

Hyperparameter optimization, quantum-assisted model performance prediction, and benchmarking of AI-based High Energy Physics workloads using HPC

Eric Wulff, Maria Girone, David Southwick, Juan Pablo García Amboage and Eduard Cuba

CERN, Esplanade des Particules 1, 1211 Geneva 23, Switzerland

E-mail: eric.wulff@cern.ch

Abstract. Training and Hyperparameter Optimization (HPO) of deep learning-based AI models are often compute resource intensive and calls for the use of large-scale distributed resources as well as scalable and resource efficient hyperparameter search algorithms. This work studies the potential of using model performance prediction to aid the HPO process carried out on High Performance Computing systems. In addition, a quantum annealer is used to train the performance predictor and a method is proposed to overcome some of the problems derived from the current limitations in quantum systems as well as to increase the stability of solutions. This allows for achieving results on a quantum machine comparable to those obtained on a classical machine, showing how quantum computers could be integrated within classical ML tuning pipelines.

Furthermore, results are presented from the development of a containerized benchmark based on an AI-model for collision event reconstruction that allows us to compare and assess the suitability of different hardware accelerators for training deep neural networks.

1. Introduction

Current state-of-the-art hyperparameter (HP) search algorithms such as Hyperband [1], the Asynchronous Successive Halving Algorithm (ASHA) [2] and Bayesian Optimization Hyperband (BOHB) [3] rely on a method of early termination, where badly performing trials are automatically terminated to free up compute resources for new trials to be started.

Basing the stopping criterion on the relative ranking of trials according to a chosen metric, often accuracy or validation loss, can be problematic. Since the training process is non-linear, the ranking of trials at the decision point does not necessarily hold at the target point, see figure 1. A potential solution to this problem is to use a non-linear stopping criterion, e.g., using Support Vector Regression (SVR), to predict the final model performance from a partially trained model [4].

There are two main motivations for why to use quantum annealing (QA)-based SVRs over classical SVRs. Firstly, the training of so called Quantum-SVRs (QSVRs) is of time complexity $\mathcal{O}(N)$ while for classical SVRs it is $\mathcal{O}(N^3)$ [5], where N is the number of data points, i.e., the size of the training dataset. Secondly, the ensemble of solutions produced by the QA could potentially result in better performance compared to the single solution provided by classical SVRs. This has been empirically shown for several use cases in [6–8] and is explained by the

fact that the many different, but all relatively close to optimal, solutions given by the QA often generalize better to previously unseen data than the single solution given by a classical SVR [6].

Large-scale High Performance Computing (HPC) systems are especially suited to run Hyperparameter Optimization (HPO) algorithms such as ASHA due to the superlinear scaling that can be achieved, as demonstrated in tests performed on the JURECA-DC [9] system at the Jülich Supercomputer Centre (JSC), presented in figure 2.




Figure 1: Sketch of two learning curves resulting from different HP configurations.




Figure 2: Scaling of HPO workflow on the JURECA-DC-GPU system.

2. Quantum support vector regression for model performance prediction

The potential to speed up the HPO process via performance prediction as well as the use of a quantum annealer (QA) to train the performance predictor is investigated. The Quantum-SVR (QSVR) performance achieved is comparable to classical SVR, showing that quantum computers are good candidates to be integrated within classical ML tuning workflows in the future.

A Graph Neural Network (GNN)-based algorithm, developed for the task of machine learned particle flow (MLPF) reconstruction [10, 11] in High Energy Physics (HEP), acts as the base model for which studies are performed. A dataset consisting of learning curves and HP configurations was generated, see figure 3, by training 296 different configurations of MLPF on the publicly available Delphes dataset [12]. Configurations were drawn randomly from the HP space defined in table 1. The trainings were run in a distributed data-parallel mode on compute nodes with four NVIDIA A100 GPUs each. Training one model for 100 epochs on such a node required roughly 16 hours of wall time. To speed up the generation of learning curves, trainings were run in parallel on 24 compute nodes on the JURECA-DC-GPU system at JSC.

To establish a baseline, classical SVR was studied by fitting 1000 models on different train/test splits with the same constraint on dataset size as required by the QSVR. Due to the limited number of qubits on the QA, only 20 training samples could be used for fitting. Results vary depending on the training split and statistics are shown in the figure 4 and table 2.

