

# Pointing Stabilization of a 1Hz High Power Laser via Machine Learning

Alessio Amodio<sup>\*1</sup>, Dan Wang<sup>\*1</sup>, Curtis Berger<sup>1</sup>, Hai-En Tsai<sup>1</sup>, Samuel K Barber<sup>1</sup>, Jeroen van Tilborg<sup>1</sup>, Alexander Picksley<sup>1</sup>, Zakchary Eisentraut<sup>1</sup>, Neel Rajeshbhai Vora<sup>1</sup>, Mahek Logantha<sup>1</sup>, Qing Ji<sup>1</sup>, Qiang Du<sup>1</sup>, Russell Wilcox<sup>1</sup>, and Anthony Gonsalves<sup>1</sup>

<sup>1</sup>Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA, United States

## 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 like X-ray free-electron lasers and high-energy colliders. Achieving this stability is especially challenging for 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 Terawatt (30 mJ amplification), our method achieved RMS pointing stabilization of 0.34 and 0.59  $\mu\text{rad}$  in the x and y directions, reducing jitter by 65% and 47%. 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:** Petawatt/1Hz laser, Terawatt Operation, laser-plasma accelerator (LPA), stabilization, machine learning, deterministic control

## 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<sup>[1]</sup>. 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<sup>[2]</sup>. LPAs have a variety of applications including novel light sources, X-ray free-electron lasers, 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<sup>[3,4]</sup>. This challenge intensifies in high-energy LPA configurations with pre-formed plasma

waveguides<sup>[5]</sup>, 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<sup>[6]</sup>. Since the stability of LPA-generated electron beams is dictated by the stability of the drive laser<sup>[4]</sup>, 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 (full width half maximum, FWHM) micrometers<sup>[7]</sup>, 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<sup>[8,9]</sup>. 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 data in terms of distance, then translate to angle and also position normalized to beam size at the

<sup>\*</sup>The first two authors contributed equally to this work.

Correspondence to: Qiang Du, Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA, United States. Email: qdu@lbl.gov

This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

10.1017/hpl.2025.63

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 OAP by dividing the position error by the OAP focal length, following a similar definition as in Ref.<sup>[8,10]</sup>, 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  $\mu\text{rad}$  (RMS value) stability at our BELLA PW beamline and 1.5  $\mu\text{rad}$  (RMS value) stability over 90 minutes at a 200TW/1Hz Ti: sapphire laser system<sup>[11]</sup>. 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, 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<sup>[12]</sup>.

To date, limited results have been reported regarding actively stabilized, 100-TW-class,  $\leq 10$  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  $\mu\text{rad}$  (RMS) pointing stability using 80 MHz unamplified / pilot beam for feedback sampling<sup>[8]</sup>. At one of our other facility, the BELLA hundred terawatt undulator experimental beamline<sup>[10]</sup>, 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  $\mu\text{m}$  position error<sup>[3]</sup>. Additionally, an optimized Fourier filter has been developed and tested on 1kHz beam corrections<sup>[13]</sup>, and machine learning is under development for longitudinal focus stabilization<sup>[14]</sup>. 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<sup>[15]</sup>.

In this work, we present an integrated machine learning (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.<sup>[16]</sup>, where a neural network was employed as a time series forecaster for non-amplified laser data from the BELLA

Hundred Terawatt Undulator (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/1Hz achieved sub-microradian 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 Freerun Analysis

