
Unpaired Translation of Point Clouds for Modeling Detector Response

Mingyang Li¹ Michelle Kuchera¹ Raghuram Ramanujan¹
Adam Anthony² Curtis Hunt³ Yassid Ayyad⁴

¹ Davidson College ² High Point University

³ Michigan State University ⁴ University of Santiago de Compostela

{mili, mikuchera, ramananujan}@davidson.edu

{aanthon2}@highpoint.edu

{huntc}@frib.msu.edu

{yassid.ayyad}@usc.es

Abstract

Modeling detector response is a key challenge in time projection chambers. We cast this problem as an unpaired point cloud translation task, between data collected from simulations and from experimental runs. Effective translation can assist with both noise rejection and the construction of high-fidelity simulators. Building on recent work in diffusion probabilistic models, we present a novel framework for performing this mapping. We demonstrate the success of our approach in both synthetic domains and in data sourced from the Active-Target Time Projection Chamber.

1 Introduction

Time projection chambers (TPCs) are used for full 3-D reconstruction of events in nuclear and particle physics experiments. The high angular coverage and sensitivity of TPCs allow for the possibility of capturing and analyzing nearly all reaction particles within the chamber. This necessarily increases the complexity of the device, and often leads to detector signals with both electronic and physical noise effects that are difficult to model accurately, especially in the case of gas-filled TPCs, which we study in this paper. Recently, there has been a surge of interest in using machine learning, and particularly generative models, to create theory-agnostic, data-driven simulators in nuclear and particle physics [1–5].

Despite the promise of generative models, their effectiveness hinges on the nature of the data used for training, presenting a key trade-off. Models trained on experimental data can capture the marginal distribution of the observations, including nuances such as the detector response. However, such data requires unfolding techniques to recover true particle information, limiting their utility in targeted analyses on simulated data. Conditional generative models trained on labeled simulated data offer precise control over the sampling process. But this comes at the cost of simulation fidelity, as such models cannot accurately capture the detector response. Addressing this tension—between the accuracy of the model trained on the experimental data and the controllability of the model trained on the simulations—remains a key problem.

In this work, we address this challenge by presenting a framework for *unpaired translation of point clouds*: given a pair of datasets \mathcal{D}_X and \mathcal{D}_Y comprising point clouds, we train models $G : \mathcal{D}_X \mapsto \mathcal{D}_Y$ and $H : \mathcal{D}_Y \mapsto \mathcal{D}_X$ that can transform events from one domain into events that resemble members of the other domain. In the scenario where \mathcal{D}_X is simulated data and \mathcal{D}_Y is experimental data, G models the detector response for simulated events, while H subtracts noise from experimental events so that they resemble the idealized versions generated by simulators. Crucially, no one-to-

one correspondence between the elements of \mathcal{D}_X and \mathcal{D}_Y is required to learn these translational mappings. We evaluate our approach on synthetic point clouds, as well as data sourced from an experiment conducted using the Active-Target Time Projection Chamber (AT-TPC) [6] at the National Superconducting Cyclotron Laboratory (NSCL).

2 Background and Related Work

Diffusion Models Diffusion probabilistic models (DPMs) have recently gained significant attention due to their impressive results in generative tasks, particularly in areas such as image [7–9] and text generation [10, 11]. A stochastic diffusion process is a Markov chain defined by a series of noise-addition steps that gradually distort the data. The central idea behind diffusion models is to use another Markov chain to convert a sample drawn from a Gaussian distribution back into a sample from the original distribution, by applying learned transformations. In particle physics, there have been significant efforts in using diffusion models for simulation of point cloud data [12–19]. Additionally, image-based diffusion models were used for simulation in LArTPCs [20, 15]. In this work, we adapt a diffusion-based approach proposed by Luo and Hu for generating point clouds, and use it as a building-block in our system [21].

