

Scholarly Knowledge Extraction from Published Software Packages

Muhammad Haris¹[0000–0002–5071–1658], Markus Stocker^{2,1}[0000–0001–5492–3212],
and Sören Auer^{2,1}[0000–0002–0698–2864]

¹ L3S Research Center, Leibniz University Hannover 30167, Hannover, Germany
`haris@l3s.de`

² TIB—Leibniz Information Centre for Science and Technology, Germany
`{markus.stocker,auer}@tib.eu`

Abstract. A plethora of scientific software packages are published in repositories, e.g., Zenodo and figshare. These software packages are crucial for the reproducibility of published research. As an additional route to scholarly knowledge graph construction, we propose an approach for automated extraction of machine actionable (structured) scholarly knowledge from published software packages by static analysis of their (meta)data and contents (in particular scripts in languages such as Python). The approach can be summarized as follows. First, we extract metadata information (software description, programming languages, related references) from software packages by leveraging the Software Metadata Extraction Framework (SOMEF) and the GitHub API. Second, we analyze the extracted metadata to find the research articles associated with the corresponding software repository. Third, for software contained in published packages, we create and analyze the Abstract Syntax Tree (AST) representation to extract information about the procedures performed on data. Fourth, we search the extracted information in the full text of related articles to constrain the extracted information to scholarly knowledge, i.e. information published in the scholarly literature. Finally, we publish the extracted machine actionable scholarly knowledge in the Open Research Knowledge Graph (ORKG).

Keywords: Analyzing Software Packages · Open Research Knowledge Graph · Code Analysis · Abstract Syntax Tree · Scholarly Communication · Machine Actionability

1 Introduction

A variety of general and domain-specific knowledge graphs have been proposed to represent (scholarly) knowledge in a structured manner [7,26]. General purpose knowledge graphs include DBpedia³ [18], Wikidata⁴ [25], YAGO [23], etc., whereas domain-specific infrastructures include approaches in Cultural

³ <https://www.dbpedia.org>

⁴ https://www.wikidata.org/wiki/Wikidata:Main_Page

Heritage [13], KnowLife in Life Sciences [9], Hi-Knowledge in Invasion Biology⁵ [10,8], COVID-19 Air Quality Data Collection⁶, Papers With Code in Machine Learning⁷, Cooperation Databank in Social Sciences⁸ [21], among others. In addition, knowledge graph technologies have also been employed to describe software packages in a structured manner [16,2].

Extending the state-of-the-art, we propose an approach for scholarly knowledge extraction from published software packages by static analysis of package contents, i.e., (meta-)data and software (in particular, Python scripts), and represent the extracted knowledge in a knowledge graph. The main purpose of this knowledge graph is to capture information about the materials and methods used in scholarly work described in research articles.

We address the following research question: Can structured scholarly knowledge be automatically extracted from published software packages? Our approach consists of the following steps:

1. *Mining software packages* deposited in Zenodo⁹ using its REST API¹⁰ and analyzing the API response to extract the linked metadata information, i.e., associated scholarly articles. We complement the approach by leveraging the Software Metadata Extraction Framework (SOMEF) to parse the README files and extract other related metadata information (i.e., software name, description, used programming languages).
2. *Perform static code analysis* to extract information about the procedures performed on data. We utilize Abstract Syntax Tree (AST) representations to statically analyze program code and identify operations performed on data.
3. *Identify scholarly knowledge* by performing keyword-based search of extracted information in article full text. Thus, among all the information extracted from software packages we identify that which is scholarly knowledge.
4. *Construct a knowledge graph* of scholarly knowledge extracted from software packages. For this purpose, we leverage the Open Research Knowledge Graph (ORKG)¹¹ [14], a production research infrastructure that supports producing and publishing machine actionable scholarly knowledge.

