

Generative Deep Learning-Guided Control of Non-Classical Colloidal Crystallization for Programmable Functional Materials and Energy Applications

Longfan Lon

Department of Computer Science, George Mason University, Fairfax, VA, USA.
longfangl@gmu.edu

Yiminhui Wei

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
yiminhuiw@uab.edu

Darren Sims

Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
darren.sims472@missouri.edu

Abstract

The programmable assembly of colloidal building blocks into crystalline architectures with tailored optical, mechanical, and electrochemical properties holds transformative potential for functional materials and energy technologies. Traditional crystallization paradigms emphasize classical nucleation and growth, but a growing body of evidence reveals non-classical pathways involving multi-step particle attachment, transient intermediate phases, and complex free-energy landscapes that elude simple thermodynamic control. Generative deep learning offers a new approach to navigate these landscapes by learning latent representations of crystallization trajectories and synthesizing control policies that drive the system toward desired polymorphs and microstructures. This paper presents a system-level framework for generative deep learning-guided control of non-classical colloidal crystallization, with emphasis on architectural design, infrastructure requirements, governance, robustness, and sustainability. We examine the structural trade-offs inherent in coupling high-dimensional perception, stochastic generative models, and real-time feedback control. The integration of in situ microscopy, physics-informed neural surrogates, and reinforcement learning loops is analyzed in terms of data fidelity, latency, and resilience. Governance and fairness issues arise from the potential for biased training datasets that overlook minority building blocks or environmentally suboptimal pathways, as well as dual-use concerns when programmable crystallization is exploited for harmful material fabrication. Deployment scalability demands modular cyber-physical instrumentation, secure data pipelines, and adaptive model updating that account for drift in colloidal suspensions. Through cross-domain comparisons with chemical synthesis and autonomous laboratories, we identify key policy implications for open data sharing, safety verification, and equitable access to programmable matter technologies. The paper concludes by outlining a roadmap toward energy-relevant applications including photonic crystals, battery electrodes, and catalytic supports, where generative control of non-classical crystallization can enhance efficiency, reduce waste, and accelerate materials discovery in alignment with sustainable development goals.

Keywords

non-classical crystallization, colloidal self-assembly, generative deep learning, programmable materials, inverse design, cyber-physical systems, energy materials, materials governance, sustainability.

1. Introduction

The capacity to direct the organization of colloidal particles into periodic lattices with custom symmetry and composition underpins many frontier applications in photonics, energy storage, catalysis, and advanced sensing. Colloidal crystallization, however, rarely follows the idealized pathways described by classical nucleation theory. Instead, assembly frequently proceeds through a sequence of intermediate, often metastable states in which particles attach, rearrange, and undergo phase transitions that are profoundly sensitive to local chemical environments, hydrodynamic fields, and surface forces. This non-classical behavior manifests as cycles of aggregation, densification, and crystallographic reorientation that cannot be captured by monotonic free-energy functions [1,2]. The recognition that prenucleation clusters, amorphous precursors, and particle-attachment events constitute essential stages of crystallization has reshaped the design of inorganic and hybrid materials [3]. In the domain of colloids, shape and interaction anisotropy enable lock-and-key recognition and directional bonding that amplify the diversity of accessible crystalline structures, making colloidal systems ideal testbeds for programmable matter [4]. Meanwhile, the energy sector has seen analogous non-classical behavior in electrochemical processes, where two-step oxidation pathways and intermediate electronic configurations trigger catalytic water oxidation under operando conditions, demonstrating that non-classical mechanisms are not a curiosity but a functional design principle [5].

Concurrently, machine learning has permeated materials science, offering data-driven predictions of phase stability, synthesis recipes, and property optimization. Early efforts established that deep neural networks could interpolate across vast composition–structure–property spaces, while generative models such as variational autoencoders and generative adversarial networks have enabled inverse design of molecules and crystals [6,7,8]. More recently, Boltzmann generators have demonstrated that thermodynamic sampling of complex many-body systems can be accelerated by learning invertible transformations between configurational space and latent representations [9]. These foundational generative architectures [10,11] have set the stage for moving beyond static prediction toward dynamic control of non-equilibrium assembly processes.

