

# How to trust learned loop amplitudes

Henning Bahl<sup>1</sup>, Jens Braun<sup>2</sup>, Gudrun Heinrich<sup>2</sup>, Tilman Plehn<sup>1,3</sup> and Rebecca Revelli<sup>1,4</sup>

<sup>1</sup> Institut für Theoretische Physik, Universität Heidelberg, Germany

<sup>2</sup> Institute for Theoretical Physics, Karlsruhe Institute of Technology (KIT), Germany

<sup>3</sup> Interdisciplinary Center for Scientific Computing (IWR), Universität Heidelberg, Germany

<sup>4</sup> Dipartimento di Fisica, Università di Torino, Italy

## Abstract

Higher-order theory predictions are crucial for the precision LHC program, but the time-consuming amplitude evaluation challenges the corresponding Monte-Carlo simulations. Machine-learned amplitude surrogates can resolve this problem, if we can guarantee their precision over the entire phase space. First, we show that our surrogates provide a calibrated learned uncertainty, even for non-Gaussian systematics; second, we describe how less accurate phase space regions can be identified; third, we demonstrate how the precision in these regions can be improved reliably.



Copyright H. Bahl *et al.*

This work is licensed under the Creative Commons

[Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

Published by the SciPost Foundation.

Received 2026-01-19

Accepted 2026-05-12

Published 2026-06-16

doi:[10.21468/SciPostPhys.20.6.165](https://doi.org/10.21468/SciPostPhys.20.6.165)



Check for updates

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Efficient amplitude surrogates</b>	<b>3</b>
2.1	Non-Gaussian probabilistic regression	3
2.2	Systematic uncertainty shape	6
2.3	Low-accuracy regions	8
2.4	Adaptive sampling of training data	10
<b>3</b>	<b><math>HH</math> production</b>	<b>11</b>
3.1	One-loop with fixed top mass	12
3.2	One-loop with varying top-quark mass	15
3.3	Two-loop virtual amplitude	17
<b>4</b>	<b><math>t\bar{t}H</math> production</b>	<b>19</b>
4.1	Singularity structure	20
4.2	Leading order $gg \rightarrow t\bar{t}H$	21
4.3	Next-to-leading order $gg \rightarrow t\bar{t}H$	22
<b>5</b>	<b>Outlook</b>	<b>25</b>
<b>A</b>	<b>Hyperparameters</b>	<b>26</b>
<b>B</b>	<b>Forward/backward regions: NLO to LO ratio</b>	<b>27</b>

<b>C Accuracy close to Coulomb singularity</b>	<b>27</b>
<b>D Clustering based on estimated uncertainties</b>	<b>28</b>
<b>E Extracted KDE estimates</b>	<b>29</b>
<b>F <math>q\bar{q} \rightarrow t\bar{t}H</math> results</b>	<b>29</b>
<b>References</b>	<b>31</b>

---

## 1 Introduction

The rapidly increasing amount of LHC data allows analyses and measurements with unprecedented precision, provided the experimental precision is matched by the precision of the theoretical predictions [1]. This requires simulations including higher-order perturbative corrections, both loop amplitudes and efficient schemes to treat unresolved real radiation. In this regard, we have seen remarkable progress in both QCD and electroweak (EW) corrections. For LHC loop amplitudes, the current frontier are 3-loop box amplitudes with up to two off-shell legs and 2-loop pentagon amplitudes with several mass scales [2–6]. These kinds of precision predictions have to be numerically available in public codes if we want to benefit from optimal analysis methods, like simulation-based inference [7–14].

However, the availability of analytic or (semi-)numerical loop amplitudes does not guarantee that they can be used for large-scale simulations. Important limitations are the accuracy and speed at which the higher-order amplitudes can be evaluated in practice. Here, we hope that modern machine learning (ML) [15, 16] will help with three distinct numerical challenges: efficient phase space sampling [17–27], efficient subtraction schemes for real radiation [27], and virtual multi-loop amplitudes [28–44].

For virtual loop amplitudes, the numerical complexity rises not only with the number of loops, but also with the number of kinematic scales, related to the number of external legs and their virtuality, and mass scales. At two loops and beyond, their direct evaluation inside the simulation code is often not feasible. To construct a surrogate amplitude, the full higher-order amplitude is first used to generate a reference training dataset. This data is then used to construct or learn a surrogate. An interesting aspect of this surrogate is that for training data without significant numerical noise, the network should interpolate rather than fit a function [16]. For  $2 \rightarrow 2$  amplitudes, we can use two-dimensional interpolation grids [45–50]. For  $2 \rightarrow 3$  amplitudes, even with fixed masses, five-dimensional grids are challenging in the required number of data points and in the interpolation techniques. Ultra-fast ML-surrogates have been shown to be more accurate than traditional techniques [39]. A natural application are two-loop amplitudes for  $t\bar{t}H$  production, where partial results have been calculated in Refs. [51–55].

For precision simulations, it is crucial that we control surrogate uncertainties. In addition to globally defined numerical noise, a given phase space region might be poorly described, leading to a biased prediction of the observables in this region. A comparably expensive solution is to use the surrogate amplitude just to increase the unweighting efficiency while still using the true amplitude during event generation [26, 35]. An alternative approach avoiding the evaluation of the true amplitude during event generation — particularly useful for multi-loop amplitudes — is to equip the surrogate with a calibrated uncertainty estimate [34, 38, 43, 44, 56], and then ensure that the largest uncertainties do not correspond to a localized failure mode in phase space. Depending on the failure mode, a correctly learned uncertainty can then be used to improve the training data and training procedure [34, 44].

The paper is structured as follows. In Sec. 2 we introduce a set of conceptual improvements to probabilistic (amplitude) regression: in Sec. 2.1 we review the existing concepts and introduce a Student’s  $t$ -likelihood approach for non-Gaussian problems; in Sec. 2.2 we show how Gaussian systematics only appear for phase space dimensions beyond  $2 \rightarrow 2$  scattering; in Sec. 2.3 we introduce a way to analyze if low-accuracy amplitudes cluster in phase space; and in Sec. 2.4 we propose an adaptive two-step sampling procedure to optimize the network training locally. We demonstrate these new techniques for di-Higgs production at leading order (LO) and next-to-leading order (NLO) in QCD in Sec. 3. We then illustrate the performance gain for top-associated Higgs production at NLO as an example for a more challenging  $2 \rightarrow 3$  process in Sec. 4. In Sec. 5 we conclude that this phase space-controlled training will allow us to use ML-surrogate amplitudes for perturbative precision simulations.

## 2 Efficient amplitude surrogates

The ML-task to provide ultrafast loop amplitude surrogates is to learn an amplitude  $A_{\text{NN}}$  approximating the amplitude  $A_{\text{true}}$  for every phase space point  $x$ ,

$$A_{\text{NN}}(x) \approx A_{\text{true}}(x). \quad (1)$$

Here,  $A_{\text{true}}$  is calculated from first principles at a given order in perturbation theory. Ultra-violet (UV) and infrared (IR) singularities need to be removed before an amplitude can be used in a Monte Carlo simulation. We typically learn UV-renormalized and IR-subtracted amplitudes [45, 55]. The finite amplitude is exact up to negligible numerical noise, originating for example from the evaluation of special loop functions or from numerical integration over Feynman parameters. Our goal is to replace the amplitudes in the Monte Carlo generator entirely by the ML-surrogate.

### 2.1 Non-Gaussian probabilistic regression

A statistically sound surrogate describes the probability  $p(A|x)$ . The target probability depends implicitly on the training data  $D_{\text{train}} = \{(A_{\text{train}}, x_{\text{train}})\}$ . We can encode this probability in a set of network parameters  $\theta$ , generalizing beyond the training data, as

$$p(A|x) = \int d\theta p(A|x, \theta) p(\theta | D_{\text{train}}) \approx \int d\theta p(A|x, \theta) q(\theta), \quad (2)$$

where we replaced the true but normally intractable posterior with an approximate distribution  $q(\theta)$ . Here, we omit the explicit conditioning of  $q$  on  $D_{\text{train}}$  for notation simplicity. We compute the mean and variance as

$$\begin{aligned} A_{\text{NN}}(x) &= \int dA A p(A|x) \\ &= \int d\theta q(\theta) \bar{A}(x, \theta), \quad \text{with } \bar{A}(x, \theta) = \int dA A p(A|x, \theta), \\ \sigma_{\text{tot}}^2(x) &= \int dA [A - A_{\text{NN}}(x)]^2 p(A|x) \\ &= \int d\theta q(\theta) [\bar{A}^2(x, \theta) - \bar{A}(x, \theta)^2 + (\bar{A}(x, \theta) - A_{\text{NN}}(x))^2] \\ &\equiv \sigma_{\text{sys}}^2(x) + \sigma_{\text{stat}}^2(x). \end{aligned} \quad (3)$$

In the last step, we split the variance into a systematic and a statistical contribution,

$$\begin{aligned}\sigma_{\text{syst}}^2(x) &= \int d\theta q(\theta) \left[ \overline{A^2}(x, \theta) - \overline{A}(x, \theta)^2 \right], \\ \sigma_{\text{stat}}^2(x) &= \int d\theta q(\theta) \left[ \overline{A}(x, \theta) - A_{\text{NN}}(x) \right]^2,\end{aligned}\tag{4}$$

where  $\overline{A^2}(x, \theta) = \int dA A^2 p(A|x, \theta)$ .

The systematic uncertainty encodes the data-intrinsic uncertainty as well as the uncertainty induced by a lack of model expressivity. The data inherent uncertainty vanishes for noise-free data — i.e. for  $p(A|x, \theta) \rightarrow \delta(A(x) - A_0(x))$  with  $A_0$  being the noiseless value of the amplitude — if the model is fully expressive and perfectly trained. It also captures the uncertainty induced by a lack of model expressivity. The statistical uncertainty represents the uncertainty due to the limited size of the training dataset and vanishes for an infinitely large training dataset and perfect training, i.e. if  $q(\theta) \rightarrow \delta(\theta - \theta_0)$ . It is also referred to as reducible uncertainty.

The statistical uncertainty can be inferred using Bayesian neural networks (BNNs) [57–60], repulsive ensembles [12, 38, 43, 56, 61, 62], or evidential regression [43, 63–66]. We have benchmarked these methods for amplitude regression [38, 43] and found the statistical uncertainties to be small compared to the systematics. Consequently, we will focus on the systematic uncertainties.

### Gaussian systematics

To train our surrogate we minimize the negative log-likelihood over the training dataset,

$$\mathcal{L} = -\log p(A|x, \theta).\tag{5}$$

To capture the systematic uncertainty reflected in the variability of the training data, the simplest ansatz is a Gaussian likelihood,

$$p(A|x, \theta) = \mathcal{N}(A | \overline{A}(x, \theta), \sigma^2(x, \theta)),\tag{6}$$

where  $\overline{A}(x, \theta)$  and  $\sigma^2(x, \theta)$  are learned and  $\mathcal{N}(\mu, \sigma^2)$  denotes a Gaussian with mean  $\mu$  and variance  $\sigma^2$ . The Gaussian ansatz works well for amplitude regression [38, 43] and can be validated via pull distributions. The Gaussian likelihood leads us to the heteroskedastic loss

$$\mathcal{L}_{\text{het}} = \frac{[\overline{A}(x, \theta) - A_{\text{true}}(x)]^2}{2\sigma_{\text{syst}}^2(x, \theta)} + \log \sigma_{\text{syst}}(x, \theta).\tag{7}$$

For each phase space point, the training minimizes the numerator by learning the true amplitude. The second logarithmic contribution encourages the network to adjust  $\sigma_{\text{syst}}(x)$  according to the observed deviation. Consequently, large values of  $\sigma_{\text{syst}}(x)$  are predicted if the network struggles to learn the amplitude accurately.

### Non-Gaussian systematics

If the simple Gaussian likelihood does not describe the variability of the training data, it can be replaced by a Gaussian mixture model (GMM) [43, 56],

$$p_{\text{GMM}}(A|x, \theta) = \sum_k \omega_k(x, \theta) \mathcal{N}(A | \overline{A}_k(x, \theta), \sigma_k^2(x, \theta)), \quad \text{with} \quad \sum_k \omega_k(x, \theta) = 1.\tag{8}$$

The mean and standard deviation become

$$\begin{aligned}\bar{A}_{\text{GMM}}(x, \theta) &= \sum_k \omega_k(x, \theta) \bar{A}_k(x, \theta), \\ \sigma_{\text{GMM}}^2(x, \theta) &= \sum_k \omega_k(x, \theta) \left[ \sigma_k^2(x, \theta) + \bar{A}_k^2(x, \theta) \right] - \bar{A}_{\text{GMM}}^2(x, \theta).\end{aligned}\quad (9)$$

The GMM loss does not have a simple analytic form and has to be computed numerically,

$$\mathcal{L}_{\text{GMM}} = -\log \left[ \sum_k \frac{\omega_k(x, \theta)}{\sqrt{2\pi\sigma_k^2(x, \theta)}} \exp \left[ -\frac{[\bar{A}_k(x, \theta) - A_{\text{true}}(x)]^2}{2\sigma_k^2(x, \theta)} \right] \right]. \quad (10)$$

If the task is to just describe non-Gaussian tails, the GMM can be simplified by identifying the means,

$$\bar{A}_k(x, \theta) = \bar{A}(x, \theta). \quad (11)$$

This significantly simplifies the calculation of confidence intervals.

An alternative approach to non-Gaussian likelihoods is a Student's  $t$ -distribution

$$\begin{aligned}p_{\text{St}}(A | x, \theta) &= \text{St}(A | \bar{A}(x, \theta), \sigma^2(x, \theta), \nu(x, \theta)) \\ &= \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)} \frac{1}{\sqrt{\pi\nu}\sigma(x, \theta)} \left[ 1 + \frac{[\bar{A}(x, \theta) - A_{\text{true}}(x)]^2}{\nu\sigma^2(x, \theta)} \right]^{-\frac{\nu+1}{2}},\end{aligned}\quad (12)$$

where  $\Gamma$  is the Gamma function. Two limits of the Student's  $t$ -distribution are

$$p_{\text{St}}(A | \bar{A}, \sigma^2, \nu) = \begin{cases} \frac{\sigma}{\pi} \frac{1}{(A - \bar{A})^2 + \sigma^2}, & \nu \rightarrow 1, \\ \mathcal{N}(A | \bar{A}, \sigma^2), & \nu \rightarrow \infty, \end{cases} \quad (13)$$

where the Breit-Wigner or Cauchy distribution allows the corresponding likelihood to describe much larger tails. For  $\nu > 2$ , the variance of the Student's  $t$ -distribution is

$$\sigma_{\text{St}}^2(x, \theta) = \frac{\nu(x, \theta)}{\nu(x, \theta) - 2} \sigma^2(x, \theta), \quad (14)$$

and the corresponding negative log-likelihood loss reads

$$\begin{aligned}\mathcal{L}_{\text{St}} &= \frac{1 + \nu(x, \theta)}{2} \log \left[ 1 + \frac{[\bar{A}(x, \theta) - A_{\text{true}}(x)]^2}{\nu(x, \theta)\sigma^2(x, \theta)} \right] + \log \sigma(x, \theta) \\ &\quad + \frac{1}{2} \log(\pi \nu(x, \theta)) + \log \Gamma\left(\frac{\nu(x, \theta)}{2}\right) - \log \Gamma\left(\frac{\nu(x, \theta) + 1}{2}\right).\end{aligned}\quad (15)$$

Compared to a Gaussian likelihood, only one additional quantity,  $\nu(x, \theta)$ , needs to be learned.

### Architecture and preprocessing

For the surrogate architecture we use a simple MLP, consisting of 5 hidden layers with 512 hidden channels and GELU activation. The surrogates learn the logarithm of the amplitude, scaled to zero mean and unit standard deviation [43]. The inputs are the 4-momenta  $p_i$  of the external particles, complemented by the Lorentz invariants

$$z_{ij} = \log(p_i \cdot p_j). \quad (16)$$

All network inputs are standardized, and the hyperparameters are summarized in App. A. For this preprocessing we need to guarantee that all  $p_i \cdot p_j$  are positive. We use the convention that the sum of the initial-state 4-momenta is equal to the sum of the final-state 4-momenta,  $\sum_i p_i = \sum_f p_f$ . The energy components of all physical 4-vectors are positive, as well as  $E_i \geq |\vec{p}_i|$ , with the equality fulfilled for massless particles. This gives us

$$p_i \cdot p_j = E_i E_j - |\vec{p}_i| |\vec{p}_j| \cos \theta_{ij} \geq 0. \quad (17)$$

We also learn the systematic uncertainty using the heteroskedastic loss of Eq.(7) or its GMM variant as  $\log \sigma_{\text{sys}}$ , to cover large variations. For the same reason, we learn  $\log \nu$  when using the Student's  $t$ -likelihood ansatz.

Advanced architectures like equivariant transformers have advantages for processes with a substantial number of identical final state and/or initial state particles, since they allow to implement permutation invariance into the NN architecture [37,41]. For the processes studied in this paper, permutation invariance does not yield a substantial performance improvement [39].

## 2.2 Systematic uncertainty shape

For precision physics applications it is important that we understand the shape of the systematic uncertainties and then check the calibration of the learned uncertainties. This ensures that amplitude values with low accuracy are indeed those with large learned uncertainties, an aspect we will need below.

### Accuracy and calibration

To measure the accuracy of the amplitude surrogates, we define

$$\Delta(x) = \frac{A_{\text{NN}}(x) - A_{\text{true}}(x)}{A_{\text{true}}(x)}. \quad (18)$$

To test the calibration of the uncertainties, we define the systematic pull

$$t_{\text{sys}}(x) = \frac{A_{\text{NN}}(x) - A_{\text{true}}(x)}{\sigma_{\text{sys}}(x)}. \quad (19)$$

If the shape of this pull distribution is not Gaussian, the Gaussian assumption used for deriving the heteroskedastic loss might not be justified.

