ReproducedPapers.org

Openly Teaching and Structuring Machine Learning Reproducibility

Yildiz, Burak; Hung, Hayley; Krijthe, Jesse H.; Liem, Cynthia C.S.; Loog, Marco; Migut, Gosia; Oliehoek, Frans A.; Panichella, Annibale; Pawelczak, Przemysław; Picek, Stjepan

DOI

[10.1007/978-3-030-76423-4_1](https://doi.org/10.1007/978-3-030-76423-4_1)

Publication date

2021

Document Version

Final published version

Published in

Reproducible Research in Pattern Recognition - 3rd International Workshop, RRPR 2021, Revised Selected Papers

Citation (APA)

Yildiz, B., Hung, H., Krijthe, J. H., Liem, C. C. S., Loog, M., Migut, G., Oliehoek, F. A., Panichella, A., Pawelczak, P., Picek, S., de Weerd, M., & van Gemert, J. (2021). ReproducedPapers.org: Openly Teaching and Structuring Machine Learning Reproducibility. In B. Kerautret, M. Colom, A. Krähenbühl, D. Lopresti, P. Monasse, & H. Talbot (Eds.), *Reproducible Research in Pattern Recognition - 3rd International Workshop, RRPR 2021, Revised Selected Papers* (pp. 3-11). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 12636 LNCS). Springer. https://doi.org/10.1007/978-3-030-76423-4_1

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

ReproducedPapers.org: Openly Teaching and Structuring Machine Learning Reproducibility

Burak Yildiz^(✉) , Hayley Hung, Jesse H. Krijthe , Cynthia C. S. Liem ,
Marco Loog , Gosia Migut, Frans A. Oliehoek , Annibale Panichella ,
Przemysław Pawełczak , Stjepan Picek , Mathijs de Weerd ,
and Jan van Gemert

Delft University of Technology, Postbus 5, 2600 AA Delft, The Netherlands
b.yildiz@tudelft.nl

Abstract. We present [ReproducedPapers.org](https://reproducedpapers.org): an open online repository for teaching and structuring machine learning reproducibility. We evaluate doing a reproduction project among students and the added value of an online reproduction repository among AI researchers. We use anonymous self-assessment surveys and obtained 144 responses. Results suggest that students who do a reproduction project place more value on scientific reproductions and become more critical thinkers. Students and AI researchers agree that our online reproduction repository is valuable.

Keywords: Machine learning · Reproducibility · Online repository

1 Introduction

Reproducibility is a cornerstone of science: if an experiment is not reproducible, we should question its conclusions. Yet, machine learning papers are lacking reproductions [7, 12]. Possible reasons may include a misaligned incentive between reproducing results and the short-term measures of career success associated with more ‘wins’ [26] and publishing ‘novel’ work [15]. Nevertheless, high-impact can be achieved, for instance, when a reproduction fails spectacularly, e.g. [6, 8, 10, 11, 14, 16, 18, 19, 24]. Yet, these take colossal amounts of manual effort [1, 2, 9, 22] or massive resources [16, 23]. There are venues for publishing reproductions [3, 4, 25], which are typically peer-reviewed and thus uphold various selection standards to guarantee quality. We argue that this emphasis on quality is a hurdle for sharing light-weight reproductions. Important and useful examples of light-weight reproductions include partial results, small variants on the algorithm, hyperparameter sweeps, etc. Low-barrier options are indeed available in workshop challenges [13, 21] organized at conferences such as ICPR, NeurIPS, ICLR, or ICML. However, such avenues are hard to maintain on a long-term basis, as a workshop may or may not be organized. We argue that there is a need for a low-barrier and long-term venue for machine learning reproductions.

A complementary angle on low-barrier reproductions is to improve university student training. We should teach the next generation of machine learning practitioners the importance of the reproducibility of research work, as done in other computer science domains such as computer networking, where results reproduction is the means to learn new material [30]. Doing a reproduction project in a course aligns with several important learning objectives for machine learning students. Among others, students (1) should be able to read, critique, and explain a scientific paper; (2) implement a method; (3) run, evaluate, investigate, and extend existing research or code; and (4) write clearly and concisely about code and methods. A reproduction project also lets students experience differences between published results and an implementation, which stimulates a critical attitude and allows reflections on the scientific process.

Fig. 1. A screenshot of [ReproducedPapers.org](https://reproducedpapers.org). We allow multiple reproductions of the same original paper and investigations of several aspects, such as *Reproduced*, *Replicated*, *Hyperparameter check*, etc. Our online repository is user-centered: its sufficient if a user sees value in uploading some form of reproduction. Having such a repository is well-suited for students and adds structure to reproducibility in machine learning.

