Neural cellular automata for solidification microstructure modelling

Tang, Jian; Kumar, Siddhant; De Lorenzis, Laura; Hosseini, Ehsan

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Neural cellular automata for solidification microstructure modelling

Jian Tang[a,b], Siddhant Kumar[c], Laura De Lorenzis[b,*], Ehsan Hosseini[a,*]

[a] *Empa Swiss Federal Laboratories for Materials Science & Technology, 8600 Dübendorf, Switzerland*
[b] *Department of Mechanical and Process Engineering, ETH Zürich, 8092 Zürich, Switzerland*
[c] *Department of Materials Science and Engineering, Delft University of Technology, 2628 CD Delft, The Netherlands*

Dataset link: https://github.com/HighTempIntegrity/JianTang-NCA01

Abstract

We propose Neural Cellular Automata (NCA) to simulate the microstructure development during the solidification process in metals. Based on convolutional neural networks, NCA can learn essential solidification features, such as preferred growth direction and competitive grain growth, and are up to six orders of magnitude faster than the conventional Cellular Automata (CA). Notably, NCA deliver reliable predictions also outside their training range, e.g. for larger domains, longer solidification duration, and different temperature fields and nucleation settings, which indicates that they learn the physics of the solidification process. While in this study we employ data produced by CA for training, NCA can be trained based on any microstructural simulation data, e.g. from phase-field models.

© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Keywords: Microstructure modelling; Convolutional neural networks; Computational speed; Cellular automata

1. Introduction

The microstructure of metals collectively denotes their grains and metallurgical phases, their size, relative amount, shape/morphology and crystallographic orientation (texture). The microstructure strongly influences the physical and mechanical response of metallic materials, and a wide range of properties can be achieved by manipulating the microstructural features through thermomechanical processing. Therefore, correlating processing conditions and the resulting microstructure has always been an important research field. Understanding the influence of the solidification process conditions on the features of the resulting microstructure has attracted even more attention since the advent of metal additive manufacturing (MAM) processes that allow site-specific control of the solidification conditions [1,2]. However, achieving the goal of microstructure control, e.g. fabricating functionally graded materials, via MAM requires reliable solidification microstructure models that enable quantitative prediction of the process-microstructure relationship.

Microstructure formation during the solidification process comprises two phenomena: nucleation and grain growth [3]. Nucleation is the formation of stable nuclei in the molten metal and typically requires that the melt

---

cools down below its melting temperature (undercooled). Growth includes enlargement of the nuclei (or existing grains) depending on the temperature field and microstructure state in the vicinity of the solidification interface [3–6]. The formation of solid material during growth is accompanied by the release of latent heat of fusion which should be extracted from the solidification interface to allow further growth. From a thermal point of view, this indicates that the growth direction aligns with the maximum temperature gradient direction, where the fastest heat dissipation through conduction takes place [3]. While crystallographic orientation considerations are not relevant for nucleation and formed nuclei have random orientations [7,8], the growth phenomenon is anisotropic and nuclei/grains grow faster along certain crystallographic orientations (i.e. *preferred growth directions*, <100> for BCC/FCC crystals) [3]. Thermal and crystallographic factors collectively result in the selective growth of nuclei/grains or the so-called *competitive grain growth* phenomenon, which means that the nuclei/grains whose preferred growth directions are aligned with the maximum temperature gradient grow faster and become the dominant feature of the microstructure [3]. Consideration of the above phenomena is the essential part of the solidification microstructure modelling strategies [3,9–17].

Besides simple analytical/empirical models [3,9,10], the two 'physics-based' microstructure modelling techniques are Phase Field Modelling (PFM) and Cellular Automata (CA). The main focus of these two modelling strategies is on representing the grain growth phenomena, while they often adopt empirical probabilistic or ad-hoc considerations for the nucleation phenomena [7,18]. PFM for microstructure modelling [19,20] represents the microstructure by a set of conserved, e.g. for solute atom concentration, and non-conserved phase field parameters, e.g. for the state (solid/liquid) and grain ID (grains with various crystallographic orientation). PFM defines the free energy of the system as a function of the conserved and non-conserved phase field parameters, their gradients, temperature, and possibly additional variables. An anisotropic grain boundary energy is considered to account for the preferred growth direction. This energy term is maximum for grain boundaries perpendicular to the crystal preferred growth directions so that grains tend to grow anisotropically along these directions [20,21]. The Cahn–Hilliard [22,23] and Allen–Cahn [19,24,25] equations are respectively used to govern the evolution of conserved and non-conserved phase-field parameters such that the free energy of the system tends to its minimum. The kinetics of the evolution depends on the reduction rate of the system's free energy and on temperature-dependent mobility factors. PFM is a very detailed microstructural modelling strategy that not only well represents important phenomena such as preferred growth directions and competitive grain growth, but also can provide predictions about microscale features such as secondary dendrite growth morphology and solute atom distribution [11,12,26]. However, the need for space and time resolutions in the orders of 0.01–1 μm and 0.1–10 μs leads to a very high computational cost for PFM, which limits the simulation domain size to hundreds of micrometres [11–14,26–29]. As an example, Yan et al. [13] report a computational time of 13 days for the solidification microstructure simulation of a $350 \times 350 \times 150 \ \mu m^3$ domain using an NVIDIA Tesla M2090 GPU.

The CA method, developed by Gandin and Rappaz [15,16], evaluates the microstructure and temperature state locally in the vicinity of the solidification interface and combines crystallographic considerations and experimentally or theoretically driven dendrite tip growth rate equations to predict the grain growth phenomenon [4–6]. Numerous examples of the application of CA for investigating the microstructure development during casting [16,17,30] and MAM [31–33] demonstrate the experimental relevance of the CA method and its ability for representing the preferred grain growth and competitive grain growth phenomena. CA is often used to predict grain size, morphology, and crystallographic texture [31,32,34], although there are a few studies that extend CA for predicting solute atom distribution [35,36]. The typical space and time resolutions for CA simulations are 0.1–10 μm and 0.5–100 μs, respectively. In comparison with PFM, the coarser discretization and the rather local and simple calculations lead to lower computational cost for CA, such that the simulation domain size can be up to several cubic millimetres for a few days of computations [31–34]. Still, it is obvious that even CA cannot be used for large-scale simulations or, importantly, for optimization problems where numerous simulation runs are required to find the optimum set of process conditions for a desired type of microstructure. Therefore, there is a substantial need to develop novel microstructure simulation techniques which retain physical relevance and reliability but reduce the computational cost by orders of magnitude.

Deep learning has recently emerged as a central tool for scientific computing, and there are many successful examples where a deep neural network (DNN) is trained to serve as a highly efficient simulation tool [37–39]. A common DNN-based strategy for microstructure modelling is to project the microstructure onto a low-dimensional space by utilizing statistical descriptors and/or dimensionality reduction techniques, and to use DNNs to describe

their evolution in this latent space [40–42]. However, the employment of simple statistical descriptors, such as spatial correlation functions of phases or concentration distributions, may compromise the accuracy of predictions and impede the back transformation to microstructure images [40–44]. Alternatively, advanced auto-encoders can be used to reduce the dimensionality and reconstruct from the latent space back to the original microstructure with minimal loss of information [45,46]. However, modelling the development of the microstructure in a latent space precludes the direct integration of physical constraints and laws into the DNN, resulting in black-box solutions with limited applicability beyond their trained domain.

In this paper, we propose a new DNN-based approach, denoted as neural cellular automata (NCA), to predict microstructure development during the solidification process. It is similar to the classical microstructure modelling frameworks (e.g. CA and PF) that 'graphically' describe the evolution of microstructure, and allows for the direct incorporation of physical concepts. Our NCA idea is inspired by the work of Mordvintsev et al. [47,48], who integrated a fully connected neural network into the CA framework to reconstruct or repair dedicated image patterns seen during the training. Similar to their work [47,48], our NCA is an extension of the conventional CA, where a convolutional neural network (CNN) acts as a highly efficient and flexible 'brain' for governing the evolution of cells.

2. From cellular automata (CA) to neural cellular automata (NCA)

The CA microstructural modelling framework discretizes the simulation domain into a series of cubic/square cells, each containing a set of microstructure-related parameters such as phase state, crystallographic orientation, temperature, and possibly more [31,32]. The phase state can be either solid (S), liquid (L), or growing (G). The (G) cells are located between the (S) and (L) cells, representing the solidification interface. Three Euler angles describe the crystallographic orientation of (S) and (G) cells; these are angles between the global reference coordinate system and the principal directions of the crystal at the cells (i.e. <100> crystallographic orientations, which for BCC/FCC coincide with their preferred growth directions). A number of simple rules then govern the evolution of the parameters of each cell, based on information from the cell itself and from those in its local neighbourhood. The rules are devised such that the physics of grain growth during solidification is reproduced. For example, a (G) or (S) cell with a temperature higher than the melting point transforms to the (L) state, or a (G) cell with no (L) cell in the neighbourhood changes to the (S) state. As essential constituent, CA employ the so-called 'decentred octahedron method' [15,16] to consider the preferred growth directions. For BCC/FCC systems, the 'decentred octahedron' method assumes an octahedron for each (G) cell whose diagonals are parallel to the <100> directions of the cell's crystal. A theoretically or empirically driven equation for dendrite tip growth rate takes the cell's temperature and defines the growth of the octahedron along its diagonals. The octahedron envelope finally grows into neighbouring cells, and when it captures the centre of a neighbour liquid cell, its state (L) transforms to (G). The new growing cell inherits the same crystallographic orientation as the parent cell and starts to grow its own octahedron. The growth of octahedra along <100> crystallographic orientations and the dependency of their growth rate on the local temperature ensure that CA reproduce the phenomena of preferred growth directions and competitive grain growth in solidification microstructure modelling. The 'decentred octahedron method' schematic and the detailed algorithm of CA are given in Appendix A.