Via the The European Center of Excellence in Exascale Computing "Research on AI- and Simulation-Based Engineering at Exascale" (CoE RAISE), access to JUPSI, a QA at JSC, was leveraged to train a QSVR model for model performance prediction. Due to the probabilistic nature of quantum processes the annealing is run multiple times and returns multiple solutions which can then be combined in different ways to create the final QSVR model. Different methods to combine the QSVR solutions are described in [13]. The predictions of the best performing QSVR are plotted against the true values in figure 5.

Hyperparameter	Sample distributions
Learning rate	$\log lr \sim U[10^{-6}, 3 \cdot 10^{-2}]$
Classification graph layers	{0, 1, 2, 3, 4}
Regression graph layers	{0, 1, 2, 3, 4}
dropout	$U[0.0, 0.5]$
bin size	{8, 16, 32, 64, 128}
Output dimensions	{8, 16, 32, 64, 128, 256}
Weight decay	$\log wd \sim U[10^{-6}, 10^{-1}]$

Table 1: Hyperparameter space used to create the learning curve dataset.




Figure 3: The best learning curves from the learning curve dataset, zoomed in on epochs 25 to 100.




Figure 4: Resulting R^2 distribution after training and testing an SVR on 1000 different splits on a classical computer.




Figure 5: Predicted vs true validation loss values for the single best performing QSVR model.

With the aim of stabilizing QSVR performance, different techniques were tested by combining several QSVRs trained on 20 samples each. The most successful approach found was to split the training set into disjoint subsets of 20 points and train a QSVR for each subset. The final ensemble prediction was calculated as the average of all the individual predictions.

To evaluate the QSVR combination technique, 80 training and 150 testing points were used. This experiment was repeated for 10 random train/test splits, where the results are presented in Figure 6. Note that each prediction is made by combining 4 QSVRs trained on 20 points each. Although this approach did not improve the maximum R^2 score, it did produce more

stable results by significantly improving the worst performing split and reducing the standard deviation of R^2 scores between splits, as can be seen in table 2. This is interesting as one of the problems faced while using the Q-SVR has been the instability of the results. Further improvement could be made by modifying the weights in the weighted average that combine the QSVRs, as currently a simple average is used. A comparison between the statistics of the R^2 scores for different models described in this work are presented in table 2.




Figure 6: Best QSVR ensemble model predictions vs true values after running the QSVR combination technique 10 times for 10 different train/test splits using 80 training samples.

	Best	Worst	Mean	Std	Number of trainings
SVR	0.959	0.318	0.889	0.050	1000
Sim-QSVR	0.949	0.383	0.901	0.045	100
QSVR	0.948	0.742	0.880	0.056	10
QSVR Ensemble	0.927	0.857	0.899	0.019	10

Table 2: R^2 values and statistics for the different regression models. The SVR model is a classical SVR trained on a classical computer. The Sim-QSVR model is a QSVR model trained on a classical computer using simulated quantum annealing. The QSVR model is a QSVR trained on the QA and the QSVR Ensemble is the combined predictions from 4 QSVRs trained on the QA using 20 training samples each.

3. Containerized AI Benchmark

The promise of better accuracy and reconstruction scalability at inference time for AI-based algorithms is preceded by resource-intensive training. A self-contained, reproducible and containerized benchmark based on the training of deep learning models is proposed to explore the feasibility of deploying AI-driven HEP applications to different HPC environments. The metric shown in figure 7 is the training throughput, or training samples processed per second, on a subset of the publicly available Delphes dataset [12] in HEPscore [14] format.

4. Conclusions

HPC systems are essential for running large-scale HPO and distributed training and can significantly increase model performance as well as speed up the iteration of model development and training. It has been shown that superlinear scaling in speed-up can be achieved for HPO workflows, indicating the benefits of HPC for such use cases.

The strong potential of using performance prediction for HPO was demonstrated, encouraging the use of this technique in future HPO studies. It was also shown that, despite the current




Figure 7: Benchmark results.

limitations of quantum computers, it is possible to train QSVR models on a QA while achieving prediction performance comparable to those obtained on a classical SVR. This encourages further studies in utilizing hybrid quantum/HPC workflows for HPO as well as in other use cases.

Finally, the development of a containerized benchmark application with an AI use case from HEP allows for quick and easy benchmarking of new hardware accelerators in the HEP Score format.