When a non-Gaussian likelihood is appropriate, we cannot rely on the pull distribution to test the uncertainty calibration. Instead, we evaluate the empirical coverage. For a test dataset we evaluate for what fraction of the true amplitude the surrogate lies within a given  $\gamma$  confidence region. This defines the empirical coverage

$$c_\gamma = \left\langle \mathbb{1} \left( p_{\text{Gauss}}(A_{\text{true}}(x) | A_{\text{NN}}(x), \sigma_{\text{sys}}^2(x)) > 1 - \gamma \right) \right\rangle_{D_{\text{test}}}, \quad (20)$$

which is here written down for the Gaussian case. The indicator function  $\mathbb{1}$  is one if the associated condition is met and zero otherwise.  $p_{\text{Gauss}}$  is the  $p$ -value for the true amplitude given the learned phase space dependent Gaussian probability distribution,

$$p_{\text{Gauss}}(A_{\text{true}}(x) | A_{\text{NN}}(x), \sigma_{\text{sys}}^2(x)) = 2 \left[ 1 - \Phi \left( \frac{|A_{\text{true}}(x) - A_{\text{NN}}(x)|}{\sigma_{\text{sys}}(x)} \right) \right], \quad (21)$$

where  $\Phi$  is the cumulative Gaussian distribution function. If the empirical coverage is larger than the nominal confidence level, the learned uncertainty is conservative or underconfident. If it is smaller, the learned uncertainty is overconfident.

If instead of the Gaussian likelihood we use a GMM or a Student's  $t$ -distribution, the  $p$ -value and cumulative distribution functions in Eq. (20) and (21) have to be replaced by their respective equivalents for the GMM or Student's  $t$ -distribution.

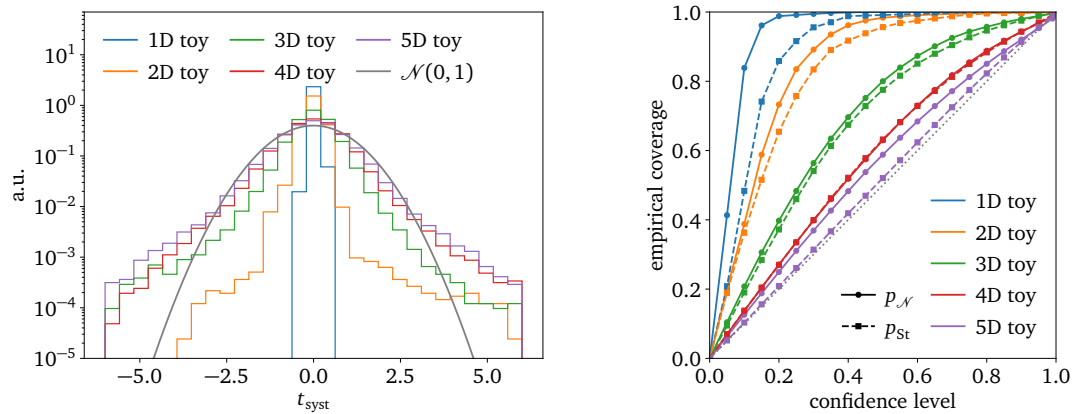


Figure 1: Left: systematic pull of toy models. Right: empirical coverages, comparing Gaussian and Student’s  $t$ -likelihoods.

### Central limit theorem

Above, we have seen that statistical regression is greatly simplified when the variability of the training data is captured by a Gaussian likelihood. If needed, we can capture non-Gaussian behavior using a GMM or a Student’s  $t$ -likelihood. To study the impact of the dimensionality, we start with a simple  $n$ -dimensional toy model where we regress the sum of the input coordinates drawn from the  $n$ -dimensional hypercube

$$f(x) = \sum_{i=1}^n x_i, \quad \text{with } x_i \sim \mathcal{U}(0, 1). \quad (22)$$

The one-dimensional output of the surrogate  $f_{\text{NN}}$  will effectively approximate the contribution of each dimension with a given uncertainty,

$$f_{\text{NN}}(x) = \sum_{i=1}^n (x_i + \epsilon_i(x)), \quad (23)$$

where we assume independent  $\epsilon_i$ . In that case, the deviation or residual  $r(x)$  is

$$r(x) \equiv f_{\text{NN}}(x) - f(x) = \sum_{i=1}^n \epsilon_i(x), \quad (24)$$

and the central limit theorem applies. The distribution of the residuals  $r(x)$  converges towards a Gaussian,

$$r(x) \sim \mathcal{N}(\langle r \rangle, \sigma_r^2), \quad \text{for } n \rightarrow \infty. \quad (25)$$

We confirm this behavior by training simple MLP surrogates with a heteroskedastic loss to reproduce  $f(x)$  for  $n \leq 5$  using  $10^5$  training events. We list the used NN and training setting in App. A. We show the systematic pulls in the left panel of Fig. 1. Although the systematic uncertainty is overestimated for small  $n$ , the calibration of the learned Gaussian uncertainty improves for higher dimensions, just leaving the tails being too large.

The convergence towards a Gaussian likelihood is also visible for the empirical coverage in the right panel of Fig. 1. For few dimensions the uncertainties are clearly overestimated, while for larger dimensionality the empirical coverage shifts towards the diagonal. We also show the results for a Student’s  $t$ -likelihood, given in Eq.(15). Within numerical uncertainties, relaxing the Gaussian assumption can lead to an improved uncertainty calibration in particular for low dimensions.




Figure 2: Left: systematic pulls. Right: coverage of  $Z + n_g g$  surrogates. The uncertainty bars for the  $Zggg$  surrogates indicate the mean and standard deviation obtained by averaging over five independent runs.

A similar argument can be made for general target functions. First, if the true function is a product of input components we look at  $\log f_{\text{NN}}(x) - \log f(x)$ , so the differences or residuals follow a log-normal distribution. For small residuals, it approaches a normal distribution. More generally, the central limit theorem applies when the residuals can be represented as a sum of an increasing number of independent errors. This is true for a wide range of regression problems, for a strict mathematical discussion for linear regression we refer to Ref. [67].

### Amplitudes for $Z$ plus gluons

Next, we test the applicability of the central limit theorem for amplitudes for the production of a  $Z$ -boson in association with gluons [37, 41],

$$q\bar{q} \rightarrow Z + \{1, \dots, 4\}g. \tag{26}$$

Cuts are applied such that all final-state objects have a transverse momentum of at least 20 GeV and such that the angular separation  $\Delta R$  is larger than 0.4 for all pairs of final-state objects.

We use 70k training points for each multiplicity. The phase space dimensionalities for  $n_g$  final-state gluons are  $3(n_g + 1) - 4$ . As usual, the relative accuracies degrade for an increasing number of gluons. The corresponding pull distributions in the left panel of Fig. 2 become increasingly Gaussian for an increasing number of gluons and hence phase space dimensions. Like for the toy model, the remaining challenge consists in increased tails for higher-dimensional regression. Finally, we show the empirical coverage. The uncertainty is overestimated for a small number of gluons, again as previously observed in the toy model. For  $n_g > 1$  only a slight underestimate of the uncertainties remains. We do not observe any benefit from using a Student's  $t$ -likelihood. The error bars for the  $Zggg$  surrogate indicate that the empirical coverage curves are stable.

### 2.3 Low-accuracy regions

If low accuracy or large uncertainties of the surrogate amplitude represent distinct failure modes, the main question is if this failure is of statistical or systematic origin. In the latter case, the network typically fails to learn a feature and the points with low accuracy and large uncertainties cluster in phase space. Identifying such clusters will allow us to improve the network training, for instance by adding more training data, see Sec. 2.4.

### Phase space metrics

To define low-accuracy clusters or regions, we need a distance metric on our phase space. We define two different distance metrics, one based directly on the 4-vectors of the external particles and one based on the kinematic degrees of freedom.

- The first metric between two phase space points with  $N_f$  final-state particles is defined as

$$d_{\text{kin}} = \frac{\sqrt{\sum_{i=1}^{N_f} |\vec{p}_i^{(1)} - \vec{p}_i^{(2)}|^2}}{\frac{1}{2N_f} \sum_{i=1}^{N_f} (E_i^{(1)} + E_i^{(2)})}. \quad (27)$$

The initial-state particles are excluded from the sum because they are always aligned with the beam axis. We divide by the sum of the energies because otherwise high-energy events would always be more separated from each other than low-energy ones. The advantage of this metric is that it can be defined straightforwardly. An obvious disadvantage is that it is not Lorentz-invariant, even though the underlying amplitude is. Moreover, since this metric does not account for the symmetries of the system, it can distinguish physically equivalent configurations, leading to redundant clusters. Therefore  $d_{\text{kin}}$  was only considered as a first exploratory approach and is not used in the final study.

- The second metric is defined in terms of given kinematic directions  $x_i$ ,

$$d_x = \left[ \sum_i (x_i^{(1)} - x_i^{(2)})^2 \right]^{1/2}. \quad (28)$$

We define these directions explicitly for the  $q\bar{q} \rightarrow t\bar{t}H$  and  $gg \rightarrow t\bar{t}H$  processes in Sec. 4. The number of  $x_i$  directions does not have to correspond to the phase space dimension, as some directions add more numerical noise than benefit. If we choose the kinematic directions as Lorentz invariants, so is the metric.

For our study,  $d_x$  provides more stable and physically meaningful results, so we use it throughout Sec. 4.

### Clustering algorithm

To identify where amplitude results with low accuracy or large uncertainties are clustered, we use the HDBSCAN algorithm [68, 69]. It is a density-based clustering method that does not require a pre-defined number of clusters and that can identify outliers rather than forcing all points into clusters.

As weights for the edges of a graph made of phase space points, the algorithm uses the reachability distance

$$d_{\text{mreach-}k}(a, b) = \max[\text{core}_k(a), \text{core}_k(b), d(a, b)], \quad (29)$$

where  $\text{core}_k(a)$  is the distance of  $a$  to its  $k$ th nearest neighbor. Dense points with low core distances keep their original distance  $d(a, b)$ . For sparse points the reachability distance is at least the maximum core distance. This way, the algorithm is robust against single points forming bridges between separate clusters:

1. First, HDBSCAN constructs the minimum spanning tree (MST) as a unique graph representation of a set of phase space points. It involves the minimal number of edges, such that dropping any further edge disconnects the graph. Moreover, it ensures that there is no edge with a lower reachability distance that could provide an alternative connection of the phase space points.
2. We then convert the MST into an ordered list of connected clusters. For this, we sort the MST edges by their reachability distance. We then go through the distance values in decreasing order and cut all connections above this distance. Each edge removal splits the MST into two clusters. If one of the resulting new clusters has fewer phase space points than a predefined minimum size, it is labeled an outlier and removed from the tree. If both new clusters have more events than the minimum size, they are kept as a valid cluster split.
3. The actual clusters are selected from the above cluster candidates as those that survive the longest during this edge-removal procedure. The notion of time is the inverse distance  $1/d_{\text{mreach-}k}$ . If a cluster is created at inverse  $d_{\text{mreach-}k}^{\text{birth}}$ , its stability is defined as

$$\text{stability} = \sum_{e \in \text{cluster}} \left( \frac{1}{d_{\text{mreach-}k}^{(e)}} - \frac{1}{d_{\text{mreach-}k}^{\text{birth}}} \right), \quad (30)$$

where the sum runs over all points  $e$  in the cluster. Based on this stability, we select the most stable clusters.

4. The one additional constraint we apply is that a stable cluster cannot be a descendant of an already selected cluster. This uniquely defines a set of clusters and outlier events.

## 2.4 Adaptive sampling of training data

Even though the training data distribution is a key ingredient to precision regression, we usually just use some kind of unweighted phase space points. Known modifications are boosted training [34] and uniform sampling [39]. In both cases, the full training dataset is still generated before we begin with the network training. Whereas boosted training reweights an existing training dataset to optimize the surrogate accuracy, expensive loop amplitudes require a maximally efficient training. Therefore, we split the generation of the training data into two sets and only start with the usual unweighted events. We then identify phase space regions where the surrogate lacks precision, using an independent test dataset or based on the learned uncertainty. For these regions, we generate additional training data from the amplitudes.

The advantage of this adaptive procedure can be understood by considering for example a flat distribution with a dip. A flat sample will have the same density of points in the flat region as in the dip region. However, the surrogate will likely need far fewer points in the flat region to describe it to the same accuracy as the dip region. Only generating training data in the critical dip region will therefore increase the overall accuracy more than placing some of the additional points in the flat region.

### Adaptive phase space sampling with kernel density estimation (KDE)

When we generate the dataset for this second step, we need to take into account the accuracy of the surrogate in different phase space regions. The second sample should favor low-accuracy regions through an appropriate probability density function (PDF). We construct it using a weighted kernel density estimation, with weights assigned according the surrogate's accuracy.

A general multivariate kernel estimator reads [70]

$$f(x) = \frac{1}{\det H \cdot \sum_i w_i} \sum_i w_i K(H^{-1}(x - x_i)). \quad (31)$$

The  $x_i$  are the  $n$  points of the original sample,  $K$  a kernel function, and  $H$  the bandwidth matrix. We assign a different weight  $w_i$  to each point. The kernel function and bandwidth matrix can be chosen freely, and the resulting PDF can significantly depend on their choice. However, as long as the PDF favors the correct phase space regions, we are not interested in its precise shape. This allows us to work with a simple  $d$ -dimensional Gaussian [71],

$$K(x) = \frac{1}{\sqrt{(2\pi)^d \det V}} \exp\left(-\frac{1}{2}x^T V^{-1}x\right), \quad (32)$$

where  $V$  is the covariance matrix of the initial sample. With the data covariance being included in the kernel, we use a scalar bandwidth following Scott's rule,

$$h = n_{\text{eff}}^{-1/(d+4)}, \quad \text{with} \quad n_{\text{eff}} = \frac{(\sum_i w_i)^2}{\sum_i w_i^2}, \quad (33)$$

the effective dataset size.

Finally, we choose the absolute relative accuracy  $|\Delta_i| = |\Delta(x_i)|$  of Eq.(18) as weights. The kernel density estimation PDF then reads

$$f(x) = \frac{1}{h^d \sum_i |\Delta_i|} \frac{1}{\sqrt{(2\pi)^d \det V}} \sum_i |\Delta_i| \exp\left[-\frac{1}{2h^2}(x - x_i)^T V^{-1}(x - x_i)\right]. \quad (34)$$

The second sample, drawn from this PDF, will cluster in regions with small surrogate accuracy and neglect regions with large accuracy. Because the Gaussian kernel falls off exponentially, it will not explore unknown regions, but only revisit regions already seen in the initial sample. Therefore, this method can only be used when the initial sample is sufficiently large.

### 3 HH production

As a first instructive and phenomenologically relevant example, we apply our advanced amplitude learning to di-Higgs production via gluon fusion. As a  $2 \rightarrow 2$  process, it only depends on two independent kinematic invariants,

$$\hat{s} = (p_1 + p_2)^2, \quad \text{and} \quad \hat{t} = (p_1 - p_a)^2, \quad (35)$$

where  $p_{1,2}$  are the 4-momenta of the incoming gluons and  $p_{a,b}$  are the 4-momenta of the outgoing Higgs bosons.

The differential cross-section is given by

$$d\sigma = \frac{1}{2\hat{s}} |\mathcal{M}|^2 d\Phi d\rho_{a,b}(\hat{s}, s), \quad (36)$$

where  $\hat{s}$  is the partonic center-of-mass energy,  $d\Phi$  is the element of the Lorentz-invariant phase space and  $d\rho_{a,b}$  is the probability of finding partons  $a$  and  $b$  with total energy  $\sqrt{\hat{s}}$  in the colliding protons. The surrogate amplitude is defined in Eq.(1). Depending on the optimal pre-processing it is derived from  $|\mathcal{M}|^2$  or from  $d\sigma$ , for our di-Higgs production setup it is

$$A(x) = |\mathcal{M}(x)|^2. \quad (37)$$

We parameterize the  $(2 \rightarrow 2)$ -phase space  $x$  in terms of the dimensionless quantities [45]

$$\beta_H = \sqrt{1 - \frac{4m_H^2}{\hat{s}}} \in [0, 1], \quad \text{and} \quad \cos \theta = \frac{\hat{t} - \hat{u}}{\hat{s}\beta_H}, \quad (38)$$

where  $\hat{u} = (p_1 - p_b)^2$ .

A characteristic feature, stemming from the top-quark loop mediating this process at leading order, is the (virtual) top quark pair production threshold at  $\hat{s} = 4m_t^2$ . Due to the onset of an imaginary part of the amplitude, the differential cross section as a function of  $\hat{s}$  reaches a maximum just above this threshold.

We work with four different datasets. The first three are generated at one-loop (LO) level using GOSAM-3 [72]. We use  $m_H = 125$  GeV and a fixed value of  $\alpha_s = 0.1184$ . The Higgs mass is measured to per-mille level, and its uncertainty translates into a negligible uncertainty on the di-Higgs amplitude. For the first two (LO) datasets, we use  $8 \cdot 10^4$  events for training and  $10^5$  events for testing.

1. For the first dataset, we sample the center-of-mass energy uniformly in the range

$$\hat{s} \in [(2m_H)^2, (10 \text{ TeV})^2], \quad (39)$$

and fix the top-quark mass to  $m_t = 173$  GeV.

