

Machine Learning-Driven Prediction of Dielectric Properties in Advanced Electronic Ceramics

Eustaquio Cifuentes

Department of Computer Science, University of New Hampshire, Durham, NH, USA.
eustaquio2001@unh.edu

Lajita Borde

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

Donatello Gatto

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.
donatellowork@missouri.edu

Abstract

Advanced electronic ceramics are central to capacitors, sensors, actuators, and high-frequency communication systems, where dielectric permittivity, loss, breakdown strength, and temperature stability determine device performance and reliability. Machine learning has emerged as a powerful approach for predicting these properties, yet the most consequential challenges are not merely algorithmic. They concern data infrastructure, experimental integration, reproducibility, deployment constraints, robustness under distribution shift, fairness across laboratories and material systems, and governance across industrial and public research ecosystems. This paper examines machine learning-driven prediction of dielectric properties in advanced electronic ceramics from a system-level perspective. It first situates ceramic dielectrics within the broader context of ferroelectric physics, multilayer ceramic capacitor manufacturing, lead-free material development, and high-temperature polymer-ceramic composites. It then analyzes model architectures, including descriptor-based regression, graph representations, neural network potentials, and active learning, while emphasizing that predictive accuracy must be balanced against interpretability, computational cost, and experimental validation. The discussion extends to data infrastructure, metadata standards, model documentation, auditing, and concept drift, arguing that dielectric property prediction should be treated as a socio-technical infrastructure rather than a standalone computational task. Sustainability and policy implications are considered, including raw material criticality, sintering energy, open data, and equitable access to high-throughput facilities. Case illustrations from solid electrolytes, polymer dielectrics, and multilayer ceramic capacitors clarify how cross-domain lessons can inform ceramic dielectric design. The paper concludes that robust, fair, and sustainable machine learning for electronic ceramics requires integrated architectures that couple models, instruments, materials databases, human expertise, and governance mechanisms from the outset.

Keywords

machine learning; electronic ceramics; dielectric properties; multilayer ceramic capacitors; materials informatics; robustness; governance; sustainability.

1. Introduction

Advanced electronic ceramics form the material backbone of modern electrical infrastructure. They enable charge storage in multilayer ceramic capacitors, frequency control in resonators and filters, sensing in automotive and industrial systems, and actuation in precision devices. Their performance is governed by interrelated dielectric properties, including permittivity, dielectric loss, breakdown strength, leakage current, temperature coefficient, and fatigue resistance. Predicting these properties is difficult because they emerge from coupled phenomena across electronic, crystallographic, microstructural, and device scales. Machine learning offers a route to navigate this complexity by learning patterns from experimental and simulated data, but its value depends on how well it is embedded within material science workflows and industrial deployment systems. Early materials informatics research demonstrated that data-driven models can accelerate property prediction and guide candidate selection [1]. Subsequent work expanded machine learning across molecular and solid-state materials science, showing that flexible statistical models can complement first-principles simulation and experimental screening [2]. In solid-state materials science, machine learning has been applied to structure-property relationships, synthesis planning, and characterization analysis [3]. These advances motivate the present inquiry, but they also expose a central tension: dielectric ceramics are manufactured systems whose performance depends on processing history, interfaces, dopants, grain boundaries, and electrode integration. A model that predicts a scalar property from a chemical formula may be technically impressive yet insufficient for factory qualification, reliability engineering, or regulatory acceptance.

The system-level view adopted here treats prediction as one component in a longer chain that includes data generation, curation, feature construction, model training, uncertainty estimation, experimental validation, scale-up, quality control, and lifecycle monitoring. This chain is socio-technical because it involves laboratory protocols, vendor relationships, intellectual property, safety constraints, and institutional incentives. Materials intelligence ecosystems have been proposed as integrated environments in which data, simulation, machine learning, and experimentation reinforce one another [18]. Reviews of machine learning for energy materials similarly emphasize that predictive performance alone does not resolve problems of data scarcity, inconsistent reporting, or transferability across chemistries [19]. For electronic ceramics, these concerns are especially acute because dielectric properties are sensitive to minor compositional changes, sintering atmospheres, electrode materials, and measurement conditions. Consequently, machine learning-driven prediction must be evaluated not only by test-set error but by its ability to support decisions under uncertainty, across manufacturing scales, and over time.

