Check for updates

High Power Laser Science and Engineering, (2026), Vol. 14, e35, 7 pages.
doi:10.1017/hpl.2026.10129

HIGH POWER LASER SCIENCE AND ENGINEERING

RESEARCH ARTICLE

Compensation of carrier envelope phase slip using machine learning

Sung In Hwang[1], Jeong Moon Yang[1], Dongyoon Yoo[1], Jin Woo Yoon[1,2], Bin Kim[1,3], Seong Ku Lee[1,2], Kyung Taec Kim[1,3], and Jae Hee Sung[1,2]

[1]*Center for Relativistic Laser Science, Institute for Basic Science, Gwangju, Republic of Korea*

[2]*Advanced Photonics Research Institute, Gwangju Institute of Science and Technology, Gwangju, Republic of Korea*

[3]*Department of Physics and Photon Science, Gwangju Institute of Science and Technology, Gwangju, Republic of Korea*

(Received 3 September 2025; revised 3 March 2026; accepted 17 March 2026)

Abstract

We demonstrate the effective compensation of carrier envelope phase (CEP) slip in a high-power femtosecond laser using a machine-learning (ML)-based control scheme. The compensation is achieved through fine dispersion tuning guided by a recurrent neural network (RNN) that predicts the temporal evolution of CEP slip, combined with a reinforcement-learning (RL) scheme that determines the optimal corrective actions. With this RNN+RL framework, the integrated CEP noise is reduced by more than a factor of two compared with a conventional proportional–integral–derivative controller. The proposed ML-based control methodology provides a versatile tool for stabilizing and optimizing various parameters in high-power laser systems.

Keywords: carrier envelope phase; high-power laser; machine learning

1. Introduction

Recently, machine learning (ML) has been actively adopted not only in industrial areas, such as quality inspection, predictive maintenance and production process optimization, but also in a wide range of scientific fields, including protein folding structure analysis, particle acceleration and nuclear fusion research[1–6]. In these applications, various data characteristics are learned and identified using ML models to enable accurate prediction and analysis. Likewise, ML has begun to play a significant role in laser science, where it has been applied to reconstruct, predict and enhance various performance characteristics of laser systems[7–18].

The carrier envelope phase (CEP) is a critical parameter in high-power ultrashort lasers used for high-precision measurements. A CEP-locked laser source is necessary for achieving high signal-to-noise ratios in experiments involving atomic dynamics, particularly when probing ultrafast transitions. Because these dynamics often occur on timescales comparable to a single optical cycle, a CEP-locked laser enables direct observation and accurate characterization of such sub-cycle processes[19–24].

CEP locking in high-power femtosecond lasers is typically achieved using dedicated CEP control devices. Generally, a high-power femtosecond laser comprises an oscillator followed by amplification stages. At the oscillator stage, the CEP is first stabilized by adjusting the pump power or finely tuning the dispersion using wedge insertion under proportional–integral–derivative (PID) control[25–29]. With this approach, the CEP slip can normally be reduced to below a root-mean-square (RMS) value of 100 milliradians (mrad). However, during subsequent amplification, the CEP slip often increases significantly because of various environmental perturbations. To mitigate the additional CEP slip introduced during amplification, several dispersion-adjustment strategies have been implemented. Among these, PID control incorporating optical components has been widely employed for CEP stabilization[30,31]. To date, most dispersion-tuning devices used for CEP slip compensation have relied on PID-based control schemes.

From a control-theory perspective, high-performance control is ultimately evaluated by the minimization of the accumulated tracking error between the system state and its desired target over time. In principle, this error can be

Correspondence to: J. H. Sung, Advanced Photonics Research Institute, Gwangju Institute of Science and Technology, Gwangju 61005, Republic of Korea. Email: sungjh@gist.ac.kr

© The Author(s), 2026. 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 licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/hpl.2026.10129 Published online by Cambridge University Press

reduced either by applying corrective actions at an extremely high update rate or by minimizing the instantaneous deviation at each control step. However, practical systems are constrained by actuator dynamics and computational throughput, which fundamentally limit the achievable control-loop frequency. As a result, arbitrarily increasing the update rate is not a viable strategy. Under such constraints, superior control performance relies on the controller's ability to consistently suppress instantaneous deviations even in the presence of complex and time-varying disturbances.

Conventional PID control employs a simple predictive mechanism based on proportional, integral and derivative terms. Its corrective action is determined primarily from the most recent measurements or averaged values over a short temporal window, with the control gains typically tuned through iterative, trial-and-error procedures to mitigate overshoot and reduce response delays. Due to this structure, PID control exhibits limited adaptability to nonlinear behaviors, evolving system dynamics and unanticipated environmental changes. In contrast, the ML-based controller autonomously learns from historical and real-time data, enabling it to extract complex temporal patterns that are difficult to model analytically. By leveraging this data-driven representation, the controller can anticipate future system behavior and optimize control actions proactively. This predictive capability provides a pathway to maintaining minimal instantaneous error despite limited update rates, thereby offering a distinct performance advantage over conventional PID strategies.

