman’s rho correlation is widely used [2, 3]. However, Pearson correlation, alone or in combination with Spearman’s, is reported in several papers on word similarity and relatedness [31, 32, 29] as a way to evaluate the strength of such similarity and relatedness between word pairs and the correct ranking. In this paper, we adhere to the latter and report both correlations’ metrics.

To work with the concept embeddings produced by Vecsgrafo, we use a concept-based word similarity measure. Essentially, each pair of words is mapped to concepts, using the most frequent word meaning as reference, and similarity is calculated on the corresponding concept embeddings.

Word analogy [33] is another intrinsic evaluation method where the goal is to find x such that the relation $x:y$ resembles a sample relation $a:b$ by operating on the corresponding word vectors.

Word prediction, introduced in [12], proposes the use of a test corpus to model how accurately the resulting embeddings can predict a linguistic annotation (t , sf , l , c or g). This approach essentially emulates the word2vec skipgram loss function on a test corpus unseen during training, iterating through the sequence of tokens in a corpus using a context window. For each center word or linguistic annotation, we calculate the (weighted) average of the embeddings for the context tokens and compare it to the embedding for the center word using cosine similarity. If the cosine similarity is

close to 1, the center word is being correctly predicted based on its context. We use this to plot the similarities for the whole embedding vocabulary, providing insights about possible disparities in the quality of common and uncommon words, as well as the applicability of the resulting embeddings to a different corpus.

Text classification: Since word embeddings are usually at the input of neural text classifiers, the performance of the latter can be used as an indicator of the fitness of the former. In contrast to models frequently used for text classification [34], such as Naive Bayes and Support Vector Machines (SVM) [35] which work in a highly dimensional space, deep neural networks work on a compressed space thanks to the use of fixed length embeddings. Neural networks with convolutional layers CNN [36] using a single layer [37, 38] or multiple-layers [39, 40] have shown state of the art performance for NLP tasks without the need of hand-crafted features [37], one of the main limitation of traditional approaches.

4. Experimental work

This section describes the corpus, the embeddings learned from the linguistic annotations, and the experiments carried out to gain insights on their usefulness in the evaluation tasks.

4.1. Learning scientific embeddings with Vecsgrafo

The scientific corpus used in the following experiments is extracted from SciGraph [9], a knowledge graph for scientific publications. From it, we extract titles and abstracts of articles and book chapters published between 2001 and 2017. The resulting corpus consists roughly of 3.2 million publications, 1.4 million distinct words, and 700 million tokens.

Next, we use Expert.ai’s Natural Language API to parse and disambiguate the text and add the linguistic annotations associated to each token or group of tokens. To this purpose, the API relies on its own knowledge graph, called Sensigrafo, which encodes linguistic information in a way similar to WordNet and follows a rule-based approach to word-sense disambiguation. Sensigrafo contains approximately 400K lemmas and 300K concepts interlinked via 61 relation types and is available in 14 languages. In addition to Sensigrafo, Vecsgrafo has been successfully trained on several knowledge graphs², including WordNet. For convenience, in this work we use Sensigrafo as the reference knowledge graph in our

²Tutorial on Hybrid Techniques for Knowledge-based NLP
<http://hybridnlp.expertsystemlab.com/tutorial>

Table 4: Account of token types and embeddings learned from the text in the title and abstracts from research articles and book chapters in English published between 2001 to 2017 and available in SciGraph.

Algorithm	Types	Total	Distinct	Embeddings
Swivel	<i>t</i>	707M	1,486,848	1,486,848
	<i>sf</i>	805M	5,090,304	692,224
Vecsgrafo	<i>l</i>	508M	4,798,313	770,048
	<i>g</i>	804M	25	8
	<i>c</i>	425M	212,365	147,456

experiments. Furthermore, it should be noted that any other NLP toolkit would have been equally suitable as long as the generation of the linguistic annotations used in this work are supported.

The corpus parsing and annotations generated by the Natural Language API are reported in Table 4. For space-separated tokens, we learned an initial set of embeddings using Swivel, which is the reference algorithm for Vecsgrafo. For the remaining linguistic annotations (*sf*, *l*, *g*, *c*), as well as their combinations of size 2 and 3, we learned their embeddings using Vecsgrafo. The quantitative difference between the number of distinct annotations and actual embeddings is due to filtering. Following previous findings [12], we filtered out elements with grammar type articles, punctuation marks or auxiliary verbs and generalized tokens with grammar type entity or person proper noun, replacing the original token with special tokens *grammar#ENT* and *grammar#NPH*, respectively.

4.2. Word similarity and relatedness

In Table 5, we report the 16 datasets used in the word similarity evaluation task. Note that we only applied concept similarity in those combinations where *sf* or *l* were not present. Table 6 shows the word similarity evaluation results for the Vecsgrafo embeddings in the 16 similarity and relatedness datasets.

From the data, we observe clear evidence that the embeddings derived from linguistic annotations, individually or jointly learned, perform better in the similarity task than using token embeddings. An individual analysis of each linguistic annotation shows that *l* embeddings outperform *c* embeddings, which in turn are superior to *sf* embeddings. Most of the similarity datasets in Table 5 contain mainly nouns and lemmatized verbs. Semantic and lemma embeddings seem to be better representations for these datasets, since they condense different words and morphological variations in one representation. Surface forms, on the other hand, may generate several embeddings for the different morphological vari-

Table 5: Similarity and relatedness datasets used for embeddings evaluation.

