

DEVELOPMENT OF A SURROGATE MODEL OF THE T09 BEAMLINE AT THE CERN PS EAST AREA

D. Banerjee *, P. Alexaki, J. Bernhard, N. Charitonidis, G. Dal Maso,
M. van Dijk, S. Esen, A. M. Goillot, F. Metzger, L. J. Nevay,
B. Rae, S. Schuh, F. W. Stummer, CERN, Geneva, Switzerland,
E. Barber, University of Surrey, Surrey, UK
R. Jakob Tietgen, Johannes Gutenberg Universitaet, Mainz, Germany

Abstract

The T09 beamline of the Proton Synchrotron (PS) East Experimental Area delivers mixed secondary particle beams for diverse experiments and test beams. This requires the possibility of generating flexible beam conditions. Predicting beam parameters at the experimental area, such as particle composition, spatial profiles, and momentum spread, typically demands time-intensive simulations with tools like Beam Delivery Simulation (BDSIM) and FLUKA. This work presents the development and first studies of a surrogate model of the T09 beamline that reproduces key beam characteristics with greatly reduced computational cost, making it a promising candidate for future real-time applications. Trained on a database of detailed BDSIM simulations spanning multiple collimator configurations, the model employs machine-learning-based regression to predict beam transmission, spot size, and phase-space parameters. Preliminary validation against full Monte Carlo simulations shows agreement at a level of 15 % for the tested configurations, supporting rapid parameter scans and indicating potential for future real-time beam optimisation. This approach is also scalable for other CERN beamlines.

INTRODUCTION

The East Area at CERN has served the physics community for over 50 years, delivering secondary beams produced from the 24 GeV/c proton beam extracted from the Proton Synchrotron (PS) [1]. Following the East Area renovation [2] during Long Shutdown 2, the updated facility began operation in 2021. The T09 beamline accepts secondary particles at a 30 mrad vertical production angle, is approximately 50 m long, and includes three bending dipoles, eight quadrupoles, and three four-jaw collimators, providing a selectable relative momentum spread ($\Delta p/p$) between 0.7 % and 15 %. Each experiment requires flexible operation with regular adjustment of collimator and quadrupole settings to tailor beam flux, momentum spread, spot size, and phase-space properties. Predicting the resulting beam properties requires full Monte Carlo runs with Beam Delivery Simulation (BDSIM) [3], which can take hours per configuration due to the low yield of the secondary beam necessitating large event statistics. A fast surrogate model, trained to emulate simulation output, will enable real-time beam property prediction and efficient parameter optimisation. This paper summarises

the first studies performed to confirm the efficacy of such a model for the T09 beamline.

SIMULATION CAMPAIGN AND DATASET

In this study, detailed simulations of the beamline were performed using BDSIM [3], a Geant4-based [5] Monte Carlo code tracking particles through the magnetic lattice while modelling various physics processes including particle interactions with matter. The T09 BDSIM model was validated against MAD-X (Methodical Accelerator Design) [4] as shown in Fig. 1. Then a split-simulation strategy was employed: a short run to a marker downstream of the T09 target produced a filtered particle distribution, which was reused as input for all full-beamline runs, substantially reducing the per-configuration computational cost.




Figure 1: MAD-X vs. BDSIM validation. The optics of the beamline is also shown here with the different elements.

Five diagnostic markers were placed at the end of the beamline (exp1–exp5) 1 m apart to record the beam parameters. Only 26 configurations were simulated, defined by six input parameters: the horizontal and vertical half-apertures of collimators XCHV011 and XCHV014, ranging between ± 5 mm and ± 50 mm, and XCHV026, ranging between ± 1.5 mm and ± 45 mm. For each configuration, 1×10^9 primary particles on target equivalent of secondary beam were tracked and processed to produce phase-space histograms in $Y-X$, $X'-X$, $Y'-Y$, and $p-X$ at the marker locations. The beamline was simulated with the reference momentum at 15 GeV/c. An 80 %/20 % train–test split yields 21 configurations for training and 5 for evaluation. The training was performed only for the horizontal parameters X , X' , and momentum p to validate the approach.

* dipanwita.banerjee@cern.ch

SURROGATE MODEL ARCHITECTURE

Histogram Encoding

Rather than predicting scalar optics parameters, the surrogate model was trained to predict full histogram distributions. For each combination of observable (x -position, x' -divergence, momentum p) and position marker (exp1–exp5), the 1000-bin histogram from each training configuration is assembled into a matrix. Principal Component Analysis (PCA) reduces this to a compact set of scores. A regressor then maps the six collimator inputs to these scores, and the predicted histogram is recovered via inverse PCA transform.

