

Contribution ID: 711 Contribution code: **contribution ID 711**Type: **Poster**

ROB: Benchmarking on the Cloud

The Reproducible Open Benchmarks for Data Analysis Platform (ROB)[1][2] is a platform developed to help evaluate data analysis workflows in a controlled competition-style environment. ROB was inspired by the Top Tagger Comparison analysis (2019)[3] that compared multiple different top tagger neural networks. ROB has two main goals: (1) reduce the amount of time required to organize and evaluate such benchmarks, and (2) ensure reproducibility of benchmark results. Towards the first goal, ROB provides a platform where the benchmark coordinator defines the benchmark workflow template and provides input data (e.g. training data) for the benchmarked task. The benchmark participants provide their implementations for individual steps in the benchmark workflow. ROB orchestrates the execution of the individual workflows and ranks their results to compare the performance of all benchmark submissions.

In this work, we address the second goal. ROB relies on existing workflow engines to execute the benchmark workflows. To achieve reproducibility of benchmark results we integrate ROB with the Reproducible research data analysis platform (REANA). REANA allows users to run their containerized workflows and archive the results on a remote server[4]. Using REANA as the execution backend for ROB makes it easier to archive a workflow run and reproduce the results since REANA archives all the required files and workflow logs for reproducing the results on the cloud. By using REANA, we also achieve a significant scalability improvement for ROB over the previously used native workflow engine. In future work, we are further planning to extend ROB to be able to run workflows on other services such as Google Cloud Platform and AWS.

References:

- [1] Scailfin. “Reproducible and Reusable Data Analysis Workflow Server.”GitHub, github.com/scailfin/flowserv-core.
- [2] Scailfin. “Web-based user interface for the Reproducible Open Benchmarks for Data Analysis Platform (ROB)”GitHub, <https://github.com/scailfin/rob-ui>
- [3] Kasiczka, G., et al. “The Machine Learning Landscape of Top Taggers.”ArXiv.org, 23 July 2019, arxiv.org/abs/1902.09914.
- [4] “REANA: Reproducible Analysis Platform.”Documentation, docs.reana.io/

Significance

By Integrating REANA, ROB can now archive the run and the benchmarks on the cloud. This benchmark can now easily be referenced and reproduced, which is one of the primary goals of ROB. By running workflows on the cloud, ROB is no longer limited to the server it is running on and can run multiple workflows and benchmark them parallelly.

References

<https://indi.to/ck4k2>

Speaker time zone

Compatible with America

Primary authors: RAWAT, Ajay (University of Washington (US)); HSU, Shih-Chieh (University of Washington Seattle (US)); MUELLER, Heiko

Presenter: RAWAT, Ajay (University of Washington (US))

Session Classification: Posters: Raspberry

Track Classification: Track 1: Computing Technology for Physics Research