Article

Deep learning with plasma plume image sequences for anomaly detection and prediction of growth kinetics during pulsed laser deposition

Check for updates

Sumner B. Harris ✉, Christopher M. Rouleau, Kai Xiao & Rama K. Vasudevan

Materials synthesis platforms that are designed for autonomous experimentation are capable of collecting multimodal diagnostic data that can be utilized for feedback to optimize material properties. Pulsed laser deposition (PLD) is emerging as a viable autonomous synthesis tool, and so the need arises to develop machine learning (ML) techniques that are capable of extracting information from in situ diagnostics. Here, we demonstrate that intensified-CCD image sequences of the plasma plume generated during PLD can be used for anomaly detection and the prediction of thin film growth kinetics. We develop multi-output (2 + 1)D convolutional neural network regression models that extract deep features from plume dynamics that not only correlate with the measured chamber pressure and incident laser energy, but more importantly, predict parameters of an auto-catalytic film growth model derived from in situ laser reflectivity experiments. Our results demonstrate how ML with in situ plume diagnostics data in PLD can be utilized to maintain deposition conditions in an optimal regime. Further, the predictive capabilities of plume dynamics on the kinetics of film growth or other film properties prior to deposition provides a means for rapid pre-screening of growth conditions for the non-expert, which promises to accelerate materials optimization with PLD.

The recent advent of autonomous synthesis platforms which include multiple in situ or automated diagnostics and characterization techniques[1–5] drives the need to develop machine learning (ML) models that take advantage of this multimodal synthesis data. In situ diagnostics and characterizations are employed in autonomous materials synthesis to derive optimization metrics or correlate growth kinetics with experimental controls[6–9]. In addition, they can be utilized to predict material properties before synthesis and detect anomalous (erroneous) conditions that may arise during long, unsupervised experiments. Pulsed laser deposition (PLD) is a promising physical vapor deposition technique for autonomous synthesis[10] of numerous materials systems due to its compatibility with various optical, electrical, and electron-based diagnostic measurements. Thus, the need to develop ML methods that leverage PLD diagnostic advantages is nascent.

Several types of in situ diagnostics are used during PLD synthesis to monitor either the growing film or the plasma plume: Reflection high-energy electron diffraction (RHEED)[11,12], reflectivity and ellipsometry[12–14], intensified-CCD (ICCD) imaging[15–17], and ion/Langmuir probes[18–20]. Although ML for synthesis has been available for a long time, its applications to in situ diagnostics data is limited. Seminal work was performed by the May[21] group wherein neural networks were used with RHEED patterns to predict film properties or forecast RHEED intensity during growth with molecular beam epitaxy. Vasudevan et al. explored the use of other ML algorithms to better understand RHEED image sequences[22], which was then furthered by the Comes[23] group. Haotong et al. also recently demonstrated machine learning with RHEED images to construct structural phase maps with respect to PLD synthesis conditions[24]. Despite the widespread adoption of deep learning across the physical sciences, the adoption in this domain remains highly limited, likely due to sparse or undisciplined datasets. Here, utilizing our autonomous PLD platform, we are able to generate comparatively large amounts of highly-disciplined PLD synthesis data that enables us to explore ML methods with in situ diagnostics for thin film growth.

Center for Nanophase Materials Sciences, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA. ✉e-mail: harrissb@ornl.gov

Plume imaging data from ICCD cameras (or even just visual inspection or common CMOS camera images) has promise to increase the reproducibility of PLD experiments or be used to predict characteristics of the film such as growth rate, domain size, crystallinity, or defect density which alter material properties. Key features of the PLD plume and its expansion dynamics such as brightness and color (optical emission)[25], angular distribution, maximum kinetic energy[16], the presence of "fast" or "slow" components[26,27] can be correlated to properties of the thin film or be indicators that the plume conditions are consistent among experiments. Moreover, if plume dynamics alone can act as a reasonable predictor of film properties after synthesis, quick ICCD imaging experiments can be used as a low-fidelity surrogate[28,29] to more rapidly explore the synthesis parameter space and reduce time/resources wasted on full growth and characterization with conditions that would produce poor quality outcomes. However, analysis of plume image sequences can be a time consuming—i.e., an "offline" effort—and is likely a major contributing factor to the lack of widespread adoption as an on-the-fly, quantitative diagnostic for synthesis. Also, it is generally unknown how certain aspects of the plume are correlated to materials properties which is another major barrier for utilizing ICCD imaging in real-time. Thus, ML with plume images may provide a method of extracting key features of the plume and correlating them with quantities of interest.

Here, we seek to prove the viability of deep learning with ICCD images for monitoring plume conditions or predicting film properties in real-time during PLD synthesis. We use a $(2+1)$D convolutional neural network (CNN) to extract deep features from ICCD image sequences to correlate plume dynamics to PLD conditions and film growth kinetics. To detect anomalies in the anticipated conditions during unsupervised autonomous synthesis routines, we show that ICCD image sequences alone can predict the chamber pressure and laser energy—two typically measured qualities—but more importantly, predict and improve the accuracy of parameters of a thin film growth model. Our results demonstrate the potential of routine plume imaging for effective anomaly detection and prediction of growth kinetics or other film properties that can be used to improve and accelerate PLD synthesis experiments.

Methods
Data generation
The ICCD image sequences, growth parameters, and growth kinetics parameters were generated during an autonomous PLD synthesis campaign which is separate from this work. Complete details of the experimental results, PLD chamber design, and diagnostic methods can be found elsewhere[8,10]. In the experiment, WSe$_2$ was chosen as the material of choice to develop the autonomous synthesis methodology because the challenge of depositing a material composed of elements with disparate vapor pressures presented the opportunity to tune Se compensation by co-deposition and the Raman spectrum of WSe$_2$ is highly sensitive to defects and crystallinity which could act as an in situ metric for optimization. However, the optical emission from the PLD plumes as recorded by the ICCD camera were not used for predictions or to aid optimization at the time.

