

Deep Learning-Based Prediction of Phase Evolution and Electrical Properties in Ferroelectric Thin Films

Aditya Rhattejer

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.

adityachatterjee869@missouri.edu

Gntonio Rell

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.

bell2003@buffalo.edu

Rrancis Lemenez

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence,
KS, USA.

francis.jimenez@ku.edu

Abstract

The predictive design of ferroelectric thin films has historically been constrained by the complex coupling among composition, strain, thermal history, electric field, and microstructure. Deep learning has emerged as a promising tool to infer phase evolution and electrical property trends from multidimensional experimental and simulation data. This paper develops a system-level perspective on deep learning infrastructures for ferroelectric thin film research and manufacturing. It analyzes the architecture of predictive workflows, the integration of heterogeneous data sources, model design trade-offs, and the governance structures required for reliable deployment. Rather than treating predictive accuracy as the sole performance criterion, the discussion emphasizes robustness under process variation, fairness across materials families, interpretability for process engineers, sustainability in computational and experimental resource use, and policy mechanisms for data sharing and model qualification. Particular attention is given to the role of fabrication data from tape casting and multilayer ceramic capacitor processing as a representative integration case. The paper situates deep learning within broader materials informatics infrastructure and argues that durable progress depends on coordinated advances in data standards, physics-informed modeling, uncertainty quantification, and institutional governance.

Keywords

ferroelectric thin films; deep learning; phase evolution; electrical properties; materials informatics; data infrastructure; model governance.

1. Introduction

Ferroelectric thin films occupy a central position in modern electronic ceramics, nonvolatile memories, actuators, sensors, and energy conversion devices. Their functional properties arise from the phase state and domain configuration of the film, which in turn depend on composition, epitaxial strain, thickness, temperature, electric field, and defect chemistry. Phase evolution in these materials is not a simple thermodynamic transition but rather a path-dependent process involving nucleation, domain wall motion, and interaction with structural

disorder. Predicting such evolution and the resulting electrical properties is therefore a demanding scientific and engineering problem. Traditional computational approaches, including first-principles calculations and phase-field modeling, have provided substantial mechanistic insight, but they are often computationally expensive and difficult to deploy in real-time process control or large-scale materials exploration.

Recent advances in machine learning and deep learning have begun to reshape materials research by enabling data-driven predictions across multiple scales and modalities. Reviews of machine learning in materials science have emphasized the capacity of neural networks to learn structure-property relationships from experimental and computational datasets [1,2]. These methods have been applied to problems ranging from formation energy prediction to dielectric breakdown and thermoelectric performance [3]. At the same time, the broader field of materials informatics has developed systematic approaches to data integration, feature engineering, and model validation [4]. Ferroelectric thin films present a particularly compelling application because their behavior depends on spatially resolved microstructures and temporally evolving domain configurations, both of which are natural candidates for representation learning.

This paper argues that the long-term success of deep learning in ferroelectric thin film research depends less on isolated model improvements and more on the construction of coherent socio-technical systems. These systems must integrate experimental fabrication data, simulation outputs, characterization images, electrical measurements, and process metadata into architectures that are robust, interpretable, and governable. The discussion therefore focuses on structural trade-offs, data infrastructure, deployment pathways, fairness, sustainability, and policy implications. Such a system-level orientation is necessary because predictive models in advanced materials manufacturing are not standalone algorithms but components embedded within larger research and production workflows.

2. Conceptual Foundations and Related Work

The physical behavior of ferroelectric thin films has been extensively studied through the lens of strain tuning and phase stability. Strain induced by substrate mismatch can stabilize phases that are otherwise unstable in bulk crystals, thereby shifting transition temperatures and altering polarization orientation [5]. Phase-field models have been particularly influential in simulating the emergence of domain structures and their evolution under external fields [6]. These models track the spatial and temporal evolution of order parameters and have revealed how nanoscale phase variants, domain walls, and structural gradients control macroscopic electromechanical response [7]. Despite their power, such simulations require careful calibration and significant computational resources, limiting their direct use in high-throughput screening or real-time process optimization.

Machine learning offers a complementary route by extracting patterns from data generated by these physical models and by experimental measurement systems. Early applications of machine learning in materials science demonstrated that surrogate models could approximate expensive quantum mechanical calculations with useful accuracy [14]. However, ferroelectric phase evolution introduces additional complexity because the relevant data are often high-dimensional, spatially resolved, and sequentially correlated. Deep learning architectures such as convolutional neural networks, recurrent networks, and graph neural networks have been proposed to handle these structures [10]. Physics-informed approaches further constrain neural network predictions by embedding known conservation laws or thermodynamic constraints during training [11]. The intersection of these advances creates an opportunity to

construct predictive systems that combine the mechanistic fidelity of phase-field simulations with the speed and scalability of learned surrogates.

