

Reproducible Research is more than Publishing Research Artefacts: A Systematic Analysis of Jupyter Notebooks from Research Articles

Max Schröder^{1,2}, Frank Krüger¹, and Sascha Spors¹

¹ *Institute of Communications Engineering, University of Rostock*

² *University Library, University of Rostock*

E-Mail: {max.schroeder, frank.krueger, sascha.spors}@uni-rostock.de

1 Introduction

With the advent of Open Science, researchers have started to publish their *research artefacts* (i.e., data, software, and other products of the investigations) in order to allow others to reproduce their investigations. While this publication is beneficial for science in general [1], it often lacks a comprehensive documentation and completeness with respect to the artefacts. This, in turn, prevents the successful reproduction of the analyses.

Jupyter notebooks recently have gained increased attention as a method for publishing research investigations [2, 3]. The term ‘Jupyter notebook’ is used for both: (1) a JSON document (*.ipynb files) and (2) a corresponding web service for reading, writing, and executing these notebooks. The notebooks encapsulate both documentation and source code along with the source code output i.e., the results in so-called *cells* inside a single streamlined document (see Fig. 1). The web application enables text editing, but also the interactive execution of the source code cells and, thus, the re-execution of the previous investigations. For this purpose, different programming kernels such as Python, R, and Octave exists.

Despite the increasing use of Jupyter notebooks for the publication of research investigations, their reproducibility is not automatically guaranteed by the document format and the web service. A crucial reproducibility requirement is the detailed documentation of the *computing environment*. As new software releases are being developed, the dependencies have to explicitly state the software version used for the study. Ideally, the environment is ready-to-use for other researchers by containerisation techniques.

We have been *systematically* analysing research publications that also published their investigations as Jupyter notebooks. In this paper, we present preliminary results of this analysis for five publications. The results show, that the quality of the published research artefacts must be improved in order to assure reproducibility.

The next section introduces how relevant publications were identified and how the analysis was performed. The results are described in Section 3 which is followed by our conclusion (Section 4).

2 Method

In order to identify recent publications including their published Jupyter notebooks, we used the PubMed Cen-

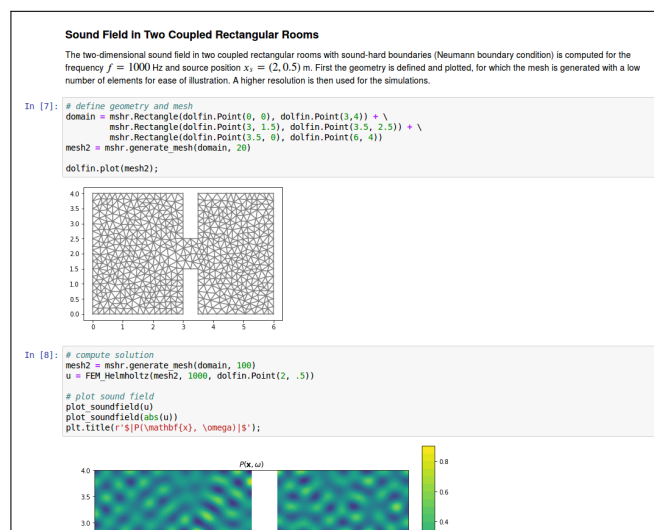


Figure 1: Screenshot of a Jupyter notebook [4] containing three ‘cells’: the first one contains textual documentation and the latter two contain python source code followed by the corresponding output in the form of plots.

tral database (PMC) [5]. This database has been chosen, because its articles are freely accessible and it contains a large number of journals. The notebook file type ‘ipynb’ was used as search term. Additionally, we restricted the results to publications from the year 2018 as we expect that software dependencies are more likely to be available and that better practices are implemented in more recent publications. From 43 publications in the PMC database seven publications were excluded as the jupyter notebooks mentioned do not refer to the original analysis of the publication (n=6) or are no longer available in the repository (n=1). For the preliminary results in this paper, we randomly selected five publications from the resulting 36 publications. Figure 2 (left) illustrates the overall workflow.

Each of the five publications was analysed with respect to the reproducibility of its jupyter notebooks as specified in Figure 2 (right). If the computing environment is not available as virtual machine or container image, we attempted to reconstruct it based on the documentation within the repository. For this purpose, we use the containerisation solution ‘Docker’ and an official Jupyter notebook image that is incrementally extended by installing requirements. The reconstruction attempt was limited to 3 hours and a single environment per article, as we with unlimited resources the reconstruction is

