



City Research Online

City St George's, University of London

Citation: Araz, J. Y., Beck, A., Reboud, M., Spannowsky, M. & van Dyk, D. (2026). Communicating likelihoods with normalising flows. The European Physical Journal C, 86(7), 805. doi: 10.1140/epjc/s10052-026-16045-9

This is the published version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/37973/>

Link to published version: <https://doi.org/10.1140/epjc/s10052-026-16045-9>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).



Communicating likelihoods with normalising flows

Jack Y. Araz^{1,2,3,a} , Anja Beck^{4,b} , M  ril Reboud^{5,c} , Michael Spannowsky^{6,d} , Danny van Dyk^{6,e}

¹ Department of Physics and Astronomy, University College London, London WC1E 6B, UK

² School of Science & Technology, City St. George's, University of London, London EC1V 0HB, UK

³ Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794, USA

⁴ Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁵ Universit   Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France

⁶ Department of Physics, Institute for Particle Physics Phenomenology, Durham University, Durham DH1 3LE, UK

Received: 24 September 2025 / Accepted: 23 June 2026
  The Author(s) 2026

Abstract We present a machine-learning-based workflow to model an unbinned likelihood from its samples. A key advancement over existing approaches is the validation of the learned likelihood using rigorous statistical tests, such as the Kolmogorov–Smirnov test of the joint distribution. Our method enables the reliable communication of experimental and phenomenological likelihoods for subsequent analyses. We demonstrate its effectiveness through three case studies in high-energy physics. To support broader adoption, we provide an open-source reference implementation, *nabu*.

1 Introduction

The reinterpretability of experimental results is a cornerstone of progress in high-energy physics, where the sheer expense and effort required to generate data preclude the luxury of rerunning experiments. In collider experiments, for instance, the complexity of the data and the sophisticated workflows needed for its analysis mean that sharing results in a way that enables reinterpretation is both a scientific imperative and a technical challenge [1–15]. Traditional approaches, such as sharing either the raw experimental data or likelihoods, present significant limitations. Raw data sharing is resource-intensive and raises concerns about data management and accessibility [16, 17]. Providing full likelihoods [18, 19] is computationally expensive and often impractical due to large file sizes and the high number of experimental nuisance

parameters, as evidenced by the substantial efforts behind LHC combinations of Higgs measurements [20–22]. Simplified likelihoods mitigate these issues but introduce assumptions that can overly constrain their usability [23].

We suggest using Likelihood Models (LMs) instead. A LM is a mathematical representation that quantifies the probability of observing a specific data set given underlying parameters. In high-energy physics, likelihood models are essential for interpreting experimental results and comparing theoretical predictions against observed data. These models strive to encapsulate the complete statistical information of an experiment, making them a cornerstone for parameter estimation [24], hypothesis testing [25, 26], and model comparison [27–30]. However, constructing and utilising these models can be computationally challenging, mainly when dealing with high-dimensional parameter spaces or complex workflows that include detector effects and systematic uncertainties.

ML has emerged as a transformative tool in high-energy physics in recent years. ML-based approaches offer potent solutions for complex, high-dimensional problems arising in detector simulation, event reconstruction, and parameter inference. Among these, Normalizing Flows (NFs) stand out as a promising avenue for constructing surrogate models of likelihoods. By learning the underlying probability distribution of experimental data, Normalizing Flows (NFs) can provide flexible and accurate approximations of the likelihood function that are computationally efficient to evaluate. These surrogate models enable data sharing in a compact, interpretable format and facilitate downstream tasks such as inference and model comparison. Previous work in this direction (see, for example, Refs. [31–33]) has demonstrated the potential of Normalizing Flows (NFs) to encapsulate the essential

^a e-mail: j.araz@ucl.ac.uk

^b e-mail: anja.beck@cern.ch

^c e-mail: merilreboud@gmail.com

^d e-mail: michael.spannowsky@durham.ac.uk

^e e-mail: danny.van.dyk@gmail.com (corresponding author)

features of high-dimensional distributions while remaining computationally tractable.

A parallel challenge in high-energy physics lies in the computational cost of the simulation pipeline. Monte Carlo simulations, augmented by parton showering and detector simulations, are indispensable for modelling experimental interactions. However, their expense often limits their use, especially in scenarios requiring iterative evaluation, such as Bayesian inference or optimisation of theoretical models. Differentiable programming workflows [34–42], which integrate gradients throughout the simulation pipeline, offer a potential solution. These workflows allow efficient navigation of the parameter space but are currently limited by the monumental effort required to adapt existing simulations to a fully differentiable framework.

Our work bridges these gaps by employing Normalizing Flows (NFs) to build surrogate models that approximate both the likelihood function and the simulation workflow, aligning well with simulation-based inference [43–49]. This dual application accelerates parameter inference and extends the utility of surrogate models to data generation and pseudo-experimental analysis tasks. By designing the surrogate models to capture the essential features of the experimental data while being accessible to a wide range of researchers, we ensure that our approach supports the community’s broader goals of reproducibility and reinterpretability [50–53].

