



Ontology Enrichment from Texts: A Biomedical Dataset for Concept Discovery and Placement

Hang Dong
University of Oxford
hang.dong@cs.ox.ac.uk

Yuan He
University of Oxford
yuan.he@cs.ox.ac.uk

Jiaoyan Chen
The University of Manchester & University of Oxford
jiaoyan.chen@manchester.ac.uk

Ian Horrocks
University of Oxford
ian.horrocks@cs.ox.ac.uk

Abstract

Mentions of new concepts appear regularly in texts and require automated approaches to harvest and place them into Knowledge Bases (KB), e.g., ontologies and taxonomies. Existing datasets suffer from three issues, (i) mostly assuming that a new concept is pre-discovered and cannot support out-of-KB mention discovery; (ii) only using the concept label as the input along with the KB and thus lacking the contexts of a concept label; and (iii) mostly focusing on concept placement w.r.t a taxonomy of atomic concepts, instead of complex concepts, i.e., with logical operators. To address these issues, we propose a new benchmark, adapting MedMentions dataset (PubMed abstracts) with SNOMED CT versions in 2014 and 2017 under the Diseases subcategory and the broader categories of Clinical finding, Procedure, and Pharmaceutical / biologic product. We provide usage on the evaluation with the dataset for out-of-KB mention discovery and concept placement, adapting recent Large Language Model based methods.¹

CCS Concepts

• Computing methodologies → Ontology engineering; Language resources.

Keywords

Ontology Enrichment, Text Mining, Entity Linking, Concept Placement, Language Models, Biomedical Ontologies, SNOMED CT

ACM Reference Format:

Hang Dong, Jiaoyan Chen, Yuan He, and Ian Horrocks. 2023. Ontology Enrichment from Texts: A Biomedical Dataset for Concept Discovery and Placement. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*, October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3583780.3615126>

¹The dataset, data construction scripts, and baseline implementation are available at <https://zenodo.org/record/8228005> and <https://github.com/KRR-Oxford/OET>.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CIKM '23, October 21–25, 2023, Birmingham, United Kingdom

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0124-5/23/10.

<https://doi.org/10.1145/3583780.3615126>

1 INTRODUCTION

Identifying new concepts and placing them into Knowledge Bases (KBs, e.g., ontologies and taxonomies) from texts such as a vast amount of publications is a key application of KB construction and AI for scientific discovery [13]. Emerging concepts are particularly common in the biomedical domain and KB can easily be outdated. For example, new variants of SARS-CoV-2 have kept emerging since 2020; “Curry-Jones syndrome” was not added to SNOMED CT ontology [8] until 2017.

Existing datasets on using texts to enrich ontologies are relevant to several tasks, but each of them only reflects a part of the whole picture. In *Taxonomy Completion* [20, 29], a pre-specified out-of-KB (a.k.a. NIL) concept is used to enrich a taxonomy. In *Ontology Extension* [11], this is extended into description logic based ontologies, which include complex concepts which can be considered [6, 15] but not yet for the existing datasets [11]. In *Concept Post-coordination* [5, 22], an out-of-KB concept is defined with several existing concepts and attributes, i.e., placed under a complex concept. Datasets from all the three tasks above assume that the input concept term is already specified and non-contextual (e.g., without contexts in a corpus), which does not reflect the real-world situation. In *Out-of-KB Mention and Entity Discovery* [7, 16, 17], out-of-KB mentions and their clustering are discovered from texts, but their placement in KBs has not been fully investigated.

In this study, we propose a new benchmark for new entity discovery and placement, supporting two sequential tasks: (i) *Out-of-KB Mention and Entity Discovery*: identifying new mentions of concepts from texts which are not included in a KB; (ii) *Concept Placement*: given a new entity expressed as a mention in the text, placing it into a KB, either an ontology with complex concepts or a taxonomy with only atomic concepts. Our new dataset and task setting are different from previous work in terms of the characteristics below:

- **Out-of-KB or NIL discovery**: inclusion of *out-of-KB mentions* from texts to support their concept discovery and placement.
- **Contextual terms**: inclusion of *contexts* for mentions, distinct from only using concept labels as in the previous work.
- **Complex concepts**: placement of concepts under logic-equipped *complex concepts*, instead of atomic concepts alone.

