

# OmniJet- $\alpha_C$ : Learning point cloud calorimeter simulations using generative transformers

Joschka Birk,<sup>1,\*</sup> Frank Gaede,<sup>2,†</sup> Anna Hallin,<sup>1,‡</sup>  
Gregor Kasieczka,<sup>1,§</sup> Martina Mozzanica,<sup>1,¶</sup> and Henning Rose<sup>1,\*\*</sup>

<sup>1</sup>*Institute for Experimental Physics, Universität Hamburg  
Luruper Chaussee 149, 22761 Hamburg, Germany*

<sup>2</sup>*Deutsches Elektronen-Synchrotron DESY,  
Notkestr. 85, 22607 Hamburg, Germany*

We show the first use of generative transformers for generating calorimeter showers as point clouds in a high-granularity calorimeter. Using the tokenizer and generative part of the OMNIJET- $\alpha$  model, we represent the hits in the detector as sequences of integers. This model allows variable-length sequences, which means that it supports realistic shower development and does not need to be conditioned on the number of hits. Since the tokenization represents the showers as point clouds, the model learns the geometry of the showers without being restricted to any particular voxel grid.

## I. INTRODUCTION

Machine learning (ML) methods have been a common ingredient in particle physics research for a long time, with neural networks being applied to object identification already in analyses at LEP [1]. Since then, the range of applications has grown drastically, with ML methods being developed and used for example in tagging [2–4], anomaly detection [5–8], individual reconstruction stages like particle tracking [9–11] or even full event interpretation and reconstruction [12]. Another important use case for ML in high energy physics (HEP) is detector simulation. With the increasing luminosity of the large-scale experiments in HEP, the computational cost of high-precision Monte-Carlo (MC) simulations is going to exceed the available computing resources [13]. Generative methods have the potential to significantly reduce this resource requirement, which is why a considerable amount of research has been spent on exploring architectures for detector simulation [14, 15]. Examples include GANs [16–27], variational autoencoders (VAEs) and their variants [28–33], normalizing flows and various types of diffusion models [33–51].

Most ML methods in HEP are designed, developed and trained for very specific tasks. The focus on specialized models means that the full potential of the vast datasets we have access to is not being utilized. Furthermore, while these models may be more resource efficient than the traditional methods they seek to enhance or replace, developing and training each model from scratch still requires significant amounts of both human and computational resources. For reasons like these, there has been a growing interest in developing foundation models for

particle physics [52–60] in the past couple of years. A foundation model is a machine learning model that has been pre-trained on a large amount of data, and can then be fine-tuned for different downstream tasks [61]. The idea behind utilizing pre-trained models is that their outputs can significantly enhance the performance of downstream tasks, yielding better results than if the model were to be trained from scratch. While the models mentioned above have focused on exploring different tasks in specific subdomains, like jet physics, a more ambitious goal eventually would be to develop a foundation model for all tasks in all subdomains, including for example tracking, shower generation and anomaly detection in general (not restricted to jets). The hope would be that it could then utilize the full amount of diverse data from our experiments, to boost the performance of all possible downstream tasks. The first step towards such a model must be to be able to handle tasks from different subdomains in the same computational framework.

In this work, we apply the generative part of OMNIJET- $\alpha$  [55], originally developed for jet physics, to a completely different subdomain: electromagnetic shower generation in collider calorimeters. We show that the OMNIJET- $\alpha$  architecture and workflow also works for generating showers, opening up the possibility of exploring transfer learning for showers in a setting that has already proved successful in the context of jet physics. This is the first example of an autoregressive generative model utilizing the GPT architecture for calorimeter point clouds (as opposed to the fixed calorimeter geometries of Ref. [62]). We denote this extended model capable of handling showers as OMNIJET- $\alpha_C$  (OMNIJET- $\alpha$  Calorimeter). Showing that we can use the same framework for two very different subdomains is an important step towards developing a foundation model for all computing and data analysis tasks in particle physics.

This paper is organized as follows. Section II describes the dataset used, section III the experimental setup, and section IV presents the results. Finally, we offer our conclusions in section V.

---

\* [joschka.birk@uni-hamburg.de](mailto:joschka.birk@uni-hamburg.de)

† [frank.gaede@desy.de](mailto:frank.gaede@desy.de)

‡ [anna.hallin@uni-hamburg.de](mailto:anna.hallin@uni-hamburg.de)

§ [gregor.kasieczka@uni-hamburg.de](mailto:gregor.kasieczka@uni-hamburg.de)

¶ [martina.mozzanica@uni-hamburg.de](mailto:martina.mozzanica@uni-hamburg.de)

\*\* [henning.rose@studium.uni-hamburg.de](mailto:henning.rose@studium.uni-hamburg.de)

## II. DATASET

The International Large Detector (ILD) [63] is one of two detector concepts proposed for the International Linear Collider (ILC) [64], an electron-positron collider that is initially operated at 250 GeV center-of-mass energy and extendable to higher energies up to 1 TeV. ILD is optimized for the Particle Flow Algorithm [65] that aims at reconstructing every individual particle. The detector therefore combines precise tracking and vertexing capabilities with good hermiticity, and highly granular sandwich calorimeters. The electromagnetic calorimeter of ILD (the Si-W ECAL [66]) consists of 20 layers with 2.1 mm thick W-absorbers followed by 10 layers with 4.2 mm W-absorbers, all interleaved with 0.5 mm thick Si-sensors that are subdivided into  $5 \text{ mm} \times 5 \text{ mm}$  cells.

The dataset used in this work was originally created for Ref. [67], where more details on the detector and simulation can be found. Showers of photons with initial energies uniformly distributed between 10 – 100 GeV are simulated with Geant4 [68] using a detailed and realistic detector model implemented in DD4hep [69]. The resulting showers are projected into a regular 3D grid with  $30 \times 30 \times 30 = 27\,000$  voxels. The 3D-grid data is converted into a point cloud format, where each point has four features: the  $x$ - and  $y$ -position (transverse to the incident particle direction), the  $z$ -position (parallel to the incident particle direction), and the energy. The incoming photon enters the calorimeter at perpendicular incident angle from the bottom at  $z = 0$  and traverses along the  $z$ -axis, hitting cells in the center of the  $x$ - $y$  plane. A staggered cell geometry results in small shifts between the layers.

We preprocess the four input features ( $x$ ,  $y$ ,  $z$  and energy) by standardization. The energy feature is log-transformed before being scaled and shifted, which has the additional advantage that generated energies are by design non-negative.

The dataset has 950 000 samples, of which 760 000 are used for training, 95 000 for validation, and 95 000 as test samples.

