

Patterns

Self-supervised learning and prediction of microstructure evolution with convolutional recurrent neural networks

Highlights

- A convolutional recurrent neural net is trained to predict microstructure evolution
- The model accurately predicts short-term dynamics and long-term statistical properties
- It is significantly more efficient than the traditional PDE-based modeling approaches
- It can predict microstructure evolution based on incomplete information of systems

Authors

Kaiqi Yang, Yifan Cao,
Youtian Zhang, ..., Daniel Aberg,
Babak Sadigh, Fei Zhou

Correspondence

mt20@rice.edu (M.T.),
zhou6@llnl.gov (F.Z.)

In brief

Material microstructure plays a key role in the processing-structure-property relationship of engineering materials. Microstructure evolution is commonly simulated by computationally expensive continuum models. Yang et al. apply convolution recurrent neural networks to learn and predict several microstructure evolution phenomena of different complexities. The method is significantly faster than the traditional approach and capable of predicting the evolution process in systems with unknown material parameters. It provides a useful data-driven alternative to microstructure simulation.

Article

Self-supervised learning and prediction of microstructure evolution with convolutional recurrent neural networks

Kaiqi Yang,¹ Yifan Cao,¹ Youtian Zhang,¹ Shaoxun Fan,¹ Ming Tang,^{1,*} Daniel Aberg,² Babak Sadigh,² and Fei Zhou^{2,3,*}

¹Department of Materials Science and NanoEngineering, Rice University, Houston, TX 77005, USA

²Physical and Life Sciences Directorate, Lawrence Livermore National Laboratory, Livermore, CA 94550, USA

³Lead contact

*Correspondence: mt20@rice.edu (M.T.), zhou6@llnl.gov (F.Z.)

<https://doi.org/10.1016/j.patter.2021.100243>

THE BIGGER PICTURE Material microstructures are small structural features intermediate between atoms and macroscopic products with often decisive impact on the performance of engineering materials. A major goal of modern materials science is to improve material properties through the control of microstructure evolution during material processing and service. Microstructure evolution is traditionally simulated by continuum models based on partial differential equations. Here we demonstrate that convolutional recurrent neural networks, a type of machine-learning method, can be trained to predict various microstructure evolution phenomena with significantly improved efficiency. The method can learn the evolution rules from microstructure image sequences and make reliable predictions even with incomplete information about the systems or underlying mechanisms. This work illustrates the increasing power of data-driven approaches to address the computational challenges in microstructure modeling.



Proof-of-Concept: Data science output has been formulated, implemented, and tested for one domain/problem

SUMMARY

Microstructural evolution is a key aspect of understanding and exploiting the processing-structure-property relationship of materials. Modeling microstructure evolution usually relies on coarse-grained simulations with evolution principles described by partial differential equations (PDEs). Here we demonstrate that convolutional recurrent neural networks can learn the underlying physical rules and replace PDE-based simulations in the prediction of microstructure phenomena. Neural nets are trained by self-supervised learning with image sequences from simulations of several common processes, including plane-wave propagation, grain growth, spinodal decomposition, and dendritic crystal growth. The trained networks can accurately predict both short-term local dynamics and long-term statistical properties of microstructures assessed herein and are capable of extrapolating beyond the training datasets in spatiotemporal domains and configurational and parametric spaces. Such a data-driven approach offers significant advantages over PDE-based simulations in time-stepping efficiency and offers a useful alternative, especially when the material parameters or governing PDEs are not well determined.

INTRODUCTION

Material microstructures are mesoscale structural features that serve as an indispensable link between atomistic building blocks and macroscopic properties, leading to direct impacts on the processing-structure-property relationship of engineered materials. Tailoring material properties through controlled microstruc-

ture evolution under nonequilibrium conditions during material processing or service, including ubiquitous phenomena such as solidification, solid-state phase transformations, and grain growth, is arguably a cornerstone of modern materials science. The ability to understand and predict microstructure evolution has therefore long been a pivotal goal of computational materials design.

Due to time and length scales well beyond the capability of molecular dynamics, simulations of microstructure evolution often rely on coarse-grained models such as partial differential equations (PDEs) as employed in the phase-field method.^{1–4} Nevertheless, this approach also faces notable challenges. First of all, microstructure simulations employing PDEs remain fairly expensive. In the temporal dimension, strict upper limits on the minimum time-step size are dictated by the stability of numerical schemes that employ explicit time integration for nonlinear PDEs. Likewise, implicit time-integration methods handle larger time steps at the expense of additional inner iteration loops at each step. In addition, while in principle governing PDEs can be derived from the underlying thermodynamic and kinetic considerations, identifying, parametrizing, and validating PDEs in practice require significant efforts. For complicated or less studied materials, the evolution rules might be either not fully understood or too complex to be described by tractable PDEs.

We propose a machine-learning (ML) method as an alternative to microstructure evolution modeling. Recent progress in ML, and deep neural networks⁵ in particular, enables a data-driven approach to solving PDEs in place of the traditional numerical method.^{6–18} Based on statistical learning with big datasets, ML models can be applied without explicit prior knowledge of the physical mechanisms. With proper training, it is possible for ML algorithms to infer “hidden” parameters from the input microstructure images and identify the correct evolution trajectory. Moreover, ML models allow much larger time stepping to achieve significant speed in the temporal domain. For example, Raissi and coworkers used a single four-layer neural network^{6,7} to obtain the solutions to the Burgers equation, which otherwise require 500 Runge-Kutta iterations. Breen et al. tackled the notoriously difficult three-body problem with a 10-layer neural net, skipping thousands of smaller time steps.¹⁵ Similarly, coarser spatial grids may be used in ML models, as will be shown in this work. Although previous studies reveal the power of neural nets in rediscovering and solving different types of differential equations, deep learning of microstructure evolution, which can be described by PDEs in $2 + 1$ (i.e., 2 spatial and 1 temporal) or $3 + 1$ dimensions, remains a challenging subject.

In this work, we apply the convolutional recurrent neural network (RNN) to predict the spatiotemporal evolution of microstructure represented by two-dimensional (2D) image sequences. RNNs are neural nets designed to predict temporal data sequences with a hidden memory unit.^{19,20} With the development of effective variants such as the long short-term memory (LSTM) to address the vanishing gradient problem during back-propagation,²¹ RNNs have found widespread success in natural language processing,^{22,23} speech recognition,²⁴ and computer vision.^{25–27} Recently, LSTM combined with convolutional neural nets (CNNs) has been proposed for predictive learning of spatiotemporal sequences.²⁸ Compared with other neural network architectures designed to emulate $1 + 1$ and $2 + 1$ PDE solutions, convolutional LSTM employs CNN to efficiently extract latent spatial features of the system, which is advantageous in capturing the spatial correlation inherent in the evolution dynamics. Among several variants of convolutional RNNs developed in recent years,^{29–31} the Eidetic 3D LSTM (E3D-LSTM) model³¹ goes a step further by applying convolution to both the spatial and the temporal dimensions (i.e., 3D convolution

[3D-Conv] for $2 + 1$ systems) for integrated spatiotemporal feature extraction. Such a design facilitates a deeper coupling between the spatial and the temporal domains and enables improved performance in image sequence prediction in both short and long times. We choose the E3D-LSTM model for this work and use the terms E3D-LSTM and RNN interchangeably hereafter.

We assessed RNN’s learning ability and predictive power in the context of four well-known evolution phenomena with increasing levels of complexity: plane-wave propagation, grain growth, spinodal decomposition, and dendritic crystal growth. To facilitate comparison with physics-based models, the training datasets were generated from PDE-based simulations or explicit mathematical functions, whose behavior is well understood. A focus of our study is to examine to what degree RNNs can grasp and extract the evolution rules from the microstructure images it sees. To this end, extensive and stringent tests are devised to evaluate how well RNNs generalize and extrapolate the learning within the spatiotemporal domain and configurational and parametric spaces. We find that the properly trained RNN is able to extend the predictions up to 10 times the time spans of the training data, with significantly larger time-step sizes than used in the PDEs, and to systems of larger dimensions. It can forecast the evolution of systems with parameters not included in the training sets or initial configurations that exhibit significantly different statistical distributions from the training images. In addition to the excellent pixel-wise comparison between the ground truth and short-term predictions, the RNN accurately captures the statistical properties of microstructures in the examples considered (e.g., average size, grain, particle size, or interface curvature distribution) in the long term. The satisfactory performance of the RNN in these tests provides compelling evidence that it is capable of “emulating” the physical principles underlying diverse microstructure evolution phenomena, which explains why it is able to make reliable predictions well beyond the scope of training data. Such extrapolation capability further improves the RNN’s efficiency by allowing it to be trained with a relatively small data size. Our work illustrates the promise of ML approaches in general as a useful alternative to physics-based simulations of microstructure evolution.

Broadly speaking, the use of ML algorithms has grown very rapidly in materials science in recent years.^{32–35} They have seen diverse applications ranging from the discovery of new materials^{36–40} to the predictions of materials’ properties,^{41–45} the development of accurate and efficient potentials for atomistic simulations,^{46–49} microscopic and spectroscopic data analysis and processing,^{50–62} and effective inference of a material’s properties from a limited experimental dataset.^{63,64} A large number of these works are devoted to material microstructure, with encouraging results, including microstructure classification and quantification,^{50–54,65–67} image segmentation,^{55,56} predictions of microstructure-property relations,^{57,68–70} mapping processing-microstructure relations,^{71–74} microstructure optimization,^{75–77} and equilibrium configuration prediction.⁷⁸ Datasets in these works are mainly in the form of static microstructure images. This work focuses on revealing the important temporal correlation between images of microstructures along their evolution trajectory.




