arXiv:2411.18721v1 [cond-mat.mtrl-sci] 27 Nov 2024

A Machine Learning Approach Capturing Hidden Parameters in Autonomous Thin-Film Deposition

Yuanlong Zheng[1,2], Connor Blake[2], Layla Mravac[2], Fengxue Zhang[3], Yuxin Chen[3], and Shuolong Yang[2]
[1] *Department of Physics, University of Chicago, Chicago, IL 60637, USA*
[2] *Pritzker School of Molecular Engineering, University of Chicago, Chicago, IL 60637, USA and*
[3] *Department of Computer Science, University of Chicago, Chicago, IL 60637, USA*
(Dated: December 2, 2024)

Abstract

The integration of machine learning and robotics into thin film deposition is transforming material discovery and optimization. However, challenges remain in achieving a fully autonomous cycle of deposition, characterization, and decision-making. Additionally, the inherent sensitivity of thin film growth to hidden parameters such as substrate conditions and chamber conditions can compromise the performance of machine learning models. In this work, we demonstrate a fully autonomous physical vapor deposition system that combines in-situ optical spectroscopy, a high-throughput robotic sample handling system, and Gaussian Process Regression models. By employing a calibration layer to account for hidden parameter variations and an active learning algorithm to optimize the exploration of the parameter space, the system fabricates silver thin films with optical reflected power ratios within 2.5% of the target in an average of 2.3 attempts. This approach significantly reduces the time and labor required for thin film deposition, showcasing the potential of machine learning-driven automation in accelerating material development.

I. INTRODUCTION

In physical vapor deposition (PVD) of thin film materials, the traditional human-led process encompasses numerous cycles of selecting deposition parameters, performing deposition, characterizing film properties, and re-adjusting the parameters accordingly. The recent development of machine learning (ML) [1], coupled with advancements in robotics [2], now potentially enables fully automating this process, liberating researchers from the repetitive cycles and accelerating the optimization of material properties [3].

Several studies have sought to integrate ML with the PVD process. Typically, these approaches involved training ML models that map deposition parameters, such as substrate temperature, deposition rate, and flux ratio, to material properties, such as stoichiometry [4–6], electrical conductivity [7–9], surface morphology [10, 11], crystallinity [12–17], and superconducting critical temperatures [18]. The trained models are then used to predict material properties, with Bayesian optimization (BO) frequently employed to autonomously determine the deposition parameters for subsequent samples [9, 19, 20]. Such optimization has the potential to replace human decision-making, efficiently exploring the deposition parameter space to optimize the material properties [21].

However, ML models are known for their sensitivity to noise in the training data. In the meantime, the thin film deposition process can be prone to irreproducibility due to factors such as different initial substrate conditions and chamber environments [5, 8, 18]. The inconsistency in the growth outcome may be treated by the model as noise which undermines the model's performance. These "hidden parameters" that lead to the irrepeatability of thin film properties are often difficult to capture and incorporate into the model. The field of ML-assisted thin film deposition currently lacks a systematic approach to effectively account for these hidden parameters.

Beyond algorithmic challenges, the full potential of ML-assisted PVD hinges on the complete automation of hardware systems. PVD systems require high vacuum (HV) or ultra-high vacuum (UHV) environments, which present significant challenges for fully automating sample transfer and characterization processes. As such, most studies in the field of ML-assisted thin film deposition still rely on traditional manual handling of samples, which limits sample throughput and hinders the realization of the fully autonomous ML-assisted PVD [9, 18, 22]. Shimizu et al. demonstrated a system fully automating the deposition of Nb-doped TiO_2 films and the minimization of its resistance [23]. However, the system requires complex multi-chamber systems with sophisticated transfer mechanisms, thereby increasing the complexity of the setup and limiting its scalability. Achieving a scalable method to fully automate PVD systems is a critical hurdle to unlocking the full promise of ML integration in thin film deposition.

II. METHODS AND RESULTS

In this work, we address inefficiencies in current ML-driven thin film deposition by developing a fully automated system with high sample throughput and a mechanism to account for hidden parameter variations. Our system integrates a UHV chamber with a 72-slot robotic sample handling system optimized for maximizing throughput, in-situ optical characterization, and ML-powered prediction of characterization outputs and optimal growth conditions. To account for the varying hidden parameters, we introduce a calibration layer that provides critical information about the initial condition

