



Classifier surrogates: sharing AI-based searches with the world

Sebastian Bieringer^{1,a} , Gregor Kasieczka¹ , Jan Kieseler² , Mathias Trabs³ 

¹ Institut für Experimentalphysik, Universität Hamburg, Luruper Chaussee 149, 22761 Hamburg, Germany

² Institut für Experimentelle Teilchenphysik, Karlsruher Institut für Technologie, Wolfgang-Gaede-Str. 1, 76131 Karlsruhe, Germany

³ Institut für Stochastik, Karlsruher Institut für Technologie, Englerstr. 2, 76131 Karlsruhe, Germany

Received: 27 February 2024 / Accepted: 9 September 2024
© The Author(s) 2024

Abstract In recent years, neural network-based classification has been used to improve data analysis at collider experiments. While this strategy proves to be hugely successful, the underlying models are not commonly shared with the public and rely on experiment-internal data as well as full detector simulations. We show a concrete implementation of a newly proposed strategy, so-called Classifier Surrogates, to be trained inside the experiments, that only utilise publicly accessible features and truth information. These surrogates approximate the original classifier distribution, and can be shared with the public. Subsequently, such a model can be evaluated by sampling the classification output from high-level information without requiring a sophisticated detector simulation. Technically, we show that continuous normalizing flows are a suitable generative architecture that can be efficiently trained to sample classification results using conditional flow matching. We further demonstrate that these models can be easily extended by Bayesian uncertainties to indicate their degree of validity when confronted with unknown inputs by the user. For a concrete example of tagging jets from hadronically decaying top quarks, we demonstrate the application of flows in combination with uncertainty estimation through either inference of a mean-field Gaussian weight posterior, or Monte Carlo sampling network weights.

1 Introduction

Current experimental work in particle physics, for example by the ATLAS and CMS collaborations, uses deep learning-based taggers to great success [1–4]. Such models often define unique and essential quantities in the analysis chain, which are hard to understand in terms of physical quanti-

ties. While the performance benefit is apparent, best practices for sharing the analysis as for traditional cut-based analyses [5,6] are not yet established. This especially hinders the re-interpretation of experimental results. Recently, a first set of proposals on sharing neural network-based results has been published [7]. On the purely technical side, solutions exist for sharing serialized networks [8,9] and some first searches shared with serialized models have been made public [10–13].

However, when the model inputs contain features which are not available outside the collaborations or can only be simulated at high computational cost within the collaboration, the benefit of sharing the network weights is limited as results still can either not be reproduced at all, or are very expensive. Costly and unavailable input features include detector level quantities, such as hits, or highly detector dependent quantities, such as soft jet-substructure variables. For example, both b -taggers of ATLAS and CMS use detector dependent information [14,15] and current research shows the best classification performance is achieved when using detector-level data, rather than only high-level observables [4,16]. For these cases, sharing a surrogate model trained to reproduce the classification results from truth-, parton- or reconstruction-level inputs has recently been proposed in discussions at the LHC Reinterpretation Forum and the 2023 PhysTeV workshop at Les Houches [7]. We will follow the newly introduced terminology and refer to such models as *Classifier Surrogates*. In this work

- we demonstrate for a concrete example how such a Classifier Surrogate could be constructed and evaluated
- and present a novel combination of Continuous Normalizing Flows with Monte Carlo-based Bayesian Neural Networks (BNN) for this purpose.

^a e-mail: sebastian.guido.bieringer@uni-hamburg.de (corresponding author)

Complementary to sharing the full likelihood or the full statistical model [17]

$$p(\text{data} \mid \mu),$$

a Classifier Surrogate can be used to model dependencies on parameters $\tilde{\mu}$ that were not explicitly included in the statistical model at the time of the release and are hard to model with public fast simulation tools like Delphes [18]. Altering the parameters requires that the released model includes intermediate information, for example distributions of observables that are used in a template fit. These might stem from the output x of a complex neural network classifier. For such distributions, the application of a classification surrogate can be beneficial.

In practice, a Classifier Surrogate

$$p(x \mid c)$$

can be used to predict classifier output from any single-event surrogate input

$$c \sim p(c \mid \tilde{\mu}).$$

This simulation of truth-, parton- or reconstruction-level data allows an arbitrary choice of parameters $\tilde{\mu}$. If the simulated event is out-of-distribution (OOD) of the training data of the classifier, the surrogate will predict large uncertainties and thus prevent the practitioner from interpreting the analysis where the classification can not be applied reliably. For simulated events within the classifiers input range, the surrogate predicts samples from the distribution of viable classifier output. This output prediction in turn can be used to estimate expectation values in histogram bins of derived observables in full analogy to the processing of the classification results from observed data. A statistical model for the new parameters

$$p(\text{data} \mid \tilde{\mu})$$

can again be derived from the processed and possibly histogrammed surrogate output, for example by assuming Poisson-distributed bin values. The surrogate strategy therefore is a truly “open-world” approach to sharing a classifier-aided analysis.

The uncertainties from the statistical limitation of the dataset, as well as the the smearing introduced by the detector simulation and reduced information of the input c may also be absorbed into an additional nuisance parameter of the new statistical model.

Depending on the nuisance handling strategy used for classifier training [19], the dependence on the nuisance

parameters needs to be included in the surrogate

$$p(x \mid c) \rightarrow p(x \mid c, \vartheta)$$

for nuisance-parameterized classifiers or in the corresponding input model

$$p(c \mid \tilde{\mu}) \rightarrow p(c \mid \tilde{\mu}, \vartheta)$$

for nuisance-invariant approaches.