A further motivation is the growing pressure to shorten development cycles while reducing environmental and economic costs. Electronic ceramic components are produced in enormous volumes, and even small improvements in dielectric performance can affect energy efficiency, device miniaturization, and material consumption. At the same time, the transition toward lead-free dielectrics, base-metal electrodes, and high-temperature operation introduces new constraints that cannot be captured by simple interpolation from historical data. Governance mechanisms such as algorithmic auditing and model documentation become relevant because predictive systems may influence material selection, supplier qualification, and safety-critical component design [21]. Model cards, for example, provide structured reporting of intended use, performance conditions, and limitations, which can make machine learning artifacts more accountable in engineering contexts [22]. This paper therefore examines machine learning for dielectric property prediction as an interdisciplinary infrastructure problem, not merely as a modeling exercise.

2. Materials and System Context of Advanced Electronic Ceramics

The dielectric behavior of electronic ceramics is rooted in crystal chemistry and polarization dynamics. Barium titanate and related perovskites exhibit ferroelectricity because competing long-range and short-range forces produce spontaneous polarization that can be reoriented by electric fields [12]. This microscopic mechanism gives rise to high permittivity near phase transitions, but it also creates strong temperature and field dependence. Ferroelectric ceramics are used in capacitors, sensors, and actuators because their polarization response can be engineered through composition, doping, grain size, and defect chemistry [13]. In practical devices, however, the relevant property is not a single intrinsic constant. It is a distributed response shaped by grain boundaries, secondary phases, porosity, internal stresses, electrode interfaces, and degradation under bias. Machine learning models that ignore microstructure or processing history may capture average trends but fail in the tails of the distribution, where reliability problems often reside.

Multilayer ceramic capacitors illustrate the coupled nature of material and device design. They consist of many thin dielectric layers co-fired with metal electrodes, which demands compatibility between dielectric shrinkage, electrode oxidation, and sintering atmosphere [14]. Base-metal electrode systems required major advances in reduction-resistant dielectrics and controlled oxygen partial pressure during processing. The resulting property targets include high permittivity, low loss, high breakdown strength, stable capacitance over temperature, and resistance to insulation resistance degradation. A brief introduction to ceramic capacitors highlights how dielectric formulation, layer thickness, and electrode geometry jointly determine volumetric efficiency and reliability [15]. For machine learning, these dependencies mean that labels are not purely material properties; they are process-conditioned and device-conditioned outcomes. A dataset compiled from literature may mix measurement frequencies, bias voltages, sintering schedules, and electrode configurations, creating hidden variables that undermine naïve regression.

Lead-free dielectric development adds another layer of system complexity. Environmental and regulatory pressures have encouraged the replacement of lead-based perovskites in energy storage and actuator applications, but candidate systems often trade high polarization for lower breakdown strength or worse temperature stability [16]. High-temperature polymer nanocomposites offer a parallel case in which organic and inorganic phases are combined to achieve flexible, high-energy-density capacitors [17]. These examples show that dielectric property prediction must handle heterogeneous material classes, composite morphologies, and interface-dominated behavior. Aqueous tape casting for BaTiO₃ films for ultrathin multilayer ceramic capacitor applications further demonstrates how slurry chemistry, binder burnout, lamination, and sintering affect film quality and dielectric performance [11]. Such processing routes create pathways for machine learning to link formulation and manufacturing parameters to final dielectric properties, but only if data capture is sufficiently granular and metadata-rich. Without that granularity, models risk learning laboratory-specific artifacts rather than transferable materials physics.