However, translating generative models into actionable control commands for colloidal crystallization requires a holistic socio-technical systems perspective. The design and deployment of such control loops involve not only algorithmic innovation but also decisions about data governance, infrastructure scalability, and fairness in prioritizing which materials are explored and for whose benefit. The dual role of artificial intelligence as both an enabler of sustainable development and a potential source of environmental burden and inequity has been articulated in policy frameworks that call for responsible innovation [12,13]. A generative deep learning system that decides which crystallization pathways to pursue must be assessed on criteria beyond yield or lattice perfection, including energy consumption of the computation, interpretability of the control policy, and the equitable distribution of resulting technological gains.

This article develops a system-level architecture for generative deep learning-guided control of non-classical colloidal crystallization, structured to address the entire lifecycle from data acquisition through governance and energy deployment. We do not present equations or algorithmic pseudocode; rather, we analyze structural trade-offs, infrastructure requirements, fairness and robustness dimensions, and policy implications that arise when advanced machine intelligence is woven into the fabric of materials manufacturing.

2. Non-Classical Crystallization in Colloidal Systems and the Need for Intelligent Control

Classical nucleation theory portrays crystallization as a single-barrier activated process in which a critical nucleus emerges stochastically and then grows by monomer addition. In colloidal suspensions, however, interparticle potentials are tunable over a wide range, and the relaxation timescales of density fluctuations often exceed the characteristic times for particle reconfiguration, giving rise to dense liquid precursors, amorphous clusters, and gel-like networks that subsequently reorganize into ordered lattices. Experimental characterization using in situ scattering and electron microscopy has revealed multi-step pathways where amorphous-to-crystalline transitions occur via oriented attachment, mesocrystal formation, and particle-mediated ripening [2]. The free-energy landscape of such systems is rugged and reconfigurable by external stimuli such as temperature, ionic strength, electric fields, or depletion forces, making the control problem intrinsically high-dimensional and history-dependent.

The energy materials community has independently identified the importance of non-classical pathways. For example, operando X-ray spectroscopy has uncovered a two-step oxidation sequence that stabilizes Zhang-Rice singlet states in cobalt-based catalysts, triggering efficient water oxidation through a transient electronic structure not predicted by equilibrium thermodynamics [5]. This observation reinforces the notion that intermediate states, far from being obstacles, can be exploited to achieve desired functionalities if the kinetic trajectory is properly guided. In colloidal crystallization for photonic bandgap materials or battery electrodes, analogous intermediate phases may enable the formation of defect-suppressed, large-area single crystals or hierarchical porosity that enhances ion transport. Therefore, the control challenge is not merely to suppress non-classical effects but to channel them toward targeted outcomes.

Conventional feedback control, even when augmented with microscopic imaging, struggles to manage the combinatorial explosion of possible trajectories. Recent experimental advances have, however, provided unprecedented direct observation and control of non-classical crystallization pathways in binary colloidal systems, yielding time-resolved, real-space data that capture the sequence of attachment, relaxation, and phase selection [14]. These observations confirm that small perturbations in local density can steer the system toward distinct polymorphs, but they also underscore the enormous dimensionality of the control space. The data streams generated by such experiments are precisely the kind that can feed deep generative models, enabling the extraction of low-dimensional reaction coordinates that compress the complexity of the free-energy landscape into learned latent variables.

3. Generative Deep Learning Framework and Control Architecture

A generative deep learning architecture for colloidal crystallization control must integrate perception, representation learning, policy generation, and actuation in a closed loop. The perception layer ingests high-throughput in situ imaging and scattering data, which are

mapped via convolutional or graph neural encoders into a latent space that captures the essential order parameters of the colloidal configuration. A generative component, realized as a variational autoencoder, generative adversarial network, or normalizing flow, then learns to produce plausible next states or full trajectories conditioned on a desired target crystal phase and the current observed state. This enables the system to act as a virtual simulator that can be queried in real time to evaluate the outcomes of candidate actions, such as changing the temperature ramp rate, adjusting the ionic strength, or applying an electric field gradient.

The integration of generative models with reinforcement learning forms a model-based control loop where the policy network selects actions that maximize a reward function aligned with crystal quality, phase purity, or energy efficiency. Unlike direct reinforcement learning on physical hardware, which would be sample-inefficient and potentially catastrophic, the generative world model allows thousands of simulated rollouts per second, dramatically accelerating policy learning. This is conceptually analogous to the way Boltzmann generators enable sampling of equilibrium states for molecular systems [9], but adapted to the non-equilibrium, time-dependent setting of colloidal crystallization. The architecture must address the inherent stochasticity of non-classical pathways: multiple trajectories from the same initial condition may lead to different final crystals, so the generative model should capture the full conditional distribution over futures, not just the most likely outcome. Quantifying this epistemic and aleatoric uncertainty is essential for safe exploration and for identifying when the model's predictions become unreliable due to out-of-distribution conditions [20].