If trained on truth- or parton-level, generating surrogate input events $c \sim p(c \mid \tilde{\mu})$ does not require detector simulation and can thus significantly improve the computational cost of any re-interpretation. Furthermore, eliminating the detector simulation also removes a major bottle-neck for sharing the results with colleagues, that do not have access to collaboration internal simulation-settings.

We introduce the strategy on the concrete example of a classifier derived from the particle transformer [16]. This setup is introduced in Sect. 2. In Sect. 3, we then discuss why a Classifier Surrogate needs to employ a generative architecture and introduce a possible architecture in Sect. 4. To model increased uncertainty for unknown inputs, we develop two BNN implementations of the architecture in Sect. 5. In Sect. 6, we discuss the performance of the surrogate both for data within the distribution of the training data, as well as for data new to the model. We evaluate calibration and scaling to the tails of the distribution, as well as OOD indication.

2 Particle transformer and JetClass dataset

As internal taggers of the big collaborations are not available for public study, we choose to emulate the state-of-the-art jet tagger, the Particle Transformer (ParT) [16]. ParT is an attention-based model trained to distinguish 10 different types of jets using per-particle information and trained on the 100 M JetClass dataset [20]. The features include kinematics, particle identification, and trajectory displacement of every particle in the jet.

From the large initial JetClass dataset as stand-in for the internal collaboration datasets, we distill our toy dataset by calculating transverse momentum, pseudorapidity, scattering angle, jet energy, number of particles, soft drop mass [21] and N-subjettiness [22] for $N = 1, \dots, 4$, as well as the output of the full ParT for the regarding event. For the first studies we will restrain the experiments to the first five jet-observables as well as the true top or QCD label as surrogate input.

While learning a surrogate of a multiclassifier is possible by using a generative architecture with a multidimensional output space, we restrict the setup to finding a surrogate for binary classification of top jets. The toy train and validation datasets contain 1M jet events each from Z-events and




Fig. 1 Histograms of the ParT classification results for the 1000 jet events of the training data, closest to the two arbitrary jets indicated by the dotted lines. Although being classified with varying confidence from detector-level data, the high-level observables p_T , E_{jet} and n_{const} appear identical

hadronic decay of $t\bar{t}$. To reduce the 10-dimensional ParT output to a binary classification result, we rescale

$$p_{t \rightarrow bqq'} = \frac{p_{t \rightarrow bqq'}^{ParT}}{p_{t \rightarrow bqq'}^{ParT} + p_{Z \rightarrow qq'}^{ParT}}$$

3 Detector smearing distribution

Due to the stochasticity of the detector simulation, jets with the same high-level observables can differ a lot on detector-level. Similarly, jets simulated for identical truth- or parton-level events, would vary significantly. These jets will thus result in different ParT outputs defining the likelihood per set of high-level observables

$$p(\underbrace{ParT}_x | \underbrace{p_T, \eta, \phi, E_{jet}, n_{const}, \dots}_c). \tag{1}$$

Based on its physical origin, we will also refer to this distribution as the detector smearing distribution.

We can generate a first approximation of this distribution by generating a histogram of the ParT output for the closest points in p_T , E_{jet} and n_{const} . In Fig. 1 we show this histogram for the 1000 nearest jet events in the training sample for two arbitrary jet events in the bulk of the transverse momentum distribution at $p_T \approx 530$ GeV. The imperfect ParT classification introduces an output distribution with tails for events indistinguishable from the high-level features. Employing a generative architecture as introduced in Sect. 4, allows us to infer this distribution from the high-level observables.

For the toy setup, we assume the classifier to be constructed invariant for the relevant nuisance parameters [19]. Whenever a nuisance-parameterized classifier is applied, the nuisance parameters need to be included into the likelihood as well.

4 Neural density estimation

While all flavours of generative models have found numerous applications in high-energy physics, for example in [23] and [24], normalizing flows can easily and efficiently be applied to infer complex, low-dimensional conditional distributions [25, 26]. For an early application to particle physics, see for example MadMiner [27] and Bayesflow [26, 28]. In our tests, coupling block-based Normalizing Flows exhibit great performance for dense phase space regions, but larger deviations when modelling tails of distributions. To boost the performance of the model we employ Continuous Normalizing Flows (CNF), a generalization of coupling block Flows based on ordinary differential equations (ODE) introduced in Sect. 4.1.

In Classifier Surrogates, the deficiency of coupling block-based normalizing flows to model distribution tails is masked to large extend by the softmax-normalization employed on the classifier, and thus also surrogate, when calculating class probabilities. We do observe similar performance between both architectures. However, CNFs are also much more parameter efficient allowing us to reduce the number of parameters needed by a factor of ≈ 20 at the cost of slower inference time. As the weights of the surrogate are designed to be shared, and we do expect their use in case studies rather than evaluating on millions of jets, we believe that CNFs are best suited for the application.

