

This work was written as part of one of the author's official duties as an Employee of the United States Government and is therefore a work of the United States Government. In accordance with 17 U.S.C. 105, no copyright protection is available for such works under U.S. Law.

Public Domain Mark 1.0

<https://creativecommons.org/publicdomain/mark/1.0/>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Cyber-Security Knowledge Graph Generation by Hierarchical Nonnegative Matrix Factorization

Ryan Barron^{†‡}, Maksim E. Eren^{*‡}, Manish Bhattarai[†], Selma Wanna^{*},
Nicholas Solovyev[†], Kim Rasmussen[†], Boian S. Alexandrov[†], Charles Nicholas[‡] and Cynthia Matuszek[‡]

[†]Theoretical Division, Los Alamos National Laboratory. Los Alamos, USA.

^{*}Advanced Research in Cyber Systems, Los Alamos National Laboratory. Los Alamos, USA.

[‡]Department of Computer Science and Electrical Engineering, University of Maryland, Baltimore County. Maryland, USA.

Abstract—Much of human knowledge in cybersecurity is encapsulated within the ever-growing volume of scientific papers. As this textual data continues to expand, the importance of document organization methods becomes increasingly crucial for extracting actionable insights hidden within large text datasets. Knowledge Graphs (KGs) serve as a means to store factual information in a structured manner, providing explicit, interpretable knowledge that includes domain-specific information from the cybersecurity scientific literature. One of the challenges in constructing a KG from scientific literature is the extraction of ontology from unstructured text. In this paper, we address this topic and introduce a method for building a multi-modal KG by extracting structured ontology from scientific papers. We demonstrate this concept in the cybersecurity domain. One modality of the KG represents observable information from the papers, such as the categories in which they were published or the authors. The second modality uncovers latent (hidden) patterns of text extracted through hierarchical and semantic non-negative matrix factorization (NMF), such as named entities, topics or clusters, and keywords. We illustrate this concept by consolidating more than two million scientific papers uploaded to arXiv into the cyber-domain, using hierarchical and semantic NMF, and by building a cyber-domain-specific KG.

Index Terms—non-negative matrix factorization, cybersecurity, knowledge graph, topic modeling

I. INTRODUCTION

With the ever-growing volumes of text data, organizing large quantities of scientific papers and identifying patterns within the scientific community continues to be a challenging problem. Knowledge Graphs (KGs) are one of the key techniques for storing factual knowledge in a structured manner. However, building KGs from scientific literature remains a challenge since much of the useful context in data, such as entities, relationships between entities, and properties of entities, is hidden in an unstructured manner. At the same time, domains such as cybersecurity are highly diverse, containing knowledge from a multidisciplinary set of fields. This necessitates a more granular or highly domain-specific organization to discover sub-topic areas. In this paper, we introduce a concept for building domain-specific, multi-modal KGs using hierarchical and semantic Non-negative Matrix Factorization (NMF) for topic modeling. In our concept framework, we consolidate generic scientific corpora into a highly-specific domain and

use the extracted latent (hidden) patterns as an ontology in the KG.

Here, for extracting highly-domain-specific clusters, we introduce a hierarchical approach named HSNMFk-SPLIT that can extract fine-grained sub-topics and their semantic sub-structures from large text corpora, using NMF with automatic model determination (NMFk [1]). For semantic NMFk, given a text-document matrix $\mathbf{X} \in \mathbb{R}_+^{F \times N}$ where F is the number of tokens in the vocabulary and N is the number of documents and a word-context (co-occurrence) matrix $\mathbf{S} \in \mathbb{R}_+^{F \times F}$, the number of hidden topics, k , can be automatically estimated through joint factorization. Here we enact joint factorization with the SPLIT method, which allows combination of patterns from separate factorizations. Incorporating the semantic structure of the text with the ability to estimate the number of topics enables a coherent separation of the latent topics and accurate document clustering [2], [3]. We apply the SPLIT method in a hierarchical manner to expand and separate the main topics into sub-topics for a chosen domain, where the extracted sub-semantic structures serve as narrow vocabularies or scientific-jargon seeds for Name Entities Recognition (NER). Furthermore, to enhance the semantic clustering in each topic, we also jointly factorize, using SPLIT, the category-text matrix $\mathbf{C} \in \mathbb{R}_+^{F \times L}$ (L is the number of distinct categories), values of which represent the TF-IDF of tokens per document category. Finally, HSNMFk-SPLIT addresses the computation overhead that arises from factorizing a large TF-IDF matrix \mathbf{X} by separating N documents into an arbitrary number of distinct matrices, factorizing each smaller TF-IDF matrix separately in a distributed manner [1], and finally combining the patterns from each matrix with the SPLIT method.