Deploying such a framework at scale requires careful infrastructure design. Data streaming from electron microscopes or optical tomography must be compressed and transferred with low latency to GPU clusters, where models are trained continuously in a federated or incremental manner to prevent catastrophic forgetting when colloidal formulations change. Edge computing nodes located close to the synthesis chambers can run lightweight inference policies, while full model retraining occurs in the cloud. This separation introduces trade-offs between inference speed and model fidelity. Moreover, the generative models themselves must be auditable: regulatory frameworks increasingly demand that decisions affecting material safety or environmental release be explainable, which motivates the use of interpretable latent spaces and attention mechanisms that highlight which particle configurations most influence the policy's choices.

4. Governance, Fairness, and Societal Implications

The power to program colloidal crystallization through generative deep learning raises profound governance questions. Training datasets are inevitably biased by the history of materials research, which has favored certain compositions, particle shapes, and synthesis protocols over others, often due to economic incentives or geopolitical priorities. A generative model trained predominantly on data from spherical polymer colloids will exhibit degraded performance when transferred to anisotropic metal-organic framework particles or bio-derived nanocrystals, potentially excluding these sustainable materials from the accelerated discovery pipeline. Fairness in the context of materials informatics means ensuring that generative models do not perpetuate a narrow focus on high-performance but environmentally damaging materials; instead, they should be designed to explore diverse regions of the design space, including those aligned with green chemistry principles.

Algorithmic bias in generative models can manifest in subtle ways. If the reward function for the reinforcement learning policy overweights crystallinity metrics while neglecting energy consumption or toxicity of the intervening chemical agents, the system will converge on

pathways that are efficient for the laboratory but unsustainable at industrial scale. Surveys of bias and fairness in machine learning have documented analogous problems in social domains, where optimizing a single metric without constraints leads to unintended discrimination [15]. The same rigor applied to fairness-aware machine learning must be brought to bear on materials control, requiring multi-objective design that includes environmental and social indicators from the outset.

Generative control of crystallization also carries dual-use potential. The capability to rapidly design and fabricate colloidal crystals with precise optical or catalytic properties could be repurposed to create components for chemical weapons, surveillance devices, or unauthorized high-energy-density materials. Foresight analyses of malicious uses of artificial intelligence have emphasized that dual-use risks are amplified when the generative model can hide dangerous design rules within an opaque latent representation, making detection difficult [16]. Security measures, such as anomaly detection in generated crystal structures, access controls on cloud-based generative model APIs, and integration with material safety databases, must be incorporated into the system architecture [17]. Furthermore, international standards for tracking provenance of machine-learning-designed materials could deter misuse and support accountability.

The policy landscape for autonomous materials laboratories is still nascent. Existing ethical frameworks for artificial intelligence, such as those emphasizing beneficence, non-maleficence, autonomy, justice, and explicability, provide a starting point but are insufficiently specific to the physical synthesis domain [13]. Governance mechanisms should include mandatory safety case documentation before a generative model is authorized to drive physical experiments at scale, continuous monitoring of out-of-distribution behavior that might indicate model drift or adversarial manipulation, and transparent reporting of the model's training data and algorithmic methodology to enable peer review and reproducibility. Equity considerations also extend to access: if only a few well-resourced institutions can build the necessary cyber-physical infrastructure, the benefits of programmable colloidal materials will be concentrated, exacerbating global divides in energy technology and healthcare diagnostics that rely on these materials.

5. Robustness, Scalability, and Deployment Infrastructure

Real-world deployment of generative deep learning-guided crystallization presents a cascade of systems engineering challenges. The physical environment of colloidal synthesis is subject to stochastic fluctuations in temperature, humidity, and impurity levels that are not fully captured in the training data. The generative model must be robust to these covariate shifts, and the control policy must degrade gracefully rather than issuing destabilizing commands. Robustness can be enhanced through domain randomization during training, where the simulated environment is intentionally perturbed, and through online adaptation where local models are fine-tuned with a small amount of fresh data collected from the specific reactor. Such adaptive mechanisms introduce feedback loops between the learning algorithm and the physical system that can, if not carefully regulated, lead to oscillatory behavior analogous to control instabilities in adaptive optics.

