







femto-PIXAR: a self-supervised neural network method for reconstructing femtosecond X-ray free electron laser pulses

GESA GOETZKE,^{1,†} RAJAN PLUMLEY,^{2,3,4,†}  GREGOR HARTMANN,⁵ TIM MAXWELL,⁴ FRANZ-JOSEF DECKER,⁴ ALBERTO LUTMAN,⁴  MIKE DUNNE,⁴ DANIEL RATNER,^{4,6}  AND JOSHUA J. TURNER^{3,4,7} 

¹*Deutsches Elektronen-Synchrotron (DESY), Notkestrasse 85, 22607 Hamburg, Germany*

²*Department of Physics, Carnegie Mellon University, Pittsburgh, PA, USA*

³*Stanford Institute for Materials and Energy Sciences, Stanford University, Menlo Park, CA, USA*

⁴*SLAC National Accelerator Laboratory, Menlo Park, CA, USA*

⁵*Helmholtz-Zentrum Berlin für Materialien und Energie GmbH, Albert-Einstein-Strasse 15, 12489 Berlin, Germany*

⁶*dratner@slac.stanford.edu*

⁷*joshuat@slac.stanford.edu*

[†]The authors contributed equally to this work.

Abstract: X-ray free electron lasers (X-FELs) produce ultrafast pulses in a wide range of lasing configurations, supporting a wide variety of scientific applications, including structural biology, materials science, and atomic and molecular physics. Shot-by-shot characterization of the X-FEL pulses is crucial for the analysis of experiments as well as for tuning the X-FEL performance. However, for the weak pulses found in advanced configurations, e.g., those needed for monochromatic, two-pulse studies of quantum materials, there is no current method for reliably resolving pulse profiles. Here, we show that an interpretable neural network (NN) model can reconstruct the individual pulse power profiles for sub-picosecond pulse separation without the need for simulations. Using experimental data from low-signal X-FEL pulse pairs, we demonstrate a NN can learn the pulse characteristics on a shot-by-shot basis when conventional methods fail. This new method enables the characterization of weak pulses—a condition expected to dominate future experimental configurations such as at the Linac Coherent Light Source-II—and opens the door to a wide range of new experiments.

© 2025 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

Fluctuations in microscopic systems are directly related to their fundamental excitations – a concept dating back to Einstein’s description of the Brownian motion of pollen particles. This relationship has motivated the need for measuring fluctuations within microscopic systems to directly compare with first-principles theoretical models. However, the spatial and temporal scales of quantum systems far exceed those accessible by classical microscopy, and despite the advances in modern light sources, the energy scales pertinent to these fluctuations remain challenging to probe. One notable exception to this has been in inelastic neutron scattering, where small energy changes of scattered neutrons can be observed, and has led to understanding of dispersion in a myriad of areas, from frustrated magnets and high temperature superconductors, to biological tissue and exotic topological structures, such as skyrmions. However, the reliance of inelastic neutron scattering on large crystals and its limitations with strongly absorbing elements restrict its applicability. The advent of the X-ray free electron laser has changed all of this. Its capability to generate femtosecond (fs) laser pulses with extremely high pulse energies at X-ray

wavelengths has ushered in a new era of science, with new developments rapidly evolving in many fields [1,2] including structural biology [3,4], enzyme catalysis [5], and astrophysics [6,7]. Furthermore, the development of monochromatic, two-pulse configurations brings new potential for fluctuation studies within quantum materials at their relevant timescales, i.e. at the meV energy scale and below [8–13].

