

Single interpolated generative model for anomalies

Ranit Das^{*} and David Shih[†]

NHETC, Department of Physics and Astronomy, Rutgers University, Piscataway, New Jersey 08854, USA



(Received 8 April 2025; accepted 25 September 2025; published 29 October 2025)

A key step in any resonant anomaly detection search is accurate modeling of the background distribution in each signal region. Data-driven methods like CATHODE accomplish this by training separate generative models on the complement of each signal region, and interpolating them into their corresponding signal regions. Having to retrain the generative model on essentially the entire dataset for each signal region is a major computational cost in a typical sliding window search with many signal regions. Here, we present SIGMA, a new, fully data-driven, computationally efficient method for estimating background distributions. The idea is to train a single generative model on all of the data and interpolate its parameters in sideband regions in order to obtain a model for the background in the signal region. The SIGMA method significantly reduces the computational cost compared to previous approaches, while retaining a similar high quality of background modeling and sensitivity to anomalous signals.

DOI: [10.1103/rj53-2x6j](https://doi.org/10.1103/rj53-2x6j)

I. INTRODUCTION

There has been a rapid growth in new methods for resonant anomaly detection in recent years [1–31]. These methods are starting to move from proof-of-concept to actual analyses applied to ATLAS [32] and CMS data [33] in the model-agnostic search for new physics. They have also been successfully applied to astronomical data from the Gaia Space Telescope in the search for stellar streams [34–37].

In resonant anomaly detection, one uses additional features x to enhance the sensitivity of a “bump hunt” in a primary feature m where the signal is assumed to be localized. The overall goal is to accurately estimate the Neyman-Pearson classifier

$$R_{\text{optimal}}(x) = \frac{p_{\text{data}}(x)}{p_{\text{background}}(x)} \quad (1)$$

in a signal region (SR), which is defined as a window in m . By cutting on this “optimal anomaly score,” one can preferentially select signal over background events in a model-agnostic way, i.e., without having to specify the signal model in advance.

In many approaches, a key step is to learn a fully data-driven estimate for $p_{\text{background}}(x)$ by interpolating from neighboring control regions (CR) in m . In methods like ANODE [3] or CATHODE [7], one trains a conditional density estimator or conditional generative model, respectively, on the complement of the signal region, which for a narrow signal region is almost the entire dataset. For each signal region this must be performed separately, and this leads to extremely expensive trainings for a full sliding-window search.

There have been efforts to mitigate this computational cost. In CURTAINS4F [26], one can train a single flow on the entire dataset, not masking out the SR, and another conditional flow on short control regions to learn the transformation between them. This can then be used to interpolate into the signal region and saves significantly on computational cost since the all-data flow need only be learned once, and the short control region flow is usually learned on significantly less data.

Recently in RAD-OT [30], the idea was to dispense with the flow altogether and just use OT to transport sidebands to sidebands, thereby interpolating into the SR. This was shown to be computationally very inexpensive but at the cost of sacrificing some quality in the interpolation.

In this paper, we present a single interpolated generative model for anomalies (SIGMA), a new approach for computationally efficient resonant anomaly detection, which is much simpler than CURTAINS4F and maintains the high quality of the more expensive methods like CATHODE. Our idea is to train a single flow on all the data, as in CURTAINS4F, and simply interpolate the parameters of the flow in the resonant feature. We find the simplest implementation of this is using flow-matching, where the

^{*}Contact author: ranit@physics.rutgers.edu

[†]Contact author: shih@physics.rutgers.edu

vector field itself can be easily interpolated into the signal region. We will demonstrate that SIGMA produces high quality background samples while saving significantly on training cost—even more than CURTAINS4F since even the short sideband flows can be dispensed with.

This paper is organized as follows: in Sec. II A we describe the dataset used to demonstrate our method. The architecture of the generative model used for our background template is described in II B 2. In II C 1 we describe the slower interpolation method from [3,7] using flow-matching. The new interpolation method SIGMA is described in Sec. II C 2. The model hyperparameters and the training are described in Sec. II D. In Sec. III, we show the results for anomaly detection, and the background template obtained from different methods. Finally in Sec. IV we present our conclusions and further discussion. In Appendix A, we discuss the importance of using frequency embedding for SIGMA. In Appendix B we analyze the background templates using the classifier test, developed in [38]. In Appendix C, we discuss the background template quality for SIGMA for different signal regions.

II. SETUP

A. Dataset

We use the LHC R&D Dataset [39,40] for this study with dataset and train-val-test splits similar to [3,7]. The Standard Model (SM) background consists of QCD dijet events, and $W' \rightarrow X(\rightarrow qq)Y(\rightarrow qq)$ events with $m_{W'} = 3.5$ TeV, $m_X = 500$ GeV and $m_Y = 100$ GeV form the signal. These are simulated using Pythia 8 [41,42] and Delphes 3.4.1 [43–45]. The reconstructed particles are clustered into jets using the anti- k_T algorithm [46,47] with $R = 1$ using Fastjet [48]. Events that satisfy the $p_T > 1.2$ TeV jet trigger are kept.

The training features we use are

$$m = m_{JJ}, \quad x = [m_{J_1}, \Delta m_J, \tau_{21}^{J_1}, \tau_{21}^{J_2}, \Delta R], \quad (2)$$

where invariant masses of the subjects satisfy $m_{J_1} < m_{J_2}$, and $\Delta m_J = m_{J_2} - m_{J_1}$. We chose dijet invariant mass m_{JJ} as the resonant variable m , with the signal region (SR) defined as $m \in [3.3, 3.7]$ and $m \notin [3.3, 3.7]$ form the control regions (CR). The n-subjettiness ratios are $\tau_{ij} = \tau_i / \tau_j$ [49,50]. ΔR is the angular separation of the leading two jets in $\eta - \phi$ space. Including it does not enhance the performance for anomaly detection, compared to just baseline features. But it is a feature that is strongly correlated with the resonant feature, m_{JJ} , and hence including it provides a stringent test for any interpolation method [51].

