HIGH-PRECISION REGRESSORS FOR PARTICLE PHYSICS

Anonymous authors
Paper under double-blind review

ABSTRACT

Monte Carlo simulations of physics processes at particle colliders like the Large Hadron Collider at CERN take up a major fraction of the computational budget. For some simulations, a single data point takes seconds, minutes, or even hours to compute from first principles. Since the necessary number of data points per simulation is on the order of 10^9 – 10^{12}, machine learning regressors can be used in place of physics simulators to significantly reduce this computational burden. However, this task requires high precision regressors that can deliver data with relative errors less than 1% or even 0.1% over the entire domain of the function. In this paper, we develop optimal training strategies and tune various machine learning regressors to satisfy the high-precision requirement. We leverage symmetry arguments from particle physics to optimize the performance of the regressors. Inspired by ResNets, we design a Deep Neural Network with skip connections that outperform fully connected Deep Neural Networks. We find that at lower dimensions, boosted decision trees far outperform neural networks while at higher dimensions neural networks perform better. We show that these regressors can speed up simulations by a factor of 10^3 – 10^6 over the first-principles computations currently used in Monte Carlo simulations. Additionally, using symmetry arguments derived from particle physics, we reduce the number of regressors necessary for each simulation by an order of magnitude. Our work can significantly reduce the training and storage burden of Monte Carlo simulations at current and future collider experiments.

1 INTRODUCTION

Particle physics experiments like those at the Large Hadron Collider at CERN, are running at progressively higher energies and are collecting more data than ever before. As a result, the experimental precision of the measurements they perform is continuously improving. However, to infer what these measurements mean for the interactions between the fundamental constituents of matter, they have to be compared with, and interpreted in light of, our current theoretical understanding. This is done by performing first-principles computations for these high energy processes order by order in a power series expansion. After the computation, the resulting function is used in Monte Carlo simulations. The successive terms in the power series expansion, simplistically, become progressively smaller. Schematically, this can be written as:

$$F(\boldsymbol{x}) = f_{00}(\boldsymbol{x}) + \alpha\, f_{01}(\boldsymbol{x}) + \alpha^2\, \{f_{11}(\boldsymbol{x}) + f_{02}(\boldsymbol{x})\} + \dots\,. \tag{1}$$

where $\alpha \ll 1$ is the small expansion parameter. The term of interest to our current work is the one enclosed by the curly braces in equation 1 which we will refer to as the second-order term[1]. The function, $F(\boldsymbol{x})$, must be evaluated on the order of 10^{10} times or so for each simulation. However, for many processes, evaluating the second-order term, specifically, f_{02}, is computationally space- and time-intensive and could take several seconds to compute a single data point. Moreover, these samples cannot be reused leading to an overall high cost of computation for the entire process under consideration.

A simple solution to speed up the computation of the functions is to build a regressor using a representative sample. However, to achieve the precision necessary for matching with experimental

[1]Here *order* refers to the power of the expansion coefficient α.

results, the regressors need to produce very-high accuracy predictions over the entire domain of the function. The requirements that we set for the regressors, and in particular what we mean by *high precision*, are:

High precision: prediction error $< 1\%$ over more than 90% of the domain of the function
Speed : prediction time per data point of $< 10^{-4}$ seconds
Lightweight : the disk size of the regressors should be a few megabytes at the most for portability

In this work we explore the following novel concepts:

- With simulated data from real physics processes occurring in particle colliders, we study the error distributions over the entire input feature spaces for multi-dimensional distributions when using Boosted Decision Trees (BDT), Deep Neural Networks (DNN) and Deep Neural Networks with skip connections (sk-DNN).
- We study these regressors for 2, 4, and 8 dimensional (D) data making comparisons between the performance of BDTs, DNN and sk-DNNs with the aim of reaching errors smaller than 10^{-2} – 10^{-3} over at least 90% of the input feature space.
- We outline architectural decisions, training strategies and data volume necessary for building these various kinds of high-precision regressors.

In what follows, we will show that we can reduce the compute time of the most compute-intensive part of f_{02} (defined in equation 1) by several orders of magnitude, down from several seconds to sub-milliseconds without compromising the accuracy of prediction. We show that physics-motivated normalization strategies, learning strategies, and invocation of physics symmetries will be necessary to achieve the goal of high precision. In our experiments, the BDTs outperform the DNNs for lower dimensions while the DNNs give comparable or better accuracy at higher dimensions. DNNs with skip connections outperform fully connected DNNs even with half the number of parameters and can approach very close to the high-precision criteria with 8D data while fully connected DNNs fail. Our goal will be to make the most lightweight regressor that can be used for real-time prediction facilitating the speed-up of the Monte Carlo generators used for simulating collider processes.

2 Related Work

Building models for regression of amplitudes has been a continued attempt in the particle physics literature in the recent past. Boosted Decision Trees (BDTs) have been the workhorse of particle physics for a long time but mostly for performing classification of tiny signals from dominating backgrounds (Radovic et al., 2018). However, the necessity to use BDTs as a regressor for theoretical estimates of experimental signatures has only been advocated recently (Bishara & Montull, 2019) and has been shown to achieve impressive accuracy for 2D data.