Importantly, with X-FELs it is now possible to deliver X-ray pulses at *sub-picosecond* spacing for experiments, utilizing the transverse coherence of each X-FEL pulse to resolve instantaneous snapshots in structures with sub-angstrom resolution. These stroboscopic scattering events are repeated and compared to measure fluctuations within the sample, offering unprecedented access to dynamic processes at ultrafast timescales. This highlights the urgent need for reliable characterization of X-FEL laser pulse intensities with femtosecond resolution. Notably, X-FELs are highly sensitive to operational parameters which govern the durations, energies, and sequences of X-ray pulses [14,15]. Compounding this complexity, X-FELs exhibit intrinsic intensity fluctuations, causing stochastic variations in the initial pulse energy spectrum and photon density from shot to shot [16,17]. These beam intensity fluctuations from the X-FEL present significant challenges for experiments where accurate characterization of pulse properties, including pulse duration and energy density, are needed for shot-to-shot corrections in order to extract meaningful results from the measurement series. Existing non-invasive diagnostics for absolute pulse energies like the gas energy monitor [18] or the X-ray gas monitor [19] cannot resolve the individual pulses at this few fs timescale. Other diagnostics, such as angular streaking [20,21] can predict X-ray pulse profiles in the attosecond to few fs regime while THz streaking [22] has shown similar capabilities in the fs regime. However, these methods require sophisticated setups and can be considered complex experiments themselves rather than a standard diagnostic at this time. Sanchez-Gonzalez et al. [23] utilized machine learning to combine information from different detectors in order to improve pulse diagnostics, including the temporal distance of two electron bunches, but did not provide pulse intensity information on these timescales. A method that has been successfully used to analyze pulses in the few fs timescale is the XTCAV method [24,25,30,31], described in more detail in both the setup, and the data and methodology sections below. The XTCAV pulse monitoring system works reliably for typical self-amplified spontaneous emission (SASE) experiments and has become a standard operational pulse diagnostic for X-FELs such as the Linac Coherent Light Source (LCLS). However, advanced X-FEL modes, e.g. self-seeding to provide a monochromatic beam, results in weak pulses that the conventional XTCAV analysis cannot resolve. Moreover, these modes are expected to become the standard for next generation high-repetition-rate XFELs, including LCLS-II. For these modes, the existing XTCAV analysis fails to resolve the signal over the background. These challenges have so far hindered use of this new capability to deliver new science in the observation of fluctuations for understanding properties in materials.

We address this need by developing a novel self-supervised machine learning (ML) framework called femto-PIXAR: *femtosecond Power Inference of X-ray pulses using AI References*, which provides sufficient sensitivity to resolve microjoule-scale pulses with femtosecond pulse separation. We show our approach accurately predicts pulse energies for an experimental dataset using a self-seeded beam, verified against total pulse energy monitors in real X-FEL data, and confirms that the predictions match the expected pulse delays. This demonstration will make ultrafast fluctuation studies possible in the sub-picosecond regime, and will also provide a tool for a wide array of X-FEL experiments in many other fields as well.

2. Setup

For this experiment, the X-FEL was set up to produce two monochromatic X-ray pulses starting from a single, long electron bunch, itself generated from two overlapping injector lasers. The two X-ray pulses were created using a double-slotted V-shaped foil (for time domain control)

[24,26], and hard X-ray self-seeding [27–29] (for longitudinal coherence of each individual pulse and energy selection). To our knowledge this configuration, depicted in Fig. 1(a), has never been used before. The experiment was conducted at a photon energy of 8.35 keV (Ni *K*-edge), with slotted-foil pulse separations of 0 fs (single pulse), ≈ 20 fs, and ≈ 30 fs. Different stages of the electron bunch shaping sequence are labeled e_{I}^- , e_{II}^- , e_{III}^- , e_{IV}^- , representing the same electron distribution at different instances as it proceeds in time. A 3 μm thick aluminum foil with two slots of variable separation from 1 mm to 5 mm was placed in the middle of the dispersive section of the bunch compressor BC2 [26]. This is because a bunch compressor in an X-FEL works similarly to a chirped pulse amplification laser system: an initially chirped pulse is compressed due to different path lengths for different energies, producing a spatial expansion of the electron beam at the position of the slotted foil. At the locations where the electron bunch e_{I}^- passes through the foil the electron emittance was ‘spoiled’ (i.e. increased) due to coulomb scattering. The emittance growth suppresses X-FEL amplification in the undulator, so that only the electrons e_{II}^- that pass through the slits of the V-shaped foil will later produce X-ray pulses. Downstream of the bunch compressor the relative distance between the unspoiled electron distributions e_{III}^- is in the direction of travel, resulting in an arrival time difference Δt , that can be changed by shifting the position of the slotted foil with respect to the beam center line. e_{IV}^- passes through the undulator magnets resulting in X-ray lasing. Between the first and second undulator segment a hard X-ray self-seeding scheme is applied in order to obtain monochromatic photon pulses (not shown here). Downstream of the undulators, in the X-band Transverse-deflecting-mode CAVity (XTCAV) [30,31], a radio-frequency electromagnetic field streaks the electron beam. The XTCAV is combined with an energy spectrometer that provides energy dispersion perpendicular to the streaking direction. The electron bunch is deflected onto a fluorescence target screen, dispersing the electron energy across the vertical axis while the arrival time is dispersed along the horizontal axis. The fluorescence is recorded with a CCD camera, allowing measurement of the electron bunch’s longitudinal energy profile with femtosecond resolution. Examples of XTCAV images are displayed in Fig. 1. This transient electron energy profile can be analyzed with or

