Ref. Ares(2026)10187210 - 29/10/2026

THRILL

TECHNOLOGY FOR HIGH-REPETITION-RATE
INTENSE LASER LABORATORIES

Deliverable Data	
Deliverable number	D5.5 (D16)
Deliverable name	Release of an open source machine learning code for beam stability investigation
Work Package	WP 5.4
Lead WP/deliverable beneficiary	Helmholtz-Zentrum Dresden-Rossendorf (HZDR)
Type and dissemination level	Report, public
Deliverable status Achieved	
Submitting author	Jeffrey Kelling
Verified (WP leader)	Patrick Audebert
Approved (Coordinator)	Zsuzsanna Major
Due date of deliverable	31.10.2026

Funded by the European Union

This project has received funding by the European Union's HORIZON-INFRA-2022-TECH-01 call under grant agreement number 101095207

Table of Contents

About THRILL 1
Executive summary 2
Abbreviations 3
1 Introduction and objectives 4
2 Delivered software 5
2.1 Software package and release 5
2.2 Delivered artefacts and access 5
2.3 Intended use and reuse 5
3 Beam pointing analysis 6
3.1 PHELIX beamline and jitter sources 6
3.2 Pointing data and preprocessing 6
3.3 Analysis methods and metrics 6
4 Prediction and compensation system modelling 7
4.1 Patch-MLP prediction model 7
4.2 Compensation loop and mirror dynamics 7
4.3 Simulation workflow and latency budget 8
5 Results and discussion 8
5.1 Short-term compensation performance 8
5.2 Long-term stability 9
5.3 Failure cases and limitations 9
6 Conclusion 10
7 References 10
8 Annexes 11
8.1 Data and code availability 11
8.2 Software contents 11

Funded by the European Union

This project has received funding by the European Union's HORIZON-INFRA-2022-TECH-01 call under grant agreement number 101095207

List of tables

This report contains no numbered tables.

List of figures

This report contains no numbered figures.

Disclaimer

This document is part of the deliverables from the project THRILL, which has received funding from the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

About THRILL

The THRILL project deals with providing new schemes and devices for pushing forward the limits of research infrastructures (RI) of European relevance and ESFRI landmarks. To do so, the project partners have identified several technical bottlenecks in high-energy high-repetition-rate laser technology that prevent it from reaching the technical readiness level required to technically specify and build the needed devices, and guaranteeing sustainable and reliable operation of such laser beamlines at the partnering RIs. Advancing the technical readiness of these topics is strategically aligned with the long-term plans and evolution of the ESFRI landmarks FAIR, ELI (-BL) and Eu-XFEL, and RI APOLLON, bringing them to the next level of development and strengthening their leading position.

The project is focused and deliberately restricted to three enabling technologies, which require the most urgent efforts and timely attention by the community: high-energy high-repetition-rate amplification, high-energy beam transport and optical coating resilience for large optics. To reach our goals, the major activity within THRILL will be organized around producing several prototypes demonstrating a high level of technical readiness. Our proposal is addressing not yet explored technical bottlenecks - such as transport over long distances of large-aperture laser beams via relay imaging using all-reflective optics - and aims at proposing concrete steps to increase the performances and effectiveness of the industrial community through the co-development of advanced technologies up to prototyping in operational environments.

The project is not only pushing technology, it is also offering an outstanding opportunity to train a qualified work force for RIs and industry. With this in mind, the structure of THRILL promotes synergetic work, fast transfer to industry and integrated research activities at the European level. Access to the RIs will be granted as in-kind contribution.

Executive summary

This report documents the delivery of D16, the release of an open source machine learning code for beam stability investigation by Helmholtz-Zentrum Dresden-Rossendorf (HZDR, M46). The deliverable comprises the open-source software release, the accompanying scientific paper by Wang et al. [1] and the PHELIX beam-pointing dataset used in the study.

The software implements a **hybrid predictive-feedback stabilisation scheme** for the PHELIX laser system at GSI. A patch-based multilayer perceptron (Patch-MLP) forecasts short-term horizontal and vertical beam-pointing errors from diagnostic time series, and a proportional-integral-derivative (PID) controller converts the predictions into correction commands for an upstream tip-tilt mirror. **This pre-compensation offsets the latency of a conventional feedback loop**.

The report describes the delivered software, the beam pointing analysis, and the prediction and compensation system modelling. Over a ten-hour simulated test, the predictive controller reduced residual jitter by approximately 17-20 % in RMS, 11-13 % in peak-to-valley and 18-21 % in FWHM compared with a conventional PID controller. The controller remained stable without drift across 110 one-minute test segments. **The code, paper and data are openly available** under a Creative Commons Attribution 4.0 International licence.