## III. METHODS

This work uses the workflow of OMNIJET- $\alpha$  [55], which is a foundation model originally developed for jet physics. OMNIJET- $\alpha$  uses a VQ-VAE [54, 70–72] to tokenize the input features. The constituents of the jets, or in this case the voxel hits of the showers, are represented as a sequence of integers. These sequences are used as input for the generative model, which is a GPT-style [73] model. Since the model only expects integers, it is not dependent on a specific type of data as input as long as it can be represented in this format. Moreover, the model accepts variable-length sequences, which means that it can be used equally well for jets with a variable number of constituents as for showers with a variable num-

ber of hits. The training target of the model is next token prediction, that is, it learns the probability of each token given a sequence of previous tokens,  $p(x_i|x_{i-1}, \dots, x_0)$ . This means that it is straightforward to use the trained model for autoregressive generation, where each new token is generated conditioned on the previous ones in the sequence. While OMNIJET- $\alpha$  also has classification capabilities, this work only focuses on the generative part. One key feature of OMNIJET- $\alpha$  is that it learns the sequence length from context. This removes the need for specifying the number of elements in the sequence beforehand.

The VQ-VAE and generative model were trained using the hyperparameters described in Appendix A. For the VQ-VAE, the best epoch was selected via lowest validation loss. After training, the VQ-VAE was frozen. The input data was tokenized using this model, and then fed into the generative model for training. Here again the epoch with the lowest validation loss was chosen as the best epoch. New showers in the form of integer sequences were then generated using this final generative model, and the frozen VQ-VAE was used to decode these integer sequences back into physical space.

## IV. RESULTS

In the following we will present the results of the training of the VQ-VAE and the generative model. For comparison we use the test dataset, which the models never saw during training. As a benchmark for shower generation the performance of OMNIJET- $\alpha_C$  is compared to two state-of-the-art generative networks: one point cloud model, CALOCLOUDS II [74], and one fixed-grid model, L2LFLAWS [75]. CALOCLOUDS II is a continuous time score-based diffusion model that has been further distilled into a consistency model (CM), whereas L2LFLAWS is a flow-based model using coupling flows with convolutional layers. L2LFLAWS has already been trained on this dataset in [75], and the showers were provided to us directly by the authors. For CALOCLOUDS II however, no such training was available. Instead we ran this training ourselves, using the same hyperparameters as in [74] with the exception of training the diffusion model for 3.5 M iterations instead of 2 M, and the consistency model for 2.5 M iterations instead of 1 M. This is the first time CALOCLOUDS II has been trained on a dataset in which the granularity matches the one of the calorimeter.

### A. Token quality

We first investigate the encoding and decoding capabilities of the VQ-VAE. To judge the effect of the tokenization and potential loss of information, we compare the original showers with the reconstructed showers on hit-level. A perfect reconstruc-

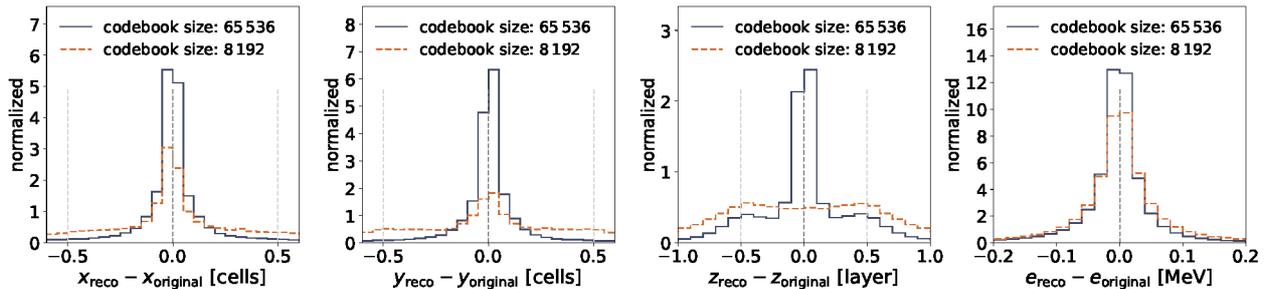


Figure 1: Reconstruction resolution for the input features ( $x$ ,  $y$ ,  $z$ , energy) for different codebook sizes.

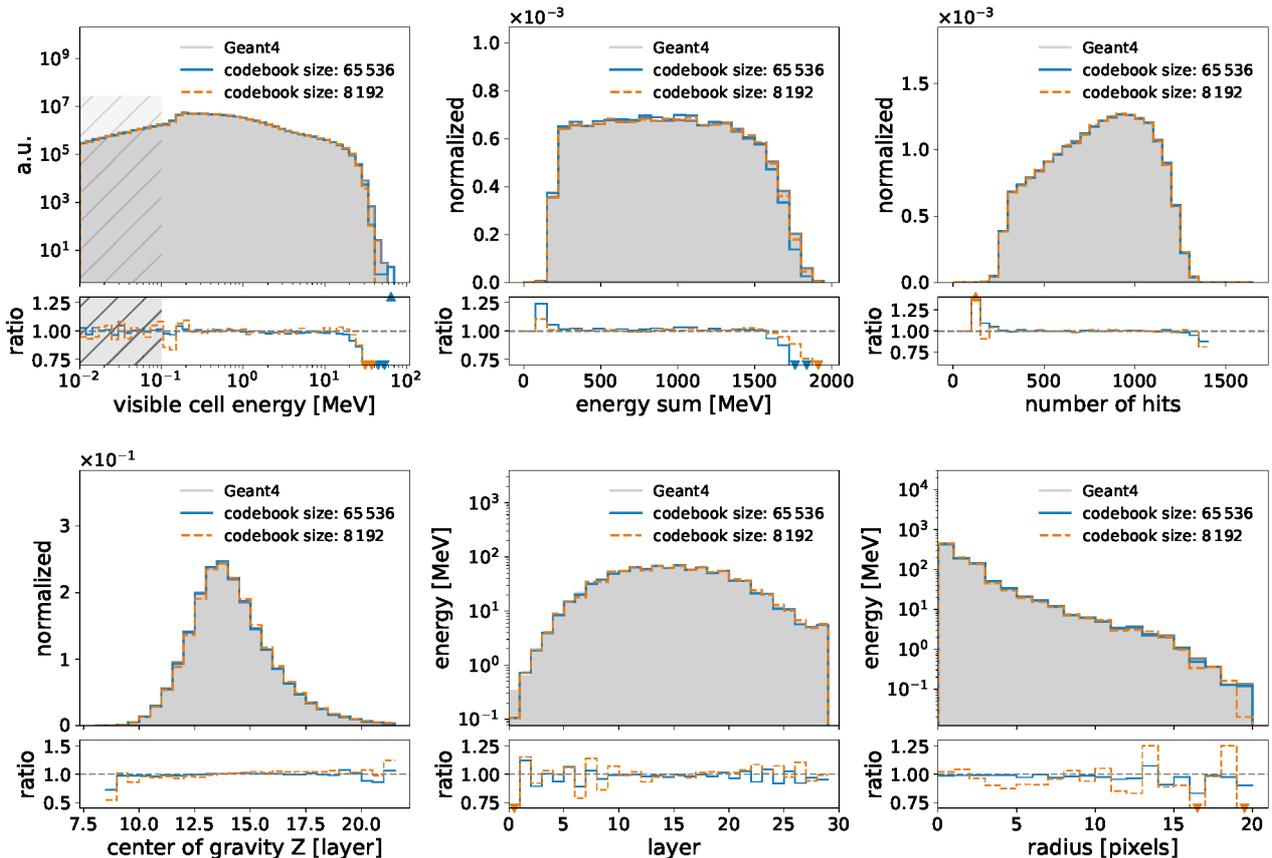


Figure 2: Distributions of physical observables between Geant4 (grey, filled) with the codebook size of 65 536 (blue) and codebook size of 8 192 (orange). Hits that were below the MIP threshold (0.1 MeV), i.e. those in the shaded region of the visible cell energy plot, were not considered for the comparison in the remaining distributions. This cutoff can affect the number of hits for reconstructed showers.