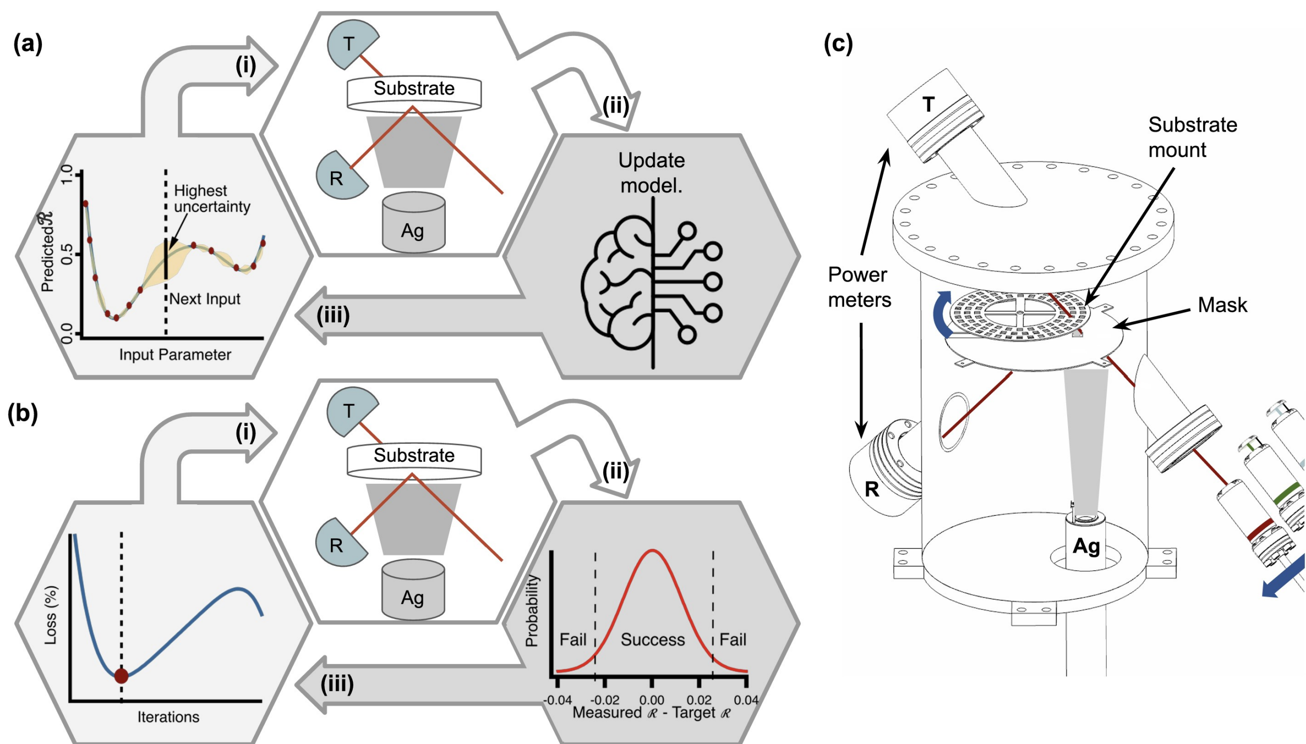

FIG. 1: A self-learning autonomous physical vapor deposition system for silver thin film growth. (a) The autonomous learning cycle incorporates (i) identification of the growth condition with the highest model uncertainty, (ii) sample growth, and (iii) updating the model with new data. (b) The autonomous testing cycle incorporating (i) prediction of the optimal growth condition that minimizes the loss function, (ii) sample growth, and (iii) assessment of success. (c) Schematic illustration of the autonomous deposition setup featuring robotic sample handling and in-situ optical characterization.

of each sample and the chamber environment. Furthermore, we employ an autonomous learning approach using Gaussian Process Regression (GPR) models to efficiently explore the multidimensional input space and accurately predict the target properties of the films. By integrating these techniques, we have engineered an ML-driven, closed-loop deposition system, thereby eliminating the need for human intervention at any stage of the growth cycle. We demonstrate the autonomous synthesis of silver thin films with optical properties within 2.5% of the user-specified target in an average of 2.3 attempts, underscoring the potential of our ML-driven system for significantly accelerating the optimization of material properties.

1. Introduction of the System

To demonstrate the principles of ML-driven autonomous PVD, we seek to fabricate silver thin films with user-specified optical properties. The growth conditions of silver films such as the growth rate, substrate temperature, and film thickness [24] impact the films' porosity [25], grain size [26], and electron-phonon interactions [27]. All of these aspects are correlated with the real and imaginary parts of the optical constants for silver films. It is difficult to model all of these mechanisms using simple physics laws, which warrants an ML-driven material optimization.

Our autonomous PVD system incorporates a shadow mask beneath a 72-slot sample handling system, ensuring that only one sample is exposed to the molecular beam at a time (Figure 1(c)). Silver (99.999%, Thermo Fisher) is deposited onto double-side polished BK7 substrates using an effusion cell (MBE-Komponenten) at a base pressure of $< 5 \times 10^{-9}$ mbar and a deposition pressure of 1×10^{-8} mbar. The reflectivity and absorptivity of the silver thin films are characterized using five p-polarized lasers with wavelengths of 443, 514, 689, 781, and 817 nm (Coherent StingRay). The lasers are mounted on a computer numerically controlled (CNC) linear rail (Figure 1(c)) pointing at the substrate with an incident angle of 45 degrees. The transmitted (P_t) and reflected power (P_r) are measured, and we define the reflected power ratio $\mathscr{R}$ and absorptivity $\mathscr{A}$:

$$\mathscr{R} = \frac{P_r}{P_r + P_t}, \quad \mathscr{A} = \frac{P_i - P_r - P_t}{P_i}$$

where P_i denotes the incident power. For convenience, the reflected power ratio and the absorptivity at 443 nm

are denoted as $\mathscr{R}_{443}$ and $\mathscr{A}_{443}$, respectively, and similarly for other wavelengths.

All hardware control and data acquisition are managed by MATLAB scripts which allow the system to grow and characterize up to 72 samples consecutively without human intervention. The optical characterization data are fed into the ML algorithm, which predicts the growth condition required for the sample to attain user-specified $\mathscr{R}$ values at targeted wavelengths. The process is divided into the autonomous learning and autonomous testing stages. In the learning stage, optical characterization data are used to train ML models, which then determine the next point to explore for further learning (Figure 1(a)). In the testing stage, the trained model guides the system to fabricate silver thin films with user-specified optical properties (Figure 1(b)).

2. Calibration Layer

For thin film deposition, slight differences in substrate and chamber conditions can significantly affect growth dynamics, especially for the first few atomic layers [28]. Parameters such as substrate surface roughness, and residual pressure of various elements in the chamber are unknown to us. Although some of these parameters can be measured with various difficulties, it is infeasible for one to exhaust all the factors in the deposition process. To quantify the effects of these hidden parameters on the sample properties, we first grow 24 samples at an effusion cell temperature (T) of 875°C and a deposition time (t) of 5000 seconds. After 5000 seconds, the $\mathscr{R}_{817}$ of the samples has a distribution with a standard deviation of 0.061 (Figure 2(c)).

