

A parametrically-Conditioned Deep Learning Surrogate for Coherent Spinodal Decomposition

Original

A parametrically-Conditioned Deep Learning Surrogate for Coherent Spinodal Decomposition / Fantasia, A.; Lanzoni, D.; Di Eugenio, N.; Monteleone, A.; Bergamaschini, R.; Montalenti, F.. - In: ADVANCED THEORY AND SIMULATIONS. - ISSN 2513-0390. - 9:2(2026), pp. 1-14. [10.1002/adts.202502144]

Availability:

This version is available at: 11583/3011373 since: 2026-05-25T13:22:01Z

Publisher:

Wiley

Published

DOI:10.1002/adts.202502144

Terms of use:







This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

RESEARCH ARTICLE OPEN ACCESS

A parametrically-Conditioned Deep Learning Surrogate for Coherent Spinodal Decomposition

Andrea Fantasia¹  | Daniele Lanzoni^{1,2}  | Niccolò Di Eugenio¹  | Angelo Monteleone¹  | Roberto Bergamaschini¹  | Francesco Montalenti¹ 

¹Dept. of Materials Science, University of Milano-Bicocca, Milano, Italy | ²Dept. of Physics, University of Genova, Genova, Italy

Correspondence: Roberto Bergamaschini (roberto.bergamaschini@unimib.it)

Received: 6 November 2025 | **Revised:** 21 January 2026 | **Accepted:** 22 January 2026

Keywords: coherent spinodal decomposition | convolutional neural network | recurrent neural network | surrogate model

ABSTRACT

A machine-learning framework is designed to tackle the spinodal decomposition of coherently strained alloy microstructures, with cubic anisotropy in elastic constants. A database of phase-field simulations, leveraging a Green's function approximation of the elastic field, is constructed for randomly chosen composition fields and widely variable misfit strain, producing a continuum variation in the phase morphology from smooth to strongly anisotropic domains. A convolutional recurrent neural network is then trained to accurately predict the full time-evolution sequence under the explicit conditioning of the known misfit parameter, at a reduced computational cost. Extensive error analysis at the pixel level and for global descriptors is used to assess the model accuracy and evaluate its generalization capability on longer timescales and larger computational domains. The model returns a one-to-one match of the ground-truth simulations over the temporal range of training sequences. Moreover, it can reliably predict average behaviors for sequences several times longer the training ones, albeit losing one-to-one accordance. As a proof, the NN trained model is used for reconstructing the full phase diagram of the system, achieving a 98% degree of accuracy. The proposed framework is general and can be applied beyond the specific, prototypical system here considered, enabling high-throughput parametric studies.

1 | Introduction

Modeling microstructural features and their dynamics is key for materials science and engineering due to their fundamental impact on functional and mechanical properties [1, 2]. A full characterization of these aspects is generally demanding from an experimental point of view, especially when time-dependent properties are involved, making computational methods particularly useful. Indeed, a plethora of simulation approaches have been developed for decades on this topic, with phase-field (PF) methods being particularly successful [3]. Still, the problem remains challenging for several reasons. First, high-accuracy methods usually bring with them also high computational costs. Second, the connection between models and experiments often requires difficult, indirect comparisons.

Finally, understanding microstructure properties is inherently a multi-scale and multi-physics problem, thus involving many parameters and conditions that are hard to control and identify.

In recent years, ML and, most prominently, Neural Networks (NN), have revolutionized materials science, opening new ways for materials discovery [4], design [5–7], and property prediction [8–11]. ML-driven interatomic potentials changed the rules of atomistic modeling, delivering accuracy on par with first-principle approaches at a fraction of the computational costs [12–14]. Recent studies have also highlighted the potential for optimizing material fabrication processes through real-time feedback control [15], showcasing the potential synergy between ML and experimental techniques.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Advanced Theory and Simulations* published by Wiley-VCH GmbH

Concerning mesoscale dynamics of materials, NNs have demonstrated exceptional power for accelerating or fully surrogating conventional numerical approaches (for a recent review see Ref. [16]). This has been possible using tools imported from modern deep learning, e.g., based on dimensionality reduction [17–19], convolutional NNs [20–22], and time-series specific architectures such as recurrent networks [23–26]. Moreover, physics-specific novel methods such as Physics Informed Neural Networks [27, 28] have been developed. Once trained, these models allow for the prediction of time-evolutions of microstructures without resorting to the explicit solution of the underlying partial differential equations. Substantial speed-ups in the solution time, frequently reported in the order of $100 - 1000\times$ [16, 21, 23], can be achieved, thus promising the feasibility of high-throughput studies, eventually giving access to large-scale domains or long-times [23, 26, 29] that exceed the training dataset characteristics.

The typical workflow in these studies has been to train a NN model, capable of recognizing spatial- and time-correlations within suitably processed time-series of data, either pixelated images of the relevant fields or some latent-space representation of them, and then predict the evolution frames starting from a known state/microstructure. In the simplest approaches, the NN model is trained to learn the solution for an assigned parameter set and to predict the specific evolution sequence when changing the initial configuration only. More advanced studies extended the training over more varied datasets, comprising cases generated by the same dynamics but sampling a “broad” range of values of some constitutive parameters, responsible for qualitatively different outcomes. In such a case, the NN is expected to implicitly identify the leading parameters by processing a short initial evolution sequence and then continuing it at later stages [18, 23]. Importantly, the NN does not provide any estimation of the actual value of the conditioning parameters as they remain encoded in the latent representation of the input data in these approaches. A few recent works in fluid dynamics [30] and in materials growth from physical vapor deposition [31] implemented explicit conditioning by external parameters (or boundary conditions [21]) within different NN architectures to control the predictive process by supplying them in input along with the initial configuration. Our work fits into this trend, adapting a previously established Convolutional Recurrent Neural Networks (CRNN) method [26, 32] to explicitly include a conditioning parameter as an input for the network while leaving other factors implicitly inferred from the system configuration itself.

We here consider the phenomenon of spinodal decomposition, a second-order phase transition by which a homogeneous mixture spontaneously separates into two distinct phases after cooling below a critical temperature [33–35]. This mechanism plays a critical role in various systems of metallurgical interest, such as Al-Zn alloys [36, 37], Ni-based superalloys [38], in certain oxides [39, 40], and also in polymer blends [41]. Recent interest in spinodal behavior has also been reignited for the design and study of spinoid metamaterials [42–44]. In crystalline materials, the two emerging phases typically differ in molar volume, leading to the development of elastic strain during the transformation. This strain can give rise to characteristic microstructural patterns, particularly in systems with anisotropic or spatially inhomogeneous elastic properties [35, 45–47]. Representative examples of

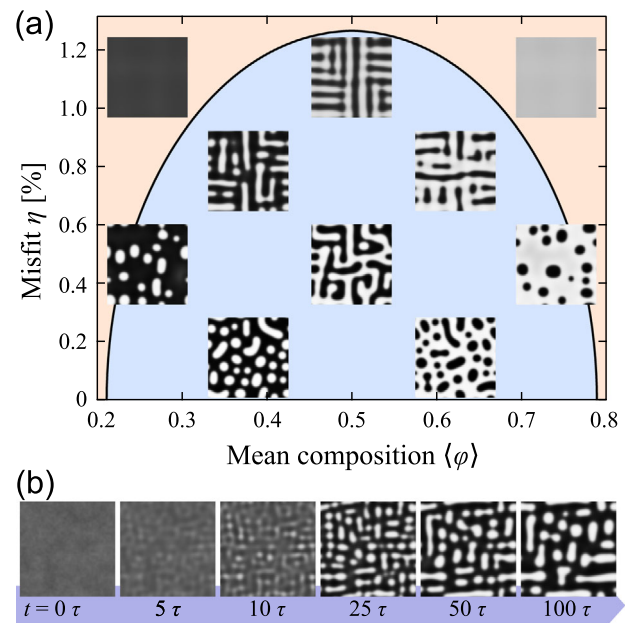


FIGURE 1 | (a) Phase diagram and examples of the different morphologies of the spinodal decomposition of a binary alloy with cubic elastic anisotropy, as a function of mean composition $\langle\varphi\rangle$ and misfit strain η . The reported images correspond to representative evolution stages obtained by the numerical integration of the Cahn–Hilliard Equation (8) starting from the homogeneous phase slightly perturbed by Perlin noise around the average composition $\langle\varphi\rangle$. (b) Example of a phase-field evolution from the dataset, starting from a Perlin noise initial condition and corresponding to a misfit strain of $\eta \approx 0.48\%$.

microstructural evolutions, obtained from our simulations using the standard PF approach discussed in the Methods section, for a binary mixture with cubic elastic constants [46, 48], are shown for different mean composition and misfit η in Figure 1a. A continuum variation from rounded to elongated and clearly anisotropic domains is found for increasing misfit.

Several machine learning approaches have taken spinodal decomposition as a prototypical case study, although neglecting strain effects, either resorting to latent-space dynamics [17, 18, 49] or exploiting architectures sensitive to spatial and time correlations directly in pixel space (see e.g. Refs. [31, 50, 51]), among which CRNN proved very convenient [23–26].

Developing a NN model able to cope with the broad variability of morphologies for strained systems is clearly a more demanding task compared to the reported studies, as a standard NN setup would hardly be usable unless providing a huge amount of examples. Additionally, the resulting NN architecture may require a very large number of parameters, making the approach inefficient from a computational standpoint and failing to provide an acceleration with respect to traditional schemes. The here developed CRNN architecture tackles such complexity by taking η as an explicit conditioning parameter. Dependencies from composition and geometry are instead implicitly learned from the evolving concentration field, encoded as an image. In addition, we introduce physics-based information as inductive biases and regularization [16]. In particular, we specifically enforce physical consistency by mimicking the underlying evolution laws with

specialized NN layers and physics-inspired regularization terms in the training objective. This way, we show that the NN model yields accurate predictions, capturing the typical features set by the misfit parameter. Moreover, the trained model shows good generalization capabilities on both temporal and spatial scales, thus giving the possibility of scouting cases beyond the training conditions.