tion would yield a Dirac delta function for the difference between reconstructed and original values for each feature. However, as shown in Fig. 1, while the distributions surrounding the center are indeed narrow, they do have some spread. A codebook size of 65 536 shows a narrower resolution distribution than a codebook size of 8 192. In particular, the reconstruction of  $z$  for the latter has a larger spread of  $\sigma_{8192}^z = 0.66$  layers compared to  $\sigma_{65536}^z = 0.4$  layers with the larger codebook size. For the energy, the respective spread values are  $\sigma_{8192}^{\text{energy}} = 0.11$  MeV and  $\sigma_{65536}^{\text{energy}} = 0.07$  MeV. Furthermore, the reconstructed  $z$  distribution demonstrates a broader spread and a more complex reconstruction relative to the transverse coordinates  $x$  and

$y$ , which exhibit similar and narrower distributions. This difference in reconstruction accuracy can be attributed to a broader spatial extent of the showers along the longitudinal axis  $z$ . However, because voxels are discrete, the three spatial features need to be rounded to integers. Perfect resolution is achieved if these values remain within  $\pm 0.5$  before rounding, the region indicated by the light gray lines in Fig. 1.

To accurately compare the reconstructed showers with the original showers on hit and shower level, we need to apply postprocessing. This step is explained in Appendix B and essentially projects hits back into the voxel grid and processes duplicate hits (hits that are identical in all of the three spatial features).

The quality of the tokenization is also evaluated

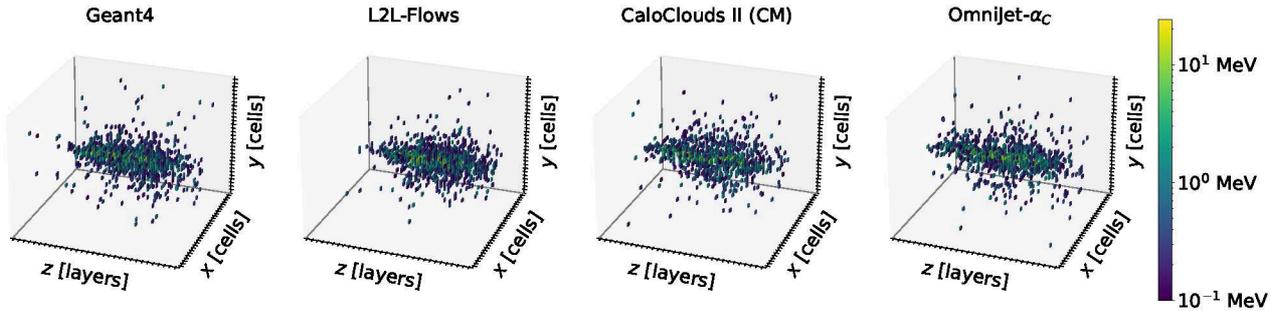


Figure 3: Examples of individual photon showers with a total energy sum of 1000 MeV generated by Geant4 (left), L2L-Flows (center left), the CALOCLOUDS II (CM) (center right) and OMNIJET- $\alpha_C$  (right).

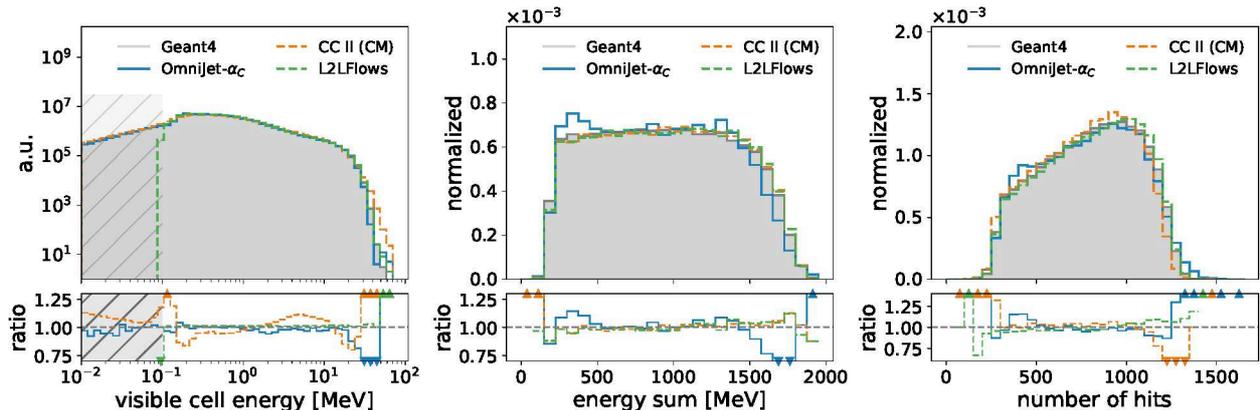


Figure 4: Distributions of per-cell energy (left), total energy sum (middle) and the number of hits above 0.1 MeV (right) between Geant4 (grey, filled) and the generative models: OMNIJET- $\alpha_C$  (blue), the CALOCLOUDS II (CM) (orange, dashed) and L2L-Flows (green, dashed).

on hit- and shower-level. For this analysis, showers are converted to tokens and then back to physical space. Fig. 2 shows different feature distributions. Generally we observe good agreement with the original distributions. Rare tokens, such as those located at the edges of the shower or tokens associated with high-energy hits, exhibit the lowest reconstruction quality. Again the VQ-VAE with the larger codebook size performs better and has the smallest loss of information.

## B. Shower generation

Following training, OMNIJET- $\alpha_C$  generates point clouds autoregressively. Initialized with a start token (a special token that initiates the autoregressive generation process), the model predicts the probability distribution for the next token based on the preceding sequence. OMNIJET- $\alpha_C$  then samples from this distribution, appending the chosen token to the growing sequence. This process continues until a stop token (a special token that represents the end of the generated sequence) is generated or the maximum sequence length of 1700 tokens is reached. Unlike most ML-based shower generators, OMNIJET- $\alpha_C$  is not trained to generate showers for specific incident photon energies. Instead, the model learns to generate showers with a variety of energies.

