

Mapping the Past: Geographically Linking an Early 20th Century Swedish Encyclopedia with Wikidata

Axel Ahlin[†], Alfred Myrne[†], Pierre Nugues

Lund University

Lund, Sweden

{ax5047ah-s, al5247my-s}@student.lu.se, Pierre.Nugues@cs.lth.se

Abstract

In this paper, we describe the extraction of all the location entries from a prominent Swedish encyclopedia from the early 20th century, the *Nordisk Familjebok* ‘Nordic Family Book.’ We focused on the second edition called *Uggleupplagan*, which comprises 38 volumes and over 182,000 articles. This makes it one of the most extensive Swedish encyclopedias. Using a classifier, we first determined the category of the entries. We found that approximately 22 percent of them were locations. We applied a named entity recognition to these entries and we linked them to Wikidata. Wikidata enabled us to extract their precise geographic locations resulting in almost 18,000 valid coordinates. We then analyzed the distribution of these locations and the entry selection process. It showed a higher density within Sweden, Germany, and the United Kingdom. The paper sheds light on the selection and representation of geographic information in the *Nordisk Familjebok*, providing insights into historical and societal perspectives. It also paves the way for future investigations into entry selection in different time periods and comparative analyses among various encyclopedias.

Keywords: entity annotation, entity linking, named entity recognition

1. Introduction

The Swedish *Nordisk Familjebok* encyclopedia ‘Nordic Family Book’ was first published in the late 19th century. It holds a significant place in the history of Swedish literature and is widely regarded as one of the most comprehensive and authoritative encyclopedic works in Swedish. The *Nordisk Familjebok* was originally created with the purpose of providing accurate knowledge and information to the Swedish public.

1.1. The *Uggleupplagan* Edition

The inaugural edition, commonly referred to as the “First Edition,” was published from 1876 to 1899, spanning 20 volumes. The encyclopedia aimed to encompass a wide array of subjects, including history, geography, science, literature, arts, and more. It featured thoroughly written articles authored by numerous experts in their respective fields. As such, it can be regarded as a text that reflects the dominating worldview in Sweden in the early 20th century.

The second edition of the encyclopedia, also called *Uggleupplagan*, ‘Owl Edition’, was published between 1904 and 1926 and served as a revised and expanded version of the original encyclopedia. This edition comprised 38 volumes and boasted an extensive range of articles, illustrations, maps, and photographs, providing a comprehensive and visually captivating resource. It is

the most extensive encyclopedia ever printed in Swedish (Aronsson, 2003).

The encyclopedia also received a third, fourth, and fifth edition, in the years 1923–1939, 1951–1957, and 1993, respectively. However, the second edition was the most influential, and is therefore the focus of this study.

In this work, we extracted the named entities from the *Uggleupplagan* encyclopedia and we analyzed the qualitative trends in the selection of entries. More specifically, we focused on the geographic entries.

1.2. Contributions

The contributions of this paper are:

1. We cleaned a 20th century encyclopedia from a raw OCR text and we structured it in the JSON format;
2. We extracted the entries corresponding to named geographic entities using a combination of pre-trained models and fine-tuned classifiers;
3. We mapped these entities to Wikidata items using the entry text and the Swedish description in Wikidata. We first fetched candidates from Wikidata; we ranked them using a semantic cosine similarity; and we selected the candidate with the highest score;
4. We extracted the coordinates from the chosen candidates and we visualized these entries on

[†] Equal contribution

a map;

5. We published our dataset on GitHub¹.

We found that most geographic entities in the encyclopedia are located in Sweden, as well as in other places in Europe. We also noted a specific focus on Germany and the United Kingdom. This corresponds to countries where Sweden had strong ties at the time of the encyclopedia's production, including historical, scientific, cultural, and economic relations.

2. Previous Work

In this work, we used the digital version of an encyclopedia. We applied a binary categorization to extract the location entries, and we then linked these entries to Wikidata.

2.1. Digitized Encyclopedias

A complete digitized version of *Uggleupplagan* is available from Projekt Runeberg². This project focuses on making older Nordic literature freely available to the public through digital means and *Nordisk Familjebok* is one of them. It is operated as a non-profit organization within the University of Linköping (Nationalencyklopedin, 2023).