Dataset	Main Content	Rel. Type	Word Pairs	Raters per Pair
The RG-65 [41]	Nouns	Sim.	65	51
MC-30 [42]	Nouns	Sim.	30	51
WS-353-ALL [43]	Nouns	Mix.	353	13-16
WS-353-REL [44]	Nouns	Rel.	252	13-17
WS-353-SIM [44]	Nouns	Sim.	203	13-18
YP-130 [45]	Verbs	Rel.	130	6
VERB-143 [46]	Conjugated Verbs	Rel.	143	10
SIMVERB3500 [47]	lemmatized Verbs	Sim.	3500	10
MTurk-287 [48]	Nouns	Rel.	287	23
MTurk-771 [49]	Nouns	Rel.	771	20
MEN-TR-3K [50]	Nouns	Rel.	3000	10
SIMLEX-999-Adj [51]	Adjectives	Sim.	111	50
SIMLEX-999-Nou [51]	Nouns	Sim.	666	50
SIMLEX-999-Ver [51]	Verbs	Sim.	222	50
RW-STANFORD [52]	Noun, Verbs, Adj	Rel.	2034	10
SEMEVAL17 [53]	Named Entities, MWE	Rel.	500	3

ations of words, which might not be necessary for this task and can hamper the similarity estimation.

Regarding Vecsgrafo embeddings jointly learned for two or more linguistic annotations, semantic embeddings *c* improve the results in all the combinations where they were used. Concept embeddings enhance the lexical representations based on surface forms and the grouped morphological variations represented by lemmas in the similarity task. On the other hand, grammatical information *g* seems to hamper the performance of Vecsgrafo embeddings except when they are jointly learned with the semantic embeddings *c*. If we focus the analysis on the top 5 results in Table 6, *l* and *c* embeddings, jointly or individually learned, are the common elements in the embedding combinations that perform best in the similarity task. On the other hand, the performance of the embeddings from surface forms *sf* and related combinations is at the bottom of the table. In general, the embeddings based on *l* and *c* seem to be better correlated to the human notion of similarity than embeddings based on *sf* and *g*.

To summarize, the results obtained in the similarity task show evidence that the embeddings benefit from the additional semantic information contributed by the linguistic annotations. Lexical variations presented at the level of surface forms are better representations in this regard than space-separated tokens. The conflation

Table 6: Spearman’s rho for a subset of all the similarity datasets and the final learned vecisgrafo and swivel embeddings. Results are sorted in descending order on the harmonic mean of Spearman’s rho and Pearson correlation.

Embeddings	MC-30	MEN-TR-3k	MTurk-287	MTurk-771	RG-65	RW-STANFORD	SEMEVAL17	SIMLEX-999-Adj	SIMLEX-999-Nou	SIMLEX-999-Ver	SIMVERB3500	VERB-143	WS-353-ALL	WS-353-REL	WS-353-SIM	YP-130	Avg. Spearman	Avg. Pearson	Harmonic mean
l_c	0.642	0.645	0.556	0.526	0.561	0.359	0.553	0.421	0.330	0.078	0.162	0.265	0.538	0.539	0.553	0.439	0.448	0.443	0.445
sf_l_c	0.670	0.612	0.473	0.487	0.614	0.299	0.568	0.366	0.298	0.184	0.152	0.189	0.536	0.535	0.580	0.447	0.438	0.446	0.442
sf_c	0.630	0.635	0.478	0.486	0.573	0.328	0.549	0.417	0.289	0.117	0.148	0.248	0.541	0.526	0.588	0.357	0.432	0.440	0.436
l_g_c	0.596	0.629	0.518	0.517	0.521	0.281	0.544	0.455	0.308	0.176	0.174	0.138	0.473	0.478	0.515	0.458	0.424	0.426	0.425
l	0.649	0.607	0.519	0.496	0.697	0.292	0.520	0.307	0.298	0.096	0.129	0.186	0.486	0.455	0.553	0.332	0.414	0.411	0.413
c	0.653	0.612	0.558	0.449	0.565	0.299	0.583	0.316	0.344	0.119	0.248	-0.243	0.471	0.442	0.571	0.507	0.406	0.402	0.404
sf_l	0.539	0.585	0.483	0.423	0.545	0.251	0.501	0.382	0.258	0.159	0.138	0.169	0.510	0.509	0.519	0.363	0.396	0.400	0.398
sf_g_c	0.532	0.620	0.521	0.456	0.469	0.280	0.542	0.406	0.305	0.141	0.146	0.156	0.452	0.447	0.476	0.266	0.388	0.399	0.394
sf	0.515	0.582	0.475	0.439	0.440	0.282	0.529	0.384	0.303	0.130	0.153	0.121	0.472	0.472	0.490	0.363	0.384	0.380	0.382
g_c	0.446	0.584	0.490	0.405	0.480	0.275	0.571	0.323	0.350	0.201	0.267	-0.228	0.420	0.409	0.509	0.508	0.376	0.384	0.380
sf_l_g	0.510	0.564	0.494	0.401	0.327	0.237	0.446	0.364	0.249	0.144	0.110	0.124	0.429	0.462	0.462	0.318	0.353	0.360	0.356
l_g	0.536	0.552	0.535	0.430	0.523	0.219	0.454	0.228	0.251	0.095	0.119	0.001	0.364	0.420	0.385	0.385	0.344	0.346	0.345
sf_g	0.502	0.524	0.460	0.395	0.345	0.197	0.436	0.239	0.231	0.091	0.102	0.109	0.400	0.421	0.464	0.332	0.328	0.329	0.328
t	0.200	0.332	0.323	0.220	0.483	0.218	0.051	0.337	0.164	0.180	0.172	0.107	0.155	0.101	0.187	0.308	0.221	0.279	0.247

of these variations in a base form at the lemma level and their link to an explicit concept increases performance in the prediction of word relatedness. Jointly learned embeddings for surface forms, lemmas, and concepts achieve the best overall similarity results.

