Optimizing the IFMIF-DONES Particle Accelerator with Differentiable Deep Learning Surrogate Models

Galo Gallardo Romero[1] * † **Guillermo Rodríguez Llorente**[1,2,3] * †
Lucas Magariños Rodríguez[1] **Rodrigo Morant Navascués**[1]
Nikita Khvatkin Petrovsky[1] **Roberto Gómez-Espinosa Martín**[1] †
[1]HI Iberia [2]University Carlos III of Madrid [3]Gregorio Millán Barbany Institute
{ggallardo, grodriguez, robertogemartin}@hi-iberia.es

Abstract

In this work, Deep Learning Surrogate Models are employed to optimize the quadrupole values in the initial section of the High Energy Beam Transport Line of the IFMIF-DONES accelerator. Two Fourier Neural Operator models were trained: one for predicting two-dimensional beam profiles and another for forecasting one-dimensional beam statistics along the accelerator's longitudinal axis. These models offer up to 3 orders of magnitude speedup compared to traditional simulations, with a trade-off of maintaining accuracy within percentage errors below 6%. Moreover, their differentiability allows seamless integration with optimization algorithms, enabling efficient tuning of quadrupole values to achieve specific beam objectives. This approach offers a robust solution for enhancing the performance of IFMIF-DONES accelerator and other scientific experiments.

1 Introduction

One of the key challenges in advancing nuclear fusion power plants is understanding the effects of neutron irradiation on reactor materials [1]. To tackle this, the IFMIF-DONES (International Fusion Materials Irradiation Facility DEMO Oriented Neutron Source) facility is developing a linear accelerator that will produce a neutron spectrum resembling that of fusion reactions, achieved through the collision of deuterons with a lithium target [2]. This research, conducted as part of the DONES-FLUX project, explores the application of differentiable Deep Learning Surrogate Models (DLSMs) in order to optimize the quadrupoles of the first section of the High Energy Beam Transport (HEBT-S1) [3] in the IFMIF-DONES accelerator. DLSMs are deep learning-based models designed to approximate complex physical, chemical, or biological processes, providing computationally efficient substitutes for resource-intensive simulations [4]. Their differentiable nature enables their integration with optimization algorithms, such as Gradient Descent (GD) [5], to identify parameter sets that satisfy specific beam objectives. Neural operators (NOs) are a branch of DLSMs that extend neural network architectures to approximate operators mapping between infinite-dimensional function spaces [6]. This extension is particularly valuable for solving systems of partial differential equations, which model physical phenomena. In these NO models, input functions $a(x)$ are processed to predict output functions $u(x)$, as illustrated in Appendix A. The architecture includes a layer P that transforms the input function into its first hidden representation. This is then processed through a series of kernels with non-linear activation functions, culminating in the output representation generated by layer Q. Notably, the Fourier Neural Operator (FNO), which is the chosen architecture for this work, leverages the Fourier transform to compute these kernels [7]. This FNO architecture is selected for three key reasons: it is purely data-driven, it is discretization invariant (unlike standard neural networks), and it

*Joint first authorship
†Corresponding authors

D3S3: Data-driven and Differentiable Simulations, Surrogates, and Solvers @ NeurIPS 2024.

has demonstrated success across various fields with positive results [8, 9, 10, 11, 12]. The trained models are able to successfully predict deuteron beam functions along the longitude of the HEBT-S1, which is displayed in the first panel of Fig. 1. For details on the methodology, refer to Section 2, while the results are shown in Section 3 and the discussion and outlook are presented in Section 4.

2 Methodology

The workflow, outlined in Fig. 1, begins by constructing a lattice with HEBT-S1 measurements from [3] (Appendix F). Following this, beam simulations were performed with OPAL [13]. Data was generated from 8192 simulations of deuteron bunches with 100k particles[3] traveling through the lattice, with varying normal quadrupole magnetic field values k_j ranging from 4 to 8 Tm^{-1}. Given the six quadrupoles in the system, the distances and signs of each element were fixed to prevent an exponential increase in the number of possible combinations, and thus computational resources.

Two FNO models were trained with 80% of the dataset, leaving the rest for testing. The first model, a 2D FNO, predicts the normalized two-dimensional beam profile deuteron distributions $\rho_x(x, z)$ and $\rho_y(y, z)$, where z (in m) represents the longitudinal position along the accelerator. As the discretization of these two output functions is the same in all axes, they are treated as a general output function $\rho(x, z)$ with two channels, one for each transversal profile in x and y (in mm). The second model, a 1D FNO, predicts the root mean square and maximum beam envelope size for both transversal components ($\mathrm{RMS}_x(z)$, $\mathrm{RMS}_y(z)$, $\mathrm{MAX}_x(z)$, $\mathrm{MAX}_y(z)$), as well as the particle count $n(z)$ divided by the total number of particles. Finally, the parameter space was explored using the GD optimization algorithm Adam [14] to address the inverse problem, determining the optimal quadrupole settings that achieve desired beam configurations at the end of the lattice. This is only possible because the trained models are differentiable with respect to the input parameters, k_j. The loss function is described in C, and minimizes the distance between the target configuration and the current one. The frameworks used are NVIDIA Modulus [15] and Pytorch [16]. Diagrams for both models are displayed in Appendix A, while their hyperparameters are detailed in Appendix B.