Unpaired Translation The task of image translation using conditional GANs was first proposed and studied in the seminal Pix2Pix paper [22]. This work assumed access to paired data, where each image in the source domain was aligned with a corresponding, translated image in the target domain – for example, the same scene as viewed during the day and at night. To address the challenge of collecting such data at scale, Zhu et al. employed a novel cycle consistency loss to perform image-to-image translation in an unpaired setting [23]. Follow-up work has improved upon unpaired translation performance on images using both GANs [24, 25] and diffusion models [26, 27]. More recently, Wu and De la Torre proposed CycleDiffusion for unpaired image translation with diffusion models [28]. Their key contribution was the DPM-Encoder, a novel method for encoding input images into a latent space that permitted perfect reconstruction via reverse diffusion. In physics, several applications of unpaired translation have been proposed in recent years [29–33]. Notably, Huang et al. used unpaired image translation to map between LArTPC simulations and experiments. Prior unpaired point cloud translation work using diffusion models has focused on the application of 3-D object completion [34, 35]. Zheng et al. and Li and Lin proposed non-diffusion based unpaired point cloud translation architectures for objects such as furniture [36, 37]. Our work focuses on the unpaired translation of point cloud track data from experimental and simulated AT-TPC data using diffusion models.

AT-TPC The AT-TPC is a detector that is used to probe the structure of atomic nuclei. In this work, we used data from an AT-TPC experiment studying the fission properties of nuclei in the neutron deficient lead region [38]. A heavy ion beam was produced via ^{208}Pb fragmentation, and the species of interest separated using the A1900 fragment separator [39, 40]. The resulting cocktail beam was identified using the HEIST particle identification system [41] before the beam was directed into the AT-TPC where fission was induced through a fusion-fission reaction using He gas as the target. By collecting the charge produced as the beam and fission fragments travel through the AT-TPC on its highly segmented pad plane, the particle tracks are reconstructed as four-dimensional (three space, one charge) point clouds. The fission experiment produced only two fission products, creating “Y”-shaped events. This dataset was chosen out of recent AT-TPC experiments as a useful first application because the relative visual simplicity of the events allow for straightforward qualitative evaluation of our results.

3 Methods

We begin with an informal, high-level description of our system which combines CycleDiffusion [28], an unpaired image translation framework, with a specialized diffusion-based generative model for point clouds [21]. CycleDiffusion relies on a striking property of diffusion models: that they produce “*uniquely identifiable encodings*” [42], i.e., diffusion models trained on the same data produce close to identical output images when primed with the same randomly sampled input, regardless of model architecture, or training and sampling procedures [42]. In fact, this alignment of latent representations occurs even when the models are trained on different datasets — Khruikov et al. found that characteristics such as color and texture, as well as high-level semantic features such as pose and gender are preserved when the same latent code is decoded by networks modeling

completely different distributions [43]. This suggests one way to perform unpaired translation between domains: we independently train two diffusion models on the two domains of interest \mathcal{D}_X and \mathcal{D}_Y . We encode a source point cloud using its in-domain encoder, but then decode the resultant latent representation with the other diffusion model. This, in effect, is CycleDiffusion [28]. We leverage Luo and Hu’s point cloud diffusion architecture to adapt this idea to the translation of point clouds [21]. We discuss some of the relevant technical details in the rest of this section.