We use all 1-million SM background events from the original R&D dataset and inject different amounts of signal (3000 or lower) from the first 70k signal events of the R&D

dataset. This is meant to represent actual un-labeled data from the experiment. For reference, the signal-injection of $N_{\text{sig}} = 1000$ corresponds to $\approx 2.2\sigma$ significance for the inclusive bump hunt in m_{JJ} .

We split this dataset into training and validation sets, with 85-15 split for training all generative models. For IAD, following [7], we classify between data in SR and an additional 272k QCD dijet events in the SR [52] (which represents true background). For CATHODE, we classify between data in SR, and 360k events sampled from the 10 best instances of the trained generative model. We use the remaining 30k signal events from the original R&D dataset, along with the additional 340k QCD dijet events in the SR from [52] for evaluating SIC curves.

For each N_{sig} , we produce ten different datasets by injecting different randomly selected signal events from the R&D dataset. This is used to obtain the error bars on the performance curves in the figures.

B. Flow-matching and network architecture

We use conditional flow matching (CFM) [53–57] to model the background template as a proof of concept, but our idea could be extended to be used with diffusion models [58–64] as well.¹ HEP has seen a widespread use of both CFM [67–77] and diffusion models [76–87] as generative models. Besides providing a better quality background template than ordinary normalizing flows, we found that flow matching also trains faster—in approximately 30 mins, as compared to a normalizing flow which trains in ~ 2.8 h. In this section, we give a brief description of conditional flow-matching and then describe the neural network architecture used to model the flow.

1. Conditional flow matching

Conditional flow matching is an efficient way to train continuous normalizing flows (CNFs) [88]. In CNFs, one aims to learn the vector field $u_t(x_t): [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ which implements a continuous transformation of data x_t via

$$\frac{dx_t}{dt} = u_t(x_t), \quad (3)$$

At $t = 0$, x_0 follows the data distribution $p_{\text{data}}(x_0)$, and at $t = 1$, x_1 follows a known distribution $p_{\text{base}}(x_1)$. Here we take $p_{\text{base}}(x)$ to be the unit normal distribution $\mathcal{N}(x|0, 1)^d$. In between, x_t is sampled from a density $p_t(x_t)$ which is generated by the vector field; the trajectory of these densities is referred to as the probability density path.

Training CNFs with maximum likelihood is extremely computationally expensive, because each evaluation of the

¹Diffusion models have lead to state-of-the-art image generators [65,66], hence we speculate that their use could lead to superior template quality.




FIG. 1. ResNet architecture.

likelihood requires solving an ODE. The breakthrough insight of conditional flow matching is that one can instead regress the desired vector field $u_t(x_t)$ with an MSE-type loss, by matching it to a *conditional* vector field $u_t(x_t|x_0)$. This generates a conditional probability path $p_t(x|x_0)$ where at $t = 0$, we have $p_0(x|x_0) = \mathcal{N}(x|x_0, \sigma^2 I)$ where σ^2 is very small, and at $t = 1$, we have $p_1(x|x_0) = \mathcal{N}(x|0, 1)^d$. Marginalizing this path over the data distribution $p_{\text{data}}(x_0)$ gives us the unconditional probability $p_t(x)$:

$$p_t(x) = \int dx_0 p_t(x|x_0) p_{\text{data}}(x_0). \quad (4)$$

In CFM, this conditional vector field is regressed with a neural network $v_\theta(x_t|t)$ by minimizing the CFM loss

$$\mathcal{L}(\theta) = \|v_\theta(x_t|t) - u_t(x_t|x_0)\|^2, \quad (5)$$

In [53], it was shown that the CFM objective, when averaged over $t \sim \mathcal{U}[0, 1]$, $x_0 \sim p_{\text{data}}(x_0)$ and $x_t \sim p_t(x|x_0)$, is minimized by $v_\theta(x|t) = u_t(x)$. In addition to t , our models also have an additional conditional feature: m , the resonant feature, which allows us to model the vector field $u_t(x|m)$ corresponding to $p(x|m)$. After integrating Eq. (3) to obtain x_t , $p(x_0|m)$ can be obtained using

$$\log p(x_0|m) = \log(x_1) - \int_1^0 dt \nabla_{x_t} \cdot u_t(x_t|m). \quad (6)$$

2. Network architecture

Similar to [24], we use a ResNet-style [89] architecture to model the vector field $u_t(x|m)$, since it is more expressive than an MLP. Rather than pass m and t directly to the ResNet, it is beneficial to first map them to an embedding space. For t , we use a frequency embedding $\beta: [0, 1] \rightarrow \mathbb{R}^{2(L'+1)}$, similar to those used in diffusion models [58]

$$\begin{aligned} \beta(t) = & (\sin(\pi t), \cos(\pi t), \dots, \\ & \sin((L'+1)\pi t), \cos((L'+1)\pi t)). \end{aligned} \quad (7)$$

The resonant feature m is first rescaled to $[-1, 1]$, and then mapped to an embedding space $\alpha: [-1, 1] \rightarrow \mathbb{R}^{2L}$. Depending on the interpolation method, we find different embeddings are preferable, see Sec. II C for more details. The t and m embeddings are concatenated and passed to an MLP to produce the full context vector

$$\gamma_c(m, t) = \phi_c(\alpha(m), \beta(t)), \quad (8)$$

where $\phi_c: \mathbb{R}^{2(L+L')} \rightarrow \mathbb{R}^n$ is an MLP.

