



# NASA Science Mission Directorate Knowledge Graph Discovery

Roelien C. Timmer  
University of New South Wales  
Australia

Fech Scen Khoo  
University of Oldenburg  
Germany

Megan Mark  
Florida Institute of Technology  
United States

Marcella Scoczynski Ribeiro  
Martins  
Federal University of Technology  
Brazil

Anamaria Berea  
George Mason University  
United States

Gregory Renard  
The Applied AI Company (AAICO)  
United States

Kaylin Bugbee  
NASA Marshall Space Flight Center  
United States

## ABSTRACT

The size of the National Aeronautics and Space Administration (NASA) Science Mission Directorate (SMD) data catalog is growing exponentially, allowing researchers to make discoveries. However, making discoveries is challenging and time-consuming due to the size of the data catalogs, and as many concepts and data are indirectly connected. This paper proposes a pipeline to generate knowledge graphs (KGs) representing different NASA SMD domains. These KGs can be used as the basis for dataset search engines, saving researchers time and supporting them in finding new connections. We collected textual data and used several modern natural language processing (NLP) methods to create the nodes and the edges of the KGs. We explore the cross-domain connections, discuss our challenges, and provide future directions to inspire researchers working on similar challenges.

## CCS CONCEPTS

• Information systems → Content analysis and feature selection; Language models.

### ACM Reference Format:

Roelien C. Timmer, Fech Scen Khoo, Megan Mark, Marcella Scoczynski Ribeiro Martins, Anamaria Berea, Gregory Renard, and Kaylin Bugbee. 2023. NASA Science Mission Directorate Knowledge Graph Discovery. In *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*, April 30–May 04, 2023, Austin, TX, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3543873.3587585>

## KEYWORDS

Knowledge Graphs (KGs), National Aeronautics and Space Administration (NASA) Science Mission Directorate (SMD), Named-entity recognition (NER), Knowledge Representation, Natural Language Processing (NLP)

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

WWW '23 Companion, April 30–May 04, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-9419-2/23/04...\$15.00 <https://doi.org/10.1145/3543873.3587585>

## 1 INTRODUCTION

As the interest in space exploration has increased, along with digitalization, the number of datasets in NASA's data catalog is growing exponentially. By 2023, the National Aeronautics and Space Administration (NASA) Science Mission Directorate (SMD) foresees the generation of at least 400 petabytes per year and expects this number to increase as new models are run, and new missions are launched [10]. NASA has expressed the need for a search engine that enables scientists to search across the five NASA SMD scientific pursuits and allows them to discover new connections [5, 11].

Search engines are often built with KGs as their backbone. KGs are a popular method to structure big and complex information. KGs organize information in a graph structure. There is no fixed definition of a KG, but in general, the nodes represent entities, and the edges represent the type of relationships [12]. KGs make it easier to digest complex information and make it easier to find connections between different concepts.

There are initiatives specifically for creating KGs for specific NASA SMD domains, such as the Heliophysics KNOWledge Network (Helio-KNOW) project [23]. Helio-KNOW is a community-build collection of software and systems to organize Heliophysics information<sup>1</sup>. However, there are no KGs for all five of the NASA SMD domains. In addition, research is needed to connect the KGs of the five different NASA SMD scientific pursuits: (i) Heliophysics, (ii) Astrophysics, (iii) Planetary Science, (iv) Earth Science, and (v) Biological & Physical Science. There has been an increased interest in multidisciplinary research between these different domains. For example, heliophysicists and astrophysicists can benefit from studying the high-energy solar flaring activity together to understand stellar flares [7]. Before merging the KGs graphs with, for example, entity alignment [29, 31], there is a need for a better understanding of the domain overlap.

In this paper, we outline how we approached building KGs for the NASA SMD, of which the main steps are: (i) the collection of sufficient relevant textual data to train the NLP models, (ii) the selection of NLP algorithms to create the nodes and the edges of the KGs, (iii) exploring the overlap of the domains, (iv) validation, and, (v) visualization. While a significant contribution of this paper

<sup>1</sup>Link to Helio-KNOW project Github: <https://github.com/rmcgranaghan/Helio-KNOW>

is providing a pipeline for generating KGs for the NASA SMD domains, our discussion of the challenges and recommendations for future work is valuable to other researchers. We also show our initial results of analyzing the overlap of the NASA SMD scientific pursuits.