We demonstrate our hierarchical method by performing topic modeling of titles and abstracts from papers posted on arXiv (<https://arxiv.org/>), which number more than two million. We then select a main topic of Deep Learning, followed by a sub-topic of Security, and then a sub-sub topic of Cybersecurity and Adversarial ML to build a domain specific KG of observable data (titles, authors, etc.) and the latent patterns (NER, keywords, and clusters/topics). This KG can be used to identify emerging trends and facilitate the discovery of relevant research papers for cybersecurity professionals in highly specific domains such as Adversarial ML. Our initial

results when demonstrating this concept show the ability and practicality of our HSNMFk-SPLIT to extract meaningful information from topics and their semantic sub-structures from large datasets.

II. RELEVANT WORK

A. Knowledge Graphs

Recent studies have demonstrated the application of knowledge graphs in identifying emerging trends, uncovering hidden connections, and predicting future research directions within scientific domains [4]–[11]. The integration of semantic analysis with graph-based structures allows for a more nuanced understanding of the thematic evolution in fields such as deep learning and cybersecurity [12], [13]. These advances highlight the potential of knowledge graphs to transform the exploration and exploitation of scientific literature, moving beyond traditional search mechanisms to enable a deeper, contextually rich engagement with content [14], [15]. To achieve these results, care must be given to the specific process for constructing and analyzing knowledge graphs, as studied and explained in [16]–[20].

Similar to our work, [21], scrapes the web for information relevant to the domain specific knowledge. However their domain is in medicine, and ours is in cyber security. Further, they do not use decompositions to construct the graph, but rather vocabularies. Similarly, [22], uses dates, topics, and events as entities to make the knowledge graph but omits the use of decomposition to organize large data. Our decompositions are enhanced by incorporating the semantic structure of the text with the ability to estimate the number of topics, which enables a coherent separation of the latent topics and accurate document clustering. This method is supported by the findings in [2], [3]. Building on [2], [3], we expand our approach through hierarchical decomposition, as well as incorporating semi-supervised labels of arXiv categories in joint factorization to refine our capabilities.

B. Hierarchical Topic Modeling

A common method to hierarchical topic modeling is through a probabilistic approach, in which the corpus is treated as a mixture of topics that can be modeled with a probability distribution. One such method is hierarchical Latent Dirichlet Allocation (hLDA) using the nested Chinese Restaurant Process (nCRP) beforehand [23]. This model allows for arbitrarily large branching factors in topic hierarchies with the nCRP aiding in inferring the appropriate depth of the topic tree. Another LDA-based approach is Hierarchical Dirichlet Processes (HDP). In this method, a set of Dirichlet processes are coupled via their base measure, which is itself distributed according to a Dirichlet process [24]. The authors introduce two Markov Chain Monte Carlo (MCMC) sampling schemes for posterior inference under hierarchical Dirichlet process mixtures and test the approaches on multiple datasets.

More recent approaches to hierarchical topic modeling make use of variational auto-encoders, recurrent neural networks, and transformers [25]–[29]. While successful, these methods

face a problem common to deep learning: successful models depend heavily on the quantity and quality of the training data. HSNMFk-SPLIT is a favorable approach to hierarchical topic modeling since it is instance-based. Only the data that is being modeled is required for computation.

III. DATASET AND PRE-PROCESSING

Using HSNMFk-SPLIT, we perform topic modeling on the titles and abstracts of literature uploaded to arXiv¹ [30]. Our pre-processing pipeline includes removal of common English stop-words, stop-words that are often present in scientific literature (such as ‘doi’, ‘preprint’, ‘copyright’, ‘figure’, and ‘demonstrate’), stop-phrases (such as ‘All rights reserved.’), non-ASCII characters, symbols, next-line characters, tags (such as HTML and LaTeX), and e-mail addresses. We also join hyphenated words into single tokens, make the tokens lower-case, and lemmatize the tokens using the Python package *NLTK* [31] (NLTK is used with the allowed postags ‘NOUN’, ‘ADJ’, ‘VERB’, ‘ADV’, and ‘PROPN’). The non-English abstracts are excluded from analysis using the Python implementation of *language-detection* software [32]. Finally, after removing those abstracts with less than 10 tokens, we are left with 2,178,187 documents in our corpus represented with a TF-IDF matrix $\mathbf{X} \in \mathbb{R}_+^{10,280 \times 2,178,187}$. To reduce noise, we remove the tokens present in more than 80% and less than 500 documents when building \mathbf{X} .

