

Collaborative-AI Knowledge Graph Generation: Taxonomization of IATE, the EU Terminology

Michael Wetzel^{1*}, Stelios Mathioulakis², Jochen Hummel^{1,2}, and Alena Vasilevich¹[0000-0002-9769-1885]

¹ Coreon GmbH, Berlin, Germany
michael@coreon.com alena@coreon.com
<https://www.coreon.com>

² ESTeam AB, Gothenburg, Sweden
stelios@esteam.se jochen@esteam.se
<https://www.esteam.se/>

Abstract. Formalized knowledge is a powerful resource for AI projects, but it is usually created at great expense. Taxonomization is linking a flat set of concepts into a hierarchical knowledge graph, and in this work, we present our approach to semi-automatic generation of such concept maps, elevating a sub-domain of IATE terminology into a multilingual knowledge graph. We taxonomized a flat list of concepts within the COVID sub-domain, benchmarking two approaches to tackle this task: automatic concept map creation using an enhanced ML-powered language model and manual creation of the graph by a linguist expert. We dwell on advantages of the collaborative method, made easy by a user-friendly UI, and show how the achieved productivity rate can make taxonomization of large terminology databases viable.

Keywords: Collaborative taxonomization · Knowledge Graph UI · KG Construction Methods · Machine Learning

1 Introduction

In the realm of data-driven businesses, structured data is a valuable resource. IATE (Interactive Terminology for Europe)¹, with almost one million concepts storing multilingual terms and metadata, holds a large part of the textual knowledge of the EU (European Union). However, while IATE database has domain crosslinks with the EuroVoc thesaurus², its concepts are not linked within: they stand alone and can only be accessed lexically.

If IATE were taxonomized, i.e. related concepts linked up into knowledge graphs yielding a full-fledged taxonomy, its data would not only be consumed

* All authors contributed equally.

Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹ Available at <https://iate.europa.eu/>

² <http://europa.eu/eurovoc/>

by linguists but also become machine-readable. This would convert IATE into a powerful resource for AI projects, particularly within the segment of small and medium-sized enterprises (SMEs) that rarely have the means to create multilingual formalized knowledge – building and maintenance of big taxonomies require not only vast financial support but also supervision of domain-specific experts.

State-of-the-art techniques for automatic taxonomy generation utilize information extraction and pattern-based methods [5, 8], combinations of tags and Wikipedia Category Hierarchy [2, 7], already existing knowledge graph’s triple representations [10], and strict domain-specific properties: having a list of terms for taxonomization, they rely on a corpus containing these terms, assuming that the corpus represents the domain adequately [6, 12]. To our knowledge, none of the published studies solved the ‘out-of-context’ terms problem, which is exactly our case since we were dealing with a list of concepts without any established relations or domain corpora at hand.

The objective of this study was therefore to investigate if, given a flat list of Coronavirus (COVID) sub-domain concepts without relations between them, we could facilitate a speedy creation of a deeply-structured taxonomy, turning the input concepts into a mono-hierarchical knowledge graph.

2 Methodology and Data

To test our hypothesis, we chose two scenarios: manual and semi-automatic taxonomization. In the first scenario, the setup was the following: a linguist took a flat list of concepts and turned them into a structured taxonomy, establishing broader-narrower relations between the concepts using Coreon UI³. In the second scenario, an ML-powered algorithm drafted the taxonomy, which was later curated by another linguist with the help of the same tool. In both cases, the taxonomies feature not only initial input concepts but also the established relations between them and higher-level parent nodes that bring a meaningful structure into the taxonomy. When given the same set of concepts and asked to create a taxonomy, people are likely to produce different results due to our natural diversity of views on how concepts shall be categorized. We therefore did not aim for an in-depth review of semantic coherence within the created concept clusters, focusing rather on comparison of measurable parameters, such as working hours, amount of “transactions” in the software, number of relations created. To make the results comparable to some degree, we established that the resulting taxonomies should share 5 identical top-level nodes (i.e. “entities”, “instruments”, “medicine”, “social aspects”, and “miscellaneous”) and that parent nodes may not have more than 20-25 children. It should also be noted that both curators are not domain-knowledge experts, yet they were well familiar with the used software. Figure 1 illustrates the steps executed for both scenarios. Exported 424 IATE COVID concepts, used in both cases as initial input, were identical.