The following sections outline a framework for NF-based surrogate models to represent experimental likelihoods. We provide practical methodologies for their training and validation and demonstrate their application to various high-energy physics use cases. Additionally, we discuss how these surrogate models can serve as a stepping stone towards a fully differentiable simulation pipeline, enabling efficient, iterative exploration of theoretical and experimental parameter spaces. By intentionally designing our framework to be user-friendly and broadly applicable, we ensure that it can be adopted across the community for diverse reinterparation and inference tasks.

2 Methods and step-by-step recipe

As part of our workflow, we pursue the following objectives:

- constructing a LM from pre-existing samples¹;
- testing the compatibility of this LM with respect to the samples;

¹ Throughout, “sample” refers to a random D -dimensional variate of a probability distribution, e.g. as obtained by a particle-physics experiment from a decay or scattering process, or obtained in a Bayesian data analysis from an a-posteriori probability density; concrete examples are discussed in Sect. 3.

- drawing approximate samples from this LM; and
- storing the LM for future use, e.g. the evaluation of the model density or generation of additional approximate samples.

A Python implementation of our workflow is available through the open source software *nabu* [54],² which is built on *flowjax* [55], *equinox* [56], and *jax* [57]. The documentation of *nabu* will be made available [54].

Preparations In many cases, a “standardisation” of the samples (i.e. transformation of the samples to a parameter space with standard properties) is useful before training. We support such standardisation, e.g., through transforming each variable in the D -dimensional data set to new variables with zero mean and unit variance. In high-energy physics, Dalitz analyses yield a special case of non-Cartesian bounded parameter space, as shown in one of the examples below. For such cases, we provide the means to transform the samples to a Cartesian parameter space. After the optional “standardisation” step, the workflow splits the available data set of size N into a training set ($N_{\text{tr}} \simeq 72\% \times N$ by default), a validation set ($N_{\text{v}} \simeq 8\% \times N$ by default), and a testing set ($N_{\text{test}} \simeq 20\% \times N$ by default).

Model We train a LM $\tilde{L}(\vec{\Theta})$ on the training data set, comprised of samples $\vec{\vartheta}_1$ to $\vec{\vartheta}_{N_{\text{tr}}}$. Here $\vec{\Theta}$ represent a D -dimensional random variable, and each sample $\vec{\vartheta}_i \in T \equiv \mathbb{R}^D$ is distributed $\vec{\vartheta}_i \sim L$. Following the convention of Ref. [33], we refer to the vector space T as the *target space*. If the LM $\tilde{L}(\vec{\Theta})$ is trained sufficiently and successfully, it can serve as a proxy for the true likelihood $L(\vec{\Theta})$.

We propose to use the following type of D -dimensional LM,

$$\tilde{L}(\vec{\Theta}) = |J_f| \mathcal{N}^D(f(\vec{\Theta}) | \vec{\mu} = \vec{0}, \Sigma = \mathbf{1}), \quad (1)$$

where $\mathcal{N}^D(\vec{\beta} | \vec{\mu}, \Sigma)$ represents a multivariate Gaussian PDF centred around μ with covariance matrix Σ . The combination of a suitable optimisable bijective transformation f applied to $\vec{\Theta}$ and its Jacobian J_f make it possible for $\tilde{L}$ to mimic the true likelihood L [58, 59] for the purpose of evaluation or generation of additional samples. Such a bijection is commonly called a NF from the target space into the base space B ,

$$f(\vec{\Theta}) \equiv f(\vec{\Theta}; \vec{v}) = \vec{\beta}(\vec{v}) \in B \equiv \mathbb{R}^D. \quad (2)$$

Here $\vec{v}$ represent a set of model-specific nuisance parameters that need to be inferred to ensure agreement between the LM, $\tilde{L}$, and the genuine likelihood L . Commonly, f is constructed as a sequence of N_L flow layers of individual bijections,

² Named after the Assyrian patron god of scribes and wisdom.

$f \equiv f_1 \circ f_2 \circ \dots \circ f_{N_L}$. Our default choice of Normalizing Flows (NFs) employs MAF [60] alternated with configurable permutations. Our implementation [54] also supports further types of Normalizing Flows (NFs), including e.g. non-volume preserving transformations (RealNVP) [61]. We keep the number of flow layers configurable. Each flow layer is associated with an Multilayer Perceptron (MLP) to learn the functions used in the underlying bijections. The width, depth, and choice of activation functions of these MLPs are configurable. As a transformation function, we employ either a simple affine transformation or RQS [62] where the transformation parameters have been set via an MLP. The trainable nuisance parameter set $\vec{v}$ is comprised of the union of the parameters of all MLPs.

