

Taxonomy Enrichment with Text and Graph Vector Representations

Irina Nikishina^{a,*}, Mikhail Tikhomirov^b, Varvara Logacheva^a, Yuriy Nazarov, Alexander Panchenko^a, and Natalia Loukachevitch^b

^a *Skolkovo Institute of Science and Technology, Moscow, Russia*

E-mails: irina.nikishina@skoltech.ru, v.logacheva@skoltech.ru, a.panchenko@skoltech.ru

^b *Research Computing Center, Lomonosov Moscow State University, Moscow, Russia*

E-mails: tikhomirov.mm@gmail.com, nazarov.yuriy.pavlovich@gmail.com, louk_nat@mail.ru

Editors: Mehwish Alam, FIZ Karlsruhe - Leibniz Institute for Information Infrastructure, Germany; Davide Buscaldi, LIPN, Université Sorbonne Paris Nord, France; Michael Cochez, Vrije University of Amsterdam, the Netherlands; Francesco Osborne, Knowledge Media Institute, (KMi), and The Open University, UK; Diego Reforgiato Recupero, University of Cagliari, Italy; Harald Sack, FIZ Karlsruhe - Leibniz Institute for Information Infrastructure, Germany

Solicited review: Six anonymous reviewers

Abstract.

Knowledge graphs such as DBpedia, Freebase or Wikidata always contain a taxonomic backbone that allows the arrangement and structuring of various concepts in accordance with hypo-hypernym (“class-subclass”) relationship. With the rapid growth of lexical resources for specific domains, the problem of automatic extension of the existing knowledge bases with new words is becoming more and more widespread. In this paper, we address the problem of *taxonomy enrichment* which aims at adding new words to the existing taxonomy.

We present a new method which allows achieving high results on this task with little effort. It uses the resources which exist for the majority of languages, making the method universal. We extend our method by incorporating deep representations of graph structures like node2vec, Poincaré embeddings, GCN etc. that have recently demonstrated promising results on various NLP tasks. Furthermore, combining these representations with word embeddings allows us to beat the state of the art.

We conduct a comprehensive study of the existing approaches to taxonomy enrichment based on word and graph vector representations and their fusion approaches. We also explore the ways of using deep learning architectures to extend taxonomic backbones of knowledge graphs. We create a number of datasets for taxonomy extension for English and Russian. We achieve state-of-the-art results across different datasets and provide an in-depth error analysis of mistakes.

Keywords: taxonomy enrichment, graph vector representations, word embeddings, graph convolutional auto-encoder

1. Introduction

The central idea of Semantic Web is to make the content of the Internet pages machine-interpretable. For that, web pages should be linked to *ontologies* [1, 2] — databases which contain the information on classes of objects, their properties, and relations be-

tween the classes. The relations between objects are particularly important in an ontology, because they form its structure. They can be of different types corresponding to different types of relationships between real-world objects. One of the most important relationships is the class-subclass relation. It allows organising entities into a *taxonomy* — a tree structure where entities are represented as nodes and the edges between them denote *subclass-of* or *instance-of* relationship. The class-subclass relations and taxonomies built from

*Equal contribution with Mikhail Tikhomirov. Corresponding author: E-mail: irina.nikishina@skoltech.ru.

1 them are crucial for understanding the place and the
2 purpose of an object or a concept in the world. This
3 relation and the hierarchical structure created by it are
4 also a basis of many knowledge bases and *knowledge*
5 *graphs* — a particular type of knowledge bases where
6 objects are organised in a graph structure. There, the
7 nodes of a graph are objects, and the edges of a graph
8 are relations between the objects.

9 To support the use of specific knowledge base
10 structures, such as thesauri or taxonomies, there exist
11 specifications and standards classification schemes.
12 For instance, Simple Knowledge Organization Sys-
13 tems (SKOS)¹ define several types of lexical-semantic
14 relations in the terms of semantic web, e.g. “has
15 broader/narrower”, “has exact match”, and “is in map-
16 ping relation with”. SKOS are used to create structures
17 like thesauri.

18 However, SKOS does not fully comply with our tax-
19 onomy. The class-subclass relation which is the base
20 of taxonomies is a relation between objects X and Y
21 such that the sentence “X is a kind of Y” is accept-
22 able for native speakers. On the other hand, the closest
23 SKOS analogue of taxonomic class-subclass relation
24 is the “broader term relation”. It is different from the
25 class-subclass relation, because it is less specific. For
26 example, it can include the part-whole relations.

27 The usefulness of an ontology or a knowledge base
28 depends largely on its completeness and its ability to
29 fully reflect the real world. However, since the world
30 is changing, the ontologies need to be constantly up-
31 dated to stay relevant. There currently exist compre-
32 hensive knowledge bases such as Freebase, DBPedia
33 or Wikidata as well as ontologies for specific domains.
34 Many areas of knowledge require their own knowledge
35 bases, and all of them need to be maintained and ex-
36 tended. This is an expensive and time-consuming pro-
37 cess which can only be conducted by an expert who is
38 proficient in the discipline and understands the struc-
39 ture of a knowledge base. Thus, in order to speed up
40 and simplify this task, it becomes more and more im-
41 portant to develop systems that could automatically en-
42 rich the existing knowledge bases with new words or at
43 least facilitate the manual extension process. The task
44 of automatically or semi-automatically adding new en-
45 tities to hierarchical structures is referred to as *taxon-*
46 *omy enrichment*.

47 In this work, we aim at reviewing the existing tax-
48 onomy enrichment models and propose new methods
49

1 which address their drawback. We also aim at evalu-
2 ating the scalability of different methods to new lan-
3 guages and datasets.

4 The state-of-the-art taxonomy enrichment methods
5 have two main drawbacks. First of all, they often
6 use unrealistic formulations of the task. For exam-
7 ple, SemEval-2016 task 14 [3] which was the first ef-
8 fort to evaluate this task in a controlled environment,
9 provided definitions of the query words (words to be
10 added to a taxonomy). This is very informative re-
11 source, so the majority of the presented methods heav-
12 ily depended on those definitions [4, 5]. However, in
13 the real-world scenarios, such information is usually
14 unavailable, which makes the developed methods in-
15 applicable. We tackle this problem by testing our new
16 methods and the state-of-the-art methods in a realistic
17 setting.

18 Another gap in the existing research is that the ma-
19 jority of methods use the information from only one
20 source. Namely, some researchers use the information
21 from distributional word embeddings, whereas others
22 consider graph-based models which represent a word
23 based on its position in a taxonomy. Our intuition is
24 that the information from these two sources is com-
25plementary, so combining them can improve the per-
26formance of taxonomy enrichment models. Therefore,
27 we propose a number of ways to incorporate various
28 sources of information.

29 First, we propose the new **DWRank** method which
30 uses only distributional information from pre-trained
31 word embeddings and is similar to other existing meth-
32 ods. We then enable this method to incorporate the
33 different sources of graph information. We compare
34 the various ways of getting the information from a
35 knowledge graph. Finally, another modification of our
36 method successfully combines the information from
37 different sources, beating the current state of the art.

38 To place our models in the context of the research
39 on taxonomy enrichment, we compare them with a
40 number of state-of-the-art models. To the best of our
41 knowledge, this is the first large-scale evaluation of
42 taxonomy enrichment methods. We are also the first to
43 evaluate the methods on datasets of different sizes and
44 in different languages.

45 This work is an extended version of the work de-
46 scribed in [6–8]. The novelty of this particular article
47 as compared to the previous publications is as follows:
48

- 49 1. We present a new taxonomy enrichment method
50 **DWRank** which combines distributional infor-
51

51 ¹<https://www.w3.org/2004/02/skos/intro>

1 information and the information extracted from Wik-
 2 tionary.

- 3 2. We present an extension of DWRank called
 4 **DWRank-Graph** which uses various graph-
 5 based representations via a common interface.
- 6 3. We present **DWRank-Meta** — an extension of
 7 DWRank which combines the information from
 8 different sources and beats the state-of-the-art
 9 models.
- 10 4. We present **WBSR** — a method for taxonomy ex-
 11 tension which leverages the information from the
 12 Web. This approach is the current state of the art
 13 in the task.
- 14 5. We conduct a large-scale computational study
 15 of various approaches to taxonomy enrichment,
 16 which features multiple methods (including ours
 17 as well as state-of-the-art approaches), multiple
 18 datasets and languages.
- 19 6. We present datasets for studying the evolution
 20 of wordnets for English and Russian, extending
 21 the monolingual setup of the RUSSE’2020 shared
 22 task [9] with a larger Russian dataset and similar
 23 English versions.
- 24 7. We explore the benefits of meta-embeddings
 25 (combinations of embeddings) and graph embed-
 26 dings for the task of taxonomy enrichment.
- 27 8. We provide an in-depth error analysis of different
 28 types of mistakes of the state-of-the-art models.
- 29 9. We provide mappings to WordNet Linked Open
 30 Data (LOD) Inter-Lingual Index (ILI) from Rus-
 31 sian to English synsets. A dataset of multilingual
 32 hypernyms opens possibilities for various use-
 33 cases for cross-lingual operations on taxonomies.

34 We release all our code and datasets to enable fur-
 35 ther research in this direction.²

36 The rest of the paper is organised as follows. In
 37 Section 2 we formulate the task and provide the rel-
 38 evant definitions. Section 3 is devoted to the existing
 39 work on taxonomy enrichment. Then, in Section 4 we
 40 present our datasets and describe their creation pro-
 41 cess and features. Sections 5, 6, and 7 contain the de-
 42 scriptions of our methods: in Section 5 we introduce
 43 our core method DWRank and in Sections 6 and 7 de-
 44 scribe its extensions DWRank-Graph and DWRank-
 45 Meta. We report the performance of our models and
 46 compare them with the other approaches in Section 8.
 47 Finally, in Section 9 we analyse the potential limita-
 48 tions of our methods.

50 ²<https://github.com/skoltech-nlp/diachronic-wordnets>

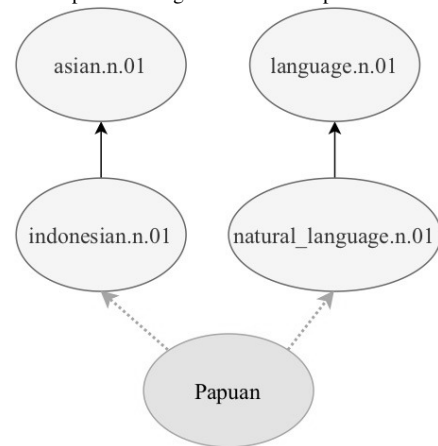
2. Task Formulation and Definitions1

1 In this subsection we formulate the task and explain
 2 the main concepts and task-related terms.

2.1. Task Formulation2

3 Let us consider an example of taxonomy. Figure 1
 4 demonstrates a subgraph for the word “Papuan” re-
 5 trieved from WordNet. There, each concept has one
 6 or more *parents* (concepts which it is derived from).
 7 The parents in their turn have their own parents, and
 8 one can trace the word affiliation until the most ab-
 9 stract *root* concepts of the taxonomy. Here, the word
 10 “Papuan” is attached to the synsets “indonesian” and
 11 “natural_language” as it can mean both an ethnicity
 12 and a language. Note that words can have multiple par-
 13 ents for different reasons. Two or more parents can
 14 point to the fact that a word has multiple meanings
 15 (as in the case with “Papuan”). Alternatively, a word
 16 meaning can be a combination of meanings of two
 17 higher-level concepts.
 18
 19
 20
 21
 22
 23

24 Figure 1. Example of adding a new word “Papuan” to the taxonomy



25 Therefore, in order to add a word to a taxonomy, we
 26 need to find its hypernym among the entities (synsets
 27 in case of wordnets) of this taxonomy. Here, we refer
 28 to a word absent from the taxonomy (a word which
 29 we would like to add) as a *query word*. Our task is to
 30 attach query words to an existing taxonomic tree.
 31
 32
 33
 34
 35
 36
 37
 38
 39

40 The task of finding a single suitable hypernym
 41 synset is difficult for a machine, because the number
 42 of nodes in the existing taxonomy can be very large
 43 (e.g. in WordNet the number of noun synsets is around
 44 29,000). Thus, a model trained to solve this task will
 45 inevitably return many false answers if asked to pro-
 46
 47
 48
 49
 50
 51

vide only one synset candidate. Moreover, as we can see from Figure 1, a word may have multiple hypernyms.

Thus, in the majority of works on taxonomy enrichment the requirement of providing a single correct answer is relaxed. Instead of that, a common approach is to provide k (typically 10 to 20) most suitable candidates. This list is more likely to contain correct synsets. This setting is also consistent with the manual computer-assisted annotation. Presenting an annotator with a small list of candidates will facilitate the taxonomy extension process: annotators will only have to look through a short list of high-probability hypernym candidates instead of searching hypernyms in a list of all synsets of the taxonomy. Thus, we formulate the task of *taxonomy extension* as a soft ranking problem, where the synsets are ranked according to their suitability for a given word.

This is an established approach to this task, the same formulation was also used in research works on taxonomy enrichment and in taxonomy enrichment shared tasks [3, 9].

2.2. Definitions

In this section, we provide several key definitions crucial for understanding of our work.

Knowledge base is a data structure used to store knowledge which comprise of concepts and relations (known as TBox-es) and individuals, concept, and relation instantiations (also known as ABox-es) [10].

Knowledge graph is a knowledge base that uses a graph-structured data model to represent data. Formally, a knowledge graph G is a set of triplets $\{h, r, t\} \subseteq E \times R \times E$ where E is the entity set and R is the relation set. Normally, knowledge graphs comprise multiple types of relations R .

Synset is a set $S \subseteq \{s_1, \dots, s_n\}$ of words or phrases corresponding to a concept or a knowledge base entity in the set E . Synset is the major element which denotes a node in a knowledge graph G , therefore, in our case S equals E .

Hypernymy relation is a relation $r \in R$. According to [11], hypernymy relation exists between objects X and Y if native speakers accept sentences constructed using such patterns as “An X is a (kind of) Y ”. Hypernymy is transitive and asymmetrical. Such relations are central organizing relations for nouns and verbs in WordNet [11]. In fact, hypernymy relations comprise

class (or subsumption) relations and instantiation relations considered in ontology studies and description logics [12].

Taxonomy is a special case of a knowledge graph G . It is a tree-based structure where nodes from the set E are connected with the hypernymy relation r . Thus, the set of relations $R = \{r\}$. Elements included in the set E are words or concepts. In a taxonomy G they are arranged in a hierarchical structure from the most abstract concept (*root*) to the most specific concepts (*leaves*).

Query words (new words) — words to be added to the taxonomy T , usually manually collected by experts. In this paper, we use the following notation for this set $Q \subseteq \{q_1, \dots, q_n\}$. Note that words may be ambiguous: one q_i can correspond to more than one synset s_i .

Knowledge graph completion is a task of completing a Knowledge Graph G by finding a set of missing triples $\{h, r, t | h \in E, r \in R, t \in E, h, r, t \notin G\}$.

Taxonomy enrichment can be considered as a special case of the knowledge base completion task. This task aimed at associating each new word $q \in Q$, which is not yet included in the taxonomy T , with the appropriate hypernyms from it.

3. Related Work

There exist numerous approaches to knowledge base completion, which make use of various neural network architectures described in review papers [13, 14]. Most of them apply low-dimensional graph embeddings [15, 16], deep learning architectures like autoencoders [17] or graph convolutional networks [18–20]. Another group of approaches makes use of tensor decomposition approaches: Tucker decomposition [21] and Canonical Polyadic decomposition (CANDECOMP/PARAFAC) [22, 23]. Another task, which is closely related to the taxonomy creation is the Knowledge base construction. There exists multiple approaches solving the problem: Text2Onto [24] a pillar language-independent approach which applies user interaction or fully automated methods [25, 26] which apply text-mining tools and external sources, which are mostly applied for the scholarly domain. However, knowledge base completion task assumes a generic graph while in the taxonomy enrichment task deals with tree structures and specific methods of tree pro-

1 censing are commonly used in this field. In this arti-
2 cle we focus on this specific task and narrow down
3 our scope to enrichment and population of taxonomic
4 structures. It should be noted however that the majority
5 of the ontologies and knowledge bases possess some
6 kind of taxonomic backbone and therefore the task of
7 construction and maintaining such a semantic structure
8 is fundamental.