³ <https://www.coreon.com/>

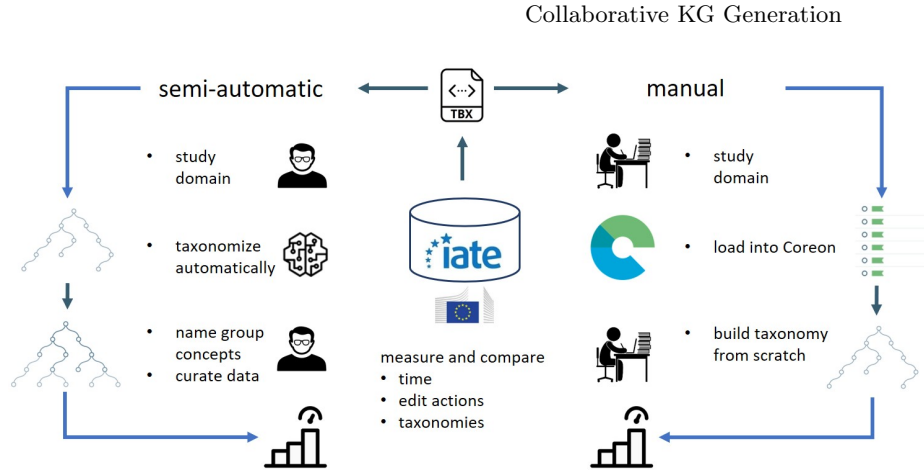


Fig. 1. Comparing two taxonomization approaches.

2.1 Semi-automatic Taxonomization

Inspired by the work done on hierarchical clustering, collaborative tagging systems, unsupervised taxonomy construction, and graph-based exploration of semantic spaces [4, 13, 15], we have developed a novel ML approach for automatic taxonomy induction, using context-free tokens, i.e. terms and synonyms comprising each IATE COVID concept, as input data. First, input is vectorized, so text tokens representing each concept are converted into dense vectors of numbers. Since our dataset contains many rare disease-specific tokens, we used a pre-trained and publicly available FastText word embedding (WE) model to convert terms into vectors [9]. Next, we calculate a pairwise cosine similarity for all WE combinations, forming a symmetric matrix of similarities. To create a graph, each node is represented by an input concept, and to form the vertices/edges of the graph, a threshold is used, i.e. a median cosine similarity of all the given similarities [11]. If two concepts have a similarity higher than the determined threshold, an edge is formed between the nodes that correspond to these two concepts. The formed edge will then have the similarity score as a weight attribute. Ultimately, all combinations of concept nodes are checked to see whether a connection between them will be formed, hence resulting in an undirected weighted graph. Subsequently, Louvain algorithm is used on the aforementioned graph to reveal nodes that are more closely related than others [1, 3]. The process is executed recursively in order to further break down clusters into subclusters; with every such division, an intermediate node is created until there is nothing to break down. The produced taxonomy starts with a root node and expands until the leaf nodes reach the input concepts. The automatically generated intermediate branch nodes are labelled with temporary IDs that are to be named manually. The algorithm is essentially forming concept clusters, as concept nodes form clusters with a common parent node. Also, neighbouring cluster concepts are semantically closer than clusters farther apart in the taxon-

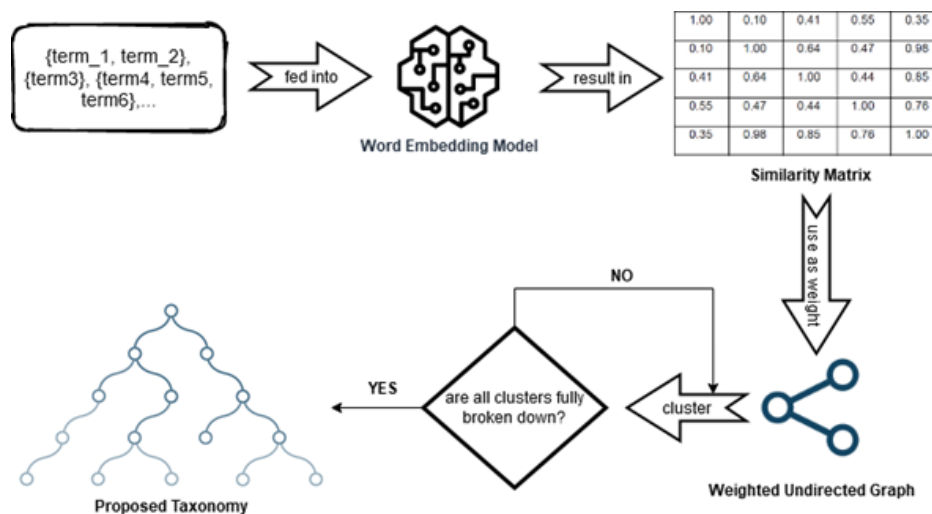


Fig. 2. ML-driven taxonomization process.

omy. Figure 2 summarizes the described taxonomization process; input concepts are represented by sets of terms in curly brackets.