We reserve a study of how to condition the model on the incident energy for future work. This would allow the user to request showers of a specific energy. In this first version however, we will only compare the full spectrum of showers.

We see in Fig. 3 that OMNIJET- $\alpha_C$ , CALOCLOUDS II (CM) and L2L-Flows generate showers that appear to be visually acceptable compared to Geant4. Next, we compare the performance of OMNIJET- $\alpha_C$  to CALOCLOUDS II (CM) and L2L-Flows for three different quantities<sup>1</sup>.

Fig. 4 (left) compares the cell energies. We observe an accurate performance of OMNIJET- $\alpha_C$  across almost the entire energy range, on par with L2L-Flows. For the higher energies we see some deviations for both OMNIJET- $\alpha_C$  and CALOCLOUDS II (CM). As seen in Fig. 2, the mis modeling for OMNIJET- $\alpha_C$  is introduced by the VQ-VAE. The behavior of CALOCLOUDS II (CM) is consistent with what was seen in the original paper. The shaded area in the histogram corresponds to the region below half the energy of a minimal ionizing particle (MIP). In real detectors, read-outs at such small energies are dominated by noise. Therefore, cell energies below

<sup>1</sup> Note that compared to the original training of CALOCLOUDS II in Ref. [74], this training is done at physical, i.e. lower, resolution.

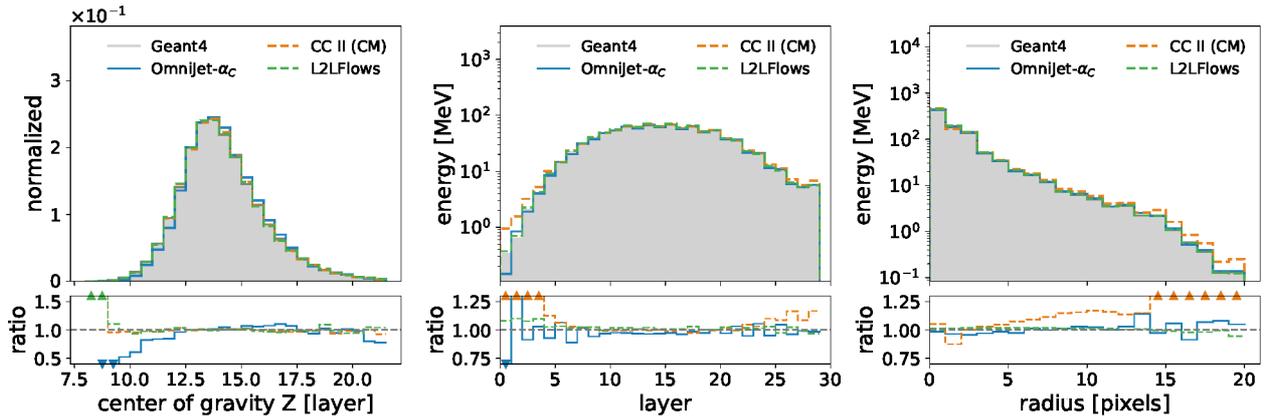


Figure 5: Distributions of center of gravity (left), mean energy per layer (middle) and the mean energy per layer between Geant4 (grey, filled) and the generative models: OMNIJET- $\alpha_C$  (blue), the CALOCLOUDS II (CM) (orange, dashed) and L2L-Flows (green, dashed).

0.1 MeV will not be considered in the following discussion, and the remaining plots and distributions only include cells above this cut-off.

Fig. 4 (center) shows the distribution of the total energy sum of showers. For this calculation, the energy of all hits surpassing half the MIP energy are added up for each shower. This distribution is strongly correlated to the incident photon energy on which L2LFlows and CALOCLOUDS II (CM) are conditioned. OMNIJET- $\alpha_C$  has to learn this distribution on its own.

Finally, Fig. 4 (right) shows the number of hits. While the L2L-Flows and CALOCLOUDS II (CM) are conditioned on this distribution, OMNIJET- $\alpha_C$  is able to achieve good agreement with the Geant4 distribution without this conditioning. The discrepancies we see are a small peak at a shower length of around 400 to 500, and also some showers that are too long.

In Fig. 5 we compare the spatial properties of the shower. The left plot shows that the Geant4 distribution of the center of gravity along the  $z$ -axis is well modeled by all three architectures. OMNIJET- $\alpha_C$  performs better in the center of the peak than at the edges.

The longitudinal energy distribution, depicted in the middle plot of Fig. 5, reveals a comparatively weaker performance of the OMNIJET- $\alpha_C$  model and CALOCLOUDS II (CM) compared to L2LFlows in the initial 10 layers. However, OMNIJET- $\alpha_C$  outperforms CALOCLOUDS II (CM) in the first 4 layers. The mismodeling of OMNIJET- $\alpha_C$  in the initial layers is likely attributable to the tokenization process (see Fig. 2), where these layers, being less common, are represented by a limited number of tokens. A similar degradation is observed in the outer regions of the radial energy distribution (right plot of Fig. 5), although OMNIJET- $\alpha_C$  still outperforms CALOCLOUDS II (CM).

Another important aspect for comparing generative models is the single-shower generation time. Generating 1000 showers, randomly sampled across

all incident energies, resulted in a mean and standard deviation of  $(3 \pm 1)$ s per shower. The generation was performed with a batch size of 2 on an NVIDIA<sup>®</sup> A100 GPU. In contrast, Geant4 on a CPU required  $(4.1 \pm 0.2)$ s per shower [28]. Therefore, our model demonstrates a speedup factor of 1.39 in this case. On identical hardware and with a batch size of 1000, L2LFlows achieves per-shower generation times of  $(3.24 \pm 0.05)$ ms and a speedup factor of 1260. CALOCLOUDS II on identical hardware but with a batch size of 100 generates one shower in  $(16 \pm 6)$ ms and achieves a speedup factor of 255. The comparatively slow performance of OMNIJET- $\alpha_C$  is attributable to the generation being autoregressive. Since this study did not prioritize generation speed, optimizations such as multi-token generation are left for future work.

## V. CONCLUSION

In this work, we take a first important step towards building a foundation model for several subdomains of particle physics. We show that we are able to use the architecture and workflow of a foundation model originally developed for jet physics to generate electromagnetic showers in a calorimeter, a fundamentally different problem. This is a notable difference to previous efforts for foundation models in HEP, which so far focused on tasks within one subdomain, mostly different tasks within jet physics. It is also the first implementation of a GPT-style autoregressive generative model for calorimeter shower point cloud generation.

The next immediate step will be to investigate whether this model can be used for transfer learning between different types of showers. In the long term, we aim to develop a joint model that can work with both jets and showers. Combining tasks from different subdomains in one single framework is a necessary step towards a foundation model for particle physics that can handle a variety of data types and tasks.