3. Machine Learning Architectures for Dielectric Property Prediction

The earliest materials machine learning models often used handcrafted descriptors derived from composition, elemental properties, and crystal structure. For dielectric materials, descriptor-based regression can predict permittivity, loss, and breakdown-related quantities when sufficient high-quality data exist. Accelerated property prediction using machine learning demonstrated that relatively simple models can outperform empirical rules and guide

screening [8]. In polymer dielectrics, machine learning strategies have been used to explore chemical space for high permittivity and low loss, showing that data-driven design can identify promising repeat units and composites [6]. For dielectric breakdown, organized high-throughput data has been used to develop phenomenological understanding rather than only black-box prediction, illustrating the value of interpretable features [7]. These approaches are attractive for electronic ceramics because they can incorporate domain knowledge such as tolerance factors, electronegativity, ionic radius, and valence stability. However, they also inherit the limitations of descriptor choice: descriptors that ignore defect chemistry or grain boundary composition may miss the mechanisms that dominate real ceramic performance.

More expressive architectures include kernel methods, ensemble trees, Gaussian processes, and deep neural networks. Gaussian processes are especially useful when uncertainty quantification is required, because they provide predictive variance that can guide experimental selection. Neural networks can capture nonlinear interactions among composition and processing variables, but they typically require larger datasets and careful regularization. Graph neural networks and representation learning offer a way to encode local atomic environments without predefined descriptors, which is valuable for perovskites and disordered ceramics. Nevertheless, model complexity must be justified by data availability and deployment context. In industrial settings, a model that cannot be explained to a process engineer or validated with a small number of physical experiments may be rejected regardless of benchmark accuracy. The trade-off between predictive power and interpretability is therefore not philosophical; it is operational.

Machine learning interatomic potentials represent a different layer of prediction. High-dimensional neural network potentials can approximate quantum-mechanical energy landscapes with near first-principles accuracy at reduced cost [4]. Machine learning interatomic potentials are emerging as tools for simulating thermal transport, phase stability, defects, and polarization dynamics in materials [5]. For dielectric ceramics, such potentials could enable large-scale molecular dynamics or Monte Carlo simulations of domain wall motion, oxygen vacancy migration, and interface polarization. Yet the development of reliable potentials requires diverse training data from quantum calculations, careful validation on properties not included in training, and attention to long-range electrostatics. A potential trained on bulk BaTiO₃ may fail at surfaces, grain boundaries, or under high electric fields. Thus, even the most advanced simulation architecture remains dependent on data infrastructure and validation protocols. The integration of potentials with continuum models and device-level simulations remains an open systems challenge.

Active learning and adaptive sampling provide a bridge between prediction and experimentation. Instead of training once on a fixed dataset, active learning selects the most informative experiments or simulations to reduce uncertainty and improve decision-making [20]. This is particularly relevant for electronic ceramics, where each synthesis and measurement cycle is costly and where composition-processing spaces are vast. Active learning has been used in materials science to guide targeted design under uncertainty, but its success depends on accurate uncertainty estimates, fast surrogate models, and flexible laboratory workflows. Failed experiments can also be valuable training data, as demonstrated in materials discovery where negative results improved predictive models [9]. High-throughput screening has been used to evaluate thousands of candidate solid electrolytes, showing how computational filters can prioritize experimental investigation [10]. For dielectric ceramics, similar screening could rank dopants, sintering aids, or layered

architectures, but the property targets are multi-objective and often conflicting. A material with high permittivity may have high loss, poor breakdown strength, or unacceptable temperature stability. Multi-objective optimization and Pareto analysis therefore become essential.