Briefly, 127 WSe$_2$ films were grown to approximately 1 monolayer thickness on 90 nm SiO$_2$/Si substrates using PLD by co-ablating two targets (WSe$_2$ and Se). Each sample had a combination of background pressure P, substrate temperature T, laser energy on the WSe$_2$ target E_1, and laser energy on the Se target E_2. During growth, a sequence of 50 ICCD images was taken with delay times (relative to the excimer laser pulse) from 2–150 μs and, simultaneously, laser reflectivity was used to measure the reflected contrast change vs. time of the sample to detect sub-monolayer nucleation and growth of the WSe$_2$ film. Therefore, each individual deposition has a corresponding sequence of 50 ICCD images, a set of (P, T, E_1, E_2), and a laser reflectivity curve. Here, we use the laser energy E_1, E_2 because it is the quantity that is directly reported in the dataset but it could be converted to fluence (J cm^{-2}) if desired because the laser spot sizes are known. Since laser energy and fluence are directly proportional, the choice will not greatly affect the outcome of the trained models. For reference, the spot size of the excimer

laser was fixed throughout the experiments to be 0.0256 cm^2 and 0.0410 cm^2 on the WSe$_2$ and Se targets, respectively.

Each ICCD image is 1 channel (16-bit) with 1024×1024 pixels. The images are Gaussian filtered to reduce noise, resized to 40×40 pixels, and the log of the intensity is taken to enhance weak characteristics in each image (because the luminous intensity of the plume spatiotemporally varies by orders of magnitude). The intensity for each sequence is normalized for machine learning which concurrently mitigates issues related to the attenuation of the absolute intensity of the plume as material coats on the ICCD viewport over time. This gives each image sequence a depth, height, and width $D \times H \times W = 50 \times 40 \times 40$. Lastly, the laser reflectivity curves are fit with an auto-catalytic growth model that can represent the nucleation and growth of the film in terms of fractional monolayer coverage[13]. The three growth kinetics parameters returned by this model are denoted s_0 and s_1 for the sticking coefficients of plume species to the substrate and film, respectively, and J which is the flux of arriving species. For the purposes of this work, we do not claim that this is the correct growth model or mechanism but rather use it as a fitting equation that effectively reproduces the observed laser reflectivity curves so that the growth kinetics curves can be predicted by deep learning.

Machine learning models and training
Figure 1 shows a schematic of the ML networks used for this this study. We used a $(2+1)$D CNN[30] to extract deep features from the ICCD image sequences for multi-ouput regression. A $(2+1)$D convolution splits a 3D convolution into successive 2D spatial and 1D temporal convolutions with the benefit of a reduced number of parameters and increased nonlinearities relative to a standard 3D convolution[30]. To make predictions from (P, T, E_1, E_2) alone for comparison, we use a simple multilayer perceptron (MLP). For a mixed input model which includes both ICCD images and (P, T, E_1, E_2), we combine the output from the ICCD and growth parameter MLP into a final, combined MLP.

The $(2+1)$D convolution layers consist of 2 sequential 3D convolutions with kernel sizes of $(1 \times 3 \times 3)$ and $(3 \times 1 \times 1)$ for the spatial and temporal components, respectively, followed by a 3D batch normalization layer and a leaky ReLU activation[31]. The ICCD features are extracted using 3 successive $(2+1)$D layers (with 64, 128, and 256 filters) with an average pooling layer after each with a $(2 \times 2 \times 2)$ kernel to downsample the data. The output of the last $(2+1)$D layer is flattened into a $256 \times 6 \times 5 \times 5$ tensor and passed through 2 final fully connected layers (Linear 1, Linear 2) with 64 and 32 nodes each. Features are extracted from the growth parameters by feeding the 4 inputs into a simple MLP, 2 fully connected layers with 48 and 32 nodes, respectively. For predictions of P, E_1, and E_2 using only ICCD images, the Linear 2 layer is reduced to 3 outputs for regression and the same is done for predictions of s_0, s_1, and J with the growth parameters MLP. Finally, the mixed input model for predictions of growth kinetics concatenates the outputs of each sub-network and feeds through a 3 layer MLP with 16, 24, and 32 nodes in each layer.

