

Data-driven and model-agnostic approach to solving combinatorial assignment problems in searches for new physics


Anthony Badea^{1,2,*} and Javier Montejo Berlingen^{3,4,†}

¹*Harvard University, Cambridge, Massachusetts 02138, USA*

²*University of Chicago, Chicago, Illinois 60637, USA*

³*Instituto de Física de Altas Energías, Campus UAB, 08193 Bellaterra (Barcelona), Spain*

⁴*QUP, KEK, 1-1 Ōho, Tsukuba, Ibaraki 300-3256, Japan*

 (Received 13 September 2023; accepted 18 December 2023; published 9 January 2024)

We present a novel approach to solving combinatorial assignment problems in particle physics. The correct assignment of decay products to parent particles is achieved in a model-agnostic fashion by introducing a neural network architecture, *Passwd-ABC*, which combines a custom layer based on attention mechanisms and dual autoencoders. We demonstrate how the network, trained purely on background events in an unsupervised setting, is capable of reconstructing correctly hypothetical new particles regardless of their mass, decay multiplicity, and substructure, and produces simultaneously an anomaly score that can be used to efficiently suppress the background. This model allows the extension of the suite of searches for localized excesses to include nonresonant particle pair production where the reconstruction of the two resonant masses is thwarted by combinatorics.

DOI: [10.1103/PhysRevD.109.L011702](https://doi.org/10.1103/PhysRevD.109.L011702)

Searches for beyond the Standard Model (BSM) physics at the LHC are producing stringent bounds on the most theoretically appealing models [1–7]. This has led to a paradigm shift in the program toward more signature-based rather than model-based searches. The archetype of this approach is the search for an excess in the diobject mass spectrum, which has provided a fruitful history of discoveries, from the J/ψ meson [8,9] to the Higgs boson [10,11]. An important shortcoming of this approach is that it is restricted to models with resonant production of a single particle, from which all the relevant final-state objects are assumed to originate. However, many BSM models lead to nonresonant pair production which faces the additional challenge of assigning the decay products to parent particles in order to reconstruct masses. This challenge is further exacerbated in decays to high-multiplicity final states. Heuristic approaches have been applied in final states with low object multiplicities, such as dijet pairs or di- b -jet pairs [12–15], where all possible combinations can be tested. Some common choices are minimizing the sum of distances or the mass asymmetry. These approaches do not scale easily to higher object multiplicity as the amount

of combinations grows in a factorial way, leading to a worsening in the resolution of the reconstructed mass together with a rapid increase in complexity beyond CPU capacity. Though machine-learning approaches have been proposed to tackle this problem, they always start with a fixed signal model assumption [16–21]. This restricts the model to learn the specific multiplicity, structure and intermediate resonances of the decay. Furthermore, by imposing the signal hypothesis in the final state topology, the background spectrum is sculpted to become strongly signal-like.

In this paper, we present a methodology to reconstruct the four-momenta of pair-produced particles by assigning decay products to parent particles without knowledge of the decay mode. The only assumption embedded in the model is pair-production of new particles with identical decays. This is accomplished by training a custom neural network (NN) directly on data in an unsupervised approach, outputting an object-to-parent assignment and an anomaly score. Beyond this, the significant computational innovations are:

- (i) A custom permutation equivariant and easily scalable NN layer tailored to address combinatorial assignment problems based on the attention mechanism [22], dubbed *attention-based combinatorial layer*, or *ABC layer*. The ABC layer provides a per-object probability to be assigned to a certain category.
- (ii) A model based on ABC layers, named *particle assignment for unknown decays*, *Passwd-ABC*. It is

* anthony.badea@cern.ch

† jmontejo@cern.ch

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

trained in an unsupervised setting directly on data events, without any signal hypothesis or input from simulation, and is capable of assigning accurately the decay products of pair-produced particles. The performance of the model is shown to be insensitive to the BSM particle masses and decay structure, and capable of reconstructing high-multiplicity final states.

