High Power Laser Science and Engineering, (2025), Vol. 13, e35, 9 pages.
doi:10.1017/hpl.2025.41

RESEARCH ARTICLE

Pointing stabilization of a 1 Hz high-power laser via machine learning

Alessio Amodio†, Dan Wang†, Curtis Berger, Hai-En Tsai, Samuel K. Barber, Jeroen van Tilborg, Alexander Picksley, Zachary Eisentraut, Neel Rajeshbhai Vora, Mahek Logantha, Qing Ji, Russell Wilcox, Qiang Du, and Anthony Gonsalves

Lawrence Berkeley National Laboratory, Berkeley, California, USA

(Received 30 November 2024; revised 9 April 2025; accepted 16 April 2025)

Abstract

High-power lasers are vital for particle acceleration, imaging, fusion and materials processing, requiring precise control and high-energy delivery. Laser plasma accelerators (LPAs) demand laser positional stability at focus to ensure consistent electron beams in applications such as X-ray free-electron lasers and high-energy colliders. Achieving this stability is especially challenging for the low-repetition-rate lasers in current LPAs. We present a machine learning method that predicts and corrects laser pointing instabilities in real-time using a high-frequency pilot beam. By preemptively adjusting a correction mirror, this approach overcomes traditional feedback limits. Demonstrated on the BELLA petawatt laser operating at the terawatt level (30 mJ amplification), our method achieved root mean square pointing stabilization of 0.34 and 0.59 μrad in the x and y directions, reducing jitter by 65% and 47%, respectively. This is the first successful application of predictive control for shot-to-shot stabilization in low-repetition-rate laser systems, paving the way for full-energy petawatt lasers and transformative advances across science, industry and security.

Keywords: deterministic control; laser–plasma accelerator; machine learning; petawatt/1 Hz laser; stabilization; terawatt operation

1. Introduction

High-power lasers have seen remarkable advancements, enabling scientists to explore new and exciting frontiers in research such as particle acceleration, advanced imaging, nuclear fusion, medical therapies, materials processing, astrophysics and defense applications[1]. Among the most promising applications are laser plasma accelerators (LPAs), which offer the potential to significantly reduce the cost and size of accelerators while delivering comparable energy levels[2]. LPAs have a variety of applications including novel light sources, X-ray free-electron lasers (XFELs) and future high-energy colliders. In particular, the stability of laser parameters is crucial for investigating laser–plasma interactions, especially when working with tightly focused, short-pulse laser beams. Beam positional stability is of paramount importance for the drive laser, as any laser pointing instability directly translates into instability in the generated electron beam[3,4]. This challenge intensifies in high-energy LPA configurations with pre-formed plasma waveguides[5], impacting the precision required across a wide range of LPA applications.

Despite advances in LPA drive laser technology, current systems do not meet the stringent requirements of future applications, such as colliders and XFELs, which demand electron beam transverse positional uncertainty to be a fraction of the beam size[6]. Since the stability of LPA-generated electron beams is dictated by the stability of the drive laser[4], laser position instability must be controlled to be a fraction of the laser beam size at the focal point. Given that our laser focus is typically 50–60 μm (full width at half maximum, FWHM)[7], the transverse laser position error must be limited to a few micrometers. Position error at focus translates to pointing error at the focusing optic, so the error is typically reported as an angle referenced to that optic[8,9]. This angle is corrected by tilting a mirror upstream. In our case, the final optics are too large to reposition quickly, so a mirror before magnification is used, increasing the required angular range by the magnification factor. In this paper, we will refer to the measured beam position error

Correspondence to: Q. Du, Lawrence Berkeley National Laboratory, 1 Cyclotron Rd., Berkeley, CA 94720-8099, USA. Email: qdu@lbl.gov
†These authors contributed equally to this work.

© The Author(s), 2025. Published by Cambridge University Press in association with Chinese Laser Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike licence (https://creativecommons.org/licenses/by-nc-sa/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the same Creative Commons licence is used to distribute the re-used or adapted article and the original article is properly cited. The written permission of Cambridge University Press must be obtained prior to any commercial use.

https://doi.org/10.1017/hpl.2025.41 Published online by Cambridge University Press

data in terms of distance, then translate it to angle and also position normalized to beam size at the end. Note that we present the position error or jitter at the target/focus plane in micrometers, then translate this to pointing angle at the off-axis parabolic (OAP) mirror by dividing the position error by the OAP focal length, following a similar definition as in Refs. [8,10], so numbers are comparable. While this work does not yet achieve stabilization at the LPA target, this will be the focus of future work with full-energy operation of the BELLA PW system.

Perturbations within the system, such as mechanical vibrations of mounted optics, thermal variations within individual optics and thermal drifts of the overall laser setup, lead to long-term drifts in the propagating laser field of the drive laser and its instability. Improved performance has been demonstrated in laser systems with passive controls, achieving an outstanding performance value of 1.3 μrad (root mean square (RMS) value) stability at our BELLA PW beamline and 1.5 μrad (RMS value) stability over 90 min at a 200 TW/1 Hz Ti:sapphire laser system[11]. Active feedback control systems have been applied to stabilize the beam, but their effectiveness is limited by the pulse repetition rate. According to the Nyquist–Shannon theorem, the maximum controllable frequency without aliasing is half the sampling rate. Practically, the control bandwidth is 10%–20% of the sampling rate, and the frequency content of fluctuations can exceed 100 Hz in position and 10 Hz in the angular domain, posing significant challenges for low-repetition-rate systems[12].