2 Related Work

Several approaches have been suggested to retrieve metadata from software repositories. Mao et al. [19] proposed the Software Metadata Extraction Framework (SOMEF) to extract metadata from software packages published on

⁵ <https://hi-knowledge.org>

⁶ <https://covid-aqs.fz-juelich.de>

⁷ <https://paperswithcode.org>

⁸ <https://cooperationdatabank.org>

⁹ <https://zenodo.org>

¹⁰ <https://developers.zenodo.org>

¹¹ <https://www.orkg.org/orkg/>

GitHub. Specifically, the framework employs machine learning-based methods to extract repository name, software description, citations, reference URLs, etc. from README files and to represent the metadata in structured formats (JSON-LD, JSON and RDF). SOMEF was later extended to extract additional metadata and auxiliary files (e.g., Notebooks, Dockerfiles) from software packages [16]. Moreover, the extended work also supports creating a knowledge graph of parsed metadata, thus improving search of software deposited in repositories. Abdelaziz et al. [1] proposed CodeBreaker, a knowledge graph with information about 1.3 million Python scripts published on GitHub. The graph was embedded in an IDE to recommend code functions while writing software. Similarly, Graph-Gen4Code [2] is a knowledge graph with information about software included in GitHub repositories. It was generated by analyzing the functionalities of Python scripts and linking them with the natural language artefacts (documentation and forum discussions on StackOverflow and StackExchange). The knowledge graph contains 2 billion triples. Several other machine learning-based approaches for searching [11] software scripts and summarization [3,12] have been proposed. The Pydriller [22] and GitPython¹² frameworks were proposed to mine information from GitHub repositories, including source code, commits, branch differences, etc. Similarly, ModelMine [20] mines and analyzes models included in repositories. Vagavolu et al. [24] presented an approach that leverages Code2vec [17] and includes semantic graphs with Abstract Syntax Tree (AST) for performing different software engineering tasks. [4] presented an AST based-approach for code representation and considered code data flow mechanisms to suggest code improvements.

3 Methodology

In this section, we present our methodology for automatically extracting *scholarly* knowledge from software packages and building a knowledge graph from the extracted meta(data). Figure 1 provides an overview of the key components.

3.1 Mining Software Packages

We mine software packages from the Zenodo repository by leveraging its REST API. The metadata of each package is analyzed to retrieve its DOI and metadata about related versions and associated scholarly articles. The versions of software packages are retrieved by interpreting `relation: isVersionOf` metadata, whereas the DOI of the linked article, if available, is fetched using the `relation: cites` or `relation: isSupplementTo` metadata. We also leverage the Software Metadata Extraction Framework (SOMEF) and GitHub API to extract additional metadata from software packages, in particular software name, description, used programming languages, GitHub URL. Since not all software

¹² <https://github.com/gitpython-developers/GitPython>

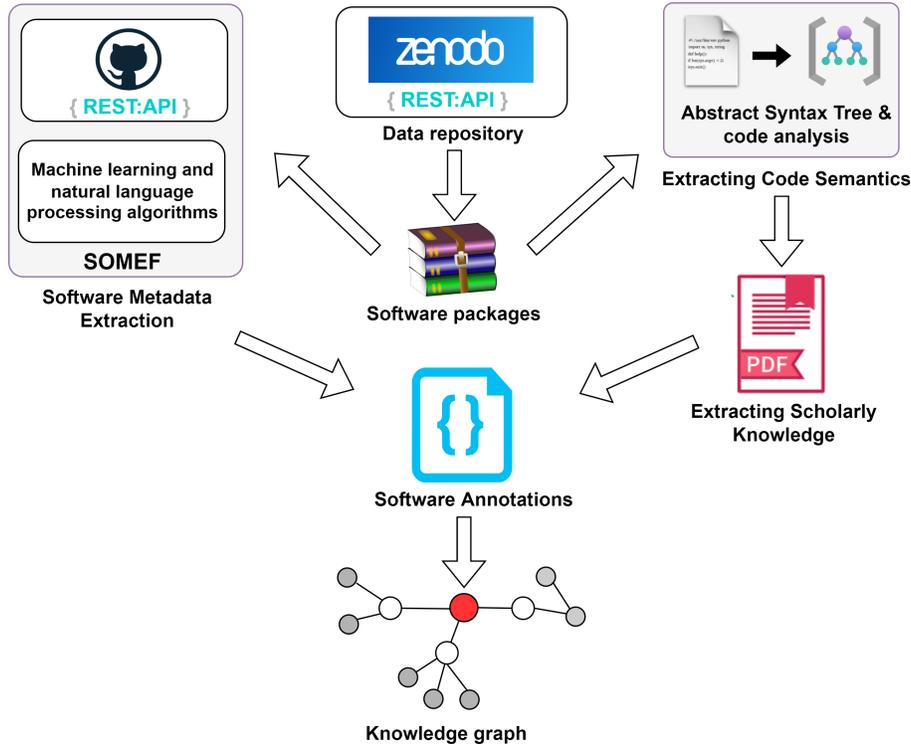


Fig. 1: Pipeline for constructing a knowledge graph of scholarly knowledge extracted from software packages: 1) Mining software packages from the Zenodo repository using its REST API; 2) Extracting software metadata by analyzing the Zenodo API results as well as the GitHub API, using SOMEF; 3) Performing static code analysis using AST representations of software to extract code semantics, in particular operations on data; 4) Performing keywords-based search in article full texts to identify scholarly knowledge; 5) Knowledge graph construction with scholarly knowledge extracted from software packages.