3.1 Architecture Details

We begin by first training two independent point cloud DPMs, one for each domain \mathcal{D}_X and \mathcal{D}_Y . These models use the encoder-decoder architecture described by Luo and Hu, modified to accept our 4-D point cloud data [21]. The encoders are standard PointNet encoders [44], while the decoders are neural networks implementing a reverse diffusion process [21]. The values of key hyperparameters used for training these models are given in Table 1. Once training is complete, the components of the DPMs are rewired as shown in the schematic in Figure 1. Here, Enc_X (Enc_Y) and Dec_X (Dec_Y) denote the trained encoder and decoder for domain \mathcal{D}_X (\mathcal{D}_Y) respectively. The $DPMEnc_X$ module is a DPM-Encoder which is described in greater detail in Section 3.2. This component contains no trainable parameters. To translate an event $\mathbf{X}^{(0)}$ from domain \mathcal{D}_X to domain \mathcal{D}_Y , we start by processing the event through both encoders, Enc_X and Enc_Y , to obtain its latent representations z_X and z_Y within each respective domain. Subsequently, we pass the latent representation z_X along with the event $\mathbf{X}^{(0)}$ to $DPMEnc_X$, which simulates the forward diffusion process for T steps producing the event $\mathbf{X}^{(T)}$ along with the estimated noise added at each step ϵ_X . Note that $\epsilon_X = \epsilon_T \oplus \dots \epsilon_2 \oplus \epsilon_1$, where \oplus denotes concatenation. In the final step, the decoder for the \mathcal{D}_Y domain, Dec_Y , uses $\mathbf{Y}^{(T)} = \mathbf{X}^{(T)}$ as the initial point cloud and starts the decoding process, guided by z_Y and ϵ_X . One single decoding step is computed as

$$\mathbf{Y}^{(t-1)} = \mu_Y(\mathbf{Y}^{(t)}, t, z_Y) + \sigma_t \epsilon_t \quad t = T, \dots, 1 \quad (1)$$

where μ_Y is the learned mean estimator in Dec_Y [21], which estimates the mean point cloud $p(\mathbf{Y}^{(t-1)}|\mathbf{Y}^{(t)})$, and stochastically at time step t is added by $\sigma_t \epsilon_t$, where $\epsilon_t \sim \mathcal{N}(0, \mathbf{I})$ and σ_t is the variance of the added noise at time step t which increases linearly with time. The result of the final decoding step $\mathbf{Y}^{(0)}$ is the translation of the original event $\mathbf{X}^{(0)}$ into domain \mathcal{D}_Y . To perform the translation in the opposite direction, we can simply swap X and Y in the preceding discussion. The code to reproduce our model and results can be accessed at the following repository: https://github.com/alpha-davidson/attpc_unpaired_trans.




Figure 1: Schematic depicting our unpaired translation model at inference time. Arrows in red denote our modifications to the standard CycleDiffusion architecture to incorporate a point cloud diffusion model.

Hyperparameter	Value
Initial Learning Rate	0.001
Final Learning Rate	0.0001
Optimization Algorithm	Adam [45]
Batch Size	128
Number of Training Iterations	1,000,000
Latent Dimension Size	256
Diffusion Steps	256

Table 1: Hyperparameters used during training the DPMs.

3.2 Modified DPM-Encoder

Wu and De la Torre proposed the DPM-Encoder as a means to deterministically map images to a latent code in stochastic diffusion models [28]. Their key insight was to retain the noise that was added to the samples during each step of the forward diffusion process, together with the final output $\mathbf{X}^{(T)}$, and to treat the concatenation of these quantities as the latent encoding of the original input

image. While this is a relatively straightforward idea, it leads to impressive results in a variety of tasks, including unpaired translation of images [28]. Unlike image data, however, point cloud events are permutation invariant, and thus, we cannot use the DPM-Encoder directly in our framework. We augment the DPM-Encoder with an additional shape latent that is computed using a PointNet encoder (see Figure 1). This enables the encoding of topological information about the original point cloud $\mathbf{X}^{(0)}$ in the latent representation, prior to decoding and translation via the reverse diffusion process in Dec_Y .

3.3 Data

We evaluate our approach on three different domains, in each case translating between a “clean” and a “noisy” point cloud.