The digitization had four main steps:

1. The Runeberg editors first scanned a paper copy of the encyclopedia;
2. They applied an optical character recognition (OCR) to extract the text from the image scans;
3. They structured the text into entries and made them available by headwords. Readers can then access the content from a web browser;
4. Finally, volunteers proofread the machine-generated text, but a majority of entries has not yet been proofread. The scanned pages are openly available for verification.

2.2. Entry Categorization

The *Nordisk Familjebok* entries consist of common as well as proper nouns and the first step is to classify them in these two categories: Whether they are a location or not.

Transformers architectures (Vaswani et al., 2017), and notably the Bidirectional Encoder Representations from Transformers model (BERT)

¹<https://github.com/axelahlin/uggleupplagan>

²<https://runeberg.org/nf/>

(Devlin et al., 2019) have produced the highest performances on classification tasks such as those of the GLUE benchmark (Wang et al., 2018).

The BERT architecture makes it possible to build large pre-trained models that users can then further fine-tune on classification tasks as well as on entity recognition.

2.2.1. BERT

BERT utilizes a transformer-based architecture restricted to its encoder part. From an input consisting of a sequence of words, this architecture captures the contextual relationships between them bidirectionally and maps each word to a contextual dense vector. The bidirectional context allows BERT to have a deeper understanding of language and to handle tasks like text classification and named entity recognition more efficiently.

BERT is pre-trained on a large corpus of English text data and, during this process, it learns to predict masked words within sentences and determines the relationships between sentence pairs.

2.2.2. KB-BERT

The BERT architecture has been subsequently pre-trained on many languages outside English. Malmsten et al. (2020) pre-trained BERT models on Swedish text from the National Library of Sweden that they called KB-BERT. They gathered their corpus from diverse sources, including books, news articles, government publications, Swedish Wikipedia, and internet forums.

The KB-BERT pre-trained models served as a source of multiple application tasks for Swedish text. Remmer et al. (2021) used them to classify patient records in Swedish while Johansson (2022) and Nyqvist (2021) applied them to named entity recognition. Bridal et al. (2022) fine-tuned the base KB-BERT model to evaluate the de-identification of Swedish clinical data. Finally, Nielsen (2023) showed that one of the KB-BERT models outperformed all other Swedish models for NER tasks.

2.3. Entity Linking and Disambiguation

One task in this work was to link encyclopedia entities to Wikidata items. Wikidata is a multilingual knowledge base associated to Wikipedia. It contains more than 100 million entries at the time of this study and it has become a central repository for authority data disambiguation and linking (Tharani, 2021). It contains useful properties for each entry. For instance, entities with geographical locations have the P625 property, which contains the longitude and latitude of the entity.

Examples of studies using Wikidata include [Pratapa et al. \(2022\)](#), which linked event descriptions in 44 languages to this repository; [Hamdi et al. \(2021\)](#) has a similar goal with proper nouns in historical documents such as newspapers.

2.3.1. Historical Data

The tasks of classification and entity linking on historical data present unique challenges. This includes OCR errors, different grammar and language rules, and more. [Pontes et al. \(2020\)](#) discusses some of these challenges. [Nugues \(2022\)](#) linked the named entities of a French dictionary from the turn of the 20th century to Wikidata. The author then visualized the geographic coordinates of the location entities on a map.

2.3.2. Named Entity Linking

Named entity linking often proceeds in three steps ([Mrini et al., 2022](#)):

1. The first step is to identify the mentions of named entities in text;
2. For a given mention M , extract a list of candidate entities that could match it: $\{E_1, E_2, \dots, E_n\}$;
3. Finally, rank the candidates and select the first one: $\arg \max_i P(E_i|M)$.

To link entities, [Hoffart et al. \(2011\)](#) used Wikipedia disambiguation links and context similarity scores based on keywords. [Francis-Landau et al. \(2016\)](#) computed a cosine similarity between the embeddings of the mention context and of an entity description taken from a dictionary of entities. [Logeswaran et al. \(2019\)](#) also used an entity mention and an entity description as input. They separated both texts by a [SEP] token and they trained a BERT-based classifier to decide if the mention corresponded to the entity. The authors used the classifier score to rank the candidates.

2.3.3. SBERT

In our work, the list of candidates consisted of Wikidata items and we ranked the Wikidata descriptions from the definition of the encyclopedia entry.

We followed [Francis-Landau et al. \(2016\)](#) and we computed an embedding vector for both descriptions, one from Wikidata and the other from the encyclopedia. Nonetheless, instead of convolutional neural networks, we used transformers.