The 127 sample dataset for this study can be considered quite small in the context of typical deep-learning studies. In the context of materials synthesis with physical vapor deposition techniques, this is a typical size compared to other similar autonomous synthesis studies, typically numbering a few 10s to 100s of samples[32-34]. During training, the data is shuffled and split 70/30 into training and validation sets, 88 and 39 samples, respectively, and training is done with batch gradient decent. To increase the variety in this small dataset and improve model generalization, we augment both the image and growth parameter data. The image sequences are transformed with random rotations, translations, shear, and scale in each epoch such that the same transformation is applied only in the spatial dimensions for each temporal frame. Similarly, P, T, E_1, E_2 are augmented by adding Gaussian noise with an appropriate variance to simulate the measurement accuracy of P and T and natural shot-to-shot stability of the excimer laser. We found that training with data augmentation improved model predictions. We used the mean square error (MSE) loss function and the Adam[35] optimizer. Hyperparameter tuning was done using Ray Tune[36]
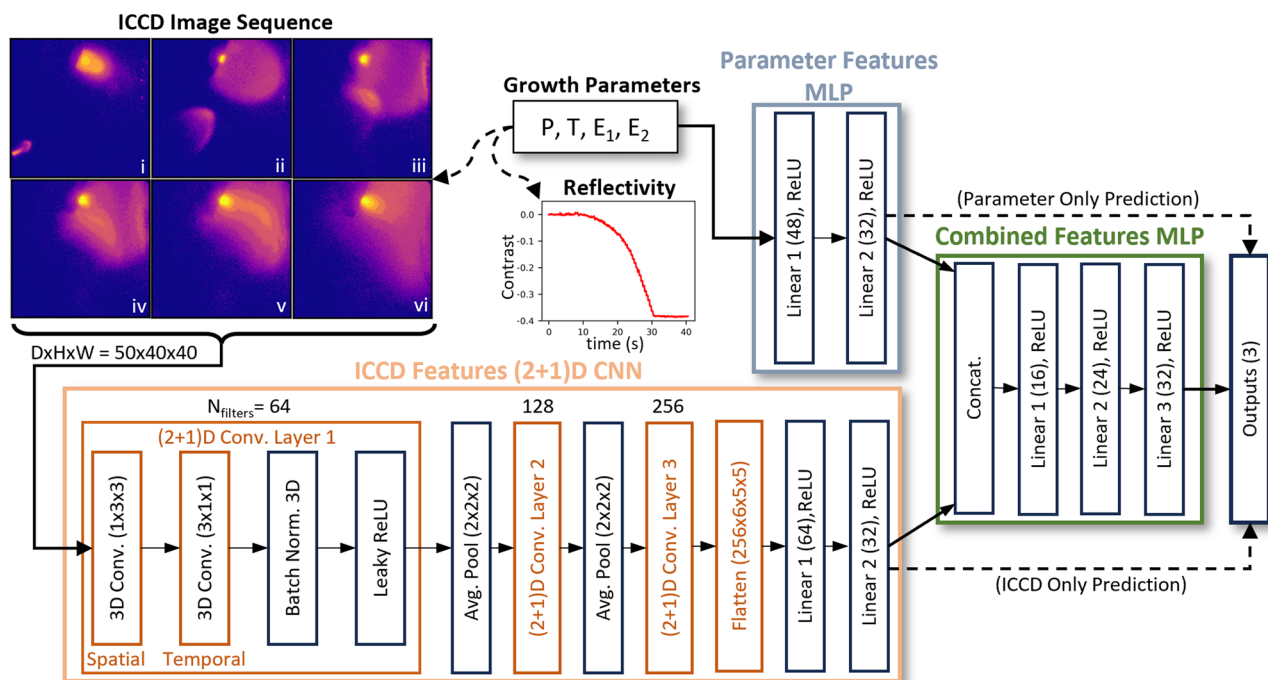
Fig. 1 | Schematic diagram of the neural networks used for multi-output regression with intensified-CCD (ICCD) image sequences and/or growth parameters as inputs. Pulsed laser deposition (PLD) growth for each sample is carried out by selecting a chamber pressure, substrate temperature, and laser energy on each PLD target (P, T, E_1, E_2). During the growth, an ICCD image sequence and in situ laser reflectivity curve are collected. The image sequence is fed into a $(2 + 1)$D convolutional neural network (CNN) to extract deep features from the plume dynamics. For experimental anomaly detection, P, E_1, E_2 are predicted from the ICCD features alone. The growth kinetics parameters (s_0, s_1, J) derived from reflectivity measurements are predicted by either a multilayer perceptron (MLP) using the growth parameters, the $(2 + 1)$D CNN, or a by combining the features from both in a final MLP.

with the Optuna alogrithm[37] and asynchronous successive halving (ASHA)[38] for early-stopping of poor performing runs. The hyperparameters that were optimized were the Adam learning rate, L2 regularization, and number of nodes in the linear layers. Hyperparameter tuning was done for each model discussed in this paper.

Results and discussion
Anomaly detection
In PLD, several usually unmonitored factors can affect the reproducibility between depositions. One well-known issue in PLD is that material is deposited on the chamber's laser window over time which attenuates the laser energy and changes the plume dynamics[39,40]. Another common unknown is current state of the PLD target where changes during ablation or between similar targets can cause a variation in plume dynamics and film stoichiometry[41]. Because all of these effects cause changes in the plume dynamics, deep learning can be used to encode the essential features of the plume dynamics under different P and E for a single experimental campaign and be used for feedback in future experiments. For example, when attempting to replicate depositions that are separated by significant amounts of time or other interim experiments—as would be the case in a User Facility, for example—an ICCD image sequence can be fed to a model to compare its predictions to the current conditions. If the model consistently predicts parameters that are far from current values, this is a clear indicator of anomalous, unreplicated plume conditions. Moreover, during a rapid sequence of depositions like in an autonomous synthesis campaign, model predictions can be used for compensatory purposes, such as mitigating the effects of laser entry window coating or resolving problems in the process gas flow control system.

For anomaly detection, we use only the ICCD image sequence and train the model to predict P, E_1, and E_2 (ICCD Features $(2 + 1)$D CNN in Fig. 1). The plume dynamics within the 2–150 μs delay time showed no correlation with the substrate temperature T, as anticipated. Figure 2a shows

the coefficient of determination r^2 for P, E_1, E_2, and the mean of all three as a function of training epoch. The model was trained for 2000 epochs and a checkpoint was taken at the epoch with the highest mean r^2. The pressure prediction has the greatest performance with $r^2_{val} = 0.963$ on the validation set with E_1 and E_2 having $r^2_{val} = 0.904$ and $r^2_{val} = 0.895$, respectively, giving an average $r^2_{val} = 0.921$. For comparison, we also trained a model using only a single ICCD image (2 μs delay) rather than the full sequence, using the same model architecture but without the temporal convolution. This single image model performed significantly worse than the image sequence, with $r^2_{val} = 0.850$, 0.814, and 0.810 for P, E_1, and E_2, respectively, a mean $r^2_{val} = 0.825$. These results indicate that deep learning with ICCD image sequences can be used to effectively predict PLD processing parameters for anomaly detection and that including both the spatial and temporal components of the plume dynamics in deep learning greatly increases model accuracy.