4.1 Continuous Normalizing Flows and conditional flow matching

First introduced in [29], CNFs define a transformation $\phi_t : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ called *flow* dependent on a time variable t . The time variable is the continuous equivalent to the number of a coupling blocks in a coupling block-flow [30]. Instead of having multiple flow instances, the dependence of ϕ on t is defined through the ODE

$$\frac{d}{dt} \phi_t(x) = v_t(\phi_t(x)), \quad \phi_0(x) = x, \tag{2}$$

by the time dependent *vector-field* $v_t : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, which itself is approximated by a deep neural network

$$\tilde{v}_t(\cdot, \theta) \approx v_t.$$

While this network can be arbitrarily complex, we stick to fully-connected architectures due to the low dimensionality of the task. In our case, the flow transforms data from a Gaussian distribution $\mathcal{N}(0, 1)$ for $t = 0$ into ParT output at $t = 1$. This choice sets the boundaries of the *probability path* $p_t : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}_{>0}$ induced by the vector-field trough

Equation (2) and

$$p_t(x) = p_0 \left(\phi_t^{-1}(x) \right) \det \left(\frac{\partial \phi_t^{-1}(x)}{\partial x} \right). \tag{3}$$

A standard CNF is trained by solving the ODE Eq. (2) in reverse and minimizing the negative log-likelihood (NLL) of input data at $t = 1$. The computation of this loss objective is expensive, especially for higher dimensional models.

Thus, the authors of [31] introduce the conditional flow matching (CFM) objective

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t, q(x_1), p_t(x|x_1)} \|u_t(x|x_1) - \tilde{v}_t(x; \theta)\|^2 \tag{4}$$

It reduces the calculation of the optimization criterion to the calculation of a mean-squared error between the network output $\tilde{v}_t(x; \theta)$ and an analytical solution u_t for sampled $t \sim \mathcal{U}(0, 1)$, $x_1 \sim q$ and $x \sim p_t(\cdot|x_1)$. Here, q is the probability distribution of the input data. A good choice of u_t and corresponding p_t is a Gaussian conditional probability path with mean and variance changing linear in time (optimal transport) [31]. The CFM-loss (4) then reduces even further

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t, q(x_1), p(x_0)} \left\| (x_1 - (1 - \sigma_{\min}) x_0) - \tilde{v}_t(\sigma_t x_0 + \mu_t; \theta) \right\|^2, \tag{5}$$

where $\mu_t = tx_1$, $\sigma_t = 1 - (1 - \sigma_{\min})t$, $p(x_0) = \mathcal{N}(0, 1)$ and σ_{\min} a small parameter, that can be chosen to match the noise level of the training data.

4.2 Conditional density estimation

Following the coupling-block flow based example of [26], we can extend CNFs to approximate a conditional density

$$p_t(x | c) = p_0 \left(\phi_t^{-1}(x, c) | c \right) \det \left(\frac{\partial \phi_t^{-1}(x, c)}{\partial x} \right), \tag{6}$$

where the noise distribution is independent of the condition $p_0(\cdot | c) = p_0(\cdot)$, by appending the vector of conditions to every layer of the vector field model $\tilde{v}_t(x, c; \theta)$. For our surrogate, x will be the ParT output and c will be the vector of jet-observables.

5 Bayesian Neural Networks

To indicate the application of the surrogate on data not included in tagger and thus surrogate training, we employ

Bayesian deep learning. Through modeling of (or sampling from) a posterior weight distribution

$$\pi(\theta | \mathcal{D}),$$

these methods give a large spread of predictions for data not included in the loss objective during training. This posterior distribution is the distribution of weights θ of the network $\tilde{v}_t(\cdot, \theta)$ given the training data

$$\mathcal{D} = \left\{ (x^{(1)}, c^{(1)}), (x^{(2)}, c^{(2)}), \dots \right\}.$$

Multiple instances from the weight posterior will form an ensemble of networks with differing weights. With both being conditional probability distributions, the weight posterior has to be distinguished from the likelihood of classifier output (1) that is to be inferred by every CNF in the ensemble. Sections 5.1 and 5.2 introduce two different approaches to connect both distributions.

5.1 Mean-field Gaussian variational inference (VIB)

A first way to relate the the weight posterior $\pi(\theta | \mathcal{D})$ to a CNF is to approximate it with an uncorrelated Normal distribution $\tilde{\pi}(\theta)$ [32]. This approximation is usually inferred during optimization of the network, by minimizing the Kullback–Leibler divergence (D_{KL}) between the posterior and its approximation

$$\begin{aligned} \mathcal{L}_{\text{VIB}} &= D_{\text{KL}} \left[\tilde{\pi}(\theta), \pi(\theta | \mathcal{D}) \right] \\ &= - \int d\theta \tilde{\pi}(\theta) \log \pi(\mathcal{D} | \theta) \\ &\quad + D_{\text{KL}} \left[\tilde{\pi}(\theta), \pi(\theta) \right] + \text{constant}, \end{aligned} \tag{7}$$

where $\pi(\theta)$ is the prior imposed on the network weights. Following the construction in [33], we bridge the gap between the CFM-loss (5) and the log-likelihood of the data in (7) by introducing a factor k that can be optimized to account for the difference

$$\mathcal{L}_{\text{VIB-CFM}} = \mathbb{E}_{\tilde{\pi}(\theta)} \mathcal{L}_{\text{CFM}} + k D_{\text{KL}} \left[\tilde{\pi}(\theta), \pi(\theta) \right]. \tag{8}$$

5.2 AdamMCMC

While the derivation of the loss (8) lacks theoretic backing and its optimization can take considerably longer than that of the CFM-loss (5) alone, the low dimensionality of the Classifier Surrogate problem allows us to directly sample the weight posterior distribution through Markov chain Monte Carlo (MCMC).