As transformer, we chose Sentence-BERT (SBERT) ([Reimers and Gurevych, 2019](#)). SBERT is a modified version of BERT that uses both Siamese and triplet network architectures. SBERT

takes a sentence as input and outputs the sentence embedding vector. We can then compare two sentence embeddings using, for example, a cosine similarity.

3. Method

To the best of our knowledge, there exists no corpus of dictionary entries in Swedish annotated with location and links. We could not then rely on using fully-supervised methods. As we wanted to avoid a completely manual annotation procedure, we designed a semi-automatic pipeline, where we used classification and linking.

This pipeline consists of five steps shown in Figure 1:

1. We first scraped the encyclopedia text from the Runeberg website and we organized the dataset as a JSON file of entries;
2. We trained a classifier to determine whether an entry is a location or not. As training set, we annotated manually a set of positive and negative samples; we applied this classifier and we discarded the nonlocations;
3. For each location, we used the encyclopedia text and extracted the headword to query Wikidata for candidate items having the “geographic location” property: $- P_{625} -$;
4. We computed the cosine similarity between each Wikidata item description in Swedish and the encyclopedia text;
5. Finally, we extracted and plotted the geographic coordinates of the item with the highest cosine similarity to the encyclopedia text.

The classification step is necessary to the extraction of geographical entries. Instead, one could be tempted to use the headword to fetch all items with the coordinate property P_{625} . However, this would not be ideal as it would include a vast number of items that may not be relevant.

Specifically, some Wikidata items, such as the *Mona Lisa* item, have coordinates assigned for the P_{625} property. These coordinates represent the geographic location of the item itself, not necessarily meaning that it actually is a location. Hence, we cannot establish an equivalence between an actual geographic location and a Wikidata item having the P_{625} property.

3.1. Scraping

We scraped a digitized version of the *Uggleupplagan* from the Projekt Runeberg website³. It comprises the 38 volumes available as text files. The

