

Morphological evolution via surface diffusion learned by convolutional, recurrent neural networks: Extrapolation and prediction uncertainty

Daniele Lanzoni ^{*}, Marco Albani , Roberto Bergamaschini , and Francesco Montalenti 
 Materials Science Department, University of Milano-Bicocca, Via R. Cozzi 55, I-20125 Milano, Italy



(Received 17 June 2022; accepted 13 September 2022; published 3 October 2022)

We use a convolutional, recurrent neural network approach to learn morphological evolution driven by surface diffusion. To this aim we first produce a training set using phase field simulations. Intentionally, we insert in such a set only relatively simple, isolated shapes. After proper data augmentation, training and validation, the model is shown to correctly predict also the evolution of previously unobserved morphologies and to have learned the correct scaling of the evolution time with size. Importantly, we quantify prediction uncertainties based on a bootstrap-aggregation procedure. The latter proved to be fundamental in pointing out high uncertainties when applying the model to more complex initial conditions (e.g., leading to the splitting of high aspect-ratio individual structures). The automatic smart augmentation of the training set and design of a hybrid simulation method are discussed.

DOI: [10.1103/PhysRevMaterials.6.103801](https://doi.org/10.1103/PhysRevMaterials.6.103801)

I. INTRODUCTION

The intensive use of machine learning (ML) methods in recent years has proven a valuable tool in materials science and condensed matter physics [1–5]. ML techniques have proven to be comparable or superior to traditional methods in tasks such as phase identification [6], experimental images processing [7] and device property estimation [8]. One of the areas that is attracting the most interest is the application of regression techniques to the development of interatomic potentials [9,10]. The leverage of ML tools such as neural networks (NN) [11–14] and Gaussian processes [15–19] allowed for the construction of force fields far more accurate than classical semi-empirical potentials which however can be evaluated at a fraction of the computational effort of full *ab initio* calculations [20]. This computational speed-up makes it possible to tackle systems intractable by standard density functional theory approaches [21,22].

Despite these improvements, the possibility of simulating complex mesoscale phenomena is still a hard task, both for the spatial and the time scales involved. In this respect, resorting to continuum scale methods is often the only solution.

The smoothing of discrete atomistic behavior in continuum fields typically leads to the formulation of evolution laws in the form of partial differential equations (PDEs). It should nevertheless be remembered that even within continuum approaches the computational cost can be remarkable. For instance, fine meshes are sometimes required by finite-element method (FEM) solvers, and/or stiff equations involving high-order derivatives might be involved (typical examples include Navier-Stokes equations in fluid dynamics [23] or degenerate Cahn-Hilliard equations in the presence of strong anisotropy [24]). In the last few years, indeed, the interest in applying ML techniques for treating PDEs problems grew substantially

[25–30]. Several of these new methods exploit the possibility of encoding spatial information of PDE solutions in images (when on 2D domains) or 3D volumetric data. This permits to leverage well-established ML techniques which have proven extremely successful in computer vision tasks. Among the different ML approaches that can be applied to predict image sequences, we find particularly intriguing the convolutional, recurrent neural network (CRNN) approach introduced in 2015 by Shi *et al.* [31] who modified the original long-short term memory (an instance of recurrent neural networks) replacing fully connected units with convolution operations. This variant, originally applied in the context of precipitation nowcasting, allows for pixel-wise prediction of the next image in a sequence, exploiting an efficient synergy of computer vision and recurrent NN approaches.

More recently, Young *et al.* [30] applied CRNNs (including the refinement introduced in Ref. [32]) to some materials-science relevant phenomena, notably spinodal decomposition and dendritic growth as obtained from phase-field models. In this paper, instead, we apply CRNNs to another key problem in materials science, i.e., morphological evolution via surface diffusion [33,34]. Additionally, we implement a simple but effective scheme based on an ensemble model capable of providing an on-the-fly estimate of the prediction uncertainty. This allows one to understand under which conditions the model reliably handles extrapolation and when, instead, generated sequences are untrustworthy. With this respect, after building a convenient dataset based on phase-field simulations, we train, validate, and test the CRNN model on progressively harder tasks also involving heavy extrapolation.

The paper is organized as follows. In Sec. II, we present the phase-field model of surface diffusion and the selection criteria of the evolution produced to construct the datasets. Section III is devoted to the discussion of the specific implementation of our NN architectures, details the training procedure and presents how the prediction uncertainty estimation scheme is implemented. In Sec. IV, we report

^{*}Corresponding author: d.lanzoni@campus.unimib.it

performances on challenging settings, involving, for example, generalization to longer evolutions and more complex initial conditions. Finally, Sec. V summarizes the current work and presents some of the possible future extensions.

II. DATA GENERATION

A simple curvature-driven surface diffusion model [33,34] is considered for the evolution of a two-dimensional profile Γ . Local material fluxes J are set along the surface coordinate s as a function of the surface gradients of the profile curvature κ so that $\vec{J}_s = -M_s \nabla_s \kappa$, with M_s a mobility constant. The resulting PDE problem is conveniently solved by a phase-field approach, tracking the profile implicitly by means of an order parameter φ , equal to 1 and 0 in the inner and outer domains delimited by the curve Γ and with a diffuse interface along it, with finite, possibly small thickness ϵ . By convention, the sharp profile is identified as the $\varphi = 0.5$ isoline. The surface dynamics can then be rewritten in terms of the time evolution of the phase-field itself, resulting in the Cahn-Hilliard (CH) equation set [34]:

$$\begin{aligned} \frac{\partial \varphi}{\partial t} &= \nabla \cdot M(\varphi) \nabla \kappa \\ g(\varphi) \kappa &= -\epsilon \nabla^2 \varphi + \frac{1}{\epsilon} W'(\varphi) \end{aligned} \quad (1)$$

with $W(\varphi) = 18\varphi^2(1 - \varphi)^2$ a double-well potential with minima in the two “bulk phases” $\varphi = 0$ and 1 and $g(\varphi) = (5/3)W(\varphi)$ a stabilizing function introduced for numerical convenience [35–37]. The degenerate mobility function $M(\varphi) = M_s(36/\epsilon)\varphi^2(1 - \varphi)^2$ is used to restrict the dynamics within the interface region, thus returning the proper surface diffusion behavior in the sharp-interface limit for $\epsilon \rightarrow 0$. The numerical solution of the PDE system (1) is performed by finite-element method, using the adaptive multidimensional simulations (AMDiS) [38,39] toolbox. Both a semi-implicit time-integration scheme and adaptive local mesh refinement are exploited for numerical efficiency. The length scale in the simulations has been scaled with respect to the interface thickness ϵ . A 400×400 domain is considered to embed all test geometries with finest mesh resolution of $\sim 1/7$ to properly represent the diffuse interface. Time is scaled with respect to the surface mobility coefficient M_s . A fixed time step of 10^{-5} is set and a total evolution time of 0.6 is considered for all simulations, as sufficient to reach a convex shape for all initial conditions considered. The choice of a fixed integration time step over more sophisticated adaptive schemes has been chosen since the CRNN structure considers sequence elements as equally spaced in time. This parameter set has been tested to ensure numerical convergence and stability for all geometries considered in the present study, offering the best tradeoff between accuracy and execution speed, the latter being important to produce a sufficiently large dataset for the ML training. As an example, the evolution of a simple rectangular domain is shown in Fig. 1(a).

