

IMPROVING PEAK PREDICTION IN MEBT SURROGATE MODELS USING FOURIER FEATURE MAPPING*

J. Cha^{†1}, X. Liu¹, J. Thiyagalingam¹, K. Baker²

¹STFC, Swindon, United Kingdom

²ISIS Neutron and Muon Source, Didcot, United Kingdom

Abstract

Accurate prediction of localised beam features such as peak loss and rapid envelope oscillations is essential for reliable operation of the Medium Energy Beam Transport (MEBT) at ISIS. Standard neural network surrogate models often struggle to capture these effects due to spectral bias, favouring smooth variations over sharp changes. In this work, we investigate the use of Fourier feature mapping to improve surrogate model performance on high-frequency beam behaviour. The method projects spatial inputs into a sinusoidal basis, enabling the network to represent high-frequency variations more effectively. We train two surrogate models on MEBT simulation data, one using Fourier feature mapping and the other without, and compare their performance in capturing peak regions. Results show that Fourier feature mapping significantly improves peak prediction accuracy while maintaining comparable performance on global beam trends. This demonstrates a straightforward and computationally efficient approach for enhancing surrogate modelling of beam dynamics in low-energy transport systems.

INTRODUCTION

Particle accelerators, including light sources and neutron facilities, aim to deliver a high-quality and reliable beam to experimental beamlines. This requires tuning machine parameters such as radio-frequency (RF) cavities and magnetic elements based on available diagnostics and operational experience. While particle tracking simulations provide accurate predictions during the design phase, their computational cost limits their use in real-time operation, where beam characterisation is further constrained by limited and often disruptive diagnostics.

Beam loss is a key operational diagnostic, providing an indication of beam degradation and playing a critical role in machine protection. Its spatial distribution is particularly important, but limited diagnostic coverage makes detailed beam behaviour difficult to infer, motivating the use of fast predictive models.

Particle tracking codes such as PARMELA [1], BMAD [2], elegant [3], and ASTRA [4] provide high-fidelity simulations of beam dynamics, but their computational expense prevents their use in time-sensitive applications such as online optimisation and control. Neural network-based surrogate models have therefore been

adopted to approximate these simulations with reduced computational cost [5, 6].

Despite their success as universal function approximators, neural networks exhibit spectral bias, favouring low-frequency components and struggling to represent high-frequency variations [7, 8]. This limits their ability to capture localised features such as sharp beam loss peaks, where inaccuracies can significantly affect cumulative quantities.

In this work, we investigate Fourier feature mapping to improve the representation of high-frequency beam behaviour in surrogate models. The method augments the input with sinusoidal basis functions, enabling improved capture of rapid spatial variations. It is evaluated on a surrogate model of the Medium Energy Beam Transport (MEBT) at the ISIS Neutron and Muon Source [9], demonstrating improved peak prediction accuracy while maintaining performance on global beam trends.

MEBT MODEL

To investigate the impact of Fourier feature mapping on high-frequency beam prediction, we consider a surrogate model of the planned Medium Energy Beam Transport (MEBT) at the ISIS Neutron and Muon Source.

In the current ISIS configuration, a significant fraction of the beam from the ion source is lost before entering the synchrotron due to transport inefficiencies [10]. The MEBT is designed to improve matching between the radio-frequency quadrupole (RFQ) and the first drift tube linac (DTL) tank, increasing transmission efficiency from approximately 70 % to 96 %.

The MEBT consists of eight quadrupoles for transverse focusing and four rebunching cavities for longitudinal control. The beam dynamics depend on the field strengths of these elements and the incoming beam conditions, including beam size, position, and current.

Beam transport is simulated using the ASTRA particle tracking code [4], implemented via LUME [11]. Simulations track particles along the 2.04 m beamline and evaluate beam properties at 200 points, including beam size, emittance, energy, bunch length, and particle loss.

The local beam loss profile exhibits sparse, high-frequency behaviour with sharp peaks, while cumulative loss forms a step-like function. Accurate prediction of these peaks is critical, as errors can propagate and affect cumulative quantities relevant for machine protection and optimisation. As the MEBT is not yet operational, simulation-based modelling provides the only means to study this behaviour, motivating the development of fast surrogate models.

* This work was supported by the STFC Facilities Fund through the Ada Lovelace Centre.