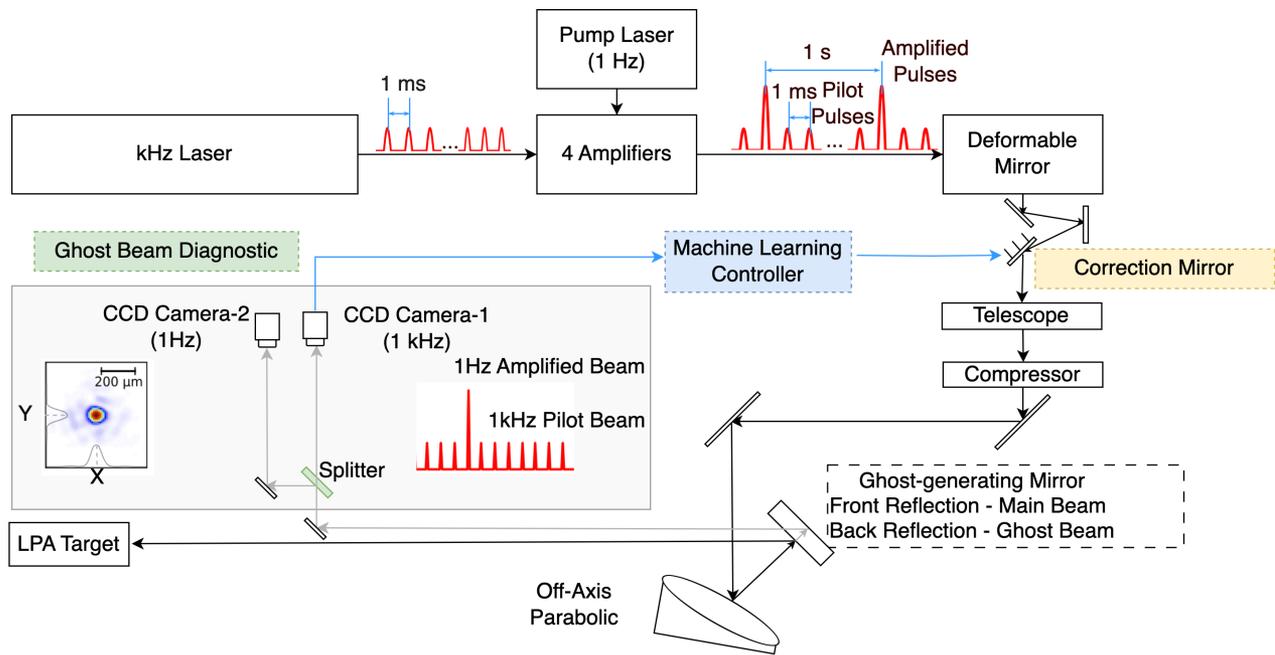
The experimental setup at the Petawatt BELLA laser facility is illustrated in Fig. 1<sup>[10]</sup>. The BELLA laser can deliver over 40J of infrared energy per pulse at 800 nm in about 30 fs, achieving a peak optical power exceeding 1.3 PW at a repetition rate of 1 Hz. The system starts with a front-end 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 reflects 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, which we control to correct pointing at the target. Its dimension is 100 mm  $\times$  150 mm, using commercially-produced mount (4  $\times$  6 inches) with closed-loop controls inside, driving an S-340 Piezo Tip/Tilt Platform<sup>[17]</sup>.

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 inch), transported with mirrors of 12 inch to 19 inch diameters, which are too heavy to control.

The laser beam is then directed to an off-axis parabolic (OAP) mirror, with focal length of 13.5 meters, which focuses it to a spot size of  $\sim 50$ -60  $\mu\text{m}$  FWHM. After the OAP, there is a Ghost-generating mirror 14 inch in diameter and 2 inches thick. The second surface of this Ghost-



**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.

generating mirror is coated to reflect, providing a “ghost” beam with enough intensity for the position sensing cameras. The ghost beam diagnostics setup, shown in the inset of Fig. 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 CCD cameras triggered at 1 Hz and 1 kHz. The former is used to monitor the position at focus of the main amplified pulse and the latter to provide the data for machine learning-based stabilization.

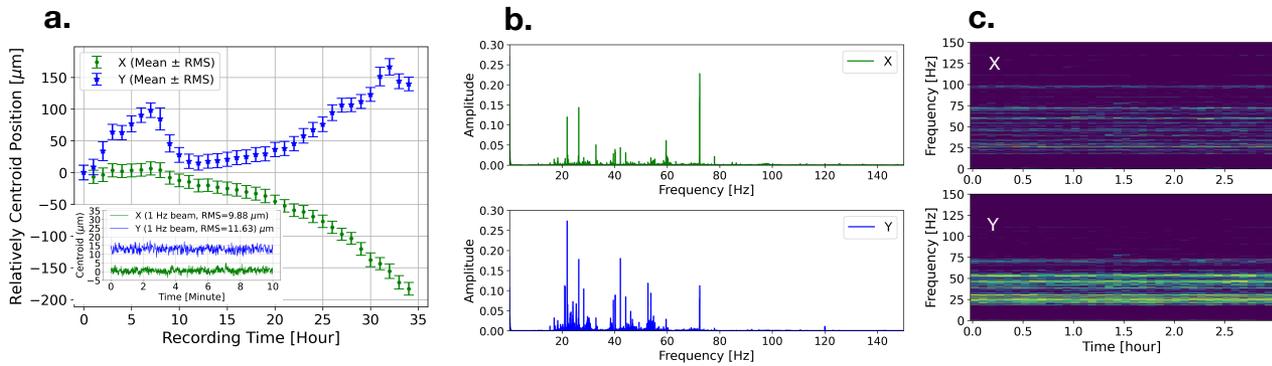
We have found that the pilot beam is highly correlated with the petawatt amplified beam, as both traverse the same optical paths<sup>[3,12]</sup>. The 1 kHz pilot beam enables the analysis and the precise monitoring of high-frequency errors, ensuring that noise up to 500 Hz is effectively measured for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