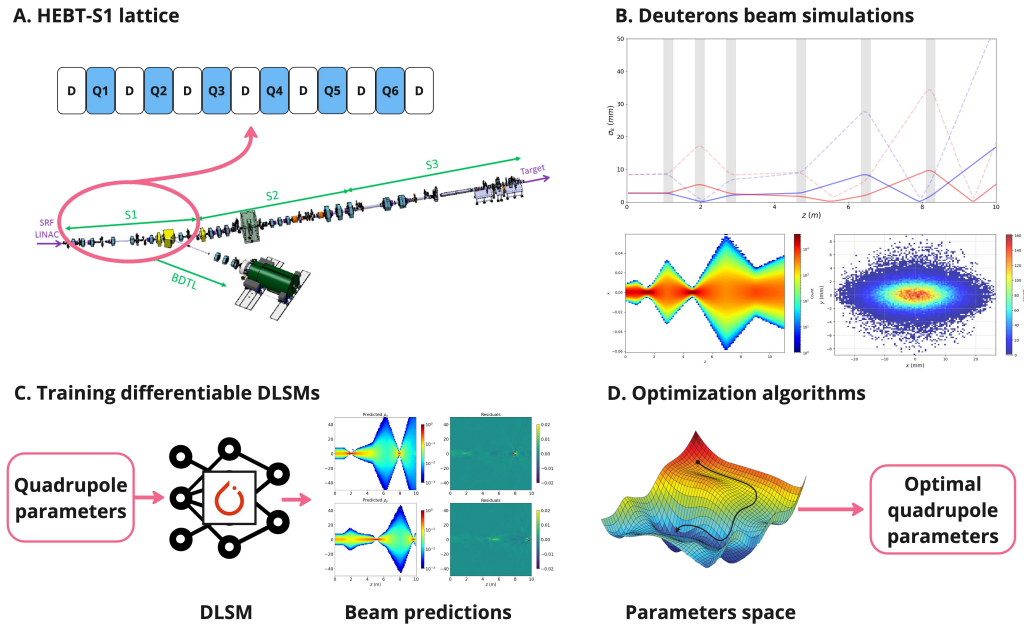

Figure 1: (A) IFMIF-DONES HEBT sections and diagram of the S1 lattice with drift zones and quadrupoles. (B) Deuteron beams simulation results from OPAL, with the beam statistics, profile and footprints. (C) Training DLSMs for fast inference. (D) Optimization algorithms applied to the differentiable models in order to solve inverse problems and obtain optimal parameters for design.

[3]These particles are injected as gaussian distributions with parameters seen in Table 5.

3 Results

An example of prediction results for the 1D model is displayed in Fig. 2, where the predicted beam statistics are compared to the simulated values. Overall, this 1D model has a maximum mean average percentage error (MAPE) of 5.63%. For the 2D model, Fig. 3, shows a test case comparing the predicted two-dimensional beam profile functions with simulated results. This model demonstrates a mean absolute error (MAE) of $4.7{\cdot}10^{-3}$ for $\rho_x(x, z)$ and $5.6{\cdot}10^{-3}$ for $\rho_y(y, z)$. The speed up factor is the same for both models, with inference times 10^3 times faster than OPAL simulations (Each simulation takes approximately 20 seconds to complete). A comprehensive summary of the error metrics (MAE and MAPE) for both models, covering all relevant variables, is provided in Appendix I, allowing for a detailed assessment of their performance in the test set. Furthermore, a few $n(z)$ test predictions can be seen in Appendix J, as well as other test examples that are available in Appendix G, offering a broader view of model performance under varied conditions.

Finally, Table 1 presents the results of the optimization process for five different examples, detailing the target and achieved configurations for RMS and MAX evaluated at $z = 10$ m for both axes. The table includes the optimized quadrupole settings, the time t in minutes required to reach a solution, and the normalized number of particles as reported by OPAL during the evaluation. A fixed point of $n(10)/10^5 = 0.98$ was used as optimization threshold for the losses, as can be seen in Appendix E, so the solutions never lose more than 2% of the particles. The optimization process effectively meets the predefined objectives in approximately 10 minutes, achieving minimal particle loss and ensuring optimal beam performance. These results highlight the efficiency and effectiveness of the proposed models, demonstrating their capability to deliver desired outcomes within practical time constraints.

4 Discussion

The trained FNO models demonstrate remarkable accuracy and speed in predicting beam statistics and profile distributions, being up to 10^3 times faster than traditional OPAL simulations while maintaining minimal error rates. This acceleration enables near-instantaneous inference of beam properties, which, combined with the differentiable nature of the models, makes them ideal for integration into optimization loops. These loops can be used to solve inverse problems, such as identifying the optimal set of quadrupole values needed to achieve a desired beam configuration.

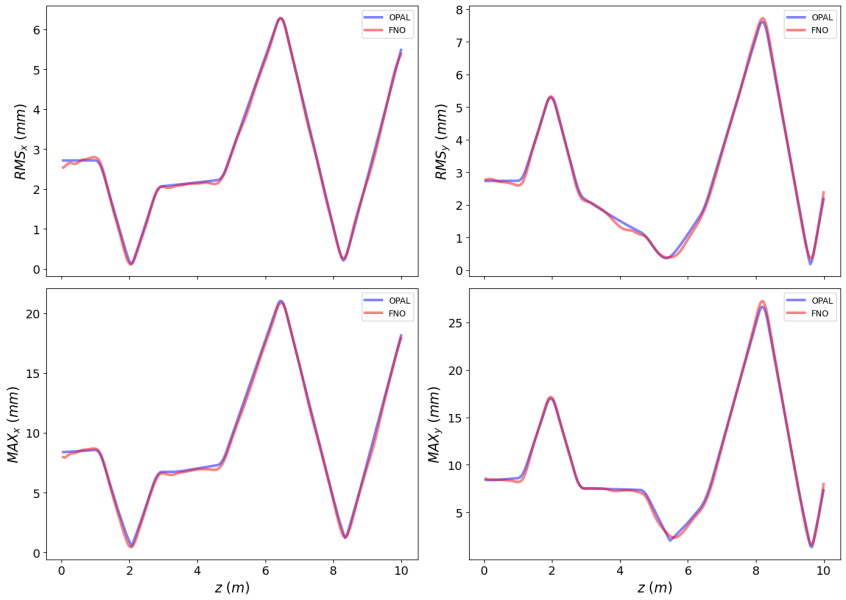

Figure 2: Test prediction results for the 1D beam statistics functions along the accelerator longitude for the axes x and y. Blue: OPAL simulations. Red: FNO predictions. This example was obtained by solving the inverse problem for the first objective in Table 1.

