

Online electron beam trajectory correction in free-electron lasers via adaptive dictionary learning

Bowen Zhang^{1,2,3}, Nengyuan Zhang^{1,2,3}, Nanshun Huang^{1,*},
Chao Feng^{1,†} and Zhentang Zhao¹

¹*Shanghai Advanced Research Institute, Chinese Academy of Sciences, Shanghai 201210, China*

²*University of Chinese Academy of Sciences, Beijing 100049, China*

³*Zhangjiang Laboratory, Shanghai 201210, China*



(Received 23 February 2025; accepted 8 April 2026; published 8 June 2026)

During accelerator operation, the beam trajectory in free-electron laser (FEL) facilities frequently drifts from its predefined path due to machine instabilities and environmental factors, thereby degrading beam quality. In this work, we propose an online trajectory correction method based on dictionary learning, which significantly reduces data requirements and allows real-time, continuous updates of the response matrix. Simulation studies under static and drifting conditions demonstrate that the dictionary learning method achieves a root-mean-square error of 0.2 m/rad or lower for response matrix elements, using only ~1000 trajectory-excitation data pairs. Experimental validation at the Shanghai Soft X-ray Free Electron Laser (SXFEL) further demonstrates its robustness, achieving a correction precision better than 0.01 mm and confirming its suitability for online implementation in advanced XFEL facilities.

DOI: [10.1103/4hx5-x1n7](https://doi.org/10.1103/4hx5-x1n7)

I. INTRODUCTION

In modern particle accelerator facilities, including synchrotron light sources, colliders, and x-ray free-electron lasers (XFELs), precise control of the electron beam trajectory is essential for stable operation and high-quality output. Over time, however, various perturbations (such as temperature fluctuations, ground vibrations, equipment aging, and external electromagnetic interference) can cause the beam orbit to gradually deviate from its design path [1]. This drift necessitates continuous orbit correction using steering magnets to maintain the desired trajectory.

Historically, orbit correction in circular accelerators has been studied and implemented since the 1970s [2]. With advancements in beam position monitor (BPM) and corrector magnet technology, modern global orbit correction methods rely on linear response matrices that map corrector strengths to BPM readings. Early algorithms such as MICADO [3] and HARMON laid the foundation for model-based correction, and singular value decomposition (SVD) later became a standard tool for inverting the often ill-conditioned orbit response matrix [4,5]. The SVD-based

approach addresses the underlying least-squares problem by decomposing the response matrix into singular values and corresponding singular vectors, enabling a stable pseudoinversion even when the matrix is nearly degenerate. The effectiveness of SVD in handling ill-conditioned systems has led to its widespread adoption in orbit correction algorithms [6,7].

In a typical orbit-feedback system, an orbit response matrix is either measured directly or derived from a machine model; SVD is then used to compute the corrector magnet excitations (currents) that minimize orbit errors in a least-squares sense. One challenge, however, is that any mismatch between the nominal response matrix and the actual machine response will degrade the performance of the feedback system [8]. Moreover, as machine conditions evolve, an initially measured response matrix can become outdated, necessitating periodic remeasurement and calibration. Such recalibration procedures consume dedicated beam time.

To address this issue, quasiparastic methods have been developed to rapidly measure the response matrix by injecting low-amplitude, controlled excitations that minimally perturb the beam [9,10]. Recently, the recursive least squares (RLS) method was applied to continuously refine the response matrix by leveraging data from the closed-loop orbit-feedback system, eliminating the need for dedicated beam studies [11,12]. While effective, this approach requires many iterations to converge, limiting its real-time applicability to high-repetition-rate large-scale facilities such as storage rings. In parallel, data-driven approaches such as neural networks and reinforcement learning have been explored for

*Contact author: huangnanshun@sari.ac.cn

†Contact author: fengc@sari.ac.cn

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International license*. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

adaptive orbit correction, but their current implementations face practical constraints in real accelerator operations [13,14].

Within the broader context of machine learning, dictionary learning (DL) has emerged as a promising technique for handling inverse problems with structured data. This adaptability makes dictionary learning a compelling choice for refining the orbit response matrix, which may possess complex structure beyond what a simple global decomposition can represent.

Another important consideration is the structural difference between orbit feedback in storage rings and those in linear accelerators like FELs [15]. Single-pass FELs have a lower-triangular response matrix, and it may be ill-conditioned, making SVD correction prone to noise amplification. Furthermore, the low repetition rate (typically 10–100 Hz for normal conducting linacs) and limited beam availability in FEL facilities restrict the amount of data that can be collected for response matrix updates, complicating efforts to maintain an accurate model of the machine over time.