The paper is structured as follows. In Section 2, we first define the physical model for the microstructure, based on a Phase-Field (PF) description (2.1), detail the dataset preparation process (2.2), and then outline the CRNN architecture, training methodology, and evaluation framework (2.3). The results are then discussed in Section 3, inspecting the reliability of the NN with respect to the ground truth numerical simulations, both from a pixel-level perspective (3.1) and with respect to the reproduction of mean global properties (3.2). The analysis is meant to assess the nature of prediction errors and the NN generalization performances for longer time and larger domain sizes (this latter inspected in 3.3). Finally, the trained model is applied to numerically reconstruct the phase diagram (3.4). In the Conclusions section, we review the work achievements and their potential uses.

2 | Methods

2.1 | Phase Field Model of Coherently Strained Spinodal Decomposition

The standard way for modeling spinodal decomposition dynamics is by the PF approach [3]. For isothermal conditions, the system configuration is specified by an order parameter φ , tracing the local composition of the binary mixture, and by the total strain tensor field $\varepsilon(x)$. Linear elasticity is considered throughout the paper. The system free energy F is then expressed by the Ginzburg–Landau functional [52]:

$$F[\varphi, \varepsilon] = \int_D \left(\frac{\varepsilon}{2} |\vec{\nabla}\varphi|^2 + w(\varphi) + \rho(\varphi, \varepsilon) \right) d\vec{x} \quad (1)$$

where D is the physical domain and ε sets the width of the interface between the two phases. The $|\vec{\nabla}\varphi|^2$ term accounts for the energy cost of such interfaces, while the remaining “bulk” contribution divides into a double-well potential

$$w(\varphi) = \frac{18}{\varepsilon} \varphi^2(1 - \varphi)^2 \quad (2)$$

with minima at $\varphi = 0$ and $\varphi = 1$, corresponding to the equilibrium compositions of the two separated phases in the absence of strain [53], and an elastic energy density $\rho(\varphi, \varepsilon)$:

$$\rho(\varphi, \varepsilon) = \frac{1}{2} C_{ijkl} (\varepsilon_{kl} - \varepsilon_{kl}^*(\varphi)) (\varepsilon_{ij} - \varepsilon_{ij}^*(\varphi)) \quad (3)$$

where C_{ijkl} is the fourth-order tensor of elastic constants, and ε^* is the eigenstrain tensor, i.e. the zero-stress strain [54]. Einstein notation, implying summation of repeated indices, is adopted. The ρ term accounts for the different lattice parameters of the pure materials composing the alloy. Due to this additional energy cost, the bulk energy minima tend to move inward of the $[0,1]$ range.

In the present study, we consider (elastically) homogeneous materials, i.e. with C_{ijkl} independent of φ , and purely dilational eigenstrain:

$$\varepsilon_{ij}^*(\varphi) = \eta(\varphi - \langle\varphi\rangle)\delta_{ij} \quad (4)$$

where η is the lattice mismatch between the two phases, and a simple linear approximation for the dependence of ε_{ij}^* from φ has been used, taking the homogeneous phase at mean composition $\langle\varphi\rangle$ as reference state. Notice, also, that this choice explicitly disregards the possible presence of other elastic contributions, such as dislocations and grain boundaries, which are therefore not considered in the present study. Since the time-scale of strain relaxation is much faster than the diffusive one of the phase separation, it is customary [55] to assume that mechanical equilibrium holds at any time, so that the strain field ε can be directly computed by solving the (static) equilibrium condition. For a periodic and elastically homogeneous material, this can be conveniently achieved in Fourier space. Following Refs. [35, 54, 55], the Fourier transform of the equilibrium stress $\hat{\sigma}$ field can be directly obtained from the Fourier transform of the order parameter $\hat{\varphi}$ as:

$$\hat{\sigma} = -\eta \mathbf{B}(\vec{q}) \hat{\varphi} \quad (5)$$

$$B_{ij}(\vec{q}) = C_{ijkl} (\delta_{kl} - \Omega_{km} C_{mnop} \delta_{op} q_n q_l) \quad (6)$$

where \vec{q} is the wavevector and $\Omega_{kn}(\vec{q})$ is the inverse of the Green tensor for acoustic displacement, i.e., $\Omega_{kn}^{-1} = C_{kmno} q_m q_o$.

Thanks to the quasi-equilibrium assumption, the free-energy functional (1) just depends parametrically on strain, so that the dynamics of spinodal decomposition can be fully described by the well-known Cahn–Hilliard equation [52] as

$$\frac{\partial \varphi}{\partial t} = M \nabla^2 \mu = M \nabla^2 (-\varepsilon \nabla^2 \varphi + w'(\varphi) + \rho'(\varphi, \varepsilon)) \quad (7)$$

where $\mu = \delta F / \delta \varphi$ is the local chemical potential, M is a mobility constant and the ' symbol indicates derivation with respect to φ . From Equations (3) and (4), for elastically homogeneous materials, it can be found that $\rho'(\varphi, \varepsilon) = -\eta \text{tr}(\sigma)$, being $\text{tr}(\sigma)$ the trace of the stress tensor set from Equation (5). Numerically, the time-integration of Equation (7) is here conveniently achieved by using a semi-implicit spectral scheme [56] as

$$\hat{\varphi}_{t+\delta t} = \frac{\hat{\varphi}_t - \delta t M q^2 \widehat{w'(\varphi_t)}}{1 + \delta t M (\varepsilon q^4 + \eta^2 \text{tr}(\mathbf{B}) q^2)} \quad (8)$$

where $\widehat{w'}$ indicates the Fourier transform of the first derivative of the double-well potential and δt is the integration time-step. Spatial discretization is instead set on a uniform 2D square grid.

From a linear-stability analysis, it can be demonstrated [35] that the mixed phase is unstable and undergoes spontaneous phase separation when the following criterion is met:

$$w''(\varphi) + \eta^2 \min_{\vec{q}} (\text{tr}(\mathbf{B}(\vec{q}))) < 0 \quad (9)$$

In the strain-free case, the stability of the homogeneous mixture is determined by the curvature of $w(\varphi)$, which indeed defines the spinodal lines in the (φ, T) -phase diagram. Elastic strain, on the other hand, always penalizes decomposition, thus stabilizing the homogeneous phase and extending the metastability within a part of the spinodal region. Moreover, by considering anisotropic materials, $\mathbf{B}(\vec{q})$ changes according to the crystallographic orientation, thus making phase separation direction-dependent. In particular, this can lead to a suppression of the phenomenon along hard crystallographic directions and to an enhancement along soft ones, thus resulting in anisotropic phase patterns.

In this work, we focus on the case of cubic symmetry and consider an arbitrary material with a high Zener anisotropy ratio $Z = 2C_{44}/(C_{11} - C_{12}) = 4$ and a wide misfit range $\eta \in [0, 1.2\%]$. Elastic constants are fixed as $C_{11} = C_{44} = 2C_{12} = 3 \times 10^4$ in the arbitrary units of the double-well potential and interface term. All values have been chosen to maximize the variability in the phase-separation patterns. By this choice, softer directions are aligned with the $\langle 10 \rangle$ and $\langle 01 \rangle$ axes so that, for high strains, striped domains are expected to form along them [46, 57].

A collection of representative microstructures, obtained by performing PF evolutions for varying combinations of $\langle \varphi \rangle$ and η , is shown in Figure 1a. Consistently with theory and other simulation studies [46], configurations having $(\langle \varphi \rangle, \eta)$ values outside of the spinodal region (orange region) are stable against small fluctuations around the mean composition, thus remaining in the homogeneous single-phase state. An example of a PF time-evolution sequence for a case of moderate misfit $\eta \approx 0.482\%$, serving to visually illustrate the type of microstructural transformation considered in this work, is shown in Figure 1b. As discussed, domains aligned with the $\langle 10 \rangle$ and $\langle 01 \rangle$ elastically soft directions rapidly emerge from the initial random configuration, although in this case the elastic contributions are not strong enough to prevent the formation of more rounded inclusions.

2.2 | Dataset Generation

A dataset of 2000 time-sequences of spinodal decomposition has been constructed by the numerical integration of Equation (8) for different initial random configurations, average composition $\langle \varphi \rangle$, and misfit strain η . All simulations have been performed on a 128×128 uniform square grid of collocation points with periodic boundary conditions naturally enforced by the spectral description. To ensure a smooth resolution of the PF interface, we set $\epsilon = 5$, using grid point distance as the unit for length. Since the mobility M just acts as a scaling factor for the time scale (see Equation (7)), we take it unitary and use a time-step $\delta t = 0.01$ to achieve a stable numerical solution. Each sequence in the dataset comprises 100 snapshots, taken at fixed time intervals τ , corresponding to 100 integration time-steps δt , in such a way that consecutive frames are sufficiently different to provide effective information for the NN training. In the following, we use τ as the unit of time.

