

Intelligent Microstructure–Property Modeling of Perovskite Oxide Dielectrics for Miniaturized Electronic Devices

Raj Hegde

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.

raj73@unr.edu

Yash A. Joshi

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.

yashjoshi17@uab.edu

Zhengling Yin

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.

zhenglingyin@uc.edu

Tobias Rarker

School of Computing, Clemson University, Clemson, SC, USA.

tobiasbarker540@clemson.edu

Abstract

The escalating demand for miniaturized electronic devices has placed severe constraints on dielectric materials, particularly perovskite oxide ceramics used in multilayer capacitors, tunable filters, and embedded passive components. The performance of these materials depends on a complex interplay among grain size, domain structure, defect chemistry, interface quality, and processing history. Traditional microstructure-property modeling approaches often fail to capture these multiscale nonlinear interactions with sufficient speed and fidelity for modern device development. Intelligent modeling frameworks that integrate machine learning, high-throughput experiments, and domain-aware representations offer a promising route toward accelerated discovery and deployment. This paper presents a systems-level examination of intelligent microstructure-property modeling for perovskite oxide dielectrics. It considers data infrastructure, multiscale representation, model architecture, structural trade-offs, deployment constraints, robustness, fairness, sustainability, and governance. The discussion emphasizes that successful implementation requires more than algorithmic sophistication. It demands coordinated infrastructure for data provenance, interoperable metadata, reliable uncertainty quantification, and lifecycle-aware design. The paper also examines how advanced fabrication data for ultrathin film dielectrics can inform model development while cautioning against overreliance on narrow datasets. By treating the modeling pipeline as a socio-technical system, the paper identifies structural tensions between edge-level inference and cloud-scale training, between model complexity and interpretability, and between open innovation and proprietary material data. The concluding perspective outlines future directions including self-driving laboratories, multimodal foundation models,

and adaptive governance frameworks that can support sustainable miniaturization of electronic components.

Keywords

perovskite oxide dielectrics, microstructure-property modeling, machine learning, materials informatics, multilayer ceramic capacitors, data governance, robust deployment, sustainability.

1. Introduction

The miniaturization of electronic devices has become a defining pressure across consumer electronics, automotive systems, medical instrumentation, and communication infrastructure. Passive components, especially multilayer ceramic capacitors, are increasingly expected to deliver higher capacitance density, lower equivalent series resistance, and stable behavior under severe thermal and electrical stress. Perovskite oxide dielectrics such as barium titanate based systems occupy a central position in this landscape because their polarization mechanisms, grain boundary effects, and defect landscapes can be engineered across length scales. However, the very features that make these materials attractive also make their microstructure-property relationships difficult to model. Grain size, domain wall mobility, oxygen vacancy distribution, dopant segregation, and interface roughness interact in ways that are strongly nonlinear and processing sensitive. Deep learning has introduced powerful tools for extracting patterns from complex data spaces, but the deployment of these tools in dielectric materials research requires more than predictive accuracy measured on static benchmarks [1]. It requires a broader architectural perspective that connects material synthesis, characterization, property measurement, device testing, and uncertainty-aware decision making.

Materials informatics has steadily matured from isolated regression tasks to integrated workflows that combine experimental data, computational simulations, and domain knowledge [2]. In multilayer ceramic capacitor research, the need for such workflows is especially acute because the thickness of individual dielectric layers continues to shrink, approaching regimes where grain-scale heterogeneity dominates macroscopic behavior [3]. Historical advances in ferroelectric ceramic processing relied heavily on empirical adjustments and phenomenological understanding [4]. Those approaches produced remarkable technologies, but they are now reaching their practical limits when the number of composition and processing variables expands dramatically. Intelligent modeling systems can navigate these high-dimensional spaces more effectively if they are designed with careful attention to data semantics, physical priors, and lifecycle constraints.

A systems-level perspective is therefore necessary. The challenge is not simply to select a more expressive neural network or to collect a larger dataset. It is to construct a stable infrastructure in which data generation, model training, validation, deployment, and recalibration operate coherently. This includes laboratory information management, metadata standards, ontologies for microstructural features, and mechanisms for tracking provenance across research groups. Machine learning in materials informatics has demonstrated that well-curated data can reveal process-structure-property linkages that are otherwise obscured by variability and sparse sampling [5]. Yet the path from a successful laboratory model to robust industrial deployment remains understudied. The present paper addresses that gap by examining intelligent microstructure-property modeling from a system architecture and governance standpoint, with perovskite oxide dielectrics as a representative and technologically important case.