In this work, we propose a novel orbit correction framework that employs dictionary learning to continuously and efficiently refine the response matrix during routine operation at conventional XFELs. In this scheme, the response matrix is iteratively updated to progressively reflect the machine's true, time-varying state. The refined matrix is then used to compute the necessary corrector adjustments, thus enabling an adaptive feedback loop for orbit control. Our approach effectively integrates physics-based modeling with online machine learning: it retains the high accuracy of an initial physics model while gradually incorporating data-driven corrections for slow drifts [16]. Importantly, this method is data-efficient, requiring only modest amounts of new data, and can be deployed without interrupting normal x-ray FEL user operations. In this way, slow hardware parameter drifts in low-repetition-rate linacs can be tracked and compensated in real time, allowing the beam trajectory to remain stable despite gradual changes in machine conditions.

This paper is structured as follows: Sec. II describes the mechanism of trajectory correction and the corresponding algorithm. Section III presents the application of these methods, including numerical simulation results. Experiments conducted at the Shanghai soft x-ray FEL facility (SXFEL) are discussed in Sec. IV, and conclusions are offered in Sec. V.

II. TRAJECTORY CORRECTION IN DICTIONARY LEARNING FRAMEWORK

A. Basic trajectory correction

Generally, to accomplish trajectory correction, four essential elements are required: diagnostic components for the beam, such as BPMs, to measure the beam position; steering magnets, also known as correction magnets, to

adjust the beam trajectory; a reference trajectory, which can be established based on previous experimental data; and an effective algorithm to calculate the necessary magnet excitations.

The response matrix $R^{n \times m}$ reflects the relationship between changes in the trajectory and the excitation of the correction magnets. In a system comprising n BPMs and m correction magnets, we utilize a vector $\Delta x = (\Delta x_1, \Delta x_2, \dots, \Delta x_n)$ to denote the trajectory changes across all BPMs from one time step to the next, and a vector $u = (u_1, u_2, \dots, u_m)$ to denote the excitations of each correction magnet. Consequently, we have

$$\Delta x = Ru + w, \quad (1)$$

where w represents the noise in the system, encompassing both machine noise, including BPM noise and beam instability, and environmental noise, such as ground vibrations and earth tides.

Neglecting measurement noise, the required corrector excitation u for trajectory feedback can be determined from the known response matrix R and the measured orbit deviation Δx . Depending on the matrix dimensions, u is computed as

$$\text{if } (m > n), \text{ then } u = R^+ \Delta x, \quad (2)$$

$$\text{if } (m = n), \text{ then } u = R^{-1} \Delta x, \quad (3)$$

$$\text{if } (m < n), \text{ then } u = (R^T R)^{-1} R^T \Delta x. \quad (4)$$

In Eq. (2), R^+ denotes the pseudoinverse of R . When the response matrix R is degenerate (ill-conditioned), an exact inverse may not exist and the solution may be nonunique. In such cases, we compute R^+ using SVD to obtain a stable correction. If R has full column rank, Eq. (4) corresponds to a least-squares-type expression, yielding a correction that minimizes $\|\Delta x - Ru\|_2$.

In general, standard correction methods rely on a static R (either measured or modeled). However, the machine state drifts over time, causing the actual response matrix to deviate from the model, thereby degrading correction performance. To address this, we propose an iterative approach to update R using dictionary learning.

B. Dictionary learning from noise data

Following Ref. [16], we construct a linear system to estimate the response matrix R by aggregating BPM readings and corresponding corrector excitations over time, in the absence of noise. This is expressed as

$$\Delta x = Ru, \quad (5)$$

where u represents the matrix containing the corrector excitations, and Δx is the matrix of BPM-measured orbit




FIG. 1. Schematic illustration of a dictionary learning algorithm framework for reconstructing the response matrix $R^{n \times m}$ from noisy measurements. The matrix $R^{n \times m}$ is partitioned into overlapping $a \times a$ patches (here $a = 4$), each approximated as a linear combination of atoms from a learned overcomplete dictionary $D^{(a \times a) \times K}$ ($K = 900$). The coefficient tensor $A^{K \times h_b \times w_b}$ encodes the active atoms per patch, with at most T_0 nonzero entries per patch. Reassembly via overlapping patch averaging yields $R_{\text{rebuild}}^{n \times m}$, which inherits the lower-triangular structure of the original $R^{n \times m}$.

changes induced by the correctors. With sufficiently long acquisition, Eq. (5) becomes highly overdetermined. This overdetermination provides a robust basis for accurately estimating the system matrix. The objective of the study is to update the response matrix R , which constitutes an inverse problem.