Analyzing system data is essential for control modeling, especially for developing data-driven machine learning solutions. Figure 2(a) shows the beam position data for both the x and y axes over a 35-hour 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 average centroid curves for both the x and y axes, indicating gradual changes in laser alignment over time. We include an inset in Fig. 2(a) showing a 10-minute segment of the 1 Hz beam position data in both X and Y. 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  $\sigma_{x,free} = 10.15\mu\text{m}$  and  $\sigma_{y,free} = 11.82\mu\text{m}$ . 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<sup>[18,19]</sup>. 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-minute 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 the machine learning models with the 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-hour period, showing the temporal variations in noise frequency components derived in Fig. 2(b). This short-time Fourier transform analysis highlights transient changes and periodic noise behavior, helping to understand fluctuations over time. The data presented here



**Figure 2.** Freerun noise analysis of PW beamline. (a) Centroid position of the pilot beam over 35 hours showing long-term drift. (b) Fourier analysis of noise frequency components in X and Y directions based on a 10-minute subset of centroid data from the 1 kHz pilot beam. (c) Temporal evolution of frequency components over 3 hours.

provides 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 machine learning 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 Fig. 1. The mirror assembly includes its own PID (Proportional-Integral-Derivative) controller, which is operated through an external setpoint.

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

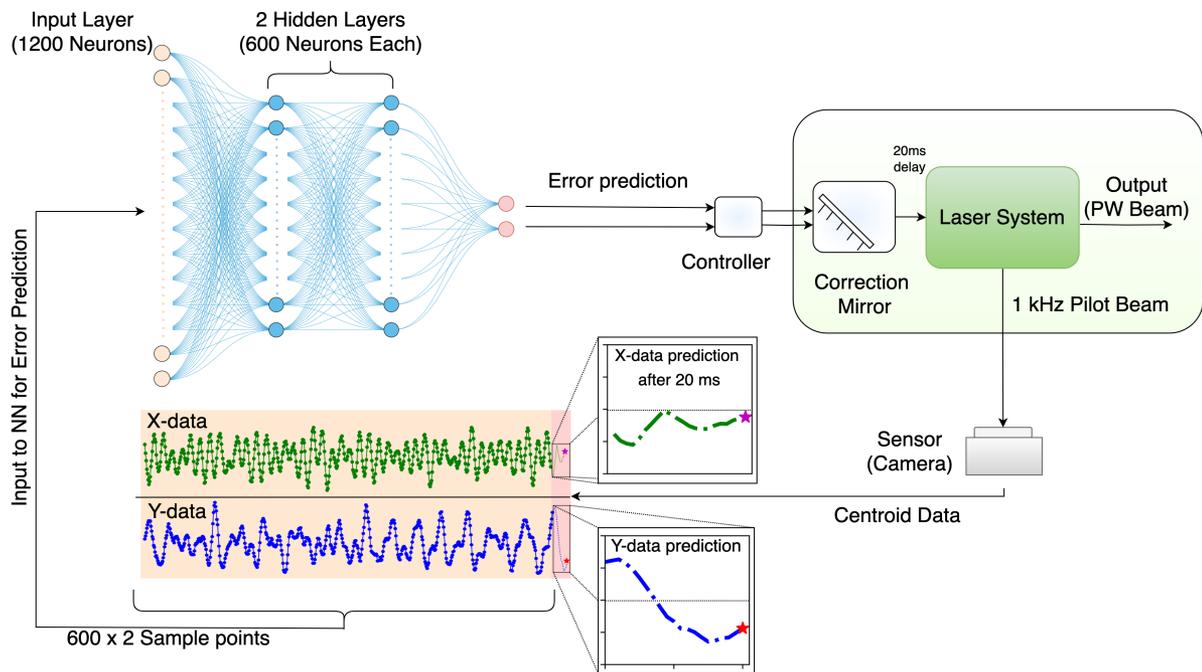
While both Ref.<sup>[3]</sup> and Ref.<sup>[8]</sup> show long-term and shot-to-shot oscillation suppression, they highlight the fundamental bandwidth limitations of PID-based feedback due to hardware constraints<sup>[20]</sup>. In particular, Ref.<sup>[3]</sup> reported that a large 4-inch correction mirror resulted in a measured feedback bandwidth of only 20 Hz. As shown in Supplementary Material Fig.S2-1, the spectrum comparison of experimental data from the BELLA HTU facility includes both the Aligna system on (traditional feedback stabilization) and the Aligna system off. 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, Fig.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.

