Efficient Probabilistic Modeling of Crystallization at Mesoscopic Scale

Pol Timmer
Eindhoven University of Technology
l.c.timmer@student.tue.nl

Koen Minartz
Eindhoven University of Technology
k.minartz@tue.nl

Vlado Menkovski
Eindhoven University of Technology
v.menkovski@tue.nl

Abstract

Crystallization processes at the mesoscopic scale, where faceted, dendritic growth, and multigrain formation can be observed, are of particular interest within materials science. These processes are highly nonlinear, stochastic, and sensitive to small perturbations of system parameters and initial conditions. Traditional numerical models of these systems are computationally demanding. To address this, we introduce the Crystal Growth Neural Emulator (CGNE), a machine learning emulator that efficiently models crystallization using autoregressive latent variable models. For our experiment on ice crystal growth, CGNE improves inference time by a factor of 11 compared to numerical simulations. To validate simulation quality, we compare morphological properties of crystals to those from numerical simulations, and find that CGNE substantially improves simulation fidelity and diversity over existing probabilistic models.

1 Introduction

Solidification is the process of a substance transitioning from a liquid or gaseous state to a solid state, which is evident in various systems, including the crystallization of metals and the growth of snow crystals from water vapor, illustrated in Figure 1. The interplay of highly sensitive nonequilibrium and nonlinear processes as well as stochastic effects results in complex structures emerging during the solidification process. As such, studying these processes not only has significant practical implications in fields like materials science and metallurgy [15], but can also help us to understand the mechanisms behind emergent complexity in general. In particular, this complexity manifests itself at the mesoscopic scale: while at the atomic scale one might see atoms arrange themselves into a well-organised lattice, and at the macroscopic scale alloys solidify to an ingot, intricate pattern formations caused by the polycrystalline solidification process arise at the mesoscopic scale.

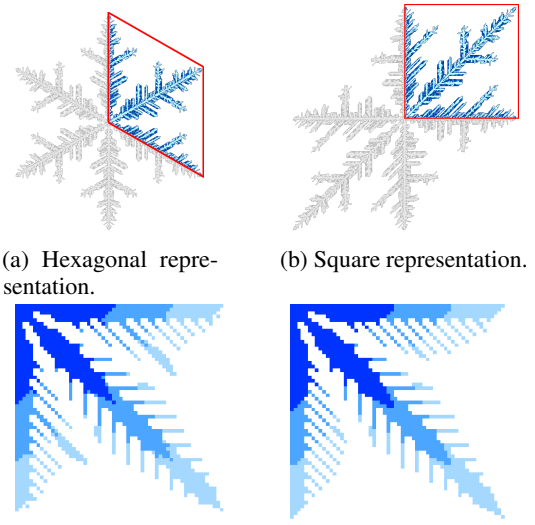

(a) Hexagonal representation.

(b) Square representation.

(c) Crystal growth simulation (LCA).

(d) Crystal growth simulation (CGNE).

Figure 1: Illustration of CGNE: The hexagonal domain is converted to a square grid. Exploiting the crystal's symmetries, only a quadrant of this grid is simulated.

D3S3: Data-driven and Differentiable Simulations, Surrogates, and Solvers @ NeurIPS 2024.

Numerical simulation is a powerful tool to enhance our understanding of these processes, and various methods modeling the stochastic growth of complex solidification processes have been proposed. Front-tracking [36, 22] and phase-field [18, 4, 13, 3] have provided valuable insights into multigrain and dendritic growth simulations, but occasionally struggle with dynamic and numerical instabilities, especially in cases with both faceted and dendritic growth, such as in snow crystals [1, 23]. More recently, local cellular automaton (LCA) based models [10, 29, 11, 23, 12, 19, 33] have advanced the field, capturing previously unseen morphological features of such growth [24]. However, these models demand high computational resources due to the fine spatiotemporal resolution required to uphold simulation accuracy in the sensitive crystallization process.