Dictionary learning, an unsupervised learning technique, learns a data-adaptive dictionary from empirical data that provides structured priors for inverse problems and has been widely applied in image and signal processing, including MRI [17]. The schematic diagram of dictionary learning is shown in Fig. 1. This dictionary comprises a collection of base vectors, known as atoms, where each data sample can be expressed as a linear combination of these atoms. Dictionary learning can extract intrinsic features from data, which is highly beneficial for data representation and processing, and can reduce data requirements. In our study, we will employ dictionary learning to reconstruct the response matrix incorporating the data of trajectory and correction magnets. By leveraging the learned structure of the response matrix, accurate reconstruction can be achieved even with limited sampling.

The goal of dictionary learning is to solve the associated optimization problem

$$\begin{aligned} \min_{D, \Gamma} \sum_{i,j} \|F_{ij}R - D\alpha_{ij}\|_2^2, \\ \text{s.t. } \|\alpha_{ij}\|_0 \leq T_0 \quad \forall i, j. \end{aligned} \quad (6)$$

We partition R into overlapping $a \times a$ patches (each patch contains a^2 elements) to reduce computational complexity while preserving local structure. The operator F_{ij} extracts the (i, j) th patch (flattened to $a^2 \times 1$), i.e., $F_{ij}R = R_{ij}$, where (i, j) indexes the patch position. The dictionary D comprises K atoms, and each atom can be

viewed as an $(a^2 \times 1)$ column vector. For each patch, α_{ij} denotes its coefficient vector with respect to D . Γ denotes the collection of coefficient vectors α_{ij} . T_0 limits the number of active coefficients in each patch by retaining only the T_0 entries with the largest absolute values in each α_{ij} .

As illustrated in Fig. 1, we consider a 24×24 response matrix. From this matrix, we learn an overcomplete dictionary comprising 900 atoms, each of size $(4^2 \times 1)$, while the patch size is likewise $(4^2 \times 1)$. The matrix is partitioned into overlapping patches using a sliding window with horizontal stride $s_h = 1$ and vertical stride $s_v = 1$, yielding a total of $h_b \times w_b$ patches, where:

$$h_b = \left\lceil \frac{n-a}{s_h} \right\rceil + 1 = 21, \quad w_b = \left\lceil \frac{m-a}{s_h} \right\rceil + 1 = 21. \quad (7)$$

Here, h_b and w_b denote the number of patches along the vertical and horizontal axes, respectively—i.e., the spatial grid dimensions of the patch decomposition. Each patch (i, j) (with $i = 1, \dots, h_b$, $j = 1, \dots, w_b$) is then approximated as a linear combination of atoms from the dictionary, encoded by the coefficient vector α_{ij} . These coefficients are stored in 3D tensor $A^{K \times h_b \times w_b}$, where the (i, j) th slice corresponds to the coefficient vector α_{ij} for patch (i, j) . Finally, the reconstructed matrix is obtained by reassembling all patches via overlapping patch averaging: for each element in the final matrix, its value is computed as the weighted average of all patches that cover it, with weights proportional to their contribution in the reconstruction. This process ensures global consistency while preserving local structure.

While the dictionary learning framework encodes the structural prior of the response matrix, a reliable

reconstruction must also incorporate actual measurement data to ensure fidelity to observations. Now we aim to solve the optimization problem

$$\begin{aligned} \min_{R,D,\Gamma} \sum_{i,j} & \left\| F_{ij}R - D\alpha_{ij} \right\|_2^2 + \nu \left\| Ru - \Delta x \right\|_2^2, \\ \text{s.t. } & \left\| \alpha_{ij} \right\|_0 \leq T_0 \quad \forall i, j. \end{aligned} \quad (8)$$

Specifically, the data fidelity term $\left\| Ru - \Delta x \right\|_2^2$ ensures consistency between the predicted output and the observed measurements. The weight ν balances its influence within the objective.

We solve (8) using an alternating minimization (block coordinate descent) scheme. First, with R fixed, it runs inner dictionary-learning iterations to update the structural prior, obtaining coefficient vector α_{ij} and updating dictionary D [18–20]. Then, with α_{ij} and D fixed, updating R reduces to a smooth quadratic objective combining the patch reconstruction term and the data-fidelity term $\left\| Ru - \Delta x \right\|_2^2$, employing the accelerated proximal gradient method [21] to efficiently optimize this subproblem.

The online dictionary learning framework is implemented based on the open-source DICTLEARN package [22], with modifications tailored for beam orbit correction. The iteratively updated response matrix, obtained using dictionary learning, is periodically fed into the orbit feedback algorithm to enhance trajectory correction precision.