9 The existing studies on the taxonomies can be di-
10 vided into three groups. The first one addresses the Hy-
11 pernym Discovery problem [27]: given a word and a
12 text corpus, the task is to identify hypernyms in the
13 text. However, in this task the participants are not given
14 any predefined taxonomy to rely on. The second group
15 of works deals with the Taxonomy Induction problem
16 [28–30], in other words, creation of a taxonomy from
17 scratch. Finally, the third direction of research is the
18 Taxonomy Enrichment task: the participants extend a
19 given taxonomy with new words. In this article, we fo-
20 cus on the taxonomy enrichment task and explore var-
21 ious approaches based on text and graph embeddings
22 and their combinations to the solution of this problem.
23 The following sections provide overview of the various
24 strains of research related to the given problem.

25 3.1. Prior Art on Taxonomy Enrichment2

26
27
28 Until recently, the only dataset for the taxonomy en-
29 richment task was created under the scope of SemEval-
30 2016. It contained definitions for the new words, so the
31 majority of models solving this task used the defini-
32 tions. For instance, *Deftor* team [4] computed defini-
33 tion vector for the input word, comparing it with the
34 vector of the candidate definitions from WordNet using
35 cosine similarity. Another example is *TALN* team [5]
36 which also makes use of the definition by extracting
37 noun and verb phrases for candidates generation.

38 This scenario may be unrealistic for manual anno-
39 tation because annotators are writing a definition for
40 a new word and adding new words to the taxonomy
41 simultaneously. Having a list of candidates would not
42 only speed up the annotation process but also iden-
43 tify the range of possible senses. Moreover, it is pos-
44 sible that not yet included words may have no defini-
45 tion in any other sources: they could be very rare (“ap-
46 paratchik”, “falanga”), relatively new (“selfie”, “hash-
47 tag”) or come from a narrow domain (“vermiculite”).

48 Thus, following RUSSE-2020 shared task [9], we
49 stick to a more realistic scenario when we have no def-
50 initions of new words, but only examples of their us-
51 age. For this shared task we provide a baseline as well

1 as training and evaluation datasets based on RuWord-
2 Net [31] which will be discussed in the next section.
3 The task exploited words which were recently added to
4 the latest release of RuWordNet and for which the hy-
5 pernym synsets for the words were already identified
6 by qualified annotators. The participants of the compe-
7 tition were asked to find synsets which could be used
8 as hypernyms.

9 The participated systems mainly relied on vector
10 representations of words and the intuition that words
11 used in similar contexts have close meanings. They
12 cast the task as a classification problem where words
13 need to be assigned one or more hypernyms [32] or
14 ranked all hypernyms by suitability for a particular
15 word [33]. They also used a range of additional re-
16 sources, such as Wiktionary, dictionaries, additional
17 corpora [34]. Interestingly, only one of the well-
18 performing models [35] used context-informed em-
19 beddings (BERT) or external tools such as online Ma-
20 chine Translation (MT) and search engines (the best-
21 performing model denoted as *Yuriy* in the workshop
22 description paper).

23 In this paper, we would like to work out methods
24 which depend on graph based structures and combine
25 them with the approaches applying word embeddings.
26 At the same time, we want our methods to benefit
27 from the existing data (e.g. corpora, pre-trained em-
28 beddings, Wiktionary).

29 3.2. Word Vector Representations for Taxonomies2

30
31
32 Approaches using word vector representations are
33 the most popular choice for all tasks related to
34 taxonomies. When solving the *Hypernym Discovery*
35 problem in SemEval-2018 Task 9 [27] most of par-
36 ticipants use word embeddings. For instance, Bernier-
37 Colborne and Barriere [36] predict the likelihood of
38 the relationship between an input word and a candi-
39 date using word2vec embeddings. Berend et al. [37]
40 use Word2vec vectors as features of a logistic regres-
41 sion classifier. Maldonado and Klubicka [38] simply
42 consider top-10 closest associates from the Skip-gram
43 word2vec model as hypernym candidates. Pre-trained
44 GloVe embeddings [39] are also used to initialize em-
45 beddings for an LSTM-based Hypernymy Detection
46 model [40].

47 Participants also solve the SemEval-2016 Task 13
48 on taxonomy induction with word embeddings [41]:
49 they compute the vector offset as the average offset of
50 all the pairs generated and exploit it to predict hyper-
51 nyms for the new data. Afterwards, in [42] the authors

1 apply word2vec embeddings similarity to improve the
2 approaches of the SemEval-2016 Task 13 participants.

3 The vast majority of participants of SemEval-2016
4 task 14 [3] and RUSSE’2020 [9] also apply word
5 embeddings to find the correct hypernyms in the exist-
6 ing taxonomy. For instance, the participants compute
7 a definition vector for the input word by comparing
8 it with the definition vectors of the candidates from
9 the wordnet using cosine similarity [4]. Another op-
10 tion is to train word2vec embeddings from scratch and
11 cast the task as a classification problem [32]. Some
12 participants compare the approach based on XLM-
13 R model [43] with the word2vec “hypernyms of co-
14 hyponyms” method [34]. It considers nearest neigh-
15 bours as co-hyponyms and takes their hypernyms as
16 candidate synsets.

17 Summing up, the usage of distributed word vector
18 representations is a simple yet efficient approach to
19 the taxonomy-related tasks and should be considered a
20 strong baseline [9, 27].

22 3.3. Meta-Embedding Approaches to Word 23 Representation2

24 Vector representations can be learned on various
25 datasets and using various models. It has been shown
26 that combining word embeddings is beneficial for NLP
27 tasks, e.g. dependency parsing [44], and in medical do-
28 main [45].

29 Coates et al. [46] show that simple vector combin-
30 ing approaches, such as concatenation or averaging,
31 can significantly improve the overall performance for
32 several tasks. For instance, singular value decomposi-
33 tion (SVD) demonstrates good results with the ability
34 to control the final dimension of vectors [47]. Autoen-
35 coders [48] further promote the idea of creating meta-
36 embeddings. The authors propose several algorithms
37 for combining various word vectors into one vector
38 by encoding initial vectors into some meta-embedding
39 space and then decoding backwards.

40 As for the CAEME approach, all word vectors
41 are encoded into meta-vectors and then concatenated.
42 Then, the decoding step uses a concatenated repre-
43 sentation to predict the original vector representations.
44 The AAEME approach is similar to CAEME, except
45 that each vector is mapped to a fixed-size vector and all
46 encoded representations are averaged, but not concate-
47 nated. An obvious advantage of this approach is the
48 ability to control the dimension of meta-embeddings.

49 For any AEME approach, different loss functions
50 can be used at the decoding stage: MSE loss, KL-

1 divergence loss, cosine distance loss and also their
2 combinations. In [49] the authors investigated the per-
3 formance of the autoencoders depending on the loss
4 function. They discover that there is no evident win-
5 ner across tasks and that different loss functions are
6 defined by different tasks.

7 Meta-embeddings are already used in such machine
8 learning tasks as dependency parsing [44], classifica-
9 tion in healthcare [45], named-entity recognition [49,
10 50], sentiment analysis [49], word similarity and anal-
11 ogy tasks [46–48]. To the best of our knowledge, meta-
12 embeddings have not been applied to the taxonomy en-
13 richment task, especially for the fusion of texts and
14 graph embeddings.

16 3.4. Graph-based Representations for Taxonomies2

17 Taxonomies can be represented as graphs and there
18 exist various approaches to learn graph-based repre-
19 sentations. They are thoroughly compared in [51] on
20 the link prediction task, which is closely related to Tax-
21 onomy Enrichment. The paper also demonstrates that
22 the combination of text and graph embeddings gives a
23 boost on the link prediction task. Most of those meth-
24 ods listed in [51] have also been tested on tasks related
25 to the taxonomy enrichment.

26 For instance, node2vec embeddings [52] are used
27 for taxonomy induction among other network embed-
28 dings [53]. In [42], the authors perform the same task.
29 They use hyperbolic Poincaré embeddings to enhance
30 automatically created taxonomies. The SemEval-2016
31 subtask of reattaching query words to the taxon-
32 omy is quite similar to taxonomy enrichment which
33 we perform. However, the datasets of the SemEval-
34 2016 Task 13 are restricted to specific domains,
35 which leaves an open question of the efficiency of
36 Poincaré embeddings for the general domain and
37 larger datasets. Moreover, [42] use Hearst Patterns to
38 discover hyponym-hypernym relationships. This tech-
39 nique operates on words, and cannot be transferred to
40 word-synset relations without extra manipulation.

41 As for the Knowledge Graph Construction task,
42 which is a more general task in relation to the Tax-
43 onomy Induction, the vast majority of approaches
44 also use word embeddings as node representations.
45 Several approaches like [54] and apply ELMO em-
46 beddings [55] to predict entities and their relations
47 for the knowledge graph. Other approaches [56] uti-
48 lize a combination of ELMO, Poincaré and node2vec
49 embeddings to enhance knowledge graph build upon
50 Wikipedia.

Graph convolutional networks (GCNs) [57] as well as graph autoencoders [58] are mostly applied to the link prediction task on large knowledge bases. For example, in [59] the authors present an expanded review of the field and compare a wide variety of existing approaches. Graph embeddings are also often used for other taxonomy-related tasks, e.g. entity linking [60]. As for the taxonomy enrichment task, we are only aware of a recent approach TaxoExpan [61] which applies position-enhanced graph neural networks (GCN [62] and GAT [63]) that we also evaluate on our datasets³.

Thus, to the best of our knowledge, our work is the first computational study of Taxonomy enrichment task which aggregates and considers different existing and new approaches for taxonomy enrichment. We compare graph- and word-based representations computed from the synsets and hypo-hypernym relations for hypernym prediction demonstrating state-of-the-art results.

4. Diachronic WordNet Datasets¹

The important part of our study is the observation that one can learn from the history of the development of lexical resources through time. More specifically, we make use of the various historic snapshots (versions) of WordNet lexical graphs and setup a task of their automatic completion assuming the manual update of the ground truth. This diachronic analysis – similar to diachronic lexical analysis of word meanings – is used to build two datasets in our study: one for English, another one for Russian based respectively on Princeton WordNet [64] and RuWordNet taxonomies. It is important to mention that by using the word “diachronic” we do not imply lexical diachrony, e.g., semantic shifts of words over time [65], but the temporal extension of Wordnet stored in its versions. Each dataset consists of a taxonomy and a set of novel words to be added to this resource. The statistics are provided in Table 1.

4.1. English Dataset²

To compile dataset, we choose two versions of WordNet and then select words which appear only in a newer version. For each word, we get its hypernyms from the newer WordNet version and consider them as

³The results achieved on our datasets are significantly lower than the baseline, probably because of the incorrect model launching.

Table 1

Statistics of two diachronic WordNet datasets used in this study.

Dataset	Nouns	Verbs
<i>WordNet1.6 - WordNet3.0</i>	17 043	755
<i>WordNet1.7 - WordNet3.0</i>	6 161	362
<i>WordNet2.0 - WordNet3.0</i>	2 620	193
<i>RuWordNet1.0 - RuWordNet2.0</i>	14 660	2 154
<i>RUSSE'2020</i>	2 288	525

gold standard hypernyms. We add words to the dataset if only their hypernyms appear in both versions. We do not consider adjectives and adverbs, because they often introduce abstract concepts and are difficult to interpret by context. Besides, the taxonomies for adjectives and adverbs are worse connected than those for nouns and verbs making the task more difficult.

In order to find the most suitable pairs of releases, we compute WordNet statistics (see Table 2). New words demonstrate the difference between the current and the previous WordNet version. For example, it shows that the dataset generated by “subtraction” of WordNet 2.1 from WordNet 3.0 would be too small, they differ by 678 nouns and 33 verbs. Therefore, we create several datasets by skipping one or more WordNet versions. The statistics for the datasets we select for our study are provided in Table 1.

As gold standard hypernyms, we use not only the immediate hypernyms of each *lemma* (initial form of a word — infinitive for a verb, single number and nominative case for a noun, etc.) but also the second-order hypernyms: hypernyms of the hypernyms. We include them in order to make the evaluation less restricted. According to our empirical observations, the task of automatically identifying the exact hypernym might be too challenging, and finding the “region” where a word belongs (“parents” and “grandparents”) can already be considered a success.

This method of dataset construction does not use any language-specific or database-specific features, so it could be transferred to other wordnets or taxonomies with timestamped releases.

All datasets⁴ created for this research and the code⁵ for their construction are publicly available.

⁴<https://zenodo.org/record/4279821>

⁵<https://github.com/skoltech-nlp/diachronic-wordnets>

Table 2

Statistics of the English WordNet taxonomies used in this study.

Taxonomy	Synsets		Lemmas		New words	
	Nouns	Verbs	Nouns	Verbs	Nouns	Verbs
<i>WordNet 1.6</i>	66 025	12 127	94 474	10 319	-	-
<i>WordNet 1.7</i>	75 804	13 214	109 195	11 088	11 551	401
<i>WordNet 2.0</i>	79 689	13 508	114 648	11 306	4 036	182
<i>WordNet 2.1</i>	81 426	13 650	117 097	11 488	2 023	158
<i>WordNet 3.0</i>	82 115	13 767	117 798	11 529	678	33

Table 3

Examples of Russian nouns with translation mapped to English WordNet.

Word	Translation	RuWordNet hypernyms	RuWordNet hypernym names	WordNet hypernyms
абсентеизм	absenteeism	["147309-n", "117765-n", "117017-n"]	['неучастие, отказ от участия', 'уклониться (отказаться)', 'отказаться, не согласиться']	["non-engagement.n.01", "evasion.n.03", "rejection.n.01"]
кибертерроризм	cyber terrorism	["7334-n", "4590-n", "2400-n"]	['преступление против общественной безопасности', 'компьютерное преступление', 'терроризм']	[null, "cybercrime.n.01", "terrorism.n.01"]
метропоезд	subway train	["141975-n", "7133-n"]	['электрическое транспортное средство', 'электropоезд']	[null, null]

4.2. Russian Datasets²

In order to create an analogous version to English dataset for Russian, we use the RuWordNet taxonomy [31]. RuWordNet comprises *synsets* — sets of synonyms expressing a particular concept. A synset consists of one or more *senses* — words or multi-word constructions in the initial form. Therefore, we use the current version of RuWordNet and the extended version of RuWordNet which has not been published yet to compile the dataset (cf. Table 1).

The RUSSE'2020 dataset was created for the Dialogue Evaluation [9] and can be viewed as a restricted subset of the Russian dataset. In the RUSSE'2020 the following categories of words were excluded:

- all three-symbol words and the majority of four-symbol words;
- diminutive word forms and feminine gender-specific job titles;
- words which are derived from words which are included in the published RuWordNet;
- words denoting inhabitants of cities and countries;

- geographic and personal names;
- compound words that contain their hypernym as a substring.

4.3. WordNet ILI mapping (ru-en)²

In order to connect wordnets in different languages the Inter-Lingual Index (ILI) is used [66]. This mapping is designed to make possible coordination between wordnet projects. For the Russian test sets we also provide mapping from RuWordNet to WordNet⁶. For each hypernym synset of each query word we present the corresponding WordNet synset index. Table 3 demonstrates several examples of this kind of mapping. As we can see, datasets for the Russian nouns and verbs are extended with additional column called "WordNet synsets", where the corresponding WordNet3.0 synsets are listed in accordance with the Russian synset list.

However, not all synsets have an equivalent in the other language, as there exist untranslatable concepts

⁶<https://doi.org/10.5281/zenodo.4969267>

and lacunae. Therefore, we present in the Table 4 the coverage of the WordNet synsets for the hypernyms of query words from the test set. This mapping can be further used for multilingual experiments.

Table 4

Coverage of the WordNet synsets for the hypernyms in the Russian test sets.

Input type	Total	Have ILI mapping
Non-restricted nouns		
All synsets	41,694	28,425
Unique synsets	4,777	2,791
Query words	17,475	15,251
Restricted (private) nouns		
All synsets	4,456	3,087
Unique synsets	1,376	885
Query words	1,920	1,720
non-restricted verbs		
All synsets	6,783	4,860
Unique synsets	1,473	931
Query words	2,872	2,606
restricted (private) verbs		
All synsets	1,110	821
Unique synsets	611	419
Query words	477	440

4.4. Evaluation Metric2

The goal of diachronic taxonomy enrichment is to build a newer version of a wordnet by attaching the new given terms to the older wordnet version. We cast this task as a soft ranking problem and use Mean Average Precision (MAP) score for the quality assessment:

$$MAP = \frac{1}{N} \sum_{i=1}^N AP_i, \quad (1)$$

$$AP_i = \frac{1}{M} \sum_{j=1}^n prec_j \times I[y_j = 1],$$

where N and M are the number of predicted and ground truth values, respectively, $prec_i$ is the fraction of ground truth values in the predictions from 1 to i , y_i is the label of the i -th answer in the ranked list of predictions, and I is the indicator function.

This metric is widely acknowledged in the Hypernym Discovery shared tasks, where systems are also

evaluated over the top candidate hypernyms [27]. The MAP score takes into account the whole range of possible hypernyms and their rank in the candidate list.