1. $(\mathcal{L}_X, \mathcal{L}_Y)$: a pair of synthetic point cloud datasets comprising clean and noisy line segments, respectively, oriented along the z axis. \mathcal{L}_X is a proxy for simulated data, while \mathcal{L}_Y is a stand-in for experimental data. Each line is composed of points with $x = 0$, a fixed y sampled from $U(0, 2)$ and 256 z values sampled from $U(0, 2)$. We corrupt the elements of \mathcal{L}_Y by adding noise sampled from $\mathcal{N}(0, 0.1y)$ to each point. This task is designed to evaluate the model’s ability to capture simple non-uniform noise distributions. \mathcal{L}_X and \mathcal{L}_Y each consist of 1,000 events.
2. $(\mathcal{G}_X, \mathcal{G}_Y)$: a pair of synthetic datasets containing clean and noisy geometric shapes (triangular prisms and cuboids), respectively, of different sizes and orientations. The shapes are created by sampling points uniformly at random along the edges of the solid. We add Gaussian noise to each point in each example in \mathcal{G}_Y to create the noisy examples. This task is intended to assess the model’s ability to perform translations between two domains in a multi-class setting. \mathcal{G}_X comprises 2,000 examples, evenly distributed between 1,000 triangular prisms and 1,000 cuboids; \mathcal{G}_Y contains noisy versions of the same examples.
3. $(\mathcal{A}_X, \mathcal{A}_Y)$: simulated and experimental fission events from the AT-TPC, respectively. These events have a characteristic “Y” shape, where the base of the Y is the incoming beam and the other two legs are outgoing the fission fragments. Our simulated dataset consists of 5,000 events, while the experimental dataset includes 11,500 events selected by applying a cut on the fission angle to filter out events where no fission occurred.

4 Results

We evaluate the ability of our model to learn a conditional noise distribution in translation on the $(\mathcal{L}_X, \mathcal{L}_Y)$ dataset by calculating the fitted standard deviation and error after translation. The quantitative results can be seen in Tab. 2, and qualitative translation results are shown in Fig. 2.

y	$\sigma(y)$	$\sigma_T \pm \sigma_T/\sqrt{N}$	MAE
0.04	.004	0.010 ± 0.0004	0.005
0.63	.063	0.076 ± 0.003	0.013
1.32	.132	0.121 ± 0.005	0.009
1.99	.199	0.181 ± 0.008	0.015

Table 2: Standard deviations of noise added by the $\mathcal{L}_X \rightarrow \mathcal{L}_Y$ translation process using our model. We report the ground truth standard deviation $\sigma(y) = 0.1y$, the standard deviation, σ_T , of translated events with estimated standard error σ_T/\sqrt{N} , and the mean absolute error between $\sigma(y)$ and σ_T . The consistently small MAE values demonstrate that the translation model is capable of learning conditional noise relationships.

We evaluate the translation and reconstruction quality of our model on the $(\mathcal{G}_X, \mathcal{G}_Y)$ and $(\mathcal{A}_X, \mathcal{A}_Y)$ datasets. Given a domain D , we calculate the Jensen-Shannon divergence (JSD) between events translated into domain D and the original events in D to quantify the quality of the translation (JSD(trans)). We compare this to two baseline JSD scores: one computed between different batches of events drawn from within the same domain D (JSD(in-domain)), and one computed between a batch of randomly generated point clouds and events from D (JSD(rand)). We use Chamfer Distance



Figure 2: Samples of translation results on on the $(\mathcal{L}_X, \mathcal{L}_Y)$ dataset in two directions.

(CD) to measure reconstruction accuracy after performing bidirectional translations (for example, $\mathcal{G}_X \rightarrow \mathcal{G}_Y \rightarrow \mathcal{G}_X$). Quantitative results are shown in Tab. 3, and visual results are shown in Fig. 3.