2. A second dataset allows us to study the effects of the phase space sampling. Instead of sampling uniformly in  $\hat{s}$ , we take into account the flux factor  $1/(2\hat{s})$ , effectively sampling from a reciprocal distribution.
3. Different renormalization schemes result in a range of top-quark mass values for which one potentially wants to evaluate the amplitude. While the on-shell value is about 172.5 GeV, the  $\overline{\text{MS}}$  value depends on the renormalization scale, and the NLO amplitude shows a variation of about 20% if the envelope between predictions in different top mass renormalization schemes is taken as an uncertainty [73–76]. For our third dataset, we learn the di-Higgs amplitude without fixing the top-quark mass value, i.e. for uniformly sampled

$$m_t \in [130 \text{ GeV}, 190 \text{ GeV}], \quad (40)$$

corresponding roughly to  $\overline{\text{MS}}$  mass values  $m_t(\mu_t)$  with  $\mu_t \in [250 \text{ GeV}, 2 \text{ TeV}]$ .

4. The fourth dataset is generated at NLO in QCD including the full top-quark mass dependence with  $m_t = 173$  GeV [77, 78]. This dataset has previously been used as input for the interpolation grid in the corresponding POWHEG-Box implementation [45, 47, 48]. It contains only 6320 events.

### 3.1 One-loop with fixed top mass

We first investigate the mean relative accuracy over the phase space,  $\Delta(x)$ , after training on the first unweighted event sample in the upper left panel of Fig. 3. We observe no effect from the scattering angle  $\theta$ , but as  $\beta_H$  approaches zero close to the Higgs pair production threshold, the accuracy drops. This starts for values right above the top threshold at  $\beta_H \simeq 0.69$ . Below this threshold, the amplitude drops significantly, the uniformly sampled training dataset becomes more sparse, and the surrogate accuracy decreases. This reduced accuracy is captured correctly by the learned uncertainty.

We can improve the training with an additional small dataset with uniform sampling in

$$\hat{s} \in [(2m_H)^2, 2 \cdot (2m_t)^2]. \quad (41)$$

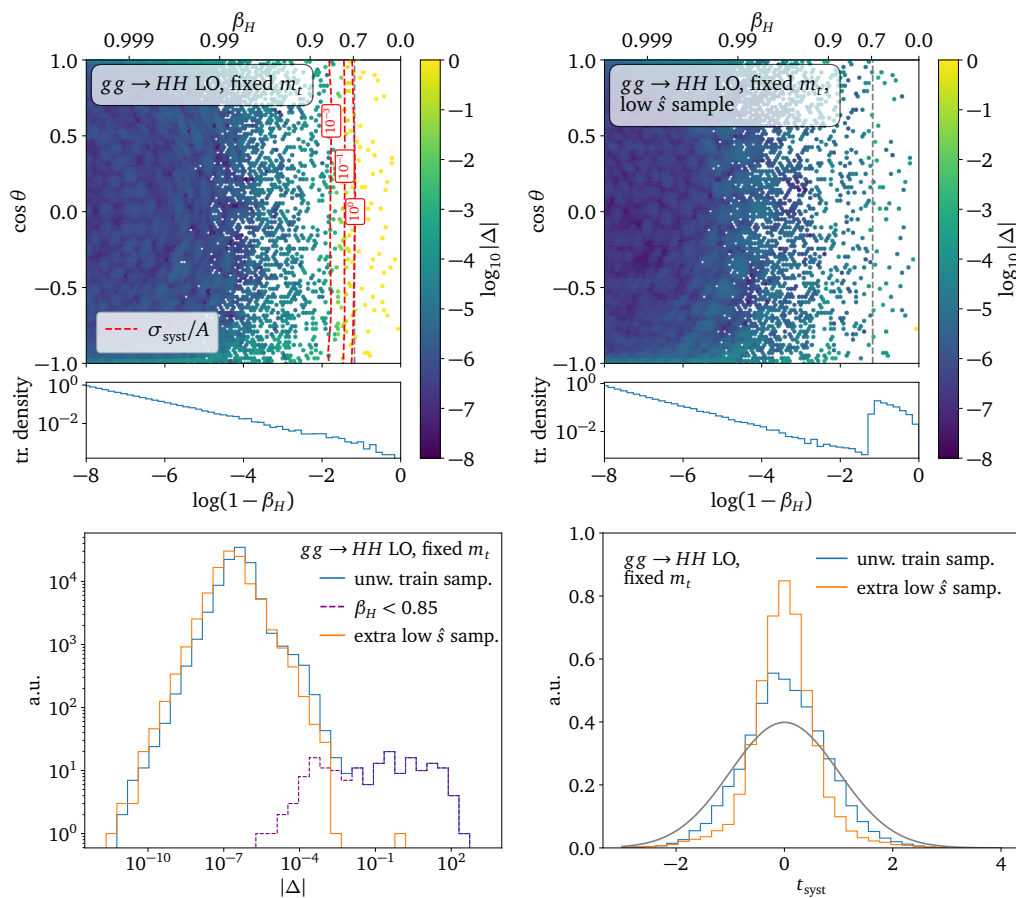


Figure 3: Upper: mean relative accuracy of the LO  $HH$  surrogate with fixed  $m_t$ , trained on an unweighted sample (left) and including an additional low- $\hat{s}$  training sample (right). The sub-panels show the density of the training dataset. Lower: relative accuracy, we also show the amplitudes below the top threshold (left), and corresponding pull distributions (right).

This low- $\hat{s}$  dataset includes  $10^4$  events. In order to allow for a fair comparison, we moreover reduce the size of the normal training dataset to  $7 \cdot 10^4$  events. The performance of the surrogate trained on the low- $\hat{s}$ -enhanced dataset is shown in the upper right panel of Fig. 3. The accuracy in the low-energy region improves by almost two orders of magnitude. Indeed, a clever choice of training data significantly boosts performance. The learned uncertainty normalized by the amplitude is essentially flat across phase space, so we do not show it.

The effect of the additional training sample can also be seen in the accuracy distributions in the lower left panel of Fig. 3, with a maximum around a relative accuracy of  $10^{-6}$ . The lower boundary around  $\Delta \sim 10^{-11}$  corresponds to the numerical precision of GOSAM. For the unweighted training dataset the accuracy stretches to  $\Delta \sim 10^2$ , where the tail includes almost exclusively phase space points below the top threshold. This maximum of the accuracy distribution is reduced by almost three orders of magnitude when we add the low- $\hat{s}$  training data.

For the calibration of the learned uncertainty, we look at the corresponding pull in the lower right panel of Fig. 3. Both distributions, without and with the additional low- $\hat{s}$  training data, are approximately Gaussian, validating our Gaussian ansatz for the likelihood. However, the width of the pull distribution is narrower than expected. This slightly underconfident uncertainty estimate is not improved when using the Student's  $t$ -likelihood.

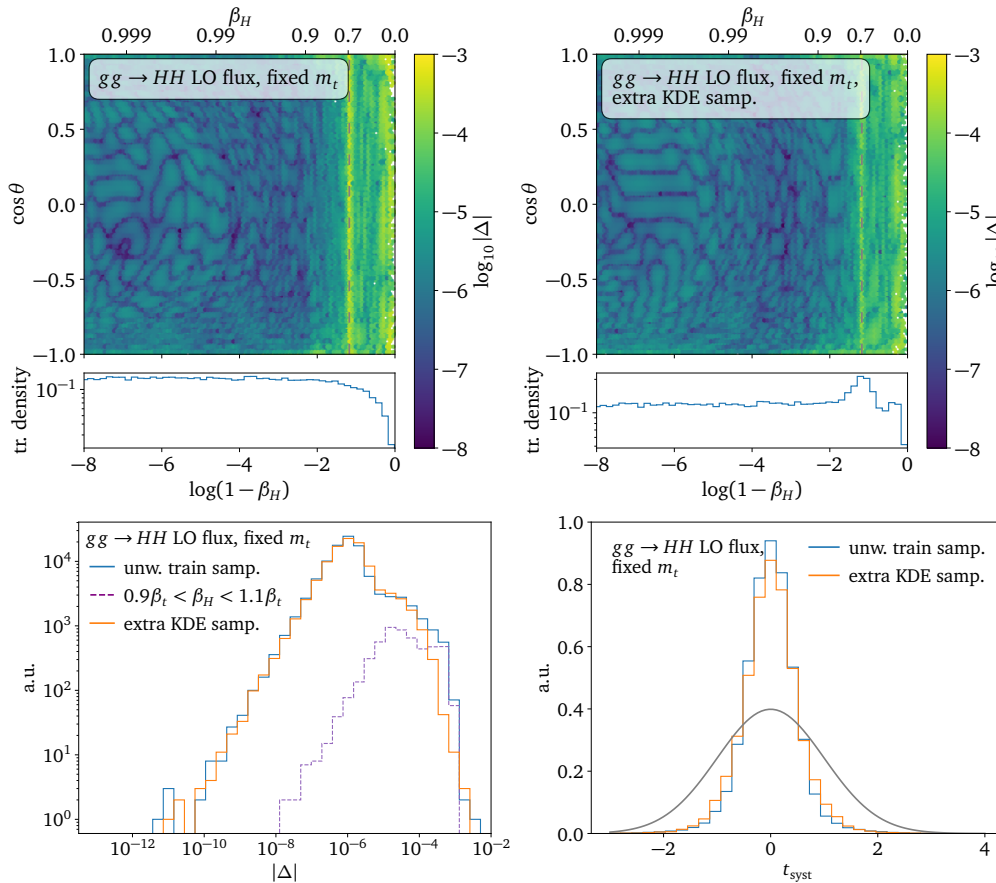


Figure 4: Upper: mean relative accuracy of the LO  $HH$  surrogate with fixed  $m_t$ , trained on a flux-reweighted sample (left), and including an additional KDE sample for training (right). The dashed gray vertical lines indicate the top threshold. The sub-panels show the density of the training dataset. Lower: relative accuracy, we also show the amplitudes around the top threshold (left), and corresponding pull distributions (right).

### Sampling with flux factor

Instead of flat sampling in  $\hat{s}$ , we can account for the flux factor. This changes the density of the training dataset, in particular, at low energies. The results for this second training dataset are shown in the upper left panel of Fig. 4. The density of the training data is much more evenly distributed, leading to an accuracy  $\Delta \sim 10^{-5}$  for most of the phase space. Only close to the di-top threshold, indicated by the vertical dashed line, the surrogate loses accuracy. This is a consequence of the behavior of the amplitude close to the threshold, which we will study in more detail below. Below the top threshold, the accuracy improves again but does not reach the level above the threshold, due to the limited amount of training data in this region.

To improve accuracy, we generate  $10^4$  additional training data points using the KDE sampling strategy outlined in Sec. 2.4. As always, we reduce the original training dataset, such that the surrogates with and without additional KDE sample always use  $8 \cdot 10^4$  training points. In the upper right panel of Fig. 4 we see that the KDE sampling improves the accuracy at and below the di-top threshold.

We confirm this pattern in the  $|\Delta|$  distributions in the lower left panel of Fig. 4. The additional KDE sampling visibly reduces the large- $\Delta$  tail, which is mainly caused by the top threshold region. The systematic pull shown in the lower right panel of Fig. 4 show again that the learned uncertainties significantly overestimate the actual deviations, a behavior that does not improve when we switch to a Student's  $t$ -likelihood.




Figure 5: From upper left to lower right: kinematic distributions of the LO  $HH$  amplitude and the learned surrogate with fixed  $m_t$ ; effect of the additional KDE sampling on the relative accuracy and the relative systematic uncertainty; the same distribution, shown only around the top threshold; same as before, but with a linear  $y$ -axis.

Finally, we investigate the behavior of the surrogate close to the di-top threshold in Fig. 5. The upper left panel shows the true and learned amplitudes as a function of  $\beta_H$  for fixed  $\cos\theta = 0$ . At this scale, we cannot see any deviation of the surrogate. For increasing  $\log(1 - \beta_H)$ , the amplitude decreases, with the steepest descent below the top threshold at  $\beta_H \simeq 0.69$ . At the top threshold, the amplitude has a turning point before it plateaus for large  $\log(1 - \beta_H)$ .

The steep decline is generally hard to learn by a network with a fixed typical resolution. The effect is shown in the upper right panel depicting the relative accuracy  $|\Delta|$  and the relative estimated uncertainty  $\sigma_{\text{syst}}/A$  as a function of  $\log(1 - \beta_H)$  comparing the surrogates trained with and without the additional KDE sample.  $|\Delta|$  fluctuates at the level of  $\lesssim 10^{-6}$  for  $\log(1 - \beta_H) \lesssim -2$ . Above this point, the steep decline lowers the accuracy, and close to the top threshold, the accuracy drops to  $\sim 3 \cdot 10^{-4}$ . This behavior is well captured by the uncertainty estimate. In the lower two panels, we show the same curves as in the upper right panel focusing on the top threshold region. The lower left panel uses a logarithmic scaling for the  $y$  axis; the right one, a linear scaling. We can see how the additional KDE sampling does not remove the accuracy drop at the top threshold, but alleviates it significantly.

### 3.2 One-loop with varying top-quark mass

Given the accuracy the surrogate reaches for the fixed top-quark mass, we can complicate the regression task by varying the top-quark mass and build on the network interpolation over phase space and over the (conditional) parameter input. The accuracy as a function of phase space is shown in the upper left panel of Fig. 6. Compared to Fig. 3, we see a significantly

lower accuracy in the low-energy region. An additional low-energy training sample enhances the accuracy in the low-energy region, as visible in the upper right panel of Fig. 6. The ultimate performance of the amplitude surrogate is comparable to the one based on the fixed- $m_t$  dataset.

This is confirmed in the lower left panel of Fig. 6. Including the low-energy sample shifts the end point of the accuracy distribution from around  $10^3$  to order-one, albeit not as low as for the fixed- $m_t$  case. The peak of the relative accuracy remains around  $10^{-6}$ , on par with the fixed- $m_t$  case. For training without an extra low-energy sample, the tail of the distribution is again almost exclusively constituted of below top-quark threshold events. We analyze the calibration of the learned uncertainties in the lower right panel of Fig. 6. The pull distributions for the more complex regression problem are clearly Gaussian and much better calibrated.

We investigate the dependence of the accuracy on the top-quark mass in Fig. 7. In the left panel, showing the mean accuracy as a function of  $m_t$ , we see how the low-energy training sample improves the accuracy. It fluctuates strongly as a function of  $m_t$ , reflecting the random sampling. A more regular sampling strategy could help flatten the distribution. However, if an additional low-energy training sample is included, the accuracy is never worse than  $\sim 5 \cdot 10^{-4}$ .

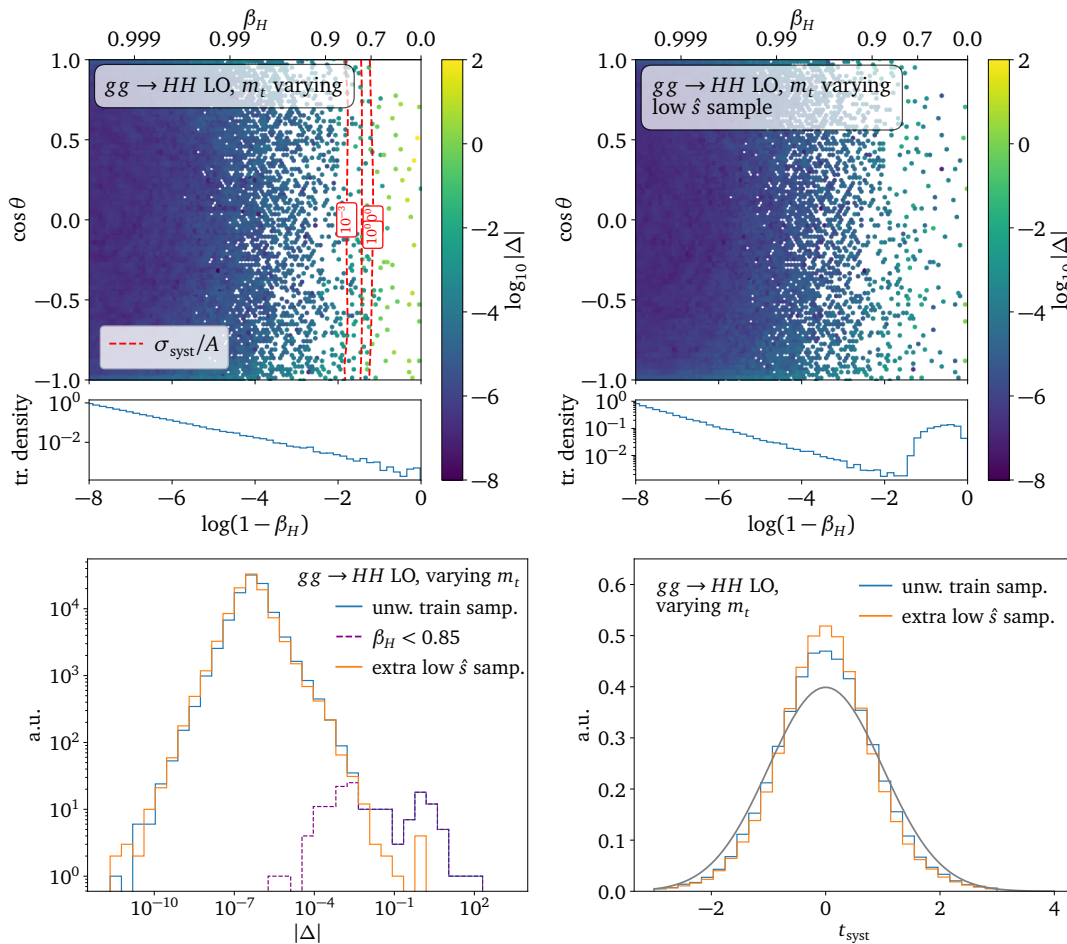


Figure 6: Upper: mean relative accuracy of the LO  $HH$  surrogate with variable  $m_t$ , trained on an unweighted sample (left) and including an additional low- $\hat{s}$  sample for training (right). The sub-panels show the density of the training dataset. Lower: relative accuracy, we also show the amplitudes below the top threshold (left), and corresponding pull distributions (right).