Recently, scaling numerical simulations with neural networks has gained traction in the scientific machine learning community, with successful applications in climate science [26, 8, 5], nuclear fusion [28, 9], and fluid dynamics [14, 21]. Similarly, various neural approaches have been proposed for crystal growth modeling. One branch focuses on Molecular Dynamics (MD) simulations of crystallization [37, 17]. While such MD simulations that operate at the microscopic scale have impressively scaled up to 1 million atoms, a simulation at the mesoscopic scale would require many orders of magnitude more atoms to be simulated. Other approaches accelerate phase field simulations directly at the mesoscopic scale [30, 27, 34]. However, these methods are deterministic and unsuitable for stochastic crystallization processes. More broadly, probabilistic neural simulation methods have been proposed, but these do not straightforwardly apply to the discrete nature of the crystal growth problem [7, 35, 2], or fail to reliably capture the sensitive and stochastic relationship between environmental parameters and crystal morphology [25]. As such, no effective method exists for efficient probabilistic crystal growth simulation at the mesoscopic scale. In this work, we introduce a neural simulator model that greatly accelerates the simulation of mesoscopic scale crystallization processes. Our main contributions are as follows:

- We propose the Crystal Growth Neural Emulator (CGNE)[1], an autoregressive probabilistic model for simulating crystal growth processes at the mesoscopic scale.
- We observe that existing training techniques fail to enable the model to reliably capture the joint distribution over environmental parameters and crystal morphologies. We diagnose Latent Variable Neglect (LVN) as the root cause of the problem and introduce a technique, called samplewise decoder dropout, that systematically prevents LVN.
- We construct a new snow crystal growth dataset that captures the intricate relationships between environmental parameters and crystal morphology.
- Our evaluations show that CGNE outperforms a state-of-the-art probabilistic neural emulator [25]. Further, we show that decoder dropout significantly enhances simulation diversity and sample quality by preventing LVN, offering broader implications for neural simulation.

2 Approach

At time t, the system state is denoted as (x_t, y), where x_t is a time-varying, binary function on a grid and y is a real-valued vector of static environmental parameters. The system's evolution is represented by the sequence $(x_{0:T}, y)$. We aim to fit a simulation model p by maximizing the log-likelihood for the training set $\{(x_{0:T}, y)_i\}_{i=1}^N$. Given the Markovian nature of LCA-based numerical simulators and the known distributions of initial conditions and environmental parameters, our task simplifies to optimizing the autoregressive formulation:

$$\frac{1}{N}\sum_{i=1}^{N} \log p\left((x_{0:T}, y)_i\right) = \frac{1}{N}\sum_{i=1}^{N}\sum_{t=0}^{T-1} \log p\left((x_{t+1} \mid x_t, y)_i\right). \tag{1}$$

After training on Equation 1, we can generate new samples by applying $p\left((x_{t+1} \mid x_t, y)_i\right)$ autoregressively.

Model Design. We aim to develop a model capable of producing realistic simulation trajectories with good sample efficiency. We base our model on the PNS framework [25], which employs a Conditional Variational Autoencoder (CVAE) structure [32]. The model's prior is conditioned on the current state x_t, while the posterior considers both the current and next states x_t, x_{t+1}. Unlike PNS, our model also integrates environmental parameters y into both the prior and posterior, resulting in a conditional prior $p(z \mid x_t, y)$ and a posterior $q(z \mid x_t, x_{t+1}, y)$, as shown in Figure 2. CGNE is

[1] Source code & dataset: https://github.com/poltimmer/CGNE

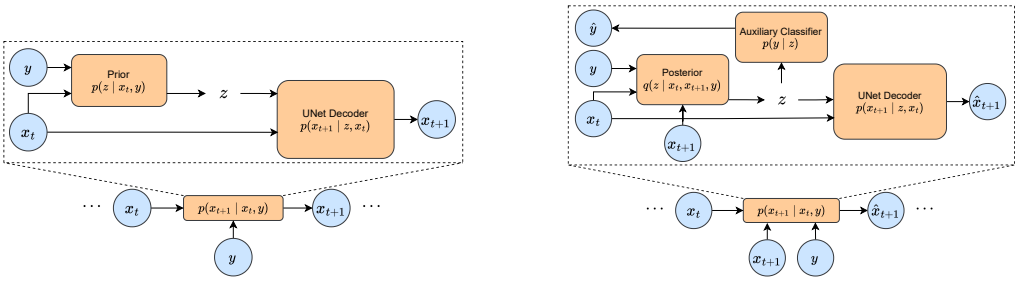

(a) Sample pass of the model.

(b) Training pass of the model.

Figure 2: CGNE model overview.

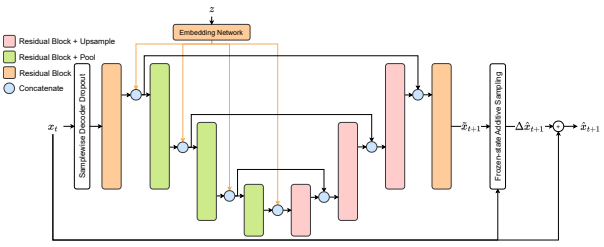

Figure 3: Overview of CGNE's decoder, including samplewise decoder dropout and frozen-state additive sampling.