4. Data Infrastructure, Experimental Integration, and Deployment Pipelines

The quality of machine learning prediction is bounded by the quality of data infrastructure. In electronic ceramics, relevant data are generated by X-ray diffraction, electron microscopy, impedance spectroscopy, ferroelectric hysteresis, breakdown testing, thermal analysis, and manufacturing process logs. These data are heterogeneous, often stored in laboratory notebooks, spreadsheets, instrument-specific formats, and proprietary databases. Materials intelligence ecosystems require interoperable schemas, persistent identifiers, provenance tracking, and versioned datasets [18]. Without such infrastructure, models cannot be reliably reproduced or updated. Metadata must capture not only composition but also synthesis route, powder source, milling time, binder system, sintering profile, electrode material, measurement frequency, temperature, humidity, and electrical bias. A model trained on permittivity values without these contextual variables may appear accurate in cross-validation but fail when transferred to a new furnace or supplier. Data infrastructure is therefore not a back-office concern; it is a primary determinant of predictive validity.

Experimental integration is equally important. A prediction pipeline that is disconnected from laboratory information management systems, instrument control, and quality assurance procedures will struggle to influence real decisions. Self-driving laboratories aim to close the loop by using machine learning to propose experiments, robotic systems to execute them, and automated characterization to generate new labels. For ceramic dielectrics, full automation is difficult because powder processing, tape casting, sintering, and electrode application involve many manual or semi-manual steps. Nevertheless, partial automation can be valuable, especially for combinatorial synthesis and high-throughput electrical testing. Model documentation practices such as model cards can clarify intended use, training data, evaluation conditions, and limitations before a model is deployed in an engineering workflow [22]. Algorithmic auditing frameworks can further assess whether a model performs consistently across material classes, suppliers, and operating conditions [21]. These governance tools are not replacements for physical validation, but they help prevent silent failures when models are reused outside their original context.

Deployment pipelines must also address concept drift. Material systems, manufacturing equipment, and market requirements change over time. A model trained on historical BaTiO₃ formulations may degrade when the supplier changes powder morphology or when regulations restrict a dopant. Concept drift research shows that statistical relationships between inputs and outputs can change gradually or abruptly, requiring monitoring and adaptation [23]. In dielectric ceramics, drift can arise from new sintering aids, modified electrode compositions, or revised test standards. A robust deployment architecture should therefore include drift detection, periodic retraining, uncertainty thresholds, and fallback rules. It should also document when a prediction is outside the model's validated domain. For safety-critical applications such as automotive or medical electronics, the ability to abstain from prediction may be more valuable than a confident but wrong estimate. This is a system design principle: reliability is achieved through redundancy, monitoring, and human oversight, not solely through model accuracy.

5. Structural Trade-offs in Model Design and Ceramic Processing

Every machine learning architecture embeds structural trade-offs. Deep models can capture complex nonlinearities but require large datasets and computational resources. Kernel methods and Gaussian processes offer uncertainty estimates but scale poorly with large data. Descriptor-based models are interpretable but may miss emergent phenomena. Physics-informed models can improve extrapolation but require assumptions that may not hold across processing conditions. In solid-state materials science, these trade-offs are well documented, and no single approach dominates all tasks [3]. For dielectric ceramics, the choice of model should follow the decision problem. Early-stage screening may prioritize recall, accepting some false positives to avoid missing promising candidates. Late-stage qualification may prioritize precision and calibrated uncertainty because experimental validation is expensive and failure is costly. A system-level architecture may combine multiple models, using fast filters for broad screening and slower, more accurate models for final ranking.

Processing trade-offs are equally significant. Aqueous tape casting for ultrathin multilayer ceramic capacitors involves balancing slurry stability, green sheet strength, binder burnout, shrinkage matching, and final dielectric density [11]. Machine learning can model these relationships, but the data are often sparse because experiments are expensive and industrial process windows are proprietary. Transfer learning from related materials or processes can help, but negative transfer is a risk when the underlying chemistry differs. Multilayer ceramic capacitor manufacturing further imposes constraints from co-firing with base-metal electrodes, where reduction resistance and oxygen partial pressure control are critical [14]. Lead-free dielectric development adds constraints related to temperature stability and energy storage density [16]. These processing realities mean that a predicted material with excellent intrinsic permittivity may be unmanufacturable or unreliable. The model must therefore be coupled to manufacturability models, cost models, and reliability models.