To account for this variability, we find that the $\mathscr{R}$'s of the 24 samples measured at $t = 1000$ seconds can be used to quantify the fluctuations in the growth dynamics and its correlation with $\mathscr{R}$'s at later times. Here we plot $\mathscr{R}_{443}$, $\mathscr{R}_{689}$, $\mathscr{R}_{817}$ at 5000 seconds against their $\mathscr{R}_{689}$ at 1000 seconds and fit them with a simple linear function (Figure 2 (a-c)). The standard deviation of the $\mathscr{R}_{817}$ at 5000 seconds is 0.061. Meanwhile, its root-mean-square deviation (RMSD) from the linear fit is 0.032. This suggests that $\mathscr{R}_{689}$ at 1000 seconds is moderately correlated with the variation of $\mathscr{R}$'s at later times. Therefore, by initially growing a thin layer at $T = 875$°C for 1000 seconds and measuring its $\mathscr{R}_{689}$, we partially captures the effect of the hidden parameters on the growth outcome. We define this layer as the *calibration layer*, and its $\mathscr{R}_{689}$ is denoted as $\mathscr{R}_c$. We choose 689 nm as the wavelength for $\mathscr{R}_c$ due to its median position among the five wavelengths, thereby providing the most representative information.

For all samples in the subsequent experiments, we first grow the calibration layer and record $\mathscr{R}_c$. The sample is then grown for time t at temperature T. $\mathscr{R}$ and $\mathscr{A}$ are collected for all wavelengths throughout the growth (Figure 3). To maximize data acquisition efficiency, we divide the total deposition time into 98-second blocks. Within each block, the reflected and transmitted power at each wavelength are sequentially measured starting at

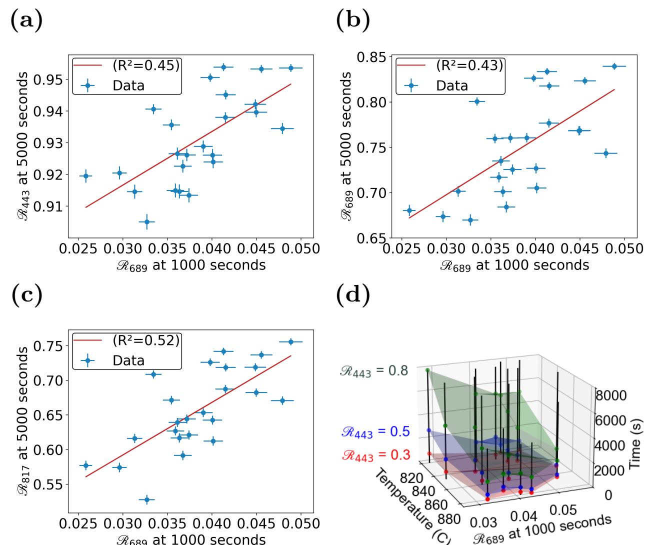

FIG. 2: Calibration layer accounts for variations in optical properties of silver films. (a-c) A linear fit between $\mathscr{R}_{443}$, $\mathscr{R}_{689}$, $\mathscr{R}_{817}$ at 5000 seconds and $\mathscr{R}_{689}$ at 1000 seconds. The moderate R-squared values indicate that the variances of $\mathscr{R}$'s at later times can be partially captured by $\mathscr{R}$ at 1000 seconds. (d) Translucent planes of time required to reach specified $\mathscr{R}_{443}$ values for given effusion cell temperatures and $\mathscr{R}_c$'s. $\mathscr{R}_c$ is negatively correlated with the time required to reach a given $\mathscr{R}_{443}$.

the time points 0, 16, 32, 48, and 64 seconds (Table 1). The intervals between measurements are limited by the speed of the linear rail and the collection period of 5 seconds per wavelength. Note that $\mathscr{R}_c$, though a measured value, is treated as an input parameter of the model in the dataset.

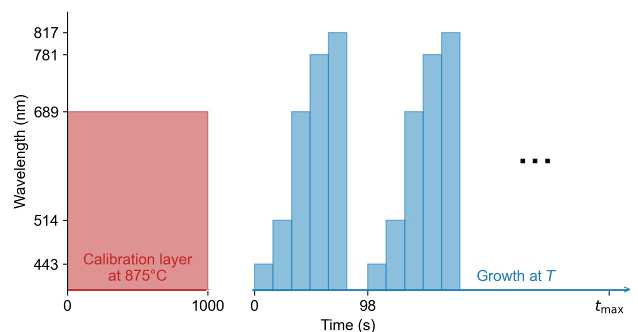

FIG. 3: Illustration of the data acquisition process for each iteration in the autonomous learning stage: the calibration layer is first grown to determine the $\mathscr{R}_c$, followed by measuring optical properties at all 5 wavelengths in cycles.

3. Autonomous Learning

GPR models are a form of probabilistic supervised learning models derived from non-parametric statistics.

Model Inputs				Model Outputs	
Temperature (°C)	Time (s)	$\mathscr{R}_c$	Wavelength (nm)	$\mathscr{R}$	$\mathscr{A}$
872.5	0	0.038	443	0.356	0.422
872.5	16	0.038	514	0.158	0.194
872.5	32	0.038	689	0.044	0.114
872.5	48	0.038	781	0.039	0.048
872.5	64	0.038	817	0.034	0.014
872.5	98	0.038	443	0.394	0.442
872.5	...	0.038	...	...	...
872.5	$t_{\max}$	0.038	817	0.824	0.363

TABLE I: An exemplary data set taken at the effusion cell temperature 872.5 ℃, with model inputs and outputs.