III. SIMULATIONS

A. Static system

Simulations were conducted using OCELOT [23] with parameters representative of the SXFEL, as outlined in Table I. These simulations were successfully executed within a FODO structure. The complete experimental setup consisted of six identical FODO cells, as illustrated in Fig. 2. In each cell, a corrector magnet was positioned 0.06 m downstream of the quadrupole magnet, followed by a monitor located 0.06 m from the corrector. The distance between the monitor and the subsequent quadrupole magnet was set at 2 m.

The beta functions of 6 FODO cells are shown in Fig. 3. We take into account both the x and y directions simultaneously. Consequently, our simulations are conducted using 24 corrector and 24 monitor values. Each FODO

TABLE I. Main parameters of the beam used in the simulation.

Beam parameter	Value	Unit
Average energy	1.38	GeV
Initial horizontal coordinate x	1×10^{-4}	m
Initial horizontal coordinate y	-2×10^{-4}	m
Horizontal (generalized) momentum p_x	2×10^{-5}	rad
Vertical (generalized) momentum p_y	-1×10^{-5}	rad




FIG. 2. Schematic representation of a FODO cell configuration. The components, arranged from left to right, include a focusing quadrupole, a BPM, a corrector, a defocusing quadrupole, another BPM, another corrector, and a focusing quadrupole.

cell has horizontal and vertical phase advances of $\mu_x/2\pi = 0.113$ and $\mu_y/2\pi = 0.113$, respectively.

Initially, a static system devoid of slow drift effects was simulated, with fixed quadrupole K values, yet subject to measurement deviations. Through computational simulations, the response matrix, R_{init} was derived at this stage and used as the initial value for the response matrix, corresponding to the matrix either measured or obtained from the model. Subsequently, random perturbations, within a range of 5%, were introduced into the K values of the quadrupole magnets to simulate the actual operating conditions of the machine.

In a FEL facility, downstream correctors are unable to influence upstream BPM readings. As a result, the response matrix for each direction is expected to exhibit a triangular structure. Figure 4 illustrates the response matrices before and after the lattice modifications, along with their difference. The two triangular regions represent the response matrices for the x and y directions, respectively. If the overall response matrix deviates from this triangular pattern, it typically reflects physical sources, such as magnet tilt, misalignments, or solenoidal fields, which introduce x - y coupling. By contrast, BPM calibration errors, measurement noise, and energy jitter do not create true coupling, but may exhibit couplinglike signatures or disrupt the triangular pattern in the measured response matrix. Then, trajectory noise with an rms amplitude of $10 \mu\text{m}$ was introduced into the system, while a round-robin dithering excitation of $10 \mu\text{rad}$ was applied to the correctors. The dictionary was




FIG. 3. Beta functions of 6 FODO cells.




FIG. 4. (a) Initial response matrix. (b) Response matrix with 5% random perturbation added to the K values of quadrupoles. (c) The difference between two response matrices.

initialized with $K = 900$ atoms derived from a discrete cosine transform basis, a commonly used initialization for dictionary-learning methods. The initial coefficient codes were then obtained by solving a linear inverse problem to sparsely represent each patch with a small subset of atoms.

In each iteration, 100 trajectories were collected through 100 consecutive feedback cycles with continuous round-robin dithering applied to the 24 correctors, and the response matrix was subsequently refined using a dictionary-learning algorithm. After each iteration, the updated response matrix served as the initialization for the next iteration. This procedure was repeated for a total of 100 iterations.

To quantify the difference between the response matrix before and after the iterations, the error, $d_T = R_T - R$, is defined as follows:

$$|d_T|_{\text{rms}} = \sqrt{\frac{\sum_{i=1}^n \sum_{j=1}^m |a_{ij} - b_{ij}|^2}{nm}}, \quad (9)$$

where discrepancy $|d_T|_{\text{rms}}$ denotes the root mean square error deviation of matrix elements. Here, T denotes the number of iterations. By analyzing the successive updates to R_T , the accuracy of the response matrix could be assessed through the reduction in the discrepancy parameter.

The algorithm was executed in its default configuration on a dual-socket AMD EPYC 7T83 server and exhibited a computational footprint of approximately 4 CPU cores (observed CPU usage $\sim 400\%$). The total wall-clock time for 100 iterations was approximately 10 min, including data generation and all preprocessing steps. The results are shown in Fig. 5. This figure depicts the evolution of the discrepancy value throughout the iterative process. It can be observed that, under the assumption of a static system, the discrepancy rapidly decreases and stabilizes at a very low level within the first 10 iterations. This behavior indicates that the amount of data required for effective trajectory correction is relatively small. In each iteration, the algorithm consumes a fresh batch of 100 trajectory-excitation

data pairs. Although 100 iterations would require 10,000 such pairs in total, the algorithm converged after approximately 10 iterations, resulting in a total usage of only around 1000 datasets. Figure 6 provides a comparison of the response matrix before and after the iterative process, highlighting the improvements achieved through the algorithm.