packages include the `cites` or `isSupplementTo` relations in metadata, we utilize SOMEF to parse the README files of software packages as an additional approach to extract the DOI of the related scholarly article.

Static Code Analysis We utilize Abstract Syntax Tree (AST) representations for static analysis of Python scripts included in software packages. AST provides structured representations of scripts, omitting unnecessary syntactic details (e.g., semicolons, commas, and comments). Our goal is to extract information about the data used in scripts and the procedures performed on that data. Our developed Python-based module sequentially reads the scripts contained in software packages and generates the AST. The implemented procedures and variables are

```

1 test_data = pd.read_csv("Sample.csv")
2 reference = pd.read_csv("Reference.csv")
3 train = reference.drop("MF_name",1)
4 test_data = test_data.drop("gene",1)
5
6 score_adj = []
7 for o in range(len(test_data.columns)):
8     test = test_data.loc[:,test_data.columns[o]]
9     im_name = train.columns
10    svr = LinearSVR(random_state=0)
11    model = svr.fit(train, test)
12    score = model.coef_
13    score[np.where(score<0)] = 0
14    score_adj.append((score/sum(score)))
15 score_adj = pd.DataFrame(score_adj)
16 score_adj.to_csv("score.csv")

```

 **Input datasets**
 **Output dataset**

Fig. 2: Static code analysis: Exemplary Python script (shortened) included in a software package. The script lines highlighted with same color show different procedural changes that a particular variable has undergone.

tokenized and represented as nodes in the tree, which facilitates the analysis of the code flow. Thus, by traversing the tree we extract the information about the data used in the scripts, the procedures performed on the data and, if available, the output data. Fig. 2 shows the Python script included in the software package¹³. The script shows an example in which `Sample.csv` and `Reference.csv` used as input data, then the operation `LinearSVR` is performed on the data, and finally the resulting data `score.csv` is generated.

Fig. 3 shows the AST of the Python script (Fig. 2) created using a suitable Python library¹⁴. For simplicity, we show the AST of lines 1, 10, and 16. In the tree structure, the name of the node represents the functionality of each line of the script. For example, line 1 performs a task that reads data and assigns it to a variable. Therefore, the relevant node in the tree is labelled `Assign`. We retrieve all leaf nodes since they represent variables, their values, and procedures. Analyzing these script semantics, we can then find the flow of data between

¹³ <https://zenodo.org/record/5874955>

¹⁴ <https://docs.python.org/3/library/ast.html>

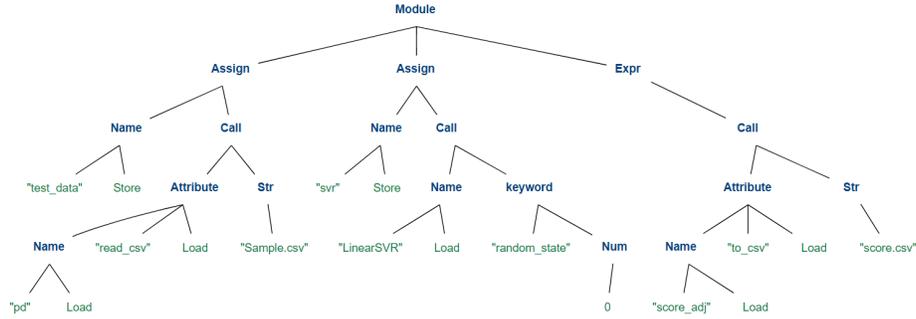


Fig. 3: Abstract Syntax Tree (AST) of the script shown in Fig. 2. For simplicity, the AST is shown only for Lines 1, 10 and 16.

procedures. We investigate the flow of variables that contain the input data, i.e., examining which operations used a particular variable as a parameter.