Full Hamiltonian Monte Carlo (HMC) is still often considered the gold-standard for inferring weight posteriors [34].

The large size of the training data however forces us to use stochastic MCMC algorithms. As one instance of this class, we choose AdamMCMC [35] due to its easy implementation. Competing algorithms, such as stochastic gradient HMC [36] or symmetric splitting HMC [37], will likely produce similar results.

We initialize the AdamMCMC-chain with a network optimized using the CFM-loss objective (5) and solve the ODE (2) to determine the negative log-likelihood \mathcal{L}_{NLL} of the data for every step of the MCMC from there on.

To ensure detailed balance we employ a Metropolis–Hastings (MH) correction with acceptance rate

$$\alpha = \frac{\exp(-\lambda \mathcal{L}_{\text{NLL}}(\tau_i)) q(\theta_i | \tau_i)}{\exp(-\lambda \mathcal{L}_{\text{NLL}}(\theta_i)) q(\tau_i | \theta_i)} \tag{9}$$

for all steps of the chain. Through the proportionality $\pi(\theta | \mathcal{D}) \propto -\mathcal{L}_{\text{NLL}}$ (Bayes formula) the acceptance step guarantees sampling from the weight posterior. Here, the parameter λ gives the inverse temperature of the tempered posterior distribution sampled from. The proposed weights τ_i are drawn from a proposal distribution centered on a gradient descent step

$$\tilde{\theta}_{i+1} = \text{Adam}(\theta_i, \mathcal{L}_{\text{NLL}}(\theta_i)) \tag{10}$$

calculated using the Adam algorithm [38]. We can use the momentum terms of the update to ensure high acceptance rates by smearing the proposal distribution in the direction of the last update

$$\tau_i \sim q(\cdot | \theta_i) = \mathcal{N}(\tilde{\theta}_{i+1}, \sigma^2 \mathbf{1} + (\tilde{\theta}_{i+1} - \theta_i)(\tilde{\theta}_{i+1} - \theta_i)^\top). \tag{11}$$

To efficiently run this algorithm, we evaluate the NLL on batches of data. For proofs on convergence and invariant distribution of this algorithm, we refer to [35].

6 Results

To learn the detector smearing distribution from data, we found a CNF with only 3 multi-layer perceptrons (MLPs) with 3 layers of dimension 64 and ELU activation to be sufficient. The condition and time variable t are concatenated to every MLP input, totaling in 31617 network parameters. Converting to VIB as in [32], doubles the number of parameters. We train on a balanced set of 4M jets in batches of 131,072 for 4000 epochs using Adam [38] with a constant learning rate of 10^{-3} . As loss objective, we use the CFM-loss as introduced in Eqs. (5) and (8) respectively. To achieve good coverage, we choose $c = 100$ and $\lambda = 50$ from a course grid search over multiple orders of magnitude.




Fig. 2 Histograms of 50,000 samples drawn from the detector smearing distributions learned with a CFM-model. Uncertainties are generated by drawing the samples from 11 points sampled from the network posterior approximation or chain. The ParT-output for the arbitrary QCD and top jet used as condition is indicated with dotted lines. Both jet events are the same as for Fig. 1

We run the AdamMCMC chain for another 1000 epochs with the learning rate reduced to $5 \cdot 10^{-6}$ and $\sigma = 0.05$. For the sampled posterior we always report the results from CFM-optimization in solid lines and the uncertainty calculated as the min-max-envelope of 10 drawings and for the learned approximation (VIB) we give the mean and the min-max-envelope over 11 sets of weights.

Using a fully-connected architecture, the sampled networks, either from the VIB-approximation or MCMC, can be easily exported as a serialized file using ONNX [9] at only 0.3 MB per instance. The the ODE defined by the network remains to be solved at inference time.

6.1 In-distribution

We can use the trained CNFs to generate another approximation of the detector smearing distribution by performing the forward direction starting at different points in latent space but for the same high-level features. Figure 2 shows histograms of the generated data for the same arbitrary jet events as Fig. 1.

We can see similar distributions for the approximation with CNFs as for the histograms of the closest events. The biggest discrepancy occurs between the distribution for the QCD jet obtained using AdamMCMC and VIB. It can be attributed to the difference between the model at initialization of the AdamMCMC chain and the posterior mean output of VIB. The initialization can be adapted to accommodate desired behaviours, if well defined, by choosing between different epochs of the CFM-optimization. Furthermore, increasing the chain length decreases the dependence of the ensemble output on the initialization overall.




Fig. 3 Empirical over nominal coverage calculated by taking 1000 samples from the learned detector smearing distribution for 10,000 jet events each. Uncertainties again are calculated from 11 points of the network posterior

6.1.1 Uncertainty calibration

To find out whether surrogate predictions using AdamMCMC are in general conservative, we need to look at the calibration of the estimated detector smearing distributions for multiple events, here 10,000. Per event we take 1000 samples from the inferred distribution and calculate q -quantiles for 50 values of q (nominal coverage) linearly spaced between 0 and 1. We then evaluate the empirical coverage, that is the fraction of corresponding ParT output within the respective quantile of the inferred distribution. The calibration is perfect when nominal and empirical coverage agree. Figure 3 shows very good calibration for both methods, where AdamMCMC in fact tends to be slightly more confident than VIB approximations.