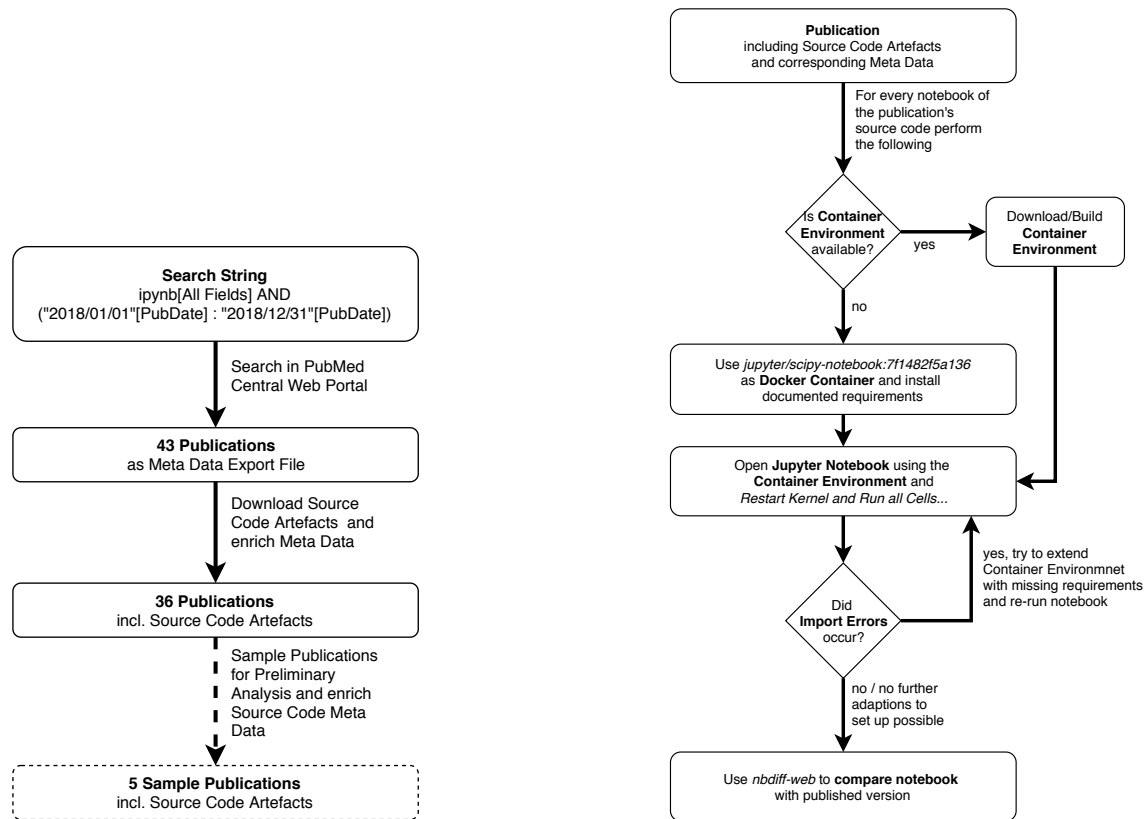


Figure 2: The identification workflow of relevant publications from the PubMed Central database (left); the reproducibility analysis performed for every identified publication (right).

always possible and authors use a homogeneous environment during their investigations if not stated otherwise. Afterwards, all cells of the Jupyter notebook are executed and the result compared to the published version using the notebook differentiation tool ‘nbdime’. The analysis in this paper is performed on a Dell Latitude 7490 with Ubuntu¹ v16.04.6, Docker² v18.09.3, and conda³ v4.6.8.

3 Results

The meta data of the overall set of relevant publications is analysed in the next section. Section 3.2 contains the preliminary results of a detailed evaluation of the reproducibility of the Jupyter notebooks.

3.1 Meta Data Analysis

The meta data has been analysed for the full set of identified publications (n=36) according to Fig. 2 (left). We evaluated two aspects that are crucial for the accessibility of the source code artefacts: the repository that is used and the source code license (consecutively):

GitHub is the most frequently used repository for the upload of source code artefacts (see Fig. 3). Most mentions of GitHub repositories in the publications, however, lack in specifying a concrete version of the source code that is used for the investigations and mention the base url only.

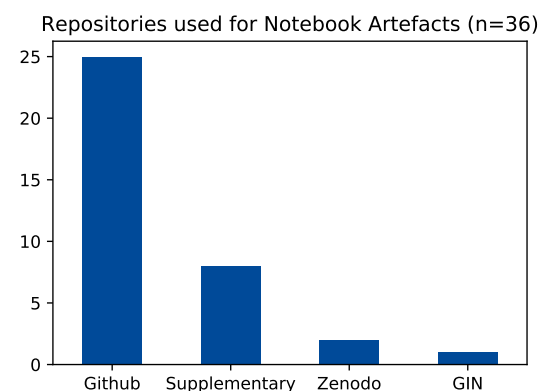


Figure 3: Which repositories are used to publish source code artefacts i. e., Jupyter notebooks?