To date, limited results have been reported regarding actively stabilized, 100-TW-class, less than or equal to 10 Hz laser systems, using high-repetition-rate low-power 'pilot' laser beams as proxies for the main amplified beam. A multi-TW/2 Hz system achieved 2.6 μrad (RMS) pointing stability using an 80 MHz unamplified/pilot beam for feedback sampling[8]. At one of our other facilities, the BELLA hundred terawatt undulator (HTU) experimental beamline[10], similar stability improvements were achieved using a 1 kHz pilot beam, significantly reducing the uncertainty in the generated electron beam, particularly the low-frequency components associated with long-term drift, with an estimated pointing stability of 3 μm position error[3]. In addition, an optimized Fourier filter has been developed and tested on 1 kHz beam corrections[13], and machine learning (ML) is under development for longitudinal focus stabilization[14]. However, challenges persist, especially with the large optical components required for high-power lasers, which introduce significant lag because of their inertia, limiting the control bandwidth. Traditional feedback control corrects errors based on previous observations, but the inherent lag prevents full compensation, causing relatively large shot-to-shot error[15].

In this work, we present an integrated ML-based approach to predict and mitigate system errors in laser stabilization. The initial proof-of-concept for using ML in this context was introduced in Ref. [16], where a neural network was employed as a time-series forecaster for non-amplified laser data from the BELLA HTU experimental setup. This early work demonstrated the potential of ML to address the bandwidth limitations of existing stabilization systems. Building on this foundation, we demonstrate here that ML enables preemptive movement of slow correction mirrors, effectively compensating for system errors and achieving beam stabilization on a shot-by-shot basis. Our tests at the BELLA PW/1 Hz beamline operating at TW/1 Hz achieved sub-microradiant transverse stabilization, marking the first successful implementation of this technique in high-power, low-repetition-rate laser facilities. Our demonstration was conducted at terawatt-level operation (30 mJ amplification); the BELLA laser system is capable of petawatt-level operation, and this work lays the groundwork for future implementations at full energy. To the best of our knowledge, this work represents the first demonstration of shot-to-shot stabilization in such a challenging environment, paving the way for feedback control systems that overcome traditional bandwidth constraints.

2. Beamline setup and free run analysis

The experimental setup at the petawatt BELLA laser facility is illustrated in Figure 1[10]. The BELLA laser can deliver over 40 J of infrared energy per pulse at 800 nm in about 30 fs, achieving a peak optical power exceeding 1.3 PW at a repetition rate of 1 Hz. The system starts with a front-end low-power seed laser operating at 1 kHz. This beam is then passed through a series of amplifiers powered by 1 Hz pump lasers, producing an amplified 1 Hz beam along with an unamplified 1 kHz pilot beam.

In the laser amplification chain, there is a telescope between each stage to expand the beam size as the energy increases, maintaining fluence lower than the optics' damage threshold. The beam size is 70 mm in diameter after the amplifiers and before the compressor.

After amplification, the beam is reflected from a deformable mirror and the 'correction mirror', which is used to make fine adjustments to beam pointing. It is the last mirror that can be used for correction; this mirror, actuated with piezoelectric transducers, is controlled to correct pointing at the target. Its dimensions are 100 mm × 150 mm, using commercially produced mount (4 inches × 6 inches) with closed-loop controls inside, driving an S-340 Piezo Tip/Tilt Platform[17].

The beam is then expanded in a telescope before it enters an adjustable pulse compressor, which allows fine-tuning of the laser pulse duration. After the telescope, the beam diameter is 210 mm (8.27 inches), transported with mirrors of 12–19 inches diameter, which are too heavy to control.

https://doi.org/10.1017/hpl.2025.41 Published online by Cambridge University Press

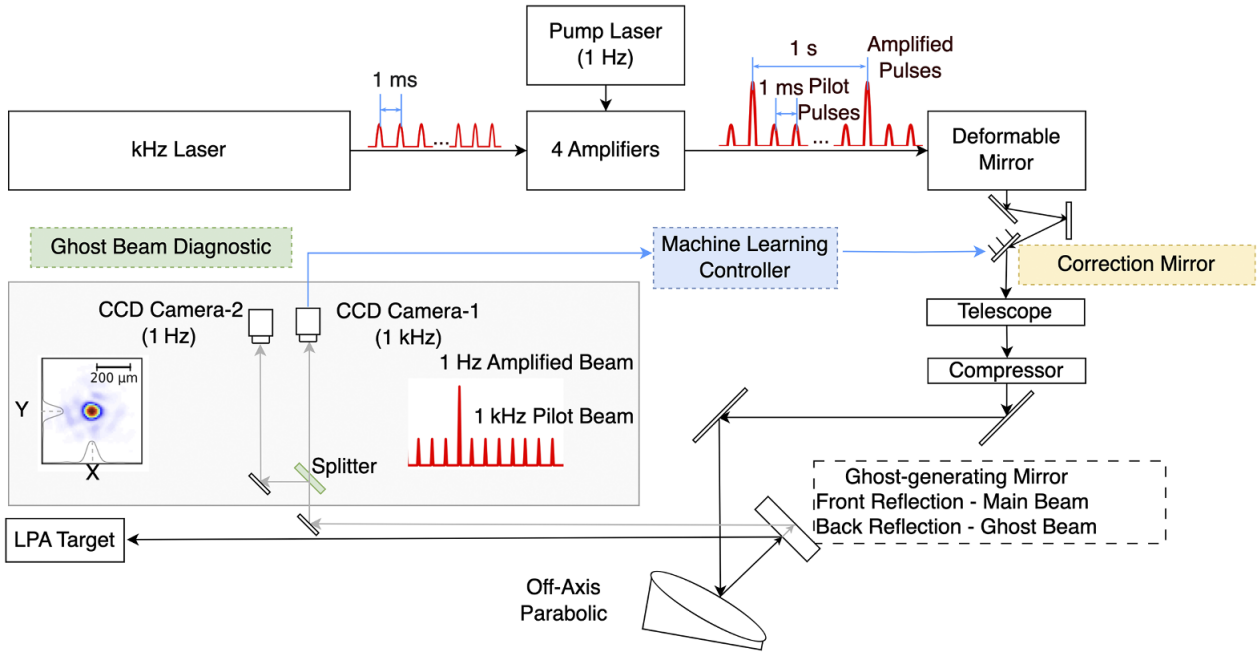


Figure 1. Schematic of the BELLA PW laser system, including the pilot beam diagnostics, correction mirror and focus optics. The setup enables high-resolution monitoring and control of the amplified laser beam.