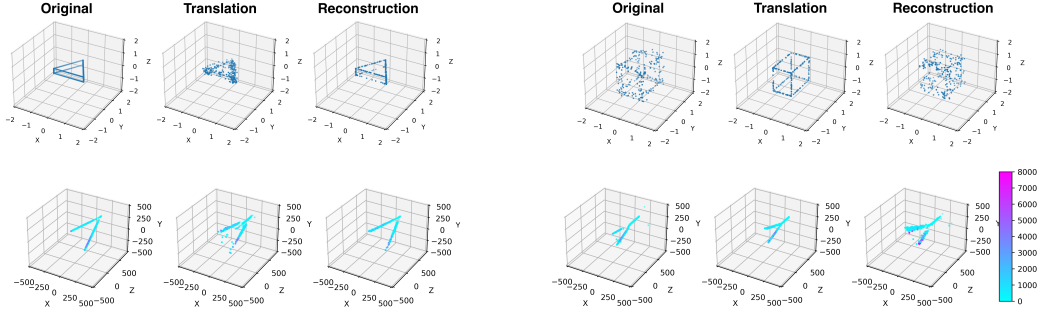


Figure 3: Translation and reconstruction results on $(\mathcal{G}_X, \mathcal{G}_Y)$ (top) and $(\mathcal{A}_X, \mathcal{A}_Y)$ data (bottom).

Model	JSD(trans)	JSD(in-domain)	JSD(rand)	CD(reco)	$\sigma_{CD(reco)}$	CD(clean)
$\mathcal{G}_X \rightarrow \mathcal{G}_Y$	0.323	0.314	0.711	0.163	0.179	—
$\mathcal{G}_Y \rightarrow \mathcal{G}_X$	0.382	0.270	0.701	0.153	0.157	—
$\mathcal{A}_X \rightarrow \mathcal{A}_Y$	0.044	0.005	0.701	0.053	0.066	—
$\mathcal{A}_Y \rightarrow \mathcal{A}_X$	0.005	0.004	0.703	0.269	0.523	0.079

Table 3: Evaluation metrics for the models trained using the $(\mathcal{G}_X, \mathcal{G}_Y)$ and $(\mathcal{A}_X, \mathcal{A}_Y)$ datasets. JSD is used to evaluate translation quality, and is compared with JSDs of in-domain and random data. CD is used to evaluate reconstruction quality. For experimental data reconstruction, we remove the extreme outliers, which comprised $\sim 1\%$ of the reconstructed events, to create a clean reco subdataset. Our results on the $(\mathcal{G}_X, \mathcal{G}_Y)$ dataset provide a baseline performance, and our comparable results on the $(\mathcal{A}_X, \mathcal{A}_Y)$ datasets provide support that this approach is suitable for our application, aside from a small percentage of outliers.

5 Conclusions

Wu and De la Torre proposed a DPM-Encoder that maps images into an exactly reconstructable latent space. We used DPMs with a modified DPM-Encoder to address the challenge of unpaired translation in AT-TPC events. By effectively translating between simulated and experimental datasets, our model offers an approach for generating realistic detector responses and reducing noise in experimental data.

Acknowledgments and Disclosure of Funding

This material is based on work supported by the National Science Foundation under grants OAC-296101 and PHY-296009.