The features $x \in \mathbb{R}^d$ are also concatenated with the frequency embedding for time and conditional features to form the input vector $\gamma_i(x, m, t)$

$$\gamma_i(x, m, t) = \phi_i(x, \alpha(m), \beta(t)), \quad (9)$$

where $\phi_i: \mathbb{R}^{2(L+L')+d} \rightarrow \mathbb{R}^n$ is another MLP.

Finally, the context and input vectors are passed through K ResNet blocks as shown in Fig. 1. The k th-ResNet Block $h^k: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ takes as input h^{k-1} and the context vector γ_c , and outputs

$$h^k = h^{k-1} + \phi^k(h^{k-1}) * \gamma_c, \quad (10)$$

where $\phi^k: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the MLP corresponding to the ResNet block. For $k = 0$, $h^0 = \gamma_i$.

We use the ResNet model defined in the NFLOWS package [90]. The output of the final ResNet block h^K is passed through a linear layer (to change dimensions from the n -dimensional input vector space to the d -dimensional feature space) in order to obtain the final vector field $v_\theta(x|t, m)$.

C. Interpolation

In this section we discuss the two different types of interpolation: First we discuss a flow-matching version of background interpolation used in [3, 7, 22, 91]. Then we discuss our main result, the SIGMA method, where we first train a model, v_θ^{data} to learn all data, and then obtain v_θ^{int} for each SR through a fast interpolation procedure.

1. Interpolation using $v_\theta^{\text{CR}}(x|t, m)$

Similar to [3, 7], we use $\alpha(m) = m$, for the embedding for the resonant feature. For every SR, a model is trained to learn $v_\theta^{\text{CR}}(x|t, m \in \text{CR})$ for its corresponding CR data. Since one needs to train a separate $v_\theta^{\text{CR}}(x|t, m)$ for each CR, this method can be computationally expensive. Once trained, the interpolated vector field into SR is simply $v_\theta^{\text{CR}}(x|t, m \in \text{SR})$. Since the CR-only training did not ever see the signal in the SR, the interpolation into the SR cannot have learned the signal. Therefore the CR-only interpolation method should give a model for the background-only distribution.

2. SIGMA: Interpolation using $v_\theta^{\text{data}}(x|t, m)$

In SIGMA, the idea is to first learn the full data distribution (including the signal region) as well as possible,² and then to interpolate the *parameters* of the vector field from m outside the SR into m into the SR, so that the interpolated background model does not include signal. Since the interpolation step does not require any additional training time, the background template for different SRs can be obtained trivially from the single, common data training.

²One could attempt to deliberately make v_θ^{data} less expressive, in order to not learn the signal in the SR, and protect the interpolation against signal contamination. But this would be difficult to control and likely would result in a worse fit to the background overall.

To learn the full data distribution optimally, including the more localized, higher frequency modes corresponding to signal, we found that a frequency embedding for m was beneficial [92, 93]. The resonant feature m is first rescaled to $(-1, 1)$, and then $\alpha(m)$ is defined as

$$\alpha(m) = (\sin(2^0 \pi m), \cos(2^0 \pi m), \dots, \sin(2^{L-1} \pi m), \cos(2^{L-1} \pi m)). \quad (11)$$

There are in principle many ways to interpolate the vector field from the CR into the SR. We will compare two in this work and see which one works better empirically. In real-world applications where there is no access to truth labels (background vs signal), one could still determine the better interpolation method using simulations, or using signal injection tests on top of real data, similar to how ATLAS and CMS currently set limits on searches using weak supervision [32, 33].

The two interpolation methods will explore are

- (1) Linearly interpolating the vector field itself

$$v_\theta^{\text{int}}(x|t, m) = \xi * v_\theta^{\text{data}}(x|t, m_1) + (1 - \xi) * v_\theta^{\text{data}}(x|t, m_2) \quad (12)$$

where $\xi = (m - m_2)/(m_1 - m_2)$, with $m \in \text{SR}$, $m_1, m_2 \in \text{CR}$, and $m_2 < m < m_1$.

- (2) Linearly interpolating the context and input vectors

$$\begin{aligned} \gamma_c^{\text{int}}(m, t) &= \xi * \gamma_c(m_1, t) + (1 - \xi) * \gamma_c(m_2, t) \\ \gamma_i^{\text{int}}(x, m, t) &= \xi * \gamma_i(x, m_1, t) \\ &\quad + (1 - \xi) * \gamma_i(x, m_2, t), \end{aligned} \quad (13)$$

And then the interpolated vector field v_θ^{int} is just given by v_θ^{data} evaluated at $\gamma_c = \gamma_c^{\text{int}}$ and $\gamma_i = \gamma_i^{\text{int}}$.

D. Model hyperparameters and training

For training $v_\theta^{\text{data}}(x|t, m)$, we use $L = 9$, $L' = 3$. The MLPs ϕ^c and ϕ^i in (8) and (9) consist of a single layer of $\text{dim} = 256$, with ReLU activation. We use 4 ResNet blocks, with MLPs ϕ^k consisting of 2 layers of $\text{dim} = 256$, with $\text{dropout} = 0.2$ and $\text{BatchNormalization}$ in between the layers. We train the models using Pytorch [94] to a maximum of 2000 epochs, with a batch-size of 4096, using the AdamW optimizer [95] with a learning rate of 3×10^{-4} , and ReduceLROnPlateau learning rate scheduler. After the first 100 epochs, we evaluate the negative log-likelihood [using Eq. (6)] for the validation data every 5 epochs. The negative log-likelihood is used for selecting the best 10 models. It is also used as a metric for the learning rate scheduler, which reduces the learning rate




FIG. 2. For $N_{\text{sig}} = 3000$ and $m \in \text{SR}$, we show, Top panel: the histograms of the samples of the features. The shaded blue histogram is the background present in the data. The shaded yellow is the signal added to the background. Bottom panel: the pulls with respect to the background distribution.