More specifically, the study uses a SNOMED CT subset as the ontology, and the time difference of two versions (in 2014 and 2017) to synthesise new entities, then uses MedMentions Entity Linking dataset [25] (from PubMed abstracts to UMLS) to construct in-KB and out-of-KB mentions. The study further introduces the usage

Table 1: Comparison of relevant datasets and tasks on KB (e.g., ontology and taxonomy) enrichment from texts. *NIL Discovery* denotes whether the task can support discovering out-of-KB mentions (cf. in-KB mentions). *Contextual Term* denotes whether the input term has a context window in a text corpus. *Concept Placement* denotes whether the task finally places (or can be used to place) the term in the KB. *Complex Concepts* denote whether the placement position in the KB includes complex concepts. The asterisk (*) denotes that only data construction scripts are available instead of the dataset itself.

Datasets (with public access link in citations)	Task	NIL Discovery	Contextual Term	Concept Placement	Complex Concepts
MAG, WordNet [20, 29, 34]; OSConcepts, DroneTaxo, MeSH, SemEval [31, 33]	Taxonomy Completion	No	No	Yes	No
ChEBI ₅₀₀ , ChEBI ₅₀₀ ⁺ [11]	Ontology Extension	No	No	Yes	No
SNOMED CT (English, manual, small-scale) [22]; SNOMED CT (Spanish, automated) [5]*	Concept Post-coordination	No	No	Yes	Yes
NILK [17]; ShARe/CLEF 2013 [30]; CLEF HIPE 2020 [10]; EDIN [21]; NEEL 2015-2016 [26, 28]	Out-of-KB Mention and Entity Discovery	Yes	Yes	No	No
MedMentions-SNOMED-CT-14 (Disease) [this work]	(-CPP, -) Concept Discovery and Placement	Yes	Yes	Yes	Yes

of the data with evaluation for out-of-KB mention discovery and concept placement. We provide benchmarking results adapting rule-based and BERT-based Entity Linking [7, 32] and prompting with GPT-3.5-turbo. Results show that the dataset well differentiates the performance between rule-based and BERT-based methods, and the Pre-trained and Large Language Model (LLM) based methods are still yet to achieve satisfying results.

2 RELATED WORK

The representative datasets are summarised in Table 1 based on the four tasks introduced in Section 1. We only list the public and accessible datasets. We next discuss the related work of each task. **Taxonomy Completion and Ontology Extension.** Studies in taxonomy completion [20, 29, 31, 33, 34] and ontology extension [11] aim to enrich KB using the concept labels and the concept graph structure. However, the studies usually assume that the new term (or concept label) is pre-discovered, which is not the case in the real-world scenario, where new mentions of concepts need to be discovered from corpora. Also, from the perspectives of OWL (Web Ontology Language) [2, 12], most of these studies focus only on atomic concepts and do not place the new concept under a *complex concept*, e.g., with existential restrictions used in SNOMED CT (e.g., [24] focuses on placement under only atomic concepts in SNOMED CT). Also, datasets in both tasks use concept terms as input and do not consider contexts.

Concept Post-coordination. The studies aim to place a new concept by describing it with existing concepts and attributes in the ontology [5, 22]. Dataset construction steps in both works [5, 22] assume that the new concepts or terms are pre-discovered and without context windows from a corpus.

Out-of-KB Mention and Entity Discovery. The studies aim to discover new mentions from texts, w.r.t. to a KB [7, 30] and group them into entity clusters [17, 18, 21, 27]. There is a growth of datasets in this area recently, constructed through Manual Labelling, KB pruning, and/or KB versioning [7]. The studies, however, do not place the newly discovered entities into a KB.

In this work, we present dataset construction for *Concept Discovery and Placement* to support a comprehensive set of characteristics (Table 1), with usage for benchmarking, e.g., with Pre-trained and Large Language Models.

3 PROBLEM DEFINITION

The task of *Concept Discovery and Placement* inputs contextual, in-KB and out-of-KB mentions in a corpus and a KB (more formally as an OWL ontology [2, 12]) and outputs an enriched KB where each out-of-KB mention is inserted into a *directed edge*, i.e., $\langle \text{parent}, \text{child} \rangle$, of the KB, when the out-of-KB mention is the child of the parent and the parent of the child. The child is considered NULL when the mention corresponds to a leaf concept. The parent can be a *complex concept*.

