




ELSEVIER

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/infosys

Editorial

A collaborative approach to computational reproducibility

Reproducing and verifying experiments contributes to science in several ways. First, reproducibility enables reviewers (and readers in general) to verify the outcomes presented in papers, which is crucial for science to be self-correcting. Second, it allows new methods to be objectively compared against methods presented in previous publications. Third, studies indicate that reproducibility increases impact, visibility, and research quality [1,4,9,11,13,15], and helps defeat self-deception [12]. Finally, computational reproducibility, by creating an executable artifact, enables researchers to build on top of previous work directly by simply extending the software.

Although a standard in natural science, reproducibility has been only episodically applied in experimental computer science. Scientific papers often present a large number of tables, plots and pictures that summarize the obtained results, but that loosely describe the steps taken to derive them. Not only can the methods and the implementation be complex, but also their configuration may require setting many parameters and/or depend on particular system configurations. As a consequence, reproducing the results from scratch is time-consuming, error-prone, and sometimes just infeasible. This has led to a credibility crisis in computational science [8]. Studies have shown that the fraction of papers in various computer science conferences that can be effectively reproduced is discouragingly low [3,6,10,15].

While many researchers recognize the importance of reproducibility, they are often held back by the challenge of making it happen. Authors must describe and encapsulate the entire experiment, which includes data, parameters, source code, dependencies and environment, so that the results can be properly verified and explored. If the experiment has not been systematically documented and made reproducible from the start, it may be hard to track all the necessary components to include in such compendium, and important aspects may be mistakenly left out. As an example, some numerical models, if not fully described, may lead to different implementations that are mathematically equivalent, but numerically different, rendering them irreproducible [7]. Even when the original researchers have tried to make their results reproducible, follow-on researchers may not be able to

reproduce the results for several reasons: insufficient documentation; the experiment may not run in their operating system; there may be missing libraries; library versions may be different; and the inability to install all the required dependencies.

The practical difficulties that reproducibility entails end up overshadowing its clear benefits to science. The process is often seen as a burden to both authors and reviewers. In fact, different surveys from various domains indicate that the effort and time required to make an experiment reproducible is one of the main reasons why authors do not do so [3,6,14].

Overall, to increase the practice of reproducibility in computational science, we have identified two main goals:

1. *Usability*: development of tools that make it easier and significantly less time-consuming for authors to do reproducible research, and for reviewers to execute computational artifacts (and modify them) corresponding to published results.
2. *Incentives*: a new publication model that recognizes the efforts of making experiments reproducible (for authors) and verifying published scientific results (for reviewers).

Fortunately, a plethora of reproducibility solutions that address the first goal have been recently designed and implemented by the community. In particular, packaging tools (e.g., *ReproZip*¹) and virtualization tools (e.g., *Docker*²) are promising solutions towards facilitating reproducibility for both authors and reviewers. For example, with *ReproZip* [5], authors can automatically track the dependencies of their existing experiments and create a self-contained package using only two commands. Reviewers can then use *ReproZip* to unpack, reproduce, and vary experiments using as few as three commands, even in a system that differs from the operating system of the authors' original environment. *Docker* [2] complements *ReproZip* in the sense that it is able to create a virtualization environment for the experiment that is lightweight and easy to use and deploy.

¹ <https://vida-nyu.github.io/reprozip/>.

² <https://www.docker.com/>.

To address the second goal, we have implemented a new publication model for the Reproducibility Section of Information Systems Journal. In this section, *authors submit a reproducibility paper* that explains in detail the computational assets from a previous published manuscript in Information Systems. Submission is by invitation only.

A reproducibility paper describes all the software and data used to derive the published results, as well as provides instructions on how to reproduce and validate such results. In addition, authors can use this paper to discuss the benefits and challenges they encountered in making their experiment reproducible. Using Mendeley Data,³ authors also submit their code, data, and optionally a ReproZip package or a Docker container to make the review process easier. Reviewers not only review the reproducibility paper, but also validate the results and claims published in the original manuscript. Once the paper is accepted, *reviewers also become co-authors* and are encouraged to add a section in the paper that states the extent to which the software is portable, is robust to changes, and is likely to be usable as a subcomponent or as a basis for comparison by future researchers. The review is not blinded, so authors and reviewers are encouraged to engage in a discussion about the validity of the experimental results as many times as necessary.

This model addresses our second goal in two ways. First, it incentivizes authors to make their experiment reproducible by creating a second publication from their effort. Second, it recognizes the difficulties of the review's job—which, even when using ReproZip or Docker, may be time-consuming—thus allowing reviewers to become co-authors of the same publication.

Our first reproducibility paper was published online in January 2016 [16]. It validates the results and claims presented in the authors' original manuscript [17]. The experimental setup in this paper is complex, involving many details and configuration parameters. Using the manuscript alone, it would not be possible to reproduce its results. The corresponding reproducibility paper, together with the published computational assets (which were derived using ReproZip and Docker), greatly increases the likelihood that readers will reproduce the results and reuse the approach in future research.