This paper is structured as follows. We first outline related literature and work in Section 2. In Section 3, we propose a pipeline to generate KGs for the different NASA SMD domains. We present initial results in Section 4, including an overview of the data gathered and a cross-domain analysis. Finally, we have an extensive discussion of the challenges we faced and suggestions for future work in Section 5 followed by a concise conclusion in Section 6.

## 2 RELATED WORK

In this section, we describe the main relevant existing research related to this work. We introduce the main concepts of KGs (Section 2.1) and discuss methods used for scientific information processing (Section 2.2).

### 2.1 Knowledge Graphs

The goal of KGs is to organize information in a graph structure. KGs have different definitions and forms, but in general, the nodes of a KG represent entities, i.e., concepts, objects, situations, or events, and the edges represent the relationship between the entities, which can both be qualitative and quantitative [12]. KGs gained popularity when big tech companies, such as Google, Microsoft, Facebook, and Yahoo!, created their own KGs and used them as the core of their semantic search engines. The three most well-known KGs are DBPedia [2], Wikidata [30], and Google KG [26].

### 2.2 Scientific Information Processing

In science, there are large volumes of complex information. KGs are popular [3, 20] to structure all this information. One approach to generating scientific KGs is extracting entities from scientific documents [14, 16, 33]. There are multiple named-entity recognition (NER) algorithms to extract scientific entities, such as AstroBERT [13], and SciBERT [4]. These scientific entities can serve as the nodes of the KGs. AstroBERT and SciBERT are based on Google's Bidirectional Encoder Representations from Transformers (BERT) [9].

## 3 METHODOLOGY

We propose a pipeline to generate KGs for the five NASA SMD domains. The pipeline consists of five major components, pre-processing (Section 3.1), entity extraction (Section 3.2), entity linkage (Section 3.3), validation (Section 3.4), and visualization (Section 3.5), which are visualized in Figure 1. We describe all five components below.

### 3.1 Pre-processing

We gather textual data from different sources. We need a high amount of textual data as this will be the input of the NLP models.

The NASA SMD provided a list with definitions from the five different Scientific Pursuits<sup>2</sup>. These definitions are a valuable source because they are created by NASA and used internally by researchers. We also downloaded terms from the Unified Astronomy Thesaurus (UAT), which is a community-supported, open-source project from the American Astronomical Society and builds upon the International Astronomical Union (IAU) Thesaurus<sup>3</sup>[1]. The UAT consists of astronomical concepts and their inter-relationships. The strength of the UAT is that anyone in the community can contribute by providing additions, refinements, and revisions; therefore, the UAT's quality and quantity are constantly being improved. To have sufficient data to train NLP models, we want to gather more textual data. We opt to web scrape Wikipedia summaries and arXiv abstracts based on the UAT terms. We web scrape Wikipedia because of the broad coverage of relevant concepts and arXiv abstracts because they contain a lot of relevant technical terms. In Section 4.1, we provide more details about the web scrape method we applied and how much textual data we gathered from all four sources.

The textual data is stored in different formats, such as Portable Document Format (PDF), Comma-Separated Values (CSV), Computable Document Format (CDF), and HyperText Markup Language (HTML), and therefore needs to be normalized to ensure that all the texts have the same format. We tokenize and lemmatize all the texts. With tokenization, the texts are divided into tokens, and with lemmatization, all the words of the texts are converted to their roots. We also remove stopwords and identify common bigrams and trigrams.

### 3.2 Extraction

To generate the nodes of the KGs, we extract the keywords from the texts. We use Yet Another Keyword Extractor, better known as Yake!<sup>4</sup> [6]. Yake! is an unsupervised automatic keyword algorithm that returns a list of keywords and a corresponding quantitative score indicating the relevance of each keyword<sup>5</sup>. Examples of keywords extracted with Yake! from the UAT are *Tauri stars*, *Kerr black holes*, *Classical novae*, and *Hubble Space Telescope*<sup>6</sup>. In Section 5.4, we discuss other methods that we tried and tested to extract entities to generate the nodes of the KGs.