In this paper, we align the benefits of an online reproduction repository with those of teaching reproducibility. We introduce ReproducedPapers.org: an open, light-weight repository of reproduced papers which flexibly allows any sort of reproduction work, see Fig. 1. This repository benefits the research community while at the same time being well-equipped at accepting contributions from students. Although the standard of student reproductions might be lower than those required for peer reviewed reproductions, they can still give valuable insights such as clarifying which parts are difficult to implement or identifying the reproducibility level of elements. Such online reproductions are a low-threshold portfolio-building opportunity, which in turn may prove a valuable incentive to start doing more reproductions, as well as an opportunity to facilitate sharing reproductions that otherwise would not have been shared.

Our online repository shares traits with other light-weight, bottom-up, grass-roots community efforts such as *ArXiv* [5], *Open Review* [28], and *Papers with Code* [29]. Other efforts on facilitating reproducibility include software for reproducible and reusable experiments [20], open specification neural network diagrams [17], and a framework for automatic parsing of deep learning research paper to generate the implementation [27]. Similar to these approaches, in our work, we combine the traits from online repositories with those of tools facilitating reproducibility by providing an online repository that facilitates teaching as well as structuring reproducibility.

We make the following contributions. 1. We propose a new online reproduction repository; 2. We conduct a proof of concept with students from an MSc Deep Learning course to perform a reproduction project and populate the repository; 3. We evaluate the usefulness of the repository among AI researchers and the learning objectives among students by anonymous surveys.

2 The Online Reproduction Repository

We performed a proof of concept experiment with a reproducibility project for students of the MSc Deep Learning course taught by this paper’s last author at Delft University of Technology (TU Delft). We solicited relevant papers among university staff and ensured that (i) data is available, (ii) it is clear which table or figure to reproduce, and (iii) the computational demands are reasonable. Students were also allowed to themselves suggest a paper to reproduce. On their paper of choice, they worked in groups of 2 to 4, for 8 weeks, for approximately one-third of their studying time (i.e., about 13 h a week). For grading, students submitted a blog in PDF and also the URL of an online version of their blog to ReproducedPapers.org to populate the repository. For students who do not wish to share a blog with the world, we offer a private option, which is only visible to course administrators. The option to publicly blog about reproducing machine learning provides an simple opportunity for students to build an online portfolio while simultaneously incentivizing making reproductions.

Table 1. Different aspects of reproduction which are highlighted as badges (see Fig. 1).

Aspect	Description
• Replicated	A full implementation from scratch without using any pre-existing code
• Reproduced	Existing code was evaluated
• Hyperparams check	New evaluation of hyperparameter sensitivity
• New data	Evaluating new datasets to obtain similar results
• New algorithm variant	Evaluating a different variant
• New code variant	Rewrote/ported existing code to be more efficient/readable
• Ablation study	Additional ablation studies

We explicitly allow for light-weight reproduction efforts such as evaluating existing code, checking only certain parts of the paper, proposing minor variations, doing hyperparameter sweeps, etc. Our current options (aspects) are shown in Table 1, and we will add others as the need arises. Authors label their reproduction with the relevant aspects themselves.

We developed [ReproducedPapers.org](https://reproducedpapers.org) in-house as a simple web application. It is implemented by this paper’s first author, and its source code is available on GitHub¹. Registering is necessary only when adding reproductions. Currently, the repository has 90 registered users and hosts 24 unique papers and 57 paper reproductions. Most papers have multiple reproductions, and only five reproductions are marked as private. The top-3 most-used aspects are *Replicated* (32 times); *Reproduced* (29 times) and *Hyperparams check* (17 times). Figure 2 whose data is derived from self-reported blog posts by users shows both success and failure rates to be around 40%.

Fig. 2. Current [ReproducedPapers.org](https://reproducedpapers.org) statistics. (a) Reproduction success rates; (b) Number of reproductions per paper ID.

¹ <https://github.com/CVLab-TUdelft/reproduced-papers>.

Fig. 3. Responses to survey questions from students who contributed to [ReproducedPapers.org](https://reproducedpapers.org). Letting students themselves do a reproduction promotes a critical mindset (a and b), and teaches the value of scientific reproductions (c). In addition, the students considered it a positive experience (d, e). We conclude that these traits align with our learning objectives.

3 Survey Analysis

We evaluate student learning objectives and how AI researchers perceive our online reproduction repository by analyzing the results of small anonymous surveys for two groups: (i) students who recently added their reproduction to our repository and (ii) anybody identifying her/himself working in AI. The second group was invited to the survey through social media and emails. Both groups share the same questions, where the students have five additional questions to evaluate education. The survey data is available at [ReproducedPapers.org](https://reproducedpapers.org)².

We received a total of 144 responses: 43 from course students and 101 from third-party AI researchers all over the world. Of the latter, 87 identify themselves as a junior or senior researcher, and 14 as a student.