## ACKNOWLEDGEMENTS

The authors would like to thank William Korcari for support with the dataset, as well as Thorsten Buss for providing the L2LFlows samples.

JB, AH, GK, MM and HR are supported by the DFG under the German Excellence Initiative – EXC 2121 Quantum Universe – 390833306, and

by PUNCH4NFID – project number 460248186. This work has used the Maxwell computational resources at Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany.

## CODE AVAILABILITY

The code for this paper can be found at [https://github.com/uhh-pd-ml/omnijet\\_alpha.c](https://github.com/uhh-pd-ml/omnijet_alpha.c).

- 
- [1] T. Behnke and David G. Charlton, “Electroweak measurements using heavy quarks at LEP,” *Phys. Scripta* **52**, 133–157 (1995).
- [2] Maxence Draguet (ATLAS), *Flavour Tagging with Graph Neural Network at ATLAS*, Tech. Rep. (CERN, Geneva, 2024).
- [3] Maja Karwowska, Lukasz Graczykowski, Kamil Deja, Miłosz Kasak, and Małgorzata Janik (ALICE), “Particle identification with machine learning from incomplete data in the ALICE experiment,” *JINST* **19**, C07013 (2024), arXiv:2403.17436 [hep-ex].
- [4] Spandan Mondal and Luca Mastrolorenzo, “Machine learning in high energy physics: a review of heavy-flavor jet tagging at the LHC,” *Eur. Phys. J. ST* **233**, 2657–2686 (2024), arXiv:2404.01071 [hep-ex].
- [5] ATLAS collaboration, “Dijet resonance search with weak supervision using  $\sqrt{s} = 13$  TeV  $pp$  collisions in the ATLAS detector,” *Phys. Rev. Lett.* **125**, 131801 (2020), arXiv:2005.02983 [hep-ex].
- [6] ATLAS collaboration, “Anomaly detection search for new resonances decaying into a Higgs boson and a generic new particle  $X$  in hadronic final states using  $\sqrt{s} = 13$  TeV  $pp$  collisions with the ATLAS detector,” *Phys. Rev. D* **108**, 052009 (2023), arXiv:2306.03637 [hep-ex].
- [7] ATLAS collaboration, “Search for New Phenomena in Two-Body Invariant Mass Distributions Using Unsupervised Machine Learning for Anomaly Detection at  $\sqrt{s}=13$  TeV with the ATLAS Detector,” *Phys. Rev. Lett.* **132**, 081801 (2024), arXiv:2307.01612 [hep-ex].
- [8] CMS Collaboration, “Model-agnostic search for dijet resonances with anomalous jet substructure in proton-proton collisions at  $\sqrt{s} = 13$  TeV,” (2024), arXiv:2412.03747 [hep-ex].
- [9] Jared Dynes Bursleson, Sylvain Caillou, Paolo Calafiura, Jay Chan, Christophe Collard, Xiangyang Ju, Daniel Thomas Murnane, Mark Neubauer, Minh Tuan Pham, Charline Rougier, Jan Stark, Heberth Torres, and Alexis Vallier (ATLAS), *Physics Performance of the ATLAS GNN4ITk Track Reconstruction Chain*, Tech. Rep. (CERN, Geneva, 2023).
- [10] ATLAS collaboration, *Computational Performance of the ATLAS ITk GNN Track Reconstruction Pipeline*, Tech. Rep. (CERN, Geneva, 2024).
- [11] Anthony Correia, Fotis Giasemis, Nabil Garroum, Vladimir Vava Gligorov, and Bertrand Granado, “Graph Neural Network-Based Pipeline for Track Finding in the Velo at LHCb,” in *Connecting The Dots 2023* (2024) arXiv:2406.12869 [physics.ins-det].
- [12] Julián García Pardinás, Marta Calvi, Jonas Eschle, Andrea Mauri, Simone Meloni, Martina Mozzanica, and Nicola Serra, “GNN for Deep Full Event Interpretation and Hierarchical Reconstruction of Heavy-Hadron Decays in Proton–Proton Collisions,” *Comput. Softw. Big Sci.* **7**, 12 (2023), arXiv:2304.08610 [hep-ex].
- [13] Andreas Adelman *et al.*, “New directions for surrogate models and differentiable programming for High Energy Physics detector simulation,” in *Snowmass 2021* (2022) arXiv:2203.08806 [hep-ph].
- [14] Simon Badger *et al.*, “Machine learning and LHC event generation,” *SciPost Phys.* **14**, 079 (2023), arXiv:2203.07460 [hep-ph].
- [15] Claudius Krause, Michele Fucci Giannelli, Gregor Kasieczka, Benjamin Nachman, Dalila Salamani, David Shih, Anna Zaborowska, Oz Amram, Kerstin Borrás, Matthew R. Buckley, Erik Buhmann, Thorsten Buss, Renato Paulo Da Costa Cardoso, Anthony L. Caterini, Nadezda Chernyavskaya, Federico A. G. Corchia, Jesse C. Cresswell, Sascha Diefenbacher, Etienne Dreyer, Vijay Ekambaram, Engin Eren, Florian Ernst, Luigi Favaro, Matteo Franchini, Frank Gaede, Eilam Gross, Shih-Chieh Hsu, Kristina Jaruskova, Benno Käch, Jayant Kalaganam, Raghav Kansal, Taewoo Kim, Dmitrii Kobylanski, Anatolii Korol, William Korcari, Dirk Krücker, Katja Krüger, Marco Letizia, Shu Li, Qibin Liu, Xiulong Liu, Gabriel Loaiza-Ganem, Thandikire Madula, Peter McKeown, Isabell-A. Melzer-Pellmann, Vinicius Mikuni, Nam Nguyen, Ayodele Ore, Sofia Palacios Schweitzer, Ian Pang, Kevin Pedro, Tilman Plehn, Witold Pokorski, Huilin Qu, Piyush Raikwar, John A. Raine, Humberto Reyes-Gonzalez, Lorenzo Rinaldi, Brendan Leigh Ross, Moritz A. W. Scham, Simon Schnake, Chase Shimmin, Eli Shlizerman, Nathalie Soybelman, Mudhakar Srivatsa, Kalliopi Tsolaki, Sofia Vallecorsa, Kyongmin Yeo, and Rui Zhang, “Calochallenge 2022: A community challenge for fast calorimeter simulation,” (2024), arXiv:2410.21611 [cs.LG].
- [16] Michela Paganini, Luke de Oliveira, and Benjamin Nachman, “Accelerating Science with Generative Adversarial Networks: An Application to 3D Particle Showers in Multilayer Calorimeters,” *Phys. Rev. Lett.* **120**, 042003 (2018), arXiv:1705.02355 [hep-ex].
- [17] Michela Paganini, Luke de Oliveira, and Benjamin Nachman, “CaloGAN : Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks,” *Phys. Rev. D* **97**, 014021 (2018), arXiv:1712.10321 [hep-ex].