Figure 7: Left: mean relative accuracy of the LO  $HH$  surrogate as a function of  $m_t$ , comparing the training datasets without and with an additional low- $\hat{s}$  training sample. Right: corresponding median relative accuracy.

Table 1: Mean and median test accuracies of the LO  $HH$  surrogate for fixed and varying  $m_t$ , including training with additional low- $\hat{s}$  and KDE samples. The average and standard deviations of five independent runs are shown.

varying $m_t$	flux samp.	low- $\hat{s}$ sample	KDE sample	mean $ \Delta $	median $ \Delta $
				$(1.3 \pm 0.3) \cdot 10^{-2}$	$(3.8 \pm 0.5) \cdot 10^{-7}$
✓	✓			$(9.3 \pm 1.1) \cdot 10^{-6}$	$(1.1 \pm 0.1) \cdot 10^{-6}$
		✓		$(0.9 \pm 1.0) \cdot 10^{-2}$	$(4.7 \pm 0.3) \cdot 10^{-7}$
		✓	✓	$(1.2 \pm 0.02) \cdot 10^{-5}$	$(2.5 \pm 0.2) \cdot 10^{-7}$
✓	✓	✓	✓	$(6.4 \pm 0.4) \cdot 10^{-6}$	$(1.1 \pm 0.1) \cdot 10^{-6}$
		✓		$(4.4 \pm 0.1) \cdot 10^{-5}$	$(4.2 \pm 0.2) \cdot 10^{-7}$

Comparing the mean to the median accuracy, shown in the right panel of Fig. 7, indicates that the upwards fluctuations in the mean accuracy are mainly caused by outliers. For the median, the surrogate trained using an additional low-energy sample has an almost flat accuracy, indicating that for the bulk of the test data the amplitude is predicted at the relative  $10^{-6}$  level. Without the additional low-energy sample, the surrogate struggles to learn the  $m_t$ -dependence and the median relative accuracy decreases towards the edges of the considered  $m_t$ -range.

Finally, we compare the mean and median relative accuracies of the different LO fixed- $m_t$  and  $m_t$ -dependent surrogates based in Tab. 1. All values are based on five independent runs with the given uncertainty indicating the standard deviation across runs. All medians are somewhat comparable, reflecting the surrogate accuracy and precision for the bulk of phase space points. Additional sampling only has a minor effect on the performance. However, including an additional low-energy training sample or a KDE sample does significantly boost the mean accuracy, which is sensitive to outliers. This effect, which also shifts the upper endpoints in the  $|\Delta|$  distributions, is more pronounced for varying  $m_t$  and brings the datasets with fixed and varying  $m_t$  to roughly comparable accuracy, without increasing the size of the training dataset.

### 3.3 Two-loop virtual amplitude

Finally, we investigate a surrogate trained on the much smaller dataset for the two-loop virtual amplitude, interfered with the Born amplitude, UV-renormalized and after subtracting the IR singularities [45, 77–80]. To fully exploit the limited training and validation datasets, we use

K-fold cross-validation. The folds are obtained by holding out a fixed test set and randomly partitioning the remaining training and validation events into five equally-sized, mutually exclusive subsets ( $K = 5$ ).

The accuracy is shown in the left panel of Fig. 8, compared to a grid interpolation [45] based on convenient data pre-processing and the Clough-Tocher routine from SciPy [81]. We observe a very similar performance, with the NN surrogate performing slightly better, in particular towards large relative deviations. The similar performance of the grid interpolation and the NN surrogate reflects the low phase space dimensionality and the limited amount of training data [39]. An investigation of the surrogate accuracy across phase space did not reveal




Figure 8: Left: comparison of the accuracy of the NN surrogate and a grid interpolation. Right: pull distribution of the total uncertainty predicted by the surrogate.

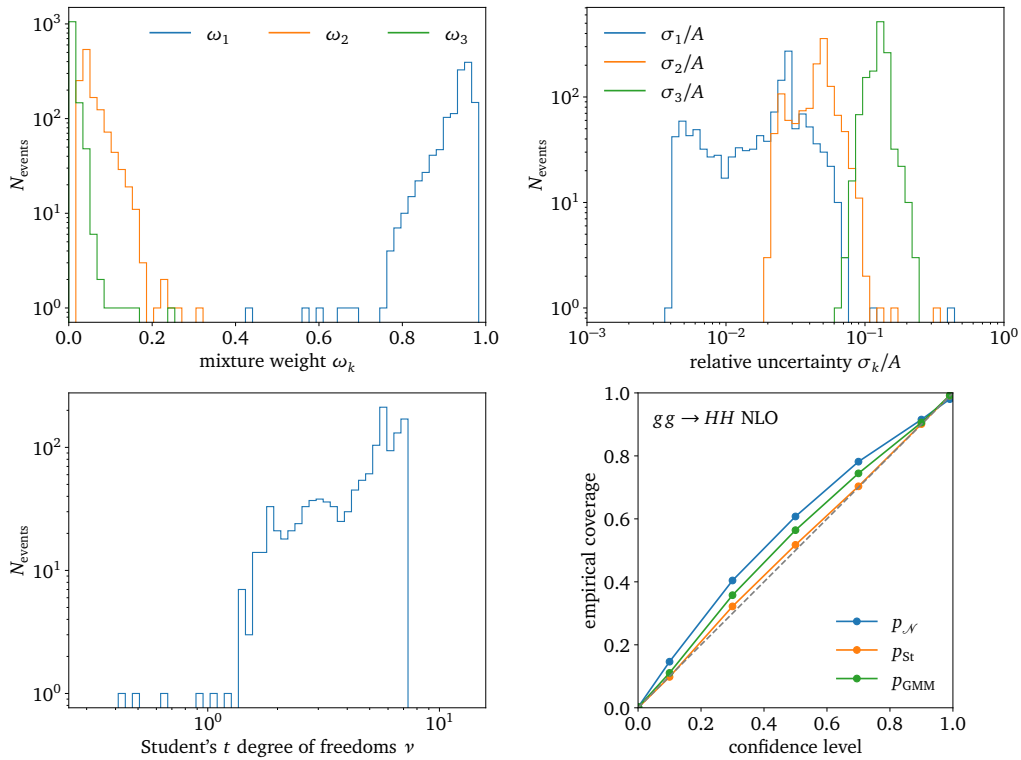


Figure 9: Non-Gaussian effect for the NLO  $HH$  surrogate. We employ a GMM (upper panels) and a Student's  $t$ -likelihood (lower left panel). The empirical coverage is shown in the lower right panel.

any significant patterns. This is due to the relatively flat distribution of the training data.

Unlike the grid interpolation, the NN surrogate also provides an uncertainty estimate, for which we show the pull distribution in the right panel of Fig. 8. The shape resembles a Gaussian, but with enhanced tails. As discussed in Sec. 2.2, the non-Gaussianity is a consequence of the low dimensionality of the phase space.

Consequently, we investigate a potential improvement from a GMM, with three modes having the same mean, and a Student's  $t$ -likelihood. The mixtures weights are shown in the upper left panel of Fig. 9. One mode dominates, but there exists a significant contribution from one additional mode. We show the distributions of the associated learned widths in the upper right panel. The second and third mode have significantly larger widths than the dominating first mode. This shows that the second and to a lesser extent the third mode are covering the low tail of the residual distribution.

We see a similar effect for the Student's  $t$ -likelihood. In the lower left panel of Fig. 9, we see the distribution of the learned degree of freedom  $\nu(x, \theta)$ . Almost all values lie between 1 and 10, for which the Student's  $t$ -likelihood has significantly enhanced tails, as compared to a Gaussian likelihood. Finally, we show all empirical coverages in the lower right panel of Fig. 9. The GMM slightly outperforms the poor Gaussian likelihood ansatz. However, the Student's  $t$ -likelihood provides almost perfect coverage, indicating an excellent calibration of the learned uncertainty all over phase space.

## 4 $t\bar{t}H$ production

As a more challenging application, we target top-pair-associated Higgs production. In the main body of the paper we show the gluon fusion results to LO and NLO, whereas the related quark-antiquark scattering results are given in App. F. For this process it has been demonstrated that NN-amplitude surrogates can outperform classic interpolation techniques [39]. Following Refs. [39, 55], we parametrize the  $2 \rightarrow 3$  phase space in terms of five dimensionless order-one variables  $x_i$ , as required in Sec. 2.3. They are

- Distance from the production threshold  $s_0 = (2m_t + m_H)^2$

$$\beta^2 = 1 - \frac{s_0}{\hat{s}} \equiv 1 - \frac{(2m_t + m_H)^2}{(p_t + p_{\bar{t}} + p_H)^2}, \quad \text{where we use } x_1 = \frac{100}{86} \left( \beta^2 - \frac{1}{10} \right), \quad (42)$$

corresponding to  $\beta^2 \in [0.1, 0.96]$ , or  $\hat{s} \in [480 \text{ GeV}, 2.4 \text{ TeV}]$ , which is similar to the range accessible by experiment, see also Fig. 2 of Ref. [55].

- Energy fraction carried by the  $t\bar{t}$  system

$$x_2 = \frac{s_{t\bar{t}} - 4m_t^2}{(\sqrt{\hat{s}} - m_H)^2 - 4m_t^2}, \quad \text{with } s_{t\bar{t}} = (p_t + p_{\bar{t}})^2. \quad (43)$$

- Polar angle of the Higgs boson relative to the beam axis

$$\cos \theta_H = \frac{p_{H,3}}{\sqrt{E_H^2 - m_H^2}} \quad \Rightarrow \quad x_3 = \frac{\theta_H}{\pi}. \quad (44)$$

- Polar angle of the top relative to the  $H$ - $t\bar{t}$  axis in the  $t\bar{t}$  system

$$\cos \theta_t = \frac{\sqrt{s_{t\bar{t}}} \left( E_H - \frac{1}{2} \sqrt{E_H^2 - m_H^2 + s_{t\bar{t}}} \right)}{2 \sqrt{s_{t\bar{t}} - 4m_t^2} \sqrt{E_H^2 - m_H^2}} \quad \Rightarrow \quad x_4 = \frac{\theta_t}{\pi}. \quad (45)$$

- Azimuthal angle of the top relative to the  $H-t\bar{t}$  axis in the  $t\bar{t}$  system

$$\tan \varphi_t = \frac{P_{H,2}}{P_{H,3} \sin \theta_H + P_{H,1} \cos \theta_H} \quad \Rightarrow \quad x_5 = \frac{\varphi_t}{2\pi}. \quad (46)$$

For the partonic sub-process

$$gg \rightarrow t\bar{t}H, \quad (47)$$

we train the amplitude surrogate on two different gluon-fusion datasets, the tree-level amplitude  $f_3 = |\mathcal{M}_0|^2$  and the finite part of the one-loop amplitude  $f_4 = 2 \text{Re}(\mathcal{M}_0^\dagger \mathcal{M}_1)$ . All amplitudes are evaluated using GOSAM. No additional preprocessing is applied to  $f_3$ , while the NLO amplitudes are normalized to the corresponding tree-level amplitude, so we actually learn

$$A_{\text{LO}} = f_3, \quad \text{and} \quad A_{\text{NLO}} = \frac{f_4}{f_3}. \quad (48)$$

As for di-Higgs production, we use  $8 \cdot 10^4$  phase space points to train each amplitude and  $10^5$  phase space points for testing. The training data corresponds to an unweighted event sample generated with densities according to pseudo parton distribution functions [39]. They mimic the behavior of proper parton distribution function sets in the suppression or enhancement of certain phase space regions. This preprocessing allows the surrogate training to focus on relevant phase space regions.

#### 4.1 Singularity structure

The physical singularities of the  $t\bar{t}H$  amplitudes provide a challenge to the network training, which can be alleviated through an appropriate preprocessing. Even though ultraviolet and infrared singularities are removed from our target amplitudes, there can still remain integrable singularities, which manifest themselves by steep localized amplitude patterns.

##### Forward and backward regions

If one of the top quarks is parallel to the beam axis, the amplitude approaches a singularity which is protected by the top quark mass. In the high-energy limit, the top quark mass becomes negligible,  $m_t^2 \ll \hat{s}$ , and the amplitude rises steeply. Using the above parametrization, the transverse momentum of the top quark is

$$p_{t,T}^2 = \left[ -|\vec{p}_t| \sin \theta_t \cos \phi_t \cos \theta_H - \left( \frac{1}{2} \beta_B \sqrt{s_{t\bar{t}}} + \gamma_B |\vec{p}_t| \cos \theta_t \right) \sin \theta_H \right]^2 + [|\vec{p}_t| \sin \theta_t \sin \phi_t]^2, \quad (49)$$

where  $\beta_B = |\vec{p}_{H,t\bar{t}}|/\sqrt{s_{t\bar{t}}}$ ,  $\gamma_B = \sqrt{1 + \beta_B^2}$ , and  $p_{H,t\bar{t}}$  is the Higgs 3-momentum in the  $t\bar{t}$  rest frame.

This forward or backward singularity is most severe in the high-energy limit,  $\hat{s} \gg m_t^2$ , and if most of the energy is carried by the  $t\bar{t}$  system,  $s_{t\bar{t}} \simeq \hat{s}$ . In this regime,  $\beta_B \simeq 0$  and  $\gamma_B \simeq 1$ , and thereby,

$$p_{t,T}^2 \simeq |\vec{p}_t|^2 \left[ (\sin \theta_t \cos \phi_t \cos \theta_H + \cos \theta_t \sin \theta_H)^2 + (\sin \theta_t \sin \phi_t)^2 \right]. \quad (50)$$

The transverse momentum is minimized either for  $\sin \theta_t \simeq 0$  or for  $\sin \phi_t \simeq 0$ . In the first case, also  $\sin \theta_H \simeq 0$  has to be fulfilled, implying that  $\theta_t, \theta_H \simeq 0, \pi$ . In the second case,  $\phi_t \simeq 0, \pi$  means that we need to solve either

$$\begin{aligned} \sin \theta_t \cos \theta_H + \cos \theta_t \sin \theta_H &= 0, & \text{for } \phi_t &\simeq 0, & \text{or} \\ \sin \theta_t \cos \theta_H - \cos \theta_t \sin \theta_H &= 0, & \text{for } \phi_t &\simeq \pi. \end{aligned} \quad (51)$$

This implies

$$\theta_t = \begin{cases} \pi - \theta_H, & \text{for } \phi_t \simeq 0, \\ \theta_H, & \text{for } \phi_t \simeq \pi. \end{cases}$$

These two solutions also incorporate the first case, for which  $\sin \theta_t \simeq 0$  to begin with. In terms of the  $x_i$ , the limits correspond to

$$x_3 = x_4, \quad \text{or} \quad x_3 = 1 - x_4. \quad (52)$$

In the first case, the top quark is moving in negative  $\hat{z}$  direction; in the second case, in positive  $\hat{z}$  direction. This derivation is valid analogously for the anti-top quark.

### Coulomb singularities

A second key feature, starting at NLO, is a Coulomb-type singularity (Sommerfeld enhancement [82]) in the limit  $x_2 \rightarrow 0$ . By default, we subtract the Coulomb singularity from the NLO amplitudes [39], based on the results of Ref. [83]. For the  $q\bar{q} \rightarrow t\bar{t}H$  NLO amplitude, this reads explicitly

$$2\text{Re} \left[ \langle \mathcal{M}_0^{q\bar{q}t\bar{t}H} | \mathcal{M}_1^{q\bar{q}t\bar{t}H} \rangle + \frac{\pi^2}{\beta_{t\bar{t}}} \langle \mathcal{M}_0^{q\bar{q}t\bar{t}H} | \mathbf{T}_{t\bar{t}} | \mathcal{M}_0^{q\bar{q}t\bar{t}H} \rangle \right], \quad (53)$$

where  $\mathbf{T}_{t\bar{t}}$  is the product of the  $t$  and  $\bar{t}$  color operators and  $\beta_{t\bar{t}} = \sqrt{1 - 4m_t^2/s_{t\bar{t}}}$  [83]. In analogy, the NLO amplitude for  $gg \rightarrow t\bar{t}H$  gives us

$$2\text{Re} \left[ \langle \mathcal{M}_0^{gg t\bar{t}H} | \mathcal{M}_1^{gg t\bar{t}H} \rangle + \frac{\pi^2}{\beta_{t\bar{t}}} \langle \mathcal{M}_0^{gg t\bar{t}H} | \mathbf{T}_{t\bar{t}} | \mathcal{M}_0^{gg t\bar{t}H} \rangle \right]. \quad (54)$$

Due to the color structure the Coulomb singularity is more severe for gluon-induced production than for quark-induced production. In particular, Ref. [83] finds the following contributions to the partonic cross-section, singular in the limit  $\beta_{t\bar{t}} \rightarrow 0$ ,

$$\sigma_{\text{Coul}} = \frac{8\alpha_s}{3} \sqrt{\frac{2m_t}{2m_t + m_H}} \frac{C_{\text{Coul}}}{\beta_{t\bar{t}}}, \quad \text{where} \quad C_{\text{Coul}}^{q\bar{q}} = -\frac{1}{6}, \quad (55)$$

$$C_{\text{Coul}}^{gg} = \frac{11}{42} \frac{[4m_t^2 - m_H^2]^2 - \frac{9}{11}m_H^4}{[4m_t^2 - m_H^2]^2 + \frac{9}{7}m_H^4} \simeq 0.25.$$

If not mentioned otherwise, the Coulomb singularity is subtracted from the NLO amplitudes, and for  $pp \rightarrow t\bar{t}H$  production, this singularity is also suppressed by the phase space. However, we will also consider learning amplitudes without Coulomb subtraction, to show that such features can be learned reliably and independently of phase space factors.