Several key definitions are as follows. Formally, an *OWL ontology* is a Description Logic KB that contains a set of axioms [2, 12]. We focus on the TBox (or the terminology part) in an ontology, which mainly consists of General Concept Inclusion axioms of the form $A \sqsubseteq B$, where A (and B) are either atomic or complex concepts [1]. A TBox can be reduced to a *taxonomy* (or a subsumption hierarchy) after the classification process [1]. Borrowing the definition of taxonomy in [19, 29], we use a simple definition of ontology as a set of concepts and directed edges, where both can be atomic or complex. *Directed Edges* are edges in an ontology (or a taxonomy [31]) which contain a direct parent and a direct child. *Complex edges* are edges which have a complex concept as the parent². *Complex concepts* mean concepts that involve at least one logical operators, e.g., negation (\neg), conjunction (\sqcap), disjunction (\sqcup), existential restriction ($\exists r.C$), universal restriction ($\forall r.C$), etc. [1].³





An ideal dataset for *Concept Discovery and Placement* requires a real-world text corpus and a large OWL ontology (reducible to a taxonomy), with gold-standard directed edges (possibly complex) for each out-of-KB concept linked to the mentions in the corpus.

4 DATASET CONSTRUCTION

Step 0: KB and Subset Selection. We consider SNOMED CT [8], one of the most important OWL ontology in the biomedical domain, and choose a subset by selected categories of concepts. We focus on the second level category, *Disease* (disorder) and the first level categories, *Clinical finding*, *Pharmaceutical / biologic product*, and *Procedure*, abbreviated to *CPP* as the initials of the categories.

²It is not likely to have complex concepts as direct (asserted) children in an ontology.

³The ontology language OWL [12], based on Description Logic, has expressiveness beyond RDF [2]; an example complex concept as a parent concept is $Arthritis \sqsubseteq Arthropathy \sqcap \exists hasMorphology.Inflammatory$ [15].

Step 0: KB and Subset Selection**Step 1: KB Versioning to Synthesize Out-of-KB Entities****Step 2: Edge Extraction for In-KB and Out-of-KB Entities****Step 3: Mention-to-Edge Data Creation with Sources of Entity Linking and Alignment****Figure 1: Data construction pipeline: KB and Subset Selection, KB Versioning, Edge Extraction, and Mention-Edge Data Creation.**

CPP categories have the most important types of complex edges for placement (or post-coordination), according to Kate [22].

The subset selection has two steps: (i) transforming equivalence axioms into subsumption axioms by outermost conjunctions, e.g., from $A \equiv C_1 \sqcap \exists R.(C_2 \sqcap C_3)$ to $A \sqsubseteq C_1$ and $A \sqsubseteq \exists R.(C_2 \sqcap C_3)$; (ii) pruning the ontology to only keep the selected categories.⁴

The implementation is as follows. First, we transform SNOMED CT files⁵ into OWL format with `snomed-owl-toolkit`⁶. Then, step (i) is implemented through ontology processing with `DeepOnto`⁷ [14]; and step (ii) through `Protégé`⁸ to remove the other categories.

Step 1: KB Versioning. We follow a KB versioning strategy [7, 17] to synthesise out-of-KB entities for the older KB. The concept gap between the two versions of SNOMED CT subsets (ver 20140901 and 20170301) is considered. The numbers of concepts in the older and the newer sub-KB are 64,900 and 72,595, resp., for the *Disease* sub-category and are 175,895 and 188,988, resp., for *CPP* categories.

Step 2: Edge Extraction. We extract the directed edges in the older sub-KB for both in-KB and out-of-KB entities. For in-KB entities, this is achieved by querying all the direct parents and children in the older KB. For out-of-KB entities, this is achieved by querying all the *most* direct, in-KB (older KB) parents and children from the newer sub-KB, given that edges for the out-of-KB entities are not available in the older KB. If the entity is a leaf node, we set the direct child as NULL, as in Zhang et al. [34]. The querying process is based on ontology processing module in `DeepOnto` [14].