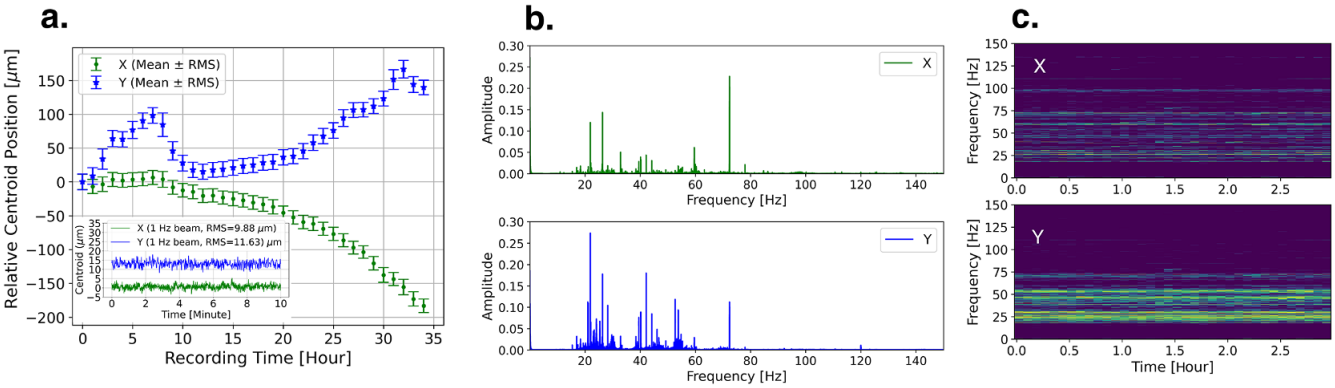


Figure 2. Free run noise analysis of the PW beamline. (a) Centroid position of the pilot beam over 35 h showing long-term drift. (b) Fourier analysis of noise frequency components in the x and y directions based on a 10-min subset of centroid data from the 1 kHz pilot beam. (c) Temporal evolution of frequency components over 3 h.

The laser beam is then directed to an OAP mirror, with focal length of 13.5 m, which focuses it to a spot size of approximately 50–60 μm FWHM. After the OAP, there is a ghost-generating mirror 14 inches in diameter and 2 inches in thickness. The second surface of this ghost-generating mirror is coated for high reflection, providing a 'ghost' beam with enough intensity for the position sensing cameras. The ghost beam diagnostics setup, shown in the inset of Figure 1, includes a 50:50 beamsplitter that creates two copies of the ghost beam for two measurements.

The reflected ghost beam is directed to two charge-coupled device (CCD) cameras triggered at 1 Hz and 1 kHz. The former is used to monitor the position at the focus of the main amplified pulse and the latter to provide the data for ML-based stabilization.

We have found that the pilot beam is highly correlated with the petawatt amplified beam, as both traverse the same optical paths[3,12]. The 1 kHz pilot beam enables the analysis and the precise monitoring of high-frequency errors, ensuring that noise up to 500 Hz is effectively measured for system diagnostics and performance optimization.

Analyzing system data is essential for control modeling, especially for developing data-driven ML solutions. Figure 2(a) shows the beam position data for both the x- and y-axes over a 35-h free (unstabilized) run without any existing feedback that can move the long-term slow drift. In this figure, each point represents the average centroid position. The system stability is quantified by the RMS deviation of the centroid positions, providing a clear measure of position error over time. A slow drift can be observed in the

https://doi.org/10.1017/hpl.2025.41 Published online by Cambridge University Press

average centroid curves for both the x- and y-axes, indicating gradual changes in laser alignment over time. We include an inset in Figure 2(a) showing a 10-min segment of the 1 Hz beam position data in both the x and y directions. This inset highlights the rapid, short-term shot-to-shot fluctuations that are superimposed on the slower drift. Here, we focus on short-term laser position instability, or shot-to-shot jitter, as illustrated in the inset plot, rather than long-term drift, which is low frequency and can be mitigated by slow, existing active feedback systems. The average short-term instability, calculated as the mean of the 34 data points (one data point per hour), shows the average instability of $\sigma_{x,\text{free}} = 10.15\ \mu\text{m}$ and $\sigma_{y,\text{free}} = 11.82\ \mu\text{m}$. It is typical for the x-axis to exhibit slightly better stability than the y-axis in the BELLA system, as well as in other high-power lasers[18,19]. This discrepancy can be attributed to greater vertical (y-axis) vibration of optical mounts in response to ground vibration.

Figure 2(b) displays a typical spectrum obtained through Fourier analysis of a 10-min subset of centroid data from the 1 kHz pilot beam, providing a frequency-domain view of noise characteristics. The results show that the predominant frequency components for both the x- and y-axes are in the tens of hertz range. The horizontal axis (x) exhibits a simpler frequency profile compared to the vertical axis (y), which shows a broader range of frequency components between 20 and 60 Hz, along with an additional peak at 120 Hz. This greater complexity in the y-axis data explains why ML models with training processes tend to perform better in predicting the behavior of the x-axis.

Figure 2(c) further illustrates how the frequency spectrum evolves over a 3-h period, showing the temporal variations in noise frequency components derived in Figure 2(b). This short-time Fourier transform analysis highlights transient changes and periodic noise behavior, helping one to understand fluctuations over time. The data presented here provide a comprehensive assessment of both the spatial and temporal noise characteristics, which is crucial for developing control models, as a foundation for offline training and testing of ML algorithms.

3. Machine learning control diagram

The main limitation on feedback control loop bandwidth is primarily due to the inertia of the correction mirror (100 mm × 150 mm) shown in Figure 1. The mirror assembly includes its own PID (proportional–integral–derivative) controller, which is operated through an external setpoint.

