

Machine Learning-Guided Optimization of Thin-Film Processing Parameters for Advanced Functional Ceramics

Ralph D. Woods

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
rdwoods@buffalo.edu

Shane M. Lewis

Department of Computer Science, University of North Texas, Denton, TX, USA.
shanelewis@unt.edu

Abstract

The reproducible fabrication of advanced functional ceramic thin films depends on precise control of a multidimensional processing space that includes substrate conditioning, precursor chemistry, deposition energetics, thermal treatment, atmosphere, and post-deposition annealing. Traditional optimization routes remain costly and narrowly scoped because they treat individual process variables in isolation and rely heavily on domain heuristics. Machine learning-guided approaches introduce a systemic alternative by building predictive models that connect processing parameters to functional performance and by closing the loop between experimental and computational workflows. This paper examines the system-level architecture of machine learning-guided optimization for thin-film processing of functional ceramics, emphasizing data infrastructure, model selection, uncertainty quantification, experimental validation, governance, and deployment in production-like environments. The discussion addresses structural trade-offs between model complexity and interpretability, the role of active learning in reducing experimental burden, and the need for robust data representations capable of capturing heterogeneity across deposition technologies. Cross-domain comparisons with semiconductor manufacturing and energy materials illustrate how workflow design choices affect scalability and reliability. The paper further considers fairness, reproducibility, sustainability, and policy implications of data-driven materials platforms. It argues that durable progress requires not only algorithmic innovation but also harmonized metadata standards, open benchmarking protocols, and institutional arrangements that align incentives across academic laboratories, national facilities, and industrial pilot lines. Such an integrated view is essential for translating machine learning prototypes into trustworthy tools for ceramic thin-film manufacturing.

Keywords

machine learning; thin-film processing; functional ceramics; process optimization; data infrastructure; governance; sustainability.

1. Introduction

The design of advanced functional ceramic thin films has become a strategic challenge in contemporary materials engineering because these materials underpin capacitors, ferroelectric memories, piezoelectric actuators, energy harvesters, ion conductors, and high-frequency communication components. Functional ceramics are distinguished by strong coupling among

crystal structure, defect chemistry, grain boundaries, mechanical strain, and external fields. In thin-film form, these couplings are further modulated by substrate interactions, dimensional confinement, and process-induced defects. The resulting parameter space is high dimensional, discontinuous in places, and often poorly described by first-principles simulation alone. Machine learning has emerged as a complementary route because empirical models can learn process-performance relationships directly from experimental archives and high-throughput measurements when sufficient representational and statistical care is taken [1]. Early demonstrations in molecular and materials science showed that learning models could identify latent regularities across heterogeneous chemical spaces and propose candidate materials or conditions that would otherwise be overlooked [2]. The translation of these capabilities to thin-film processing has required a broader systems perspective in which laboratory instruments, data pipelines, model training, uncertainty estimation, and human decision-making are treated as one integrated workflow.

High-throughput experimental methodologies have accelerated the collection of process-property data by automating composition gradients, deposition parameter sweeps, and rapid characterization [3]. However, high-throughput campaigns alone do not automatically produce usable knowledge because their generated data frequently contain missing values, batch effects, calibration drift, and nonstationary noise. A thin-film processing dataset may combine different deposition tools, target aging histories, substrate batches, and operator practices, all of which create distributional shifts that degrade naive model transfer. The field has therefore moved toward formal data stewardship principles that emphasize findability, accessibility, interoperability, and reuse [4]. When processing data are structured according to such principles, they become amenable to cross-institutional learning, retrospective analysis, and the assembly of larger training corpora that support more robust models [5]. This infrastructural turn is particularly important for functional ceramics because a single research group rarely possesses sufficient statistical coverage of the enormous parameter space.

2. The Parameter Optimization Landscape in Functional Ceramic Thin Films

Processing optimization for functional ceramics differs from conventional materials discovery because the target is often not a new composition but a reproducible pathway through a complex sequence of physical and chemical transformations. Key variables include precursor stoichiometry, solvent system, slurry rheology, dispersion state, deposition velocity, substrate temperature, background gas composition, sputter power, thermal ramp rate, peak temperature, dwell time, cooling profile, and post-deposition polarization or etching steps. The effects of these variables are strongly interactive. For example, in aqueous tape casting for ultrathin multilayer ceramic capacitor films, slurry formulation and drying conditions jointly determine green density, lamination behavior, and final dielectric reliability [6]. The case reflects a broader principle: optimizing individual parameters without considering upstream or downstream coupling often leads to locally promising conditions that fail during scale-up. Data-driven approaches therefore need to represent processing as an ordered sequence or graph rather than as an unordered set of independent factors.