In this paper, we demonstrate a significant improvement in CEP slip compensation by employing a dispersion-adjustment device controlled through ML. Unlike previous ML approaches that primarily utilized convolutional neural networks for image-based pulse-width diagnostics or Bayesian neural networks for uncertainty estimation in laser–plasma interactions, the present work targets a fundamentally different task of modeling and controlling the time-dependent CEP slip. Because CEP slip exhibits drift-like temporal behavior, an architecture capable of learning sequential dependencies is essential. Accordingly, we adopt a recurrent neural network (RNN), which offers an effective balance between predictive accuracy and computational efficiency, which is required for real-time control. On top of this predictive model, we incorporate a reinforcement-learning (RL) agent that determines optimal corrective actions based on the RNN-forecasted CEP evolution. This integrated RNN+RL framework enables both accurate prediction of long-term CEP drift and reward-optimized feedback control, leading to faster convergence and enhanced stability relative to previous ML-based methods that lack temporal modeling or dynamic policy learning. As a result, the CEP drift decreased by almost two times compared with a conventional PID control method.

2. Experiment

CEP locking was implemented using a 1-kHz high-power femtosecond laser system delivering 10 mJ, 25 fs pulses at a central wavelength of 800 nm (Femtolasers, Femtopower X), as shown in Figure 1. The CEP value was first stabilized at the oscillator stage using an acousto-optic frequency shifter operated under PID control, yielding a low residual CEP noise below 100 mrad (RMS)[29]. The incident laser beam entering the amplifier had a diameter of approximately 5 mm. Dispersion tuning was achieved using a fused-silica wedge pair with an anti-reflection coating optimized for 700–900 nm. Each wedge had an apex angle of 8°, and one wedge was mounted on a motorized translation stage to provide fine dispersion adjustment. The wedge size of 30 mm × 20 mm ensured that the full beam passed through the wedge pair without clipping. For measurement of the CEP value of an amplified beam, its partial beam with a low energy of approximately 5 μJ was directed to an f-$2f$ interferometry setup. In this setup, the beam was first spectrally broadened after passing through a 1-mm-thick sapphire plate, resulting in an octave-spanning spectrum. The spectrally broadened pulse was then passed through a beta-barium crystal to obtain an f-$2f$ interference signal. For acquisition of the f-$2f$ interference signal used for CEP drift measurement, a spectrometer (Ocean Optics, FLAME-S-VIS-NIR-ES) with a minimum integration time of 20 ms was employed. This spectrometer provides a detection range spanning 350–1000 nm with an optical resolution of 1.5 nm. For the actual CEP retrieval, the spectral interference signal within the 440–580 nm range was utilized. Although the measured CEP value represents an integrated value imposed by the hardware limitations of the spectrometer, it reliably captures the trend of CEP drift driven by the slow variations in the amplifiers[32–35].

The CEP values of high-power femtosecond lasers are strongly affected by environmental fluctuations. When the CEP values were measured over 1.5 h without any CEP slip compensation, a large deviation of 623 mrad (RMS) was observed, primarily due to factors such as ambient temperature variations, as shown in Figure 2(a). Notably, both the CEP drift and the ambient temperature exhibited a 25-min periodicity, corresponding to the operating cycle of the laboratory air-conditioning system. Fourier analysis of the measured CEP trace further revealed that the major contributions to the CEP slip originated from several unknown low-frequency disturbances below 0.5 Hz, as shown in Figure 2(b). These observations indicate that the CEP slip can be substantially reduced if the control bandwidth exceeds 0.5 Hz.

An RNN model was first applied to accurately predict the dispersion-adjustment values required for CEP slip compensation. As a deep learning algorithm designed to handle sequential data, an RNN is well suited for time-series

https://doi.org/10.1017/hpl.2026.10129 Published online by Cambridge University Press

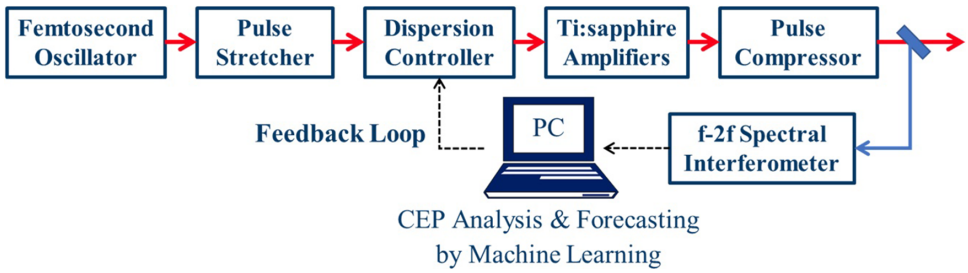


Figure 1. Schematic of the experimental setup for compensation of CEP drift using machine learning.