### 3.3 Linkage

To link the nodes in the KGs, we considered two different techniques. First, count the frequency of the keywords co-occurring in a sentence as subjects and objects. The higher the frequency, the stronger the edge between the two nodes. Second, we quantify the strength of the edges by calculating the semantic similarity between NEs with the help of large pre-trained NLP models. For example, we finetune a model of the BERT model based on all the textual

<sup>2</sup>Example of the definition of Coronal Mass Ejection: *An eruption in the outer solar atmosphere that sends billions of tons of magnetized plasma clouds into interplanetary space. When traveling at high speeds these ejections create shocks in the solar wind. Earth-intercept of a CME is often followed by a geomagnetic storm.*

<sup>3</sup>Link to the UAT glossary: <https://astrothesaurus.org>

Link to the UAT glossary in CSV format: <https://github.com/astrothesaurus/UAT/blob/master/UAT.csv>

<sup>4</sup>Link to Yake!: <https://github.com/LIAAD/yake>

<sup>5</sup>The Yake! keyword is more relevant with a lower score.

<sup>6</sup>More examples of keywords extracted with Yake! can be found in [22].

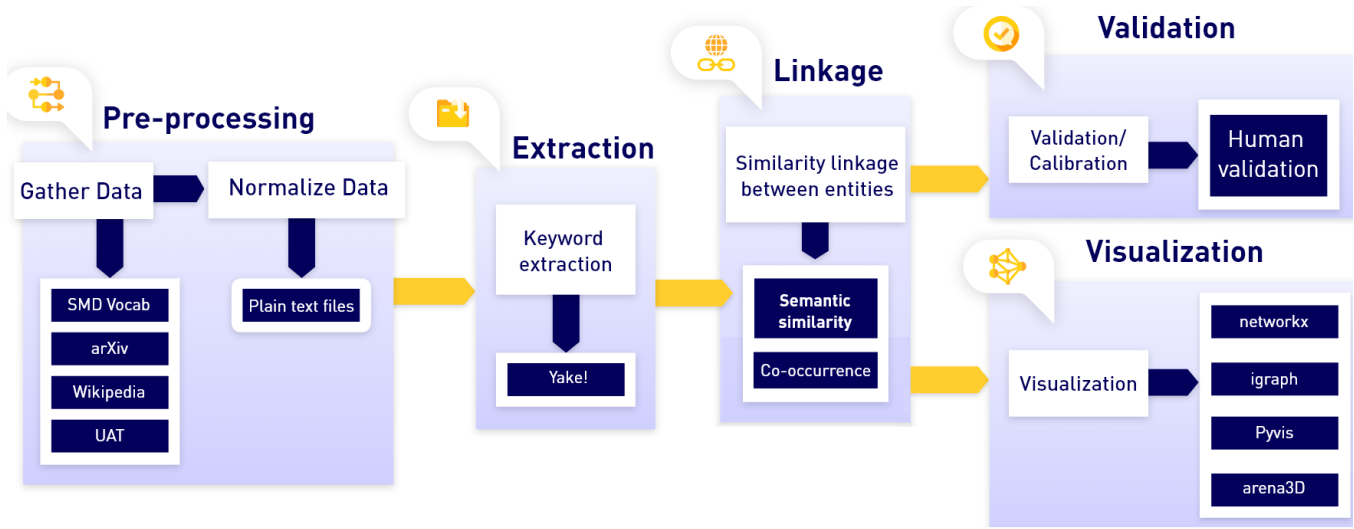


Figure 1: Proposed Pipeline for the Generation of Knowledge Graphs for the NASA SMD.

data we collected and calculate the semantic similarity between the nodes [9].