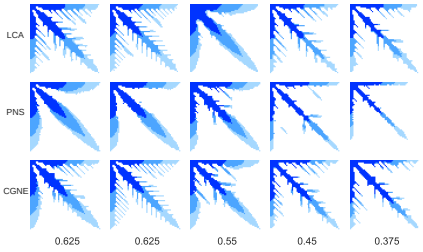

Figure 4: Qualitative comparison of simulated trajectories for different ρ values.

fully convolutional, preserving the data domain's local translation symmetries. This includes a UNet decoder [31] and a spatial latent space, detailed in Figure 3. The latent variable z passes through an embedding network before concatenation with the input of each residual block in the decoder. As we only consider solidification, as opposed to melting and sublimation, we introduce frozen-state additive sampling to guarantee monotonic crystal growth. This method preserves the solidified portion from the previous step and merges it with predictions for only the non-solidified areas of the prior step, ensuring physical accuracy. An auxiliary classifier $p(y \mid z)$ reconstructs environmental parameters y from the latent variable during training, improving their representation in the latent space [16].

Training Strategies. As the decoder is conditioned on the previous simulation state x_t, and differences between consecutive states (x_t and x_{t+1}) are incremental, inferring a reasonable estimate for x_{t+1} is possible to a large degree without using z. Consequently, gradients flowing through z, which have a low signal-to-noise ratio, are overshadowed by the gradients flowing through the direct path to x_t. This is amplified by the UNet's skip connections, causing the model to neglect the latent variables. Although the one-step-ahead estimates of such a deterministic model might be reasonable, the lack of accumulation of randomness results in a severe underestimation of sample diversity at the end of the trajectory. While with posterior collapse the posterior $q(z \mid x_t, x_{t+1}, y)$ collapses to the prior $p(z \mid x_t, y)$, with LVN the environmental parameter y is also ignored, causing the conditional prior to be uninformative.

To counteract this, we implement samplewise dropout from x_t to the decoder, similar to word dropout [6]. We anneal the dropout rate linearly from 0.9 to 0.1 during training, reinforcing the information flow through latent variables, and thereby enriching trajectory diversity. Further, we incorporate beta annealing [6] and the free bits trick [20] to prevent posterior collapse. Another issue arises when differences between subsequent states are gradual, in which case the model can struggle to learn the dynamics and instead tends to resort to the local minimum of an identity function. To improve simulation efficiency and prevent this failure mode, we downsample the temporal resolution, increasing the observed differences between subsequent simulation states.

Experiments. We validate our model on snow crystal growth, a complex crystallization process stemming from the combined effects of faceted and dendritic growth, using a dataset we generated via Janko Gravner and David Griffeath's stochastic LCA algorithm [11]. This algorithm simulates growth from a single seed crystal under seven tuneable environmental parameters. We fix six parameters and vary ρ, the ambient vapor saturation parameter sampled from a uniform distribution between 0.35 and 0.65, which significantly influences crystal morphology.

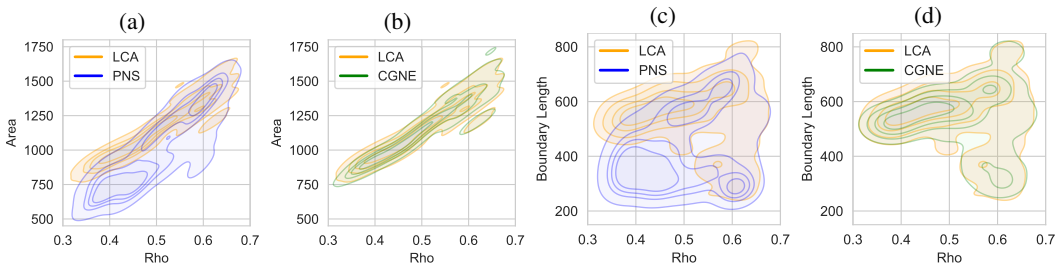

Figure 5: Joint distributions over morphological crystal properties and environmental parameters for PNS and CGNE, compared with the actual joint distribution produced by LCA.