FIG. 3. Left: SIC curves for a signal injection of 1000 (which corresponds to $S/\sqrt{B} \approx 2.2$). Right: SIC at $\epsilon_B = 0.001$ for different signal injections. While v_θ^{CR} retains the best signal sensitivity, the much more computationally efficient v_θ^{int} (context) is only a little bit worse by comparison. Meanwhile, v_θ^{int} (linear) shows considerably worse performance at lower signal strengths.

by a factor of 0.3 if the metric does not decrease for 50 epochs. We use early stopping to stop the training if the validation negative log-likelihood does not improve for 100 epochs. We use the same set of hyperparameters for the CR model as well.

To generate samples from the vector field, we first sample noise from the base distribution, which is used as the initial condition to integrate the vector field from $t = 1$ to $t = 0$, using `odeint` from Zuko [96]. This function uses the adaptive Dormand-Prince method [88,97] to perform the integration. For v_θ^{int} , we choose $m_1 = 3.2$ TeV and $m_2 = 3.8$ TeV.

Finally, for the classifier, we use boosted decision trees (specifically a `HistGradientBoostingClassifier` from `scikit-learn`), which has been shown in [19] to be robust under noisy features. We train this BDT with default hyperparameters for 200 epochs.

III. RESULTS

A. Samples

In Fig. 2, we compare the samples in the SR from different models to background and data using histograms³ and pull distributions (with respect to the background),⁴ for $N_{\text{sig}} = 3000$. We find that pulls for the samples from v_θ^{data} show the presence of signal, similar to the pulls for the data distribution. For the samples from the other three models (v_θ^{int} (linear), v_θ^{int} (context) and v_θ^{CR}), there are clear differences in the pull distributions. All the interpolation

methods can smoothen out the signal, but in general v_θ^{CR} seems to produce the best background template, followed by v_θ^{int} (context), and finally v_θ^{int} (linear) being the worst in quality. This is especially clear in the ΔR distribution, a feature strongly correlated with m . Here we see the best pulls from v_θ^{CR} , significantly exaggerated pulls from v_θ^{int} (linear), and moderate pull values from v_θ^{int} (context). This shows that v_θ^{int} (context) better models the correlation in m as compared to v_θ^{int} (linear). In Appendix B, we further analyze the samples using the classifier test. In Appendix C, we show the results for the sample quality for different SRs.

We emphasize that these differences in background modeling quality stem primarily from different interpolation methods, and not density estimation quality. Both new interpolation models, v_θ^{int} (linear) and v_θ^{int} (context), are the same model v_θ^{data} in the SB. They differ only in how they are interpolated into the SR, as defined in Eq. (12) and (13). Meanwhile, v_θ^{CR} differs from the rest of the models in that it was trained only on CR data (the SR was masked out), and through the absence of the frequency embedding. Appendix A shows that the frequency embedding does not effect density estimation in the absence of signal. Hence the difference in the pulls is again primarily due to the interpolation method rather than the density estimation.

B. Anomaly detection performance

The main performance metric we use for anomaly detection is the significance improvement characteristic (SIC). In terms of the signal efficiency (ϵ_S or TPR) and background efficiency (ϵ_B or FPR), one has

$$\text{SIC} = \frac{\epsilon_S}{\sqrt{\epsilon_B}}. \quad (14)$$

³Samples from v_θ^{data} have the same normalization as the data distribution. Whereas the samples from v_θ^{int} (linear), v_θ^{int} (context) and v_θ^{CR} have the same normalization as the background distribution.

⁴We assume Poisson uncertainty for the pulls.

TABLE I. Timing comparison of different methods.

Method	Model	Timing
CATHODE/ANODE	Normalizing flows	3 hours per SR
CATHODE/ANODE	Flow matching	30 mins per SR
CURTAINS4F4	Normalizing flows	3 hours (base model) +7 mins per SR
RAD-OT	Optimal transport	10 mins per SR
SIGMA	Flow matching	30 mins (training) +30 secs per SR

The SIC characterizes the improvement to the nominal significance, using a cut on the anomaly score from a given method. SIC is dependent on the working point, and we present this dependence by plotting SIC vs the background rejection ($1/\epsilon_B$). At lower signal efficiencies, there are large statistical uncertainties in the SIC values, so we introduce a cutoff in the SIC curves when the relative statistical error on the background efficiency exceeds 20%. We also compare SIC at FPR = 0.001 to represent a typical, fixed working point that an experiment might choose in practice.

In Fig. 3 (left), we compare the SIC curves for $N_{\text{sig}} = 1000$, showing error bars for 10 different signal injections. The SIC curves for samples from $v_{\theta}^{\text{data}}(x|t, m)$ have large error bars. This is a result of the model being able to learn the full data distribution better for some signal injections (resulting in worse SIC), and so not so well in other cases (resulting in better SIC).⁵

The weakly supervised classifier identifies both the signal and the mismodeling artifacts, evident in the pull distributions in Fig. 2. At higher signal strengths the classifier is able to better identify the signal along with the mismodeling effects. At lower signal strengths, the mismodeling effects overcome the signal, which worsens the signal sensitivity. Hence v_{θ}^{CR} , with the best pulls in Fig. 2, results in the best SIC at all signal strengths in Fig. 3 (right). Both v_{θ}^{int} methods are able to interpolate out the signal, leading to an improvement over v_{θ}^{data} for $N_{\text{sig}} > 1500$. At lower N_{sig} values, the performance of v_{θ}^{int} (context) hews fairly closely to the benchmark set by v_{θ}^{CR} . This is significantly better than v_{θ}^{int} (linear), whose performance suffers much more, due to the worse pulls shown in Fig. 2. As mentioned in Sec. III A, this difference in performance is due to different background interpolation in the SR.

C. Timing comparison

In Table I, we compare the total time taken for generating background templates for different methods.