The noise level and the magnitude of corrector perturbations have been shown to influence the convergence performance of the algorithm. In this study, we refer to the method developed for the noninvasive correction of electron orbit in a synchrotron radiation light source [11,12]. This approach, based on the recursive least-squares method that iterates once for each new pair of data points collected, enables simultaneous determination of the response matrix and orbit correction in a storage ring.

For the simulations, the same lattice structure, initial response matrix, and perturbations applied to the quadrupole magnet K values were employed. The chosen noise level, perturbation magnitude, and dither amplitude were set at $10 \mu\text{m}$, $20 \mu\text{rad}$, and $10 \mu\text{rad}$, respectively. The results, presented in Fig. 7, demonstrate a strong correlation between




FIG. 5. Variation of the value of $|d_T|_{\text{rms}}$ over 100 iterations. The inset shows the results of the first 10 iterations.




FIG. 6. (a) Response matrix obtained through an iterative dictionary learning algorithm. (b) The difference with the actual response matrix.

the investigated variables, thereby supporting the hypothesis. Additionally, Fig. 8 shows the difference between the response matrix obtained after iterations using the recursive least-squares method and the actual response matrix, highlighting the accuracy of the approach.

The RLS method converges smoothly without oscillations or rebounds because it solves a convex optimization problem at each step, and with good-quality data, it steadily approaches the correct solution. However, the iteration speed of the recursive least-squares method significantly decreases as the noise level and the magnitude of corrector perturbations are reduced, due to its reliance on “learning from noise.”

In contrast, the dictionary learning approach starts from a more structured initial guess of the response matrix. In the early iterations, the dictionary learning step captures useful patterns in the matrix, which leads to a rapid initial drop in discrepancy. As a result, the optimizer can find the right direction more quickly, and convergence is faster at the beginning of iterations. Compared to the RLS method, the dictionary learning approach requires 2 orders of magnitude less data while achieving an accurate reconstruction of the response matrix. This makes it better suited for FEL




FIG. 7. Variation of the value of $|d_T|_{\text{rms}}$ over 10^6 iterations when using RLS method.




FIG. 8. (a) The response matrix after RLS iteration. (b) The difference with the actual response matrix.

devices operating under low-noise conditions and thus a more efficient and practical choice for such systems.

This advantage stems from the underlying sparsity of the representation: in dictionary learning, the sparsity constraint reduces the effective degrees of freedom by activating only T_0 atoms per patch (typically 30 in our case), which improves sample efficiency compared with estimating a full response matrix in RLS.

B. Dynamical system simulation

In practice, XFEL systems often operate under conditions of slow drift. The slow drift of quadrupole strengths in FELs caused by power supply aging or thermal effects is modeled as a linear ramp. To simulate this behavior, we define the initial $\mathbf{K}$ value of all quadrupoles $\mathbf{K}_0$ as changing with the drift coefficient f , resulting in $\mathbf{K}_{\text{real}}$:

$$\mathbf{K}_{\text{real}}(t) = \mathbf{K}_0 \times f(t) \quad (10)$$

$$f(t) = 1 + r_k \times t \quad (11)$$

where r_k is the drift rate term and t is the time parameter. A linear increase of the drift coefficient with time yields a continuously unidirectional drifting system. If the drift rate r_k is set to 1×10^{-5} , then the quadrupole magnets $\mathbf{K}$ values increase by 0.001% after each orbit feedback cycle and corresponding data pair acquisition. After $t = 10^4$ effective correction steps, the cumulative drift reaches 10%.

In time-drifting systems, the orbit correction performance is observed to degrade if the response matrix measured at the initial state is continuously applied without updates. To simulate a system undergoing sustained deterioration, a drift coefficient f is introduced by scaling the $\mathbf{K}$ values of all quadrupoles, assuming an identical initial error in the response matrix. During actual accelerator operation, a gain of less than unity is commonly employed to achieve under-correction:

$$u = \text{gain} \cdot \mathbf{R}^{-1} \Delta x \quad (12)$$

where the gain is set to 0.5 and the orbit is corrected with the response matrix measured at the beginning of the




FIG. 9. Evolution of the rms beam position over 100 iterations, showing both raw trajectory data (gray) and a five-point moving average (red).

iteration. Figure 9 presents the orbit statistics recorded over 100 iterations under conditions of $10 \mu\text{rad}$ dithering excitation and $10 \mu\text{m}$ rms BPM noise, with the response matrix held constant at its initial value. To better visualize the underlying upward trend, the red line represents a five-point moving average of the raw rms beam position data, computed by averaging each point with its two neighbors on either side, which effectively suppresses high-frequency noise while preserving the long-term drift. The smoothed trend reveals a gradual upward drift in orbit deviation, indicating progressive degradation due to uncorrected quadrupole parameter drift. This confirms that even slow, linear drifts, when left uncorrected, lead to measurable beam quality loss over thousands of feedback cycles, highlighting the need for adaptive response matrix updates.