The semantic structure of the documents is represented with matrix $\mathbf{S} \in \mathbb{R}_+^{10,280 \times 10,280}$ where the values represent the number of times two words co-occur in a predetermined window length of $w = 100$ tokens. We normalize \mathbf{S} with Shifted Positive Point-wise Mutual Information (SPPMI) [33], with shift $s = 4$. The categorical semantic structure is represented with the matrix $\mathbf{C} \in \mathbb{R}_+^{10,280 \times 172}$ values of which represents the TF-IDF of tokens per document category. Here the category of the document is assigned by the author in arXiv². While a given document can be assigned to more than one category, we use the primary category assignment.

IV. METHODS

In this section, we first describe the HSNMFk-SPLIT method, then explain how we use the clusters found by HSNMFk-SPLIT to build our KG.

A. HSNMFk-SPLIT

We use a publicly available implementation of NMFk for its automatic model determination capability, which avoids over/under-fitting to topics, and apply it in a hierarchical manner [1]³. HSNMFk-SPLIT is a method that automatically estimates the number of latent topics and extracts coherent topics by exploiting the semantic representation encoded in the word-context matrix, categorical information encoded in the word-category matrix, together with the text-document matrix. Given a TF-IDF text-document matrix $\mathbf{X} \in \mathbb{R}_+^{F \times N}$, SPPMI

¹arXiv Dataset: <https://www.kaggle.com/datasets/Cornell-University/arxiv>

²arXiv Category Taxonomy: https://arxiv.org/category_taxonomy

³NMFk is available in <https://github.com/lanl/T-ELF>.

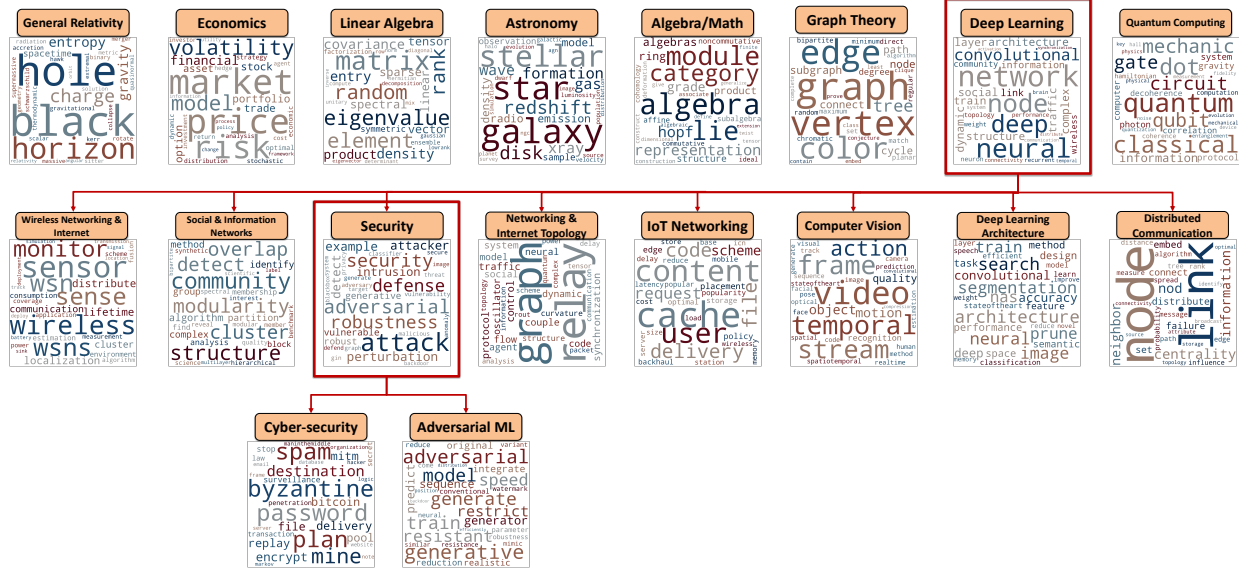


Fig. 1: Word cloud of selected topics and their interpreted labels. Method is hierarchically applied three times, selecting a topic to extract its sub-topics at each level. Selected topics from each level are shown in each row of the figure.