⁵One could possibly improve this error bar by ensembling multiple generative models on the same dataset, which we leave for future studies.

The background template generation method in [3,7,22,91] uses normalizing flows, which takes 3 hours to train. The flow-matching version introduced in this paper, v_{θ}^{CR} , offers a significant speedup to around 30 mins per SR. The training time for the base model of CURTAINS4F4, which also uses normalizing flows, is around 3 hours, and the lighter flow for each SR trains in 7 mins [26]. Using optimal transport, RAD-OT offered a faster alternative to CURTAINS4F4 at 10 mins per SR. Finally, our method SIGMA, trains the initial flow-matching model for 30 mins and subsequent sampling takes 30 seconds per SR. Hence SIGMA offers significant gains in computational efficiency compared to other methods.

IV. CONCLUSIONS

In this work we developed SIGMA, a new and efficient interpolation method for generative modeling of backgrounds in resonant anomaly detection. Rather than needing to train separate generative models in the complement of each signal region, SIGMA re-uses a single generative model trained on all of the data, with a subsequent interpolation of its parameters from control regions into the signal region of interest. This reduces the computational cost of SIGMA significantly relative to previous approaches such as ANODE and CATHODE [3,7], while preserving the high quality of their background templates and signal sensitivity.

The question of choosing the best interpolated background template in a data-driven way is an open one. We showed that the SIC, described in Sec. III, is very sensitive to bad background templates. For the LHCO R&D dataset, the v_{θ}^{int} (context) method showed the best performance, only a little worse than v_{θ}^{CR} . However, with real data we obviously do not have access to signal and background labels to calculate the SIC with. Nevertheless, it should be possible to perform signal injection tests, as ATLAS and CMS currently do to set limits on signal models [32,33], with artificial Gaussian or other signals added to real data. One could then see first-hand whether the “all data” model learns the data plus artificial signal, and whether the interpolated versions correctly smooth out the artificial signal.

There are many interesting future directions suggested by our work. Given that there was still a small performance gap between the previous, expensive interpolation method (masking out the CR) and SIGMA, it is worthwhile to explore further improvements to our method, such as using diffusion models instead of flow-matching or performing some kind of nonlinear interpolation using more than two mass points in the control region.

It would also be interesting to see how SIGMA applies to noisy features introduced in [19]. Finally, it should be totally straightforward to apply SIGMA to higher-dimensional SRs (multiple resonant features), which previously would have been even more computationally prohibitive.

ACKNOWLEDGMENTS

We are grateful to Darius Faroughy and Sung Hak Lim for helpful discussions and Gregor Kasieczka for comments on the draft. This work was supported by DOE Grant No. DE-SC0010008. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC award No. HEP-ERCAP0027491. The authors acknowledge the Office of Advanced Research Computing (OARC) at Rutgers, The State University of New Jersey [98] for providing access to the Amarel cluster and associated research computing resources that have contributed to the results reported here.

DATA AVAILABILITY

The data that support the findings of this article are openly available [99].

APPENDIX A: IMPORTANCE OF USING FREQUENCY EMBEDDING

The frequency embedding is necessary for v_{θ}^{data} to learn the small signal component in data. In this discussion, we denote the data model without a frequency embedding (i.e., using $\alpha(m) = m$ with m being an unscaled mass) as $\tilde{v}_{\theta}^{\text{data}}$. The data model with a frequency embedding, given by Eq. (11), is denoted as v_{θ}^{data} .




FIG. 4. The SIC curve for background template obtained v_{θ}^{data} (in black, solid) and $\tilde{v}_{\theta}^{\text{data}}$ (in black, dotted), trained on data without signal injection are similar to IAD. However, comparing the same models when trained with signal injection, show that the SIC for v_{θ}^{data} (with frequency embedding), is worse than $\tilde{v}_{\theta}^{\text{data}}$ (without frequency embedding). This is because v_{θ}^{data} learns the small signal component in data, better than $\tilde{v}_{\theta}^{\text{data}}$. Hence, the frequency embedding is crucial for the model to best learn the data.

As seen in Fig. 4, without signal injection, both v_{θ}^{data} and $\tilde{v}_{\theta}^{\text{data}}$ give a SIC very close to the IAD. This also implies that the frequency embedding does not affect the density estimation in the absence of any signal. However with the signal injection, this SIC worsens for both models. v_{θ}^{data} learns the signal better and hence shows worse SIC values than $\tilde{v}_{\theta}^{\text{data}}$.

For our method to be fully signal-model agnostic, we do not want the data model to smoothen out parts of the data distribution that could be relevant to the background template. Hence it is crucial that we start with the best possible data model, which in this case is the model with a frequency embedding.

APPENDIX B: CLASSIFIER TEST

To better understand what the models learned, we study the weight distributions from the classifier test introduced in [38]. We train an ensemble of 50 `HistGradientBoostingClassifier` classifiers to distinguish the samples from each model vs the true background samples. Hence the weight distributions show

$$\text{weight}(x) = \frac{P_{\text{samples}}(x)}{P_{\text{background}}(x)}. \quad (\text{B1})$$

If model v_{θ}^{data} does learn the signal, we expect the samples after a cut $\text{weight}(x) > 1$ to look like the signal. Whereas a similar cut on samples from v_{θ}^{CR} or v_{θ}^{int} should not look like the signal and should correspond to modeling between the interpolated samples and the background.