Machine learning (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 Fig. 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 Fig. 3) limits the PID controller to correcting slow drifts but 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 disturbance frequency sources shown in Fig. 2. With a control bandwidth limitation [referenced in Ref.<sup>[13]</sup>], 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). LSTMs 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 lay-



**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.

ers 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 LSTMs 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 the MLP model is driven by its simplicity and ease of integration into an 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 testing and correction phases.

Hyperparameters refer to the configuration settings of a machine learning 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 1,200 neurons in the input layer (accounting for  $600 \times 2$  sample points for both the X and Y axes) and two hidden layers with 600 neurons, optimized through hyperparameter tuning using Optuna<sup>[21]</sup>. 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 milliseconds, 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 milliseconds 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 milliseconds and an input window of 600 milliseconds, we collect data from 380 milliseconds to 980 milliseconds (within a cycle length of 1000 ms). This data is then used to make predictions and adjust the mirror position. After 20 milliseconds, 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 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.

#### 4. Simulation Results

We present ML predictions versus the actual measured values for the focused PW beam position over a 20-minute period, as shown in Fig. 4. This 20-minute timeframe includes a total of 1,200 pulses, which allows for direct comparison between the experimental recordings and the predicted values. The ML model was trained using data from a 100-second 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  $i$  milliseconds to  $i + 599$  milliseconds as the input, combined the X and Y coordinates into a single array, and labeled the centroid position at  $i + 619$  milliseconds 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 milliseconds to 980 milliseconds 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 freerun (uncorrected) and corrected RMS values, divided by the freerun RMS value. In our case, as shown in Fig. 4(b), we achieved an error reduction of 77.4% in the X direction and 57.5% in the Y direction.

Fig.5 shows the results of a parameter scan. Fig.5(a) illustrates how the input window (the number of sample points in the X and Y dimensions) affects control performance, while Fig. 5(b) demonstrates the influence of delay time caused by the lag in mirror response. The analysis is based on 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 Fig. 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.

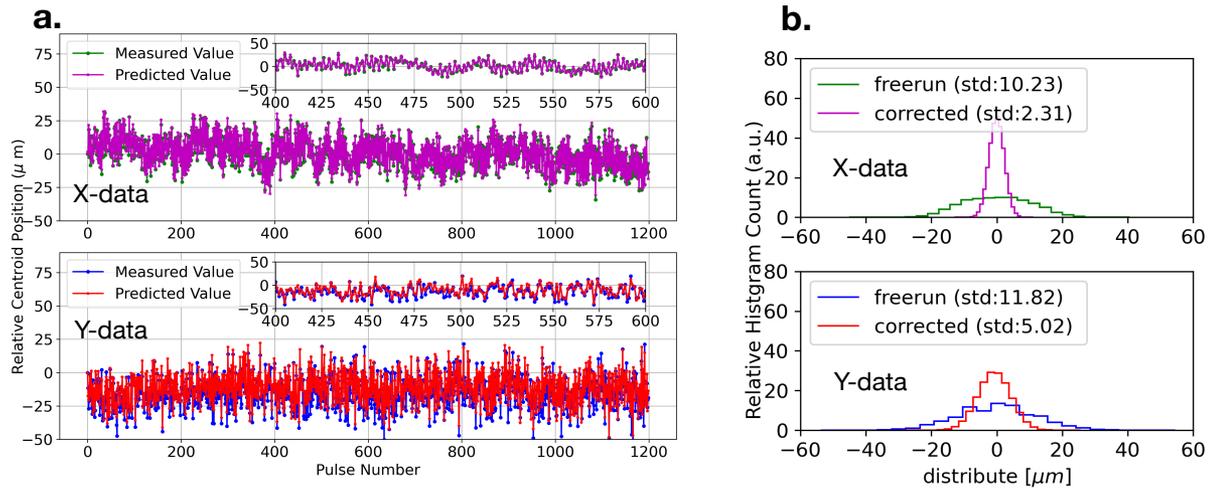
We evaluated 10 datasets, each representing one hour of data (the same as in Fig. 2). For each of these 10 hours, we calculated the mean and root mean square (RMS) values for statistical analysis. The results align with our expectations and reveal several key observations: 1. Both X and Y axes exhibit 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 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 (Field-Programmable Gate Array) technology in future studies; 4. For our experiment, the optimal parameters considering both accuracy and calculation latency were an input window of 600 milliseconds and a delay time of 20 milliseconds, which resulted in an error reduction of 74.6% for the X-axis and approximately 61.8% to 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 Fig. S2-2. This figure presents the centroid distribution from the 1 Hz HTU data, comparing the experimental results with the Aligna system off, 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.