The initial configuration for each simulation is designed to represent a homogeneous mixed phase with local composition fluctuations, eventually seeding the phase separation process and resulting in the typical microstructural evolution patterns of

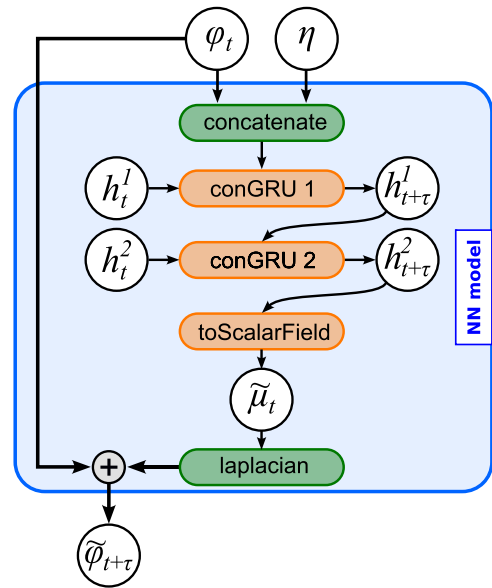


FIGURE 2 | Schematic representation of the Convolutional Recurrent NN model for evolution prediction. At each time step, the model predicts the time-evolved state $\tilde{\varphi}_{t+\tau}$ from the previous state φ_t , the hidden states h_t^i and (i enumerating the convolutional GRU submodule), and the parameter η provided as a constant field. Rather than predicting $\tilde{\varphi}_{t+\tau}$ directly, the network outputs an auxiliary scalar field $\tilde{\mu}_t$, from which the updated state is computed. Green boxes represent fixed operations, while orange ones are learnable transformations.

Figure 1a. This is obtained by initializing the composition field in the form of Perlin noise [58], a gradient noise algorithm that generates structured correlated random profiles. Unlike white noise, Perlin noise consists of smoother patterns with features that the NN can properly distinguish from actual numerical noise.

To achieve robust generalization during model training, a sufficiently diverse set of initial configurations is needed. To this goal, a randomized shifting and rescaling of the Perlin noise profiles is implemented to vary the average composition $\langle \varphi \rangle$ in the range $[0.2, 0.8]$. Since with the present parameters $\varphi \approx 0.21$ and $\varphi \approx 0.79$ are the limiting compositions for spinodal decomposition at zero misfit (see also Figure 1a), this ensures the presence of both homogeneous and phase-separated configurations. The extent of the fluctuations in the Perlin noise profile is instead randomly sampled from a normal distribution with a standard deviation equal to 0.1

2.3 | CRNN Architecture

In this work, we extend the CRNN architecture proposed in Refs. [26, 32] to achieve explicit conditioning to an assigned parameter, as sketched in the scheme of Figure 2. The goal of the CRNN is to predict the entire time evolution of the composition field φ solely based on an initial configuration φ_0 and the known value of the misfit parameter η .

The NN takes as input the composition field at each time step φ_t , in the form of a 128×128 pixel image. Due to the recurrent nature of the NN, in principle the prediction may be based

on a short subsequence; however, we focus here on the more challenging possibility of providing the model with the initial condition φ_0 only. To condition the evolution on the misfit parameter, the φ_t map is concatenated to η at each timestep. The pair (φ_t, η) is then internally represented as a single 128×128 image with two channels, one of which is uniform and equal to η . Despite its simplicity, this approach is particularly convenient for its ease of implementation and can be straightforwardly extended to include additional parameters and, possibly, non-uniform fields too. At variance with other approaches in the literature [25], the used architecture does not employ pooling and flattening operations, so to preserve the fully-convolutional character of the CRNN enabling direct generalization to arbitrary domain sizes [26]. Circular padding [59] is used to enforce periodic boundary conditions by construction, consistently with the dataset examples.

Input data are processed through two stacked convolutional gated recurrent unit (convGRU in Figure 2) blocks [60]. Each convGRU block employs 24 channels and 5×5 convolutional kernels, chosen according to a preliminary hyperparameters scan to ensure the best performance and stability of the model. This choice, moreover, guarantees a fairly contained number of NN parameters ≈ 150000 , which translates into lower computational costs both at training and evaluation time.

As a performance figure of merit, we measured the ratio in wall time for evaluating a time frame τ by the NN trained model with respect to the semi-implicit numerical scheme used to generate the dataset. The former takes full-profit of CUDA acceleration on a NVIDIA RTX A5000 GPU while the latter were executed on a Xeon W-2133 CPU on the same workstation. Speed-ups of approximately $40\times$ on a 128×128 domain, $90\times$ on a 256×256 domain, and $145\times$ on a 512×512 domain are estimated. A substantially linear scaling with the number of collocation points [26] is found for the NN approach, making it especially well-suited for simulations on large domains. It must be however noted that the previous comparison represents the best-case scenario, as the speed-up just reduces to a factor 3–4 \times if evaluating the NN model on the aforementioned CPU or exploiting GPU acceleration of the numerical solver. Nonetheless, we remark that the model formulation here considered is particularly convenient for numerical evaluations, first because of the use of a Green's function approximation for elasticity, and easily transferable to GPU thanks to the availability of efficient Fast-Fourier-Transform algorithms. This would not be the case if considering more quantitative approaches, e.g. Finite-Element Method computations of the elastic field, or tackling more complex, highly non-linear problems for which the largest speed-ups first reported are expected to be well representative.

In principle, the CRNN could be devised to directly produce the next system state $\tilde{\varphi}_{t+\tau}$ (tilde distinguishing predicted quantities from dataset ones), as typically done in the literature studies [23–25]. As already discussed in Ref. [26], however, this approach lacks any constraint on the conservation of φ , which is inherent to the Cahn–Hilliard dynamics. A significant improvement in the consistency of the NN predictions with the conservative flow dynamics prescribed by Equation (7), especially for time-extrapolation, is achieved by introducing an additional physics-inspired output layer. The network is indeed designed to output

the scalar field $\tilde{\mu}_t$, taking the role of the local chemical potential (multiplied by M) in Equation (7) so that the new frame can be generated as (see Figure 2):

$$\tilde{\varphi}_{t+\tau} = \tilde{\varphi}_t + \nabla^2 \tilde{\mu}_t \quad (10)$$

This is an alternative version of the approach implemented in Ref. [26], where the flow current $\vec{J} = -M\nabla\mu$ was evaluated to enforce exact conservation of φ by construction and is here found to provide superior results, as detailed below and in the Supporting Information. From a NN perspective, the presence of this physics-inspired layer acts as a strong regularizer, e.g. by making it impossible for the model to produce non-conservative dynamics, mitigating extrapolation errors when the generated sequences exceed the length of those used in training.

We define the loss function used to train the model as the spatial and temporal average of the weighted sum of three components: the mean squared error (MSE) between the PF ground truth and NN predicted sequences (L_φ), the MSE between the squared gradients of these sequences (L_∇), and the MSE on the mean free-energy density $\langle f \rangle \approx 2\langle w \rangle + \langle \rho \rangle$ (set from Equation (1) with the approximation $\epsilon|\nabla\varphi|^2 \approx w(\varphi)$, see e.g. in Ref. [61]) predicted by NN with respect to its true PF value (L_f):

$$L(\theta) = \frac{1}{N_{\text{ts}}T} \sum_{i=1}^{N_{\text{ts}}} \sum_{t=1}^T (L_\varphi(\theta) + \lambda_\nabla L_\nabla(\theta) + \lambda_f L_f(\theta)) \quad (11)$$

with

$$\begin{aligned} L_\varphi(\theta) &= \langle [\varphi_i(t) - \tilde{\varphi}_i(t|\theta)]^2 \rangle \\ L_\nabla(\theta) &= \langle [|\nabla\varphi_i(t)|^2 - |\nabla\tilde{\varphi}_i(t|\theta)|^2]^2 \rangle \\ L_f(\theta) &= \langle f(\varphi_i(t)) - f(\tilde{\varphi}_i(t|\theta)) \rangle^2 \end{aligned} \quad (12)$$

Here, θ represents the set of NN parameters, i indexes the N_{ts} elements of the training set, t is the time step, ranging from one to the total sequence length T , and $\langle \cdot \rangle$ indicates the spatial average. The parameters λ_∇ and λ_f are the two weights controlling the relative importance of the gradient and free-energy terms, respectively. Based on preliminary analysis, which evaluated the impact of these terms in the loss function on the temporally evolved profiles predicted by the model, they were set to 60 and 150, respectively.

While the $L_\varphi(\theta)$ term, evaluating the pixel-by-pixel correspondence between the predicted and true fields, provides a generic definition of loss function, the additional two terms $L_\nabla(\theta)$ and $L_f(\theta)$ introduce physics-inspired regularizations tailored for the model at hand. $L_\nabla(\theta)$ penalizes local gradients in the composition field φ , thus suppressing spurious oscillations within the segregated domains. $L_f(\theta)$ instead compares the free energies of predicted and true configurations, thus enforcing their match in terms of such a thermodynamic global quantity. As shown in Figures SI–S3, the addition of these two loss terms significantly enhances the NN model performances, improving both the stability and quality of the prediction, especially in extrapolation conditions.

The full dataset of 2000 simulations is used for the training of the NN, with a 4:1 random partitioning between training and validation. Data augmentation techniques, including reflections, 90° rotations, and the \mathbb{Z}_2 symmetry transformation $\varphi \rightarrow (1 - \varphi)$, are employed to further improve the robustness and diversity of the training dataset, encoding symmetries not already present by construction into the model architecture. Additionally, to enhance the model's generalization capabilities, noise injection regularization is applied during training by introducing small random perturbations such that $\varphi \rightarrow \varphi + \delta$, where δ is sampled from a Gaussian white noise, as in our previous work in Ref. [26] (δ standard deviation is 1×10^{-3}).

The training of the NN parameters is performed using the standard implementation of the Adam optimizer [62]. Thanks to the recurrent character of the model, the so-called curriculum learning technique [63] has been exploited. In practice, at the first epoch, the model is provided with 99 snapshots from a sequence and tasked with predicting only the final step, with the loss evaluated solely on that last prediction. As training progresses, the number of input snapshots is gradually reduced, which requires the model to predict increasingly longer portions of the sequence. By the end of the ramp, the model has only the initial snapshot and must predict the remaining sequence, with the loss evaluated on all of its predictions. Although this procedure may lead to an apparent initial increase in loss, it facilitates more stable and efficient learning by progressively introducing the complexity of the task.