Formally, given an input set $\mathcal{X} = \{x_1, \dots, x_N\}$ and a category set $\mathcal{Y} = \{y_1, \dots, y_C\}$, the goal is to assign to each input $x \in \mathcal{X}$ a label $y \in \mathcal{Y}$, $l(x_i) = y_j$. This defines C exclusive subsets $P_k = \{x \in \mathcal{X} : l(x) = y_k, \forall x \in \mathcal{X}\}$. A combinatorial assignment is a function $f: \mathcal{X} \rightarrow A = \{0, 1\}^{N \times C}$, and for the correct assignment $a_{ij} = \mathbb{I}[l(x_i) = y_j]$, where $\mathbb{I}[\cdot]$ is an indicator function that is one when \cdot is true and zero otherwise. Our goal is then to build a model that provides a differentiable approximation to $f \sim f': \mathcal{X} \rightarrow A' = [0, 1]^{N \times C}$, which can be mapped back to the hard combinatorial assignment if desired via $a_{ij} = \text{argmax}_{j \in C} (a'_{ij})$. In our setup, each of the two parent particles is considered a category (P_1, P_2). An additional category can be introduced to account for objects not originating from the decay products, such as initial-state radiation (ISR), underlying event, or pileup. Setups with a higher number of categories, for example targeting four-top production, are also possible but not considered here for simplicity. The parent particles are reconstructed as $p_j = \sum_i^N a'_{ij} \cdot x_i$, which also enables a differentiable approximation to the assignment of objects to parent candidates by introducing a' instead. Other differentiable approximations for the object assignment such as Gumbel-Softmax [23,24] were tested and found to yield worse performance.

If a signal model is assumed, the assignment problem can be framed in a supervised learning approach. Objects can be assigned a label based on truth-level information and the loss function is simply the binary or multiclass cross-entropy loss summed over all objects. This strategy, adopted by previous works [16–21], requires training the model on signal simulation, inducing a strong model dependence and shaping of background distributions.

In order to adapt the problem to an unsupervised learning strategy, the problem is reframed as finding the combinatorial assignment such that the *similarity* between the two particle sets is maximized, $P_1 \approx P_2$. The choice of similarity metric is therefore the key ingredient to this setup. Different heuristic approaches have been used in the past in analyses where the number of combinations allows for a brute-force approach, such as the grouping with minimum mass asymmetry, $\mathcal{A} = \frac{|m_1 - m_2|}{m_1 + m_2}$. Such similarity metrics can also be introduced in our setup which would lead to a CPU-efficient way of approximating the result from iterating over all combinations, but the performance of the reconstructed variable would never surpass the heuristic approach. Instead, as detailed below, the chosen similarity

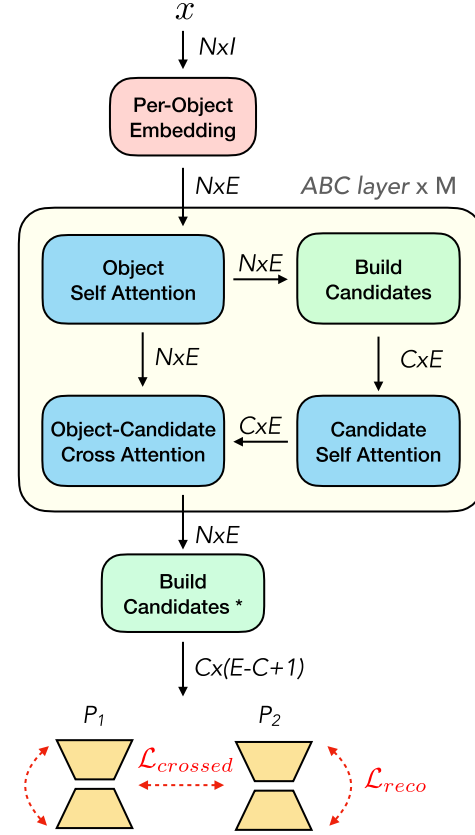


FIG. 1. Sketch of the passwd-ABC model and components of the ABC layer. Multiple ABC layers are stacked followed by a final candidate-building operation. The C first features encoding the category probability are swapped for the mass of the candidate. This final candidate building block is notated with an asterisk. Red lines indicate the terms that are used to build the reconstruction and similarity loss.

metric is taken to be the distance between both parent particles in a learned feature space, allowing the model to identify and build better features to quantify the similarity.