3.2 Identifying Scholarly Knowledge

Not all information extracted from software packages and AST-analyzed program code is scholarly knowledge. Information is scholarly knowledge if it is included in a scholarly article. Hence, we filter the information extracted from software packages for information referred to in the article citing the software package. For this, we employ keyword-based search. Specifically, we search for the terms extracted in AST-analyzed program code in the related article full text. Assuming that the DOI of the related article has been identified, we fetch the PDF version of the article by utilizing the Unpaywall REST API¹⁵. We make use of the Unpaywall API because, contrary to DOI metadata, it provides the URL to the PDF version of scholarly articles. In our example (Fig. 2), the extracted terms (`Sample`, `Reference`, `read_csv`, `LinearSVR`, `svr.fit`, and `to_csv`) are searched in the PDF and we find `Sample`, `Reference` and `LinearSVR` are cited in the scholarly article. We thus assume that the extracted information is scholarly knowledge.

3.3 Knowledge Graph Construction

We now construct the knowledge graph with the scholarly knowledge obtained in the analysis of software packages. For this, we leverage the Open Research Knowledge Graph (ORKG) [14]. The ORKG aims to represent scholarly articles in a machine actionable and structured form. Abstractly speaking, the ORKG represents research contributions describing key results, the materials and methods used to obtain the results, and the addressed research problem.

¹⁵ https://api.unpaywall.org/v2/10.1186/s12920-019-0613-5?email=unpaywall_01@example.com

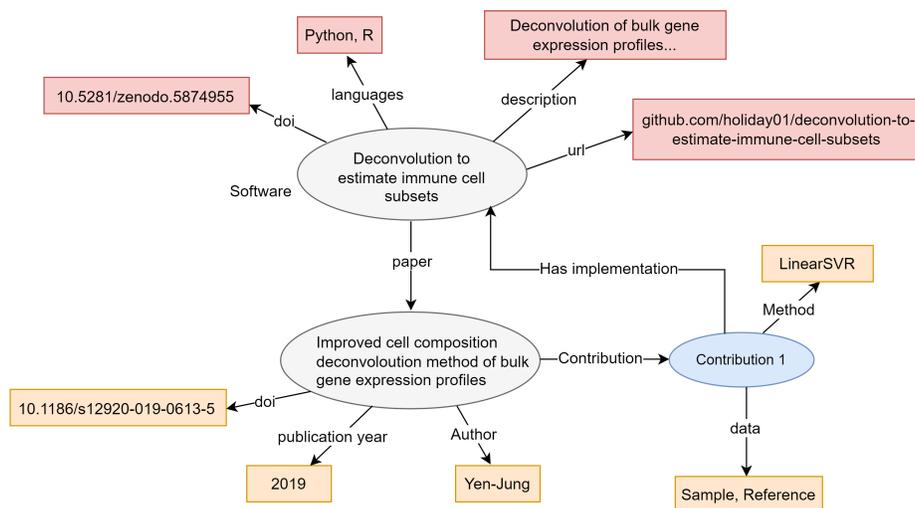


Fig. 4: Knowledge graph depicting the scholarly knowledge extracted from a software package related to an article, describing key aspects (e.g., method used) of a research contribution of the work described in the article.

The scholarly information extracted from software packages is organized in triples and ingested into ORKG using its REST API. Fig. 4 shows the resulting knowledge graph for a paper and its research contribution¹⁶. The figure also shows the metadata of corresponding software package¹⁷.

4 Results and Discussion

At the time of writing, there are more than 80,000 software packages available on Zenodo. To expedite the execution process, we discard packages larger than 400 MB. We thus consider 52,236 software packages. We further process only those software packages that are also available on GitHub, that is 40,239 packages. We analyze the metadata of the software packages and the respective README files and find a total of 6221 research articles, of which 642 articles are associated with the related software packages in metadata through the `cites` or `isSupplementTo` relations. The remaining 5579 articles are extracted by analyzing the README files of the software packages using SOMEF. We only analyze software packages that include Python scripts and have linked scholarly articles, that is 2172 packages. Table 1 summarizes the statistics.

Out of 6221 articles, 4328 are described in ORKG because for the remaining articles the DOIs in README files are not parsed correctly. The articles added to ORKG include ORKG research contribution descriptions linking the software

¹⁶ <https://orkg.org/paper/R209873>

¹⁷ <https://orkg.org/content-type/Software/R209880>

Table 1: Statistics about the (scholarly) information extracted from software packages.