- [18] Luke de Oliveira, Michela Paganini, and Benjamin Nachman, “Controlling Physical Attributes in GAN-Accelerated Simulation of Electromagnetic Calorimeters,” *J. Phys. Conf. Ser.* **1085**, 042017 (2018), arXiv:1711.08813 [hep-ex].
- [19] Martin Erdmann, Lukas Geiger, Jonas Glombitza, and David Schmidt, “Generating and refining particle detector simulations using the Wasserstein distance in adversarial networks,” *Comput. Softw. Big Sci.* **2**, 4 (2018), arXiv:1802.03325 [astro-ph.IM].
- [20] Pasquale Musella and Francesco Pandolfi, “Fast and Accurate Simulation of Particle Detectors Using Generative Adversarial Networks,” *Comput. Softw. Big Sci.* **2**, 8 (2018), arXiv:1805.00850 [hep-ex].
- [21] Martin Erdmann, Jonas Glombitza, and Thorben Quast, “Precise simulation of electromagnetic calorimeter showers using a wasserstein generative adversarial network,” *Computing and Software for Big Science* **3** (2019), 10.1007/s41781-018-0019-7.
- [22] Dawit Belayneh, Federico Carminati, Amir Farbin, Benjamin Hooberman, Gulrukh Khattak, Miaoyuan Liu, Junze Liu, Dominick Olivito, Vitória Barin Pacela, Maurizio Pierini, Alexander Schwing, Maria Spiropulu, Sofia Vallecorsa, Jean-Roch Vlimant, Wei Wei, and Matt Zhang, “Calorimetry with deep learning: particle simulation and reconstruction for collider physics,” *The European Physical Journal C* **80** (2020), 10.1140/epjc/s10052-020-8251-9.
- [23] Anja Butter, Sascha Diefenbacher, Gregor Kasieczka, Benjamin Nachman, and Tilman Plehn, “GANplifying event samples,” *SciPost Phys.* **10**, 139 (2021), arXiv:2008.06545 [hep-ph].
- [24] Martina Javurkova (ATLAS), “The Fast Simulation Chain in the ATLAS experiment,” *EPJ Web Conf.* **251**, 03012 (2021).
- [25] Sebastian Bieringer, Anja Butter, Sascha Diefenbacher, Engin Eren, Frank Gaede, Daniel Hundhausen, Gregor Kasieczka, Benjamin Nachman, Tilman Plehn, and Mathias Trabs, “Calomplification — the power of generative calorimeter models,” *JINST* **17**, P09028 (2022), arXiv:2202.07352 [hep-ph].
- [26] G. et all Aad, “Deep generative models for fast photon shower simulation in atlas,” *Computing and Software for Big Science* **8** (2024), 10.1007/s41781-023-00106-9.
- [27] Baran Hashemi, Nikolai Hartmann, Sahand Sharifzadeh, James Kahn, and Thomas Kuhr, “Ultra-high-granularity detector simulation with intra-event aware generative adversarial network and self-supervised relational reasoning,” *Nature Communications* **15** (2024), 10.1038/s41467-024-49104-4.
- [28] Erik Buhmann, Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Anatolii Korol, and Katja Krüger, “Getting high: High fidelity simulation of high granularity calorimeters with high speed,” *Computing and Software for Big Science* **5** (2021), 10.1007/s41781-021-00056-0.
- [29] Erik Buhmann, Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Anatolii Korol, and Katja Krüger, “Decoding photons: Physics in the latent space of a bib-ae generative network,” *EPJ Web of Conferences* **251**, 03003 (2021).
- [30] Erik Buhmann, Sascha Diefenbacher, Daniel Hundhausen, Gregor Kasieczka, William Korcari, Engin Eren, Frank Gaede, Katja Krüger, Peter McKeown, and Lennart Rustige, “Hadrons, better, faster, stronger,” *Mach. Learn. Sci. Tech.* **3**, 025014 (2022), arXiv:2112.09709 [physics.ins-det].
- [31] G. Aad, “Deep generative models for fast photon shower simulation in atlas,” *Computing and Software for Big Science* **8** (2024), 10.1007/s41781-023-00106-9.
- [32] Jesse C. Cresswell, Brendan Leigh Ross, Gabriel Loaiza-Ganem, Humberto Reyes-Gonzalez, Marco Letizia, and Anthony L. Caterini, “CaloMan: Fast generation of calorimeter showers with density estimation on learned manifolds,” in *36th Conference on Neural Information Processing Systems: Workshop on Machine Learning and the Physical Sciences* (2022) arXiv:2211.15380 [hep-ph].
- [33] Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Anatolii Korol, Katja Krüger, Peter McKeown, and Lennart Rustige, “New angles on fast calorimeter shower simulation,” *Mach. Learn. Sci. Tech.* **4**, 035044 (2023), arXiv:2303.18150 [physics.ins-det].
- [34] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” (2015), arXiv:1503.03585 [cs.LG].
- [35] Yang Song and Stefano Ermon, “Generative modeling by estimating gradients of the data distribution,” (2020), arXiv:1907.05600 [cs.LG].
- [36] Yang Song and Stefano Ermon, “Improved Techniques for Training Score-Based Generative Models,” (2020), arXiv:2006.09011 [cs.LG].
- [37] Jonathan Ho, Ajay Jain, and Pieter Abbeel, “Denoising diffusion probabilistic models,” (2020), arXiv:2006.11239 [cs.LG].
- [38] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole, “Score-based generative modeling through stochastic differential equations,” (2021), arXiv:2011.13456 [cs.LG].
- [39] Vinicius Mikuni and Benjamin Nachman, “Score-based generative models for calorimeter shower simulation,” *Phys. Rev. D* **106**, 092009 (2022), arXiv:2206.11898 [hep-ph].
- [40] 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,” (2023), arXiv:2305.04847 [physics.ins-det].
- [41] 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,” (2023), arXiv:2307.04780 [cs.LG].
- [42] Vinicius Mikuni and Benjamin Nachman, “CaloScore v2: Single-shot Calorimeter Shower Simulation with Diffusion Models,” (2023), arXiv:2308.03847 [hep-ph].
- [43] Oz Amram and Kevin Pedro, “CaloDiffusion with GLaM for High Fidelity Calorimeter Simulation,” (2023), arXiv:2308.03876 [physics.ins-det].
- [44] C. Chen, O. Cerri, T. Q. Nguyen, J. R. Vlimant, and M. Pierini, “Analysis-Specific Fast Simulation at the LHC with Deep Learning,” *Computing and Software for Big Science* **5**, 15 (2021).
- [45] Claudius Krause and David Shih, “Fast and accurate simulations of calorimeter showers with normalizing flows,” *Phys. Rev. D* **107**, 113003 (2023), arXiv:2106.05285 [physics.ins-det].