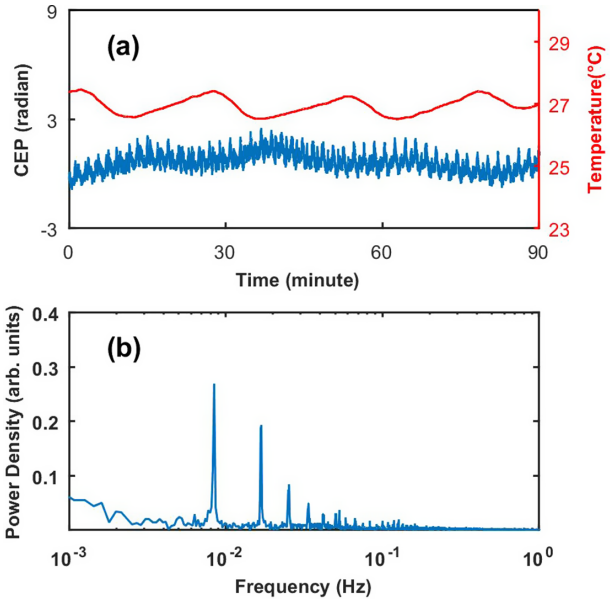


Figure 2. (a) CEP value and ambient temperature measured for 1.5 h and (b) Fourier analysis of the CEP value.

prediction tasks[36]. By processing current inputs alongside information from previous time steps, the network effectively captures temporal dependencies. The model is trained via backpropagation through time, during which the predicted values are compared with the measured values to compute an error. The network's internal weights are then updated to minimize this error. Through repeated iterations of this process, the RNN progressively improves its forecasting capability and ultimately provides highly accurate predictions of the dispersion adjustments needed for CEP stabilization.

The RNN model was customized to compensate for the CEP slip in the 1-kHz high-power femtosecond laser. Its architecture was designed to effectively capture the temporal dependencies inherent in the CEP data. At each time step t, the hidden state is updated according to the following:

$$H_t = \tanh\left(W_{\mathrm{hh}} \cdot H_{t-1} + W_{\mathrm{xh}} \cdot x_t + b_{\mathrm{h}}\right), \tag{1}$$

where the previous hidden-state vector (H_{t-1}) has a size of 200 × 1, and the recurrent weight matrix (W_{hh}) with a size of 200 × 200 enables the network to incorporate information from earlier time steps. The input weight matrix (W_{xh}) with a size of 200 × 1 maps the scalar input (x_t) into the hidden-state space, and the bias vector (b_{h}) with a size of 200 × 1 provides an offset. The hyperbolic tangent (tanh) activation function constrains the hidden-state output within the range from –1 to 1, enabling stable dynamics. Through this recurrent update process, the hidden state serves as the network's internal memory, retaining information from previous inputs and thereby enabling the model to represent the long- and short-term temporal evolution of the CEP slip.

The predicted value (x_{t+1}) was obtained from the processed hidden state (H_t) according to the following:

$$x_{t+1} = W_{\mathrm{hy}} \cdot H_t + b_y, \tag{2}$$

where W_{hy} is the output weight matrix with a size of 1 × 200 and b_y is an output bias. The predicted value was compared with the measured value ($\widehat{x}_{t+1}$) using the mean squared error, represented by the following:

$$\mathrm{Error} = \frac{1}{2}(x_{t+1} - \widehat{x}_{t+1})^2. \tag{3}$$

The RNN model was trained to minimize this error using backpropagation, during which all weight matrices and bias terms ($W_{\mathrm{hh}}, W_{\mathrm{xh}}, W_{\mathrm{hy}}$, b_{h} and b_y) were updated by gradient descent:

$$W \leftarrow W - \eta \frac{\partial \mathrm{Error}}{\partial W},\ b \leftarrow b - \eta \frac{\partial \mathrm{Error}}{\partial b}, \tag{4}$$

with a learning rate $\eta = 0.001$. Through repeated iterations of this optimization process, the RNN progressively refined its internal parameters and achieved increasingly accurate predictions of the required dispersion-adjustment values[37].