normalized word-context matrix $\mathbf{S} \in \mathbb{R}_+^{F \times F}$, and word-category matrix $\mathbf{C} \in \mathbb{R}_+^{F \times L}$ at each node of the hierarchical tree, HSNMfK extracts the topics - the columns of matrix \mathbf{W} and coordinates of the documents, i.e., columns of matrix \mathbf{H} corresponding to the \mathbf{X} for that node in the tree. HSNMfK is performed by solving a joint optimization problem at each node in the tree:

$$\begin{aligned} & \underset{\mathbf{W} \in \mathbb{R}_+^{F \times k}, \mathbf{H} \in \mathbb{R}_+^{k \times N}, \mathbf{G} \in \mathbb{R}_+^{k \times F}, \mathbf{J} \in \mathbb{R}_+^{k \times L}}{\text{minimize}} \\ & \frac{1}{2} \|\mathbf{X} - \mathbf{WH}\|_F^2 + \alpha \|\mathbf{S} - \mathbf{WG}\|_F^2 + \beta \|\mathbf{C} - \mathbf{WJ}\|_F^2 \end{aligned} \quad (1)$$

where $\|\cdot\|_F^2$ is the Frobenius norm, and α and β are regularization parameters controlling the weight of \mathbf{S} and \mathbf{C} in the decomposition. F is the number of words in the vocabulary, N is the number of documents, and L is the number of categories. One way of evaluating the above expression is by concatenating the TF-IDF matrix \mathbf{X} with \mathbf{S} , and \mathbf{C} , and finally applying NMFk [1] on the concatenation [34]. However, this is computationally challenging as the matrices are significantly larger after concatenation. Hence, we incorporate the SPLIT technique to minimize the computation and be able to perform the decomposition of each matrix in parallel (HSNMfK-SPLIT):

- **Factorizing large matrices via SPLIT:** First, split $\mathbf{X} \in \mathbb{R}_+^{F \times N}$ into m chunks so that the i^{th} chunk given by $\mathbf{X}_i \in \mathbb{R}_+^{F \times \frac{N}{m}}$ can be factorized with NMFk as $\mathbf{X}_i \approx \mathbf{W}_i \mathbf{H}_i$ where $\mathbf{W}_i \in \mathbb{R}_+^{F \times k_i}$ and $\mathbf{H}_i \in \mathbb{R}_+^{k_i \times \frac{N}{m}}$. Here m chunks can be factorized in parallel. Next, concatenate m \mathbf{W}_i matrices obtained from previous step to construct $\tilde{\mathbf{W}} \in \mathbb{R}_+^{F \times K}$ where $\tilde{\mathbf{W}} = [\mathbf{W}_1 | \mathbf{W}_2 | \dots | \mathbf{W}_m]$ and $K = k_1 + k_2 + \dots + k_m$. Now factorize $\tilde{\mathbf{W}}$ with NMFk such that

$\tilde{\mathbf{W}} \approx \mathbf{W}_x \mathbf{M}_x$ where $\mathbf{W}_x \in \mathbb{R}_+^{F \times p}$, $\mathbf{M}_x \in \mathbb{R}_+^{p \times K}$ and $\mathbf{M}_x = [\mathbf{M}_1 | \mathbf{M}_2 | \dots | \mathbf{M}_m]$ where $\mathbf{M}_j \in \mathbb{R}_+^{p \times k_j}$. Finally, from previous steps, multiply \mathbf{M}_i and \mathbf{H}_i to obtain $\mathbf{H}_i^{(*)}$ where $\mathbf{H}_i^{(*)} = \mathbf{M}_i \times \mathbf{H}_i$ and $\mathbf{H}_i^{(*)} \in \mathbb{R}_+^{p \times \frac{N}{m}}$