References

- [1] Anja Butter and Tilman Plehn. *Generative Networks for LHC Events*, chapter Chapter 7, pages 191–240. World Scientific Publishing Company, 2022. doi: 10.1142/9789811234033_0007. URL https://www.worldscientific.com/doi/abs/10.1142/9789811234033_0007.
- [2] Yasir Alanazi, Nobuo Sato, Pawel Ambrozewicz, Astrid Hiller-Blin, Wally Melnitchouk, Marco Battaglieri, Tianbo Liu, and Yaohang Li. A survey of machine learning-based physics event generation. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4286–4293. International Joint Conferences on Artificial Intelligence Organization, 8 2021. doi: 10.24963/ijcai.2021/588. URL <https://doi.org/10.24963/ijcai.2021/588>. Survey Track.
- [3] Anja Butter, Tilman Plehn, Steffen Schumann, Simon Badger, Sascha Caron, Kyle Cranmer, Francesco Armando Di Bello, Etienne Dreyer, Stefano Forte, Sanmay Ganguly, Dorival Gonçalves, Eilam Gross, Theo Heimel, Gudrun Heinrich, Lukas Heinrich, Alexander Held, Stefan Höche, Jessica N. Howard, Philip Ilten, Joshua Isaacson, Timo Janßen, Stephen Jones, Marumi Kado, Michael Kagan, Gregor Kasieczka, Felix Kling, Sabine Kraml, Claudius Krause, Frank Krauss, Kevin Kröninger, Rahool Kumar Barman, Michel Luchmann, Vitaly Magerya, Daniel Maitre, Bogdan Malaescu, Fabio Maltoni, Till Martini, Olivier Mattelaer, Benjamin Nachman, Sebastian Pitz, Juan Rojo, Matthew Schwartz, David Shih, Frank Siegert, Roy Stegeman, Bob Stienen, Jesse Thaler, Rob Verheyen, Daniel Whiteson, Ramon Winterhalder, and Jure Zupan. Machine learning and LHC event generation. *SciPost Phys.*, 14:079, 2023. doi: 10.21468/SciPostPhys.14.4.079. URL <https://scipost.org/10.21468/SciPostPhys.14.4.079>.
- [4] Baran Hashemi and Claudius Krause. Deep generative models for detector signature simulation: A taxonomic review. *Reviews in Physics*, 12:100092, 2024. ISSN 2405-4283. doi: <https://doi.org/10.1016/j.revip.2024.100092>. URL <https://www.sciencedirect.com/science/article/pii/S2405428324000029>.
- [5] Farzana Yasmin Ahmad, Vanamala Venkataswamy, and Geoffrey Fox. A comprehensive evaluation of generative models in calorimeter shower simulation, 2024. URL <https://arxiv.org/abs/2406.12898>.
- [6] J. Bradt, D. Bazin, F. Abu-Nimeh, T. Ahn, Y. Ayyad, S. Beceiro Novo, L. Carpenter, M. Cortesi, M.P. Kuchera, W.G. Lynch, W. Mittig, S. Rost, N. Watwood, and J. Yurkon. Commissioning of the active-target time projection chamber. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 875:65–79, 2017. ISSN 0168-9002. doi: <https://doi.org/10.1016/j.nima.2017.09.013>. URL <https://www.sciencedirect.com/science/article/pii/S0168900217309683>.
- [7] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models, 2022. URL <https://arxiv.org/abs/2112.10741>.
- [8] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.
- [9] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- [10] Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto. Diffusion-lm improves controllable text generation. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9781713871088.
- [11] Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *International Conference on Machine Learning*, 2023. URL <https://api.semanticscholar.org/CorpusID:264451832>.
- [12] Matthew Leigh, Debajyoti Sengupta, Guillaume Quétant, John Andrew Raine, Knut Zoch, and Tobias Golling. PC-JeDi: Diffusion for particle cloud generation in high energy physics. *SciPost Phys.*, 16:018, 2024. doi: 10.21468/SciPostPhys.16.1.018. URL <https://scipost.org/10.21468/SciPostPhys.16.1.018>.
- [13] Erik Buhmann, Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Anatolii Korol, William Korcari, Katja Krüger, and Peter McKeown. Caloclouds: fast geometry-independent highly-granular calorimeter simulation. *Journal of Instrumentation*, 18(11):P11025, nov 2023. doi: 10.1088/1748-0221/18/11/P11025. URL <https://dx.doi.org/10.1088/1748-0221/18/11/P11025>.