3 | Results

3.1 | Training, Validation and Testing

As a first assessment of the NN model, we inspect its training and validation performances. Figure 3a reports the decay of the loss L as a function of the number of training epochs. In about 100 epochs, the value of L settles, albeit with some fluctuations, as expected because of the stochastic nature of the optimization procedure (see, e.g., previous work in Ref. [26]). The spike in the training loss (blue line) in the initial epochs is instead associated with the curriculum learning procedure discussed in the previous section.

The model returning the best performance on the validation set was then selected for all subsequent analysis. To get a visual confirmation that it properly predicts the phase separation dynamics, in Figure 3b and Figure 3c we compare true and predicted evolutions from the validation set for low $\eta \approx 0.07\%$ ($\langle \varphi \rangle \approx 0.396$) and relatively high $\eta \approx 0.76\%$ ($\langle \varphi \rangle \approx 0.388$) misfit conditions. As shown, a nearly one-to-one match between the sequences is achieved, but for minor local differences in the connectivity of the segregated domains.

We next check the NN model's predictive capability on an independent test set to ensure robustness and generalizability. The test set comprises 500 sequences generated with the same characteristics as the training dataset but reaching a longer time of 500τ , i.e. 5 times longer than the sequences used during training, allowing us to assess the model's temporal generalization capability. As an illustrative example, in Figure 4a we

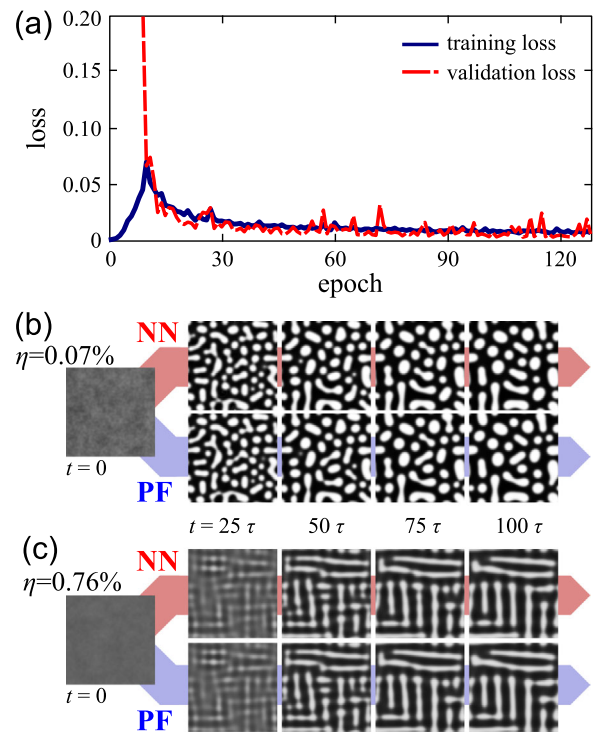


FIGURE 3 | (a) Training and validation losses during training. The spike in the first 10 epochs is due to the curriculum learning procedure. (b, c) Comparison between NN predicted and PF true evolutions for two cases from the validation cases, with low ($\eta \approx 0.07\%$) and high ($\eta \approx 0.76\%$) misfit, respectively.

report a comparison between the NN predicted and ground truth dynamics for a sequence with $\eta \approx 0.45\%$ taken at random from the test set. As far as considering interpolation within the 100τ range of training, the same almost one-to-one match found on the validation cases (Figure 3b,c) is still achieved. However, due to the accumulation of small prediction errors across the time steps and to critical bifurcation events, such as domain splitting and coalescence, the pixel-level match progressively degrades, and local discrepancies in the morphologies are observed. At time $t = 500\tau$, the predicted profile still looks plausible from a qualitative point of view, but fails in reproducing several features of the ground truth as highlighted by the mapping of absolute error $|\Delta\varphi| = |\varphi - \tilde{\varphi}|$.

To provide a quantitative evaluation of the accuracy of the model predictions, we consider the Mean Absolute Error (MAE), i.e. $\langle |\Delta\varphi| \rangle$, at each evolution time step. These metrics provides a measure of the degree of pixel-to-pixel overlap between predicted and true frames.

In Figure 4b, we report the statistical distribution of the MAE evaluated on the whole test set sequences. The median MAE and the inter-quartile range (IQR) indicate a sublinear error accumulation, which remains as low as 0.025 within the full $[0, 100\tau]$ range of time interpolation, growing up to a median value of ≈ 0.09 when extrapolating to 500τ . These limits of confidence are in line with the typical extrapolation errors reported by other works in the literature [16, 21] and are inherently bound to the numerical stability against bifurcations in the

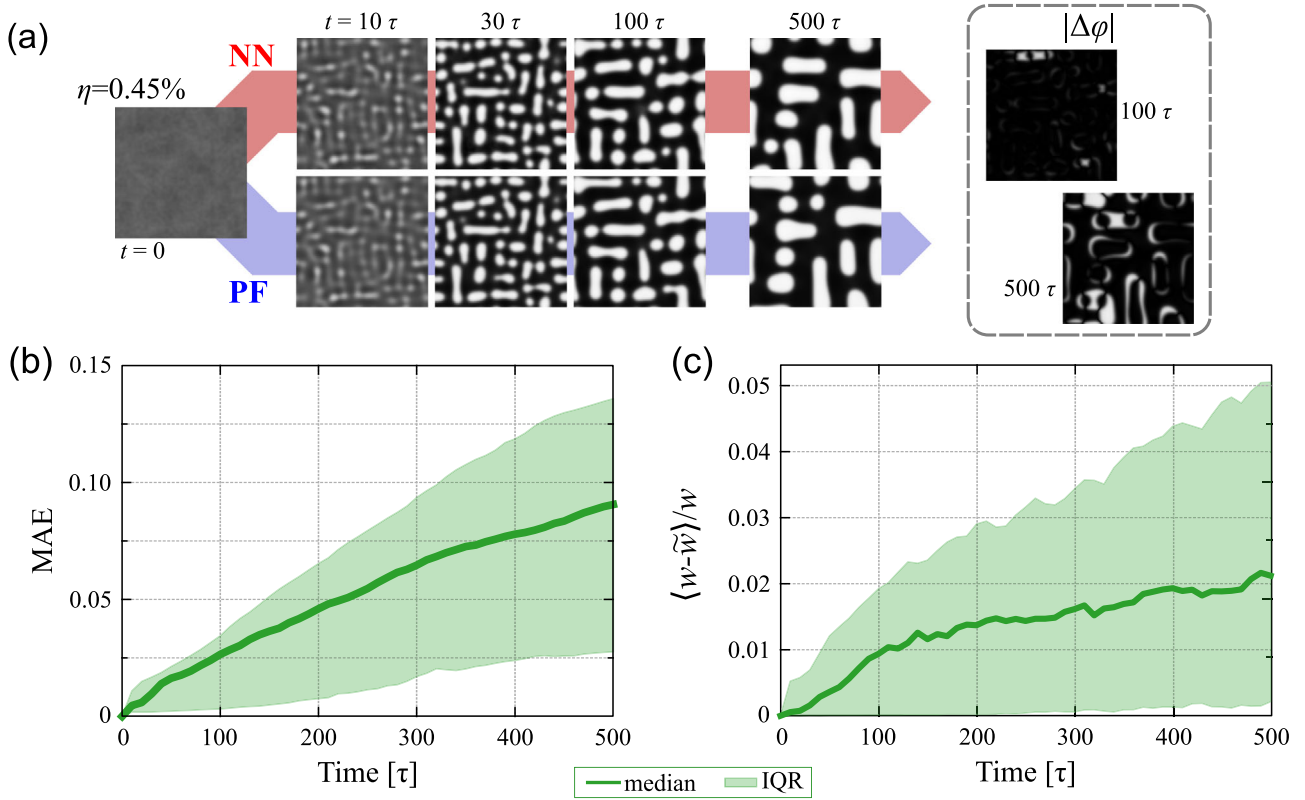


FIGURE 4 | (a) Test sequence of spinodal decomposition for misfit $\eta \approx 0.45\%$, comparing the NN prediction with the ground-truth PF evolution, initiated from a Perlin noise profile with $\langle \varphi \rangle = 0.5$. The map of absolute error, evaluated as the point-wise difference between the two profiles $|\Delta\varphi| = |\varphi - \tilde{\varphi}|$, is reported for $t = 100\tau$ and 500τ . Time evolution (b) of the mean absolute error (MAE) and (c) of the relative error on the interface length, measured by $\langle \omega(\varphi) \rangle$, between NN and true PF simulations evaluated on a test set of 500 sequences, up to 500τ . The median value (solid line) and inter-quartile range (IQR, shaded region) are shown.

time-dependent dynamics, also affecting the PF simulations (e.g. when varying spatial and temporal discretizations or introducing noise). Indeed, even small discrepancies injected at some time during the evolution sequence may progressively amplify as time advances and eventually lead to distinct morphological outcomes. Other error metrics, evaluating the model predictions beyond the pixel-level, are hence to be considered in order to assess how the trained NN complies with the underlying physical model.

Among these, a key global quantity that can serve the purpose is the total interface length, which can be estimated, normalized on the domain size, as $2\langle \omega(\varphi) \rangle$ [61]. In Figure 4c, we show the trend of the relative error on the prediction of interface length on the same test set, highlighting the median and inter-quartile range. With respect to MAE, the error for this metric grows more slowly, remaining well below 2% for the first 100τ , and still under $\approx 5\%$ even at 500τ for most of the considered cases, as shown by the inter-quartile range. These quantities are of the same order of magnitude as reported in previous studies [25, 26, 51], where the relative errors in characteristic lengths is reported to be in the order of few percent. Notice, however, that in the present case the interface length only serves as a geometric characterization, as the evolution of the system is also driven by elastic contribution. Still, the results indicate that the NN model provides a reasonable representation of the overall domain evolution even in the extrapolation regime where pixel-level match is lost.