- **Incorporate semantic and category structure:** Given word SPPMI matrix \mathbf{S} and category matrix \mathbf{C} , utilize NMFk to decompose $\mathbf{S} \approx \mathbf{W}_s \mathbf{H}_s$ and $\mathbf{C} \approx \mathbf{W}_c \mathbf{H}_c$ where $\mathbf{W}_s \in \mathbb{R}_+^{F \times s}$, $\mathbf{H}_s \in \mathbb{R}_+^{s \times F}$, $\mathbf{W}_c \in \mathbb{R}_+^{F \times c}$ and $\mathbf{H}_c \in \mathbb{R}_+^{c \times L}$. Now concatenate \mathbf{W}_x with \mathbf{W}_s and \mathbf{W}_c to obtain \mathbf{W}_+ where $\mathbf{W}_+ = [\mathbf{W}_x | \mathbf{W}_s | \mathbf{W}_c]$ and $\mathbf{W}_+ \in \mathbb{R}_+^{F \times Z}$ where $Z = p + s + c$. Then factorize \mathbf{W}_+ with NMFk as $\mathbf{W}_+ = \mathbf{WY}$ where $\mathbf{W} \in \mathbb{R}_+^{F \times t}$, $\mathbf{Y} \in \mathbb{R}_+^{t \times Z}$, $\mathbf{Y} = [\mathbf{Y}_x | \mathbf{Y}_s | \mathbf{Y}_c]$ and $\mathbf{Y}_x \in \mathbb{R}_+^{t \times p}$, $\mathbf{Y}_s \in \mathbb{R}_+^{t \times s}$ and $\mathbf{Y}_c \in \mathbb{R}_+^{t \times c}$. Finally, multiply \mathbf{Y}_x with $\mathbf{H}^{(*)}$ to obtain \mathbf{H} such that $\mathbf{H} = \mathbf{Y}_x \mathbf{H}^{(*)}$ where $\mathbf{H} \in \mathbb{R}_+^{t \times N}$ and $\mathbf{W} \in \mathbb{R}_+^{F \times t}$
- **Hierarchical topic modeling:** For each node in the tree, once \mathbf{W} and \mathbf{H} are obtained, \mathbf{W} models the topics, and \mathbf{H} models the document's topic assignment. Here we compute the maximum indices for each column of \mathbf{H} (also known as H-clustering) [35] to estimate the best topic representation of each document given as $\text{topic}(\mathbf{X}_j) = \arg\max \mathbf{H}_j$. The documents in each topic (super-topic) can then be factorized again following the above steps to estimate the underlying sub-topics along the tree depth.

B. Building the Knowledge Graph

The knowledge graph, a Neo4j graph [37], is composed of nodes and relations from several forms of data. First and foremost, the documents are inserted into the graph, where they are assigned observable attributes in the original data. These observable data include titles, authors, DOI, year, publication information, citations, references, affiliations, etc. Next, the cleaned text decompositions are used to extract




Fig. 2: Data pipeline from starting from arXiv data (1), text cleaning (2), running HSNMFk-SPLIT for topics (3), extracting the keywords for the decomposition (4), extracting the Named Entities per document (5), and then structurally aggregating the data into the knowledge graph (6). Images generated with DALL-E [36].

HSNMFk-SPLIT topics and keywords from a highly domain-specific leaf decomposition about security. The words and topics are injected into the graph, connected to their documents from the *argmax* of the topic probability for each document. The keywords are connected to the topics where the edges contain the probabilities as weights. Next, the text of each document is examined with spaCy’s NER pipeline [38], which produces 18 labels. We insert the 6 most useful labels of the 18 total for the purpose of cyber-security ontological exploration: Organization, Event, Person, Location, Product, and Geopolitical Entity. Each of the NERs is connected to each source document that produces it; some of documents may produce the same NER entities. Documents are further connected into communities where NERs are shared.

V. RESULTS

Our results are shown in Figure 1 with the word clouds (most prominent words corresponding to each column of \mathbf{W}), and with the distribution of categories in each topic and sub-

topic in Figure 3 (the top ten categories are shown for the given topic). While word clouds highlight the interpretability and quality of the extracted topics, the distribution of the categories shows the specificity of the documents as we look at the sub-categories. We apply our method hierarchically three times, at each level selecting one topic to expand further to identify its sub-topics. Since the initial \mathbf{X} is large, we chunk the documents into 20 separate matrices and apply semantic joint factorization with automatic determination of the number of latent topics. Individually, the 20 matrices revealed between 30 and 80 latent topics. After joint factorization, we have identified 24 total super-topics in 2 million+ documents.