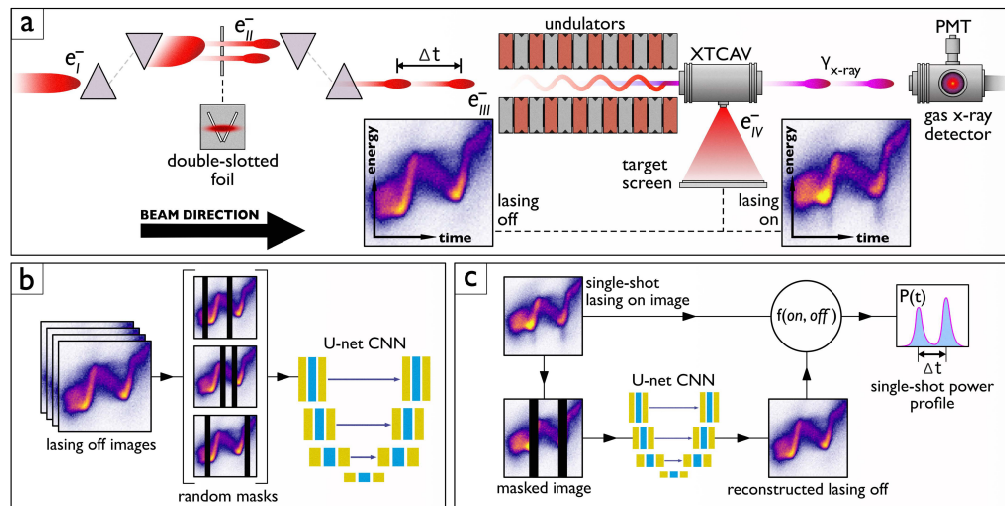


Fig. 1. (a) Diagram of the generation and shaping of a single-shot X-ray pulse pair at the LCLS, beginning with the transmission of a single electron bunch through the double-slotted foil. Both, the slotted foil method for pulse separation and the XTCAV diagnostic are described in the text. (b) Diagram of the ML training process. (c) Diagram of the ML-based reconstruction of the transient X-ray power profile $P(t)$ from a single-shot lasing-on image.

without X-ray lasing suppressed in order to determine the power profile of the outgoing X-ray laser as a function of time. This highly sensitive comparison of lasing-on and lasing-off beams is at the center of the analysis presented in this article. The X-ray laser pulses emitted by the electron bunch are unaffected by the XTCAV deflector magnet and continue downstream to the X-ray gas energy monitor (GEM), where the total X-ray pulse energy is transmissively monitored on a shot-to-shot basis using a photomultiplier tube (PMT) to measure the fluorescence induced in a small volume of N₂ gas [18]. After the GEM, the X-rays continue down the beamline to the experimental endstations. This entire process repeats at a 120 Hz repetition rate.

3. Data and methodology

As described in the previous section, the XTCAV system images the longitudinal (time-energy) distribution of the electron bunch with or without lasing suppressed. By comparison with a lasing-off reference, it is possible to infer the temporal profile of the X-ray pulse. However, because each shot is unique and the lasing-on and lasing-off references cannot be measured simultaneously, this poses a fundamental challenge in determining the change in the electron phase space due to the lasing process. The “classical” approach has been to record a set of images while suppressing lasing, cluster these images, and then select the most similar cluster to reference against each lasing-on image [30–32]. While the classical approach has been demonstrated for various two-pulse modes [14,24,25], it struggles for small pulse energies, when shot-to-shot differences between the true lasing-off electron bunch and the selected reference cluster exceed the lasing-induced changes. As a result, when combining self-seeding with two-pulse modes, the classical method breaks down.

A deep learning method proposed by Ren et al. [33] sidesteps the need for matched lasing-off references by training a neural network (NN) directly to predict power from a lasing-on image; however, the NN in [33] relies on simulated labels for training, and it is particularly challenging to simulate the complex electron beam distribution produced by the slotted foil. For the slotted foil setup used in this work, only a narrow segment of the beam lases. We can exploit this feature to approximate a reference in the short lasing region by using information in the nearby non-lasing regions. For example, Zeng et al. [34] suggested to use a polynomial regression across the lasing region as a reference. However, the polynomials cannot model the complex phase space of this setup and fail to reproduce the weak lasing pulses.

Here, we instead use a NN to generate the full corresponding lasing-off reference image. We employ a U-net architecture [35], originally developed for biomedical image segmentation, and now broadly used in different fields [36–39]. A U-net uses an encoder–decoder architecture with symmetric skip connections, letting it capture the global structure within the data while preserving fine details, making it ideal for resolving low signal-to-background features. We trained the NN using self-supervision, eliminating the need for simulations. Specifically, we masked regions of lasing-off pulses, and trained the NN to regenerate those regions. Once trained, the network receives a lasing-on image with masked lasing regions and reconstructs the best lasing-off reference. We then calculated the center of mass (COM) to retrieve the X-ray pulse distribution [31]. Our method provides high-fidelity lasing-off references without relying on the accuracy of computationally-expensive simulations while retaining interpretability through direct observation of the generated reference.