Our previous work[3] on a similar setup, the HTU, demonstrated stability improvements using a commercial feedback controller (ALIGNA-4D provided by TEM Messtechnik) operating at 1 kHz. Reference [8] also shows the effectiveness of active stabilization based on the traditional feedback system.

While both Refs. [3] and [8] show long-term and shot-to-shot oscillation suppression, they highlight the fundamental bandwidth limitations of PID-based feedback due to hardware constraints[20]. In particular, Ref. [3] reported that a large 4-inch correction mirror resulted in a measured feedback bandwidth of only 20 Hz. We measured the performance of a commercially available stabilizer, the Aligna system, at the BELLA HUT facility. The resulting spectrum clearly demonstrates that the Aligna system suppresses frequency components primarily in the 10–20 Hz range. This behavior is consistent with the known bandwidth limitations of traditional feedback control systems. However, Figure 2(b) shows that our laser system exhibits significantly higher frequency components beyond 20 Hz, presenting a major challenge for implementing traditional feedback control effectively. This challenge is further exacerbated by our specific setup, which uses a large 4-inch by 6-inch correction mirror – making it even more susceptible to bandwidth limitations – and serves as a key motivation for the new approach presented in this work.

ML is a widely considered a valuable approach in time-series forecasting topic, which includes prediction of laser pointing oscillation. The control system diagram is shown in Figure 3. In this feedback loop, ML is employed to predict the behavior of the 1 Hz/PW laser based on information from the 1 kHz pilot signal captured by the ghost beam diagnostic camera. This allows the controlled mirror to be repositioned in advance, effectively compensating for any predicted errors and enabling shot-to-shot stabilization of the 1 Hz/PW beam.

The standard feedback method using a 1 kHz pilot beam relies on a simple PID controller, which treats the 1 kHz beam data as the error signal for correction. However, due to the slow response of the mirror (20 ms), the error signal used for correction does not reflect the actual system error when the mirror reaches the corrected position. This 20 ms delay (illustrated in the zoomed-in plots in Figure 3) limits the PID controller to correcting slow drifts, but it makes it incapable of addressing high-frequency components in the shot-to-shot jitters. In the frequency domain, the mismatch between the correction signal and the actual error 20 ms later is primarily influenced by the disturbance frequency sources shown in Figure 2. With a control bandwidth limitation (referenced in Ref. [13]), the PID controller cannot address high-frequency system errors, unlike our ML approach, which is not constrained by such bandwidth limitations.

We evaluated several ML models with different neural network architectures, including long short-term memory (LSTM) networks, simple recurrent neural networks (RNNs), and multi-layer perceptrons (MLPs). LSTM networks are specialized for handling sequential data, such as time series, by learning long-term dependencies and effectively capturing patterns over extended periods. Simple RNNs, while similar, have a more basic architecture and are typically used for shorter sequences due to their limitations in retaining long-term memory. MLPs, in contrast, are traditional feedforward networks consisting of multiple

https://doi.org/10.1017/hpl.2025.41 Published online by Cambridge University Press

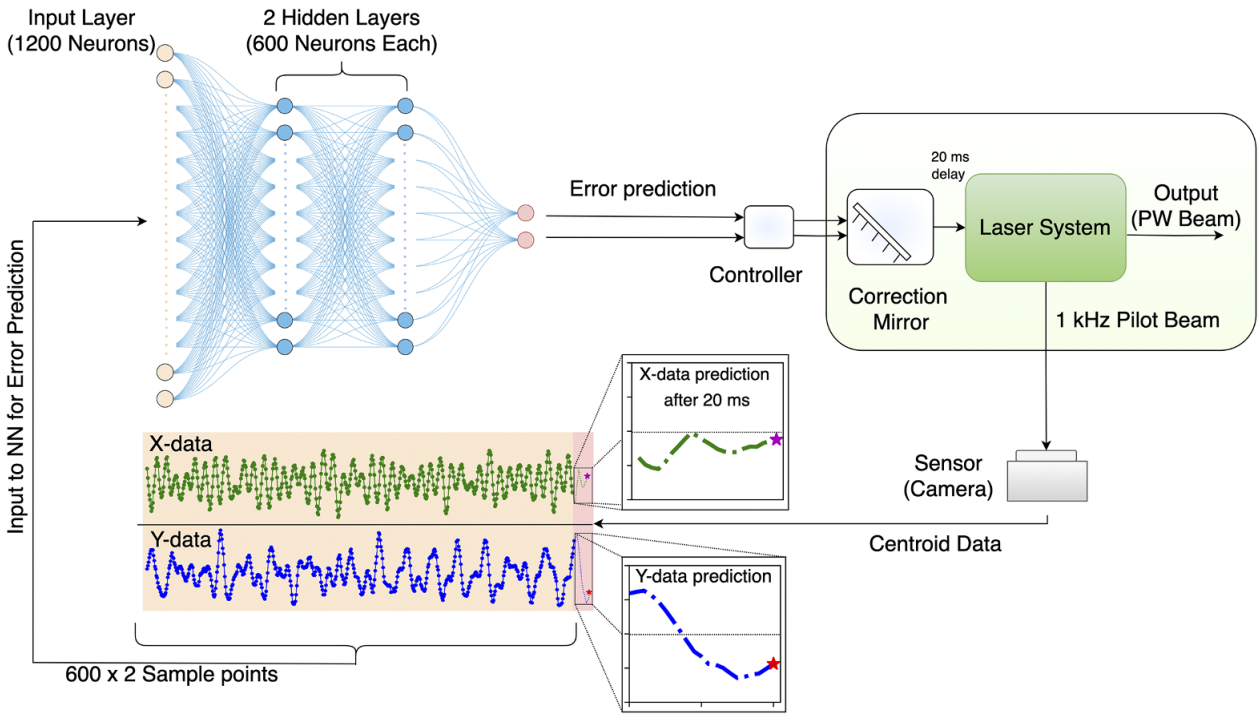


Figure 3. Machine learning control diagram: machine learning feedback loop for predictive control of the BELLA PW laser. The model uses pilot beam data to adjust the correction mirror preemptively, compensating for system noise.

layers of nodes where each layer is fully connected to the next. Although MLPs are generally used for non-sequential data, our study found that they performed comparably to RNNs and LSTM networks on sequential data after appropriate preprocessing and feature engineering. This indicates that MLPs, with the right data preparation, can effectively capture essential patterns in sequential datasets.