2. Background and System Context

The emergence of data-intensive materials science has been described as part of a broader fourth paradigm in which scientific discovery is driven by large-scale data analysis rather than solely by theory or experiment [6]. This paradigm shift has been particularly visible in chemical synthesis and materials discovery, where even failed experiments have been converted into predictive resources that guide subsequent exploration [7]. In the domain of dielectric ceramics, however, the translation has been slower because the relevant data are often multimodal, heterogeneous, and tightly coupled to proprietary fabrication conditions. Machine learning methods developed for solid-state materials have shown promise in predicting formation energies, band gaps, and other physical quantities, but extending these methods to microstructure-sensitive properties demands a different level of representational sophistication [8]. Unlike a crystal structure that can be summarized by symmetry operations and atomic positions, a sintered ceramic microstructure contains grains, pores, secondary phases, dislocations, and grain boundary films whose statistical and topological descriptions are not fully standardized.

Perovskite oxide dielectrics illustrate this complexity clearly. Their functional response depends on ferroelectric domain configurations, defect dipoles, space charge zones near electrodes, and residual stresses imposed by thermal expansion mismatch. These characteristics arise from a sequence of powder synthesis, tape casting, lamination, binder burnout, sintering, and electroding steps. Each step can alter the final defect chemistry and microstructural state. From a systems perspective, this means that an intelligent modeling environment must integrate process logs, raw materials certificates, in-line metrology, electron microscopy, X-ray diffraction, impedance spectroscopy, and failure analysis records. The information is temporally distributed and institutionally fragmented. A model trained only on final property measurements may learn spurious correlations because it cannot observe the intermediate states that causally determine device performance. Thus, the background context for intelligent modeling is not purely mathematical; it is deeply infrastructural and organizational.

The socio-technical character of this problem becomes clearer when one considers how material datasets are produced. Experimental laboratories may use different thermal profiles, electrode materials, and measurement protocols. A model that performs well on one group's data may fail when transferred to another because the input distribution is subtly shifted by unreported processing details. This distribution shift is a central concern for robust deployment. Governance mechanisms such as standardized reporting, data versioning, and community-level ontology alignment are therefore as important as model selection. In this environment, the role of intelligent systems is not to replace materials engineers but to surface high-value experimental designs, identify anomalies, and provide calibrated uncertainty estimates that allow practitioners to balance exploration and exploitation under constrained resources.

3. Data Infrastructure and Multiscale Representation

A foundational requirement for intelligent microstructure-property modeling is the construction of data infrastructure that can represent materials information at multiple scales. At the electronic and atomic scale, composition, doping, and defect formation energies are important. At the nanoscale, domain walls and grain boundary structures dominate. At the micrometer scale, grain size distributions, porosity, and connectivity become decisive. At the component scale, layer thickness, electrode roughness, and mechanical integrity interact with

electrical performance. Transformer-style sequence models have enabled significant progress in learning from variable-length and multimodal data streams, and their attention mechanisms can be adapted to process heterogeneous processing logs and time-series sensor data [9]. However, applying such models to microstructure-property prediction requires more than passing images through a neural network. It requires an explicit strategy for aligning representations across scales so that causal influences are not lost.

The data infrastructure must support the capture of micrographs, diffraction patterns, impedance spectra, thermal profiles, and chemical analyses with sufficient contextual metadata to make them reusable. This is a nontrivial systems challenge because many legacy instruments produce outputs that are difficult to parse automatically. Intelligent infrastructure involves automated ingestion, segmentation, and registration of microstructure images, as well as semantic annotation of grain boundaries and phase distributions. Without these capabilities, the data remain siloed in individual instruments or laboratory notebooks. The modeling layer then inherits biases and gaps that cannot be corrected downstream. In this sense, data infrastructure is not a precursor to modeling but a constitutive part of the modeling system itself.