Step 3: Mention-Edge Data Creation. A corpus with mentions linked to entities in SNOMED CT is needed to synthesise contextual mentions and gold-standard edges. Following Dong et al. [7], we use `MedMentions`⁹ [25], containing around 5,000 biomedical paper abstracts in PubMed¹⁰ where the mentions were manually and

exhaustively linked to UMLS [4] (version 2017AA)¹¹. We leverage the alignment of entities in UMLS to obtain SNOMED CT entities (ver 20170301)¹². Using the output from the previous steps, we can then map each in-KB and out-of-KB mention to edges¹³ in the older sub-KB (SNOMED CT ver 20140901).

We thus create mention-to-edge datasets with the information for each mention, each rendered in a JSON format. The information includes the left and the right contexts of the mention ($ctxt_l$ and $ctxt_r$), the mention or concept itself and its SNOMED CT ID, the parent and child in the older SNOMED CT ID (and with expression for complex concepts) and their labels.¹⁴ We use `DeepOnto`'s verbaliser [14, 15] to form labels of the complex concepts.

Regarding data splitting for the benchmark, for *out-of-KB mention and concept discovery*, the dataset follows the original splits of training, validation, and testing sets from `MedMentions`; for *concept placement*, the setting is unsupervised for out-of-KB mentions, i.e., training (and validating) with in-KB mentions but testing on out-of-KB mentions (and in-KB mentions)

We provide two formats of the data, *mention-level*, with edges grouped for each mention; and *mention-edge-pair-level*, where each mention-edge pair occupies a row and mentions are repeated if there are multiple edges. Statistics of the datasets are in Table 2.

5 DATA USAGE FOR BENCHMARKING

5.1 Evaluation with the Data

Metrics for Out-of-KB Mention Discovery The dataset supports the metrics in [7], including overall accuracy for all in-KB and out-of-KB mentions (A); *out-of-KB* precision (P_o), recall (R_o), and F_1 score (F_{1o}) to measure how well out-of-KB mentions are detected; and *in-KB* precision (P_{in}), recall (R_{in}), and F_1 score (F_{1in}).

⁴Given that many subsumption relations are in fact described as equivalence axioms with conjunction with concepts from to-be-removed categories, simply using step (ii) without step (i) will result in many broken hierarchies instead of a connected one.

⁵<https://www.nlm.nih.gov/healthit/snomedct/archive.html>

⁶<https://github.com/IHTSDO/snomed-owl-toolkit>

⁷<https://github.com/KRR-Oxford/DeepOnto>

⁸<http://protegeproject.github.io/protege/>

⁹<https://github.com/chanzuckerberg/MedMentions>

¹⁰<https://pubmed.ncbi.nlm.nih.gov/>

¹¹There are also other datasets available, e.g. COMETA [3], which directly links mentions in social media posts on Reddit to SNOMED CT. Still, we consider scientific publications a more reliable source and leave social media for a future study.

¹²If mention of one UMLS entity matches multiple SNOMED CT entities, we create a mention row in the data for each SNOMED CT entity.

¹³We also filtered the edges to one-hop (including leaf concept to NULL) and two-hop from any paths in the ontology.

¹⁴The keys in the JSON format are available in <https://github.com/KRR-Oxford/OET>.

Table 2: Statistics for datasets for *Concept Discovery and Placement*, for SNOMED CT (ver 20140901, “S14”) under different categories: “Disease” and “CPP”, i.e., *Clinical finding*, *Procedure*, and *Pharmaceutical / biologic product*. A mention-edge pair denotes a mention (in a corpus) and one of its directed edges in the KB. Mentions are from the MedMentions dataset (“MM”). * The numbers of edges are those having one hop (including leaf nodes to NULL) and two hops from any paths in the ontology.

	Ontology: # all (# complex)		Corpus: # Mentions / # Mention-edge pairs / # Mention-edge pairs with <i>complex</i> edges			
	concepts	edges*	train, in-KB	valid, in-KB	test, in-KB	out-of-KB
MM-S14-Disease	64,900 (824)	237,826 (4,997)	11,812 / 887,840 / 917	4,248 / 383,457 / 203	3,970 / 316,319 / 393	605 / 1,637 / 13
MM-S14-CPP	175,895 (2,718)	625,994 (19,401)	34,704 / 1,398,111 / 9,475	11,707 / 548,295 / 4,305	11,564 / 478,424 / 4,129	1,000 / 2,131 / 22