Deployment trade-offs include latency, memory, and maintainability. A model intended for cloud-based research screening can be large and computationally heavy. A model embedded in a production line or edge diagnostic device must operate with limited resources and deterministic timing. High-throughput screening of solid electrolytes demonstrates the value of computational filters, but those filters are typically run offline or on clusters [10]. Ceramic capacitor manufacturing may benefit from rapid inline prediction of dielectric layer quality, but such models must be robust to sensor noise and environmental variation. A brief introduction to ceramic capacitors emphasizes that device performance depends on geometry and materials integration, not only on intrinsic dielectric constants [15]. Consequently, deployment architectures should separate concerns: data acquisition, feature extraction, inference, uncertainty estimation, decision logic, and human review. This modularity supports maintainability and allows individual components to be updated as new data arrive.

6. Robustness, Fairness, and Governance in Materials Intelligence

Robustness is a first-order requirement for machine learning in materials engineering. Models must perform reliably under distribution shift, measurement noise, missing data, and adversarial or accidental outliers. In dielectric ceramics, distribution shift can occur when a model trained on laboratory-scale pellets is applied to tape-cast films, or when a model trained on one measurement frequency is used at another. Concept drift research provides methods for detecting changes in data distributions over time [23]. Robustness also requires stress testing on minority classes, such as compositions with unusual dopants or processing routes. A model that performs well on average but fails for a critical high-temperature application may be unacceptable. Robustness should therefore be evaluated through worst-

case analysis, sensitivity analysis, and domain-specific validation. In solid-state materials science, the diversity of structures and properties makes robustness especially challenging [3].

Fairness in materials intelligence is less often discussed but increasingly relevant. Fairness can refer to equitable access to data, computational resources, and experimental facilities across institutions and regions. It can also refer to model performance across material classes, suppliers, or application domains. If a model is trained primarily on data from well-funded laboratories using high-purity precursors, it may perform poorly for manufacturers using recycled or lower-grade raw materials. Algorithmic auditing frameworks can help identify such disparities by examining performance across subgroups and documenting decision impacts [21]. Model cards can disclose training data composition and known limitations, enabling users to assess whether a model is appropriate for their context [22]. In electronic ceramics, fairness also intersects with environmental justice, because extraction and processing of critical metals can impose local burdens. A governance system that ignores these dimensions may optimize technical metrics while worsening social and environmental outcomes.

Governance mechanisms must span the lifecycle of models, data, and decisions. At the dataset level, governance includes consent, licensing, provenance, and quality control. At the model level, it includes documentation, validation, versioning, and access control. At the deployment level, it includes monitoring, incident reporting, and rollback procedures. Materials intelligence ecosystems require coordination among universities, national laboratories, manufacturers, and regulators [18]. Reviews of machine learning for energy materials highlight the need for standardized benchmarks and reproducible workflows [19]. For dielectric ceramics, governance should also address safety and reliability standards, especially when predictions influence component qualification for automotive, aerospace, or medical devices. A model that recommends a new dielectric composition may trigger costly synthesis and testing; if the model is opaque or unvalidated, it may erode trust. Governance is therefore not bureaucratic overhead but a condition for adoption.

7. Sustainability and Policy Implications

Sustainability is a growing concern in electronic ceramics because manufacturing involves high-temperature sintering, solvent use, and critical raw materials. Lead-free dielectric development is motivated partly by toxicity and regulatory pressure, but lead-free alternatives may require different processing conditions or rare dopants [16]. Multilayer ceramic capacitor manufacturing has improved through base-metal electrode technology, but it still consumes significant energy and materials [14]. Machine learning can contribute to sustainability by reducing the number of failed experiments, optimizing sintering profiles, and identifying compositions that use less critical elements. However, the computational cost of training large models and running high-throughput simulations must also be considered. A sustainable materials intelligence system should evaluate the environmental footprint of both the model and the material it discovers.