However, the design of our dataset disagrees with MAP metric. As we described in Section 4, the gold-standard hypernym list is extended with second-order hypernyms (parents of parents). This extension can distort MAP. If we consider all gold standard answers as compulsory for the maximum score, it means that we demand models to find both direct and second-order hypernyms. This disagrees with the original motivation of including second-order hypernyms to the gold standard — it was intended to make the task easier by allowing a model to guess a direct *or* a second-order hypernym.

On the other hand, if we decide that guessing *any* synset from the gold standard yields the maximum MAP score, we will not be able to provide an adequate evaluation for words with multiple direct hypernyms. There exist two cases thereof:

1. the target word has two or more hypernyms which are co-hyponyms or one is a hypernym of the other — this word has a single sense, but the annotator decided that multiple related hypernyms are needed to reflect all shades of the meaning,
2. the target word has two or more hypernyms which are not directly connected in the taxonomy and neither are their hypernyms. This happens if:
 - (a) the word’s sense is a composition of senses of its hypernyms, e.g. “impeccability” possesses two components of meaning: (“correctness”, “propriety”) and (“morality”, “righteousness”);
 - (b) the word is polysemous and different hypernyms reflect different senses, e.g. “pop-up” is a book with three-dimensional pages (“book, publication”) and a baseball term (“fly, hit”).

While the case 2a corresponds to a monosemous word and the case 2b indicates polysemy, this difference does not affect the evaluation process. We propose that in both these cases in order to get the maximum MAP score a model should capture all the unrelated hypernyms which correspond to different components of sense. At the same time, we should bear in mind that guessing a direct hypernym or a second-order hypernym are equally good options. Therefore, following [9], we evaluate our models with modified MAP. It transforms a list of gold standard hyper-

nyms into a list of connected components. Each of these components includes hypernyms (both direct and second-order) which form a connected component in a taxonomy graph. (According to graph theory, connected component is a subgraph, in which there is a path between any two nodes.) Thus, in the case 1 we will have a single connected component, and a model should guess *any* hypernym from it to get the maximum MAP score. In the cases 2a and 2b we will have multiple components, and a model should guess *any* hypernym from *each* of the components.

5. Base Methods

Here we first describe our baseline model which is a method of synset ranking based on distributional embeddings and hand-crafted features (the method was proposed as a baseline for RUSSE-2020 shared task [6]). We then propose extending it with new features extracted from Wiktionary and use the alternative sources of information about words (e.g. graph representations) and their combinations.

5.1. Baseline

We consider the approach by Nikishina et al. [6] as our baseline. There, we first create a vector representation for each synset in the taxonomy by averaging vectors (pretrained embeddings) of all words from this synset. Then, we retrieve top 10 synsets whose vectors are the closest to that of the *query word* (we refer to these synsets as *synset associates*). For each of these *associates*, we extract their immediate hypernyms and hypernyms of all hypernyms (second-order hypernyms). This list of the first- and second-order hypernyms forms our *candidate set*. We need to rank the candidates by their relevance for the query word. Note that the lists of candidates for different associates can have intersections. When forming the overall candidate set, we make sure that each candidate occurs in it only once.

The intuition behind the method is the following. We propose that if a synset of a taxonomy is a *parent* of a word which is similar to our query word, it can also be a parent of this query word.

To rank the candidate set of synsets we train a Linear Regression model with L2-regularisation on the training dataset formed of the words and synsets of WordNet. Candidate hypernyms are ranked by their model

output score. We limit the output to the $k = 10$ best candidates.

We rank the candidate set using the following features:

- $n \times \text{sim}(v_i, v_{h_j})$, where v_x is a vector representation of a word or a synset x , h_j is a hypernym, n is the number of occurrences of this hypernym in the merged list, $\text{sim}(v_i, v_{h_j})$ is the cosine similarity of the vector of the input word i and hypernym vector h_j ;
- the candidate presence in the Wiktionary hypernyms list for the input word (binary feature);
- the candidate presence in the Wiktionary synonyms list (binary feature);
- the candidate presence in the Wiktionary definition (binary feature);
- the average cosine similarity between the candidate and the Wiktionary hypernyms of the input word.

5.2. DWRank2

We present a new method of taxonomy enrichment — Distributional Wiktionary-based synset Ranking (**DWRank**). It combines distributional features with features from Wiktionary. DWRank builds up on the baseline described in Section 5.1. We extend the baseline Logistic Regression model with the new features which mainly account for the number of occurrences of a synset in the candidate lists of different synset associates (nearest neighbours) of the query word. We introduce the following new features:

- the number of occurrences (n) of the synset in the merged candidate list and the quantity $\log_2(2+n)$ which serves for smoothing,
- the minimum, average, and maximum proximity level of the synset in the merged candidate list:
 - * the level is 0 if the synset was added based on similarity to the query word,
 - * the level of 1 is for the immediate hypernyms of the query word,
 - * the level of 2 is for the hypernyms of the hypernyms,
- the minimum, average, and maximum similarities of the query word to all words of the synset,
- the features based on hyponyms of a candidate synset (“children-of-parents”):
 - * we extract all hyponyms (“children”) of the candidate synset,

- * for each word/phrase in each hyponym synset we compute their similarity to the query word,
- * we compute the minimum, average, and maximum similarity for each hyponym synset,
- * we form three vectors: a vector of minimums of similarities, average similarities, and maximum similarities of hyponym synsets,
- * for each of these vectors we compute minimum, average, and maximum. We use these resulting 9 numbers as features.

These features account for different aspects of similarity of the candidate’s children to the query word and help defining if these children can be the query word’s co-hyponyms (“siblings”).

Moreover, in this approach we use cross-validation and feature scaling when training the Logistic Regression model.

This methods could be easily extended to other languages that possess a taxonomy, a wiki-based open content dictionary (Wiktionary) and text embeddings like fastText or/and word2vec and GloVe.

5.3. Web-based Synset Ranking (WBSR)²

In this section, we describe **WBSR** (Web-based Synset Ranking) a method which leverages the power of the existing general-purpose services. It makes use of the two famous search engines: Google (for both English and Russian datasets) and Yandex⁷ (for Russian only). According to our hypothesis, the search results for a word are likely to contain its hypernyms or co-hyponyms as they are often used to define a word via generalisation or by providing synonyms (co-hyponyms). For instance, if we do not know what “abdominoplasty” is, searching for it with a search engine can yield its definition “a cosmetic surgery procedure”.

Another source that we could probably benefit from is another taxonomy, preferably larger than the one we work with. However, there might be no other taxonomies available in the same language. Therefore, in this case we can resort to Machine Translation and automatically translate query words into a rich-resource language (e.g. English) in order to use an existing taxonomy (e.g. English Princeton WordNet). In this study

we use Yandex Machine Translation system⁸ to translate query words into English and then translate hypernyms (if they are found) back into Russian.

The main drawback of using external sources such as search engines and machine translation systems is their weak reproducibility. The search results are dependent on the search history, so reproducing the experiment on a different account or after a relatively long period of time is problematic. However, since the method greatly improves the performance even with trivial handling of the collected data, we use it despite its drawback. To make our results reproducible, we release all data from the external sources used in our approach.⁹

Similarly to the approaches described above, here we also make use of Wiktionary and fastText embeddings cosine similarity. However, we treat synsets and words/phrases that they consist of in a different way. In the previously described approaches we computed embeddings for multiword phrases by averaging word embeddings of individual words in them. Here we treat them as sentences — we compute their embeddings using the `get_sentence_vector` method from fastText Python library. There, fastText vectors are divided by their norms and then averaged, so that only vectors with the positive L_2 -norm value are considered. Secondly, we do not combine the word/phrase vectors into a synset vector but operate with the word/phrase embeddings directly.

Similarly to DWRank, the algorithm consists of two steps: candidate generation and candidate ranking. Our candidate list is formed of the following synsets:

- synsets which contain words/phrases from the list of top-10 nearest neighbours of the query word;
- hypernyms and second-order hypernyms of those synsets;
- Wiktionary-based candidates:
 - * synsets that contain words/phrases listed in Wiktionary as the hypernyms of the query word;
 - * hypernym synsets of these synsets;
- cross-lingual candidates (for the Russian language only):
 - * synsets that contain words/phrases listed in the English WordNet as the hypernyms of the query word;

⁷<https://yandex.com>

⁸<https://translate.yandex.ru/>

⁹<https://doi.org/10.5281/zenodo.4540717>

* hypernym synsets of these synsets.

Analogously to DWRank, we then rank all candidates by a logistic regression model which uses the following features:

- the candidate synset contains a word/phrase from the list of query word’s nearest neighbours;
- the candidate synset is a hypernym of one of the nearest neighbours;
- the candidate is a second-order hypernym of one of the nearest neighbours;
- the candidate synset contains a hypernym of the query word from Wiktionary;
- the candidate synset is a hypernym of the synset which contains a hypernym of the query word from Wiktionary;
- the candidate synset contains a word which is present in the definition of the query word from Wiktionary;
- the candidate synset is present in the list of English WordNet synset candidates (for the Russian language only);
- the candidate synset is present in the list of hypernyms of the WordNet candidates (for the Russian language only);
- the candidate synset contains words which occur on the Google results page;
- the candidate synset contains words which occur on the Yandex results page (for the Russian language only).

The training set for the logistic regression model is formed from a wordnet in the relevant language as follows. For each query word in the list of query words we first find the most similar lemma which is contained in the wordnet. We know hypernyms for these lemmas and use them to generate the training set. We generate the candidate list as described before. First- and second-order hypernyms in this candidate list are used as positive examples for the corresponding lemmas, and synsets from the candidate list which are not hypernyms are considered negative examples.

This approach participated in the RUSSE’2020 Taxonomy Enrichment task for the Russian Language. The method achieved the best result on the nouns track. Therefore, we consider it as the-state-of-the-art method for Russian.

5.4. WordNet Path Prediction2

A completely different approach to make use of fastText embeddings is presented in the work of Cho et

al. [67]. The authors experiment with encoder-decoder models in order to solve the task of the direct hypernym prediction. They use a standard LSTM-based sequence-to-sequence model [68] with Luong attention [69]. First, they average fastText embeddings for the input word or phrase and put it through the encoder. The decoder sequentially generates a chain of synsets from the encoder hidden state, conditioned on the previously generated ones. The authors consider two different setups:

- *hypo2path* — given the input word, generate a sequence of synsets starting from the root synset and going down the taxonomy to the closest hypernym;
- *hypo2path reverse* — given the input word, generate a sequence of synsets starting from the closest hypernym up to the root entity.

To be able to apply this sequence-to-sequence architecture to our data, we build new datasets similar to the ones described in [67]. We generate a path from the WordNet starting from the root node to the target synset or word. Analogously to the original work, we include multiple paths from the root to the parents of the query word. We filter the validation set to only include queries that do not occur anywhere in the full taxonomy paths of the training data. To sort candidates generated by the decoder, we enumerate the generated *hypo2path* sequence from the right to the left or the *hypo2path reverse* from the left to the right and get the first 10 synsets.

Additionally, we extend this approach by replacing the LSTM+attention architecture with the Transformer architecture [70]. During training we provide an embedding of a synset as input to the Transformer and expect the model to generate a sequence of synsets starting from the hypernym of the input synset. During inference we provide embedding of query words as input expect the model to output sequences of synsets starting with the direct hypernyms.

5.5. Word Representations for DWRank2

We test our baseline approach and DWRank with different types of embeddings: fastText [71], word2vec [72] embeddings for English and Russian datasets and also GloVe embeddings [39] for the English dataset.

We use the fastText embeddings from the official website¹⁰ for both English and Russian, trained on

¹⁰<https://fasttext.cc/docs/en/crawl-vectors.html>

Common Crawl from 2019 and Wikipedia CC including lexicon from the previous periods as well. For word2vec we use models from [73, 74] for both English¹¹ and Russian.¹² We lemmatise words and synsets for both languages with the same UDPipe [75] model which was used while training the representations. For the out-of-vocabulary (OOV) words we find all words in the vocabulary with the longest prefix matching this word and average their embeddings like in [33]. As for the GloVe embeddings, we also use them from the official website¹³ trained on Common Crawl, the vocabulary size is 840 billion tokens.

6. DWRank-Graph1

The DWRank method extracts a set of candidate synsets based on the similarities of word vectors. So far we used only distributional word vectors (fastText, GloVe, etc.) to represent words. On the other hand, graph-based representations can contain the taxonomic information which is absent in distributional embeddings [51].

Here, we present **DWRank-Graph**. This is the same DWRank method where the distributional embeddings are replaced with graph representations. The score prediction model and the features it uses do not change. Below we describe the graph representations and their combinations we applied in DWRank-graph.

6.1. Poincaré Embeddings2

Poincaré embeddings is an approach for “learning hierarchical representations of symbolic data by embedding them into hyperbolic space — or more precisely into an “ n -dimensional Poincaré ball” [76]. Poincaré models are trained on hierarchical structures and simultaneously capture hierarchy and similarity due to the underlying hyperbolic geometry. According to the authors, hyperbolic embeddings are more efficient on the hierarchically structured data and may outperform Euclidean embeddings in several tasks, e.g. in Taxonomy Induction [42].

Therefore, we use Poincaré embeddings of our wordnets for the taxonomy enrichment task. We train Poincaré ball model for our wordnets using the default

parameters and the dimensionality of 10, which yields the best results on the link prediction task [76].

However, applying these embeddings to the task is not straightforward, because Poincaré model’s vocabulary is non-extensible. It means that new words that we need to attach to the existing taxonomy will not have any Poincaré embeddings at all and we cannot make use of the embeddings similarity. To overcome this limitation, we compute top-5 fastText nearest synsets (analogously to the procedure described in Section 5.1) and then aggregate embeddings in hyperbolic space using Einstein midpoint, following [77]. The resulting vector is considered as an embedding of the input word in the Poincaré space.

Then, we use vectors from this vector space to generate candidates for the DWRank approach. As the model we present in Section 5.2 does not depend on the types of input embeddings, we are able to provide Poincaré embeddings as input.

6.2. Node2vec Embeddings2

The hierarchical structure of the taxonomy is a graph structure, and we may also consider taxonomies as graphs and apply random walk approaches to compute embeddings for the synsets. For this purpose we apply node2vec [52] approach which represents a “random walk of a fixed length l ” with “two parameters p and q which guide the walk in breadth or in depth”. Node2vec randomly samples sequences of nodes and then applies the skip-gram model [78] to train their vector representations. We train node2vec representations of all synsets in our wordnets with the following parameters: *dimensions* = 300, *walk_length* = 30, *num_walks* = 200. The other parameters are taken from the original implementation.

However, analogously to Poincaré vector space, node2vec model has no techniques for representing out-of-vocabulary words. Thus, it is unable to map new words to the vector space. To overcome this limitation, we apply the same technique of averaging top-5 nearest neighbours from fastText and considering their mean vector as the new word embedding and search for the most similar synsets.

6.3. Graph Neural Networks2

The models described above have a major shortcoming: the resulting vectors for the input words heavily depend on their representations in the fastText model. This can lead to the incorrect results if the word’s near-

¹¹<http://vectors.npl.eu/repository/20/29.zip>

¹²<http://vectors.npl.eu/repository/20/185.zip>

¹³<https://nlp.stanford.edu/projects/glove/>

est neighbour list is noisy and does not reflect its meaning. In this case the noise will propagate through the graph embedding (Poincaré or node2vec) model and result in inaccurate output even if the graph embedding model is of high quality.

Therefore, we test different graph neural network (GNN) architectures — Graph Convolutional Network [57], Graph Attention Network [63] and GraphSAGE [79] (SAmple and aggregatE) which make use of both fastText embeddings and the graph structure of the taxonomy.

All the above mentioned models work similarly. According to [51], “GCN works similarly to the fully-connected layers for neural networks. It multiplies weight matrices with the original features but masking them with an adjacency matrix. Such a method allows to account not only node representation, but also representations of neighbours and two-hop neighbours.” The GraphSAGE model addresses the problem of unseen nodes representation by training a set of aggregator functions that learn to aggregate feature information from a node’s local neighborhood. GCN, on the contrary, learns a distinct embedding vector for each node. In GAT, the convolution operation from GCN is replaced with the attention mechanism. It uses the self-attention mechanism of Transformers [80] to aggregate the information from the one-hop neighbourhood.

FastText embeddings are used as input node features for all models, which is definitely an advantage of the model over Poincaré and node2vec, as they do not use word embeddings for training. Even though new words are not connected to the taxonomy, it is still possible to compute their embeddings according to their input node features.

We get the vector representations of query words from one of the pre-trained GNN models and then use them as the input to DWRank. Even though all methods work similarly, they demonstrate different performance on different datasets.