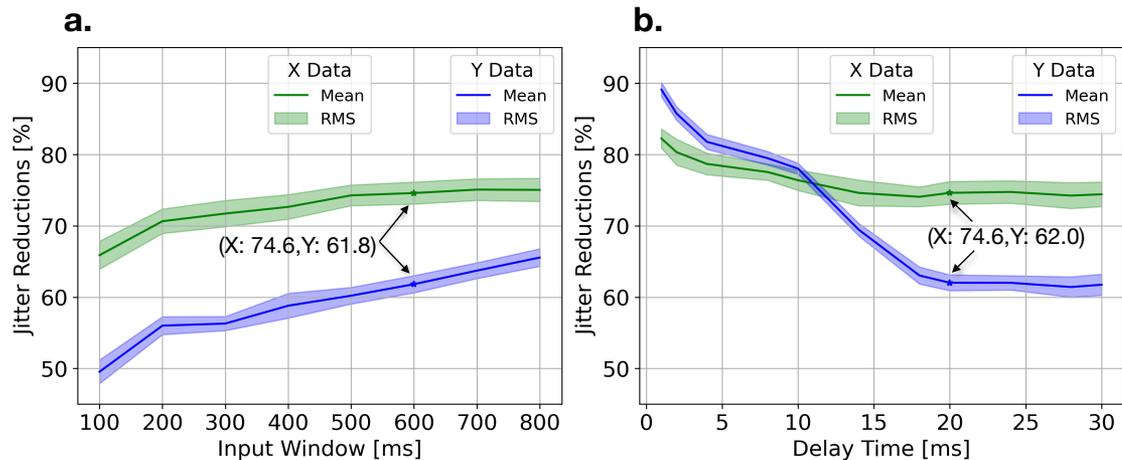
#### 5. Experimental Results and Discussion

Experimental results, presented in Fig. 6(a), show the effectiveness of our approach over a recording period of approximately one hour, consisting of 30 minutes of freerun operation followed by 30 minutes of machine learning (ML) correction. During the freerun period, the RMS jitter in the X direction was approximately  $13.1 \mu\text{m}$ , and in the Y direction, it was around  $14.8 \mu\text{m}$ . With ML correction enabled, we observed a drop of the RMS values to  $4.6 \mu\text{m}$  in X and  $7.9 \mu\text{m}$  in Y. This corresponds to jitter reductions of 64.9% in the X direction and 46.9% in the Y direction.

Fig. 6(b) shows the measured results in terms of shot-to-shot error compared with the beam size, with the focal