### 3.4 Validation

Validating the KGs is a big challenge as the nodes of the KGs mainly contain technical terms that only experts are fully familiar with. Examples of the nodes are: *Fra Mauro*, *Galilean Moons*, *Chirality*, *YORP Effect*, *Nectarian* and *Caldera*. Therefore, to validate the quality of the KGs, we validated the quality ourselves and invited domain experts to provide feedback. In Section 5.1, we discuss alternative validation techniques.

### 3.5 Visualization

To gain a deeper understanding of the KGs, we create visualizations with NetworkX<sup>7</sup> [15], iGraph<sup>8</sup> [8], Pyvis<sup>9</sup> [25], and arena3D<sup>10</sup> [24]. All these four software packages have their own benefits and shortcomings. These visualizations assist us in validating the quality of the KGs but can also assist researchers make discoveries.

## 4 INITIAL RESULTS

In this section, we would like to discuss some initial results. This includes describing the type of data we collected (Section 4.1) and the results of exploring the overlap between some of NASA SMD’s scientific pursuits (Section 4.2).

### 4.1 Data Collection

We collected data from four different sources: (i) SMD, (ii) arXiv, (iii) Wikipedia, and (iv) UAT. We used a definition list provided

Table 1: A Summary of all the Textual Data Collected.

Source	Size	Type
Science Mission Directorate	3,096,448	definitions
arXiv	168,084	abstracts
Wikipedia	695	summaries
Unified Astronomy Thesaurus	2,826	terms

by the NASA SMD consisting of more than three million definitions and the UAT consisting of 2,826 terms. We web scraped relevant Wikipedia summaries and arXiv abstracts based on these UAT terms.

We used a collection of all the arXiv abstracts up to the end of 2021, a total of approximately two million abstracts<sup>11</sup>. We filtered for abstracts that are related to the UAT terms. Specifically, for each abstract of this collection, we extracted the main 20 keywords with Yake! and selected the abstracts for which the semantic similarity between the 20 keywords and the UAT terms surpasses a certain threshold. We opted for 20 keywords as this was a balanced trade-off between the quantity and the quality of keywords. While we wanted to collect a high number of keywords, we also needed the keywords to be relevant. In total, we ended up with 168,084 relevant arXiv abstracts. We followed a similar procedure for selecting relevant Wikipedia abstracts, selecting 695 summaries<sup>12</sup>. Table 1 shows the main statistics, source, size, and type of all the textual data we gathered.

### 4.2 Overlapping Domains

In the future, we want to merge the KGs of the different NASA SMD domains with entity alignment [29, 31]. Before merging the KGs, we

<sup>7</sup>Link to NetworkX: <https://networkx.org/>

<sup>8</sup>Link to iGraph: <https://igraph.org/>

<sup>9</sup>Link to Pyvis: <https://pyvis.readthedocs.io/en/latest/>

<sup>10</sup>Link to Arena3D: <https://arena3d.org/>

<sup>11</sup>Link to the arXiv abstracts: <https://huggingface.co/datasets/gfissore/arxiv-abstracts-2021/blob/main/README.md>

<sup>12</sup>We used the *wikipedia* module of Python to parse summaries from Wikipedia. Link to the *wikipedia* module: <https://pypi.org/project/wikipedia/>

want to better understand how the different domains overlap. We experimented with k-means, an unsupervised clustering algorithm where all observations are partitioned into k clusters where each observation belongs to the cluster with nearest mean<sup>13</sup> [19].

We applied k-means to cluster a random selection of 1,500 Planetary Science and 1,500 Heliophysics arXiv abstracts without providing domain labels. We set k, the number of clusters, equal to two, as we were interested in whether the k-means algorithm can broadly distinguish between the two domains. K-means managed to classify the abstracts with greater than 80% accuracy. This high accuracy indicates a distinct difference between these two domains. Contrarily, we suspect the accuracy was not perfect due to a significant overlap in terminology between the domains.

## 5 CHALLENGES AND FUTURE DIRECTIONS

Many challenges arose during the effort and may show to be a continuous challenge as the work continues to improve the KGs. In this section, we will discuss our challenges concerning validation (Section 5.1), access to validated ontologies (Section 5.2), computational restraints (Section 5.3), generating KG nodes (Section 5.4), domain overlap (Section 5.5), and visualization (Section 5.6).

