

---

# Pay Attention to Mean-Fields for Point Cloud Generation

---

**Benno Käch**

Deutsches Elektronen-Synchrotron DESY,  
Germany  
benno.kaech@desy.de

**Isabell-A. Melzer-Pellmann**

Deutsches Elektronen-Synchrotron DESY,  
Germany  
isabell.melzer@desy.de

**Dirk Krücker**

Deutsches Elektronen-Synchrotron DESY,  
Germany  
dirk.kruecker@desy.de

## Abstract

Collider data generation with machine learning has become increasingly popular in particle physics due to the high computational cost of conventional Monte Carlo simulations, particularly for future high-luminosity colliders. We propose a generative model for point clouds that employs an attention-based aggregation while preserving a linear computational complexity with respect to the number of points. The model is trained in an adversarial setup, ensuring input permutation equivariance and invariance for the generator and critic, respectively. To stabilize known unstable adversarial training, a feature matching loss is introduced. We evaluate the performance on two different datasets. The former is the top-quark JETNET150 dataset, where the model outperforms the current state-of-the-art GAN-based model, despite having significantly fewer parameters. The latter is dataset 2 of the CaloChallenge, which comprises point clouds with up to  $30\times$  more points compared to the first dataset. The model and its corresponding code are available at <https://github.com/kaechb/MDMA/tree/NeurIPS>.

## 1 Introduction

Machine Learning (ML) has a longstanding presence in High Energy Physics (HEP) with its applications becoming standard across all stages of HEP data analysis [1]. While supervised learning has predominantly been used, there is a growing interest in generative models due to the detailed Monte Carlo simulations (MC) typically utilised in HEP. They provide near-perfect representations of experimental measurements, but they demand vast computing resources [2]. As higher and higher luminosities are reached (e.g. at the High Luminosity LHC) [3] this becomes a more and more pressing problem. A potential solution could be provided by generative models, as they are orders of magnitude faster during inference and have proven to perform well, e.g. in image generation [4; 5]. However, there is no universal solution for assessing whether the multidimensional features that the model is sampling contain the same correlations between the different features as the true model. Attention-based aggregation methods have displayed competitive performance for a wide array of problems. The inclusion of a mechanism in the proposed model was motivated by this factor. However, the quadratic scaling of self-attention with the number of inputs does not allow direct application to large point clouds. Thus, our information aggregation relies on cross-attention to an artificial mean-field, resulting in a linear scaling of computational complexity. The performance is

first evaluated on the JETNET150 [6] top-quark dataset, which comprises of PYTHIA simulated particles [7], clustered to jets using the anti- $k_T$  algorithm [8]. While PYTHIA jet generation is not the main computational bottleneck that needs resolution, there are known complex high-level analytical correlations (e.g. the relative invariant jet mass) which are difficult to capture. The metrics discussed in [9] also allow for convenient benchmarking with other models from the field. Subsequently, the application of the same model to calorimeter simulation is explored on the CaloChallenge Dataset 2 [10]. The calorimeter energy deposits are first converted to point clouds to make use of the sparsity which is present. Although this requires extra positional coordinates of the hits, this efficiently represents the data by excluding empty cells.

## 2 Related Work

### 2.1 JETNET150

The JETNET150 datasets, introduced by Kansal et al. [6], contain five datasets of about 180,000 samples each. For the sake of brevity, we will only discuss the dataset containing jets initiated by top quarks. However, it is anticipated that the latter contains the most intricate substructure among all the available jet datasets. For further clarification, information regarding particle features and high-level correlations, such as the invariant mass of the jet, can be found in [6]. Key metrics, Wasserstein distances, Frechet Physics Distance (FPD) and Kernel Physics Distance (KPD), are employed as defined in [9]. Multiple models [11; 12; 13; 14; 15; 16; 9] exhibit noteworthy performance on these datasets. We compare our model to EPiC-GAN by Buhmann et al. [16] as it is the only model that scales up to the JETNET150 dataset [17] using only particle multiplicity as a condition, and no other jet observables as input to the generator.

### 2.2 CaloChallenge

The CaloChallenge dataset contains calorimeter showers simulated by GEANT4 [18], paired together with the energy of the incoming particle. To transform these to point clouds, the voxel index is used as a coordinate. The spatial location is dequantised by adding noise distributed as  $p(x) = \frac{(n_2 - n_1)}{n_1 + n_2} \cdot x$ , where  $n_2, n_1$  are the total number of hits in neighbouring bins in the training set and  $x \sim U([0, 1])$ . Other models have been proposed for these datasets [19; 20], but their adaptability to non-idealised detectors with non-regular geometry remains uncertain. Due to the unavailability of general metrics that allow meaningful comparison, the primary focus will be on presenting and briefly discussing our results. On this dataset there are on average  $\sim 1600$  hits, with a maximum of up to  $\sim 4000$  hits per shower, and a total of 6480 cells, scaling up the number of points to generate in a cloud by a factor of up to  $\sim 30$ .