For the present study, the dataset was constructed considering only prototypical shapes which could be obtained by rectangles and simple intersection of rectangles: simple bars, symmetrical crosses, asymmetrical crosses and 45 crosses. An example for all possible initial topologies is reported

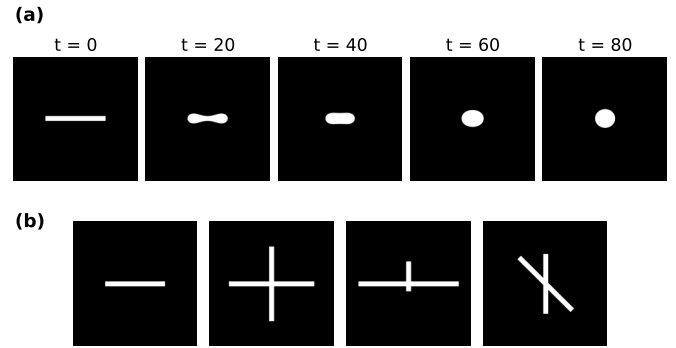


FIG. 1. (a) Prototypical evolution of a rectangular domain by isotropic surface diffusion. Intermediate steps present a characteristic hourglass shape, while the final state is circular. (b) Examples of shapes present in the training set (from left to right: Simple bars, symmetrical crosses, asymmetrical crosses and 45 crosses).

in Fig. 1(b). Different dataset elements were constructed by scanning rectangles aspect ratios between 1 and 20, keeping the short side at 8 units of ϵ . The limitation in the aspect ratios prevents domain splitting due to Rayleigh-like instability [40]. All the constraints in the dataset construction were intentional, as one of the objectives of our work was to understand to which extent a machine learning algorithm is able to generalize to unobserved configurations given a simple and controlled set of examples.

III. MACHINE LEARNING METHODS

A. Convolutional, recurrent neural network

As stated in the introduction, PDEs in 2D can be conveniently encoded as images. This approach is particularly straightforward in the case of phase field models, since the value of the order parameter φ can be restricted between 0 and 1. Successive states of the morphological evolution of an initial shape under surface diffusion, therefore, can be represented by sequences of gray-scale images.

In recent years, convolutional neural networks (CNN) have proven to be one of the most effective machine learning architectures for tasks involving data with spatial structures. Invented in the late eighties [41,42], they are at present one of the standard tools in tasks such as audio data analysis [43,44], image classification [45,46] and inference from volumetric data [47,48]. The main building block of CNNs consists in the convolution operation. Starting from raw data, hierarchically more complex “abstract” features may be extracted by successive application of convolutional kernels (typically in the form of square matrices) containing learnable parameters. Recently, convolutional structures revealed their potential also in condensed matter and materials science research. For instance, they have proven effective in analyzing experimental images [7], in accelerating or reducing the dimensionality of complex continuum model simulations [25–28] and in tackling mesoscale simulations [30].

The CNNs’ capability of identifying spatial correlations is effective in tasks involving PDEs because they encode several mathematical properties by design [49,50]. First, convolutions are local operations: pixels which are farther apart than the

kernel size are not allowed to “communicate.” If CNNs are used to approximate PDEs, this directly translates in the locality of (partial) differential operators. In the present case, for instance, the isotropic Cahn-Hilliard equation predicting the evolution of a domain only depends on its local curvature.

A second important property of convolution transformations is that they are equivariant with respect to the group of spatial translations. This is beneficial to the training procedure, both in terms of training time and quality of the final result, as the NN does not need to learn the underlying symmetry group. Correspondingly, predicted evolution will not depend on the absolute position of the initial shape by construction. The use of convolution operators also allows for a straightforward implementation of periodic boundary conditions, which can be simply obtained by using a so-called “circular padding” [51,52].

Last, if all transformations in the NN are either convolutions or pixelwise (fully convolutional NNs), the model can operate on images of arbitrary size. In practice, this means that the CNN can be trained on images representing small simulation cells and, once parameters are optimized, the learned evolution rules can be applied on much bigger domains, considering the pixel dimension as fixed [49,50]. This allows one to tackle mesoscale systems with simulation cells several times the size of the original dataset, an advantage clearly demonstrated in [30].

While CNNs are suited for dealing with local spatial correlations, the goal of the current work is to approximate morphological time evolution. This kind of task involves analyzing and generating temporal sequences. In this field, RNNs emerged as one of the most effective machine learning strategies. Recurrent structures usually exploit a hidden state, which can be considered as a sort of “memory” of previous states of the system and is capable of capturing time correlations [53]. There are several possible implementations for RNNs, but in the last years long-short term memory (LSTM, [54]) and gated recurrent units (GRU, [55]) have emerged as the most used and effective. In our case, we opted for the latter, as GRUs often offer performances comparable to LSTMs with a lower number of parameters [56], which translates to faster evaluation and a reduced possibility of overfitting.

Similarly to fully convolutional NNs, which can transform images of any size, RNNs can elaborate sequences of arbitrary length. This means that a RNN based model can use a partial evolution of any length (possibly a single initial state) as an input and then generate sequences of arbitrary length as well [49,50].

As the objective of the current work is approximating the time evolution of systems described by PDEs, the spatial features learning capabilities of CNNs and the temporal ones of RNNs should be used simultaneously. Following Refs. [31,57], we implemented a convolutional GRU (ConvGRU) cell as the building block of our NN structure. The main idea consists in having a hidden state which is itself structured as a (multiple channel) image and in replacing linear transformations of regular GRUs with convolutions. The result is a neural network capable of learning spatiotemporal correlations on a pixel level. Notice that our approach, while similar to the one of Ref. [30], does not require convolutions involving temporal variables, thus using a smaller number of

parameters. Additionally, we note that, from a purely technical point of view, extension of the present approach to three spatial dimensions is straightforward, but the actual problem stems in GPU requirements and in the effort needed to generate the dataset.

Here, for the sake of reproducibility, we list all hyperparameters. The paragraph can be skipped by uninterested readers. The overall CRNN structure used in the current work is composed of two stacked ConvGRU cells, for a total of $\sim 2.8 \times 10^5$ parameters. Both hidden states have 35 channels, convolution kernels are set to 5×5 , stride 1 (i.e., the convolution runs over all pixels), and circular padding encoding periodic boundary conditions by construction are used. The hidden layer is converted to the next state by a 3×3 convolution, stride 1 and circular padding followed by a sigmoid activation function reducing the output image in the range $[0,1]$. This set of hyperparameters produced the most satisfying results while containing computational costs in our tests.