Entity	Total
<i>Software package</i>	52236
<i>Paper</i>	Explicit links in metadata: 642; SOMEF-based link extraction: 5579 (Total: 6221)
<i>GitHub URL</i>	40,239
<i>Python-based software packages, linked with articles</i>	2172
<i>Analyzed Python scripts</i>	67,936

package and including information about computational methods and data used in research extracted by analyzing the software packages.

Software semantics and Named Entity Recognition (NER) models There exist numerous approaches for the extraction of scholarly knowledge from articles using machine learning and natural language processing, including scientific named entity recognition [15,6] and sentence classification [5]. These approaches process the entire text to extract the essential entities in scholarly articles, which is costly in terms of data collection and training. Moreover, the approaches require large training data to achieve acceptable performance. We argue that extracting scholarly knowledge from software packages as proposed here is a significant step towards automated and cheap construction of scholarly knowledge graphs. Instead of extracting scholarly entities from full texts using machine learning models, the scholarly knowledge is extracted from related software packages with more structured data.

Future directions In future work, we aim to develop a pipeline that will automatically execute the software packages that contain scholarly knowledge. Such an approach can be integrated into software repositories (zenodo, figshare) to automatically execute the published software and determine whether the (extracted) scholarly knowledge is reproducible.

5 Conclusions

Our work is an important step towards automated and scalable mining of scholarly knowledge from published software packages and creating the knowledge graph using the extracted data. The resulting knowledge graph holds the links between articles and software packages, as well as and most interestingly descriptions of the computational methods and materials used in research work presented in articles. Evaluated on zenodo, our approach can be extended to other repositories, e.g., figshare, as well as software in languages other than Python, e.g., R, Java, Javascript, and C++—potentially further increasing the number of articles and related scholarly knowledge added to ORKG.

Acknowledgment

This work was co-funded by the European Research Council for the project ScienceGRAPH (Grant agreement ID: 819536) and TIB–Leibniz Information Centre for Science and Technology.

References

1. Abdelaziz, I., Srinivas, K., Dolby, J., McCusker, J.P.: A demonstration of code-breaker: A machine interpretable knowledge graph for code. In: SEMWEB (2020)
2. Abdelaziz, I., Dolby, J., McCusker, J., Srinivas, K.: A toolkit for generating code knowledge graphs. In: Proceedings of the 11th on Knowledge Capture Conference. p. 137–144. K-CAP '21, Association for Computing Machinery, New York, NY, USA (2021). <https://doi.org/10.1145/3460210.3493578>, <https://doi.org/10.1145/3460210.3493578>
3. Ahmad, W., Chakraborty, S., Ray, B., Chang, K.W.: A transformer-based approach for source code summarization. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 4998–5007. Association for Computational Linguistics (Jul 2020). <https://doi.org/10.18653/v1/2020.acl-main.449>
4. Allamanis, M., Brockschmidt, M., Khademi, M.: Learning to represent programs with graphs. arXiv preprint arXiv:1711.00740 (2017)
5. Brack, A., Hoppe, A., Buschermöhle, P., Ewerth, R.: Cross-domain multi-task learning for sequential sentence classification in research papers. In: Proceedings of the 22nd ACM/IEEE Joint Conference on Digital Libraries. JCDL '22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3529372.3530922>, <https://doi.org/10.1145/3529372.3530922>
6. Brack, A., Müller, D.U., Hoppe, A., Ewerth, R.: Coreference resolution in research papers from multiple domains. In: Hiemstra, D., Moens, M.F., Mothe, J., Perego, R., Potthast, M., Sebastiani, F. (eds.) *Advances in Information Retrieval*. pp. 79–97. Springer International Publishing, Cham (2021)
7. Chao, K., Tao, L., Li, M., Guoyu, J., yuchao, W., Yu, Z.: Construction and application research of knowledge graph in spacecraft launch. *Journal of Physics: Conference Series* **1754**, 012180 (02 2021). <https://doi.org/10.1088/1742-6596/1754/1/012180>
8. Enders, M., Havemann, F., Ruland, F., Bernard-Verdier, M., Catford, J.A., Gómez-Aparicio, L., Haider, S., Heger, T., Kueffer, C., Kühn, I., et al.: A conceptual map of invasion biology: Integrating hypotheses into a consensus network. *Global Ecology and Biogeography* **29**(6), 978–991 (2020)
9. Ernst, P., Meng, C., Siu, A., Weikum, G.: Knowlife: A knowledge graph for health and life sciences. pp. 1254–1257 (03 2014). <https://doi.org/10.1109/ICDE.2014.6816754>
10. Heger, T., Pahl, A.T., Botta-Dukát, Z., Gherardi, F., Hoppe, C., Hoste, I., Jax, K., Lindström, L., Boets, P., Haider, S., et al.: Conceptual frameworks and methods for advancing invasion ecology. *Ambio* **42**(5), 527–540 (2013)
11. Husain, H., Wu, H.H., Gazit, T., Allamanis, M., Brockschmidt, M.: Codesearchnet challenge: Evaluating the state of semantic code search (June 2020), <https://www.microsoft.com/en-us/research/publication/codesearchnet-challenge-evaluating-the-state-of-semantic-code-search/>