Due to the stochastic nature of growth, crystals from the same parameters can vary drastically at the pixel level. Low ρ values produce thin crystals with few side branches, whereas high values yield thicker crystals with numerous branches; however, the branch attachment points can differ even with consistent ρ. This variability makes pixel-wise accuracy an inadequate metric for validation. Instead, we assess whether the distribution of morphological properties of snowflakes generated by CGNE matches that of the LCA simulator. Our primary evaluation metric is the expected value of the type-2 Wasserstein distance (Equation 2), between the joint distributions of snowflake area (a) and boundary length (b), conditioned on ρ, of the model ($p(a, b \mid \rho)$) and the LCA simulator ($p_{\text{LCA}}(a, b \mid \rho)$):

$$\text{EWD} = \mathbb{E}_{\rho \sim p(\rho)} \left[W_2 \left(p(a, b \mid \rho),\, p_{\text{LCA}}(a, b \mid \rho) \right) \right]. \tag{2}$$

EWD values are smaller when model distributions closely match the ground-truth distributions. Additionally, we use the Evidence Lower Bound (ELBO) as a secondary metric to gauge model fit against held-out test data. As a baseline, we take PNS [25], a recent state-of-the-art model for probabilistic simulation of dynamical systems.

3 Results

Figure 4 shows simulated snowflakes for five ρ values using CGNE, LCA, and PNS. LCA produces thinner, less voluminous snowflakes as ρ decreases, reflecting lower water vapor saturation. Notably, LCA can yield diverse snowflake morphologies for the same ρ, highlighting the stochastic nature of the ground truth process. PNS struggles to capture this diversity, and generates very similar snowflakes for the same ρ, evident at $\rho = 0.625$. This indicates a tendency towards a deterministic model, caused by the lack of samplewise decoder dropout leading to LVN. PNS also struggles with forming thin branches and finer details, often resulting in snowflakes lacking thin side branches and containing physically impossible floating crystal pieces for lower ρ values. Conversely, CGNE produces complete and realistic snowflakes across all ρ values, maintaining finer details without distortions.

Figure 5 compares the joint distributions over ρ, crystal area, and boundary length of PNS, CGNE, and LCA. CGNE's distributions align closely with LCA, capturing all relevant modes and the overall correlation (Figure 5b and Figure 5d). In contrast, PNS's distributions fail to show correct correlations and modes, particularly at low ρ values (Figure 5a and Figure 5c), where PNS struggles with branch formation. Quantitatively, CGNE shows a significantly better expected Wasserstein distance and ELBO compared to PNS (Table 1), indicating a more accurate probability distribution fit. CGNE also enhances inference efficiency, being **11x** faster than our GPU-accelerated LCA model on average. While LCA takes a median of **4.11s** per simulation trajectory, CGNE completes in just **0.36s**. This speed advantage is flexible and inversely proportional to temporal resolution, and thus can be even more significant depending on the application.

Table 1: Quantitative results: Expected Wasserstein Distance and the Evidence Lower Bound.

Method	EWD ↓	ELBO ↑
PNS	202.8	-0.0670
CGNE	**43.8**	**-0.0428**

4 Conclusions and Future Work

We introduce the CGNE architecture, significantly enhancing the speed and accuracy of probabilistic simulations for crystallization at the mesoscopic scale. We address Latent Variable Neglect, a common issue in such models, using samplewise decoder dropout, which markedly improves both the accuracy and diversity of simulations compared to PNS, a recent state-of-the-art model. For future research we consider implementing mechanisms to prevent floating crystal pieces, and expanding CGNE's application to datasets of other crystallization processes, such as in metallurgy.

References

[1] John W. Barrett, Harald Garcke, and Robert Nürnberg. Numerical computations of faceted pattern formation in snow crystal growth. *Physical Review E*, 86(1), jul 2012. URL https://doi.org/10.1103%2Fphysreve.86.011604.

[2] Federico Bergamin, Cristiana Diaconu, Aliaksandra Shysheya, Paris Perdikaris, José Miguel Hernández-Lobato, Richard E. Turner, and Emile Mathieu. Guided autoregressive diffusion models with applications to PDE simulation. In *ICLR 2024 Workshop on AI4DifferentialEquations In Science*, 2024. URL https://openreview.net/forum?id=1avNKFEIOL.

[3] Sudipta Biswas, Dehao Liu, Larry K Aagesen, and Wen Jiang. Solidification and grain formation in alloys: a 2d application of the grand-potential-based phase-field approach. *Modelling and Simulation in Materials Science and Engineering*, 30(2), January 2022. URL http://dx.doi.org/10.1088/1361-651X/ac46dc.