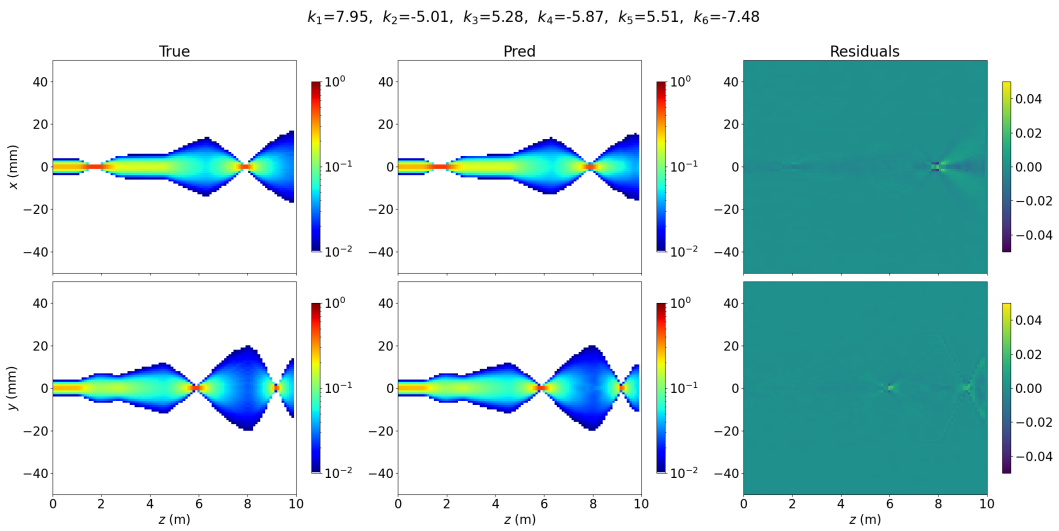

Figure 3: Test prediction results for the 2D beam profile distributions along the accelerator longitude for the axes x and y. Left: OPAL simulations. Middle: FNO predictions. Right: residual errors.

Furthermore, as highlighted in other studies, such as [17], the significantly reduced inference times allow DLSMs to serve as environments for Deep Reinforcement Learning agents [18], where the sheer number of training steps required would not be feasible with traditional simulators. These results highlight the powerful synergy between DLSMs and other deep learning techniques, offering new possibilities for enhancing optimization and control in complex scientific facilities, such as the IFMIF-DONES accelerator, that will benefit future nuclear fusion power plants. The next phase of this research will address the limitation of quadrupole range by tackling the exponential growth in computational resources needed as the number of magnetic elements increases. This is the reason why only the HEBT-S1 section has been optimized. Future work will expand these techniques to other parts of the accelerator and explore models with additional channels, incorporating particle momentum data to provide the optimization algorithm with information of the phase space, widely used in particle accelerators [19].

Table 1: Optimization summary. These are just a few varied examples, but the algorithm is able to find solutions, as long as they exist in the space of possible solutions. Some of the footprints achieved with these solutions can be seen in Appendix D.

Target (mm)				Achieved (mm)				k_j (Tm^{-1})						t (min)	$n(10)/10^5$
RMS_x	RMS_y	MAX_x	MAX_y	RMS_x	RMS_y	MAX_x	MAX_y	k1	k2	k3	k4	k5	k6		
5.4	2.4	18.0	8.0	5.4	2.4	17.9	8.0	5.9	-7.1	6.9	-5.2	5.0	-6.3	10.42	1.00
4.7	2.6	16.0	7.0	4.7	2.6	15.6	6.9	5.1	-5.7	7.6	-6.4	5.1	-6.0	3.42	0.99
2.5	3.4	9.0	12.0	2.6	3.1	8.9	12.0	5.0	-8.0	7.4	-7.7	5.0	-6.9	10.44	1.00
3.0	3.0	11.0	10.0	3.0	3.0	10.5	9.8	5.3	-6.7	7.9	-8.0	5.0	-6.0	1.52	0.99
6.4	8.5	22.0	30.0	6.3	8.5	21.5	27.0	5.4	-6.2	7.9	-6.3	5.4	-7.2	3.34	0.99
11.8	10.4	38.9	33.0	11.7	10.3	39.0	32.5	5.0	-6.3	7.9	-6.6	5.9	-7.5	10.45	0.99
16.5	1.0	50.0	3.4	16.5	1.0	49.8	3.3	6.5	-7.5	7.1	-5.6	6.5	-5.8	2.30	0.99
9.8	10.2	33.4	37.3	9.8	10.2	32.5	31.3	5.0	-6.0	7.6	-6.3	5.6	-7.4	3.7	0.99
7.0	6.6	23.5	19.5	7.0	6.0	23.2	19.5	6.4	-7.8	6.2	-5.0	5.0	-7.8	10.43	0.99
13.3	4.3	48.6	14.2	13.3	4.3	45.7	13.6	5.5	-6.5	7.7	-6.5	6.2	-6.4	1.48	0.99
7.7	11.6	25.9	29.9	7.7	11.6	25.9	29.4	6.0	-5.0	5.1	-6.0	5.0	-7.3	4.31	0.98
4.2	13.1	14.5	37.5	4.1	13.1	14.5	36.7	5.0	-5.5	5.0	-6.9	5.0	-7.6	10.47	0.99

Acknowledgments

We express our gratitude to CDTI for DONES-FLUX project funding, IFMIF-DONES and UGR.

References

[1] A. Ibarra, F. Arbeiter, D. Bernardi, W. Krolas, M. Cappelli, U. Fischer, R. Heidinger, F. Martin-Fuertes, G. Micciché, A. Muñoz, F.S. Nitti, T. Pinna, A. Aiello, N. Bazin, N. Chauvin, S. Chel, G. Devanz, S. Gordeev, D. Regidor, F. Schwab, and the full IFMIF-DONES team. The european approach to the fusion-like neutron source: the ifmif-dones project. *Nuclear Fusion*, 59(6):065002, may 2019.

[2] Davide Bernardi, Angel Ibarra, Frederik Arbeiter, F. Arranz, Mauro Cappelli, Philippe Cara, J. Castellanos, Hervé Dzitko, A. García, J. Gutiérrez, W. Królas, Francisco Martin-Fuertes, G. Micciché, Antonio Muñoz Roldan, F. Nitti, Tonio Pinna, Ivan Podadera, J. Pons, Yuefeng Qiu, and Raquel Román. The ifmif-dones project: Design status and main achievements within the eurofusion fp8 work programme. *Journal of Fusion Energy*, 41:24, 10 2022.

[3] Ivan Podadera, Angel Ibarra, David Jimenez Rey, Joaquin Molla, Concepcion Oliver, Rodrigo Varela, Victor Villamayor, Oriol Nomen, and Daniel Sanchez Herranz. Beam diagnostics for the multi-mw high energy beam transport line of dones. In *8th International Beam Instrumentation Conference*, 11 2019.

[4] Zhongkai Hao, Songming Liu, Yichi Zhang, Chengyang Ying, Yao Feng, Hang Su, and Jun Zhu. Physics-informed machine learning: A survey on problems, methods and applications, 2023.

[5] Jiawei Zhang. Gradient descent based optimization algorithms for deep learning models training. *CoRR*, abs/1903.03614, 2019.

[6] Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

[7] Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *CoRR*, abs/2010.08895, 2020.

