

# **Machine Learning-Enabled Process–Structure–Property Optimization of Ceramic Materials for Additive Manufacturing**

Anish M. Mahajan

Department of Computer Science, University of North Texas, Denton, TX, USA.  
anish1974@unt.edu

Ruben J. Reynolds

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

Arthur Willis

Department of Computer Science, Binghamton University, Binghamton, NY, USA.  
willisarthur@binghamton.edu

Jeffrey R. Rhodes

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis,  
OR, USA.  
jeffreymail@oregonstate.edu

## **Abstract**

The additive manufacturing of ceramic components offers transformative opportunities for energy storage, biomedical implants, aerospace thermal protection, and microelectronic packaging. However, the strong coupling between processing conditions, multiscale microstructure, and final performance remains a persistent barrier to reliable industrial deployment. Machine learning has emerged as a powerful methodology for navigating this complexity, but its value depends on the broader system in which it is embedded. This paper presents a systems-oriented analysis of machine learning-enabled process-structure-property optimization for ceramic additive manufacturing. Rather than focusing narrowly on model accuracy, it examines the architectural, infrastructural, and governance dimensions required to integrate predictive models into materials development workflows. The paper reviews ceramic-specific process constraints, including rheology, debinding, drying, and sintering, and discusses how heterogeneous data from in situ sensing, characterization, and simulation can be organized into learning systems. Structural trade-offs are explored among model interpretability and predictive power, centralized and distributed data architectures, and short-term optimization and long-term robustness. The analysis extends to sustainability, fairness, and accountability concerns that arise when data-driven tools influence experimental design and industrial qualification. Policy implications are considered through emerging artificial intelligence governance frameworks. The paper argues that durable progress depends not only on better predictive models but also on institutional mechanisms, shared data protocols, and socio-technical safeguards. A holistic systems architecture can enable machine learning to support ceramic additive manufacturing from powder formulation and printing through microstructure control and property certification. This perspective connects materials

informatics, advanced manufacturing, and responsible artificial intelligence governance for academic and industrial research communities.

## **Keywords**

ceramic additive manufacturing; machine learning; process-structure-property optimization; data infrastructure; governance; robustness; policy.

## **1. Introduction**

Ceramic additive manufacturing has attracted sustained interest because ceramics possess high hardness, thermal stability, corrosion resistance, and dielectric or piezoelectric functionality that are difficult to realize with metals and polymers. However, the same properties that make ceramics attractive also complicate their layerwise fabrication. High melting temperatures and low thermal shock resistance mean that laser-based processes often generate cracks, porosity, and anisotropic microstructures. Binder-based and slurry-based methods introduce additional challenges associated with rheology, drying, organic removal, and densification. These challenges make the relation between processing conditions, as-printed structure, and final properties highly nonlinear and context-dependent [1]. In this setting, the classical process-structure-property paradigm becomes indispensable because a systematic mapping among process parameters, microstructural descriptors, and performance attributes is needed to move beyond trial-and-error experiments [2].

Machine learning has become a central tool in materials science because it can approximate high-dimensional relationships without requiring complete first-principles insight. Early materials informatics efforts demonstrated that appropriate descriptors and curated datasets allow predictive models to guide experimental exploration across large composition and processing spaces [3]. More recent work has extended these methods from molecular and crystalline systems to complex manufacturing environments where data are noisy, multimodal, and generated at multiple length scales and time scales [4]. In additive manufacturing, machine learning models have been used for process monitoring, defect classification, parameter recommendation, and property prediction [5]. These applications suggest that data-driven optimization can accelerate the qualification of ceramic materials, but they also reveal that the underlying system architecture is at least as important as the choice of learning algorithm.

This paper adopts a systems-level perspective on machine learning-enabled process-structure-property optimization for ceramic additive manufacturing. It does not propose a single model or algorithm. Instead, it analyzes the structural trade-offs, data infrastructures, governance arrangements, deployment pathways, and policy constraints that shape whether learning systems become reliable and sustainable in industrial practice. Such a perspective is necessary because ceramic additive manufacturing data are often sparse, heterogeneous, and proprietary, while the required predictions span slurry formulation, printing parameters, drying stresses, sintering shrinkage, grain growth, porosity, and final electromechanical response. By treating machine learning as part of a larger socio-technical infrastructure, the paper connects materials engineering with emerging debates about robustness, fairness, and accountability in artificial intelligence.