### 5.1 Validation

Validating KGs is our biggest challenge. We asked domain experts to assess the global quality of the KGs and provided feedback ourselves. Manually validating was time-consuming and due to time and resource constraints we only managed to validate parts of the KGs. The main feedback was that the KGs with the Yake! keywords were more appropriate. In Section 5.4, a deeper discussion about the quality of the nodes is given. In the future, we would like to include a human in the loop, so the quality of the KGs can continuously be improved. One option would be that users of the KGs can suggest changes, directly apply changes, or changes are made indirectly based on the users' behaviour [21]. In addition, we aim to implement automatic metrics such as Corroborative Fact Validation (COPAAL) [27], Deep Fact Validation (DeFacto) [18], or FactCheck [28]. COPAAL is easy to implement as it implements path scoring solely based on the provided KG. COPAAL can be complemented by DeFacto and FactCheck as these two metrics validate the quality of the KG based on external knowledge resources.

### 5.2 Access to validated ontologies

We intended to use validated ontologies as the backbone of our KGs. Unfortunately, we did not find relevant existing validated ontologies. We considered building ontologies from scratch and integrating them across the different NASA SMD domains, but eventually, we did not create them due to time constraints. In the future, we want to build ontologies manually and embed them in the KGs.

### 5.3 Computational Restraints

Creating the KGs and experimenting with the data is computationally expensive and time-consuming. The big data sets, in combination with the different NLP algorithms, result in complex and

intense processes. Due to time limitations, we limited the data size for some experiments. For example, for our domain overlap experiment (see Section 4.2), we limited the number of abstracts to 3,000 and only ran these experiments for two out of the five NASA SMD scientific pursuits.

### 5.4 Generating nodes

To generate nodes for the KGs, we considered multiple methods. We started by extracting the subject, predicate, and object for each sentence of our dataset. Instead, we found that the nodes generated by Yake! were the most suitable. For example for a sentence from the NASA SMD definition list *They are made of cooler solar material, or plasma, supported in the sun's atmosphere by magnetic fields.*, the former extract *they* and *magnetic fields* as the subject and object while the latter extracts *solar materials* and *magnetic fields* as keywords. We do believe that for the generation of the nodes, there is still the opportunity for improvement as Yake! is not specifically designed to recognize terms related to the NASA SMD. In the future, we aim to create NER algorithms similar to AstroBERT [13], and SciBERT [4]. In line with these two NER algorithms, we want to fine-tune models from the BERT family [9].

### 5.5 Domain Overlap

To promote cross-domain research, we need to generate five KGs, one for each of NASA's SMD domains, and we need to merge them. To merge the KGs, we need to apply entity alignment [29, 31]. Entity alignment is a technique where entities of different KGs are identified to be referring to the same real-world object [32]. Entity alignment is challenging, as we face many homonyms. Homonyms are words that share the same spelling but have different meanings [17]. For example, the term *storms* could point to the geomagnetic storms for solar physicists or ice storms on earth.

### 5.6 Visualization

The KGs we created were too complex to visualize. Clustering the nodes helped a little, but we had too many nodes to create a visually enticing image. We visualized KGs based on the data selections, but we found that these KGs were too limited. In addition, generating interactive KG required too many computational resources and led to slowly working interactive KGs. In the future, we would like to build a robust Application Programming Interface (API) for 3-D visualization of the knowledge graphs to assist researchers in finding new connections.

## 6 CONCLUSION

The size of the NASA SMD data catalog is growing exponentially due to digitalization and the launch of new models. The NASA SMD promotes cross-domain opportunities, while historically, the data management of these domains has been in silos. The advancement of NLP allows us to process large amounts of textual data and generate cross-domain discovery KGs. This work presents a pipeline to generate KGs for the NASA SMD domains. We explore the domain overlap and discuss how the KGs can be merged in the future.