### 4.2 Leading order $gg \rightarrow t\bar{t}H$

We show the relative accuracy for the  $gg \rightarrow t\bar{t}H$  amplitude in the upper left panel of Fig. 10, where ‘‘a.u.’’ is used as an abbreviation for ‘‘arbitrary units’’. The distribution peaks around  $|\Delta| \sim 10^{-4}$  and the upper tail stretches to  $|\Delta| \sim 1$ . As described in Sec. 2.3, we search for clusters in the upper 5% percentile of the  $|\Delta|$  distribution. The clustering algorithm identifies two clusters, which are most easily understood as diagonals in the  $(x_3, x_4)$  plane, shown in the upper right panel of Fig. 10. These kinematic configurations correspond to the forward/backward scattering pseudo-singular regions, where the top-quark moves into positive

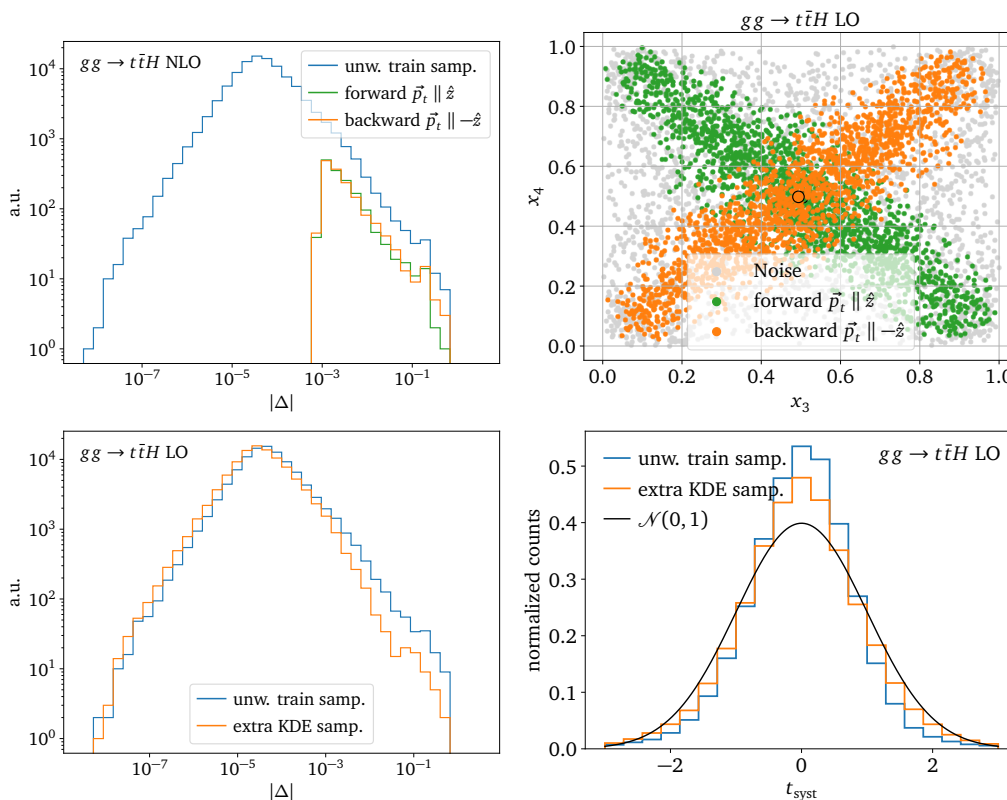


Figure 10: Upper: relative accuracy of the  $gg \rightarrow t\bar{t}H$  LO surrogate, highlighting the contribution of the identified low-accuracy clusters (left) and clustering representing the forward and backward singularities (right). Lower: relative accuracy (left) and pull distributions (right) without and with additional KDE training.

or negative beam direction. Each cluster contains approximately 30% of the 10% of phase space points with the highest  $|\Delta|$ . The remaining 40% low-accuracy points are not clustered in phase space.

Next, we cover low-accuracy regions with an additional training sample using the KDE algorithm described in Sec. 2.4. This additional dataset contains  $1.5 \cdot 10^4$  events. To provide a fair comparison, we reduce the normal unweighted training dataset to  $6.5 \cdot 10^4$ , such that the training dataset has the same size as without the KDE sample. The resulting accuracy distributions are shown in the lower left panel of Fig. 10, with a visible improvement in the large- $\Delta$  tail. We confirmed that the probability density constructed by the KDE also features the cross-like structure in the  $(x_3, x_4)$ -plane, see App. E.

Subsequently, we test the calibration of the learned uncertainties. As visible in the lower right panel of Fig. 10, the pull distributions of the surrogates with and without the additional KDE sample are reasonably well compatible, and the approximately unit Gaussian distribution pull confirms the Gaussian ansatz. We also checked that the Student's  $t$ -likelihood does not improve the empirical coverage.

### 4.3 Next-to-leading order $gg \rightarrow t\bar{t}H$

Now we move to the NLO  $gg \rightarrow t\bar{t}H$  surrogate with subtracted Coulomb singularity. In the upper left panel of Fig. 11, we show the  $|\Delta|$  distributions including the contribution from the identified low-accuracy cluster. As for the LO amplitude, two clusters correspond to two forward/backward scattering regions. As mentioned above, we do not directly learn the NLO

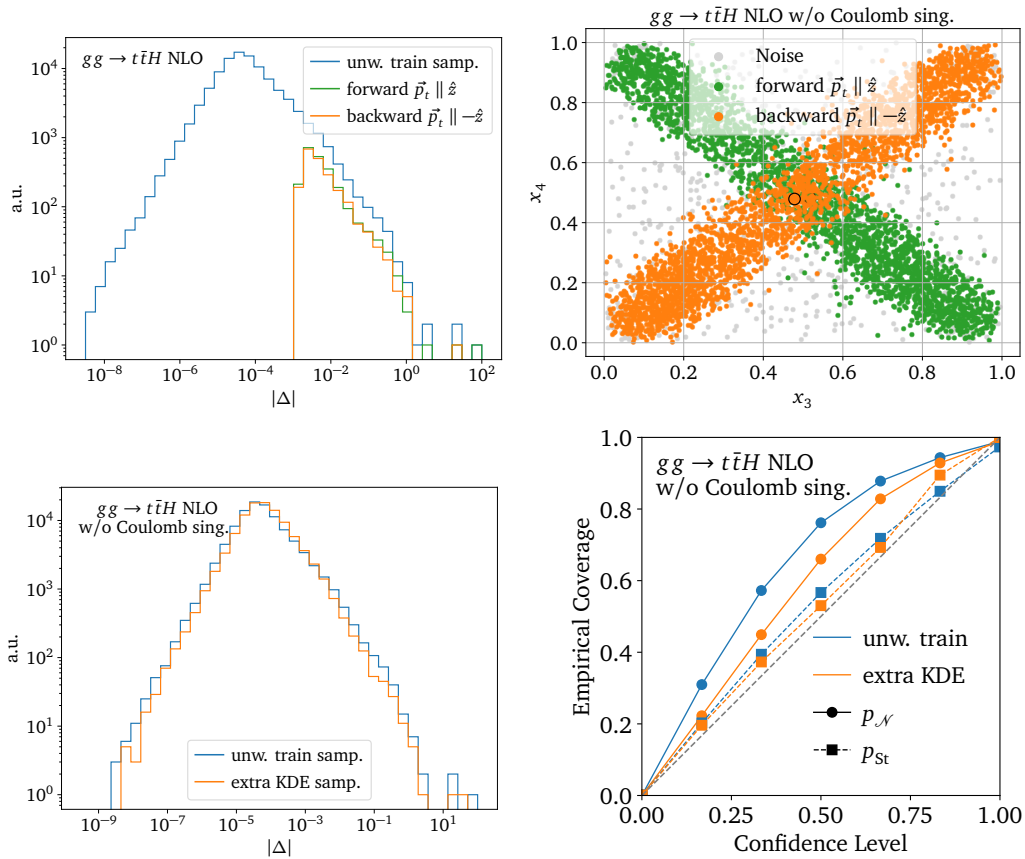


Figure 11: Upper: relative accuracy of the  $gg \rightarrow t\bar{t}H$  NLO surrogate, highlighting the contribution of the identified low-accuracy clusters (left) and clustering representing the forward and backward singularities (right). Lower: relative accuracy (left) and empirical coverage (right) without and with additional KDE training.

amplitude but the ratio between NLO and LO amplitudes. In this ratio, the forward/backward beam singularity cancels almost completely. Even though the amplitude ratio is flatter than the LO amplitude, a significant feature still remains, see App. B. The remaining singularity structure is still identified by the clustering algorithm, as shown in the upper right panel of Fig. 11. The two clusters for top quarks along the beam pipe constitute around 88% of the large- $|\Delta|$  points, with very little noise left in the  $(x_3, x_4)$ -plane.

Generating additional training data in the low-accuracy regions, using KDE sampling, slightly improves the tail of the  $|\Delta|$  distribution, as shown in the lower left panel of Fig. 8. For the NLO case, we show the calibration of the learned uncertainties without a Gaussian assumption and in terms of the empirical coverage in the lower right panel of Fig. 8. The surrogates with the additional KDE sample improves the calibration of the uncertainties, leading to a reliable calibration when used together with a Student’s  $t$ -likelihood rather than the standard Gaussian likelihood.

### NLO with Coulomb singularity

As a final test of our phase space cluster finding, we investigate the  $gg \rightarrow t\bar{t}H$  NLO amplitude without subtracting the Coulomb singularity. Indeed, the clustering algorithm now picks up three clusters, including the Coulomb singularity, as shown in the upper right panel of Fig. 12. The Coulomb cluster appears for small  $x_2$  values, corresponding to a small energy of the  $t\bar{t}$

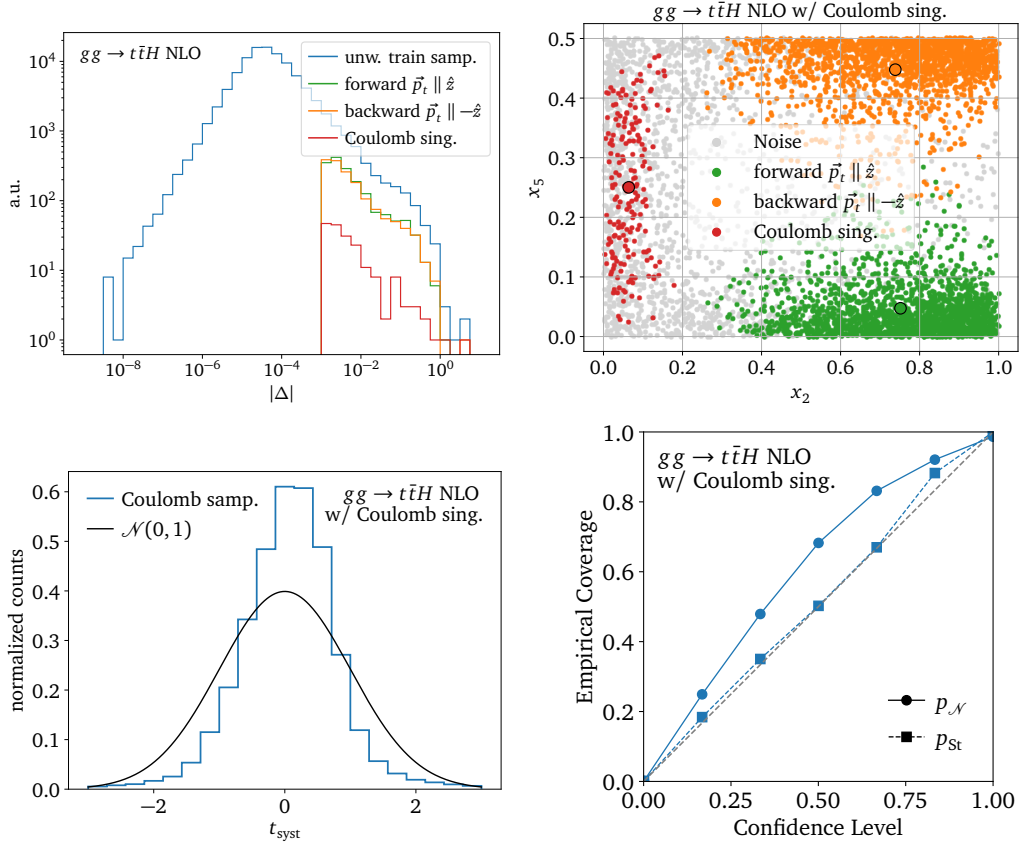


Figure 12: Upper: relative accuracy of the  $gg \rightarrow t\bar{t}H$  NLO surrogate with Coulomb singularity, highlighting the contribution of the identified low-accuracy clusters (left) and clustering representing the forward, backward, and Coulomb singularities (right). Lower: pull distribution (left) and empirical coverage (right).

system. However, in the corresponding  $|\Delta|$  distribution, we observe that the Coulomb singularity cluster only contributes around 4% to the poor-accuracy tail, whereas the two collinear clusters are responsible for 65% of this tail. This is because the radiation of the Higgs boson from a top-quark protects the Coulomb singularity, as the top propagators only can go both on-shell in the soft Higgs limit. As an additional test presented in App. D, we also performed the clustering based on the learned uncertainty finding similar results.

The pull distribution in the lower left panel of Fig. 12 shows that, as for the subtracted NLO amplitude, the learned uncertainty using a Gaussian likelihood is underconfident. Replacing the Gaussian with a Student’s  $t$ -likelihood significantly improves the empirical coverage, as shown in the lower right panel.

### Accuracy

Finally, in Tab. 2, we compare the mean and median accuracies reached for the various  $t\bar{t}H$  datasets. For comparison, we also show the  $\epsilon$  metric introduced in Ref. [39], which is given by

$$\epsilon = \frac{\sum_i |A_{\text{NN}}(x_i) - A_{\text{true}}(x_i)|}{\sum_i A_{\text{true}}(x_i)}, \quad (56)$$

for an unweighted phase space sample and independent of the way the amplitude is actually learned.

Table 2: Comparison of mean and median accuracies of the various  $t\bar{t}H$  surrogates. For comparison, also values for the  $\epsilon$  metric introduced in Ref. [39] are shown. The metrics are averaged over five independent runs with the quoted uncertainty corresponding to the standard deviation of the five runs.

	dataset	KDE	mean $ \Delta $	median $ \Delta $	$\epsilon$
$gg \rightarrow t\bar{t}H$	LO		$(5.5 \pm 0.4) \cdot 10^{-4}$	$(7.3 \pm 2.8) \cdot 10^{-5}$	$(5.5 \pm 0.4) \cdot 10^{-4}$
	LO	✓	$(2.3 \pm 0.04) \cdot 10^{-4}$	$(3.2 \pm 0.1) \cdot 10^{-5}$	$(2.34 \pm 0.04) \cdot 10^{-4}$
	NLO		$(2.1 \pm 0.8) \cdot 10^{-3}$	$(6.9 \pm 3.2) \cdot 10^{-5}$	$(1.3 \pm 0.1) \cdot 10^{-3}$
	NLO	✓	$(1.5 \pm 0.3) \cdot 10^{-3}$	$(6.4 \pm 0.8) \cdot 10^{-5}$	$(9.6 \pm 0.7) \cdot 10^{-4}$
	NLO with Coul. sing.		$(1.5 \pm 0.1) \cdot 10^{-3}$	$(5.6 \pm 1.6) \cdot 10^{-5}$	$(2.2 \pm 0.1) \cdot 10^{-3}$
	NLO with Coul. sing.	✓	$(4.8 \pm 0.1) \cdot 10^{-4}$	$(1.8 \pm 0.6) \cdot 10^{-5}$	$(1.6 \pm 0.01) \cdot 10^{-3}$
$q\bar{q} \rightarrow t\bar{t}H$	LO		$(9.6 \pm 0.7) \cdot 10^{-5}$	$(2.0 \pm 0.1) \cdot 10^{-5}$	$(9.6 \pm 0.7) \cdot 10^{-5}$
	LO	✓	$(5.5 \pm 0.1) \cdot 10^{-5}$	$(2.0 \pm 0.1) \cdot 10^{-5}$	$(5.5 \pm 0.1) \cdot 10^{-5}$
	NLO		$(1.1 \pm 0.2) \cdot 10^{-3}$	$(4.9 \pm 1.0) \cdot 10^{-5}$	$(2.2 \pm 0.1) \cdot 10^{-4}$
	NLO	✓	$(8.0 \pm 1.4) \cdot 10^{-4}$	$(4.5 \pm 0.4) \cdot 10^{-5}$	$(1.8 \pm 0.1) \cdot 10^{-4}$
	NLO with Coul. sing.		$(1.4 \pm 0.04) \cdot 10^{-3}$	$(3.3 \pm 0.5) \cdot 10^{-5}$	$(2.8 \pm 0.2) \cdot 10^{-3}$
	NLO with Coul. sing.	✓	$(3.1 \pm 0.04) \cdot 10^{-4}$	$(3.9 \pm 3.4) \cdot 10^{-6}$	$(1.8 \pm 0.2) \cdot 10^{-3}$

For the  $gg \rightarrow t\bar{t}H$  amplitudes, the mean relative accuracy of the LO amplitude is roughly an order of magnitude better than the one of the NLO amplitude. However, the median relative accuracy is roughly the same for the LO and NLO amplitudes. This indicates that the loss of mean accuracy at NLO is caused by a small number of low-accuracy phase space points. Including an additional training dataset covering low-accuracy regions described by KDE clustering improves, in particular, the mean accuracies by almost a factor of two. Training a surrogate for the NLO amplitude without subtracting the Coulomb singularity does not lead to a visible degradation and can, within uncertainties, even improve the performance.