The analysis in [12] evaluates, using the same experimental settings, Vecisgrafo embeddings for lemmas and concepts learned from a 2018 dump of Wikipedia, achieving a Spearman’s rho of 0.566. This contrasts with the 0.448 achieved by the same embeddings learned from the SciGraph corpus. These results may not be directly comparable due to the difference in size between the two corpora (3B tokens in Wikipedia vs. 700M in SciGraph). Another factor is that Wikipedia is a general-purpose resource, while SciGraph is specialized in the scientific domain.

4.3. Analogical reasoning

For the analogy task, we use the Google analogy test set³, which contains 19,544 question pairs (8,869 semantic questions and 10,675 syntactic questions) and 14 types of relations (9 syntactic and 5 semantic). The accuracy of the embeddings in the analogy task is reported in Table 7.

Along the same lines as the word similarity and relatedness results, in this task the linguistic annotations also generate better embeddings than space-separated tokens. However, surface forms are more relevant than lemmas in this case. Our results indicate that surface

Table 7: Results of the analogy tasks sorted in descending order on the overall results.

Embeddings	Analogy Task		
	Semantic	Syntactic	Overall
sf_c	11.72%	14.17%	12.94%
sf_g_c	14.13%	11.65%	12.89%
sf_l_c	9.27%	12.10%	10.69%
l_g_c	14.83%	3.09%	8.96%
sf_l	7.18%	8.63%	7.91%
l_c	10.47%	4.06%	7.26%
sf	6.69%	6.43%	6.56%
sf_l_g	6.60%	5.74%	6.17%
sf_g	6.07%	3.79%	4.93%
l	4.07%	1.27%	2.67%
l_g	2.33%	0.79%	1.56%
t	0.64%	0.09%	0.37%

form embeddings achieve a positive result in both semantic and syntactic analogies while the performance of lemma embeddings is high for semantic analogies and really low in the case of syntactic analogies. In the latter case, the embeddings generated taking into account the morphological variations of words produce better results than those derived from the base forms. Indeed, some syntactic analogies actually require taking into account the information contained in surface forms. For example, capturing analogies like "bright/brighter" or "cold/colder" needs embedding representations for each individual term, whereas in a lemma-based approach only the terms "cold" and "bright" would be represented.

The highest accuracy is achieved by jointly learning *sf* and *c* embeddings. In fact, these two linguistic annotations are among the top 3 results along with *g* or *l*. As in word similarity, semantic embeddings *c* improve every

³Analogy dataset available at: [https://aclweb.org/aclwiki/Google_analogy_test_set_\(State_of_the_art\)](https://aclweb.org/aclwiki/Google_analogy_test_set_(State_of_the_art))

combination in which they are used. If we focus on the semantic analogies, we can see that l , g and c reach the highest accuracy. Nevertheless, the performance of this combination is very poor on syntactic analogies given that l do not include morphological variations of words that are heavily represented in the syntactic analogies. In general, the worst results are obtained when the combinations do not include c .

In Shazeer et al. [3], Swivel embeddings learned from Wikipedia and Gigaword achieved an analogy accuracy in the same dataset used in these experiments of 0.739, while the best result reported in our analysis is 0.129. As in the similarity and relatedness evaluation, the smaller size and domain specificity of the SciGraph corpus seem to have a negative impact on the results obtained in these generic evaluation tasks.

4.4. Word prediction

We select 17,500 papers from Semantic Scholar, which are not in our SciGraph corpus, as the unseen text on which we try to predict the embeddings for certain words (or linguistic annotations) according to the context. We applied the same annotation pipeline to the test corpus so that we can use the embeddings learned from SciGraph. As baselines we used pre-trained embeddings learned for linguistic annotations that we derived from *generic*, i.e. non-scientific, corpora: a dump of Wikipedia from January 2018 containing 2.89B tokens and UMBC [54], a web-based corpus of 2.95B tokens.

The overall results are shown in Table 8. We also plot the full results for some of the prediction tasks in Figure 2. The results show a clear pattern: embeddings learned from linguistic annotations significantly outperform the plain space-separated token embeddings (t). That is, the cosine similarity between the predicted embedding for the linguistic annotation on Semantic Scholar and the actual embeddings learned from SciGraph is higher. Recall that the predicted embedding is calculated by averaging the embeddings of the words or their annotations in the context window.

In general, embeddings learned from linguistic annotations on SciGraph are better at predicting embeddings in the Semantic Scholar corpus. In this case, it was easier to predict surface form embeddings, followed by lemma and concept embeddings. We assume this is because sf annotations also contain morphological information. Similarly, concept information may be more difficult to predict due to possible disambiguation errors or because the specific lemma used to refer to a concept may still provide useful information for the prediction task.