We selected the topic related to Deep Learning. The first row of Figure 1 shows the word cloud for this topic, and the first row of Figure 3 displays the distribution of categories that are in this topic. The categories are mainly represented in diverse sub-fields of Computer Science, ranging from Cyber (*cs.CR*) to Computer Vision (*cs.CV*). At the second level, the sub-topic of Deep Learning produced 29 total sub-topics, which




Fig. 3: Distribution of the top ten document categories, based on the arXiv author assignments. Each topic corresponds to the word clouds from Figure 1. Selected topics from each level are shown in each row of the figure.

include Computer Vision and Security. The sub-topic that can be interpreted as Computer Vision based on the tokens in the word cloud is now narrowed down to the *cs.CV*. On the other hand, the Security sub-topic is mainly distributed between the categories *cs.CR* (Cyber), *cs.CV* (Computer Vision), and *cs.LG* (Language); therefore, we expand the Security sub-topic to the third level, further separating it into three sub-sub-topics, two of which can be interpreted as Cyber-security and Adversarial ML based on the word clouds. Similarly at the third level, the Cyber-security topic category distribution is narrowed down to *cs.CR*. The Adversarial ML topic contains documents from three main categories of *cs.CR*, *cs.CV*, and *cs.LG*, which is reasonable based on this field and how research focuses on adversarial computer vision and language models.

In the domain specific leaf decomposition, 1383 vocabulary words were used to discover three latent topics, which are the counts of the entities in the graph. These three topics

contain their own unique, non-overlapping set of documents. Specifically, topic 0 had 597 documents, topic 1 had 114 documents, and topic 2 had 565 documents. These topics were composed from the 1383 vocabulary words, but the probability of the words occurring in each topic changes according to the documents contained. The knowledge graph contained NER entities such as 245 organizations, 7 events, 64 persons, 3 locations, 679 products, and 22 geopolitical entities. Overall the graph produced 3758 node entities and 9428 edge relationships. Queries could be asked of the knowledge graph, which have to first be translated from natural language into Cypher language. For example, 'Which papers mention MNIST?' returns all of documents that have mnist as a NER, a count of 50 documents. Similarly, 'Which documents in the security decomposition are in the 'cs.SE' (software engineering) arXiv category' finds 7 documents, which are further linked to 2 organizations, an event, topics 0 and 2, a product, and 7 additional categories.

VI. CONCLUSION

In this paper, we have presented a concept for building a domain specific multi-model knowledge graph (KG) from unstructured scientific text data. Our KG includes observable entities from data such as paper titles, authors, publication year, as well as latent (hidden) patterns such as named entities, keywords, and topics/clusters that are extracted with a new semantic hierarchical NMF method, named HSNMFk-SPLIT, for document organization. HSNMFk-SPLIT is designed for large corpora using distributed joint factorization, and has the capability to perform automatic selection of the number of latent topics. The semantic structure and automatic model selection allows extraction of the coherent topics. We have demonstrated the feasibility of HSNMFk-SPLIT in performing topic and sub-topic extraction on the abstracts of more than two million papers posted on arXiv, and consolidating these generic scientific papers to cyber-specific data to build a highly-domain-specific KG.

ACKNOWLEDGMENT

This research was funded by the Los Alamos National Laboratory (LANL) Laboratory Directed Research and Development (LDRD) grant 20230067DR and LANL Institutional Computing Program, supported by the U.S. Department of Energy National Nuclear Security Administration under Contract No. 89233218CNA000001.

REFERENCES

- [1] M. Eren, N. Solovyev, R. Barron, M. Bhattarai, D. Truong, I. Boureima, E. Skau, K. Rasmussen, and B. Alexandrov, "Tensor Extraction of Latent Features (T-ELF)," Los Alamos National Laboratories, Tech. Rep., Oct. 2023. [Online]. Available: <https://github.com/lanl/T-ELF>
- [2] R. Vangara, E. Skau, G. Chennupati, H. Djidjev, T. Tierney, J. P. Smith, M. Bhattarai, V. G. Stanev, and B. S. Alexandrov, "Semantic nonnegative matrix factorization with automatic model determination for topic modeling," in *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2020, pp. 328–335.