GPR models are defined as a continuous set of distributions $f(\mathbf{x})$ where f maps values $\mathbf{x} \in \mathbb{R}^n$ to the space of Gaussian probability density functions on y. For any two data points $(\mathbf{x}_1, y_1)$ and $(\mathbf{x}_2, y_2)$, the covariance between y_1 and y_2 is assumed to be purely a function of $\mathbf{x}_1$ and $\mathbf{x}_2$, $k(\mathbf{x}_1, \mathbf{x}_2)$.

A GPR model is fully characterized by its prior function $g(\mathbf{x})$, mean function $m(\mathbf{x}) \equiv \mathbb{E}[f(\mathbf{x})]$, and covariance function $k(\mathbf{x}_1, \mathbf{x}_2)$. The prior function $g(\mathbf{x})$ describes a "guess" for how the output space should appear in the absence of data. The mean function describes the expected y value for a given input $\mathbf{x}$ based on the set of data points $\{(\mathbf{x}_i, y_i)\}$. The covariance function, referred to as a kernel, describes how well one should be able to predict $\tilde{y}$ given $\tilde{\mathbf{x}}$ and $(\mathbf{x}, y)$. Here, we used a radial basis function (RBF) kernel.

$$k_{\mathrm{RBF}}(\mathbf{x}_1, \mathbf{x}_2) = \exp\left(-\frac{1}{2}(\mathbf{x}_1 - \mathbf{x}_2)^T \mathbf{\Theta}^{-\mathbf{2}} (\mathbf{x}_1 - \mathbf{x}_2)\right) \tag{1}$$

chosen for its ability to act as a universal approximator [29]. The length scale matrix $\mathbf{\Theta} = \mathrm{diag}(\theta_1, \theta_2, \ldots, \theta_n)$ contains the length-scale parameters which are optimized during learning and describe the width of the Gaussian kernel used for prediction.

Mapping the data and prior distribution to the predicted output (the posterior distribution) is known as conditioning and is computed via Bayes' rule. As such, the predicted posterior distribution at a point $\tilde{\mathbf{x}}$ based on observed data $\{(\mathbf{x}_i, y_i)\}$ is given by a normal distribution:

$$f(\tilde{\mathbf{x}}) \sim \mathcal{N}(k_i (K^{-1})_{ij} y_j, k(\tilde{\mathbf{x}}, \tilde{\mathbf{x}}) - k_i (K^{-1})_{ij} k_j) \tag{2}$$

where $k_i = k(\tilde{\mathbf{x}}, \mathbf{x}_i)$ and $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ for symmetric kernels like RBF. Here we have used the prior $g(\mathbf{x}) = 0$ for simplicity. From here, this distribution will be written as $\mathrm{GPR}_y(\mathbf{x}) = (\mu_y, \sigma_y)$ with μ_y and σ_y referred to as the "predicted mean" and "predicted uncertainty" respectively. While this distribution appears to be computationally intractable for large datasets, the function $m(\mathbf{x})$ has a closed-form representation for the RBF kernel and is readily computed by libraries such as GPyTorch [30].

We employ two GPR models for $\mathscr{R}$ and $\mathscr{A}$ with RBF kernels, both having input parameters $\mathbf{x} = (T, t, \lambda, \mathscr{R}_c)$, and outputs (μ_R, σ_R) and (μ_A, σ_A), respectively. Here, we define two modes of generating training data for models: predefined learning and autonomous learning. In predefined learning, the growth parameters of the training samples are set to divide the parameter space into equal increments. In autonomous learning, the model autonomously determines these growth parameters in a way that optimally learns the parameter space.

A series of predefined learning is used to initialize the model during which 9 samples are grown between 820 and 880°C with 7.5°C increments and for t of

$$t_{\max}(T) = 3.22 \times 10^{14} \times e^{-0.0285 T(^\circ C)} (\text{seconds}) \tag{3}$$

where $t_{\max}$ is determined such that each sample reaches a thickness yielding an $\mathscr{R} > 0.8$ for all wavelengths. This ensures the generation of data over a wide range of $\mathscr{R}$ values for training the model, while also avoiding the unnecessary time spent growing films until $\mathscr{R}$ asymptotically approaches 1. The exponential relationship between $t_{\max}$ and T assumes that silver's vapor pressure, and consequently the growth rate, increases exponentially with T, ensuring that all samples at $t_{\max}$ have approximately the same thickness. During the growth of each sample, $\mathscr{R}$ and $\mathscr{A}$ are repeatedly measured at each wavelength to obtain data between $t = 0$ and $t_{\max}$ (Table 1).

The system proceeds to the autonomous learning stage. After each sample's calibration layer is grown and $\mathscr{R}_c$ is measured, the effusion cell temperature T for the remainder of the growth is selected within the interval of [820, 880] °C according to:

$$T_{\text{selected}} = \arg\max_T \left\{ \frac{\sqrt{\sum_{t=0}^{t_{\max}(T)} \sum_\lambda \sigma_R(T, t, \lambda, \mathscr{R}_c)^2}}{t_{\max}(T)} \right\} \tag{4}$$

Here the T with the highest prediction uncertainty for all relevant times and wavelengths is selected for the growth. The sample is then grown and measured at a series of times, as shown in Figure (3), until $t = t_{\max}(T_{\text{grow}})$ specified by Eqn. (3). Such a complete deposition and characterization process, until the growth time reaches $t_{\max}(T_{grow})$, is termed *one iteration* (Figure 3). Note that the model only needs to choose T, among all input parameters (T, t, λ, $\mathscr{R}_c$), because data are collected for each sample repeatedly over time and for all wavelengths, while $\mathscr{R}_c$ is a result of the current chamber and substrate initial condition and cannot be controlled by the algorithm. After each iteration of deposition and characterization is finished, the GPR model is updated with the new data. This process is repeated for 8 iterations until the maximum uncertainty in the parameter space converges (Figure 4 (b)). The total learning stage consists