² <https://reproducedpapers.org/survey-data.zip>.

3.1 Evaluating Student Learning Objectives

The survey questions and results can be found in Fig. 3. We evaluate the following objectives.

Doing a Reproduction Project Increases Critical Thinking. Results in Fig. 3(a) show that doing a reproduction taught most students something new about the scientific process. Figure 3(b) suggests that students become more critical to published results.

Doing a Reproduction Project Makes Students Value Reproductions More. The results in Fig. 3(c) indicate that after doing a reproduction, a great majority of students place more value on scientific reproductions.

Students Find a Reproduction Project a Positive Experience. The results in Fig. 3(d,e) demonstrates that students valued the work and prefer to do a reproduction more often. Results suggest that having a reproducibility project teaches skills considered important by both student and teacher.

3.2 Evaluating the AI Researcher Survey Respondents

Figure 4 shows results for the third party AI researchers. We found the following.

The AI Researcher Survey Respondents Find Online Reproductions Valuable. Results in Fig. 4(a,d,g) show that students and, especially, researchers find an online reproduction valuable and useful. According to Fig. 4(i), there is no clear preference for doing a reproduction or writing a paper. Figure 4(e) suggests that the perceived value of reproduction by the community is smaller for researchers than for students.

The AI Researcher Survey Respondents Find an Online Reproduction Repository Valuable. Results in Fig. 4(b,c) show that students and researchers appreciate an online reproduction repository. Figure 4(f) shows that researchers are less likely than students to help contribute by doing reproductions.

The AI Researcher Survey Respondents See an Educational Role for Courses Where Students Do a Reproduction Project. Results in Fig. 4(h) show that researchers and students agree that reproduction projects should be used more often in courses.

Additionally, we make the following observations from Fig. 4:

- (i) When compared to students, the *researchers think the community values reproductions less (e) and want their own team to work on reproductions less (f)*. This may suggest an inverse relationship between perceived value and willingness to contribute. Yet, when comparing researchers against themselves, most think the community values reproductions, and most researchers would like to contribute.
- (ii) More researchers *want their work reproduced (g) than that they are willing to contribute (f)*. Can we place our hope on the students as future researchers, as they are much more willing to contribute?

Fig. 4. Responses to survey questions by 57 students and 87 self-identified AI researchers. The survey question is in the sub-caption. Researchers and students agree that: Reproductions are valuable (a, d, g), that an online repository adds value (b, c), and that more courses should use a reproduction project (h). Researchers differ from students in that researchers more strongly find a reproduction valuable (a), and would consult online reproductions more (d). Researchers think a reproduction is valued less by the community (e) and are less likely to contribute with reproductions (f). Students and researchers both do not agree among themselves if a new paper is more valuable than a reproduction (i), suggesting that the answer is ‘it depends’. We conclude that the respondents welcome an online repository for teaching and structuring reproducibility.

- (iii) There is a clear consensus that *reproductions are valuable* (a, d, g, i) but *some researchers feel that the community does not reward it enough* (e). Therefore, an important question is how we can change the perception of a low reward for doing reproductions, beyond repositories as reported on here.

4 Discussion and Conclusions

It should be clear that our results and corresponding analysis are rather preliminary. We are convinced, however, that they warrant low-barrier and long-term solutions accommodating research reproduction. Our [Reproduced Papers.org](https://ReproducedPapers.org) provides one such outlet. We hope that future analysis of the further accumulated survey data may sketch an even clearer picture. We hope others consider reproducing our effort.

The main conclusions that we draw at present are the following three. 1. Doing a reproduction course project aligns well with learning objectives, and students find it a positive experience. 2. A reproducibility project improves the perceived value of reproductions, and allowing students to blog online about their reproduction project offers an extra incentive to do a reproduction. 3. AI researcher survey respondents are positive about online reproductions and a reproduction repository.

We finally call on the community to add their reproductions to the website ReproducedPapers.org and deploy it in courses: may the next generation of machine learners be reproducers.