- [46] Claudius Krause and David Shih, “Accelerating accurate simulations of calorimeter showers with normalizing flows and probability density distillation,” *Phys. Rev. D* **107**, 113004 (2023), [arXiv:2110.11377 \[physics.ins-det\]](#).
- [47] Simon Schnake, Dirk Krücker, and Kerstin Borrás, “Generating Calorimeter Showers as Point Clouds,” (2022).
- [48] Claudius Krause, Ian Pang, and David Shih, “CaloFlow for CaloChallenge Dataset 1,” (2023), [arXiv:2210.14245 \[physics.ins-det\]](#).
- [49] Allison Xu, Shuo Han, Xiangyang Ju, and Haichen Wang, “Generative Machine Learning for Detector Response Modeling with a Conditional Normalizing Flow,” (2023), [arXiv:2303.10148 \[hep-ex\]](#).
- [50] Matthew R. Buckley, Claudius Krause, Ian Pang, and David Shih, “Inductive CaloFlow,” (2023), [arXiv:2305.11934 \[physics.ins-det\]](#).
- [51] Manjunath Omana Kuttan, Kai Zhou, Jan Steinheimer, and Horst Stöcker, “Towards a foundation model for heavy-ion collision experiments through point cloud diffusion,” (2024), [arXiv:2412.10352 \[hep-ph\]](#).
- [52] Tomoe Kishimoto, Masahiro Morinaga, Masahiko Saito, and Junichi Tanaka, “Pre-training strategy using real particle collision data for event classification in collider physics,” in *37th Conference on Neural Information Processing Systems* (2023) [arXiv:2312.06909 \[hep-ex\]](#).
- [53] Huilin Qu, Congqiao Li, and Sitian Qian, “Particle Transformer for Jet Tagging,” (2022), [arXiv:2202.03772 \[hep-ph\]](#).
- [54] Tobias Golling, Lukas Heinrich, Michael Kagan, Samuel Klein, Matthew Leigh, Margarita Osadchy, and John Andrew Raine, “Masked particle modeling on sets: towards self-supervised high energy physics foundation models,” *Mach. Learn. Sci. Tech.* **5**, 035074 (2024), [arXiv:2401.13537 \[hep-ph\]](#).
- [55] Joschka Birk, Anna Hallin, and Gregor Kasieczka, “OmniJet- $\alpha$ : the first cross-task foundation model for particle physics,” *Mach. Learn. Sci. Tech.* **5**, 035031 (2024), [arXiv:2403.05618 \[hep-ph\]](#).
- [56] Philip Harris, Michael Kagan, Jeffrey Krupa, Benedikt Maier, and Nathaniel Woodward, “Re-Simulation-based Self-Supervised Learning for Pre-Training Foundation Models,” (2024), [arXiv:2403.07066 \[hep-ph\]](#).
- [57] Vinicius Mikuni and Benjamin Nachman, “OmniLearn: A Method to Simultaneously Facilitate All Jet Physics Tasks,” (2024), [arXiv:2404.16091 \[hep-ph\]](#).
- [58] Andrew J. Wildridge, Jack P. Rodgers, Ethan M. Colbert, Yao yao, Andreas W. Jung, and Miaoyuan Liu, “Bumblebee: Foundation Model for Particle Physics Discovery,” in *38th conference on Neural Information Processing Systems* (2024) [arXiv:2412.07867 \[hep-ex\]](#).
- [59] Oz Amram, Luca Anzalone, Joschka Birk, Darius A. Faroughy, Anna Hallin, Gregor Kasieczka, Michael Krämer, Ian Pang, Humberto Reyes-Gonzalez, and David Shih, “Aspen Open Jets: Unlocking LHC Data for Foundation Models in Particle Physics,” (2024), [arXiv:2412.10504 \[hep-ph\]](#).
- [60] Joshua Ho, Benjamin Ryan Roberts, Shuo Han, and Haichen Wang, “Pretrained Event Classification Model for High Energy Physics Analysis,” (2024), [arXiv:2412.10665 \[hep-ph\]](#).
- [61] Rishi Bommasani *et al.*, “On the opportunities and risks of foundation models,” (2022), [arXiv:2108.07258 \[cs.LG\]](#).
- [62] Qibin Liu, Chase Shimmin, Xiulong Liu, Eli Shlizerman, Shu Li, and Shih-Chieh Hsu, “Calo-VQ: Vector-Quantized Two-Stage Generative Model in Calorimeter Simulation,” (2024), [arXiv:2405.06605 \[physics.ins-det\]](#).
- [63] Halina Abramowicz *et al.* (ILD Concept Group), “International Large Detector: Interim Design Report,” (2020), [arXiv:2003.01116 \[physics.ins-det\]](#).
- [64] “The International Linear Collider Technical Design Report - Volume 1: Executive Summary,” (2013), [arXiv:1306.6327 \[physics.acc-ph\]](#).
- [65] Mark A. Thomson, “Particle flow calorimetry,” *J. Phys. Conf. Ser.* **293**, 012021 (2011).
- [66] T. Suehara *et al.*, “Performance study of SKIROC2/A ASIC for ILD Si-W ECAL,” *JINST* **13**, C03015 (2018), [arXiv:1801.02024 \[physics.ins-det\]](#).
- [67] Erik Buhmann, Sascha Diefenbacher, Engin Eren, Frank Gaede, Gregor Kasieczka, Anatolii Korol, and Katja Krüger, “Getting High: High Fidelity Simulation of High Granularity Calorimeters with High Speed,” *Comput. Softw. Big Sci.* **5**, 13 (2021), [arXiv:2005.05334 \[physics.ins-det\]](#).
- [68] S. Agostinelli *et al.* (GEANT4), “GEANT4: A simulation toolkit,” *Nucl. Instrum. Meth.* **A506**, 250–303 (2003).
- [69] Markus Frank, F. Gaede, C. Grefe, and P. Mato, “DD4hep: A Detector Description Toolkit for High Energy Physics Experiments,” *J. Phys. Conf. Ser.* **513**, 022010 (2014).
- [70] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu, “Neural discrete representation learning,” (2018), [arXiv:1711.00937 \[cs.LG\]](#).
- [71] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei, “BEiT: BERT Pre-Training of Image Transformers,” (2022), [arXiv:2106.08254 \[cs.CV\]](#).
- [72] Minyoung Huh, Brian Cheung, Pulkit Agrawal, and Phillip Isola, “Straightening out the straight-through estimator: Overcoming optimization challenges in vector quantized networks,” (2023), [arXiv:2305.08842 \[cs.LG\]](#).
- [73] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever, “Improving language understanding by generative pre-training,” (2018).
- [74] 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,” *JINST* **19**, P04020 (2024), [arXiv:2309.05704 \[physics.ins-det\]](#).
- [75] Thorsten Buss, Frank Gaede, Gregor Kasieczka, Claudius Krause, and David Shih, “Convolutional L2LFlows: generating accurate showers in highly granular calorimeters using convolutional normalizing flows,” *Journal of Instrumentation* **19**, P09003 (2024).
- [76] Michael R. Zhang, James Lucas, Geoffrey Hinton, and Jimmy Ba, “Lookahead optimizer: k steps forward, 1 step back,” (2019), [arXiv:1907.08610 \[cs.LG\]](#).
- [77] Hongwei Yong, Jianqiang Huang, Xiansheng Hua, and Lei Zhang, “Gradient centralization: A new optimization technique for deep neural networks,” (2020), [arXiv:2004.01461 \[cs.CV\]](#).