Figure 1(b and c) shows the core steps for the network training and inference process. A technical description of the steps is included in the supplemental material. The network was trained on about 30k measured lasing-off samples with an additional 1.5k frames being used as validation to check for overfitting. Additionally, we have reserved 460 lasing off samples as a test dataset, that were never seen by the network. The inputs to the network were cropped, background subtracted and normalized lasing off images, where all values within the two random mask-regions were set to zero (black bars in Fig. 1(b) and (c)). The width of an individual mask

was chosen to be 16 pixels and the distance range for the two masks was set from 0 to 40 pixels, corresponding to approximately 0 to 38 fs. We then trained a U-net to reconstruct the electron phase space of the lasing-off data (Fig. 1(b)). As the COM is the critical value for evaluation, we used a weighted combination of the mean squared error (MSE) of pixel values and the absolute error of the COM. The COM has a relative low weight of $1e-4$ as we found that higher weights significantly worsen the reconstruction. After training, we masked the lasing regions in a lasing-on data set. As the network is trained only on the lasing-off data, it reconstructs the best matching lasing-off electron phase space exactly for a given image (Fig. 1(c)). This lasing-off image will also exactly match the original x and y position, so no shifting in the time axis or normalization with the GEM (as in the classical approach) is necessary.

A crucial step of this process was to determine the correct mask positions in the lasing data. The separation of the two pulses was approximately fixed by the slit separation, but due to shot-to-shot differences in electron bunch energy and compression of the electron bunch, the position of the features relative to the COM of the electron bunch was not constant. Our method scans the image to find two locations where input and reconstruction differ most, centering the masks on the maximum signal. This procedure is described in more detail in the supplementary material. We then use the methods from Behrens et. al [31], marked as $f(\text{on, off})$ in Fig. 1 which calculates the transient X-ray power profile by using the COM from the lasing-on image “on” and subtracting the COM of the network reconstructed lasing-off image “off”.

4. Results

Figure 2 shows example reconstructions for the three different configurations together with one lasing-off example. All images were randomly drawn from the top 1% of shots with the largest GEM values, ensuring a clearly visible signal; none were used for network training. The reconstructions reliably reproduce the two-dimensional structure, and the main difference between lasing on and lasing off is within the masked regions. We note that the U-Net effectively de-noises the reconstructed masked region, as the random noise signals cannot be inferred from the neighboring regions. However, this has no impact on the pulse reconstruction. While most of the difference between lasing on and lasing off is within the masked region, there is some residual signal in the regions outside the masks, where the network has full information and should be able to reconstruct the beam perfectly. We find this residual increases when the center-of-mass (COM) loss is introduced or assigned greater weight. In principle the effect could be removed through additional hyper-parameter tuning; however, because these minor artifacts do not affect the extracted power profiles, no further optimizations were necessary. To demonstrate the efficacy of the proposed method, we benchmark the predicted power and delay against existing diagnostics. First, we compare the integral of the predicted profile with the total energy of both pulses measured by the GEM, which cannot resolve individual pulses on the fs scale but measures the total energy for each pair. In Fig. 3 one can see the correlation of the integral of the reconstructed power profile with the GEM for the three different slotted foil configurations. We note that while there is a strong linear correlation between the predicted power and GEM monitor, there is a disagreement in proportionality. Further experiments are necessary to determine if the disagreement is due to either the GEM or XTCAV calibration, or bias from the method. However, if the absolute pulse energy is a measurement of interest, it would be possible to correct for any proportional bias from the XTCAV diagnostic by using the GEM, though this change would not affect the experimental analysis outlined here.