The LM is trained by minimising a loss function ℓ_{tr} with respect to the nuisance parameters $\vec{v}$. This training is performed on a configurable number of batches of equal sample size N_{batch} . For each batch, we define its loss ℓ as the mean value of the unbinned standard-normal likelihood of the samples in base space,

$$\ell_{\text{tr}}(\vec{v}) = -\frac{1}{N_{\text{batch}}} \sum_i^{N_{\text{batch}}} \ln \left[\mathcal{N}^D(\vec{\beta}_i(\vec{v}) | \vec{0}, \mathbb{1}) \right]. \quad (3)$$

Here, the sum iterates over all samples in a given batch, and we abbreviate $\vec{\beta}_i(\vec{v}) \equiv f(\vec{v}_i; \vec{v})$. We minimise this loss for each batch, using the “ADAM” [63] minimiser by default, with an adjustable learning rate. A training epoch is completed once every batch has been used for minimisation. At this point, we record the mean of the loss value across all training batches.

We repeat the loss computation using the validation set, yielding ℓ_v . The training procedure stops either after a configurable number of epochs have passed or if the validation loss ℓ_v stops decreasing for a configurable number of epochs.

Testing After completion of the training, the base-space samples should follow a standard normal distribution if the LM has adequately learned the features inherent to the training sample set. We test if this is true by using a suitable test statistic. A training outcome is rejected if its p -value is smaller than a pre-defined threshold (default: 3%). We find this to be essential, since we encountered training outcomes that passed visual inspection despite a p -value below this threshold.

Our primary choice of test statistic is the unbinned KS [64] test on the PDF for the square of the two-norm of the transformed samples $\vec{\beta}_i$. For a D -dimensional standard normal PDF of the samples, the PDF of the squared two-norm is a

$$\chi^2 = ||\vec{\beta}||^2 = \sum_{i=1}^D |\beta_i|^2$$

distribution with D degrees of freedom. As a consequence of having a known joint PDF, we can efficiently perform the KS test on our training and validation sets after every epoch, if so requested by the user. Our approach has benefits compared to other works in the literature [32]. First, our approach utilises a single *known* distribution that tests the joint distribution rather than each marginal PDF. Second, the known CDF increases the performance of the KS test and removes the Monte-Carlo uncertainty for the p -value, which arises from the limited size of the generated testing sets in two-sample tests, e.g. for unknown distributions or distributions without an evaluable CDF. Further checks of the marginal distributions remain possible and are supported by nabu. However, given the combination of expecting convergence toward a multivariate Gaussian distribution in $\vec{\beta}$ (due to the likelihood used in Eq. (1)) and passing the p value threshold for the KS tests in χ^2 , no such further checks uncovered problems in our examples.


Our secondary choice of test statistic is primarily intended for visualisation: we split the two-norm samples $||\vec{\beta}_i||_2^2$ into K bins of equal probability content. For a large number of samples, $N/K \gtrsim 100$, the number of samples in each bin follows a normal distribution with $\mu = N/K$ and $\sigma = \sqrt{N/K}$. We use this expected mean and standard deviation to perform a binned χ^2 test on combining the first $K - 1$ bins.

Storage We provide the means to store an LM in a single file representing the NF and all coordinate transformations applied to the sample set before training. The file format makes use of JSON and the serialization routines provided by the equinox package [56]. This approach facilitates third parties’ dissemination and usage of the LM.


3 Concrete examples

We illustrate the performance of our workflow using real-world examples. Our selection of examples covers experimental and phenomenological results with a variety of different dimensionalities and complexities of distribution. Figure 1 contains a summary plot for each example. The upper part of each plot shows the distribution of the two-norms of the transformed samples after successful training. Overlaid is the corresponding expected χ^2 distribution. The calculated p -values are presented at the top, and the residuals part of each plot illustrates the similarity between the empirical and expected distributions. The significance of each residual is shown using colour coding, with yellow bins indicating 1σ to 2σ fluctuations and red bins indicating a fluctuation beyond 2σ .

In example (a), we train a LM on $pp \rightarrow Z(\rightarrow \mu\mu) + \text{jet}$ events [65] published by the ATLAS collaboration [66]. The original data set features 24 dimensions. However, in the



(a) ATLAS data [65] from Ref. [66].



(b) LHCb Dalitz plot data from Ref. [67].


(c) WET $b \rightarrow ul^- \bar{\nu}$ samples [68] from Ref. [69].

Fig. 1 Summary plots for the three training examples. We show the p -value corresponding to the KS and binned χ^2 tests. The “Density” part of each plot shows the 10 bins in $\|\vec{\beta}\|^2$ and the expected χ^2 PDF as a blue curve. Each bin is expected to contain 10% of the testing set for a perfect model. In the “Residuals” part of each plot, the grey, gold, and red coloured bins indicate deviations of less than 1σ , within $[1, 2)\sigma$, and more than 2σ , respectively

underlying $2 \rightarrow 3$ -body process, most physical information is expected to be encoded in only five independent observables. We choose $p_T^{\mu\mu}$, $y^{\mu\mu}$, $p_T^{\mu 1}$, $\Delta\eta_{12}^2 = \eta_{\mu 1}^2 - \eta_{\mu 2}^2$, $\Delta\phi_{12}^2 = \phi_{\mu 1}^2 - \phi_{\mu 2}^2$ as representative observables for this physical information. The data set has been standardised by scaling and shifting all parameters to the range $[0, 1]$. This is followed by applying a logarithmic transformation to the modified $p_T^{\mu\mu}$ and $p_T^{\mu 1}$ distributions and rescaling them to $[0, 1]$ once more. This approach mitigates challenges in learning distributions that extend to very large values.