Pooling and batch normalization are common techniques used in image processing [49,50], to reduce the dimensionality of the feature map and re-scale the intermediate representation respectively. Preliminary tests showed a negligible impact on the quality of the learned models for the present, simple application. Therefore they were not used in the current work. It is still worth noticing that pooling layers, due to their intrinsic capability of compressing the spatial resolution of the representations, may become important for dynamics involving long-range interactions. Composition of pooling, convolutions and transposed convolutions may indeed allow the flow of information from and to pixels which are farther apart than the kernel size, similarly to a U-net architecture [58].

We conclude this section by emphasizing the fact that all computational operations involved in a trained CRNN are typically much cheaper computationally than a FEM solver (by orders of magnitude for all examples discussed in the paper). Additionally, they can easily be implemented on parallel machines (or GPUs). Lastly, the time steps between one state and the following in the evolution obtained by machine learning is much longer than the one used in traditional schemes. All these advantages contribute in compressing the computational times of the predicted evolution, allowing to tackle configurations which would be inconvenient to simulate by traditional methods (e.g., long-time evolutions, large, or complex systems).

B. Equation symmetries and training

As specified in the previous section, one of the main advantages of CRNN is the encapsulation of some symmetries of the underlying equation. In particular, Cahn-Hilliard equation is equivariant with respect to spatial translations i.e., if $\varphi(\vec{x}, t)$ is a solution of the equation with a given initial condition $\varphi_0(\vec{x})$, then $\varphi(\vec{x} + \vec{R}, t)$ is the solution given the initial condition $\varphi_0(\vec{x} + \vec{R})$ for any vector \vec{R} . This property is maintained by convolution operations. Moreover, the CH equation is local, as the evolution of the phase field at some point \vec{x} only depends on the value of φ and its spatial derivatives at that point. In a picture in which convolutions provide finite difference approximations of derivatives, there is a finite kernel size

which is capable of providing the evolution law. Standard convolutions, however, do not encode other symmetries which are nonetheless present in the degenerate Cahn-Hilliard equation. In order to leverage these additional features, we employed two strategies: Modification of the loss function and data augmentation.

In supervised machine learning tasks, the objective of the training procedure consists in the minimization of a loss function L , which can be considered as the deviation between the predicted output and the correct one present in the training set. The base for the loss function used in this work was the commonly used in regression tasks mean squared error loss:

$$L(\theta) = \frac{1}{N_{TS}T} \sum_{i=1}^{N_{rs}} \sum_{t=1}^T \langle (\varphi_i(t) - \hat{\varphi}_i(t|\theta))^2 \rangle_x. \quad (2)$$

Here, θ represents the set of NN parameters, i indexes the training set elements, t indexes the time step and runs from 1 to the total length of the sequence T , φ is the true sequence of the system temporal evolution in the training set, $\hat{\varphi}$ is the time evolution as predicted by the CRNN and $\langle \cdot \rangle_x$ denotes spatial average.

As the CH equation encodes the continuity law [34], it implies the global conservation of the order parameter φ . It must therefore hold that at each timestep the spatial integral of the approximated phase field is equal to the integral predicted by FEM. This condition is equivalent to the equality of the spatial mean of the phase fields, at any time t and for every evolution in the training set i , as the two quantities are proportional. In order to enforce this global conservation condition, the original loss function has been therefore modified by adding an L2 loss term with respect to this spatial average of the phase field:

$$\begin{aligned} \tilde{L}(\theta) = \frac{1}{N_{TS}T} \sum_{i=1}^{N_{rs}} \sum_{t=1}^T \langle (\varphi_i(t) - \hat{\varphi}_i(t|\theta))^2 \rangle_x & \quad (3) \\ + \alpha (\langle \varphi_i(t) \rangle_x - \langle \hat{\varphi}_i(t|\theta) \rangle_x)^2, & \quad (4) \end{aligned}$$

where α sets the penalization strength. In our tests, we observed that a value of $\alpha = 2$ led to the best performances.

Another symmetry to be taken into account is the equivariance of the evolution with respect to rotations of the initial condition. In order to enforce it, we adopted a data augmentation strategy, in which every time sequence is rotated by a random angle before being passed to the neural network. This approach regularizes parameters in convolutional kernels, forcing them to approximately respect rotational equivariance.

A last data augmentation procedure has been used to encode in the neural network the symmetry in the double well potential and mobility function with respect to the transformation $\varphi \rightarrow 1 - \varphi$ (see Sec. II). This means that the negative version of evolution in Fig. 1(a) is still a valid solution to the Cahn-Hilliard equation. The physical interpretation of this is that holes in a material matrix will undergo the same morphological evolution as a material domain with an identical initial shape. Being a discrete symmetry, it is possible to expose the CRNN to all sequences both in the positive (φ) and negative ($1 - \varphi$) version. This way, gradients in the backpropagation algorithm are never computed on positive or negative only variants, reducing biases with respect to this symmetry.

C. Prediction uncertainty estimation

One of the main pitfalls that can be encountered during the training of NNs and other machine learning models is that it is not clear if the produced answer is reliable or the model is being used in extrapolation conditions [59]. This problem is particularly critical in the current case, as the training set was restricted on purpose to simple initial conditions, in order to avoid evolution which require computationally intensive simulations and specialized algorithms (see discussion in Sec. IV). In general, a prediction uncertainty quantification is essential to detect such issues. A common strategy to assess this problem is through testing: model performances are monitored on an ad-hoc test set, representing typical tasks the machine learning model will encounter once deployed [49,50]. In our case, however, the construction of a proper test set presents several challenges. For example, it is not clear how a representative set of initial conditions should be built. Furthermore, such set should comprise also configurations presenting complex interactions which, as it was already mentioned, are inconvenient to simulate in high number because of computational effort.

An alternative solution is leveraging a resampling procedure known as bootstrap aggregation (bagging) [60]. This technique allows for the construction of an ensemble of models, whose prediction capabilities are higher than the individual models' one. In addition, the dispersion of the ensemble predictions represents an estimation of the prediction uncertainty [61].

In our case, we opted for a simple implementation of bagging. Starting from the original training set composed of N sequences, a number of new sets with the same size are generated by random extraction with replacement. These new sets are called bootstrap samples (BS). Extraction with replacement guarantees that BSs, though drawn from the same original set, are different. For each BS an independent model is fitted, thus generating the neural network ensemble. It can be proven from a statistical point of view that, if the learned models present an unstable enough training procedure, the aggregate prediction of the ensemble is more robust and reliable than individual ones [62]. Another key advantage is that variance and/or standard deviation between individual predictions of the models in the ensemble can be calculated. In the current framework, this allows for the definition of an on-the-fly, pixelwise prediction uncertainty, thus the identification of regions and instants in which the predicted evolution is unreliable.