Table 3: Results on out-of-KB mention discovery

MM-S14-Disease	A	P_o	R_o	F_{1o}	P_{in}	R_{in}	F_{1in}
Sieve-based	55.9	6.4	47.0	11.2	88.1	56.2	68.6
BLINKout	67.2	14.6	17.4	15.9	69.0	68.6	68.8
MM-S14-CPP	A	P_o	R_o	F_{1o}	P_{in}	R_{in}	F_{1in}
Sieve-based	49.7	3.3	59.3	6.3	85.7	49.6	62.8
BLINKout	65.9	22.5	32.6	26.6	66.8	66.4	66.6

Table 4: Results on out-of-KB concept placement

MM-S14-Disease	P / R @1	P / R @5	P / R @10	P / R @50
Edge-Bi-encoder	4.5 / 1.6	6.0 / 11.0	5.4 / 19.9	2.1 / 38.4
+GPT-3.5	4.3 / 1.6	-	-	-
MM-S14-CPP	P / R @1	P / R @5	P / R @10	P / R @50
Edge-Bi-encoder	2.2 / 1.0	2.2 / 5.2	2.0 / 9.4	1.4 / 32.4
+GPT-3.5	2.5 / 1.2	-	-	-

Metrics for Concept Placement The dataset supports the metrics used in taxonomy completion, to evaluate the ranking of edges for a given mention [19, 29, 34]. The metrics mainly include Precision at k ($P@k$), Recall at k ($R@k$), F_1 score at k ($F_1@k$), Mean Rank (MR), and Mean Reciprocal Rank (MRR). We report $P@k$ and $R@k$ for different top- k values.

5.2 Experimenting with the Data

We experiment with the data w.r.t the two tasks using a rule-based method and recent, LLM-based methods.

For Out-of-KB Mention Discovery Existing methods are supervised, i.e., require a certain amount of NIL in the training data. Thus, we split the “out-of-KB” mentions in Table 2 based on the NIL mentions in the original MedMentions data split. The number of training, validation, and testing NIL mentions are 568, 260, and 172, resp., for CPP (in total 1,000 mentions); and 329, 161, and 115, resp., for Disease sub-categories (in total 605 mentions). For the rule-based method, we use Sieve-based approach, which uses rules designed for biomedical texts and predicts a mention as out-of-KB if no in-KB entity can be linked to [9]. For the LLM-based method, we follow BLINKout [7] to detect out-of-KB mentions from texts adapting a two-step BERT-based approach [32]: candidate generation with bi-encoder and candidate selection with cross-encoder. Out-of-KB mentions are discovered through NIL entity representation and classification in the cross-encoder [7]. We used default parameters with top- k value as 50 and domain-specific model, SapBERT [23].

Results on Out-of-KB Mention Discovery Table 3 shows that BLINKout performs much better than the Sieve-based approach in

terms of the overall accuracy and out-of-KB F_1 scores. However, it is still challenging to achieve satisfying performance to identify out-of-KB mentions (with out-of-KB F_1 between 15% and 30%).

For Concept Placement We use the mention-edge pairs (see Table 2) to train and validate a model to match an *in-KB* mention to its gold-standard directed edges in a KB and then test on *out-of-KB* mentions, following the unsupervised setting. The model architecture includes edge candidate generation with an optional step of edge selection. For edge candidate generation, we adapt the bi-encoder [32], with the input of a contextual mention and an edge (i.e., edge-bi-encoder), to match a contextual mention to a directed edge in an ontology using their concept names¹⁵. Top- k edges rankings are selected after this step. For an optional edge selection among the top- k , we test the capability of zero-shot prompting of an LLM, GPT-3.5 (“gpt-3.5-turbo”), where k is set as 50. The prompt includes a header, the mention with contexts, and the top- k candidate edges to query the LLM to select the correct edges¹⁶.

Results on Concept Placement Table 4 suggests that concept placement as edge prediction is very challenging. Also, using GPT-3.5 to select top-1 from the top-50 edge candidates does not improve, or only improves marginally, the results with the prompts. This may suggest the limitation of the state-of-the-art LLM interacting with formal, domain-specific knowledge using zero-shot prompting.

6 CONCLUSION AND FUTURE STUDIES

This work introduced a new benchmark for Ontology Enrichment from Texts by Concept Discovery and Placement. The dataset focuses on enriching OWL ontologies as formal KBs, which are reducible to and thus compatible with taxonomies. Compared to the prior art, the dataset supports a more comprehensive set of characteristics, including NIL Discovery, Contextual Term, Concept Placement, and Complex Concepts. We propose a pipeline to construct this resource and release a dataset using MedMentions corpus (PubMed abstracts), UMLS and SNOMED CT ontologies. We provide usage of the data by evaluating recent LLM-based methods.