## **2. Process-Structure-Property Chain and Ceramic-Specific Constraints**

The process-structure-property chain in ceramic additive manufacturing is more difficult to characterize than in many metallic systems because multiple phase transformations and

transport phenomena occur sequentially. A ceramic part may begin as a powder bed or a viscoelastic slurry, undergo layer shaping through selective binding, photopolymerization, extrusion, or laser interaction, and then require debinding and sintering to achieve a dense microstructure. Each stage introduces distinct physical mechanisms and failure modes. Binder chemistry affects powder packing and green strength; printing speed and layer thickness influence surface roughness and internal interfaces; heating rate and atmosphere determine cracking and grain size [6]. A parallel challenge appears in ceramic film casting, where aqueous slurry formulation, drying conditions, and sintering schedules must be co-optimized to achieve high density and uniform dielectric performance in ultrathin multilayer devices [7]. Although tape casting is not an additive manufacturing process, its dependence on the same processing sequence illustrates why predictive models must account for coupled composition, transport, and microstructural evolution.

Robocasting and other direct ink writing methods have shown that rheological properties such as yield stress, storage modulus, and recovery time govern print fidelity and layer stacking [8]. Selective laser sintering and melting of ceramics, by contrast, require careful control of local thermal gradients because rapid solidification and solid-state transformations can produce residual stress and microcracking [9]. These examples demonstrate that a single ceramic additive manufacturing platform cannot be described by one universal model. A machine learning system for one process may rely heavily on powder size distribution and laser scan parameters, while another may depend on ink viscoelasticity and drying kinetics. The structure of the data and the causal proximity of predictors to final properties therefore vary across process families. Consequently, system designers must explicitly represent multiple processing stages and their interactions instead of assuming that a monolithic model can capture all relevant physics.

Moreover, the property targets in ceramics are often highly specialized. For energy storage or microelectronic applications, dielectric constant, breakdown strength, leakage current, and thermal stability must be optimized together. For structural applications, fracture toughness, hardness, creep resistance, and thermal shock resistance are important. These requirements create multi-objective optimization problems in which improving one property can degrade another. Machine learning can help by mapping process parameters to these properties through microstructural descriptors such as grain size distribution, porosity, phase fractions, crack density, and crystallographic texture. However, reliable mapping depends on a consistent representation of microstructure across imaging modalities and on the ability to track how these descriptors evolve through sintering and cooling. The next section therefore considers the architectural role of machine learning in linking process variables, microstructural states, and property outcomes.

### **3. Machine Learning Architectures for Multiscale Optimization**

Machine learning models for ceramic additive manufacturing must integrate information from multiple scales because molecular interactions, particle packing, layer-level thermal history, mesoscale defects, and component-level mechanical behavior all influence final performance. A common architectural strategy is to use surrogate models that replace expensive simulations or experiments for selected subproblems, such as predicting sintering shrinkage from initial density and thermal schedule. Another strategy is to use convolutional neural networks on microstructural images to predict effective properties or to identify process-induced defects. Reviews of machine learning in additive manufacturing emphasize that successful implementations depend heavily on feature engineering, data quality, and alignment between

model outputs and actionable process controls [10]. In ceramic systems, these requirements are especially severe because many process variables are continuous, strongly correlated, and constrained by equipment limits.

One important design choice is whether to adopt physics-informed models, purely data-driven models and hybrid formulations. Purely data-driven models often fail when extrapolating beyond the training distribution, while physics-informed models can incorporate conservation laws, thermodynamic constraints, or known sintering kinetics. Hybrid architectures can use physics-derived features, such as normalized energy density or capillary pressure, as inputs to neural networks, thereby limiting the search space and improving generalization across equipment settings. A systems design must decide where physics knowledge enters the learning pipeline, how model outputs are transformed into process parameter recommendations, and how uncertainty propagates through the multiscale chain [11]. This is particularly important for ceramic additive manufacturing because simulation-based training data are expensive, experimental data are often scarce, and the cost of failure during sintering or qualification is high.