Figure 1. Application of the RNN to predicting plane-wave propagation

(A) Examples of output frames predicted by the trained RNN (P) based on 10 input frames in comparison with the ground truth (G).

(B) RMSE (black) and SSIM (blue) of the predictions averaged over 200 testing cases as a function of the frame index j .

(C) Relative errors of the wave-propagation parameters ($|k|$, ω , and β) inferred from the predicted images.

sequence was divided into staggered 20-frame training clips (i.e., frames 1–20, 11–30, etc.), each of which represented a training data point. For testing, the RNN was used to predict the next 50 frames based on an input of 10 consecutive frames. A total of 1,500 tests were performed.

Figure 1A illustrates two representative tests, which visually show little difference between the ground truth and the predictions. Figure 1B shows the pixel-wise comparison based on the root-mean-square error (RMSE) and structural similarity index measure (SSIM)⁷⁹ averaged over the 1,500 tests. Both RMSE and SSIM vary between 0 and 1, and lower

RMSE or higher SSIM scores indicate better agreement between the predictions and the ground truth. It can be seen that the RNN exhibits high pixel-wise accuracy in the short term within the length of training clips, where RMSE stays below 0.5% and SSIM above 99%. In the longer term, both RMSE and SSIM vary with time at a greater rate, but remain below 5% (or above 93%) for up to 50 output frames. As a more revealing measurement of how well the RNN recognizes the wave-propagation rules, the parameters in Equation 1 were extracted from the predicted images and compared with their ground truth values. As shown in Figure 1D, the predicted $|\vec{k}|$ and ω differ from the ground truth by less than 2%, but β shows a larger deviation up to 20%. A probable reason for the less accurate prediction of β is that β characterizes a slower decaying mode of wave motion and may require longer training sequences to learn precisely its temporal behavior.

Overall, the RNN exhibits excellent performance when applied to the simple plane-wave propagation problem. Next, we test it against more realistic microstructure evolution problems.

RESULTS

We employed numerical simulations to generate sequences of 64×64 -pixel images as training datasets for four classical examples of evolution phenomena, i.e., plane-wave propagation, grain growth, spinodal decomposition, and dendritic crystal growth. With varied complexity, they represent a good combination of testing problems for evaluating the capability of the RNN in predicting microstructure evolution.

Plane-wave propagation

Before delving into problems pertinent to real materials, we first tested the RNN with a simple toy model: plane-wave propagation dynamics of a scalar field c explicitly described by the following expression:

$$c(x, y, t) = \frac{1}{2} \sin(k_x x + k_y y + \omega t + \theta_0) \exp(-\beta t) + \frac{1}{2}, \quad (\text{Equation 1})$$

where $\vec{k} = (k_x, k_y)$ is the wave vector, θ_0 is a random phase, and β is a decay exponent. We used Equation 1 to generate image sequences, each of which consisted of 200 frames at a time interval of 0.005 between two adjacent frames starting at $t = 0$. The parameters in Equation 1 were randomly chosen for each sequence: $2\pi/|\vec{k}| \in [0.3, 0.6]$, $2\pi/\omega \in [0.03, 0.06]$, $2\pi/\beta \in [1.5, 6]$, and $\theta_0 \in [0, 2\pi]$. Among the generated sequences, 80 were used for training, 20 for validation to evaluate model convergence during training, and 100 for testing. Each simulation

Grain growth

Grain growth describes the increase in the average grain size in polycrystals with time to reduce the excess energy associated with grain boundaries. During the process, some grains grow, while others shrink and disappear, leading to a persisting drop in the number of grains in the system. The growth or shrinkage rate of a grain in 2D polycrystals is determined by its number of sides N according to the famous von Neumann-Mullins or “N-6” rule.^{80,81}

$$\frac{dA}{dt} = M\gamma \frac{\pi}{3} (N - 6), \quad (\text{Equation 2})$$

where A is the grain area, and M and γ are the grain boundary mobility and energy, respectively. Equation 2 states that any grains with fewer than six neighbors will shrink, and those with more than six sides will grow at a rate proportional to $N - 6$.

We generated the training data by performing isotropic 2D grain growth simulations with a phase-field mode⁸² (see [experimental procedures](#)). Simulations were performed on a 256×256 grid with periodic boundary conditions to accommodate a sufficient number of grains. Subsequently, the simulation images were downsampled to 64×64 pixels by averaging. Each simulation employed the same parameters but started with a different initial configuration constructed by Voronoi tessellation with 100 random seeds. It output a 20-frame clip after a relaxation period, which was intended to remove the artifacts in the polycrystalline structure. The time interval between two adjacent frames corresponded to 80 PDE time steps. The first frame in a clip contained ~ 75 grains and the last one had ~ 45 grains. A total of 2,400 clips were prepared for training and 600 for validation during training.

After training, the RNN was subjected to a set of more challenging extrapolation tests than in the wave-propagation problem. First, we applied the trained model to predict longer image sequences with less input information. The RNN was required to predict 199 frames based on only one input frame. Theoretically, this was feasible, as grain growth obeys the dissipation dynamics described by PDEs of the first order in time (Equation 8). Here the length of the test sequences was 10 times that of the training clips, and more significantly, 90% of the output frames (frame index $j = 21\text{--}200$) depicted coarsened polycrystalline states never seen by the RNN during training. Figure 2A presents two representative tests, which show that the RNN does a very good job in the temporal extrapolation. The predictions and ground truth were difficult to distinguish visually in the short term, e.g., at frame index $j = 30$, but visible local structure difference emerged at the later stage. Figure 2C shows that the average RMSE of 1,000 tests rises and stabilizes around 20%, while SSIM decreases to ~ 0.4 at the 200th frame. Despite the increasing difference, the predicted polycrystalline structures were free of any noticeable artifacts throughout the sequences. We note that the accumulation of the discrepancy between the ground truth and the predictions is inevitable in the long term. This is because the grain boundary connectivity bifurcates upon grain disappearance (see examples in Figure S1), which leads two initially identical configurations on to divergent evolution pathways. As such, statistical measurement of the similarity between the two polycrystalline configurations is more meaningful than pixel-wise comparison, and the RNN performs very well in this regard. As shown in Figure 2D, the error in the predicted average grain area $\langle A \rangle$ of 1,000 testing cases remained below 5%, while $\langle A \rangle$ had a 5-fold increase. Figure 2E shows that the predictions and the ground truth also have very good agreement in the grain size distribution. The Euclidean distance between them is only 0.71% at $j = 50$ and still has a low value of 1.61% at $j = 200$. The RNN thus faithfully reproduced the statistical char-

acteristics of polycrystals even after a 10-fold extrapolation in time. Further extension of the prediction to 1,000 frames is presented in Figures S2A and S2B, which show that the RNN still managed to produce realistic looking microstructure and exhibited good agreement with the ground truth in grain size distribution. A nonnegligible portion of the trials had only one grain remaining in the system beyond $j > 1,000$. While we intentionally limited the training data to a small time span here to examine the RNN's extrapolating capability, its long-time prediction accuracy can be improved by including data from the later evolution stage into the training sets.

Next, we subjected the RNN to spatial extrapolation tests by asking it to predict grain growth in a system much larger than the training images. The 3D-Conv in E3D-LSTM operates on inputs and internal states with a fixed filter size ($5 \times 5 \times 2$ used in this work) that is independent of the input image size. After the weights in the 3D-Conv filters in the network are trained by 64×64 images to learn the spatiotemporal correlation of the system, the same filters can slide over larger images to predict their evolution. Therefore, the evolution rules learned by the model are expected to be extendable to larger domains. Figure S3 presents the results of the grain growth kinetics on a 256×256 mesh predicted by the RNN trained on 64×64 -pixel images. The predictions exhibit similar RMSE and SSIM compared with those for the smaller 64×64 -pixel domain. The spatial extensibility of the RNN means that there is no need for retraining the model when applying it to problems of different sizes, which is a very appealing feature for practical applications.

As the third type of extrapolation test, the RNN was applied to predict the evolution of artificial polycrystalline configurations qualitatively different from the training data. Figure 2B shows such an example, in which the system contains four orderly arranged four-sided grains embedded within four larger grains.

Its statistical difference from the training configurations is quantified by their distinct two-point correlation functions $g(r)$ as shown in Figures S4A and S4B. The individual grains in the artificial polycrystal have a strong spatial correlation as reflected by the sharp peaks in $g(r)$. Despite the notable morphological difference from those generated by random Voronoi tessellation, its evolution is accurately captured by the RNN.

The above tests demonstrate the RNN's capability to generalize and extrapolate its learning in the spatiotemporal and configurational spaces. This is a strong indication that it has grasped the evolution rules, which is further supported by other evidence. Grain growth consists of two elementary processes: the continuous shrinkage or expansion of grains without changing their number of sides, N , and the discontinuous changes in the grain boundary connectivity when grains switch edges or disappear. The former process is governed by the $N - 6$ rule (Equation 2) resulting from the curvature-driven boundary movement. In Figure 3A, we show the average growth rates for grains with different N using data from all 1,000 tests. The predictions faithfully reproduce the N dependence of the ground truth. On the other hand, Figure 3B illustrates all four possible topological events that could occur to the grain boundary network upon grain disappearance or edge switching in a 2D system. Many instances of these events exist in the training dataset and are




Figure 2. Application of the RNN to predicting grain growth

(A) Examples of RNN output frames (P) based on one input frame in comparison with the ground truth (G).

(B) RNN prediction of the evolution of an artificial polycrystalline configuration, in which four small four-sided grains are embedded in larger six-sided grains.