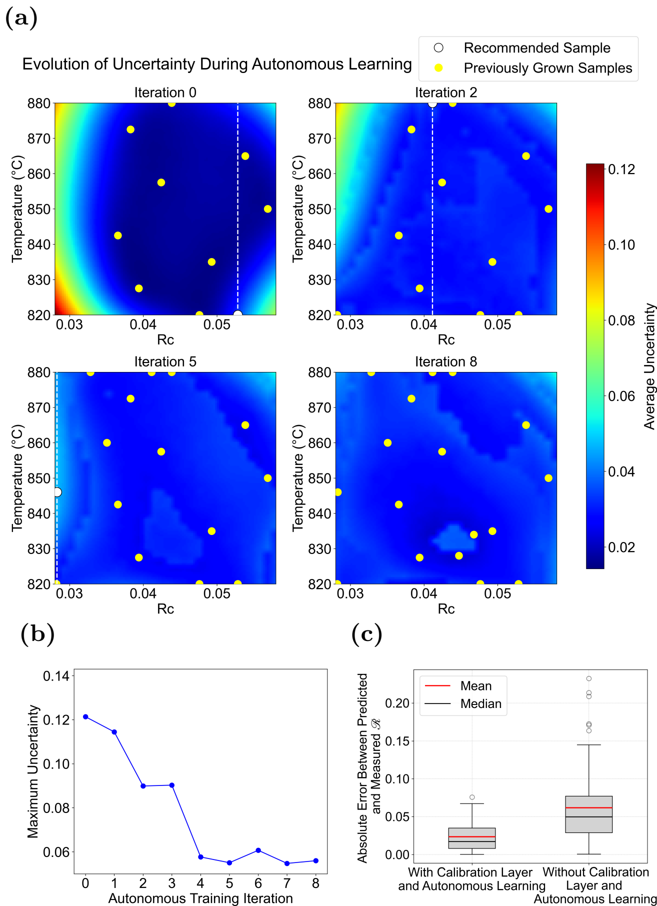

FIG. 4: Performance of autonomous learning in the fabrication of silver thin films. (a) Evolution of model predicted uncertainty, averaged over all $\mathscr{R}_\lambda$'s, during autonomous learning. (b) Convergence of the maximum uncertainty during autonomous learning. (c) Comparing the model prediction errors, defined as the differences between all measured and predicted $\mathscr{R}_\lambda$'s. The error distribution when adopting the calibration layer and autonomous learning is benchmarked against the case without these techniques.

of 17 samples, including 9 from predefined learning and 8 from autonomous learning.

Figure 4(a) shows the evolution of the average predicted uncertainty for all $\mathscr{R}$'s during autonomous learning. After 8 iterations, the uncertainty over the entire parameter space becomes more uniformly distributed and has an average value of 0.032. The uncertainty does not further decrease below this level. The remaining uncertainty is likely due to the hidden parameters unaccounted by the calibration layer and measurement noises (e.g. uncertainty in laser power measurements).

Unlike typical Bayesian optimizations aimed at optimizing a single output [31], our model is trained to acquire comprehensive information over the entire parameter space, enabling the system to respond to arbitrary $\mathscr{R}$ requests at given wavelengths.

4. Autonomous Testing

4.1 Single-wavelength $\mathscr{R}$ Targets

We select 5 random single-wavelength targets $\mathscr{R}_\lambda^{\text{target}}$, one for each wavelength of the lasers. Given a certain $\mathscr{R}_c$, there exist infinitely many (T, t) that can achieve the requested $\mathscr{R}$ at the specified wavelength. The degeneracy is removed by also aiming for the minimum $\mathscr{A}_\lambda$. The loss function for each single-wavelength target is defined as

$$\begin{aligned}\mathscr{L}_\lambda &= (\mu_{\mathscr{R},\lambda} - \mathscr{R}_\lambda^{\text{target}})^2 + 4\sigma^2_{\mathscr{R},\lambda} + \mu^2_{\mathscr{A},\lambda} + 4\sigma^2_{\mathscr{A},\lambda} \\ &\text{if } |\mu_{\mathscr{R},\lambda} - \mathscr{R}_\lambda^{\text{target}}| > 0.01 : \\ &\quad \mathscr{L}_\lambda \mapsto \mathscr{L}_\lambda + 100(\mu_{\mathscr{R},\lambda} - \mathscr{R}_\lambda^{\text{target}})^2 \end{aligned} \tag{5}$$

The uncertainties of predicted $\mathscr{R}$ and $\mathscr{A}$ are added to the loss function to penalize growth conditions with high uncertainties. The loss increases rapidly when the difference between the target and predicted $\mathscr{R}$ exceeds 0.01 to prioritize proximity to the target value of $\mathscr{R}$.