To form the aggregate prediction, we used simple mean. On a qualitative level, when models inside the ensemble provide diverging predictions, aggregation of the individual predictions will produce a gray "halo," as the mean of the predicted values of φ is different from 0 or 1. At the same time, the ensemble standard deviation allows for the definition of a confidence interval on a quantitative level.

In order to maximize the independence in the training procedure, bootstrapping has been applied on both training (~ 1200 sequences) and validation (~ 250 sequences) sets and an ensemble of 15 models was constructed. This value was chosen as a reasonable compromise between the ensemble size, training times and prediction computational costs, the latter two scaling linearly with the number of models.

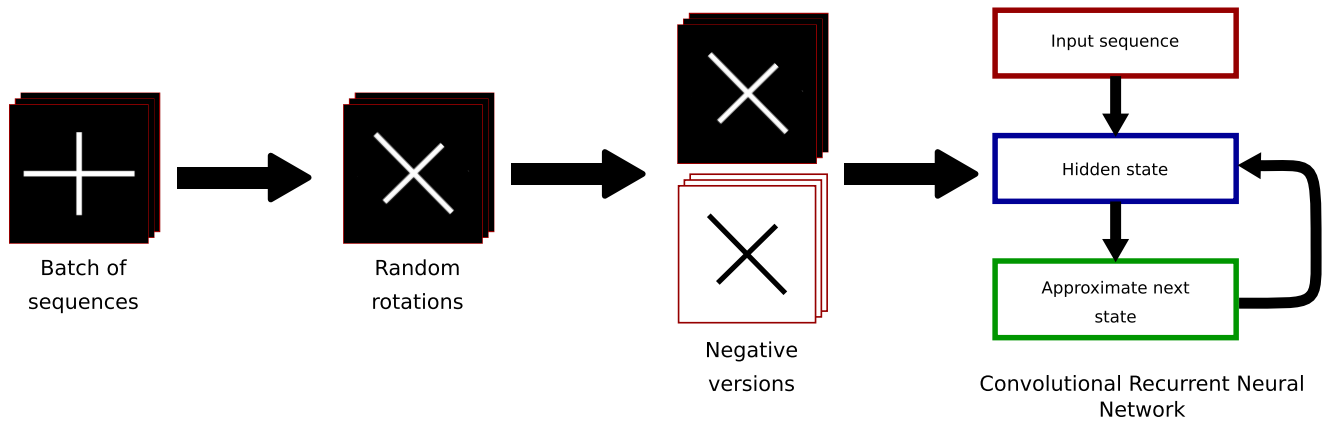


FIG. 2. Scheme of the training procedure. A single parameters optimization step consist in a random extraction of a batch of sequences which are converted to the right resolution size, randomly rotated and passed to the CRNN together with their negative versions.

D. Training, validation, and testing

As anticipated, training has been performed with the modified mean squared error loss \tilde{L} reported in Eq. (3). NNs' parameters were optimized using the standard ADAM algorithm [63], as implemented in the PYTORCH library [64].

Training and validation sets are composed by sequences of 150 snapshots, starting 0, 25 and 50 time steps from the initial state of the system. This choice allows for the exposure of the models to different initial conditions. The time difference between each timestep is 250 time integration steps of the high-fidelity phase-field FEM solver (see Sec. II). Sequences originating from the same initial condition were exclusively inserted in either the training or the validation set, in order to remove correlations which would affect validation loss.

From a practical standpoint, training is summarized in Fig. 2 and proceeds as follows. A batch of sequences is extracted from the training set, converted to a 90×90 pixels format and randomly rotated. The batch is then augmented with the negative version of the evolution through the transformation $\varphi \rightarrow 1 - \varphi$. Next, truncated subsequences are passed to the Convolutional Recurrent structure, whose task is to generate its completion. The truncation point is selected at random for every batch. This choice allows for a compromise between the full, more demanding task of generating a sequence starting from a single initial state and the easier (but far from the actual application mode) task of completing a partial sequence. The produced sequence is then compared with the one in the training set, loss function is evaluated and parameters are updated.

A training plot is reported for an individual model in the bootstrap aggregate in Fig. 3(a). Validation loss is always calculated by providing the model only the initial state of the evolution, as this is the actual task the network is required to tackle with new, unobserved initial conditions. Although noisy, after ~ 75 epochs (in one epoch all elements in the training set are presented once to the NN model) the validation loss settles, indicating the training procedure has converged. No increasing trends can be observed, confirming the absence of overfitting. Oscillations in the loss value as a function of epochs are induced by the randomness of the optimization procedure. The evolution of a validation initial condition as

predicted by the model is reported in figure Fig. 3(b), confirming the model parameters have converged.

As an additional confirmation of the training procedure results, we inspected evolution predicted starting from previously unobserved and qualitatively different Y shapes. Arms of these new initial conditions have the same thickness as rectangles and crosses present in the original dataset. Inspection of the example evolution reported in Fig. 4 reveals that predicted evolution exhibits only negligible deviations from the true one. At the same time, prediction uncertainty is low, confirming that the aggregate model is not in extrapolation conditions even though the total number of frames predicted is 300, double the 150 used during training, suggesting good generalization capabilities to long-time evolutions. We emphasize that the close agreement on this specific test configurations cannot be considered as an indication of performances on general initial conditions, as a single class of shapes has been tested. Our concern, here, is instead to check the ensemble model behavior for progressively harder tasks.

IV. RESULTS AND DISCUSSION

In the following section, we test the capabilities of the presented machine learning approach in predicting long term morphological evolution (i.e., for times much longer than the 150 time steps used in training) and in capturing important scale-dependent behaviors which were not explicitly inserted in the training set. We conclude the section by inspecting generalization capabilities and limitations in the case of complex initial conditions with many interacting shapes.