6.4. Text-Associated Deep Walk2

Text-Associated Deep Walk (TADW) [81] is another approach that incorporates text and graph information into one vector representation. The method is based on the DeepWalk algorithm [82] which learns feature representations by simulating uniform random walks. To be specific, the sampling strategy in DeepWalk can be seen as a special case of node2vec with $p = 1$ and $q = 1$.

The authors prove that the DeepWalk approach is equivalent to matrix factorization. They incorporate text features of vertices into network representation learning within the framework of matrix factorization. First, they define matrix $M \in \mathbb{R}^{|V| \times |V|}$ where each entry M_{ij} is the logarithm of the average probability that vertex v_i randomly walks to vertex v_j in a fixed number of steps. In comparison to DeepWalk, where the goal is to factorize matrix M into the product of two low-dimensional matrices $W \in \mathbb{R}^{k \times |V|}$ and $H \in \mathbb{R}^{k \times |V|}$ ($k \ll |V|$), TADW aims to factorize matrix M into the product of three matrices: $W \in \mathbb{R}^{k \times |V|}$, $H \in \mathbb{R}^{k \times f_i}$ and text features $T \in \mathbb{R}^{f_i \times |V|}$. As text features, in this work we apply fastText embeddings.

After having learnt the factorisation of matrix M , we use rows of matrix W as node (synset) embeddings in DWRank.

6.5. High-Order Proximity preserved Embeddings2

High-Order Proximity preserved Embeddings (HOPE) [83] is yet another approach that embeds a graph into a vector space preserving information about graph properties and structure. Unfortunately, most structures cannot preserve the asymmetric transitivity, which is a critical property of directed graphs. To solve the problem, the authors employ matrix factorization to directly reconstruct asymmetric distance measures like Katz index, Adamic-Adar or common neighbors. This approach is scalable — it preserves high-order proximities of large scale graphs, and capable of capturing the asymmetric transitivity. HOPE outperforms state-of-the-art algorithms in tasks of reconstruction, link prediction and vertex recommendation.

As for our Taxonomy Enrichment task, we also apply HOPE to generate graph embeddings to be used as input in the DWRank model. The main difference with the other embeddings is that HOPE does not incorporate textual information from the nodes.

7. DWRank-Meta1

In DWRank we employed only distributional information, i.e. pre-trained word embeddings, whereas in DWRank-Graph we represented words using the information from the graph structure of the taxonomy and usually ignoring their distributional properties. Meanwhile, taxonomy enrichment models may benefit from combining these two types of information. Therefore, we present **DWRank-Meta** — an exten-

sion of DWRank which combines multiple types of input word representations.

As in DWRank-Graph, the process of candidates selection, the feature set and the algorithm of synset ranking stay intact. Here we change the input representations of words and synsets.

7.1. Base Meta-Embeddings2

The easiest way of combining embeddings of different types is to concatenate them and use the concatenated vector as an input. We refer to this method as **Concat** and combine different subsets of distributional (fastText, word2vec, GloVe) and graph-based (Poincaré, node2vec, GCN, GAT, GraphSAGE, TADW, HOPE) embeddings. In addition to that, we perform Singular value decomposition (SVD) over this concatenation as proposed in [47]. This approach is referred to as **SVD**.

7.2. Autoencoded Meta-Embeddings2

We propose using two variants of autoencoders for the generation of meta-embeddings: Concatenated Autoencoded Meta-Embeddings (CAEME) and Averaged Autoencoded Meta-Embeddings (AAEME) [48]. They have shown good results on the task of evaluating lexical similarity. However, they have never been applied to taxonomy enrichment.

We generate meta-embeddings as follows. Let us consider an embedding model $s(w)$. For each of such embedding models we train an autoencoder consisting of an encoder and a decoder:

$$\begin{aligned} E(s(w)) &= h(w), \\ D(h(w)) &= \hat{s}(w), \\ L_{ED} &= \text{dist}(s(w), \hat{s}(w)), \end{aligned} \quad (2)$$

where E and D are the encoder and the decoder, and L is the loss used for training of the autoencoder. The loss is implemented as the distance (dist) between the original and the reconstructed embeddings. The dist can be any distance or similarity measure such as MSE, KL-divergence, or cosine distance. In our preliminary study, the cosine distance showed the best results, so we use it in our experiments.

Let us consider two embedding models $s_1(w)$ and $s_2(w)$. In such a case, the input to each decoder is

not the result of the corresponding encoder, but meta-embeddings, which depends on the both encoders. Depending on the approach, meta-embeddings can be built in different ways, we construct the meta-embeddings as follows. In case of CAEME, we take an L_2 -normalised concatenation of the two source embeddings encoded with respective encoders $E_1(s_1(w))$ and $E_2(s_2(w))$:

$$m(w) = \frac{E_1(s_1(w)) \oplus E_2(s_2(w))}{\|E_1(s_1(w)) \oplus E_2(s_2(w))\|_2}, \quad (3)$$

where \oplus is the concatenation operation.

The drawback of this model is the growing dimensionality of meta-embeddings for cases where we combine multiple source embeddings. To fight that, we can replace the concatenation operation with averaging, yielding AAEME. It computes meta-embedding of a word w from its two source embeddings $s_1(w)$ and $s_2(w)$ as the L_2 -normalised sum of internal representations $E_1(s_1(w))$ and $E_2(s_2(w))$:

$$m(w) = \frac{E_1(s_1(w)) + E_2(s_2(w))}{\|E_1(s_1(w)) + E_2(s_2(w))\|_2}. \quad (4)$$

In CAEME, the dimensionality of the meta-embedding space is the sum of the dimensions of the source embeddings, whereas in AAEME it stays the same. The AAEME encoder can be seen as a special case of the CAEME encoder where the meta-embedding is computed by averaging the two encoded sources in equation 3 instead of their concatenation.

7.3. Training of Autoencoders2

We can impose additional restrictions on *AEME models during training. One of such restrictions is the use of triplet loss. We restrict a word to be closer to the words that are semantically related to it according to the taxonomy than to a randomly chosen word with some *margin*:

$$L(w_a, w_p, w_n) = \max(\|m(w_a) - m(w_p)\| - \|m(w_a) - m(w_n)\| + \text{margin}, 0), \quad (5)$$

where $\|\cdot\|$ is a distance function, w_a is the target word, w_p and w_n are positive and negative words, respectively.

The algorithm of calculating triplet loss is as follows:

1. for each word presented in the taxonomy, we compile a list of semantically related words which includes synonyms, hyponyms and hypernyms;
2. at each epoch, we randomly select K positive words from this related words set and form a set of K negative words by selecting them randomly from the vocabulary;
3. if the word is not presented in the taxonomy, then we cannot form a list of related words for it. In this case, we generate positive vectors for it by adding random noise to its vector;
4. next, we calculate the triplet margin loss by combining the triplet loss with the original loss as $\alpha * loss + (1 - \alpha) * triplet_loss$.

We use the following parameters for the triplet loss: $K = 5$, $margin = 0.1$, $alpha = 0.005$. These parameters were selected via grid search with AAEME algorithm on the English 1.7 dataset.

8. Experiments

In this section, we report and discuss the performance of our models in the taxonomy enrichment task. We experiment with our DWRank approach and its modifications DWRank-Graph and DWRank-Meta. In addition to that, we compare them with the baseline introduced in RUSSE’2020 shared task and with a number of state-of-the-art methods. We conduct the experiments with English and Russian wordnets.

8.1. Experimental Setup

For each result we add the standard deviation values. We calculate them as follows. We randomly sample 80% of the test data and calculate the MAP scores on that part of the test set. We repeat the same procedure 30 times and then calculate the standard deviation on those 30 MAP values.

The MAP metric should be interpreted as follows: the higher the score, the better the results. For instance, $MAP@k = 1.0$ means that all of the N correct hypernyms are present in the first $top-N$ positions in the ranked list of candidates. Moreover, with $MAP@k = 1.0$ the first candidate is always correct. $MAP@k \geq 0.5$ means that at least one of the correct hypernyms is present in the two first positions ($top-2$) in the ranked lists of candidates. $MAP@k \geq 0.3$ means that at least one of the correct hypernyms is present in

the first three positions ($top-3$) in the ranked lists of candidates.

We show the performance of different methods on the nouns attribution task for English (nouns 1.6) in Figure 2 and for Russian (non-restricted nouns) in Figure 3. The X axis shows the MAP score for each method, the methods are listed along the Y axis. For DWRank-Meta models, we list the embeddings used in the model in brackets. The colors of bars in figures correspond to different types of input embeddings. The orange color stands for the vanilla DWRank – DWRank which uses only distributional embeddings. Purple denotes the DWRank-Graph variants. For the DWRank-Meta approaches exploiting only distributional embeddings we use pink, and DWRank-Meta on word and graph embeddings is denoted with the bright green color. Previous SOTA approaches are shown in yellow.

The full results for all experiments on the English and Russian datasets can be seen in Appendix A in Tables 12 and 13, respectively. Here, when listing embeddings used in DWRank-Meta models, we use the shortcut “words” to denote the combination of fastText, word2vec, and GloVe embeddings.

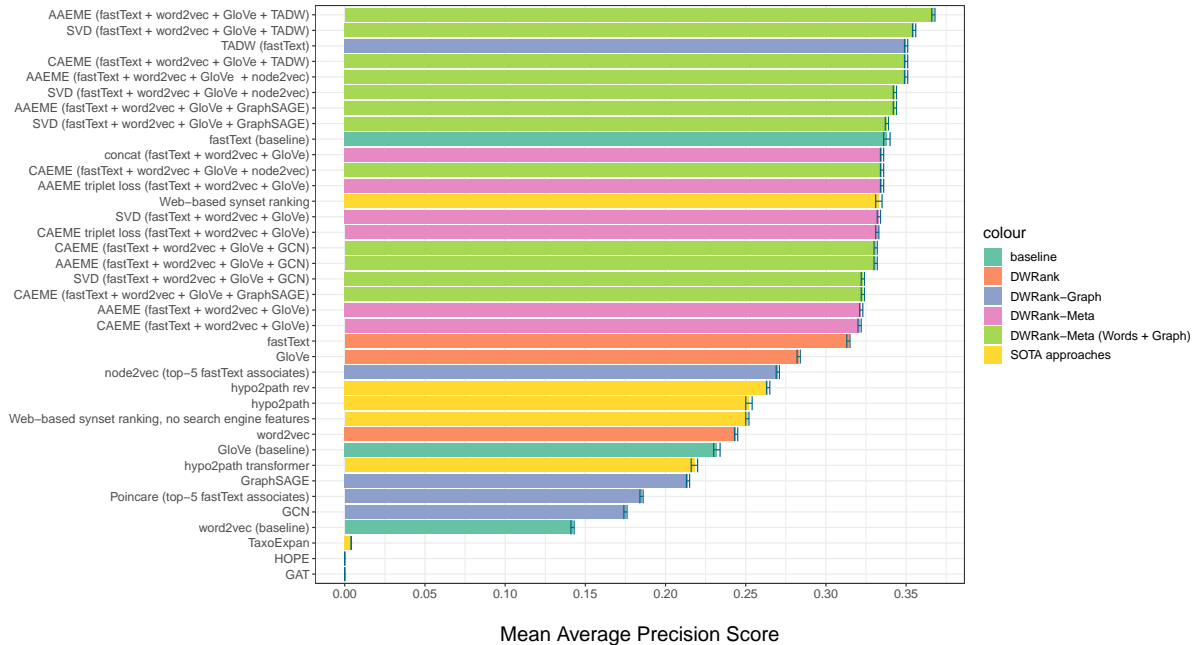
8.2. Results

DWRank-Meta Figures 2 and 3 show that the leaderboard for both English and Russian nouns is dominated by DWRank-Meta models. While English benefits from the union of distributional and graph embeddings, for Russian distributional embeddings alone perform on par with their combinations with graph embeddings. Besides that, high-performing variants of DWRank-Meta for English feature TADW, node2vec, and GraphSAGE, whereas for Russian TADW is the only graph embedding model which does not decrease the scores of DWRank-Meta.

We see that triplet loss significantly improves the results for DWRank-Meta models (cf. AAEME/CAEME with and without triplet loss) for both English and Russian.

DWRank-Graph On the other hand, DWRank-Graph fails in the task of taxonomy extension for all datasets. TADW model is the only graph embedding model which can compete with DWRank-Meta models. This can be explained by the fact that TADW is an extended version of DeepWalk and applies the skip-gram model with the pre-trained fastText representations. In contrast to that, the other graph models suffer from the noisy representations of OOV query words.

Figure 2. Comparison of the method performance on *nouns_1.6* dataset for English. Each colour denotes the method type and the embeddings type used.



At the same time, despite the success of TADW, it does not outperform models based solely on distributional embeddings, showing that graph representations apparently do not contribute any information which is not already contained in distributional word vectors.

Baselines⁴ We also notice that for both languages the baselines are quite competitive. They are substantially worse than the best-performing models, but they are much simpler to implement and is fast and easy to train. Therefore, we suggest that it should be preferred in the situation of limited resources and time.

However, the choice of embedding models is crucial for the baselines (as well as for the vanilla DWRank which performs closely). We see that fastText outperforms word2vec and GloVe embeddings for almost all languages and datasets. The low scores of GloVe and word2vec embeddings on baseline and DWRank methods can be explained by data coverage issues. Fixed vocabularies of word2vec and GloVe do not allow generating any representation for missing query words, whereas fastText can handle them.

SOTA models⁴ Neither of SOTA models managed to outperform the fastText baseline or approach the best DWRank-Meta variants. Web-based synset ranking (WBSR) model shows that the information from online search engines and Machine Translation mod-

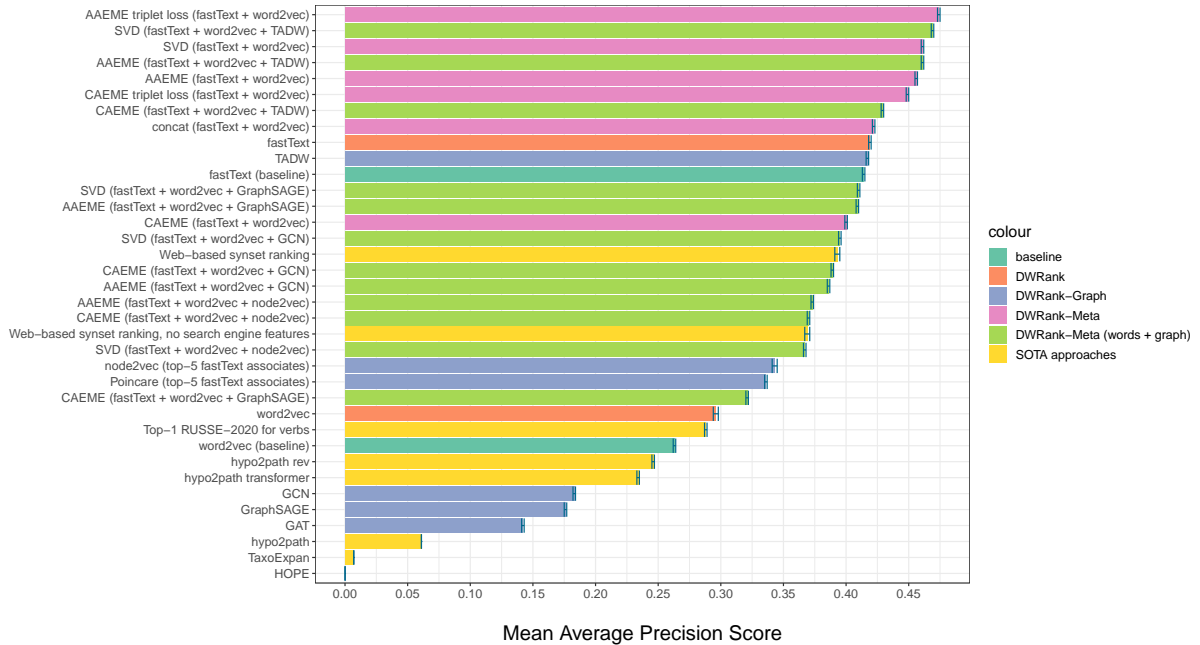
els is beneficial for the task – its performance without this information drops dramatically. However, this information is not enough to outperform the word embedding-based models.

The performance of hypo2path model is even lower than that of WBSR. Being an autoregressive generative model, it is very sensitive to its own mistakes. Generating one senseless hypernym can ruin all the following chain. Conversely, when starting with the root hypernym “entity.n.01”, it often takes a wrong path. Finally, TaxoExpan model relies on definitions of words which we did not provide in this task. Therefore, its results are close to zero. We do not consider them credible and provide them in italics.