(C) RMSE (black) and SSIM (blue) of the predictions averaged over 1,000 cases as a function of the frame index j .

(D) Time evolution of the average grain area in 1,000 testing cases predicted by the RNN versus ground truth.

(E) Grain size distribution at $j = 50$ and 200 predicted by the RNN versus ground truth. Effective grain radius was calculated by $\sqrt{A/\pi}$.

the phenomenon of spontaneous phase separation in unstable binary or multi-component systems widely found in alloys and polymer blends.⁸³ Mathematically, the spatiotemporal evolution during spinodal decomposition is described by the Cahn-Hilliard (C-H) equation⁸⁴ (Equation 9 in experimental procedures), which is numerically solved to generate the ground truth in this work. Compared with grain growth, spinodal decomposition is a more complex evolution phenomenon, since it involves not only curvature-driven interface migration but also coupled long-range diffusion of chemical species. The complexity is also reflected by the fourth-order nonlinear C-H equation versus the second-order phase-field PDEs for grain growth.

Spinodal decomposition consists of two distinct stages: a fast composition modulation growth stage, followed by a slower coarsening stage, at which the length scale of the phase-separation pattern gradually increases due to the Gibbs-Thomson effect.⁸³ Due to their very different time scales, image sequences with a fixed time interval cannot effectively resolve both stages at the same time. Here we chose to train the RNN to recognize the system evolution in the second coarsening

observed by the RNN model during training. The numerical examples in Figure 3B show that the trained RNN is able to correctly predict each one of them. Therefore, the satisfactory performance of the RNN derives from its learning of the elementary steps of the grain growth process.

Spinodal decomposition

As a third example of microstructure evolution phenomena, we trained the RNN to predict spinodal decomposition, which is

stage. Training and validation data were generated from 480 and 120 simulations, respectively, which employed the same parameters but different initial states. The system started from a uniform binary mixture with one of three compositions at $c_0 = 0.25, 0.5$, and 0.75 , which produced different types of domain morphologies. A random noise of amplitude $\Delta c = 0.01$ was added to the initial configurations to trigger phase separation. Each simulation produced 100 images, and the system became phase separated after 2 or 3 frames. Similar to the




Figure 3. Evidence of the RNN capturing the evolution rules of grain growth

(A) The RNN accurately predicts the dependence of the grain growth rate $\langle dA(N)/dt \rangle$ on the number of grain sides N . $\langle dA(N)/dt \rangle$ is averaged over grains of the same N in all of the testing cases.

(B) Examples from testing cases show that the RNN correctly predicts the four possible topological events when a grain disappears or loses an edge to its neighbors. Red circles highlight where the events occur in the predicted images.

wave-propagation problem, these frames were divided into staggered 20-frame training clips (i.e., frames 1–20, 11–30, ..., 81–100). The time interval between 2 adjacent frames corresponded to 370 time steps on average in phase-field simulations, which employed an implicit PDE solver with variable time-step size. To ensure conservation of mass during evolution, the E3D-LSTM model was modified to enforce that the average of the image pixel values remained unchanged after passing through the neural net.

We performed temporal extrapolation tests on the trained model in a way similar to the case of grain growth. The RNN was asked to output 200 frames, or 10 times the training clip length, given 1 input frame that was taken from the 50th frame of a simulation starting from a uniform mixture. Seventy-five percent of the output frames ($j = 51$ –201) thus fell outside the time span of the training sets. In addition, predictions based on 10 input frames were also tested. The results are presented in Figure S5 and show similar performance compared with those with 1 input frame only, which indicates that the information contained in one initial frame is sufficient for the RNN to correctly project the evolution trajectory. Figure 4A showcases several examples from a total of 510 tests with 170 each having $c_0 = 0.25$, 0.5, or 0.75. The short-term predictions up to $j \sim 50$ closely

resemble the ground truth, which was quantified by the low RMSE (<0.06) and high SSIM (>0.97) in Figure 5A. While the discrepancy gradually accumulates with time and visible differences appear at the later stage, the long-term predictions are realistic looking and no artifacts can be discerned. In addition, Figure S6 confirms that the conservation of c is strictly obeyed in the output frames. Morphology-wise, it is difficult to tell by human eyes whether the images are generated by the RNN or simulations. Such similarity is corroborated by the statistical analysis of the microstructure. In Figure 5B, we compare the interface curvature distributions in the predicted versus ground truth images of 170 testing cases with $c_0 = 0.5$, which have a bi-continuous two-phase morphology. The agreement is very good in both short and long terms, which can be quantified by the Euclidean distance between the two distributions: 0.00936 at frame $j = 26$ and 0.05811 at $j = 201$. On the other hand, systems with $c_0 = 0.25$ or 0.75 contain individual particles of the minority phase ($c = 1$ or 0) dispersed within the majority phase. Figure 5C shows the time dependence of the average particle size $\langle R \rangle$ for 170 tests with $c_0 = 0.25$. The corresponding particle size distributions are presented in Figure 5D. The comparison is again satisfactory. The predicted $\langle R \rangle$ has a maximal error of 1.04% within the test period, and the Euclidean distance between the predicted and the true size distributions is only 0.01 at $j = 26$ and 0.034 at $j = 201$. In Figure S2C, we show an example of the RNN prediction up to 2,000 frames. While its pixel-wise difference from the ground truth becomes larger, the predicted structure remains realistic and does not show any image blurring. The predicted interface curvature distribution also agrees well with the ground truth as shown in Figure S2D.

We next performed the spatial extrapolation test by applying the trained model to a larger 256×256 -pixel domain. As shown in Figure S7, the RNN performs equally well in the extended system, with RMSE and SSIM comparable to those in the smaller domain. Furthermore, Figure 4B shows an example in which the RNN's ability to predict the evolution of configurations "foreign" to the training set was tested. The initial configuration in the test was created by placing circular particles of $c = 1$ with random radii on square lattice sites in the matrix of $c = 0$. As revealed by the two-point correlation functions in Figures S4C and S4D, the particles in this microstructure exhibited strong spatial correlation, while those in the training images were spatially uncorrelated. Although it never saw such a configuration during training, the RNN captured its evolution very well.

The impressive extrapolation capability of the RNN when applied to spinodal decomposition implies its understanding of the physical rules of this phenomenon. The coarsening of the spinodal structure is thermodynamically driven by the interface curvature dependence of chemical potentials (i.e., the Gibbs-Thomson effect) and kinetically limited by the diffusion of the species in the system. Figures 5B and 5D show that the RNN grasps the Gibbs-Thomson effect, which causes the fraction of low-curvature interface segments to increase with time, and Figure 5C confirms that the diffusion-controlled coarsening kinetics is captured by the model. Apart from the accurate statistical representation, the examples in Figures 5E and 5F illustrate that RNN is also capable of predicting subtle local morphological changes. The fate of the particle highlighted by red in Figure 5E is determined by the relative sizes of its neighbor particles and itself,




Figure 4. Application of the RNN to predicting spinodal decomposition

(A) Comparison between predictions (P) and ground truth (G) from two testing cases, in which the RNN outputs 200 frames based on 1 input frame of spinodal structure generated from random perturbation to a system of uniform composition.

(B) The RNN prediction of the evolution of an artificial biphasic configuration, in which second-phase particles ($c = 1$) of randomly chosen radii are arranged in an orderly manner within the primary phase ($c = 0$).

which exchange mass between one another via diffusion due to the size-dependent chemical potential. The red particle first grows at the expense of a smaller neighbor, but subsequently shrinks by losing mass to the other two bigger particles nearby. In Figure 5F, the particle in red receives an incoming diffusion flux from two smaller adjacent particles. Its growth rate exhibits two bursts, which coincide with the complete dissolution of the two particles. The RNN's ability to predict detailed evolution features as demonstrated in these examples further inspires confidence in its comprehension of the underlying physics.

Dendrite growth

In the last example, we gave the RNN a more challenging task to predict dendritic crystallization patterns. During crystal growth, dendritic structures, like beautiful snowflakes, often form due to the morphological instability of the growth front, which is promoted by the negative temperature and/or species concentration gradient(s) ahead of the phase boundary and the interface energy anisotropy. Such instability phenomena are intrinsically difficult to predict. In addition, dendrite growth is a multi-physical process coupling phase transformation, long-range mass/heat transport, and interface instability. As a result, microstructure images fed to the RNN do not contain the complete information of

the system state, which further increases the difficulty of making accurate predictions.

Here we generated training data using a phase-field model of solidification in pure systems by Kobayashi.⁸⁵ As described in the experimental procedures, the spatiotemporal evolution of the system state is described by two coupled PDEs for the temperature (T) and phase-field (ϕ) variables. ϕ distinguishes between the solid ($\phi = 1$) and the liquid ($\phi = 0$) phases during solidification. We use $\phi(t, x, y)$ to create the microstructure images. T and other parameters in the governing equation (Equation 15), such as the normalized latent heat K , are thus hidden from the learning process. We performed phase-field simulations on a 64×64 mesh, in which a small solid nucleus was placed at or near the center and surrounded by the supercooled liquid phase. The training and validation sets contained 800 and 200 simulations, respectively. To enrich the training data, each simulation had a different nucleus, crystal orientation θ_0 , and K . Specifically, K was randomly chosen from (1.2, 2) and θ_0 from $(0, \pi/3)$ (crystal was assumed to have six-fold symmetry). The nucleus was given a random shape (circle, rectangle, or ellipse), size (2–6 pixels), and off-center distance (± 5 pixels in the x and y directions). Similar to the case of spinodal decomposition, 100 image frames with equal time intervals were obtained from a simulation and divided into eight staggered 20-frame training clips.