The current model is trained for a relatively atypical PLD scenario using co-ablation of two targets, which creates complex plume-plume interactions, during growth of a transition metal dichalcogenide. The more typical scenario would be ablation of a single target during growth of an oxide material. In order to apply the model to a more standard PLD scenario, the model should be retrained on imaging data for the specific target material with the single-target geometry. During our initial exploration of data preprocessing and model selection, we tested the much simpler, single-target scenario by training on single images of only the WSe_2 target, cropped from the full images, which were still reasonably effective for predicting P. We believe that training on ICCD image sequences of single target ablation would also be as successful as the current co-ablation case. Other experimental parameters that cause changes to the plume dynamics could also be included for prediction, such as the excimer laser spot size. In addition, the current $(2 + 1)$D model has 3.09 million parameters but trains quickly on GPUs, taking ~0.3 s/epoch during training when using an Nvidia A100 GPU, with 1 batch/epoch of size 88. Therefore, the data collection and retraining of similar models for different materials and PLD geometries is a
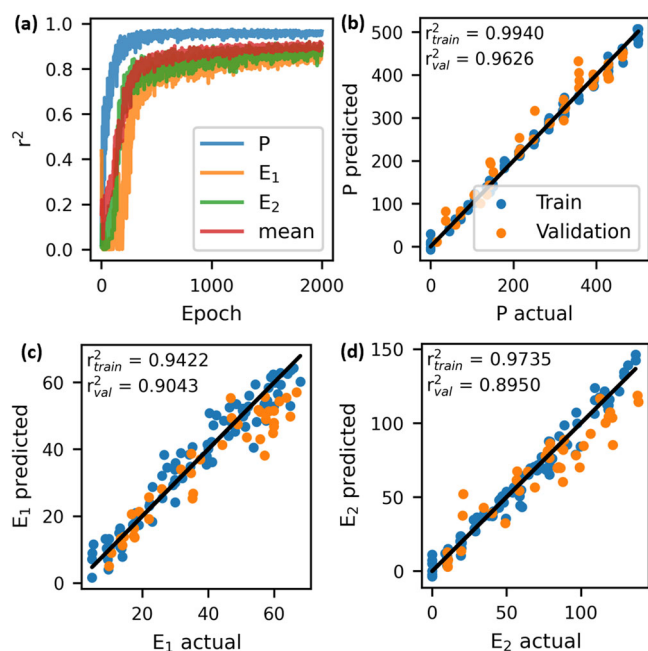
Fig. 2 | Training results for the anomaly detection model which predicts the chamber pressure P and laser energies E_1 and E_2 from intensified-CCD image sequences using a $(2 + 1)$D convolutional neural network. a Validation set coefficient of determination r^2 vs. epoch for each predicted quantity and the mean r^2. Predicted vs actual **b** P, **c** E_1, and **d** E_2 with a mean $r^2 = 0.921$ on the validation set.

reasonably fast process that we envision could be easily incorporated as part of the setup for a series of PLD experiments on different PLD systems to help ensure reproducible film growth.

Prediction of growth kinetics

Prediction of film properties or growth kinetics from plume diagnostics is attractive because it enables partial exploration of the PLD parameter space with a lower time and material cost, which accelerates the rate of discovery or optimization of materials. While some film properties may be dominated by thermodynamics, features of the plume may be able to explain some portion of the variance within a dataset. Different species within a multielement PLD plume (e.g., WSe_2) take on different angular distributions which can affect film stoichiometry[42] thus altering film properties. Depending on the laser and pressure conditions, the plume is composed of a mix of atomic and molecular species, nanoparticles, and large particulates[43,44]. These "components" affect the plume expansion dynamics and optical emission which can be detected by ICCD imaging. In the context of autonomous synthesis with PLD, ICCD image sequences could be used to provide additional information to the ML models that determine the sequence of experiments. For instance, if ICCD data can be encoded to predict at least some of the variance in an optimization variable, it shows promise to be utilized for deep kernel learning[45].

To demonstrate the predictive capabilities of ICCD imaging on quantities related to the growing film, we train several models to predict the growth kinetics of the WSe_2 films as determined from laser reflectivity. An example reflectivity curve is shown in Fig. 1 and the 3 target quantities are s_0, s_1, and J (described in the section "Data generation"). We trained three different models to compare performance between ICCD features only, growth parameter (P, T, E_1, E_2) features only, and both features combined (Fig. 1). Each model was trained for 2000 epochs and the checkpoint with the highest mean r^2_{val} was saved. For the combined model, the optimized model states from the ICCD $(2 + 1)$D CNN and the growth parameter MLP were loaded and frozen so that only the combined MLP (see Fig. 1) was trained and optimized. Unfreezing the ICCD and growth parameter feature subnetworks and retraining did not lead to significantly different results.

Figure 3a shows the learning curve of the validation set r^2 vs epoch for each quantity along with the mean value and Fig. 3b–d shows the predicted vs. actual values. Interestingly, the ICCD image sequence alone has reasonable predictive power for the growth kinetics, with a mean $r^2_{val} = 0.815$. Since thin film growth is essentially dominated thermodynamics, it is unexpected to achieve such good predictions without knowledge of the substrate temperature. We believe that the images are correlated to the growth kinetics in this case because the relative plume intensity (brightness) from each target along with the temporal dynamics (kinetic energy) are directly, but not trivially, correlated to the flux term J and the sticking coefficients s_0 and s_1, respectively. While s_0 and s_1 should be strongly correlated to the substrate temperature, the flux ratio (brightness) from the targets will also affect the growth rate and high kinetic energy may cause resputtering[46]. While these factors can be intuitively rationalized, the specific features of the plume dynamics that play a driving role in the growth kinetics are not clear beforehand, which underscores the value of deep learning with plume diagnostics.

By comparison, an MLP trained with only the growth parameters performs slightly better than the ICCD model with $r^2_{val} = 0.835$. Figure 3e–h shows the training results of the MLP. However, the MLP predictions are less stable during training, evident by the large variation in r^2 values. The best model captured during training has poor performance on the training set for s_0 and s_1. While this model accounts for temperature in the growth kinetics, the unstable performance is likely caused by experimental uncertainties related to E_1 and E_2. While this MLP has values for the laser energy, the real dynamics of the plume (partially determined the laser) are captured in the ICCD model which may explain why it has more stable predictions. Finally, we train a combined model and the results are shown in Fig. 3i–l). The combined model outperforms both previous models with an $r^2_{val} = 0.847$ with more stable predictions than the MLP. This result shows great promise ML models using combined multimodal PLD diagnostic measurements with basic process parameters for use in autonomous workflows.