Jointly learning embeddings for sf and other annotations (c or l) produces sf embeddings which are slightly

Table 8: Prediction results of embeddings for linguistic annotations.

Corpus	Embed.	pred	Cosine Sim.	#tokens	Out-of-voc oov	oov %
SciGraph	sf	sf	0.765	3,841,378	617,094	16.064
SciGraph	sf_c	sf	0.751	3,841,378	617,144	16.066
SciGraph	sf_l	sf	0.741	3,841,378	617,121	16.065
SciGraph	l_c	l	0.733	2,418,362	81,087	3.353
SciGraph	l	l	0.730	2,418,362	81,041	3.351
SciGraph	l_c	c	0.690	2,012,471	1,481	0.074
SciGraph	c	c	0.690	2,012,471	1,572	0.078
Wiki	l_c	l	0.676	2,418,362	116,056	4.799
UMBC	l_c	l	0.674	2,418,362	93,169	3.853
UMBC	l_c	c	0.650	2,012,471	1,102	0.055
Wiki	l_c	c	0.650	2,012,471	2,216	0.110
SciGraph	t	t	0.576	3,396,730	205,154	6.040

harder to predict than when trained on their own. However, jointly learning l and c embeddings produces better results; i.e., jointly learning lemmas and concepts has a synergistic effect.

When comparing to the baselines, we see that SciGraph-based embeddings outperform both the Wikipedia and UMBC-based embeddings. The out-of-vocabulary (oov) numbers provide an explanation: both baselines produce embeddings which miss many terms in the test corpus, but which are included in the SciGraph embeddings. Wikipedia misses 116K lemmas and UMBC misses 93K, but SciGraph only misses 81K of the lemmas in the test corpus. Wikipedia misses the highest number of concepts (2.2K), followed by SciGraph (1.5K) and UMBC (1.1K); however, despite missing more concepts than UMBC, SciGraph-based embeddings outperform both baselines.

Manual inspection of the missing words shows that UMBC is missing nearly 14K lemmas (mostly for scientific terms) which the SciGraph-based embeddings contain, such as "negative bias temperature instability", "supragranular", "QHD". Inversely, the SciGraph vocabulary is missing nearly 7K lemmas for non-scientific entities, such as "Jim Collins" and "Anderson School of Management", but also for misspelled words. However, most of the missing words (around 42K) are not found in either UMBC or SciGraph and include very specific metrics such as "40 kilopascal", acronyms ("DB620") and named entities ("Victoria Institute of Forensic Medicine"). We observe a similar pattern when comparing missing concepts.

4.5. Classification of scientific documents

Leveraging SciGraph metadata, we formulate a multi-label classification task that aims at predicting one or more of 22 first level categories such as Mathematical Sciences, Engineering or Medical and Health Sciences.

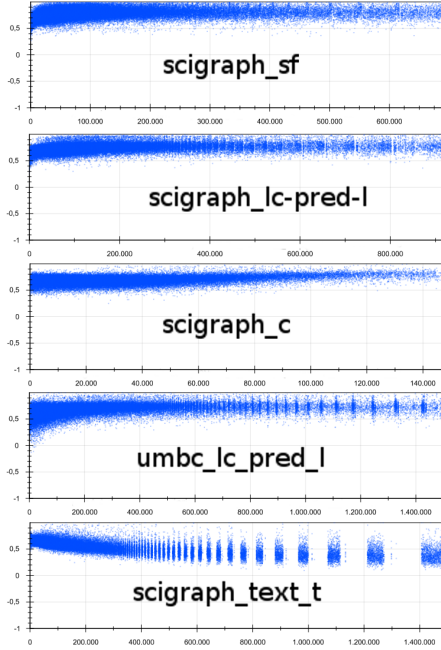


Figure 2: Embedding prediction plots. The horizontal axis aligns with the terms in the vocabulary sorted by frequency. The vertical axis is the average cosine similarity (which can range from 0 to 1) between the predicted embedding on the Semantic Scholar corpus and the actual learned embedding from SciGraph. According to this task the closest the curve draw by the plotted dots to 1 the better. Blank spaces indicates lack of coverage

To train the classifier, we use a vanilla CNN implementation available in Keras. Such neural network consists of 3 convolutional layers with 128 filters and a 5-element window size, each followed by a max pooling layer, a fully-connected 128-unit ReLU layer and a sigmoid output. Our goal is not to design the best classifier, but to evaluate the different types of embeddings through its performance.

To evaluate the classifiers we select articles published in 2011 and use ten-fold cross-validation. We define a vocabulary size of 20K elements, and max sequence size of 1000. As baseline, we train a classifier that learns from embeddings generated randomly, following a normal distribution. As upper bound we learn a classifier that is able to optimize the embeddings during task-specific training. In addition, we train classifiers using embeddings generated for t , sf , l and c .

The performance of the classifiers measured in terms of precision, recall and f-measure is shown in Table 9. l embeddings contribute to learn the best classifier, followed by sf and t embeddings. If we focus on f-measure we can see again that grammatical analysis (l , sf) performs better than the lexical level (t). However, the

Table 9: Learning results using token-based embeddings generated randomly (baseline) and learned by the convolutional neural network (upper bound), and Swivel embeddings for t , sf , l and c .