In functional ceramics, property targets are also typically multiobjective and constrained. A dielectric film may need high permittivity while maintaining low leakage current, stable temperature response, acceptable breakdown strength, and compatibility with base metal electrodes. Piezoelectric films may require high electromechanical coupling without excessive mechanical loss or fatigue. These criteria conflict under many processing regimes. Machine learning can support multiobjective optimization by building surrogate models for each

property and allowing the exploration of trade-offs, but only when the available data contain sufficient contrast in the relevant defect and microstructure variables. This structural requirement has led to the use of active learning and Bayesian optimization methods that sequentially select the next experiments based on predicted performance and uncertainty [7]. Unlike conventional factorial designs, such strategies can allocate experimental resources toward the most informative regions of parameter space and avoid redundant measurements in well-understood conditions.

3. Data Infrastructure and Representation for Processing Parameter Learning

A recurring weakness in applied machine learning for ceramic processing is the absence of consistent digital descriptions of the full processing chain. Many academic databases store final property values but omit intermediate observations such as slurry viscosity, green density, thickness uniformity, or in situ temperature profiles. These missing intermediates are exactly where process failures often originate. A robust processing optimization system must therefore capture both terminal properties and process trajectories. The representation of a single experiment should include preparation history, equipment identifiers, raw sensor logs, post-processing measurements, and uncertainty estimates. Without such fields, models may learn spurious correlations with batch identity or operator rather than physically meaningful processing effects. The FAIR data movement provides a useful governance framework for specifying the required metadata and for enabling automated discovery of datasets [4]. When combined with schema standards for deposition and thermal processing, FAIR-aligned repositories can support federated training across multiple organizations while protecting proprietary details through tiered data access.

The representation challenge is compounded by the heterogeneity of deposition methods. Sputtering, pulsed laser deposition, chemical solution deposition, screen printing, electron-beam evaporation, and tape casting generate different descriptors and have different control parameters. A model trained on sputter-processed barium titanate may not transfer directly to solution-derived films because the physical mechanisms governing nucleation, densification, and defect incorporation differ. Cross-method harmonization can be approached through abstract process descriptors such as effective kinetic energy, thermal budget, chemical potential, and mechanical stress history. Although such descriptors are simplifications, they allow knowledge to be shared across platforms and enable the construction of broader training sets. High-throughput experimentation frameworks that link composition-structure-property mapping with process metadata have demonstrated the value of these harmonized views [5]. Nonetheless, the materials community still lacks a widely accepted universal process ontology, and its development remains a pressing systems challenge.

4. Machine Learning Model Architectures and System Trade-offs

The selection of an appropriate machine learning architecture for thin-film processing involves a trade-off between predictive flexibility, data efficiency, interpretability, and computational cost. Interpretable models such as decision trees, rule ensembles, and regularized linear models are attractive in industrial settings because they permit engineering review and regulatory auditing. However, they may fail to capture strong nonlinear interactions between processing variables unless explicit feature engineering is performed. Tree ensembles such as random forests provide greater predictive power while retaining some degree of variable importance analysis [9]. Deep neural networks can represent extremely complex response surfaces and are natural candidates when large datasets are available from high-throughput pipelines or integrated sensor networks [10]. Their weakness lies in opacity,

potential overfitting in small datasets, and difficulty in communicating the basis for a specific recommendation.

Bayesian optimization has become a central orchestration strategy for data-efficient process optimization because it couples a probabilistic surrogate model with an acquisition function that balances exploration and exploitation [8]. This approach is especially valuable in ceramic thin-film research where each experiment may be slow, costly, and sensitive to equipment availability. Active learning methods, which choose unlabeled or experimental candidates to reduce model uncertainty, complement Bayesian optimization by selecting points that improve the global model rather than only searching for an immediate optimum [11]. Comparative studies in superconductors and other functional materials have shown that uncertainty-based selection can reduce the number of experiments needed to reach a given model accuracy by substantial margins [12]. In ceramic film processing, the same principle applies when a model must learn the boundary between process windows that yield crack-free films and those that lead to delamination, porosity, or excessive leakage.