Performance for different datasets⁴ Figures 2 and 3 as well as the results in the appendix show that there is no single best-performing model. While DWRank-Meta is almost always the best, for different datasets different variations of this model are the most successful. The results are usually consistent for the same language and part of speech (e.g. for different versions of English nouns datasets the best-performing model is the same), but there are exceptions to this regularity.

Interpretable Evaluation⁴ MAP metric which is the standard way of evaluating taxonomy enrichment

Figure 3. Performance of different models on the Russian non-restricted dataset. Each colour denotes the method type and the embeddings type used.



models has a serious drawback. Namely, it is not interpretable, which hampers the understanding of the models' performance.

Therefore, in addition to MAP we report the Precision@k score which can be interpreted as the ratio of correct synsets in the *top-k* outputs of a model. We evaluate our best systems automatically with the Precision@k ($k = 1, 2, 3$) score. The choice of the values of k is explained by the fact that the average number of true ancestors is 2 for the English words and 3 for the Russian words. Thus, Precision@k for $k > 3$ will be unfairly understated, because there will always be at most 3 correct answers out of k . This means that for the $k = 4$ the maximum is 0.75, for $k = 5$ it is 0.6, etc.

Table 5 shows the Precision@k scores for the best performing English and Russian models on the nouns datasets. Both of them are DWRank-Meta models with AAEME autoencoders. The English model uses three types of distributional embeddings and TADW graph embeddings, while the Russian model uses only fastText and word2vec but benefits from the triplet loss. We see that Precision@k is particularly high for Russian. There, over a half of generated lists contain a correct synset in the first position. This shows that DWRank-Meta can successfully be used as a helper tool for taxonomy extension.

The results for English are lower. However, this should not be considered as a sign of lower performance of models for English. The Russian and English datasets consist of different words, so they cannot be directly compared.

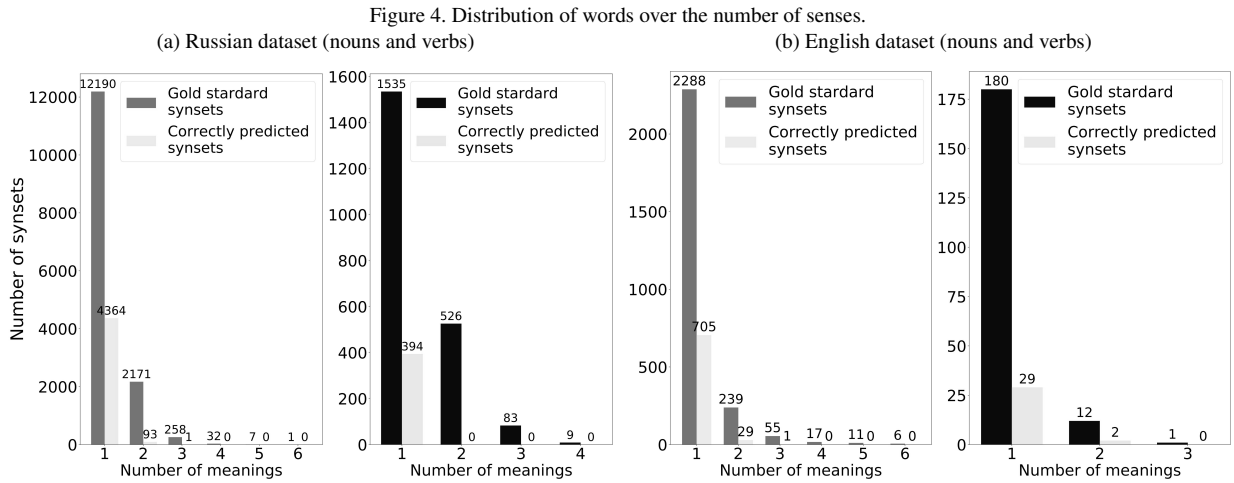
9. Error Analysis1

To better understand the difference in systems performance and their main difficulties, we made a quantitative and qualitative analysis of the results.

9.1. Comparison of Graph-based Approaches with Word Vector Baselines2

First of all, we wanted to know to what extent the set of correct answers of graph-based models overlaps with the one of fastText-based models. In other words, we would like to know if the graph representations are able to discover hypernymy relations which could not be identified by word embeddings.

Therefore, for each new word we computed average precision (*AP*) score and compared those scores across different approaches. We found that at least 90% words for which fastText failed to identify correct hypernyms (i.e. words with $AP = 0$) also have



Method	Pr@1	Pr@2	Pr@3
English nouns 1.6-3.0			
baseline	0.260	0.184	0.144
DWRank	0.245	0.170	0.135
DWRank-Meta	0.288	0.185	0.143
DWRank-Graph	0.278	0.189	0.150
DWRank-Meta (Words + Graph)	0.311	0.199	0.154
English verbs 1.6-3.0			
baseline	0.173	0.126	0.101
DWRank	0.260	0.169	0.126
DWRank-Meta	0.238	0.168	0.131
DWRank-Graph	0.238	0.158	0.118
DWRank-Meta (Words + Graph)	0.259	0.164	0.124
Russian non-restricted nouns			
baseline	0.346	0.228	0.171
DWRank	0.347	0.228	0.172
DWRank-Meta	0.396	0.257	0.196
DWRank-Graph	0.347	0.224	0.168
DWRank-Meta (Words + Graph)	0.397	0.255	0.192
Russian non-restricted verbs			
baseline	0.251	0.181	0.139
DWRank	0.282	0.196	0.154
DWRank-Meta	0.368	0.245	0.190
DWRank-Graph	0.274	0.191	0.149
DWRank-Meta (Words + Graph)	0.341	0.231	0.180

Table 5

Precision@k for the best-performing models for the English and Russian nouns datasets.

the *AP* of 0 in all the graph-based models. This means that if fastText cannot provide correct hypernyms for a word, other models cannot help either. Moreover, all words which are correctly predicted by graph-based approaches, are also correctly predicted by fastText. Moreover, only 8% to 55% words correctly predicted by fastText are also correctly predicted by any of the graph-based models. At the same time, the number of cases where graph-based models perform better than fastText is very low (3–5% cases). Thus, combining them cannot improve the performance significantly. This observation is corroborated by the scores of the combined models.

To contrast the performance of the text and graph embeddings and to demonstrate the input and the output formats of the models we present Tables 7 and 8 in Appendix A. They demonstrate the main features of the tested approaches. The examples do not pretend to be the general case example, however, they do provide the idea about ranking of the results and the performance of text, graph and fusion embedding types.

9.2. Performance on Polysemous Words²

The differences in word semantics make the dataset uneven. In addition to that, we would also like to understand whether the performance of models depends on the number of connected components (possible meanings) for each word. Thus, we examine how many words with more than one meaning can be predicted by the system.

Figure 4 depicts the distribution of synsets over the number of senses they convey. As we can see, the

vast majority of words are monosemous. For Russian nouns, the system correctly identifies almost half of them, whereas for other datasets the share of correctly predicted monosemous words is below 30%. This stems from the fact that for distributional models it is difficult to capture multiple senses in one vector. They usually capture the most widespread sense of a word. Therefore, the number of predicted synsets with two or more senses is extremely low. A similar power law distribution would be obtained using BERT embeddings, as we are still averaging embeddings from all contexts. This may be one of the reasons why contextualised models did not perform better than the fast-Text models which capture the main meaning only but do it well.

9.3. Error Types2

In order to understand why a large number of word hypernyms (at least 60%) are too difficult for models to predict, we turn to manual analysis of the system outputs. We find out that errors can be divided into two groups: system errors caused by distributional models limitations and taxonomy inaccuracies. Therefore, we come across five main error types:

Type 1. Extracted nearest neighbours can be semantically related words but not necessary co-hyponyms:

- delist (WordNet); expected senses: get rid of; predicted senses: remove, delete;
- хэштег (hashtag, RuWordNet); expected senses: отличительный знак, пометка (tag, label); predicted senses: символ, короткий текст (symbol, short text).

Type 2. Distributional models are unable to predict multiple senses for one word:

- latakia (WordNet); expected senses: tobacco; municipality city; port, geographical point; predicted senses: tobacco;
- запорожец (zaporozhets, RuWordNet); expected senses: житель города (citizen, resident); марка автомобиля, автомобиль (car brand, car); predicted senses: автомобиль, мототранспортное средство, марка автомобиля (car, motor car, car brand).

Type 3. System predicts too broad / too narrow concepts:

- midweek (WordNet); expected senses: day of the week, weekday; predicted senses: time period, week, day, season;

- медянка (smooth snake, RuWordNet); expected senses: неядовитая змея, уж (non-venomous snake, grass snake); predicted senses: змея, рептилия, животное (snake, reptile, animal).

Type 4. Incorrect word vector representation: nearest neighbours are semantically far from the meaning of the input word:

- falanga (WordNet); expected senses: persecution, torture; predicted senses: fish, bean, tree, wood.;
- кубокилометр (cubic kilometer, RuWordNet); expected senses: единица объема, единица измерения (unit of capacity, unit of measurement); predicted senses: город, городское поселение, кубковое соревнование, спортивное соревнование (city, settlement, competition, sports contest).

Type 5. Unaccounted senses in the gold standard datasets, inaccuracies in the manual annotation:

- emeritus (WordNet); expected senses: retiree, non-worker; predicted senses: professor, academician;
- сепия (sepia, RuWordNet); expected senses: морской моллюск “sea mollusc”; predicted senses: цвет, краситель (color, dye).

In order to check how useful the predicted synsets are for a human annotator (i.e. if a short list of possible hypernyms can speed up the manual extension of a taxonomy), we conduct the manual evaluation of 10 random nouns and 10 random verbs for both languages (the words are listed in Table 6). We focus on worse-quality cases and thus select words whose MAP score is below 1. Annotators with the expertise in the field and the knowledge of English and Russian were provided with guidelines and asked to evaluate the outputs from our best-performing system. Each word was labelled by 4 expert annotators, Fleiss’s kappa is 0.63 (substantial agreement) for both datasets.

We compute Precision@k score (the share of correct answers in the generated lists from position 1 to k) for k from 1 to 10, shown in Figure 5. We can see that even for words with MAP below 1 our model manages to extract useful hypernyms.

10. Conclusions1

In this work, we performed a large-scale computational study of various methods for taxonomy

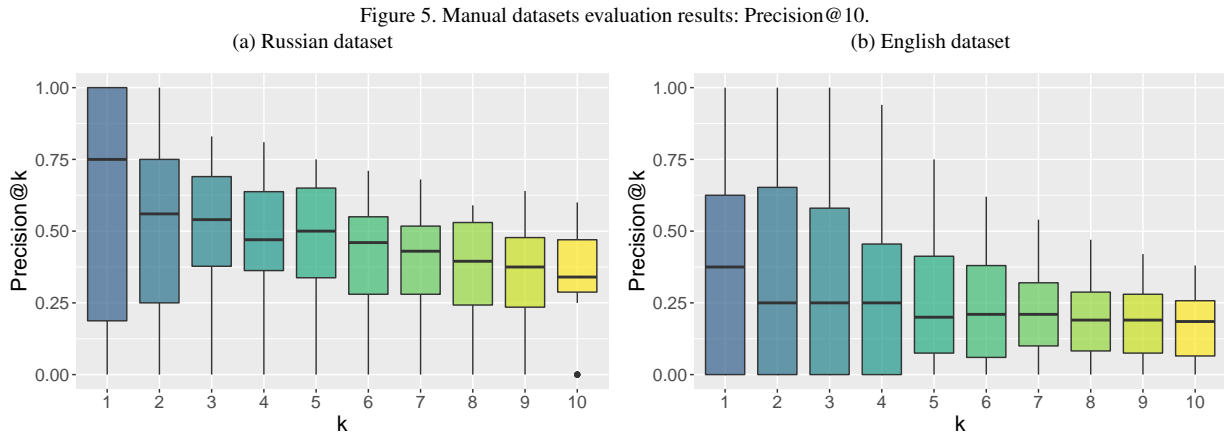


Table 6
Words selected for the manual evaluation.

Language	Word List
English	falanga, venerability, ambulatory, emeritus, salutatory address, eigenvalue of a matrix, liposuction, moppet, dinette, snoek, to fancify, to google, to expense, to porcelainize, to junketeer, to delist, to podcast, to deglaze, to shoetree, to headquarter
Russian	барабашка, листинг, стихосложение, аукционист, точилка, гиперреализм, серология, огрызок, фен, марикультура, уломать, отфотошопить, тяпнуть, растушевать, завраться, леветь, мозолить, загоститься, распеваться, оплавить

enrichment. We also presented datasets for studying diachronic evolution of wordnets for English and Russian, extending the monolingual setup of the RUSSE’2020 shared task [9] with a larger Russian dataset and similar English versions.

We presented a new taxonomy enrichment method called DWRank, which combines distributional information and the information extracted from Wiktionary outperforming the baseline method from [6] on the English datasets. We also presented its extensions: DWRank-Graph and DWRank-Meta which use graph and meta-embeddings via a common interface.

We also explored the benefits of meta-embeddings (combinations of embeddings) and graph embeddings for the task of taxonomy enrichment. On the Russian datasets DWRank-Meta performed best using fastText and word2vec word embeddings. For the English dataset the combination of word (fastText, word2vec and GloVe) and graph (TADW) embeddings demonstrated the best performance.

In this paper we also presented WBSR — a method for taxonomy extension which leverages the informa-

tion from the Web. This approach was the current state of the art in the task for the Russian language. Now it lags significantly behind the new DWRank-meta approaches.

According to our experiments, word vector representations are simple, powerful, and extremely effective instrument for taxonomy enrichment, as the contexts (in a broad sense) extracted from the pre-trained word embeddings (fastText, word2vec, GloVe) and their combination are sufficient to attach new words to the taxonomy. TADW embeddings are also useful and efficient for the taxonomy enrichment task and in combination with the fastText, word2vec and GloVe approaches demonstrate SOTA results for the English language and compatible results for the Russian language.

Error analysis also reveals that the correct synsets identified by graph-based models are usually retrieved by the fastText-based model alone. This makes graphs representations mostly irrelevant and excessive. Nonetheless, there exist cases where graph representa-

tions were able to identify correctly some hypernyms which were not captured by fastText.

Despite the mixed results of the application of graph-based methods, we propose further exploration of the graph-based features as the existing resource contains principally different and complementary information to the distributional signal contained in text corpora. One way to improve their performance, may be to use more sophisticated non-linear projection transformations from word to graph embeddings. Another promising way in our opinion is to explore other types of meta-embeddings to mix word and graph signals, e.g. GraphGlove [84]. Moreover, we find it promising to experiment with temporal embeddings such as those of [85] for the taxonomy enrichment task.

Last but not least, we plan to explore methods that do not rely on the set of pre-defined candidates for inclusion in a taxonomy, as generation and mining of such a set may be a challenging problem in its own.

Acknowledgments

The participation of Mikhail Tikhomirov in the reported study (experiments with meta-embeddings) was funded by RFBR, project number 19-37-90119. The work of Natalia Loukachevitch (preparation of RuWordNet data for the experiments, mappings of the RuWordNet synsets to English WordNet Linked Open Data (LOD) Inter-Lingual Index (ILI)) is supported by the Russian Science Foundation (project 20-11-20166). The work of Alexander Panchenko, Varvara Logacheva, and Irina Nikishina was conducted in the framework of joint MTS-Skoltech laboratory.