Next, we compare our predictions to two existing state-of-the-art XTCAV analysis methods. In order to evaluate the correlation with the GEM, in the classical approach (see the previous section), the energy jitter (shift in the y -axis) cannot be compensated for by normalization with the GEM but must instead be corrected using a scheme that matches the head and tail of the COM profiles for the lasing-on and lasing-off images. We also compare our results with the polynomial

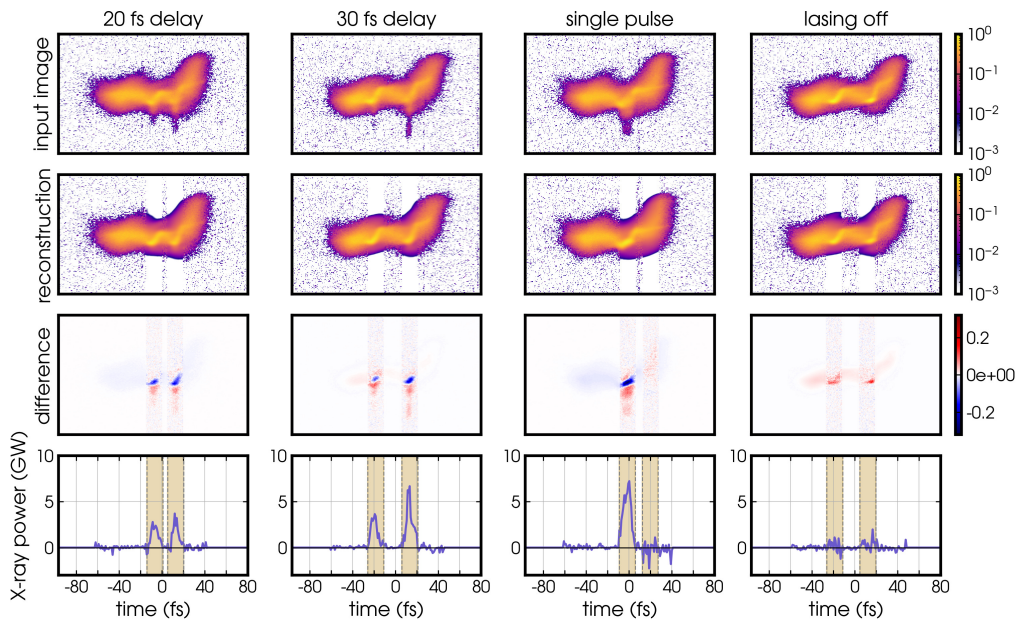


Fig. 2. Demonstration of femto-PIXAR reference creation: The XTCAV images show the respective longitudinal electron densities of the FEL bunch, with the x-axis corresponding to time and the y-axis corresponding to the electrons' kinetic energy. The color density in each panel represents the amount of electron charge per pixel. The top row displays background-subtracted XTCAV images for three different slotted-foil setups (lasing-on) plus a lasing-off sample at a 30 fs delay. The second row shows the predicted lasing-off reconstruction; the third row the difference between the measured XTCAV data and that reconstruction; and the bottom row the resulting reconstructed power profiles.




Fig. 3. Correlations with the GEM for the three different delays using the U-net approach. Both the classical method (clas.) [30] and polynomial regression (poly.) [34] reconstruction methods displayed in the insets show poor correlation between the predicted power-profile integrals and the total GEM pulse energy especially on the weaker double pulses (≈ 20 fs and ≈ 30 fs delay).

regression method of Zeng et al. [34]. Some example predictions of this method can be found in the supplemental material. Fig. 3 compares the GEM correlation for all three methods across all three delays, and neither of the existing methods results in the expected correlation.

Figure 4 shows the temporal and energy characteristics of the X-ray laser pulses as retrieved by the U-net reconstruction scheme. The time assignments of each shot were obtained by fitting the

X-ray power trace with two Gaussian functions. For these plots we took only the data where both Gaussian curve fits indicated a signal larger than $15 \mu\text{J}$. We find that the predicted time difference between both pulses closely matches the delay given by the slotted foil condition. The temporal separation jitter aligns well with the 9% jitter observed by Ding et al. [24]. As an additional benchmark we checked the predicted pulse energies for lasing-off data with randomly placed masks and find that 80% of the data has a deviation from zero of less than $15 \mu\text{J}$, and 99.8% less than $25 \mu\text{J}$. We find that the pulse energy difference between the two pulses is centered well around zero for lasing-off-data, indicating that our approach does not introduce a bias in the ratio between the first (p1) and the second (p2) pulse. We also compute the relative squared error between the input frames and the lasing off frames for this data where we have the ground truth (true signal). More information on these additional benchmarks showcasing the validity of the method can be found in the supplemental material.