We see that v_{θ}^{CR} has the best AUC, and v_{θ}^{int} (linear) has the worst AUC. The difference in the width of the weight distributions of v_{θ}^{int} (linear) and v_{θ}^{int} (context) is due to different ways of interpolating v_{θ}^{data} into the SR, as noted in Section III A. v_{θ}^{int} (context) and v_{θ}^{data} has similar AUCs, but in Fig. 5 for samples from v_{θ}^{data} , the weight distribution has a long tail for $\text{weight} > 1$. The resulting samples from v_{θ}^{data} after a cut $\log(\text{weight}) > 0.1$ in Fig. 6 show the peaks where the signal is localized, as expected. The interpolated samples after the same cut do not show the signal.⁶ This shows that v_{θ}^{int} is able to remove the signal. For ΔR samples after a cut, v_{θ}^{int} (linear) shows the mismodeling artifacts which are also visible in pull distributions in Fig. 2.

APPENDIX C: SAMPLE QUALITY ACROSS DIFFERENT SRS

In the main body of the paper we focused on the SR $m_{JJ} \in [3.3, 3.7]$ TeV that is centered on the LHCO R&D signal with $m_{JJ} = 3.5$ TeV. Here we explore the

⁶Samples with similar AUCs result in samples with different anomaly detection performance. This shows that AUC by itself can be misleading, and one needs to study the weight distributions to get a better understanding, as was shown in [38].




FIG. 5. The weight distributions for classifier trained on samples from v^{data} (top left)/ v^{CR} (bottom left)/ v^{int} (context) (top right)/ v^{int} (linear) (bottom right) vs true background. The AUC of v^{int} (context) is slightly worse than v^{data} , which produces a long tail for $\log(\text{weight}) > 0$ corresponding to the signal.

quality of generated samples interpolating into other SRs that do not contain the signal. Taking the same FM model trained on the data which contains 1k signal events as in the main body of the paper, we interpolated it into overlapping SRs of width 400 GeV, centered at $m_{JJ} = (3.5 - 0.1 \cdot (5 - n))$ TeV, for $n = 1, \dots, 10$. $n = 5$ corresponds to the SR considered in the main body of the paper.

For each SR, we trained classifiers between the true background samples (in SR) vs the background template obtained from v^{int} (context). We used the classifier AUC as a metric to assess the background template quality. As shown in Fig. 7, we find that the classifier AUCs for all the SRs are close to ~ 0.5 . This shows that SIGMA is robust across different SRs.




FIG. 6. For $N_{\text{sig}} = 3000$ and $m \in \text{SR}$, this figure shows the histograms of the samples of the features after a classifier cut $\log \text{weight} > 0.1$. For ΔR , we chose a higher cut $\log \text{weight} > 0.2$. After a cut, we see that the samples from v^{data} (in orange) are localized around the signal. This shows that the model was able to learn the signal. The samples from $v^{\text{int}}(\text{context})$, $v^{\text{int}}(\text{linear})$ and v^{CR} do not show the signal. The samples from v^{int} after a cut look similar to the other background templates, except for ΔR , where the classifier finds the mismodeling artifacts.




FIG. 7. The figure shows the classifier AUC for background template quality in signal regions centered at $m_{JJ} = (3.5 - 0.1 \cdot (5 - n))$ TeV, for $n = 1, \dots, 10$. Classifiers are trained between true background samples in SR, and background template obtained from v_θ^{int} (context). SIGMA is robust across different SRs in the resonant feature.

-
- [1] J. H. Collins, K. Howe, and B. Nachman, *Phys. Rev. Lett.* **121**, 241803 (2018).
 - [2] J. H. Collins, K. Howe, and B. Nachman, *Phys. Rev. D* **99**, 014038 (2019).
 - [3] B. Nachman and D. Shih, *Phys. Rev. D* **101**, 075042 (2020).
 - [4] A. Andreassen, B. Nachman, and D. Shih, *Phys. Rev. D* **101**, 095004 (2020).
 - [5] G. Stein, U. Seljak, and B. Dai, in *Proceedings of the 34th Conference on Neural Information Processing Systems* (2020).
 - [6] O. Amram and C. M. Suarez, *J. High Energy Phys.* **01** (2021) 153.
 - [7] A. Hallin, J. Isaacson, G. Kasieczka, C. Krause, B. Nachman, T. Quadfasel, M. Schlaffer, D. Shih, and M. Sommerhalder, *Phys. Rev. D* **106**, 055006 (2022).
 - [8] J. H. Collins, P. Martín-Ramiro, B. Nachman, and D. Shih, *Eur. Phys. J. C* **81**, 617 (2021).
 - [9] K. Benkendorfer, L. L. Pottier, and B. Nachman, *Phys. Rev. D* **104**, 035003 (2020).
 - [10] G. Kasieczka, B. Nachman, and D. Shih, *arXiv:2107.02821*.
 - [11] A. Hallin, G. Kasieczka, T. Quadfasel, D. Shih, and M. Sommerhalder, *Phys. Rev. D* **107**, 114012 (2023).
 - [12] M. F. Chen, B. Nachman, and F. Sala, *J. High Energy Phys.* **07** (2023) 188.
 - [13] J. F. Kamenik and M. Szewc, *Phys. Lett. B* **840**, 137836 (2023).
 - [14] D. Sengupta, S. Klein, J. A. Raine, and T. Golling, *SciPost Phys.* **17**, 046 (2024).
 - [15] J. A. Raine, S. Klein, D. Sengupta, and T. Golling, *Front. Big Data* **6**, 899345 (2023).
 - [16] T. Golling, G. Kasieczka, C. Krause, R. Mastandrea, B. Nachman, J. A. Raine, D. Sengupta, D. Shih, and M. Sommerhalder, *Eur. Phys. J. C* **84**, 241 (2024).
 - [17] T. Golling, S. Klein, R. Mastandrea, and B. Nachman, *Phys. Rev. D* **107**, 096025 (2023).
 - [18] G. Bickendorf, M. Drees, G. Kasieczka, C. Krause, and D. Shih, *Phys. Rev. D* **109**, 096031 (2024).
 - [19] T. Finke, M. Hein, G. Kasieczka, M. Krämer, A. Mück, P. Prangaikul, T. Quadfasel, D. Shih, and M. Sommerhalder, *Phys. Rev. D* **109**, 034033 (2024).
 - [20] M. Freytsis, M. Perelstein, and Y. C. San, *J. High Energy Phys.* **02** (2024) 220.
 - [21] E. Buhmann, C. Ewen, G. Kasieczka, V. Mikuni, B. Nachman, and D. Shih, *Phys. Rev. D* **109**, 055015 (2024).
 - [22] R. Das, G. Kasieczka, and D. Shih, *arXiv:2312.11629*.
 - [23] H. Beauchesne, Z.-E. Chen, and C.-W. Chiang, *J. High Energy Phys.* **02** (2024) 138.
 - [24] D. Sengupta, M. Leigh, J. A. Raine, S. Klein, and T. Golling, *J. High Energy Phys.* **04** (2024) 109.
 - [25] G. Kasieczka *et al.*, *Rep. Prog. Phys.* **84**, 124201 (2021).
 - [26] D. Sengupta, S. Klein, J. A. Raine, and T. Golling, *SciPost Phys.* **17**, 046 (2024).
 - [27] R. Mastandrea and B. Nachman, in *Proceedings of the 36th Conference on Neural Information Processing Systems* (2022).
 - [28] C. L. Cheng, G. Singh, and B. Nachman, *Phys. Rev. Lett.* **135**, 021801 (2025).
 - [29] K. Benkendorfer, L. Le Pottier, and B. Nachman, *Phys. Rev. D* **104**, 035003 (2021).
 - [30] M. Leigh, D. Sengupta, B. Nachman, and T. Golling, *arXiv:2407.19818*.
 - [31] C. Li *et al.*, *arXiv:2405.12972*.
 - [32] G. Aad *et al.* (ATLAS Collaboration), *Phys. Rev. Lett.* **125**, 131801 (2020).
 - [33] CMS Collaboration Model-agnostic search for dijet resonances with anomalous jet substructure in proton-proton collisions at $\sqrt{s} = 13$ TeV, Technical Report No. CMS-EXO-22-026, CERN-EP-2024-291 (CERN, Geneva, 2024).
 - [34] D. Shih, M. R. Buckley, L. Necib, and J. Tamanas, *Mon. Not. R. Astron. Soc.* **509**, 5992 (2021).