- [14] Erik Buhmann, Frank Gaede, Gregor Kasieczka, Anatolii Korol, William Korcari, Katja Krüger, and Peter McKeown. Caloclouds ii: ultra-fast geometry-independent highly-granular calorimeter simulation. *Journal of Instrumentation*, 19(04):P04020, apr 2024. doi: 10.1088/1748-0221/19/04/P04020. URL <https://dx.doi.org/10.1088/1748-0221/19/04/P04020>.
- [15] Matthew Leigh, Debajyoti Sengupta, John Andrew Raine, Guillaume Quétant, and Tobias Golling. Faster diffusion model with improved quality for particle cloud generation. *Phys. Rev. D*, 109:012010, Jan 2024. doi: 10.1103/PhysRevD.109.012010. URL <https://link.aps.org/doi/10.1103/PhysRevD.109.012010>.
- [16] Fernando Torales Acosta, Vinicius Mikuni, Benjamin Nachman, Miguel Arratia, Bishnu Karki, Ryan Milton, Piyush Karande, and Aaron Angerami. Comparison of point cloud and image-based models for calorimeter fast simulation. *Journal of Instrumentation*, 19(05):P05003, may 2024. doi: 10.1088/1748-0221/19/05/P05003. URL <https://dx.doi.org/10.1088/1748-0221/19/05/P05003>.
- [17] Erik Buhmann, Cedric Ewen, Darius A. Faroughy, Tobias Golling, Gregor Kasieczka, Matthew Leigh, Guillaume Quétant, John Andrew Raine, Debajyoti Sengupta, and David Shih. Epic-ly fast particle cloud generation with flow-matching and diffusion, 2023. URL <https://arxiv.org/abs/2310.00049>.
- [18] Cheng Jiang, Sitian Qian, and Huilin Qu. Choose your diffusion: Efficient and flexible ways to accelerate the diffusion model in fast high energy physics simulation, 2024. URL <https://arxiv.org/abs/2401.13162>.
- [19] Guillaume Quétant, John Andrew Raine, Matthew Leigh, Debajyoti Sengupta, and Tobias Golling. Pippin: Generating variable length full events from partons, 2024. URL <https://arxiv.org/abs/2406.13074>.
- [20] Zeviel Imani, Taritree Wongjirad, and Shuchin Aeron. Score-based diffusion models for generating liquid argon time projection chamber images. *Phys. Rev. D*, 109:072011, Apr 2024. doi: 10.1103/PhysRevD.109.072011. URL <https://link.aps.org/doi/10.1103/PhysRevD.109.072011>.
- [21] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2837–2845, June 2021.
- [22] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5967–5976, 2017. doi: 10.1109/CVPR.2017.632.
- [23] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251, 2017. doi: 10.1109/ICCV.2017.244.
- [24] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks, 2020.
- [25] Amélie Royer, Konstantinos Bousmalis, Stephan Gouws, Fred Bertsch, Inbar Mosseri, Forrester Cole, and Kevin Murphy. *XGAN: Unsupervised Image-to-Image Translation for Many-to-Many Mappings*, pages 33–49. Springer International Publishing, Cham, 2020. ISBN 978-3-030-30671-7. doi: 10.1007/978-3-030-30671-7_3. URL https://doi.org/10.1007/978-3-030-30671-7_3.
- [26] Hiroshi Sasaki, Chris G. Willcocks, and Toby P. Breckon. Unit-ddpm: Unpaired image translation with denoising diffusion probabilistic models, 2021.
- [27] Min Zhao, Fan Bao, Chongxuan Li, and Jun Zhu. Egsde: Unpaired image-to-image translation via energy-guided stochastic differential equations. *ArXiv*, abs/2207.06635, 2022. URL <https://api.semanticscholar.org/CorpusID:250526607>.
- [28] Chen Henry Wu and Fernando De la Torre. A latent space of stochastic diffusion models for zero-shot image editing and guidance. In *ICCV*, 2023.
- [29] Abid Khan, Chia-Hao Lee, Pinshane Y. Huang, and Bryan K. Clark. Leveraging generative adversarial networks to create realistic scanning transmission electron microscopy images. *npj Computational Materials*, 9(1), 5 2023. doi: 10.1038/s41524-023-01042-3.
- [30] Romane Scherrer, Thomas Quiniou, Thierry Jauffrais, Hugues Lemonnier, Sophie Bonnet, and Nazha Selmaoui-Folcher. Holographic reconstruction enhancement via unpaired image-to-image translation. *Appl. Opt.*, 61(33):9807–9816, Nov 2022. doi: 10.1364/AO.471131. URL <https://opg.optica.org/ao/abstract.cfm?URI=ao-61-33-9807>.