<sup>13</sup>For k-means, we use the *scikit-learn*, also known as *sklearn*, a module of Python. Link to *scikit-learn* module: <https://scikit-learn.org/>

## ACKNOWLEDGMENTS

This work has been enabled by the Frontier Development Lab<sup>14</sup>. FDL is a collaboration between SETI Institute and Trillium Technologies Inc., in partnership with NASA, Google Cloud, Intel, NVIDIA, and many other public and private partners. The material is based upon work supported by NASA. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Aeronautics and Space Administration. The authors would like to thank the FDL organizers, the SETI institute, Trillium, the FDL partners and sponsors, and the reviewers for their constructive comments during the research sprint.

## REFERENCES

- [1] A Accomazzi, N Gray, C Erdmann, C Biemesderfer, K Frey, and J Soles. 2014. The Unified Astronomy Thesaurus. *Astronomical Data Analysis Software and Systems XXIII* 485 (2014), 461.
- [2] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. In *The Semantic Web*. Springer, 722–735.
- [3] Sören Auer, Viktor Kovtun, Manuel Prinz, Anna Kasprzik, Markus Stocker, and Maria Esther Vidal. 2018. Towards a knowledge graph for science. In *Proceedings of the 8th International Conference on Web Intelligence, Mining and Semantics*. 1–6.
- [4] Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A Pretrained Language Model for Scientific Text. 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)*. 3615–3620.
- [5] Kaylin Bugbee, Rahul Ramachandran, Ashish Acharya, Dai-Hai Ton That, John Hedman, Ahmed Eleish, Charles Driessnack, Wesley Adams, and Emily Foshee. 2022. Selecting Approaches for Enabling Enterprise Data Search: NASA's Science Mission Directorate (SMD) Catalog. In *IGARSS 2022-2022 Institute of Electrical and Electronics Engineers (IEEE) International Geoscience and Remote Sensing Symposium*. IEEE, 6836–6839.
- [6] Ricardo Campos, Vitor Mangaravite, Arian Pasquali, Alípio Jorge, Célia Nunes, and Adam Jatowt. 2020. YAKE! Keyword extraction from single documents using multiple local features. *Information Sciences* 509 (2020), 257–289.
- [7] Thomas Chen. 2022. Bridging astrophysics and heliophysics: observing high-energy solar flaring activity to understand stellar flares. *AAS/High Energy Astrophysics Division* 54, 3 (2022), 108–61.
- [8] Gabor Csardi, Tamas Nepusz, et al. 2006. The igraph software package for complex network research. *International Journal of Complex Systems* 1695, 5 (2006), 1–9.
- [9] 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 (NAACL-HLT) 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, 4171–4186.
- [10] NASA's Science Mission Directorate. 2019. *Science Mission Directorate's Strategy for Data Management and Computing for Groundbreaking Science 2019 – 2024*. Technical Report.
- [11] Ruth Duerr, Ahmed Eleish, Mark Parsons, Daniel Berrios, Kaylin Bugbee, and Peter Fox. 2021. Developing the Cross-Disciplinary Information Model for NASAs Science Mission Directorate. In *American Geophysical Union (AGU) Fall Meeting Abstracts*, Vol. 2021. IN45H–0520.
- [12] Lisa Ehrlinger and Wolfram Wöfl. 2016. Towards a Definition of Knowledge Graphs. In *International Conference on Semantic Systems*.
- [13] Felix Grezes, Sergi Blanco-Cuarema, Alberto Accomazzi, Michael J Kurtz, Golnaz Shapurian, Edwin Henneken, Carolyn S Grant, Donna M Thompson, Roman Chyla, Stephen McDonald, et al. 2021. Building astroBERT, a language model for Astronomy & Astrophysics. *arXiv preprint arXiv:2112.00590* (2021).
- [14] Sonal Gupta and Christopher D Manning. 2011. Analyzing the dynamics of research by extracting key aspects of scientific papers. In *Proceedings of 5th international joint conference on natural language processing*. 1–9.