Abbreviations

Abbreviation	Definition
FWHM	Full width at half maximum
MAS	Main Amplifier Sensor
ML	Machine learning
MLP	Multilayer perceptron
P2V	Peak-to-valley
Patch-MLP	Patch-based multilayer perceptron
PHELIX	Petawatt High-Energy Laser for Heavy Ion Experiments
PID	Proportional-integral-derivative
PSD	Power spectral density
RMS	Root mean square

1 Introduction and objectives

High-power laser facilities such as PHELIX at GSI require excellent beam pointing stability to ensure reproducible experiments. Mechanical vibrations, thermal drifts and air turbulence propagate through the beamline and cause pointing jitter at the target. Pointing errors are particularly critical for small targets and for steep focusing optics, where they can be converted into higher-order aberrations.

The PHELIX[1] pre-amplifier bridge is the dominant source of pointing jitter, while pointing is diagnosed downstream at the Main Amplifier Sensor (MAS). Conventional PID stabilisation suppresses slow drift but is limited by the delay in obtaining the diagnostic signal and by the inertia of large steering mirrors. Deliberate pre-compensation upstream of the jitter source is therefore attractive, provided that the controller can act before the disturbance reaches the target.

The objective of D16 is to release an open-source machine learning code that supports beam stability investigation and to document its use for beam pointing analysis, prediction and compensation. The deliverable provides a reproducible workflow that can be adapted to other high-energy laser facilities and to future real-time implementations. The scientific basis and the complete results are presented in the published paper by Wang et al. [1].

[1] Zs. Major et al.,"High-energy laser facility PHELIX at GSI: Latest advances and extended capabilities." High Power Laser Sci. Eng. **12** (2024): e39. DOI:10.1017/hpl.2024.17

2 Delivered software

2.1 Software package and release

The deliverable implements the D16 description: an open source machine learning code for beam stability investigation (HZDR, M46), released under an open source licence on a suitable platform, with GitLab used for development and RODARE for archival publication. The delivered software is an open-source Python/Jupyter-based package for beam stability investigation. It contains data-loading and preprocessing notebooks, a training pipeline for the Patch-MLP, a pretrained model and a prediction-and-simulation notebook. The release is archived on RODARE and is available under the Creative Commons Attribution 4.0 International (CC BY 4.0) licence.

The development repository is hosted on GitLab, as required for an open-source release on a suitable platform. The package is designed to be reusable: file paths, training parameters and simulation settings are exposed and annotated in the notebooks. The release includes the pretrained model weights used for the reported simulations.

The code can be used to reproduce the published simulations and to retrain the predictor on new PHELIX data. Its modular structure allows the predictor to be replaced or retrained without changing the data-handling and simulation scripts, and the simulation environment can be adapted to other beamlines by supplying new pointing time series and actuator parameters.

2.2 Delivered artefacts and access

The paper is available as Wang et al. (2026), "Patch-MLP-based predictive control: simulation of upstream pointing stabilization for PHELIX laser system", *Mach. Learn.: Sci. Technol.* **7** 015023, **DOI** 10.1088/2632-2153/ae393d [1]. The software release is archived as Patch-MLP and simulation at https://doi.org/10.14278/rodare.4035 [2], and the PHELIX pointing data are archived as PHELIX laser beam center position (THRILL) at https://doi.org/10.14278/rodare.4033 [3].

The code DOI resolves to all versions of the software release, and the data DOI resolves to all versions of the PHELIX pointing dataset. Both are openly accessible and are intended to support reproduction of the reported simulations and reuse of the workflow on other laser systems.

2.3 Intended use and reuse

The release is intended as a basis for future real-time implementation and for further studies of predictive laser beam stabilisation. The data-loading and visualisation scripts document the expected input format, while the training and prediction notebooks expose the model and simulation

parameters.

Because the model can be retrained in a few minutes with about one minute of new data, the workflow is suitable for repeated application during experimental campaigns or after changes in beamline behaviour. Because the model performance indirectly quantifies predictability of the observed pointing jitter, the training and evaluation workflow **can be used to evaluate whether the given beamline configuration can be stabilized using active means and whether deterministic sources for pointing jitter are present in upstream of the observed position**.

3 Beam pointing analysis

3.1 PHELIX beamline and jitter sources

PHELIX is a high-energy laser system at GSI. Its pre-amplifier bridge is the dominant source of ground-vibration-induced pointing jitter, while the Main Amplifier Sensor provides the experimental pointing diagnostic downstream of the main amplifier chain. The spatial separation between the jitter source and the diagnostic motivates upstream pre-compensation with a tip-tilt mirror.