[8] Tingtao Zhou, Xuan Wan, Daniel Zhengyu Huang, Zongyi Li, Zhiwei Peng, Anima Anandkumar, John F. Brady, Paul W. Sternberg, and Chiara Daraio. Ai-aided geometric design of anti-infection catheters, 2023.

[9] Vignesh Gopakumar, Stanislas Pamela, Lorenzo Zanisi, Zongyi Li, Anima Anandkumar, and MAST Team. Fourier neural operator for plasma modelling, 2023.

[10] Wenhui Peng, Shaoxiang Qin, Senwen Yang, Jianchun Wang, Xue Liu, and Liangzhu (Leon) Wang. Fourier neural operator for real-time simulation of 3d dynamic urban microclimate. *Building and Environment*, 248:111063, 2024.

[11] Thorsten Kurth, Shashank Subramanian, Peter Harrington, Jaideep Pathak, Morteza Mardani, David Hall, Andrea Miele, Karthik Kashinath, and Animashree Anandkumar. Fourcastnet: Accelerating global high-resolution weather forecasting using adaptive fourier neural operators, 2022.

[12] Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier neural operator with learned deformations for pdes on general geometries. *J. Mach. Learn. Res.*, 24(1), March 2024.

[13] Andreas Adelmann, Pedro Calvo, Matthias Frey, Achim Gsell, Uldis Locans, Christof Metzger-Kraus, Nicole Neveu, Chris Rogers, Steve Russell, Suzanne Sheehy, Jochem Snuverink, and Daniel Winklehner. Opal a versatile tool for charged particle accelerator simulations, 2019.

[14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.

[15] Modulus Contributors. Nvidia modulus: An open-source framework for physics-based deep learning in science and engineering, February 2023. Version 1.2.0.

[16] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. *CoRR*, abs/1912.01703, 2019.

[17] Guillermo Rodríguez-Llorente, Galo Romero, and Roberto G.E. Martin. Applications of fourier neural operators in the ifmif-dones accelerator. In *ICLR 2024 Workshop on AI4DifferentialEquations In Science*, 03 2024.

[18] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.

[19] C.A. Lindstrøm and M. Thévenet. Emittance preservation in advanced accelerators. *Journal of Instrumentation*, 17(05):P05016, May 2022.

A FNO architecture

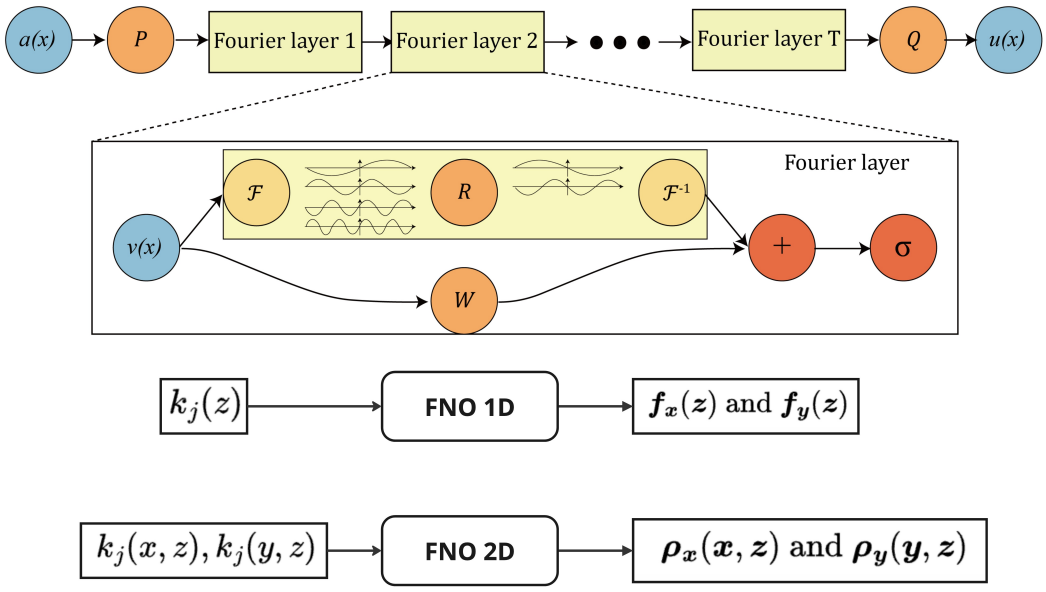

Figure 4: Top: FNO architecture from [7]. It processes input functions $a(x)$ to predict output functions $u(x)$, where x is the discretization space of the problem. The input is first transformed into a hidden representation by a layer P. This hidden representation is then passed through a series of layers where kernels are computed using the Fourier transform. After processing, the final layer Q produces the output function $u(x)$. The Fourier transform enables the FNO to efficiently compute these kernels, making it a key feature of this architecture.Bottom: inputs and outputs of the trained models. The quadrupole values k_j (where j stands for the quadrupole number, from 1 to 6) are taken as constants formatted as tensors with the dimensions of the output functions (with the same discretization). The functions $f(z)$ represent the beam statistics one-dimensional functions presented in the main text and the number of particles, while ρ are the 2D beam profile distributions for both axes x and y.

B Hyperparameters

Table 2: NVIDIA Modulus FNO architecture hyperparameters for both models. If there are two values, the first corresponds to the one dimensional model. The hyperparameters were chosen based in positive results from previous works.

Hyperparameter	Value
scheduler	tf_exponential_lr
optimizer	adam
loss	sum
decoder.nr_layers	1
decoder.layer_size	256-512
fno.dimension	1-2
fno.nr_fno_layers	4
fno.fno_modes	12
scheduler.decay_rate	0.95
scheduler.decay_steps	1000
training.max_steps	10000
batch_size.grid	32-8
batch_size.validation	32-8

Table 3: Optimization hyperparameters for GD.

Hyperparameter	Value
scheduler	exponential_lr
optimizer	adam
loss	custom distance
scheduler.gamma	0.999
max_steps	100000
early_stop_loss	0.005
learning_rate	1

C Loss function

The custom loss function is defined as the euclidean distance between the RMS_i target configuration (subindex T) and the current one, plus penalization terms that punish out of bound solutions. These are computed for the quadrupoles, for the MAX_i variables and for the number of particles n to establish a threshold. The described loss components are displayed below:

$$d(RMS_{T,i}, RMS_i) = \sqrt{\sum_{i=1}^{N} \left(RMS_{T,i} - RMS_i\right)^2} \tag{1}$$

$$f(x, u, l, w) = \begin{cases} (x - u) \times w & \text{if } x > u \\ (l - x) \times w & \text{if } x < l \\ 0 & \text{if } l \leq x \leq u \end{cases} \tag{2}$$

Where x describes k_j, MAX_i and n. The values l and u stand for the lower and upper bounds of each variable, respectively. The quadrupoles ranges are set between 4 and 8 Tm^{-1}, the MAX_i ranges are set between 0 and the target to reach and the n range is set from 0.98 to 1. The weights w are 10 for both MAX variables and 100 for the quadrupoles and n. The index i defines the x and y coordinates and j stands for the number of quadrupoles.