The ABC layer and passwd-ABC model composition are shown in Fig. 1. The input to the model, $x \in \mathbb{R}^{N \times I}$, consist of N objects with I input features. The input objects are fed through NNs to produce learned embeddings per object, $e = (\text{NN}(x_1), \dots, \text{NN}(x_N)) \in \mathbb{R}^{N \times E}$, with E the embedding dimension. At various points in the model, the C first embedded features per object will be softmax-ed to represent the probability for each object to be assigned to a given category: $a' = \sigma(e' \cdot e) \in \mathbb{R}^{N \times C}$, with σ the softmax function acting over the category dimension.

The ABC layer consists of three individual attention blocks. The first is an object self-attention block, after which category probabilities are computed and candidate parents are built in the embedded space by summing all the input objects weighted by their category probabilities, $p_j^{(e)} = \sum_i^N a'_{ij} \cdot e_i$. The parent candidates are input to the second self-attention block, followed by an object-candidate

cross-attention block. This structure allows the layer to learn relations among objects, candidates, and object-candidate pairs. The input and output of the ABC layer have the same (N, E) dimension, which allows multiple such layers to be stacked. The ABC layer is permutation equivariant. This avoids the need to arbitrarily sort the input objects and alleviates problems in supervised training applications due to varying class imbalances at each input position.

Depending on the use case, another candidate-building block can be appended as a last step, summing objects in the learned feature space or Lorentz space to obtain parent candidate four-momenta. This architecture allows great flexibility as the training target can be defined at the per-object level, e.g., classification of truth-labeled objects; at the particle candidate level in Lorentz space, e.g., regressing the BSM particle mass; or even mixtures of learned features and Lorentz features as described below.

The ABC layer is used as the foundation block for `Passwd-ABC`. The additional ingredient is an autoencoder block which is applied identically to each particle candidate. More advanced configurations of autoencoders can easily be adapted to the network structure but are left for future works. Particle candidates are built using the same approach as in the ABC layer. In addition, the reconstructed mass is added as an extra feature of the parent particle and the category scores are dropped. The autoencoder compresses the dimensionality of the embedded feature space to a latent space, $z = \{AE^{\text{enc}}(p_j) : j \in C\} \in \mathbb{R}^{C \times L}$, summarizing each particle candidate's properties, and then attempts to reconstruct the original inputs, $\hat{p} = \{AE^{\text{dec}}(z_j) : j \in C\}$. The distance in latent space between the two candidates, $\|z_1 - z_2\|$, is used as the metric for similarity that is required to solve the problem. Other approaches such as minimizing the energy mover's distance [25,26] between the two candidates were also examined but found to have worse performance.

The loss function used to train `Passwd-ABC` is

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_{\text{reco}} + \mathcal{L}_{\text{crossed}} + \mathcal{L}_{\text{random}} (+\mathcal{L}_{\text{ISR}}) \\ \mathcal{L}_{\text{reco}} &= \lambda_{\text{reco}} \cdot (\|p_1 - \hat{p}_1\| + \|p_2 - \hat{p}_2\|) \\ \mathcal{L}_{\text{crossed}} &= \lambda_{\text{crossed}} \cdot \|z_1 - z_2\| \\ \mathcal{L}_{\text{rand}} &= \lambda_{\text{rand}} \cdot \max(0, 1 + \|z_1 - z_2\| - \|z_1^{\text{rand}} - z_2^{\text{rand}}\|) \\ \mathcal{L}_{\text{ISR}} &= \lambda_{\text{ISR}} \cdot E_T(P_{\text{ISR}})/\text{GeV} \end{aligned}$$

Where $\mathcal{L}_{\text{reco}}$ is the typical autoencoder reconstruction loss and $\mathcal{L}_{\text{crossed}}$ is the distance between both reconstructed particles in the autoencoder latent space. Since all of the candidate features are learned, a simple minimization of these two terms will lead to a collapse of zero-valued candidate input features and latent space. This is avoided by adding an additional term, $\mathcal{L}_{\text{random}}$, based on the triplet loss [27]. Two additional parent particles are built using a random combination of objects, and the distance between

TABLE I. Benchmark production modes and decays that are considered within the context of R-parity-violation SUSY. Hadronic decays of bosons and top quarks are considered as they have the highest branching fractions. No jet-flavor information is used in the model, therefore b -quarks are not explicitly notated.