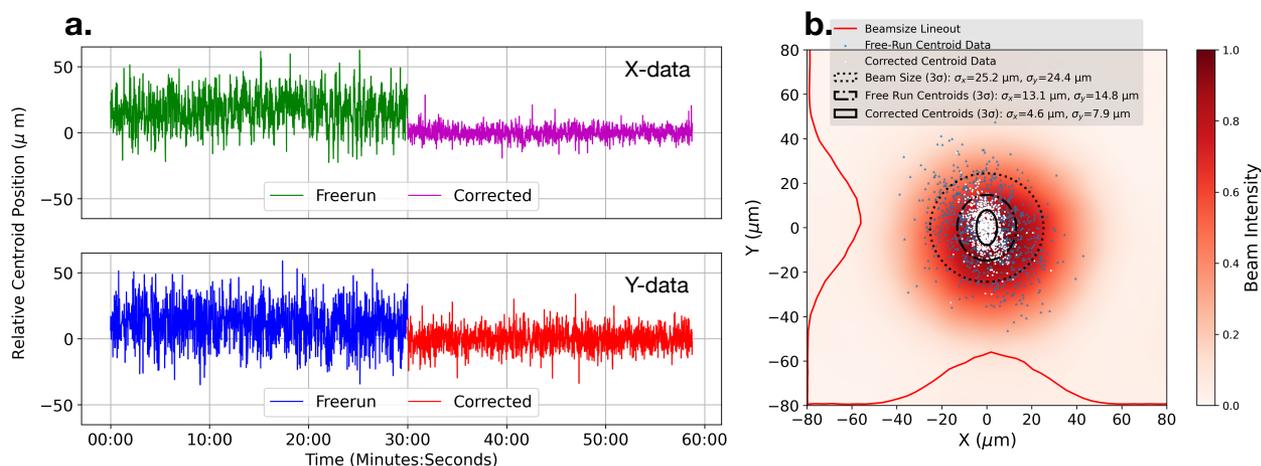


**Figure 4.** Simulation results show one case of ML model predictions versus measured centroid values for the PW beam, given a data set of a 20-min (1200 data points) of 1 Hz beam) from Fig. 2(a). (b) Statistics of the 1200 data points before and after ML correction show jitter reduction of 77.4% in X and 57.5% in Y, demonstrating the model's effectiveness in simulated conditions.



**Figure 5.** Simulation on Parameter Scan with 10 hours data (datasets are same as in Fig. 2). (a): Examining (a) the impact of the 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 X and 62% in Y.



**Figure 6.** Experimental Validation (a). Comparison of freerun and ML-corrected jitter over one hour in time domain, (b). centroid distribution in 2D X-Y plane compared with the focused laser beam spot as background, each dots is the centroid of each pulse. The ellipses represented the  $\sigma$  of each distribution.

spot displayed as a red background. The raw image of the focal spot is also shown in the supplemental material 3. The centroid distribution is plotted on top of this image as blue dots (freerun) and white dots (corrected). In addition, the plot shows the  $3\sigma$  ellipse representing the beam size (red curve), the  $3\sigma$  ellipse for the freerun centroid distribution (blue curve), and the  $3\sigma$  ellipse for the corrected centroid distribution (white curve). The standard deviation values, labeled as  $\sigma$ , are provided in the plot legend. For the freerun case, the jitter-to-beam-size ratio in the x-direction is  $13.1 \mu\text{m}/25.2 \mu\text{m}$  (0.52), and in the y-direction, it is  $14.8 \mu\text{m}/24.4 \mu\text{m}$  (0.61). After applying ML correction, the ratio in the x-direction is reduced to  $4.6 \mu\text{m}/25.2 \mu\text{m}$  (0.18), and in the y-direction, it becomes  $7.9 \mu\text{m}/24.4 \mu\text{m}$  (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 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  $4.6 \mu\text{m}$  in X and  $7.9 \mu\text{m}$  in Y correspond to angular deviations at the OAP focusing mirror of  $0.34 \mu\text{rad}$  in X and  $0.59 \mu\text{rad}$  in Y, given the OAP focal length of 13.5 meters.

The laser energy used in Fig. 6 is amplified but not fully, reaching around 30 mJ. This serves as a proof-of-principle for the entire system close-loop correction. We 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 machine learning-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 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 value of instability reductions of  $\sim 65\%$  in X and up to  $\sim 47\%$  in Y validate the efficacy of our machine learning 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 laser plasma accelerators, X-ray free-electron lasers, and high-energy colliders.

## Acknowledgement

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

## References

1. Colin N Danson, Constantin Haefner, Jake Bromage, Thomas Butcher, Jean-Christophe F Chanteloup, Enam A Chowdhury, Almantas Galvanauskas, Leonida A Gizzi, Joachim Hein, David I Hillier, Danson,