References

- [1] T. Berners-Lee, J. Hendler and O. Lassila, The semantic web, *Scientific american* **284**(5) (2001), 34–43.
- [2] A. Gómez-Pérez and O. Corcho, Ontology languages for the semantic web, *IEEE Intelligent systems* **17**(1) (2002), 54–60.
- [3] D. Jurgens and M.T. Pilehvar, SemEval-2016 Task 14: Semantic Taxonomy Enrichment, in: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, Association for Computational Linguistics, San Diego, California, 2016, pp. 1092–1102. doi:10.18653/v1/S16-1169. <https://www.aclweb.org/anthology/S16-1169>.
- [4] H. Tanev and A. Rotondi, Defcor at SemEval-2016 Task 14: Taxonomy enrichment using definition vectors, in: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, Association for Computational Linguistics, San Diego, California, 2016, pp. 1342–1345. doi:10.18653/v1/S16-1210. <https://www.aclweb.org/anthology/S16-1210>.
- [5] L. Espinosa-Anke, F. Ronzano and H. Saggion, TALN at SemEval-2016 Task 14: Semantic Taxonomy Enrichment Via Sense-Based Embeddings, in: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, Association for Computational Linguistics, San Diego, California, 2016, pp. 1332–1336. doi:10.18653/v1/S16-1208. <https://www.aclweb.org/anthology/S16-1208>.
- [6] I. Nikishina, A. Panchenko, V. Logacheva and N. Loukachevitch, Studying Taxonomy Enrichment on Diachronic WordNet Versions, in: *Proceedings of the 28th International Conference on Computational Linguistics*, Association for Computational Linguistics, Barcelona, Spain, 2020.
- [7] I. Nikishina, N. Loukachevitch, V. Logacheva and A. Panchenko, Exploring Graph-based Representations for Taxonomy Enrichment, in: *Proceedings of the 11th Global Wordnet Conference*, Global Wordnet Association, University of South Africa (UNISA), 2021, pp. 126–136. <https://www.aclweb.org/anthology/2021.gwc-1.15>.
- [8] M. Tikhomirov and N. Loukachevitch, Meta-Embeddings in Taxonomy Enrichment Task, in: *Computational Linguistics and Intellectual Technologies: papers from the Annual conference "Dialogue"*, 2021.
- [9] I. Nikishina, V. Logacheva, A. Panchenko and N. Loukachevitch, RUSSE'2020: Findings of the First Taxonomy Enrichment Task for the Russian Language, in: *Computational Linguistics and Intellectual Technologies: papers from the Annual conference "Dialogue"*, 2020.
- [10] G. De Giacomo and M. Lenzerini, TBox and ABox Reasoning in Expressive Description Logics, 1996, pp. 37–48.
- [11] G.A. Miller, *WordNet: An electronic lexical database*, MIT press, 1998.
- [12] D.L. McGuinness and A. Borgida, Explaining subsumption in description logics, *IJCAI (1)* **3** (1995).
- [13] M. Nickel, K. Murphy, V. Tresp and E. Gabrilovich, A Review of Relational Machine Learning for Knowledge Graphs, *Proceedings of the IEEE* **104** (2015). doi:10.1109/JPROC.2015.2483592.
- [14] H. Paulheim, Knowledge graph refinement: A survey of approaches and evaluation methods, *Semantic web* **8**(3) (2017), 489–508.
- [15] A. Bordes, N. Usunier, A. Garcia-Durán, J. Weston and O. Yakhnenko, Translating Embeddings for Modeling Multi-Relational Data, in: *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13*, Curran Associates Inc., Red Hook, NY, USA, 2013, pp. 2787–2795–.
- [16] B. Shi and T. Weninger, ProjE: Embedding Projection for Knowledge Graph Completion, in: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, AAAI'17, AAAI Press, 2017, pp. 1236–1242–.
- [17] W. Yu, C. Zheng, W. Cheng, C.C. Aggarwal, D. Song, B. Zong, H. Chen and W. Wang, Learning Deep Network Representations with Adversarially Regularized Autoencoders, in: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, Association for Computing Ma-

- chinery, New York, NY, USA, 2018, pp. 2663–2671–. ISBN 9781450355520. doi:10.1145/3219819.3220000.
- [18] N. Li, Z. Bouraoui and S. Schockaert, Ontology Completion Using Graph Convolutional Networks, in: *The Semantic Web – ISWC 2019*, C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. Cruz, A. Hogan, J. Song, M. Lefrançois and F. Gandon, eds, Springer International Publishing, Cham, 2019, pp. 435–452. ISBN 978-3-030-30793-6.
- [19] T. Dettmers, P. Minervini, P. Stenetorp and S. Riedel, Convolutional 2d knowledge graph embeddings, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32, 2018.
- [20] C. Shang, Y. Tang, J. Huang, J. Bi, X. He and B. Zhou, End-to-end structure-aware convolutional networks for knowledge base completion, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, 2019, pp. 3060–3067.
- [21] I. Balazevic, C. Allen and T. Hospedales, TuckER: Tensor Factorization for Knowledge Graph Completion, 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)*, Association for Computational Linguistics, Hong Kong, China, 2019, pp. 5185–5194. doi:10.18653/v1/D19-1522. <https://www.aclweb.org/anthology/D19-1522>.
- [22] T. Lacroix, N. Usunier and G. Obozinski, Canonical Tensor Decomposition for Knowledge Base Completion, in: *Proceedings of the 35th International Conference on Machine Learning*, J. Dy and A. Krause, eds, Proceedings of Machine Learning Research, Vol. 80, PMLR, Stockholmsmässan, Stockholm Sweden, 2018, pp. 2863–2872. <http://proceedings.mlr.press/v80/lacroix18a.html>.
- [23] T. Lacroix, G. Obozinski and N. Usunier, Tensor Decompositions for Temporal Knowledge Base Completion, in: *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, OpenReview.net, 2020. <https://openreview.net/forum?id=rke2P1BFwS>.
- [24] P. Cimiano and J. Völker, Text2Onto: A Framework for Ontology Learning and Data-Driven Change Discovery, in: *Proceedings of the 10th International Conference on Natural Language Processing and Information Systems, NLDB’05*, Springer-Verlag, Berlin, Heidelberg, 2005, pp. 227–238–. ISBN 3540260315. doi:10.1007/11428817_21.
- [25] D. Dessi, F. Osborne, D. Reforgiato Recupero, D. Buscaldi and E. Motta, Generating knowledge graphs by employing Natural Language Processing and Machine Learning techniques within the scholarly domain, *Future Generation Computer Systems* **116** (2021), 253–264. doi:<https://doi.org/10.1016/j.future.2020.10.026>. <https://www.sciencedirect.com/science/article/pii/S0167739X2033003X>.
- [26] F. Osborne and E. Motta, Klink-2: Integrating Multiple Web Sources to Generate Semantic Topic Networks, 2015. ISBN 978-3-319-25006-9. doi:10.1007/978-3-319-25007-6_24.
- [27] J. Camacho-Collados, C. Delli Bovi, L. Espinosa-Anke, S. Oramas, T. Pasini, E. Santus, V. Shwartz, R. Navigli and H. Saggion, SemEval-2018 Task 9: Hypernym Discovery, in: *Proceedings of The 12th International Workshop on Semantic Evaluation*, Association for Computational Linguistics, New Orleans, Louisiana, 2018, pp. 712–724. doi:10.18653/v1/S18-1115. <https://www.aclweb.org/anthology/S18-1115>.
- [28] G. Bordea, P. Buitelaar, S. Faralli and R. Navigli, SemEval-2015 Task 17: Taxonomy Extraction Evaluation (TEEval), in: *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, Association for Computational Linguistics, Denver, Colorado, 2015, pp. 902–910. doi:10.18653/v1/S15-2151. <https://www.aclweb.org/anthology/S15-2151>.
- [29] G. Bordea, E. Lefever and P. Buitelaar, SemEval-2016 Task 13: Taxonomy Extraction Evaluation (TEEval-2), in: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, Association for Computational Linguistics, San Diego, California, 2016, pp. 1081–1091. doi:10.18653/v1/S16-1168. <https://www.aclweb.org/anthology/S16-1168>.
- [30] P. Velardi, S. Faralli and R. Navigli, OntoLearn Reloaded: A Graph-Based Algorithm for Taxonomy Induction, *Computational Linguistics* **39**(3) (2013), 665–707. <https://www.aclweb.org/anthology/J13-3007>.
- [31] N.V. Loukachevitch, G. Lashevich, A.A. Gerasimova, V.V. Ivanov and B.V. Dobrov, Creating Russian wordnet by conversion, in: *Computational Linguistics and Intellectual Technologies: papers from the Annual conference “Dialogue”, 2016*, pp. 405–415.
- [32] M. Kunilovskaya, A. Kutuzov and A. Plum, Taxonomy Enrichment: Linear Hyponym-Hypernym Projection vs Synset ID Classification, in: *Computational Linguistics and Intellectual Technologies: papers from the Annual conference “Dialogue”, 2020*.
- [33] D. Dale, A simple solution for the Taxonomy enrichment task: Discovering hypernyms using nearest neighbor search, in: *Computational Linguistics and Intellectual Technologies: papers from the Annual conference “Dialogue”, 2020*.
- [34] N. Arefyev, M. Fedoseev, A. Kabanov and V. Zizov, Word2vec not dead: predicting hypernyms of co-hyponyms is better than reading definitions, in: *Computational Linguistics and Intellectual Technologies: papers from the Annual conference “Dialogue”, 2020*.
- [35] M. Tikhomirov, N. Loukachevitch and E. Parkhomenko, Combined Approach to Hypernym Detection for Thesaurus Enrichment, in: *Computational Linguistics and Intellectual Technologies: papers from the Annual conference “Dialogue”, 2020*.
- [36] G. Bernier-Colborne and C. Barrière, CRIM at SemEval-2018 Task 9: A Hybrid Approach to Hypernym Discovery, in: *Proceedings of The 12th International Workshop on Semantic Evaluation*, Association for Computational Linguistics, New Orleans, Louisiana, 2018, pp. 725–731. doi:10.18653/v1/S18-1116. <https://www.aclweb.org/anthology/S18-1116>.
- [37] G. Berend, M. Makrai and P. Földiák, 300-sparsans at SemEval-2018 Task 9: Hypernymy as interaction of sparse attributes, in: *Proceedings of The 12th International Workshop on Semantic Evaluation*, Association for Computational Linguistics, New Orleans, Louisiana, 2018, pp. 928–934. doi:10.18653/v1/S18-1152. <https://www.aclweb.org/anthology/S18-1152>.
- [38] A. Maldonado and F. Klubička, ADAPT at SemEval-2018 Task 9: Skip-Gram Word Embeddings for Unsupervised Hypernym Discovery in Specialised Corpora, in: *Proceedings of The 12th International Workshop on Semantic Evaluation*, Association for Computational Linguistics, New Or-

- leans, Louisiana, 2018, pp. 924–927. doi:10.18653/v1/S18-1151. <https://www.aclweb.org/anthology/S18-1151>.
- [39] J. Pennington, R. Socher and C. Manning, GloVe: Global Vectors for Word Representation, in: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Doha, Qatar, 2014, pp. 1532–1543. doi:10.3115/v1/D14-1162. <https://www.aclweb.org/anthology/D14-1162>.
- [40] V. Shwartz, Y. Goldberg and I. Dagan, Improving Hypernymy Detection with an Integrated Path-based and Distributional Method, in: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Berlin, Germany, 2016, pp. 2389–2398. doi:10.18653/v1/P16-1226. <https://www.aclweb.org/anthology/P16-1226>.
- [41] J. Pocostales, NUIG-UNLP at SemEval-2016 Task 13: A Simple Word Embedding-based Approach for Taxonomy Extraction, in: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, Association for Computational Linguistics, San Diego, California, 2016, pp. 1298–1302. doi:10.18653/v1/S16-1202. <https://www.aclweb.org/anthology/S16-1202>.
- [42] R. Aly, S. Acharya, A. Ossa, A. Köhn, C. Biemann and A. Panchenko, Every Child Should Have Parents: A Taxonomy Refinement Algorithm Based on Hyperbolic Term Embeddings, in: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, 2019, pp. 4811–4817. doi:10.18653/v1/P19-1474. <https://www.aclweb.org/anthology/P19-1474>.
- [43] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer and V. Stoyanov, Unsupervised Cross-lingual Representation Learning at Scale, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Online, 2020, pp. 8440–8451. doi:10.18653/v1/2020.acl-main.747. <https://www.aclweb.org/anthology/2020.acl-main.747>.
- [44] M. Bansal, K. Gimpel and K. Livescu, Tailoring continuous word representations for dependency parsing, in: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2014, pp. 809–815.
- [45] S. Chowdhury, C. Zhang, P.S. Yu and Y. Luo, Mixed Pooling Multi-View Attention Autoencoder for Representation Learning in Healthcare, *arXiv preprint arXiv:1910.06456* (2019).
- [46] J. Coates and D. Bollegala, Frustratingly Easy Meta-Embedding—Computing Meta-Embeddings by Averaging Source Word Embeddings, *arXiv preprint arXiv:1804.05262* (2018).
- [47] W. Yin and H. Schütze, Learning word meta-embeddings, in: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2016, pp. 1351–1360.
- [48] D. Bollegala and C. Bao, Learning word meta-embeddings by autoencoding, in: *Proceedings of the 27th international conference on computational linguistics*, 2018, pp. 1650–1661.
- [49] J.O. Neill and D. Bollegala, Meta-embedding as auxiliary task regularization, *arXiv preprint arXiv:1809.05886* (2018).
- [50] G.I. Winata, Z. Lin and P. Fung, Learning multilingual meta-embeddings for code-switching named entity recognition, in: *Proceedings of the 4th Workshop on Representation Learning for NLP (ReplANLP-2019)*, 2019, pp. 181–186.
- [51] I. Makarov, M. Makarov and D. Kiselev, Fusion of text and graph information for machine learning problems on networks, *PeerJ Computer Science* 7 (2021).
- [52] A. Grover and J. Leskovec, node2vec: Scalable Feature Learning for Networks, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
- [53] N. Liu, X. Huang, J. Li and X. Hu, On interpretation of network embedding via taxonomy induction, in: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 1812–1820.
- [54] Y. Luan, L. He, M. Ostendorf and H. Hajishirzi, Multi-Task Identification of Entities, Relations, and Coreference for Scientific Knowledge Graph Construction, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 3219–3232. doi:10.18653/v1/D18-1360. <https://www.aclweb.org/anthology/D18-1360>.
- [55] M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee and 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, New Orleans, Louisiana, 2018, pp. 2227–2237. doi:10.18653/v1/N18-1202. <https://www.aclweb.org/anthology/N18-1202>.
- [56] K. Han, P. Yang, S. Mishra and J. Diesner, WikiCSSH: extracting computer science subject headings from Wikipedia, in: *ADBS, TPDL and EDA 2020 Common Workshops and Doctoral Consortium*, Springer, 2020, pp. 207–218.
- [57] T.N. Kipf and M. Welling, Semi-supervised classification with graph convolutional networks, *arXiv preprint arXiv:1609.02907* (2016).
- [58] T.N. Kipf and M. Welling, Variational Graph Auto-Encoders, *NIPS Workshop on Bayesian Deep Learning* (2016).
- [59] A. Rossi, D. Firmani, A. Matinata, P. Merialdo and D. Barbosa, Knowledge Graph Embedding for Link Prediction: A Comparative Analysis, *arXiv preprint arXiv:2002.00819* (2020).
- [60] D. Pujary, C. Thorne and W. Aziz, Disease Normalization with Graph Embeddings, in: *Proceedings of SAI Intelligent Systems Conference*, Springer, 2020, pp. 209–217.
- [61] J. Shen, Z. Shen, C. Xiong, C. Wang, K. Wang and J. Han, TaxoExpan: Self-supervised Taxonomy Expansion with Position-Enhanced Graph Neural Network, in: *Proceedings of The Web Conference 2020*, 2020, pp. 486–497.
- [62] T.N. Kipf and M. Welling, Semi-Supervised Classification with Graph Convolutional Networks, in: *International Conference on Learning Representations (ICLR)*, 2017.
- [63] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Liò and Y. Bengio, Graph Attention Networks, *ICLR* (2018).
- [64] G.A. Miller, WordNet: a lexical database for English, *Communications of the ACM* 38(11) (1995), 39–41.
- [65] D. Schlechtweg, B. McGillivray, S. Hengchen, H. Dubossarsky and N. Tahmasebi, SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection, in: *Proceedings*