- [15] Aric Hagberg, Pieter Swart, and Daniel S Chult. 2008. *Exploring network structure, dynamics, and function using NetworkX*. Technical Report. Los Alamos National Lab.(LANL), Los Alamos, NM (United States).
- [16] Yufang Hou, Charles Jochim, Martin Gleize, Francesca Bonin, and Debasish Ganguly. 2021. TDMSci: A Specialized Corpus for Scientific Literature Entity Tagging of Tasks Datasets and Metrics. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. 707–714.
- [17] James R Hurford, Brendan Heasley, and Michael B Smith. 2007. *Semantics: a coursebook*. Cambridge university press.
- [18] Jens Lehmann, Daniel Gerber, Mohamed Morsey, and Axel-Cyrille Ngonga Ngomo. 2012. Defacto-deep fact validation. In *International semantic web conference*. Springer, 312–327.
- [19] Stuart Lloyd. 1982. Least squares quantization in PCM. *Institute of Electrical and Electronics Engineers (IEEE) transactions on information theory* 28, 2 (1982), 129–137.
- [20] Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. 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*. 3219–3232.
- [21] Emaad Manzoor, Jordan Tong, Sriniketh Vijayaraghavan, and Rui Li. 2022. Expanding Knowledge Graphs with Humans in the Loop. *arXiv preprint arXiv:2212.05189* (2022).
- [22] Megan Mark, Fech Scen Khoo, Marcella Scoczynski Ribeiro Martins, Roelien C Timmer, Anamaria Berea, and Gregory Renard. 2022. A Knowledge Graph Connecting Multiple NASA Science Mission Directorate Domains. In *Fall Meeting 2022*. American Geophysical Union (AGU).
- [23] Ryan McGranaghan, Samuel J Klein, and Agnes Cameron. 2021. The NASA Heliophysics KNOWledge Network (Helio-KNOW) project (an essay for the Space Data Knowledge Commons). (2021).
- [24] Georgios A Pavlopoulos, Seán I O'Donoghue, Venkata P Satagopam, Theodoros G Soldatos, Evangelos Pafilis, and Reinhard Schneider. 2008. Arena3D: visualization of biological networks in 3D. *BMC systems biology* 2, 1 (2008), 1–7.
- [25] Giancarlo Perrone, Jose Unpingco, and Haw-minn Lu. 2020. Network visualizations with Pyvis and VisJS. *arXiv preprint arXiv:2006.04951* (2020).
- [26] Amit Singhal et al. 2012. Introducing the knowledge graph: things, not strings. *Official Google blog* 5 (2012), 16.
- [27] Zafar Habeeb Syed, Michael Röder, and Axel-Cyrille Ngonga Ngomo. 2019. Un-supervised discovery of corroborative paths for fact validation. In *International Semantic Web Conference*. Springer, 630–646.
- [28] Zafar Habeeb Syed, Michael Röder, and Axel-Cyrille Ngonga Ngomo. 2018. Factcheck: Validating rdf triples using textual evidence. In *Proceedings of the 27th Association for Computing Machinery (ACM) International Conference on Information and Knowledge Management*. 1599–1602.
- [29] Bayu Distiawan Trisedya, Jianzhong Qi, and Rui Zhang. 2019. Entity alignment between knowledge graphs using attribute embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 297–304.
- [30] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the Association for Computing Machinery (ACM)* 57, 10 (2014), 78–85.
- [31] Yuting Wu, Xiao Liu, Yansong Feng, Zheng Wang, Rui Yan, and Dongyan Zhao. 2019. Relation-Aware Entity Alignment for Heterogeneous Knowledge Graphs. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19)*. International Joint Conferences on Artificial Intelligence Organization, 5278–5284. <https://doi.org/10.24963/ijcai.2019/733>
- [32] Kaisheng Zeng, Chengjiang Li, Lei Hou, Juanzi Li, and Ling Feng. 2021. A comprehensive survey of entity alignment for knowledge graphs. *AI Open* 2 (2021), 1–13.
- [33] Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli. 2021. Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*.

<sup>14</sup>Link to Frontier Development Lab website: <https://frontierdevelopmentlab.org/>