- [31] Stefano Carrazza and Frédéric A. Dreyer. Lund jet images from generative and cycle-consistent adversarial networks. *Eur. Phys. J.*, C79(11):979, 2019. doi: 10.1140/epjc/s10052-019-7501-1.
- [32] D. Torbunov, Y. Huang, H. Yu, J. Huang, S. Yoo, M. Lin, B. Viren, and Y. Ren. Uvcgan: Unet vision transformer cycle-consistent gan for unpaired image-to-image translation. In *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 702–712, Los Alamitos, CA, USA, jan 2023. IEEE Computer Society. doi: 10.1109/WACV56688.2023.00077. URL <https://doi.ieeecomputersociety.org/10.1109/WACV56688.2023.00077>.
- [33] Yi Huang, Dmitrii Torbunov, Brett Viren, Haiwang Yu, Jin Huang, Meifeng Lin, and Yihui Ren. Unpaired image translation to mitigate domain shift in liquid argon time projection chamber detector responses, 2024. URL <https://arxiv.org/abs/2304.12858>.
- [34] Xin Wen, Zhizhong Han, Yan-Pei Cao, Pengfei Wan, Wen Zheng, and Yu-Shen Liu. Cycle4completion: Unpaired point cloud completion using cycle transformation with missing region coding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [35] Wenxiao Zhang, Hossein Rahmani, Xun Yang, and Jun Liu. Reverse2complete: Unpaired multimodal point cloud completion via guided diffusion. In *ACM Multimedia 2024*, 2024. URL <https://openreview.net/forum?id=fheUSqLh7J>.
- [36] Jia-Wen Zheng, Jhen-Yung Hsu, Chih-Chia Li, and I-Chen Lin. Characteristic-preserving latent space for unpaired cross-domain translation of 3d point clouds. *IEEE Transactions on Visualization and Computer Graphics*, 30(8):5212–5226, 2024. doi: 10.1109/TVCG.2023.3287923.
- [37] Chih-Chia Li and I-Chen Lin. Unpaired translation of 3d point clouds with multi-part shape representation. *Proc. ACM Comput. Graph. Interact. Tech.*, 6(1), may 2023. doi: 10.1145/3585508. URL <https://doi.org/10.1145/3585508>.
- [38] Adam K Anthony. *Fission in the Lead Region*. PhD thesis, Michigan State University, 2023.
- [39] DJ Morrissey, BM Sherrill, M Steiner, A Stolz, and I Wiedenhoever. Commissioning the A1900 projectile fragment separator. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 204:90–96, 2003.
- [40] A Gade and BM Sherrill. NSCL and FRIB at Michigan State University: Nuclear science at the limits of stability. *Physica Scripta*, 91(5):053003, 2016.
- [41] A. K. Anthony, C. Y. Niu, R. S. Wang, J. Wieske, K. W. Brown, Z. Chajecski, W. G. Lynch, Y. Ayyad, J. Barney, T. Baumann, D. Bazin, S. Beceiro-Novo, J. Boza, J. Chen, K. J. Cook, M. Cortesi, T. Ginter, W. Mittig, A. Pype, M. K. Smith, C. Soto, C. Sumithrarachchi, J. Swaim, S. Sweany, F. C. E. Teh, C. Y. Tsang, M. B. Tsang, N. Watwood, and A. H. Wuosmaa. Beam particle identification and tagging of incompletely stripped heavy beams with HEIST. *Review of Scientific Instruments*, 93(1):013306, 01 2022. ISSN 0034-6748. doi: 10.1063/5.0068180. URL <https://doi.org/10.1063/5.0068180>.
- [42] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=PxtTIG12RRHS>.
- [43] Valentin Khruikov, Gleb Ryzhakov, Andrei Chertkov, and Ivan Oseledets. Understanding ddpn latent codes through optimal transport. *arXiv preprint arXiv:2202.07477*, 2022.
- [44] R. Qi Charles, Hao Su, Mo Kaichun, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, 2017. doi: 10.1109/CVPR.2017.16.
- [45] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.