

Hyperparameter optimization of data-driven AI models on HPC systems

Eric Wulff¹, Maria Girone¹ and Joosep Pata²

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

²NICPB, Rävåla pst 10, 10143 Tallinn, Estonia

E-mail: eric.wulff@cern.ch

Abstract. In the European Center of Excellence in Exascale Computing "Research on AI- and Simulation-Based Engineering at Exascale" (CoE RAISE), researchers develop novel, scalable AI technologies towards Exascale. This work exercises High Performance Computing resources to perform large-scale hyperparameter optimization using distributed training on multiple compute nodes. This is part of RAISE's work on data-driven use cases which leverages AI- and HPC cross-methods developed within the project. In response to the demand for parallelizable and resource efficient hyperparameter optimization methods, advanced hyperparameter search algorithms are benchmarked and compared. The evaluated algorithms, including Random Search, Hyperband and ASHA, are tested and compared in terms of both accuracy and accuracy per compute resources spent. As an example use case, a graph neural network model known as MLPF, developed for Machine Learned Particle-Flow reconstruction, acts as the base model for optimization. Results show that hyperparameter optimization significantly increased the performance of MLPF and that this would not have been possible without access to large-scale High Performance Computing resources. It is also shown that, in the case of MLPF, the ASHA algorithm in combination with Bayesian optimization gives the largest performance increase per compute resources spent out of the investigated algorithms.

1. Introduction

One of the primary goals in the European Center of Excellence in Exascale Computing "Research on AI- and Simulation-Based Engineering at Exascale" (CoE RAISE) [1] is the development and expansion of Artificial Intelligence (AI) and High-Performance Computing (HPC) methods along representative use cases from research and industry. While Work Package 3 (WP3) "Compute-Driven Use-Cases at Exascale" covers use cases that are compute-driven, WP4 "Data-Driven Use-Cases at Exascale" has a strong focus on data-driven technologies, i.e., analyzing data-rich descriptions of physical phenomena. Example use cases vary widely and range from fundamental physics and remote sensing to 3D printing and acoustics.

The work of WP4 is highly integrated with WP2 "AI- and HPC-Cross Methods at Exascale". Experts in WP2 provide support on HPC and AI methods to use cases in WP4. This support manifests itself in porting code to new HPC architectures and machines, in performance analyses and engineering of codes, and in the development of AI solutions for the individual use cases.

In the work presented here, HPC resources are leveraged to perform large-scale hyperparameter optimization (HPO) using distributed training on multiple nodes as part of WP4. As an example use case from the field of High Energy Physics (HEP), the AI-based

particle-flow reconstruction algorithm called Machine-Learned Particle-Flow (MLPF) [2] acts as the base model for which HPO is performed. MLPF is developed in the Compact-Muon-Solenoid (CMS) Collaboration [3] at CERN and combines information from tracks and calorimeter clusters to reconstruct particle candidates.

Further developments of AI in RAISE have the potential to greatly impact the field of High Energy Physics by efficiently processing the large amounts of data that will be produced by particle detectors in the coming decades. Moreover, HPO is model agnostic and could be widely applied in other sciences using AI, e.g., in the fields of seismic imaging, remote sensing, defect-free additive manufacturing and sound engineering that are part of Work Package 4.

HPO, sometimes referred to as hyperparameter tuning or hypertuning, is the process of tuning the hyperparameters of the model to optimize its performance. Hyperparameters are the parameters that are not learned during the model training but must be defined by the user. Some examples are model-architecture-related parameters such as the number of layers in the model and the number of nodes in each layer, or optimization-related parameters such as the batch size and the learning rate.

Hypertuning deep learning-based AI models is often compute resource intensive, partly due to the high cost of training a single hyperparameter configuration to completion and partly because of the infinite set of possible hyperparameter combinations to evaluate. There is therefore a need for large-scale, parallelizable and resource efficient hyperparameter search algorithms.