3.2 | Time Evolution of Mean Properties

In the previous section, we concluded that the NN model is capable of producing pixel-wise accurate predictions within the training time range, while at longer stages, larger errors were obtained. This section, instead, aims at assessing to what extent such long-time predictions remain overall physically meaningful, despite abandoning the claim of reproducing the exact sequence. For this purpose, we then focus here on the time trends of global descriptors, representative of the system state and hence less affected by local bifurcation effects.

Since global behaviors differ according to the average composition $\langle \varphi \rangle$ and misfit parameter η , a dedicated test set of 50 independent evolutions for fixed $\langle \varphi \rangle = 0.5$ and $\eta = 0.6\%$ is here devised (an analogous analysis for a case with $\langle \varphi \rangle = 0.5$ and $\eta = 0.3\%$ is reported in Figure S6). Each sequence extends up to a time of 1000τ , i.e. 10 times longer than the sequences used in training.

As the Cahn–Hilliard equation drives the system evolution toward minimization of the free energy defined in Equation (1), this quantity can be taken as a primary descriptor to assess how well the NN predicted evolution obeys the underlying thermodynamics [26]. Figure 5a reports the time evolution of the mean free-energy density $\langle f \rangle$, also used as a regularizing parameter in the loss. Solid lines represent ensemble medians, while shaded areas correspond to the inter-quartile range (IQR).

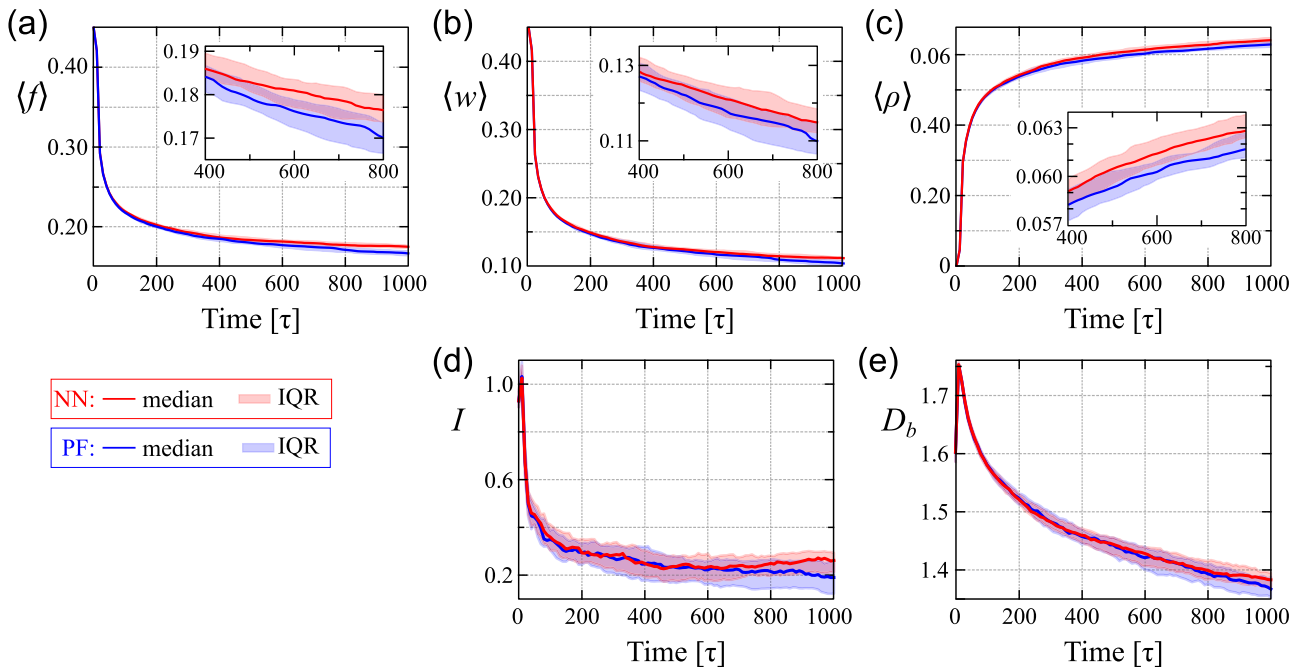


FIGURE 5 | Time evolution of (a) mean free energy density $\langle f \rangle$, (b) mean interface energy density $\langle w \rangle$, (c) mean elastic energy density $\langle \rho \rangle$; (d) isotropy degree I and (e) boundary fractal dimension D_b , obtained from 50 independent simulations for fixed $\langle \varphi \rangle = 0.5$ and $\eta = 0.6\%$ and different initial Perlin noise configurations. The median value is traced by a solid line while shaded areas represent the inter-quartile range (IQR). NN predictions are shown in red while the corresponding PF ground truth ones are in blue.

Overall accordance is achieved, with median values falling within the IQR of each other for most time. Deviations become more evident only at the later stages, at $\approx 900\tau$, marking the loss of model reliability also for global properties, even if not yet corresponding to an actual breakdown in the model predictions.

For a more in-depth analysis, in Figure 5b,c, we consider separately the two contributions to f , i.e. the mean interface energy density $\approx 2\langle w(\varphi) \rangle$ and the mean elastic energy density $\langle \rho \rangle$. Despite the NN model slightly underestimating the decay rate of the first and overestimating the rise of the second, NN predictions and PF ground truth values still fall within the IQR of each other. Importantly, the relative errors on the medians for the two contributions are similar, showing that the model is not biased toward one of the two.

Two additional descriptors, sensitive to the geometrical features of the segregated domains, are also inspected in Figure 5d,e. The first, in panel (d), is the isotropy degree I , defined as a measure of the directional balance of the domain boundary profile, evaluated from the histogram of oriented gradients [64] computed on each snapshot. In detail, each image is divided into local cells of 4×4 pixels, grouped into 2×2 blocks for local normalization, and gradient orientations are binned into 10 angular intervals between 0° and 180° . The average gradient magnitude in each orientation bin, obtained after spatial averaging, reflects the relative abundance of interfaces aligned along that direction. I is then defined as the ratio between the mean magnitude of non-axial orientations $\langle M_{(-10)} \rangle$ and that along the $\langle 10 \rangle$ horizontal and vertical directions $\langle M_{(10)} \rangle$:

$$I = \frac{\langle M_{(-10)} \rangle}{\langle M_{(10)} \rangle} \quad (13)$$

Values close to unity indicate isotropic patterns, whereas values close to zero reflect the preferential alignment of structures along the horizontal and vertical directions.

The second descriptor, in Figure 5e, is the boundary fractal dimension [65] D_b , which instead characterizes the complexity of the domain interfaces by estimating how the number of edge pixels scales with the observation length scale in a box-counting sense, computed as the (negative) slope of the log-log relationship between the number of occupied boxes $N(L)$ and the box size L :

$$D_b = -\frac{d \log N(L)}{d \log L} \quad (14)$$

Higher values correspond to more irregular boundaries. Both metrics are well reproduced by the NN predictions, despite not being included in the loss function, thus giving further confirmation about the model's ability to capture the geometrical features of the evolving microstructures.

According to the previous analyses, the NN model is capable of producing physically appropriate global behaviors even for extrapolation up to $\approx 8 - 9$ times the training sequence length.

We find it interesting to rationalize the amplitude of the observed prediction error by comparing it to the effects of random perturbations in the PF dynamics. To this purpose, the integration scheme of Equation (8) is modified so that at each time step a random Gaussian noise is added to the value of φ_i . Exact conservation of mass is enforced by subtracting the noise's mean value over all pixels. We then inspect the variability of the solution profiles induced by such noise by comparing the profiles obtained from PF simulations with noise with the corresponding ground truth




FIGURE 6 | (a) Time evolution of MAE between PF simulations with noise and ground truth deterministic ones from the test set of Figure 4, $MAE = \langle |\varphi_{\text{noise}} - \varphi| \rangle$. (b) Time Evolution of the mean free energy density $\langle f \rangle$ comparing PF simulations with noise and corresponding NN predictions from the test set of Figure 5. Solid curves report median values, while shaded areas are inter-quartile ranges. Blue is used for ground-truth and red for the perturbed dynamics.

ones from the deterministic PF evolution (see Figure S4). For a quantitative comparison, we consider as ground truth the same 500 PF sequences from the testing set used for Figure 4 and for each one, perform a new PF simulation starting from the same initial configuration but including noise. Figure 6a shows the MAE trend computed on this statistical ensemble as $MAE = \langle |\varphi_{\text{noise}} - \varphi| \rangle$. Given the accumulation of deviations introduced by noise, the MAE grows over time and, as shown in Figure S5, it also increases with the standard deviation of the injected noise. By comparing the observed trends with those of Figure 4b, it is found that a noise with standard deviation of 7.5×10^{-4} returns a variability range very similar to the prediction error of the trained NN model. Notably, the inter-quartile range for the perturbed PF dynamics is quite wider, meaning that the NN model has a narrower error distribution for the same median value.

As discussed in the previous sections, the MAE only measures pixel-level similarity between two evolutions, and cannot show consistency between more coarse properties, such as the free energy decay. For a direct comparison with the NN error analysis of Figure 5, the same set of 50 independent evolutions with $\langle \varphi \rangle =$

0.5 and $\eta = 0.6\%$ was considered and compared with perturbed dynamics as discussed in the previous paragraph. Figure 6b shows the evolution of the mean free energy density $\langle f \rangle$ for the PF noisy ensemble and for the NN test set already reported in Figure 4a. The compelling match of the two curves, both for median values and IQR range, clearly indicates how the NN error accumulation is indeed equivalent to having a random error, or possibly physical fluctuations, in the PF solution. Since these do not compromise the ensemble thermodynamic behavior, the comparison proves the physical consistency of the NN trajectories, even when they miss the one-to-one correspondence with the ground-truth due to bifurcations in topological events (e.g., domain pinching and merging).