In testing, the trained RNN model was required to predict 50 frames from 10 consecutive input frames, which were taken from the first half of a simulation. Predictions were not extended to longer times because the dendrite tips already approached the domain boundaries after 50 output frames in many tests, and growth stagnated subsequently. Instead, we focused on conducting the extrapolation tests in the model parameter space. Specifically, K was randomly and uniformly selected from (0.8, 2.4) to generate ground truth data in the testing cases. This means that half of the selected K values fell outside its range in the training data, i.e., (1.2, 2). θ_0 and the solid nucleus shape were also randomized. Figure 6 presents several examples from a total of 600 testing cases. The predicted dendritic structure matched the ground truth well in all the cases, even at $K = 1.161$ and 2.106, which were outside the scope of training data. In particular, the RNN captured the fine features of the dendrites, such as the locations of secondary side branches. It can be seen that the crystal growth pattern depended strongly on K . Smaller K resulted in thicker primary branches and more compact morphology. The RNN managed to recognize the correct evolution trajectory based on the input images without prior knowledge of the underlying K value. Figure 7A shows the RMSE and SSIM of the predictions averaged over the 600 testing cases. The RNN fared well in pixel-wise comparisons, although the prediction error increased faster with time than in the cases of grain growth and spinodal decomposition, which can be attributed to the more complex physics of the dendrite growth process.

As a more revealing indicator of the RNN's performance, we used several shape descriptors (Feret diameter d_{feret} , convexity, solidity) to characterize the dendrite morphology. Feret diameter, which is defined as the maximum distance between two parallel tangent lines touching the shape, provides a measure of the linear dendrite dimension. Convexity and solidity quantify the degrees of concavity and compactness of the crystal.




Figure 5. Accuracy of the RNN in predicting spinodal decomposition

(A) RMSE and SSIM of RNN predictions averaged over 510 testing cases as a function of the frame index j .

(B) Distribution of the interface segment curvature κ at $j = 26$ and 201 in 170 testing cases with $c_0 = 0.5$. (C and D) Evolution of the average second-phase particle radius $\langle R \rangle$ (C) and particle area A_p distribution (D) in 170 testing cases with $c_0 = 0.25$. $\langle R \rangle$ was calculated as $\sqrt{\langle A_p \rangle / \pi}$.

(E and F) (Top) Examples of local morphological evolution predicted by the RNN from two testing cases with $c_0 = 0.25$. (Bottom) Size evolution of the red particle in the images as predicted by the RNN versus ground truth.

spatial coarsening is also demonstrated. Because of the diffuse-interface representation, a grain boundary typically needs to be resolved by 5 or 6 pixels in phase-field simulations to maintain desired numerical accuracy.³ However, the RNN is not subject to the same spatial resolution requirement and can predict system evolution on a coarser mesh (64×64) than used in phase-field simulations (256×256). We benchmarked the computational performance of E3D-LSTM on both a GPU (Nvidia GeForce GTX 1080-Ti) and a CPU (Intel i7 3.2 GHz) and compared them with the performance of phase-field simulations run on a CPU. The results are summarized in Table S1. Averaged over more than 500 trials, the RNN accelerates the predictions

Figure 7B shows the time evolution of these descriptors from one test, while their average errors for all 600 tests are plotted in Figure 7C. It can be seen that the RNN accurately predicted the dendritic shape evolution with an average error less than 7% throughout the tests. In addition to global metrics, we also examined how well the RNN reproduced local dendritic structural features. In Figure 7D, the number of secondary branches formed on a primary branch in a test is plotted as a function of time. It shows that the RNN performed very well in predicting the frequency of the side-branching events occurring near the dendrite tip.

DISCUSSION

In addition to prediction accuracy, we compared the computational efficiency of using the RNN for microstructure evolution predictions with that of PDE-based simulations. During training and testing, the time interval between two RNN output frames is 80 times that of the average time-step size used in the grain growth simulation, 370 times in spinodal decomposition, and 7 times in dendritic crystal growth. This illustrates that the RNN can improve efficiency by using larger time steps than PDE solvers, for which the time-step size is limited by the stability of the numerical schemes. In the grain growth example, the RNN's advantage in

by 92 times for spinodal decomposition and 79 times for dendrite growth when run on a GPU, and 7.6 and 8.6 times on a CPU, respectively. The speedup is more significant in the grain growth example (718-fold on the GPU, 87-fold on the CPU) because of the spatial coarse graining exploited by the RNN. In our tests, it takes 130–450 s to load and initialize pretrained RNN models. Therefore, the RNN is very efficient, especially when applied to a large number of cases in one run so that the overhead associated with initialization is small.

The overall efficiency of the RNN in predicting microstructure evolution also depends on the training data size and the efforts and resources required for data collection. Figure S8 shows the dependence of the validation error on the number of training clips N_{clip} for plane-wave propagation and grain growth. In both cases, the improvement in model performance becomes negligible after N_{clip} goes beyond $\sim 2,000$. Tests show that the optimal accuracy can be further tuned, e.g., with the number of layers or the number of hidden features. In principle, a large enough neural net could be arbitrarily accurate, but in practice, training such a model becomes infeasible. Our current model is a decent compromise between accuracy and computational cost. On the other hand, we find that increasing the length of training clips beyond 20 frames does not significantly improve the prediction accuracy. For all of the examples in this work, the time spent




Figure 6. Application of the RNN to predicting dendritic crystal growth

RNN predictions (P) versus ground truth (G) from five testing cases with different K values, in which the RNN outputs 50 frames based on 10 input frames.

evolution rules are localized, which could be relatively easily recognized by E3D-LSTM through the 3D-Conv operation that specializes in remembering local motion. In contrast, training the RNN to predict spinodal decomposition is more challenging because the long-range mass transport inherent in the process creates longer and stronger spatiotemporal correlation, which requires more convolution operations and longer-term memory states to extract the essential features. In fact, the model can be successfully trained to predict grain growth with only two E3D-LSTM layers, but four layers are needed for spinodal decomposition to reach similar prediction accuracy. We also find it necessary to include longer image sequences (100 frames) into the training datasets for spinodal decomposition to better inform the RNN of the evolution trajectories. Predicting dendrite growth presents additional challenges due to the interface instability and the existence of hidden variables (7) not directly seen by the RNN.

The PDEs underlying the three microstructure examples investigated here (grain growth, spinodal decomposition, dendrite growth) describe dissipative dynamics, in which the system's evolution rate decreases with time. For example, the growth rate of the average grain or particle size varies as $t^{-1/2}$ and $t^{-2/3}$ in grain growth and spinodal decomposition, respectively. Such behavior is common in microstructure phenomena, as the thermodynamic driving force continues to diminish during evolution. Compared with

on generating the training datasets is comparable to the model training time. The plane-wave propagation and dendrite growth examples also demonstrate the transferability of the trained model, which can robustly interpolate or even extrapolate predictions to parameter values not included in the training data. Therefore, the data requirement of the RNN should not present a major obstacle to its applications.

Despite the overall very impressive performance, our tests show that the learning rate and predictive power of the RNN vary with the nature of the microstructure evolution phenomena it is applied to. Among all the examples, the RNN demonstrates the best learning ability in predicting grain growth, because its

other PDEs such as the wave equation and chaotic PDEs that have been successfully emulated by neural networks, ML of these problems faces a new challenge, i.e., to train the neural nets to predict the long-time behavior based on only short-time data with much faster evolution dynamics. Nevertheless, we find that E3D-LSTM can utilize microstructure images from the early stage to reliably predict the much slower evolution at 10-fold larger times. To our best knowledge, this impressive long-term predictive ability has not been demonstrated for similar PDE systems. We attribute it to the novel architecture of E3D-LSTM, which integrates 3D CNN into the LSTM to better capture the long-term spatiotemporal correlation.




Figure 7. Accuracy of the RNN in predicting dendritic crystal growth

(A) RMSE (black) and SSIM (blue) of the predictions averaged over 600 testing cases.

(B) Time evolution of the Feret diameter d_{feret} , convexity, and solidity of a growing crystal from a testing case. Solid lines are ground truth (G), and dashed lines are predictions (P). The shape descriptors were calculated in ImageJ after image binarization. Convexity is defined as L_h/L_c , where L_c is the crystal perimeter and L_h is the perimeter of the convex hull of the crystal. Solidity is defined as A_c/A_h , where A_c is the crystal area and A_h is the area surrounded by the convex hull of the crystal shape.

(C) Relative errors of predicted d_{feret} , convexity, and solidity of crystals averaged over 600 testing cases as a function of image frame index j .

(D) Development of secondary branches on the dendritic crystal in a testing case. N_s is the number of secondary branches on a primary branch of the dendrite. Insets above and below the curves show the crystal shape from the ground truth and predictions, respectively, at times marked by the black squares.

Another distinction between this work and previous NN-based PDE emulators lies in that we intentionally require the RNN to be trained with only partial information of the PDE solutions or to make predictions with incomplete knowledge of the underlying PDEs. In the grain growth case, training images are generated from the 3-norm of 100 grain-orientation order parameters $\{\eta_1, \eta_2, \dots, \eta_{100}\}$, which are solved from the governing equations (Equation 8). Information on grain orientation is lost in the images. The PDEs for dendrite growth (Equations 14 and 15) involve both the phase field ϕ and the temperature field T , but only ϕ is used to create microstructure images. Furthermore, the RNN trained for dendrite growth is not explicitly given the value of the latent heat parameter (K) in the PDEs during prediction and instead needs to infer its value from the input image sequence to identify the correct evolution trajectory. We made these choices for training and testing because they reflect how the RNN may be potentially used in real applications, in which missing information is often the norm rather than the exception. Like our training sets, microstructure images obtained from optical or electron microscopy are usually grayscale images of phase contrast and do not capture all the physical fields relevant to the evolution processes. It is also common for some properties of a material system not to be known or accurately characterized so that the corresponding PDE parameters are ill determined. Our study demonstrates that a well-trained RNN can not only serve as a PDE emulator, but also infer implicit material properties from spatiotemporal data and provide a “reduced order” representation of the targeted problems to lower the data demand and improve the training and prediction efficiency.