The simpler structure of MLPs can be advantageous in terms of computational efficiency and ease of implementation, especially when real-time performance is crucial. The choice of MLP model is driven by its simplicity and ease of integration into a field-programmable gate array (FPGA)-based control loop, which is critical for achieving precise timing in future applications. An important observation from our study is that proper scaling of the training data is essential; the training data must reflect the range of values expected during the testing and correction phases. Hyperparameters refer to the configuration settings of an ML model that are not learned from the data during training but are set before the learning process begins. Examples of hyperparameters include the number of neurons in a neural network layer, the learning rate and the number of training epochs. Tuning these parameters is necessary for optimizing the model's performance.

The chosen layout of the MLP model has 1200 neurons in the input layer (accounting for 600×2 sample points for both the x- and y-axes) and two hidden layers with 600 neurons, optimized through hyperparameter tuning using Optuna[21]. This configuration provides a robust and efficient solution for the control loop, balancing performance and implementation complexity.

The input to the ML model consists of sequential data from the 1 kHz pilot beam in both the x and y directions. The 'input window' refers to the duration of data used for making predictions. In our setup, the input window is 600 ms, but this duration can be adjusted as needed. The 'delay time' is the interval between the last sample point in the input window and the prediction point, which is the arrival time of the PW pulse. The delay time depends on the response time of the controller mirror. For our system, the delay time must be at least 20 ms to ensure the mirror has enough time to reach a stable position. In our case, for the first corrected PW beam, with a delay time of 20 ms and an input window of 600 ms, we collect data from 380 to 980 ms (within a cycle length of 1000 ms). The data are then used to make predictions and adjust the mirror position. After 20 ms, the mirror reaches the desired position, correcting the system noise just as the 1 Hz PW beam is delivered to the target.

We have conducted a comprehensive study on the effects of varying the input window and delay time in the simulation described in Section 4. This analysis helps to determine the optimal values for these parameters, ensuring the precise timing and accuracy needed for effective system noise compensation.

https://doi.org/10.1017/hpl.2025.41 Published online by Cambridge University Press

4. Simulation results

We present ML predictions versus the actual measured values for the focused PW beam position over a 20-min period, as shown in Figure 4(a). This 20-min timeframe includes a total of 1200 pulses, which allows for direct comparison between the experimental recordings and the predicted values. The ML model was trained using data from a 100-s pilot beam run, yielding 100,000 data points in total. To optimize sample collection, we applied a sliding window technique to generate a sufficient number of training samples. Specifically, we labeled the centroid positions from i to $i+599$ ms as the input, combined the x and y coordinates into a single array, and labeled the centroid position at $i+619$ ms as the output to predict. Using this labeled data, the ML model learned the relationship between the input and output in a supervised learning framework.

Then, for testing, we used data collected from 380 to 980 ms in each pulse cycle as input to predict the centroid position at the end of the pulse cycle, which occurs every second. This approach allowed us to evaluate the model's ability to predict the centroid position for each 1 Hz PW pulse. The results showed a strong agreement between the ML predictions and the actual measured values, demonstrating the accuracy of the ML model in offline testing.

The remaining error after correction is calculated as the difference between the predicted and actual values. To evaluate the effectiveness of the ML model, we compared the centroid positions before and after ML correction to determine the error reduction percentage. This percentage is calculated as the difference between the free run (uncorrected) and corrected RMS values, divided by the free run RMS value. In our case, as shown in Figure 4(b), we achieved an error reduction of 77.4% in the x direction and 57.5% in the y direction.

Figure 5 shows the results of a parameter scan. Figure 5(a) illustrates how the input window (the number of sample points in the x and y dimensions) affects control performance, while Figure 5(b) demonstrates the influence of delay time caused by the lag in mirror response. The analysis is based on the PW beamline data shown in Figure 1 and utilizes the control model depicted in Figure 3. From a learning perspective, increasing the duration of the 'input window' provides the ML model with more time-series information, improving prediction accuracy until the input data contains sufficient information, at which point accuracy saturates. Conversely, shorter 'delay times' make it easier for the ML model to predict trends, also enhancing accuracy. The optimal values for these parameters depend on the specific system characteristics, as illustrated by the information in Figure 2. We conducted training and testing across various datasets collected over different time periods, calculating the centroid positions before and after ML correction to assess the percentage reduction in position error.

We evaluated 10 datasets, each representing 1 h of data (the same as in Figure 2). For each of these 10 h, we calculated the mean and RMS values for statistical analysis. The results align with our expectations and reveal several key observations. (1) Both the x- and y-axes exhibit improved performance with a longer input window and a shorter delay time, highlighting the importance of using an adequate amount of time-sequential data for accurate predictions. (2) The x-axis outperforms the y-axis in terms of control precision, likely due to its simpler dynamics, which make