References

1. Anand, K., Wang, Z., Loog, M., van Gemert, J.: Black magic in deep learning: how human skill impacts network training. In: British Machine Vision Conference (BMVC) (2020)
2. Bonneel, N., Coeurjolly, D., Digne, J., Mellado, N.: Code replicability in computer graphics. *ACM Trans. Graph.* **39**(4) (2020)
3. Colom, M., Kerautret, B., Krähenbühl, A.: An overview of platforms for reproducible research and augmented publications. In: Kerautret, B., Colom, M., Lopresti, D., Monasse, P., Talbot, H. (eds.) *RRPR 2018*. LNCS, vol. 11455, pp. 25–39. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-23987-9_2
4. Colom, M., Kerautret, B., Limare, N., Monasse, P., Morel, J.M.: IPOL: a new journal for fully reproducible research; analysis of four years development. In: 2015 7th International Conference on New Technologies, Mobility and Security (NTMS), pp. 1–5. IEEE (2015)
5. Cornell University Library: arxiv. <https://arxiv.org>, September 2011. Accessed 20 June 2020
6. Dacrema, M.F., Cremonesi, P., Jannach, D.: Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In: *Proceedings of the 13th ACM Conference on Recommender Systems* (2019)
7. Drummond, C.: Replicability is not reproducibility: nor is it good science. In: *Evaluation Methods for Machine Learning Workshop at the 26th ICML* (2009)
8. Engstrom, L., et al.: Implementation matters in deep policy gradients: a case study on PPO and TRPO. arXiv preprint [arXiv:2005.12729](https://arxiv.org/abs/2005.12729) (2020)
9. Fursin, G., Moreau, T., Reddi, V.: ASPLOS 2020 artifact evaluation report. In: *Proceedings of the ASPLOS*, pp. vi–vii. ACM (2020)
10. Gorman, K., Bedrick, S.: We need to talk about standard splits. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2786–2791 (2019)

11. Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., Meger, D.: Deep reinforcement learning that matters. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
12. Hutson, M.: Artificial intelligence faces reproducibility crisis. *Science* **359**(6377), 725–726 (2018)
13. Kerautret, B., Colom, M., Lopresti, D., Monasse, P., Talbot, H. (eds.): RRPR 2018. LNCS, vol. 11455. Springer, Cham (2019). <https://doi.org/10.1007/978-3-030-23987-9>
14. Lin, J.: The neural hype and comparisons against weak baselines. *ACM SIGIR Forum* **52**(2), 40–51 (2019)
15. Lipton, Z.C., Steinhardt, J.: Research for practice: troubling trends in machine-learning scholarship. *Commun. ACM* **62**(6), 45–53 (2019)
16. Lucic, M., Kurach, K., Michalski, M., Gelly, S., Bousquet, O.: Are GANs created equal? a large-scale study. In: *Advances in Neural Information Processing Systems*, pp. 700–709 (2018)
17. Marshall, G., Freitas, A.: The Diagrammatic AI Language (DIAL): Version 0.1. arXiv preprint [arXiv:1812.11142](https://arxiv.org/abs/1812.11142) (2018)
18. Melis, G., Dyer, C., Blunsom, P.: On the state of the art of evaluation in neural language models. In: *International Conference on Learning Representations* (2018). <https://openreview.net/forum?id=ByJHuTgA->
19. Musgrave, K., Belongie, S., Lim, S.N.: A metric learning reality check. arXiv preprint [arXiv:2003.08505](https://arxiv.org/abs/2003.08505) (2020)
20. Paganini, M., Forde, J.Z.: dagger: A Python Framework for Reproducible Machine Learning Experiment Orchestration (2020)
21. Pineau, J., Sinha, K., Fried, G., Ke, R.N., Larochelle, H.: ICLR reproducibility challenge 2019. *ReScience C* **5**(2), 5 (2019)
22. Raff, E.: A step toward quantifying independently reproducible machine learning research. In: *NeurIPS*, pp. 5486–5496 (2019)
23. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do imagenet classifiers generalize to imagenet? In: *ICML*, pp. 5389–5400 (2019)
24. Riquelme, C., Tucker, G., Snoek, J.: Deep Bayesian bandits showdown: an empirical comparison of Bayesian deep networks for Thompson sampling. In: *International Conference on Learning Representations* (2018). <https://openreview.net/forum?id=SyYe6k-CW>
25. Rougier, N.P., Hinsen, K.: ReScience C: a journal for reproducible replications in computational science. In: Kerautret, B., Colom, M., Lopresti, D., Monasse, P., Talbot, H. (eds.) RRPR 2018. LNCS, vol. 11455, pp. 150–156. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-23987-9_14
26. Sculley, D., Snoek, J., Wiltschko, A., Rahimi, A.: Winner’s curse? On pace, progress, and empirical rigor. In: *ICLR Workshop* (2018). <https://openreview.net/forum?id=rJWF0Fywfwf>
27. Sethi, A., Sankaran, A., Panwar, N., Khare, S., Mani, S.: DLPaper2Code: auto-generation of code from deep learning research papers. In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018)
28. Soergel, D., Saunders, A., McCallum, A.: Open scholarship and peer review: a time for experimentation. In: *Proceedings of the ICML* (2013)
29. Stojnic, R., Taylor, R.: Papers with code—a Facebook AI project, July 2018. <https://paperswithcode.com>. Accessed 20 June 2020
30. Yan, L., McKeown, N.: Learning networking by reproducing research results. *SIGCOMM Comput. Commun. Rev.* **47**(2), 19–26 (2017)