Code for Research Papers (C4RP)

Keziah Naggita
Toyota Technological Institute at Chicago
knaggita@ttic.edu

Abstract

Mentorship not only helps students discover their hidden potential, push beyond boundaries, learn new skills and grow both professionally and personally but also helps mentors share their expertise, and grow as individuals. Mentorship, especially long-distance, is slowed down by several obstacles ranging from communication issues, loss of morale from both parties, laws and rules deterring researchers from collaborating with people outside the lab, among other things. Code for Research Papers (C4RP) aims at fostering research entry for early researchers from underserved populations through initiating collaboration between mentors both abroad and local, lecturers and students in some of the major universities in Africa. The students will reproduce an already published machine learning paper with the supervision of a mentor, and a lecturer. And then publish their findings in a selected conference, journal or workshop and be awarded credit at the end.

Introduction

There is an exponential growth in the number of Machine Learning (ML) research papers released annually on different platforms and avenues for example arXiv, IEEEXplore, OpenReview, Google Scholar, ResearchGate, ML conferences and workshops, journals among others [1]. Several ML researchers are concerned about the inability to reproduce [2, 3] most of these research papers [4–7]. Many deep learning sceptics are unable to reproduce deep learning models [8]. The hyperparameters used are usually presented like magic numbers without a detailed explanation of the desired choice. Failure to reproduce the metrics like accuracy someone else claimed to have gotten on a similar dataset with the same methodology is quite disturbing for lots of researchers. And this is usually attributed to factors like the choice of; computing resources, libraries, programming languages, GPU / CPU, and frameworks used. [1] Additionally, decisions like weights initialization, number of times the researcher ran the model on the train set and possibly on the test set, shuffling or not shuffling the datasets can also significantly affect the results, cherry-picking among others. Some researchers may or may not disclose this information, which is very crucial for Reproducibility in ML (RML). Even though this could be a learning experience for early researchers,[2] it's a time-consuming task for researchers aimed at making advancements, fairly compare existing models or improve upon existing research. Reproducing already existing research is a great opportunity for novice researchers and those looking to enter the field. This is mainly because it helps them get acquainted with the research area and fully understand the underlying concepts.

Several undergraduate students in developing countries, especially in Africa, are faced with the challenge of finding great ML research projects and mentors to supervise them conduct research in the different ML areas, for example, Computer Vision, Natural Language Processing, Learning Theory, Kernel Methods, Fairness and Ethics, Convex Optimization among others [9]. This is because the ratio of ML researchers in Africa to the enthusiastic prospective mentees is very big. And although lots of initiatives to form mentorship relationships have been created for example Deep Learning Indaba,[3], Data Science Africa,[4] and BlackinAI [5], the gap is still very big.

Even though mentorship has proven to improve the lives of both mentors and mentees [10–12], long-distance mentorship has proven very hard to conduct, and in most cases, neither effective nor efficient. Factors like; morale loss on either end, intermittent internet, timezone differences and cultural differences fail long-distance mentorships. And sometimes, the desire to mentor students outside the lab to collaborate on lab projects is constrained by lots of factors, for example, grant limitations, university and or company policies, remuneration challenges, among other things. This well-meaning task ends up in most cases being time-wasting

[1] https://determined.ai/blog/reproducibility-in-ml/

[2] https://sites.google.com/view/icml-reproducibility-workshop/home

[3] http://www.deeplearningindaba.com/

[4] http://www.datascienceafrica.org/

[5] https://blackinai.github.io/

and frustrating for both the mentors and mentees [13, 14].

Code for Research Papers (C4RP) aims at fostering research entry for undergraduate students in African universities. C4RP seeks to support lecturers in some of the major universities in Africa, for example, the University of Cape Town and Makerere University to introduce RML as a creditworthy course. Within the course, the students will reproduce an already published paper for credit. And to help the students learn as much as possible and finish the course in scheduled time, they will pick a research paper published at conference and or workshops and reproduce the published findings. Following will be required to write a reproducibility challenge paper expounding on what they completed, failed to achieve, challenges faced, lessons learnt, and recommendations. Finally, students will submit the report to a reproducibility challenges workshop at a major conference, for example, the International Conference on Machine Learning (ICML)[6] and the Association for the Advancement of Artificial Intelligence (AAAI)[7]. And in addition to the lecturers overlooking the course, mentors, will supervise, advise and guide the selected students in conducting the research.

Several mentorship programs face difficulty finding enthusiastic students and keep the passion burning for those already passionate despite all the unforeseen circumstances. C4RP solves this problem because it's a creditworthy course setting mentorship program, students will push on until they finish the project. The program will equip students with research, personal and professional skills, and helps them develop a deep understanding of the selected ML area. It will also ensure the mentors and lecturers efficiently and effectively share their expertise with students. The acquired skills will make it easier for the students to pursue graduate studies in ML, and also increase the output of research findings from Africa. Additionally, C4RP will enable advanced researchers to focus more on the task at hand instead of spending time to first reproducing the paper.

How it will work

The mentorship process will involve a sign-up from the university to participate in the program, a match of mentors to different students in different universities based on the research interests and availability of the involved parties, and finally a report from the lecturers and the mentors detailing the progress of students, and everyone's experience with the program; students, mentors and lecturers.