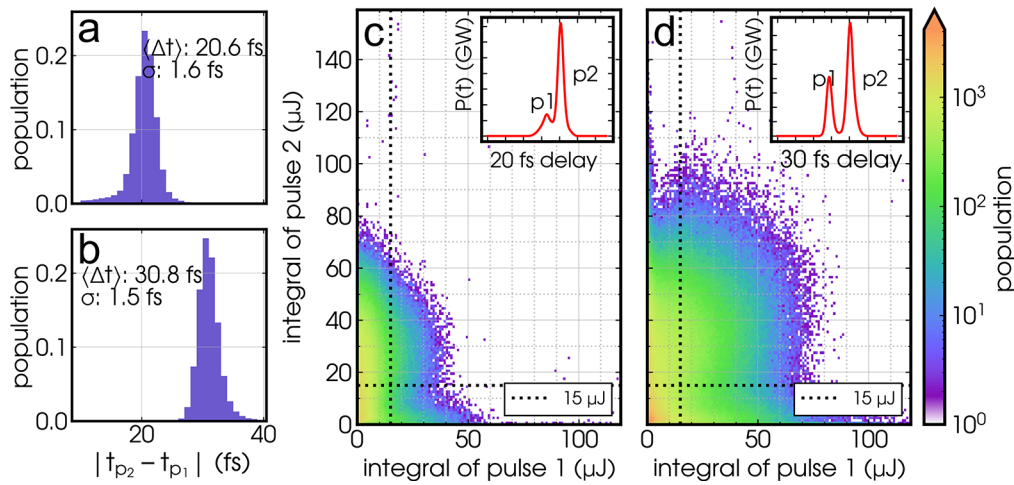


Fig. 4. (a) & (b): Histograms of temporal separation $\Delta t = |t_{p2} - t_{p1}|$ between each pulse for the 20 fs and 30 fs delay setups, respectively. (c) & (d): Histograms of the integrated energy in each pulse for both setups. The dashed line indicates the $15 \mu\text{J}$ minimum energy threshold used to select shots for the time-separation analysis in (a) & (b). The insets show the average of the transient X-ray power profiles as a function of time t (fs), where the relative average amplitudes of the pulse 1 (p1) and pulse 2 (p2) are indicated.

5. Discussions and summary

In conclusion, femto-PIXAR is a novel X-FEL diagnostic method which we demonstrate by reconstructing X-ray power profiles from real X-FEL data. For the configuration used to deliver monochromatic pulse pairs at few fs distance, the classical method of XTCAV analysis, which uses a matching algorithm to subtract similar lasing-off shots for the energy loss calculation, fails to reproduce meaningful correlations with the X-ray gas energy monitor. By contrast, our method is able to distinguish the relative amplitude of each X-ray pulse on a shot-to-shot basis which is necessary for X-ray photon fluctuation spectroscopy experiments [40] and other future experiments at X-FEL facilities that utilize non-standard beam setups. Our approach has various benefits, avoiding the need for energy calibration with an external reference. A neural network based scheme also scales well for the large datasets expected from high-throughput experiments enabled by next generation X-FEL facilities such as the LCLS-II [41]. As opposed to previous deep-learning work, our approach uses self-supervision, avoiding the need for high-fidelity simulations. Our method also reconstructs the full 2D phase space, providing

interpretability of the network and giving crucial information about when the network might not be applicable for a specific dataset. Future work will consider extensions to more general X-FEL configurations beyond the slotted foil. For example, using other diagnostics in addition to the XTCAV could enable analysis of configurations with lasing across the full electron bunch [23], e.g. to provide the femtosecond resolution needed for temporal ghost imaging [42] or help overcome hardware limitations that restrict temporal resolution to a few femtoseconds, enabling access to the attosecond regime. The analysis chain after reconstructing the lasing-off reference can also make use of machine learning, e.g. incorporating spectral measurements to refine the X-ray pulse profile [43].

Funding. U.S. Department of Energy (DE-AC02-76SF00515., DE-SC0022216); IVF project InternLabs (HIR3X). (0011193).

Acknowledgment. Use of the Linac Coherent Light Source (LCLS), SLAC National Accelerator Laboratory, is supported by the US Department of Energy, Office of Science, Office of Basic Energy Sciences under Contract No. DEAC02-76SF00515. This work has in part been funded by the IVF project InternLabs-0011 (HIR3X). R. Plumley acknowledges support from the US Department of Energy, Office of Science, Basic Energy Sciences, for the Materials Sciences and Engineering Division under contract DE-AC02-76SF00515. J. J. Turner acknowledges support from the U.S. Department of Energy (DOE), Office of Science, Basic Energy Sciences under Award No. DE-SC0022216. D. Ratner acknowledges support from the U.S. Department of Energy (DOE), Office of Science, Basic Energy Sciences under Award No. DE-AC02-76SF00515.

Disclosures. The authors declare no conflicts of interest.

Data availability. The data presented in this study are available on request from the corresponding authors.

Supplemental document. See [Supplement 1](#) for supporting content.