- of the Fourteenth Workshop on Semantic Evaluation, International Committee for Computational Linguistics, Barcelona (online), 2020, pp. 1–23. <https://www.aclweb.org/anthology/2020.semeval-1.1>.
- [66] F. Bond, P. Vossen, J. McCrae and C. Fellbaum, CILI: the Collaborative Interlingual Index, in: *Proceedings of the 8th Global WordNet Conference (GWC)*, Global Wordnet Association, Bucharest, Romania, 2016, pp. 50–57. <https://www.aclweb.org/anthology/2016.gwc-1.9>.
- [67] Y. Cho, J.D. Rodriguez, Y. Gao and K. Erk, Leveraging WordNet Paths for Neural Hypernym Prediction, in: *Proceedings of the 28th International Conference on Computational Linguistics*, International Committee on Computational Linguistics, Barcelona, Spain (Online), 2020, pp. 3007–3018. doi:10.18653/v1/2020.coling-main.268. <https://www.aclweb.org/anthology/2020.coling-main.268>.
- [68] I. Sutskever, O. Vinyals and Q.V. Le, Sequence to sequence learning with neural networks, *Advances in neural information processing systems* **27** (2014), 3104–3112.
- [69] T. Luong, H. Pham and C.D. Manning, Effective Approaches to Attention-based Neural Machine Translation, in: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Lisbon, Portugal, 2015, pp. 1412–1421. doi:10.18653/v1/D15-1166. <https://www.aclweb.org/anthology/D15-1166>.
- [70] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L.u. Kaiser and I. Polosukhin, Attention is All you Need, in: *Advances in Neural Information Processing Systems*, Vol. 30, I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan and R. Garnett, eds, Curran Associates, Inc., 2017. <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- [71] P. Bojanowski, E. Grave, A. Joulin and T. Mikolov, Enriching Word Vectors with Subword Information, *Transactions of the Association for Computational Linguistics* **5** (2017), 135–146.
- [72] T. Mikolov, I. Sutskever, K. Chen, G.S. Corrado and J. Dean, Distributed Representations of Words and Phrases and their Compositionality, in: *Advances in Neural Information Processing Systems* **26**, C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani and K.Q. Weinberger, eds, Curran Associates, Inc., 2013, pp. 3111–3119.
- [73] M. Fares, A. Kutuzov, S. Oepen and E. Velldal, Word vectors, reuse, and replicability: Towards a community repository of large-text resources, in: *Proceedings of the 21st Nordic Conference on Computational Linguistics*, Association for Computational Linguistics, Gothenburg, Sweden, 2017, pp. 271–276. <https://www.aclweb.org/anthology/W17-0237>.
- [74] A. Kutuzov and E. Kuzmenko, *WebVectors: A Toolkit for Building Web Interfaces for Vector Semantic Models*, in: *Analysis of Images, Social Networks and Texts: 5th International Conference, AIST 2016, Yekaterinburg, Russia, April 7-9, 2016, Revised Selected Papers*, D.I. Ignatov, M.Y. Khachay, V.G. Labunets, N. Loukachevitch, S.I. Nikolenko, A. Panchenko, A.V. Savchenko and K. Vorontsov, eds, Springer International Publishing, Cham, 2017, pp. 155–161. ISBN 978-3-319-52920-2. doi:10.1007/978-3-319-52920-2_15.
- [75] M. Straka and J. Straková, Tokenizing, POS Tagging, Lemmatizing and Parsing UD 2.0 with UDPipe, in: *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 88–99. <http://www.aclweb.org/anthology/K/K17/K17-3009.pdf>.
- [76] M. Nickel and D. Kiela, Poincaré Embeddings for Learning Hierarchical Representations, in: *Advances in Neural Information Processing Systems* **30**, I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan and R. Garnett, eds, Curran Associates, Inc., 2017, pp. 6341–6350.
- [77] C. Gülcühre, M. Denil, M. Malinowski, A. Razavi, R. Pascanu, K. Hermann, P. Battaglia, V. Bapst, D. Raposo, A. Santoro and N.D. Freitas, Hyperbolic Attention Networks, *arXiv preprint arXiv:1805.09786* (2019).
- [78] T. Mikolov, K. Chen, G. Corrado and J. Dean, Efficient Estimation of Word Representations in Vector Space, in: *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*, Y. Bengio and Y. LeCun, eds, 2013. <http://arxiv.org/abs/1301.3781>.
- [79] W.L. Hamilton, R. Ying and J. Leskovec, Inductive representation learning on large graphs, in: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 1025–1035.
- [80] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L.u. Kaiser and I. Polosukhin, Attention is All you Need, in: *Advances in Neural Information Processing Systems*, Vol. 30, I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan and R. Garnett, eds, Curran Associates, Inc., 2017. <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- [81] C. Yang, Z. Liu, D. Zhao, M. Sun and E. Chang, Network representation learning with rich text information, in: *Twenty-fourth international joint conference on artificial intelligence*, 2015.
- [82] B. Perozzi, R. Al-Rfou and S. Skiena, Deepwalk: Online learning of social representations, in: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2014, pp. 701–710.
- [83] M. Ou, P. Cui, J. Pei, Z. Zhang and W. Zhu, Asymmetric transitivity preserving graph embedding, in: *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 1105–1114.
- [84] M. Ryabinin, S. Popov, L. Prokhorenkova and E. Voita, Embedding Words in Non-Vector Space with Unsupervised Graph Learning, *arXiv preprint arXiv:2010.02598* (2020).
- [85] R. Goel, S.M. Kazemi, M. Brubaker and P. Poupart, Diachronic Embedding for Temporal Knowledge Graph Completion, *Proceedings of the AAAI Conference on Artificial Intelligence* **34**(04) (2020), 3988–3995. doi:10.1609/aaai.v34i04.5815. <https://ojs.aaai.org/index.php/AAAI/article/view/5815>.
- [86] H. Hanke and D. Knees, A phase-field damage model based on evolving microstructure, *Asymptotic Analysis* **101** (2017), 149–180.
- [87] E. Lefever, A hybrid approach to domain-independent taxonomy learning, *Applied Ontology* **11**(3) (2016), 255–278.

- [88] P.S. Meltzer, A. Kallioniemi and J.M. Trent, Chromosome alterations in human solid tumors, in: *The Genetic Basis of Human Cancer*, B. Vogelstein and K.W. Kinzler, eds, McGraw-Hill, New York, 2002, pp. 93–113.
- [89] P.R. Murray, K.S. Rosenthal, G.S. Kobayashi and M.A. Pfaller, *Medical Microbiology*, 4th edn, Mosby, St. Louis, 2002.
- [90] E. Wilson, Active vibration analysis of thin-walled beams, PhD thesis, University of Virginia, 1991.
- [91] A.V. Aho and J.D. Ullman, *The Theory of Parsing, Translation and Compiling*, Vol. 1, Prentice-Hall, Englewood Cliffs, NJ, 1972.
- [92] G. Andrew and J. Gao, Scalable training of L1-regularized log-linear models, in: *Proceedings of the 24th International Conference on Machine Learning*, 2007, pp. 33–40.
- [93] M.S. Rasooli and J.R. Tetreault, Yara Parser: A Fast and Accurate Dependency Parser, *Computing Research Repository arXiv:1503.06733* (2015), version 2. <http://arxiv.org/abs/1503.06733>.
- [94] R.K. Ando and T. Zhang, A Framework for Learning Predictive Structures from Multiple Tasks and Unlabeled Data, *Journal of Machine Learning Research* **6** (2005), 1817–1853.
- [95] P. Bojanowski, E. Grave, A. Joulin and T. Mikolov, Enriching Word Vectors with Subword Information, *Transactions of the Association for Computational Linguistics* **5** (2017), 135–146.
- [96] J. Devlin, M.-W. Chang, K. Lee and K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, 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. 4171–4186. doi:10.18653/v1/N19-1423. <https://www.aclweb.org/anthology/N19-1423>.
- [97] M. Schlichtkrull and H. Martínez Alonso, MSejrKu at SemEval-2016 Task 14: Taxonomy Enrichment by Evidence Ranking, in: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, Association for Computational Linguistics, San Diego, California, 2016, pp. 1337–1341. doi:10.18653/v1/S16-1209. <https://www.aclweb.org/anthology/S16-1209>.
- [98] Z.S. Harris, Distributional structure, *Word* **10**(2–3) (1954), 146–162.
- [99] X. Cai, Y. Luo, Y. Zhang and X. Yuan, Improving Word Embeddings by Emphasizing Co-hyponyms, in: *International Conference on Web Information Systems and Applications*, Springer, 2018, pp. 215–227.
- [100] G.A. Miller, WORDNET: A LEXICAL DATABASE FOR ENGLISH, in: *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*, 1992. <https://www.aclweb.org/anthology/H92-1116>.
- [101] A. Moro and R. Navigli, SemEval-2015 Task 13: Multilingual All-Words Sense Disambiguation and Entity Linking, in: *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, Association for Computational Linguistics, Denver, Colorado, 2015, pp. 288–297. doi:10.18653/v1/S15-2049. <https://www.aclweb.org/anthology/S15-2049>.
- [102] M. Negri and B. Magnini, Using wordnet predicates for multilingual named entity recognition, in: *Proceedings of The Second Global Wordnet Conference*, 2004, pp. 169–174.
- [103] J. Herrera, A. Penas and F. Verdejo, Textual entailment recognition based on dependency analysis and wordnet, in: *Machine Learning Challenges Workshop*, Springer, 2005, pp. 231–239.
- [104] S.P. Ponzetto and M. Strube, Exploiting Semantic Role Labeling, WordNet and Wikipedia for Coreference Resolution, in: *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, Association for Computational Linguistics, New York City, USA, 2006, pp. 192–199. <https://www.aclweb.org/anthology/N06-1025>.
- [105] Y. Kuratov and M. Arkipov, Adaptation of Deep Bidirectional Multilingual Transformers for Russian Language, *arXiv preprint arXiv:1905.07213* (2019).
- [106] B. Shi and T. Weninger, ProjE: Embedding projection for knowledge graph completion, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31, 2017.
- [107] Y. Lin, Z. Liu, M. Sun, Y. Liu and X. Zhu, Learning entity and relation embeddings for knowledge graph completion, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29, 2015.
- [108] M. Maziarz, M. Piasecki, E. Rudnicka and S. Szpakowicz, plWordNet as the Cornerstone of a Toolkit of Lexico-semantic Resources, in: *Proceedings of the Seventh Global Wordnet Conference*, University of Tartu Press, Tartu, Estonia, 2014, pp. 304–312. <https://www.aclweb.org/anthology/W14-0142>.
- [109] M. Asgari, A. Hadian and B. Minaei-Bidgoli, Farsbase: The persian knowledge graph, *Semantic Web-Interoperability, Usability, Applicability* (2018).
- [110] N. Jain, Domain-Specific Knowledge Graph Construction for Semantic Analysis, in: *European Semantic Web Conference*, Springer, 2020, pp. 250–260.
- [111] P.A. Bonatti, S. Decker, A. Polleres and V. Presutti, Knowledge graphs: New directions for knowledge representation on the semantic web (dagstuhl seminar 18371), in: *Dagstuhl Reports*, Vol. 8, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [112] J. McCrae, C. Fellbaum and P. Cimiano, Publishing and Linking WordNet using lemon and RDF, in: *Proceedings of the 3rd Workshop on Linked Data in Linguistics*, 2014.
- [113] S.M. Harabagiu and D.I. Moldovan, Enriching the wordnet taxonomy with contextual knowledge acquired from text, *Natural language processing and knowledge representation: language for knowledge and knowledge for language* (2000), 301–333.
- [114] R. Snow, D. Jurafsky and A.Y. Ng, Semantic taxonomy induction from heterogenous evidence, in: *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, 2006, pp. 801–808.
- [115] D. Jurgens and M.T. Pilehvar, Reserating the awesometastic: An automatic extension of the WordNet taxonomy for novel terms, in: *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2015, pp. 1459–1465.

Appendix A. Ranked Examples and All Results

In this Appendix we provide Tables 7, 8 and 9 with the examples that demonstrate the input and the output formats of the models as well as Tables 12 and 13 that show the performance of all models for both English and Russian datasets.

From both tables 7 and 8 (underlined bold text denotes predictions of the model from the ground truth) we can see that at least one of the correct candidates usually appears in the list of candidates from word embeddings (first part of the table), whereas among candidates from graph embeddings we do not see any decent synsets. Poincaré embeddings retrieved by aggregating words from fastText provide too broad concepts which are clearly too far from the correct answers (“activity.n.01”, “exposure.n.03”, “action.n.02”). Node2vec embeddings are both semantically far and abstract. GraphSAGE is sticking to the word “play” and are too far from the correct answers in general. TADW manages to predict the correct synset “therapy.n.01” in the list of candidates, however, its position is much lower than the positions of the same synsets among the candidates provided by word embeddings-based systems.

The candidates for the words “play therapy” and “eyewitness” provided by models based on text embeddings do contain at least one true answer in the list. The position of such words varies from the fourth to the sixth position.

DWRank-Meta on both word and graph embeddings may provide the results that improve the ranking, e.g. “therapy.n.01” is the correct candidate on the first place in the list for both AAEME triple loss (fastText, word2vec and GloVe) approach and for the AAEME (fastText, word2vec, Glove, TASDW) approach. For the word “eyewitness” the DWRank-Meta models on words and graphs the true candidates are placed in the worse positions, however, they still win before the DWRank-Graph approach.

Table 7
Examples of predictions for noun from the English v 1.6-3.0 dataset with various models.

play therapy psychotherapy.n.02, therapy.n.01			
fastText (baseline)	fastText (DWRank)	AAEME triplet loss (fastText + word2vec + GloVe)	AAEME (fastText + word2vec + GloVe + TADW)
play.n.03	play.n.03	<u>therapy.n.01</u>	<u>therapy.n.01</u>
play.n.01	activity.n.01	activity.n.01	medical_care.n.01
baseball_play.n.01	plan_of_action.n.01	medical_care.n.01	play.n.03
<u>therapy.n.01</u>	<u>therapy.n.01</u>	diversion.n.01	play.n.01
activity.n.01	play.n.01	play.n.01	activity.n.01
diversion.n.01	dramatic_composition.n.01	act.n.02	dramatic_composition.n.01
plan_of_action.n.01	outdoor_game.n.01	<u>psychotherapy.n.02</u>	plan_of_action.n.01
action.n.01	medical_care.n.01	play.n.03	show.n.04
action.n.02	golf.n.01	dramatic_composition.n.01	treatment.n.01
dramatic_composition.n.01	diversion.n.01	behaviour_therapy.n.01	act.n.02
GraphSAGE	TADW	Poincaré	node2vec
baseball_play.n.01	play.n.03	activity.n.01	presentation.n.03
play.n.03	play.n.01	exposure.n.03	presentation.n.01
play.n.01	plan_of_action.n.01	rejection.n.01	operation.n.01
squeeze_play.n.02	dramatic_composition.n.01	agreement.n.06	performance.n.02
activity.n.01	activity.n.01	light_unit.n.01	contact.n.01
diversion.n.01	baseball_play.n.01	blessing.n.01	exposure.n.03
dramatic_composition.n.01	show.n.04	vulnerability.n.02	activity.n.01
attempt.n.01	diversion.n.01	action.n.02	exposure.n.08
plan_of_action.n.01	use.n.01	influence.n.02	union.n.04
play.n.17	action.n.02	assent.n.01	exposure.n.06

Table 8

Examples of predictions for nouns from the English v 1.6-3.0 dataset with various models.

Ramadan			
islamic_calendar_month.n.01, calendar_month.n.01, fast.n.01, abstinence.n.02			
fastText (baseline)	fastText (DWRank)	AAEME triplet loss (fastText + word2vec + GloVe)	AAEME (fastText + word2vec + GloVe + TADW)
islamic_calendar_month.n.01	islamic_calendar_month.n.01	calendar.n.01	islamic_calendar_month.n.01
calendar_month.n.01	calendar_month.n.01	islamic_calendar_month.n.01	calendar_month.n.01
holiday.n.02	time_period.n.01	calendar_month.n.01	time_period.n.01
religious_holiday.n.01	calendar.n.01	lunar_calendar.n.01	calendar.n.01
abstinence.n.02	islam.n.01	time_period.n.01	muslim.n.01
hindu_calendar_month.n.01	lunar_calendar.n.01	religionist.n.01	religionist.n.01
day.n.04	asian.n.01	muslim.n.01	religious_holiday.n.01
sacred_text.n.01	religion.n.02	religion.n.02	holiday.n.02
god.n.01	muslim.n.01	islam.n.01	lunar_calendar.n.01
place_of_worship.n.01	holiday.n.02	person.n.01	islam.n.01
graphSAGE	TADW	Poincaré	node2vec
islamic_calendar_month.n.01	islamic_calendar_month.n.01	time_period.n.01	calendar_month.n.01
calendar_month.n.01	calendar_month.n.01	islamic_calendar_month.n.01	islamic_calendar_month.n.01
arab.n.01	time.n.02	religion.n.02	revolutionary_calendar_month.n.01
muslim.n.01	religious_holiday.n.01	time.n.02	islam.n.01
semitic.n.01	calendar.n.01	measure.n.03	time_period.n.01
religionist.n.01	time_period.n.01	calendar_month.n.01	hindu_calendar_month.n.01
god.n.01	holiday.n.02	term.n.02	muharram.n.01
saint.n.02	lunar_calendar.n.01	year.n.01	shawwal.n.01
zoysia.n.01	god.n.01	time_off.n.01	rabi_i.n.01
islam.n.01	jewish_holy_day.n.01	leisure.n.01	rajab.n.01

Table 9

Examples of predictions for verbs from the English v 1.6-3.0 dataset with various models.