Once 100 data pairs have been collected, the response matrix is updated by the algorithm and orbit correction is performed with the revised matrix. Additionally, trajectory

noise and perturbations were introduced. The trajectory noise was maintained at an rms amplitude of $10 \mu\text{m}$ and the corrector magnets were excited in a round-robin manner with a perturbation of $10 \mu\text{rad}$ rms throughout the simulation. This simulation setup effectively replicated the realistic behavior of XFEL systems under slow drift, noise, and perturbation conditions.

The response-matrix update was carried out for a total of 100 iterations using 10,000 collected data pairs. Figure 10 illustrates the final real matrix and the differences in the response matrices with and without applying the algorithm when $r_k = 1 \times 10^{-5}$, demonstrating that the matrices obtained through iterative processing closely approximate the true response matrix at the final time step.

Alternatively, Fig. 11 intuitively compares the algorithm's performance using discrepancy data under different drift rates. It contrasts the deviation between the real, time-drifting response matrix and its estimated counterpart in two scenarios: one with no updates after the initial measurement, and one with continuous updates via the dictionary learning algorithm. Take $r_k = 2 \times 10^{-6}$ as an example. Initially, the deviation is about 0.21. Without updates, the error increases to about 0.27 monotonically due to the system's progressive degradation. This trend is consistent across all tested drift rates: in the absence of updates (dashed lines), the discrepancy grows monotonically with iteration number, and the growth rate scales with the drift magnitude—higher r_k leads to faster degradation. Conversely, with dictionary learning updates, the deviation remains stable at approximately the initial level throughout. Although the discrepancy exhibits moderate fluctuations when using dictionary learning, it remains consistently lower than the true model mismatch—staying around 0.2 throughout the simulation. This confirms that the method not only effectively suppresses drift-induced error accumulation and maintains correction fidelity over extended operation but also rapidly tracks system drift—making it particularly advantageous in scenarios with faster parameter variations.




FIG. 10. (a) The real response matrix at the end of the last iteration. (b) The difference between the measured response matrix and the real response one. (c) The difference between the response matrix obtained after iterations and the real matrix.




FIG. 11. Discrepancies between the response matrix used for trajectory correction and the actual matrix over 100 iterations, with and without the dictionary learning algorithm under different drift rates: $r_k = 1 \times 10^{-5}$, 5×10^{-6} , 2×10^{-6} , and 1×10^{-6} .

Effective updates to the response matrix also contribute to enhanced orbit stability, and Fig. 12 illustrates the BPM rms values for each iteration when dictionary learning is applied. Compared to Fig. 9, under the same drift rate, the overall orbit rms shows a reduction, while the BPM rms stays consistently low and exhibits a modest decreasing tendency throughout the iteration sequence. This demonstrates that the dictionary learning algorithm enables accurate and robust orbit correction in the presence of parameter drift.

In actual accelerator operation, slow parameter drift is not the only concern; abrupt changes such as a sudden quadrupole failure or power supply fault can also occur and




FIG. 12. Evolution of the rms beam position over 100 iterations, showing both raw trajectory data (gray) and a five-point moving average (red) when using the dictionary learning algorithm.




FIG. 13. Discrepancies between the response matrix used for trajectory correction and the actual matrix over 30 iterations, with and without the dictionary learning algorithm, following a sudden 50% reduction in the K value of the second quadrupole magnet at iteration 14.

present distinct challenges to orbit correction systems. To emulate such a realistic fault scenario while maintaining numerical feasibility, we introduce an abrupt 50% reduction in the K value of a single quadrupole magnet when $t = 1333$ —representing a severe but realistic malfunction that induces a significant and immediate perturbation to the beam dynamics. All other system parameters, including BPM noise level and dithering amplitude, remain unchanged throughout the simulation.

Figure 13 illustrates the rms orbit deviation over 30 feedback iterations. As expected, the uncorrected system—operating with a fixed response matrix—exhibits a large transient spike in orbit error immediately following the fault, followed by sustained degradation that remains elevated for the remainder of the simulation. In contrast, the DL-enabled system demonstrates rapid adaptation: within just three iterations after the perturbation, the rms error recovers to near its pre-fault level and remains stable thereafter. This behavior highlights the algorithm's robustness to abrupt parameter changes and its ability to maintain high-quality orbit correction even under severe, unexpected disturbances.