Several other machine learning algorithms have been used for speeding up sample generation for Monte Carlo simulations. Winterhalder et al. (2022) proposed the use of Normalizing Flows (Jimenez Rezende & Mohamed, 2015) with Invertible Neural Networks to implement importance sampling (Müller et al., 2018; Ardizzone et al., 2018). Recently, neural network surrogates have been used to aid Monte Carlo Simulations of collider processes (Danziger et al., 2022). Badger et al. (2022) used Bayesian Neural networks for regression of particle physics amplitudes with a focus on understanding error propagation and estimation. Chen et al. (2021) attempted to reach the high-precision regime with neural networks and achieved 0.7% errors integrated over the entire input feature space. Physics-aware neural networks were studied by Maître & Truong (2021) in an attempt to handle singularities in the regressed functions. In the domain of generative models, GANs (Goodfellow et al., 2014; Springenberg, 2016; Brock et al., 2018) and VAEs (Brock et al., 2018) have been used for sample generation (Butter et al., 2021; Otten et al., 2021).

Similar applications have surfaced in other domains of physics where Monte Carlo simulations are used. Self-learning Monte Carlo methods have been explored by Liu et al. (2017). Applications of Boltzmann machines (Huang & Wang, 2017), deep neural networks (Shen et al., 2018) and autoregressive neural networks (Wu et al., 2021) have been seen recently. Stratis et al. (2022) use neural networks in Quantum Monte Carlo simulations to learn eigenvalues of Hamiltonians and the free

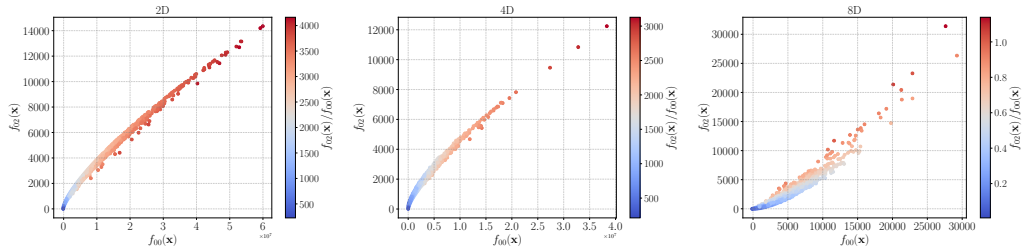

Figure 1: The effects of normalizing second-order term $f_{02}(\boldsymbol{x})$, with the zeroth-order term, $f_{00}(\boldsymbol{x})$. The two functions are highly correlated ($\rho = 0.87$ for 8D, $\rho = 0.91$ for 4D and $\rho = 0.96$ for 2D) and the resulting normalized functions, $f(\boldsymbol{x}) = f_{02}(\boldsymbol{x})/f_{00}(\boldsymbol{x})$ are much less peaked.

energy of spin configurations, an application that lies outside the domain of particle physics. However, the primary goal of all these efforts has been to avoid first-principles computation and, hence, reduce compute time while staying below credible error budgets that are set in a problem-specific manner.

In contrast to prior works (Bishara & Montull, 2019; Winterhalder et al., 2022; Butter et al., 2021; Otten et al., 2021), the novelty of our contribution is that we try to attain high precision in the entire domain of the function being regressed with fast and efficient regressors. For that, we compare BDTs and neural networks for functions with 2D, 4D, and 8D input feature spaces. We propose a new architecture which is a DNN with skip connections to avoid the problem of vanishing gradients for deeper neural networks and show that they perform better than fully connected DNNs. We also propose novel methods, derived from physics domain knowledge, for scaling the function being regressed with another function that is computationally inexpensive to calculate and is highly correlated with the function being regressed. We leverage symmetry properties of the physical process under consideration for reduction of the number of regressors required to be trained. The applicability of our work goes beyond the domain for which it has been developed and can be used for any application that requires high precision in speeding up simulations or sample generation.

3 Models: Decisions Trees and Neural Networks

In this section, we will develop several methods that will enable us to achieve the high-precision requirements that we set. As a benchmark, we will use the condition: $|\delta| < 1\%$ in over 90% of the domain of the function being regressed. Here, δ is defined as the L1 distance between the predicted value of the function, $y_{\text{predicted}}$, and its true value, y_{true}, normalized by y_{true}, or,

$$\delta = \frac{y_{\text{predicted}} - y_{\text{true}}}{y_{\text{true}}} \tag{2}$$

knowing a priori that y_{true} is positive definite. Usually, the performance of a regressor and model comparison in machine learning is done using a single accuracy measure which is a statistical average of the distribution for that accuracy measure over the entire test sample. This, however, does not provide a complete picture of the accuracy of the regressor in high-precision applications. Looking at error distributions instead of a single number leads to a better criterion for model selection and enhances the interpretability of the model.