Active learning and Bayesian optimization have become influential tools for guiding adaptive experiments in materials science [10]. Bayesian optimization provides a principled framework for balancing the exploration of uncertain regions with the exploitation of promising compositions and processing conditions [11]. When coupled with microstructure-level descriptors, these methods can direct scarce experimental resources toward regions of the design space that are likely to reveal new physical mechanisms or improve device-level objectives. Active learning strategies that use uncertainty estimates from surrogate models have been applied to targeted design in materials science, allowing researchers to reduce the number of costly experiments required to achieve a given performance target [12]. In dielectric ceramics, such approaches are especially valuable because sintering and electrical testing are capital intensive. However, the benefits of active learning depend on the quality of the underlying descriptors and the fidelity of the uncertainty model. If the representation cannot capture grain boundary chemistry or domain wall contributions, the optimization loop may confidently explore irrelevant regions. Multiscale representation and data infrastructure are therefore tightly linked to the effectiveness of adaptive experimentation.

4. Intelligent Modeling Architectures

The architectural choices for intelligent microstructure-property modeling must reflect the physical and logistical constraints of dielectric materials development. Recent advances in film fabrication for ultrathin multilayer ceramic capacitors have produced high-density processing-property records that can support microstructure-level learning [13]. Such datasets allow models to connect tape casting conditions, binder chemistry, sintering schedules, and final dielectric performance more tightly than was previously possible. At the same time, ferroelectric thin film research has established a rich body of knowledge about thickness-dependent domain structures, interfacial layers, and electrode effects [14]. The applications of modern ferroelectrics extend far beyond simple capacitance, encompassing tunable resonators, energy storage, infrared sensing, and logic-compatible memory elements [15]. These diverse application contexts create conflicting requirements for dielectric response, which the modeling architecture must accommodate.

Graph-based approaches have been particularly promising for representing crystal structures and local chemical environments in a way that is interpretable and physically informed [16].

Crystal graph convolutional networks, for example, encode atomic connectivity and can predict material properties with competitive accuracy while exposing learned features that correlate with known chemistry. Extending graph-based methods to microstructures is more difficult because microstructural graphs may need to capture grain adjacency, triple junctions, and phase boundaries rather than only atomic bonds. Nevertheless, the conceptual foundation is useful: a material system can be represented as a network of interacting entities whose local properties determine global behavior. This perspective aligns with the physics of polycrystalline dielectrics, where grain boundary barriers and domain interactions govern macroscopic permittivity and loss. Transformer-based models can complement graph-based methods by handling long-range correlations in processing histories and multimodal inputs. Hybrid architectures that combine graph encoders for microstructural topology with sequence encoders for time-dependent processing data are therefore a compelling direction.

The choice of architecture also has implications for interpretability and trust. In dielectric materials, domain experts often need to know why a model predicts a certain failure probability or temperature coefficient. A black-box neural network may provide high accuracy but limited actionable insight. Architectures that separate physical baseline terms from learned correction terms, or that incorporate known scaling laws for grain size effects, can preserve interpretability while still benefiting from data-driven flexibility. Such physics-informed modeling reduces the risk of unphysical extrapolation and improves the acceptance of model outputs by process engineers. The architecture is thus not merely a computational arrangement; it is a governance artifact that shapes how humans interact with predictions.

5. Structural Trade-offs and Deployment Constraints

The deployment of intelligent modeling systems in miniaturized electronic device manufacturing involves structural trade-offs that are often overlooked in academic studies. A model trained on a large cloud-scale dataset can achieve high predictive accuracy, but the latency and confidentiality constraints of industrial fabrication may demand edge-level inference. Graph network models that learn physical interactions can support simulation-like reasoning, yet they may be computationally expensive when embedded in real-time process control loops [17]. These trade-offs force system architects to decide whether the model should run close to fabrication tools, where data remain localized and decisions can be made rapidly, or in centralized research computing environments, where more powerful retraining and cross-facility learning are possible. Both options have merit, and the optimal solution may be a federated architecture in which local models are periodically aggregated without transferring raw proprietary data.