The combined model can be used to predict the growth kinetics in a larger synthesis parameter space than the current dataset explored and identify conditions where the growth times or nucleation rates fall within specified values. Notably, the standard Bayesian optimization (BO) that was used in the growth study that generated this dataset[10] did not incorporate any plume imaging data. In experiments utilizing BO for microscopy, for example, it has been shown that utilizing the local image structure is highly beneficial for optimizing the targeted material properties, which is done by constructing a deep kernel learning (dKL) model[45] to map the local image patches to the functional property measured by a spectroscopic measurement. In this case, the equivalent is the multimodal diagnostics data that accompany each deposition. However, in the microscopy case, the high-resolution image is acquired prior to the spectroscopy, so all the image patches associated with individual pixels in the scan region are available to train the model at all times. Here, we note that it may be possible to utilize a similar dKL approach by for instance, performing ICCD-only experiments, which can be much faster (and cheaper) than completing an entire film deposition, to acquire this information. More approaches utilizing multifidelity Bayesian optimization are also possible, where this model would be considered the low-fidelity approximator, whereas the full film growth would be considered the high-fidelity portion.

Feature map analysis

Analyzing the feature maps (FMs) that are learned by the convolutional layers can potentially provide insight into the features of the plume dynamics that relate to the target. To explore what has been learned by the network, we visualize the activations of the growth kinetic model's $(2 + 1)$D CNN FMs with an example ICCD image sequence. Many of the FMs are difficult to interpret or look similar with subtle differences between them. Here, we show selected FMs which we can intuitively interpret. Figure 4a shows 5 frames from a processed input image with delay times from 2 to 26.2 µs. In the experiment, the WSe_2 and Se targets were mounted in the top and bottom target holder positions, respectively. Thus, the individual WSe_2 and
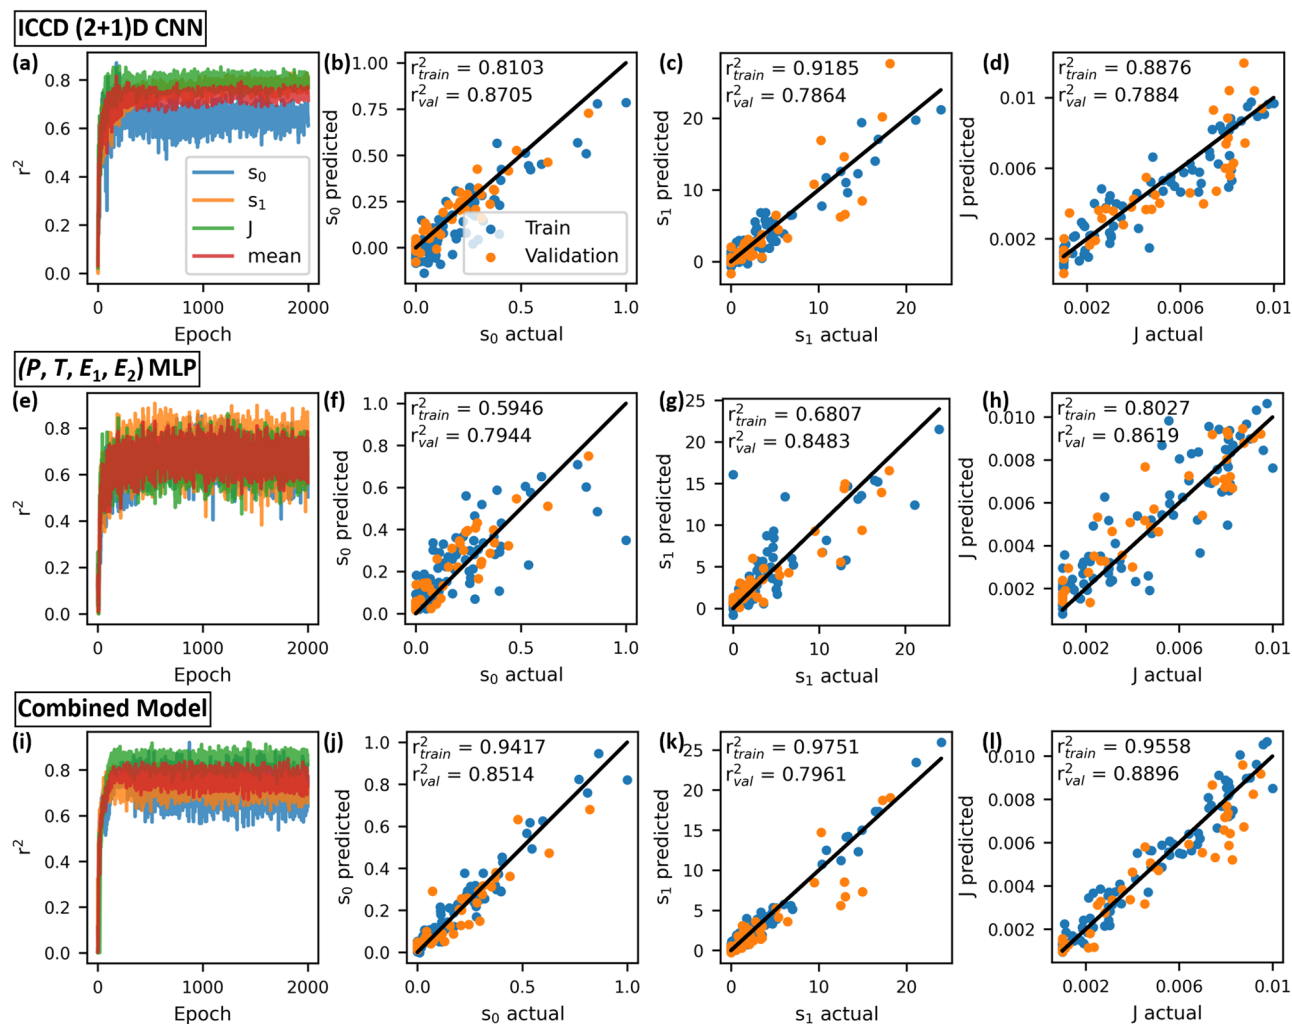
Fig. 3 | Training results for the WSe$_2$ growth kinetics models which predict the kinetic parameters s_0, s_1, and J using either a $(2+1)$D convolutional neural network (CNN) with intensified-CCD (ICCD) image sequences, a multilayer perceptron (MLP) with the pressure, substrate temperature, and laser energies (P, T, E_1, E_2), or both combined. **a–d** ICCD $(2+1)$D CNN model results have a mean $r^2_{val} = 0.815$. **e–h** Growth parameter MLP results with a mean $r^2_{val} = 0.835$. **i–l** Combined model using both ICCD and growth parameter features performs the best with mean $r^2_{val} = 0.847$.