### 3 Architecture

The model consists of a generator and a critic in a GAN setup. Both generator and critic accept a variable number of points as input and are built to be equivariant respective invariant under permutation of the input, as architectures that respect the underlying symmetries are a good treatment for the curse of dimensionality when optimising a function over a high-dimensional space [21]. To aggregate information, a mean-field is introduced and initialized as the mean of all points. The interaction of all particles with each other is approximated by the interaction with this mean-field point only, allowing for scalability to large point clouds. The mean-field aggregates information via cross-attention to the points. The attention-based aggregation enables the critic and generator to dynamically select important points and disregard unimportant ones<sup>1</sup> The minimal building block for the generator as well as the critic is depicted in Figure 1. Multiple such blocks are stacked to build the main body of the architecture. Prior to the first block, a point-wise linear map is applied, mapping them to an  $l$ -dimensional latent space, in which further operations take place. The difference between the generator and the critic lies in the final layers of the architecture.

For the critic, a fully-connected neural network with one hidden layer is applied to the final mean-field to yield the score of the given input, whereas the generator maps the points down to the input dimension with a single point-wise layer. Different methods of transferring information from the mean-field to the points were experimented with. The simplest technique is adding the mean-field to all other points. While this approach works to a certain extent, issues arise as this aligns all points with the mean-field making the calculation of the dot product attention problematic. A point-wise layer, that takes the points and the mean-field as input, comes at the cost of introducing more parameters, but its introduction is necessary for a competitive performance. Supplying the number of points in a shower explicitly to the model is crucial for the performance of the generator and critic. The model is implicitly conditioned on the number of points in a cloud  $n$  by choosing the dimension of the input noise as  $n \times f$  (where  $f$  are the number of features per point). But since attention is a weighted sum, where the weights add up to unity, the mean-field is by construction agnostic to the number of elements during information exchange. To condition the showers with the incoming energy for the CaloChallenge dataset, a gated linear unit is utilised in the generator and critic.

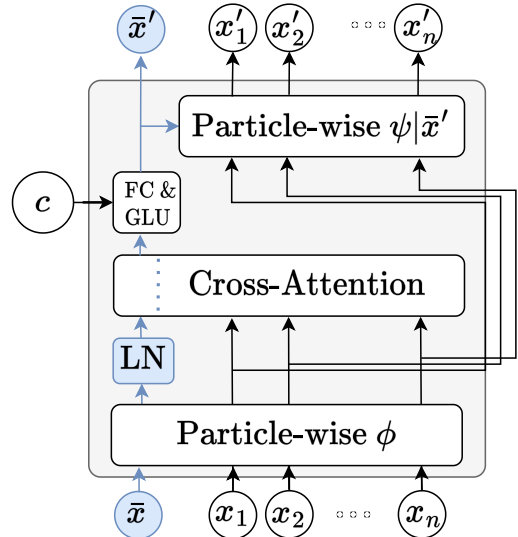


Figure 1: Minimal building block in the architecture comprises point-wise fully connected layers,  $c$  are the variables that are used for conditioning using a gated linear unit, LN is a layer norm and FC stands for a fully-connected layer which is applied to the mean-field. The cross-attention layer is the sole means of interaction between different points.

#### 3.1 Training and Evaluation

For the JETNET150 dataset, we found that the LSGAN [22] training performs best, while the Wasserstein GAN together with gradient penalty [23] performs best for the CaloChallenge dataset. Optimization of the network is done using Adam [24] with momentum set to zero. The learning rate is varied with Cosine Annealing after a linear warmup to  $10^{-4}$ , which is commonly used with attention-based models [25]. The critic for the CaloChallenge uses spectral normalisation [26]<sup>2</sup>, while the one for JETNET150 uses weight normalisation [28].

<sup>1</sup>Note that although attention is more commonly known from sequential NLP tasks, it is actually a permutation equivariant operation. It is therefore well suited to unordered sets such as point clouds.

<sup>2</sup>Note that this is done to enforce the Lipschitz constraint - however, it is shown in [27] that the usual dot-product attention is not Lipschitz and they propose an L2-based attention, which is also applied in this work.




Figure 2: Comparison between the JETNET150 dataset and generated samples. The first 3 plots shows the distribution of the marginal features ( $\eta^{rel}$ ,  $\phi^{rel}$ ,  $p_T^{rel}$ ) of all particles. The plot on the right depicts the relative invariant jet mass which provides a difficult to model high-level correlation. A ratio between the generated data and the ground truth is shown below every plot.