6.1.2 Epistemic uncertainty

In contrast to uncertainty due to noisy data resulting in the detector smearing distributions, epistemic uncertainty is the uncertainty encoded in the variations within the ensembles of network weights induced by data sparsity. For a further dive into the behaviour of the epistemic uncertainty, we calculate the mean distance of the maximum discrepancy between instances of the network posterior

$$\delta_{\text{epis}} = \frac{1}{n_{\text{stat}}} \sum_{i=0}^{n_{\text{stat}}} \max_{p(\theta|\mathcal{D})} \phi_{1,\theta}(x_i) - \min_{p(\theta|\mathcal{D})} \phi_{1,\theta}(x_i) \quad (12)$$

for a total of $n_{\text{stat}} = 1000$ points drawn from the Gaussian latent space $x_1, \dots, x_{n_{\text{stat}}} \sim \mathcal{N}(0, 1)$. Ideally, this error estimate is large for sparsely populated areas of the high-level

feature space and small in the bulk of the distribution. To investigate this behaviour, we plot a histogram of the high-level features of the training data as well as δ_{epis} for 10,000 jet events chosen at random from a test set for both methods in Fig. 4.

The most instructive panels show the dependence of the error estimate on the number of constituents in the jet n_{const} , which is the most descriptive input feature. We can see high uncertainties occurring in the regions where the distributions for QCD and top jets overlap in the training data. These are events that can not easily be attributed to one of the two classes by the five high-level observables alone, resulting in high uncertainties. These events also make up the high-error bulk when plotted over the other high-level features.

For every tailed distribution, we can also see an increase of the error estimate for top jet predictions towards the edges of the data. This behaviour is stronger for AdamMCMC than for VIB at the cost of higher uncertainties overall.

The same behaviour is not observed for QCD jets. This again can be traced back to the distribution of n_{const} . The distribution of the number of particles of top jets is fully within the support of the one for QCD jets inducing high epistemic uncertainties for both highly and lowly populated jets. On the other hand, the distribution of top jets does not include events with as few particles as for QCD jets, allowing a perfect classification of these jets that dominates the low uncertainty edge of the plotted cloud.

6.1.3 Adding informative features

Another measure for the informative value of a detector smearing distribution generated by a Classifier Surrogate is the predicted accuracy

$$\hat{a} = \frac{1}{n_{\text{stat}}} \sum_{i=0}^{n_{\text{stat}}} \begin{cases} \mathbf{1}_{[0.5,1]}(\phi_{1,\theta}(x_i)) & \text{for top jets} \\ \mathbf{1}_{[0,0.5]}(\phi_{1,\theta}(x_i)) & \text{for QCD jets} \end{cases} \quad (13)$$

per jet event, with $\mathbf{1}_A(x)$ the indicator function of set A . The cut value of 0.5 is arbitrary and can be chosen in line with the experimental analysis. Our choice reflects the requirement to yield symmetric output distributions in case of uninformative high-level input.

Figure 5 shows histograms of the predicted accuracy for 10,000 jet events chosen at random from the full balanced test set. The distributions are generated from surrogates using the five high-level jet features from before, as well as for surrogates including the soft drop mass m_{SD} and the N -subjettiness for $N \in \{1, \dots, 4\}$. Naively, we assume that adding more information will lead to more certain predictions and thus will shift the distributions towards high accuracy values. In the highest value bin, the information hierarchy is well reproduced, with the highest number of input

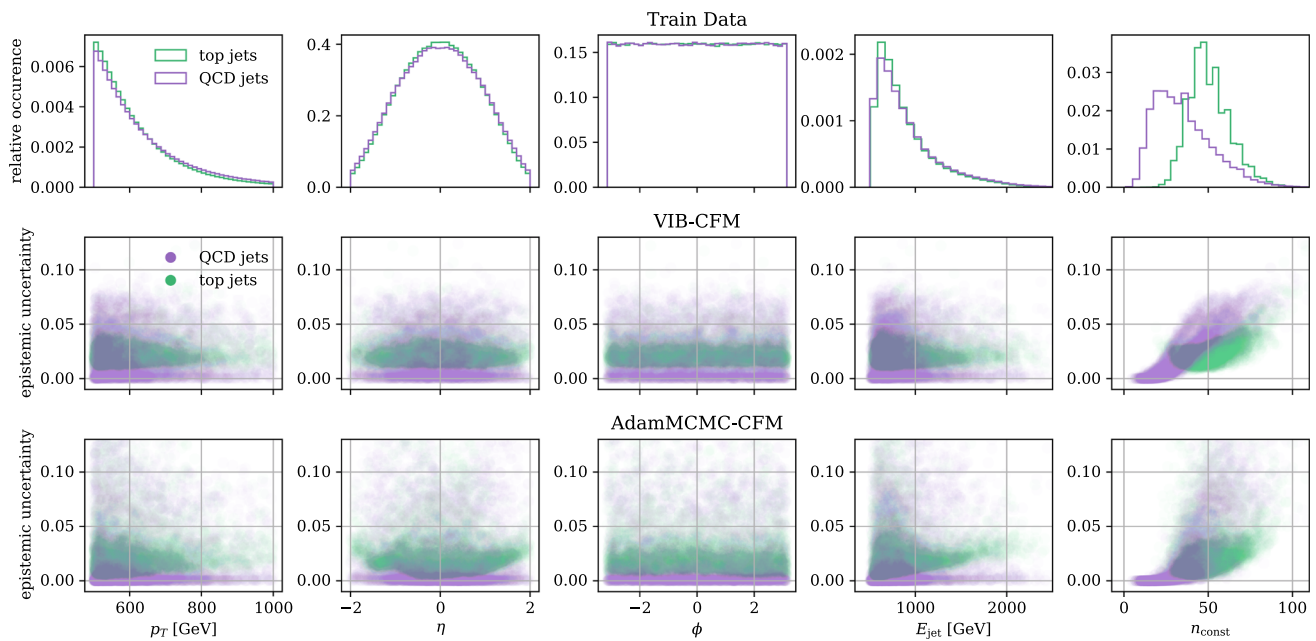


Fig. 4 Epistemic uncertainty calculated from the mean difference between 11 points from the network posterior over 1000 samples drawn from the respective learned detector smearing distribution for each event

of the validation set. The uncertainty shows a clear scaling towards the edges of the train data, as well as in regions where n_{const} is uninformative