There are factors which constrain the computational efficiency of CA. The incremental progress of CA simulations should be small to ensure that the 'capture' of the neighbouring cells – i.e. the transformation of the (L) neighbouring cells into the (G) state due to growth – is accurately modelled. The time increment is often limited to $\Delta t = \alpha d / v_{max}$, where d: cell size, α: a constant between $(0, 1)$, and v_{max}: maximum growth rate among all growing cells in the simulation domain. Additionally, the consideration of information from neighbouring cells is typically limited to their immediate surroundings and is done through a 'for loop' method. NCA enhances the computational efficiency and, potentially, the flexibility and physical relevance of CA through the power of CNNs.

Fig. 1 illustrates a schematic representation of the NCA modelling strategy. Similar to CA, NCA discretize the simulation domain into cubic/square pixels and incrementally (i.e. frame after frame) update the microstructure based on the temperature and microstructure states in the local neighbourhoods of individual pixels. The required input data to initialize NCA are the locations of the nuclei, their size and their crystallographic orientation. Also required is the temperature evolution within the simulation domain during the solidification process; the temperature data are fed to the network as undercooling magnitude (i.e. $\Delta T = T_{melting} - T$). A continuous variable $P \in [0,1]$ is considered for each pixel to represent its phase state, where $P = 0$ and $P = 1$ indicate fully liquid and solid
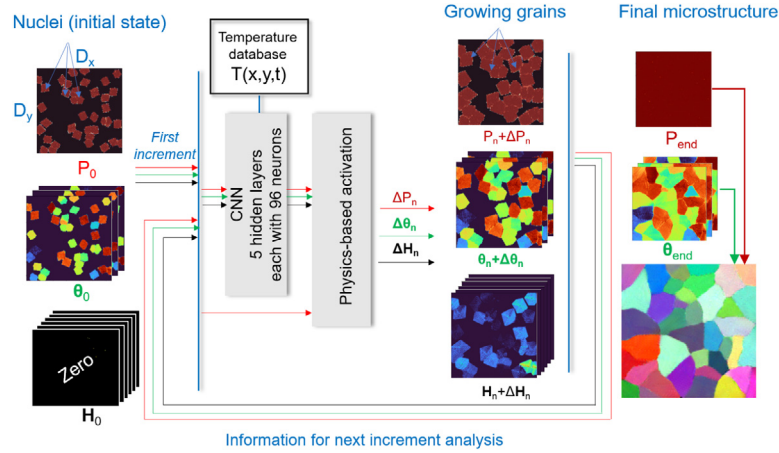
Fig. 1. Schematic of the NCA modelling strategy.

states, respectively. Similar to CA, the crystallographic orientations of the pixels are defined with three Euler angles $\boldsymbol{\theta} = (\alpha_1, \alpha_2, \alpha_3)$, i.e. the angles between the principal directions of the solidified pixel and the reference coordinate system.

NCA predicts the sequential microstructure evolution by recurrently employing a trained CNN. The CNN takes the temperature, phase state, three Euler angles and six hidden channels as inputs and delivers pixel-wise increments of the phase state, Euler angles and hidden channels (i.e. ΔP, $\Delta \theta$ and ΔH, respectively) as output. Apart from the phase state and Euler angles, the hidden channels allow the CNN to exchange information between different frames and better track the evolution of microstructural features. The physical meaning of the hidden channels is not explicitly defined and they have zero value at the start of the analysis. The number of hidden channels and the CNN architecture are decided based on a hyperparameter tuning exercise (see Section 3.4). Accordingly, six hidden channels and 1+5+1 convolutional layers with ReLU activation function and 96 neurons per hidden layer, including 3×3 input/hidden-layer kernels and 1×1 output-layer kernels, are utilized in the NCA. The details of the adopted CNN and NCA training are given in Section 3.2 and Section 3.3, respectively.

To improve the physical relevance of the algorithm, the NCA integrates a 'physics-based activation function' at the output layer to prevent the CNN from violating explicit known physical laws of solidification. This 'activation function' directly encodes the knowledge that the grain growth mainly takes place at the solidification interface; it enforces no change of microstructure states for a pixel when: (i) $P < 0.1$ and $T > T_{melting}$, i.e. the liquid pixel is at a temperature above the melting point, (ii) no neighbour pixel with $P \geq 0.1$, i.e. not in the vicinity of the solidification interface, (iii) $P > 0.99$ and $T < T_{melting}$, i.e. the solid pixel temperature is below the melting point. Values of 0.1 and 0.99 are the thresholds for the phase state P indicating liquid and solid. They can be seen as network hyperparameters and are taken from [47,48].

The adoption of multi-layer convolution enables NCA to implicitly consider information from large neighbourhoods and integrate complex and highly nonlinear rules for governing the evolution of the microstructural parameters of the pixels. This leads to an advantage of NCA over CA in encoding the complexity of the solidification process, e.g. when an NCA model is trained based on data from PFM. Section 5 also discusses the significantly higher computational speed of NCA over CA.

3. Methods

3.1. Materials

In this work, we choose Hastelloy X as a model alloy to investigate the applicability of NCA in predicting its microstructure development under various solidification conditions. The generation of NCA training and validation data through CA necessitates information about the preferred growth direction and dendrite growth kinetics. Hastelloy X is a nickel-based alloy with an FCC crystal structure (solid-solution strengthened) and hence has the
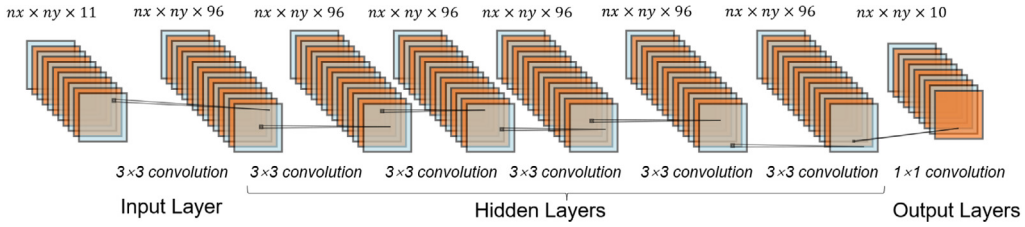
Fig. 2. Schematic of selected CNN architecture for NCA (image generated by the online tool NNSVG). The numbers at the top represent the dimensions of the input/output for each convolutional layer, and the bottom numbers indicate the size of each convolutional kernel.

preferred growth direction along the <001> orientation [3,49,50]. Meanwhile, the Kurz–Giovanola–Trivedi (KGT) model [5] is used to define the dendrite tip growth rate of Hastelloy X as a function of undercooling as:

$$v(\Delta T) = 4.60 \times 10^{-10} \times \Delta T^2 + 3.15 \times 10^{-6} \times \Delta T^3 \tag{1}$$

Further details about the derivation of the above KGT model for Hastelloy X are described in Appendix B.

3.2. Convolutional neural networks for NCA

The architecture of the CNN in NCA is composed of fully connected convolutional layers, as depicted in Fig. 2. Each convolutional layer consists of a kernel that performs pixel-wise convolution. The mathematical transformation in a convolutional layer is:

$$\mathbf{x}_{ijk}^{l+1} = f(\sum_o \sum_p \sum_q \mathbf{w}_{opqk}^l \mathbf{x}_{i\pm o, j\pm p, q}^l + \mathbf{b}_k^l) \tag{2}$$

where $\mathbf{w}$ and $\mathbf{b}$ denote the trainable convolutional kernel weights and bias, f represents a non-linear activation function, $\mathbf{x}^l$ is the input of the current convolutional layer and $\mathbf{x}^{l+1}$ is its output, which serves as the input for the next layer.

In brief, CNN can be considered a neural network 'function' that decides the evolution in the state of a pixel by analysing the information from its neighbouring pixels. The weights and biases of the network are determined through network training, where a high-dimensional optimization problem is solved to minimize the difference between the CNN predictions and the ground truth data generated by the CA (as explained in Section 3.3). To achieve the best performance, the hyperparameters of CNN, such as the number of hidden layers and the activation function type, need to be optimized (as discussed in Section 3.4).

3.3. NCA training & testing

Table 1 shows that different variants of NCA are trained to represent Hastelloy X solidification microstructure under different conditions. Ground truth data are created using CA simulations with a domain size of 55×55 μm^2 for each solidification condition. The data are divided into a 9:1 ratio for training and validation [51,52]. Additionally, ten further CA simulations are performed to evaluate the generalization error of NCA variants under unseen cases, e.g. larger domain sizes, longer solidification durations, different temperature fields and nucleation settings. The orientation and location of nuclei are randomly assigned in the CA simulations for generating training, validation, and testing data except explicitly mentioned otherwise (i.e. single grain growth). The CA simulations adopt a time increment of $\Delta t = 8$ μs, except for quasi-3D cases, where $\Delta t = 40$ μs.