D Footprints

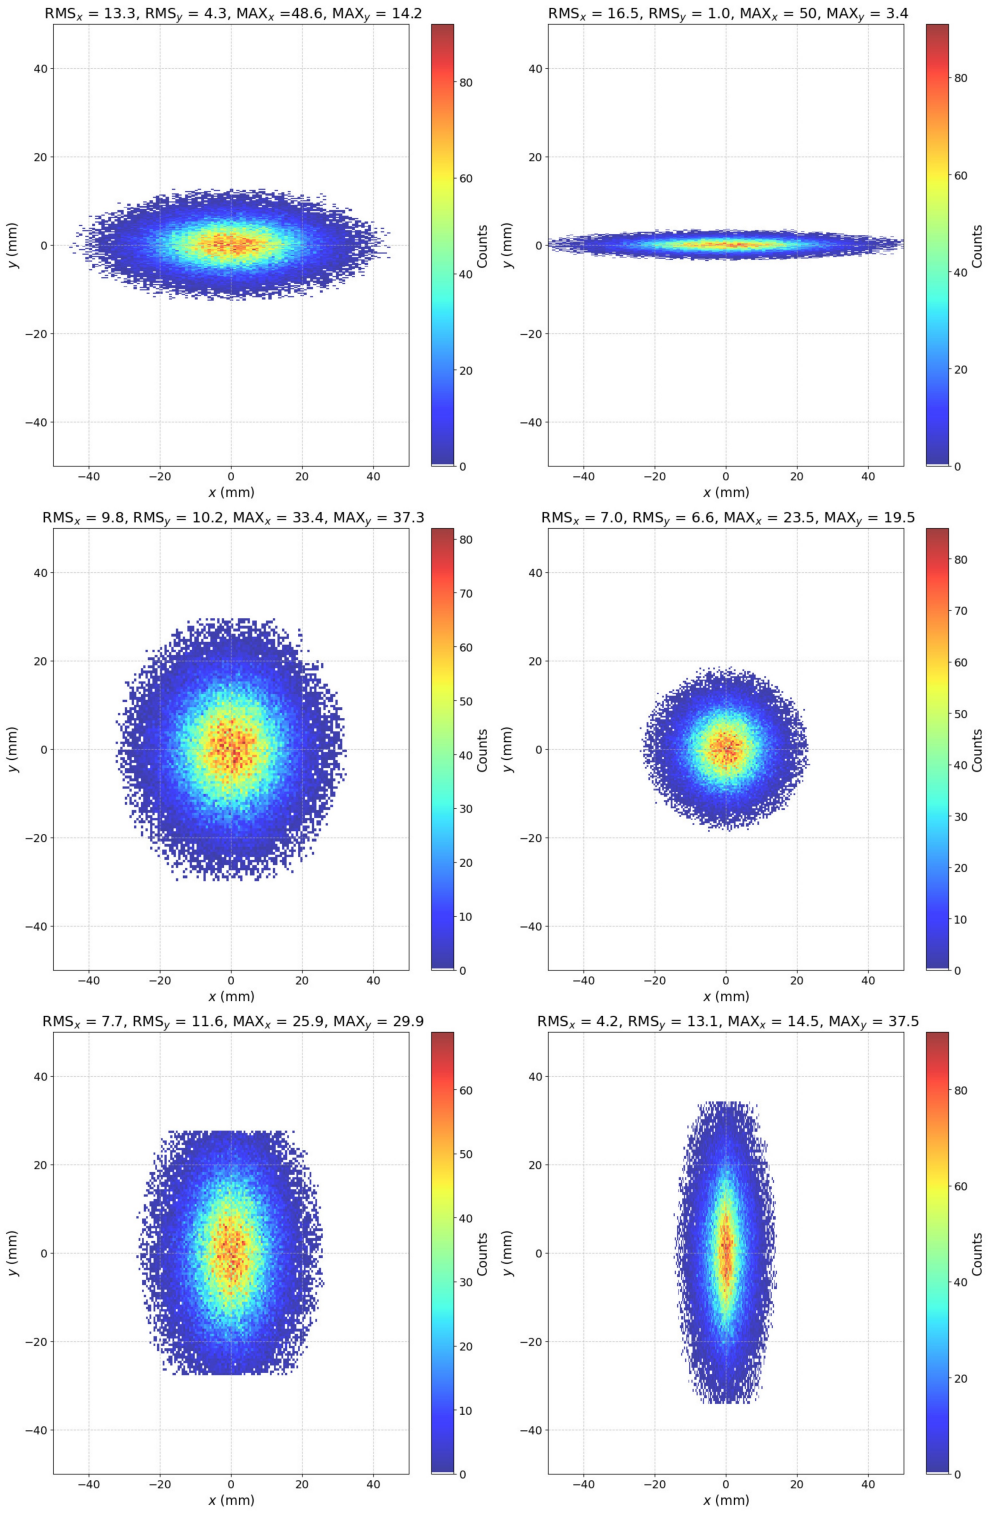

Figure 5: Examples of footprints achieved in OPAL with the optimal quadrupole parameters found for 6 of the targets from Table 1.