This work makes use of a distributed computing tool called Ray [4], and more specifically the part of Ray called Tune [5]. Tune is an open-source tool for multi-node distributed hypertuning which integrates well with modern machine learning frameworks like e.g., TensorFlow [6] and PyTorch [7]. It also supports integration with many other hypertuning tools such as Scikit-Optimize [8], HyperOpt [9], Optuna [10], SigOpt [11], and more.

In the following, section 2 describes the example use case for which HPO is performed, section 3 describes how HPO can be used in a wide variety of applications and highlight synergies within CoE RAISE, and finally, section 4 presents the conclusions.

2. Example use case: Event reconstruction and classification at the CERN HL-LHC

With the upcoming upgrade of the Large Hadron Collider (LHC) to the High Luminosity LHC (HL-LHC), the HEP community will face a significant increase in data production. This motivates efforts to optimize the speed and efficiency with which data is collected, processed, and analyzed, and is one of the major challenges that must be solved by the time the HL-LHC starts operation at the end of 2027.

One of the many different approaches that are being investigated to tackle this challenge is to replace traditional HEP algorithms with faster, parallelizable AI-driven approaches. These approaches promise to deliver similar or even better physics performance and can relatively easily be accelerated by hardware such as Graphics Processing Units (GPUs) or Field Programmable Gate Arrays (FPGAs).

One such traditional algorithm that could potentially be replaced by an AI-based version is the so-called Particle-Flow (PF) reconstruction algorithm [12]. It processes signals from different sub-detectors and combines them to construct higher-level physics objects. These objects are used for downstream workflows and are important for physics analyses involving hadronic jets and missing transverse energy. An effort to construct a machine-learned PF algorithm is the so-called MLPF algorithm, which is based on a deep neural network implemented using a Graph Neural Network (GNN) formalism. A detailed description of its first iteration can be found in [2] while a more recent version is described in [13]. The code to build, train, and evaluate the model is publicly available online [14].

The best performing MLPF hyperparameters were found after two stages of hypertuning.

The first stage was performed on the Jülich Wizard for European Leadership Science (JUWELS) Booster [15] at the Jülich Supercomputer Centre in Jülich, Germany, and required 19,574 core-hours to complete. Each compute node on the JUWELS Booster has two AMD EPYC Rome 7402 CPUs with 48 cores clocked at 2.8 GHz and four NVIDIA A100-SXM4-40GB GPUs. The so-called Bayesian Optimization Hyperband (BOHB) [16] algorithm was used to tune parameters of the optimizer such as the `lr` and the learning rate schedule as well as the `dropout` and other model-specific internal hyperparameters. The BOHB search space is summarized in table 1.

The second hypertuning stage was performed on CoreSite at the Flatiron Institute in New York, NY, USA, using twelve compute nodes, each equipped with a 64-core Intel Icelake Platinum 8358 CPU clocked at 2.6 GHz and four NVIDIA A100-SXM4-40GB GPUs. The best hyperparameter values found from the first search were fixed and stage two instead tuned various architecture parameters such as the number of graph layers and the number of graphs in each layer, as well as the number of nodes and layers used for decoding, and a few other model-specific parameters. The search space of the second stage is summarized in table 2. In addition, a different hypertuning algorithm called Asynchronous Successive Halving Algorithm (ASHA) [17] was used in combination with Bayesian optimization. The ASHA algorithm allows for an efficient use of compute resources when performing distributed multi-node hypertuning by early stopping trials that underperform relative to others. The second stage of hypertuning consumed approximately 56,730 core-hours. This work would not have been possible without access to HPC resources, as can be illustrated by a back-of-the-envelope calculation to compute that the two hypertuning stages would have taken roughly 6 months to complete using a single GPU, compared to about 83 hours using supercomputers.

Table 1: Search space used in the hypertuning run using the BOHB algorithm.

Hyperparameter	Search space
<code>lr</code>	$\log lr \sim U(10^{-4}, 3 \cdot 10^{-2})$
<code>dropout</code>	(0, 0.5)
<code>clip_value_low</code>	(0, 0.2)
<code>dist_mult</code>	(0.01, 0.2)

Table 2: Search space used for ASHA in combination with Bayesian optimization.