IV. ONLINE EXPERIMENTS

The simulation results indicate that the dictionary learning algorithm can effectively track response-matrix variations and improve orbit correction performance under model mismatch. To validate this capability experimentally, proof-of-principle tests were carried out at the SXFEL. The experiment was performed in the FODO drift section of the linac, involving 13 BPMs (SBPM30–SBPM42) and 13 pairs of horizontal and vertical corrector magnets.




FIG. 14. Evolution of the mean beam orbit during the online experiment. Panels (a) and (b) show the mean BPM readings in the x and y directions, respectively. The orbit remains nearly constant throughout the experiment, indicating that the closed-loop orbit feedback remains operational during the DL-enabled online matrix update and that the adaptive update does not disrupt routine orbit feedback.

During the experiment, the electron beam energy and repetition rate were maintained at 1.227 GeV and 10 Hz, respectively. The orbit-feedback system was initialized previously under a reference lattice configuration. The lattice was then intentionally changed by adjusting the quadrupole settings, so that the response matrix used by the feedback no longer matched the actual machine response. In this way, a controlled response-matrix mismatch was introduced to emulate realistic optics drift or lattice variation during operation. After the lattice change, the modified optics was kept fixed, while the feedback loop remained closed and the response matrix was periodically updated using the dictionary learning algorithm.

To improve data reliability and suppress stochastic fluctuations, each data point was obtained by averaging 40 consecutive orbit readings. After each algorithm iteration, a 300-s interval was allowed for data accumulation and to mitigate lag effects during continuous corrector operation. The maximum correction kick was limited to 0.1 mrad as a safety bound for routine operation. Owing to experimental time constraints, a total of 15 iterations were performed over 90 min. In each iteration, 150 trajectory-corrector data pairs were collected. The orbit-change data and corrector-magnet data were recorded separately, and the corresponding absolute orbit values were retrieved from the SXFEL historical database.

Figures 14 and 15 summarize the corresponding experimental observations. Figure 14 shows that the mean beam orbit remains stable during the online update, while Fig. 15 shows a reduction in the rms orbit jitter, demonstrating improved orbit stabilization performance. Specifically, Fig. 14 presents the evolution of the mean beam orbit during the online experiment after the intentional lattice change. The mean orbit remained nearly constant in both transverse planes throughout the test. This result indicates that the basic closed-loop orbit-feedback function remained intact during the DL-enabled online response matrix update

and that the adaptive update did not compromise routine beam stabilization.

The improvement is summarized in Fig. 15, which presents the standard deviation of the BPM readings over successive iterations. A clear decrease in the average BPM fluctuation is observed during the early iterations, followed by saturation at a lower level. This behavior demonstrates that the dictionary-learning-based update can recover the degraded correction performance caused by the intentionally introduced lattice change and maintain stable beam transport thereafter. Therefore, the proposed method recovered the orbit correction performance degraded by response-matrix mismatch in FEL operation.




FIG. 15. Evolution of the BPM fluctuation level over successive iterations during the online experiment. The solid line shows the average rms beam position over all BPMs, while the shaded region indicates the spread among individual BPMs. The gradual reduction and subsequent saturation of the fluctuation level indicate that the DL-enabled response matrix update progressively compensates for the mismatch introduced by the lattice change.

V. CONCLUSIONS

In this work, we introduced a dictionary learning approach to adaptively reconstruct the response matrix for XFEL beam trajectory correction. Unlike traditional methods that rely on a static, premeasured response matrix, our algorithm updates the matrix online without disrupting normal facility operation, thereby reducing both data requirements and correction latency. Through simulations under static, slowly drifting, and abrupt fault conditions, we demonstrated that the dictionary learning technique not only achieves high precision comparable to physics-based models but also exhibits high data efficiency. Furthermore, experiments at SXFEL demonstrated that, when the lattice was changed, the proposed DL-based update successfully compensated for the resulting model mismatch and restored stable orbit correction. Future research directions include exploring the influence of noise level and excitation perturbations on convergence speed, as well as expanding the application to additional sections of the facility beyond the drift region.

ACKNOWLEDGMENTS

The authors thank Hang Luo and Kaiqing Zhang for helpful discussions. This work was supported by the National Natural Science Foundation of China (12435011, 12125508, 12541501), the CAS Project for Young Scientists in Basic Research (YSBR-115), the Strategic Priority Program of the CAS (XDB0530300), and the Innovation Program of Shanghai Advanced Research Institute, CAS (2025CP006).