Over the last 3 decades, many types of microstructure from different materials systems have been successfully simulated by various phase-field models.^{2–4} The phase-field method has become a proven and versatile computational technique for capturing complex microstructural morphology and coupling multiple physical processes within a unified framework.

Microstructures from phase-field simulations faithfully reproduce experimental observations in a diverse set of materials systems,^{86–92} and therefore can serve as reliable training data to train the RNN for predicting a wide variety of microstructure evolution phenomena.

The rapid advancement in *in situ* and *operando* characterization techniques in recent years has fueled the collection of experimental spatiotemporal microstructure data from a wide range of materials systems at an ever-increasing rate.^{93–97} At the same time, however, it becomes an increasing challenge to efficiently analyze the acquired data to generate useful insights. The deep learning approach examined in this work provides a valuable tool for extracting high-level, quantitative information from such data. Many types of experimental digital images could be used as input to the RNN, including those from microscopy and tomographic reconstruction, after standard preprocessing (resizing, denoising, segmentation, etc.). For microstructure evolution governed by known evolution rules, the RNN could be trained with numerical simulation results with parameter values randomly sampled from the relevant parameter space. It can then be applied to the experimental data of a specific system to learn its parameters and augment *operando* experiments, which often require sophisticated instruments (e.g., synchrotron X-ray beamlines) with limited availability, by extending the observations in time and space. When dealing with new microstructure evolution phenomena that do not have physical models yet, the RNN could be trained directly with experimental image sequences and subsequently be used as a surrogate model to probe the system behavior more conveniently and/or under conditions that are difficult to access experimentally.

In conclusion, we trained a convolutional RNN (E3D-LSTM) to predict the spatiotemporal evolution of material microstructure. Using training data from four distinct evolution processes (plane-wave propagation, grain growth, spinodal decomposition, and dendritic crystal growth), the RNN, which is composed

of the same network architecture, is able to adapt efficiently to different evolution rules. The ability of the RNN to generalize learning beyond the training datasets was systematically examined by a series of extrapolation tests. In addition to performing very well in pixel-wise comparison with ground truth in short-term predictions, the RNN accurately described the statistical properties of microstructures assessed herein over long periods up to 10 times the training data's time span. Without additional training, neural nets trained on small-size images could be straightforwardly applied to larger systems with comparable accuracy. The method can reliably predict the evolution of microstructures whose morphology or underlying material parameters differ qualitatively from the training data. The spatiotemporal, configurational, and parametric extensibility demonstrated by the RNN suggest that it is capable of learning the evolution rules of the microstructure phenomena considered here, which provides the physical basis for its practical applications. Computationally, the RNN is not restricted by the numerical stability of PDE solvers and can employ time-step size 1–2 orders of magnitude larger than PDE-based simulations in our tests. Beyond accelerating simulations shown in this study, the ML approach may provide a valuable pathway toward prediction of microstructure evolution in situations where the material parameters or evolution principles are not completely resolved or only partial information of the system state is available. The current E3D-LSTM model can also be extended to predicting the spatiotemporal evolution of 3D systems in a straightforward manner by replacing the 3D-Conv operation with 4D convolution^{98,99} (3 spatial + 1 temporal) in the network. Given the lack of optimized 4D convolution subroutines in current deep learning frameworks, a practical solution is to extend from 2D to 3D spatial convolutions in alternative RNNs such as PredRNN. Its application to the learning of molecular dynamics trajectory will be reported elsewhere.

EXPERIMENTAL PROCEDURES

Resource availability

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, Fei Zhou (zhou6@lnl.gov).

Materials availability

This study did not generate new unique reagents.

Data and code availability

Original data have been deposited to Mendeley Data: <https://data.mendeley.com/datasets/xdnjy9p5zn/1>.

Convolutional recurrent neural network

Unlike static data without temporal context, sequential data such as the microstructure evolution trajectories in the form of image sequences require special treatment for deep neural networks to learn efficiently and accurately. Designed to take advantage of the temporal information of sequential inputs, the RNN along with its LSTM variants was first successfully employed in voice recognition and natural language processing. Recently, Shi et al.²⁸ proposed a convolutional RNN model to make full use of features in both spatial and temporal domains for image sequence prediction. Figure S9 compares the structures of CNN, RNN, and convolutional RNN. Unlike the vanilla RNN, a convolutional RNN uses a CNN instead of fully connected layers to extract latent features from input images and represent the system state in the convolutional latent space. When the system state is updated by the network at new times, it is passed through a decoder, which is typically a single convolutional layer without bias and activation, to generate predicted images in the real space.

More recent studies replace the initially stacked chain structure²⁸ with sophisticated neural nets to improve information flow and reach better performance.

For example, Yunbo Wang and coworkers developed a series of neural networks for spatiotemporal predictive learning.^{29–31} The latest E3D-LSTM model was employed in our study. Compared with other state-of-the-art models that use 2D convolution operations, the E3D-LSTM integrates 3D (one temporal and two spatial dimensions) convolution (3D-Conv) deep into RNNs, which is effective for modeling local representations in a consecutive manner. As shown in Figure 1C of Wang et al.,³¹ successive input frames are encoded by 3D-Conv encoders before being fed to the E3D-LSTM units. Outputs of E3D-LSTM units are decoded with a 3D-Conv layer to obtain the real-space predictions. In addition to adopting 3D-Conv as its basic operations, E3D-LSTM exploits a self-attention mechanism to memorize long-term interactions in addition to short-term motions. This is achieved by implementing two distinct memory states in E3D-LSTM: spatiotemporal memory and eidetic 3D memory. The former is designed to capture the short-term motion,²⁹ while the latter computes the relation between local patterns and the whole memory space to distinguish and revoke temporally distant memories.

Model setup

Each data point in the training sets is a sequence of N_t 2D images generated by a scalar field $c(t, x, y)$ ($0 \leq t \leq 1$, $t = 1 \dots N_t$, $x = 1 \dots N_x$, $y = 1 \dots N_y$). The spatial dimensions N_x and N_y are 64 unless otherwise stated. For each problem considered, the training dataset is a 4D array $c_i(t, x, y)$ with N_{total} image sequences ($i = 1 \dots N_{total}$). Following Wang et al.,³¹ four E3D-LSTM layers are stacked together in the model (only two layers in the case of grain growth), each with 64 hidden features. For spinodal decomposition, a normalization layer was added at the end to enforce mass conservation. The model is implemented in TensorFlow¹⁰⁰ and trained on four NVidia V100 or 1080-Ti GPUs. Typical training time is 36 h, with an initial learning rate of 10^{-3} that gradually decays to 10^{-5} . The training image size is chosen to be 64×64 because it is large enough to accommodate sufficient variation in microstructure configurations and also provides adequate resolution to resolve interfaces in microstructure with at least 1–2 pixels. During training, we started from about 400 image clips and gradually increased the number of clips until the model accuracy reached a plateau. The validation set is 1/4 the size of the training set, which is typical in NN training.

Data usage and augmentation

The whole dataset was randomly partitioned into three subsets: training, validation, and testing/prediction (e.g., at a ratio of 70:15:15). The validation set was used to monitor the convergence during training, while the testing or prediction set was completely withheld from training. The latter may also include customized sequences with different spatial/temporal dimensions and/or initial configurations. Training data were augmented by performing symmetry operations of the 2D point group $4mm$ on the original images, which transform $c(t, x, y)$ to $c(t, \bar{x}, y)$, $c(t, x, \bar{y})$, $c(t, \bar{x}, \bar{y})$, $c(t, y, x)$, $c(t, \bar{y}, x)$, $c(t, y, \bar{x})$, and $c(t, \bar{y}, \bar{x})$ ($\bar{x} \equiv N_x + 1 - x$, $\bar{y} \equiv N_y + 1 - y$). Such data transformations can be achieved by array rearrangements and do not require additional float-point calculations.

Analysis methods

RMSE and SSIM were used in pixel-wise comparison between ground truth and predictions. RMSE is defined as:

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^{N_x} \sum_{j=1}^{N_y} (p_g(i, j) - p_p(i, j))^2}{N_x N_y}}, \quad (\text{Equation 3})$$

where $p_g(i, j)$ and $p_p(i, j)$ are the pixel values of ground truth and predictions, respectively. SSIM⁷⁹ is defined as:

$$\text{SSIM} = \frac{(2\bar{p}_g\bar{p}_p + c_1)(2\sigma_{gp} + c_2)}{(\bar{p}_g^2 + \bar{p}_p^2 + c_1)(\sigma_g^2 + \sigma_p^2 + c_2)}, \quad (\text{Equation 4})$$

where $\bar{p}_k$ and σ_k ($k = g, p$) are the average pixel value and variance of ground truth or predictions, respectively, and σ_{gp} is their covariance. c_1 and c_2 are small constants and chosen to be $c_1 = (0.01L)^2$ and $c_2 = (0.03L)^2$, where L is the range of pixel values. The Euclidean distance between the distributions of quantity q from RNN predictions and ground truth is defined as:

$$d = \sqrt{\sum_{i=1}^n (q_i^g - q_i^p)^2}, \quad (\text{Equation 5})$$

where n is the number of bins within the interval between the minimum and the maximum of q , and q_i^g and q_i^p are normalized counts in the i -th bin of the ground truth and predictions, respectively. $n=20$ was used for all the calculations.