- [35] D. Shih, M. R. Buckley, and L. Necib, *Mon. Not. R. Astron. Soc.* **529**, 4745 (2024).
- [36] M. Pettee, S. Thanvantri, B. Nachman, D. Shih, M. R. Buckley, and J. H. Collins, *Mon. Not. R. Astron. Soc.* **527**, 8459 (2023).
- [37] D. Sengupta, S. Mulligan, D. Shih, J. A. Raine, and T. Golling, *Mon. Not. R. Astron. Soc.* **536**, 1104 (2024).
- [38] R. Das, L. Favaro, T. Heimel, C. Krause, T. Plehn, and D. Shih, *SciPost Phys.* **16**, 031 (2024).
- [39] G. Kasieczka, B. Nachman, and D. Shih, R&D Dataset for LHC Olympics 2020 Anomaly Detection Challenge (2022) <https://doi.org/10.5281/zenodo.6466204>.
- [40] G. Kasieczka, B. Nachman, D. Shih *et al.*, *Rep. Prog. Phys.* **84**, 124201 (2021).
- [41] T. Sjöstrand, S. Mrenna, and P. Skands, *J. High Energy Phys.* **05** (2006) 026.
- [42] T. Sjöstrand, S. Mrenna, and P. Skands, *Comput. Phys. Commun.* **178**, 852 (2008).
- [43] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaitre, A. Mertens, and M. Selvaggi, *J. High Energy Phys.* **02** (2014) 057.
- [44] A. Mertens, *J. Phys. Conf. Ser.* **608**, 012045 (2015).
- [45] M. Selvaggi, *J. Phys. Conf. Ser.* **523**, 012033 (2014).
- [46] M. Cacciari and G. P. Salam, *Phys. Lett. B* **641**, 57 (2006).
- [47] M. Cacciari, G. P. Salam, and G. Soyez, *J. High Energy Phys.* **04** (2008) 063.
- [48] M. Cacciari, G. P. Salam, and G. Soyez, *Eur. Phys. J. C* **72**, 1896 (2012).
- [49] J. Thaler and K. Van Tilburg, *J. High Energy Phys.* **03** (2011) 015.
- [50] J. Thaler and K. Van Tilburg, *J. High Energy Phys.* **02** (2012) 093.
- [51] J. A. Raine, S. Klein, D. Sengupta, and T. Golling, *Front. Big Data* **6**, 89934 (2023).
- [52] D. Shih, Additional QCD Background Events for LHCO2020 R&D (signal region only) (2021).
- [53] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, [arXiv:2210.02747](https://arxiv.org/abs/2210.02747).
- [54] A. Tong, K. Fatras, N. Malkin, G. Huguette, Y. Zhang, J. Rector-Brooks, G. Wolf, and Y. Bengio, [arXiv:2302.00482](https://arxiv.org/abs/2302.00482).
- [55] A.-A. Pooladian, H. Ben-Hamu, C. Domingo-Enrich, B. Amos, Y. Lipman, and R. T. Q. Chen, [arXiv:2304.14772](https://arxiv.org/abs/2304.14772).
- [56] X. Liu, C. Gong, and Q. Liu, [arXiv:2209.03003](https://arxiv.org/abs/2209.03003).
- [57] M. S. Albergo and E. Vanden-Eijnden, [arXiv:2209.15571](https://arxiv.org/abs/2209.15571).
- [58] J. Ho, A. Jain, and P. Abbeel, [arXiv:2006.11239](https://arxiv.org/abs/2006.11239).
- [59] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 10684–10695.
- [60] T. Karras, M. Aittala, T. Aila, and S. Laine, *Adv. Neural Inf. Process. Syst.* **35**, 26565 (2022), [arXiv:2206.00364](https://arxiv.org/abs/2206.00364).
- [61] P. Dhariwal and A. Nichol, *Adv. Neural Inf. Process. Syst.* **34**, 8780 (2021), [arXiv:2105.05233](https://arxiv.org/abs/2105.05233).
- [62] Y. Song and S. Ermon, *Adv. Neural Inf. Process. Syst.* **32** (2019), [arXiv:1907.05600](https://arxiv.org/abs/1907.05600).
- [63] V. Mikuni and B. Nachman, *Phys. Rev. D* **106**, 092009 (2022).
- [64] V. Mikuni and B. Nachman, *J. Instrum.* **19**, P02001 (2024).
- [65] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, in *International Conference on Machine Learning* (Pmlr, 2021), pp. 8821–8831.
- [66] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, [arXiv:2204.06125](https://arxiv.org/abs/2204.06125).