The analysis treats the horizontal and vertical beam-centre coordinates separately. This is necessary because the dominant disturbance components and the residual jitter can differ between the two axes.

3.2 Pointing data and preprocessing

The dataset contains 111 one-minute files recording the x-y beam-centre position and timestamps, with a sampling interval of 0.244 ms, corresponding to a sampling rate of about 4.1 kHz. The first file is used for training and the remaining 110 files form the test set, covering approximately ten hours of operation.

Acquisition was intermittent, with mostly five-minute intervals and occasional ten-minute gaps. The training and test sets are temporally disjoint, and no test information is used for training or hyperparameter selection. Input windows contain 2000 samples of the two pointing channels and are normalised to the interval [0, 1] with min-max scaling.

3.3 Analysis methods and metrics

Beam-pointing stability is quantified with the root-mean-square (RMS), peak-to-valley (P2V) and full-width-at-half-maximum (FWHM) of the beam-centre distribution. Residual jitter is analysed in the time domain and in the frequency domain using power spectral density estimates. Gain metrics are

defined as the ratio of the compensated residual to the original residual, with values below one indicating improvement.

Allan deviation and cumulative RMS analyses characterise stability over different time scales. Failure segments are investigated with fast Fourier transform analysis to identify spectral components that are absent from the training data.

4 Prediction and compensation system modelling

4.1 Patch-MLP prediction model

The predictor is a patch-based multilayer perceptron designed for short-horizon time-series forecasting. The input sequence is divided into non-overlapping patches of 200 samples, corresponding to 2 x 100 samples for the two axes. Each patch is flattened and linearly projected to a 256-dimensional embedding.

The embeddings are concatenated and processed by an MLP with hidden layers of 512 and 128 neurons and ReLU activations. The output layer produces the predicted horizontal and vertical beam positions at a future time step. The model has about 2,000,000 trainable parameters and is trained with a mean-squared-error loss.

Training uses one minute of data, about 240,000 samples, an 80/20 temporal split and 50 epochs on two NVIDIA V100 GPUs. The prediction offset is five samples, corresponding to a horizon of about 1.22 ms. The offset is chosen to exceed the model inference time plus the system latency, and a sliding-window inference scheme produces a continuous stream of predictions.

The Patch-MLP architecture was selected because it is lightweight and captures local temporal patterns relevant to short-term beam jitter. Transformer-based architectures were avoided because of their higher computational cost and because their permutation-invariant self-attention is less reliable for short-horizon physical time series.

4.2 Compensation loop and mirror dynamics

The predicted pointing error is passed to a discrete-time PID controller whose output drives a piezo-actuated tip-tilt mirror upstream of the pre-amplifier bridge. The mirror is modelled as a second-order mass-spring-damper system with a natural frequency of 200 Hz and a damping ratio of 0.35. The second-order transfer function captures the dominant resonance of a large, for example four-inch, mirror.

The simulated correction is superimposed on the measured beam-centre trajectory to obtain the

corrected signal. The conventional PID loop includes a 0.6 ms system delay, representative of the latency reported for the Apollon ARTAO system. The predictive controller applies the forecast correction in advance, so the simulated response is effectively delay-free.

The prediction horizon is selected so that the forecast covers both the model computation time and the actuator and driver latency. This allows the controller to compensate the phase lag that limits a feedback-only PID loop in the frequency band of interest.

4.3 Simulation workflow and latency budget

The workflow compares three signals: the uncontrolled beam position, the traditional PID response and the MLP-driven PID response. The same measured trajectory is used as input for both controllers, ensuring a fair comparison. The traditional PID is tuned as a conventional feedback loop, while the predictive controller uses the forecast error.

The simulation computes residual metrics for every one-minute segment and aggregates them over the full test set. The model is trained once and then kept fixed for the entire ten-hour evaluation, without retuning on test data. The latency budget combines the measured system delay, the mirror response and the model inference time; the **estimated inference latency of about 0.3 ms leaves sufficient margin for the 1.22 ms prediction horizon**. The workflow is implemented in the released notebooks and can be rerun end-to-end. In addition and infernece pipeline has also been implemented in C++ in order to provide low-latency inference at low latency variance.

5 Results and discussion

5.1 Short-term compensation performance

In one-minute comparisons, both controllers reduce the amplitude envelope relative to the uncontrolled signal, but the MLP-driven PID produces a tighter beam-centre distribution than the traditional PID on both axes. The power spectral density shows lower residual power for the

predictive controller, particularly below 60 Hz.