Embed. Algorithm	Types	P	R	F-measure
Random Normal	t	0,7596	0,6775	0,7015 Δ
Learned by CNN	t	0,8062	0,767	0,7806 ∇
Swivel	t	0,8008	0,7491	0,7679
Vecsigrafo	sf	0,8030	0,7477	0,7684
Vecsigrafo	l	0,8035	0,7539	0,7728
Vecsigrafo	c	0,7945	0,7391	0,7583

performance of c embeddings for this task is rather low. The low number of c embeddings (see Table 4) seems to affect negatively the learning process. On the other hand, all the classifiers, including the one trained with c , improve the baseline of random embeddings. However, none of the classifiers reached the performance of the upper bound.

Next, we train classifiers using Vecsigrafo embeddings from all the 10 possible combinations of the linguistic annotations (sf , l , g , c) with size 2 and 3. The classifiers are trained using the merging techniques described in Section 3.2. We combine single embeddings per each linguistic annotation type and merge them using vector operators like average, concatenate, and PCA. The average performance of the classifiers is presented in Table 10. Each column is sorted according to the reported evaluation metric.

Table 10: Average performance of the 70 classifiers trained using different strategies to mix and merge Vecsigrafo embeddings. Each column is sorted in descending order of the corresponding metric.

Precision		Recall		F-Measure	
Types	Average \downarrow	Types	Average \downarrow	Types	Average \downarrow
sf_l_g	0,8126	sf_l	0,7605	sf_l	0,7776
sf_g	0,8091	l_c	0,7556	sf_l_g	0,7750
l_g	0,8090	sf_l_c	0,7538	l_g_c	0,7730
sf_g_c	0,8084	l_g_c	0,7524	l_c	0,7728
l_g_c	0,8073	sf_l_g	0,7524	sf_l_c	0,7720
g_c	0,8060	sf_c	0,7520	sf_g_c	0,7716
sf_l	0,8056	sf_g_c	0,7504	l_g	0,7702
sf_l_c	0,8031	l_g	0,7477	sf_c	0,7699
l_c	0,8021	sf_g	0,7374	sf_g	0,7638
sf_c	0,8006	g_c	0,7338	g_c	0,7599

Regarding precision, the top 3 classifiers are learned from different combinations of sf , l and g , indicating that, precision-wise, grammatical information is more relevant than semantic information (c). However, note that the common linguistic element in the 6 first classifiers is g , even when combining it with c . In general, removing g embeddings produced the least precise classifiers. This means that grammar information is what makes the difference, regardless of the other linguistic elements. That

said, the influence of g is enhanced when used in combination with sf and l . Note that the precision of the top 5 classifiers is better than the upper bound classifier (Table 9), where the embeddings were learned at classification training time, although the linguistic-based embeddings were not learned for this specific purpose.

The recall analysis shows a different picture since grammar embeddings g do not seem to have a decisive role on the performance of the classifier recall-wise, while c gain more relevance. sf and l help in learning the best classifier. The combination of c and l seems to benefit recall, as seen in 3 of the top 4 classifiers. In contrast when concepts are combined with sf the recall is lower. The fact that l embeddings are more directly related to c than sf seems to make a difference in the recall analysis when these three elements are involved. In general, g -based embedding combinations generate classifiers with lower recall.

Finally, the f-measure data shows more heterogeneous results since by definition it is the harmonic mean of precision and recall, and hence the combinations that generate the best f-measure also have a high precision and a high recall. The combination of sf and l is at the top (best recall), followed by their combination with g (best precision). c appears in positions 3 to 6 in the ranking; however, when combined solely with sf or g , the f-measure results are the worst. When combined with at least two other elements, g ranks high in f-measure, whereas when g is combined with a single linguistic annotation type, the resulting performance of the classifiers is the worst.

5. Discussion

Although the different evaluation tasks assess the embeddings from different perspectives, in all of them we have observed that the embeddings learned from linguistic annotations outperform space-separated token embeddings. In the intrinsic evaluation tasks, the embeddings learned from linguistic annotations have a different performance, mainly due to the vocabulary used in the evaluation datasets. For example, in the analogical reasoning datasets, syntactic analogies are more represented compared to semantic analogies. Also, most of the syntactic analogies often include morphological variations of the words, which are better covered by surface form embeddings than lemma embeddings, where morphological variations are conflated in a single base form. For example, comparative and superlative analogies like *cheap-cheaper* and *cold-colder* may require morphological variations of the adjective by adding the *-er* and *-(e)st* suffixes. Furthermore, although some adjectives

and adverbs may produce irregular forms, the only ones included in the dataset are *good-better*, *bad-worse*, *good-best* and *bad-worst*.

In contrast, in the similarity and relatedness datasets (see Table 5) most of the word pairs available for evaluation are either nouns, named entities or lemmatized verbs—only the Verb-143 dataset contains conjugated verbs—. Therefore, most of the word pairs are non-inflected word forms. As a result, concept and lemma embeddings are the most suitable representations for this task, which is in line with the evaluation results reported in Table 6.

Remarkably, in the analogy task, surface form embeddings jointly learned with concept embeddings achieve the highest performance. In the similarity and relatedness task, jointly learning concepts and lemma embeddings also improves the lemma embeddings performance in the task. Therefore, including concept embeddings in the learning process of other embeddings generally helps to learn better distributional representations.

As additional insight gained from the word prediction task, we observe that the embeddings learned from linguistic features seem to better capture the distributional context, as evaluated against unseen text. Surface form and lemma embeddings were the easiest to predict.