- [67] J. Birk, E. Buhmann, C. Ewen, G. Kasieczka, and D. Shih, *Phys. Rev. D* **111**, 052008 (2025).
- [68] E. Buhmann, C. Ewen, D. A. Faroughy, T. Golling, G. Kasieczka, M. Leigh, G. Quétant, J. A. Raine, D. Sengupta, and D. Shih, [arXiv:2310.00049](https://arxiv.org/abs/2310.00049).
- [69] L. Favaro, A. Ore, S. Palacios Schweitzer, and T. Plehn, *SciPost Phys.* **18**, 088 (2025).
- [70] A. Butter, T. Jezo, M. Klasen, M. Kuschick, S. Palacios Schweitzer, and T. Plehn, *SciPost Phys. Core* **7**, 064 (2024).
- [71] J. Birk, A. Hallin, and G. Kasieczka, *Mach. Learn. Sci. Tech.* **5**, 035031 (2024).
- [72] F. Vaselli, F. Cattafesta, P. Asenov, and A. Rizzi, *Mach. Learn. Sci. Tech.* **5**, 035007 (2024).
- [73] T. Heimel, N. Huetsch, R. Winterhalder, T. Plehn, and A. Butter, *SciPost Phys.* **17**, 129 (2024).
- [74] V. Mikuni and B. Nachman, *Phys. Rev. D* **111**, L051504 (2025).
- [75] S. Bieringer, G. Kasieczka, J. Kieseler, and M. Trabs, *Eur. Phys. J. C* **84**, 972 (2024).
- [76] A. Butter, N. Huetsch, S. Palacios Schweitzer, T. Plehn, P. Sorrenson, and J. Spinner, *SciPost Phys. Core* **8**, 026 (2025).
- [77] N. Huetsch *et al.*, *SciPost Phys.* **18**, 070 (2025).
- [78] E. Buhmann, S. Diefenbacher, E. Eren, F. Gaede, G. Kasieczka, A. Korol, W. Korcari, K. Krüger, and P. McKeown, *J. Instrum.* **18**, P11025 (2023).
- [79] E. Buhmann, F. Gaede, G. Kasieczka, A. Korol, W. Korcari, K. Krüger, and P. McKeown, *J. Instrum.* **19**, P04020 (2024).
- [80] V. Mikuni and B. Nachman, *SciPost Phys.* **16**, 062 (2024).
- [81] A. Shmakov, K. T. Greif, M. J. Fenton, A. Ghosh, P. Baldi, and D. Whiteson, *SciPost Phys.* **18**, 117 (2025).
- [82] B. Hashemi and C. Krause, *Rev. Phys.* **12**, 100092 (2024).
- [83] D. Sengupta, M. Leigh, J. A. Raine, S. Klein, and T. Golling, *J. High Energy Phys.* **04** (2024) 109.
- [84] P. Devlin, J.-W. Qiu, F. Ringer, and N. Sato, *Phys. Rev. D* **110**, 016030 (2024).
- [85] O. Amram and K. Pedro, *Phys. Rev. D* **108**, 072014 (2023).
- [86] Z. Imani, T. Wongjirad, and S. Aeron, *Phys. Rev. D* **109**, 072011 (2024).
- [87] M. Leigh, D. Sengupta, J. A. Raine, G. Quétant, and T. Golling, *Phys. Rev. D* **109**, 012010 (2024).
- [88] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, [arXiv:1806.07366](https://arxiv.org/abs/1806.07366).
- [89] K. He, X. Zhang, S. Ren, and J. Sun, [arXiv:1512.03385](https://arxiv.org/abs/1512.03385).
- [90] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, *nflows: Normalizing flows in PyTorch* (2020).
- [91] A. Hallin, G. Kasieczka, T. Quadfasel, D. Shih, and M. Sommerhalder, *Phys. Rev. D* **107**, 114012 (2023).
- [92] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, [arXiv:2003.08934](https://arxiv.org/abs/2003.08934).

- [93] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. A. Hamprecht, Y. Bengio, and A. Courville, [arXiv:1806.08734](#).
- [94] A. Paszke *et al.*, [arXiv:1912.01703](#).
- [95] I. Loshchilov and F. Hutter, [arXiv:1711.05101](#).
- [96] F. Rozet *et al.*, Zuko: Normalizing flows in pytorch (2022), <https://doi.org/10.5281/zenodo.7625672>.
- [97] J. Zhuang, N. Dvornek, X. Li, S. Tatikonda, X. Papademetris, and J. Duncan, [arXiv:2006.02493](#).
- [98] <https://it.rutgers.edu/oarc>.
- [99] R. Das, Sigma: Single interpolated generative model for anomalies (2024), <https://github.com/rd804/SIGMA>.