An eight-layer NF is employed for training, utilising MAF bijections with RQS transformations. Each MAF layer contains a single-layer MLP with 512 hidden units, while the RQS transformation is configured with 12 knots within the $[0, 1]$ support. The initial learning rate is set to 1% and decays with a half-life of 25 epochs until reaching a minimal learning rate of 10^{-6} . The NF is trained for a maximum of 600 epochs, with early stopping applied if the negative log probability of the validation set does not improve for more than 50 epochs. Testing the LM yields a KS p -value of 66.5%. The outcome of this example is illustrated in Fig. 1. The resulting model file uses ~ 4 MB of storage space, while the compressed samples use ~ 14 MB. In Fig. 2 we showcase that the LM is in good agreement with the 1D marginal distributions (using $p_T^{\mu\mu}$ as an example), despite our validation only taking the overall χ^2 into account.

In example (b), we train an LM on efficiency corrected events of the decay $B^+ \rightarrow D^+ D^- K^+$ as published by the LHCb collaboration [67]. The events have been extracted from the supplementary material. Each event corresponds to a 2-tuple of the squared $D^- K^+$ mass and the squared $D^+ D^-$ mass. Before training, we manually map the support of the Dalitz plot to the space $(-\infty, +\infty)^2$. For this two-dimensional example, we require 8 flow layers of type RQS, and an MLP with two hidden layers & 64 neurons per layer to learn the likelihood. After testing, our LM yields acceptable p -values of 10% and 79.9% for the KS test and the binned χ^2 test, respectively. The current LHCb measurement does not yet provide sufficient detail to highlight resonance bands in $m_{D^+ D^-}^2$ as expected from the presence of broad charmonium resonances above the open-charm threshold. On the other hand, we are content that the LM does not “hallucinate” non-existing structures. This example is, therefore, a good illustration of the stability of our choice of LM. The resulting model file uses ~ 300 kB of storage space, while the samples use ~ 14 kB.

In example (c), we train an LM on the posterior samples of the Weak Effective Theory parameters in the $ub\ell\nu$ sector [68], obtained from a fit to exclusive $\bar{B} \rightarrow \{\pi, \rho, \omega\}\ell^- \bar{\nu}$ decays [69]. This data set contains 5-dimensional samples following a multimodal distribution. The underlying data is the same as used for the pilot study presented in Ref. [33]. In




Fig. 2 Diagnostic plot demonstrating the agreement between the binned prediction (blue bins) of the learned LM for example (a) and the binned samples (black points) using the variable $p_T^{\mu\mu}$ as an example. For an explanation of the colour scheme of the “Residuals”, see the caption of Fig. 1

contrast to earlier efforts, our present approach is *successful* in learning the full five-dimensional distribution, as indicated by its acceptable p -values of 18.4% and 15.9% for the KS test and the binned χ^2 test, respectively. As a consequence, the results of Ref. [69] can now be used in subsequent BSM model studies or SMEFT studies, without repeating the original analysis. This represents a substantial reduction of the computation costs, given the large number of 50 hadronic nuisance parameters that are irrelevant to any BSM implications. The resulting model file uses ~ 90 kB of storage space, while the compressed samples use ~ 5 MB.

4 Conclusion and outlook

Analysis preservation, and reinterpretation of experimental and phenomenological analyses is a significant hurdle to progress in the field of high-energy physics. Here, we have presented a workflow to model unbinned likelihoods for this purpose. When constructing likelihood models, we make use of established techniques (such as normalising flows) and pieces of software that emerged in the context of machine-learning applications. We combine these techniques with robust statistical tests that ensure the likelihood model is accurate. Using three concrete real-world examples from experiment and phenomenology, we have illustrated that our proposed workflow can model likelihoods with complicated non-Gaussian features and strong correlations. Except for our two-dimensional example, the storage size of our mod-

els is much smaller than the size of the training samples. This confirms an advantage of our proposed workflow for data preservation and storage.

We will continue developing the *nabu* software [54], expanding its capabilities with additional types of normalising flows and a broader range of known transformations. Future applications of *nabu* include integration with the EOS software for flavour physics [70], enabling precise posterior constraints on Weak Effective Theory parameters inferred from flavour-changing processes [69, 71, 72]. Additionally, *nabu* will interface with *Spey* [23] to accelerate analysis reinterpretation by constructing a comprehensive catalogue of likelihood models from publicly available experimental data. Looking ahead, *nabu* is designed for extendability beyond normalising flows, including advanced machine learning techniques such as diffusion models, which we aim to explore in future work.