A. Stationary state stability

One of the main concerns that should be addressed is the stability of the NN prediction for long-time evolution. While a simple test was provided with the evolution of Fig. 4, we perform now a more in depth analysis which also considers an unobserved topology. In particular, the time evolution of an annular shape stationary state configuration is considered. The goal of this test is threefold: the first check consists in verifying that the aggregate model correctly predicts that the hole in the center of the domain is a stable configuration (as the

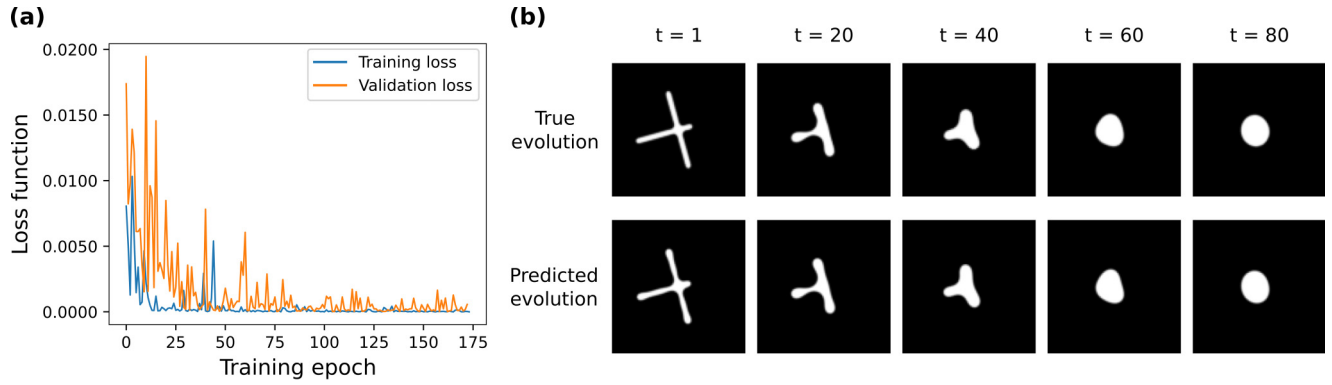


FIG. 3. (a) Training and validation loss value as a function of the epoch number for one model in the ensemble. (b) Example predicted evolution for a validation set element and the corresponding high-fidelity FEM sequence.

curvature and, thus, the chemical potential is uniform); second, we check if the stability is predicted for times exceeding 150 time steps, indicating good time generalization capabilities of the model; third, we verify that prediction uncertainty rises if and as soon as the generated evolution diverges from the correct one.

The evolution and the corresponding aggregate model standard deviation is reported in Figs. 5(top) and 5(bottom), respectively. As it can be clearly observed, the annular configuration is stable for times corresponding to approximately 500 snapshots, more than 3 times the length of the training sequences. This confirms that the NN model is both effective in recognizing this unobserved steady state configurations and in generalizing to sequences longer than the ones provided during training.

After 500 snapshots, a progressive deviation of the evolution, caused by models in the aggregate unphysically filling the hole, can be observed. As soon as this happens, however, a clear increase in the prediction uncertainty is reported by the formation of the halo visible in the predicted φ at $t = 750$. Correspondingly, an increase of the standard deviation in the hole region can be observed, confirming the extrapolation recognition capabilities of our approach.

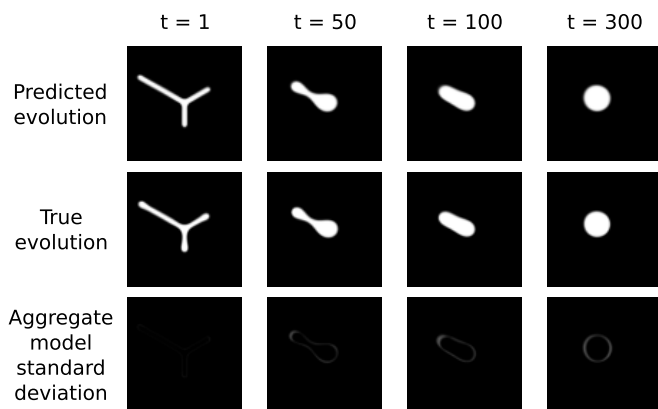


FIG. 4. Example evolution of a Y shape for testing the aggregate model performances (not included in training). The lack of high model variance regions confirms that this example is far from extrapolation conditions. Generated sequences are twice as long as the ones of the training procedure.

B. Size-dependent trends

As reported in Sec. III D, training and validation shapes, as well as the initial test on Y configurations, all shared a common characteristic arm thickness. In order to check the ability of this ML method to replicate scale-dependent phenomena, we decided to analyze to evolution of ellipsoidal domains. Self-similar shapes present a self-similar evolution under the (degenerate) Cahn-Hilliard equation [65–67]. Therefore, ellipsoidal domains will progressively reduce their aspect ratio (AR) starting from the initial value and reaching a circular equilibrium shape with a rate depending on the size of the initial condition.

To check if the CRNN was capable of capturing this trend, we compared the evolution of ellipses with an initial AR of 3 as simulated by FEM solver and by the ML approach. Figure 6(a) reports such comparison. As can be observed, the NN evolution closely follows the high-fidelity phase field one. In Figs. 6(b) and 6(c), the aspect ratio as a function of time is also reported for the smallest and biggest ellipse, as a further proof of consistency. We emphasize again that the number of snapshots in these evolution is greater (more than 3 times) that of the sequences used during training.

The sequence of Fig. 6 also highlights another characteristic of the NN approach. While PF evolutions of ellipses have been run independently, in the ML based prediction the evolution of all domains is calculated simultaneously by placing all of them in the same simulation cell. This does not increase computational costs for a cell of fixed size, as the number of

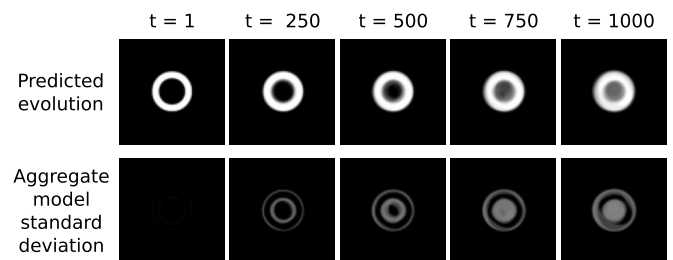


FIG. 5. Evolution (top) and aggregate model standard deviation (bottom) of a stable annular configuration. This kind of topology was not present in the training set. As soon as the evolution trajectory deviates from the correct stable configuration, aggregate uncertainty increases.

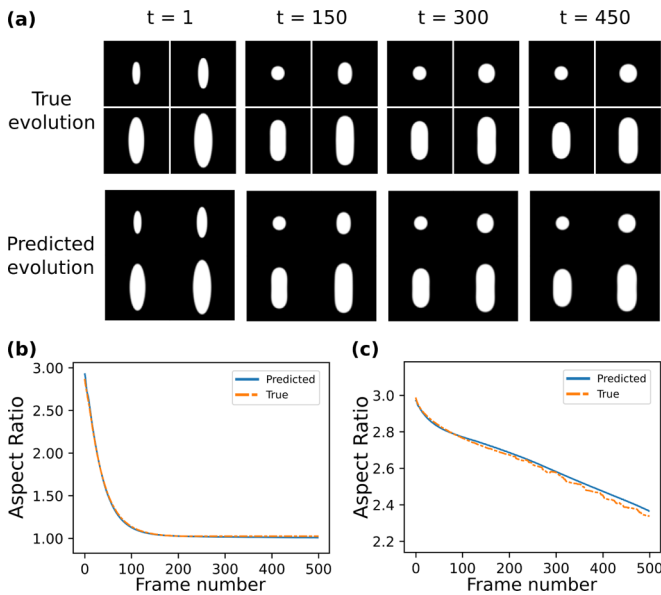


FIG. 6. (a) Evolution of elliptical domains. The learned model has implicitly learned time scaling rules with the initial shape size. As computational costs depends on the size of the image for Convolutional Recurrent Neural Networks, all domains are evolved concurrently in the same simulation cell. (b) and (c) aspect ratio evolution for the smaller and bigger ellipses respectively as a function of frame number. Blue continuous line represents the AR predicted by the machine learning approach and dashed orange line reports the "True" aspect ratio as obtained by FEM solver.

operations for a convolutional neural network only depends on the number of pixels present in the image. This is in contrast to phase field approaches based on FEM, whose cost depends on the complexity of shapes due to mesh refinement. As all ellipses can be fitted in a single image, the overall computational cost is equivalent to treating ellipses independently. Additionally, this example shows explicitly how the machine learned evolution is both equivariant with respect to the group of spatial translations and local.

