



Contribution ID: 526 Contribution code: THU 31

Type: Poster

Towards more efficient job scheduling in ALICE: predicting job execution time using machine learning

Thursday 24 October 2024 16:00 (15 minutes)

The ALICE Grid processes up to one million computational jobs daily, leveraging approximately 200,000 CPU cores distributed across about 60 computing centers. Enhancing the prediction accuracy for job execution times could significantly optimize job scheduling, leading to better resource allocation and increased throughput of job execution. We present results of applying machine learning techniques to predicting the execution time of ALICE computational jobs. To this end, we focus on the following main challenges in this prediction task:

- (1) Feature extraction and selection: extracting the relevant features from the collected data and selecting the ones that are most important for model training and inference
- (2) Model selection: identifying an ML model that is accurate and robust for our prediction problem.
- (3) Model decay: making sure that the model accuracy does not deteriorate in time as new data arrives, possibly from an evolving data distribution.
- (4) Near-real-time processing: predictions need to be made in near-real-time.

Our goal is to develop a solution capable of predicting job execution times for batches of hundreds of elements in less than 100 milliseconds, without compromising accuracy or hindering continuous learning. This requires striking a delicate balance between computational complexity and real-time performance. By addressing these challenges within the ALICE CERN experiment framework, we can enhance job scheduling efficiency and optimize resource allocation, ultimately advancing scientific research in particle physics.

Acknowledgements. This work is co-financed in part supported by the Ministry of Science and Higher Education (Agreement Nr 2023/WK/07) and by the program of the Ministry of Science and Higher Education entitled PMW

Primary authors: Dr BALIS, Bartosz (AGH University of Krakow (PL)); GRIGORAS, Costin (CERN); Dr KURDZIEL, Marcin (AGH); Mr FACISZEWSKI, Michał (AGH); Mr ZASADA, Mikołaj (AGH); Ms ŚWIĘTEK, Sara (AGH); LELEK, Tomasz Marcin (AGH University of Krakow (PL)); THE ALICE COLLABORATION, on behalf of

Presenter: LELEK, Tomasz Marcin (AGH University of Krakow (PL))

Session Classification: Poster session

Track Classification: Track 4 - Distributed Computing