Acknowledgements We thank Peter Stangl for his helpful comments on the manuscript. MS and DvD acknowledge support by the UK Science and Technology Facilities Council (STFC) through grant ST/X003167/1. DvD further acknowledges UK STFC support through grant ST/V003941/1. AB wants to acknowledge the support of MIT’s Department of Physics.

Data Availability Statement This manuscript has associated data in a data repository. [Author’s comment: The datasets analysed in the course of this study are available through references [65, 67, 68]. The datasets generated in the course of this study are available through <https://doi.org/10.5281/zenodo.14902619>.]

Code Availability Statement This manuscript has associated code/software in a data repository. [Author’s comment: The software written in support of this study is available through reference [54] and the <https://github.com/nabu-hep/nabu-hep>.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.
Funded by SCOAP³.

References

1. G. Alguero, J.Y. Araz, B. Fuks, S. Kraml, SciPost Phys. **14**, 009 (2023). <https://doi.org/10.21468/SciPostPhys.14.1.009>. [arXiv:2206.14870](https://arxiv.org/abs/2206.14870) [hep-ph]
2. J.Y. Araz, B. Fuks, M.D. Goodsell, M. Utsch, Eur. Phys. J. C **82**, 597 (2022). <https://doi.org/10.1140/epjc/s10052-022-10511-w>. [arXiv:2112.05163](https://arxiv.org/abs/2112.05163) [hep-ph]