Data scarcity is another structural constraint. Although high-throughput experiments can generate thousands of samples, the number of reliable failure records or long-term aging measurements is usually much smaller. In dielectric research, aging and degradation under temperature and voltage stress are critical for device qualification, yet these tests require extended durations and are rarely repeated many times. A model that is optimized for highly sampled compositions may be poorly calibrated for the sparse regions where reliability data matter most. This imbalance creates a tension between predictive accuracy on average and robustness in the tails of the performance distribution. Intelligent systems must therefore include uncertainty quantification, active data acquisition, and causal reasoning to prevent overconfident extrapolation into under-sampled stress conditions.

The infrastructure needed to maintain such models over time also presents trade-offs. Continuous data streams from production lines can improve model performance, but they

require robust data versioning and drift detection. A dielectric manufacturing line may gradually shift because raw powder suppliers change, furnace profiles drift, or environmental humidity varies seasonally. If these changes are not represented in the training distribution, the model may silently degrade. System architects must choose between periodic offline retraining, continuous online adaptation, and hybrid approaches that freeze certain layers while updating others. Each choice affects reproducibility, interpretability, and operational safety. The structural challenge is to build a learning system that remains stable under gradual drift while remaining flexible enough to capture genuine process improvements.

6. Robustness, Fairness, and Generalization

Robustness is a central concern when intelligent models are transferred from academic datasets to industrial deployment. Recent progress in deep learning for materials science has shown that models can achieve impressive performance on curated benchmarks, but their behavior under distribution shift is much less understood [18]. For perovskite oxide dielectrics, distribution shift arises when models trained on laboratory-scale samples are used to predict behavior in mass production, where layer thicknesses are smaller, electrode roughness is different, and thermal budgets are constrained. Robust deployment requires deliberate evaluation protocols that include held-out processing lines, unseen compositions, and simulated measurement noise. Without such protocols, high reported accuracy can mask brittleness. The field can learn from cross-domain advances in representation learning, such as the use of large-scale models that exhibit surprising generalization beyond their original training domain [19]. Although the material systems differ, the underlying lesson is that representational richness and diversity often support robustness more effectively than narrow task-specific fitting.

Fairness in materials modeling may appear less immediate than in social applications, but it is nonetheless relevant. The distribution of experimental data is shaped by institutional priorities, funding patterns, and the availability of advanced instrumentation. Well-resourced laboratories can generate large datasets for high-value dielectric compositions, while less common formulations or environmental conditions remain under-represented. An intelligent system optimized solely for average performance may direct resources away from these neglected regions, reinforcing existing biases in material knowledge. Fairness, in this context, means ensuring that the modeling infrastructure does not systematically exclude technologically important but less heavily sampled regions of the design space. It also means that model benefits should be accessible across different research and manufacturing contexts rather than concentrated among organizations that can afford proprietary data lakes and specialized computing resources.

Generalization across temperatures, frequencies, and electric fields is particularly important for dielectric components. A model that predicts room-temperature capacitance may be inadequate for automotive applications that require stable performance from subzero conditions to high-temperature engine environments. Frequency dependence, nonlinearity under bias, and aging behavior are all critical. If these dimensions are not represented in the training data, the model may appear successful in narrow validation while failing in operational settings. Generalization therefore requires deliberate data curation across operational envelopes, as well as physically motivated descriptors that encode known temperature and field scaling. The governance of data collection and the design of evaluation benchmarks are inseparable from the technical objective of building robust dielectric models.

7. Sustainability and Lifecycle Implications

Sustainability considerations increasingly influence the design and deployment of intelligent materials systems. The production of perovskite oxide dielectrics involves energy-intensive sintering steps, the consumption of high-purity raw materials, and the generation of waste streams from tape casting and electrode deposition. Intelligent modeling can reduce environmental impact by minimizing failed experiments, optimizing sintering schedules, and identifying formulations that achieve equivalent device performance with lower processing temperatures. Emerging materials intelligence ecosystems emphasize the need to integrate machine learning not only with discovery but also with lifecycle considerations such as supply chain resilience, recyclability, and energy consumption [20]. This perspective is particularly relevant for miniaturized electronic devices, where material quantities per component are small but the cumulative industrial volume is enormous.