The data construction method can be applied to other KBs in the biomedical domain and KBs in various domains.

The baseline LLM-based methods are yet to achieve satisfying performance on the benchmark. Further methods are encouraged to address this challenge.

Acknowledgements. This work is supported by EPSRC projects, including ConCur (EP/V050869/1), OASIS (EP/S032347/1), UK FIRES (EP/S019111/1); and Samsung Research UK (SRUK).

¹⁵An edge is represented as “parent tokens [P-TAG] child tokens [C-TAG]”.

¹⁶Further details, parameter settings, and prompts of the experiments are available at <https://github.com/KRR-Oxford/OET>.

References

- [1] Franz Baader, Ian Horrocks, Carsten Lutz, and Uli Sattler. 2017. *A Basic Description Logic*. Cambridge University Press, Cambridge, 10–49. <https://doi.org/10.1017/9781139025355.002>
- [2] Franz Baader, Ian Horrocks, Carsten Lutz, and Uli Sattler. 2017. *Ontology Languages and Applications*. Cambridge University Press, Cambridge, 205–227. <https://doi.org/10.1017/9781139025355.008>
- [3] Marco Basaldella, Fangyu Liu, Ehsan Shareghi, and Nigel Collier. 2020. COMETA: A Corpus for Medical Entity Linking in the Social Media. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 3122–3137. <https://doi.org/10.18653/v1/2020.emnlp-main.253>
- [4] Olivier Bodenreider. 2004. The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Research* 32, suppl_1 (01 2004), D267–D270. <https://doi.org/10.1093/nar/gkh061>
- [5] Javier Castell-Diaz, Jose Antonio Miñarro-Giménez, and Catalina Martínez-Costa. 2023. Supporting SNOMED CT postcoordination with knowledge graph embeddings. *Journal of Biomedical Informatics* 139 (2023), 104297. Data available at <https://github.com/JavierCastellD/SemanticFormalizationSNOMED>.
- [6] Jiaoyan Chen, Yuan He, Yuxia Geng, Ernesto Jiménez-Ruiz, Hang Dong, and Ian Horrocks. 2023. Contextual semantic embeddings for ontology subsumption prediction. *World Wide Web* (2023), 1–23.
- [7] Hang Dong, Jiaoyan Chen, Yuan He, Yinan Liu, and Ian Horrocks. 2023. Reveal the Unknown: Out-of-Knowledge-Base Mention Discovery with Entity Linking. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (Birmingham, United Kingdom). Association for Computing Machinery, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3583780.3615036>
- [8] Kevin Donnelly et al. 2006. SNOMED-CT: The advanced terminology and coding system for eHealth. In *Medical and Care Compunctics 3*. Studies in health technology and informatics, Vol. 121. IOS Press, Amsterdam, Netherlands, 279–290.
- [9] Jennifer D'Souza and Vincent Ng. 2015. Sieve-Based Entity Linking for the Biomedical Domain. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Association for Computational Linguistics, Beijing, China, 297–302. <https://doi.org/10.3115/v1/P15-2049>
- [10] Maud Ehrmann, Matteo Romanello, Alex Flückiger, and Simon Clematide. 2020. Extended overview of CLEF HIPE 2020: named entity processing on historical newspapers. In *Working Notes of CLEF 2020 - Conference and Labs of the Evaluation Forum*. CEUR-WS.org, RWTH Aachen, Aachen, Germany, 1–38. Data available at <https://github.com/impresso/CLEF-HIPE-2020/tree/master/data>.
- [11] Martin Glauer, Adel Memariani, Fabian Neuhaus, Till Mossakowski, and Janna Hastings. 2022. Interpretable ontology extension in chemistry. *Semantic Web Preprint* (2022), 1–22. Data available at <https://zenodo.org/record/6023497>.
- [12] Bernardo Cuenca Grau, Ian Horrocks, Boris Motik, Bijan Parsia, Peter Patel-Schneider, and Ulrike Sattler. 2008. OWL 2: The next step for OWL. *Journal of Web Semantics* 6, 4 (2008), 309–322. Semantic Web Challenge 2006/2007.