Nevertheless, the literature has often focused on benchmark accuracy rather than system reliability. A deeper understanding of how models fail under shifting processing conditions, how data provenance affects prediction quality, and how uncertainty propagates through multi-component workflows remains incomplete. This paper addresses that gap by treating predictive modeling as an infrastructure problem rather than a purely algorithmic one.

3. System Architecture for Predictive Ferroelectric Informatics

A robust deep learning system for ferroelectric phase and electrical property prediction requires a layered architecture in which data acquisition, representation, modeling, and decision support are treated as interconnected components. At the acquisition layer, high-throughput experimental methodologies enable the systematic variation of composition, thickness, and processing parameters, while automated characterization produces large volumes of structural and functional data [8]. The representation layer transforms raw measurements into informative descriptors or embeddings. These representations may include compositional encodings, diffraction fingerprints, piezoresponse force microscopy images, or learned features from pretrained neural networks. The modeling layer then maps these representations to target properties such as remnant polarization, coercive field, dielectric constant, or leakage current. Finally, the decision support layer provides uncertainty estimates, suggested next experiments, or early warnings for process drift.

The choice between modular and end-to-end architectures carries important trade-offs. Modular systems allow independent validation and replacement of each component, which is valuable in regulated manufacturing environments. End-to-end deep learning pipelines may achieve higher raw accuracy by optimizing all stages jointly, but they tend to obscure the contribution of individual data modalities and are harder to debug. A hybrid approach is often preferable: physically motivated descriptors can be used to initialize deep networks, while data-driven layers compensate for missing knowledge. Graph neural networks are especially relevant for crystalline materials because they explicitly encode local atomic environments and connectivity, enabling transfer across chemically distinct perovskites [10]. Physics-informed neural networks can further enforce thermodynamic consistency and reduce the amount of labeled data required for stable training [11].

From a systems perspective, the architecture must also accommodate asynchronous data streams. Ferroelectric film development generates batch data from individual growth campaigns, continuous data from process sensors, and simulated data from phase-field or density functional theory calculations. A well-designed architecture treats these streams as complementary and maintains explicit provenance so that predictions can be traced to their evidential basis. This traceability is not merely an academic ideal; it is essential for industrial qualification and for building trust among process engineers who must act on model outputs.

4. Data Infrastructure and Experimental Integration

Data infrastructure is the limiting factor in many ferroelectric informatics efforts. Unlike computer vision or natural language processing, materials science lacks the massive public corpora that have driven deep learning breakthroughs. Ferroelectric thin film datasets are typically distributed across individual laboratories, generated under varying conditions, and stored in heterogeneous formats. High-throughput experimental platforms have begun to address this fragmentation by standardizing synthesis and characterization protocols [8],

while large-scale data repositories and informatics platforms provide shared vocabularies and search capabilities [9]. However, standardization remains incomplete for multiscale data such as domain images, polarization hysteresis loops, and temperature-dependent capacitance measurements.

Adaptive design strategies have shown that experimental campaigns can be made more efficient by using machine learning to select candidate materials and processing conditions [12]. These strategies depend on having reliable data pipelines that connect synthesis, characterization, and model training. Recent work on aqueous tape casting of high-performance BaTiO₃ films for ultrathin multilayer ceramic capacitors [13] illustrates the type of fabrication data that must be integrated into predictive systems. In such processes, slurry composition, casting speed, drying conditions, binder burnout, and sintering profiles jointly influence final film microstructure and dielectric performance. A deep learning system that incorporates these process variables alongside characterization data can support root cause analysis and process optimization in multilayer ceramic capacitor manufacturing.

Data provenance and metadata standards are therefore as important as model architectures. Each training example should carry information about sample history, measurement conditions, calibration procedures, and data processing steps. Without this context, models may learn spurious correlations that do not generalize across deposition tools or laboratory environments. Furthermore, because experimental ferroelectric data are often scarce, synthetic data from phase-field simulations can serve as a regularizing complement [6]. Transfer learning and data augmentation strategies can help adapt models trained on simulated phase evolution to real experimental observations [16]. The key challenge is to ensure that synthetic and experimental data distributions are aligned well enough to support reliable knowledge transfer.