Simulation method

Phase-field simulations were employed to generate the ground truth for three microstructure evolution processes, i.e., grain growth, spinodal decomposition, and dendritic crystal growth. The phase-field method is a powerful computational technique for modeling microstructure evolution in diverse materials systems.^{2–4} In a phase-field model, different phases are represented by one or multiple order parameters, and their interfaces are tracked by the level sets of the order parameters. The spatiotemporal evolution of the microstructure is described by the governing equations of the order parameters derived from thermodynamic and kinetic principles.

Grain growth

Isotropic grain growth in a 2D polycrystalline structure was simulated by a multi-order-parameter phase-field model.⁸² In the model, a set of order parameters $\{\eta_1(x), \eta_2(x), \dots, \eta_N(x)\}$ is used to represent N distinct grain orientations. The free energy of the system is expressed as:

$$F = \int \left[f(\eta_1, \eta_2, \dots, \eta_N) + \frac{\nu}{2} \sum_i 1(\nabla \eta_i)^2 \right] dV, \quad (\text{Equation 6})$$

where the homogeneous free energy density f is given by:

$$f = m \left[\sum_i 1 \left(\frac{\eta_i^4}{4} - \frac{\eta_i^2}{2} \right) + \frac{3}{2} \sum_i 1 \sum_{j>i} \eta_i^2 \eta_j^2 + \frac{1}{4} \right], \quad (\text{Equation 7})$$

which has N local minima located at $(\eta_1, \eta_2, \dots, \eta_N) = (1, 0, \dots, 0), (0, 1, \dots, 0), \dots, (0, 0, \dots, 1)$. The evolution of $\eta_i(x)$ ($i = 1 \dots N$) follows the time-dependent Ginzburg-Landau or Allen-Cahn^{101,102} equation:

$$\frac{\partial \eta_i}{\partial t} = -L \frac{\delta F}{\delta \eta_i}. \quad (\text{Equation 8})$$

In all the simulations, the dimensionless parameters $N = 100$, $m = 1$, $\nu = 1$, and $L = 1$ were used. The initial polycrystalline structure was generated by Voronoi tessellation¹⁰³ with 100 grains. Equation 8 was solved by the forward Euler finite difference scheme with periodic boundary conditions and grid spacing $\Delta x = 1$ and time-step size $\Delta t = 0.2$. Single-channel images of the polycrystalline structure were generated by assigning $\sum_{i=1}^N \eta_i^3$ as the pixel value so that pixels were close to 0 in the grain boundary region and 1 inside the grains.

Spinodal decomposition

Spinodal decomposition was simulated by the C-H equation:⁸⁴

$$\frac{\partial c}{\partial t} = \nabla \cdot \left[Mc(1-c) \nabla \left(\frac{\partial f_{chem}}{\partial c} - \epsilon \nabla^2 c \right) \right], \quad (\text{Equation 9})$$

where c is the molar fraction of a species in a binary system. We used the regular solution model to describe the homogeneous free energy density:

$$f_{chem}(c) = RT[c \ln c + (1-c) \ln(1-c)] + \omega c(1-c), \quad (\text{Equation 10})$$

with a positive value assigned to the regular solution coefficient ω to favor phase separation. Equations 9 and 10 were solved with no-flux boundary conditions. The dimensionless parameter values $\omega = 0.27397$, $\epsilon = 0.1682$, and $M = 1$ and mesh spacing $\Delta x = 1$ were used in all of the simulations. Equation 9 was solved with an implicit variable-order backward differentiation formula (BDF) solver in COMSOL Multiphysics with an average dimensionless time-step size of 4.01. Images were output from simulations at a time interval of 1,500, or an average of 370 steps between two frames.

Dendrite growth

We used a phase-field model developed by Kobayashi⁸⁵ to simulate the dendritic solidification process in a pure materials system. Compared with other more quantitative models,^{4,104} this model was chosen for its simplicity, since

the purpose of this work was not to study dendritic growth but to use it as an example to evaluate the RNN. The system state was described by the temperature field T and an order parameter ϕ , which distinguishes between the solid ($\phi = 1$) and liquid ($\phi = 0$) phases. The free energy of the system is given by:

$$F[\phi, T] = \int \left[\frac{1}{2} \epsilon(\theta)^2 |\nabla \phi|^2 + f(\phi, T) \right] d\mathbf{r}, \quad (\text{Equation 11})$$

where the anisotropy of the solid/liquid interface energy is controlled by the orientation dependence of the gradient energy coefficient, $\epsilon(\theta) = \epsilon_0(1 + \delta \cos[n(\theta - \theta_0)])$, where θ represents the interface normal and is calculated from the gradient of ϕ as $\theta = \arctan(-\phi_y/\phi_x)$. We employed $n = 6$ in simulations to produce dendrites with 6-fold symmetry. f is a double-well potential:

$$f(\phi, T) = \frac{1}{4} \phi^4 - \left[\frac{1}{2} - \frac{1}{3} m(T) \right] \phi^3 + \left[\frac{1}{4} - \frac{1}{2} m(T) \right] \phi^2, \quad (\text{Equation 12})$$

$$m(T) = \frac{\alpha}{\pi} \arctan[\gamma(T_{eq} - T)], \quad (\text{Equation 13})$$

where T_{eq} is the solid/liquid equilibrium temperature. The time evolution of the coupled ϕ and T fields is governed by:

$$\tau \frac{\partial \phi}{\partial t} = -\frac{\delta F}{\delta \phi}, \quad (\text{Equation 14})$$

$$\frac{\partial T}{\partial t} = \nabla^2 T + K \frac{\partial \phi}{\partial t}, \quad (\text{Equation 15})$$

where constant K represents the latent heat. The following dimensionless parameters were used in all the simulations: $\alpha = 0.9$, $\gamma = 10$, $T_{eq} = 1$, $\tau = 0.001$, $\epsilon_0 = 0.01$, and $\delta = 0.03$. K and θ_0 varied. The system had a uniform initial temperature at $T(t = 0, x, y) = 0$. Equations 14 and 15 were solved with a variable-order BDF solver in COMSOL Multiphysics with mesh spacing $\Delta x = 1$ and average time-step size $\Delta t = 5.7 \times 10^{-4}$. Images were output from simulations at a time interval of 0.004, or an average of seven time steps between two frames.

Supplemental information description

More details can be found in the [supplemental information](#).

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.patter.2021.100243>.

ACKNOWLEDGMENTS

K.Y. and M.T. acknowledge support from the Department of Energy, Basic Energy Sciences under project DE-SC0019111. Y.C. and Y.Z. acknowledge support from the National Science Foundation under project CMMI-1929949. S.F. is supported by the ACS Petroleum Research Fund under grant 57523-DNI. The work of D.A., B.S., and F.Z. was supported by the Critical Materials Institute, an Energy Innovation Hub funded by the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, and Advanced Manufacturing Office, and performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Phase-field simulations were performed on supercomputers at the Texas Advanced Computing Center (TACC) at The University of Texas. RNN training and testing were performed on supercomputers at LLNL and TACC.

AUTHOR CONTRIBUTIONS

F.Z. conceived the project. F.Z. and M.T. supervised the project. K.Y., Y.Z., and H.C. performed phase-field simulations. K.Y., H.C., S.F., and F.Z. performed RNN training and testing. K.Y., H.C., Y.Z., S.F., M.T., D.A., B.S., and

F.Z. analyzed and discussed the results. M.T., F.Z., K.Y., and H.C. wrote the manuscript with input from other authors.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: November 15, 2020