The improvement is attributed to the prediction offset, which reduces the phase lag of the feedback loop. The predictive controller suppresses dominant residual peaks without introducing new oscillations, indicating compatibility with the mirror dynamics.

5.2 Long-term stability

The fixed model was evaluated on 110 one-minute segments covering approximately ten hours. The predictive controller consistently outperformed the traditional PID over the full test period. Averaged reductions for the x direction were 17.2 % RMS, 10.8 % P2V and 18.2 % FWHM; for the y direction they were 20.1 % RMS, 12.5 % P2V and 21.4 % FWHM. The overall reduction is acheived, by the **model suppressing dominant characteristic modes of the system**. This happens at the cost of random prediction errors increasing the lowest residual pointing jitter amplitudes over classical PID results.

Figure 1: Beam pointing distribution for uncorrected mean vs simulated PID and MLP-driven PID.

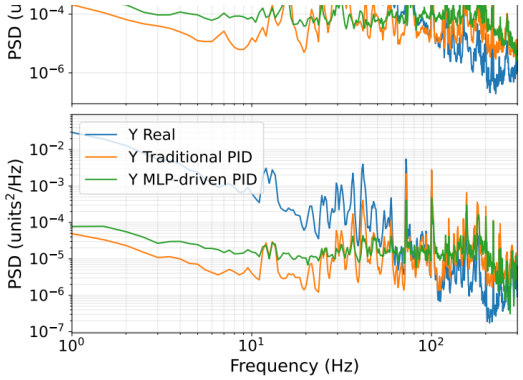


Figure 2: Power spectrum density of beam pointing; uncorrected vs. simulated PID and MLP-driven PID

No upward drift or degradation of the gain curves was observed over the ten-hour test. The model can be retrained in a few minutes with about one minute of new data if the system behaviour changes between days or experimental campaigns.

5.3 Failure cases and limitations

A small number of segments showed degraded performance for all three metrics. These segments contain a spectral peak near 67.7 Hz that is absent from the training data and from well-predicted segments. The peak is attributed to a distribution shift caused by an external vibration source, possibly a gantry near the laser system; this interpretation requires confirmation in dedicated follow-up experiments.

The current study is simulation-based. Experimental validation with a characterised steering-mirror

assembly remains future work, and real-time deployment requires access to the raw perturbed tip-tilt signal or a pseudo-open-loop reconstruction of it.

6 Conclusion

D16 delivers an open-source machine learning code for beam stability investigation, together with the associated paper and PHELIX pointing data. The software implements a Patch-MLP predictor and a PID-based compensation loop for upstream laser beam pointing stabilisation.

The beam pointing analysis identifies the pre-amplifier bridge as the dominant jitter source and defines the metrics used to quantify stability. The prediction and compensation modelling captures the predictor, the mirror dynamics and the latency budget of the control loop.

Simulations over ten hours show consistent reductions of the order of 10-20 % relative to conventional PID control, without drift. The release provides a reproducible basis for future experimental validation and real-time implementation at PHELIX and other high-energy laser facilities.

7 References

[1] Wang, J., Ohland, J. B., Chang, Y.-Y., Pandit, V., Bock, S., Okukura, A.-H., Eisenbarth, U., Irman, A., Bussmann, M., Schramm, U., and Kelling, J. (2026), “Patch-MLP-based predictive control: simulation of upstream pointing stabilization for PHELIX laser system“*, Mach. Learn.: Sci. Technol.* **7** 015023, **DOI** 10.1088/2632-2153/ae393d .

[2] Wang, J. (2025). Patch-MLP and simulation [Software]. RODARE. https://doi.org/10.14278/rodare.4035

[3] Wang, J., Ohland, J. B., Okukura, A.-H., and Eisenbarth, U. (2025). PHELIX laser beam center position (THRILL) [Dataset]. RODARE. https://doi.org/10.14278/rodare.4033

8 Annexes

8.1 Data and code availability

The software release is archived at https://doi.org/10.14278/rodare.4035 under CC BY 4.0. The PHELIX pointing dataset is archived at https://doi.org/10.14278/rodare.4033 under CC BY 4.0. The scientific paper is available at https://iopscience.iop.org/article/10.1088/2632-2153/ae393d .

The dataset contains 111 one-minute files with x position, y position and timestamp columns recorded by the PHELIX MAS. The first file is used for training and the remaining files for prediction and evaluation.

8.2 Software contents

The code package contains four files: read.ipynb for reading and inspecting the PHELIX pointing data, MLP+patch.ipynb for defining and training the Patch-MLP model, MLP prediction and simulation.ipynb for applying the model and running the control simulations, and the pretrained model weights used for the reported results.