Se plumes can be easily distinguished at the top and bottom of the first frame (2 µs). As plume expansion progresses, the Se plume collides with the WSe$_2$ plume (14.1 µs) which induces additional luminescence at the collision fronts (20.1 µs), likely caused by excitation of neutral atomic species and then the plumes intermix.

Figure 4b shows the activations of a FM from the 1st spatial convolution. Many of the initial FMs look similar to the input image which suggest that the early layers are learning simple features such as distinguishing the PLD plumes from the background as well as acting like simple edge filters. In Fig. 4b, the filter gives no weight to the background and positive values to the plume regions that is correlated with optical emission intensity. Figure 4c shows a FM from the spatial convolution in layer 2. This FM now includes some temporal dynamics, with positive and negative weights added in each frame in the direction of travel and also tends to highlight the edges of the plumes. The most complex activation we show is from the temporal convolution of layer 3, shown in Fig. 4d. In this case, positive weight seems to be given to the current location of the plume while negative weight could be interpreted as the direction of plume edge acceleration/deceleration. For instance, at 2 µs the Se plume is shown as positive weight, corresponding the the plume location in Fig. 4a. The negative region is indicating the direction of acceleration of the leading edge. At 14.1 µs, the Se plume has collided with the WSe$_2$ plume and the front begins to accelerate laterally, expanding the width of the plume font. Meanwhile 14.1 µs,

the back edge of the Se plume will begin to decelerate as it "snowplows" into the rest of the mass, indicated by the red region on the back behind the Se plume. The activations of the deepest convolutional layers clearly encode complex spatiotemporal dynamics, which is likely the origin of the increased performance relative to the model trained with a single image ("Anomaly detection").

We also visualized the saliency maps using the input gradient method[47], which computes pixel importance by multiplying the input image by the gradient of the model output. The saliency maps are also difficult to interpret but provide a method to visualize which regions of the image sequence are important for predictions. Based on the saliency maps, we find that the model is focused on features of the plume, rather than random areas or the background, which provides more support that spatiotemporal dynamics of the plume are critical features for model prediction. Saliency map visualizations are available in the code provided with this article.

In this work, we investigated the viability of utilizing deep learning with intensified-CCD (ICCD) image sequences of the plasma plume generated during pulsed laser deposition (PLD) for real-time feedback and predictions during thin film growth. Our findings indicate that deep learning, particularly the use of $(2+1)$D convolutional neural networks, can effectively extract complex spatiotemporal features from PLD plume dynamics that are correlated with synthesis conditions and growth kinetics. We demonstrated
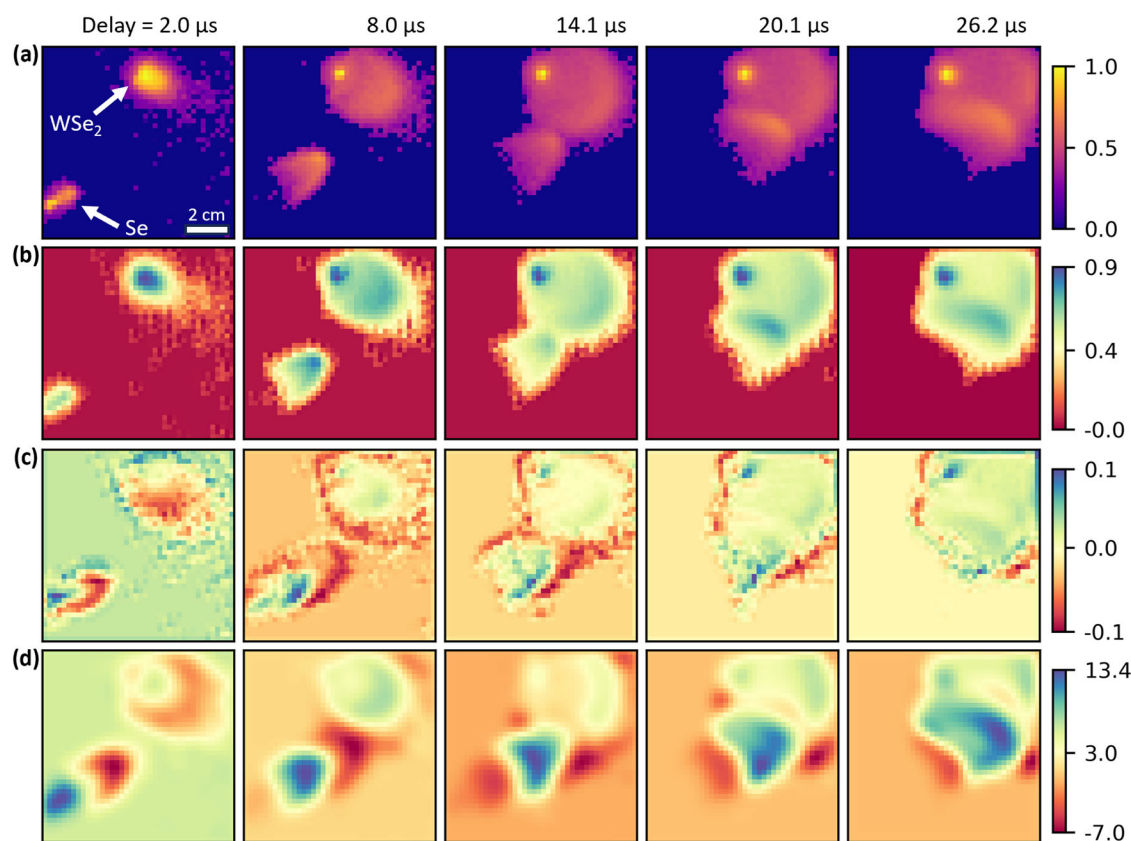
Fig. 4 | Activations for select feature maps (FM) of the convolutional layers in the (2 + 1)D convolutional neural network model for growth kinetics highlight the deep features learned from the plumes generated during pulsed laser deposition. a Five frames from the processed input image at delay times between 2.0–26.2 μs show the expansion and collision of the WSe_2 and Se plumes. **b** FM from the spatial convolution of layer 1 is similar to the input and discerns the plumes from the background. **c** FM from the 2nd layer spatial convolution appears to function like an edge filter, giving very little weight to the background and plume centers while highlighting the edges. **d** The 3rd layer's temporal convolution encodes spatio-temporal dynamics, giving positive weight to the current location of each plume body and negative weight near edges in the direction of acceleration/deceleration in each frame.