Physics informed normalization: An attempt to build regressors with the raw data from the Monte Carlo simulations results in a failure to meet the high-precision requirements that we have set. Hence, we have to appeal to a novel normalization method derived from the physics that governs the physical processes. The functions of interest in particle physics processes at colliders are often very highly peaked in one or more dimensions. This makes it quite difficult to build a regressor that will retain the desired high precision over the entire domain of the function. This problem cannot be addressed by log scaling or standardizing to zero mean and unit variance since the peaks can be quite narrow and several orders of magnitude greater than the mean value of the function. We show examples in figure 1. A regressor trained on the log scaled function provides an error distribution over the entire domain which, when exponentiated, transforms to large errors around the peak. This

behavior is not desirable. Normal scaling does not help since the standard deviation of the distribution is much smaller than the peak value, often being several orders of magnitude smaller, making the peak-values outliers.

As a solution, we normalized the regressed function with the zeroth-order contribution given by the first term of equation (1), i.e., we transform to a distribution:

$$f(\boldsymbol{x}) = \frac{f_{02}(\boldsymbol{x})}{f_{00}(\boldsymbol{x})}\,. \tag{3}$$

This first-order term, $f_{00}(\boldsymbol{x})$, also has a similar peak structure and is highly correlated with the second-order term, $f_{02}(\boldsymbol{x})$ with $\rho \sim 0.9$. Hence, this normalization yields a distribution, $f(\boldsymbol{x})$, that is more tractable to regress. Computation of the first-order term from first principles is numerically inexpensive and does not require regression. Furthermore, we standardize the distribution by removing the mean and scaling to unit variance..

3.1 DNN DESIGN DECISIONS

The DNN architectures that we used are fully connected layers with *ReLU* activation for the hidden layers and linear activations for the output layer. We use the following design decisions:

Objective function: we use the mean-square-error loss function without any regularization. While we use linear relative error, δ, to estimate the performance of the model over the entire feature space, our decision to use the mean-square-error loss function is made so as to preferentially penalize outliers and reduce their occurrence.

Learning rate: It is necessary to cool down the learning rate as a minimum of the objective function is approached. This is absolutely necessary to search out an optimum that allows for uniformly low error over the entire feature space. For both the DNN and the sk-DNN, we use the Adam optimizing algorithm. Kingma & Ba (2015) discuss an inverse square-root decay rate scaling for the Adam optimizer. We do not find this optimal for this high-precision application. The square-root cooling is quite rigid in its shape as it is invariant to scaling up to a multiplicative constant. Hence, we use an exponential cooling of the learning rate which has an asymptotic behavior similar to the inverse-square-root scaling but its shape is far more tunable. The learning rate is cooled down starting from 10^{-3} down to 10^{-6}. Much of the learning is completed during the early stages of the training, i.e. within 200 epochs. The R^2 score at this point is about 0.5% from the final R^2 score ($> 99.9\%$). However, to attain the high-precision requirements, the final stages of the training are necessary and take about 1500 – 2500 more epochs.

Training epochs and validation: An early stopping criterion based on RMSE is used to determine the number of epochs the regressors were trained for with *patience*[2] set to an unusually large number, 200 epochs. We use this large patience to allow the optimizer to possibly move to a better optimum while having a very small learning rate if a better one exists. We split the data into 60% training set and 40% validation set. The large fraction for the validation set is necessary to make sure that errors are uniformly low over the entire domain of the function being regressed.

3.2 DNN WITH SKIP CONNECTIONS

In addition to a fully connected DNN, we also experiment with a DNN with skipped connections (sk-DNN) to address the problem of vanishing gradients for deeper neural networks. The building block of the sk-DNN is illustrated in figure 2. Given an input $\boldsymbol{x}$ the output of the block is

$$\boldsymbol{y} = g\left(f(\boldsymbol{x}) + \boldsymbol{W}\boldsymbol{x}\right) \tag{4}$$

where $f(\boldsymbol{x})$ is the output of the third layer with linear activation and $\boldsymbol{W}$ is a trainable weight matrix of dimension $i \times j$ when the input dimension, i, is different from the output dimension, j, and $\boldsymbol{I}$ otherwise. The structure of this block can be derived from the Highway Network (Srivastava et al., 2015) architecture with the *transform* gate set to $\boldsymbol{I}$ and the *carry* gate set to $\boldsymbol{W}$ for $\dim(\boldsymbol{x}) \neq \dim(\boldsymbol{y})$

[2]We define patience as the number of epochs after which the training is stopped as no reduction is seen in the RMSE computed from the validation set and the weights and biases are reset to those corresponding to the lowest RMSE.

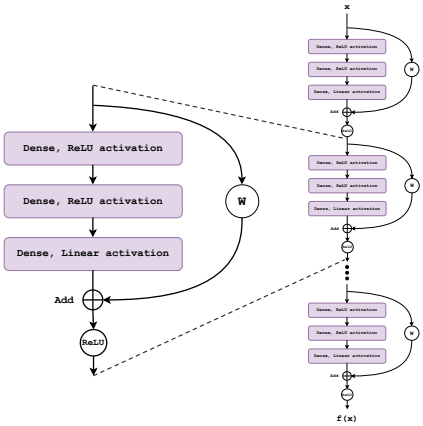

Figure 2: The building block for a DNN with skip connections. The first two layers are fully connected with ReLU activation. The last layer has linear activation and is added to the input through the skip connection before being transformed with a non-linear ReLU function. The matrix $\boldsymbol{W}$ is a weight matrix which is trainable if the input and output dimensions are different for the block and $\boldsymbol{I}$ otherwise. The blocks are stacked sequentially to form the neural network.

and $\boldsymbol{I}$ otherwise. Structurally, the sk-DNN block is similar to a ResNet block (He et al., 2015) with a different set of hidden layers.