Hyperparameter	Search space
<code>bin_size</code>	{16, 32, 40, 64, 80}
<code>distance_dim</code>	{32, 64, 128, 256}
<code>ffn_dist_hidden_dim</code>	{32, 64, 128, 256}
<code>ffn_dist_num_layers</code>	{1, 2, 3, 4}
<code>num_graph_layers_common</code>	{1, 2, 3, 4}
<code>num_graph_layers_energy</code>	{1, 2, 3, 4}
<code>num_node_messages</code>	{1, 2, 3, 4}
<code>output_dim</code>	{32, 64, 128, 256}

In both stages described above, the search algorithms were allowed to draw 200 samples from the hyperparameter search space. The best hyperparameters found according to validation loss after both stages of hypertuning are reported in table 3 and various metrics as a function of the training epoch are shown in figure 1.

To see if HPO improved the model performance, the loss and classification accuracy of the model before and after hypertuning is plotted and compared as a function of the training epoch in figure 2. Comparing these curves shows that the mean validation loss decreased by almost a factor of two (approximately 44%) and that the accuracy increased by more than the uncertainty. It is also clear from comparison of figures 2a and 2b as well as of figures 2c and 2d that the training became more stable as a result of hypertuning since the curves exhibit much less volatility after hypertuning, especially in the second half of the training.

3. Distributed training and hypertuning: synergies across sciences in RAISE

HPO algorithms are model-agnostic in their nature and could be applied in any field of science making use of AI. Hence, the benchmarking of HPO algorithms is of interest for all use cases

Table 3: Best hyperparameters found.

Hyperparameter	Value
lr	0.001129
dropout	0.016312
clip_value_low	0.001998
dist_mult	0.120898
bin_size	64
distance_dim	64
ffn_dist_hidden_dim	128
ffn_dist_num_layers	3
num_graph_layers_common	3
num_graph_layers_energy	2
num_node_messages	3
output_dim	64


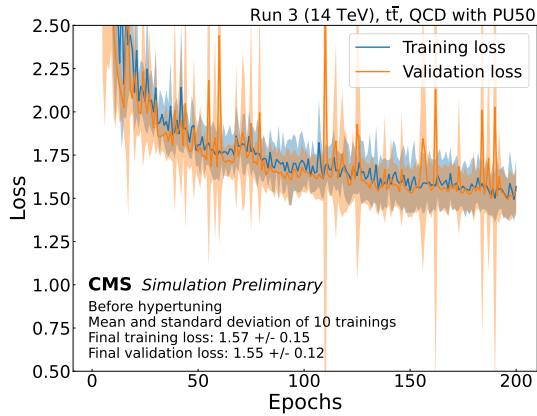
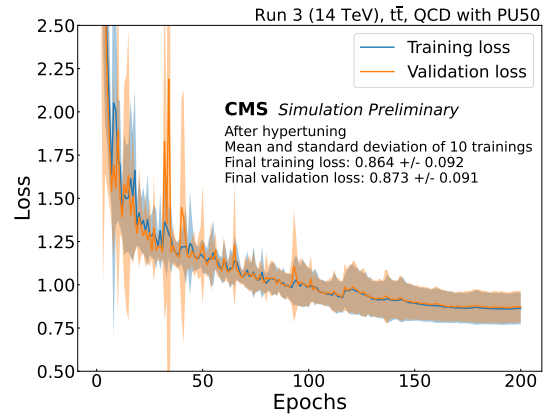


Figure 1: Loss and accuracy curves for top performing trials according to validation loss after hyper-tuning using the ASHA algorithm in combination with Bayesian optimization drawing 200 samples from the search space. From top to bottom: validation loss, validation classification loss, validation regression loss and validation classification weighted accuracy. The trials were trained for up to 200 epochs but the plots zoom in on epochs 70 and onward for better visibility.


in CoE RAISE WP4. In light of this, the hypertuning of MLPF was used as an example workflow to benchmark HPO algorithms in Ray Tune by running a variety of them using four compute nodes. The number of samples drawn were varied and results were analyzed in terms of samples drawn, compute hours spent, best achieved validation loss and best improvement per compute resources spent. The results are presented in figures 3 and 4. Figure 3 shows that both Hyperband and ASHA significantly outperforms random search in terms of core-hours spent per sample. Comparing the runs using ASHA, the combination with Bayesian optimization adds some overhead compared to the combination with random search making the ASHA + random search combination perform best from this perspective. Figure 4 gives another point of view, where the validation loss is plotted against the core-hours spent. From this viewpoint, it is clear that ASHA + Bayesian optimization gives the highest improvement per spent core-hour.