3.3 | Domain-Size Generalization Tests

A major advantage of using a fully-convolutional architecture is its independence from the domain size. In this section, we then apply the trained NN model to 512×512 domains, i.e. with an area 16 times larger than those used during training, and test its robustness and generalization capabilities beyond the training domain size. Figure 7a,b compares NN predictions and PF ground truth morphologies obtained after evolution for 100τ for two representative cases with $\langle \varphi \rangle = 0.45$ and $\eta = 0.4\%$ and 0.8% , respectively. Aside from local discrepancies related once again to domain pinching or coalescence, the predicted profiles are in excellent agreement with the true ones, accurately capturing the dominant morphological features and large-scale evolution trends.

As an additional test, we analyze the NN performances in predicting the evolution of peculiar configurations that differ qualitatively from the Perlin noise ones of the training set, while also considering domain-size and sequence-length generalization. In Figure 7c, we consider an initial condition featuring a spatially varying average composition, which is higher in the center and lower toward the edges of the domain. A 512×512 large domain and a long sequence reaching a total time of 1000τ ($10 \times$ the training set) are considered. Despite the deviations from the training distribution, the model accurately captures the overall morphological evolution.

Both analyses demonstrate generalization capabilities and the effectiveness of the learning framework in scaling to larger domains.

3.4 | NN Reconstruction of the Phase Diagram

In this section, we aim to demonstrate the potential of the previously tested model as a fast and reliable surrogate of the ground-truth PF approach to infer physical insights into the system. For this purpose, we here consider the task of reconstructing the phase diagram theoretically provided in Figure 1a. To this end, we use the same test set of 500 sequences, randomly sampling $(\langle \varphi \rangle, \eta)$ pairs, as exploited in the analysis of Figure 4. Here we establish whether each one corresponds to phase separation or to a single homogeneous phase by evaluating the difference between the maximum and minimum values of φ at 100τ and classify the system as single phase whenever it is below an

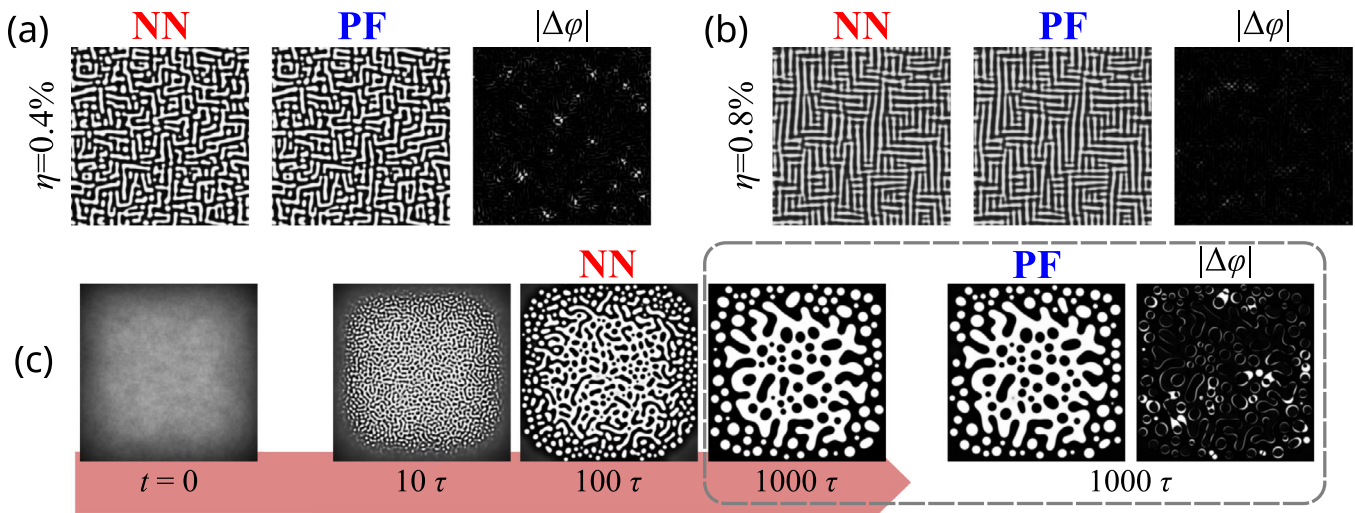


FIGURE 7 | (a,b) Comparison between predicted and true profiles for spinodal decomposition at $t = 100\tau$, starting from initial Perlin noise profile on large 512×512 domains, for (a) $\eta = 0.4\%$, and (b) $\eta = 0.8\%$. The absolute error map over the domain is also reported. (c) Example of a 1000τ -long evolution sequence on a 512×512 domain with $\eta = 0\%$ and starting from a spatially varying initial composition: higher in the center and lower at the edges. Despite local differences, the NN prediction is in accordance with the ground truth up to 1000τ , demonstrating generalization capabilities beyond the training distribution.

arbitrary threshold of 0.05. This classification is performed on both NN and PF profiles so that the model predictions can be directly compared with the ground truth PF behavior under the same conditions. Results are reported in Figure 8a, along with the analytical spinodal line. In the large majority of cases, the NN output matches the PF one, distinguishing in the same way between single and two-phase conditions. Only for a few configurations (9 out of 500), the model was unable to correctly predict the outcome of the dynamics (black crosses). Figure 8b reports three representative examples where the NN fails to reproduce the correct behavior at time 100τ . In these cases, the true evolutions remain uniform, whereas the model incorrectly predicts phase separation or introduces spurious artifacts and unphysical features. Such artifacts are not observed in other examples, except in critical conditions where the system lies near the threshold between phase separation and uniform mixing. No cases in which the NN wrongly predicts a final uniform composition state are observed.

It must be, however, noticed that $\langle\varphi\rangle$ and η alone are insufficient to unambiguously determine the onset of phase separation even in the PF simulations, as the amplitude of the initial perturbations also plays a significant role. This is the reason why in the simulated phase diagram, both from NN and PF, it is not possible to find a perfectly sharp boundary between the two conditions as expected by the theory. To further illustrate the effect of the initial conditions, Figure 8c presents two evolutions starting from the same Perlin noise profile but with different perturbation amplitudes. While the configuration with larger perturbations undergoes spinodal decomposition, the other evolves to a uniform field. The NN model is found to accurately capture both outcomes, highlighting its sensitivity to the initial perturbation even for borderline conditions between phase separation and homogeneous phase.

We stress once again that the NN model has no explicit information on the stability condition of Equation (9), hence we conclude

that the model implicitly learned the role that elastic strain plays in penalizing spinodal decomposition. This suggests the potential application of this approach to more complex systems where corresponding analytical expressions are not available, possibly requiring a high number of calculations.

4 | Conclusion

Based on a suitable convolutional recurrent NN architecture, we proposed a model able to predict the temporal evolution of alloy microstructures during coherent spinodal decomposition encompassing, the whole morphological variability of the $(\langle\varphi\rangle, \eta)$ phase-diagram, at a reduced computational cost.

Testing demonstrated high accuracy in sequence prediction as far as considering interpolation within the training sequence length, even if considering a very wide range of lattice misfit and average compositions. Generalization to long evolutions (up to five times the sampled one in training), on the other hand, leads to deviations which cause the loss of pixel-by-pixel match but still retain a good description of global properties in terms of ensemble averages, thus making the model useful also for sequences as long as 9 times the training one. Notably, the NN prediction errors have been found equivalent to noise injection into the true PF dynamics, showing that the deep learning model is consistent with moderate levels of perturbations, responsible for bifurcation events (e.g. domain pinching or coalescence), altering the specific connectivity of segregated domains but not global properties. The reliability of the NN as a surrogate of the generating PF simulations has been demonstrated by recovering the theoretical phase diagram.

The proposed framework is versatile and applicable to the broader class of continuum models beyond the specific Cahn–Hilliard equation, which was chosen since as particularly well-suited to generate complex and diverse patterns as a function of a




FIGURE 8 | (a) Phase diagram reconstruction using the trained NN. Each point corresponds to a $(\langle\phi\rangle, \eta)$ combination sampled by a NN simulation starting from a random Perlin noise configuration and evolving up to 100. The analytical spinodal line is reported as a solid black line for reference, and points are categorized according to the match of the NN prediction with the corresponding PF ground truth. (b) Representative failure cases at 200τ where the model incorrectly predicts phase separation or introduces artifacts (marked by black crosses in (a)). (c) Comparison between NN and PF profiles at 200τ for two evolutions starting from the same Perlin noise profile but with different amplitudes. The one with larger perturbations (left) undergoes phase separation while the other (right) remains mixed. Both cases are captured by the NN model. Simulations are performed on 128×128 domains.

governing parameter and, as such, the optimal testing ground for severely assessing the NN architecture properties. The low-computational cost of the NN, combined with the possibility of training the NN model on a less demanding small-size domain while still applying it to larger regions, here demonstrated for a $16\times$ upscaling, could enable the study of systems where no analytical approach exists and/or conventional numerical simulations are highly demanding.

Finally, a potential application of this workflow lies in its use on real experimental images, offering a pathway to bridge the gap between simulation and experiment. By incorporating measurements, this approach could further refine predictions, enabling targeted design and optimization in materials science and semiconductor physics. In fact, the direct, on-the-fly application of ML techniques to experimental data has already begun to emerge [15] and is likely to become relevant in the coming years.

Acknowledgements

A.F., D.L. contributed equally to this work.

Open access publishing facilitated by Università degli Studi di Milano-Bicocca, as part of the Wiley - CRUI-CARE agreement.