References

1. C. Bostedt, S. Boutet, D. M. Fritz, *et al.*, “Linac coherent light source: The first five years,” *Rev. Mod. Phys.* **88**(1), 015007 (2016).
2. J. Rossbach, J. R. Schneider, and W. Wurth, “10 years of pioneering X-ray science at the free-electron laser flash at desy,” *Phys. Rep.* **808**, 1–74 (2019).
3. S. Boutet, L. Lomb, G. J. Williams, *et al.*, “High-resolution protein structure determination by serial femtosecond crystallography,” *Science* **337**(6092), 362–364 (2012).
4. H. N. Chapman, P. Fromme, A. Barty, *et al.*, “Femtosecond X-ray protein nanocrystallography,” *Nature* **470**(7332), 73–77 (2011).
5. S. L. Rose, S. V. Antonyuk, D. Sasaki, *et al.*, “An unprecedented insight into the catalytic mechanism of copper nitrite reductase from atomic-resolution and damage-free structures,” *Sci. Adv.* **7**(1), eabd8523 (2021).
6. S. Bernitt, G. V. Brown, J. K. Rudolph, *et al.*, “An unexpectedly low oscillator strength as the origin of the fe xvii emission problem,” *Nature* **492**(7428), 225–228 (2012).
7. S. M. Vinko, O. Ciricosta, B. I. Cho, *et al.*, “Creation and diagnosis of a solid-density plasma with an X-ray free-electron laser,” *Nature* **482**(7383), 59–62 (2012).
8. Y. Sun, M. Dunne, P. Fuoss, *et al.*, “Realizing split-pulse X-ray photon correlation spectroscopy to measure ultrafast dynamics in complex matter,” *Phys. Rev. Res.* **2**(2), 023099 (2020).
9. C. Gutt, L.-M. Stadler, A. Duri, *et al.*, “Measuring temporal speckle correlations at ultrafast X-ray sources,” *Opt. Express* **17**(1), 55–61 (2009).
10. R. Plumley, S. Chitturi, C. Peng, *et al.*, “On ultrafast X-ray scattering methods for magnetism,” *Adv. Phys.: X* **9**(1), 2423935 (2024).
11. M. Seaberg, B. Holladay, J. Lee, *et al.*, “Nanosecond X-ray photon correlation spectroscopy on magnetic skyrmions,” *Phys. Rev. Lett.* **119**(6), 067403 (2017).
12. M. Seaberg, B. Holladay, S. Montoya, *et al.*, “Spontaneous fluctuations in a magnetic fe/gd skyrmion lattice,” *Phys. Rev. Res.* **3**(3), 033249 (2021).
13. W. Roseker, S. Hruszkewycz, F. Lehmkuhler, *et al.*, “Towards ultrafast dynamics with split-pulse X-ray photon correlation spectroscopy at free electron laser sources,” *Nat. Commun.* **9**(1), 1704 (2018).
14. F.-J. Decker, K. L. Bane, W. Colocho, *et al.*, “Tunable X-ray free electron laser multi-pulses with nanosecond separation,” *Sci. Rep.* **12**(1), 3253 (2022).
15. F.-J. Decker, W. Colocho, A. Halavanau, *et al.*, “Two and multiple bunches with the LCLS copper linac,” in *Proc. IPAC’22*, (2022), pp. 1089–1091.
16. R. Bonifacio, L. De Salvo, P. Pierini, *et al.*, “Spectrum, temporal structure, and fluctuations in a high-gain free-electron laser starting from noise,” *Phys. Rev. Lett.* **73**(1), 70–73 (1994).
17. Y. Sun, F.-J. Decker, J. Turner, *et al.*, “Pulse intensity characterization of the LCLS nanosecond double-bunch mode of operation,” *J. Synchrotron Radiat.* **25**(3), 642–649 (2018).