C. Extension to complex conditions

Finally, we show performances of our method in the case of complex initial conditions containing multiple interacting material domains and discuss limitations of our approach in these settings. We recall that the training set did not contain any evolution involving interacting domains, therefore these conditions represent an extrapolation test. Figure 7 shows the evolution of one of such configurations. Notice that an important computational effort would be required in order to treat the problem with traditional FEM solvers, as fine meshes are required to track the domain boundaries topology. On the other hand, NN evolution is orders of magnitude cheaper, as discussed in previous sections. As it is clear from Fig. 7, the predicted evolution quickly exhibits a strong increase in prediction uncertainty in regions where nontrivial topological changes occur. This is made evident by the formation of halos and the increase of the variance between predictions of models in the ensemble.

These topological changes arise from well known behaviors in systems evolving by surface diffusion. In particular, high aspect ratio configurations undergo a Rayleigh-like instability and split into two or more domains as a result of a phenomenon known as pinch-off [40]. Notice that, even though such configurations were not present in the training set, the implemented prediction uncertainty estimation scheme is capable of identifying regions of the simulations where topological changes are likely to happen.

A less evident source of prediction uncertainty is present in the evolution of Fig. 7. Due to the existence of a lower bound in the representable physical dimension induced by the pixel discretization of the phase field, domains separated by less than such length are spuriously connected by gray regions. This can be noticed in Fig. 7 at $t = 1$: When the distance between two domains is comparable or smaller than the pixel size, gray areas are formed due to resolution limits. This affects the subsequent evolution, as the NN model has no access to information on the high-resolution FEM counterpart. Regions containing close domains at $t = 1$ also exhibit high uncertainty at subsequent times. While this phenomenon is

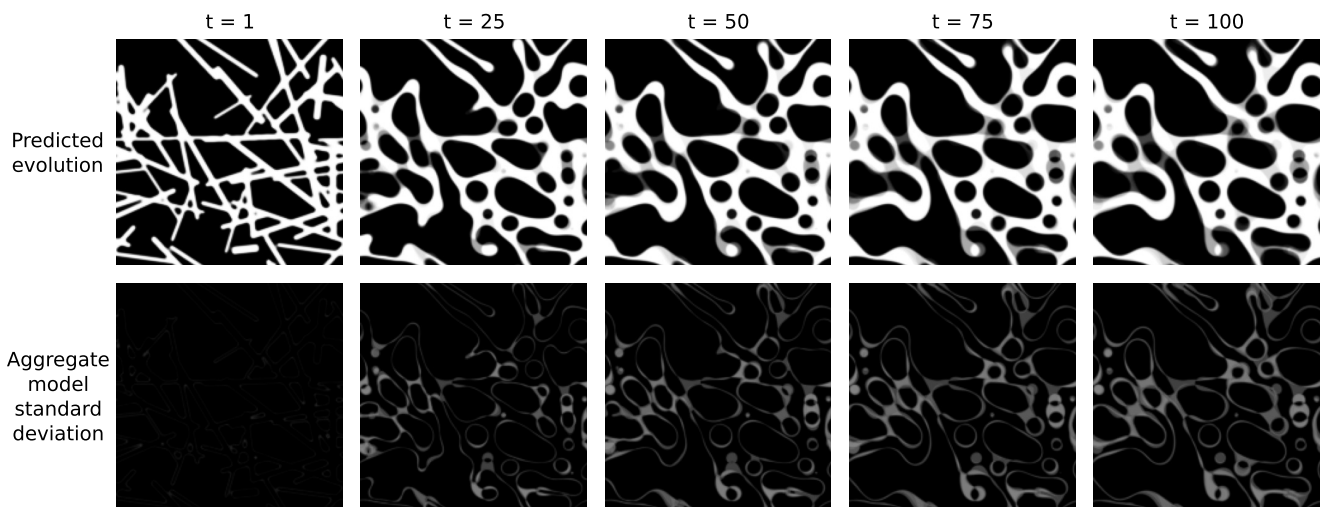


FIG. 7. Evolution of complex initial conditions comprising many interacting domains. Critical topological changes such as pinching, which was not present in the training set, are marked as high prediction error areas.

still detected by the aggregate model standard deviation as a source of uncertainty, initial conditions for full NN evolutions should be considered carefully to avoid such pathological states. A possible approach, though not tested in depth in the current work, is to exploit the CRNN ability to process subsequences and feed the ML model more than a single state of the FEM evolution as an initial condition. The information lost by resolution downscaling may therefore be recaptured in temporal correlations between successive states of the system.

V. CONCLUSIONS

In this paper, we have shown that CRNN can be conveniently used to learn morphological evolution via surface diffusion. Even by training the model over a dataset including only extremely simple evolutions (simple shapes, single domains throughout the whole evolution), reliable extrapolation to both more complex shapes and long timescales has been proven.

After introducing a convenient on-the-fly estimate of the uncertainty prediction, we showed that the model, instead, becomes less reliable during some critical evolution, involving, e.g., lateral aggregation and splitting of a single domain via

pinching. Limitations induced by the pixel-level discretization of the computational domains have also been discussed. Our results naturally open two further directions for future work: smart completion of the shape-evolution dataset by focusing solely on the high uncertainties regions, and designing a hybrid ML-PF approach where one automatically switches [29] between the (computationally expensive) high-fidelity PF approach to the RCNN one based on the uncertainty prediction. The latter would also remove pathological behaviors induced by finite resolution, as FEM solvers are capable to refine meshes as needed.

Finally, while here the methodology was tested under simple isotropic conditions and in the absence of external fields, suitable extensions to treat phenomena such as faceting [24], kinetically limited growth [68] and/or heteroepitaxy [69] are envisaged.

ACKNOWLEDGMENTS

We wish to thank M. Ceriotti for urging us to add uncertainty predictions to our procedure. We acknowledge the CINECA award under the ISCRA initiative, for the availability of high-performance-computing resources and support.