- Colin N and Haefner, Constantin and Bromage, Jake and Butcher, Thomas and Chanteloup, Jean-Christophe F and Chowdhury, Enam A and Galvanauskas, Almantas and Gizzi, Leonida A and Hein, Joachim and Hillier, David I and others. Petawatt and Exawatt class lasers worldwide. *High Power Laser Science and Engineering*, 7:e54, 2019.
2. Eric Esarey, Carl B Schroeder, Esarey, Eric and Schroeder, Carl B and Leemans, Wim P. Physics of laser-driven plasma-based electron accelerators. *Reviews of modern physics*, 81(3):1229–1285, 2009.
  3. Curtis Berger, Sam Barber, Fumika Isono, Kyle Jensen, Joseph Natal, Anthony Gonsalves, Berger, Curtis and Barber, Sam and Isono, Fumika and Jensen, Kyle and Natal, Joseph and Gonsalves, Anthony and van Tilborg, Jeroen. Active nonperturbative stabilization of the laser-plasma-accelerated electron beam source. *Physical Review Accelerators and Beams*, 26(3):032801, 2023.
  4. Andreas R. Maier, Niels M. Delbos, Timo Eichner, Lars Hübner, Sören Jalas, Laurids Jeppe, Spencer W. Jolly, Manuel Kirchen, Vincent Leroux, Philipp Messner, Matthias Schnepf, Maximilian Trunk, Paul A. Walker, Christian Werle, Andreas R. Maier and Niels M. Delbos and Timo Eichner and Lars Hübner and Sören Jalas and Laurids Jeppe and Spencer W. Jolly and Manuel Kirchen and Vincent Leroux and Philipp Messner and Matthias Schnepf and Maximilian Trunk and Paul A. Walker and Christian Werle and Paul Winkler. Decoding sources of energy variability in a laser-plasma accelerator. *Physical Review X*, 10(3):031039, 2020.
  5. A. J. Gonsalves, K. Nakamura, J. Daniels, H.-S. Mao, C. Benedetti, C. B. Schroeder, Cs. Tóth, J. van Tilborg, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, E. Esarey, A. J. Gonsalves and K. Nakamura and J. Daniels and H.-S. Mao and C. Benedetti and C. B. Schroeder and Cs. Tóth and J. van Tilborg and D. E. Mittelberger and S. S. Bulanov and J.-L. Vay and C. G. R. Geddes and E. Esarey and W. P. Leemans. Generation and pointing stabilization of multi-GeV electron beams from a laser plasma accelerator driven in a pre-formed plasma waveguide. *Physics of Plasmas*, 22(5), 2015.
  6. G. R. White G. R. White and T. O. Raubenheimer. Transverse jitter tolerance issues for beam-driven plasma accelerators. Technical report, SLAC National Accelerator Lab., Menlo Park, CA (United States), 2019.
  7. WP Leemans, R Duarte, E Esarey, S Fournier, CGR Geddes, D Lockhart, CB Schroeder, Cs Tóth, J-L Vay, Leemans, WP and Duarte, R and Esarey, E and Fournier, S and Geddes, CGR and Lockhart, D and Schroeder, CB and Tóth, Cs and Vay, J-L and Zimmermann, S. The berkeley lab laser accelerator (bella): A 10 GeV laser plasma accelerator. In *AIP Conference Proceedings*, volume 1299, pages 3–11. American Institute of Physics, 2010.
  8. Guillaume Genoud, Franck Wojda, Matthias Burza, Anders Persson, Genoud, Guillaume and Wojda, Franck and Burza, Matthias and Persson, Anders and Wahlström, C-G. Active control of the pointing of a multi-terawatt laser. *Review of Scientific Instruments*, 82(3), 2011.
  9. M. Mori, A. Pirozhkov, M. Nishiuchi, K. Ogura, A. Sagisaka, Y. Hayashi, S. Orimo, A. Fukumi, Z. Li, M. Kado, M. Mori and A. Pirozhkov and M. Nishiuchi and K. Ogura and A. Sagisaka and Y. Hayashi and S. Orimo and A. Fukumi and Z. Li and M. Kado and H. Daido. Development of beam-pointing stabilizer on a 10-TW Ti:Al<sub>2</sub>O<sub>3</sub> laser system JLITE-X for laser-excited ion accelerator research. *Laser physics*, 16:1092–1096, 2006.
  10. K. Nakamura, H.S. Mao, A. J. Gonsalves, H. Vincenti, D. E. Mittelberger, J. Daniels, A. Magana, Cs. Toth, K. Nakamura and H.S. Mao and A. J. Gonsalves and H. Vincenti and D. E. Mittelberger and J. Daniels and A. Magana and Cs. Toth and W. P. Leemans. Diagnostics, control and performance parameters for the BELLA high repetition rate Petawatt class laser. *IEEE Journal of Quantum Electronics*, 53(4):1–21, 2017.
  11. Fenxiang Wu, Zongxin Zhang, Xiaojun Yang, Jiabing Hu, Penghua Ji, Jiayan Gui, Cheng Wang, Junchi Chen, Yujie Peng, Xingyan Liu, Yanqi Liu, Xiaoming Lu, Yi Xu, Yuxin Leng, Ruxin Li, Fenxiang Wu and Zongxin Zhang and Xiaojun Yang and Jiabing Hu and Penghua Ji and Jiayan Gui and Cheng Wang and Junchi Chen and Yujie Peng and Xingyan Liu and Yanqi Liu and Xiaoming Lu and Yi Xu and Yuxin Leng and Ruxin Li and Zhizhan Xu. Performance improvement of a 200TW/1Hz Ti: sapphire laser for laser wakefield electron accelerator. *Optics & Laser Technology*, 131:106453, 2020.
  12. Fumika Isono, Jeroen van Tilborg, Samuel K Barber, Joseph Natal, Curtis Berger, Hai-En Tsai, Tobias Ostermayr, Anthony Gonsalves, Cameron Geddes, Isono, Fumika and van Tilborg, Jeroen and Barber, Samuel K and Natal, Joseph and Berger, Curtis and Tsai, Hai-En and Ostermayr, Tobias and Gonsalves, Anthony and Geddes, Cameron and Esarey, Eric. High-power non-perturbative laser delivery diagnostics at the final focus of 100-TW-class laser pulses. *High Power Laser Science and Engineering*, 9:e25, 2021.
  13. Joseph Natal, Samuel Barber, Fumika Isono, Curtis Berger, Anthony J Gonsalves, Matthias Fuchs, Natal, Joseph and Barber, Samuel and Isono, Fumika and Berger, Curtis and Gonsalves, Anthony J and Fuchs, Matthias and van Tilborg, Jeroen. High-bandwidth image-based predictive laser stabilization via optimized fourier filters. *Applied Optics*, 62(2):440–446, 2023.
  14. J Einstein-Curtis, SJ Coleman, NM Cook, JP Edelen,