18. S. P. Hau-Riege, R. M. Bionta, D. D. Ryutov, *et al.*, “Near-ultraviolet luminescence of n 2 irradiated by short X-ray pulses,” *Phys. Rev. Lett.* **105**(4), 043003 (2010).
19. K. Tiedtke, J. Feldhaus, U. Hahn, *et al.*, “Gas detectors for X-ray lasers,” *J. Appl. Phys.* **103**(9), 094511 (2008).
20. N. Hartmann, G. Hartmann, R. Heider, *et al.*, “Attosecond time–energy structure of X-ray free-electron laser pulses,” *Nat. Photonics* **12**(4), 215–220 (2018).
21. S. Li, Z. Guo, R. N. Coffee, *et al.*, “Characterizing isolated attosecond pulses with angular streaking,” *Opt. Express* **26**(4), 4531 (2018).
22. M. C. Hoffmann, I. Grguraš, C. Behrens, *et al.*, “Femtosecond profiling of shaped X-ray pulses,” *New J. Phys.* **20**(3), 033008 (2018).
23. A. Sanchez-Gonzalez, P. Micaelli, C. Olivier, *et al.*, “Accurate prediction of X-ray pulse properties from a free-electron laser using machine learning,” *Nat. Commun.* **8**(1), 15461 (2017).
24. Y. Ding, C. Behrens, R. Coffee, *et al.*, “Generating femtosecond X-ray pulses using an emittance-spoiling foil in free-electron lasers,” *Appl. Phys. Lett.* **107**(19), 191104 (2015).
25. A. Marinelli, R. Coffee, F.-J. Decker, *et al.*, “Twin-bunch two-colour FEL at LCLS,” in *Proc. IPAC’16*, vol. IPAC2016 (2016), pp. 1032–1036.
26. P. Emma, M. Borland, and Z. Huang, “Attosecond X-ray pulses in the LCLS using the slotted foil method,” in *Proc. FEL’04*, (JACoW Publishing, pp. 333–338., 2004).
27. G. Geloni, V. Kocharyan, and E. Saldin, “A novel self-seeding scheme for hard X-ray FELs,” *J. Mod. Opt.* **58**(16), 1391–1403 (2011).
28. J. Amann, W. Berg, V. Blank, *et al.*, “Demonstration of self-seeding in a hard-X-ray free-electron laser,” *Nat. Photonics* **6**(10), 693–698 (2012).
29. A. Lutman, F.-J. Decker, J. Arthur, *et al.*, “Demonstration of single-crystal self-seeded two-color X-ray free-electron lasers,” *Phys. Rev. Lett.* **113**(25), 254801 (2014).
30. Y. Ding, C. Behrens, P. Emma, *et al.*, “Femtosecond X-ray pulse temporal characterization in free-electron lasers using a transverse deflector,” *Phys. Rev. ST Accel. Beams* **14**(12), 120701 (2011).
31. C. Behrens, F.-J. Decker, Y. Ding, *et al.*, “Few-femtosecond time-resolved measurements of X-ray free-electron lasers,” *Nat. Commun.* **5**(1), 3762 (2014).
32. T. J. Maxwell, C. Behrens, Y. Ding, *et al.*, “Femtosecond-scale X-ray FEL diagnostics with the LCLS x-band transverse deflector,” in *SPIE Proceedings*, (2014).
33. X. Ren, A. Edelen, A. Lutman, *et al.*, “Temporal power reconstruction for an X-ray free-electron laser using convolutional neural networks,” *Phys. Rev. Accel. Beams* **23**(4), 040701 (2020).
34. L. Zeng, C. Feng, D. Gu, *et al.*, “Online single-shot characterization of ultrafast pulses from high-gain free-electron lasers,” *Fundamental Res.* **2**(6), 929–936 (2022).
35. O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention*, (Springer, 2015), pp. 234–241.
36. Y. Cai and J. Yuan, “A review of U-Net network medical image segmentation applications,” in *Proceedings of the 2022 5th International Conference on Artificial Intelligence and Pattern Recognition*, (Association for Computing Machinery, 2023), p. 457–461.
37. W.-J. Lee, W.-T. Hsieh, B.-H. Fang, *et al.*, “Device simulations with a U-Net model predicting physical quantities in two-dimensional landscapes,” *Sci. Rep.* **13**(1), 1 (2023).
38. V. Gorse, R. Mitteau, and J. Marot, “Using a physics constrained U-Net for real-time compatible extraction of physical features from west divertor hot-spots,” *J. Fusion Energ.* **43**(1), 13 (2024).
39. A. Kamali and K. Laksari, “Physics-informed UNets for discovering hidden elasticity in heterogeneous materials,” *J. Mech. Behav. Biomed. Mater.* **150**, 106228 (2024).
40. L. Shen, M. Seaberg, E. Blackburn, *et al.*, “A snapshot review-Fluctuations in quantum materials: from skyrmions to superconductivity,” *MRS Adv.* **6**(8), 221–233 (2021).
41. J. Galayda, “The LCLS-II: A high power upgrade to the LCLS,” in *Proc. IPAC’18*, (2018), pp. 18–23.
42. D. Ratner, J. Cryan, T. Lane, *et al.*, “Pump-probe ghost imaging with SASE FELs,” *Phys. Rev. X* **9**(1), 011045 (2019).
43. D. Ratner, F. Christie, J. Cryan, *et al.*, “Recovering the phase and amplitude of X-ray FEL pulses using neural networks and differentiable models,” *Opt. Express* **29**(13), 20336–20352 (2021).