that ICCD images are highly correlated with chamber pressure and laser energy, thereby providing a means for real-time plume monitoring to detect anomalous conditions during long, unsupervised autonomous experiments. We also showed that plume dynamics are a viable predictor of a thin film growth kinetics model parameters and that when incorporated with basic synthesis parameters, the model performance is increased. These two case studies highlight how the marriage of in situ plasma diagnostics and machine learning can provide real-time feedback for PLD synthesis with predictive power over the growth environment and materials properties. This work serves as a proof-of-principle for future applications of ICCD images for use in deep kernel learning or multi-fidelity optimization experiments with PLD. We anticipate that this work will encourage more widespread adoption of plume imaging techniques, including lower cost CMOS imaging, during PLD for increased reproducibility and accelerated optimization or discovery of materials.

Data availability
The data that support the findings of this study are openly available at https://github.com/sumner-harris/Deep-Learning-with-ICCD-Images.

Code availability
The code for this study can be accessed via this link https://github.com/sumner-harris/Deep-Learning-with-ICCD-Images.

References
1. Volk, A. A. et al. AlphaFlow: autonomous discovery and optimization of multi-step chemistry using a self-driven fluidic lab guided by reinforcement learning. *Nat. Commun.* **14**, 1403 (2023).
2. Burger, B. et al. A mobile robotic chemist. *Nature* **583**, 237–241 (2020).
3. Manzano, J. S. et al. An autonomous portable platform for universal chemical synthesis. *Nat. Chem.* **14**, 1311–1318 (2022).
4. Coley, C. W. et al. A robotic platform for flow synthesis of organic compounds informed by AI planning. *Science* **365**, eaax1566 (2019).
5. Szymanski, N. J. et al. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature* **624**, 86–91 (2023).
6. Bianchini, M. et al. The interplay between thermodynamics and kinetics in the solid-state synthesis of layered oxides. *Nat. Mater.* **19**, 1088–1095 (2020).
7. Miura, A. et al. Observing and modeling the sequential pairwise reactions that drive solid-state ceramic synthesis. *Adv. Mater.* **33**, 2100312 (2021).
8. Harris, S. B. et al. Real-time diagnostics of 2D crystal transformations by pulsed laser deposition: controlled synthesis of Janus WSSe monolayers and alloys. *ACS Nano* **17**, 2472–2486 (2023).
9. Nikolaev, P. et al. Autonomy in materials research: a case study in carbon nanotube growth. *npj Comput. Mater.* **2**, 16031 (2016).
10. Harris, S. B. et al. Autonomous synthesis of thin film materials with pulsed laser deposition enabled by in situ spectroscopy and automation. Preprint at https://arxiv.org/abs/2308.08700 (2023).
11. Tischler, J. Z. et al. Nonequilibrium interlayer transport in pulsed laser deposition. *Phys. Rev. Lett.* **96**, 226104 (2006).

12. Gruenewald, J. H., Nichols, J. & Seo, S. S. A. Pulsed laser deposition with simultaneous in situ real-time monitoring of optical spectroscopic ellipsometry and reflection high-energy electron diffraction. *Rev. Sci. Instrum.* **84**, 043902 (2013).

13. Puretzky, A. A. et al. In situ laser reflectivity to monitor and control the nucleation and growth of atomically thin 2D materials. *2D Mater.* **7**, 025048 (2020).

14. Langereis, E., Heil, S. B. S., van de Sanden, M. C. M. & Kessels, W. M. M. In situ spectroscopic ellipsometry study on the growth of ultrathin TiN films by plasma-assisted atomic layer deposition. *J. Appl. Phys.* **100**, 023534 (2006).

15. Mahjouri-Samani, M. et al. Nonequilibrium synthesis of TiO_2 nanoparticle "building blocks" for crystal growth by sequential attachment in pulsed laser deposition. *Nano Lett.* **17**, 4624–4633 (2017).

16. Lin, Y.-C. et al. Low energy implantation into transition-metal dichalcogenide monolayers to form Janus structures. *ACS Nano* **14**, 3896–3906 (2020).

17. Glavin, N. R. et al. Synthesis of few-layer, large area hexagonal-boron nitride by pulsed laser deposition. *Thin Solid Films* **572**, 245–250 (2014).

18. Lee, H. N., Ambrose Seo, S. S., Choi, W. S. & Roueau, C. M. Growth control of oxygen stoichiometry in homoepitaxial $SrTiO_3$ films by pulsed laser epitaxy in high vacuum. *Sci. Rep.* **6**, 19941 (2016).

19. Irimiciuc, S. A., Chertopalov, S., Lancok, J. & Craciun, V. Langmuir probe technique for plasma characterization during pulsed laser deposition process. *Coatings* **11**, 762 (2021).