During training, the CA-predicted phase state and Euler angles, along with undercooling values, are used as inputs for the NCA (the Euler angles and undercooling are normalized by 90° and 100 K, respectively). The CNN is initialized with zero biases and random weights and kernel parameters (except for the output layer). The training uses the Adam optimizer (learning rate of 10^{-3} decaying to 9×10^{-5}) for 7000 epochs and the L2 loss defined by Eq. (3). To achieve faster convergence and following [47,48], the weight and bias of the output layer are initialized

Table 1

Details of the NCA training for different solidification conditions.

NCA model variant	Single-grain growth[a]	Multi-grain growth	Training time[b]
Isothermal 2D single-grain growth[c]	18 CA simulations	–	0.8 h
Isothermal 2D multi-grain growth	18 CA simulations	27 CA simulations	2 h
Isothermal quasi-3D multi-grain growth	729 CA simulations	929 CA simulations	72 h
Non-isothermal 2D multi-grain growth[d]	180 CA simulations	270 CA simulations	24 h

[a]The domain size is 55×55 μm² and the nucleation density for the multi-grain simulations is 8260–11570 mm⁻². The nuclei locate at the centre of the domain for the single-grain growth samples. It is observed that including single-grain growth data enhances the training efficiency for representing multi-grain growth (See Appendix A).

[b]An NVIDIA RTX A5000 GPU is employed.

[c]A constant undercooling of 20 K is used for isothermal simulations.

[d]Non-isothermal simulations include solidification with 45 different nuclei settings (including single and multi-nuclei) under 10 different temperature fields with undercoolings in the range of 20–25 K.

with zero and the gradients of the loss are normalized by their Frobenius norms at every epoch (to prevent gradient explosion).

$$l = \sum_{Frame} \left[\sum_{Solidpixel} \left[(P^{NCA} - P^{CA})^2 + \| \boldsymbol{\theta}^{NCA} - \boldsymbol{\theta}^{CA} \|^2 \right] / N_{Solidpixel} \right] / N_{Frame} \tag{3}$$

Here N_{Frame} and $N_{Solidpixel}$ are the numbers of the time frame and solid pixel in each frame (pixel with $P > 0.99$), respectively. Finally, NCA's accuracy is evaluated by comparing its results with CA predictions:

$$Accuracy = 1 - N_{\Delta\phi(CA,NCA)>15°} / N_{total} \tag{4}$$

where N_{total} is the total number of pixels and $N_{\Delta\phi(CA,NCA)>15°}$ is the number of pixels with orientation error $\Delta\phi$ exceeding 15° (i.e. the threshold of establishing grain boundary).

3.4. Hyperparameter tuning

A hyperparameter tuning exercise is conducted to find the optimal CNN architecture. This involves training and evaluating multiple neural networks with different architectures, varying the number of hidden convolutional layers and neurons per hidden layer. We keep other details, such as kernel size, learning rate, and activation function, consistent with Refs. [47,48].

The hyperparameter tuning exercise is conducted solely for non-isothermal solidification data, and the CNN architecture that yields the best results is used for all other conditions. To reduce the computational costs of hyperparameter tuning, we limit the number of epochs for training each model to 2500 and use a fast-decaying learning rate schedule (70% decay every 1000 epochs). This approach ensures that each model is trained in less than 0.5 h.

Table 2 shows that the best training and validation accuracy is achieved with five hidden layers and 96 neurons per hidden layer and hence is adopted for this study. The impact of varying the number of input hidden channels on CNN accuracy is also examined which indicates that the adoption of six hidden channels produces the best results (Table 3).

4. NCA results

4.1. Isothermal solidification: 2D single-grain growth

Initially, the NCA model is trained with 2D CA isothermal simulation data for a domain size of 55×55 μm² (55×55 pixels). A single nucleus with a random α_1 and $\alpha_2 = \alpha_3 = 0$ is placed in the centre of the domain with a uniform and steady temperature field representing an undercooling of 20 K. The model achieves training and validation accuracies of 98.0% and 97.7%, respectively, which indicates that it learns well the kinetics of grain growth and the preferred growth direction mechanism. Furthermore, the trained model is used to predict the solidification microstructure for a 3× larger domain and a 5× longer solidification duration with the nucleus located off the centre. As presented in Fig. 3, the NCA can predict such cases with an accuracy of 99.6 ± 0.2%.

Table 2

Training and validation accuracies of the NCA model with various numbers of hidden layers and neurons per hidden layer.

Number of hidden layers	Number of neurons per hidden layer	Training accuracy (%)	Validation accuracy (%)
9	32	95.0	93.5
12	32	95.0	93.3
9	64	96.1	93.2
12	64	95.7	92.9
3	96	96.8	95.1
5	96	96.7	95.7
7	96	96.0	94.4

Table 3

Training and validation accuracies of the NCA model with various numbers of hidden channels.

| Number of hidden channels	Training accuracy (%)	Validation accuracy (%)
0	96.6	96.1
3	97.3	96.1
6	97.1	96.4
9	97.1	96.3
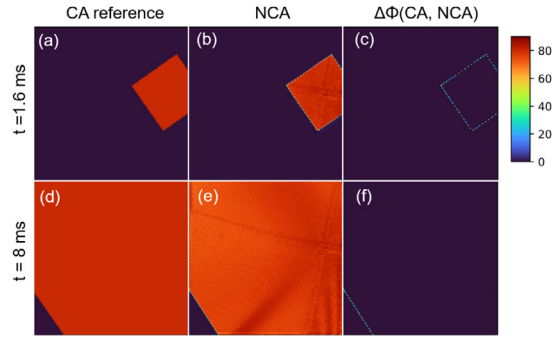
Fig. 3. Test case for 2D isothermal single-grain growth in a domain of 160×160 μm^2 with a randomly selected nucleus location: CA results (a and d), NCA results (b and e), and orientation difference $\Delta\phi$ between them (c and f). Note that the domain size and time duration of the test are 3× larger and 5× longer than those adopted for the training.

4.2. Isothermal solidification: 2D multi-grain growth

In the second stage, the data from additional multi-grain CA simulations are added to the previous ones for training and validation of the NCA. In these data nuclei are placed randomly in the 55×55 μm^2 domain with a nuclei density in the range of 8260–11,570 mm^{-2}. Similar to the previous simulations, the 2D nature of the simulations allows considering $\alpha_2 = \alpha_3 = 0$ while α_1 takes a random value. The assessment of the NCA model for multi-grain simulations indicates training and validation accuracies of 99.9% and 96.5%. The trained model further shows a test accuracy of $98.0 \pm 0.9\%$ in predicting multi-grain growth in a larger domain. Fig. 4 illustrates NCA and CA results for a test case, showing that NCA well reproduces the microstructure evolution and the effect of nuclei density on the final grain size. For both single and multi-grain growth cases, the main inconsistency between NCA and CA is observed at the grain boundaries and solid–liquid interfaces. While CA represents the grain boundaries as a sharp interface, the involvement of the convolution operator in the NCA leads to a diffusive interface at the grain boundaries.

To obtain grain size distributions from the microstructures generated by NCA (shown in Fig. 4g), it was essential to establish a clear boundary between grain domains and hence, a post-processing algorithm was devised to address this requirement. The algorithm corrects the Euler angle of each pixel by replacing it with the Euler angle of a nearby nucleus that displays the most similar orientation. This correction serves to convert the diffuse interface into a well-defined sharp boundary. For more comprehensive information about this post-processing technique, refer to E.
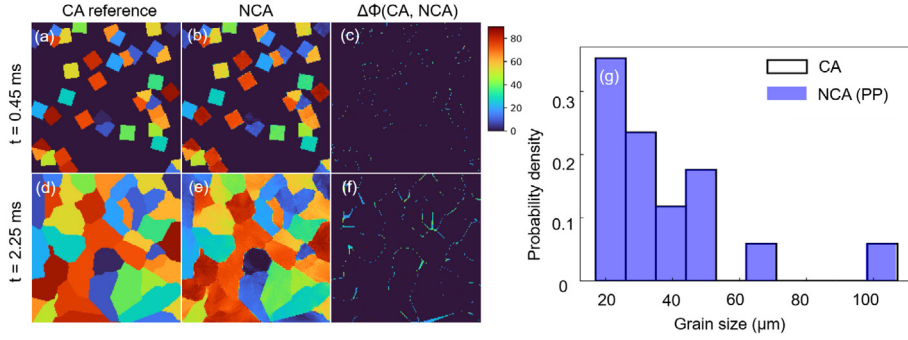
Fig. 4. Test case for 2D isothermal multi-grain growth in a domain of $160\times160\ \mu m^2$ with nucleation density of 1562 mm^{-2}: CA results (a and d), NCA results (b and e), and orientation difference $\Delta\phi$ between them (c and f), grain size distributions derived from CA and NCA results (analysed by the MATLAB toolbox MTEX [53]) (g).
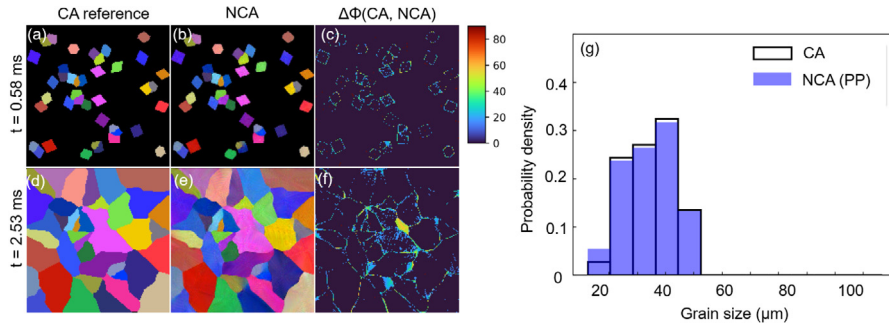
Fig. 5. Test case for conditions with three evolving Euler angles in a domain of $160\times160\ \mu m^2$ with 40 nuclei: CA results (a and d), NCA results (b and e), and orientation difference $\Delta\phi$ between them (c and f), grain size distributions derived from CA and NCA results (analysed by the MATLAB toolbox MTEX [53]) (g). The three Euler angles of the nuclei are randomly selected from [0°, 90°]. The RGB values of the pixel colour in (a), (b), (d), and (e) are the Euler angles normalized by 90°. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4.3. Isothermal solidification: towards quasi-3D multi-grain growth

Without changing the NCA architecture, we now examine its applicability for a scenario similar to quasi-3D grain growth simulation. A quasi-3D simulation, where the 3D domain is replaced by three perpendicular 2D cross-sections [54,55], serves as a computationally cheaper alternative for a 3D microstructure simulation. Such simulations require coupling between the simulations in the three perpendicular 2D cross-sections and allow for the evolution of all three Euler angles. Here, as a preliminary study, we allow for the evolution of all three Euler angles in the 2D simulations.