We keep the normalization of the target variable and the learning rate decay schedule the same as for the DNN. We also test the sk-DNN with the weight matrix, $\boldsymbol{W}$ fixed with a random initialization of the elements and see no difference in the accuracy of the model. In the final experiments, we will keep $\boldsymbol{W}$ trainable.

4 Experiments

4.1 Physics Simulations

Symmetry properties reduce the number of required functions

Dimensionality	Total functions	Independent functions	Sum is physical?
2D	18	5	Yes
4D	162	25	No
8D	8	4	Yes

Table 1: The number of total and independent functions that arise at 2, 4, and 8D and the number of independent functions, a minimal subset from which the other functions can be generated by re-mapping the feature space.

The functions in question are maps, $f_{ij}^{(n)} : \mathbb{R}^n \to \mathbb{R}$, where $n \in \{2, 4, 8\}$ and $i, j \in \{0, 1, 2\}$, cf. equation 1. The domain of the functions, i.e. the feature space, is mapped to the unit hypercube and populated from a uniform distribution. The corresponding datasets are generated using the particle physics simulation code `VVAMP` (Gehrmann et al., 2015) from first principles using building-block functions that we will refer to as form factors. Apart from the 2D dataset, which is a special case of the 4D one, the same form factors were used to generate the 4D and 8D datasets. The difference between the 4D and 8D feature spaces lies in the physics of the process in question, namely the number of external particles the functions describe. The regressor of the 4D functions, $g_{ij}^{(4)} \approx f_{ij}^{(4)}$, can be used to generate the 8D functions, $f_{ij}^{(8)}$, after multiplying by two other (exact) functions that are computationally inexpensive to calculate and summing them.

The number of resulting functions, technically called helicity amplitudes, depends on the dimension as shown in table 1. While the number of required regressors for the 4D feature space is the largest,

it also offers the most flexibility for further downstream physics analyses. Specifically, to generate the 8D functions, more details of the process have to be specified during data generation which is then frozen into the regressor. Consequently, different particle physics analyses will require different regressors. By contrast, the 4D regressors are more general-purpose and do not contain *any* frozen physics parameters.

The smaller number of necessary functions in the third column of table 1 is obtained by leveraging the symmetry properties discussed below which are derived from physics domain knowledge. The last column indicates whether summing the functions has a physics meaning; in the cases where it does, i.e. 2D and 8D, only the single regressor of the sum of the functions is truly required.

Symmetry properties: the full set of functions, $f_{ij}^{(n)}$, for any dimension, n, is over complete. Pairs of functions can be mapped into one another via particular permutations of the external particles the process describes. This translates into a linear transformation on the second coordinate, x_2, independently and in combination with the permutation of the third and fourth coordinates, x_3 and x_4, in feature space. For example, in 4D, two permutations $\pi_{12} : p_1 \leftrightarrow p_2$ and $\pi_{34} := p_3 \leftrightarrow p_4$, where p_i is a particle with label i reduces the number of independent functions from 162 to 25.

Permutation	particle symmetry	coordinate symmetry
π_{12}	$p_1 \leftrightarrow p_2$	$x_2 \to 1 - x_2$
π_{34}	$p_3 \leftrightarrow p_4$	$x_2 \to 1 - x_2$ and $x_3 \leftrightarrow x_4$

Computational burden of Monte Carlo simulations. Generating the 2D, 4D and 8D datasets required 144 hours on 96 `AMD EPYC 7402` cores for 13 million data points per set. This had to be done twice, once for the 2D dataset and once for the 4D and 8D datasets which were generated from the same computationally intensive form factors which have to be calculated from first principles. In contrast, the regressors that we build generate a million samples in a few seconds to a few minutes on any desktop computer.

4.2 BDT vs. DNN vs. SK-DNN

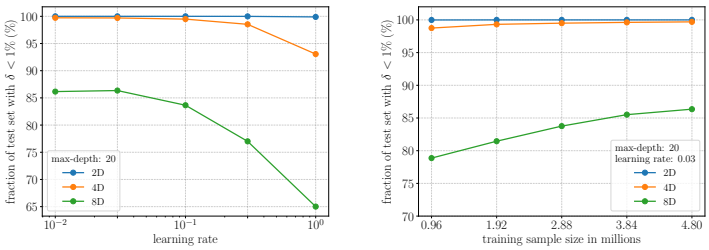

Figure 3: **Left panel:** the effect of learning rate on achieving high precision with BDTs. The learning rate is not an important parameter for low dimensions but significantly changes the precision of the regressor for higher dimensions. **Right panel:** the data volume required to train a BDT. For lower dimensions small volumes of data ($<$ 1M) is sufficient. However, for higher dimensions, a lot more data is necessary.