Funding

F.M., R.B., and D.L. acknowledge financial support from ICSC—Centro Nazionale di Ricerca in High-Performance Computing, Big Data and Quantum Computing, funded by the European Union—NextGenerationEU. D.L. acknowledges financial support from ICSC SPOKE 7 CNR, project INNOVATOR, CUP B93C22000620006, CN00000013.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data supporting the findings of this study are available upon reasonable request from the authors. The dataset used to train and validate the model are openly available in Materials Cloud Archive at <https://doi.org/10.24435/materialscloud:b4-9t>, reference number 2025.117. The code used to train the NN model is freely available on GitHub at <https://github.com/dlanzo/CRANE>.

References

1. Y. Kok, X. P. Tan, P. Wang, et al., “Anisotropy and Heterogeneity of Microstructure and Mechanical Properties in Metal Additive Manufacturing: A Critical Review,” *Materials & Design* 139 (2018): 565–586, <https://doi.org/10.1016/j.matdes.2017.11.021>.
2. X. Li, L. Lu, J. Li, X. Zhang, and H. Gao, “Mechanical Properties and Deformation Mechanisms of Gradient Nanostructured Metals and Alloys,” *Nature Reviews Materials* 5 (2020): 706–723, <https://doi.org/10.1038/s41578-020-0212-2>.
3. N. Provatas and K. Elder, *Phase-Field Methods in Materials Science and Engineering*, 1st edition (Wiley, 2010), <https://doi.org/10.1002/9783527631520>.
4. P. Lyngby and K. S. Thygesen, “Data-Driven Discovery of 2D Materials by Deep Generative Models,” *npj Computational Materials* 8 (2022): 232, <https://doi.org/10.1038/s41524-022-00923-3>.
5. Y. Zhao, E. M. D. Siriwardane, Z. Wu, et al., “Physics Guided Deep Learning for Generative Design of Crystal Materials With Symmetry Constraints,” *npj Computational Materials* 9 (2023): 38, <https://doi.org/10.1038/s41524-023-00987-9>.
6. E. T. Chenebuah, M. Nganbe, and A. B. Tchagang, “A Deep Generative Modeling Architecture for Designing Lattice-Constrained Perovskite Materials,” *npj Computational Materials* 10 (2024): 198, <https://doi.org/10.1038/s41524-024-01381-9>.
7. C. Karpovich, E. Pan, and E. A. Olivetti, “Deep Reinforcement Learning for Inverse Inorganic Materials Design,” *npj Computational Materials* 10 (2024): 287, <https://doi.org/10.1038/s41524-024-01474-5>.
8. L. Ward, A. Agrawal, A. Choudhary, and C. Wolverton, “A General-Purpose Machine Learning Framework for Predicting Properties of Inorganic Materials,” *npj Computational Materials* 2 (2016): 16028, <https://doi.org/10.1038/npjcompumats.2016.28>.
9. A. Dunn, Q. Wang, A. Ganose, D. Dopp, and A. Jain, “Benchmarking Materials Property Prediction Methods: the Matbench Test Set and Automatminer Reference Algorithm,” *npj Computational Materials* 6 (2020): 138, <https://doi.org/10.1038/s41524-020-00406-3>.
10. P.-P. De Breuck, G. Hautier, and G.-M. Rignanese, “Materials Property Prediction for Limited Datasets Enabled by Feature Selection and Joint Learning With MODNet,” *npj Computational Materials* 7 (2021): 83, <https://doi.org/10.1038/s41524-021-00552-2>.
11. Y. Ren, K. Zhang, Y. Zhou, and Y. Cao, “Phase-Field Simulation and Machine Learning Study of the Effects of Elastic and Plastic Properties of Electrodes and Solid Polymer Electrolytes on the Suppression of Li Dendrite Growth,” *ACS Applied Materials and Interfaces* 14 (2022): 30 658–30 671, <https://doi.org/10.1021/acsami.2c03000>.
12. A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, “Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, Without the Electrons,” *Physical Review Letters* 104 (2010): 136403, <https://doi.org/10.1103/PhysRevLett.104.136403>.
13. J. Behler, “Atom-Centered Symmetry Functions for Constructing High-Dimensional Neural Network Potentials,” *The Journal of Chemical Physics* 134 (2011): 074106, <https://doi.org/10.1063/1.3553717>.
14. A. Fantasia, F. Rovaris, O. Abou El Kheir, et al., “Development of a Machine Learning Interatomic Potential for Exploring Pressure-Dependent Kinetics of Phase Transitions in Germanium,” *The Journal of Chemical Physics* 161 (2024): 014110, <https://doi.org/10.1063/5.0214588>.
15. C. Shen, W. Zhan, S. Pan, et al., “Real-Time Self-Optimization of Quantum Dot Laser Emissions During Machine Learning-Assisted Epitaxy,” *Advanced Science* 12 (2025): 2503059, <https://doi.org/10.1002/adv.202503059>.
16. D. Lanzoni, F. Montalenti, and R. Bergamaschini, “Deep Learning for Simulating the Evolution of Condensed Matter Systems at the Continuum Scale: Methods and Applications,” *Journal of Physics: Condensed Matter* 37 (2025): 403003, <https://doi.org/10.1088/1361-648X/ae096d>.
17. D. Montes De Oca Zapiain, J. A. Stewart, and R. Dingreville, “Accelerating Phase-Field-Based Microstructure Evolution Predictions via Surrogate Models Trained by Machine Learning Methods,” *npj Computational Materials* 7 (2021): 3, <https://doi.org/10.1038/s41524-020-00471-8>.
18. C. Hu, S. Martin, and R. Dingreville, “Accelerating Phase-Field Predictions via Recurrent Neural Networks Learning the Microstructure Evolution in Latent Space,” *Computer Methods in Applied Mechanics and Engineering* 397 (2022): 115128, <https://doi.org/10.1016/j.cma.2022.115128>.
19. O. Ahmad, R. Maurya, R. Mukherjee, and S. Bhowmick, “Integrated Phase Field and Machine Learning Study of Microstructure Evolution During Interface-Controlled Spinodal Decomposition,” *Solid State Phenomena* 357 (2024): 101–106, <https://doi.org/10.4028/p-6W41X1>.
20. D. Lanzoni, F. Rovaris, L. Martín-Encinar, A. Fantasia, R. Bergamaschini, and F. Montalenti, “Accelerating Simulations of Strained-Film Growth by Deep Learning: Finite Element Method Accuracy Over Long Time Scales,” *APL Machine Learning* 2 (2024): 036108, <https://doi.org/10.1063/5.0221363>.
21. K. Alhada-Lahbabi, D. Deleruyelle, and B. Gautier, “Machine Learning Surrogate for 3d Phase-Field Modeling of Ferroelectric Tip-Induced Electrical Switching,” *npj Computational Materials* 10 (2024): 197, <https://doi.org/10.1038/s41524-024-01375-7>.
22. J. Y. Choi, T. Xue, S. Liao, and J. Cao, “Accelerating Phase-Field Simulation of Three-Dimensional Microstructure Evolution in Laser Powder Bed Fusion With Composable Machine Learning Predictions,” *Additive Manufacturing* 79 (2024): 103938, <https://doi.org/10.1016/j.addma.2023.103938>.
23. K. Yang, Y. Cao, Y. Zhang, et al., “Self-Supervised Learning and Prediction of Microstructure Evolution With Convolutional Recurrent Neural Networks,” *Patterns* 2 (2021): 100243, <https://doi.org/10.1016/j.patter.2021.100243>.
24. A. A. Kazemzadeh Farizhandi and M. Mamivand, “Spatiotemporal Prediction of Microstructure Evolution With Predictive Recurrent Neural Network,” *Computational Materials Science* 223 (2023): 112110, <https://doi.org/10.1016/j.commatsci.2023.112110>.
25. P. Wu, A. S. Iquebal, and K. Ankit, “Emulating Microstructural Evolution During Spinodal Decomposition Using a Tensor Decomposed Convolutional and Recurrent Neural Network,” *Computational Materials Science* 224 (2023): 112187, <https://doi.org/10.1016/j.commatsci.2023.112187>.
26. D. Lanzoni, A. Fantasia, R. Bergamaschini, O. Pierre-Louis, and F. Montalenti, “Extreme Time Extrapolation Capabilities and Thermodynamic Consistency of Physics-Inspired Neural Networks for the 3D Microstructure Evolution of Materials via Cahn–Hilliard Flow,” *Machine Learning: Science and Technology* 5 (2024): 045017, <https://doi.org/10.1088/2632-2153/ad8532>.
27. G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, “Physics-Informed Machine Learning,” *Nature Reviews Physics* 3 (2021): 422–440, <https://doi.org/10.1038/s42254-021-00314-5>.
28. J. He, X. Li, and H. Zhu, “An Adaptive Discrete Physics-Informed Neural Network Method for Solving the Cahn–Hilliard Equation,” *Engineering Analysis with Boundary Elements* 155 (2023): 826–838, <https://doi.org/10.1016/jenganabound.2023.06.031>.
29. S. Fan, A. L. Hitt, M. Tang, B. Sadigh, and F. Zhou, “Accelerate Microstructure Evolution Simulation Using Graph Neural Networks With Adaptive Spatiotemporal Resolution,” *Machine Learning: Science and Technology* 5 (2024): 025027, <https://doi.org/10.1088/2632-2153/ad3e4b>.