The input vector size should be determined carefully to achieve a high prediction accuracy (small error) and short inference time. Generally, increasing the input vector size is one of the most straightforward ways to enhance the performance of an RNN by providing longer temporal data. As indicated in Table 1, a larger input vector size initially led to a positive trend of reducing prediction error. However, beyond the size of 10, the error saturated at approximately 0.016 without further improvement, whereas the required

https://doi.org/10.1017/hpl.2026.10129 Published online by Cambridge University Press

Table 1. Inference time and error sum depending on input vector size.

Input vector size	2	5	10	100
Inference time (ms)	8	12	18	112
Error sum	0.17	0.019	0.016	0.014

computation time increased. Considering that each control cycle must accommodate not only RNN inference but also interpolation and fast Fourier transform (FFT)-based CEP extraction within the 200 ms control window, which corresponds to the 5 Hz operating speed of the motor stage discussed later, the increased processing time associated with larger input vectors imposes practical limitations. Therefore, an input vector size of 10 was selected as the optimal compromise, providing near-minimal prediction error while ensuring that the full processing pipeline operates robustly within the real-time constraints of the feedback loop.

The predictive ability of the RNN model was evaluated using premeasured CEP data. A total of 27,000 CEP samples were acquired over 1.5 h at a sampling rate of five points per second. The dataset was then divided into two subsets, with 80% used for training and the remaining 20% reserved for validation. When collecting 27,000 CEP data points, all measurements were included in the training of the RNN model without the removal of outliers. This approach was deliberately employed to enable the model to learn from the full range of noise and fluctuations present in the real experimental environment. The RNN model was trained on a conventional computing system equipped with a 3.4 GHz CPU and a 32 GB RAM. The training was performed for a fixed number of 1000 iterations, after which the overall error was confirmed to have converged to a value less than 0.02. Figure 3(a) shows that the RNN-predicted data closely follow the measured CEP values, demonstrating that the RNN model has strong predictive capability.

We compensated for the CEP slip using a dispersion adjustment device operated under the control of the trained RNN model. A wedged window installed before the amplification stages served as the dispersion-tuning element, and its insertion length was adjusted according to the values inferred by the RNN model. The maximum attainable control rate was limited to 5 Hz due to the mechanical response of the motors driving the wedge. For a fair comparison, the data acquisition and control update rates of the PID controller were matched to those of the RNN-based controller. The PID controller used for comparison employed a fixed proportional gain (K_p) obtained experimentally from the measured CEP response to controlled dispersion variations.

The time-series traces and power spectral densities of the CEP after compensation are shown in Figure 3. The integrated CEP noise was reduced from 270 mrad (RMS) under PID control to 198 mrad (RMS) under RNN control. These results indicate that the trained RNN model provided more accurate dispersion-adjustment values for CEP slip compensation than the conventional PID controller.

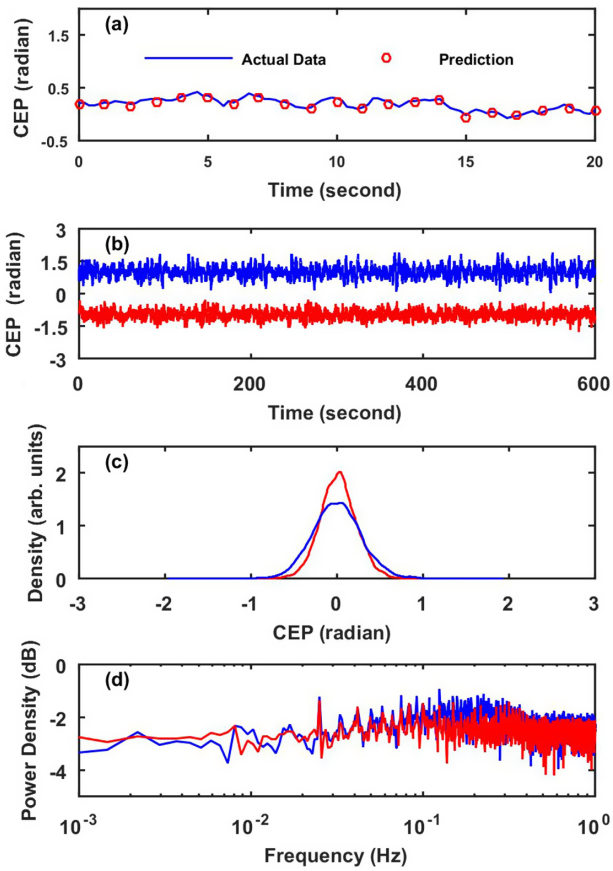


Figure 3. (a) Predicted CEP value from the trained RNN model (red circle) and actual CEP value (blue line), (b) time-series data, (c) power densities and (d) Fourier analysis of CEP compensated by the PID controller (blue) and RNN controller (red).