The computational cost of training large generative models is substantial and can undercut the very sustainability goals that programmable materials aim to advance. Trade-offs arise between model capacity and inference efficiency, and between cloud-based centralized training and edge-based distributed learning. A promising direction is the use of physics-informed embeddings that incorporate known conservation laws and symmetries, reducing the

amount of data needed and improving generalization. In soft matter systems, machine-learned descriptors have been shown to capture phase behavior and rheological properties, providing a blueprint for building data-efficient surrogates that accelerate exploration without excessive energy footprints [21]. Moreover, the deployment of autonomous self-driving laboratories must be accompanied by life-cycle assessments that account for the carbon footprint of the entire digital infrastructure, from GPU manufacturing to cooling and data transmission.

Scalability to industrial production volumes demands integration with continuous flow reactors and roll-to-roll processing. Here, the control challenge multiplies because the colloidal crystallization must be maintained homogeneously across a moving substrate or a large reactor volume. The generative model must generate spatiotemporally resolved control fields that account for spatial gradients and residence time distributions. This requires a hierarchical control architecture where a high-level generative planner sets overall process recipes, while low-level distributed controllers handle local fluctuations using model-predictive strategies. Reliability of such systems over extended operation demands fault-tolerant data pipelines, sensor fusion, and stringent cybersecurity protocols to prevent adversarial injection of false data that could drive the process into unsafe regimes. The safety of operators and the environment further mandates hardware interlocks and kill switches independent of the learning system.

6. Energy Applications and Sustainability

Colloidal crystals direct light, store energy, and catalyze reactions with efficiencies that are exquisitely sensitive to lattice perfection, defect chemistry, and crystallographic orientation. In photonic energy applications, three-dimensional colloidal photonic crystals can serve as broadband reflectors, radiative cooling coatings, or light-trapping scaffolds for solar cells. Achieving the required large-area, defect-free assemblies remains a major manufacturing hurdle, one that non-classical crystallization guided by generative models could address by steering the assembly through intermediate fluidic states that anneal out dislocations before final solidification. Similarly, in battery technology, colloidal crystal templates can be used to create hierarchical porous electrodes with enhanced ion transport, where the optimal pore architecture arises from a non-classical sequence of co-assembly and template removal [22]. The generative control framework can optimize the entire multi-step thermal and chemical treatment protocol rather than each step in isolation, balancing porosity against mechanical integrity.

The sustainability implications of such controlled crystallization are multifaceted. By reducing the number of failed synthesis attempts and material waste, generative models can minimize the consumption of precursors and solvents, many of which are toxic or resource-intensive. Furthermore, the ability to engineer metastable polymorphs that are kinetically trapped but functionally superior can unlock performance levels that reduce the amount of material needed, as in catalysts with higher turnover frequencies. However, these gains must be weighed against the resource demands of the computational infrastructure itself. A thorough systems analysis should ensure that the overall environmental balance is positive, aligning with the targets set forth by sustainable development frameworks [12]. Energy-efficient model architectures, such as sparse mixture-of-experts models or quantized neural networks, can mitigate the carbon footprint while maintaining sufficient accuracy for control tasks.

A forward-looking energy application is the programmable assembly of colloidal quantum dot superlattices for next-generation photovoltaics and light-emitting diodes. Non-classical

attachment pathways can produce epitaxially fused superlattices with coherent electronic bands, a process that is notoriously difficult to control uniformly. Generative models that learn from in situ grazing-incidence scattering data could predict the optimal temporal profile of solvent evaporation and ligand exchange, ensuring that the quantum dots fuse only along desired crystallographic directions. The potential energy savings from such precision manufacturing extend beyond the device fabrication stage to the operation of lighting and energy harvesting systems that rely on these superlattices.

7. Conclusion

Generative deep learning-guided control of non-classical colloidal crystallization sits at the intersection of statistical physics, artificial intelligence, robotics, and materials engineering. The complexity of non-classical pathways, while initially appearing as an obstacle to rational design, can be reframed as a richness that generative models are uniquely suited to harness. By learning latent representations of crystallization trajectories and coupling them with model-based reinforcement learning, it becomes possible to direct colloidal assembly toward targeted crystalline products with minimal human intervention. This article has approached the topic not as a narrow algorithmic problem but as a full-stack system design challenge that demands attention to architecture, infrastructure, governance, robustness, and sustainability.