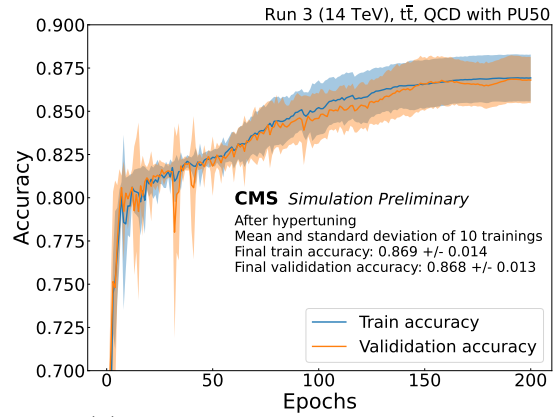
(a) Loss curves before hypertuning.



(b) Loss curves after hypertuning



(c) Accuracy curves before hypertuning.



(d) Accuracy curves after hypertuning.

Figure 2: Mean and standard deviation of loss and classification accuracy as a function of the training epoch computed from 10 trainings.

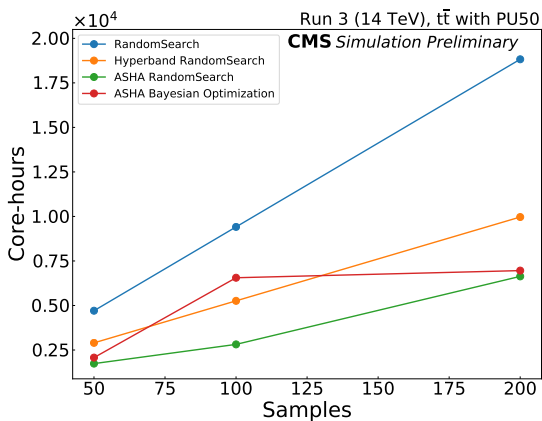


Figure 3: Comparison of hyperparameter search algorithms. Number of core-hours spent versus the number of samples drawn.

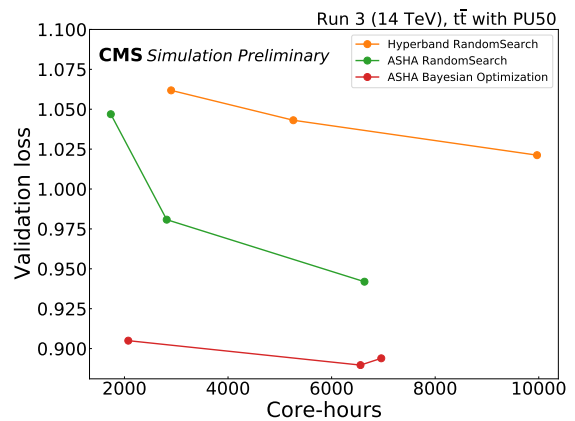


Figure 4: Comparison of hyperparameter search algorithms. Validation loss of the best trial versus number of core-hours spent.

4. Conclusion

CoE RAISE develops novel, scalable AI methods towards Exascale with use cases from a wide range of sciences and industry. HPO could benefit any data-driven AI-based algorithm and in the example use case of MLPF, large-scale distributed hypertuning significantly increased model performance. This would not have been possible without access to HPC resources since it would have taken approximately half a year of continuous hypertuning on a single GPU. Other sciences and use cases in CoE RAISE are also adopting HPC for hypertuning, including the use cases of WP4, within fields such as remote sensing, seismic imaging, defect-free additive manufacturing and sound engineering.