A new set of CA simulations including 1658 2D runs with three evolving Euler angles for isothermal solidification with undercooling of 20 K are conducted to generate the training and validation datasets for NCA. Comparison of CA and NCA for such simulations indicates training, validation and test accuracies of 93.7%, 87.9% and $89.7 \pm 1.8\%$, respectively. Fig. 5 presents CA and NCA outcomes of an examined test case, showing that the trained model is capable of simulating 2D multi-grain growth with three evolving Euler angles. Compared with the previous cases in Section 4.2, a higher level of inconsistency between CA and NCA is observed, particularly at the grain boundaries. This indicates that a more complex NCA architecture may ultimately be required for quasi-3D microstructure simulations with NCA; however, this development is outside the scope of the present study.

4.4. Non-isothermal solidification: 2D multi-grain growth

For adopting the proposed simulation strategy for practical applications, we need to demonstrate its accuracy in predicting the solidification microstructure under non-isothermal conditions. In this section, the NCA model is
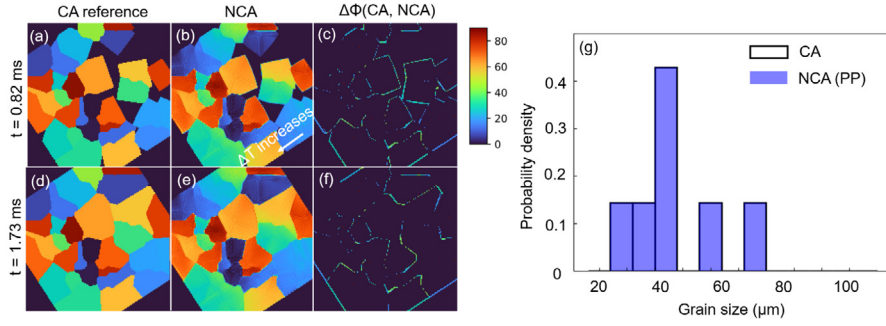
Fig. 6. Test case for 2D multi-grain growth with a random temperature gradient in a domain of 160×160 μm^2 with nucleation density of 1562 mm^{-2}: CA results (a and d), NCA results (b and e), and orientation difference $\Delta\phi$ between them (c and f), grain size distributions derived from CA and NCA results (analysed by the MATLAB toolbox MTEX [53]) (g). A temperature gradient with maximum and minimum undercooling of 25 and 20 K is applied as shown by the arrow in (b).
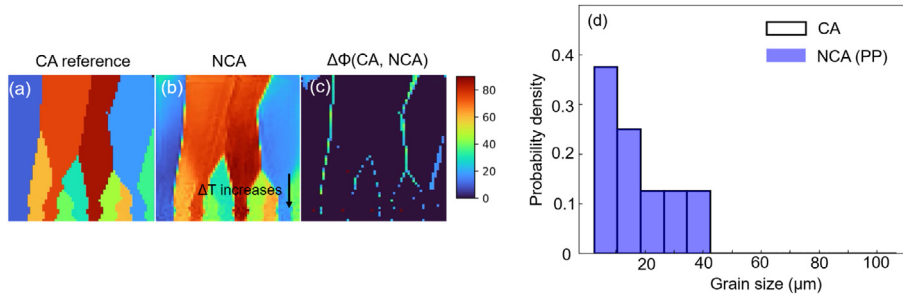
Fig. 7. Test case for 2D multi-grain growth under a vertical temperature gradient with maximum and minimum undercooling of 30 and 20 K (60×60 μm^2 domain with nucleation density of 2778 mm^{-2}): CA results (a), NCA results (b), and orientation difference $\Delta\phi$ between them (c), grain size distributions derived from CA and NCA results (analysed by the MATLAB toolbox MTEX [53]) (d). Nuclei are set 5 μm above the bottom boundary of the domain.

trained based on CA simulations for 2D steady non-isothermal solidification. Accordingly, training, validation and test accuracies of 95.8%, 95.4% and 96.2 $\pm$ 1.1% are obtained, respectively. Fig. 6 shows that the NCA model predicts slower and faster growth rates for the areas with lower and higher levels of undercooling, respectively. Furthermore and as shown in Fig. 7, NCA can also well predict the competitive growth mechanism, i.e. the grains which have one of their preferred growth directions aligned with the maximum temperature gradient grow faster than the others.

Without further training, the NCA model is finally used for predicting solidification under a transient non-isothermal temperature field, i.e. under continuous cooling. As shown in Fig. 8, although no cases with continuous cooling are included in the training, the model accurately predicts the growth kinetics and the final microstructure (test accuracy of 96.9 $\pm$ 0.8%).

In summary, the results presented in this section demonstrate the accuracy of the NCA microstructure modelling strategy and its ability to learn critical physical mechanisms of solidification such as preferred growth direction and competitive growth under various solidification conditions.

5. NCA computational speed

This section compares the computational speed of NCA and CA for a set of benchmark cases. As noted in the introduction, of the two most common physics-based microstructural modelling approaches, CA has a lower computational cost than PFM. Therefore, discussing the computational speed of NCA in comparison with CA serves as a conservative estimate for the speed increase achieved by NCA over the conventional microstructure simulation strategies. The trained NCA model in Section 4.2 is used here, and its computational speed is evaluated for simulating eight different domain sizes between 40^2 and 880^2 μm^2 with a nuclei density of 625 mm^{-2}. A
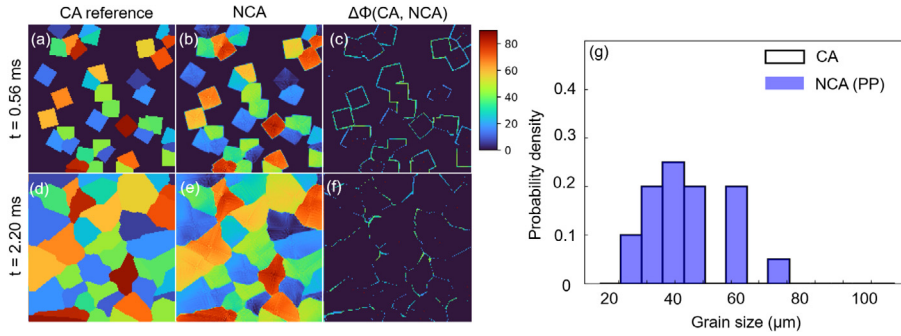
Fig. 8. Test case for 2D multi-grain growth under continuous cooling with an initial undercooling of 20 K and cooling rate of 3.125×10^3 K/s in a domain of $160 \times 160 \ \mu m^2$ with nucleation density of 1562 mm^{-2}: CA results (a and d), NCA results (b and e), and orientation difference $\Delta \phi$ between them (c and f), grain size distributions derived from CA and NCA results (analysed by the MATLAB toolbox MTEX [53]) (g). Note that the NCA did not see any data for continuous cooling during training.

constant time increment of 8 μs is adopted for both CA and NCA for the primary assessments. The measured CA and NCA runtimes on Intel(R) Xeon(R) Gold 6248R CPU for various simulation domain sizes are shown as the red and yellow curves in Fig. 9. An acceleration factor of 6 to 37 is observed for NCA, depending on the domain size.

While most CA codes for microstructure modelling are only compatible with CPUs [15,31,32], NCA can easily exploit the computational power of GPUs (without extra coding). As seen in Fig. 9, GPU computation further accelerates NCA and makes them four orders of magnitude faster than CA. Notably, the runtime of NCA on GPU is not sensitive to the simulation domain size, at least within the examined range. Therefore, even higher acceleration factors can be expected for larger domain sizes.

Adopting a larger time increment of the simulation reduces the computational cost for both NCA and CA. To evaluate how much this deteriorates accuracy, NCA is trained for a time increment increased by a factor of K in the range of 1–50, based on a subset (one frame every K) of the CA simulation database described in Section 4.2. Furthermore, the time increment for CA is similarly increased K times and the corresponding results, along with the results of NCA, are compared with the reference CA results (those for K = 1). Fig. 10 shows the accuracy of both CA and NCA for different acceleration factors in comparison with the results of the reference CA simulation (with K = 1). Expectedly, the accuracy of CA significantly decreases with the increase of K. Notably, NCA are found to be very accurate up to the acceleration factor of 20 since the multi-layer CNN implicitly considers a larger neighbourhood and utilizes more complex functions for grain growth prediction than CA.