3. J.Y. Araz, B. Fuks, G. Polykratis, *Eur. Phys. J. C* **81**, 329 (2021). <https://doi.org/10.1140/epjc/s10052-021-09052-5>. [arXiv:2006.09387](https://arxiv.org/abs/2006.09387) [hep-ph]
4. E. Conte, B. Fuks, *Int. J. Mod. Phys. A* **33**, 1830027 (2018). <https://doi.org/10.1142/S0217751X18300272>. [arXiv:1808.00480](https://arxiv.org/abs/1808.00480) [hep-ph]
5. B. Dumont, B. Fuks, S. Kraml, S. Bein, G. Chalons, E. Conte, S. Kulkarni, D. Sengupta, C. Wymant, *Eur. Phys. J. C* **75**, 56 (2015). <https://doi.org/10.1140/epjc/s10052-014-3242-3>. [arXiv:1407.3278](https://arxiv.org/abs/1407.3278) [hep-ph]
6. E. Conte, B. Dumont, B. Fuks, C. Wymant, *Eur. Phys. J. C* **74**, 3103 (2014). <https://doi.org/10.1140/epjc/s10052-014-3103-0>. [arXiv:1405.3982](https://arxiv.org/abs/1405.3982) [hep-ph]
7. E. Conte, B. Fuks, G. Serret, *Comput. Phys. Commun.* **184**, 222 (2013). <https://doi.org/10.1016/j.cpc.2012.09.009>. [arXiv:1206.1599](https://arxiv.org/abs/1206.1599) [hep-ph]
8. J.Y. Araz, A. Buckley, B. Fuks, *Eur. Phys. J. C* **83**, 664 (2023). <https://doi.org/10.1140/epjc/s10052-023-11779-2>. [arXiv:2303.03427](https://arxiv.org/abs/2303.03427) [hep-ph]
9. C. Bierlich et al., *SciPost Phys.* **8**, 026 (2020). <https://doi.org/10.21468/SciPostPhys.8.2.026>. [arXiv:1912.05451](https://arxiv.org/abs/1912.05451) [hep-ph]
10. A. Buckley, D. Kar, K. Nordström, *SciPost Phys.* **8**, 025 (2020). <https://doi.org/10.21468/SciPostPhys.8.2.025>. [arXiv:1910.01637](https://arxiv.org/abs/1910.01637) [hep-ph]
11. A. Buckley, J. Butterworth, D. Grellscheid, H. Hoeth, L. Lonnblad, J. Monk, H. Schulz, F. Siegert, *Comput. Phys. Commun.* **184**, 2803 (2013). <https://doi.org/10.1016/j.cpc.2013.05.021>. [arXiv:1003.0694](https://arxiv.org/abs/1003.0694) [hep-ph]
12. G. Alguero, J. Heisig, C.K. Khosa, S. Kraml, S. Kulkarni, A. Lessa, H. Reyes-González, W. Waltenberger, A. Wongel, *JHEP* **08**, 068 (2022). [https://doi.org/10.1007/JHEP08\(2022\)068](https://doi.org/10.1007/JHEP08(2022)068). [arXiv:2112.00769](https://arxiv.org/abs/2112.00769) [hep-ph]
13. G. Alguero, S. Kraml, W. Waltenberger, *Comput. Phys. Commun.* **264**, 107909 (2021). <https://doi.org/10.1016/j.cpc.2021.107909>. [arXiv:2009.01809](https://arxiv.org/abs/2009.01809) [hep-ph]
14. F. Ambrogio et al., *Comput. Phys. Commun.* **251**, 106848 (2020). <https://doi.org/10.1016/j.cpc.2019.07.013>. [arXiv:1811.10624](https://arxiv.org/abs/1811.10624) [hep-ph]
15. M.M. Altakach, S. Kraml, A. Lessa, S. Narasimha, T. Pascal, C. Ramos, Y. Villamizar, W. Waltenberger, *JHEP* **11**, 074 (2024). [https://doi.org/10.1007/JHEP11\(2024\)074](https://doi.org/10.1007/JHEP11(2024)074). [arXiv:2409.12942](https://arxiv.org/abs/2409.12942) [hep-ph]
16. CERN, CERN open data project, <https://opendata.cern.ch/> (2014–2024)
17. C.A. Aidala, C. Burr, M. Cattaneo, D.S. Fitzgerald, A. Morris, S. Neubert, D. Tromann, *Comput. Softw. Big Sci.* **7**, 6 (2023). <https://doi.org/10.1007/s41781-023-00099-5>. [arXiv:2302.14235](https://arxiv.org/abs/2302.14235) [hep-ex]
18. L. Heinrich, M. Feickert, G. Stark, K. Cranmer, *J. Open Source Softw.* **6**, 2823 (2021). <https://doi.org/10.21105/joss.02823>
19. K. Cranmer, G. Lewis, L. Moneta, A. Shibata, W. Verkerke, (ROOT), (2012). <https://doi.org/10.17181/CERN-OPEN-2012-016>
20. G. Aad et al. (ATLAS, CMS), *Phys. Rev. Lett.* **114**, 191803 (2015). <https://doi.org/10.1103/PhysRevLett.114.191803>. [arXiv:1503.07589](https://arxiv.org/abs/1503.07589) [hep-ex]
21. G. Aad et al. (ATLAS, CMS), *JHEP* **08**, 045 (2016). [https://doi.org/10.1007/JHEP08\(2016\)045](https://doi.org/10.1007/JHEP08(2016)045). [arXiv:1606.02266](https://arxiv.org/abs/1606.02266) [hep-ex]
22. G. Aad et al. (ATLAS, CMS), *Phys. Rev. Lett.* **132**, 021803 (2024). <https://doi.org/10.1103/PhysRevLett.132.021803>. [arXiv:2309.03501](https://arxiv.org/abs/2309.03501) [hep-ex]
23. J.Y. Araz, *SciPost Phys.* **16**, 032 (2024). <https://doi.org/10.21468/SciPostPhys.16.1.032>. [arXiv:2307.06996](https://arxiv.org/abs/2307.06996) [hep-ph]
24. H. Flacher, M. Goebel, J. Haller, A. Hocker, K. Monig, J. Stelzer, *Eur. Phys. J. C* **60**, 543 (2009) [Erratum: *Eur.Phys.J.C* **71**, 1718 (2011)]. <https://doi.org/10.1140/epjc/s10052-009-0966-6>. [arXiv:0811.0009](https://arxiv.org/abs/0811.0009) [hep-ph]