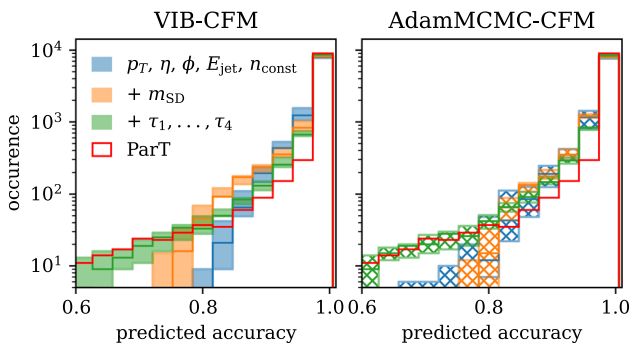


Fig. 5 Accuracy of 1000 ParT outputs predicted for each of 10,000 jet events. The different colors indicate the output of CNFs conditioned with increasing amount of features and thus provided with more information during inference. A histogram of the probabilistic ParT prediction itself is given in red

features leading to the highest number of certain outputs. In the range of 0.85 to 1, more informative input leads to fewer predictions in line with the naive assumption. For less certain predictions, a different effect can be observed. Increasing the information in the conditions allows the network to better model the ParT output, which features long tails of individual false positives and events predicted with low confidence. Thus, the Jensen-Shannon divergence between the histograms of surrogate and ParT output (Table 1) decreases with increasing number of input features.

Table 1 Jensen–Shannon-divergence between the histograms of predicted accuracies of Classifier Surrogates with different input features (Fig. 5) and the actual accuracy distribution of the ParT

JSD	VIB-CFM	AdamMCMC-CFM
$p_T, \eta, \phi, E_{\text{jet}}, n_{\text{const}}$	0.174 ± 0.018	0.147 ± 0.036
$+ m_{\text{SD}}$	0.134 ± 0.023	0.160 ± 0.013
$+ \tau_1, \dots, \tau_4$	0.080 ± 0.009	0.097 ± 0.007

6.2 Out-of-distribution

Although including an epistemic uncertainty into the evaluation this far is a nice feature to gauge uncertainties in the tail regions of the data, the true value of BNNs is indicating input that is outside the distribution of the training data by assigning high uncertainties. We use the introduced measures (12) and (13) to show the behaviour of the BNN surrogates for OOD data generated when artificially increasing the values for one jet-observable.

We produce OOD data by selecting 1000 jet events from the test set at random and increasing the values of a single jet-feature by adding a constant value. We perform this distortion for 3 dimensions, p_T, E_{jet} and n_{const} , and 10 values each. Again, we report the accuracy and error estimate calculated from $n_{\text{stat}} = 1000$ points of the learned detector smearing distribution.




Fig. 6 Behaviour of the Classifier Surrogate predictions for distorted input over the size of the distortion artificially added to the given jet-observable for 1000 random events. The first row shows the accuracy predicted per event while the second row gives the mean epistemic

The first row of Fig. 6 shows the mean accuracy predicted for the OOD data by the different drawings from the weight posterior. The envelope and solid line give the 10%- and 90%-quantile and the median over the set of events. When adding an unphysical offset to the features, we can see the mean predicted accuracy of the AdamMCMC ensemble rapidly drops. Optimally, the network predicts 0.5 when all inputs are outside the training interval to indicate equal confidence of both classes. The ensemble seems to be able to detect most outliers, but only indicates large distortions of E_{jet} for top jets and of n_{const} for QCD jets.

The predicted accuracy of the VIB samples does not exhibit any dependence on the increasing offset in the OOD data. It is sensitive only to the number of jet constituents for top jets.

In the second row, we show the error estimate based on the difference between highest and lowest proposed output in the ensemble, see Eq. 12. This measure captures the differences in the output and thus the encoded uncertainty directly. We expect increasing uncertainties for increasing offset. Only the AdamMCMC ensemble shows this behaviour, for all three disturbed input dimensions, while VIB once again is only sensitive to OOD inputs in the particle number. While the predicted accuracy did not capture the decreasing confidence for distorted E_{jet} of top jets well, the error estimate clearly indicates the unknown inputs. Similarly, distortions in n_{const} of QCD jets appear earlier in this measure.

7 Conclusion

In this paper, we proposed a first architecture for training Classifier Surrogates, which are models describing the output of a deep neural network classification based on detector-level information from high-level jet-observables and truth

uncertainty per event. Solid lines give the median of the set of events. The shaded and structured areas indicate the 10%–90%-quantile envelope of the VIB and AdamMCMC ensemble respectively

information. We show that the resulting Classifier Surrogates are well calibrated and scale with the amount of information provided. A combination with Monte Carlo generated samples from the networks Bayesian weight posterior allows for stable uncertainty quantification, that incorporates the density of the training data towards the edges. The predicted uncertainty reliably indicates unknown inputs.