- [13] J. Hastings. 2023. *AI for Scientific Discovery*. CRC Press, Milton. <https://doi.org/10.1201/9781003226642>
- [14] Yuan He, Jiaoyan Chen, Hang Dong, Ian Horrocks, Carlo Allocca, Taehun Kim, and Brahmananda Sapkota. 2023. DeepOnto: A Python Package for Ontology Engineering with Deep Learning. *arXiv preprint arXiv:2307.03067* (2023).
- [15] Yuan He, Jiaoyan Chen, Ernesto Jimenez-Ruiz, Hang Dong, and Ian Horrocks. 2023. Language Model Analysis for Ontology Subsumption Inference. In *Findings of the Association for Computational Linguistics: ACL 2023*. Association for Computational Linguistics, Toronto, Canada, 3439–3453. <https://doi.org/10.18653/v1/2023.findings-acl.213>
- [16] Nicolas Heist and Heiko Paulheim. 2023. NASTyLinker: NIL-Aware Scalable Transformer-Based Entity Linker. In *The Semantic Web - 20th International Conference, ESWC 2023, Hersonissos, Crete, Greece, May 28 - June 1, 2023, Proceedings (Lecture Notes in Computer Science, Vol. 13870)*, Catia Pesquita, Ernesto Jiménez-Ruiz, Jamie P. McCusker, Daniel Faria, Mauro Dragoni, Anastasia Dimou, Raphaël Troncy, and Sven Hertling (Eds.). Springer, Cham, 174–191.
- [17] Anastasiia Iurshina, Jiaxin Pan, Rafika Boutalbi, and Steffen Staab. 2022. NILK: Entity Linking Dataset Targeting NIL-Linking Cases. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (Atlanta, GA, USA). Association for Computing Machinery, New York, NY, USA, 4069–4073. <https://doi.org/10.1145/3511808.3557659> Data available at <https://zenodo.org/record/6607514>.
- [18] Heng Ji, Ralph Grishman, Hoa Trang Dang, Kira Griffitt, and Joe Ellis. 2011. Overview of the TAC 2011 knowledge base population track. In *Third text analysis conference (TAC 2011)*. National Institute of Standards and Technology, Gaithersburg, Maryland, USA, 1–33.
- [19] Minhao Jiang, Xiangchen Song, Jieyu Zhang, and Jiawei Han. 2022. TaxoEnrich: Self-Supervised Taxonomy Completion via Structure-Semantic Representations. In *Proceedings of the ACM Web Conference 2022* (Virtual Event, Lyon, France) (WWW '22). Association for Computing Machinery, New York, NY, USA, 925–934. <https://doi.org/10.1145/3485447.3511935>
- [20] David Jurgens and Mohammad Taher Pilehvar. 2016. 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, 1092–1102. <https://doi.org/10.18653/v1/S16-1169> Data available at <https://alt.qcri.org/semeval2016/task14/>.
- [21] Nora Kassner, Fabio Petroni, Mikhail Plekhanov, Sebastian Riedel, and Nicola Cancedda. 2022. EDIN: An End-to-end Benchmark and Pipeline for Unknown Entity Discovery and Indexing. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 8659–8673. Data available at <https://github.com/facebookresearch/EDIN>.
- [22] Rohit J Kate. 2020. Automatic full conversion of clinical terms into SNOMED CT concepts. *Journal of Biomedical Informatics* 111 (2020), 103585. Data available at <https://sites.uwm.edu/katerj/conversion/>.
- [23] Fangyu Liu, Ehsan Shareghi, Zaiqiao Meng, Marco Basaldella, and Nigel Collier. 2021. Self-Alignment Pretraining for Biomedical Entity Representations. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 4228–4238. <https://doi.org/10.18653/v1/2021.naacl-main.334>
- [24] Hao Liu, Yehoshua Perl, and James Geller. 2020. Concept Placement Using BERT Trained by Transforming and Summarizing Biomedical Ontology Structure. *J. of Biomedical Informatics* 112, C (2020), 103607. <https://doi.org/10.1016/j.jbi.2020.103607>
- [25] Sunil Mohan and Donghui Li. 2019. MedMentions: A Large Biomedical Corpus Annotated with UMLS Concepts. In *1st Conference on Automated Knowledge Base Construction, AKBC 2019, Amherst, MA, USA, May 20-22, 2019*. openreview.net, Amherst, MA, USA, 1–13. <https://doi.org/10.24432/C5G59C>