The loss function is minimized using the Adam optimizer [32], an adaptive learning rate optimization algorithm designed to handle sparse gradients on noisy problems. The set of growth parameters (T, t) at the minimum of the loss function is used for deposition. After deposition, the sample's measured $\mathscr{R}_\lambda$ is compared to $\mathscr{R}_\lambda^{\text{target}}$. If $|\mathscr{R}_\lambda - \mathscr{R}_\lambda^{\text{target}}| < 0.025$, the growth is considered successful, and the system moves to the next target. If unsuccessful, the model is updated with the obtained optical characterization data from the current sample and re-attempts the target until success. Table 2 displays the results for 5 single-wavelength targets. It took 2 attempts on average for a target to be successfully achieved. For each of the 10 samples grown during this stage, the model makes 5 predictions on its $\mathscr{R}$ for each wavelength. For these 50 total predictions, the mean absolute error (MAE) between the model predictions and measured results is 0.0246, which demonstrates the accuracy of the prediction. Moreover, the average model predicted uncertainty is 0.0267, and its proximity to the MAE demonstrates the accuracy of the model's estimate of its uncertainty.

To further demonstrate the effectiveness of the implementation of calibration layers and autonomous learning, we perform a control experiment without these methods to benchmark the model prediction performance. In the control experiment, the model only goes through the predefined learning process, in which 17 samples are grown at T between 820 and 880°C with 3.75°C increments, and for t as specified in Eqn.(3). The system is then requested to produce silver thin films that satisfy the same 5 single-wavelength target. Despite the control experiment having the same amount of training sample data as the previous experiment, it requires on average 3.6 attempts to successfully achieve each target. The MAE between the model predictions and measured results is 0.0618, also significantly increased from the previous experiment (Figure 4(c)). The control experiment hence

demonstrates the superior performance of the model using calibration layers and autonomous learning.

Target	Attempt 1	2	3	4	5	6	7
$\mathscr{R}_{443}^{\text{target}} = 0.61$	0.609						
$\mathscr{R}_{514}^{\text{target}} = 0.47$	0.413	0.452					
$\mathscr{R}_{689}^{\text{target}} = 0.88$	0.891						
$\mathscr{R}_{781}^{\text{target}} = 0.30$	0.298						
$\mathscr{R}_{817}^{\text{target}} = 0.58$	0.545	0.636	0.541	0.510	0.585		
$\mathscr{R}_{443}^{\text{target}} = 0.85$	N/A	N/A	0.914	0.845			
$\mathscr{R}_{781}^{\text{target}} = 0.47$			0.654	0.507			
$\mathscr{R}_{443}^{\text{target}} = 0.85$	N/A	N/A	0.797	N/A	0.879	0.779	0.817
$\mathscr{R}_{781}^{\text{target}} = 0.35$			0.331		0.489	0.325	0.342

TABLE II: Results of achieving specified optical properties with single and multi-wavelength targets.

4.2 Multi-wavelength Target

We also specify $\mathscr{R}$ targets spanning multiple wavelengths. The loss function is defined as

$$\mathscr{L} = \sum_{\{\lambda\}_{\text{specified}}} (\mu_{\mathscr{R},\lambda} - \mathscr{R}_{\lambda}^{\text{target}})^2 \tag{6}$$

As the number of wavelengths specified in the target increases, it is not guaranteed that there exists a set of growth parameters $(T, t, \mathscr{R}_c)$ that could produce a film with the desired optical properties. Therefore, when the loss function is minimized and any of the $\mu_{\mathscr{R},\lambda}$'s is still $>$ 0.01 away from the target, the algorithm decides that no growth parameters can be found to reach the target at the current sample's $\mathscr{R}_c$. The system aborts this sample and grows the calibration layer on the subsequent substrates until it finds one with a set of $(T, t, \mathscr{R}_c)$ that yields $\mu_{\mathscr{R},\lambda}$ within ± 0.01 of the target for all specified wavelengths.

The system proceeds to deposit the sample at the found (T, t) and subsequently measures its optical properties $\{\mathscr{R}_\lambda, \mathscr{A}_\lambda\}$. The system then updates the model with all collected data. If

$$\sum_{\{\lambda\}_{\text{specified}}} \left|\mathscr{R}_\lambda - \mathscr{R}_{\lambda}^{\text{target}}\right| < 0.025n \tag{7}$$

where n is the number of specified wavelengths in the target, the growth is considered success, and the system moves on to the next target. If unsuccessful, the system makes a new prediction using the updated model and re-attempts the target until success.

Table 2 displays the results using multi-wavelength targets. For certain samples, the model is unable to find a set of growth parameters that can yield a film to reach the specified multi-wavelength targets, denoted as N/A. These samples do not count towards the number of failed attempts. The system aborts the sample and proceeds to the next substrate. Upon switching to a sample with an $\mathscr{R}_c$ that leads to an achievable growth condition to obtain the desired optical properties, it begins the growth and can achieve the target within the allowed error. In this stage, 6 attempts are made to successfully achieve 2 targets. The success in achieving the two multi-wavelength targets demonstrates the versatility of our system to control the silver thin film's power-splitting spectrum over multiple wavelengths.

III. DISCUSSION AND CONCLUSION

In this work, we have established the ability of a fully autonomous ML-enabled system to address persisting challenges in the field of PVD thin film growth, namely the lengthy and labor-intensive human involvement and the difficulty of adapting to hidden parameters in the growth conditions. We see three components critical to a fully automated thin film deposition system: full mechanical and software automation, active learning for exploring the parameter space during data collection, and active targeting of desired outputs. Previous studies have implemented two of the three such as automation and active data collection [33, 34] or manually-executed active training and target-seeking [9], but to our knowledge, none have performed all three. Our work demonstrates the first system to implement all three in the field of thin film deposition. Using Bayesian ML, we have engineered a fully autonomous system that learns the dynamics of silver growth while accounting for variations in substrates and chamber conditions, and reliably grows samples with desired optical properties.

We choose to work with silver thin films because they are well-understood but retain the difficult aspects of thin film growth. As such, the methods that we develop are broadly applicable. The procedure we demonstrate of using predefined learning to map the parameter space, autonomous learning to minimize model uncertainty, and optimization to seek specific targets is readily generalizable to other thin film growth tasks.