Although this is common practice especially for software projects that are continuously developed, this prevents others from comprehending the original study. Instead, the other repositories Supplementary Material, Zenodo⁴, and GIN⁵ reference the original artefacts.

Even though the publication is publicly accessible by a proper license, this needs not be true for the source code artefacts that have to be declared as well. Unfortunately, almost one third of the source code artefacts lack in providing a license for their usage (see Fig. 4) and, thus,

¹see <https://www.ubuntu.com/>

²see <https://www.docker.com/>

³see <https://anaconda.org/>

⁴see <https://zenodo.org/>

⁵see <https://web.gin.g-node.org/>




Figure 4: Shows the publishing license of the Jupyter notebooks w.r.t. the publishing license of the article. The vertical line separates problematic artefact licenses (left) from standardized open access licenses (right).

impeding other researchers from re-using their investigations. For the software artefacts the most common license is the MIT⁶ license, the GPLv3⁷, and interestingly the CC0 1.0⁸ which gives all usage rights to the public domain. Interestingly, some of the articles even though listed in the PMC database are not freely accessible, but only via proprietary licenses or with no license information at all. One software artefact had no license information within the repository branch mentioned in the publication, but within the master branch the license was declared as GPLv3. In Figure 4 this is depicted as ‘Unknown/GPLv3’.

3.2 Reproducibility Analysis

For the preliminary reproducibility analysis, we randomly selected five publications from the overall set of publications (see Section 2). Each of these publications has been evaluated in detail w.r.t. the following reproducibility aspects:

- R1. How many notebooks are mentioned and how many are published?
- R2. Where are the source code artefacts documented beside the publication?
- R3. Where can the documentation of the software requirements be found?
- R4. Is the computing environment available or can it be reconstructed from the documentation?
- R5. Is the complete raw data of the study available?
- R6. Can the Jupyter notebook be completely re-executed with the same results?

Interestingly, one repository contained eleven Jupyter

⁶see <https://opensource.org/licenses/MIT>

⁷see <https://www.gnu.org/licenses/gpl-3.0>

⁸see <https://creativecommons.org/publicdomain/zero/1.0/>




Figure 5: Illustration of the reproducibility analysis showing several aspects that were analysed for every notebook from the five sample publications. Some aspects could not be rated as prerequisites are not satisfied e.g., a computing environment exists.

notebooks whereas the corresponding article mentioned only a single one (see Table 1). The other articles mentioned the majority of their notebooks which might indicate, that the publication of the source code is specifically tailored for that article (R1).

The documentation of the Jupyter notebooks was not evaluated w.r.t. their quality but only the presence of comments and explanations (R2). Table 1 shows that indeed the notebooks included some part of the documentation as well as a readme file. However, only two projects used methods such as Read the Docs respectively Sphinx⁹ and mkdocs¹⁰ that are specifically tailored for extended documentation. Interestingly, the documentation of the requirements (R3) can not always be found in the notebook itself. Other sources of documentation are also containerisation descriptions (see Table 1).

Another essential aspect for re-using Jupyter notebooks is that a computing environment is available that behaves identical w.r.t. the application compared to the environment used by the authors (R4). Environment here denotes the software dependencies that are needed to execute the notebook. Unfortunately, only one project used containerisation techniques such as Docker and Travis-CI¹¹ so that for all other projects the environment has to be reconstructed from scratch (see Table 1). We call a reconstruction of the environment successful if no import errors arose. This reconstruction was successful for 9 from the 22 notebooks from the five sample publications (see Fig. 5). The majority of environments where the reconstruction was unsuccessful results from the software artefact with 11 notebooks even though this repository

⁹see <https://readthedocs.org/>

¹⁰see <https://www.mkdocs.org/>

¹¹see <https://travis-ci.org/>

Ref.	Mentions	Documentation	Requirements	Req. Problems
[6]	1 / 2	notebook, readme	no documentation	—
[7]	1 / 11	html, notebook, readme, readthedocs	docker, notebook, readme readthedocs, travisci	docker fails build, custom image did not install successfully
[8]	1 / 2	mkdocs, notebook, readme	mkdocs	missing versions
[9]	4 / 6	notebook, notes, readme	notebook, readme	missing python packages
[10]	1 / 1	notebook, readme	no documentation	—

Table 1: Summary of the Notebook Meta Data analysis for the five sample publications: ‘Mentions’ refers to the number of jupyter notebook mentions within the publication compared to the number of jupyter notebooks within the source code repository; ‘Req. Problems’ refers to problems with documented requirements.

contained descriptions of container images. The reason was that the container could not be build and no public accessible image was found. Although the reconstruction for the other notebooks was successful in most cases, the documentation often lacks in providing a complete list of python packages including their required version requiring more effort during the reconstruction (Table 1).