³<https://runeberg.org/nf/>

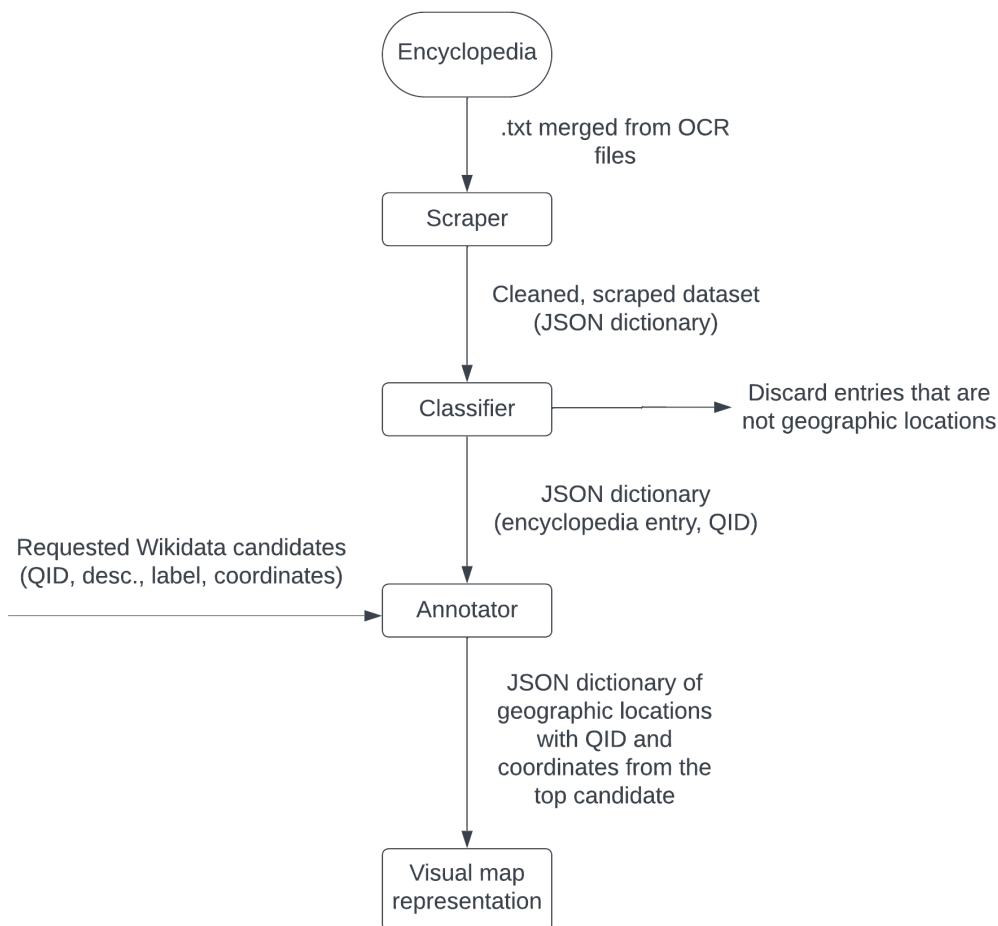


Figure 1: Architecture of the pipeline

Uggleupplagan contains over 182,000 entries (Simonsen, 2016). The entries had an average of 15.51 words, including the entry title word. This resulted in an average length of 102.19 characters per entry. For each entry, we cut the text description after 200 characters and we removed all text after the final period (“.”) character.

We noticed that longer descriptions displayed a tendency to confuse the classifier by providing superfluous context and leading to incorrect results. This is why we imposed this character limit on the entries as we believe this represents a good trade-off between necessary context for the classifier and classifier performance.

3.2. Classifier

For the geographic location classifier, we first annotated a dataset of encyclopedic entries as location or nonlocation with Boolean values. We then fine-tuned the baseline KB-BERT model to categorize the complete set of entries.

To create the classifier, we applied the pre-

trained model to the annotated entries and we used the resulting hidden states to fit a logistic regression model. We then applied the resulting model to categorize all the entries of the encyclopedia. Finally, we purged the entries that were not geographic locations.

3.3. Entity Candidate Ranking and Linking

We linked each encyclopedia entry classified as a geographic location to a Wikidata item. A Wikidata item has a unique identifier (QID) consisting of a Q prefix and a number, for instance Q1754 for *Stockholm*. Each Wikidata item has attributes describing it such as a plain text description in multiple languages. These attributes use the P prefix and a number, such as P31 for the `instance of` property, the type of the item. Wikidata location items have very often geographic coordinates with the P625 attribute.

We can search Wikidata with string queries as with a search engine. It then returns a set of items

with the Wikidata identifiers (QID).

To link the encyclopedia entries with Wikidata objects, we proceeded in two steps:

1. For each entry, we extracted the headword or first word and we used it to query Wikidata. We kept up to five items from all the candidates retrieved by the Wikidata query;
2. We then compared the encyclopedia entry text with the Swedish description of each Wikidata item:
 - We encoded the texts, encyclopedia and Wikidata, as embedded vectors with the SBERT model `all-MiniLM-L6-v2`. We believe this model represents a good trade-off between speed and performance.
 - We computed a cosine similarity for all the pairs.

We selected the candidate with the highest cosine similarity and linked the entry to the corresponding Wikidata identifier.

3.4. Querying Wikidata

For all the locations we could extract from the encyclopedia using the classifier, we fetched the corresponding geographic coordinates from Wikidata.

In addition to the string search queries, we can query Wikidata with the SPARQL graph database language. We used SPARQL to create a query for coordinate retrieval:

```
?item wdt:P625 ?coords
```

where `?item` is the QID, `wdt:P625`, the coordinate location property, and `?coords`, the coordinates to extract.

Figure 3 shows the coordinates on a world map of all the location entries in the encyclopedia.

4. Results

Our scraping resulted in 130,383 entries. We applied the classifier to these entries and we obtained 28,284 locations. This means that approximately 21.7 percent of the entries are geographic locations. Once linked, we extracted 17,793 valid coordinates. Figure 3 shows a geographic representation of the world, where we plotted all extracted coordinate points, while Figure 2 shows the coordinate distance from each retrieved coordinate to the geographic center of Sweden.

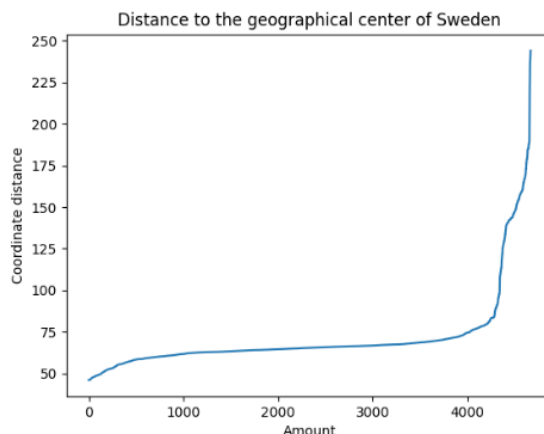


Figure 2: Distribution of entries by geographic distance to Sweden

Normalized confusion matrix

True label \ Pred. label	Location	Not location
Location	0.93	0.07
Not location	0.06	0.94

Table 1: Confusion matrix for the evaluated classifier

5. Evaluation

5.1. Classifier

We evaluated the classifier using a manually annotated test set of 200 previously unseen entries. Table 1 shows the confusion matrix and Table 2 shows normalized metrics from the statistical report of the classifier.