A second architectural consideration concerns the distinction between global and local models. A global model may attempt to predict final properties across multiple machines, powders, and geometries, but such breadth frequently sacrifices resolution. A local model trained on a single slurry system or sintering furnace can achieve higher predictive fidelity, yet it may not transfer to a new feedstock or thermal profile. Systems that organize models hierarchically, with coarse global screens followed by local refinement, can balance these tensions. For example, a global classifier might identify whether a proposed parameter set is likely to cause cracking, while local Gaussian process surrogates recommend specifics for a given print geometry. The communication between these layers, including calibration and retraining schedules, must be designed explicitly. Without such architectural clarity, machine learning systems become brittle collections of separate models whose recommendations conflict.

The role of artificial neural networks in microstructural characterization also deserves systems-level attention. Convolutional networks can segment porosity, grains, and cracks from optical or electron microscopy images, but their outputs must be transformed into quantitative descriptors that downstream property models can consume. The absence of shared descriptor schemas creates integration failures even when individual models perform well. Recent work on data-driven materials science emphasizes that standardized representations and reusable data objects are needed if the field is to move from isolated demonstrations to routine use in process optimization [12]. In ceramic additive manufacturing, microstructural descriptors must account for anisotropic features such as layer interfaces, large elongated pores, and crystallographic textures that are not easily summarized by single scalar metrics. Consequently, model architectures must be evaluated not only by regression error but also by their ability to preserve the uncertainty, spatial structure, and process relevance of microstructural information.

#### **4. Data Infrastructure and Interoperability**

A machine learning system for ceramic additive manufacturing depends on the quality and accessibility of the data that feed it. However, the data landscape remains fragmented across laboratories, companies, and instrument formats. Process logs, slurry rheology measurements, thermal images, X-ray computed tomography, electron microscopy, and mechanical test results are often stored in separate systems with incompatible metadata. This fragmentation limits the ability to construct process-structure-property datasets that span enough variability

to support robust learning. The concept of findable, accessible, interoperable, and reusable data has become widely accepted in scientific communities, yet its implementation in manufacturing settings is still uneven [13]. The systems challenge is therefore not merely technical but also organizational, involving decisions about data custody, access control, and metadata standards that must be negotiated across stakeholders.

Interoperability is particularly difficult in ceramic processing because different stages generate fundamentally different data types. A rheological measurement is a time series, a thermal camera produces image sequences, a sintering furnace log contains temperature and atmosphere histories, and a micrograph is a high-dimensional spatial dataset. To use these data jointly, the system must represent process sequences as linked objects rather than isolated files. This can be achieved through graph-based data models in which each sample, batch, layer, and characterization event has a unique identifier and explicit relationships to parent materials and processing conditions. Such models enable traceability across the lifecycle and facilitate the construction of training datasets that align inputs and outputs correctly. They also reduce the risk of silent data leakage, where information from the output stage inadvertently contaminates training features, leading to overoptimistic performance estimates.

Data infrastructure for ceramic additive manufacturing must also support versioning of both data and models. A powder lot may vary in composition or moisture content over time, a furnace may drift in temperature calibration, and an imaging system may change illumination conditions. Without versioned data and model registries, rollback and auditing become impossible. The model governance literature in industrial machine learning highlights the need for lineage tracking and controlled retraining pipelines rather than ad hoc model updates [14]. For ceramic systems, this is especially critical because safety-critical components may require certification evidence that a given model version was trained on data representative of the production condition. The broader manufacturing community has begun to adopt data platforms that support these requirements, but their application to ceramic powders, slurries, and high-temperature post-processing remains underdeveloped.