-
- [1] K. T. Butler, D. W. Davies, H. Cartwright, O. Isayev, and A. Walsh, Machine learning for molecular and materials science, *Nature (London)* **559**, 547 (2018).
 - [2] L. Himanen, A. Geurts, A. S. Foster, and P. Rinke, Data-driven materials science: Status, challenges, and perspectives, *Adv. Sci.* **6**, 1900808 (2019).
 - [3] P. Mehta, M. Bukov, C. H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to Machine Learning for physicists, *Phys. Rep.* **810**, 1 (2019).
 - [4] L. Fiedler, K. Shah, M. Bussmann, and A. Cangi, Deep dive into machine learning density functional theory for materials science and chemistry, *Phys. Rev. Mater.* **6**, 040301 (2022).
 - [5] E. H. Lee, W. Jiang, H. Alsalman, T. Low, and V. Cherkassky, Methodological framework for materials discovery using machine learning, *Phys. Rev. Mater.* **6**, 043802 (2022).
 - [6] H. W. Chung, R. Freitas, G. Cheon, and E. J. Reed, Data-centric framework for crystal structure identification in atomistic simulations using machine learning, *Phys. Rev. Mater.* **6**, 043801 (2022).
 - [7] J. Yeom, T. Stan, S. Hong, and P. W. Voorhees, Segmentation of experimental datasets via convolutional neural networks trained on phase field simulations, *Acta Mater.* **214**, 116990 (2021).
 - [8] J. Zhu, Y. Wang, Y. Huang, R. Bhushan Gopaluni, Y. Cao, M. Heere, M. J. Mühlbauer, L. Mereacre, H. Dai, X. Liu, A. Senyshyn, X. Wei, M. Knapp, and H. Ehrenberg, Data-driven capacity estimation of commercial lithium-ion batteries from voltage relaxation, *Nat. Commun.* **13**, 2261 (2022).
 - [9] T. Mueller, A. Hernandez, and C. Wang, Machine learning for interatomic potential models, *J. Chem. Phys.* **152**, 050902 (2020).
 - [10] P. Friederich, F. Häse, J. Proppe, and A. Aspuru-Guzik, Machine-learned potentials for next-generation matter simulations, *Nat. Mater.* **20**, 750 (2021).
 - [11] J. Behler and M. Parrinello, Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces, *Phys. Rev. Lett.* **98**, 146401 (2007).
 - [12] G. C. Sosso, G. Miceli, S. Caravati, J. Behler, and M. Bernasconi, Neural network interatomic potential for the phase change material GeTe, *Phys. Rev. B* **85**, 174103 (2012).
 - [13] J. Behler, Representing potential energy surfaces by high-dimensional neural network potentials, *J. Phys.: Condens. Matter* **26**, 183001 (2014).
 - [14] J. Behler, Constructing high-dimensional neural network potentials: A tutorial review, *Int. J. Quantum Chem.* **115**, 1032 (2015).
 - [15] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons, *Phys. Rev. Lett.* **104**, 136403 (2010).
 - [16] A. P. Bartók and G. Csányi, Gaussian approximation potentials: A brief tutorial introduction, *Int. J. Quantum Chem.* **115**, 1051 (2015).
 - [17] V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti, and G. Csányi, Gaussian process regression for materials and molecules, *Chem. Rev.* **121**, 10073 (2021).
 - [18] G. Imbalzano and M. Ceriotti, Modeling the Ga/As binary system across temperatures and compositions from first principles, *Phys. Rev. Mater.* **5**, 063804 (2021).
 - [19] C. Zeni, K. Rossi, T. Pavloudis, J. Kioseoglou, S. de Gironcoli, R. E. Palmer, and F. Baletto, Data-driven simulation and characterisation of gold nanoparticle melting, *Nat. Commun.* **12**, 6056 (2021).
 - [20] Y. Zuo, C. Chen, X. Li, Z. Deng, Y. Chen, J. Behler, G. Csányi, A. V. Shapeev, A. P. Thompson, M. A. Wood, and S. P. Ong, Performance and cost assessment of machine learning interatomic potentials, *J. Phys. Chem. A* **124**, 731 (2020).
 - [21] V. L. Deringer, N. Bernstein, A. P. Bartók, M. J. Cliffe, R. N. Kerber, L. E. Marbella, C. P. Grey, S. R. Elliott, and G.