25. C. Englert, R. Kogler, H. Schulz, M. Spannowsky, *Eur. Phys. J. C* **76**, 393 (2016). <https://doi.org/10.1140/epjc/s10052-016-4227-1>. [arXiv:1511.05170](https://arxiv.org/abs/1511.05170) [hep-ph]
26. C. Englert, R. Kogler, H. Schulz, M. Spannowsky, *Eur. Phys. J. C* **77**, 789 (2017). <https://doi.org/10.1140/epjc/s10052-017-5366-8>. [arXiv:1708.06355](https://arxiv.org/abs/1708.06355) [hep-ph]
27. M. Drees, H. Dreiner, D. Schmeier, J. Tattersall, J.S. Kim, *Comput. Phys. Commun.* **187**, 227 (2015). <https://doi.org/10.1016/j.cpc.2014.10.018>. [arXiv:1312.2591](https://arxiv.org/abs/1312.2591) [hep-ph]
28. D. Dercks, N. Desai, J.S. Kim, K. Rolbiecki, J. Tattersall, T. Weber, *Comput. Phys. Commun.* **221**, 383 (2017). <https://doi.org/10.1016/j.cpc.2017.08.021>. [arXiv:1611.09856](https://arxiv.org/abs/1611.09856) [hep-ph]
29. P. Athron et al. (GAMBIT), *Eur. Phys. J. C* **77**, 784 (2017) [Addendum: *Eur.Phys.J.C* **78**, 98 (2018)]. <https://doi.org/10.1140/epjc/s10052-017-5321-8>. [arXiv:1705.07908](https://arxiv.org/abs/1705.07908) [hep-ph]
30. S. Amrith, J.M. Butterworth, F.F. Deppisch, W. Liu, A. Varma, D. Yallup, *JHEP* **05**, 154 (2019). [https://doi.org/10.1007/JHEP05\(2019\)154](https://doi.org/10.1007/JHEP05(2019)154). [arXiv:1811.11452](https://arxiv.org/abs/1811.11452) [hep-ph]
31. A. Coccaro, M. Pierini, L. Silvestrini, R. Torre, *Eur. Phys. J. C* **80**, 664 (2020). <https://doi.org/10.1140/epjc/s10052-020-8230-1>. [arXiv:1911.03305](https://arxiv.org/abs/1911.03305) [hep-ph]
32. H. Reyes-Gonzalez, R. Torre, *SciPost Phys. Core* **7**, 048 (2024). <https://doi.org/10.21468/SciPostPhysCore.7.3.048>. [arXiv:2309.09743](https://arxiv.org/abs/2309.09743) [hep-ph]
33. A. Beck, M. Reboud, D. van Dyk, *Eur. Phys. J. C* **83**, 1115 (2023). <https://doi.org/10.1140/epjc/s10052-023-12294-0>. [arXiv:2309.10365](https://arxiv.org/abs/2309.10365) [hep-ph]
34. L. Heinrich, (2022). [arXiv:2203.13079](https://arxiv.org/abs/2203.13079) [stat.ME]
35. R.E.C. Smith, I. Ochoa, R. Inácio, J. Shoemaker, M. Kagan, *Phys. Rev. D* **110**, 052010 (2024). <https://doi.org/10.1103/PhysRevD.110.052010>. [arXiv:2310.12804](https://arxiv.org/abs/2310.12804) [hep-ex]
36. A. Adelmann et al., in *Snowmass 2021* (2022). [arXiv:2203.08806](https://arxiv.org/abs/2203.08806) [hep-ph]
37. T. Heimel, R. Winterhalder, A. Butter, J. Isaacson, C. Krause, F. Maltoni, O. Mattelaer, T. Plehn, *SciPost Phys.* **15**, 141 (2023). <https://doi.org/10.21468/SciPostPhys.15.4.141>. [arXiv:2212.06172](https://arxiv.org/abs/2212.06172) [hep-ph]
38. M. Kagan, L. Heinrich, (2023). [arXiv:2308.16680](https://arxiv.org/abs/2308.16680) [stat.ML]
39. A. Kofler, V. Stimper, M. Mikhasenko, M. Kagan, L. Heinrich, in *38th Conference on Neural Information Processing Systems* (2024). [arXiv:2411.16234](https://arxiv.org/abs/2411.16234) [hep-ph]
40. M. Vigl, N. Hartman, L. Heinrich, *Mach. Learn. Sci. Technol.* **5**, 025075 (2024). <https://doi.org/10.1088/2632-2153/ad55a3>. [arXiv:2401.13536](https://arxiv.org/abs/2401.13536) [hep-ex]
41. T. Heimel, N. Huetsch, F. Maltoni, O. Mattelaer, T. Plehn, R. Winterhalder, *SciPost Phys.* **17**, 023 (2024). <https://doi.org/10.21468/SciPostPhys.17.1.023>. [arXiv:2311.01548](https://arxiv.org/abs/2311.01548) [hep-ph]
42. T. Heimel, O. Mattelaer, T. Plehn, R. Winterhalder, *SciPost Phys.* **18**, 017 (2025). <https://doi.org/10.21468/SciPostPhys.18.1.017>. [arXiv:2408.01486](https://arxiv.org/abs/2408.01486) [hep-ph]
43. K. Cranmer, J. Brehmer, G. Louppe, *Proc. Natl. Acad. Sci.* **117**, 30055 (2020). <https://doi.org/10.1073/pnas.1912789117>. [arXiv:1911.01429](https://arxiv.org/abs/1911.01429) [stat.ML]
44. J. Brehmer, F. Kling, I. Espejo, K. Cranmer, *Comput. Softw. Big Sci.* **4**, 3 (2020). <https://doi.org/10.1007/s41781-020-0035-2>. [arXiv:1907.10621](https://arxiv.org/abs/1907.10621) [hep-ph]
45. J. Brehmer, K. Cranmer, G. Louppe, J. Pavez, *Phys. Rev. D* **98**, 052004 (2018). <https://doi.org/10.1103/PhysRevD.98.052004>. [arXiv:1805.00020](https://arxiv.org/abs/1805.00020) [hep-ph]
46. J. Brehmer, K. Cranmer, G. Louppe, J. Pavez, *Phys. Rev. Lett.* **121**, 111801 (2018). <https://doi.org/10.1103/PhysRevLett.121.111801>. [arXiv:1805.00013](https://arxiv.org/abs/1805.00013) [hep-ph]