Models and Training

Four algorithms were evaluated: Gaussian Process Regression (GPR) [6], Gradient Boosting Regression (GBR) [7], Support Vector Regression (SVR) [8], and Random Forest (RF) [9]. All were implemented in scikit-learn [10] with physically motivated hyperparameters (GPR: Matérn $\nu = 2.5$ kernel, 5 optimiser restarts; GBR: 200 estimators, learning rate 0.05; SVR: $C = 10$, $\epsilon = 0.01$; RF: 200 estimators), rather than from a systematic grid search. A cross-validation optimisation is planned for a larger future dataset. GPR is used here purely as a nonlinear regressor: only the predictive mean is evaluated. The predictive variance, which could serve as an out-of-distribution confidence metric, is reserved for a future extension. Each model is trained independently per (variable, marker) pair. For momentum, a parallel scalar path regresses directly on the Gaussian μ and σ of the peak region ($\pm 20\%$ around the maximum). Prediction quality is evaluated on the test set with three metrics:

- R^2 – shape residual of area-normalised distributions $\tilde{h}_i = h_i / \sum_j h_j$, computed independently of yield,

$$R^2 = 1 - \frac{\sum_i (\tilde{h}_i^{\text{true}} - \tilde{h}_i^{\text{pred}})^2}{\sum_i (\tilde{h}_i^{\text{true}} - \bar{\tilde{h}}^{\text{true}})^2};$$

the area normalisation ensures R^2 measures shape agreement only, with yield differences captured separately by $\Delta N/N$;

- $\Delta N/N$ – relative yield error,

$$\frac{\Delta N}{N} = \frac{|\sum_i h_i^{\text{pred}} - \sum_i h_i^{\text{true}}|}{\sum_i h_i^{\text{true}}};$$

- Δpeak – absolute shift of the histogram mode in physical units,

$$\Delta\text{peak} = |x_{\text{max}}^{\text{pred}} - x_{\text{max}}^{\text{true}}|.$$

MODEL COMPARISON ON THE TEST SET

Table 1 summarises mean metrics averaged over exp1–exp5 and the 5 test-set configurations. Figure 2 shows the full per-marker heatmap breakdown.

Table 1: Mean test-set metrics per model and observable, averaged over exp1–exp5 and all test configurations. $\uparrow$: higher is better ($R_{\text{max}}^2 = 1$); $\downarrow$: lower is better ($(\Delta N/N)_{\text{min}}$ and $\Delta\text{peak}_{\text{min}} = 0$). Δpeak units: mm for x , mrad for x' , GeV/c for p .

Model	Var	$R^2 \uparrow$	$\Delta N/N \downarrow$	$\Delta\text{peak} \downarrow$
GPR	x	+0.931	0.118	0.81
	x'	+0.955	0.186	0.14
	p	+0.792	0.395	0.10
GBR	x	−20.4	0.357	6.95
	x'	+0.731	0.328	0.55
	p	+0.595	0.267	0.21
SVR	x	+0.870	0.192	0.93
	x'	+0.949	0.850	0.17
	p	+0.849	0.441	0.11
RF	x	+0.962	0.896	0.77
	x'	+0.957	0.942	0.14
	p	+0.777	0.878	1.12

For transverse observables, GPR and RF achieve similar shape fidelity ($R^2 \approx 0.93$ – 0.96), but differ critically in yield: GPR keeps $\Delta N/N$ below 19% for both x and x' , whereas RF yields errors of 90–94%. SVR is competitive in shape ($R^2 \approx 0.87$ – 0.95) but incurs large yield errors for x' (85%). GBR shows generalisation failure for x ($R^2 \approx -20$), likely due to boosting instability on the small training set, and is therefore discarded. For momentum, SVR achieves the highest R^2 (+0.85), followed by GPR (+0.79) and RF (+0.78); GBR is lowest (+0.60). All models show elevated $\Delta N/N$ for p (27–44%). For the Gaussian scalar path, SVR leads on $R^2(\mu) = 0.78$; $R^2(\sigma)$ is variable across markers, with GBR showing the highest mean (≈ 0.81). GPR is therefore identified as the primary model for transverse phase-space predictions, offering the best combined shape–yield performance. SVR is preferred for the momentum PCA histogram and the scalar μ prediction.

Momentum Prediction Limitations

The momentum distribution features a sharp beam peak at approximately 15 GeV/c superimposed on a broad low-energy tail. Despite this two-component structure, only 3 PCA components capture more than 99% of the training variance. The difficulty arises not from PCA dimensionality but from the regression step: with only 21 training configurations the mapping to PCA scores encoding tail variations is poorly constrained, giving elevated $\Delta N/N$ even when the peak is correctly located. The Gaussian scalar path correctly recovers μ ($\Delta\text{peak} < 0.15$ GeV/c), but $R^2(\sigma)$ remains below 0.5 beyond exp1. This is compounded by the momentum spectrum being invariant across exp1–exp5 for a given configuration; marker-to-marker histogram variations reflect Monte Carlo statistics rather than physics, so training independent regressors per marker unnecessarily fragments the data. A larger campaign, pooling of momentum observa-



Figure 2: R^2 , $\Delta N/N$, and Δpeak heatmaps across all four models, variables, and marker positions (test set, mean $\pm$ std over configurations).

tions across markers, and a peak-plus-tail decomposition will substantially improve accuracy.