This approach should next be implemented by the large experimental collaborations to allow the statistical re-interpretation of analysis results.

Acknowledgements We thank the organisers and participants of the Les Houches – PhysTeV 2023 workshop for the stimulating environment which led to the idea for this work. Furthermore, we thank Joschka Birk for helping with data preparation in the early stages of the work.

Funding SB is supported by the Helmholtz Information and Data Science Schools via DASHH (Data Science in Hamburg – HELMHOLTZ Graduate School for the Structure of Matter) with the grant HIDSS-0002. SB and GK acknowledge support by the Deutsche Forschungsgemeinschaft (DFG) under Germany’s Excellence Strategy – EXC 2121 Quantum Universe – 390833306. JK is supported by the Alexander-von-Humboldt-Stiftung. This research was supported in part through the Maxwell computational resources operated at Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany.

Data Availability Statement response This manuscript has associated data in a data repository. [Authors’ comment: The simulated data used in this work, is publicly available [20]. The exact handling of the data can be accessed at <https://github.com/joschkabirk/jetclass-top-qcd>.]

Code Availability Statement This manuscript has associated code/software in a data repository. [Authors’ comment: The code generated during this study is available from <https://github.com/sbieringer/ClassificationSurrogates>.]

Declarations

Code The training and evaluation code is available from <https://github.com/sbieringer/ClassificationSurrogates> and an example on handel-

ing the JetClass dataset is given in <https://github.com/joschkabirk/jetclass-top-qcd>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funded by SCOAP³.