47. R. Mastandrea, B. Nachman, T. Plehn, Phys. Rev. D **110**, 056004 (2024). <https://doi.org/10.1103/PhysRevD.110.056004>. [arXiv:2405.15847](https://arxiv.org/abs/2405.15847) [hep-ph]
48. M. Dax, J. Wildberger, S. Buchholz, S.R. Green, J.H. Macke, B. Schölkopf, (2023). [arXiv:2305.17161](https://arxiv.org/abs/2305.17161) [cs.LG]
49. T. Heimel, T. Plehn, N. Schmal, (2024). [arXiv:2411.00942](https://arxiv.org/abs/2411.00942) [hep-ph]
50. J.Y. Araz, A. Buckley, G. Kasieczka, J. Kieseler, S. Kraml, A. Kvellestad, A. Lessa, T. Procter, A. Raklev, H. Reyes-Gonzalez, K. Rolbiecki, S. Sekmen, G. Unel, SciPost Phys. Community Rep. (2024). <https://doi.org/10.21468/SciPostPhysCommRep.3>
51. S. Bailey, (2022). [arXiv:2203.10057](https://arxiv.org/abs/2203.10057) [hep-ph]
52. W. Abdallah et al. (LHC Reinterpretation Forum), SciPost Phys. **9**, 022 (2020). <https://doi.org/10.21468/SciPostPhys.9.2.022>. [arXiv:2003.07868](https://arxiv.org/abs/2003.07868) [hep-ph]
53. K. Cranmer et al., SciPost Phys. **12**, 037 (2022). <https://doi.org/10.21468/SciPostPhys.12.1.037>. [arXiv:2109.04981](https://arxiv.org/abs/2109.04981) [hep-ph]
54. J.Y. Araz, D. van Dyk, et al., Nabu. <https://github.com/nabu-hep/nabu-hep>
55. D. Ward, FlowJAX: distributions and normalizing flows in jax (2025). <https://doi.org/10.5281/zenodo.14524556>
56. P. Kidger, C. Garcia, Differentiable Programming workshop at Neural Information Processing Systems (2021), p. 2021
57. J. Bradbury, R. Frostig, P. Hawkins, M.J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, Q. Zhang, <http://github.com/jax-ml/jax>. JAX: composable transformations of Python+NumPy programs (2018)
58. E. Tabak, E. Vanden-Eijnden, Commun. Math. Sci. **8**, 217 (2010). <https://doi.org/10.4310/CMS.2010.v8.n1.a11>
59. E.G. Tabak, C.V. Turner, Commun. Pure Appl. Math. **66**, 145 (2013). <https://doi.org/10.1002/cpa.21423>
60. G. Papamakarios, T. Pavlakou, I. Murray, in *Advances in Neural Information Processing Systems*, vol 30, ed. by I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Curran Associates, Inc., 2017). https://proceedings.neurips.cc/paper_files/paper/2017/file/6c1da886822c67822bcf3679d04369fa-Paper.pdf
61. L. Dinh, J. Sohl-Dickstein, S. Bengio, in *International Conference on Learning Representations* (2017). <https://openreview.net/forum?id=HkpbnH9lx>
62. C. Durkan, A. Bekasov, I. Murray, G. Papamakarios, Neural spline flows (2019). [arXiv:1906.04032](https://arxiv.org/abs/1906.04032) [stat.ML]
63. D.P. Kingma, J. Ba (2014). [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) [cs.LG]
64. Kolmogorov–Smirnov test, in *The Concise Encyclopedia of Statistics* (Springer New York, New York, NY, 2008), pp. 283–287. https://doi.org/10.1007/978-0-387-32833-1_214
65. ATLAS Collaboration, ATLAS omnifold 24-dimensional z+jets open data (2024). <https://doi.org/10.5281/zenodo.11507450>
66. G. Aad et al. (ATLAS), Phys. Rev. Lett. **133**, 261803 (2024). <https://doi.org/10.1103/PhysRevLett.133.261803>. [arXiv:2405.20041](https://arxiv.org/abs/2405.20041) [hep-ex]
67. R. Aaij et al. (LHCb), Phys. Rev. D **102**, 112003 (2020). <https://doi.org/10.1103/PhysRevD.102.112003>. [arXiv:2009.00026](https://arxiv.org/abs/2009.00026) [hep-ex]
68. D. Leljak, B. Melić, F. Novak, M. Reboud, D. van Dyk, EOS/DATA-2023-01: supplementary material for EOS/ANALYSIS-2022-05 (2023). <https://doi.org/10.5281/zenodo.8027015>
69. D. Leljak, B. Melić, F. Novak, M. Reboud, D. van Dyk, JHEP **08**, 063 (2023). [https://doi.org/10.1007/JHEP08\(2023\)063](https://doi.org/10.1007/JHEP08(2023)063). [arXiv:2302.05268](https://arxiv.org/abs/2302.05268) [hep-ph]
70. D. van Dyk et al. (EOS Authors), Eur. Phys. J. C **82**, 569 (2022). <https://doi.org/10.1140/epjc/s10052-022-10177-4>. [arXiv:2111.15428](https://arxiv.org/abs/2111.15428) [hep-ph]
71. C. Bolognani, M. Reboud, D. van Dyk, K.K. Vos, JHEP **09**, 099 (2024). [https://doi.org/10.1007/JHEP09\(2024\)099](https://doi.org/10.1007/JHEP09(2024)099). [arXiv:2407.06145](https://arxiv.org/abs/2407.06145) [hep-ph]
72. S. Meiser, D. van Dyk, J. Virto, (2024). [arXiv:2411.09458](https://arxiv.org/abs/2411.09458) [hep-ph]