In addition, an RL model was employed to further refine CEP slip compensation. Generally, an RL agent interacts with its environment and learns to make optimal decisions by maximizing long-term cumulative reward rather than focusing solely on immediate gains[38]. In our scheme, the process works as follows. Firstly, the RNN model predicts a dispersion-adjustment value based on the previous state, represented by the last 10 consecutive CEP data points. Next, the RL agent (the controller for the wedged window) chooses an action: a small perturbation to be added to the RNN-inferred value, with the aim of minimizing CEP fluctuations. The reward scheme is simple: the agent receives a reward of one if the resulting CEP stays within a predefined acceptable range, and zero otherwise. During the initial learning phase, the agent's actions are random and often lead to large CEP variations. However, actions that yield poor outcomes are progressively avoided in subsequent trials. Over time this learning process enables the RL agent to converge on a policy that consistently reduces CEP fluctuations.

https://doi.org/10.1017/hpl.2026.10129 Published online by Cambridge University Press

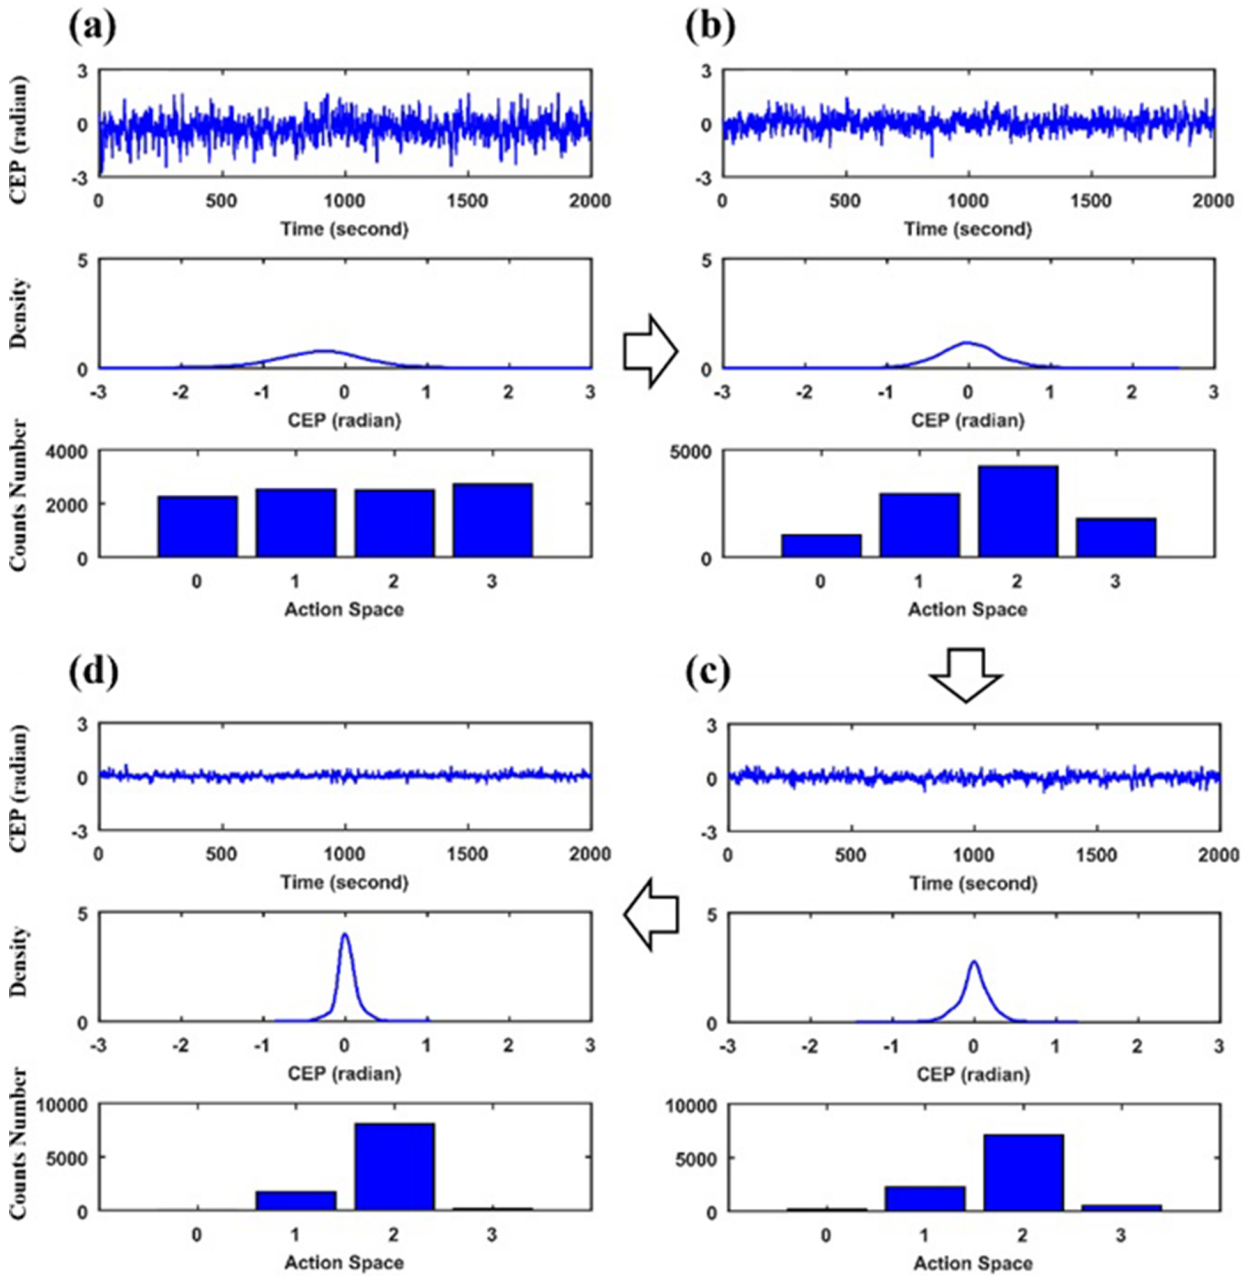


Figure 4. Training process for compensation of CEP slip using the RL model. Time-series CEP data (top), power density of the accumulated CEP (middle) and distribution of action space (bottom) in the initial stage (a), intermediate stage (b), (c) and final stage (d).