Decay		Number of jets	
\tilde{g}	$\rightarrow q\bar{q}$	$\rightarrow q(qq)$	6
\tilde{q}	$\rightarrow q\tilde{\chi}_1^0$	$\rightarrow q(qqq)$	8
\tilde{q}	$\rightarrow W/Z/H\tilde{q}'$	$\rightarrow qq(qqq)$	8
\tilde{g}	$\rightarrow qq\tilde{\chi}_1^0$	$\rightarrow qq(qqq)$	10
\tilde{t}	$\rightarrow t\tilde{\chi}_1^0$	$\rightarrow qq(qqq)$	12

them is maximized. This term is implemented in the form of a hinge loss to avoid divergences where the NN ignores the main task and focuses on maximizing the difference among random candidates. We observed that frequent trainings lead to dimensional collapse [28]. Several methods were tested to mitigate the issue [29]. None yielded significantly different results compared to the selected training which showed no collapse. An extra term, \mathcal{L}_{ISR} , is introduced when the additional ISR category is considered and is simply the transverse energy, E_T , of the ISR system in units of GeV. This term is introduced to avoid all objects being flagged as ISR which leads to perfectly identical, and empty, parent particles. All loss terms are multiplied with a corresponding hyperparameter to balance the relative weight.

In order to demonstrate the generality of the approach and to quantify the performance of the model we focus on the multi-jet final state with varying decay structures. We use a simplified R-parity-violating supersymmetry model [30,31] with a nonzero baryon-number-violating $\tilde{U}\tilde{D}\tilde{D}$ coupling [32]. We consider gluino, squark, and stop pair production (\tilde{g} , \tilde{q} , \tilde{t}) with cascade decays among themselves or to the lightest neutralino ($\tilde{\chi}_1^0$). The lightest supersymmetric particle then decays through the $\tilde{U}\tilde{D}\tilde{D}$ coupling to SM quarks, leading to final states with 6 to 12 jets and different internal substructures. An example signal is $\tilde{g}\tilde{g} \rightarrow 2 \times qq\tilde{\chi}_1^0 \rightarrow 2 \times qq(qqq)$. For the sake of generality, and to emphasize the model-agnostic focus of this letter, the following notation is used: $XX^{m_x} \rightarrow 2 \times Nj(Mj)$, to denote pair-production of particle X with mass m_x decaying to N jets plus an intermediate particle, which in turn decays to M jets, producing a final state with $2 \times (N + M)$ decay products. The signals, decays, and number of final state jets considered are given in Table I, ranging broadly over event topologies.

Simulated signal events are generated at leading-order using Madgraph [33,34] with up to one additional parton, and interfaced with Pythia 8 [35]. Background QCD events are simulated using Pythia 8. Both signal and background events are interfaced with Delphes event




FIG. 2. The average of the two reconstructed masses $m_{\text{avg}} = \frac{1}{2}(m_1 + m_2)$ (top) and reconstruction loss (bottom) of the Passwd-ABC model on a simulated sample of QCD background (black) and multiple signals with different masses and decay chains (colored). No signal sample has been used in the training of the model. The mass of the intermediate particle is chosen to be half the mass of the parent particle.

reconstruction [33,35], and required to have at least 6 jets with $p_T > 25$ GeV and $H_T > 1000$ GeV (scalar sum of jet p_T) to mimic hadronic trigger selections at the LHC experiments. A sample of about 200k background events is retained after selection, where 90% is used for training and the other 10% is reserved for evaluation. A sample of 10k events is generated for each signal and none of the signal events are used in the training.

A minimal Passwd-ABC model is implemented using a stack of two ABC layers, the first one with two categories and the second with three. This choice allows for a first splitting into two candidate particles which is then refined in the second layer by trimming some of the jets and assigning them to the ISR category. All activation functions are rectified linear units (ReLU). Embedding is a




FIG. 3. Comparison of the average of the two reconstructed masses $m_{\text{avg}} = \frac{1}{2}(m_1 + m_2)$ for the background (black) and an example signal (green) as reconstructed by passwd-ABC in its last layer (solid line), intermediate layer (dotted line), and the mass reconstructed from the minimization of the mass asymmetry (dashed line). An improvement of the signal resolution is observed after the first layer prediction is refined. The background distribution is strongly sculpted toward high masses when reconstructed by minimizing the mass asymmetry.