Although not perfect, the classification reaches high results in all metrics. Nonetheless with a large corpus such as an encyclopedia, mislabeling of entries results in locations being ignored and incorrect entries being included. The precision score indicates that a high percentage of the identified locations are indeed valid geographic locations, while the recall score indicates that a high percentage of the actual geographic locations present in the encyclopedia were correctly identified and subsequently marked on the map, as shown in Figure 3. The F1 score also indicates that we strike a good balance between accurately identifying locations (precision) and not missing many valid locations (recall).

5.2. Scraper

There is a discrepancy between our scraping and the total number of entries in the encyclopedia. This is in part due to some subentries being under

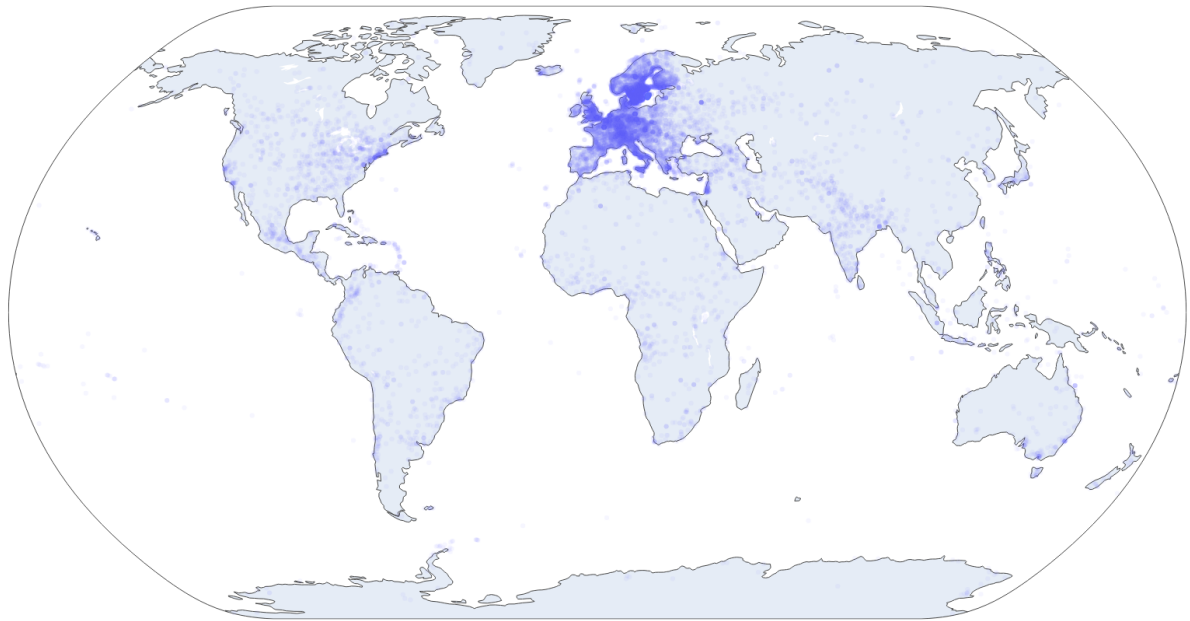


Figure 3: Visual representation of the encyclopedia’s geographic locations

Accuracy	0.935
Precision	0.939
Recall	0.930
F1-score	0.935

Table 2: Performance metrics for the classifier.

one entry, such as different members of one family all being in the surname. These are counted as one entry in our dataset. However, it may also be due to errors in the scraping process or to an incorrect optical character recognition.

5.3. Error Analysis

Our automatic annotation process removed the workload of manually annotating the encyclopedia with correct QIDs and coordinates. In certain cases, however, the annotator could miss linkages due to various factors. As an example, older spellings of location names that were valid prior to the Swedish spelling reform of 1906 were not always present as alternative name labels in Wikidata. This, in turn, led to an unsuccessful linkage between the encyclopedia entry and the corresponding Wikidata item that in reality should have been linked.