The mentorship team will ensure all participating universities can avail resources like space, internet and computers for students and lecturers to fully and effectively partake in the program. The duration of the mentorship program depends on the availability of the mentors and lecturers.

Since typically a reproducibility challenge commences when accepted papers are released, students will have at most four months to form teams, select, register and re-implement papers and submit the project report to the conference workshop. This mentorship duration is suitable for mentors and lecturers available for 4 - 6 months to guide and support the students. Students will work for a4 months.

The 6-12 months mentorship program is for students who will be working on big reproducibility projects to submit to journals, for example, proceedings of ML research[8], and ReScience C. [15] [9]

Conclusion

Code for Research Papers aims at providing both mentorship and research entry to undergraduate students in African universities through fostering a 4-12 months mentorship to students in universities without ML mentors to guide and support them to begin working on challenging and meaningful ML research projects. The students will reproduce an ML paper published at major ML conferences or workshops with the guide and support of mentors and a local lecturer. The creditworthy coursework setting of the mentorship program will motivate the students to push on a midst unforeseen failures and challenges, which is central to a successful mentorship program. The students, on producing accessible code and a write-up detailing their implementation process, successful and failed experiments, lessons learnt, recommendations and conclusions will receive academic credit. The lecturer and mentors will learn from the process, share their expertise, and together with the students reproduce a research paper of their interest, which can be a baseline for further research.

Acknowledgments

The author is grateful to the anonymous reviewers for their valuable feedback.

References

[1] Jeff Dean, David A. Patterson, and Cliff Young. A new golden age in computer architecture: Empowering the machine-learning revolution. *IEEE Micro*, 38:21–29, 2018.

[2] Lorena A. Barba. Terminologies for reproducible research. *ArXiv*, abs/1802.03311, 2018.

[6] https://openreview.net/group?id=ICLR.cc/2019/Workshop/RML

[7] https://folk.idi.ntnu.no/odderik/RAI-2019/

[8] http://proceedings.mlr.press/

[9] http://rescience.github.io/

[3] Roger D Peng. Reproducible Research in Computational Science. *Science*, 334(6060):1226–1227, 2011.

[4] André Anjos, Manuel Günther, Tiago de Freitas Pereira, Pavel Korshunov, Amir Mohammadi, and Sébastien Marcel. Continuously reproducing toolchains in pattern recognition and machine learning experiments. In *Thirty-fourth International Conference on Machine Learning*, August 2017. https://openreview.net/group?id=ICML.cc/2017/RML.

[5] Cheng Li, Abdul Dakkak, Jinjun Xiong, and Wen-Mei Hwu. Challenges and pitfalls of reproducing machine learning artifacts. *CoRR*, abs/1904.12437, 2019.

[6] Matthew B. A. McDermott, Shirly Wang, Nikki Marinsek, Rajesh Ranganath, Marzyeh Ghassemi, and Luca Foschini. Reproducibility in machine learning for health. *CoRR*, abs/1907.01463, 2019.

[7] Peter Sugimura and Florian Hartl. Building a reproducible machine learning pipeline. *CoRR*, abs/1810.04570, 2018.

[8] Gary Marcus. Deep learning: A critical appraisal. *CoRR*, abs/1801.00631, 2018.

[9] Lem Ngongalah, Wepngong Emerson, Ngwa Niba Rawlings, and James Muleme Musisi. Research challenges in Africa, an exploratory study on the experiences and opinions of African researchers. *bioRxiv*, 2018.

[10] Yasmin B. Kafai, Jean Griffin, Quinn Burke, Michelle Slattery, Deborah A. Fields, Rita Manco Powell, Michele Grab, Susan B. Davidson, and Joseph S. Sun. A cascading mentoring pedagogy in a cs service learning course to broaden participation and perceptions. In *SIGCSE*, 2013.

[11] Heather Pon-Barry, Becky Wai-Ling Packard, and Audrey St. John. Expanding capacity and promoting inclusion in introductory computer science: a focus on near-peer mentor preparation and code review. *Computer Science Education*, 27(1):54–77, 2017.

[12] Sylvia Pillon and William Osmun. Mentoring in a digital age. *Canadian family physician Médecin de famille canadien*, 59:442–4, 04 2013.

[13] Rajan Vaish, Snehalkumar (Neil) S. Gaikwad, Geza Kovacs, Andreas Veit, Ranjay Krishna, Imanol Arrieta Ibarra, Camelia Simoiu, Michael Wilber, Serge Belongie, Sharad Goel, James Davis, and Michael S. Bernstein. Crowd research: Open and scalable university laboratories. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, UIST '17, pages 829–843, New York, NY, USA, 2017. ACM.

[14] Sharon E. Straus, Mallory O Johnson, Christine Marquez, and Mitchell D. Feldman. Characteristics of successful and failed mentoring relationships: a qualitative study across two academic health centers. *Academic medicine : journal of the Association of American Medical Colleges*, 88 1:82–9, 2013.

[15] Nicolas P. Rougier and Konrad Hinsen. ReScience C: A Journal for Reproducible Replications in Computational Science. In Bertrand Kerautret, Miguel Colom, Daniel Lopresti, Pascal Monasse, and Hugues Talbot, editors, *Reproducible Research in Pattern Recognition*, volume 11455 of *Lecture Notes in Computer Science*, pages 150–156. Springer, June 2019.