Finally, in the text classification task, single embeddings from lemma and surface form annotations were more useful than space-separated token embeddings. Nevertheless, in this task, concept embeddings perform worse mainly due to the low coverage offered by such embeddings with respect to the whole vocabulary. This indicates that the knowledge graph used to annotate the text possibly provides limited coverage of the domain. We also show that jointly learning lemma and surface form embeddings helps to train the best classifier in terms of f-measure and recall. Furthermore, adding grammar embeddings produced the best overall precision.

6. Conclusions and future work

In this paper, we present a novel experimental study on the use of linguistic annotations for learning embeddings from a scientific corpus. Our experiments show in all cases that embeddings learned from linguistic annotations perform better than conventional space-separated tokens in all evaluation tasks. Moreover, we demonstrate that the vocabulary, and its grammatical characteristics, used in each evaluation task is directly related to the quality of the linguistic annotated embeddings.

Based on the results of our study, we can conclude that there is not a single type of linguistic annotation that always produces the best results consistently. However,

we identify some patterns that may help practitioners to select the most suitable combinations for each specific task and dataset.

In the similarity task, learning embeddings based on lemmas generally works better given the high rate of nouns, named entities and multi-word expressions contained in those datasets. On the contrary, in the analogy dataset there is a majority of syntactic analogies containing morphological variations of words that were better represented by single surface form embeddings. Jointly learning embeddings for concepts and either lemmas or surface forms improve the quality of each individual type of embeddings across all the evaluation tasks.

Regarding the extrinsic evaluations, in the word prediction task the embeddings learned from linguistic annotations provide evidence to better capture the distributional context of words in an unseen text than space-separated embeddings. In the classification tasks, jointly learned lemma and surface form embeddings help to train the best classifier and, if grammar embeddings are also used, then the highest precision is achieved.

As future work, we want to understand the type of linguistic information that is being encoded in the Vecsgrafo embeddings beyond the intrinsic evaluation tasks that we used in this paper. Recently, a set of probes has been released to identify whether the models encode different types of linguistic information [16, 15], including parse trees, constituent labelling and co-reference to name a few, that can be applied to Vecsgrafo. Another research path we want to explore is to use linguistic annotations to enhance neural language models like BERT and assess whether such linguistic annotations can also help to learn better language models, improving performance when fine-tuned for various downstream tasks. The idea is to extend the current transformer-based architectures used to learn language models, so that not only word pieces, but also other types of linguistic annotations, such as those discussed in this paper, can be ingested and their effect propagated across the model.

Acknowledgements

We gratefully acknowledge the EU Horizon 2020 research and innovation programme under grant agreement No. 825627 (ELG).

References

- [1] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, *CoRR abs/1301.3781* (2013). [arXiv:1301.3781](https://arxiv.org/abs/1301.3781)
URL <http://arxiv.org/abs/1301.3781>
- [2] J. Pennington, R. Socher, C. D. Manning, Glove: Global vectors for word representation., in: *EMNLP*, Vol. 14, 2014, pp. 1532–1543 (2014).
- [3] N. Shazeer, R. Doherty, C. Evans, C. Waterson, Swivel: Improving Embeddings by Noticing What’s Missing, *arXiv preprint* (2016). [arXiv:1602.02215](https://arxiv.org/abs/1602.02215).
- [4] P. Bojanowski, E. Grave, A. Joulin, T. Mikolov, Enriching word vectors with subword information, *arXiv preprint* [arXiv:1607.04606](https://arxiv.org/abs/1607.04606) (2016).
- [5] M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, L. Zettlemoyer, Deep contextualized word representations, in: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, Association for Computational Linguistics, 2018, pp. 2227–2237 (2018). doi:10.18653/v1/N18-1202.
URL <http://aclweb.org/anthology/N18-1202>
- [6] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint* [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018).
- [7] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, Q. V. Le, XLNet: Generalized Autoregressive Pretraining for Language Understanding (jun 2019). [arXiv:1906.08237](https://arxiv.org/abs/1906.08237).
URL <http://arxiv.org/abs/1906.08237>
- [8] B. Heinzerling, M. Strube, BPEmb: Tokenization-free pre-trained subword embeddings in 275 languages, in: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, European Language Resources Association (ELRA), Miyazaki, Japan, 2018 (May 2018).
URL <https://www.aclweb.org/anthology/L18-1473>
- [9] T. Hammond, M. Pasin, E. Theodoridis, Data integration and disintegration: Managing springer nature scigraph with shacl and owl., in: N. Nikitina, D. Song, A. Fokoue, P. Haase (Eds.), *International Semantic Web Conference (Posters, Demos and Industry Tracks)*, Vol. 1963 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2017 (2017).
URL <http://dblp.uni-trier.de/db/conf/semweb/iswc2017p.html#HammondPT17>
- [10] W. Ammar, D. Groeneveld, C. Bhagavatula, I. Beltagy, M. Crawford, D. Downey, J. Dunkelberger, A. Elgohary, S. Feldman, V. Ha, R. Kinney, S. Kohlmeier, K. Lo, T. Murray, H. Ooi, M. Peters, J. Power, S. Skjonsberg, L. Wang, C. Wilhelm, Z. Yuan, M. van Zuylen, O. Etzioni, Construction of the literature graph in semantic scholar, in: *NAACL-HTL*, 2018 (2018).
- [11] L. Han, A. L. Kashyap, T. Finin, J. Mayfield, J. Weese, UMBC EBIQUITY-CORE: Semantic Textual Similarity Systems, in: *Proceedings of the Second Joint Conference on Lexical and Computational Semantics*, Association for Computational Linguistics, 2013 (June 2013).
- [12] R. Denaux, J. M. Gómez-Pérez, Vecsgrafo: Corpus-based word-concept embeddings, *Semantic Web* 10 (5) (2019) 881–908 (2019). doi:10.3233/SW-190361.
URL <https://doi.org/10.3233/SW-190361>
- [13] O. Levy, Y. Goldberg, Neural word embedding as implicit matrix factorization, in: *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’14, MIT Press, Cambridge, MA, USA, 2014, pp. 2177–2185 (2014).
URL <http://dl.acm.org/citation.cfm?id=2969033.2969070>
- [14] G. Jawahar, B. Sagot, D. Seddah, What does BERT learn about the structure of language?, in: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, 2019, pp.