### 3.1.1 Deep Mean-field Matching

To stabilise the generator training, we introduce a deep mean-field matching loss. In this approach the generator additionally optimizes the mean-squared error between the mean-field in the last layer of the critic for real and fake inputs. A similar loss is proposed by Goodfellow et al. [29], but they propose to minimize the squared difference of the mean. We found that this generally stabilises the training, however, the loss sometimes diverges later during training. Consequently, we limit the application of this loss for 50, 000 gradient steps.

## 4 Results & Discussion

In the following we summarize the most important aspects of the model, which we refer to as Matching Deep Mean-field Attentive (MDMA) GAN.

### 4.1 JETNET150

During evaluation a kernel density estimate (KDE) is sampled to provide the condition of the number particles in a jet. The KDE is fit to the particle multiplicity distribution present in the training set. In Table 1, the results are compared to the ones for EPiC-GAN. The in-sample metrics (IN) calculated between the training and testing sets are given for reference as well. To obtain the results from EPiC-GAN, the checkpoint provided in their repository was used. The MDMA-GAN model outperforms EPiC-GAN on every metric, albeit our generator/critic only has 111 000/55 100 parameters which is about a quarter compared to EPiC-GAN. Figure 2 depicts the marginal distributions and the relative invariant jet mass distribution for the true and generated data.

Table 1: Comparison between the model from Ref. [16] (EPiC) and our model (MDMA) on the top-quark dataset. Bold font has been used to indicate scores where one model performs significantly better. The in-sample distance between the training and testing sample (IN) is also given for reference. Our model performs better on all metrics that were evaluated. The uncertainty on the metrics is given as the standard distance by bootstrapping the distances five times with a sample size of 50'000.

| Jet Class | Model | $W_1^M(\times 10^3)$            | $W_1^P(\times 10^3)$            | $W_1^{EFP}(\times 10^5)$        | $KPD(\times 10^4)$                 | $FPD(\times 10^4)$              |
|-----------|-------|---------------------------------|---------------------------------|---------------------------------|------------------------------------|---------------------------------|
|           | MDMA  | <b><math>0.4 \pm 0.1</math></b> | <b><math>2.5 \pm 0.1</math></b> | <b><math>2.0 \pm 0.3</math></b> | <b><math>-0.08 \pm 0.07</math></b> | <b><math>3.0 \pm 0.4</math></b> |
| Top Quark | EPiC  | $0.60 \pm 0.07$                 | $3.79 \pm 0.09$                 | $2.9 \pm 0.3$                   | $2 \pm 1$                          | $23 \pm 1$                      |
|           | IN    | $0.3 \pm 0.1$                   | $0.32 \pm 0.09$                 | $1.0 \pm 0.3$                   | $-0.1 \pm 0.1$                     | $0.2 \pm 0.2$                   |

### 4.2 CaloChallenge

Similarly, as for the JETNET150 dataset, the marginal distributions of all hits are shown in Figure 3. Since ultimately the voxel-based representation of the data is needed, the shown results need to be considered with caution. When remapping the points to the voxel structure, multiple generated hits are often mapped back to the same cell as they are too close, which highlights a clear limitation of




Figure 3: Distribution of the energy and position for MC simulated showers from the CaloChallenge dataset and for the ML generated showers.

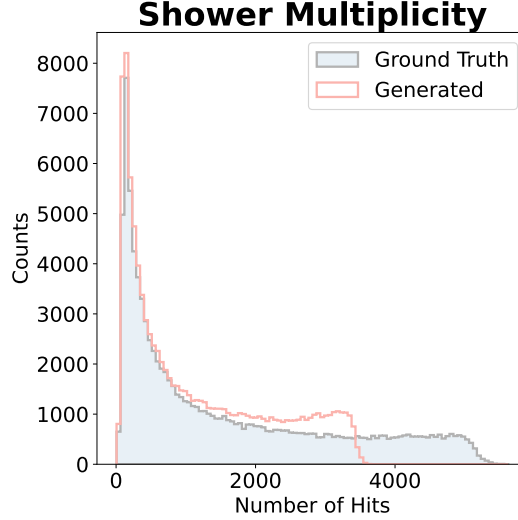


Figure 4: Limitation of point-cloud-based model: When transforming back to the voxel-based representation, numerous hits get mapped to the same cell leading to an incorrect distribution of the number of hits per shower.

the model. This occurs predictably often for clouds containing a high number of points, as can be seen in Figure 4.

## 5 Conclusion

We presented and evaluated a generative model for point clouds on two datasets. The model scales linearly with the number of points/tokens, thanks to the use of cross-attention to an artificial mean-field. To ensure stable adversarial training, the generator minimizes an L2 loss between mean-fields deep in the critic and the mean-fields from ground truth data. On the JETNET150 top-quark dataset, the proposed model exhibits state-of-the-art performance using significantly fewer parameters than another state-of-the-art model with similar performance.