Lifecycle-aware modeling requires data that extend beyond dielectric constant and loss tangent. It includes embodied energy, carbon intensity, water usage, and the availability of substitute materials. Such data are rarely collected in conventional dielectric research. Intelligent infrastructure can help by linking process parameters to sustainability metrics and enabling multi-objective optimization that balances electrical performance against environmental footprint. For example, a model might identify a sintering profile that achieves acceptable capacitance density while reducing peak furnace temperature and dwell time, thereby lowering energy use. However, these benefits depend on the willingness of manufacturers to share process data and on the existence of reliable lifecycle databases. Governance mechanisms that incentivize data sharing for sustainability purposes could amplify the positive impact of intelligent modeling.

The lifecycle perspective also highlights the importance of durability and reliability. A dielectric component that degrades prematurely creates electronic waste and undermines the environmental benefits of miniaturization. Intelligent models that predict aging under operational stress can guide the design of more durable components, reducing replacement frequency and conserving resources. This connects robustness, sustainability, and governance in a single systems-level objective. The model is valuable not merely because it predicts a property, but because it enables decisions that extend component lifetime and reduce the overall material and energy burden of electronic systems. The intersection of lifecycle assessment and machine learning in dielectric ceramics remains underexplored and represents an important avenue for future interdisciplinary work.

8. Governance, Policy, and Standards

The governance of intelligent microstructure-property modeling extends beyond individual laboratories to encompass standards, intellectual property, reproducibility, and regulatory acceptance. For perovskite oxide dielectrics used in safety-critical applications such as medical implants and automotive control systems, model outputs may influence qualification decisions. This raises questions about how to validate a continuously learning model, how to audit its failure modes, and how to ensure that predictions are traceable to underlying data. Traditional material qualification relies on standardized test methods and physical inspection. Data-driven models introduce a different form of evidence that must be integrated with existing engineering workflows. Without clear standards for reporting model uncertainty, dataset composition, and validation protocols, the adoption of intelligent systems may be slowed by regulatory caution.

Policy mechanisms can support the development of shared data resources while protecting legitimate commercial interests. Materials data are often valuable because they encode years

of process optimization and supplier relationships. A governance framework that mandates open sharing may reduce the incentive to invest in data collection. Conversely, a completely closed environment prevents the community from learning across facilities and slows the development of robust models. Intermediate models such as federated learning, data trusts, and tiered access agreements can balance these tensions. The systems research community has relevant experience in designing such arrangements for software and network infrastructure, and that experience can inform the governance of materials intelligence. The required reference position is not first but rather integrated within a broader architecture of shared and proprietary knowledge [13].

Standards for data provenance and metadata are especially important in dielectric ceramics. A micrograph without magnification, etching conditions, grain boundary phase information, and thermal history has limited reusability. A prediction without a timestamped model version and a declared applicability domain cannot be safely integrated into a manufacturing line. Standardization efforts should therefore focus on the minimum information required to reproduce a microstructure-property prediction. These standards must be developed collaboratively across academic, industrial, and governmental stakeholders to avoid excessive bureaucracy while ensuring technical adequacy. The governance challenge is to create protocols that are lightweight enough for practical adoption yet rigorous enough to support safety-critical deployment.

9. Future Directions

The future of intelligent microstructure-property modeling for perovskite oxide dielectrics will likely be shaped by three convergent trends. The first is the rise of self-driving laboratories that integrate automated synthesis, characterization, and machine learning into closed-loop experimental platforms. Such systems can generate rich, internally consistent datasets that are particularly valuable for understanding process-structure-property linkages. The second trend is the development of multimodal foundation models for materials that can absorb textual knowledge, microstructural images, and numerical measurements into shared latent spaces. These models may lower the cost of knowledge transfer across different dielectric compositions and device architectures. The third trend is the increasing emphasis on sustainability and lifecycle performance, which will require novel objective functions that go beyond immediate electrical properties.

These trends also introduce new risks. Self-driving laboratories may optimize rapidly toward narrow objectives and produce datasets that are overly concentrated in particular regions of processing space. Multimodal foundation models may inherit biases from uneven training corpora and overstate confidence in under-represented material classes. Sustainability-driven optimization may conflict with short-term performance targets and require governance structures that align incentives across supply chains. Addressing these risks will require interdisciplinary collaboration among materials scientists, computer engineers, data stewards, and policy scholars. The systems perspective advanced in this paper can help ensure that technical advances are accompanied by robust infrastructure, fair access, and responsible deployment.