With the fully autonomous system, we grow a total of 38 samples in the training and testing stage and collect more than 20,000 data points without human intervention. This represents a significant advantage over human-led experimentation. This high-throughput operation is primarily enabled by our use of a 72-slot sample handing system and in-situ quasi-continuous optical characterization. By designing chambers for high-capacity sample transfer and measurement, similar throughput can be achieved for other thin film fabrication and in-situ characterization tasks.

In contrast to other ML techniques such as artificial neural networks, Bayesian optimization allows the quantification of prediction uncertainties. This aspect has been used to decide optimal points of parameter space exploration [9, 18, 19]. Moreover, because of this ability

to quantify uncertainty, our model sets numerical thresholds to determine both when to stop the learning process and when to conclude that a target is unattainable. By exploring the variety of decisions that the model is able to make, we further reduce the amount of human involvement in the process.

The high throughput of the system also enables the quantitative examination of the fluctuations in $\mathscr{R}$'s given the same growth parameters (T, t). This fluctuation is attributed to the hidden parameters in the growth process. We introduce the concept of the calibration layer that allows the model to partially capture the effect of these hidden parameters. With the information on the otherwise unknown hidden parameters, significant improvement in model prediction performance is achieved. The fluctuations in the growth process have been found in previous studies utilizing ML for thin film deposition [5, 18], though its effect on the material property and model performance is often neglected. Shrivastava et al.[8] have made an effort to deliberately avoid growing at growth parameters that are particularly sensitive to such fluctuations. Our work represents an important step toward quantifying these hidden parameters by incorporating the calibration layer measurement $\mathscr{R}_c$ into the model's parameter space.

One can extend the calibration layer approach by incorporating additional checkpoints along the deposition trajectories and optimize in a higher-dimensional parameter space. This could eventually lead to a quasi-continuous adaptive control algorithm where the system adjusts in real time based on feedback. However, implementing such adaptive control using reinforcement learning in complex experimental settings poses challenges, as it typically requires substantial amount of training trials [35]. Our method balances elements of adaptive control with efficient model training, accounting for both the limited data availability and the inherent variability in the material growth process. With more advanced reinforcement learning algorithms [36] and higher sample throughput, one can envision that the quasi-continuous adaptive control may be implemented in thin film deposition in the future.

Acknowledgments

This work is supported by the National Science Foundation (NSF CNS-2019131) and University of Chicago Big Idea Generator Seed Grant. This work made use of the shared facilities at the University of Chicago Materials Research Science and Engineering Center, supported by National Science Foundation under award number DMR-2011854. This work made use of the Pritzker Nanofabrication Facility at the Pritzker School of Molecular Engineering at the University of Chicago, which receives support from Soft and Hybrid Nanotechnology Experimental (SHyNE) Resource (NSF ECCS-2025633), a node of the National Science Foundation's National Nanotechnology Coordinated Infrastructure.

REFERENCES

[1] M. I. Jordan and T. M. Mitchell, Machine learning: Trends, perspectives, and prospects, Science **349**, 255 (2015).

[2] E. Garcia, M. A. Jimenez, P. G. De Santos, and M. Armada, The evolution of robotics research, IEEE Robotics I& Automation Magazine **14**, 90 (2007).

[3] D. Morgan and R. Jacobs, Opportunities and challenges for machine learning in materials science, Annual Review of Materials Research **50**, 71 (2020).

[4] Y. K. Wakabayashi, T. Otsuka, Y. Krockenberger, H. Sawada, Y. Taniyasu, and H. Yamamoto, Stoichiometric growth of srtio3 films via bayesian optimization with adaptive prior mean, APL Machine Learning **1** (2023).

[5] D. M. Fébba, K. R. Talley, K. Johnson, S. Schaefer, S. R. Bauers, J. S. Mangum, and A. Zakutayev, Autonomous sputter synthesis of thin film nitrides with composition controlled by bayesian optimization of optical plasma emission, APL Materials **11**, 7 (2023).

[6] N. S. Johnson, A. A. Mishra, D. J. Kirsch, and A. Mehta, Active learning for rapid targeted synthesis of compositionally complex alloys, Materials **17**, 4038 (2024).

[7] T. Ishiyama, K. Nozawa, T. Nishida, T. Suemasu, and K. Toko, Bayesian optimization-driven enhancement of the thermoelectric properties of polycrystalline iii-v semiconductor thin films, NPG Asia Materials **16**, 17 (2024).

[8] A. Shrivastava, M. Kalaswad, J. O. Custer, D. P. Adams, and H. N. Najm, Bayesian optimization for stable properties amid processing fluctuations in sputter deposition, Journal of Vacuum Science & Technology A **42** (2024).

[9] Y. K. Wakabayashi, T. Otsuka, Y. Krockenberger, H. Sawada, Y. Taniyasu, and H. Yamamoto, Machine-learning-assisted thin-film growth: Bayesian optimization in molecular beam epitaxy of srruo3 thin films, APL Materials **7**, 10 (2019).

[10] A. S. Messecar, S. M. Durbin, and R. A. Makin, Quantum and classical machine learning investigation of synthesis–structure relationships in epitaxially grown wide band gap semiconductors, MRS Communications **14**, 660 (2024).

[11] C. Shen, W. Zhan, K. Xin, M. Li, Z. Sun, H. Cong, C. Xu, J. Tang, Z. Wu, B. Xu, and Z. Wei, Machine-learning-assisted and real-time-feedback-controlled growth of inas/gaas quantum dots, Nature Communications **15**, 2724 (2024).