20. Doggett, B. & Lunney, J. G. Langmuir probe characterization of laser ablation plasmas. *J. Appl. Phys.* **105**, 033306 (2009).

21. Lee, K. K. et al. Using neural networks to construct models of the molecular beam epitaxy process. *IEEE T. Semicond. M.* **13**, 34–45 (2000).

22. Vasudevan, R. K., Tselev, A., Baddorf, A. P. & Kalinin, S. V. Big-data reflection high energy electron diffraction analysis for understanding epitaxial film growth processes. *ACS Nano* **8**, 10899–10908 (2014).

23. Provence, S. R. et al. Machine learning analysis of perovskite oxides grown by molecular beam epitaxy. *Phys. Rev. Mater.* **4**, 083807 (2020).

24. Liang, H. et al. Application of machine learning to reflection high-energy electron diffraction images for automated structural phase mapping. *Phys. Rev. Mater.* **6**, 063805 (2022).

25. Glavin, N. R. et al. Temporally and spatially resolved plasma spectroscopy in pulsed laser deposition of ultra-thin boron nitride films. *J. Appl. Phys.* **117**, 165305 (2015).

26. Geohegan, D. B. & Puretzky, A. A. Dynamics of laser ablation plume penetration through low pressure background gases. *Appl. Phys. Lett.* **67**, 197–199 (1995).

27. Schou, J., Amoruso, S. & Lunney, J. G. *Plume Dynamics*, 67–95 (Springer US, 2007).

28. Kuya, Y., Takeda, K., Zhang, X. & Forrester, A. I. J. Multifidelity surrogate modeling of experimental and computational aerodynamic data sets. *AIAA J.* **49**, 289–298 (2011).

29. Tran, A., Tranchida, J., Wildey, T. & Thompson, A. P. Multi-fidelity machine-learning with uncertainty quantification and Bayesian optimization for materials design: application to ternary random alloys. *J. Chem. Phys.* **153**, 074705 (2020).

30. Tran, D. et al. A closer look at spatiotemporal convolutions for action recognition. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6450–6459 (IEEE Computer Society, 2018).

31. Maas, A. L. et al. Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, Vol. 30, 3 (Atlanta, 2013).

32. Shimizu, R., Kobayashi, S., Watanabe, Y., Ando, Y. & Hitosugi, T. Autonomous materials synthesis by machine learning and robotics. *APL Mater.* **8**, 111110 (2020).

33. MacLeod, B. P. et al. Self-driving laboratory for accelerated discovery of thin-film materials. *Sci. Adv.* **6**, eaaz8867 (2020).

34. Chang, J. et al. Efficient closed-loop maximization of carbon nanotube growth rate using Bayesian optimization. *Sci. Rep.* **10**, 9040 (2020).

35. Kingma, D. & Ba, J. Adam: a method for stochastic optimization. Preprint at https://arxiv.org/abs/1412.6980 (2014).

36. Liaw, R. et al. Tune: a research platform for distributed model selection and training. Preprint at https://arxiv.org/abs/1807.05118 (2018).

37. Akiba, T., Sano, S., Yanase, T., Ohta, T. & Koyama, M. Optuna: a next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2019).

38. Li, L. et al. A system for massively parallel hyperparameter tuning. In *Proceedings of Machine Learning and Systems 2020, MLSys 2020, Austin, TX, USA, March 2–4, 2020* (eds. Dhillon, I. S., Papailiopoulos, D. S. & Sze, V.) (mlsys.org, 2020).

39. Cazzaniga, A. et al. Ultra-thin Cu_2ZnSnS_4 solar cell by pulsed laser deposition. *Sol. Energ. Mat. Sol. C.* **166**, 91–99 (2017).

40. Ohnishi, T., Lippmaa, M., Yamamoto, T., Meguro, S. & Koinuma, H. Improved stoichiometry and misfit control in perovskite thin film formation at a critical fluence by pulsed laser deposition. *Appl. Phys. Lett.* **87**, 241919 (2005).

41. Cesaria, M. et al. Pulsed laser deposition of $CsPbBr_3$ films: impact of the composition of the target and mass distribution in the plasma plume. *Nanomaterials* **11**, 3210 (2021).

42. Ojeda-G-P, A., Döbeli, M. & Lippert, T. Influence of plume properties on thin film composition in pulsed laser deposition. *Adv. Mater. Interfaces* **5**, 1701062 (2018).

43. Geohegan, D. B., Puretzky, A. A., Duscher, G. & Pennycook, S. J. Photoluminescence from gas-suspended SiOx nanoparticles synthesized by laser ablation. *Appl. Phys. Lett.* **73**, 438–440 (1998).

44. Kroto, H. W., Heath, J. R., O'Brien, S. C., Curl, R. F. & Smalley, R. E. C_{60}: Buckminsterfullerene. *Nature* **318**, 162–163 (1985).

45. Liu, Y. et al. Experimental discovery of structure–property relationships in ferroelectric materials via active learning. *Nat. Mach. Intell.* **4**, 341–350 (2022).

46. Fähler, S., Sturm, K. & Krebs, H.-U. Resputtering during the growth of pulsed-laser-deposited metallic films in vacuum and in an ambient gas. *Appl. Phys. Lett.* **75**, 3766–3768 (1999).

47. Simonyan, K., Vedaldi, A. & Zisserman, A. Deep inside convolutional networks: Visualising image classification models and saliency maps. Preprint https://arxiv.org/abs/1312.6034 (2014).

Acknowledgements

Author contributions

S.B.H.: Conceptualization (equal); writing—original draft (lead); writing—review & editing (equal); data curation (lead); methodology (lead); software (lead); visualization (lead). C.M.R.: Resources (equal); conceptualization (supporting); writing—review & editing (equal). K.X.: Resources (equal); writing—review & editing (equal). R.K.V.: Conceptualization (equal); writing—review & editing (equal); software (supporting). All authors read and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Sumner B. Harris.

Reprints and permissions information is available at http://www.nature.com/reprints