The integration of intelligent modeling with manufacturing execution systems also represents an important frontier. Real-time data from production lines can inform predictive maintenance of furnaces, detection of slurry inhomogeneities, and early warning of electrode delamination. These applications demand low-latency inference and high reliability, placing new constraints on model architecture and hardware. Edge computing devices must be able to execute

compact models while preserving enough contextual information to flag anomalies for human review. In addition, the models themselves must be updated without disrupting production. This calls for modular architectures with clear contracts among data ingestion, feature extraction, inference, and alerting layers. Such modularity is not only good software engineering; it is a necessary condition for safe and auditable deployment in regulated manufacturing environments.

10. Conclusion

Intelligent microstructure-property modeling has the potential to transform the development of perovskite oxide dielectrics for miniaturized electronic devices. The performance of these materials depends on a multiscale network of compositional, microstructural, and processing variables that traditional empirical methods struggle to navigate. Machine learning can capture nonlinear interactions and guide adaptive experiments, but only when embedded within a well-designed system that addresses data infrastructure, representation, model architecture, deployment constraints, robustness, fairness, sustainability, and governance. The structural trade-offs between edge and cloud computing, between model expressiveness and interpretability, and between data sharing and proprietary control are as important as the choice of neural network. The discussion has emphasized that successful deployment requires a socio-technical approach in which data provenance, uncertainty quantification, and standardization support reliable predictions. Advanced fabrication datasets for ultrathin multilayer ceramic capacitors provide valuable opportunities to connect processing history with dielectric behavior [13], yet such datasets must be integrated with broader knowledge of ferroelectric thin films, domain physics, and device-level requirements. As self-driving laboratories and multimodal models mature, the field must ensure that these technologies are governed by principles that promote robustness, fairness, and sustainability. The resulting intelligent modeling systems will not only accelerate materials development but also contribute to a more resilient and environmentally responsible electronic infrastructure.

References

1. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
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. Hong, K., Lee, T. H., Suh, J. M., Yoon, S. H., & Jang, H. W. (2019). Perspectives and challenges in multilayer ceramic capacitors for next generation electronics. *Journal of Materials Chemistry C*, 7(32), 9782–9802.
4. Haertling, G. H. (1999). Ferroelectric ceramics: history and technology. *Journal of the American Ceramic Society*, 82(4), 797–818.
5. Ramprasad, R., Batra, R., Pilania, G., Mannodi-Kanakkithodi, A., & Kim, C. (2017). Machine learning in materials informatics: recent applications and prospects. *npj Computational Materials*, 3(1), 54.
6. Agrawal, A., & Choudhary, A. (2016). Perspective: Materials informatics and big data: Realization of the fourth paradigm of science in materials science. *APL Materials*, 4(5), 053208.

7. Raccuglia, P., Elbert, K. C., Adler, P. D. F., Falk, C., Wenny, M. B., Mollo, A., ... & Norquist, A. J. (2016). Machine-learning-assisted materials discovery using failed experiments. *Nature*, 533(7601), 73–76.
8. 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(1), 83.
9. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
10. Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems*, 25.
11. Frazier, P. I. (2018). A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*.
12. 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(1), 21.
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. Setter, N., Damjanovic, D., Eng, L., Fox, G., Gevorgian, S., Hong, S., ... & Waser, R. (2006). Ferroelectric thin films: Review of materials, properties and applications. *Journal of Applied Physics*, 100(5), 051606.
15. Scott, J. F. (2007). Applications of modern ferroelectrics. *Science*, 315(5814), 954–959.
16. 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.
17. Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., & Battaglia, P. (2018). Graph networks as learnable physics engines for inference and control. *Proceedings of the 35th International Conference on Machine Learning*, 80, 4470–4479.
18. Choudhary, K., DeCost, B., Chen, C., Jain, A., Tavazza, F., Cohn, R., ... & Wolverton, C. (2022). Recent advances and applications of deep learning methods in materials science. *npj Computational Materials*, 8(1), 59.
19. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., ... & Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589.
20. Batra, R., Song, L., & Ramprasad, R. (2021). Emerging materials intelligence ecosystems propelled by machine learning. *Nature Reviews Materials*, 6(8), 655–678.