- Csányi, Realistic atomistic structure of amorphous silicon from machine-learning-driven molecular dynamics, *J. Phys. Chem. Lett.* **9**, 2879 (2018).
- [22] F. C. Mocanu, K. Konstantinou, T. H. Lee, N. Bernstein, V. L. Deringer, G. Csányi, and S. R. Elliott, Modeling the phase-change memory material, $\text{Ge}_2\text{Sb}_2\text{Te}_5$, with a machine-learned interatomic potential, *J. Phys. Chem. B* **122**, 8998 (2018).
- [23] R. Glowinski and O. Pironneau, Finite element methods for Navier-Stokes equations, *Annu. Rev. Fluid Mech.* **24**, 167 (1992).
- [24] 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**, 2787 (2015).
- [25] B. Kim, V. C. Azevedo, N. Thuerey, T. Kim, M. Gross, and B. Solenthaler, Deep fluids: A generative network for parameterized fluid simulations, *Computer Graphics Forum* **38**, 59 (2019).
- [26] L. Fulton, V. Modi, D. Duvenaud, D. I. Levin, and A. Jacobson, Latent-space dynamics for reduced deformable simulation, *Computer Graphics Forum* **38**, 379 (2019).
- [27] X. Zhang and K. Garikipati, Machine learning materials physics: Multi-resolution neural networks learn the free energy and nonlinear elastic response of evolving microstructures, *Comput. Methods Appl. Mech. Eng.* **372**, 113362 (2020).
- [28] G. H. Teichert and K. Garikipati, Machine learning materials physics: Surrogate optimization and multi-fidelity algorithms predict precipitate morphology in an alternative to phase field dynamics, *Comput. Methods Appl. Mech. Eng.* **344**, 666 (2019).
- [29] 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 Comput. Mater.* **7**, 3 (2021).
- [30] K. Yang, Y. Cao, Y. Zhang, S. Fan, M. Tang, D. Aberg, B. Sadigh, and F. Zhou, Self-supervised learning and prediction of microstructure evolution with convolutional recurrent neural networks, *Patterns* **2**, 100243 (2021).
- [31] X. Shi, Z. Chen, H. Wang, D. Y. Yeung, W. K. Wong, and W. C. Woo, Convolutional LSTM network: A machine learning approach for precipitation nowcasting, in *Advances in Neural Information Processing Systems 28*, edited by C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Red Hook, NY, 2015).
- [32] Y. Wang, L. Jiang, M.-H. Yang, L.-J. Li, M. Long, and L. Fei-Fei, Eidetic 3d lstm: A model for video prediction and beyond, in *International Conference on Learning Representations* (2018).
- [33] W. W. Mullins, Theory of thermal grooving, *J. Appl. Phys.* **28**, 333 (1957).
- [34] B. Li, J. Lowengrub, A. Rätz, and A. Voigt, Geometric evolution laws for thin crystalline films: modeling and numerics, *Commun. Comput. Phys.* **6**, 433 (2009).
- [35] C. Gugenberger, R. Spatschek, and K. Kassner, Comparison of phase-field models for surface diffusion, *Phys. Rev. E* **78**, 016703 (2008).
- [36] M. Salvalaglio, A. Voigt, and S. M. Wise, Doubly degenerate diffuse interface models of surface diffusion, *Math. Meth. Appl. Sci.* **44**, 5385 (2021).
- [37] M. Salvalaglio, M. Selch, A. Voigt, and S. M. Wise, Doubly degenerate diffuse interface models of anisotropic surface diffusion, *Math. Meth. Appl. Sci.* **44**, 5406 (2021).
- [38] S. Vey and A. Voigt, AMDiS: Adaptive multidimensional simulations, *Comput. Vis. Sci.* **10**, 57 (2007).
- [39] T. Witkowski, S. Ling, S. Praetorius, and A. Voigt, Software concepts and numerical algorithms for a scalable adaptive parallel finite element method, *Adv. Comput. Math.* **41**, 1145 (2015).
- [40] M. Salvalaglio, P. Zaumseil, Y. Yamamoto, O. Skibitzki, R. Bergamaschini, T. Schroeder, A. Voigt, and G. Capellini, Morphological evolution of Ge/Si nano-strips driven by Rayleigh-like instability, *Appl. Phys. Lett.* **112**, 022101 (2018).
- [41] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, Backpropagation applied to handwritten zip code recognition, *Neural Comput.* **1**, 541 (1989).
- [42] Y. Le Cun, L. Jackel, B. Boser, J. Denker, H. Graf, I. Guyon, D. Henderson, R. Howard, and W. Hubbard, Handwritten digit recognition: Applications of neural network chips and automatic learning, *IEEE Commun. Mag.* **27**, 41 (1989).
- [43] S. Dieleman and B. Schrauwen, End-to-end learning for music audio, in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing* (IEEE, Piscataway, NJ, 2014), pp. 6964–6968.
- [44] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman, 1D convolutional neural networks and applications: A survey, *Mech. Syst. Signal Process.* **151**, 107398 (2021).
- [45] Z. Li, F. Liu, W. Yang, S. Peng, and J. Zhou, A survey of convolutional neural networks: Analysis, applications, and prospects, *IEEE Trans. Neural Netw. Learn. Syst.* **1** (2021).
- [46] W. Rawat and Z. Wang, Deep convolutional neural networks for image classification: A comprehensive review, *Neural Comput.* **29**, 2352 (2017).
- [47] F. Milletari, N. Navab, and S. A. Ahmadi, V-Net: Fully convolutional neural networks for volumetric medical image segmentation, in *Proceedings of the 2016 Fourth International Conference on 3D Vision* (IEEE, Piscataway, NJ, 2016).
- [48] D. Maturana and S. Scherer, VoxNet: A 3D Convolutional Neural Network for real-time object recognition, in *Proceedings of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, Piscataway, NJ, 2015).
- [49] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [50] C. C. Aggarwal, *Neural Networks and Deep Learning* (Springer International Publishing, Cham, 2018), pp. 1–492.
- [51] S. Schubert, P. Neubert, J. Poschmann, and P. Pretzel, Circular convolutional neural networks for panoramic images and laser data, in *Proceedings of the 2019 IEEE Intelligent Vehicles Symposium* (IEEE, Piscataway, NJ, 2019).
- [52] O. S. Kayhan and J. C. van Gemert, On translation invariance in CNNs: Convolutional layers can exploit absolute spatial location, in *Proceedings of the Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society, Washington, DC, 2020).
- [53] Z. C. Lipton, J. Berkowitz, and C. Elkan, A critical review of recurrent neural networks for sequence learning, [arXiv:1506.00019](https://arxiv.org/abs/1506.00019).
- [54] S. Hochreiter and J. Schmidhuber, Long short-term memory, *Neural Comput.* **9**, 1735 (1997).

- [55] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, On the properties of neural machine translation: Encoder-decoder approaches, in *Proceedings of SSTS 2014-8th Workshop on Syntax, Semantics and Structure in Statistical Translation*, 103 (2014).
- [56] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, Empirical evaluation of gated recurrent neural networks on sequence modeling, [arXiv:1412.3555](https://arxiv.org/abs/1412.3555).
- [57] N. Ballas, L. Yao, C. Pal, and A. Courville, Delving deeper into convolutional networks for learning video representations, [arXiv:1511.06432](https://arxiv.org/abs/1511.06432).
- [58] O. Ronneberger, P. Fischer, and T. Brox, U-net: Convolutional networks for biomedical image segmentation, in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Springer, Berlin, 2015), pp. 234–241.
- [59] F. Musil, M. J. Willatt, M. A. Langovoy, and M. Ceriotti, Fast and accurate uncertainty estimation in chemical machine learning, *J. Chem. Theory Comput.* **15**, 906 (2019).
- [60] L. Breiman, Bagging predictors, *Mach. Learn.* **24**, 123 (1996).
- [61] R. Tibshirani, A comparison of some error estimates for neural network models, *Neural Comput.* **8**, 152 (1996).
- [62] B. Efron, Bootstrap methods: another look at the jack-knife, in *Breakthroughs in Statistics* (Springer, Berlin, 1992), pp. 569–593.
- [63] D. P. Kingma and J. L. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [64] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai *et al.*, Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett (Curran Associates, Red Hook, NY, 2019), pp. 8024–8035.
- [65] C. Herring, Effect of change of scale on sintering phenomena, *J. Appl. Phys.* **21**, 301 (1950).
- [66] A. J. Bray, Theory of phase-ordering kinetics, *Adv. Phys.* **51**, 481 (2002).
- [67] W. B. Andrews, K. L. M. Elder, P. W. Voorhees, and K. Thornton, Effect of transport mechanism on the coarsening of bicontinuous structures: A comparison between bulk and surface diffusion, *Phys. Rev. Mater.* **4**, 103401 (2020).
- [68] M. Albani, R. Bergamaschini, A. Barzaghi, M. Salvalaglio, J. Valente, D. J. Paul, A. Voigt, G. Isella, and F. Montalenti, Faceting of Si and Ge crystals grown on deeply patterned Si substrates in the kinetic regime: Phase-field modelling and experiments, *Sci. Rep.* **11**, 18825 (2021).
- [69] F. Rovaris, R. Bergamaschini, and F. Montalenti, Modeling the competition between elastic and plastic relaxation in semiconductor heteroepitaxy: From cyclic growth to flat films, *Phys. Rev. B* **94**, 205304 (2016).