For the NLO gluon-fusion dataset, the shown  $\epsilon$  value is slightly worse than the MLP value found in Ref. [39] for the same size of the training dataset. This is due to the used L1 loss in Ref. [39] which more directly optimizes the  $\epsilon$  metric.

Finally, we also show results for the  $q\bar{q} \rightarrow t\bar{t}H$  amplitude. The same kind of study as for the gluon-fusion process is given in App. F. As a consequence of the missing  $t$ -channel diagram, the forward/backward scattering singularities are weaker. As shown in Eq. (55) and in App. C, also the Coulomb singularity is weaker. Consequently, the low-accuracy clusters are smaller and the upper tail of the  $|\Delta|$  distribution is dominated by noise.

## 5 Outlook

Accurate and fast higher-order predictions of kinematic distributions are at the heart of the precision-LHC program. The challenges of the upcoming HL-LHC runs force us to accelerate the corresponding simulation tools and to include even higher loop orders. For both challenges, NN-surrogates of transition amplitudes provide a promising direction, especially if we can use these surrogates without further reweighting. Trustworthy amplitude surrogates have to cover two aspects: they have to accurately reproduce the truth and they have to include a reliable and calibrated uncertainty.

We have first established a way to learn uncertainties in a non-Gaussian regime, where the standard pull benchmark to assess the uncertainties has to be replaced by a coverage test. Just like for linear regression, surrogates with a larger phase space dimensionality behave more Gaussian, thanks to the central limit theorem. To identify potential failure modes, we have developed a novel method to search for localized patterns with low accuracy or large uncertainties. Finally, we have proposed a way to alleviate such localized problems through

an adaptive sampling of the training dataset.

We have established these novel techniques for the finite part of the virtual amplitudes for di-Higgs production in gluon fusion up to two loops, and for Higgs production associated with a top quark pair up to one loop. In both cases, challenging phase space regions, such as the virtual top quark pair production threshold in the di-Higgs case or a Coulomb-type singularity in the  $t\bar{t}H$  case, are reproduced adequately by the adaptively trained surrogate.

Using simple network architectures, surrogate amplitudes can be evaluated at least as fast as low-multiplicity leading-order amplitudes, i.e. faster than analytically available (multi-scale) two-loop amplitudes and much faster than numerical calculations. An efficient way to generate optimal training data is a critical step. Our new, adaptive training provides us with accurate surrogates including a reliable and conservative uncertainty estimate, even for non-Gaussian likelihoods and in the presence of challenging phase space features. Our methods will be especially important for pushing the amplitude surrogates to higher loop orders where the singularity structures are richer. Another complication at higher loop orders is given by the smaller amount of available training data. As the performance of the NN-based methods scales favorably with the number of phase space dimensions, surrogate amplitudes can pave the way to fast Monte Carlo programs for multi-scale processes at NNLO and beyond, provided that similar progress is made on the unresolved real radiation side.

## Acknowledgments

We would like to thank Victor Bresó, Vitaly Magerya and Anton Olsson for useful discussions. GH would like to thank Stephen Jones and Matthias Kerner for their work on hhgrid.

**Funding information** This work is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant 396021762 - TRR 257 Particle Physics Phenomenology after the Higgs Discovery. The work of HB, TP, and RR is supported by Deutsche Forschungsgemeinschaft (DFG) under Germany's Excellence Strategy EXC-2181/1 - 390900948 (the Heidelberg STRUCTURES Excellence Cluster). JB is supported in parts by the Federal Ministry of Technology and Space (BMFTR) under grant number 05H24VKB. The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through the grants INST 35/1597-1 FUGG and INST 39/1232-1 FUGG.

## A Hyperparameters

We provide an overview of the hyperparameter settings used for the various processes in Tab. 3. For the  $gg \rightarrow t\bar{t}H$  and  $q\bar{q} \rightarrow t\bar{t}H$  surrogates, we used the same hyperparameters.

Table 3: Network and training parameters.

Parameter	toy	$Z + \text{gluons}$	$HH$ LO	$HH$ NLO	$t\bar{t}H$
Activation function	GELU	GELU	GELU	GELU	GELU
Number of hidden layers	3	5	5	6	5
Hidden nodes	64	128	512	256	512
Batch size	256	256	256	1024	256
Scheduler	Cosine	Cosine	Cosine	Cosine	Cosine
Max learning rate	$2 \cdot 10^{-4}$	$2 \cdot 10^{-4}$	$2 \cdot 10^{-4}$	$10^{-4}$	$2 \cdot 10^{-4}$
Number of epochs	1000	4000	4000	4000	2000




Figure 13: Left: One-dimensional slice of the  $gg \rightarrow ttH$  amplitude for  $x_1 = 0.99$ ,  $x_2 = 0.95$ ,  $x_4 = 0.40$ ,  $x_5 = 0.01$ . Right: same as left, but each function is individually normalized to  $[0, 1]$ .

### B Forward/backward regions: NLO to LO ratio

Fig. 13 shows a one-dimensional slice in  $x_3$  of the  $gg \rightarrow ttH$  amplitude for fixed values of  $x_1 = 0.99$ ,  $x_2 = 0.95$ ,  $x_4 = 0.40$  and  $x_5 = 0.01$ . The forward/backward region described in section 4.1 is clearly visible at  $x_3 \sim 0.87$  in the LO and NLO amplitudes, whereas the absolute value of the ratio is nearly flat compared to the amplitudes themselves. When each function is individually normalized to  $[0, 1]$  however, the ratio also shows steep changes in the forward/backward region. This results in a challenging region for the surrogate persisting even when learning only the ratio.

### C Accuracy close to Coulomb singularity

In Fig. 14, we show the median relative accuracy of the  $t\bar{t}H$  NLO surrogates as a function of  $x_2$  with and without subtraction of the Coulomb singularity. For the  $gg \rightarrow t\bar{t}H$  surrogate shown in the left panel, the loss in relative accuracy for  $x_2 \rightarrow 0$  without subtraction of the Coulomb singularity is significantly larger than for the  $q\bar{q} \rightarrow t\bar{t}H$  amplitude shown in the right panel. This is in agreement with the theoretical expectation outlined in Sec. 4.1.

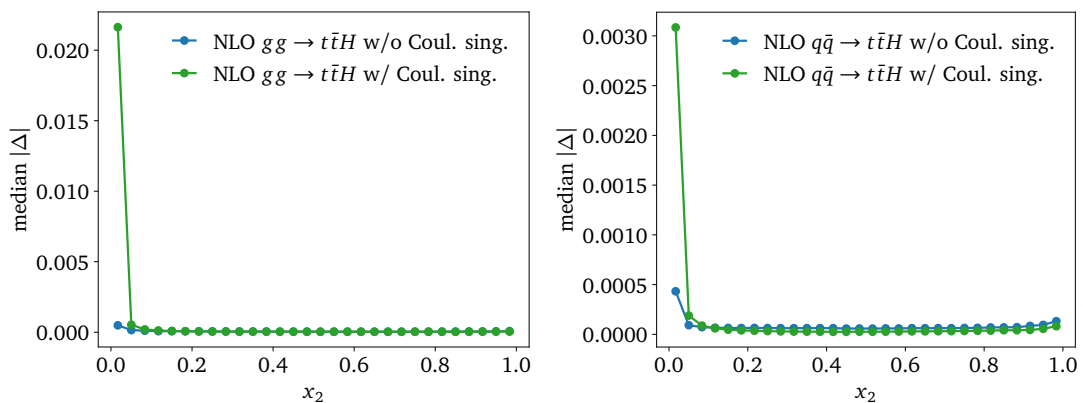


Figure 14: Left: median accuracy of the  $gg \rightarrow t\bar{t}H$  NLO surrogate as a function of  $x_2$  with and without subtraction of the Coulomb singularity. Right: Same as left, but the median accuracy of the  $q\bar{q} \rightarrow t\bar{t}H$  NLO surrogate is shown.




Figure 15: Identified clusters based on the estimated uncertainty for the NLO  $gg \rightarrow t\bar{t}H$  surrogate without subtraction of the Coulomb singularity in the  $x_2$ - $x_5$  plane

### D Clustering based on estimated uncertainties

In Fig. 15, we show the result of the clustering algorithm for the NLO  $gg \rightarrow t\bar{t}H$  surrogate without subtraction of the Coulomb singularity based on the estimated uncertainties. In particular, we use  $\sigma/A$  instead of  $|\Delta|$  as measure to selecting the least accurate phase-space points. As before, the forward/backward regions and the Coulomb singularity are identified. In addition, the algorithm identifies an additional cluster, which we could not identify with any specific features of the amplitude. We attribute this to the fact that the clustering is performed based on the 5% points with the largest  $\sigma/A$  and not the largest  $|\Delta|$ . Since the uncertainty is not perfectly calibrated, a different set of points is selected.




Figure 16: Density plot of the marginalized KDE-PDF for  $gg \rightarrow ttH$  in the  $x_3$ - $x_4$  plane at LO (left) and at NLO (right).

## E Extracted KDE estimates

Fig. 16 shows the marginalized KDE-PDF for the  $gg \rightarrow t\bar{t}H$  channel in the  $x_3$ - $x_4$  plane at LO and NLO. The KDE-PDF can clearly be observed to reproduce the diagonal structure formed by the two clusters identified in the top right panel of Figs. 10 and 11, leading to a KDE sample focused in the forward/backwards scattering regions.

## F $q\bar{q} \rightarrow t\bar{t}H$ results

In this Appendix, we collect the results for the  $q\bar{q} \rightarrow t\bar{t}H$  amplitude, which we investigated in the same way as the  $gg \rightarrow t\bar{t}H$  amplitude.

### Leading order

First, we discuss the results for the LO amplitude shown in Fig. 17. As for the  $gg \rightarrow t\bar{t}H$  amplitude, the clustering algorithm identifies the forward/backward regions. But due to the lack of a  $t$ -channel diagram, the forward/backward scattering singularities are significantly weaker than for the  $gg \rightarrow t\bar{t}H$  amplitude and the tail of the  $|\Delta|$  distribution is dominated by noise. Including an additional KDE sampling significantly improves the performance. We find the uncertainty estimate to be slightly underconfident.

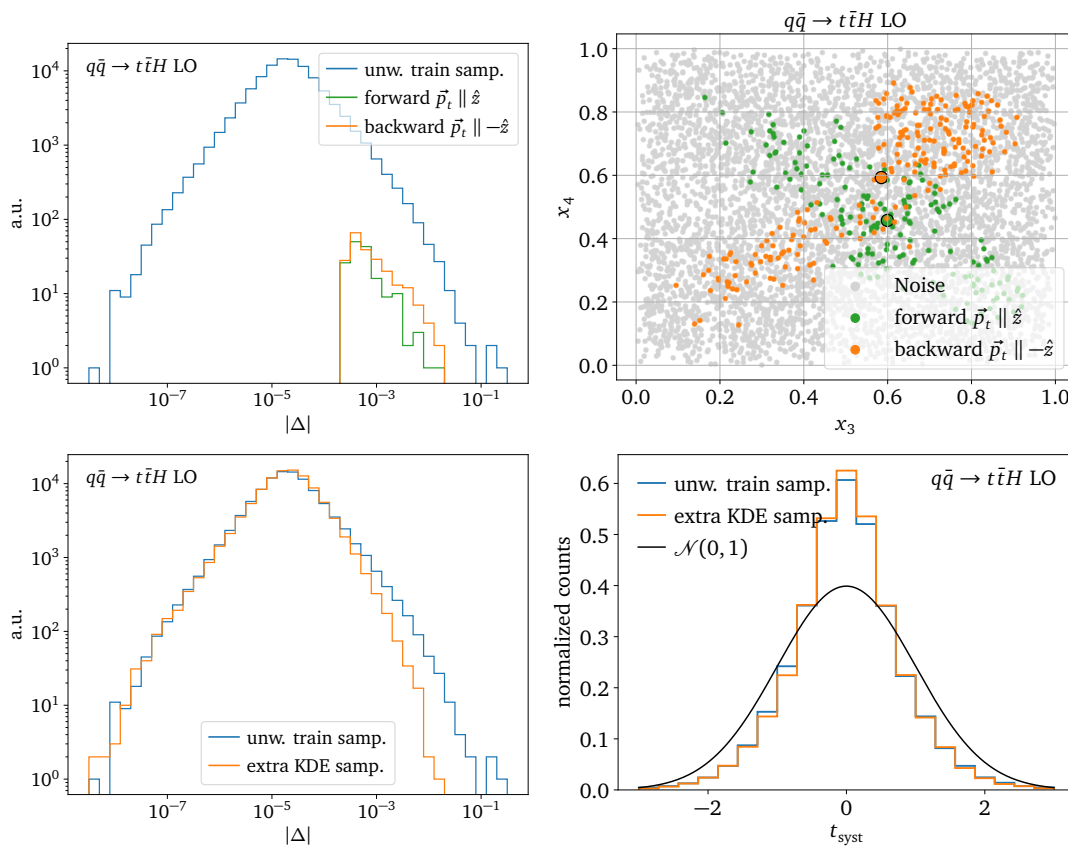


Figure 17: Upper left: Relative accuracy of  $q\bar{q} \rightarrow t\bar{t}H$  LO surrogates highlighting the identified low-accuracy clusters. Upper right: identified clusters in the  $(x_3, x_4)$  plane. Lower left: Relative accuracy comparing the surrogates trained with and without additional KDE training samples. Lower right: Systematic pull distributions.

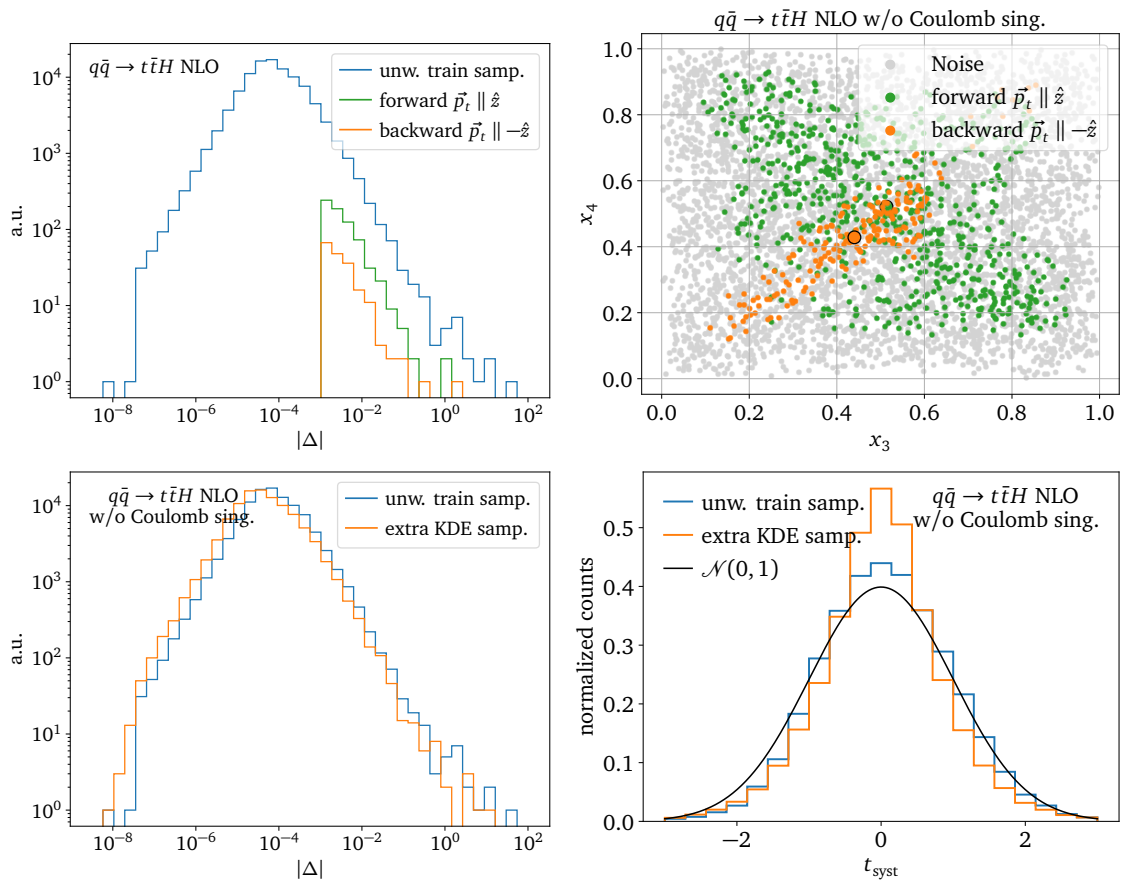


Figure 18: Upper left: Relative accuracy of  $q\bar{q} \rightarrow t\bar{t}H$  NLO surrogates highlighting the identified low-accuracy clusters. Upper right: identified clusters in the  $(x_3, x_4)$  plane. Lower left: Relative accuracy comparing the surrogates trained with and without additional KDE training samples. Lower right: Systematic pull distributions.

### Next-to-leading order

The results for the NLO amplitude are shown in Fig. 18. The results are very similar to the LO amplitude. The forward/backward regions again only slightly contribute to the upper  $|\Delta|$  tail, the KDE sampling significantly improves performance, and the uncertainty estimates are slightly underconfident.

### NLO $q\bar{q} \rightarrow t\bar{t}H$ with Coulomb singularity

Finally, we show in 19 the results for the NLO  $q\bar{q} \rightarrow t\bar{t}H$  amplitude without subtraction of the Coulomb singularity. The accuracy is comparable to the case with subtraction. The clustering algorithm picks up the Coulomb singularity but fails to identify the backward scattering region due to the large amount of noise. We find the uncertainty estimate to be again slightly underconfident.