Another important issue is the scarcity of openly available ceramic additive manufacturing datasets. While metallic additive manufacturing has benefited from public competitions and benchmark data, ceramic datasets are smaller and more guarded due to the specialized nature of the feedstocks and the commercial value of formulations. This scarcity encourages models that overfit to narrow conditions and underrepresent the effects of raw material variability. Shared data consortia, precompetitive data pools, and federated learning arrangements offer potential solutions. Federated approaches allow participants to train shared models without transferring raw data across organizational boundaries, thereby preserving confidentiality while expanding statistical coverage. However, these approaches require substantial coordination and governance to ensure that contributions are not free-ridden and that model updates do not inadvertently leak sensitive information. Systems design must therefore address incentive structures alongside technical protocols.

## **5. Structural Trade-offs in Model Design and Deployment**

A central tension in machine learning-enabled materials optimization is between predictive accuracy and interpretability. Highly flexible models such as deep neural networks can capture subtle nonlinearities in ceramic processing, but their decision logic is often opaque. In contrast, linear models or decision trees offer clearer reasoning but may miss critical interactions among variables such as slurry pH, drying rate, and sintering ramp. Interpretability matters for process engineers who must decide whether to accept a model

recommendation and for certification authorities who must assess risk [15]. Interpretable machine learning techniques, including local explanations, feature importance measures, and surrogate models, provide partial visibility into black-box behavior, but they also introduce additional computational and conceptual burdens [16]. The systems-level question is not whether interpretability should be maximized but where it is most needed: at early screening, final qualification, or operational troubleshooting.

A second structural trade-off concerns centralization versus distribution. A centralized data repository can enforce uniform quality control and enable large-scale training, but it may concentrate power and create bottlenecks. A distributed architecture preserves local autonomy and can accommodate proprietary constraints, but it complicates model consistency and evaluation. Hybrid models can use centralized metadata catalogs with distributed raw data, with computation moving to the data rather than data moving to a central location. Such architectures are familiar in large-scale scientific computing and are increasingly relevant for advanced manufacturing. For ceramic additive manufacturing, the choice depends on the number of partner sites, the sensitivity of feedstock formulations, and the need for rapid model updates across multiple production lines. An architecture that works for a single university laboratory may fail for a networked industrial supply chain with competing business units.

Deployment trade-offs also involve the timing and placement of predictions. A model can be deployed before printing to recommend parameters, during printing to detect defects from sensor streams, or after sintering to predict final properties from microstructural images. Each placement requires different latency, precision, and reliability characteristics. Pre-print recommendation can tolerate more computational time, while in situ defect detection must operate near real time and integrate with control loops. Post-process prediction may be used for acceptance decisions and thus requires calibrated uncertainty. The systems perspective suggests that the same underlying feature set and model family may not be optimal across all placements. Instead, a modular pipeline with shared feature engineering and task-specific heads can be more maintainable. Modularity also enables independent validation and incremental improvement without retraining the entire system.

Finally, a governance-aware design must balance autonomy and human oversight. Automated parameter recommendation can accelerate workflows, but fully automated decision-making may hide assumptions from operators and reduce opportunities for tacit knowledge capture. A well-designed system can treat machine learning as an advisory layer that proposes alternatives, explains the expected consequences, and logs the human decision for future learning. This human-in-the-loop structure is not only ethically preferable but also technically advantageous because operator feedback becomes a source of new labeled data. In ceramic manufacturing, where experienced engineers often notice subtle slurry or furnace behavior, such feedback loops can capture knowledge that would otherwise remain implicit. The deployment architecture must therefore include mechanisms for annotation, overrides, and periodic review.

## **6. Robustness, Uncertainty, and Failure Modes**

Robustness is a fundamental concern for machine learning systems operating in ceramic additive manufacturing because process distributions shift over time. Raw material lots vary, ambient humidity changes, nozzle wear alters flow behavior, and furnace calibration drifts. A model trained on historical data may silently degrade when these shifts occur. Concept drift adaptation is an established problem in machine learning, yet its integration into materials

workflows remains sparse [18]. Robust systems require drift detection, monitoring of prediction residuals, and scheduled retraining. In high-temperature ceramic processing, data collection is expensive, so waiting for large new labeled datasets may be impractical. Methods that efficiently incorporate small batches of new observations, such as active learning and transfer learning, can help maintain performance under distribution shift.