The accuracy of the NCA model drops for larger K values (e.g. K = 50). However, as shown in Fig. 10, employing a more complex CNN architecture (e.g. with a 9×9 convolution kernel for the input layer or with 9 hidden convolution layers) leads to an excellent accuracy even for K = 50. The ability of NCA to accurately reproduce the microstructure evolution using significantly larger time increments than CA originates from adopting a multi-layer convolution operator in NCA, which delivers the change in the pixel states based on information from a larger neighbourhood. These results prove that NCA is able to release the requirement for small time increments which limits the computational speed of conventional microstructure modelling techniques; hence, NCA can significantly outperform the computational speed of such simulations.

To summarize, NCA exhibits up to six-orders-of-magnitude improvement in computational efficiency compared to CA (see Fig. 9). The fast computational speed of NCA comes from three contributions. Firstly, NCA replaces the 'for' loop in conventional CA with efficient convolutions, leading to over 10 times speed-up without GPUs. Secondly, the high efficiency of GPUs for parallel computation further accelerates the NCA by around three orders of magnitude. Finally, NCA can benefit from the possibility of adopting larger time increments, up to 50 times larger than the conventional CA in the shown cases.

6. Conclusions

In this work, we have presented NCA, a deep-learning-based approach for modelling solidification microstructures. NCA leverages a multi-layer CNN to replace the simple rules of CA, providing higher computational efficiency
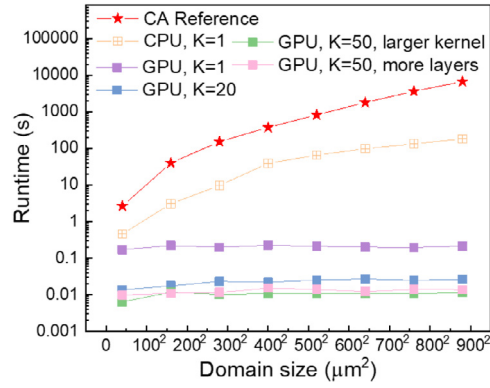
Fig. 9. Runtime of CA and NCA models for simulating microstructures with eight different domain sizes (initial nucleation density of 625 mm^{-2}). The time for I/O processing is not included in the runtime. The maximum error for all the presented NCA simulations remains below 8%. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)
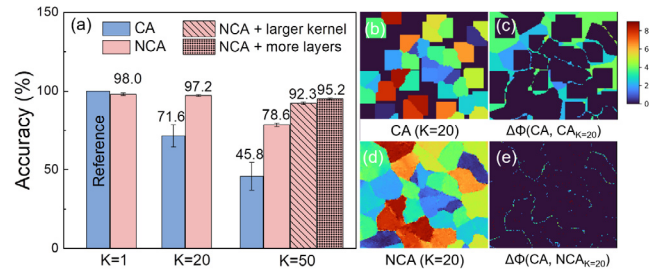
Fig. 10. (a) Accuracy of CA and NCA results with different acceleration factors of K = 1, 20, 50; microstructure of 2D multi-grain growth with K = 20 from CA (b) and NCA (c); orientation difference $\Delta\phi$ of CA (c) and NCA results (e) with K = 20 with respect to CA reference (K = 1).

and flexibility. The convolution operator in NCA is particularly relevant to the physics of solidification, which is a time-progressive phenomenon governed by temperature and microstructure state in the vicinity of the solidification interface. Importantly, NCA allows for encoding physical knowledge into the framework by adopting a 'physics-based activation function', ensuring the satisfaction of solidification physics while easing the training process and reducing the required data size.

We have demonstrated that NCA effectively learns from a relatively small number of CA simulations, since each CA simulation generates a very large dataset by exploring the spatiotemporal evolution of microstructure states for each cell. Interestingly, the trained NCA is independent of boundaries/initial conditions and can accurately predict solidification microstructures even for unseen scenarios during training. This is evident from the high test accuracy when testing the models for larger domain sizes, longer solidification durations, various initial nuclei settings and diverse temperature fields. Moreover, NCA is proven to successfully reproduce the solidification with continuous nucleation (see Appendix C), where nuclei are activated at various time steps, while all the nuclei in the training data are activated at the initial step.

NCA is shown to maintain a high level of accuracy and is up to six orders of magnitude faster than conventional CA, thanks to the efficiency of the adopted computational algorithm, compatibility with GPU computing, and ability to adopt larger time increments for simulations.

Several future extensions of the NCA framework are worth exploring. Firstly, NCA can be trained by data from other high-fidelity microstructure simulations, such as those based on phase-field models, to tackle more complex problems, e.g. the solute atom distribution during phase transition. Training such an NCA model can be achieved by introducing solute concentration as additional input and output channels and implementing a physics-informed loss constraint based on solute atom mass conservation.

Secondly, optimizing the training data and CNN architecture can further improve the predictability and efficiency of NCA. For instance, incorporating larger domain sizes in the training data or using a CNN with larger

neighbourhood sizes can capture grain growth under larger undercoolings or time increments, as demonstrated in Appendix D and Section 5. Adopting a 3D-CNN in NCA can extend the current model into 3D microstructure simulations. Finally, further computational acceleration might be achieved by replacing the CNN with a Fourier Neural Operator (FNO) [56], which employs the Fast Fourier Transform to compute convolutions more efficiently. In summary, we believe NCA bears great potential as a universal method for fast and accurate microstructure simulation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The models and scripts used in this study are available on Github (https://github.com/HighTempIntegrity/ JianTang-NCA01).

Acknowledgements

Appendices

The NCA simulation strategy is developed to capture the physics of the solidification microstructure formation, in which a CNN is trained by the outcomes of CA simulations. The details of the CA method and the generated data for training and validation of NCA are presented in Appendix A. Appendix B explains the derivation of the dendrite growth rate of Hastelloy X for consideration in the CA simulations. Appendix C and Appendix D demonstrate the applicability of NCA for simulating consecutive nuclei activation and solidification with larger undercooling ranges. Finally, Appendix E elaborates a post-processing method for NCA and its effect on the predicted microstructure.

Appendix A. CA simulations for training & validation data

CA simulations are used to generate data for training, validation and ultimately testing NCA. Algorithm 1 describes the pseudocode for the performed CA simulation in this study. As mentioned in Section 2 of the main text, the CA domain is discretized into cells, each containing information such as cell state P, Euler angles θ, temperature T, etc. The incremental evolution of the information of the cell at every increment δt is governed by the information of the cells within a 3×3 neighbourhood in the previous increment.

At the beginning of the simulation, the nuclei at locations nuc_pos are activated in a $nx \times ny$ liquid domain by assigning them the cell state of 'growing' and dedicated Euler angles nuc_ea [31]. Based on the 'decentred octahedron method' developed by Gandin and Rappaz [15], a growing octahedron with the size of 0 is considered at the centre of the nuclei cells whose diagonals make angles as nuc_ea with the reference coordinate system. The octahedron centre location Oct and size λ are tracked during the CA simulations. The incremental growth of the nuclei is simulated through expanding diagonals of the octahedra at a rate given by the 'dendrite growth velocity'. The dendrite growth velocity is calculated based on the cell undercooling (i.e. $\Delta T = T_{melting} - T$) at each increment and as described in Appendix B.

Fig. 11 illustrates a schematic of the decentred octahedron method for representing the movement of the solidification interface. For each increment, it is mathematically checked if the centre of the liquid cells in the neighbourhood of each growing cell falls inside the octahedron and if so, the state of the liquid cell is changed to growing, and a new octahedron is assigned for it. The new octahedron shares one corner with the previous octahedron, hence has the same spatial orientation (i.e. the Euler angles of the new growing cell are inherited from the parent cell). The size of the new octahedron is derived based on the cell length l and the distance a between the cell centre and the corner of the octahedra, see Fig. 11. The new octahedron then grows independently of the parent and based on its local undercooling.
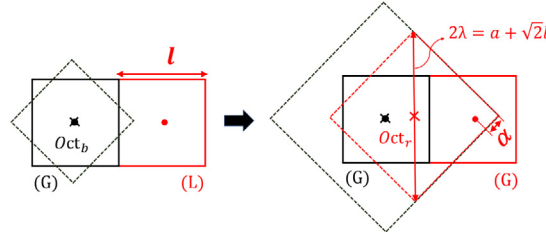
Fig. 11. Schematic of the 'decentred octahedron method' for grain growth in CA simulations.

The transition of the state growing to solid occurs when the growing cell has no neighbour in the liquid state. When there is no liquid cell in the domain, the simulation is finished. Our developed Python CA code can be accessed from Github and readers are referred to Ref. [31] for more details about the CA solidification microstructure modelling.

Various CA simulations were performed for generating training, validation and test data of the NCA. Except for the single-grain growth study (Section 4.1), the training of the NCA uses CA simulation data from both single-nucleus and multi-nuclei growth. It is observed that the accuracy of NCA in predicting multi-grain growth is improved, when the training data includes both single-nucleus and multi-nuclei CA simulations, especially for cases with a time increment factor K>1. For example, when K = 5, the NCA model trained with the dataset in Section 4.2 obtains a better test accuracy of 96.8 ± 0.8% compared with the model trained by the same amount of data with only multi-nuclei CA simulations (test accuracy of 95.3 ± 0.9%). However, the performance of the NCA model trained with only multi-nuclei data can be improved if the training data size increases, e.g. the NCA model trained with 80 multi-nuclei CA simulations obtains a test accuracy of 97.7 ± 0.5% for K = 5.