- S Barber, C Berger, Einstein-Curtis, J and Coleman, SJ and Cook, NM and Edelen, JP and Barber, S and Berger, C and van Tilborg, J. Online correction of laser focal position using FPGA-based ML models. In Journal of Physics: Conference Series, volume 2420, page 012074. IOP Publishing, 2023.
15. SJD Dann, CD Baird, N Bourgeois, O Chekhlov, S Eardley, CD Gregory, J-N Gruse, J Hah, D Hazra, SJ Hawkes, Dann, SJD and Baird, CD and Bourgeois, N and Chekhlov, O and Eardley, S and Gregory, CD and Gruse, J-N and Hah, J and Hazra, D and Hawkes, SJ and others. Laser wakefield acceleration with active feedback at 5 Hz. Physical review accelerators and beams, 22(4):041303, 2019.
  16. Curtis Berger, Anthony Gonsalves, Kyle Jensen, Jeroen van Tilborg, Dan Wang, Alessio Amodio, Curtis Berger and Anthony Gonsalves and Kyle Jensen and Jeroen van Tilborg and Dan Wang and Alessio Amodio and Sam Barber. Artificial intelligence time series forecasting for feed-forward laser stabilization. Nuclear Instruments and Methods in Physics Research A, 2024 (submitted).
  17. Physik Instrumente (PI). S-340 piezo tip / tilt platform – fast steering mirrors for active optics. <https://www.pi-usa.us/en/products/fast-steering-mirrors/s-340-piezo-tip-tilt-platform-300811>, 2024.
  18. Yi Xu, Jun Lu, Wenkai Li, Fenxiang Wu, Yanyan Li, Cheng Wang, Zhaoyang Li, Xiaoming Lu, Yanqi Liu, Yuxin Leng, Ruxin Li, Yi Xu and Jun Lu and Wenkai Li and Fenxiang Wu and Yanyan Li and Cheng Wang and Zhaoyang Li and Xiaoming Lu and Yanqi Liu and Yuxin Leng and Ruxin Li and Zhizhan Xu. A stable 200TW/1Hz Ti: sapphire laser for driving full coherent XFEL. Optics & Laser Technology, 79:141–145, 2016.
  19. Zongxin Zhang, Fenxiang Wu, Jiabing Hu, Xiaojun Yang, Jiayan Gui, Penghua Ji, Xingyan Liu, Cheng Wang, Yanqi Liu, Xiaoming Lu, Yi Xu, Yuxin Leng, Ruxin Li, Zongxin Zhang and Fenxiang Wu and Jiabing Hu and Xiaojun Yang and Jiayan Gui and Penghua Ji and Xingyan Liu and Cheng Wang and Yanqi Liu and Xiaoming Lu and Yi Xu and Yuxin Leng and Ruxin Li and Zhizhan Xu. The 1 PW/0.1Hz laser beamline in SULF facility. High Power Laser Science and Engineering, 8, 2020.
  20. Gene F Franklin, J David Powell, Abbas Emami-Naeini, Franklin, Gene F and Powell, J David and Emami-Naeini, Abbas and Powell, J David. Feedback control of dynamic systems, volume 4. Prentice hall Upper Saddle River, 2002.
  21. Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, Akiba, Takuya and Sano, Shotaro and Yanase, Toshihiko and Ohta, Takeru and Koyama, Masanori. Optuna: A next-generation hyperparameter optimization framework. In Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining, pages 2623–2631, 2019.