E Losses distributions

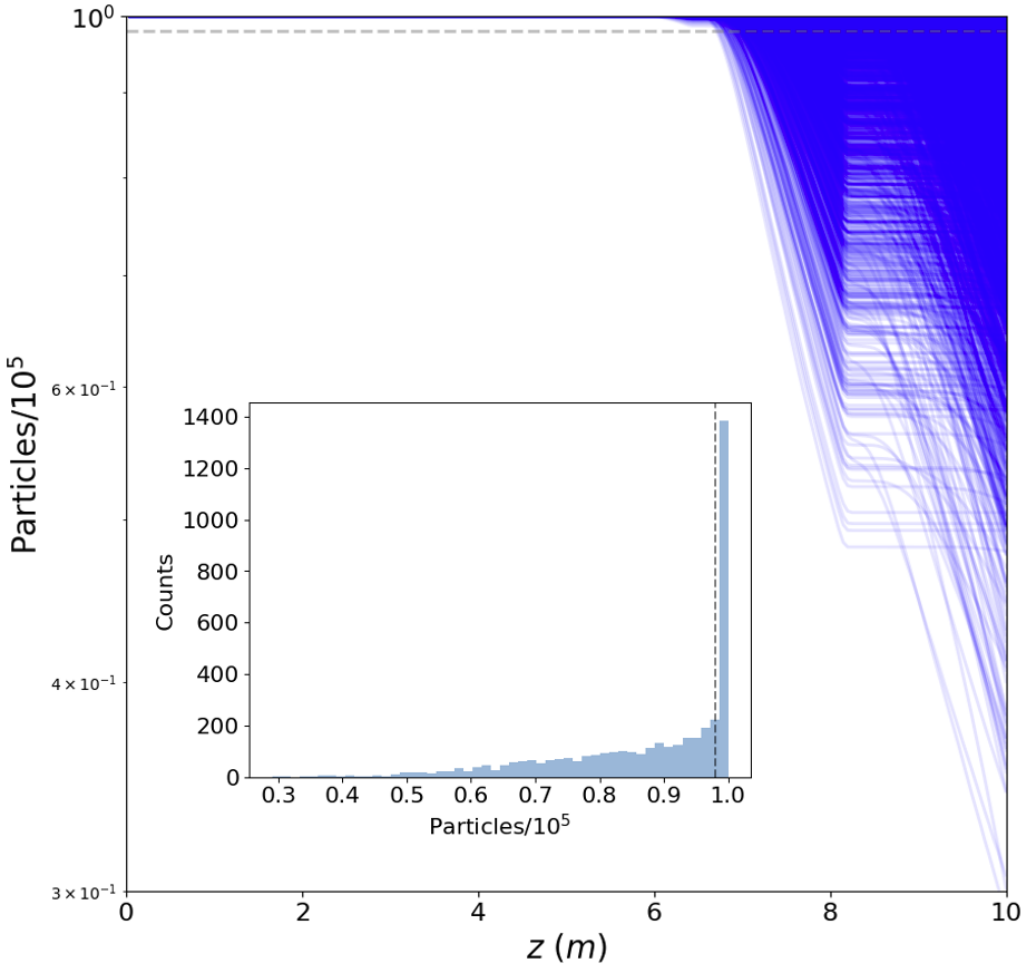

Figure 6: Normalized number of particles $n(z)$ for some of the simulations, which indicate losses in the accelerator. The dashed line corresponds to the 0.98 threshold used in the optimization algorithm, to exclude solutions with losses above that number. The histogram represents the particles distribution at the end of the lattice.

F OPAL lattice parameters and dataset creation

The following tables summarize the parameters for the lattice.in file that OPAL requires to perform the simulations. As stated in the main text, the quadrupole values range from 4 to 8 $T \cdot m^{-1}$. 8192 simulations are performed varying these values and keeping the remaining fixed. Roughly 20% of the instances are used for training, that is 6560 elements for training and 1632 for testing. Once each simulation is finished, its data functions are transformed to tensors for creating HDF5 files. The FNO implementation of nvidia-modulus requires input tensors of shape [Variable, N, channel, discretization]. The variables are explained in Section 2 and Appendix A. Each instance of the 1D model has 199 values for z and each instance of the 2D model has shape (200, 201) for (x, z) or (y, z). The (x, y) coordinates are confined within a circular aperture of radius 50 mm, defining the two-dimensional spatial domain. The z coordinate ranges from 0 to 10 m, representing the depth or longitudinal extent of the lattice.

Table 4: Magnetic elements dimensions, length and edge where it starts relative to the beginning of the lattice. The aperture radius of all the elements is r= 50 mm.

Element	Length (m)	Edge (m)
DR1 (Drift)	1.00	0.00
Q1 (Quadrupole)	0.25	1.00
DR2 (Drift)	0.60	1.25
Q2 (Quadrupole)	0.25	1.85
DR3 (Drift)	0.60	2.10
Q3 (Quadrupole)	0.25	2.70
DR4 (Drift)	1.65	2.95
Q4 (Quadrupole)	0.25	4.60
DR5 (Drift)	1.50	4.85
Q5 (Quadrupole)	0.25	6.35
DR6 (Drift)	1.50	6.60
Q6 (Quadrupole)	0.25	8.10
DR7 (Drift)	1.65	8.35

Table 5: Important parameters for the OPAL-T configuration.in file. Variables are given as they appear in the file. Fields are: particles information, injection, integration and solver.

Variable	Value
Particle	Deuteron
Number of particles	100000
Energy	40 MeV
Injection type	Gauss
$\sigma_{x,y,z}$	28 mm
σ_{p_x,p_y}	$0.023{\cdot}10^{-3}$
σ_{p_z}	$0.039{\cdot}10^{-3}$
MAXSTEPS	10000
DT	$1.65{\cdot}10^{-10}$
ZSTOP	10
METHOD	PARALLEL-T
FIELDSOLVER	FFT
$M_{X,Y,T}$	16
PARFFT$_{X,Y,T}$	true
BCFFT$_{X,Y,T}$	open
BBOXINCR	1
GREENSF	STANDARD

G Predictions examples (1D)