- [3] M. E. Eren, N. Solovyev, M. Bhattarai, K. Rasmussen, C. Nicholas, and B. Alexandrov, "Sennmf-split: Large corpora topic modeling by semantic non-negative matrix factorization with automatic model selection," in *Proceedings of the ACM Symposium on Document Engineering 2022*, ser. DocEng '22. New York, NY, USA: Association for Computing Machinery, 2022.
- [4] P. Manghi, A. Mannocci, F. Osborne, D. Sacharidis, A. Salatino, and T. Vergoulis, "New trends in scientific knowledge graphs and research impact assessment," pp. 1296–1300, 2021.
- [5] A. Pingle, A. Piplai, S. Mittal, A. Joshi, J. Holt, and R. Zak, "Relext: Relation extraction using deep learning approaches for cybersecurity knowledge graph improvement," in *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2019, pp. 879–886.
- [6] 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*, E. Riloff, D. Chiang, J. Hockenmaier, and J. Tsujii, Eds. Brussels, Belgium: Association for Computational Linguistics, Oct.-Nov. 2018, pp. 3219–3232. [Online]. Available: <https://aclanthology.org/D18-1360>
- [7] C. Wang, X. Ma, J. Chen, and J. Chen, "Information extraction and knowledge graph construction from geoscience literature," *Computers & Geosciences*, vol. 112, pp. 112–120, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0098300417309020>
- [8] S. Wang, A. Ororbia, Z. Wu, K. Williams, C. Liang, B. Pursel, and C. L. Giles, "Using prerequisites to extract concept maps from textbooks," in *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*, ser. CIKM '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 317–326. [Online]. Available: <https://doi.org/10.1145/2983323.2983725>
- [9] B. Shao, X. Li, and G. Bian, "A survey of research hotspots and frontier trends of recommendation systems from the perspective of knowledge graph," *Expert Systems with Applications*, vol. 165, p. 113764, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0957417420305881>
- [10] B. Abu-Salih, "Domain-specific knowledge graphs: A survey," *Journal of Network and Computer Applications*, vol. 185, p. 103076, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1084804521000990>
- [11] P. Chen, Y. Lu, V. W. Zheng, X. Chen, and B. Yang, "Knowedu: A system to construct knowledge graph for education," *IEEE Access*, vol. 6, pp. 31 553–31 563, 2018.
- [12] A. L. Opdahl, T. Al-Moslimi, D.-T. Dang-Nguyen, M. Gallofré Ocaña, B. Tessem, and C. Veres, "Semantic knowledge graphs for the news: A review," *ACM Computing Surveys*, vol. 55, no. 7, pp. 1–38, 2022.
- [13] L. F. Sikos, "Cybersecurity knowledge graphs," *Knowledge and Information Systems*, vol. 65, no. 9, pp. 3511–3531, 2023.
- [14] S. Auer, A. Oelen, M. Haris, M. Stocker, J. D'Souza, K. E. Farfar, L. Vogt, M. Prinz, V. Wiens, and M. Y. Jaradeh, "Improving access to scientific literature with knowledge graphs," *Bibliothek Forschung und Praxis*, vol. 44, no. 3, pp. 516–529, 2020.
- [15] Y. Zhu, E. Yan, and I.-Y. Song, "The use of a graph-based system to improve bibliographic information retrieval: System design, implementation, and evaluation," *Journal of the Association for Information Science and Technology*, vol. 68, no. 2, pp. 480–490, 2017.
- [16] M. Iannacone, S. Bohn, G. Nakamura, J. Gerth, K. Huffer, R. Bridges, E. Ferragut, and J. Goodall, "Developing an ontology for cyber security knowledge graphs," in *Proceedings of the 10th Annual Cyber and Information Security Research Conference*, ser. CISR '15. New York, NY, USA: Association for Computing Machinery, 2015. [Online]. Available: <https://doi.org/10.1145/2746266.2746278>
- [17] H. Chen and X. Luo, "An automatic literature knowledge graph and reasoning network modeling framework based on ontology and natural language processing," *Advanced Engineering Informatics*, vol. 42, p. 100959, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1474034619302642>
- [18] Y.-J. Yang, B. Xu, J.-W. Hu, M.-H. Tong, P. Zhang, and L. Zheng, "Accurate and efficient method for constructing domain knowledge graph," *Ruan Jian Xue Bao/Journal of Software*, vol. 29, pp. 2931–2947, 10 2018.