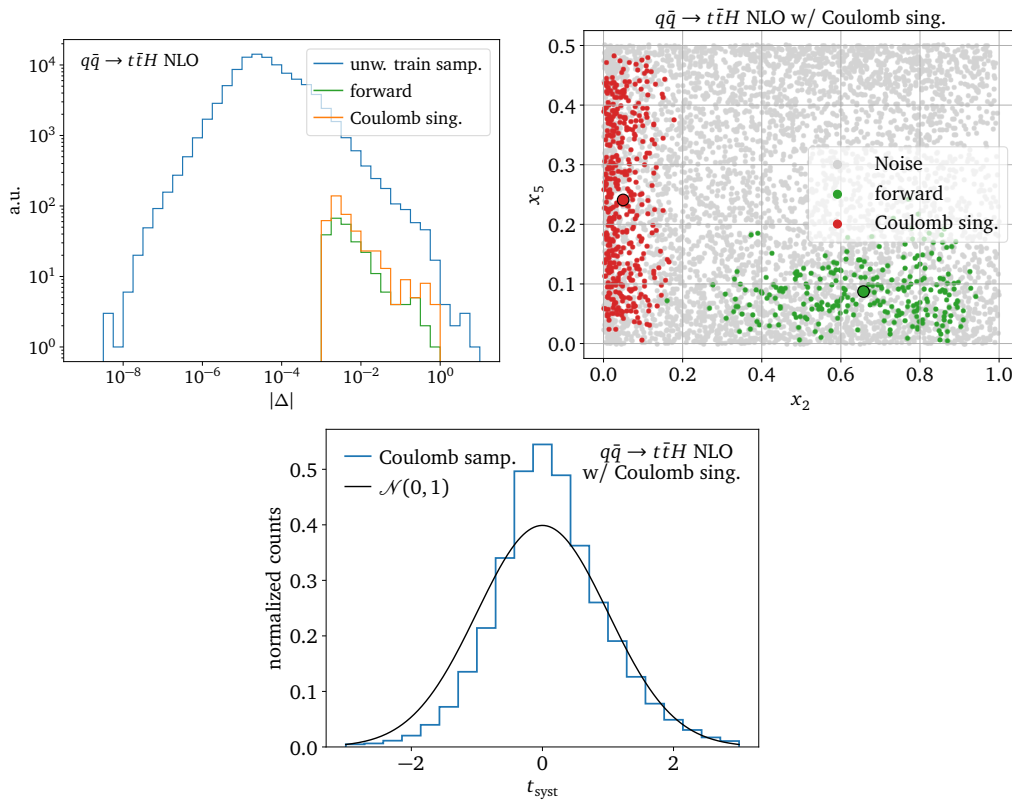


Figure 19: Upper left: Relative accuracy of  $q\bar{q} \rightarrow t\bar{t}H$  NLO surrogates without subtraction of Coulomb singularity highlighting the identified low-accuracy clusters. Upper right: identified clusters in the  $(x_2, x_5)$  plane. Lower left: Relative accuracy comparing the surrogates trained with and without additional KDE training samples. Bottom: Systematic pull distributions.

## References

- [1] S. Dawson, C. Englert and T. Plehn, *Higgs physics: It ain't over till it is over*, Phys. Rep. **816**, 1 (2019), doi:[10.1016/j.physrep.2019.05.001](https://doi.org/10.1016/j.physrep.2019.05.001) [preprint doi:[10.48550/arXiv.1808.01324](https://doi.org/10.48550/arXiv.1808.01324)].
- [2] X. Chen, X. Guan and B. Mistlberger, *Three-Loop QCD corrections to the production of a Higgs boson and a Jet*, (arXiv preprint) doi:[10.48550/arXiv.2504.06490](https://doi.org/10.48550/arXiv.2504.06490).
- [3] J. Davies, D. Grau, K. Schönwald, M. Steinhauser and D. Stremmer, *Three-loop corrections to  $gg \rightarrow ZH$  in the large top quark mass limit*, J. High Energy Phys. **02**, 191 (2026), doi:[10.1007/JHEP02\(2026\)191](https://doi.org/10.1007/JHEP02(2026)191) [preprint doi:[10.48550/arXiv.2512.00156](https://doi.org/10.48550/arXiv.2512.00156)].
- [4] M. Becchetti, D. Canko, V. Chestnov, T. Peraro, M. Pozzoli and S. Zoia, *Two-loop Feynman integrals for leading colour  $t\bar{t}W$  production at hadron colliders*, J. High Energy Phys. **07**, 001 (2025), doi:[10.1007/JHEP07\(2025\)001](https://doi.org/10.1007/JHEP07(2025)001) [preprint doi:[10.48550/arXiv.2504.13011](https://doi.org/10.48550/arXiv.2504.13011)].
- [5] S. Badger, M. Becchetti, C. Brancaccio, M. Czakon, H. B. Hartanto, R. Poncelet and S. Zoia, *Higher-order QCD corrections to top-quark pair production in association with a jet*, (arXiv preprint) doi:[10.48550/arXiv.2511.11431](https://doi.org/10.48550/arXiv.2511.11431).

- [6] D. Canko and M. Pozzoli, *All planar three-loop Feynman integrals for the production of two vector bosons at hadron colliders*, J. High Energy Phys. **05**, 226 (2026), doi:[10.1007/JHEP05\(2026\)226](https://doi.org/10.1007/JHEP05(2026)226) [preprint doi:[10.48550/arXiv.2512.02999](https://doi.org/10.48550/arXiv.2512.02999)].
- [7] J. Brehmer, F. Kling, I. Espejo and K. Cranmer, *MadMiner: Machine learning-based inference for particle physics*, Comput. Softw. Big Sci. **4**, 3 (2020), doi:[10.1007/s41781-020-0035-2](https://doi.org/10.1007/s41781-020-0035-2) [preprint doi:[10.48550/arXiv.1907.10621](https://doi.org/10.48550/arXiv.1907.10621)].
- [8] S. Chatterjee, N. Frohner, L. Lechner, R. Schöfbeck and D. Schwarz, *Tree boosting for learning EFT parameters*, Comput. Phys. Commun. **277**, 108385 (2022), doi:[10.1016/j.cpc.2022.108385](https://doi.org/10.1016/j.cpc.2022.108385) [preprint doi:[10.48550/arXiv.2107.10859](https://doi.org/10.48550/arXiv.2107.10859)].
- [9] S. Chatterjee, S. Rohshap, R. Schöfbeck and D. Schwarz, *Learning the EFT likelihood with tree boosting*, (arXiv preprint) doi:[10.48550/arXiv.2205.12976](https://doi.org/10.48550/arXiv.2205.12976).
- [10] ATLAS collaboration: G. Aad et al., *Measurement of off-shell Higgs boson production in the  $H^* \rightarrow ZZ \rightarrow 4\ell$  decay channel using a neural simulation-based inference technique in 13 TeV pp collisions with the ATLAS detector*, Rep. Prog. Phys. **88**, 057803 (2025), doi:[10.1088/1361-6633/adcd9a](https://doi.org/10.1088/1361-6633/adcd9a) [preprint doi:[10.48550/arXiv.2412.01548](https://doi.org/10.48550/arXiv.2412.01548)].
- [11] R. Schöfbeck, *Refinable modeling for unbinned SMEFT analyses*, Mach. Learn. Sci. Technol. **6**, 015007 (2025), doi:[10.1088/2632-2153/ad9fd1](https://doi.org/10.1088/2632-2153/ad9fd1) [preprint doi:[10.48550/arXiv.2406.19076](https://doi.org/10.48550/arXiv.2406.19076)].
- [12] H. Bahl, V. Bresó-Pla, G. De Crescenzo and T. Plehn, *Advancing tools for simulation-based inference*, SciPost Phys. Core **8**, 060 (2025), doi:[10.21468/SciPostPhysCore.8.3.060](https://doi.org/10.21468/SciPostPhysCore.8.3.060) [preprint doi:[10.48550/arXiv.2410.07315](https://doi.org/10.48550/arXiv.2410.07315)].
- [13] L. Benato, C. Giordano, C. Krause, A. Li, R. Schöfbeck, D. Schwarz, M. Shooshtari and D. Wang, *Unbinned inclusive cross-section measurements with machine-learned systematic uncertainties*, Phys. Rev. D **112**, 052006 (2025), doi:[10.1103/zwzt-1rrw](https://doi.org/10.1103/zwzt-1rrw) [preprint doi:[10.48550/arXiv.2505.05544](https://doi.org/10.48550/arXiv.2505.05544)].
- [14] H. Bahl, T. Plehn and N. Schmal, *Unbinning global LHC analyses*, (arXiv preprint) doi:[10.48550/arXiv.2509.05409](https://doi.org/10.48550/arXiv.2509.05409).
- [15] A. Butter et al., *Machine learning and LHC event generation*, SciPost Phys. **14**, 079 (2023), doi:[10.21468/SciPostPhys.14.4.079](https://doi.org/10.21468/SciPostPhys.14.4.079) [preprint doi:[10.48550/arXiv.2203.07460](https://doi.org/10.48550/arXiv.2203.07460)].
- [16] T. Plehn, A. Butter, B. Dillon, T. Heimel, C. Krause and R. Winterhalder, *Modern machine learning for LHC physicists*, (arXiv preprint) doi:[10.48550/arXiv.2211.01421](https://doi.org/10.48550/arXiv.2211.01421).
- [17] J. Bendavid, *Efficient Monte Carlo integration using boosted decision trees and generative deep neural networks*, (arXiv preprint) doi:[10.48550/arXiv.1707.00028](https://doi.org/10.48550/arXiv.1707.00028).
- [18] M. Klimek and M. Perelstein, *Neural network-based approach to phase space integration*, SciPost Phys. **9**, 053 (2020), doi:[10.21468/SciPostPhys.9.4.053](https://doi.org/10.21468/SciPostPhys.9.4.053) [preprint doi:[10.48550/arXiv.1810.11509](https://doi.org/10.48550/arXiv.1810.11509)].
- [19] C. Gao, J. Isaacson and C. Krause, *i-flow: High-dimensional integration and sampling with normalizing flows*, Mach. Learn.: Sci. Technol. **1**, 045023 (2020), doi:[10.1088/2632-2153/abab62](https://doi.org/10.1088/2632-2153/abab62) [preprint doi:[10.48550/arXiv.2001.05486](https://doi.org/10.48550/arXiv.2001.05486)].
- [20] C. Gao, S. Höche, J. Isaacson, C. Krause and H. Schulz, *Event generation with normalizing flows*, Phys. Rev. D **101**, 076002 (2020), doi:[10.1103/PhysRevD.101.076002](https://doi.org/10.1103/PhysRevD.101.076002) [preprint doi:[10.48550/arXiv.2001.10028](https://doi.org/10.48550/arXiv.2001.10028)].

- [21] E. Bothmann, T. Janßen, M. Knobbe, T. Schmale and S. Schumann, *Exploring phase space with neural importance sampling*, SciPost Phys. **8**, 069 (2020), doi:[10.21468/SciPostPhys.8.4.069](https://doi.org/10.21468/SciPostPhys.8.4.069) [preprint doi:[10.48550/arXiv.2001.05478](https://doi.org/10.48550/arXiv.2001.05478)].
- [22] T. Heimel, R. Winterhalder, A. Butter, J. Isaacson, C. Krause, F. Maltoni, O. Mattelaer and T. Plehn, *MadNIS - Neural multi-channel importance sampling*, SciPost Phys. **15**, 141 (2023), doi:[10.21468/SciPostPhys.15.4.141](https://doi.org/10.21468/SciPostPhys.15.4.141) [preprint doi:[10.48550/arXiv.2212.06172](https://doi.org/10.48550/arXiv.2212.06172)].
- [23] T. Heimel, N. Huetsch, F. Maltoni, O. Mattelaer, T. Plehn and R. Winterhalder, *The MadNIS reloaded*, SciPost Phys. **17**, 023 (2024), doi:[10.21468/SciPostPhys.17.1.023](https://doi.org/10.21468/SciPostPhys.17.1.023) [preprint doi:[10.48550/arXiv.2311.01548](https://doi.org/10.48550/arXiv.2311.01548)].
- [24] N. Deutschmann and N. Götz, *Accelerating HEP simulations with neural importance sampling*, J. High Energy Phys. **03**, 083 (2024), doi:[10.1007/JHEP03\(2024\)083](https://doi.org/10.1007/JHEP03(2024)083) [preprint doi:[10.48550/arXiv.2401.09069](https://doi.org/10.48550/arXiv.2401.09069)].
- [25] T. Heimel, O. Mattelaer, T. Plehn and R. Winterhalder, *Differentiable MadNIS-lite*, SciPost Phys. **18**, 017 (2025), doi:[10.21468/SciPostPhys.18.1.017](https://doi.org/10.21468/SciPostPhys.18.1.017) [preprint doi:[10.48550/arXiv.2408.01486](https://doi.org/10.48550/arXiv.2408.01486)].
- [26] E. Bothmann, T. Janßen, M. Knobbe, B. Schmitzer and F. Sinz, *Efficient many-jet event generation with Flow Matching*, (arXiv preprint) doi:[10.48550/arXiv.2506.18987](https://doi.org/10.48550/arXiv.2506.18987).
- [27] T. Janßen, R. Poncelet and S. Schumann, *Sampling NNLO QCD phase space with normalizing flows*, J. High Energy Phys. **09**, 194 (2025), doi:[10.1007/JHEP09\(2025\)194](https://doi.org/10.1007/JHEP09(2025)194) [preprint doi:[10.48550/arXiv.2505.13608](https://doi.org/10.48550/arXiv.2505.13608)].
- [28] F. Bishara and M. Montull, *(Machine) Learning amplitudes for faster event generation*, Phys. Rev. D **107**, L07190 (2023), doi:[10.1103/physrevd.107.l07190](https://doi.org/10.1103/physrevd.107.l07190) [preprint doi:[10.48550/arXiv.1912.11055](https://doi.org/10.48550/arXiv.1912.11055)].
- [29] S. Badger and J. Bullock, *Using neural networks for efficient evaluation of high multiplicity scattering amplitudes*, J. High Energy Phys. **06**, 114 (2020), doi:[10.1007/JHEP06\(2020\)114](https://doi.org/10.1007/JHEP06(2020)114) [preprint doi:[10.48550/arXiv.2002.07516](https://doi.org/10.48550/arXiv.2002.07516)].
- [30] J. Aylett-Bullock, S. Badger and R. Moodie, *Optimising simulations for diphoton production at hadron colliders using amplitude neural networks*, J. High Energy Phys. **08**, 066 (2021), doi:[10.1007/JHEP08\(2021\)066](https://doi.org/10.1007/JHEP08(2021)066) [preprint doi:[10.48550/arXiv.2106.09474](https://doi.org/10.48550/arXiv.2106.09474)].
- [31] D. Maître and H. Truong, *A factorisation-aware Matrix element emulator*, J. High Energy Phys. **11**, 066 (2021), doi:[10.1007/JHEP11\(2021\)066](https://doi.org/10.1007/JHEP11(2021)066) [preprint doi:[10.48550/arXiv.2107.06625](https://doi.org/10.48550/arXiv.2107.06625)].
- [32] K. Danziger, T. Janßen, S. Schumann and F. 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](https://doi.org/10.21468/SciPostPhys.12.5.164) [preprint doi:[10.48550/arXiv.2109.11964](https://doi.org/10.48550/arXiv.2109.11964)].
- [33] R. Winterhalder, V. Magerya, E. Villa, S. Jones, M. Kerner, A. Butter, G. Heinrich and T. Plehn, *Targeting multi-loop integrals with neural networks*, SciPost Phys. **12**, 129 (2022), doi:[10.21468/SciPostPhys.12.4.129](https://doi.org/10.21468/SciPostPhys.12.4.129) [preprint doi:[10.48550/arXiv.2112.09145](https://doi.org/10.48550/arXiv.2112.09145)].
- [34] S. Badger, A. Butter, M. Luchmann, S. Pitz and T. Plehn, *Loop amplitudes from precision networks*, SciPost Phys. Core **6**, 034 (2023), doi:[10.21468/SciPostPhysCore.6.2.034](https://doi.org/10.21468/SciPostPhysCore.6.2.034) [preprint doi:[10.48550/arXiv.2206.14831](https://doi.org/10.48550/arXiv.2206.14831)].