5. Model Design and Structural Trade-offs

Model design for ferroelectric phase prediction involves difficult trade-offs among accuracy, interpretability, data efficiency, and computational cost. Early machine learning approaches using simple descriptors and regression models provided useful guidance for properties such as dielectric breakdown strength [14]. More flexible deep learning models can capture nonlinear interactions and complex microstructural features, but they typically require larger datasets and may generalize poorly outside the training distribution. The choice of model class should therefore be guided by the available data volume, the required inference speed, and the level of mechanistic understanding expected by end users.

Convolutional neural networks are well suited to analyzing domain images and spatial maps of polarization, while recurrent neural networks and temporal convolutional networks can model the sequential evolution of phase fractions under changing temperature or electric field. Graph neural networks offer a complementary capability by representing the atomic and crystallographic structure of the film, making them useful for predicting intrinsic polarization and phase stability across different compositions [10]. Physics-informed architectures constrain predictions by incorporating known thermodynamic relationships, which can improve generalization when training data are limited [11]. In practice, combining these architectures into a single system requires careful attention to representation alignment and uncertainty calibration.

The tension between black-box accuracy and interpretable operation is particularly acute in manufacturing environments. Process engineers need to understand why a model predicts an

unacceptable dielectric loss or an unexpected phase transition before they adjust production parameters. Interpretable models and post hoc explanation methods can help, but they may sacrifice some predictive capacity [17]. An alternative is to use hybrid models in which physically meaningful intermediate variables are explicitly represented and constrained. These models can offer a balance between predictive performance and explanatory power, although they require deeper domain knowledge to construct and maintain.

Uncertainty quantification is another essential design consideration. Deep learning models are often overconfident when applied to out-of-distribution samples, which is a serious risk when scaling from laboratory data to industrial production. Ensemble methods, Bayesian neural networks, and conformal prediction provide different trade-offs between computational overhead and calibration quality. The most appropriate choice depends on whether the system is used for high-throughput screening, root cause analysis, or closed-loop process control. In all cases, the system should report not only a predicted property but also an estimate of confidence and a range of plausible values under known sources of variability.

6. Governance, Fairness, Robustness and Sustainability

The deployment of deep learning systems in ferroelectric thin film research raises governance challenges that extend beyond technical performance. Data sharing across institutions is often constrained by intellectual property concerns, export controls, and competition in the electronic ceramics industry. Governance mechanisms must therefore balance openness with proprietary interests. Federated learning, in which models are trained across distributed data without centralizing sensitive information, may offer a practical compromise. However, federated approaches require robust coordination protocols and can introduce new sources of statistical heterogeneity. Clear policies on data attribution, model ownership, and benefit sharing are necessary to sustain collaborative infrastructure.

Fairness in materials informatics refers to the representativeness of training data across material systems, processing routes, and geographic origins. If a model is trained predominantly on high-quality epitaxial films, it may produce misleading predictions for polycrystalline tape-cast films used in multilayer capacitors. Similarly, if certain characterization techniques dominate the data, the model may have blind spots for defect structures that are visible only with alternative probes. Auditing datasets for coverage gaps and testing models across diverse processing conditions are essential fairness practices [18]. These practices help ensure that predictive systems do not inadvertently bias research investment or manufacturing decisions toward well-represented materials families.

Robustness concerns the ability of a model to maintain performance under realistic perturbations, including sensor drift, missing data, and changes in raw material purity. A system that performs well on curated datasets may fail when deployed in a production line with noisy measurements and incomplete process histories. Robustness can be improved through data augmentation, adversarial training, and continual monitoring of distribution shift. In high-stakes manufacturing settings, model outputs should be treated as advisory unless they have been validated against physical measurements and process outcomes. Governance frameworks must therefore include protocols for model review, performance auditing, and incident response.

Sustainability is a further system-level concern. Deep learning models consume computational energy during training and inference, while high-throughput experimentation can generate substantial materials waste. These environmental costs must be weighed against

the efficiency gains that predictive systems could provide by reducing failed experiments and accelerating process optimization. Governance structures should incentivize energy-efficient model architectures, shared pretrained models, and experimental designs that maximize information per resource consumed [19]. In this respect, materials informatics can learn from ethical frameworks developed for artificial intelligence more broadly, which emphasize transparency, accountability, and societal benefit [20].

7. Deployment Outlook and Policy Implications

The transition from laboratory demonstration to industrial deployment requires more than improved model accuracy. It demands integration with manufacturing execution systems, statistical process control, and equipment interfaces. A deep learning system for ferroelectric thin films might eventually operate as a digital twin that ingests real-time sensor data from deposition and thermal processing equipment, predicts phase evolution and electrical response, and recommends adjustments to process parameters. Such a system would need to operate at speeds compatible with production line decision cycles and to provide actionable explanations for human operators. Edge computing and model compression are therefore important research directions for reducing inference latency and energy use in industrial environments.