Policy implications include support for open data infrastructure, shared experimental facilities, and standards for data and model reporting. Ceramic capacitors are ubiquitous, yet much manufacturing data are proprietary. A balance must be struck between intellectual property protection and the public benefits of shared benchmarks. Government-funded research often requires data management plans, but these plans may not produce interoperable or reusable datasets. Policymakers could encourage domain-specific metadata standards for dielectric measurements, synthesis protocols, and device testing. They could also fund independent

validation centers that test machine learning predictions under industrial conditions. In the polymer dielectrics domain, high-temperature nanocomposites illustrate how academic, industrial, and defense-related interests intersect [17]. Similar coordination is needed for electronic ceramics, where reliability and supply chain resilience are national and economic concerns.

Sustainability also intersects with fairness. If advanced machine learning tools are available only to large firms with extensive data and computing resources, smaller manufacturers and researchers may be disadvantaged. Open-source models, federated learning, and shared benchmarks can reduce barriers, but they require governance to protect data contributors and ensure quality. Policy should therefore support not only model development but also the human capacity to use models responsibly. Training programs for materials scientists, process engineers, and regulators should include data literacy, uncertainty communication, and ethical considerations. The goal is not to replace domain expertise but to augment it. A policy environment that treats machine learning as a tool within a broader innovation ecosystem is more likely to produce durable benefits than one that treats it as a standalone solution.

8. Case Illustrations and Cross-Domain Comparisons

The screening of solid electrolytes offers a useful cross-domain comparison. Researchers used computational structure screening to evaluate thousands of candidates for lithium-ion conductors, demonstrating how high-throughput filters can prioritize experimental work [10]. The analogy to dielectric ceramics is clear: composition and structure spaces are vast, and experimental validation is expensive. However, solid electrolytes are judged primarily by ionic conductivity and stability, whereas dielectric ceramics must satisfy multiple electrical, thermal, and mechanical criteria. The multi-objective nature of dielectric design makes Pareto optimization more complex. Failed experiments in materials discovery provide another lesson: negative results can improve models if they are captured systematically [9]. In ceramic processing, failed batches, abnormal grain growth, and breakdown failures are often discarded rather than curated. A learning infrastructure that treats failure as data could improve robustness and reduce waste.

Polymer dielectrics illustrate how machine learning can accelerate chemical design. Machine learning strategies have been used to explore polymer dielectrics for high permittivity and low loss [6]. High-temperature polymer nanocomposites show that interface engineering and composite morphology can dramatically affect breakdown strength and energy density [17]. Electronic ceramics share the need to control interfaces, but they also involve grain boundaries, pores, and dopant segregation. Aqueous tape casting for BaTiO₃ films for ultrathin multilayer ceramic capacitors demonstrates how processing innovations can push device performance [11]. Multilayer ceramic capacitor manufacturing provides a mature industrial context in which machine learning could support process optimization, defect detection, and reliability prediction [14]. The lesson from these cases is that domain-specific data and validation matter more than generic model sophistication.

Cross-domain comparisons also reveal governance differences. Energy materials research often emphasizes open databases and shared benchmarks [19]. Medical artificial intelligence has developed model cards and auditing frameworks [21], [22]. Materials science can adapt these practices, but it must account for physical experimentation, safety, and long product lifecycles. Concept drift methods from machine learning operations can inform monitoring of material models as suppliers and processes change [23]. Active learning from materials science provides a principled way to select experiments under uncertainty [20]. Taken

together, these cross-domain insights suggest that electronic ceramic prediction should be built as a federated, auditable, and adaptive system. It should connect simulation, experiment, manufacturing, and policy without assuming that any single model will remain valid indefinitely.