We use `XGBoost` (Chen & Guestrin, 2016) to implement the BDTs. In varying the architecture of the regressors, we focus on the *max-depth* of the BDT which is a hyperparameter that controls the maximum depth to which a tree grows in a boosted ensemble. If figure 3 we show how changing the learning rate and the training data volume changes the accuracy of the trained BDT models. In the final version of our experiments, we use a learning rate of 0.01 and 10 million data points of which 48% is used for training, 32% is used for validation and early stopping and 20% is used for testing. The data strategy remains the same for the neural networks.

For the neural networks, we focus on the depth, width and number of trainable parameters in the regressor (denoted as width-depth (trainable parameters) in the tables and figures). The depth of the

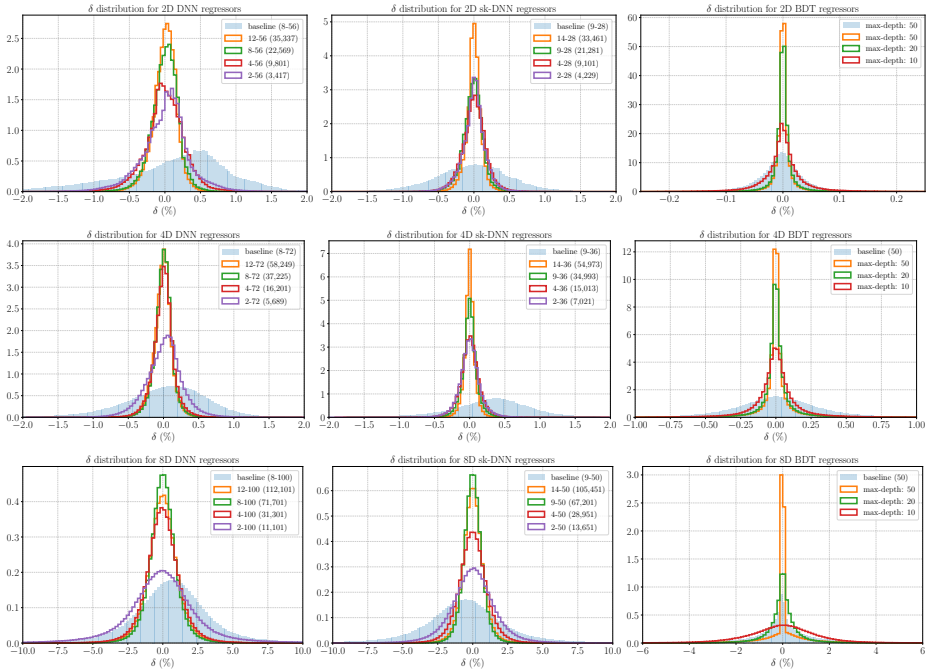

Figure 4: The $\delta = (y_{\text{predicted}} - y_{\text{true}})/y_{\text{true}}$ distribution for the various 2D (**upper panels**), 4D (**middle panels**), and 8D (**lower panels**) regressors. The labels for the DNN and sk-DNN designate *depth-width (number of parameters)* where depth is the number of layers for the DNN and the number of blocks for the sk-DNN. For the BDTs, the labels denote *max_depth:* N with max_dept being the maximum depth of the trees. Detailed analyses can be found in the text.

sk-DNN denotes the number of sequential sk-DNN blocks in the regressor and not the total number of layers. The width of the sk-DNNs is chosen to be half the width of the DNNs and the depth of the sk-DNN is adjusted so that they have approximately the same number of parameters as the DNNs with similar depth. One exception is that the sk-DNN with 2 blocks has more parameters than the corresponding DNN with 2 layers. The data strategy remains the same as for the BDTs.

To compare the performance of the regressor we use the distribution of δ as defined in equation 2. We focus on this distribution as it is important for the high-precision requirement to identify the fraction of test data that has large errors. We will identify the following statistics:

$|\delta| < 1\%$: the fraction of the test set that has δ less than 1%
μ_δ: the mean of the δ distribution
σ_δ: the standard deviation of the δ distribution

Baselines: To understand the efficacy of the optimization strategies that we developed, we build a baseline without any optimization for BDTs, DNNs and sk-DNNs. For all the regressors, we do not normalize the data as described in section 3, rather, we only log scale the data. We set the train-validation split to 80%-20%. For the BDTs, we use an ensemble with max-depth = 50, set the learning rate to 0.1. For the DNNs and sk-DNNs, we fix the learning rate of the Adam optimizer at 10^{-3}, lower the patience to 10 rounds, and use the most effective architecture chosen from amongst the high-precision regressors. The results are presented in table 2 and figure 4. We see that without the optimizations the DNNs fail to regress the function and the BDTs perform very poorly.