- 3651–3657 (Jul. 2019). doi:10.18653/v1/P19-1356.
URL <https://www.aclweb.org/anthology/P19-1356>
- [15] I. Tenney, D. Das, E. Pavlick, BERT rediscovers the classical NLP pipeline, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 4593–4601 (Jul. 2019). doi:10.18653/v1/P19-1452.
URL <https://www.aclweb.org/anthology/P19-1452>
- [16] J. Hewitt, C. D. Manning, A structural probe for finding syntax in word representations, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4129–4138 (Jun. 2019). doi:10.18653/v1/N19-1419.
URL <https://www.aclweb.org/anthology/N19-1419>
- [17] K. Clark, U. Khandelwal, O. Levy, C. D. Manning, What does bert look at? an analysis of bert’s attention, ArXiv abs/1906.04341 (2019).
- [18] G. A. Miller, Wordnet: A lexical database for english, Commun. ACM 38 (11) (1995) 39–41 (Nov. 1995). doi:10.1145/219717.219748.
URL <http://doi.acm.org/10.1145/219717.219748>
- [19] A. Bordes, N. Usunier, J. Weston, O. Yakhnenko, Translating Embeddings for Modeling Multi-Relational Data, Advances in NIPS 26 (2013) 2787–2795 (2013). arXiv:arXiv:1011.1669v3, doi:10.1007/s13398-014-0173-7.2.
- [20] M. Nickel, L. Rosasco, T. A. Poggio, Holographic embeddings of knowledge graphs, in: AAAI, 2016 (2016).
- [21] P. Ristoski, H. Paulheim, RDF2Vec: RDF graph embeddings for data mining, in: International Semantic Web Conference, Vol. 9981 LNCS, 2016, pp. 498–514 (2016). doi:10.1007/978-3-319-46523-4_30.
- [22] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. van den Berg, I. Titov, M. Welling, Modeling Relational Data with Graph Convolutional Networks (2018). arXiv:1703.06103.
- [23] S. Cavallari, E. Cambria, H. Cai, K. C. C. Chang, V. W. Zheng, Embedding Both Finite and Infinite Communities on Graphs [Application Notes], IEEE Computational Intelligence Magazine 14 (3) (2019) 39–50 (2019). doi:10.1109/MCI.2019.2919396.
- [24] S. Ji, S. Pan, E. Cambria, P. Marttinen, P. S. Yu, A Survey on Knowledge Graphs: Representation, Acquisition and Applications (feb 2020). arXiv:2002.00388.
URL <http://arxiv.org/abs/2002.00388>
- [25] J. Camacho-Collados, M. T. Pilehvar, R. Navigli, NASARI: Integrating explicit knowledge and corpus statistics for a multilingual representation of concepts and entities, Artificial Intelligence 240 (2016) 36–64 (2016). doi:10.1016/j.artint.2016.07.005.
- [26] M. Mancini, J. Camacho-Collados, I. Iacobacci, R. Navigli, Embedding words and senses together via joint knowledge-enhanced training, in: CoNLL, 2017 (2017).
- [27] S. Thoma, A. Rettinger, F. Both, Towards holistic concept representations: Embedding relational knowledge, visual attributes, and distributional word semantics, in: International Semantic Web Conference, 2017 (2017).
- [28] T. Schnabel, I. Labutov, D. Mimno, T. Joachims, Evaluation methods for unsupervised word embeddings, in: EMNLP, ACL, 2015, pp. 298–307 (2015).
- [29] J. J. Lastra-Díaz, J. Goikoetxea, M. A. H. Taieb], A. García-Serrano, M. B. Aouicha], E. Agirre, A reproducible survey on word embeddings and ontology-based methods for word similarity: Linear combinations outperform the state of the art, Engineering Applications of Artificial Intelligence 85 (2019) 645 – 665 (2019). doi:https://doi.org/10.1016/j.engappai.2019.07.010.
URL <http://www.sciencedirect.com/science/article/pii/S0952197619301745>
- [30] P. Resnik, Using information content to evaluate semantic similarity in a taxonomy, in: Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 1, IJCAT’95, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995, pp. 448–453 (1995).
URL <http://dl.acm.org/citation.cfm?id=1625855.1625914>
- [31] S. Hassan, R. Mihalcea, Semantic relatedness using salient semantic analysis, in: Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI’11, AAAI Press, 2011, p. 884–889 (2011).
- [32] J. Camacho-Collados, M. T. Pilehvar, R. Navigli, NASARI: a novel approach to a semantically-aware representation of items, in: Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Denver, Colorado, 2015, pp. 567–577 (May–Jun. 2015). doi:10.3115/v1/N15-1059.
URL <https://www.aclweb.org/anthology/N15-1059>
- [33] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, J. Dean, Distributed Representations of Words and Phrases and their Compositionality., in: NIPS, 2013 (2013). arXiv:1310.4546, doi:10.1162/jmlr.2003.3.4-5.951.
- [34] C. D. Manning, P. Raghavan, H. Schütze, Introduction to Information Retrieval, Cambridge University Press, New York, NY, USA, 2008 (2008).
- [35] S. Wang, C. D. Manning, Baselines and bigrams: Simple, good sentiment and topic classification, in: Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers - Volume 2, ACL ’12, Association for Computational Linguistics, Stroudsburg, PA, USA, 2012, pp. 90–94 (2012).
URL <http://dl.acm.org/citation.cfm?id=2390665.2390688>
- [36] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, Proceedings of the IEEE 86 (11) (1998) 2278–2324 (1998).
- [37] R. Collobert, J. Weston, A unified architecture for natural language processing: Deep neural networks with multitask learning, in: Proceedings of the 25th International Conference on Machine Learning, ICML ’08, ACM, New York, NY, USA, 2008, pp. 160–167 (2008). doi:10.1145/1390156.1390177.
URL <http://doi.acm.org/10.1145/1390156.1390177>
- [38] Y. Kim, Convolutional neural networks for sentence classification, in: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25–29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL, 2014, pp. 1746–1751 (2014).
URL <http://aclweb.org/anthology/D/D14/D14-1181.pdf>
- [39] C. dos Santos, M. Gatti, Deep convolutional neural networks for sentiment analysis of short texts, in: Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers, 2014, pp. 69–78 (2014).
- [40] X. Zhang, J. Zhao, Y. LeCun, Character-level convolutional