30. J. K. Gupta and J. Brandstetter, "Towards Multi-Spatiotemporal-Scale Generalized PDE Modeling," 2022, arXiv:2209.15616, <https://doi.org/10.48550/arXiv.2209.15616>.
31. V. Oommen, K. Shukla, S. Desai, R. Dingreville, and G. E. Karniadakis, "Rethinking Materials Simulations: Blending Direct Numerical Simulations With Neural Operators," *npj Computational Materials* 10 (2024): 145, <https://doi.org/10.1038/s41524-024-01319-1>.
32. D. Lanzoni, M. Albani, R. Bergamaschini, and F. Montalenti, "Morphological Evolution via Surface Diffusion Learned by Convolutional, Recurrent Neural Networks: Extrapolation and Prediction Uncertainty," *Physical Review Materials* 6 (2022): 103801, <https://doi.org/10.1103/PhysRevMaterials.6.103801>.
33. J. W. Cahn, "On Spinodal Decomposition," *Acta Metallurgica* 9 (1961): 795–801, [https://doi.org/10.1016/0001-6160\(61\)90182-1](https://doi.org/10.1016/0001-6160(61)90182-1).
34. J. Langer and M. Bar-on, "Theory of Early-Stage Spinodal Decomposition," *Annals of Physics* 78 (1973): 421–452, [https://doi.org/10.1016/0003-4916\(73\)90266-2](https://doi.org/10.1016/0003-4916(73)90266-2).
35. P. Fratzl, O. Penrose, and J. L. Lebowitz, "Modeling of Phase Separation in Alloys With Coherent Elastic Misfit," *Journal of Statistical Physics* 95 (1999): 1429–1503, <https://doi.org/10.1023/A:1004587425006>.
36. K. Rundman and J. Hilliard, "Early Stages of Spinodal Decomposition in an Aluminum-Zinc Alloy," *Acta Metallurgica* 15 (1967): 1025–1033, [https://doi.org/10.1016/0001-6160\(67\)90266-0](https://doi.org/10.1016/0001-6160(67)90266-0).
37. W. Xu, Y. Zhong, X. Li, and K. Lu, "Stabilizing Supersaturation With Extreme Grain Refinement in Spinodal Aluminum Alloys," *Advanced Materials* 36 (2024): 2303650, <https://doi.org/10.1002/adma.202303650>.
38. D. M. Collins, N. D'Souza, C. Panwisawas, et al., "Spinodal Decomposition Versus Classical γ' Nucleation in a Nickel-Base Superalloy Powder: An In-Situ Neutron Diffraction and Atomic-Scale Analysis," *Acta Materialia* 200 (2020): 959–970, <https://doi.org/10.1016/j.actamat.2020.09.055>.
39. H. Kim and P. C. McIntyre, "Spinodal Decomposition in Amorphous Metal-Silicate Thin Films: Phase Diagram Analysis and Interface Effects on Kinetics," *Journal of Applied Physics* 92 (2002): 5094–5102, <https://doi.org/10.1063/1.1510590>.
40. M. Ban, D. Woo, J. Hwang, S. Kim, and J. Lee, "Spinodal Decomposition-Driven Structural Hierarchy of Mesoporous Inorganic Materials for Energy Applications," *Accounts of Chemical Research* 56 (2023): 3428–3440, <https://doi.org/10.1021/acs.accounts.3c00524>.
41. J. S. Higgins and J. T. Cabral, "A Thorny Problem? Spinodal Decomposition in Polymer Blends," *Macromolecules* 53 (2020): 4137–4140, <https://doi.org/10.1021/acs.macromol.0c00581>.
42. S. Kumar, S. Tan, L. Zheng, and D. M. Kochmann, "Inverse-Designed Spinodoid Metamaterials," *npj Computational Materials* 6 (2020): 73, <https://doi.org/10.1038/s41524-020-0341-6>.
43. L. Zheng, S. Kumar, and D. M. Kochmann, "Data-Driven Topology Optimization of Spinodoid Metamaterials With Seamlessly Tunable Anisotropy," *Computer Methods in Applied Mechanics and Engineering* 383 (2021): 113894, <https://doi.org/10.1016/j.cma.2021.113894>.
44. Z. Wang, R. Dabaja, L. Chen, and M. Banu, "Machine Learning Unifies Flexibility and Efficiency of Spinodal Structure Generation for Stochastic Biomaterial Design," *Scientific Reports* 13, no. 1 (2023): 5414, <https://doi.org/10.1038/s41598-023-31677-7>.
45. H. Nishimori and A. Onuki, "Pattern Formation in Phase-Separating Alloys With Cubic Symmetry," *Physical Review B* 42 (1990): 980–983, <https://doi.org/10.1103/PhysRevB.42.980>.
46. J. Zhu, L.-Q. Chen, and J. Shen, "Morphological Evolution During Phase Separation and Coarsening With Strong Inhomogeneous Elasticity," *Modelling and Simulation in Materials Science and Engineering* 9 (2001): 499, <https://doi.org/10.1088/0965-0393/9/6/303>.
47. Z. Rao, B. Dutta, F. Körmann, et al., "Beyond Solid Solution High-Entropy Alloys: Tailoring Magnetic Properties via Spinodal Decomposition," *Advanced Functional Materials* 31 (2021): 2007668, <https://doi.org/10.1002/adfm.202007668>.
48. J. W. Cahn, "On Spinodal Decomposition in Cubic Crystals," *Acta Metallurgica* 10 (1962): 179–183, [https://doi.org/10.1016/0001-6160\(62\)90114-1](https://doi.org/10.1016/0001-6160(62)90114-1).
49. S. Fetni, T. Q. D. Pham, T. V. Hoang, et al., "Capabilities of Auto-Encoders and Principal Component Analysis of the Reduction of Microstructural Images; Application on the Acceleration of Phase-Field Simulations," *Computational Materials Science* 216 (2023): 111820, <https://doi.org/10.1016/j.commatsci.2022.111820>.
50. S. Chauhan, S. Mandal, V. Yadav, P. K. Jaiswal, M. Priya, and M. D. Shrimali, "Machine Learning Based Prediction of Phase Ordering Dynamics," *Chaos: An Interdisciplinary Journal of Nonlinear Science* 33, no. 6 (2023): 061103, <https://doi.org/10.1063/5.0156611>.
51. V. Yadav, M. Priya, M. D. Shrimali, and P. K. Jaiswal, "Graph Neural Network for Prediction of Phase-Ordering Kinetics," *Chaos: An Interdisciplinary Journal of Nonlinear Science* 35, no. 6 (2025): 061101, <https://doi.org/10.1063/5.0273728>.
52. H. Garcke, S. Maier-Paape, and U. Weikard, "Spinodal Decomposition in the Presence of Elastic Interactions," in *Geometric Analysis and Nonlinear Partial Differential Equations* eds., S. Hildebrandt and H. Karcher, (Springer, 2003), 603–635, https://doi.org/10.1007/978-3-642-55627-2_32.
53. A fixed temperature is here assumed so that φ can be considered a rescaling of the full composition range c to the $[c_1, c_2]$ one of the two phases at that temperature, i.e. $\varphi = \frac{c-c_1}{c_2-c_1}$.
54. T. Mura, *Micromechanics of Defects in Solids*, ser. Mechanics of Elastic and Inelastic Solids, eds., S. Nemat-Nasser and G. Æ. Oravas, vol. 3 (Springer Netherlands, 1987), <https://doi.org/10.1007/978-94-009-3489-4>.
55. A. Khachaturyan, *Theory of Structural Transformations in Solids*, ser. Dover Books on Engineering Series (Dover Publications, 2008).
56. Y. Wang, L.-Q. Chen, and A. Khachaturyan, "Kinetics of Strain-Induced Morphological Transformation in Cubic Alloys With a Miscibility Gap," *Acta Metallurgica et Materialia* 41 (1993): 279–296, [https://doi.org/10.1016/0956-7151\(93\)90359-Z](https://doi.org/10.1016/0956-7151(93)90359-Z).
57. J. W. Cahn, "On Spinodal Decomposition in Cubic Crystals," *Acta Metallurgica* 10 (1962): 179–183, [https://doi.org/10.1016/0001-6160\(62\)90114-1](https://doi.org/10.1016/0001-6160(62)90114-1).
58. K. Perlin, "An Image Synthesizer," *SIGGRAPH Computer Graphics* 19 (1985): 287–296, <https://doi.org/10.1145/325165.325247>.
59. S. Schubert, P. Neubert, J. Poschmann, and P. Protzel, "Circular Convolutional Neural Networks for Panoramic Images and Laser Data," in *2019 IEEE Intelligent Vehicles Symposium (IV)* (IEEE, 2019), 653–660, <https://doi.org/10.1109/IVS.2019.8813862>.
60. J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling," 2014, arXiv:1412.3555, <https://doi.org/10.48550/ARXIV.1412.3555>.
61. M. Salvalaglio, R. Backofen, R. Bergamaschini, F. Montalenti, and A. Voigt, "Faceting of Equilibrium and Metastable Nanostructures: A Phase-Field Model of Surface Diffusion Tackling Realistic Shapes," *Crystal Growth & Design* 15 (2015): 2787–2794, <https://doi.org/10.1021/acs.cgd.5b00165>.
62. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," 2014, arXiv:1412.6980, <https://doi.org/10.48550/ARXIV.1412.6980>.
63. Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum Learning," in *Proceedings of the 26th Annual International Conference on Machine Learning* (Montreal Quebec Canada: ACM, 2009), 41–48, <https://doi.org/10.1145/1553374.1553380>.
64. N. Dalal and B. Triggs, "Histograms of Oriented Gradients for Human Detection," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, 1 (2005): 886–893, <https://doi.org/10.1109/CVPR.2005.177>.

65. N. Sarkar and B. Chaudhuri, "An Efficient Differential Box-Counting Approach to Compute Fractal Dimension of Image," *IEEE Transactions on Systems, Man, and Cybernetics* 24, no. 1 (1994): 115–120, <https://doi.org/10.1109/21.259692>.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Supporting File: adts70330-sup-0001-SuppMat.pdf.