VALIDATION AGAINST AN UNSEEN CONFIGURATION

The surrogate was applied to an operational collimator setting not in the training set: XCHV011= 0.018 m, XCHV014= 0.018 m, XCHV026_x = 0.007 m, XCHV026_y = 0.030 m, with a dedicated BDSIM run providing the ground truth. Figure 3 shows the predicted and true histogram overlays for exp1. The transverse parameters are shown in absolute yield as well as normalised histograms to differentiate between shape and yield prediction. The momentum panel confirms that the Gaussian μ , σ lines (dashed) correctly locate the beam peak. Table 2 compares the Gaussian σ extracted by fitting the predicted and true histograms for GPR (transverse) and SVR (momentum).

Table 2: Gaussian σ (predicted vs. truth) for the unseen configuration, GPR (transverse) and SVR (momentum). x in mm; x' in mrad; p in GeV/c.

		exp1	exp2	exp3	exp4	exp5
σ_x [mm]	Truth	8.5	11.4	13.5	15.3	16.5
	GPR	9.95	12.2	13.1	13.7	13.9
$\sigma_{x'}$ [mrad]	Truth	2.0	2.0	2.0	1.9	1.8
	GPR	1.8	1.7	1.6	1.5	1.4
σ_p [GeV/c]	Truth	0.282	0.284	0.286	0.290	0.298
	SVR	0.271	0.276	0.281	0.287	0.295

GPR overestimates σ_x by 7–16 % for less spread of the beam (exp1–exp2) and underestimates by 2–16 % as the beam widens. For $\sigma_{x'}$, GPR systematically underestimates

by 13–23 %, likely reflecting kernel regression compressing predictions toward the training mean for a quantity with low dynamic range across configurations; a denser simulation campaign is needed to isolate the cause. SVR recovers σ_p to within 1–4 %, consistent with its superior momentum R^2 . The dataset was made with a spatial cut of ± 50 mm for every marker. For the downstream markers, as the beam size increases, part of the tails are beyond this acceptance, which causes the decrease in x' and p .




Figure 3: Predicted (coloured lines) and BDSIM truth (black) histogram distributions for the unseen setting at exp1: x , x' absolute counts (top); x , x' normalised (middle); momentum (bottom).

CONCLUSIONS

A first study introducing a surrogate model predicting full beam phase-space histograms for the T09 beamline has been carried out and preliminarily validated on one unseen configuration. Trained and tested on 26 BDSIM configurations, it combines PCA encoding with regression at five beam markers. GPR achieves the best transverse performance ($R^2 = 0.93$ for x , 0.96 for x' , yield error < 19 %) and emerges as the leading candidate for future operational use, subject to validation on a larger dataset. SVR provides the best momentum accuracy ($R^2 = 0.85$, σ_p within 1–4 %); suggesting a GPR+SVR composite for future development. Momentum accuracy is constrained by the limited training set and the two-component spectral structure. Next steps include a larger simulation campaign, a systematic hyperparameter optimisation via cross-validation, a peak-plus-tail momentum decomposition, extension to magnet-current in-

puts, and application to the T10 and T11 beamlines of the CERN East Area.

ACKNOWLEDGEMENTS

The authors thank the BDSIM development group for their support.

REFERENCES

- [1] L. Gatignon, “Design and Tuning of Secondary Beamlines in the CERN North and East Areas,” CERN Accelerator Note: CERN-ACC-NOTE-2020-0043, 2020.
- [2] J. Bernhard, F. Carvalho, S. Evrard, E. Harrouch, and G. Romagnoli, “CERN Proton Synchrotron East Area Facility: Upgrades and renovation during Long Shutdown 2,” CERN Yellow Reports: Monographs, 2021.
- [3] L. J. Nevay *et al.*, “BDSIM: An accelerator tracking code with particle–matter interactions,” *Comput. Phys. Commun.*, vol. 252, p. 107200, 2020.
doi:10.1016/j.cpc.2020.107200
- [4] L. Deniau, H. Grote, G. Roy, F. Schmidt, “The MAD-X program (Methodical Accelerator Design), Version 5.08.01,” <https://madx.web.cern.ch/releases/last-rel/madxuguide.pdf>.
- [5] S. Agostinelli *et al.*, “Geant4 – a simulation toolkit,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 506, pp. 250–303, 2003.
- [6] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.
- [7] J. H. Friedman, “Greedy function approximation: A gradient boosting machine,” *Ann. Statist.*, vol. 29, no. 5, pp. 1189–1232, 2001. doi:10.1214/aos/1013203451
- [8] C. Cortes and V. Vapnik, “Support-vector networks,” *Mach. Learn.*, vol. 20, no. 3, pp. 273–297, 1995.
doi:10.1007/BF00994018
- [9] L. Breiman, “Random forests,” *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001. doi:10.1023/A:1010933404324
- [10] F. Pedregosa *et al.*, “Scikit-learn: Machine learning in Python,” *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.