9. Future Perspectives

Future progress will depend on multimodal data integration. Dielectric properties are influenced by composition, crystal structure, microstructure, interfaces, and processing history. A model that uses only compositional features will remain limited. Integrating diffraction, microscopy, spectroscopy, and electrical measurements into unified representations could improve prediction and explanation. Materials intelligence ecosystems provide a conceptual template for such integration [18]. Self-driving laboratories and active learning can then prioritize experiments that reduce uncertainty and test model boundaries [20]. However, these systems require significant investment in instrumentation, data standards, and human-machine interfaces. They also require new forms of collaboration between computer scientists, ceramic engineers, physicists, and industrial partners.

Physics-informed machine learning and machine learning interatomic potentials offer another direction. High-dimensional neural network potentials can simulate atomic-scale phenomena with increasing accuracy [4]. Machine learning interatomic potentials can bridge quantum accuracy and larger length scales [5]. For dielectric ceramics, such methods could predict polarization switching, defect migration, and interfacial charge trapping. Yet their reliability depends on training data diversity and validation against experimental observables. Hybrid models that combine physics-based simulation with data-driven correction may be more robust than purely black-box approaches. These models also need uncertainty quantification so that engineers know when to trust them. The future architecture is therefore likely to be ensemble-based, modular, and continuously updated.

Finally, governance and policy must evolve alongside technical capability. Model documentation [22], algorithmic auditing [21], and drift monitoring [23] should become routine in materials research and manufacturing. Open benchmarks and shared metadata standards can improve reproducibility. Sustainability criteria should be incorporated into objective functions, not treated as afterthoughts. Fairness should be assessed across institutions, regions, and material classes. The most advanced prediction system will be of limited value if it is inaccessible, unaccountable, or environmentally harmful. A balanced research agenda should therefore pursue accuracy, interpretability, robustness, and governance simultaneously. Electronic ceramics are too important to modern infrastructure for machine learning to be treated as a purely computational shortcut.

10. Conclusion

Machine learning-driven prediction of dielectric properties in advanced electronic ceramics has substantial potential to accelerate discovery, optimize processing, and improve reliability. However, the field requires a system-level orientation. Predictive models must be embedded in data infrastructures, experimental workflows, manufacturing systems, and governance frameworks. The material complexity of ferroelectric perovskites, multilayer ceramic capacitors, lead-free dielectrics, and polymer-ceramic composites means that no single model or dataset will suffice. Descriptor-based regression, neural network potentials, active learning, and multimodal representations each address part of the problem, but their value depends on validation, uncertainty estimation, and domain-aware deployment. Robustness under

distribution shift, fairness in access and performance, and accountability through documentation and auditing are not peripheral concerns; they determine whether machine learning can be trusted in engineering decisions. Sustainability and policy implications further shape the long-term trajectory, from critical raw materials to energy-intensive sintering and computational footprints. A successful research program will therefore combine technical innovation with infrastructural investment, interdisciplinary collaboration, and responsible governance. By treating machine learning for electronic ceramics as a socio-technical system, researchers and policymakers can improve both the performance of dielectric materials and the legitimacy of the predictive tools used to design them.