12. Iyer, S., Konstas, I., Cheung, A., Zettlemoyer, L.: Summarizing source code using a neural attention model. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 2073–2083. Association for Computational Linguistics (Aug 2016). <https://doi.org/10.18653/v1/P16-1195>
13. Jain, N.: Domain-specific knowledge graph construction for semantic analysis. In: Harth, A., Presutti, V., Troncy, R., Acosta, M., Polleres, A., Fernández, J.D., Xavier Parreira, J., Hartig, O., Hose, K., Cochez, M. (eds.) The Semantic Web: ESWC 2020 Satellite Events. pp. 250–260. Springer International Publishing, Cham (2020)
14. Jaradeh, M.Y., Oelen, A., Farfar, K.E., Prinz, M., D’Souza, J., Kismihók, G., Stocker, M., Auer, S.: Open research knowledge graph: Next generation infrastructure for semantic scholarly knowledge. In: Proceedings of the 10th International Conference on Knowledge Capture. p. 243–246. K-CAP ’19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3360901.3364435>
15. Jiang, M., D’Souza, J., Auer, S., Downie, J.S.: Improving scholarly knowledge representation: Evaluating bert-based models for scientific relation classification. In: Ishita, E., Pang, N.L.S., Zhou, L. (eds.) Digital Libraries at Times of Massive Societal Transition. pp. 3–19. Springer International Publishing, Cham (2020)
16. Kelley, A., Garijo, D.: A framework for creating knowledge graphs of scientific software metadata. *Quantitative Science Studies* **2**(4), 1423–1446 (12 2021). https://doi.org/10.1162/qss_a_00167
17. Le, Q., Mikolov, T.: Distributed representations of sentences and documents. In: International conference on machine learning. pp. 1188–1196. PMLR (2014)
18. Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P., Hellmann, S., Morse, M., Van Kleef, P., Auer, S., Bizer, C.: Dbpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web Journal* **6** (01 2014). <https://doi.org/10.3233/SW-140134>
19. Mao, A., Garijo, D., Fakhraei, S.: Somef: A framework for capturing scientific software metadata from its documentation. In: 2019 IEEE International Conference on Big Data (Big Data). pp. 3032–3037 (2019). <https://doi.org/10.1109/BigData47090.2019.9006447>
20. Reza, S.M., Badreddin, O., Rahad, K.: ModelMine: A Tool to Facilitate Mining Models from Open Source Repositories. Association for Computing Machinery, New York, NY, USA (2020), <https://doi.org/10.1145/3417990.3422006>
21. Spadaro, G., Tiddi, I., Columbus, S., Jin, S., Ten Teije, A., Team, C., Balliet, D.: The cooperation databank: Machine-readable science accelerates research synthesis. *Perspectives on Psychological Science* p. 174569162111053319 (2020)
22. Spadini, D., Aniche, M., Bacchelli, A.: Pydriller: Python framework for mining software repositories. p. 908–911. ESEC/FSE 2018, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3236024.3264598>
23. Suchanek, F., Kasneci, G., Weikum, G.: Yago: a core of semantic knowledge. pp. 697–706 (01 2007). <https://doi.org/10.1145/1242572.1242667>
24. Vagavolu, D., Swarna, K.C., Chimalakonda, S.: A mocktail of source code representations. In: 2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE). pp. 1296–1300 (2021). <https://doi.org/10.1109/ASE51524.2021.9678551>

25. Vrandečić, D., Krötzsch, M.: Wikidata: A free collaborative knowledgebase. *Commun. ACM* **57**(10), 78–85 (sep 2014). <https://doi.org/10.1145/2629489>, <https://doi.org/10.1145/2629489>
26. Zhao, Z., Han, S.K., So, I.M.: Architecture of knowledge graph construction techniques (2018)