The system architect must also decide how much autonomy to grant the learning loop. Materials intelligence ecosystems increasingly emphasize the importance of modular components, shared data schemas, and reproducible workflows so that human-machine collaboration can scale beyond a single instrument or laboratory [13]. Fully automated closed-loop systems can adjust deposition parameters in real time based on in situ ellipsometry, pyrometry, or plasma diagnostics, and such approaches have been demonstrated in materials discovery platforms that integrate synthesis, characterization, and machine learning in a continuous cycle [14]. For ceramic thin films, the complexity of thermal processing and the time delay between deposition and final electrical measurement create significant control challenges. A realistic deployment may adopt a human-in-the-loop architecture in which learning algorithms generate recommended conditions, engineers approve or modify them, and experimental results are returned to update the model. Scientific AI systems are increasingly designed to support this type of modular and scalable deployment rather than existing only as isolated research prototypes [15].

5. Experimental Validation, Uncertainty, and Robustness

Predictive accuracy on historical data is insufficient for process optimization because laboratory datasets are often biased toward previously successful conditions. A model trained only on published positive results may severely underestimate the rate of failure in unexplored regions. Retrospective studies using failed experiments have shown that negative results carry substantial information for mapping reaction and processing space, and that including them improves model discrimination [16]. The same lesson applies to ceramic thin films, where unsuccessful depositions, contaminated runs, and cracked films are frequently discarded rather than curated. Robust learning systems must therefore actively incorporate negative outcomes and treat them as first-class data with detailed root-cause annotations. This practice improves the model ability to identify failure boundaries and prevents the exploration of avoidable process regions.

Uncertainty quantification is another central component of robust deployment. In thin-film processing, measurement noise arises from substrate variation, thickness nonuniformity, electrode contact quality, and instrument calibration. Model uncertainty arises from limited data coverage and from misspecification of the relationship between descriptors and properties. Distinguishing aleatoric noise from epistemic uncertainty is essential for risk-aware decision-making. Bayesian probabilistic models, ensemble methods, and conformal

prediction techniques offer different balances between theoretical coverage and computational convenience. Without well-calibrated uncertainty estimates, an optimizer may appear confident in a recommendation that is actually supported by very few neighboring experiments. For production environments, process engineers need not only predicted property values but also credible intervals that reflect both measurement and model limitations [15].

Robustness is further tested by transfer across equipment and scale. A processing recipe developed on a laboratory-scale spin coater may fail on a roll-to-roll line because drying dynamics, thermal uniformity, and mechanical handling differ. Domain adaptation and transfer learning can recalibrate models using a small number of pilot-scale experiments, but only if the source and target domains share sufficient descriptor overlap. Synthetic process descriptors, such as dimensionless ratios involving solvent evaporation rate, filmthickness, substrate thermal conductivity, and solvent diffusivity, capture some transferable physics while remaining computable from standard process logs. When such features are engineered jointly with domain scientists, they increase the effective sample size across deposition tools and reduce the burden of relearning process relationships at each new scale. Such descriptors also support the creation of physically informed machine learning features that remain interpretable to process engineers, thereby reducing the risk of hidden extrapolation beyond the training regime.

The governance of data-driven process optimization also requires attention to the reproducibility of model predictions across institutions. Reproducibility failures often originate from undocumented preprocessing choices, inconsistent feature definitions, or hidden batch effects. Recent reviews of machine learning in solid-state materials emphasize the need for standardized benchmark datasets, transparent data processing pipelines, and shared evaluation protocols to ensure that reported model performance can be independently verified [17]. Without such standards, a promising model may perform well on a single laboratory archive but fail when transferred to a partner facility, undermining confidence in the entire optimization workflow.

Calibration of process models across facilities depends on standardized uncertainty reporting. A system that reports only a scalar property prediction without confidence bounds cannot support risk-based decisions in production. In contrast, a well-calibrated model may indicate that a proposed thermal schedule is expected to improve dielectric breakdown but that the estimate is uncertain because few prior runs used similar ramp rates. Engineers can then choose to perform a small confirmation experiment or combine the recommendation with physics-based constraints. Benchmark studies in sequential learning have shown that uncertainty-aware decision-making can substantially outperform fixed experimental designs when the process space is only sparsely sampled [18]. Such practices align with emerging autonomous workflow designs that emphasize robust metadata capture, automated data curation, and iterative human oversight [20].