[4] W. J. Boettinger, J. A. Warren, C. Beckermann, and A. Karma. Phase-field simulation of solidification. *Annual Review of Materials Research*, 32(1), 2002. URL https://doi.org/10.1146/annurev.matsci.32.101901.155803.

[5] Boris Bonev, Thorsten Kurth, Christian Hundt, Jaideep Pathak, Maximilian Baust, Karthik Kashinath, and Anima Anandkumar. Spherical Fourier neural operators: Learning stable dynamics on the sphere. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, Jul 2023. URL https://proceedings.mlr.press/v202/bonev23a.html.

[6] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, August 2016. URL https://aclanthology.org/K16-1002.

[7] Salva Rühling Cachay, Bo Zhao, Hailey James, and Rose Yu. DYffusion: A dynamics-informed diffusion model for spatiotemporal forecasting. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=WRGldGm5Hz.

[8] Zhihan Gao, Xingjian Shi, Boran Han, Hao Wang, Xiaoyong Jin, Danielle C Maddix, Yi Zhu, Mu Li, and Bernie Wang. Prediff: Precipitation nowcasting with latent diffusion models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/pdf?id=Gh67ZZ6zkS.

[9] Vignesh Gopakumar, Stanislas Pamela, Lorenzo Zanisi, Zongyi Li, Ander Gray, Daniel Brennand, Nitesh Bhatia, Gregory Stathopoulos, Matt Kusner, Marc Peter Deisenroth, Anima Anandkumar, the JOREK Team, and MAST Team. Plasma surrogate modelling using fourier neural operators. *Nuclear Fusion*, 64(5), apr 2024. URL https://dx.doi.org/10.1088/1741-4326/ad313a.

[10] Janko Gravner and David Griffeath. Modeling snow crystal growth i: Rigorous results for packard's digital snowflakes. *Experimental Mathematics*, 15(4), 2006. URL https://doi.org/10.1080/10586458.2006.10128978.

[11] Janko Gravner and David Griffeath. Modeling snow crystal growth ii: A mesoscopic lattice map with plausible dynamics. *Physica D: Nonlinear Phenomena*, 237(3), 2008. URL https://www.sciencedirect.com/science/article/pii/S0167278907003387.

[12] Janko Gravner and David Griffeath. Modeling snow-crystal growth: A three-dimensional mesoscopic approach. *Phys. Rev. E*, 79, Jan 2009. URL https://link.aps.org/doi/10.1103/PhysRevE.79.011601.

[13] László Gránásy, Tamás Pusztai, and James A Warren. Modelling polycrystalline solidification using phase field theory. *Journal of Physics: Condensed Matter*, 16(41), oct 2004. URL https://dx.doi.org/10.1088/0953-8984/16/41/R01.

[14] Jayesh K. Gupta and Johannes Brandstetter. Towards multi-spatiotemporal-scale generalized pde modeling. *arXiv preprint*, 2022. URL https://arxiv.org/abs/2209.15616.

[15] Shuhada A. Idrus-Saidi, Jianbo Tang, Stephanie Lambie, Jialuo Han, Mohannad Mayyas, Mohammad B. Ghasemian, Francois-Marie Allioux, Shengxiang Cai, Pramod Koshy, Peyman Mostaghimi, Krista G. Steenbergen, Amanda S. Barnard, Torben Daeneke, Nicola Gaston, and

Kourosh Kalantar-Zadeh. Liquid metal synthesis solvents for metallic crystals. *Science*, 378 (6624), 2022. URL https://www.science.org/doi/abs/10.1126/science.abm2731.

[16] Maximilian Ilse, Jakub M. Tomczak, Christos Louizos, and Max Welling. Diva: Domain invariant variational autoencoders. In *Proceedings of the Third Conference on Medical Imaging with Deep Learning*, volume 121 of *Proceedings of Machine Learning Research*, Jul 2020. URL https://proceedings.mlr.press/v121/ilse20a.html.

[17] P Jarry and N Jakse. Medium range ordering in liquid al-based alloys: towards a machine learning approach of solidification. *IOP Conference Series: Materials Science and Engineering*, 1274(1), jan 2023. URL https://dx.doi.org/10.1088/1757-899X/1274/1/012001.

[18] Alain Karma and Wouter-Jan Rappel. Quantitative phase-field modeling of dendritic growth in two and three dimensions. *Phys. Rev. E*, 57, Apr 1998. URL https://link.aps.org/doi/10.1103/PhysRevE.57.4323.