Key results: we present the results of the experiments in table 2 and figure 4. We show the distribution of errors over two variables, square root of the center-of-mass energy, $\sqrt{s} \mapsto x_1$, and $\cos\theta \mapsto x_2$ in figure 5. It is clear that the BDTs far outperform the DNNs in 2D and 4D. The sk-DNN comes closer to the degree of accuracy that the BDT has in 4D. However, at 8D the sk-DNN not only far outperforms the fully connected DNNs, but also is comparable to the BDTs in the error distribution. The disadvantage of the BDTs is that they take up significant disk space as the ensemble grows large,

	2D	$\lvert\delta\rvert < 1\%(\%)$	$\lvert\delta\rvert < 0.1\%(\%)$	$\mu_\delta(\%)$	$\sigma_\delta(\%)$
DNN	2-56 (3,417)	98.09	28.18	0.0057	0.39
	4-56 (9,801)	99.37	33.28	−0.001	0.29
	8-56 (22,569)	99.93	44.03	0.0017	0.19
	12-56 (35,337)	99.98	49.26	0.0005	0.17
	baseline (8-56)	77.16	9.01	0.1361	1.83
sk-DNN	2-28 (4,229)	99.85	55.26	−0.0007	0.18
	4-28 (9,101)	99.93	50.01	0.0009	0.18
	9-28 (21,281)	99.96	57.74	−0.0147	0.15
	14-28 (33,461)	99.97	73.08	−0.0022	0.12
	baseline (9-28)	90.79	15.31	0.0173	1.39
BDT	max-depth: 10	100.0	95.01	−0.0001	0.05
	max-depth: 20	100.0	99.1	0.0	0.03
	max-depth: 50	100.0	99.16	0.0	0.02
	baseline (50)	99.91	94.04	−0.0045	0.1
	4D				
DNN	2-72 (5,689)	99.48	34.22	0.008	0.28
	4-72 (16,201)	99.94	58.87	−0.0004	0.16
	8-72 (37,225)	99.98	64.3	−0.0008	0.13
	12-72 (58,249)	99.99	63.76	−0.012	0.13
	baseline (8-72)	88.67	13.63	0.0449	1.1
sk-DNN	2-36 (7,021)	99.92	57.81	0.0	0.17
	4-36 (15,013)	99.93	59.18	0.0018	0.16
	9-36 (34,993)	99.99	72.37	−0.0006	0.12
	14-36 (54,973)	100.0	84.99	−0.0016	0.08
	baseline(9-36)	84.99	10.81	0.2701	1.11
BDT	max-depth: 10	99.26	66.16	0.0014	0.22
	max-depth: 20	99.41	81.34	0.0016	0.18
	max-depth: 50	99.4	83.19	0.0017	0.18
	baseline (50)	95.85	27.8	0.0023	0.55
	8D				
DNN	2-100 (11,101)	37.79	4.01	0.125	3.85
	4-100 (31,301)	63.83	7.65	0.015	1.52
	8-100 (71,701)	73.48	9.57	0.021	1.17
	12-100 (112,101)	68.1	8.41	0.015	1.26
	baseline (8-100)	31.97	3.3	0.799	4.38
sk-DNN	2-50 (13,651)	52.45	5.9	0.044	1.93
	4-50 (28,951)	70.63	8.76	0.017	1.23
	9-50 (67,201)	87.28	13.52	0.007	0.75
	14-50 (105,451)	83.65	12.37	0.013	0.86
	baseline (9-50)	30.91	3.23	−0.547	4.85
BDT	max-depth: 10	53.62	6.21	0.086	2.69
	max-depth: 20	86.58	25.54	0.03	1.24
	max-depth: 50	84.9	63.77	0.061	1.47
	baseline (50)	63.72	16.07	0.383	6.33

Table 2: Parameters extracted from the error distributions shown in figure 4. Predictions from 1 million test samples were used to generate these statistics. The best performing models for each of 2D, 4D, and 8D are marked in bold. More details regarding the notations are available in the text.

which is necessary for high precision applications but affects their portability. Hence the sk-DNNs are a good solution for having a portable, yet accurate regressor.

5 Conclusion

With Monte Carlo simulation in Physics being time and resource intensive, a distinct necessity of building regressors for speeding up the simulations exists. We carefully examine the requirements of high-precision for these regressors and lay down design strategies to achieve the necessary benchmarks. We use domain knowledge from particle physics to determine normalization strategies, apply symmetry arguments to reduce the number of necessary regressors, and set benchmarks for high-precision regression.

We show that Boosted Decision Trees are reliable workhorses that can easily outperform DNNs at lower dimensions even when very large and complex neural networks are used. However, this edge