The RL agent was trained by interacting with a given environment. The environment was constructed virtually using the RNN model, thereby preventing potential damage to the actual laser system during learning. Specifically, the training was carried out using a vanilla policy gradient algorithm over one million interaction time steps on a conventional computing system, as specified above. The main hyperparameters for training were a learning rate of 0.0002 and a discount factor of 0.98. To reduce computational complexity and facilitate stable policy learning, we employed a discrete action space consisting of four predefined dispersion-adjustment values (–200, –20, 20 and 200 mrad). Although a finer or continuous action space could offer greater control flexibility, it would introduce a substantial increase in the search space and training time. The selected values were empirically determined from the statistical behavior of the CEP residuals following RNN-based prediction. Specifically, the coarse steps (±200 mrad) correspond to the typical maximum deviations requiring correction, whereas the fine steps (±20 mrad), approximately 10% of the coarse magnitude, provide sufficient resolution for residual stabilization near the target. This simplified action space ensures efficient real-time operation while maintaining effective compensation performance.

https://doi.org/10.1017/hpl.2026.10129 Published online by Cambridge University Press

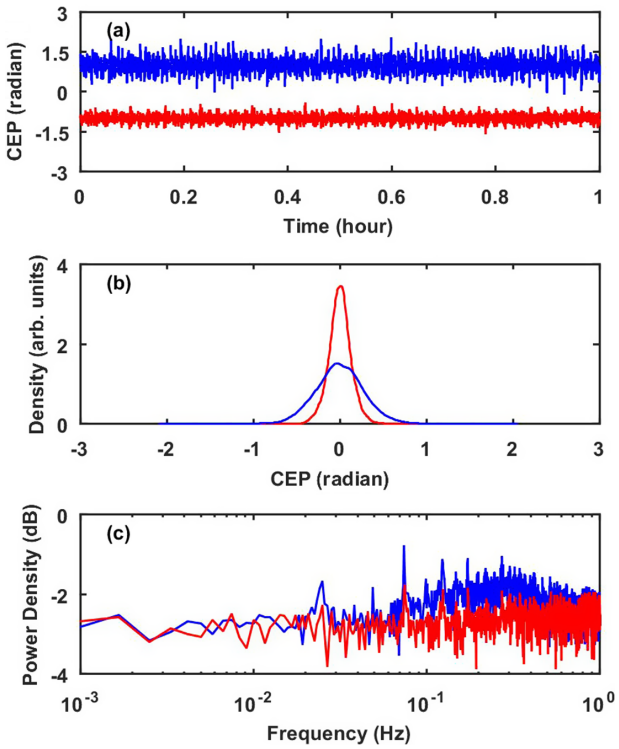


Figure 5. (a) Time-series data, (b) power densities and (c) Fourier analysis of the CEP after compensation of CEP slip using PID control (blue) and the RL model combined with the RNN model (red).