Acknowledgments

We thank our colleagues in CoE RAISE, in particular Andreas Lintermann and Marcel Aach for helpful discussions and feedback in the course of this work. We also thank our colleagues in the CMS Collaboration, especially Joosep Pata, Javier Duarte and Farouk Mokhtar for their collaboration on the MLPF studies.

Eric Wulff, David Southwick and Eduard Cuba was supported by CoE RAISE. The CoE RAISE project has received funding from the European Union’s Horizon 2020 – Research and Innovation Framework Programme H2020-INFRAEDI-2019-1 under grant agreement no. 951733.

The authors gratefully acknowledge the computing time granted through JARA on the supercomputer JURECA [9] at Forschungszentrum Jülich. The authors gratefully acknowledge the Jülich Supercomputing Centre for funding this project by providing computing time through the Jülich UNified Infrastructure for Quantum computing (JUNIQ) on the D-Wave Advantage™ System JUPSI.

References

- [1] Li L, Jamieson K, DeSalvo G, Rostamizadeh A and Talwalkar A 2016 Hyperband: A novel bandit-based approach to hyperparameter optimization URL <https://arxiv.org/abs/1603.06560>
- [2] Li L, Jamieson K G, Rostamizadeh A, Gonina E, Hardt M, Recht B and Talwalkar A 2018 Massively parallel hyperparameter tuning *CoRR* **abs/1810.05934** (*Preprint* 1810.05934)
- [3] Falkner S, Klein A and Hutter F 2018 BOHB: robust and efficient hyperparameter optimization at scale *CoRR* **abs/1807.01774** (*Preprint* 1807.01774)
- [4] Baker B, Gupta O, Raskar R and Naik N 2017 Accelerating neural architecture search using performance prediction URL <https://arxiv.org/abs/1705.10823>
- [5] Date P, Arthur D and Pusey-Nazzaro L 2021 QUBO formulations for training machine learning models *Sci Rep* **11**, 10029 URL <https://doi.org/10.1038/s41598-021-89461-4>
- [6] Willsch D, Willsch M, De Raedt H and Michielsen K 2020 Support vector machines on the d-wave quantum annealer *Computer Physics Communications* **248** 107006 ISSN 0010-4655 URL <https://www.sciencedirect.com/science/article/pii/S001046551930342X>
- [7] Li R, He Y, Zhang S, Qin J and Wang J 2018 Cell membrane-based nanoparticles: a new biomimetic platform for tumor diagnosis and treatment *Acta Pharmaceutica Sinica B* **8** 14–22 ISSN 2211-3835 URL <https://www.sciencedirect.com/science/article/pii/S2211383517303283>

- [8] Mott A, Job J, Vlimant J R, Lidar D and Spiropulu M 2017 Solving a higgs optimization problem with quantum annealing for machine learning *Nature* **550** 375–379
- [9] Jülich Supercomputing Centre 2021 JURECA: Data Centric and Booster Modules implementing the Modular Supercomputing Architecture at Jülich Supercomputing Centre *Journal of large-scale research facilities* **7** URL <http://dx.doi.org/10.17815/jlsrf-7-182>
- [10] Pata J, Duarte J, Vlimant J, Pierini M and Spiropulu M 2021 MLPF: efficient machine-learned particle-flow reconstruction using graph neural networks *The European Physical Journal C* **81** URL <https://doi.org/10.1140%2Fepjc%2Fs10052-021-09158-w>
- [11] Pata J, Duarte J, Mokhtar F, Wulff E, Yoo J, Vlimant J R, Pierini M, Girone M and on behalf of the CMS Collaboration 2023 Machine learning for particle flow reconstruction at cms *Journal of Physics: Conference Series* **2438** 012100 URL <https://dx.doi.org/10.1088/1742-6596/2438/1/012100>
- [12] Pata J *et al.* 2021 Simulated particle-level events of $t\bar{t}$ and QCD with PU200 using PYTHIA8+DELPHES3 for machine learned particle flow (MLPF) URL <https://zenodo.org/record/4559324>
- [13] Pasetto E, Riedel M, Melgani F, Michielsen K and Cavallaro G 2022 Quantum svr for chlorophyll concentration estimation in water with remote sensing *IEEE Geoscience and Remote Sensing Letters* **19** 1–5
- [14] Giordano D *et al.* 2021 HEPiX Benchmarking Solution for WLCG Computing Resources *Comput. Softw. Big Sci.* **5** 28