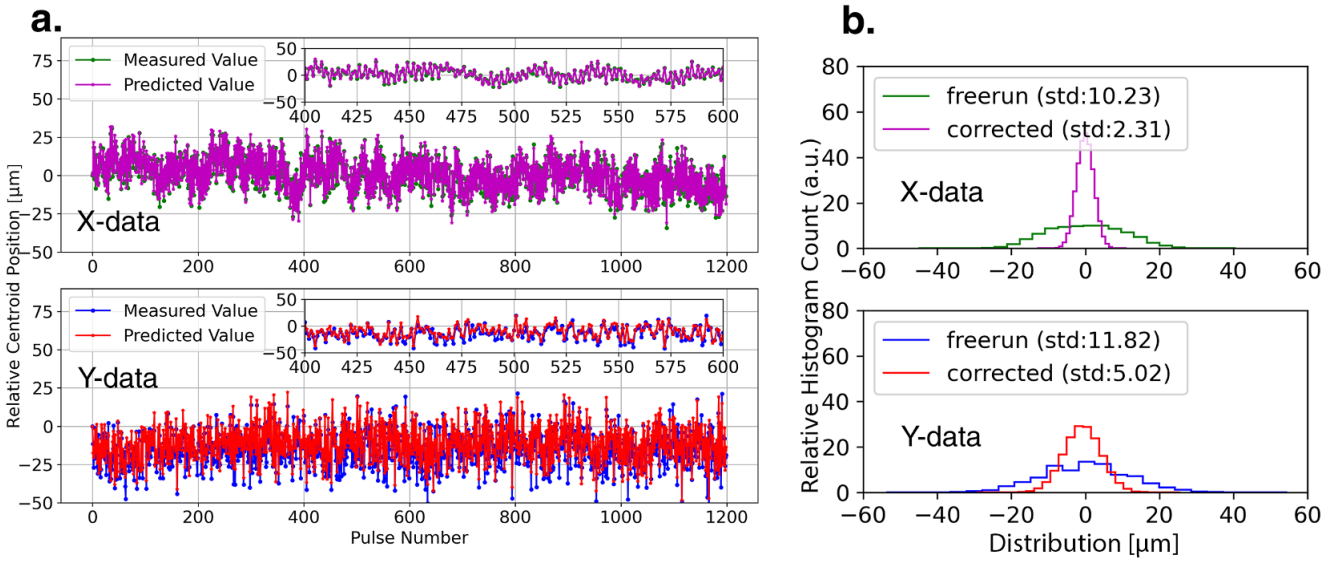


Figure 4. The simulation results show one case of ML model predictions versus measured centroid values for the PW beam, given the dataset of 20 min (1200 data points) of a 1 Hz beam from Figure 2(a). (b) Statistics of the 1200 data points before and after ML correction show jitter reduction of 77.4% in the x direction and 57.5% in the y direction, demonstrating the model's effectiveness in simulated conditions.

https://doi.org/10.1017/hpl.2025.41 Published online by Cambridge University Press

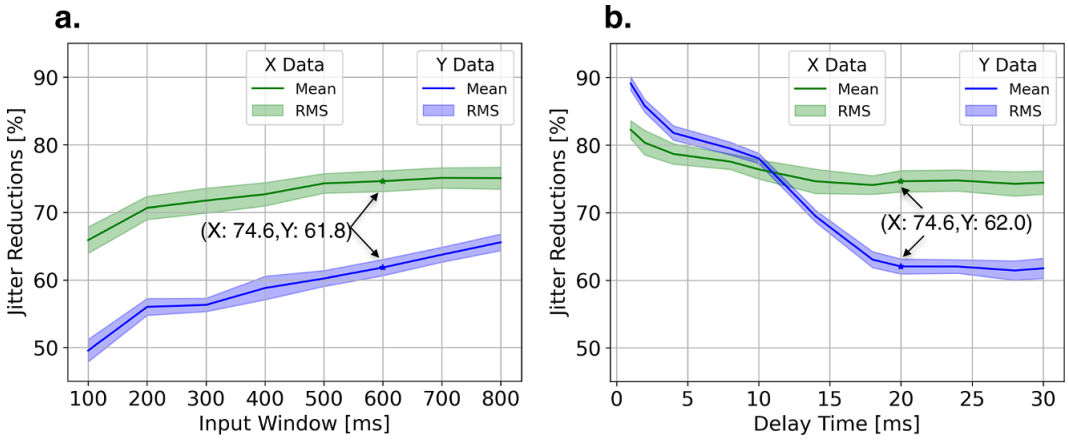


Figure 5. Simulation on a parameter scan with 10 h of data (datasets are the same as in Figure 2): examining the impact of (a) input window and (b) delay time on control performance. Results using experimental parameters (input window of 600 ms and delay time of 20 ms) indicate the average reduction in jitter is 75% in the x direction and 62% in the y direction.

it easier to model and predict. (3) Although the delay time is generally fixed for a specific mirror configuration, there is potential to fine-tune this parameter using advanced timing and synchronization systems, which could enable corrections even before the mirror reaches full stabilization. We plan to investigate this approach further using FPGA technology in future studies. (4) For our experiment, the optimal parameters considering both accuracy and calculation latency were an input window of 600 ms and a delay time of 20 ms, which resulted in an error reduction of 74.6% for the x-axis and approximately 61.8%–62.0% for the y-axis. The consistent performance improvement in the x-axis and the minor variations in the y-axis indicate a reasonable level of randomness in the ML training process and underscore the robustness of our method.

To further compare our ML-based method with the traditional Aligna system implemented at the BELLA HTU setup, we have included additional data shown in Figure S2-2 of the Supplementary Material. This figure presents the centroid distribution from the 1 Hz HTU data, comparing the experimental results with the Aligna system off and the Aligna system on, and the simulated ML-based correction. The simulated ML-based stabilization demonstrates comparable performance to our observations, reducing the standard deviation from $\sigma = 0.65$ pixels to $\sigma = 0.23$ pixels – a reduction of approximately 65%. The improved performance of the simulated ML-based approach compared to the Aligna system on ($\sigma = 0.37$ pixels) further underscores the motivation for pursuing this method.