Revised: March 2, 2021

Accepted: March 30, 2021

Published: April 22, 2021

REFERENCES

- Provatas, N., and Elder, K. (2010). *Phase-Field Methods in Materials Science and Engineering* (John Wiley & Sons, Ltd).
- Chen, L.Q. (2008). Phase-field models for microstructure evolution. *Annu. Rev. Mater. Res.* 32, 113–140.
- Moelans, N., Blanpain, B., and Wollants, P. (2008). An introduction to phase-field modeling of microstructure evolution. *Calphad* 32, 268–294.
- Steinbach, I. (2009). Phase-field models in materials science. *Model. Simulat. Mater. Sci. Eng.* 17, 073001.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep Learn. *Nat.* 521, 436–444.
- Raissi, M., Perdikaris, P., and Karniadakis, G.E. (2017). Physics Informed Deep Learning (Part I): Data-Driven Solutions of Nonlinear Partial Differential Equations. *arXiv*, arXiv:1711.10561.
- Raissi, M., Perdikaris, P., and Karniadakis, G.E. (2017). Physics Informed Deep Learning (Part II): Data-Driven Discovery of Nonlinear Partial Differential Equations. *arXiv*, arXiv:1711.10566.
- Brunton, S.L., Proctor, J.L., and Kutz, J.N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl. Acad. Sci. U S A* 113, 3932–3937.
- Rudy, S.H., Brunton, S.L., Proctor, J.L., and Kutz, J.N. (2017). Data-driven discovery of partial differential equations. *Sci. Adv.* 3, e1602614.
- Champion, K., Lusch, B., Kutz, J.N., and Brunton, S.L. (2019). Data-driven discovery of coordinates and governing equations. *Proc. Natl. Acad. Sci. U S A* 116, 22445–22451.
- Schaeffer, H. (2017). Learning partial differential equations via data discovery and sparse optimization. *Proc. R. Soc. A* 473, 20160446.
- Han, J., Jentzen, A., and Weinan, E. (2018). Solving high-dimensional partial differential equations using deep learning. *Proc. Natl. Acad. Sci. U S A* 115, 8505–8510.
- Raissi, M., Perdikaris, P., and Karniadakis, G.E. (2019). Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378, 686–707.
- Pathak, J., Hunt, B., Girvan, M., Lu, Z., and Ott, E. (2018). Model-free prediction of large spatiotemporally chaotic systems from data: a reservoir computing approach. *Phys. Rev. Lett.* 120, 024102.
- Breen, P.G., Foley, C.N., Boekholt, T., and Zwart, S.P. (2020). Newton versus the machine: solving the chaotic three-body problem using deep neural networks, *monthly notices R. Astronomical. Soc.* 494, 2465–2470.
- Moseley, B., Markham, A., and Nissen-Meyer, T. (2020). Solving the wave equation with physics-informed deep learning. *arXiv*, arXiv:2006.11894.
- Stevens, B., and Colonius, T. (2020). FiniteNet: a fully convolutional LSTM network architecture for time-dependent partial differential equations. *arXiv*, arXiv:2002.03014.
- Hu, Y., Zhao, T., Xu, Z., and Lin, L. (2020). Neural time-dependent partial differential equation. *arXiv*, arXiv:2009.03892.
- Rumelhart, D.E., Hinton, G.E., and Williams, R.J. (1988). Learning Representations by Back-Propagating Errors (MIT Press), pp. 696–699.
- Werbos, P.J. (1990). Backpropagation through time: what it does and how to do it. *Proc. IEEE* 78, 1550–1560.
- Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780.
- Sutskever, I., Martens, J., and Hinton, G.E. (2011). Generating Text with Recurrent Neural Networks (ICML).
- Cho, K., Van Merriënboer, B., Bahdanau, D., and Bengio, Y. (2014). On the properties of neural machine translation: Encoder-decoder approaches. *arXiv*, arXiv:1409.1259.
- A. Graves, N. Jaitly, Towards end-to-end speech recognition with recurrent neural networks, in: *International Conference on Machine Learning*, 2014, pp. 1764–1772.
- N. Srivastava, E. Mansimov, R. Salakhudinov, Unsupervised learning of video representations using lstms, in: *International Conference on Machine Learning*, 2015, pp. 843–852.
- J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, T. Darrell, Long-term recurrent convolutional networks for visual recognition and description, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 2625–2634.
- Finn, C., Goodfellow, I., and Levine, S. (2016). Unsupervised learning for physical interaction through video prediction. In *Advances in Neural Information Processing Systems*, pp. 64–72, NIPS'16: Proceedings of the 30th International Conference on Neural Information Processing Systems.
- Xingjian, S., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., and Woo, W.-c. (2015). Convolutional LSTM network: a machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems 28* (NIPS 2015), pp. 802–810.
- Wang, Y., Long, M., Wang, J., Gao, Z., and Yu, P.S. (2017). PredRNN: Recurrent Neural Networks for Predictive Learning Using Spatiotemporal LSTMs (NIPS).
- Wang, Y., Gao, Z., Long, M., Wang, J., and Yu, P.S. (2018). PredRNN++: Towards A Resolution of the Deep-In-Time Dilemma in Spatiotemporal Predictive Learning (ICML).
- Wang, Y., Jiang, L., Yang, M.-H., Li, L.-J., Long, M., and Li, F.-F. (2019). Eidetic 3D LSTM: A Model for Video Prediction and beyond (ICLR).
- Schmidt, J., Marques, M.R., Botti, S., and Marques, M.A. (2019). Recent advances and applications of machine learning in solid-state materials science. *NPJ Comput. Mater.* 5, 1–36.
- Rickman, J.M., Lookman, T., and Kalinin, S.V. (2019). Materials informatics: from the atomic-level to the continuum. *Acta Mater.* 168, 473–510.
- Ramprasad, R., Batra, R., Pilania, G., Mannodi-Kanakkithodi, A., and Kim, C. (2017). Machine learning in materials informatics: recent applications and prospects. *NPJ Comput. Mater.* 3, 1–13.
- Butler, K.T., Davies, D.W., Cartwright, H., Isayev, O., and Walsh, A. (2018). Machine learning for molecular and materials science. *Nature* 559, 547–555.
- Ryan, K., Lengyel, J., and Shatruk, M. (2018). Crystal structure prediction via deep learning. *J. Am. Chem. Soc.* 140, 10158–10168.
- Graser, J., Kauwe, S.K., and Sparks, T.D. (2018). Machine learning and energy minimization approaches for crystal structure predictions: a review and new horizons. *Chem. Mater.* 30, 3601–3612.
- Balachandran, P.V., Kowalski, B., Sehirlioglu, A., and Lookman, T. (2018). Experimental search for high-temperature ferroelectric perovskites guided by two-step machine learning. *Nat. Commun.* 9, 1–9.
- Ye, W., Chen, C., Wang, Z., Chu, I.-H., and Ong, S.P. (2018). Deep neural networks for accurate predictions of crystal stability. *Nat. Commun.* 9, 1–6.
- Faber, F.A., Lindmaa, A., Von Lilienfeld, O.A., and Armiento, R. (2016). Machine learning energies of 2 million elpasolite (ABC_2D_6) crystals. *Phys. Rev. Lett.* 117, 135502.