Acknowledgments

We thank our colleagues in CoE RAISE, in particular Andreas Lintermann, Morris Riedel, Marcel Aach, Eric Michael Sumner, Eray Inanc, Michael Bresser, Jennifer Lopez Barrilao, Ieva Timrote and Christina Bolanou for helpful discussions and feedback in the course of this work. We also thank our colleagues in the CMS Collaboration, especially Javier Duarte, Farouk Mokhtar, Jieun Yoo, Jean-Roch Vlimant and Maurizio Pierini for their contributions to MLPF.

Eric Wulff was supported by CoE RAISE and Joosep Pata was supported by the Mobilitas Pluss Grant No. MOBTP187 of the Estonian Research Council. 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.

References

- [1] The CoE RAISE project 2022 The CoE RAISE Website URL <https://www.coe-raise.eu>
- [2] Pata J, Duarte J, Vlimant J R, Pierini M and Spiropulu M 2021 MLPF: Efficient machine-learned particle-flow reconstruction using graph neural networks *Eur. Phys. J. C* **81** 381 (*Preprint* 2101.08578)
- [3] The CMS Collaboration *et al.* 2008 The CMS experiment at the CERN LHC *J. Instrum.* **3** S08004–S08004
- [4] Moritz P, Nishihara R, Wang S, Tumanov A, Liaw R, Liang E, Paul W, Jordan M I and Stoica I 2017 Ray: A distributed framework for emerging AI applications *CoRR* **abs/1712.05889** (*Preprint* 1712.05889)
- [5] Liaw R, Liang E, Nishihara R, Moritz P, Gonzalez J E and Stoica I 2018 Tune: A research platform for distributed model selection and training *CoRR* **abs/1807.05118** (*Preprint* 1807.05118)
- [6] Abadi M *et al.* 2015 TensorFlow: Large-scale machine learning on heterogeneous systems software available from tensorflow.org URL <https://www.tensorflow.org/>
- [7] Paszke A *et al.* 2019 Pytorch: An imperative style, high-performance deep learning library *Advances in Neural Information Processing Systems 32* ed Wallach H, Larochelle H, Beygelzimer A, d’Alché-Buc F, Fox E and Garnett R (Curran Associates, Inc.) pp 8024–8035
- [8] Head T *et al.* 2019 scikit-optimize <https://github.com/scikit-optimize/scikit-optimize>
- [9] Bergstra J, Yamins D and Cox D 2013 Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures *Proceedings of the 30th International Conference on Machine Learning (PMLR vol 28)* ed Dasgupta S and McAllester D (Atlanta, Georgia, USA: PMLR) pp 115–123
- [10] Akiba T, Sano S, Yanase T, Ohta T and Koyama M 2019 Optuna: A next-generation hyperparameter optimization framework *CoRR* **abs/1907.10902** (*Preprint* 1907.10902)
- [11] Clark S and Hayes P 2019 SigOpt Web page <https://sigopt.com> URL <https://sigopt.com>
- [12] Sirunyan A M *et al.* 2017 Particle-flow reconstruction and global event description with the CMS detector *J. Instrum.* **12** P10003–P10003 (*Preprint* 1706.04965)
- [13] Pata J for the CMS Collaboration 2022 Machine learning for particle flow reconstruction at CMS *J. Phys.: Conf. Series* **These proceedings**
- [14] Pata J, Wulff E, Mokhtar F, Duarte J and Tepper A 2021 jpata/particleflow: Baseline MLPF model for CMS DOI 10.5281/zenodo.5520559 URL <https://github.com/jpata/particleflow>
- [15] Jülich Supercomputing Centre 2021 JUWELS Cluster and Booster: Exascale Pathfinder with Modular Supercomputing Architecture at Juelich Supercomputing Centre *JLSRF* **7**
- [16] Falkner S, Klein A and Hutter F 2018 BOHB: robust and efficient hyperparameter optimization at scale *CoRR* **abs/1807.01774** (*Preprint* 1807.01774)
- [17] 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)