- [26] Giuseppe Rizzo, Amparo Elizabeth Cano Basave, Bianca Pereira, and Andrea Varga. 2015. Making Sense of Microposts (# Microposts2015) Named Entity rCognition and Linking (NEEL) Challenge. In *5th Workshop on Making Sense of Microposts (#Microposts2015)*. CEUR-WS.org, RWTH Aachen, Aachen, Germany, 44–53. Data available at https://ceur-ws.org/Vol-1395/microposts2015_neel-challenge-report/.
- [27] Giuseppe Rizzo, Bianca Pereira, Andrea Varga, Marieke Van Erp, and Amparo Elizabeth Cano Basave. 2017. Lessons learnt from the Named Entity rCognition and Linking (NEEL) challenge series. *Semantic Web* 8, 5 (2017), 667–700.
- [28] Giuseppe Rizzo, Marieke Van Erp, Julien Plu, and Raphaël Troncy. 2016. Making Sense of Microposts (# Microposts2015) Named Entity rCognition and Linking (NEEL) Challenge. In *6th Workshop on Making Sense of Microposts*. CEUR-WS.org, RWTH Aachen, Aachen, Germany, 50–59. Data available at https://ceur-ws.org/Vol-1691/microposts2016_neel-challenge-report/.
- [29] Jiaming Shen, Zhihong Shen, Chenyan Xiong, Chi Wang, Kuansan Wang, and Jiawei Han. 2020. TaxoExpan: Self-Supervised Taxonomy Expansion with Position-Enhanced Graph Neural Network. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) (WWW '20). Association for Computing Machinery, New York, NY, USA, 486–497. <https://doi.org/10.1145/3366423.3380132> Data available at <https://github.com/mickeysjm/TaxoExpan>.
- [30] Hanna Suominen, Sanna Salanterä, Sumithra Velupillai, Wendy Webber Chapman, Guergana K. Savova, Noemie Elhadad, Sameer Pradhan, Brett R. South, Danielle L. Mowery, Gareth J. F. Jones, Johannes Leveling, Liadh Kelly, Lorraine Goeriot, David Martínez, and Guido Zuccon. 2013. Overview of the ShARe/CLEF eHealth Evaluation Lab 2013. In *Information Access Evaluation. Multilinguality, Multimodality, and Visualization - 4th International Conference of the CLEF Initiative, CLEF 2013*, Pamela Forner, Henning Müller, Roberto Paredes, Paolo Rosso, and Benno Stein (Eds.), Vol. 8138. Springer, Valencia, Spain, 212–231. https://doi.org/10.1007/978-3-642-40802-1_24 Data available at <https://www.physionet.org/content/shareclefehealth2013/1.0/>.
- [31] Suyuchen Wang, Ruihui Zhao, Yefeng Zheng, and Bang Liu. 2022. QEN: Applicable Taxonomy Completion via Evaluating Full Taxonomic Relations. In *Proceedings of the ACM Web Conference 2022* (Virtual Event, Lyon, France) (WWW '22). Association for Computing Machinery, New York, NY, USA, 1008–1017. Data available at <https://github.com/sheryc/QEN>.
- [32] Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020. Scalable Zero-shot Entity Linking with Dense Entity Retrieval. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 6397–6407. <https://doi.org/10.18653/v1/2020.emnlp-main.519>
- [33] Qingkai Zeng, Jinfeng Lin, Wenhao Yu, Jane Cleland-Huang, and Meng Jiang. 2021. Enhancing Taxonomy Completion with Concept Generation via Fusing Relational Representations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (Virtual Event, Singapore). Association for Computing Machinery, New York, NY, USA, 2104–2113. <https://doi.org/10.1145/3447548.3467308> Data available at <https://github.com/DM2-ND/GenTaxo>.
- [34] Jieyu Zhang, Xiangchen Song, Ying Zeng, Jiaze Chen, Jiaming Shen, Yuning Mao, and Lei Li. 2021. Taxonomy completion via triplet matching network. In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Press, Palo Alto, California, USA, 4662–4670. Data available at <https://github.com/JieyuZ2/TMN>.