References

1. D. Guest, K. Cranmer, D. Whiteson, Deep learning and its application to LHC physics. *Ann. Rev. Nucl. Part. Sci.* **68**, 161–181 (2018). <https://doi.org/10.1146/annurev-nucl-101917-021019>. arXiv:1806.11484 [hep-ex]
2. K. Albertsson et al., Machine learning in high energy physics community white paper. *J. Phys. Conf. Ser.* **1085**(2), 022008 (2018). <https://doi.org/10.1088/1742-6596/1085/2/022008>. arXiv:1807.02876 [physics.comp-ph]
3. A. Radovic et al., Machine learning at the energy and intensity frontiers of particle physics. *Nature* **560**(7716), 41–48 (2018). <https://doi.org/10.1038/s41586-018-0361-2>
4. G. Karagiorgi, G. Kasieczka, S. Kravitz, B. Nachman, D. Shih, Machine learning in the search for new fundamental physics (2021). arXiv:2112.03769 [hep-ph]
5. S. Kraml et al., Searches for new physics: Les Houches recommendations for the presentation of LHC results. *Eur. Phys. J. C* **72**, 1976 (2012). <https://doi.org/10.1140/epjc/s10052-012-1976-3>. arXiv:1203.2489 [hep-ph]
6. W. Abdallah et al., Reinterpretation of LHC results for new physics: status and recommendations after Run 2. *SciPost Phys.* **9**(2), 022 (2020). <https://doi.org/10.21468/SciPostPhys.9.2.022>. arXiv:2003.07868 [hep-ph]
7. J.Y. Araz et al., Les Houches guide to reusable ML models in LHC analyses (2023). arXiv:2312.14575 [hep-ph]
8. D.H. Guest et al., Lwtnn/lwtnn: Version 2.13. <https://doi.org/10.5281/zenodo.6467676>
9. Open Neural Network Exchange. <https://onnx.ai>
10. ATLAS collaboration, Search for R-parity-violating supersymmetry in a final state containing leptons and many jets with the ATLAS experiment using $\sqrt{s} = 13\text{TeV}$ proton–proton collision data. *Eur. Phys. J. C* **81**(11), 1023 (2021). <https://doi.org/10.1140/epjc/s10052-021-09761-x>. arXiv:2106.09609 [hep-ex]
11. ATLAS collaboration, Search for supersymmetry in final states with missing transverse momentum and three or more b-jets in 139 fb⁻¹ of proton–proton collisions at $\sqrt{s} = 13\text{TeV}$ with the ATLAS detector. *Eur. Phys. J. C* **83**(7), 561 (2023). <https://doi.org/10.1140/epjc/s10052-023-11543-6>. arXiv:2211.08028 [hep-ex]
12. ATLAS collaboration, Search for neutral long-lived particles in pp collisions at $\sqrt{s} = 13\text{TeV}$ that decay into displaced hadronic jets in the ATLAS calorimeter. *JHEP* **06**, 005 (2022). [https://doi.org/10.1007/JHEP06\(2022\)005](https://doi.org/10.1007/JHEP06(2022)005). arXiv:2203.01009 [hep-ex]
13. ATLAS collaboration, Anomaly detection search for new resonances decaying into a Higgs boson and a generic new particle X in hadronic final states using $\sqrt{s} = 13\text{TeV}$ pp collisions with the ATLAS detector. *Phys. Rev. D* **108**, 052009 (2023). <https://doi.org/10.1103/PhysRevD.108.052009>. arXiv:2306.03637 [hep-ex]
14. ATLAS collaboration, Performance of b-Jet Identification in the ATLAS Experiment. *JINST* **11**(04), 04008 (2016). <https://doi.org/10.1088/1748-0221/11/04/P04008>. arXiv:1512.01094 [hep-ex]
15. CMS collaboration, Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV. *JINST* **13**(05), 05011 (2018). <https://doi.org/10.1088/1748-0221/13/05/P05011>. arXiv:1712.07158 [physics.ins-det]
16. H. Qu, C. Li, S. Qian, Particle transformer for jet tagging, in *Proceedings of the 39th International Conference on Machine Learning*, pp. 18281–18292 (2022)
17. K. Cranmer et al., Publishing statistical models: getting the most out of particle physics experiments. *SciPost Phys.* **12**(1), 037 (2022). <https://doi.org/10.21468/SciPostPhys.12.1.037>. arXiv:2109.04981 [hep-ph]
18. J. Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, M. Selvaggi, Delphes 3: a modular framework for fast simulation of a generic collider experiment. *J. High Energy Phys.* (2014). [https://doi.org/10.1007/jhep02\(2014\)057](https://doi.org/10.1007/jhep02(2014)057)
19. T. Dorigo, P. De Castro Manzano, Dealing with Nuisance Parameters using Machine Learning in High Energy Physics: a Review (2020). arXiv:2007.09121 [stat.ML]
20. H. Qu, C. Li, S. Qian, JetClass: A Large-Scale Dataset for Deep Learning in Jet Physics. <https://doi.org/10.5281/zenodo.6619768>
21. A.J. Larkoski, S. Marzani, G. Soyez, J. Thaler, Soft drop. *JHEP* **05**, 146 (2014). [https://doi.org/10.1007/JHEP05\(2014\)146](https://doi.org/10.1007/JHEP05(2014)146). arXiv:1402.2657 [hep-ph]
22. J. Thaler, K. Van Tilburg, Identifying boosted objects with N-subjettiness. *JHEP* **03**, 015 (2011). [https://doi.org/10.1007/JHEP03\(2011\)015](https://doi.org/10.1007/JHEP03(2011)015). arXiv:1011.2268 [hep-ph]
23. S. Badger et al., Machine learning and LHC event generation. *SciPost Phys.* **14**(4), 079 (2023). <https://doi.org/10.21468/SciPostPhys.14.4.079>. arXiv:2203.07460 [hep-ph]
24. H. Hashemi, C. Krause, Deep generative models for detector signature simulation: an analytical taxonomy (2023). arXiv:2312.09597 [physics.ins-det]
25. C. Winkler, D.E. Worrall, E. Hoogeboom, M. Welling, Learning likelihoods with conditional normalizing flows. *CoRR* (2019). arXiv:1912.00042 [cs.lg]
26. S.T. Radev, U.K. Mertens, A. Voss, L. Ardizzone, U. Köthe, Bayesflow: learning complex stochastic models with invertible neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* **33**(4), 1452–1466 (2020). arXiv:2003.06281 [stat.ML]
27. J. Brehmer, F. Kling, I. Espejo, K. Cranmer, MadMiner: machine learning-based inference for particle physics. *Comput. Softw. Big Sci.* **4**(1), 3 (2020). <https://doi.org/10.1007/s41781-020-0035-2>. arXiv:1907.10621 [hep-ph]
28. S. Bieringer, A. Butter, T. Heimel, S. Höche, U. Köthe, T. Plehn, S.T. Radev, Measuring QCD splittings with invertible networks. *SciPost Phys.* **10**(6), 126 (2021). <https://doi.org/10.21468/SciPostPhys.10.6.126>. arXiv:2012.09873 [hep-ph]
29. R.T. Chen, Y. Rubanova, J. Bettencourt, D.K. Duvenaud, Neural ordinary differential equations. *Advances in neural information processing systems*, vol. 31 (2018). arXiv:1806.07366 [cs.LG]
30. D. Rezende, S. Mohamed, Variational inference with normalizing flows, in *International Conference on Machine Learning*, pp. 1530–1538. PMLR (2015)
31. Y. Lipman, R.T.Q. Chen, H. Ben-Hamu, M. Nickel, M. Le, Flow matching for generative modeling, in *The Eleventh International Conference on Learning Representations* (2023)
32. C. Blundell, J. Cornebise, K. Kavukcuoglu, D. Wierstra, Weight uncertainty in neural network, in *International Conference on Machine Learning*, pp. 1613–1622. PMLR (2015)
33. A. Butter et al., Jet diffusion versus JetGPT—modern networks for the LHC (2023). arXiv:2305.10475 [hep-ph]

34. P. Izmailov, S. Vikram, M.D. Hoffman, A.G.G. Wilson, What are Bayesian neural network posteriors really like? in *International Conference on Machine Learning*, pp. 4629–4640. PMLR (2021)
35. S. Bieringer, G. Kasieczka, M.F. Steffen, M. Trabs, AdamMCMC: Combining Metropolis adjusted Langevin with momentum-based optimization (2023). [arXiv:2312.14027](https://arxiv.org/abs/2312.14027) [stat.ML]
36. T. Chen, E. Fox, C. Guestrin, Stochastic gradient Hamiltonian monte carlo, in *International Conference on Machine Learning*, pp. 1683–1691. PMLR (2014)
37. A.D. Cobb, B. Jalaian, Scaling hamiltonian monte carlo inference for Bayesian neural networks with symmetric splitting, in *Uncertainty in Artificial Intelligence*, pp. 675–685. PMLR (2021)
38. D.P. Kingma, J. Ba, Adam: A method for stochastic optimization. CoRR (2014). [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) [cs.LG]