[†] jaehoon.cha@stfc.ac.uk

A total of 10,000 simulations were generated by sampling the input parameter space around the design settings of the MEBT. Among these, 4,193 were discarded due to unstable beam transport with excessive loss, leaving 5,807 simulations. The dataset was split into training, validation, and test sets of size 4,065, 871, and 871 respectively, and used to evaluate surrogate models in capturing both global trends and localised high-frequency behaviour.

METHOD

Fourier Feature Mapping

Neural networks are known to exhibit spectral bias, favouring low-frequency components of a function and struggling to represent high-frequency variations [7,8]. To address this limitation, Fourier feature mapping [12] is applied to the spatial coordinate input.

Given a coordinate $\mathbf{v} \in [0, 1]^d$, the mapping is defined as

$$\gamma(\mathbf{v}) = [a_j \cos(2\pi \mathbf{b}_j^T \mathbf{v}), a_j \sin(2\pi \mathbf{b}_j^T \mathbf{v})]_{j=1}^m \quad (1)$$

where m is the embedding dimension, $\mathbf{b}_j$ defines the frequency components, and $a_j = 1/j^p$ controls the amplitude scaling, with p determining the contribution of higher-frequency components.

This transformation projects the input into a higher-dimensional space of sinusoidal basis functions, enabling improved representation of localised high-frequency behaviour.

In this work, the one-dimensional spatial coordinate along the MEBT is transformed using Fourier feature mapping and provided as input to the surrogate model.

Surrogate Model Architecture

The surrogate model is implemented as a fully connected feed-forward neural network that approximates the relationship between machine settings and beam behaviour along the MEBT. The input consists of the machine control parameters and incoming beam conditions, together with the spatial coordinate along the beamline. In the Fourier feature model, the raw spatial coordinate is replaced by its Fourier feature embedding defined in Eq. (1).

The network contains three hidden layers with 128, 256, and 128 neurons, respectively, each followed by a LeakyReLU activation. A final linear layer predicts 11 outputs at each spatial location. The first output, corresponding to the beam loss profile, is constrained to be non-negative using a ReLU activation. The cumulative loss is not predicted directly, but is reconstructed as the cumulative sum of the beam loss, resulting in 12 outputs in total.

To ensure a fair comparison, the baseline and Fourier-feature models use the same architecture and training setup, differing only in how the spatial coordinate is represented.

Training and Data Processing

All input variables are standardised using the mean and standard deviation computed from the training set. For the outputs, the first 10 values are standardised in the same

way, while the beam loss and reconstructed cumulative loss remain in their original scale. The output arrays are then transposed so that the model predicts beam descriptors as functions of spatial position along the MEBT.

The models are trained by minimising the mean squared error (MSE) between predictions and simulation outputs. Optimisation is performed using the Adam optimiser, and early stopping based on validation loss is used to avoid overfitting.

RESULTS

We evaluate the impact of Fourier feature embeddings on the prediction of beam descriptors along the ISIS MEBT. In particular, we focus on the reconstruction of sparse beam loss profiles, where standard neural networks tend to smooth or suppress sharp peaks. Models without Fourier features are compared against models with embedding dimensions $m = 25, 50$, and 100 .

Quantitative and Qualitative Comparison

The quantitative performance of the models is summarised in Table 1. We report the mean squared error (MSE) of the loss, peak reconstruction error, high-frequency (HF) error, and final cumulative error, where lower values indicate better performance.

Table 1: Quantitative Comparison of Fourier Feature Embeddings

Model	MSE ↓	Peak ↓	HF ↓	Cum. ↓
No FF	0.2912	0.3973	33.7994	4.7800
FF-25	0.1669	0.3726	33.7474	2.7414
FF-50	0.1412	0.3298	33.9781	2.5345
FF-100	0.1207	0.2580	24.3062	2.3435

While Fourier feature models improve all metrics, the differences are most evident in peak reconstruction and cumulative errors. Although the baseline model achieves a moderate MSE, this metric is misleading due to the sparsity of the loss profile, where most values are close to zero. As shown in Fig. 1 for a region containing high-frequency variations ($z = 75-175$), the baseline model fails to reproduce sharp peaks and instead predicts an almost constant response along the beamline. In contrast, Fourier-feature models accurately capture both the location and magnitude of the peaks, with performance improving as the embedding dimension increases.