The agent learned to select the optimal actions by leveraging its accumulated experience. The initial permitted CEP range was set broadly, from –2 to 2 rad, allowing the agent to explore widely and obtain rewards easily. As training progressed and the agent's performance improved, the permitted range was gradually narrowed, eventually reaching from –0.1 to 0.1 rad. Once the agent consistently achieved high rewards within this final range, its policy was saved as the optimal policy. Figure 4 shows the progress of the compensation of the CEP slip during the training of the RL agent. In the initial stage, the agent followed a random policy, resulting in uniformly distributed actions and a wide CEP distribution. As the training progressed, the agent developed a preference for specific actions that led to improved outcomes, ultimately producing a significantly narrower CEP distribution.

The RL model with the best policy, combined with the RNN model, was applied to minimize the CEP noise. As shown in Figure 5(a), the RL achieved a noticeably smaller CEP deviation compared with the PID controller. The integrated CEP noise was reduced to 130 mrad (RMS), which is approximately half of the 270 mrad (RMS) obtained with PID control, as shown in Figure 5(b). These results demonstrate that CEP slip can be effectively compensated by first generating accurate dispersion-adjustment predictions through the RNN model and then refining the control action through the RL model to optimize future performance.

3. Conclusion

In conclusion, we have demonstrated a new ML-based approach for dispersion adjustment to compensate for CEP slip in a high-power femtosecond laser. By employing an RNN model for CEP prediction together with an RL agent for optimal control, the system achieved fine dispersion tuning that reduced the CEP noise to less than half of that obtained with conventional PID control. These experimental results clearly show that ML-driven control offers an effective strategy for stabilizing CEP slip. We expect this approach to serve as a versatile framework for optimizing and fine-tuning a wide range of parameters in laser systems.

Acknowledgements

This study was supported by the Institute for Basic Science (IBS-R038-D1) and the Ultrashort Quantum Beam Facility Operation Program (No. 140011) through the Advanced Photonics Research Institute of the Gwangju Institute of Science and Technology.

References

1. J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, and D. Hassabis, Nature **596**, 583 (2021).
2. J. Kaiser, C. Xu, A. Eichler, A. S. Garcia, O. Stein, E. Bründermann, W. Kuropka, H. Dinter, F. Mayet, T. Vinatier, F. Burkart, and H. Schlarb, Sci. Rep. **14**, 15733 (2024).
3. R. Roussel, A. L. Edelen, T. Boltz, D. Kennedy, Z. Zhang, F. Ji, X. Huang, D. Ratner, A. S. Garcia, C. Xu, J. Kaiser, A. F. Pousa, A. Eichler, J. O. Lübsen, N. M. Isenberg, Y. Gao, N. Kuklev, J. Martinez, B. Mustapha, V. Kain, C. Mayes, W. Lin, S. M. Liuzzo, J. John, M. J. V. Streeter, R. Lehe, and W. Neiswanger, Phys. Rev. Accel. Beams. **27**, 084801 (2024).
4. B. K. Spears, J. Brase, P.-T. Bremer, B. Chen, J. Field, J. Gaffney, M. Kruse, S. Langer, K. Lewis, R. Nora, J. L. Peterson, J. J. Thiagarajan, B. Van Essen, and K. Humbird, Phys. Plasmas **25**, 080901 (2018).
5. J. Seo, Y.-S. Na, B. Kim, C. Y. Lee, M. S. Park, S. J. Park, and Y. H. Lee, Nucl. Fusion. **61**, 106010 (2021).
6. J. Seo, S. Kim, A. Jalalvand, R. Conlin, A. Rothstein, J. Abbate, K. Erickson, J. Wai, R. Shousha, and E. Kolemen, Nature **626**, 746 (2024).
7. A. Döpp, C. Eberle, S. Howard, F. Irshad, J. Lin, and M. Streeter, High Power Laser Sci. Eng. **11**, e55 (2023).
8. A. Gonoskov, E. Wallin, A. Polovinkin, and I. Meyerov, Sci. Rep. **9**, 7043 (2019).
9. B. Z. Djordjević, A. J. Kemp, J. Kim, R. A. Simpson, S. C. Wilks, T. Ma, and D. A. Mariscal, Phys. Plasmas **28**, 043105 (2021).

https://doi.org/10.1017/hpl.2026.10129 Published online by Cambridge University Press