This model is also the latest effort in Elsevier's history of exploring the potential for reproducibility in scientific publications in computer science and other disciplines. Beginning with the Executable Paper Grand Challenge in 2010, Elsevier has piloted several models of reproducibility in a variety of disciplines with data-rich science. The model here, *Invited Reproducibility Papers*, is an article type that joins a new class of scientific publication, creating publishable and citable artifacts to accompany reported research experiments.

The benefits of the Reproducibility Paper are threefold: proven experimental reproducibility for researchers, a collaborative full-fledged academic publication for the original authors and reproducibility reviewers, and a canonical reference point for the vetted experimental components. Further, by using the existing infrastructure of Mendeley Data and GitHub, and user-friendly tools such as ReproZip

and Docker, reproducibility papers constitute a model that is easily adoptable by other scientific journals, making reproducibility a reality in computational experiments.

We believe this new publication model will improve the degree of reproducibility in all computational sciences, thus increasing the reliability and usefulness of scientific research.

References

- [1] C.G. Begley, L.M. Ellis, Drug development: raise standards for pre-clinical cancer research, *Nature* 483 (7391) (2012) 531–533.
- [2] C. Boettiger, An introduction to docker for reproducible research, *SIGOPS Oper. Syst. Rev.* 49 (1) (2015) 71–79.
- [3] P. Bonnet, S. Manegold, M. Björling, W. Cao, J. Gonzalez, J. Granados, N. Hall, S. Idreos, M. Ivanova, R. Johnson, D. Koop, T. Kraska, R. Müller, D. Olteanu, P. Papotti, C. Reilly, D. Tsirogiannis, C. Yu, J. Freire, D. Shasha, Repeatability and workability evaluation of SIGMOD 2011, *SIGMOD Rec.* 40 (2) (2011) 45–48.
- [4] T. Brody, Evaluating research impact through open access to scholarly communication (Ph.D. thesis), University of Southampton, 2006.
- [5] F. Chirigati, R. Rampin, D. Shasha, J. Freire, ReproZip: computational reproducibility with ease, in: Proceedings of the SIGMOD'16, Demo Session, 2016, <http://dx.doi.org/10.1145/2882903.2899401>.
- [6] C. Collberg, T. Proebsting, A.M. Warren, Repeatability and Benefaction in Computer Systems Research. Tech. Rep. TR 14-04, University of Arizona, 2015.
- [7] S.M. Crook, A.P. Davison, H.E. Plesser, Learning from the past: approaches for reproducibility in computational neuroscience, in: J. M. Bower (Ed.), 20 Years of Computational Neuroscience, Springer Series in Computational Neuroscience, vol. 9, Springer, New York, 2013, pp. 73–102.
- [8] D. Donoho, A. Maleki, I. Rahman, M. Shahram, V. Stodden, Reproducible research in computational harmonic analysis, *Comput. Sci. Eng.* 11 (1) (2009) 8–18.
- [9] S. Hitchcock, The Effect of Open Access and Downloads ('Hits') on Citation Impact: A Bibliography of Studies, 2009.
- [10] J. Kovacevic, How to encourage and publish reproducible research, in: IEEE International Conference on Acoustics, Speech and Signal Processing, 2007. ICASSP 2007, vol. 4, 2007, pp. IV-1273–IV-1276.
- [11] S. Lawrence, Free online availability substantially increases a paper's impact, *Nature* 411 (6837) (2001) 521.
- [12] R. Nuzzo, How scientists fool themselves, and how they can stop, *Nature* 526 (7572) (2015) 182–185.
- [13] H.A. Piwowar, R.S. Day, D.B. Fridsma, Sharing detailed research data is associated with increased citation rate, *PLoS One* 2 (3) (2007) e308.
- [14] C. Tenopir, S. Allard, K. Douglass, A.U. Aydinoglu, L. Wu, E. Read, M. Manoff, M. Frame, Data sharing by scientists: practices and perceptions, *PLoS One* 6 (6) (2011).
- [15] P. Vandewalle, J. Kovacevic, M. Vetterli, Reproducible research in signal processing, *IEEE Signal Process. Mag.* 26 (3) (2009) 37–47.
- [16] A. Wolke, M. Bichler, F. Chirigati, V. Steeves, Reproducible experiments on dynamic resource allocation in cloud data centers, *Inf. Syst.* (2016).
- [17] A. Wolke, B. Tsend-Ayush, C. Pfeiffer, M. Bichler, More than bin packing: dynamic resource allocation strategies in cloud data centers, *Inf. Syst.* 52 (2015) 83–95.

Reproducibility Editor, Information systems
Fernando Chirigati,
New York University, USA

Publisher
Rebecca Capone,
Elsevier Inc.

³ <https://data.mendeley.com/>.

Rémi Rampin,
New York University, USA

Juliana Freire,
New York University, USA

Editor-in-Chief, Information systems
Dennis Shasha
New York University, USA