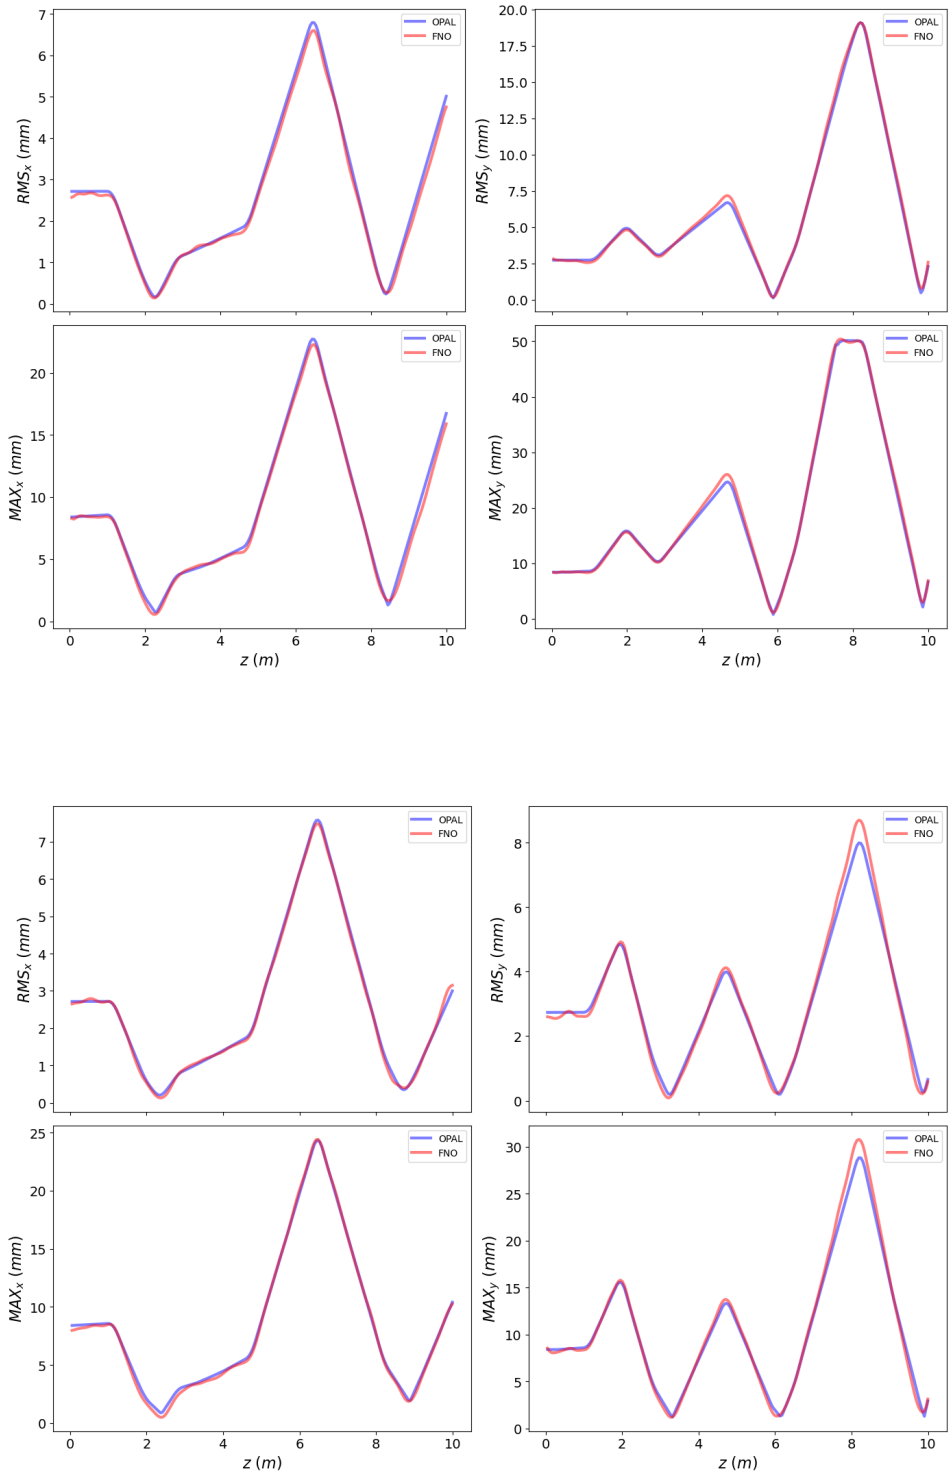

Figure 7: Two more examples for the one-dimensional beam statistics predicted by the model and compared to simulations. Blue: OPAL simulations. Red: FNO predictions. These two solutions correspond to other two configurations of Table 1.

H Predictions examples (2D)