The resulting taxonomy was revised by a curator; input concepts became leaf nodes, grouped into the generated 55 higher level nodes with unassigned temporary labels. Curator’s approach was to traverse clustered leaf nodes and assign meaningful names to each of 55 parent concepts, replacing temporary labels. Table 1 quantifies the human effort spent on the curation of the automatically created taxonomy. Even though we did not quantify the semantic soundness of the generated clusters, it is worth noting that most of them looked accurate to the curators, whereas grouping errors were likely triggered by corpora, used to pre-train ML models we utilized for taxonomization (e.g., ‘interstitial space’, space between cells, and ‘hospital pharmacy’ were wrongly clustered together, as for WE both concepts are spaces, appearing in the similar semantic neighborhoods).

2.2 Manual Taxonomization

When approaching manual taxonomization, curator’s strategy was to move from concept to concept, starting from the top level nodes and establishing broader-narrower relations, later dividing the nodes further down in sub-trees. This top-down method requires several passes through the data. Initially, concepts were sorted alphabetically which meant the taxonomist had to constantly jump from one context to another (e.g., from ‘border worker’ to ‘bronchoalveolar lavage’ to ‘budding’– concepts that are semantically far off). Table 1 provides the effort metrics of manual taxonomization (excluding all preparatory steps and domain

learning curve). While being a tedious challenge, the created taxonomy is a pure result of the intellectual human capacity, not influenced by additional obscure models or primed by pre-trained algorithms.

3 Results

Coreon UI was used for both building and curation of the presented taxonomies. Since it records any modifications in the data repository, change logs were retrieved and analyzed to determine precisely what actions the taxonomists performed and how much time was spent. Table 1 demonstrates such information for both scenarios. When comparing the numbers yielded by both approaches, we see that starting with a pre-drafted knowledge graph, the curator was five times faster, not only counting pure working hours but also when comparing the amount of transactions, i.e. events, mouse-clicks in the software.

Table 1. Taxonomization Effort.

Metrics	Manual Taxonomization	Semi-automatic Taxonomization
Curator’s time spent	40h	8h
Number of relations changed	1417	432
Number of new concepts created	115	28

When starting with an automatically pre-drafted taxonomy, curators are not faced with a long list of to-be-processed concepts; they rather work cluster by cluster. Even when clusters are off, the topic and focus are stable, and there is no jumping between contexts (see Figure 3). Also, the algorithm encourages




Fig. 3. IATE concept ‘containment measures’ becomes a group concept.

a strictly hierarchical tree-like concept system (mono-hierarchical in this case),

whereas in the manual method, the curator often hooked concepts under more than one parent. Our collaborative-AI approach can also facilitate parallel working, allowing distributed curation of separate sub-trees. Ultimately, the curator benefits from the UI that allows to comfortably restructure the resulting knowledge graph (see Figure 4).

4 Conclusion

Our study proves that it is possible to build a taxonomy while having nothing but a flat list of context-free tokens to taxonomize. We demonstrated that the suggested collaborative taxonomization approach – combining ML and human curation – can significantly bring down the effort while yielding a taxonomy of a consistent quality. Given availability of supporting textual data, the achieved performance and resource-saving advantages of our custom method makes taxonomization of even larger terminology databases viable. Every taxonomization project revolves around a specific domain, dictated by the vocabulary of the terms at hand (virology, general medical terms, etc.).

In future work, to create more accurate vector representations of domain tokens, the specific context in which they apply can be exploited: e.g., it can be used to re-train WE model instead of using a more generic pre-trained one, primed by generic Wikipedia texts. Literature suggests that such re-training usually yields better results [14]. Consistent presence of metadata in the terminology database can also aid auto-taxonomization, significantly reducing curators' toil (e.g. metadata about a semantic parent of a term or free-text concept definitions can help with parent concept identification and/or naming, extrapolating a potential parent instead of assigning concept clusters an abstract temporary ID-name).