10. M. J. V. Streeter, C. Colgan, C. C. Cobo, C. Arran, E. E. Los, R. Watt, N. Bourgeois, L. Calvin, J. Carderelli, N. Cavanagh, S. J. D. Dann, R. Fitzgarrald, E. Gerstmayr, A. S. Joglekar, B. Kettle, P. Mckenna, C. D. Murphy, Z. Najmudin, P. Parsons, Q. Qian, P. P. Rajeev, C. P. Ridgers, D. R. Symes, A. G. R. Thomas, G. Sarri, and S. P. D. Mangles, High Power Laser Sci. Eng. **11**, e9 (2023).
11. T. Zahavy, A. Dikopoltsev, D. Moss, G. I. Haham, O. Cohen, S. Mannor, and M. Segev, Optica **5**, 666 (2018).
12. Y. H. Zeng, Z. J. He, X. H. Guo, G. Z. Zhu, and X. Zhu, Opt. Lett. **49**, 3741 (2024).
13. P. V. Kolesnichenko and D. Zigmantas, Opt. Express **31**, 11806 (2023).
14. L. Meng, S. Liang, L. He, J. Hu, S. Sun, P. Lan, and P. Lu, J. Opt. Soc. Am. B **40**, 2536 (2023).
15. S. Kleinert, A. Tajalli, T. Nagy, and U. Morgner, Opt. Lett. **44**, 979 (2019).
16. M. Stanfield, J. Ott, C. Gardner, N. F. Beier, D. M. Farinella, C. A. Mancuso, P. Baldi, and F. Dollar, Sci. Rep. **12**, 5299 (2022).
17. K. Shimahara, S. Tani, H. Sakurai, and Y. Kobayashi, Commun. Eng. **2**, 1 (2023).
18. M. Veli, D. Mengu, N. T. Yardimci, Y. Luo, J. Li, Y. Rivenson, M. Jarrahi, and A. Ozcan, Nat. Commun. **12**, 37 (2021).
19. P. B. Corkum, Phys. Rev. Lett. **71**, 1994 (1993).
20. J. J. Macklin, J. D. Kmetec, and C. L. Gordon, Phys. Rev. Lett. **70**, 766 (1993).
21. A. McPherson, G. Gibson, H. Jara, U. Johann, T. S. Luk, I. A. McIntyre, K. Boyer, and C. K. Rhodes, J. Opt. Soc. Am. B. **4**, 595 (1987).
22. Z. Chang, A. Rundquist, H. Wang, M. M. Murnane, and H. C. Kapteyn, Phys. Rev. Lett. **79**, 2967 (1997).
23. N. H. B. Ch . Spielmann, S. Sartania, R. Koppitsch, M. Schnürer, C. Kan, M. Lenzner, P. Wobrauschek, and F. Krausz, Science **278**, 661 (1997).
24. P. Agostini, F. Fabre, G. Mainfray, G. Petite, and N. K. Rahman, Phys. Rev. Lett. **42**, 1127 (1979).
25. T. J. Yu and C. H. Nam, Prog. Quantum Electron. **36**, 541 (2012).
26. J. Lee, Y. S. Lee, J. Park, T. J. Yu, and C. H. Nam, Opt. Express **16**, 12624 (2008).
27. Y. S. Lee, J. H. Sung, C. H. Nam, T. J. Yu, and K.-H. Hong, Opt. Express **13**, 2969 (2005).
28. T. J. Yu, K.-H. Hong, H.-G. Choi, J. H. Sung, Y. S. Lee, and C. H. Nam, Opt. Express **15**, 8203 (2007).
29. F. Lücking, A. Assion, A. Apolonski, F. Krausz, and G. Steinmeyer, Opt. Lett. **37**, 2076 (2012).
30. A. Golinelli, X. Chen, B. Bussière, E. Gontier, P.-M. Paul, O. Tcherbakoff, P. D'Oliveira, and J.-F. Hergott, Opt. Express **27**, 13624 (2019).
31. J.-F. Hergott, O. Tcherbakoff, P.-M. Paul, P. Demengeot, O. Gobert, M. Comte, P. D'Oliveira, F. Lepetit, T. Auguste, F. Salin, R. Lopez-Martens, and J.-F. Ripoche, Opt. Express **19**, 19935 (2011).
32. T. Imran, Y. S. Lee, C. H. Nam, K.-H. Hong, T. J. Yu, and J. H. Sung, Opt. Express **15**, 104 (2007).
33. A. Thai, M. Hemmer, P. K. Bates, O. Chalus, and J. Biegert, Opt. Lett. **36**, 3918 (2011).
34. O. Gobert, N. Fedorov, O. Tcherbakoff, J. F. Hergott, M. Perdrix, F. Lepetit, D. Guillaumet, and M. Comte, Opt. Commun. **285**, 322 (2012).
35. B. E. Schmidt, A. D. Shiner, P. Lassonde, J.-C. Kieffer, P. B. Corkum, D. M. Villeneuve, and F. Légaré, Opt. Express **19**, 6858 (2011).
36. J. L. Elman, Cogn. Sci. **14**, 179 (1990).
37. D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Nature **323**, 533 (1986).
38. R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, 2018).

https://doi.org/10.1017/hpl.2026.10129 Published online by Cambridge University Press