We discuss below three examples of errors, where the system did not correctly identify the QID.

5.3.1. Aachen

The encyclopedia has an *Aachen* entry which reads:

Aachen [ak-]. 1. Regeringsområde i preussiska Rhenprovinsen, 4,155 kvkm. med 614,964 inv. (1900).
2. (Lat. Aquisgranum, fr. Aix-la-Chapelle) Hufvudort i nyssnämnda område, vid den lilla ån Worm I. Wurm, nära gränsen till Holland och Belgien.

translated as:

Aachen [ak-]. 1. Government area in the Prussian Rhine Province, 4,155 sq km with 614,964 inhab. (1900).
2. (Lat. Aquisgranum, fr. Aix-la-Chapelle) Capital in the area just mentioned, by the small river Worm I. Wurm, near the border with Holland and Belgium.

In this case, the difficulty mainly lies in the fact that the encyclopedia entry for *Aachen* has two distinct definitions: the first being the now defunct administrative district of the former German Empire (Q896929) and the second being the city itself (Q1017).

Distinguishing between these two definitions and linking them to the correct Wikidata entry automatically is not a trivial task. To classify these entries correctly, it must be ensured that the annotator is presented with sufficient context from the ambiguous entries in order to successfully choose the correct candidate for the entry. One method may be to employ a segmenter to construct sub-articles and compare the different uses of a word. This is a potential idea for further work.

5.3.2. Arktonnesos

Another entry of interest is *Arktonnesos*, where the encyclopedia entry reads:

Arktonnesos, det grekiska namnet på den i Marmarasjön utskjutande Artakihalfön.

translated as:

Arktonnesos, the Greek name of the Artaki Peninsula which projects into the Sea of Marmara.

The current name of this peninsula is *Kapıdağ Peninsula*, which can be derived only through comparing the entry with external sources. The Wikidata item (Q3780284) for Kapıdağ does not contain an alternative name that reads *Arktonnesos* or any variations on that name. Consequently, this type of cases needs external reference sources to be mapped correctly.

5.3.3. Iowa

The last example concerns the state of *Iowa* entry:

Iowa, en af Nord-Amerikas förenta stater...

translated as:

Iowa, one state of North America's United States...

It showcases the importance of adequate context.

For this entry, the linker chose Q99670857 as QID, which corresponds to a fictional analog of this state. The entity description in Wikidata reads

the federated state of Iowa in the USA as depicted in Star Trek

where as the correct Wikidata entry description (Q1546) reads:

state of the United States of America

This wrong link comes from the candidate ranking system for picking the QID. The cosine similarity between the SBERT embedding of the entry definition and the wikidata description was higher with the *Star Trek* variant than the real-life counterpart.

We believe we could improve the results with a new ranking algorithm, with more constraints on the selection of candidates, or with a better encoding model.

6. Discussion and Conclusion

As initial hypothesis for this work, we posed that the geographic locations mentioned in the encyclopedia would skew heavily towards Swedish places in particular, and European places in general. The results in Figure 3 confirms this was correct. Furthermore, the encyclopedia features a large amount of geographic locations in Germany and the United Kingdom. These are countries where Sweden historically has had strong ties, especially Germany.

While the result may seem self-evident, it is nevertheless one from which further analysis could be built. For example, this study could, with little modification of the method, be extended to include later Swedish encyclopedias. Thus it would allow a comparative analysis on the choice of locations. After the Second World War for example, how would Germany and the United States evolve on the map?

Some conclusions can be drawn in comparison to the geographic location linking and visualization of the French dictionary *Le Petit Larousse Illustré* done in Nugues (2022). The *Petit Larousse* was partially contemporary to *Uggleupplagan*, and under the assumption that the selected locations are indicative of the respective countries' zeitgeist, some comparisons can be made. There is a higher density of locations in the United States, Japan and India, possibly due to the more global nature of France's transnational relations and trade at the time, as compared to Sweden's. The French dictionary also includes many entries in the former French colonies of French North and West Africa.

7. Acknowledgments

This work was partially supported by *Vetenskapsrådet*, the Swedish Research Council, registration number 2021-04533.