[12] H. J. Kim, M. Chong, T. G. Rhee, Y. G. Khim, M. H. Jung, Y. M. Kim, H. Y. Jeong, B. K. Choi, and Y. J. Chang, Machine-learning-assisted analysis of transition metal dichalcogenide thin-film growth, Nano Convergence **10**, 10 (2023).

[13] S. R. Provence, S. Thapa, R. Paudel, T. K. Truttmann, A. Prakash, B. Jalan, and R. B. Comes, Machine learning analysis of perovskite oxides grown by molecular beam epitaxy, Physical Review Materials **4**, 083807 (2020).

[14] D. Guevarra, L. Zhou, M. H. Richter, A. Shinde, D. Chen, C. P. Gomes, and J. M. Gregoire, Materials structure–property factorization for identification of synergistic phase interactions in complex solar fuels photoanodes, npj Computational Materials **8**, 57 (2022).

[15] Z. Ni and H. Matsui, Phase control of heterogeneous $Hf_XZr_{(1-X)}O_2$ thin films by machine learning, Japanese Journal of Applied Physics **61**, SH1009 (2022).

[16] H. Liang, V. Stanev, A. G. Kusne, Y. Tsukahara, K. Ito, R. Takahashi, M. Lippmaa, and I. Takeuchi, Application of machine learning to reflection high-energy electron diffraction images for automated structural phase mapping, Physical Review Materials **6**, 063805 (2022).

[17] S. Ament, M. Amsler, D. R. Sutherland, M. C. Chang, D. Guevarra, A. B. Connolly, J. M. Gregoire, M. O. Thompson, C. P. Gomes, and R. B. Van Dover, Autonomous materials synthesis via hierarchical active learning of nonequilibrium phase diagrams, Science Advances **7**, eabg4930 (2021).

[18] I. Ohkubo, Z. Hou, J. N. Lee, T. Aizawa, M. Lippmaa, T. Chikyow, and T. Mori, Realization of closed-loop optimization of epitaxial titanium nitride thin-film growth via machine learning, Materials Today Physics **16**, 100296 (2021).

[19] Y. K. Wakabayashi, T. Otsuka, Y. Krockenberger, H. Sawada, Y. Taniyasu, and H. Yamamoto, Bayesian optimization with experimental failure for high-throughput materials growth, npj Computational Materials **8**, 180 (2022).

[20] D. Packwood, *Bayesian Optimization for Materials Science* (Springer Singapore, Singapore, 2017).

[21] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, Taking the human out of the loop: A review of bayesian optimization, Proceedings of the IEEE **104**, 148 (2015).

[22] A. G. Kusne, H. Yu, C. Wu, H. Zhang, J. Hattrick-Simpers, B. DeCost, S. Sarker, C. Oses, C. Toher, S. Curtarolo, and A. V. Davydov, On-the-fly closed-loop materials discovery via bayesian active learning, Nature Communications **11**, 5966 (2020).

[23] R. Shimizu, S. Kobayashi, Y. Watanabe, Y. Ando, and T. Hitosugi, Autonomous materials synthesis by machine learning and robotics, APL Materials **8**, 11 (2020).

[24] P. Zhao, W. Su, R. Wang, X. Xu, and F. Zhang, Properties of thin silver films with different thickness, Physica E: Low-dimensional Systems and Nanostructures **41**, 387 (2009).

[25] H. Savaloni and M. Firouzi-Arani, Dependence of the optical properties of uhv deposited silver thin films on the deposition parameters and their relation to the nanostructure of the films, Philosophical Magazine **88**, 711 (2008).

[26] H. Reddy, U. Guler, K. Chaudhuri, A. Dutta, A. V. Kildishev, V. M. Shalaev, and A. Boltasseva, Temperature-dependent optical properties of single crystalline and polycrystalline silver thin films, ACS Photonics **4**, 1083 (2017).

[27] J. Choi, F. Cheng, J. W. Cleary, L. Sun, C. K. Dass, J. R. Hendrickson, and X. Li, Optical dielectric constants of single crystalline silver films in the long wavelength range, Optical Materials Express **10**, 693 (2020).

[28] B. D. Faeth, S.-L. Yang, J. K. Kawasaki, J. N. Nelson, P. Mishra, C. T. Parzyck, C. Li, D. G. Schlom, and K. M. Shen, Incoherent cooper pairing and pseudogap behavior in single-layer FeSe/$SrTio_3$, Phys. Rev. X **11**, 021054 (2021).

[29] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, 2006).

[30] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson, Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration (2021), arXiv:1809.11165 [cs.LG].

[31] S. Greenhill, S. Rana, S. Gupta, P. Vellanki, and S. Venkatesh, Bayesian optimization for adaptive experimental design: A review, IEEE Access **8**, 13937 (2020).

[32] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization (2017), arXiv:1412.6980 [cs.LG].

[33] M. M. Noack, K. G. Yager, M. Fukuto, G. S. Doerk, R. Li, and J. A. Sethian, A kriging-based approach to autonomous experimentation with applications to x-ray scattering, Scientific Reports **9**, 11809 (2019).

[34] M. M. Noack, G. S. Doerk, R. Li, J. K. Streit, R. A. Vaia, K. G. Yager, and M. Fukuto, Autonomous materials discovery driven by gaussian process regression with inhomogeneous measurement noise and anisotropic kernels, Scientific Reports **10**, 17663 (2020).

[35] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, and S. Petersen, Human-level control through deep reinforcement learning, Nature **518**, 529 (2015).

[36] S. Kamthe and M. Deisenroth, Data-efficient reinforcement learning with probabilistic model predictive control, in *International Conference on Artificial Intelligence and Statistics* (PMLR, 2018) pp. 1701–1710.