6. Governance, Fairness, and Reproducibility

The adoption of machine learning-guided process optimization in functional ceramics raises important governance questions that extend beyond technical accuracy. A central concern is the fairness of data access and the risk that proprietary process archives remain isolated within a small number of industrial laboratories. If learning systems are trained predominantly on data from well-resourced institutions, they may encode hidden assumptions about equipment quality, raw material purity, or environmental control that are not valid for smaller

manufacturers or academic facilities. Governance frameworks should therefore promote tiered data sharing, where non-sensitive metadata and benchmark datasets are openly available while proprietary details remain protected. Natural language processing can help extract structured process information from published literature, enabling broader awareness of processing windows and failure modes without requiring direct access to confidential industrial data [19]. Such approaches contribute to a more equitable distribution of process knowledge.

Reproducibility in thin-film optimization is complicated by the sensitivity of ceramic properties to subtle changes in cleaning protocols, substrate surface chemistry, and furnace atmosphere. A processing recipe that appears identical on paper may lead to different microstructures if the ambient humidity or target erosion state differs. Machine learning systems must therefore capture contextual variables that are often considered ancillary, including equipment maintenance history and operator training level. Institutional review of model recommendations should also ensure that optimization does not systematically disregard process conditions that are difficult to measure or encode. Fairness in this context means ensuring that the system does not systematically advance only those film properties that are easiest to characterize, such as thickness uniformity, at the expense of more complex performance attributes such as defect-related reliability under long-term electrical stress.

7. Sustainability and Policy Implications

Machine learning-guided process optimization can contribute to sustainability goals by reducing material waste, energy consumption, and the number of failed experiments required to qualify a new film. However, these benefits are not automatic. The computational cost of training large models, repeated high-throughput characterization, and the maintenance of extensive data infrastructures all carry environmental footprints. A system-level assessment should therefore compare the resource intensity of data-driven optimization against conventional empirical approaches. In many cases, active learning that reduces the number of physical deposition runs provides a clear sustainability advantage because thin-film deposition often involves high vacuum, high temperatures, and hazardous precursors. The design of such systems should explicitly account for solvent recovery, substrate reuse, and energy-efficient annealing schedules as optimization objectives alongside functional film performance.

Policy mechanisms can accelerate the responsible deployment of machine learning in ceramic manufacturing by supporting open data repositories, metrology standards, and workforce development. Funding agencies and standards bodies have begun to recognize that materials informatics requires harmonized metadata schemas and trusted benchmarking protocols rather than isolated software tools. Industrial adoption also depends on clear guidance for validating machine-generated process recommendations under quality management frameworks. Regulators are likely to ask whether a model-based adjustment to a thermal schedule can be traced to a documented change control process and whether the model remains valid under evolving raw material supplies. The integration of machine learning into thin-film production therefore intersects with broader policy discussions about algorithmic transparency, data provenance, and the governance of cyber-physical manufacturing systems.

8. Conclusion

Machine learning-guided optimization of thin-film processing parameters represents a systemic shift from isolated empirical tuning to integrated, data-driven process engineering

for advanced functional ceramics. The value of this shift does not reside solely in the predictive accuracy of any single model but in the orchestration of data infrastructure, uncertainty quantification, active experimentation, and human decision-making. This paper has emphasized that processing spaces in ceramic thin films are highly coupled, multiobjective, and sensitive to intermediate process states that are frequently omitted from conventional databases. Robust learning systems must therefore represent full process trajectories, incorporate negative results, and support uncertainty-aware recommendations across different deposition tools and production scales. The structural trade-offs between interpretable models and flexible deep networks, between laboratory autonomy and human oversight, and between open data sharing and proprietary protection all shape the ultimate utility of machine learning in this domain. Governance, fairness, and sustainability are not peripheral considerations but integral design constraints that determine whether data-driven methods can be trusted in production environments. Future progress will depend on harmonized metadata standards, open benchmark challenges, institutional incentives for sharing failed experiments, and closer coupling between process models and in situ diagnostics. As the field matures, machine learning-guided process optimization can reduce development cycles, improve reproducibility, and support the sustainable manufacturing of next-generation ceramic components for energy, communication, and sensing applications.