### Appendix A: Model details and hyperparameters

Different hyperparameter configurations were tested for the individual model components of OMNIJET- $\alpha_C$ . The configurations presented in the following were found to lead to stable trainings. However, no extensive hyperparameter optimization was performed.

TABLE I: Hyperparameters used in the VQ-VAE training.

Hyperparameter	Value
Learning rate	0.001
Optimizer	Ranger
Batch size	152
Batches per epoch	1000
Number of epochs	588
Hidden dimension	128
Codebook size	65 536
$\beta$	0.8
$\alpha$	10
Replacement frequency	100

The hyperparameters used for the VQ-VAE training are shown in Tab. I. Only the codebook size, replacement frequency and the hyperparameter  $\beta$  were adjusted. The remaining hyperparameters are the same as in OMNIJET- $\alpha$ . An increase of the codebook size from 8 192 to 65 536 was found to improve the reconstruction capabilities (i.e. the resolution of the tokenized showers). The codebook utilization, i.e. the fraction of used tokens, is also monitored during the training to ensure that the resulting codebook is used completely. Unused tokens would drastically increase the number of parameters of the generative model while not adding any potential improvements in the performance of the generative model. The current setup results in a codebook utilization of the final VQ-VAE model of 99.65%. The hyperparameter  $\beta$  which defines the relative importance of how much weight should be given to updating the encoder embeddings  $z_e$  towards the codebook vectors

$z_q$  and vice versa, is decreased from 0.9 to 0.8. This leads to a higher emphasis on adapting the encoder to bring the embeddings  $z_e$  closer to the codebook vectors  $z_q$ . Furthermore, the optimization process employs a token replacement strategy based on usage frequency. The chosen replacement frequency of 100 batches (instead of 10) indicates that a token must be used at least once within the preceding 100 batches to avoid being replaced by a new token. We used the Lookahead optimizer [76] with RAdam as the inner optimizer [77].

For the hyperparameters of the backbone, no changes compared to OMNIJET- $\alpha$  were made except for the batch size. The hyperparameters used are listed in Tab. II.

TABLE II: Hyperparameters used in the generative model training.

Hyperparameter	Value
Learning rate	0.001
Optimizer	Ranger
Batch size	72
Batches per epoch	6000
Number of epochs	106
Embedding dimension	256
Number of heads	8
Number of GPT blocks	3

### Appendix B: Postprocessing

Projecting the hits of a point cloud model back onto the voxel grid can result in duplicate hits in some voxels. To resolve these duplicates, the voxels with lower energy are translated along the  $z$ -axis to the nearest unoccupied voxel position. This heuristic preserves both the total energy and the hit count while minimally impacting the  $z$ -distribution. We could also translate the voxels along the  $x$ - or  $y$ -axis, but as shown in Fig. 6 the hit energies are not invariant in these directions.

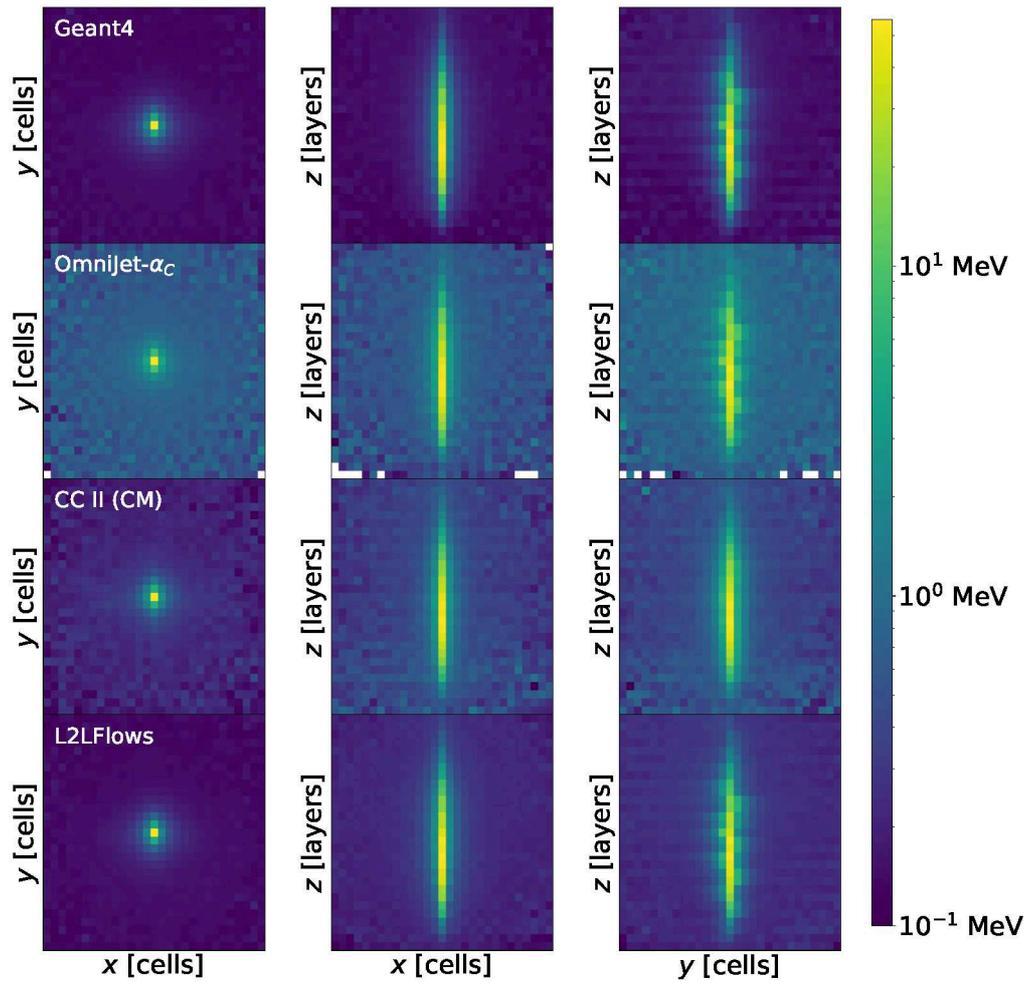


Figure 6: Overlay of 10k showers for all simulators for the full spectrum, where the voxel energies are summed along the  $z$ - (left),  $y$ - (middle) and  $x$ -axis (right). In all plots, the mean over the number of showers is taken.