- [19] R. Wang, Y. Yan, J. Wang, Y. Jia, Y. Zhang, W. Zhang, and X. Wang, "Acekg: A large-scale knowledge graph for academic data mining," in *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, ser. CIKM '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 1487–1490. [Online]. Available: <https://doi.org/10.1145/3269206.3269252>
- [20] S. M. Tiwari, F. N. Al-Aswadi, and D. Gaurav, "Recent trends in knowledge graphs: theory and practice," *Soft Computing*, vol. 25, pp. 8337 – 8355, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:234813288>
- [21] J. Costello, M. Reformat, and F. Bolduc, "Leveraging knowledge graphs and natural language processing for automated web resource labeling: Knowledge mobilization in neurodevelopmental disorders. (preprint)," *Journal of Medical Internet Research*, vol. 25, 12 2022.
- [22] D. Flocco, B. Palmer-Toy, R. Wang, H. Zhu, R. Sonthalia, J. Lin, A. L. Bertozzi, and P. J. Brantingham, "An analysis of covid-19 knowledge graph construction and applications," in *2021 IEEE International Conference on Big Data (Big Data)*. Los Alamitos, CA, USA: IEEE Computer Society, dec 2021, pp. 2631–2640. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/BigData52589.2021.9671479>
- [23] T. Griffiths, M. Jordan, J. Tenenbaum, and D. Blei, "Hierarchical topic models and the nested chinese restaurant process," in *Advances in Neural Information Processing Systems*, S. Thrun, L. Saul, and B. Schölkopf, Eds., vol. 16. MIT Press, 2003.
- [24] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei, "Hierarchical dirichlet processes," *Journal of the American Statistical Association*, vol. 101, pp. 1566 – 1581, 2006.
- [25] Y. Jin, H. Zhao, M. Liu, L. Du, and W. Buntine, "Neural attention-aware hierarchical topic model," 2021.
- [26] D. Guo, B. Chen, R. Lu, and M. Zhou, "Recurrent hierarchical topic-guided neural language models," 2020. [Online]. Available: <https://openreview.net/forum?id=Byl1W1rtvH>
- [27] F. Viegas, W. Cunha, C. Gomes, A. Pereira, L. C. da Rocha, and M. A. Gonçalves, "Cluhtm - semantic hierarchical topic modeling based on cluwords," in *Annual Meeting of the Association for Computational Linguistics*, 2020.
- [28] S. Mifrah and E. H. Benlahmar, "Topic modeling with transformers for sentence-level using coronavirus corpus," *International Journal of Interactive Mobile Technologies (iJIM)*, vol. 16, no. 17, p. pp. 50–59, Sep. 2022. [Online]. Available: <https://online-journals.org/index.php/i-jim/article/view/33281>
- [29] M. Grootendorst, "BERTopic: Neural topic modeling with a class-based tf-idf procedure," 2022.
- [30] C. B. Clement, M. Bierbaum, K. P. O'Keefe, and A. A. Alemi, "On the use of arxiv as a dataset," 2019.
- [31] S. Bird, E. Klein, and E. Loper, *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc.", 2009.
- [32] N. Shuyo, "Language detection library for java," 2010. [Online]. Available: <http://code.google.com/p/language-detection/>
- [33] O. Levy and Y. Goldberg, "Neural word embedding as implicit matrix factorization," *Advances in neural information processing systems*, vol. 27, 2014.
- [34] V. Stanev, E. Skau, I. Takeuchi, and B. S. Alexandrov, "Topic analysis of superconductivity literature by semantic non-negative matrix factorization," in *International Conference on Large-Scale Scientific Computing*. Springer, 2021, pp. 359–366.
- [35] R. Vangara, M. Bhattarai, E. Skau, G. Chennupati, H. Djidjev, T. Tierney, J. P. Smith, V. G. Stanev, and B. S. Alexandrov, "Finding the number of latent topics with semantic non-negative matrix factorization," *IEEE Access*, vol. 9, pp. 117 217–117 231, 2021.
- [36] OpenAI's DALL-E, "Visual representations of kg schema," 2024.
- [37] Neo4j, Inc., "Neo4j: The #1 platform for connected data," <https://neo4j.com/>, 2023.
- [38] Explosion AI, "spacy english core web transformer model," https://spacy.io/models/en#en_core_web_trf, 2023.