Structural trade-offs abound: the choice between high-capacity generative models and lightweight edge-deployable policies; the balance between exploring novel morphologies and exploiting known recipes; the tension between data assimilation fidelity and the latency constraints of real-time control. Addressing these trade-offs requires a convergence of expertise from computer science, chemical engineering, and social science. Governance frameworks must be developed proactively to ensure fairness in the design space, transparency in model operation, and safeguards against dual-use risks. The integration of life-cycle sustainability metrics into the reward functions of generative controllers can align the technology with global environmental goals rather than undermining them.

The research landscape is rapidly evolving. Direct observation of non-classical crystallization, as enabled by the latest microscopy techniques, provides the empirical foundation on which generative models can be trained and validated. As these data sets grow, the vision of a closed-loop autonomous platform for programmable materials moves closer to reality. The long-term impact will be measured not only by the exotic crystals produced but by the degree to which the knowledge generated is shared equitably and used to address pressing energy and environmental challenges. Realizing this vision calls for interdisciplinary collaboration, open data initiatives, and a sustained commitment to the responsible stewardship of intelligent manufacturing technologies.

References

1. Whitesides, G. M., & Grzybowski, B. (2002). Self-assembly at all scales. *Science*, 295(5564), 2418–2421.
2. De Yoreo, J. J., Gilbert, P. U. P. A., Sommerdijk, N. A. J. M., Penn, R. L., Whitlam, S., Joester, D., Zhang, H., Rimer, J. D., Navrotsky, A., Banfield, J. F., Wallace, A. F., Michel, F. M., Meldrum, F. C., Cölfen, H., & Dove, P. M. (2015). Crystallization by particle attachment in synthetic, biogenic, and geologic environments. *Science*, 349(6247), aaa6760.

3. Gebauer, D., & Cölfen, H. (2011). Prenucleation clusters and non-classical nucleation. *Nano Today*, 6(6), 564–584.
4. Aizenberg, J., Muller, D. A., Grazul, J. L., & Hamann, D. R. (2003). Direct fabrication of large micropatterned single crystals. *Science*, 299(5610), 1205–1208.
5. Peng, C. K.; Lin, Y. C.; Chiang, C. L.; Qian, Z. X.; Huang, Y. C.; Dong, C. L.; Li, J. F.; Chen, C. T.; Hu, Z. W.; Chen, S. Y. et al. Zhang-Rice singlets state formed by two-step oxidation for triggering water oxidation under operando conditions. *Nat. Commun.* 2023, 14, 529.
6. 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.
7. Sanchez-Lengeling, B., & Aspuru-Guzik, A. (2018). Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400), 360–365.
8. Gomez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernandez-Lobato, J. M., Sanchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T. D., Adams, R. P., & Aspuru-Guzik, A. (2018). Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2), 268–276.
9. Noé, F., Olsson, S., Köhler, J., & Wu, H. (2019). Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457), eaaw1147.
10. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (Vol. 27).
11. Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.
12. Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M., & Nerini, F. F. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11, 233.
13. Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.
14. Zang, S., Paul, S., Leung, C. W., Chen, M. S., Hueckel, T., Hocky, G. M., & Sacanna, S. (2025). Direct observation and control of non-classical crystallization pathways in binary colloidal systems. *Nature communications*, 16(1), 3645.
15. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
16. Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitsoff, T., Filar, B., Anderson, H., Roff, H., Allen, G. C., Steinhardt, J., Flynn, C., hEigeartaigh, S. Ó., Beard, S., Belfield, H., Farquhar, S., ... Amodei, D. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv preprint arXiv:1802.07228*.

17. Li, Y., Li, J., & Su, H. (2021). When deep learning meets security: A survey. arXiv preprint arXiv:2105.14222.
18. Tibbetts, K. M., & Geissler, P. L. (2019). Generative neural networks for inverse design of self-assembling colloidal clusters. *The Journal of Chemical Physics*, 151(16), 164102.
19. Sacanna, S., Irvine, W. T. M., Chaikin, P. M., & Pine, D. J. (2010). Lock and key colloids. *Nature*, 464(7288), 575–578.
20. Nalisnick, E., Matsukawa, A., Teh, Y. W., Gorur, D., & Lakshminarayanan, B. (2019). Do deep generative models know what they don't know? In *International Conference on Learning Representations*.
21. Spaeth, J., Daoulas, I. S., & Müller, M. (2021). Machine learning for predicting properties and morphologies of soft matter. *Soft Matter*, 17, 7748–7767.
22. Lu, J., Chen, Z., Ma, Z., Pan, F., Curtiss, L. A., & Amine, K. (2016). The role of nanotechnology in the development of battery materials for electric vehicles. *Nature Nanotechnology*, 11(12), 1031–1038.