Appendix B. Dendrite growth velocity

Although there are examples of experimental measurement of dendrite growth rate, e.g. in [57], this is often estimated from empirical models [4–6] or phase-field simulations [58]. One of the prevailing empirical models, the Kurz–Giovanola–Trivedi (KGT) model [4,5] describes the relation between the dendrite tip radius R and its growth velocity v by the following two equations:

$$P_c = \frac{vR}{2D} \tag{B.1}$$

$$G_c(P_c, v)\epsilon_c(P_c) = \frac{4\pi^2 \Gamma}{m_l R^2} + \frac{G}{m_l} \tag{B.2}$$

where G, G_c, P_c, D, m_l, Γ are the mean thermal gradient at the dendrite tip, the concentration gradient in the liquid near the dendrite tip, the Peclet number, the liquid diffusivity, the liquidus slope, and the Gibbs–Thomson coefficient, respectively. The thermal gradient G is neglected in the dendritic growth regime [7]. The concentration gradient G_c and the dimensionless variable ϵ_c are functions of the Peclet number P_c and growth rate v [5,59]:

$$G_c(P_c, v) = \frac{c_0(1 - k)v/D}{I_v(P_c)(1 - k) - 1} \tag{B.3}$$

$$I_v(x) = x \cdot exp(x) \int \frac{exp(-x)}{x} dx \tag{B.4}$$

$$\epsilon_c(P_c) = 1 - \frac{2k}{\left[1 + (2\pi/P_c)^2\right]^{0.5} + 2k - 1} \tag{B.5}$$

where k, c_0, and $I_v(x)$ are the equilibrium partition coefficient, the solute atom concentration, and the Ivantsov function, respectively. Often the constitutional undercooling, resulting from the solute atom segregation, is considered as the main contributor to the supercooling [4–6]. The constitutional undercooling with a given Peclet number P_c is:

$$\Delta T = m_l c_0 \left[1 - \frac{1}{1 - (1 - k)I_v(P_c)}\right] \tag{B.6}$$

13

Algorithm 1 Pseudocode of the CA simulation for generating training, validation, and testing data

---

Require: nx, ny, $T(x, y, t)$, nuc_pos, nuc_ea, δt
Ensure: P, θ
 1: $P \leftarrow -1, \theta \leftarrow (0, 0, 0), t \leftarrow 0$; $\triangleright$ initialize a nx×ny liquid domain
 2: $P[nuc\_pos] \leftarrow 0, \theta[nuc\_pos] \leftarrow nuc\_ea$; $\triangleright$ activate the nuclei at the initial step
 3: Half length of octahedron diagonal $\lambda[nuc\_pos] \leftarrow 0$; $\triangleright$ initialize octahedron
 4: Octahedron centre $Oct[nuc\_pos] \leftarrow Pos[nuc\_pos]$;
 5: Count the number of liquid cells ($P = -1$) Liq_num;
 6: **while** $Liq\_num \neq 0$ **do**
 7:      find the growing cell C_G;
 8:      update octahedron size $\lambda[C_G] \leftarrow \lambda[C_G] + v(T[C_G] - T_m) \times \delta t$;
 9:      ------------------ Grain growth ------------------
10:     find the liquid cells with neighbour growing cells C_{Ln};
11:     **for** each C_{Ln} **do**
12:         $d_{small} = 0$;
13:         **for** each of their growing neighbour cells C_{Gn} **do**
14:             $Vec = Pos[C_{Ln}] - Oct[C_{Gn}]$;
15:             $Vec' = Rot(\theta[C_{Gn}])^{-1} \times Vec$;
16:             $d_i = \lambda[C_{Gn}] - sum(|Vec'|)$;
17:             **if** $d_i \leq d_{small}$ **then**
18:                 $d_{small} \leftarrow d_i$;
19:                 $C_{cap} \leftarrow C_{Gn}$;
20:             **end if**
21:         **end for**
22:         **if** $d_{small} < 0$ **then**
23:             $P[C_{Ln}] \leftarrow 0, \theta[C_{Ln}] \leftarrow \theta[C_{cap}]$;
24:             assign a new octahedron for C_{Ln};
25:         **end if**
26:     **end for**
27:     --
28:     **for** each C_G **do**
29:         **if** no liquid cell in its neighbourhood **then**;
30:             $P[C_G] \leftarrow 1$;
31:         **end if**
32:     **end for**
33:     $t \leftarrow t + \delta t$
34:     Count the number of liquid cells Liq_num;
35: **end while**

The above set of equations can be solved for a range of $v \cdot R$ values and accordingly the dependence of the dendrite growth rate v on the undercooling ΔT can be derived. For the sake of simplicity, a polynomial is ultimately fitted to describe this dependence, i.e. $v(\Delta T)$.

For the application of the KGT model to multi-component systems, the diffusion field of each solute species should be superimposed, and hence, the undercooling ΔT and dendrite tip radius R are [4,5,60]:
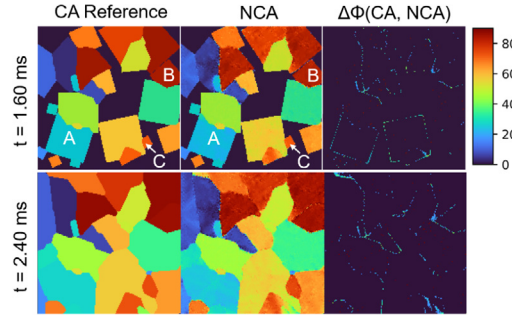
$$\Delta T = \sum_{i}^{N} \Delta T_i \tag{B.7}$$

$$R = \left[\frac{4\pi^2 \Gamma}{\sum_{i}^{N} m_l^i G_c^i(P_c^i, v)\epsilon_c^i(P_c^i) - G} \right]^{0.5} \tag{B.8}$$

Table 4
Material properties of Hastelloy X alloy for deriving its dendrite growth velocity [61–64].

System	Ni–Cr	Ni–Fe	Ni–Mo
Liquidus slope m_l (K wt%$^{-1}$)	−3.166	−1.261	−2.628
Solute composition c_0 (wt%)	20	20	10
Equilibrium partition coefficient k	0.594	0.986	0.99
Liquid diffusivity D ($10^{-9} \cdot$ m^{2}s^{-1})	3.04	3.20	8.861
Gibbs–Thomson coefficient Γ (K m)		3×10^{-7}	

Fig. 12. Test case of 2D multi-grain growth with consecutive nucleation (160×160 μm^2 domain with nucleation density of 1562 mm^{-2}). Grains A, B and C are activated consecutively during the solidification process.
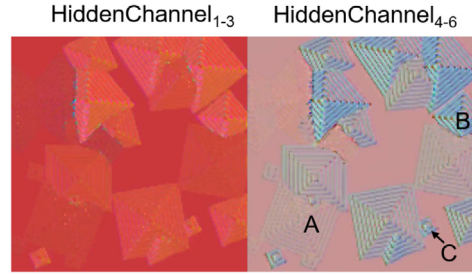
Fig. 13. Hidden channels of the test case in Fig. 12. The RGB values of the pixel colours are the values in the hidden channels normalized by their maximum and minimum values, i.e. $H_{normalized}^i = (H^i - H_{min}^i)/(H_{max}^i - H_{min}^i)$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$$P_c^i = \frac{vR}{2D_i} \tag{B.9}$$

We assume the Hastelloy X as a Ni-20Cr-20Fe-10Mo alloy and use the parameters summarized in Table 4 to derive Eq. (1) (in the main text) for its dendrite growth rate.

Appendix C. NCA for consecutive nuclei activation

The results reported in the main text discuss the application of the NCA for conditions in which all nuclei are activated at the start of the simulation. However, solidification microstructure modelling for processes such as MAM requires consideration of continuous nucleation. We examined the ability of the trained NCA in Section 4.2 to predict microstructure development for conditions where nuclei are activated consecutively during the solidification process. Fig. 12 shows the NCA model predictions in comparison with CA results, showing that the trained NCA can well represent the microstructure evolution with consecutive nucleation.

Fig. 13 represents the values of hidden channels 1–3 and 4–6 for the simulation described above. The RGB values of the pixel colours are the values in the hidden channels normalized by their maximum and minimum, i.e.
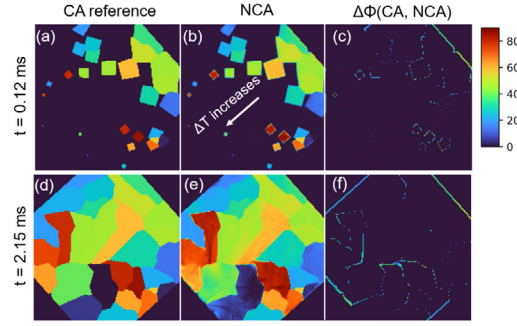
Fig. 14. Test case for 2D multi-grain growth with a random temperature gradient (160×160 μm^2 domain with nucleation density of 1562 mm^{-2}): CA results (a and d), NCA predictions (b and e), and orientation difference $\Delta\phi$ between them (c and f). A temperature gradient with maximum and minimum undercooling of 15 and 45 K is applied as shown in (b).
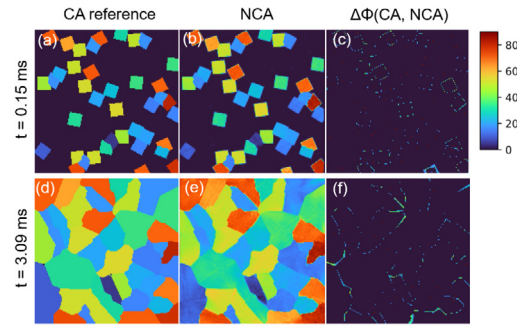
Fig. 15. Test case of 2D multi-grain growth under continuous cooling with an initial undercooling of 15 K and cooling rate of 1.5×10^4 K/s (160×160 μm^2 domain with nucleation density of 1562 mm^{-2}): CA results (a and d), NCA results (b and e), and orientation difference $\Delta\phi$ between them (c and f). Note that the NCA did not see any data for continuous cooling during training.