5. Experimental results and discussion

Experimental results, presented in Figure 6(a), show the effectiveness of our approach over a recording period of approximately 1 h, consisting of 30 min of free run operation followed by 30 min of ML correction. During the free run period, the RMS jitter in the x direction was approximately 13.1 μm, and in the y direction, it was around 14.8 μm. With ML correction enabled, we observed a drop of the RMS values to 4.6 μm in the x direction and 7.9 μm in the y direction. This corresponds to jitter reductions of 64.9% in the x direction and 46.9% in the y direction.

Figure 6(b) shows the measured results in terms of shot-to-shot error compared with the beam size, with the focal spot displayed as a red background. The raw image of the focal spot is also shown in Figure S3 of the Supplementary Material. The centroid distribution is plotted on top of this image as blue dots (free run) and white dots (corrected). In addition, the plot shows the 3σ ellipse representing the beam size (red curve), the 3σ ellipse for the free run centroid distribution (blue curve) and the 3σ ellipse for the corrected centroid distribution (white curve). The standard deviation values, labeled as σ, are provided in the plot legend. For the free run case, the jitter-to-beam-size ratio in the x-direction is 13.1 μm/25.2 μm (0.52), and in the y-direction, it is 14.8 μm/24.4 μm (0.61). After applying ML correction, the ratio in the x-direction is reduced to 4.6 μm/25.2 μm (0.18), and in the y-direction, it becomes 7.9 μm/24.4 μm (0.32). Although the reduction percentage results are slightly lower than those predicted by simulations, this discrepancy is likely due to additional sources of error, such as communication delays between the central processing unit (CPU) and hardware, which make precise timing difficult to achieve.

This is the first instance of shot-to-shot pointing error reduction beyond the bandwidth limitations typical of non-predictive control. The final errors of 4.6 μm in the x direction and 7.9 μm in the y direction correspond to angular deviations at the OAP focusing mirror of 0.34 μrad in the x direction and 0.59 μrad in the y direction, given the OAP focal length of 13.5 m.

The laser energy used in Figure 6 is amplified but not fully, reaching around 30 mJ. This serves as a proof-of-principle for the entire system closed-loop correction. We

https://doi.org/10.1017/hpl.2025.41 Published online by Cambridge University Press

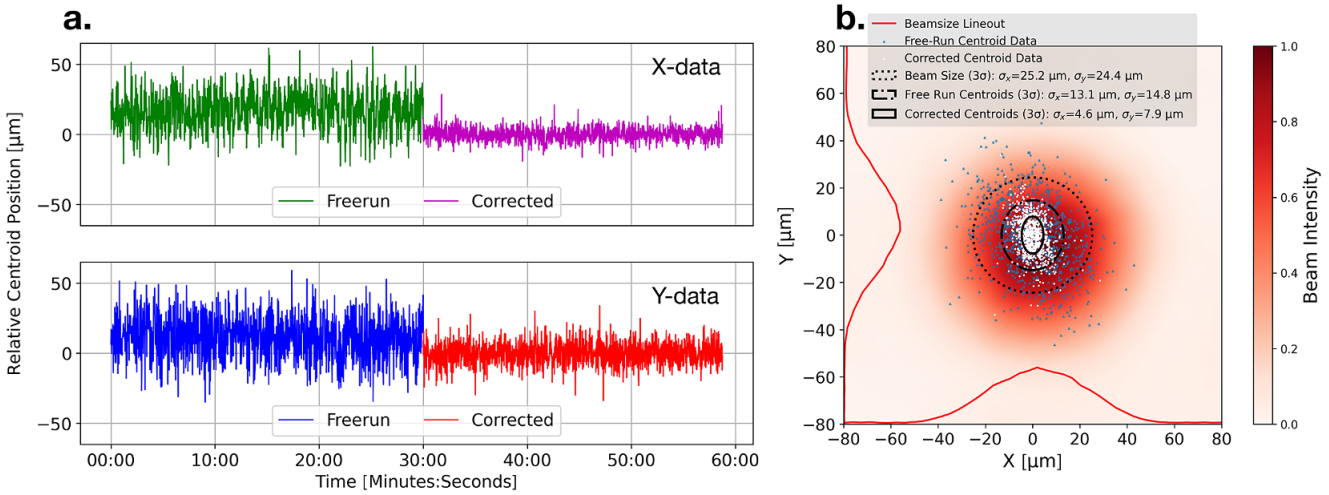


Figure 6. Experimental validation. (a) Comparison of free run and ML-corrected jitter over 1 h in the time domain. (b) Centroid distribution in the two-dimensional (2D) x–y plane compared with the focused laser beam spot as background, in which each dot is the centroid of each pulse. The ellipses represent the σ of each distribution.

plan to apply this correction method in future high-energy (30 J) LPA experiments, with a timeline integrated into the beamline schedule. This ML-based approach is also scalable to setups with additional mirrors. Our future work will involve implementing pointing stabilization at the LPA target, utilizing two mirrors for angle corrections.

6. Conclusion

We have successfully demonstrated the first implementation of ML-based predictive control for shot-to-shot pointing stabilization in a high-power, low-repetition-rate laser system. By leveraging data from the 1 kHz pilot beam, our approach anticipates system errors and preemptively adjusts the correction mirror. This method significantly reduces pointing error in the BELLA PW/1 Hz beamline, achieving sub-microradian stabilization in both the x and y directions. Compared to traditional feedback control methods, our predictive control not only overcomes bandwidth limitations but also provides a robust, scalable solution for future high-power laser applications requiring precise beam stability. The achieved RMS values of instability reductions of approximately 65% in the x direction and up to approximately 47% in the y direction validate the efficacy of our ML model in a real-time, operational environment. This establishes a new approach for laser stabilization in low-repetition-rate systems, paving the way for enhanced performance in applications such as LPAs, XFELs and high-energy colliders.

Supplementary material

The supplementary material for this article can be found at http://doi.org/10.1017/hpl.2025.41.

Acknowledgements

This work is supported by the Office of Science, Office of High Energy Physics, of the US Department of Energy, and the Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory under contract No. DE-AC02-05CH11231.