three-layer feed-forward network (FFN) with a 16-feature embedding space. All attention blocks are 4-headed attention layers, followed by a three-layer FFN, using the same 16-feature dimensionality. The autoencoder uses four layers for each of the encoder and the decoder, mapping into a 2-dimensional latent space. The model has a total of 16.7K trainable parameters. All hyperparameters in the loss function are set to 1 except for $\lambda_{\text{reco}} = 10$. Only the jet four-momenta are used as input features, parameterized as $x_i = (\log p_T, \eta, \sin \phi, \cos \phi, \log E)$. Other features such as particle type or identification scores can be added but are not considered here. Up to 12 jets are used as input and events with fewer than 12 are zero-padded. The model is trained for 10 epochs using the ADAM optimizer [36]. The batch size is 1024. The learning rate is warmed up for the first 2% of training steps to a maximal value of 10^{-3} and then decayed by 95% every further 2%.

The performance of the model is shown in Fig. 2 for a variety of signals with different masses and decay structures. The model is able to correctly reconstruct all signal distributions regardless of the mass and decay, despite having been trained purely on background. The selection of $H_T > 1000$ GeV drives the background distribution to peak at around $H_T/2 \approx 500$ GeV. The reconstruction loss of the autoencoder block provides an excellent discriminating feature to suppress the background, achieving background rejection factors of $\mathcal{O}(200)$ with 50% signal efficiency for the lowest mass points. The loss is found to be strongly correlated with the reconstructed mass but still

retains additional rejection power when analyzed in bins of reconstructed mass. Multiple possibilities exist to reduce or remove the correlation [37–42] that are left for future work.

The performance of the model is further inspected using the jet assignment right after the first ABC layer to build the parent candidates. As shown in Fig. 3, the resolution of the reconstructed mass improves from the first to the second layer as the jet classification is refined. The performance is further compared with a heuristic approach such as iterating over all possible jet combinations and selecting the one that minimizes the mass asymmetry. The reconstructed signal mass exhibits worse resolution, a positive bias, and a large high-mass tail as measured by the peak (m_{peak}) and full width at half max (Δm) values. The background spectrum is pulled toward high masses, leading to a factor 30 worse signal-to-background ratio under the signal peak, and a factor 200 worse when integrating to the highest mass.

In conclusion, this paper presents an innovative method to expand the BSM search program at the LHC. We develop the model-agnostic `Passwd-ABC`, based on ABC layers, which is suited for unsupervised learning directly on data and addresses critical limitations of current approaches. Our method not only obviates the need for

predefined signal models but also demonstrates remarkable resilience in decoding complex, high-multiplicity final states. The ability to reconstruct pair-produced particles without knowledge of the decay mode is a significant leap in the pursuit of model-independent searches. The presented model, while falling in the category of anomaly-based searches, provides handles for the inspection of a possible signal through the explicit reconstruction from its decay products. This would allow the measurement of its mass, decay structure, and other properties. We encourage the experimental collaborations to leverage this method to expand their suite of model-independent searches into final states with high object multiplicity.

We thank Lawrence Lee, Simone Francescato, and Aurelio Juste Rozas for useful discussions and suggestions. This work has been supported by the Department of Energy, Office of Science, under Grant No. DE-SC0007881 (A. B.), the Harvard Graduate Prize Fellowship (A. B.), the Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship (A. B.), Grant RYC2021-030944-I funded by MCIN/AEI/10.13039/501100011033 (J. M. B.) and by the European Union NextGenerationEU/PRTR (J. M. B.).