eyewitness			
witness.v.01, watch.v.01			
fastText (baseline)	fastText (DWRank)	AAEME triplet loss (fastText + word2vec + GloVe)	AAEME (fastText + word2vec + GloVe + TADW)
be.v.01	inform.v.01	inform.v.01	inform.v.01
be.v.03	testify.v.02	testify.v.02	testify.v.02
be.v.08	communicate.v.02	communicate.v.02	confirm.v.02
watch.v.01	announce.v.01	see.v.10	communicate.v.02
testify.v.01	confirm.v.02	confirm.v.02	watch.v.01
testify.v.02	watch.v.01	watch.v.01	be.v.01
man.v.02	testify.v.01	witness.v.02	affirm.v.03
talk.v.01	affirm.v.03	affirm.v.03	testify.v.01
guard.v.01	report.v.03	testify.v.01	reject.v.01
confirm.v.01	record.v.01	verify.v.01	experience.v.01
graphSAGE	TADW	Poincaré	node2vec
see.v.05	inform.v.01	confirm.v.02	affirm.v.03
testify.v.02	testify.v.02	examine.v.02	confirm.v.02
err.v.01	communicate.v.02	affirm.v.03	understand.v.02
confirm.v.01	declare.v.01	testify.v.02	uphold.v.03
pronounce.v.02	announce.v.01	inform.v.01	determine.v.08
idealize.v.01	testify.v.01	justify.v.02	stay_in_place.v.01
judge.v.02	record.v.01	declare.v.01	justify.v.02
negate.v.03	watch.v.01	validate.v.03	fall_asleep.v.01
reason.v.01	report.v.03	uphold.v.03	resettle.v.01
disbelieve.v.01	report.v.01	testify.v.01	settle.v.04

Table 10

Examples of predictions for verbs from the English v 1.6-3.0 dataset with various models.

theologise			
cover.v.05, broach.v.01, chew_over.v.01, think.v.03			
fastText (baseline)	fastText (DWRank)	AAEME triplet loss (fastText + word2vec + GloVe)	AAEME (fastText + word2vec + GloVe + TADW)
change.v.01	change.v.01	change.v.01	change.v.01
match.v.01	preface.v.01	preface.v.01	preface.v.01
date.v.03	make.v.03	convert.v.05	equal.v.03
reconstruct.v.01	date.v.03	match.v.01	date.v.03
determine.v.03	equal.v.03	chronologize.v.01	reconstruct.v.01
make.v.03	reason.v.01	reason.v.01	reason.v.01
speculate.v.01	chronologize.v.01	change_state.v.01	convert.v.05
convert.v.05	match.v.01	represent.v.09	formulate.v.03
reason.v.01	film.v.02	make.v.03	reflect.v.04
automatize.v.01	formulate.v.03	change.v.02	film.v.02
graphSAGE	TADW	Poincaré	node2vec
process.v.02	change.v.01	state.v.01	settle.v.04
affect.v.01	preface.v.01	speculate.v.01	stay_in_place.v.01
change.v.01	date.v.03	reason.v.01	understand.v.02
dive.v.01	film.v.02	preface.v.01	study.v.03
tame.v.01	reconstruct.v.01	match.v.01	resettle.v.01
sensitize.v.02	equal.v.03	generalize.v.01	fall_asleep.v.01
convert.v.01	convert.v.05	announce.v.02	speculate.v.01
estimate.v.01	formulate.v.03	express.v.02	discover.v.07
subject.v.01	reason.v.01	add.v.02	explicate.v.02
compound.v.05	commemorate.v.03	equal.v.01	behave.v.02
graphSAGE	TADW	Poincaré	node2vec
defend.v.02	protect.v.01	teach.v.01	receive.v.01
protect.v.01	indoctrinate.v.01	insert.v.02	insert.v.02
act.v.01	teach.v.01	inform.v.01	stay_in_place.v.01
negotiate.v.01	defend.v.02	treat.v.01	get.v.01
attack.v.03	prevent.v.02	train.v.01	inject.v.01
prevent.v.02	inject.v.01	deceive.v.02	fall_asleep.v.01
insert.v.02	prevent.v.01	indoctrinate.v.01	accept.v.02
demilitarize.v.01	insert.v.02	misinform.v.01	settle.v.04
disarm.v.02	isolate.v.01	gull.v.02	resettle.v.01
foreswear.v.02	discriminate.v.02	interact.v.01	discover.v.07

Table 11

Examples of predictions for verbs from the English v 1.6-3.0 dataset with various models.

immunise			
protect.v.01, defend.v.02, inject.v.01, administer.v.04			
fastText (baseline)	fastText (DWRank)	AAEME triplet loss (fastText + word2vec + GloVe)	AAEME (fastText + word2vec + GloVe + TADW)
indoctrinate.v.01	indoctrinate.v.01	teach.v.01	inject.v.01
inject.v.01	protect.v.01	treat.v.01	remove.v.01
defend.v.02	defend.v.02	insert.v.02	inform.v.01
remove.v.01	teach.v.01	prevent.v.01	change.v.01
treat.v.01	prevent.v.02	isolate.v.01	better.v.02
destroy.v.01	inject.v.01	inoculate.v.01	insert.v.02
prevent.v.02	insert.v.02	indoctrinate.v.01	defend.v.02
insert.v.02	defend.v.01	kill.v.01	kill.v.01
teach.v.01	kill.v.01	change.v.01	indoctrinate.v.01
administer.v.04	discriminate.v.02	enable.v.01	protect.v.01
graphSAGE	TADW	Poincaré	node2vec
defend.v.02	protect.v.01	teach.v.01	receive.v.01
protect.v.01	indoctrinate.v.01	insert.v.02	insert.v.02
act.v.01	teach.v.01	inform.v.01	stay_in_place.v.01
negotiate.v.01	defend.v.02	treat.v.01	get.v.01
attack.v.03	prevent.v.02	train.v.01	inject.v.01
prevent.v.02	inject.v.01	deceive.v.02	fall_asleep.v.01
insert.v.02	prevent.v.01	indoctrinate.v.01	accept.v.02
demilitarize.v.01	insert.v.02	misinform.v.01	settle.v.04
disarm.v.02	isolate.v.01	gull.v.02	resettle.v.01
foreswear.v.02	discriminate.v.02	interact.v.01	discover.v.07

Table 12

MAP scores for the taxonomy enrichment methods for the English datasets. Numbers **in bold** show the best model within the category, **underlined** numbers denote the best score across all the models. The combination of word embeddings (fastText, word2vec, GloVe) is denoted as *words*.

method	Nouns			Verbs		
	1.6-3.0	1.7-3.0	2.0-3.0	1.6-3.0	1.7-3.0	2.0-3.0
Baseline [6]						
fastText [71]	0.338±0.002	0.371±0.002	0.400±0.004	0.270±0.007	0.203±0.010	0.236±0.011
word2vec [72]	0.142±0.001	0.178±0.002	0.164±0.004	0.229±0.006	0.155±0.008	0.212±0.009
GloVe [39]	0.232±0.002	0.188±0.001	0.233±0.004	0.146±0.005	0.149±0.008	0.191±0.010
DWRank-Word						
fastText [71]	0.314±0.001	0.373±0.003	0.418±0.004	0.286±0.007	0.218±0.008	0.254±0.012
word2vec [72]	0.244±0.001	0.271±0.003	0.298±0.004	0.099±0.005	0.118±0.008	0.141±0.010
GloVe [39]	0.283±0.001	0.329±0.003	0.377±0.004	0.182±0.007	0.159±0.008	0.203±0.011
DWRank-Meta (Meta-embeddings based on Word Embeddings)						
concat (<i>words</i>)	0.335±0.001	0.386±0.003	0.386±0.003	0.270±0.007	0.194±0.009	0.226±0.011
SVD (<i>words</i>)	0.333±0.001	0.399±0.003	0.456±0.004	0.277±0.007	0.209±0.010	0.264±0.012
CAEMEWords	0.321±0.001	0.386±0.003	0.448±0.005	0.278±0.007	0.205±0.008	0.266±0.015
AAEMEWords	0.322±0.001	0.384±0.003	0.453±0.004	0.271±0.007	0.218±0.008	0.273±0.012
CAEME triplet loss (<i>words</i>)	0.332±0.001	0.394±0.003	0.451±0.004	0.273±0.007	0.205±0.007	0.256±0.013
AAEME triplet loss (<i>words</i>)	0.335±0.001	0.391±0.003	0.453±0.004	0.280±0.008	0.212±0.007	0.262±0.014
DWRank-Graph						
GCN [62]	0.175±0.001	0.249±0.002	0.267±0.002	0.162±0.006	0.113±0.005	0.149±0.010
GAT [63]	0.000±0.000	0.252±0.002	0.000±0.000	0.081±0.003	0.064±0.004	0.000±0.000
GraphSAGE [79]	0.214±0.001	0.282±0.002	0.224±0.003	0.127±0.004	0.114±0.004	0.090±0.008
TADW [81] (on fastText)	0.350±0.001	0.392±0.002	0.435±0.004	0.268±0.007	0.201±0.007	0.217±0.010
Poincaré [76] (top-5 fastText associates)	0.185±0.001	0.211±0.002	0.229±0.002	0.208±0.006	0.147±0.006	0.172±0.012
node2vec [52] (top-5 fastText associates)	0.270±0.001	0.312±0.002	0.341±0.004	0.175±0.006	0.128±0.007	0.118±0.012
HOPE [83]	0.000±0.000	0.000±0.000	0.000±0.000	0.000±0.000	0.000±0.000	0.000±0.000
DWRank-Meta (Meta-embeddings based on Word and Graph Embeddings)						
SVD (<i>words</i> + node2vec)	0.343±0.001	0.383±0.003	0.434±0.005	0.272±0.006	0.194±0.009	0.239±0.011
CAEME (<i>words</i> + node2vec)	0.335±0.001	0.379±0.003	0.426±0.004	0.242±0.005	0.184±0.009	0.221±0.012
AAEME (<i>words</i> + node2vec)	0.350±0.001	0.394±0.003	0.446±0.004	0.252±0.007	0.184±0.008	0.208±0.012
SVD (<i>words</i> + TADW)	0.355±0.001	0.414±0.003	0.472±0.004	0.288±0.007	0.222±0.009	0.280±0.013
CAEME (<i>words</i> + TADW)	0.350±0.001	0.404±0.003	0.458±0.004	0.267±0.007	0.212±0.007	0.247±0.011
AAEME (<i>words</i> + TADW)	0.367±0.001	0.418±0.002	0.480±0.004	0.283±0.007	0.227±0.007	0.260±0.012
SVD (<i>words</i> + GCN)	0.323±0.001	0.385±0.003	0.443±0.004	0.260±0.005	0.209±0.009	0.249±0.011
CAEME (<i>words</i> + GCN)	0.331±0.001	0.395±0.003	0.457±0.004	0.251±0.006	0.207±0.009	0.235±0.012
AAEME (<i>words</i> + GCN)	0.331±0.001	0.392±0.003	0.456±0.004	0.243±0.006	0.200±0.008	0.228±0.012
SVD (<i>words</i> + GraphSAGE)	0.338±0.001	0.401±0.003	0.464±0.004	0.239±0.006	0.194±0.009	0.221±0.011
CAEME (<i>words</i> + GraphSAGE)	0.323±0.001	0.382±0.003	0.435±0.004	0.200±0.006	0.170±0.007	0.202±0.01
AAEME (<i>words</i> + GraphSAGE)	0.343±0.001	0.406±0.003	0.468±0.004	0.238±0.007	0.178±0.008	0.209±0.011
State-of-the-art Approaches						
WBSR (Top-1 RUSSE'2020 for nouns)	0.333±0.002	0.393±0.003	0.436±0.003	0.252±0.006	0.206±0.011	0.252±0.013
WBSR, no search engine features	0.251±0.001	0.309±0.003	0.344±0.004	0.231±0.006	0.180±0.008	0.222±0.009
hypo2path rev [67]	0.264±0.001	0.283±0.003	0.238±0.007	0.173±0.005	0.104±0.008	0.118±0.009
hypo2path [67]	0.252±0.002	0.261±0.002	0.208±0.006	0.162±0.005	0.093±0.006	0.067±0.008
hypo2path transformer	0.218±0.002	0.229±0.002	0.057±0.002	0.140±0.003	0.120±0.006	0.100±0.008
TaxoExpan [61]	0.004±0.000	0.003±0.000	0.054±0.002	0.001±0.000	0.000±0.000	0.000±0.000

Table 13

MAP scores for the taxonomy enrichment methods for the Russian datasets. Numbers **in bold** show the best model within the category, **underlined** numbers denote the best score across all the models.

method	nouns		verbs	
	non-restricted	restricted	non-restricted	restricted
Baseline				
fastText [71]	0.414±0.001	0.549±0.006	0.296±0.004	0.389±0.011
word2vec [72]	0.263±0.001	0.427±0.006	0.343±0.004	0.445±0.013
DWRank (Word Embeddings)				
fastText [71]	0.419±0.001	0.572±0.005	0.337±0.003	0.428±0.007
word2vec [72]	0.296±0.002	0.569±0.005	0.250±0.003	0.284±0.011
DWRank (Meta-embeddings based on Word Embeddings)				
concat (<i>words</i>)	0.422±0.001	0.589±0.005	0.351±0.004	0.426±0.009
SVD (<i>words</i>)	0.461±0.001	0.600±0.005	0.426±0.005	0.475±0.010
CAEME (<i>words</i>)	0.400±0.001	0.561±0.005	0.342±0.003	0.416±0.008
AAEME (<i>words</i>)	0.456±0.001	0.582±0.005	0.368±0.004	0.442±0.009
CAEME triplet loss (<i>words</i>)	0.449±0.001	0.581±0.005	0.374±0.003	0.427±0.010
AAEME triplet loss (<i>words</i>)	0.474±0.001	0.593±0.006	0.399±0.004	0.449±0.010
DWRank (Graph embeddings)				
GCN [57]	0.183±0.001	0.306±0.005	0.220±0.003	0.287±0.009
GAT [63]	0.142±0.001	0.318±0.004	0.000±0.000	0.000±0.000
GraphSAGE [79]	0.176±0.001	0.348±0.005	0.181±0.003	0.226±0.008
TADW [81]	0.417±0.001	0.562±0.005	0.328±0.003	0.423±0.008
Poincaré [76] (top-5 fastText associates)	0.336±0.001	0.476±0.005	0.244±0.004	0.339±0.009
node2vec [52] (top-5 fastText associates)	0.343±0.002	0.477±0.005	0.226±0.003	0.322±0.010
HOPE [83]	0.000±0.000	0.000±0.000	0.003±0.001	0.003±0.001
DWRank (Meta-embeddings based on Word and Graph Embeddings)				
SVD (<i>words</i> + node2vec)	0.367±0.001	0.521±0.005	0.252±0.003	0.351±0.010
CAEME (<i>words</i> + node2vec)	0.370±0.001	0.533±0.005	0.267±0.003	0.362±0.010
AAEME (<i>words</i> + node2vec)	0.373±0.001	0.529±0.005	0.272±0.003	0.358±0.010
SVD (<i>words</i> + TADW)	0.469±0.001	0.604±0.006	0.394±0.005	0.455±0.010
CAEME (<i>words</i> + TADW)	0.429±0.001	0.571±0.005	0.349±0.003	0.437±0.009
AAEME (<i>words</i> + TADW)	0.461±0.001	0.584±0.005	0.362±0.004	0.439±0.009
SVD (<i>words</i> + GCN)	0.395±0.001	0.554±0.005	0.291±0.004	0.356±0.009
CAEME (<i>words</i> + GCN)	0.389±0.001	0.544±0.005	0.302±0.003	0.381±0.008
AAEME (<i>words</i> + GCN)	0.386±0.001	0.545±0.006	0.295±0.004	0.365±0.008
SVD (<i>words</i> + GraphSAGE)	0.410±0.001	0.603±0.005	0.336±0.004	0.426±0.009
CAEME (<i>words</i> + GraphSAGE)	0.321±0.001	0.541±0.005	0.266±0.004	0.345±0.007
AAEME (<i>words</i> + GraphSAGE)	0.409±0.001	0.577±0.006	0.323±0.004	0.419±0.009
State-of-the-art Approaches				
WBSR (Top-1 RUSSE'2020 for nouns)	0.393±0.002	0.552±0.005	0.293±0.004	0.428±0.010
WBSR, no search engine features	0.369±0.002	0.497±0.005	0.267±0.004	0.387±0.009
Top-1 RUSSE'2020 for verbs: [33]	0.288±0.001	0.418±0.006	0.341±0.004	0.452±0.012
hypo2path [67]	0.061±0.000	0.097±0.002	0.137±0.003	0.174±0.009
hypo2path rev [67]	0.246±0.001	0.342±0.006	0.151±0.003	0.194±0.008
hypo2path rev transformer [67]	0.234±0.001	0.331±0.004	0.152±0.003	0.201±0.008
TaxoExpan [61]	0.007±0.000	0.006±0.001	0.009±0.001	0.008±0.002