References

1. C. N. Danson, C. Haefner, J. Bromage, T. Butcher, J.-C. F. Chanteloup, E. A. Chowdhury, A. Galvanauskas, L. A. Gizzi, J. Hein, D. I. Hillier, N. W. Hopps, Y. Kato, E. A. Khazanov, R. Kodama, G. Korn, R. Li, Y. Li, J. Limpert, J. Ma, C. H. Nam, D. Neely, D. Papadopoulos, R. R. Penman, L. Qian, J. J. Rocca, A. A. Shaykin, C. W. Siders, C. Spindloe, S. Szatmári, R. M. G. M. Trines, J. Zhu, P. Zhu, and J. D. Zuegel, High Power Laser Sci. Eng. **7**, e54 (2019).
2. E. Esarey, C. B. Schroeder, and W. P. Leemans, Rev. Mod. Phys. **81**, 1229 (2009).
3. C. Berger, S. Barber, F. Isono, K. Jensen, J. Natal, A. Gonsalves, and J. van Tilborg, Phys. Rev. Accel. Beams **26**, 032801 (2023).
4. A. R. Maier, N. M. Delbos, T. Eichner, L. Hübner, S. Jalas, L. Jeppe, S. W. Jolly, M. Kirchen, V. Leroux, P. Messner, M. Schnepp, M. Trunk, P. A. Walker, C. Werle, and P. Winkler, Phys. Rev. X **10**, 031039 (2020).
5. A. J. Gonsalves, K. Nakamura, J. Daniels, H.-S. Mao, C. Benedetti, C. B. Schroeder, C. S. Tóth, J. van Tilborg, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, E. Esarey, and W. P. Leemans, Phys. Plasmas **22**, 056703 (2015).
6. G. R. White and T. O. Raubenheimer, "Transverse jitter tolerance issues for beam-driven plasma accelerators," Technical Report (SLAC National Accelerator Laboratory, 2019).
7. W. P. Leemans, R. Duarte, E. Esarey, S. Fournier, C. G. R. Geddes, D. Lockhart, C. B. Schroeder, C. Tóth, J.-L. Vay, and S. Zimmermann, AIP Conf. Proc. **1299**, 3 (2010).
8. G. Genoud, F. Wojda, M. Burza, A. Persson, and C.-G. Wahlström, Rev. Sci. Instrum. **82**, 033102 (2011).

9. M. Mori, A. Pirozhkov, M. Nishiuchi, K. Ogura, A. Sagisaka, Y. Hayashi, S. Orimo, A. Fukumi, Z. Li, M. Kado, and H. Daido, Laser Phys. **16**, 1092 (2006).
10. K. Nakamura, H. S. Mao, A. J. Gonsalves, H. Vincenti, D. E. Mittelberger, J. Daniels, A. Magana, C. Toth, and W. P. Leemans, IEEE J. Quantum Electron. **53**, 1200121 (2017).
11. F. Wu, Z. Zhang, X. Yang, J. Hu, P. Ji, J. Gui, C. Wang, J. Chen, Y. Peng, X. Liu, Y. Liu, X. Lu, Y. Xu, Y. Leng, R. Li, and Z. Xu, Opt. Laser Technol. **131**, 106453 (2020).
12. F. Isono, J. van Tilborg, S. K. Barber, J. Natal, C. Berger, H.-E. Tsai, T. Ostermayr, A. Gonsalves, C. Geddes, and E. Esarey, High Power Laser Sci. Eng. **9**, e25 (2021).
13. J. Natal, S. Barber, F. Isono, C. Berger, A. J. Gonsalves, M. Fuchs, and J. van Tilborg, Appl. Opt. **62**, 440 (2023).
14. J. Einstein-Curtis, S. J. Coleman, N. M. Cook, J. P. Edelen, S. Barber, C. Berger, and J. van Tilborg, J. Phys. Conf. Ser. **2420**, 012074 (2023).
15. S. J. D. Dann, C. D. Baird, N. Bourgeois, O. Chekhlov, S. Eardley, C. D. Gregory, J.-N. Gruse, J. Hah, D. Hazra, S. J. Hawkes, C. J. Hooker, K. Krushelnick, S. P. D. Mangles, V. A. Marshall, C. D. Murphy, Z. Najmudin, J. A. Nees, J. Osterhoff, B. Parry, P. Pourmoussavi, S. V. Rahul, P. P. Rajeev, S. Rozario, J. D. E. Scott, R. A. Smith, E. Springate, Y. Tang, S. Tata, A. G. R. Thomas, C. Thornton, D. R. Symes, and M. J. V. Streeter, Phys. Rev. Accel. Beams **22**, 041303 (2019).
16. C. Berger, A. Gonsalves, K. Jensen, J. van Tilborg, D. Wang, A. Amodio, and S. Barber, Nucl. Instrum. Methods Phys. Res. A **1078**, 170561 (2025).
17. https://www.pi-usa.us/en/products/fast-steering-mirrors/s-340-piezo-tip-tilt-platform-300811.
18. Y. Xu, J. Lu, W. Li, F. Wu, Y. Li, C. Wang, Z. Li, X. Lu, Y. Liu, Y. Leng, R. Li, and Z. Xu, Opt. Laser Technol. **79**, 141 (2016).
19. Z. Zhang, F. Wu, J. Hu, X. Yang, J. Gui, P. Ji, X. Liu, C. Wang, Y. Liu, X. Lu, Y. Xu, Y. Leng, R. Li, and Z. Xu, High Power Laser Sci. Eng. **8**, e4 (2020).
20. G. F. Franklin, J. D. Powell, A. Emami-Naeini, and J. D. Powell, *Feedback Control of Dynamic Systems*, Volume 4 (Prentice Hall, 2002).
21. T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (ACM, 2019), p. 2623.

https://doi.org/10.1017/hpl.2025.41 Published online by Cambridge University Press