Uncertainty quantification is closely related to robustness. A process engineer must know not only what property a model predicts but also how confident that prediction is. Epistemic uncertainty arises from limited data or model capacity, while aleatoric uncertainty arises from irreducible noise in the process itself. Separate treatment of these sources is important because the former can be reduced by collecting more data, while the latter may require more robust process design. Calibration and conformal prediction methods can provide prediction intervals with statistical guarantees under mild assumptions, but they are rarely deployed in ceramic additive manufacturing. The absence of rigorous uncertainty estimates limits the use of machine learning for safety-critical qualification and slows regulatory acceptance. A systems-level framework should therefore require models to report uncertainty alongside any recommendation that may affect material integrity.

Failure modes in process-structure-property optimization can be subtle. A model may achieve high average prediction accuracy while missing rare but consequential defects such as localized cracking or abnormal grain growth. Because these events are underrepresented in training data, standard loss functions may not reflect their importance. Cost-sensitive learning, anomaly detection, and causal reasoning can help address this imbalance. In ceramic systems, a missed crack prediction may lead to catastrophic part failure, while a false alarm may cause unnecessary scrap. The acceptable trade-off depends on the application context and must be specified explicitly. A model that is optimal for rapid prototyping may be unacceptable for aerospace components. Governance mechanisms should require explicit risk matrices and consequence-weighted evaluation rather than relying on generic accuracy metrics.

## **7. Governance, Fairness, and Policy Implications**

The use of machine learning in ceramic additive manufacturing raises governance challenges that connect materials science with broader artificial intelligence policy. These challenges include data provenance, model accountability, intellectual property, and liability when a data-driven recommendation leads to component failure. Existing artificial intelligence ethics frameworks emphasize principles such as transparency, justice, non-maleficence, and accountability, yet their translation to manufacturing contexts is not straightforward [19]. A model that discriminates among demographic groups is less obviously relevant in a powder processing setting, but fairness can still arise through unequal access to proprietary data, uneven deployment of advanced tools across firms, and biases introduced by representative training sets that omit certain geographic or material sources. The systems perspective must therefore consider who benefits from machine learning acceleration and who bears the costs of errors.

Policy instruments for artificial intelligence are evolving rapidly. Regulatory approaches such as risk-based classification and conformity assessment for high-risk systems have been proposed in multiple jurisdictions [20]. Ceramic additive manufacturing components used in medical devices, energy infrastructure, or transportation may fall under high-risk categories, requiring evidence of model robustness, data quality, and human oversight. A systems architecture that documents training data, version history, validation protocols, and uncertainty estimates will be better positioned to comply with such requirements than a

collection of isolated models. The integration of model reporting artifacts, similar to model cards or datasheets, can formalize this documentation [17]. These artifacts should be treated as part of the manufacturing quality system rather than as academic afterthoughts.

Algorithmic accountability also requires clarity about responsibility when multiple actors contribute to a prediction. A powder supplier, a printer operator, a model developer, and a sintering service provider may all influence the final outcome. If a machine learning system recommends a parameter set that later causes cracking, assigning liability and learning from the incident requires traceable decision logs and shared responsibility frameworks. Without such mechanisms, organizations may resist adopting machine learning systems because the legal and reputational risks are uncertain. Precompetitive standards bodies and industry consortia can help develop shared governance protocols that preserve commercial confidentiality while enabling collective learning. The ethical analysis of algorithms suggests that accountability is not a single property but a distributed practice that must be designed into organizational routines and technical systems [21].

International dimensions also matter. Ceramic additive manufacturing data may cross national borders through cloud infrastructure or distributed collaborations. Different jurisdictions may impose distinct rules on data localization, export controls, and algorithmic transparency. A machine learning system designed without attention to these legal boundaries may become unusable in global supply chains. Policy-aware system design should incorporate configurable data residency, access permissions, and auditability. It should also distinguish between raw material data, process data, and derived model parameters, because these categories may trigger different legal obligations. As artificial intelligence governance matures, the ability to demonstrate responsible data practice will increasingly become a competitive advantage rather than a pure compliance burden.