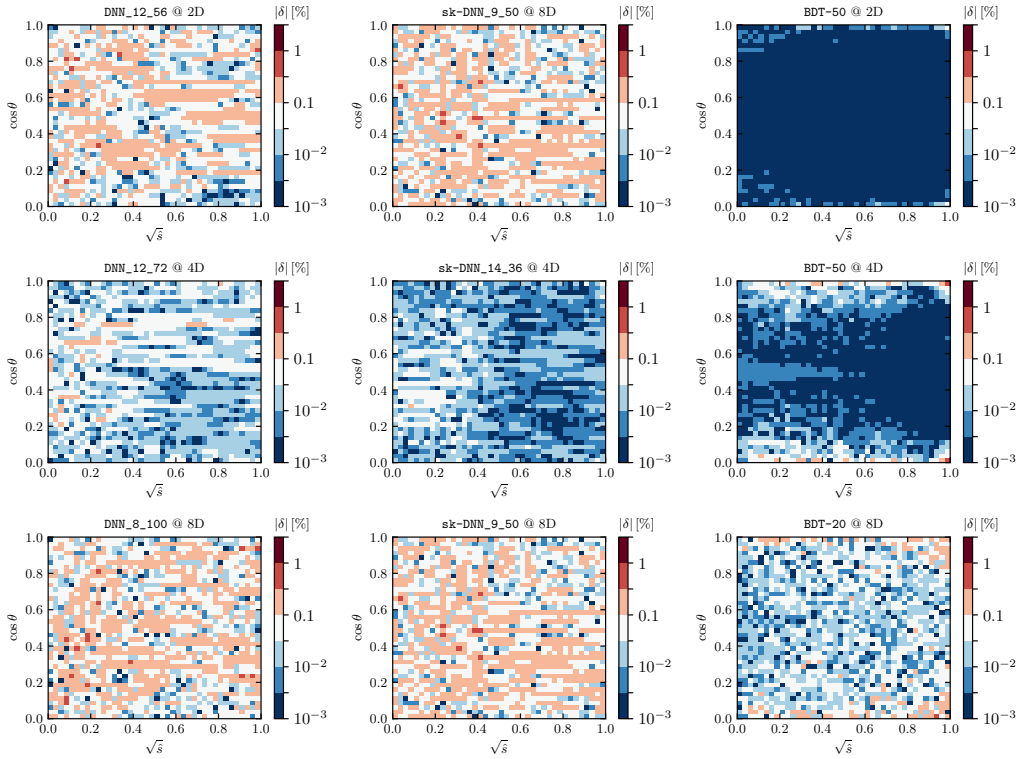

Figure 5: The $\delta = (y_{\text{predicted}} - y_{\text{true}})/y_{\text{true}}$ distribution for the various 2D (upper panels), 4D (middle panels) and 8D (lower panels) regressors. The errors are averaged over each bin.

that BDTs have over neural networks tends to fade at higher dimensions especially when DNNs with skip connections are used. In fact, for 8D data, it is imperative that sk-DNNs be used as it allows us to get close to the benchmark of $\delta < 1\%$ over 90% of the domain of the function. Moreover, once the models are sufficiently complex to learn the data well, sk-DNNs outperform DNNs that have twice the number of parameters.

The primary disadvantage of BDTs is that for higher dimensions the ensemble of trees grows large enough to take a significant amount of disk space, often > 1 GB, affecting the portability of the regressor if it is intended to be used as part of a Monte Carlo simulation package. On the other hand, the disk space occupied by a neural network stay below a few megabytes, making them a lot more portable.

In summary, the important conclusions of our work are:

- High precision regressors required to speed up Monte Carlo simulations by factors of 10^3 – 10^6 are better optimized by leveraging physics domain knowledge and symmetry arguments.
- BDTs outperform DNNs at lower dimensions but start to make large errors in predictions in parts of the function domain at higher dimensions. While fully connected DNNs perform relatively well at higher dimensions, a DNN with skip connections outperforms both BDTs and fully connected DNNs at 8D.
- sk-DNNs can outperform DNNs with twice the number of parameters. At higher dimensions, sk-DNNs are necessary to approach or meet the high-precision requirements.
- Compared to the few seconds that it takes for a single sample generation during a Monte Carlo simulation, the regressors we design can provide precise predictions in milliseconds to microseconds.

In this work we aimed at reaching the desired precision but by no means have we exhausted possibilities of achieving even higher precision. As future directions, looking into the role of activation

functions, or the modifications of the likelihood with possible physics constraints or symmetry arguments or further reducing the number of models by simultaneously predicting a set of functions from a single neural network might be directions that can be explored in detail.

CONTRIBUTION TO SUSTAINABILITY

Monte Carlo simulations of physics processes leave a very large carbon footprint. It is estimated that about 50% of the energy budget of each experiment at the Large Hadron Collider is consumed by such simulations. Hence, our work directly contributes to reducing the carbon footprint significantly through a much more efficient way of generating these events.

The regressors we build can be trained on personal computers with a few CPU threads and a single GPU in under a day as our focus has been to build lightweight models. No special hardware is required to train or test these regressors. The generation of the simulated data however requires high-performance computing clusters and we are making this data publicly and freely available for future research.

REPRODUCIBILITY STATEMENT

The code and the data will be made publicly available after the review process to maintain the double-blind requirements.

AUTHOR CONTRIBUTIONS

ACKNOWLEDGMENTS

REFERENCES

Lynton Ardizzone, Jakob Kruse, Sebastian J. Wirkert, Daniel Rahner, Eric W. Pellegrini, Ralf S. Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. *CoRR*, abs/1808.04730, 2018. URL http://arxiv.org/abs/1808.04730.

Simon Badger, Anja Butter, Michel Luchmann, Sebastian Pitz, and Tilman Plehn. Loop Amplitudes from Precision Networks. 6 2022. URL http://arxiv.org/abs/2206.14831.

Fady Bishara and Marc Montull. (Machine) Learning amplitudes for faster event generation. 12 2019. URL http://arxiv.org/abs/1912.11055.

Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *CoRR*, abs/1809.11096, 2018. URL http://arxiv.org/abs/1809.11096.

Anja Butter, Sascha Diefenbacher, Gregor Kasieczka, Benjamin Nachman, and Tilman Plehn. GANplifying event samples. *SciPost Phys.*, 10(6):139, 2021. doi: 10.21468/SciPostPhys.10.6.139. URL https://doi.org/10.21468/SciPostPhys.10.6.139.