8. Bibliographical References

- Lars Aronsson. 2003. *Nordisk familjebok, konversationslexikon och realencyklopedi*.
- Kungliga Biblioteket. 2020. *KB tillgängliggör kraftfulla modeller för språkförståelse*.
- Olle Bridal, Thomas Vakili, and Marina Santini. 2022. *Cross-clinic de-identification of Swedish electronic health records: Nuances and caveats*. In *Proceedings of the Workshop on Ethical and Legal Issues in Human Language Technologies and Multilingual De-Identification of Sensitive Data In Language Resources within the*

- 13th Language Resources and Evaluation Conference, pages 49–52, Marseille, France. European Language Resources Association.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matthew Francis-Landau, Greg Durrett, and Dan Klein. 2016. [Capturing semantic similarity for entity linking with convolutional neural networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1256–1261, San Diego, California. Association for Computational Linguistics.
- Ahmed Hamdi, Elvys Linhares Pontes, Emanuela Boros, Thi Tuyet Hai Nguyen, Günter Hackl, Jose G. Moreno, and Antoine Doucet. 2021. [A multilingual dataset for named entity recognition, entity linking and stance detection in historical newspapers](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 2328–2334, New York, NY, USA. Association for Computing Machinery.
- Johannes Hoffart, Mohamed Amir Yosef, Ilaria Bordino, Hagen Fürstenauf, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 2011. [Robust disambiguation of named entities in text](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 782–792, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Nik Johansson. 2022. Named entity recognition on transaction descriptions. Master's thesis, Lund University.
- Lajanugen Logeswaran, Ming-Wei Chang, Kenton Lee, Kristina Toutanova, Jacob Devlin, and Honglak Lee. 2019. Zero-shot entity linking by reading entity descriptions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Martin Malmsten, Love Börjesson, and Chris Hafenden. 2020. [Playing with words at the National Library of Sweden – Making a Swedish-BERT](#).
- Khalil Mrini, Shaoliang Nie, Jiatao Gu, Sinong Wang, Maziar Sanjabi, and Hamed Firooz. 2022. [Detection, disambiguation, re-ranking: Autoregressive entity linking as a multi-task problem](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1972–1983, Dublin, Ireland. Association for Computational Linguistics.
- Nationalencyklopedin. 2023. [Projekt Runeberg](#).
- Dan Nielsen. 2023. [ScandEval: A benchmark for Scandinavian natural language processing](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 185–201, Tórshavn, Faroe Islands. University of Tartu Library.
- Pierre Nugues. 2022. [Connecting a French dictionary from the beginning of the 20th century to Wikidata](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2548–2555, Marseille, France. European Language Resources Association.
- Anna Nyqvist. 2021. Bootstrapping annotated job ads using named entity recognition and swedish language models. Master's thesis, KTH, School of Electrical Engineering and Computer Science (EECS).
- Korawit Orkphol and Wu Yang. 2019. [Word sense disambiguation using cosine similarity collaborates with word2vec and wordnet](#). *Future Internet*, 11(55):114.
- Elvys Linhares Pontes, Luis Adrián Cabrera-Diego, Jose G. Moreno, Emanuela Boros, Ahmed Hamdi, Nicolas Sidère, Mickaël Coustaty, and Antoine Doucet. 2020. Entity linking for historical documents: Challenges and solutions. *Digital Libraries at Times of Massive Societal Transition*.
- Adithya Pratapa, Rishubh Gupta, and Teruko Mitamura. 2022. [Multilingual event linking to Wikidata](#). In *Proceedings of the Workshop on Multilingual Information Access (MIA)*, pages 37–58, Seattle, USA. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

Sonja Remmer, Anastasios Lamproudis, and Hercules Dalianis. 2021. [Multi-label diagnosis classification of Swedish discharge summaries – ICD-10 code assignment using KB-BERT](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 1158–1166, Held Online. INCOMA Ltd.

Maria Simonsen. 2016. *Den skandinaviske encyklopædi: Udgivelse og udformning af Nordisk familjebok & Salmonsens Konversationsleksikon*. Ph.D. thesis, Institutionen för kulturvetenskaper. Defence details Date: 2016-12-09 Time: 10:15 Place: C121, LUX, Helgonavägen 3, Lund External reviewer(s) Name: Furu-land, Gunnel Title: fil dr Affiliation: Uppsala universitet —.

Karim Tharani. 2021. Much more than a mere technology: A systematic review of wikidata in libraries. *The Journal of Academic Librarianship*, 47(2):102326.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.