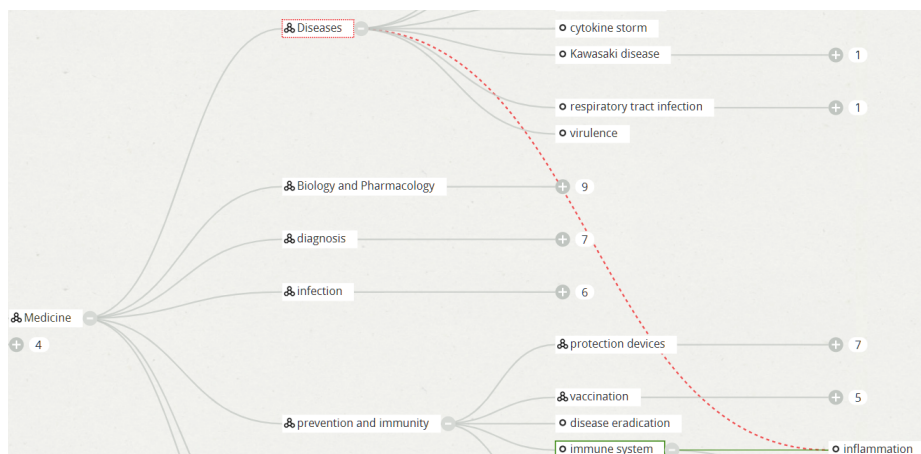


Fig. 4. New high-level node ‘Diseases’.

References

1. Blondel, V. D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* (10), P10008 (2008). <https://doi.org/10.1088/1742-5468/2008/10/p10008>
2. Carrion, B., Onorati, T., Díaz, P., Triga, V.: A taxonomy generation tool for semantic visual analysis of large corpus of documents. *Multimedia Tools and Applications*, **78**(23), 32919–32937 (2019)
3. Csardi, G., Nepusz, T.: The igraph software package for complex network research. *InterJournal (Complex Systems)*, 1695 (2006)
4. Heymann, P., Garcia-Molina, H.: Collaborative creation of communal hierarchical taxonomies in social tagging systems. Stanford (2006)
5. Iloga, S., Romain, O., Tchuente, M.: A sequential pattern mining approach to design taxonomies for hierarchical music genre recognition. *Pattern Analysis and Applications* **21**, 363–380 (2018). <https://doi.org/10.1007/s10044-016-0582-7>
6. Iloga, S., Romain, O., Tchuente, M.: An efficient generic approach for automatic taxonomy generation using HMMs. *Pattern Analysis and Applications*, **24**(1), 243–262 (2021)
7. Kotlerman, L., Avital, Z., Dagan, I., Lotan, A., Weintraub, O.: A support tool for deriving domain taxonomies from wikipedia. In: *Proceedings of the International Conference Recent Advances in Natural Language Processing*, pp. 503–508. (2011)
8. Liu, X., Song, Y., Liu, S., Wang, H.: Automatic taxonomy construction from keywords. In: *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1433–1441. (2012)
9. Mikolov, T., Grave, E., Bojanowski, P., Puhersch, Ch., Joulin, A.: *Advances in Pre-Training Distributed Word Representations*. In: *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*
10. Pietrasik M., Reformat, M.: A simple method for inducing class taxonomies in knowledge graphs. In: *European Semantic Web Conference*, vol. 12123, pp. 53–68. Springer, Cham (2020).
11. Ryu, P.M., Choi, K.S.: Taxonomy learning using term specificity and similarity. In: *Proceedings of the 2nd Workshop on Ontology Learning and Population: Bridging the Gap between Text and Knowledge*, pp. 41–48. (2006)
12. Treeratpituk, P., Khabsa, M. and Giles, C.L.: Graph-based approach to automatic taxonomy generation (grabtax). arXiv preprint arXiv:1307.1718. (2013)
13. Veremyev, A., Semenov, A., Pasiliao, E.L., Boginski, V.: Graph-based exploration and clustering analysis of semantic spaces. *Applied Network Science*, **4**(1), 1–26 (2019)
14. Yin, W., Schuetze, H.: Learning Word Meta-embeddings. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pp. 1351–1360. (2016)
15. Zhang, C., Tao, F., Chen, X., Shen, J., Jiang, M., Sadler, B., Vanni, M., Han, J.: Taxogen: Unsupervised topic taxonomy construction by adaptive term embedding and clustering. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2701–2709. (2018)