These results demonstrate that Fourier feature mapping is essential for modelling sparse, high-frequency beam loss profiles. Improvements in peak reconstruction directly translate to reduced cumulative error, which is critical for reliable machine optimisation and operation.

Embedding Dimension Ablation

The effect of embedding dimension on peak reconstruction error is summarised in Fig. 2. A clear monotonic decrease is observed as m increases, indicating that larger




Figure 1: Comparison of predicted sparse loss profiles for representative test examples. The baseline model predicts an almost constant response, while Fourier feature embeddings enable accurate reconstruction of sharp loss peaks.

Fourier embeddings improve the reconstruction of localised peaks and mitigate the low-frequency bias of the baseline network.




Figure 2: Embedding Dimension Ablation Using Peak Reconstruction Error

Overall, these results show that Fourier feature embeddings provide a simple and effective way to improve high-frequency prediction in accelerator surrogate models. In particular, they significantly enhance the reconstruction of sparse beam loss profiles while also reducing cumulative prediction error.

CONCLUSION

This work demonstrates that Fourier feature embedding significantly improves the prediction of high-frequency structures in beam profiles. While standard neural networks trained with mean squared error tend to produce overly

smooth outputs, Fourier features enable accurate reconstruction of sharp variations. Performance consistently improves with increasing embedding dimension, with FF-100 achieving the best results across all metrics, including high-frequency and peak reconstruction errors. These findings highlight Fourier feature embedding as a simple and effective approach for enhancing surrogate models, particularly in applications where capturing fine-scale variations is critical.

REFERENCES

- [1] L. Young and J. Billen, “The particle tracking code PARMELA”, in *Proc. PAC’03*, Portland, OR, USA, pp. 3521–3523, May 2003. doi:10.1109/PAC.2003.1289968
- [2] D. Sagan, “Bmad: a relativistic charged particle simulation library”, *Nucl. Instrum. Methods Phys. Res. A*, vol. 558, no. 1, pp. 356–359, 2006. doi:10.1016/j.nima.2005.11.001
- [3] M. Borland, “elegant: a flexible SDDS-compliant code for accelerator simulation”, Argonne National Laboratory, Lemont, IL, USA, Rep. LS-287, 2000. doi:10.2172/761286
- [4] K. Flöttmann, ASTRA: a space charge tracking algorithm, <https://www.desy.de/~mpyf10/>
- [5] K. Baker *et al.*, “Development of a virtual diagnostic for estimating key beam descriptors”, in *Proc. IPAC’22*, Bangkok, Thailand, pp. 969–972, Jun. 2022. doi:10.18429/JACoW-IPAC2022-TUPOST048
- [6] R. Roussel *et al.*, “Phase space reconstruction from accelerator beam measurements using neural networks and differentiable simulations”, *Phys. Rev. Lett.*, vol. 130, no. 14, p. 145001, 2023. doi:10.1103/PhysRevLett.130.145001
- [7] R. Basri, M. Galun, A. Geifman, D. Jacobs, Y. Kasten, and S. Kritchman, “Frequency bias in neural networks for input of non-uniform density”, in *Proc. ICML’20*, Vienna, Austria, pp. 685–694, Jul. 2020.
- [8] N. Rahaman *et al.*, “On the spectral bias of neural networks”, in *Proc. ICML’19*, Long Beach, CA, USA, pp. 5301–5310, Jun. 2019.
- [9] J. W. G. Thomason, “The ISIS spallation neutron and muon source: the first thirty-three years”, *Nucl. Instrum. Methods Phys. Res. A*, vol. 917, pp. 61–67, 2019. doi:10.1016/j.nima.2018.11.129
- [10] S. R. Lawrie *et al.*, “A pre-injector upgrade for ISIS, including a medium energy beam transport line and an RF-driven H[−] ion source”, *Rev. Sci. Instrum.*, vol. 90, no. 10, p. 103310, 2019. doi:10.1063/1.5127263
- [11] C. E. Mayes *et al.*, “Lightsource Unified Modeling Environment (LUME), a start-to-end simulation ecosystem”, in *Proc. IPAC’21*, Campinas, SP, Brazil, pp. 4212–4215, May 2021. doi:10.18429/JACoW-IPAC2021-THPAB217
- [12] M. Tancik *et al.*, “Fourier features let networks learn high frequency functions in low dimensional domains”, in *Proc. NeurIPS’20*, vol. 33, Virtual, pp. 7537–7547, Dec. 2020.