- networks for text classification, in: *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS'15*, MIT Press, Cambridge, MA, USA, 2015, pp. 649–657 (2015).
URL <http://dl.acm.org/citation.cfm?id=2969239.2969312>
- [41] H. Rubenstein, J. B. Goodenough, Contextual correlates of synonymy, *Communications of the ACM* 8 (10) (1965) 627–633 (1965). doi:10.1145/365628.365657.
 - [42] G. A. Miller, W. G. Charles, Contextual Correlates of Semantic Similarity, *Language and Cognitive Processes* 6 (1) (1991) 1–28 (jan 1991). doi:10.1080/01690969108406936.
 - [43] L. Finkelstein, E. Gabrilovich, Y. Matias, E. Rivlin, Z. Solan, G. Wolfman, E. Ruppín, Placing search in context: the concept revisited, *ACM Transactions on Information Systems* 20 (1) (2002) 116–131 (2002). doi:10.1145/503104.503110.
 - [44] E. Agirre, E. Alfonseca, K. Hall, J. Kravalova, M. Pas, A. Soroa, A Study on Similarity and Relatedness Using Distributional and WordNet-based Approaches, *Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the ACL (June)* (2009) 19–27 (2009). doi:10.3115/1620754.1620758.
 - [45] D. Yang, D. M. W. Powers, Verb Similarity on the Taxonomy of WordNet, *3rd International WordNet Conference* (2006) 121–128 (2006).
 - [46] S. Baker, R. Reichart, A. Korhonen, An Unsupervised Model for Instance Level Subcategorization Acquisition, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)* (2014) 278–289 (2014).
 - [47] D. Gerz, I. Vulic, F. Hill, R. Reichart, A. Korhonen, Simverb-3500: A large-scale evaluation set of verb similarity, in: *EMNLP, 2016* (2016).
 - [48] K. Radinsky, E. Agichtein, E. Gabrilovich, S. Markovitch, A word at a time, in: *Proceedings of the 20th international conference on World wide web - WWW '11*, ACM Press, New York, New York, USA, 2011, p. 337 (2011). doi:10.1145/1963405.1963455.
 - [49] P. Koehn, Europarl: A Parallel Corpus for Statistical Machine Translation, in: *Conference Proceedings: the tenth Machine Translation Summit, AAMT, AAMT, Phuket, Thailand, 2005*, pp. 79–86 (2005).
URL <http://mt-archive.info/MTS-2005-Koehn.pdf>
 - [50] E. Bruni, G. Boleda, M. Baroni, N.-K. Tran, Distributional semantics in technicolor, *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics 1 (July)* (2012) 136–145 (2012). doi:10.1109/ICRA.2016.7487801.
 - [51] F. Hill, R. Reichart, A. Korhonen, Simlex-999: Evaluating semantic models with (genuine) similarity estimation, *Computational Linguistics* 41 (2015) 665–695 (2015).
 - [52] M.-T. Luong, R. Socher, C. D. Manning, Better Word Representations with Recursive Neural Networks for Morphology, *CoNLL-2013* (2013) 104–113 (2013).
 - [53] J. Camacho-Collados, M. T. Pilehvar, N. Collier, R. Navigli, Semeval-2017 task 2: Multilingual and cross-lingual semantic word similarity, in: *SemEval@ACL, 2017* (2017).
 - [54] L. Han, A. L. Kashyap, T. W. Finin, J. Mayfield, J. Weese, Umbc ebiquity core: Semantic textual similarity systems, in: **SEM@NAACL-HLT, 2013* (2013).