41. Xie, T., and Grossman, J.C. (2018). Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.* **120**, 145301.
42. Isayev, O., Oses, C., Toher, C., Gossett, E., Curtarolo, S., and Tropsha, A. (2017). Universal fragment descriptors for predicting properties of inorganic crystals. *Nat. Commun.* **8**, 1–12.
43. Yuan, F., and Mueller, T. (2017). Identifying models of dielectric breakdown strength from high-throughput data via genetic programming. *Sci. Rep.* **7**, 1–12.
44. Kim, C., Pilania, G., and Ramprasad, R. (2016). Machine learning assisted predictions of intrinsic dielectric breakdown strength of abx3 perovskites. *J. Phys. Chem. C* **120**, 14575–14580.
45. Carrete, J., Li, W., Mingo, N., Wang, S., and Curtarolo, S. (2014). Finding unprecedentedly low-thermal-conductivity half-Heusler semiconductors via high-throughput materials modeling. *Phys. Rev. X* **4**, 011019.
46. Pun, G.P., Batra, R., Ramprasad, R., and Mishin, Y. (2019). Physically informed artificial neural networks for atomistic modeling of materials. *Nat. Commun.* **10**, 1–10.
47. Behler, J. (2016). Perspective: machine learning potentials for atomistic simulations. *J. Chem. Phys.* **145**, 170901.
48. Botu, V., and Ramprasad, R. (2015). Learning scheme to predict atomic forces and accelerate materials simulations. *Phys. Rev. B* **92**, 094306.
49. Behler, J., and Parrinello, M. (2007). Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401.
50. DeCost, B.L., Francis, T., and Holm, E.A. (2017). Exploring the microstructure manifold: image texture representations applied to ultrahigh carbon steel microstructures. *Acta Mater.* **133**, 30–40.
51. DeCost, B.L., and Holm, E.A. (2015). A computer vision approach for automated analysis and classification of microstructural image data. *Comput. Mater. Sci.* **110**, 126–133.
52. Chowdhury, A., Kautz, E., Yener, B., and Lewis, D. (2016). Image driven machine learning methods for microstructure recognition. *Comput. Mater. Sci.* **123**, 176–187.
53. Ling, J., Hutchinson, M., Antono, E., DeCost, B., Holm, E.A., and Meredig, B. (2017). Building data-driven models with microstructural images: Generalization and interpretability. *Mater. Discov.* **10**, 19–28.
54. Azimi, S.M., Britz, D., Engstler, M., Fritz, M., and Mücklich, F. (2018). Advanced steel microstructural classification by deep learning methods. *Sci. Rep.* **8**, 1–14.
55. Stan, T., Thompson, Z.T., and Voorhees, P.W. (2020). Optimizing convolutional neural networks to perform semantic segmentation on large materials imaging datasets: X-ray tomography and serial sectioning. *Mater. Characterization* **160**, 110119.
56. Bulgarevich, D.S., Tsukamoto, S., Kasuya, T., Demura, M., and Watanabe, M. (2018). Pattern recognition with machine learning on optical microscopy images of typical metallurgical microstructures. *Sci. Rep.* **8**, 1–8.
57. Rickman, J.M., Wang, Y., Rollett, A.D., Harmer, M.P., and Compson, C. (2017). Data analytics using canonical correlation analysis and Monte Carlo simulation. *NPJ Comput. Mater.* **3**, 1–6.
58. Ziatdinov, M., Maksov, A., and Kalinin, S.V. (2017). Learning surface molecular structures via machine vision. *NPJ Comput. Mater.* **3**, 1–9.
59. Ziatdinov, M., Dyck, O., Maksov, A., Li, X., Sang, X., Xiao, K., Unocic, R.R., Vasudevan, R., Jesse, S., and Kalinin, S.V. (2017). Deep learning of atomically resolved scanning transmission electron microscopy images: chemical identification and tracking local transformations. *ACS Nano* **11**, 12742–12752.
60. Ding, G., Liu, Y., Zhang, R., and Xin, H.L. (2019). A joint deep learning model to recover information and reduce artifacts in missing-wedge sinograms for electron tomography and beyond. *Sci. Rep.* **9**, 1–13.
61. Mao, Y., Wang, X., Xia, S., Zhang, K., Wei, C., Bak, S., Shadik, Z., Liu, X., Yang, Y., Xu, R., et al. (2019). High-voltage charging-induced strain, heterogeneity, and micro-cracks in secondary particles of a nickel-rich layered cathode material. *Adv. Funct. Mater.* **29**, 1900247.
62. Zheng, C., Chen, C., Chen, Y., and Ong, S.P. (2020). Random forest models for accurate identification of coordination environments from x-ray absorption near-edge structure. *Patterns* **1**, 100013.
63. Kautz, E.J., Hagen, A.R., Johns, J.M., and Burkes, D.E. (2019). A machine learning approach to thermal conductivity modeling: a case study on irradiated uranium-molybdenum nuclear fuels. *Comput. Mater. Sci.* **161**, 107–118.
64. Rickman, J., Chan, H., Harmer, M., Smeltzer, J., Marvel, C., Roy, A., and Balasubramanian, G. (2019). Materials informatics for the screening of multi-principal elements and high-entropy alloys. *Nat. Commun.* **10**, 1–10.
65. Fullwood, D.T., Niezgoda, S.R., Adams, B.L., and Kalidindi, S.R. (2010). Microstructure sensitive design for performance optimization. *Prog. Mater. Sci.* **55**, 477–562.
66. Kalidindi, S.R., Niezgoda, S.R., and Salem, A.A. (2011). Microstructure informatics using higher-order statistics and efficient data-mining protocols. *JOM* **63**, 34–41.
67. Gupta, A., Cecen, A., Goyal, S., Singh, A.K., and Kalidindi, S.R. (2015). Structure–property linkages using a data science approach: application to a non-metallic inclusion/steel composite system. *Acta Mater.* **91**, 239–254.
68. Gusenbauer, M., Oezelt, H., Fischbacher, J., Kovacs, A., Zhao, P., Woodcock, T.G., and Schrefl, T. (2020). Extracting local nucleation fields in permanent magnets using machine learning. *NPJ Comput. Mater.* **6**, 1–10.
69. Kondo, R., Yamakawa, S., Masuoka, Y., Tajima, S., and Asahi, R. (2017). Microstructure recognition using convolutional neural networks for prediction of ionic conductivity in ceramics. *Acta Mater.* **141**, 29–38.
70. Cecen, A., Dai, H., Yabansu, Y.C., Kalidindi, S.R., and Song, L. (2018). Material structure-property linkages using three-dimensional convolutional neural networks. *Acta Mater.* **146**, 76–84.
71. Kautz, E., Ma, W., Jana, S., Devaraj, A., Joshi, V., Yener, B., and Lewis, D. (2020). An image-driven machine learning approach to kinetic modeling of a discontinuous precipitation reaction. *Mater. Characterization* **166**, 110379.
72. Yang, Z., Li, X., Catherine Brinson, L., Choudhary, A.N., Chen, W., and Agrawal, A. (2018). Microstructural materials design via deep adversarial learning methodology. *J. Mech. Des.* **140**, 111416.
73. X. Li, Z. Yang, L.C. Brinson, A. Choudhary, A. Agrawal, W. Chen, A deep adversarial learning methodology for designing microstructural material systems, in: ASME 2018 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, American Society of Mechanical Engineers Digital Collection, 2018.
74. Ma, W., Kautz, E.J., Baskaran, A., Chowdhury, A., Joshi, V., Yener, B., and Lewis, D.J. (2020). Image-driven discriminative and generative machine learning algorithms for establishing microstructure–processing relationships. *J. Appl. Phys.* **128**, 134901.
75. Liu, R., Kumar, A., Chen, Z., Agrawal, A., Sundararaghavan, V., and Choudhary, A. (2015). A predictive machine learning approach for microstructure optimization and materials design. *Sci. Rep.* **5**, 1–12.
76. Liu, G., Jia, L., Kong, B., Feng, S., Zhang, H., and Zhang, H. (2017). Artificial neural network application to microstructure design of nb-si alloy to improve ultimate tensile strength. *Mater. Sci. Eng. A* **707**, 452–458.
77. Exl, L., Fischbacher, J., Kovacs, A., Oezelt, H., Gusenbauer, M., Yokota, K., Shoji, T., Hrkac, G., and Schrefl, T. (2018). Magnetic microstructure machine learning analysis. *J. Phys. Mater.* **2**, 014001.
78. Farimani, A.B., Gomes, J., Sharma, R., Lee, F.L., and Pande, V.S. (2018). Deep Learning Phase Segregation. *arXiv*, arXiv:1803.08993.

79. Wang, Z., Bovik, A.C., Sheikh, H.R., and Simoncelli, E.P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**, 600–612.
80. Von Neumann, J., and Herring, C. (1952). *Metal Interfaces* (American Society for Metals), p. 108.
81. Mullins, W.W. (1956). Two-dimensional motion of idealized grain boundaries. *J. Appl. Phys.* **27**, 900–904.
82. Moelans, N., Blanpain, B., and Wollants, P. (2008). Quantitative analysis of grain boundary properties in a generalized phase field model for grain growth in anisotropic systems. *Phys. Rev. B* **78**, 024113.
83. Balluffi, R.W., Allen, S.M., and Carter, W.C. (2005). *Kinetics of Materials* (John Wiley & Sons).
84. Cahn, J.W., and Hilliard, J.E. (1958). Free energy of a nonuniform system. i. interfacial free energy. *J. Chem. Phys.* **28**, 258–267.
85. Kobayashi, R. (1993). Modeling and numerical simulations of dendritic crystal growth. *Phys. D Nonlinear Phenomena* **63**, 410–423.
86. Vaithyanathan, V., Wolverton, C., and Chen, L. (2004). Multiscale modeling of θ' precipitation in al–cu binary alloys. *Acta Mater.* **52**, 2973–2987.
87. Gránásy, L., Pusztai, T., Börzsönyi, T., Warren, J.A., and Douglas, J.F. (2004). A general mechanism of polycrystalline growth. *Nat. Mater.* **3**, 645–650.
88. Goerler, J., Lopez-Galilea, I., Roncery, L.M., Shchyglo, O., Theisen, W., and Steinbach, I. (2017). Topological phase inversion after long-term thermal exposure of nickel-base superalloys: Experiment and phase-field simulation. *Acta Mater.* **124**, 151–158.
89. Das, S., Tang, Y., Hong, Z., Gonçalves, M., McCarter, M., Klewe, C., Nguyen, K., Gómez-Ortiz, F., Shafer, P., Arenholz, E., et al. (2019). Observation of room-temperature polar skyrmions. *Nature* **568**, 368–372.
90. Xu, R., Liu, S., Saremi, S., Gao, R., Wang, J., Hong, Z., Lu, H., Ghosh, A., Pandya, S., Bonturim, E., et al. (2019). Kinetic control of tunable multi-state switching in ferroelectric thin films. *Nat. Commun.* **10**, 1–10.
91. Tang, M., and Carter, W.C. (2013). Branching mechanisms in surfactant micellar growth. *J. Phys. Chem. B* **117**, 2898–2905.
92. Yang, K., and Tang, M. (2020). Three-dimensional phase evolution and stress-induced non-uniform Li intercalation behavior in lithium iron phosphate. *J. Mater. Chem. A* **8**, 3060–3070.
93. Liu, D., Shadike, Z., Lin, R., Qian, K., Li, H., Li, K., Wang, S., Yu, Q., Liu, M., Ganapathy, S., et al. (2019). Review of recent development of in situ/operando characterization techniques for lithium battery research. *Adv. Mater.* **31**, 1806620.
94. Wu, Y., and Liu, N. (2018). Visualizing battery reactions and processes by using in situ and in operando microscopies. *Chem* **4**, 438–465.
95. Shahani, A.J., Xiao, X., Lauridsen, E.M., and Voorhees, P.W. (2020). Characterization of metals in four dimensions. *Mater. Res. Lett.* **8**, 462–476.
96. Lim, J., Li, Y., Alsem, D.H., So, H., Lee, S.C., Bai, P., Cogswell, D.A., Liu, X., Jin, N., Yu, Y.-s., et al. (2016). Origin and hysteresis of lithium compositional spatiodynamics within battery primary particles. *Science* **353**, 566–571.
97. Hong, L., Li, L., Chen-Wiegart, Y.-K., Wang, J., Xiang, K., Gan, L., Li, W., Meng, F., Wang, F., Wang, J., et al. (2017). Two-dimensional lithium diffusion behavior and probable hybrid phase transformation kinetics in olivine lithium iron phosphate. *Nat. Commun.* **8**, 1–13.
98. Zhang, S., Guo, S., Huang, W., Scott, M.R., and Wang, L. (2020). V4D: 4D convolutional neural networks for video-level representation learning. *arXiv*, arXiv:2002.07442.
99. C. Choy, J. Gwak, S. Savarese, 4d spatio-temporal convnets: Minkowski convolutional neural networks, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3075–3084.
100. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., et al. (2016). TensorFlow: large-scale machine learning on heterogeneous distributed systems. *arXiv*, arXiv:1603.04467.
101. Allen, S.M., and Cahn, J.W. (1972). Ground state structures in ordered binary alloys with second neighbor interactions. *Acta Metall.* **20**, 423–433.
102. Allen, S.M., and Cahn, J.W. (1973). A correction to the ground state of fcc binary ordered alloys with first and second neighbor pairwise interactions. *Scripta Metall.* **7**, 1261–1264.
103. Aurenhammer, F., and Klein, R. (2000). Voronoi diagrams. *Handbook Comput. Geometr.* **5**, 201–290.
104. Karma, A., and Rappel, W.-J. (1998). Quantitative phase-field modeling of dendritic growth in two and three dimensions. *Phys. Rev. E* **57**, 4323.