-
- [1] ATLAS Collaboration, *Phys. Lett. B* **842**, 137963 (2023).
 - [2] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **98**, 052008 (2018).
 - [3] M. Aaboud *et al.* (ATLAS Collaboration), *J. High Energy Phys.* **05** (2019) 142.
 - [4] A. M. Sirunyan *et al.* (CMS Collaboration), *Phys. Lett. B* **798**, 134952 (2019).
 - [5] A. Tumasyan *et al.* (CMS Collaboration), *Eur. Phys. J. C* **81**, 970 (2021).
 - [6] ATLAS, Supersymmetry summary plots, Report No. ATLAS-PHYS-PUB-2023-025, <https://cds.cern.ch/record/2852738>.
 - [7] CMS, CMS Exotica summary plots, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/SummaryPlotsEXO13TeV>.
 - [8] J. J. Aubert *et al.* (E598 Collaboration), *Phys. Rev. Lett.* **33**, 1404 (1974).
 - [9] J. E. Augustin *et al.* (SLAC-SP-017 Collaboration), *Phys. Rev. Lett.* **33**, 1406 (1974).
 - [10] G. Aad *et al.* (ATLAS Collaboration), *Phys. Lett. B* **716**, 1 (2012).
 - [11] S. Chatrchyan *et al.* (CMS Collaboration), *Phys. Lett. B* **716**, 30 (2012).
 - [12] M. Aaboud *et al.* (ATLAS Collaboration), *Eur. Phys. J. C* **78**, 250 (2018).
 - [13] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **98**, 092002 (2018).
 - [14] ATLAS Collaboration, *Phys. Rev. D* **108**, 052003 (2023).
 - [15] G. Aad *et al.* (ATLAS Collaboration), *Phys. Rev. D* **108**, 112005 (2023).
 - [16] M. J. Fenton, A. Shmakov, T.-W. Ho, S.-C. Hsu, D. Whiteson, and P. Baldi, *Phys. Rev. D* **105**, 112008 (2022).
 - [17] S. Qiu, S. Han, X. Ju, B. Nachman, and H. Wang, *Phys. Rev. D* **107**, 114029 (2023).
 - [18] A. Badea, W. J. Fawcett, J. Huth, T. J. Khoo, R. Poggi, and L. Lee, *Phys. Rev. D* **106**, 016001 (2022).
 - [19] L. Ehrke, J. A. Raine, K. Zoch, M. Guth, and T. Golling, *Phys. Rev. D* **107**, 116019 (2023).
 - [20] A. Shmakov, M. J. Fenton, T.-W. Ho, S.-C. Hsu, D. Whiteson, and P. Baldi, *SciPost Phys.* **12**, 178 (2022).
 - [21] J. S. H. Lee, I. Park, I. J. Watson, and S. Yang, *arXiv:2012.03542*.
 - [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, *arXiv:1706.03762*.
 - [23] C. J. Maddison, A. Mnih, and Y. W. Teh, *arXiv:1611.00712*.
 - [24] E. Jang, S. Gu, and B. Poole, *arXiv:1611.01144*.
 - [25] P. T. Komiske, E. M. Metodiev, and J. Thaler, *Phys. Rev. Lett.* **123**, 041801 (2019).
 - [26] P. T. Komiske, E. M. Metodiev, and J. Thaler, *J. High Energy Phys.* **07** (2020) 006.
 - [27] F. Schroff, D. Kalenichenko, and J. Philbin, in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, Boston, MA, USA, 2015), <https://ieeexplore.ieee.org/document/7298682>.
 - [28] T. Hua, W. Wang, Z. Xue, S. Ren, Y. Wang, and H. Zhao, *arXiv:2105.00470*.
 - [29] L. Jing, P. Vincent, Y. LeCun, and Y. Tian, *arXiv:2110.09348*.
 - [30] R. Barbier *et al.*, *Phys. Rep.* **420**, 1 (2005).

- [31] R. N. Mohapatra, *Phys. Scr.* **90**, 088004 (2015).
- [32] C. Csáki, Y. Grossman, and B. Heidenreich, *Phys. Rev. D* **85**, 095009 (2012).
- [33] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *J. High Energy Phys.* **07** (2014) 079.
- [34] B. Fuks, *Int. J. Mod. Phys. A* **27**, 1230007 (2012).
- [35] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, *Comput. Phys. Commun.* **191**, 159 (2015).
- [36] D. P. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [37] G. Louppe, M. Kagan, and K. Cranmer, [arXiv:1611.01046](https://arxiv.org/abs/1611.01046).
- [38] C. Shimmin, P. Sadowski, P. Baldi, E. Weik, D. Whiteson, E. Goul, and A. Sjøgaard, *Phys. Rev. D* **96**, 074034 (2017).
- [39] O. Kitouni, B. Nachman, C. Weisser, and M. Williams, *J. High Energy Phys.* **04** (2020) 070.
- [40] G. Kasieczka and D. Shih, *Phys. Rev. Lett.* **125**, 122001 (2020).
- [41] S. Klein and T. Golling, [arXiv:2211.02486](https://arxiv.org/abs/2211.02486).
- [42] M. Algren, J. A. Raine, and T. Golling, [arXiv:2307.05187](https://arxiv.org/abs/2307.05187).