For those Jupyter notebooks where a computing environment could be reconstructed, we re-executed the notebook by using the build-in mechanism ‘Restart Kernel and Run all Cells...’. All cells run successfully for only five notebooks (see Figure 5). However, three notebooks failed as the data was not available (R5). The remaining notebooks failed due to an unset constant.

From the 22 notebooks only three notebooks could be successfully reproduced (R6) meaning that not only all cells run successfully but also the output is equal to the originally published version. One of the three notebooks, however, contained function definitions only.

4 Conclusion

In this paper we showed preliminary results of a systematic reproducibility analysis that indicate potential improvements when publishing research investigations in the form of Jupyter notebooks. Our findings are substantiated by the analysis of Chen et al. [11] who analyse reproducible research in the field of high-energy physics. While we strongly support the idea of Open Science, we think that research artefact publications must be cared as much as for the article itself in order to enable re-usage. Rule et al. [12] provide ten simple rules for reproducible research with Jupyter notebooks that is consistent with the problems we identified in this analysis and, thus, can help researchers in providing reproducible Jupyter notebooks.

Acknowledgement

This research was supported by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) within the Collaborative Research Centre 1270 ELAINE.

References

- [1] E. C. McKiernan, P. E. Bourne, C. T. Brown, S. Buck, A. Kennall, J. Lin, D. McDougall, B. A. Nosek, K. Ram, C. K. Soderberg, J. R. Spies, K. Thaney, A. Updegrave, K. H. Woo, and T. Yarkoni, “How open science helps researchers succeed,” *eLife*, vol. 5, jul 2016.
- [2] J. M. Perkel, “Why jupyter is data scientists’ computational notebook of choice,” *Nature*, vol. 563, no. 7729, pp. 145–146, oct 2018.
- [3] H. Shen, “Interactive notebooks: Sharing the code,” *Nature*, vol. 515, no. 7525, pp. 151–152, nov 2014.
- [4] S. Spors and F. Schultz, “Computational acoustics examples,” 2018. [Online]. Available: https://github.com/spatialaudio/computational_acoustics
- [5] R. J. Roberts, “Pubmed central: The genbank of the published literature,” *Proceedings of the National Academy of Sciences*, vol. 98, no. 2, pp. 381–382, 2001. [Online]. Available: <https://www.pnas.org/content/98/2/381>
- [6] A. Tambe, A. East-Seletsky, G. J. Knott, J. A. Doudna, and M. R. O’Connell, “RNA binding and HEPN-nuclease activation are decoupled in CRISPR-cas13a,” *Cell Reports*, vol. 24, no. 4, pp. 1025–1036, Jul. 2018.
- [7] B. Cummins, T. Gedeon, S. Harker, and K. Mischaikow, “DS-GRN: Examining the dynamics of families of logical models,” *Frontiers in Physiology*, vol. 9, may 2018.
- [8] J. Yang, H. Zhu, and X. Tian, “Group-level multivariate analysis in EasyEEG toolbox: Examining the temporal dynamics using topographic responses,” *Frontiers in Neuroscience*, vol. 12, jul 2018.
- [9] T. Collier and N. Manoukis, “Evaluation of predicted medfly (*ceratitis capitata*) quarantine length in the united states utilizing degree-day and agent-based models,” *F1000Research*, vol. 6, p. 1863, mar 2018.
- [10] K. H. Fisher-Wellman, M. T. Davidson, T. M. Narowski, C.-T. Lin, T. R. Koves, and D. M. Muoio, “Mitochondrial diagnostics: A multiplexed assay platform for comprehensive assessment of mitochondrial energy fluxes,” *Cell Reports*, vol. 24, no. 13, pp. 3593–3606.e10, sep 2018.
- [11] X. Chen, S. Dallmeier-Tiessen, R. Dasler, S. Feger, P. Fokianos, J. B. Gonzalez, H. Hirvonsalo, D. Kousidis, A. Lavasa, S. Mele, D. R. Rodriguez, T. Šimko, T. Smith, A. Trisovic, A. Trzcinska, I. Tsanaktsidis, M. Zimmermann, K. Cranmer, L. Heinrich, G. Watts, M. Hildreth, L. L. Iglesias, K. Lassila-Perini, and S. Neubert, “Open is not enough,” *Nature Physics*, vol. 15, no. 2, pp. 113–119, nov 2018.
- [12] A. Rule, A. Birmingham, C. Zuniga, I. Altintas, S. Huang, R. Knight, N. Moshiri, M. H. Nguyen, S. B. Rosenthal, F. Pérez, and P. W. Rose, “Ten simple rules for reproducible research in jupyter notebooks,” *CoRR*, vol. abs/1810.08055, 2018. [Online]. Available: <http://arxiv.org/abs/1810.08055>