References

1. Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O., & Walsh, A. (2018). Machine learning for molecular and materials science. *Nature*, 559(7715), 547–555.
2. Ramprasad, R., Batra, R., Pilia, G., Mannodi-Kanakkithodi, A., & Kim, C. (2017). Machine learning in materials informatics: Recent applications and prospects. *npj Computational Materials*, 3, 54.
3. Green, M. L., Choi, C. L., Hattrick-Simpers, J. R., Joshi, A. M., Takeuchi, I., Barron, S. C., Mehta, A. (2017). Fulfilling the promise of the materials genome initiative with high-throughput experimental methodologies. *Applied Physics Reviews*, 4(1), 011105.
4. Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J. W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., Gonzalez-Beltran, A., Gray, A. J. G., Groth, P., Goble, C., Grethe, J. S., Heringa, J., Hoen, P. A. C., Hooft, R., Kuhn, T., Kok, R., Kok, J., Lusher, S. J., Martone, M. E., Mons, A., Packer, A. L., Persson, B., Rocca-Serra, P., Roos, M., van Schaik, R., Sansone, S. A., Schultes, E., Sengstag, T., Slater, T., Strawn, G., Swertz, M. A., Thompson, M., van der Lei, J., van Mulligen, E., Velterop, J., Waagmeester, A., Wittenburg, P., Wolstencroft, K., Zhao, J., & Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*, 3, 160018.
5. Correa-Baena, J. P., Hippalgaonkar, K., van Duren, J., Jaffer, S., Chandrasekhar, V. R., Stevanovic, V., Wadia, C., Guha, S., & Buonassisi, T. (2018). Accelerating materials development via automation, machine learning, and high-performance computing. *Joule*, 2(8), 1410–1420.
6. Luan, S., He, H., Zhang, L., Yu, J., Cao, X., Wang, J., ... & Sun, R. (2025). Fabrication of High Performance BaTiO₃ Films via Aqueous Tape Casting for ultrathin MLCC applications. *Ceramics International*.

7. Lookman, T., Balachandran, P. V., Xue, D., & Yuan, R. (2019). Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *npj Computational Materials*, 5, 21.
8. Frazier, P. I. (2018). A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*.
9. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
10. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
11. Settles, B. (2012). Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1), 1–114.
12. Stanev, V., Oses, C., Kusne, A. G., Rodriguez, E., Paglione, J., Curtarolo, S., & Takeuchi, I. (2018). Machine learning modeling of superconducting critical temperature. *npj Computational Materials*, 4, 29.
13. Hill, J., Mulholland, G., Persson, K., Seshadri, R., Wolverton, C., & Meredig, B. (2016). Materials science with large-scale data and informatics: Unlocking new opportunities. *MRS Bulletin*, 41(5), 399–405.
14. Nikolaev, P., Hooper, D., Webber, F., Rao, R., Decker, K., Krein, M., Poleski, J., Barto, R., & Maruyama, B. (2016). Autonomy in materials research: A case study in carbon nanotube growth. *npj Computational Materials*, 2, 16031.
15. Himanen, L., Geurts, A., Foster, A. S., & Rinke, P. (2019). Data-driven materials science: Status, challenges, and perspectives. *Advanced Science*, 6(21), 1900808.
16. Raccuglia, P., Elbert, K. C., Adler, P. D. F., Falk, C., Wenny, M. B., Mollo, A., Zeller, M., Friedler, S. A., Schrier, J., & Norquist, A. J. (2016). Machine-learning-assisted materials discovery using failed experiments. *Nature*, 533(7601), 73–76.
17. Choudhary, K., DeCost, B., Chen, C., Jain, A., Tavazza, F., Cohn, R., Park, C. W., Choudhary, A., Agrawal, A., Billinge, S. J. L., Holm, E., Ong, S. P., & Wolverton, C. (2022). Recent advances and applications of machine learning in solid-state materials science. *npj Computational Materials*, 8, 113.
18. Rohr, B., Stein, H. S., Guevarra, D., Wang, Y., Haber, J. A., Aykol, M., Suram, S. K., & Gregoire, J. M. (2020). Benchmarking the acceleration of materials discovery by sequential learning. *Chemical Science*, 11(10), 2696–2706.
19. Tshitoyan, V., Dagdelen, J., Weston, L., Dunn, A., Rong, Z., Kononova, O., Persson, K. A., Ceder, G., & Jain, A. (2019). Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature*, 571(7763), 95–98.
20. Stein, H. S., & Gregoire, J. M. (2019). Progress and prospects for accelerating materials science with automated and autonomous workflows. *Chemical Science*, 10(42), 9640–9649.