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

-
- [1] X. Huang, *Beam-Based Correction and Optimization for Accelerators*, 1st ed. (CRC Press, Milton, 2020).
 - [2] D. Dinev, Closed-orbit correction in synchrotrons, *Phys. Part. Nucl.* **28**, 398 (1997).
 - [3] B. Autin and Y. Marti, Closed orbit correction of A.G. machines using a small number of magnets, Report No. CERN-ISR-MA-73-17, 1973.
 - [4] Y. Chung, G. Decker, and K. Evans, Closed orbit correction using singular value decomposition of the response matrix, in *Proceedings of the International Conference on Particle Accelerators* (IEEE, New York, 1993), p. 2263.
 - [5] V. Ziemann, Corrector ironing in the SLC final focus, in *Proceedings of the International Conference on Particle Accelerators* (IEEE, New York, 1993), p. 498.
 - [6] G. Golub, Numerical methods for solving linear least squares problems, *Numer. Math.* **7**, 206 (1965).
 - [7] G. H. Golub, C. Reinsch, F. L. Bauer, H. Rutishauser, K. Samelson, F. W. J. Olver, E. Stiefel, and A. S. Householder, *Singular Value Decomposition and Least Squares Solutions* (Springer Berlin Heidelberg, Berlin, Heidelberg, 1971), p. 134.
 - [8] S. H. Mirza, R. Singh, P. Forck, and B. Lorentz, Performance of the closed orbit feedback systems with spatial model mismatch, *Phys. Rev. Accel. Beams* **23**, 072801 (2020).
 - [9] I. P. S. Martin, M. G. Abbott, R. Bartolini, M. J. Furseman, and G. Rehm, A fast optics correction for the diamond storage ring, in *Proceedings of the 5th International Particle Accelerator Conference (IPAC'14)*, Dresden, Germany (JACoW, Geneva, Switzerland, 2014), pp. 1763–1765, [10.18429/JACoW-IPAC2014-TUPRI083](https://doi.org/10.18429/JACoW-IPAC2014-TUPRI083).
 - [10] X. Yang, V. Smaluk, L. H. Yu, Y. Tian, and K. Ha, Fast and precise technique for magnet lattice correction via sine-wave excitation of fast correctors, *Phys. Rev. Accel. Beams* **20**, 054001 (2017).
 - [11] I. Ziemann and V. Ziemann, Noninvasively improving the orbit-response matrix while continuously correcting the orbit, *Phys. Rev. Accel. Beams* **24**, 072804 (2021).
 - [12] V. Ziemann, Operational improvements for an algorithm to noninvasively measure the orbit response matrix in storage rings, [arXiv:2303.11216](https://arxiv.org/abs/2303.11216).
 - [13] D. Allen, I. Kante, and D. Bohler, Beamline steering using deep learning models, [arXiv:2408.13657](https://arxiv.org/abs/2408.13657).
 - [14] I. Kante, Machine learning for beamline steering, [arXiv:2311.07519](https://arxiv.org/abs/2311.07519).
 - [15] S. H. Mirza, R. Singh, P. Forck, and H. Klingbeil, Closed orbit correction at synchrotrons for symmetric and near-symmetric lattices, *Phys. Rev. Accel. Beams* **22**, 072804 (2019).
 - [16] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb, Solving inverse problems using data-driven models, *Acta Numer.* **28**, 1 (2019).
 - [17] S. Ravishankar and Y. Bresler, MR image reconstruction from highly undersampled k-space data by dictionary learning, *IEEE Trans. Med. Imaging* **30**, 1028 (2010).
 - [18] M. Aharon, M. Elad, and A. Bruckstein, K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation, *IEEE Trans. Signal Process.* **54**, 4311 (2006).
 - [19] T. T. Cai and W. Lie, Orthogonal matching pursuit for sparse signal recovery with noise, *IEEE Trans. Inf. Theory* **57**, 4680 (2011).
 - [20] R. Rubinstein, M. Zibulevsky, and M. Elad, Efficient implementation of the K-SVD algorithm using batch orthogonal matching pursuit, *CS Technion* **40** (2008).
 - [21] S. Ghadimi and G. Lan, Accelerated gradient methods for nonconvex nonlinear and stochastic programming, *Math. Program.* **156**, 59 (2016).
 - [22] DICTLEARN, <https://github.com/permlf/dictlearn>.
 - [23] I. Agapov, G. Geloni, S. Tomin, and I. Zagorodnov, OCELOT: A software framework for synchrotron light source and FEL studies, *Nucl. Instrum. Methods Phys. Res., Sect. A* **768**, 151 (2014).