[19] James G. Kelly and Everett C. Boyer. Physical improvements to a mesoscopic cellular automaton model for three-dimensional snow crystal growth, 2013. URL https://arxiv.org/abs/1308.4910.

[20] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems*, volume 29, 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/file/ddeebdeefdb7e7e7a697e1c3e3d8ef54-Paper.pdf.

[21] Zongyi Li, Nikola Borislavov Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Prakash Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, and Anima Anandkumar. Geometry-informed neural operator for large-scale 3d PDEs. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=86dXbqT5Ua.

[22] Kenneth G. Libbrecht. Cylindrically symmetric green's function approach for modeling the crystal growth morphology of ice. *Phys. Rev. E*, 60, Aug 1999. URL https://link.aps.org/doi/10.1103/PhysRevE.60.1967.

[23] Kenneth G. Libbrecht. Physically derived rules for simulating faceted crystal growth using cellular automata, 2008. URL https://arxiv.org/abs/0807.2616.

[24] Kenneth G. Libbrecht. *Snow Crystals: A Case Study in Spontaneous Structure Formation*. Princeton University Press, 2022. URL http://www.jstor.org/stable/j.ctv1qdqztv.

[25] Koen Minartz, Yoeri Poels, Simon Martinus Koop, and Vlado Menkovski. Equivariant neural simulators for stochastic spatiotemporal dynamics. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=CCVsGbhFdj.

[26] Tung Nguyen, Johannes Brandstetter, Ashish Kapoor, Jayesh K Gupta, and Aditya Grover. Climax: A foundation model for weather and climate. *arXiv preprint arXiv:2301.10343*, 2023. URL https://arxiv.org/abs/2301.10343.

[27] Sukeharu Nomoto, Hiroshi Wakameda, Masahito Segawa, Akinori Yamanaka, and Toshiyuki Koyama. Solidification analysis by non-equilibrium phase field model using thermodynamics data estimated by machine learning. *Modelling and Simulation in Materials Science and Engineering*, 27(8), oct 2019. URL https://dx.doi.org/10.1088/1361-651X/ab3379.

[28] Yoeri Poels, Gijs Derks, Egbert Westerhof, Koen Minartz, Sven Wiesen, and Vlado Menkovski. Fast dynamic 1d simulation of divertor plasmas with neural pde surrogates. *Nuclear Fusion*, 63(12), December 2023. URL https://iopscience.iop.org/article/10.1088/1741-4326/acf70d.

[29] S. Raghavan and Satyam S. Sahay. Modeling the grain growth kinetics by cellular automaton. *Materials Science and Engineering: A*, 445-446:203–209, 2007. URL https://www.sciencedirect.com/science/article/pii/S0921509306019836.

[30] Liangyuan Ren, Shaoning Geng, Ping Jiang, Song Gao, and Chu Han. Numerical simulation of dendritic growth during solidification process using multiphase-field model aided with machine learning method. *Calphad*, 78, 2022. URL https://www.sciencedirect.com/science/article/pii/S0364591622000554.

[31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24574-4.

[32] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, 2015. URL https://proceedings.neurips.cc/paper/2015/file/8d55a249e6baa5c06772297520da2051-Paper.pdf.

[33] Dmytro S. Svyetlichnyy. Modeling of grain refinement by cellular automata. *Computational Materials Science*, 77, 2013. URL https://www.sciencedirect.com/science/article/pii/S092702561300253X.

[34] Tianju Xue, Zhengtao Gan, Shuheng Liao, and Jian Cao. Physics-embedded graph network for accelerating phase-field simulation of microstructure evolution in additive manufacturing. *npj Computational Materials*, 8(1), September 2022. URL http://dx.doi.org/10.1038/s41524-022-00890-9.

[35] Gefan Yang and Stefan Sommer. A denoising diffusion model for fluid field prediction. *arXiv preprint*, 2023. URL https://arxiv.org/abs/2301.11661.

[36] Etsuro Yokoyama and Toshio Kuroda. Pattern formation in growth of snow crystals occurring in the surface kinetic process and the diffusion process. *Phys. Rev. A*, 41, Feb 1990. URL https://link.aps.org/doi/10.1103/PhysRevA.41.2038.

[37] Linshuang Zhang, Manyi Yang, Shiwei Zhang, and Haiyang Niu. Unveiling the crystallization mechanism of cadmium selenide via molecular dynamics simulation with machine-learning-based deep potential. *Journal of Materials Science & Technology*, 185:23–31, 2024. URL https://www.sciencedirect.com/science/article/pii/S1005030223009623.