I-Kai Chen, Matthew D. Klimek, and Maxim Perelstein. Improved neural network Monte Carlo simulation. *SciPost Phys.*, 10(1):023, 2021. doi: 10.21468/SciPostPhys.10.1.023. URL https://doi.org/10.21468/SciPostPhys.10.1.023.

Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pp. 785–794, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4232-2. doi: 10.1145/2939672.2939785. URL http://doi.acm.org/10.1145/2939672.2939785.

Katharina Danziger, Timo Janßen, Steffen Schumann, and Frank Siegert. Accelerating Monte Carlo event generation – rejection sampling using neural network event-weight estimates. *SciPost Phys.*, 12:164, 2022. doi: 10.21468/SciPostPhys.12.5.164. URL http://doi.org/10.21468/SciPostPhys.12.5.164.

Thomas Gehrmann, Andreas von Manteuffel, and Lorenzo Tancredi. The two-loop helicity amplitudes for $q\overline{q}' \to V_1 V_2 \to 4$ leptons. *JHEP*, 09:128, 2015. doi: 10.1007/JHEP09(2015)128. URL https://doi.org/10.1007/JHEP09(2015)128.

Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. URL https://arxiv.org/abs/1406.2661.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL http://arxiv.org/abs/1512.03385.

Li Huang and Lei Wang. Accelerated monte carlo simulations with restricted boltzmann machines. *Phys. Rev. B*, 95:035105, Jan 2017. doi: 10.1103/PhysRevB.95.035105. URL https://link.aps.org/doi/10.1103/PhysRevB.95.035105.

Danilo Jimenez Rezende and Shakir Mohamed. Variational Inference with Normalizing Flows. *arXiv e-prints*, art. arXiv:1505.05770, May 2015.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL http://arxiv.org/abs/1412.6980.

Junwei Liu, Yang Qi, Zi Yang Meng, and Liang Fu. Self-learning monte carlo method. *Phys. Rev. B*, 95:041101, Jan 2017. doi: 10.1103/PhysRevB.95.041101. URL https://link.aps.org/doi/10.1103/PhysRevB.95.041101.

Daniel Maître and Henry Truong. A factorisation-aware Matrix element emulator. *JHEP*, 11:066, 2021. doi: 10.1007/JHEP11(2021)066. URL https://doi.org/10.21468/10.1007/JHEP11(2021)066.

Thomas Müller, Brian McWilliams, Fabrice Rousselle, Markus Gross, and Jan Novák. Neural importance sampling. *CoRR*, abs/1808.03856, 2018. URL http://arxiv.org/abs/1808.03856.

Sydney Otten, Sascha Caron, Wieske de Swart, Melissa van Beekveld, Luc Hendriks, Caspar van Leeuwen, Damian Podareanu, Roberto Ruiz de Austri, and Rob Verheyen. Event Generation and Statistical Sampling for Physics with Deep Generative Models and a Density Information Buffer. *Nature Commun.*, 12(1):2985, 2021. doi: 10.1038/s41467-021-22616-z. URL https://doi.org/10.1038/s41467-021-22616-z.

Alexander Radovic, Mike Williams, David Rousseau, Michael Kagan, Daniele Bonacorsi, Alexander Himmel, Adam Aurisano, Kazuhiro Terao, and Taritree Wongjirad. Machine learning at the energy and intensity frontiers of particle physics. *Nature*, 560(7716):41–48, Aug 2018. ISSN 1476-4687. doi: 10.1038/s41586-018-0361-2. URL https://doi.org/10.1038/s41586-018-0361-2.

Huitao Shen, Junwei Liu, and Liang Fu. Self-learning monte carlo with deep neural networks. *Phys. Rev. B*, 97:205140, May 2018. doi: 10.1103/PhysRevB.97.205140. URL https://link.aps.org/doi/10.1103/PhysRevB.97.205140.

Jost Tobias Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. In Yoshua Bengio and Yann LeCun (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL http://arxiv.org/abs/1511.06390.

Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *CoRR*, abs/1505.00387, 2015. URL http://arxiv.org/abs/1505.00387.

Georgios Stratis, Phillip Weinberg, Tales Imbiriba, Pau Closas, and Adrian E. Feiguin. Sample generation for the spin-fermion model using neural networks. *arXiv e-prints*, art. arXiv:2206.07753, June 2022. URL https://ui.adsabs.harvard.edu/abs/2022arXiv220607753S.

Ramon Winterhalder, Vitaly Magerya, Emilio Villa, Stephen P. Jones, Matthias Kerner, Anja Butter, Gudrun Heinrich, and Tilman Plehn. Targeting multi-loop integrals with neural networks. *SciPost Phys.*, 12(4):129, 2022. doi: 10.21468/SciPostPhys.12.4.129. URL http://doi.org/10.21468/SciPostPhys.12.4.129.

Dian Wu, Riccardo Rossi, and Giuseppe Carleo. Unbiased monte carlo cluster updates with autoregressive neural networks. *Phys. Rev. Research*, 3:L042024, Nov 2021. doi: 10.1103/PhysRevResearch.3.L042024. URL https://link.aps.org/doi/10.1103/PhysRevResearch.3.L042024.