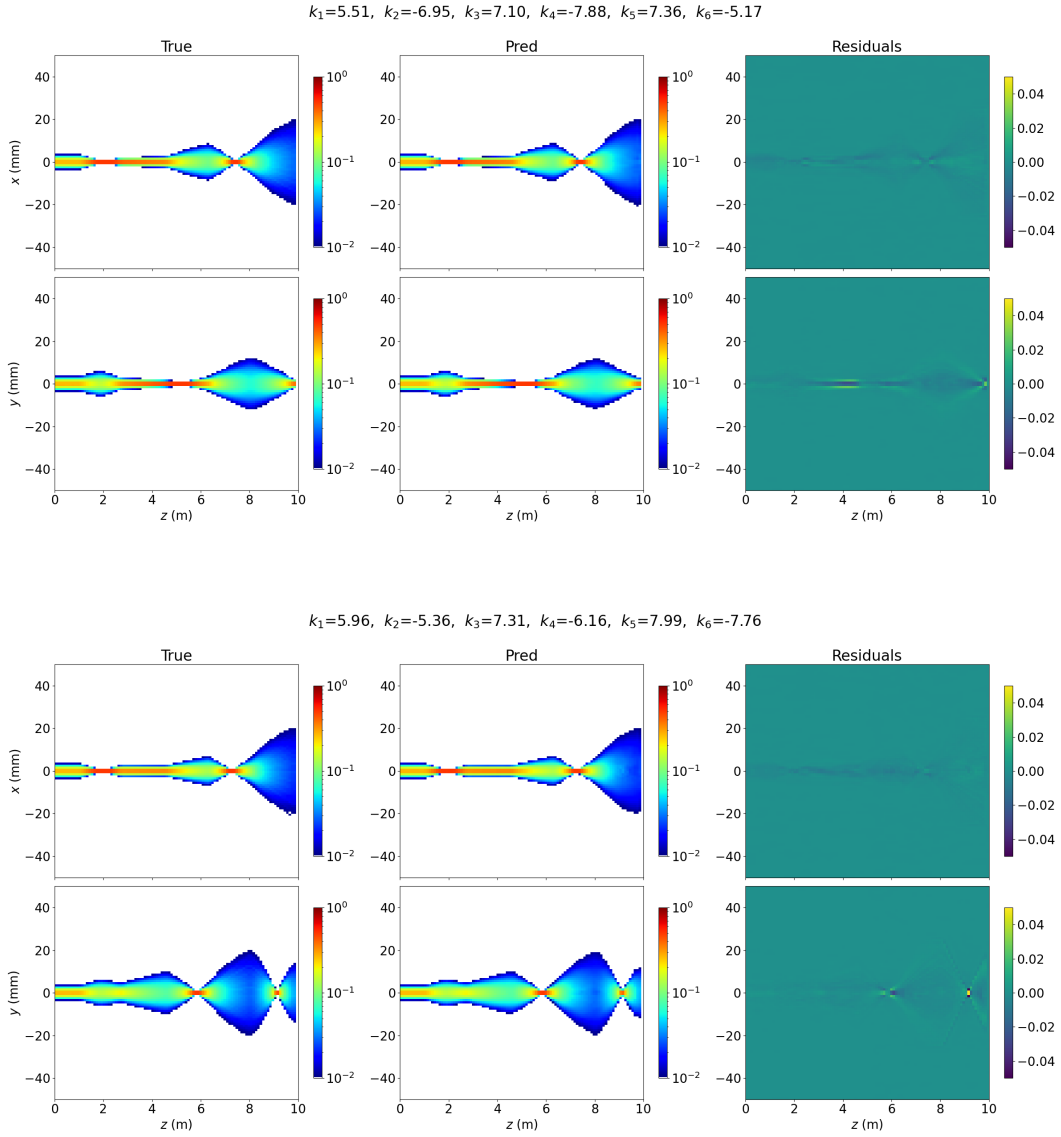

Figure 8: Two more examples for the two-dimensional beam profile distributions predicted by the model and compared to simulations. Left: OPAl simulations. Middle: FNO predictions. Right: residual errors.

I Model errors

Table 6: FNO 1D Model errors evaluated in all the test set.

Variable	MAPE (%)	MAE
$\text{RMS}_x(z)$	4.38	0.11
$\text{RMS}_y(z)$	5.63	0.16
$\text{MAX}_x(z)$	3.92	0.30
$\text{MAX}_y(z)$	3.76	0.39
$n(z)$	0.92	0.01

Table 7: FNO 2D Model errors evaluated in all the test set.

Variable	MAE
$\rho_x(x,z)$	$4.7{\cdot}10^{-4}$
$\rho_y(y,z)$	$5.6{\cdot}10^{-4}$

J FNO model predictions for the number of particles

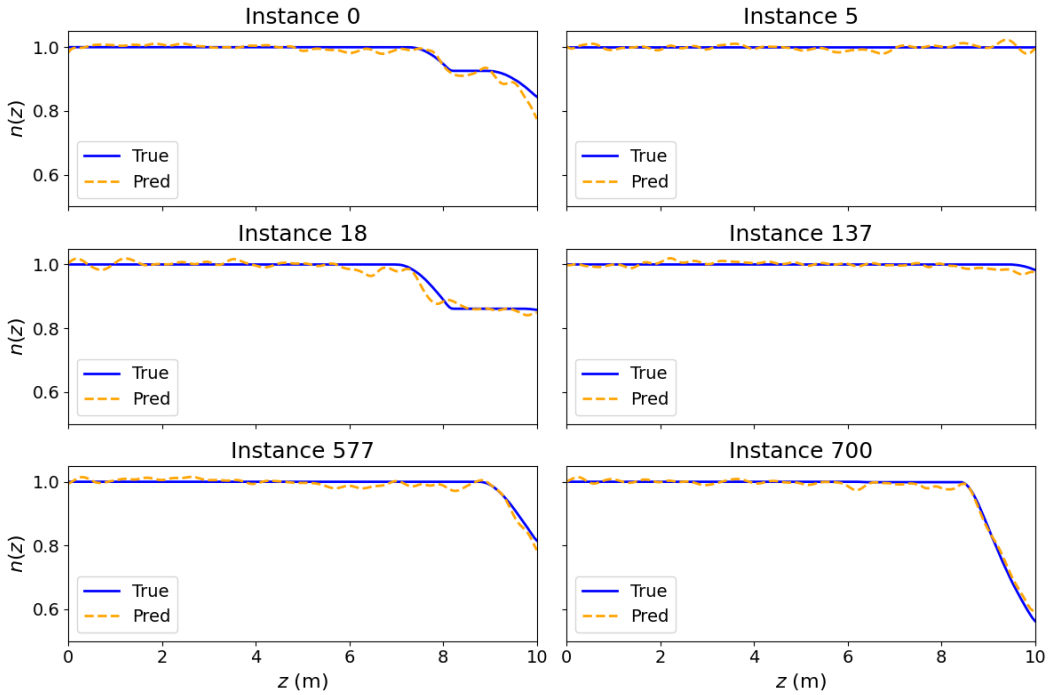

Figure 9: Test examples of the model predictions for the number of particles along z.

K System specifications

Table 8: System specifications

GPU Model	NVIDIA GeForce RTX 3060
CUDA version	12.2
CPU Model	Intel(R) Core(TM) i9-10900K CPU @ 3.70GHz
CPUs	12

Table 9: Software

Software/Code	**Version**
OS	Ubuntu 22.04.4 LTS (Docker)
nvidia-modulus.sym	1.5.0
OPAL	2022.1.0
openmpi	4.1.4
Torch	2.3.0

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: We successfully trained the FNO models to achieve the reported speed improvements and accuracy, and effectively implemented the 1D model within an optimization loop to determine the optimal quadrupole settings that satisfy the beam requirements.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The high computational cost of simulations necessitated limiting the optimization to only the S1 section of the HEBT, with fixed ranges and signs for the quadrupoles. Expanding parameters would require over 8,000 simulations, making it impractical. Therefore, we adopted a modular approach.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: no theoretical results are presented, its all applied models.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Even though the code is not provided, architecture, lattice, hyperparameters and frameworks are fully specified either in the main text or appendices.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: While all utilized tools and methods are open source and publicly available, the IFMIF-DONES consortium's policies restrict the sharing of specific data and code associated with this research.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: refer to dataset creation and hyperparameters tables in the appendices. Although hyperparameters choice is not discussed, they are based on previous works results. FNO models are very robust to hyperparameters.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: uncertainty estimation in FNO models is still a research topic. Time constraints.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: refer to system specifications appendix.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: Our work complies with the NeurIPS Code of Ethics by using only open-source tools and publicly available data, with no involvement of human subjects or sensitive data. We prioritize privacy, transparency, and responsible use of our models, avoiding potential risks such as security vulnerabilities or discrimination.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [No]

    Justification: the paper implies indirect contributions to nuclear fusion, which will have an enormous impact in the future. However, no direct impact is discussed, besides the results helping scientific facilities.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: No risk of misuse as the models are not shared, and they are only used to find optimal quadrupole values in linear accelerators.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: everything is either publicly available or open source and properly cited.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: dataset creation and model training are explained in the text.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: No human subjects nor crowdsourcing.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: No human subjects were used, therefore there is no related risk.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.