Differences between in-sample distances and generated samples are nevertheless still observed. Furthermore, the model is deployed on the CaloChallenge dataset 2 where the point clouds yield up to  $\sim 30$  more points to simulate. However, in this setting the model seems to be unable to model the spatial distribution of points correctly. It remains to be studied whether scaling up the model in terms of model size and amount of training data used resolves these remaining difficulties.

## References

- [1] HEP ML Community. A Living Review of Machine Learning for Particle Physics, 2023. URL <https://iml-wg.github.io/HEPML-LivingReview/>.
- [2] CMS Collaboration. CMS offline and computing public results, 2023. URL <https://twiki.cern.ch/twiki/bin/view/CMSPublic/CMSOfflineComputingResults>.
- [3] G Apollinari, O Brüning, T Nakamoto, and L Rossi. High luminosity large hadron collider hl-lhc, 2015. URL <https://cds.cern.ch/record/2120673>.
- [4] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022.
- [5] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks, 2019.
- [6] Raghav Kansal, Javier Duarte, Hao Su, Breno Orzari, Thiago Tomei, Maurizio Pierini, Mary Touranakou, Jean-Roch Vlimant, and Dimitrios Gunopulos. Particle Cloud Generation with Message Passing Generative Adversarial Networks. In *35th Conference on Neural Information Processing Systems*, 6 2021.
- [7] Torbjörn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands. An introduction to PYTHIA8.2. *Comput. Phys. Commun.*, 191:159, jun 2015. doi:10.1016/j.cpc.2015.01.024. URL <https://doi.org/10.1016%2Fj.cpc.2015.01.024>.
- [8] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The anti- $k_T$  jet clustering algorithm. *JHEP*, 04(04):063, apr 2008. doi:10.1088/1126-6708/2008/04/063. URL <https://doi.org/10.1088%2F1126-6708%2F2008%2F04%2F063>.
- [9] Raghav Kansal, Anni Li, Javier Duarte, Nadezda Chernyavskaya, Maurizio Pierini, Breno Orzari, and Thiago Tomei. On the evaluation of generative models in high energy physics, 2022.
- [10] Michele Fucci Giannelli, Gregor Kasieczka, Claudius Krause, Ben Nachman, Dalila Salamani, David Shih, and Anna Zaborowska. Fast Calorimeter Simulation Challenge 2022 - Dataset 2, March 2022. URL <https://doi.org/10.5281/zenodo.6366271>.
- [11] Benno Käch, Dirk Krücker, Isabell Melzer-Pellmann, Moritz Scham, Simon Schnake, and Alexi Verney-Provatas. Jetflow: Generating jets with conditioned and mass constrained normalising flows, 2022.
- [12] Matthew Leigh, Debajyoti Sengupta, John Andrew Raine, Guillaume Quétant, and Tobias Golling. Pc-droid: Faster diffusion and improved quality for particle cloud generation, 2023.
- [13] Benno Käch, Dirk Krücker, and Isabell Melzer-Pellmann. Point cloud generation using transformer encoders and normalising flows, 2022.
- [14] Vinicius Mikuni, Benjamin Nachman, and Mariel Pettee. Fast point cloud generation with diffusion models in high energy physics, 2023.
- [15] 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, 2023.
- [16] Erik Buhmann, Gregor Kasieczka, and Jesse Thaler. Epic-gan: Equivariant point cloud generation for particle jets, 2023.
- [17] Raghav Kansal et al. Jetnet150, Aug 2022. URL <https://doi.org/10.5281/zenodo.6975117>.
- [18] S. Agostinelli et al. GEANT4—a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250–303, 2003. doi:10.1016/S0168-9002(03)01368-8.
- [19] Vinicius Mikuni and Benjamin Nachman. Caloscore v2: Single-shot calorimeter shower simulation with diffusion models, 2023.

- [20] Oz Amram and Kevin Pedro. Calodiffusion with glam for high fidelity calorimeter simulation, 2023.
- [21] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, 2021.
- [22] Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, and Zhen Wang. Multi-class generative adversarial networks with the L2 loss function. 2016.
- [23] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of Wasserstein GANs. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, page 5767. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7159-improved-training-of-wasserstein-gans.pdf>.
- [24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015*, San Diego, 2015. ICLR.
- [25] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017. URL <http://arxiv.org/abs/1706.03762>.
- [26] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *6th International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=B1QRgziT->.
- [27] Hyunjik Kim, George Papamakarios, and Andriy Mnih. The lipschitz constant of self-attention, 2021.
- [28] Sitao Xiang and Hao Li. On the effects of batch and weight normalization in generative adversarial networks, 2017.
- [29] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans, 2016.