## **8. Sustainability and Life Cycle Considerations**

Machine learning-enabled process-structure-property optimization can contribute to sustainability in ceramic additive manufacturing by reducing material waste, energy consumption, and experimental iterations. Traditional ceramic process development often relies on extensive trial-and-error, consuming powders, energy, and time. Predictive models can guide experiments toward promising regions of the parameter space, reducing the number of failed builds and the associated environmental footprint. In addition, optimizing sintering schedules through data-driven prediction can lower kiln energy use because over-sintering and repeated thermal cycles are avoided. These sustainability benefits align with broader life cycle assessments of additive manufacturing, which show that environmental impacts depend heavily on process energy intensity, material efficiency, and part failure rates [22].

However, the sustainability impacts of machine learning itself must also be considered. Training large models consumes computational energy, and the deployment of monitoring systems may require additional sensors and network infrastructure. For the relatively small datasets typical of ceramic additive manufacturing, the energy cost of training is generally modest compared with large language models, but it is not negligible in aggregate across many institutions. Moreover, the data infrastructure described earlier requires storage, processing, and networking resources that have their own carbon footprint. A comprehensive sustainability evaluation should therefore include both the avoided costs from reduced physical experiments and the added costs of digital infrastructure. System architects can reduce these burdens through efficient model selection, avoiding unnecessary retraining, and using shared computational facilities.

Sustainability also intersects with fairness and equity. If advanced machine learning tools are available only to well-funded laboratories or large firms, smaller manufacturers may be further marginalized. Open-source toolkits, shared data repositories, and low-cost sensing platforms can reduce this gap. However, open sharing must be balanced against proprietary interests and the risk of misuse. Governance mechanisms that support selective sharing, such as community data trusts or escrowed test benches, can enable broader participation without undermining commercial incentives. The policy challenge is to create conditions under which the sustainability benefits of machine learning-enabled ceramic optimization are widely distributed rather than concentrated.

Finally, the circular economy perspective encourages the design of ceramic parts and processes for reuse and recycling. Machine learning can support this by predicting which feedstock compositions are most tolerant to recycled powder variability or by optimizing debinding and sintering cycles for mixed-source materials. This is an emerging area where process-structure-property models must incorporate new inputs related to powder history and contamination levels. Because recycled ceramics may exhibit bimodal particle distributions or altered surface chemistry, models trained on virgin powders may fail. Systems that include provenance-tracking and adaptive retraining will be more capable of handling post-consumer or post-industrial ceramic streams. The integration of sustainability and robustness concerns thus reinforces the need for a holistic, systems-level approach.

## **9. Conclusion**

Machine learning-enabled process-structure-property optimization holds significant promise for ceramic additive manufacturing, but its value cannot be realized through isolated predictive models alone. The complexity of ceramic processing, from slurry rheology and layer formation to debinding and sintering, requires systems that integrate heterogeneous data, multiscale models, and staged decision-making. This paper has argued that a systems perspective is essential for addressing the architectural, infrastructural, and governance challenges that determine whether machine learning tools become reliable in practice. By distinguishing among physics-informed and data-driven architectures, centralized and distributed data systems, and interpretable and opaque models, the analysis has shown that structural trade-offs must be made explicit rather than assumed.

The discussion has also emphasized that robustness and uncertainty are not secondary properties but central requirements. Process drift, rare defect modes, and distribution shifts make continuous monitoring and adaptive retraining necessary. At the same time, governance mechanisms must ensure accountability, fairness, and compliance with emerging artificial intelligence regulations. Ceramic components often operate in safety-critical applications, so model documentation, uncertainty reporting, and human oversight must be integrated into quality systems. The required reference to tape casting of barium titanate films illustrates that similar process-structure-property reasoning applies across ceramic manufacturing platforms, reinforcing the need for transferable data structures and shared modeling concepts [7].