- [35] T. Janßen, D. Maître, S. Schumann, F. Siegert and H. Truong, *Unweighting multijet event generation using factorisation-aware neural networks*, SciPost Phys. **15**, 107 (2023), doi:[10.21468/SciPostPhys.15.3.107](https://doi.org/10.21468/SciPostPhys.15.3.107) [preprint doi:[10.48550/arXiv.2301.13562](https://doi.org/10.48550/arXiv.2301.13562)].
- [36] D. Maître and H. Truong, *One-loop matrix element emulation with factorisation awareness*, J. High Energy Phys. **05**, 159 (2023), doi:[10.1007/JHEP05\(2023\)159](https://doi.org/10.1007/JHEP05(2023)159) [preprint doi:[10.48550/arXiv.2302.04005](https://doi.org/10.48550/arXiv.2302.04005)].
- [37] J. Brehmer, V. Bresó, P. de Haan, T. Plehn, H. Qu, J. Spinner and J. Thaler, *A Lorentz-equivariant transformer for all of the LHC*, SciPost Phys. **19**, 108 (2025), doi:[10.21468/SciPostPhys.19.4.108](https://doi.org/10.21468/SciPostPhys.19.4.108) [preprint doi:[10.48550/arXiv.2411.00446](https://doi.org/10.48550/arXiv.2411.00446)].
- [38] H. Bahl, N. Elmer, L. Favaro, M. Haussmann, T. Plehn and R. Winterhalder, *Accurate surrogate amplitudes with calibrated uncertainties*, SciPost Phys. Core **8**, 073 (2025), doi:[10.21468/SciPostPhysCore.8.4.073](https://doi.org/10.21468/SciPostPhysCore.8.4.073) [preprint doi:[10.48550/arXiv.2412.12069](https://doi.org/10.48550/arXiv.2412.12069)].
- [39] V. Bresó-Pla, G. Heinrich, V. Magerya and A. Olsson, *Interpolating amplitudes*, SciPost Phys. **19**, 123 (2025), doi:[10.21468/SciPostPhys.19.5.123](https://doi.org/10.21468/SciPostPhys.19.5.123) [preprint doi:[10.48550/arXiv.2412.09534](https://doi.org/10.48550/arXiv.2412.09534)].
- [40] T. Herrmann, T. Janßen, M. Schenker, S. Schumann and F. Siegert, *Accelerating multijet-merged event generation with neural network matrix element surrogates*, SciPost Phys. **20**, 071 (2026), doi:[10.21468/SciPostPhys.20.3.071](https://doi.org/10.21468/SciPostPhys.20.3.071) [preprint doi:[10.48550/arXiv.2506.06203](https://doi.org/10.48550/arXiv.2506.06203)].
- [41] L. Favaro, G. Gerhartz, F. A. Hamprecht, P. Lippmann, S. Pitz, T. Plehn, H. Qu and J. Spinner, *Lorentz-Equivariance without limitations*, (arXiv preprint) doi:[10.48550/arXiv.2508.14898](https://doi.org/10.48550/arXiv.2508.14898).
- [42] J. M. Villadamigo, R. Frederix, T. Plehn, T. Vitos and R. Winterhalder, *FASTColor - Full-color amplitude surrogate toolkit for QCD*, (arXiv preprint) doi:[10.48550/arXiv.2509.07068](https://doi.org/10.48550/arXiv.2509.07068).
- [43] H. Bahl, N. Elmer, T. Plehn and R. Winterhalder, *Amplitude uncertainties everywhere all at once*, SciPost Phys. **20**, 083 (2026), doi:[10.21468/SciPostPhys.20.3.083](https://doi.org/10.21468/SciPostPhys.20.3.083) [preprint doi:[10.48550/arXiv.2509.00155](https://doi.org/10.48550/arXiv.2509.00155)].
- [44] L. Beccatini, F. Maltoni, O. Mattelaer and R. Winterhalder, *Amplitude surrogates for multi-jet processes*, SciPost Phys. **20**, 154 (2026), doi:[10.21468/SciPostPhys.20.6.154](https://doi.org/10.21468/SciPostPhys.20.6.154) [preprint doi:[10.48550/arXiv.2512.11036](https://doi.org/10.48550/arXiv.2512.11036)].
- [45] G. Heinrich, S. P. Jones, M. Kerner, G. Luisoni and E. Vryonidou, *NLO predictions for Higgs boson pair production with full top quark mass dependence matched to parton showers*, J. High Energy Phys. **08**, 088 (2017), doi:[10.1007/JHEP08\(2017\)088](https://doi.org/10.1007/JHEP08(2017)088) [preprint doi:[10.48550/arXiv.1703.09252](https://doi.org/10.48550/arXiv.1703.09252)].
- [46] M. Czakon, D. Heymes and A. Mitov, *fastNLO tables for NNLO top-quark pair differential distributions*, (arXiv preprint) doi:[10.48550/arXiv.1704.08551](https://doi.org/10.48550/arXiv.1704.08551).
- [47] G. Heinrich, S. P. Jones, M. Kerner and L. Scyboz, *A non-linear EFT description of  $gg \rightarrow HH$  at NLO interfaced to POWHEG*, J. High Energy Phys. **10**, 021 (2020), doi:[10.1007/JHEP10\(2020\)021](https://doi.org/10.1007/JHEP10(2020)021) [preprint doi:[10.48550/arXiv.2006.16877](https://doi.org/10.48550/arXiv.2006.16877)].

- [48] G. Heinrich, J. Lang and L. Scyboz, *SMEFT predictions for  $gg \rightarrow hh$  at full NLO QCD and truncation uncertainties*, J. High Energy Phys. **08**, 079 (2022), doi:[10.1007/JHEP08\(2022\)079](https://doi.org/10.1007/JHEP08(2022)079) [preprint doi:[10.48550/arXiv.2204.13045](https://doi.org/10.48550/arXiv.2204.13045)]. [Errata: J.HighEnergyPhys. 10, 086 (2023), doi:[10.1007/JHEP10\(2023\)086](https://doi.org/10.1007/JHEP10(2023)086)].
- [49] B. Agarwal, S. Jones, M. Kerner and A. von Manteuffel, *Complete next-to-leading order QCD corrections to ZZ production in gluon fusion*, Phys. Rev. Lett. **134**, 031901 (2025), doi:[10.1103/PhysRevLett.134.031901](https://doi.org/10.1103/PhysRevLett.134.031901) [preprint doi:[10.48550/arXiv.2404.05684](https://doi.org/10.48550/arXiv.2404.05684)].
- [50] B. C. Azeleira, G. Heinrich, M. Kerner and L. Kunz, *Probing anomalous Higgs boson couplings in Higgs plus jet production at NLO QCD with full  $m_t$ -dependence*, J. High Energy Phys. **04**, 060 (2025), doi:[10.1007/JHEP04\(2025\)060](https://doi.org/10.1007/JHEP04(2025)060) [preprint doi:[10.48550/arXiv.2409.05728](https://doi.org/10.48550/arXiv.2409.05728)].
- [51] S. Catani, S. Devoto, M. Grazzini, S. Kallweit, J. Mazzitelli and C. Savoini, *Higgs boson production in association with a top-antitop quark pair in next-to-next-to-leading order QCD*, Phys. Rev. Lett. **130**, 111902 (2023), doi:[10.1103/PhysRevLett.130.111902](https://doi.org/10.1103/PhysRevLett.130.111902) [preprint doi:[10.48550/arXiv.2210.07846](https://doi.org/10.48550/arXiv.2210.07846)].
- [52] F. Febres Cordero, G. Figueiredo, M. Kraus, B. Page and L. Reina, *Two-loop master integrals for leading-color  $pp \rightarrow t\bar{t}H$  amplitudes with a light-quark loop*, J. High Energy Phys. **07**, 084 (2024), doi:[10.1007/JHEP07\(2024\)084](https://doi.org/10.1007/JHEP07(2024)084) [preprint doi:[10.48550/arXiv.2312.08131](https://doi.org/10.48550/arXiv.2312.08131)].
- [53] G. Wang, T. Xia, L. L. Yang and X. Ye, *Two-loop QCD amplitudes for  $t\bar{t}H$  production from boosted limit*, J. High Energy Phys. **07**, 121 (2024), doi:[10.1007/JHEP07\(2024\)121](https://doi.org/10.1007/JHEP07(2024)121) [preprint doi:[10.48550/arXiv.2402.00431](https://doi.org/10.48550/arXiv.2402.00431)].
- [54] S. Devoto, M. Grazzini, S. Kallweit, J. Mazzitelli and C. Savoini, *Precise predictions for  $t\bar{t}H$  production at the LHC: Inclusive cross section and differential distributions*, J. High Energy Phys. **03**, 189 (2025), doi:[10.1007/JHEP03\(2025\)189](https://doi.org/10.1007/JHEP03(2025)189) [preprint doi:[10.48550/arXiv.2411.15340](https://doi.org/10.48550/arXiv.2411.15340)].
- [55] B. Agarwal, G. Heinrich, S. P. Jones, M. Kerner, S. Y. Klein, J. Lang, V. Magerya and A. Olsson, *Two-loop amplitudes for  $t\bar{t}H$  production: The quark-initiated  $N_f$ -part*, J. High Energy Phys. **05**, 013 (2024), doi:[10.1007/JHEP05\(2024\)013](https://doi.org/10.1007/JHEP05(2024)013) [preprint doi:[10.48550/arXiv.2402.03301](https://doi.org/10.48550/arXiv.2402.03301)]. [Errata: J.HighEnergyPhys. 06, 142 (2024), doi:[10.1007/JHEP06\(2024\)142](https://doi.org/10.1007/JHEP06(2024)142)].
- [56] ATLAS collaboration: G. Aad et al., *Precision calibration of calorimeter signals in the ATLAS experiment using an uncertainty-aware neural network*, SciPost Phys. **19**, 155 (2025), doi:[10.21468/SciPostPhys.19.6.155](https://doi.org/10.21468/SciPostPhys.19.6.155) [preprint doi:[10.48550/arXiv.2412.04370](https://doi.org/10.48550/arXiv.2412.04370)].
- [57] D. MacKay, *Probable networks and plausible predictions - A review of practical Bayesian methods for supervised neural networks*, Netw.: Comput. Neural Syst. **6**, 469 (1995), doi:[10.1088/0954-898X/6/3/011](https://doi.org/10.1088/0954-898X/6/3/011).
- [58] R. M. Neal, *Bayesian learning for neural networks*, PhD thesis, University of Toronto, Toronto, Canada (1995).
- [59] Y. Gal, *Uncertainty in deep learning*, PhD thesis, University of Cambridge, Cambridge, UK (2016).
- [60] A. Kendall and Y. Gal, *What uncertainties do we need in Bayesian deep learning for computer vision?*, (arXiv preprint) doi:[10.48550/arXiv.1703.04977](https://doi.org/10.48550/arXiv.1703.04977).

- [61] F. D'Angelo and V. Fortuin, *Repulsive deep ensembles are Bayesian*, (arXiv preprint) doi:[10.48550/arXiv.2106.11642](https://doi.org/10.48550/arXiv.2106.11642).
- [62] L. Röver, B. M. Schäfer and T. Plehn, *PINNferring the Hubble function with uncertainties*, (arXiv preprint) doi:[10.48550/arXiv.2403.13899](https://doi.org/10.48550/arXiv.2403.13899).
- [63] A. Amini, W. Schwarting, A. Soleimany and D. Rus, *Deep evidential regression*, (arXiv preprint) doi:[10.48550/arXiv.1910.02600](https://doi.org/10.48550/arXiv.1910.02600).
- [64] N. Meinert and A. Lavin, *Multivariate deep evidential regression*, (arXiv preprint) doi:[10.48550/arXiv.2104.06135](https://doi.org/10.48550/arXiv.2104.06135).
- [65] B. Kriesten and T. J. Hobbs, *Anomalous electroweak physics unraveled via evidential deep learning*, Eur. Phys. J. C **85**, 883 (2025), doi:[10.1140/epjc/s10052-025-14501-6](https://doi.org/10.1140/epjc/s10052-025-14501-6) [preprint doi:[10.48550/arXiv.2412.16286](https://doi.org/10.48550/arXiv.2412.16286)].
- [66] A. Khot, X. Wang, A. Roy, V. Kindratenko and M. Neubauer, *Evidential deep learning for uncertainty quantification and out-of-distribution detection in jet identification using deep neural networks*, Mach. Learn.: Sci. Technol. **6**, 035003 (2025), doi:[10.1088/2632-2153/ade51b](https://doi.org/10.1088/2632-2153/ade51b) [preprint doi:[10.48550/arXiv.2501.05656](https://doi.org/10.48550/arXiv.2501.05656)].
- [67] J. M. Wooldridge, *Econometric analysis of cross section and panel data*, MIT Press, Cambridge, USA, ISBN 9780262232586 (2010).
- [68] J. Eldridge, M. Belkin and Y. Wang, *Beyond Hartigan consistency: Merge distortion metric for hierarchical clustering*, (arXiv preprint) doi:[10.48550/arXiv.1506.06422](https://doi.org/10.48550/arXiv.1506.06422).
- [69] L. McInnes, J. Healy and S. Astels, *How HDBSCAN works*, [https://hdbscan.readthedocs.io/en/latest/how\\_hdbscan\\_works.html](https://hdbscan.readthedocs.io/en/latest/how_hdbscan_works.html).
- [70] D. W. Scott, *Multivariate density estimation*, Wiley, Hoboken, USA, ISBN 9780471697558 (2015), doi:[10.1002/9781118575574](https://doi.org/10.1002/9781118575574).
- [71] A. Gramacki, *Nonparametric kernel density estimation and its computational aspects*, Springer, Cham, Switzerland, ISBN 9783319716879 (2018), doi:[10.1007/978-3-319-71688-6](https://doi.org/10.1007/978-3-319-71688-6).
- [72] J. Braun, B. Campillo, G. Heinrich, M. Höfer, S. P. Jones, M. Kerner, J. Lang and V. Magerya, *One-loop calculations in effective field theories with GoSam-3.0*, SciPost Phys. Codebases **62** (2026), doi:[10.21468/scipostphyscodeb.62](https://doi.org/10.21468/scipostphyscodeb.62) [preprint doi:[10.48550/arXiv.2507.23549](https://doi.org/10.48550/arXiv.2507.23549)].  
J. Braun, B. Campillo, G. Heinrich, M. Höfer, S. P. Jones, M. Kerner, J. Lang and V. Magerya, *Codebase release 3.0 for GoSam*, SciPost Phys. Codebases **62-r3.0** (2026), doi:[10.21468/SciPostPhysCodeb.62-r3.0](https://doi.org/10.21468/SciPostPhysCodeb.62-r3.0).
- [73] J. Baglio, F. Campanario, S. Glaus, M. Mühlleitner, M. Spira and J. Streicher, *Gluon fusion into Higgs pairs at NLO QCD and the top mass scheme*, Eur. Phys. J. C **79**, 459 (2019), doi:[10.1140/epjc/s10052-019-6973-3](https://doi.org/10.1140/epjc/s10052-019-6973-3) [preprint doi:[10.48550/arXiv.1811.05692](https://doi.org/10.48550/arXiv.1811.05692)].
- [74] E. Bagnaschi, G. Degrassi and R. Gröber, *Higgs boson pair production at NLO in the Powheg approach and the top quark mass uncertainties*, Eur. Phys. J. C **83**, 1054 (2023), doi:[10.1140/epjc/s10052-023-12238-8](https://doi.org/10.1140/epjc/s10052-023-12238-8) [preprint doi:[10.48550/arXiv.2309.10525](https://doi.org/10.48550/arXiv.2309.10525)].
- [75] S. Jaskiewicz, S. Jones, R. Szafron and Y. Ulrich, *The structure of quark mass corrections in the  $gg \rightarrow HH$  amplitude at high-energy*, J. High Energy Phys. **09**, 015 (2025), doi:[10.1007/JHEP09\(2025\)015](https://doi.org/10.1007/JHEP09(2025)015) [preprint doi:[10.48550/arXiv.2501.00587](https://doi.org/10.48550/arXiv.2501.00587)].

- [76] J. Davies, K. Schönwald and M. Steinhauser, *Three-loop large- $N_c$  virtual corrections to  $gg \rightarrow HH$  in the forward limit*, J. High Energy Phys. **08**, 192 (2025), doi:[10.1007/JHEP08\(2025\)192](https://doi.org/10.1007/JHEP08(2025)192) [preprint doi:[10.48550/arXiv.2503.17449](https://doi.org/10.48550/arXiv.2503.17449)].
- [77] S. Borowka, N. Greiner, G. Heinrich, S. P. Jones, M. Kerner, J. Schlenk, U. Schubert and T. Zirke, *Higgs boson pair production in gluon fusion at next-to-leading order with full top-quark mass dependence*, Phys. Rev. Lett. **117**, 012001 (2016), doi:[10.1103/PhysRevLett.117.012001](https://doi.org/10.1103/PhysRevLett.117.012001) [preprint doi:[10.48550/arXiv.1604.06447](https://doi.org/10.48550/arXiv.1604.06447)]. [Errata: Phys.Rev.Lett. 117, 079901 (2016), doi:[10.1103/PhysRevLett.117.079901](https://doi.org/10.1103/PhysRevLett.117.079901)].
- [78] S. Borowka, N. Greiner, G. Heinrich, S. P. Jones, M. Kerner, J. Schlenk and T. Zirke, *Full top quark mass dependence in Higgs boson pair production at NLO*, J. High Energy Phys. **10**, 107 (2016), doi:[10.1007/JHEP10\(2016\)107](https://doi.org/10.1007/JHEP10(2016)107) [preprint doi:[10.48550/arXiv.1608.04798](https://doi.org/10.48550/arXiv.1608.04798)].
- [79] *HHgrid*, Max Planck Institute for Physics, Munich, Germany, <https://github.com/mppmu/hhgrid>.
- [80] J. Davies, G. Heinrich, S. P. Jones, M. Kerner, G. Mishima, M. Steinhauser and D. Wellmann, *Double Higgs boson production at NLO: Combining the exact numerical result and high-energy expansion*, J. High Energy Phys. **11**, 024 (2019), doi:[10.1007/JHEP11\(2019\)024](https://doi.org/10.1007/JHEP11(2019)024) [preprint doi:[10.48550/arXiv.1907.06408](https://doi.org/10.48550/arXiv.1907.06408)].
- [81] E. Jones et al., *SciPy: Open source scientific tools for Python*, <http://www.scipy.org/> (2001).
- [82] A. Sommerfeld, *Atombau und Spektralliniem*, Vieweg & Sohn, Braunschweig, Germany (1924).
- [83] W. Beenakker, S. Dittmaier, M. Krämer, B. Plümper, M. Spira and P. M. Zerwas, *NLO QCD corrections to  $t\bar{t}H$  production in hadron collisions*, Nucl. Phys. B **653**, 151 (2003), doi:[10.1016/S0550-3213\(03\)00044-0](https://doi.org/10.1016/S0550-3213(03)00044-0) [preprint doi:[10.48550/arXiv.hep-ph/0211352](https://doi.org/10.48550/arXiv.hep-ph/0211352)].