$H^i_{normalized} = (H^i - H^i_{min})/(H^i_{max} - H^i_{min})$. Interestingly, an octahedron pattern is visible for each grain, where the size and orientation of the diagonals are related to the grain size and orientation, i.e. similar to the concept of decentred octahedron discussed in Appendix A.

Appendix D. NCA for larger undercoolings

The main text describes training and validation of NCA for up to 25 K undercooling. Application of the NCA for simulating fast solidification during MAM requires consideration of larger undercoolings. This section examines the effectiveness of NCA for solidification microstructure simulation with undercooling up to 45 K. This involves consideration of CA simulations from a larger domain size of 128×128 μm^2 for training NCA. The CA simulates solidification under a temperature gradient with undercooling in the range of 15 and 45 K and for nuclei density of 1831–3662 mm^{-2} to generate data for training and validation of NCA. NCA are trained with a fast decaying learning rate (decaying 70% every 100 steps) for 250 epochs. The model obtains accuracies of 97.7%, 97.4%, and 96.2% for training, validation and testing, respectively, see Fig. 14. The trained model can also simulate solidification under continuous cooling for the mentioned undercooling range, as presented in Fig. 15.

Appendix E. Post-processing of NCA results

Here describes a post-processing technique to remove the artefacts in the NCA predictions, e.g. the fluctuations of Euler angles within the grains (see Section 4.2). As the orientation of grains inherits from its nuclei, the main idea of this post-processing method is to replace the pixel Euler angles with that from one of the nearby nuclei that has the most similar orientations. For a pixel located at (x_i, y_i) with Euler angles $\boldsymbol{\theta}_i$ and a nucleus n_j in its vicinity,
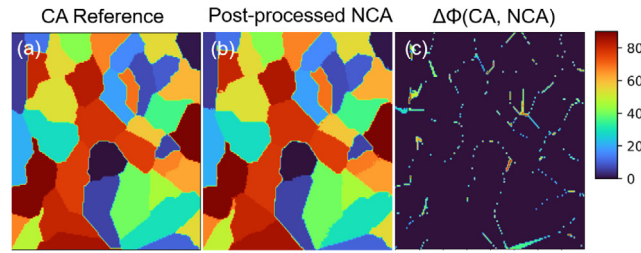
Fig. 16. Comparison of CA and post-processed NCA results from Fig. 4: (a) the CA result, (b) the NCA prediction after post-processing, (c) the orientation difference between them.

we define D_{ij} as:

$$D_{ij} = \frac{\|\boldsymbol{\theta}_i - \boldsymbol{\theta}_{n_j}\|^2}{\max_{1 \leq n_k \leq n_{tot}}\{\|\boldsymbol{\theta}_i - \boldsymbol{\theta}_{n_k}\|^2\}} + \lambda \frac{(x_i - x_{n_j})^2 + (y_i - j_{n_j})^2}{\max_{1 \leq nl \leq n_{tot}}\{(x_i - x_{n_l})^2 + (y_i - j_{n_l})^2\}} \tag{E.1}$$

D_{ij} is an index representing the differences of Euler angles and location between pixels and the nucleus. In the post-processing of NCA results, we replace the pixel's Euler angles with that of the nuclei with the smallest $D_{i,j}^{ni}$. Fig. 16 shows a comparison between the CA and the post-processed NCA results from Fig. 4. Most of the artefacts are removed after post-processing with only a few inconsistencies at the grain boundaries compared with CA results.

References

[1] Dongdong Gu, Xinyu Shi, Reinhart Poprawe, David L. Bourell, Rossitza Setchi, Jihong Zhu, Material-structure-performance integrated laser-metal additive manufacturing, Science 372 (6545) (2021) eabg1487.

[2] Tianlong Zhang, Zhenghua Huang, Tao Yang, Haojie Kong, Junhua Luan, Anding Wang, Dong Wang, Way Kuo, Yunzhi Wang, Chain-Tsuan Liu, In situ design of advanced titanium alloy with concentration modulations by additive manufacturing, Science 374 (6566) (2021) 478–482.

[3] Sindo Kou, Welding Metallurgy, Vol. 431, No. 446, New Jersey, USA, 2003, pp. 223–225.

[4] W. Kurz, D.J. Fisher, Fundamental of Solidification, Trans Tech Publications, VT, USA, 1989, pp. 45–89.

[5] W.1 Kurz, B. Giovanola, R. Trivedi, Theory of microstructural development during rapid solidification, Acta Metall. 34 (5) (1986) 823–830.

[6] J. Lipton, M.E. Glicksman, W. Kurz, Dendritic growth into undercooled alloy metals, Mater. Sci. Eng. 65 (1) (1984) 57–63.

[7] Michel Rappaz, Ch-A. Gandin, Probabilistic modelling of microstructure formation in solidification processes, Acta Metall. Mater. 41 (2) (1993) 345–360.

[8] Enzo Liotti, Carlos Arteta, Andrew Zisserman, Andrew Lui, Victor Lempitsky, Patrick S. Grant, Crystal nucleation in metallic alloys using x-ray radiography and machine learning, Sci. Adv. 4 (4) (2018) eaar4004.

[9] Amrita Basak, Suman Das, Epitaxy and microstructure evolution in metal additive manufacturing, Annu. Rev. Mater. Res. 46 (1) (2016) 125–149.

[10] Theron M. Rodgers, Jonathan D. Madison, Veena Tikare, Simulation of metal additive manufacturing microstructures using kinetic Monte Carlo, Comput. Mater. Sci. 135 (2017) 78–89.

[11] T. Pinomaa, M. Lindroos, M. Walbruhl, N. Provatas, A. Laukkanen, The significance of spatial length scales and solute segregation in strengthening rapid solidification microstructures of 316l stainless steel, Acta Mater. 184 (2020) 1–16.

[12] D. Tourret, Y. Song, A.J. Clarke, A. Karma, Grain growth competition during thin-sample directional solidification of dendritic microstructures: A phase-field study, Acta Mater. 122 (2017) 220–235.

[13] Min Yang, Lu Wang, Wentao Yan, Phase-field modeling of grain evolutions in additive manufacturing from nucleation, growth, to coarsening, npj Comput. Mater. 7 (1) (2021) 56.

[14] Tianju Xue, Zhengtao Gan, Shuheng Liao, Jian Cao, Physics-embedded graph network for accelerating phase-field simulation of microstructure evolution in additive manufacturing, npj Comput. Mater. 8 (1) (2022) 1–13.

[15] Ch-A. Gandin, Michel Rappaz, A 3D cellular automaton algorithm for the prediction of dendritic grain growth, Acta Mater. 45 (5) (1997) 2187–2195.

[16] Ch Gandin, J.-L. Desbiolles, Michel Rappaz, Ph Thevoz, et al., A three-dimensional cellular automation-finite element model for the prediction of solidification grain structures, Metall. Mater. Trans. A 30 (12) (1999) 3153–3165.

[17] Cheng Gu, Michael P. Moodispaw, Alan A. Luo, Cellular automaton simulation and experimental validation of eutectic transformation during solidification of Al-Si alloys, npj Comput. Mater. 8 (1) (2022) 134.

[18] Julia Kundin, Leslie Mushongera, Heike Emmerich, Phase-field modeling of microstructure formation during rapid solidification in Inconel 718 superalloy, Acta Mater. 95 (2015) 343–356.

[19] Nele Moelans, Bart Blanpain, Patrick Wollants, Quantitative analysis of grain boundary properties in a generalized phase field model for grain growth in anisotropic systems, Phys. Rev. B 78 (2) (2008) 024113.

[20] Alain Karma, Wouter-Jan Rappel, Quantitative phase-field modeling of dendritic growth in two and three dimensions, Phys. Rev. E 57 (4) (1998) 4323.

[21] A. Kazaryan, Y. Wang, S.A. Dregia, Bruce R. Patton, Grain growth in systems with anisotropic boundary mobility: Analytical model and computer simulation, Phys. Rev. B 63 (18) (2001) 184102.

[22] John W. Cahn, John E. Hilliard, Free energy of a nonuniform system. I. Interfacial free energy, J. Chem. Phys. 28 (2) (1958) 258–267.

[23] John W. Cahn, John E. Hilliard, Free energy of a nonuniform system. III. Nucleation in a two-component incompressible fluid, J. Chem. Phys. 31 (3) (1959) 688–699.

[24] Samuel M. Allen, John W. Cahn, A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening, Acta Metall. 27 (6) (1979) 1085–1095.