References

1. 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, 54. <https://doi.org/10.1038/s41524-017-0056-5>
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. <https://doi.org/10.1038/s41586-018-0337-2>
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. <https://doi.org/10.1038/s41524-019-0221-0>
4. Behler, J. (2021). Four generations of high-dimensional neural network potentials. *Chemical Reviews*, 121(16), 10037–10072. <https://doi.org/10.1021/acs.chemrev.0c00868>
5. Deringer, V. L., Caro, M. A., & Csányi, G. (2019). Machine learning interatomic potentials as emerging tools for materials science. *Advanced Materials*, 31(46), 1902765. <https://doi.org/10.1002/adma.201902765>
6. Mannodi-Kanakkithodi, A., Pilania, G., Huan, T. D., Lookman, T., & Ramprasad, R. (2016). Machine learning strategy for accelerated design of polymer dielectrics. *Scientific Reports*, 6, 20952. <https://doi.org/10.1038/srep20952>
7. Kim, C., Pilania, G., & Ramprasad, R. (2016). From organized high-throughput data to phenomenological theory using machine learning: The example of dielectric breakdown. *Chemistry of Materials*, 28(5), 1304–1311. <https://doi.org/10.1021/acs.chemmater.5b04109>
8. Pilania, G., Wang, C., Jiang, X., Rajasekaran, S., & Ramprasad, R. (2013). Accelerating materials property predictions using machine learning. *Scientific Reports*, 3, 2810. <https://doi.org/10.1038/srep02810>
9. 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. <https://doi.org/10.1038/nature17439>
10. Sendek, A. D., Yang, Q., Cubuk, E. D., Duerloo, K.-A. N., Cui, Y., & Reed, E. J. (2017). Holistic computational structure screening of more than 12,000 candidates for solid lithium-ion conductor materials. *Energy & Environmental Science*, 10(1), 306–320. <https://doi.org/10.1039/C6EE02697D>

11. Luan, Saiwei, et al. "Fabrication of High Performance BaTiO₃ Films via Aqueous Tape Casting for ultrathin MLCC applications." *Ceramics International* (2025).
12. Cohen, R. E. (1992). Origin of ferroelectricity in perovskite oxides. *Nature*, 358(6382), 136–138. <https://doi.org/10.1038/358136a0>
13. Scott, J. F. (2007). Applications of modern ferroelectrics. *Science*, 315(5814), 954–959. <https://doi.org/10.1126/science.1129564>
14. Kishi, H., Mizuno, Y., & Chazono, H. (2003). Base-metal electrode-multilayer ceramic capacitors: Past, present and future perspectives. *Japanese Journal of Applied Physics*, 42(1R), 1–15. <https://doi.org/10.1143/JJAP.42.1>
15. Pan, M. J., & Randall, C. A. (2010). A brief introduction to ceramic capacitors. *IEEE Electrical Insulation Magazine*, 26(3), 44–50. <https://doi.org/10.1109/MEI.2010.5482787>
16. Yang, L., Kong, X., Li, F., Hao, H., Cheng, Z., Liu, H., Li, J.-F., & Zhang, S. (2019). Perovskite lead-free dielectrics for energy storage applications. *Progress in Materials Science*, 102, 72–108. <https://doi.org/10.1016/j.pmatsci.2018.12.005>
17. Li, Q., Chen, L., Gadinski, M. R., Zhang, S., Zhang, G., Li, H. U., Iagodkine, E., Haque, A., Chen, L.-Q., Jackson, T. N., & Wang, Q. (2015). Flexible high-temperature dielectric materials from polymer nanocomposites. *Nature*, 523(7562), 576–579. <https://doi.org/10.1038/nature14647>
18. Himanen, L., Geurts, A., Foster, A. S., & Rinke, P. (2019). Data-driven materials science: Status, challenges, and perspectives. *Advanced Science*, 6(21), 1900808. <https://doi.org/10.1002/advs.201900808>
19. Chen, C., Zuo, Y., Ye, W., Li, X., Deng, Z., & Ong, S. P. (2020). A critical review of machine learning of energy materials. *Advanced Energy Materials*, 10(8), 1903242. <https://doi.org/10.1002/aenm.201903242>
20. 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. <https://doi.org/10.1038/s41524-019-0153-8>
21. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44. <https://doi.org/10.1145/3351095.3372873>
22. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–229. <https://doi.org/10.1145/3287560.3287596>
23. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 44. <https://doi.org/10.1145/2523813>