Looking forward, the field would benefit from greater investment in open benchmark datasets, interoperable metadata standards, and precompetitive data-sharing agreements. Federated learning and active learning can help address data scarcity without undermining confidentiality. Policy frameworks should recognize the specific characteristics of advanced manufacturing, where data are generated in physical production environments rather than purely digital contexts. By treating machine learning as a socio-technical infrastructure, researchers and practitioners can build systems that are not only accurate but also robust, fair,

and sustainable. Such an approach will accelerate the responsible deployment of ceramic additive manufacturing across energy, biomedical, aerospace, and microelectronic applications.

## References

1. Chen, Z., Li, Z., Li, J., Liu, C., Lao, C., Fu, Y., Liu, C., Li, Y., Wang, P., & He, Y. (2019). 3D printing of ceramics: A review. *Journal of the European Ceramic Society*, 39(4), 661–687.
2. 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.
3. 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.
4. Bock, F. E., Aydin, R. C., Huber, N., & Pahr, D. H. (2019). A review on machine learning methods for defect detection in additively manufactured parts. *Additive Manufacturing*, 29, 100798.
5. Wang, C., Tan, X. P., Tor, S. B., & Lim, C. S. (2020). Machine learning in additive manufacturing: State-of-the-art and perspectives. *Additive Manufacturing*, 36, 101538.
6. Rueschhoff, L., Costakis, W., Michie, M., Youngblood, J., & Trice, R. (2016). Additive manufacturing of dense ceramic parts via direct ink writing of aqueous alumina suspensions. *Journal of the American Ceramic Society*, 99(12), 3927–3935.
7. Luan, S., He, H., Zhang, L., Yu, J., Cao, X., Wang, J., ... & Sun, R. (2025). Fabrication of High Performance BaTiO<sub>3</sub> Films via Aqueous Tape Casting for ultrathin MLCC applications. *Ceramics International*.
8. Lewis, J. A. (2006). Direct ink writing of 3D functional materials. *Advanced Functional Materials*, 16(17), 2193–2204.
9. Travitzky, N., Bonet, A., Dermeik, B., Fey, T., Filbert-Demut, I., Schlier, L., Schlördt, T., & Greil, P. (2014). Additive manufacturing of ceramic-based materials. *Advanced Engineering Materials*, 16(6), 729–754.
10. Qi, X., Chen, G., Li, Y., Cheng, X., & Li, C. (2019). Applying neural-network-based machine learning to additive manufacturing: Current applications, challenges, and future perspectives. *Engineering*, 5(4), 721–729.
11. Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422–440.
12. Himanen, L., Geurts, A., Foster, A. S., & Rinke, P. (2019). Data-driven materials science: Status, challenges, and perspectives. *Advanced Science*, 6(21), 1900808.
13. Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J. W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., ... Mons, B. (2016). The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3(1), 160018.

14. DebRoy, T., Wei, H. L., Zuback, J. S., Mukherjee, T., Elmer, J. W., Milewski, J. O., Beese, A. M., Wilson-Heid, A., De, A., & Zhang, W. (2018). Additive manufacturing of metallic components – Process, structure and properties. *Progress in Materials Science*, 92, 112–224.
15. 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.
16. Molnar, C. (2020). *Interpretable machine learning: A guide for making black box models explainable*. Lean Publishing.
17. 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. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220–229). ACM.
18. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 1–37.
19. Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
20. Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the EU AI Act. *Computer Law & Security Review*, 43, 105662.
21. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21.
22. Kellens, K., Mertens, R., Paraskevas, D., Dewulf, W., & Duflou, J. R. (2017). Environmental impact of additive manufacturing processes: Does AM contribute to a more sustainable way of part manufacturing? *Procedia CIRP*, 61, 582–587.
23. Ziatdinov, M., Dyck, O., Maksov, A., Li, X., Sang, X., Xiao, K., Unocic, R. R., Vasudevan, R., Jesse, S., & Kalinin, S. V. (2017). Deep learning of atomically resolved scanning transmission electron microscopy images: Chemical identification and tracking local transformations. *ACS Nano*, 11(12), 12742–12752.