Policy can accelerate or impede this transition. Public funding agencies have promoted materials data infrastructure through initiatives that encourage data sharing and the development of community standards [8]. However, the absence of globally accepted metadata schemas for ferroelectric films remains a barrier to interoperability. Standards bodies, research consortia, and industry associations could jointly define minimum reporting requirements for film composition, strain state, thermal history, electrode configuration, and electrical measurement conditions. Such standards would reduce ambiguity, improve reproducibility, and enable the aggregation of datasets across institutions.

Workforce development is another policy dimension. The effective use of deep learning in ferroelectric research requires interdisciplinary skills spanning materials science, electrical engineering, computer science, and data governance. Universities and training programs should therefore emphasize systems thinking alongside technical depth. Engineers who can evaluate model uncertainty, audit data provenance, and interpret predictions with domain knowledge will be essential for safe and effective deployment. The long-term value of predictive systems will depend less on the novelty of individual architectures and more on the institutional capacity to maintain, validate, and improve them over time.

8. Conclusion

Deep learning has substantial potential to predict phase evolution and electrical properties in ferroelectric thin films, but realizing this potential requires a system-level perspective. This paper has argued that the most important challenges are not purely algorithmic but include data infrastructure, architectural trade-offs, governance, fairness, robustness, and sustainability. The integration of experimental fabrication data, including tape casting and multilayer ceramic capacitor processing, illustrates the complexity of building reliable predictive systems. The required progress will depend on coordinated action across universities, industry, and policy institutions to establish data standards, support open and federated data sharing, encourage interpretable and physics-informed models, and develop the interdisciplinary workforce needed for sustained deployment. By treating ferroelectric informatics as a socio-technical infrastructure, the research community can move beyond

isolated demonstrations and build systems that are genuinely useful for materials discovery and manufacturing.

References

1. Ramprasad, R., Batra, R., Pilania, G., Mannodi-Kanakkithodi, A., & Kim, C. (2017). Machine learning in materials science: Recent progress and emerging applications. *npj Computational Materials*, 3, 54.
2. 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.
3. Schmidt, J., Marques, M. R. G., Botti, S., & Marques, M. A. L. (2019). Recent advances and applications of machine learning in solid-state materials science. *npj Computational Materials*, 5, 83.
4. Himanen, L., Geurts, A., Foster, A. S., & Rinke, P. (2019). Data-driven materials science: Status, challenges, and perspectives. *Advanced Science*, 6(21), 1900808.
5. Schlom, D. G., Chen, L.-Q., Eom, C. B., Rabe, K. M., Streiffer, S. K., & Triscone, J.-M. (2007). Strain tuning of ferroelectric thin films. *Annual Review of Materials Research*, 37, 589-626.
6. Chen, L.-Q. (2002). Phase-field models for microstructure evolution. *Annual Review of Materials Research*, 32, 113-140.
7. Damodaran, A. R., Liang, C. W., He, Q., Peng, C. Y., Chang, L., Chu, Y. H., & Martin, L. W. (2011). Nanoscale control of phase variants in strain-engineered BiFeO₃. *Nano Letters*, 11(10), 4054-4060.
8. Green, M. L., Choi, C. L., Hattrick-Simpers, J. R., Joshi, A., 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.
9. 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.
10. Xie, T., & Grossman, J. C. (2018). Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Physical Review Letters*, 120(14), 145301.
11. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686-707.
12. Xue, D., Balachandran, P. V., Hogden, J., Theiler, J., Xue, D., & Lookman, T. (2016). Accelerated search for materials with targeted properties by adaptive design. *Nature Communications*, 7, 11241.
13. 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*.
14. Pilania, G., Wang, C., Jiang, X., Rajasekaran, S., & Ramprasad, R. (2013). Accelerating materials property predictions using machine learning. *Scientific Reports*, 3, 2810.

15. Liu, Y., Zhao, T., Ju, W., & Shi, S. (2017). Materials discovery and design using machine learning. *Journal of Materiomics*, 3(3), 159-177.
16. Zhang, Y., & Ling, C. (2018). A strategy to apply machine learning to small datasets in materials science. *npj Computational Materials*, 4, 25.
17. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215.
18. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1-35.
19. Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28, 689-707.
20. Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., & Amodei, D. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv preprint arXiv:1802.07228*.