[25] J.W. Cahn, S.M. Allen, A microscopic theory for domain wall motion and its experimental verification in fe-al alloy domain growth kinetics, Le J. de Phys. Colloques 38 (C7) (1977) C7–51.

[26] Tatu Pinomaa, Nikolas Provatas, Quantitative phase field modeling of solute trapping and continuous growth kinetics in quasi-rapid solidification, Acta Mater. 168 (2019) 167–177.

[27] P.W. Liu, Y.Z. Ji, Z. Wang, C.L. Qiu, A.A. Antonysamy, L.-Q. Chen, X.Y. Cui, L. Chen, Investigation on evolution mechanisms of site-specific grain structures during metal additive manufacturing, J. Mater Process. Technol. 257 (2018) 191–202.

[28] P.W. Liu, Z. Wang, Y.H. Xiao, Ricardo A. Lebensohn, Y.C. Liu, Mark F. Horstemeyer, X.Y. Cui, L. Chen, Integration of phase-field model and crystal plasticity for the prediction of process-structure-property relation of additively manufactured metallic materials, Int. J. Plast. 128 (2020) 102670.

[29] C. Yang, Q.Y. Xu, X.L. Su, B.C. Liu, Multiphase-field and experimental study of solidification behavior in a nickel-based single crystal superalloy, Acta Mater. 175 (2019) 286–296.

[30] Daniel Dreelan, Alojz Ivankovic, David J. Browne, Verification of a new cellular automata model of solidification using a case study on the columnar to equiaxed transition previously simulated using front tracking, Comput. Mater. Sci. 215 (2022) 111773.

[31] Yanping Lian, Zhengtao Gan, Cheng Yu, Dmitriy Kats, Wing Kam Liu, Gregory J. Wagner, A cellular automaton finite volume method for microstructure evolution during additive manufacturing, Mater. Des. 169 (2019) 107672.

[32] Johannes A. Koepf, Martin R. Gotterbarm, Matthias Markl, Carolin Körner, 3D multi-layer grain structure simulation of powder bed fusion additive manufacturing, Acta Mater. 152 (2018) 119–126.

[33] Amir Reza Ansari Dezfoli, Yu-Lung Lo, M. Mohsin Raza, 3D multi-track and multi-layer epitaxy grain growth simulations of selective laser melting, Materials 14 (23) (2021).

[34] Mohammad Sadegh Mohebbi, Vasily Ploshikhin, Implementation of nucleation in cellular automaton simulation of microstructural evolution during additive manufacturing of Al alloys, Addit. Manuf. 36 (2020) 101726.

[35] W. Wang, P.D. Lee, M. McLean, A model of solidification microstructures in nickel-based superalloys: predicting primary dendrite spacing selection, Acta Mater. 51 (10) (2003) 2971–2987.

[36] Matthew R. Rolchigo, Michael Y. Mendoza, Peyman Samimi, David A. Brice, Brian Martin, Peter C. Collins, Richard Lesar, Modeling of Ti-W solidification microstructures under additive manufacturing conditions, Metall. Mater. Trans. A 48A (7) (2017) 3606–3622.

[37] Anup Pandey, Reeju Pokharel, Machine learning based surrogate modeling approach for mapping crystal deformation in three dimensions, Scr. Mater. 193 (2021) 1–5.

[38] Peyman Saidi, Hadi Pirgazi, Mehdi Sanjari, Saeed Tamimi, Mohsen Mohammadi, Laurent K. Béland, Mark R. Daymond, Isaac Tamblyn, Deep learning and crystal plasticity: A preconditioning approach for accurate orientation evolution prediction, Comput. Methods Appl. Mech. Engrg. 389 (2022) 114392.

[39] Vivek Oommen, Khemraj Shukla, Somdatta Goswami, Rémi Dingreville, George Em Karniadakis, Learning two-phase microstructure evolution using neural operators and autoencoder architectures, npj Comput. Mater. 8 (1) (2022) 190.

[40] David Montes de Oca Zapiain, James A. Stewart, Rémi Dingreville, Accelerating phase-field-based microstructure evolution predictions via surrogate models trained by machine learning methods, npj Comput. Mater. 7 (1) (2021) 3.

[41] Sepideh Hashemi, Surya R. Kalidindi, A machine learning framework for the temporal evolution of microstructure during static recrystallization of polycrystalline materials simulated by cellular automaton, Comput. Mater. Sci. 188 (2021) 110132.

[42] C. Hu, S. Martin, R. Dingreville, Accelerating phase-field predictions via recurrent neural networks learning the microstructure evolution in latent space, Comput. Methods Appl. Mech. Engrg. 397 (2022) 115128.

[43] Yang Jiao, F.H. Stillinger, S. Torquato, Modeling heterogeneous materials via two-point correlation functions: Basic principles, Phys. Rev. E 76 (3) (2007) 031110.

[44] Yang Jiao, F.H. Stillinger, S. Torquato, Modeling heterogeneous materials via two-point correlation functions. II. Algorithmic details and applications, Phys. Rev. E 77 (3) (2008) 031135.

[45] Elizabeth J. Kautz, Predicting material microstructure evolution via data-driven machine learning, Patterns 2 (7) (2021) 100285.

[46] Vivek Oommen, Khemraj Shukla, Somdatta Goswami, Rémi Dingreville, George Em Karniadakis, Learning two-phase microstructure evolution using neural operators and autoencoder architectures, npj Comput. Mater. 8 (1) (2022) 190.

[47] Alexander Mordvintsev, Ettore Randazzo, Eyvind Niklasson, Michael Levin, Growing neural cellular automata, Distill (2020) https://distill.pub/2020/growing-ca.

[48] Eyvind Niklasson, Alexander Mordvintsev, Ettore Randazzo, Michael Levin, Self-organising textures, Distill (2021) https://distill.pub/selforg/2021/textures.

[49] Shinya Hibino, Tsubasa Todo, Takuya Ishimoto, Ozkan Gokcekaya, Yuichiro Koizumi, Kenichiroh Igashira, Takayoshi Nakano, Control of crystallographic texture and mechanical properties of Hastelloy-X via laser powder bed fusion, Crystals 11 (9) (2021) 1064.

[50] Ziliang Huang, Zirong Zhai, Wenhu Lin, Hai Chang, Yingna Wu, Rui Yang, Zhenbo Zhang, On the orientation dependent microstructure and mechanical behavior of Hastelloy X superalloy fabricated by laser powder bed fusion, Mater. Sci. Eng. A 844 (2022) 143208.

[51] Xiaoxuan Zhang, Krishna Garikipati, Machine learning materials physics: Multi-resolution neural networks learn the free energy and nonlinear elastic response of evolving microstructures, Comput. Methods Appl. Mech. Engrg. 372 (2020) 113362.

[52] George Weber, Maxwell Pinz, Somnath Ghosh, Machine learning-enabled self-consistent parametrically-upscaled crystal plasticity model for Ni-based superalloys, Comput. Methods Appl. Mech. Engrg. 402 (2022) 115384.

[53] F. Bachmann, Ralf Hielscher, Helmut Schaeben, Texture analysis with MTEX–free and open source software toolbox, in: Solid State Phenomena, Vol. 160, Trans Tech Publ, 2010, pp. 63–68.

[54] G.K. Upadhya, K.O. Yu, M.A. Layton, A.J. Paul, Probabilistic Modeling of Solidification Grain Structure in Investment Castings, Technical Report, Minerals, Metals and Materials Society, Warrendale, PA (United States), 1995.

[55] Damien Tourret, Younggil Song, Amy Jean Clarke, Alain Karma, Grain growth competition during thin-sample directional solidification of dendritic microstructures: A phase-field study, Acta Mater. 122 (2017) 220–235.

[56] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar, Fourier neural operator for parametric partial differential equations, 2020.

[57] L. Wang, L. Hu, J.F. Zhao, B. Wei, Ultrafast growth kinetics of titanium dendrites investigated by electrostatic levitation experiments and molecular dynamics simulations, Chem. Phys. Lett. 742 (2020) 137141.

[58] Jean Bragard, Alain Karma, Youngyih H. Lee, Mathis Plapp, Linking phase-field and atomistic simulations to model dendritic solidification in highly undercooled melts, Interface Sci. 10 (2) (2002) 121–136.

[59] M. Rappaz, Modelling of microstructure formation in solidification processes, Int. Mater. Rev. 34 (1) (1989) 93–124.

[60] W. Kurz, D.J. Fisher, Dendrite growth at the limit of stability: tip radius and spacing, Acta Metall. 29 (1) (1981) 11–20.

[61] Yōichi Ono, Teinosuke Yagi, Diffusion of chromium, manganese, cobalt, and nickel in molten iron saturated with carbon, Trans. Iron Steel Inst. Jpn. 11 (4) (1971) 275–279.

[62] Yao Wen-Jing, Wang Nan, Monte Carlo simulation of thermophysical properties of Ni-15% Mo alloy melt, Acta Phys. Sin. 58 (6) (2009) 4053–4058.

[63] Teng Fang, Li Wang, Yu Qi, Structural, thermodynamics and dynamics properties of fe-Ni melts with different EAM models, in: Advanced Materials Research, Vol. 750, 750, Trans Tech Publ, 2013, pp. 579–582.

[64] H. Okamoto, M.E. Schlesinger, E.M. Mueller, Ni (Nickel) Binary Alloy Phase Diagrams, ASM International, 1991.