



Advanced atomistic models for radiation damage in Fe-based alloys: Contributions and future perspectives from artificial neural networks

N. Castin^{a,*}, M.I. Pascuet^b, L. Messina^c, C. Domain^d, P. Olsson^c, R.C. Pasianot^{b,e}, L. Malerba^a

^a Studief Centrum voor Kerneenergie – Centre d'Études de l'énergie Nucléaire (SCK•CEN), NMS Unit, Boeretang 200, B2400 Mol, Belgium

^b Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), Godoy Cruz 2290 (C1425FQB) CABA, Buenos Aires, Argentina

^c KTH Royal Institute of Technology, Reactor Physics, Roslagstullsbacken 21, 106 91 Stockholm, Sweden

^d Département Matériaux et Mécanique des Composants, EDF-R&D, Les Renardières, F-77250 Moret sur Loing, France

^e Gcia. Materiales, Comisión Nacional de Energía Atómica (CNEA), Av. Gral. Paz 1499, 1650 San Martín, Argentina

ARTICLE INFO

Article history:

Received 26 October 2017

Received in revised form 16 January 2018

Accepted 9 February 2018

Keywords:

Artificial neural networks

Kinetic Monte Carlo

Irradiation damage

Multiscale modelling

ABSTRACT

Machine learning, and more specifically artificial neural networks (ANN), are powerful and flexible numerical tools that can lead to significant improvements in many materials modelling techniques. This paper provides a review of the efforts made so far to describe the effects of irradiation in Fe-based and W-based alloys, in a multiscale modelling framework. ANN were successfully used as innovative parametrization tools in these models, thereby greatly enhancing their physical accuracy and capability to accomplish increasingly challenging goals. In the provided examples, the main goal of ANN is to predict how the chemical complexity of local atomic configurations, and/or specific strain fields, influence the activation energy of selected thermally-activated events. This is most often a more efficient approach with respect to previous computationally heavy methods. In a future perspective, similar schemes can be potentially used to calculate other quantities than activation energies. They can thus transfer atomic-scale properties to higher-scale simulations, providing a proper bridging across scales, and hence contributing to the achievement of accurate and reliable multiscale models.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

It all starts with an energetic particle, often a neutron, hitting an atom. This statement could be found at the inception of any manuscript, report, or presentation at a conference, undertaking a complete characterization of the effects of irradiation in structural materials. The devil thus hides from the smallest time- and length-scales, at the atomic level, and even below, at the electronic level. Irradiation with bombarding particles (provided they are energetic enough) results in a continuous creation of point-defects in the bulk of the studied material, i.e., vacancies (Vac) and self-interstitials atoms (SIA). The increased concentration of vacancies will then enhance the kinetics of diffusion-driven process (such as precipitation of insoluble species), but might also alter the thermodynamic equilibrium, as a consequence of the establishment of fluxes of them towards sinks. Moreover, SIA are produced. They are essentially absent in unirradiated materials. SIA exhibit a profoundly different behaviour from vacancies, potentially giving rise to processes not normally observed in materials. SIA migrate

generally faster than vacancies and sometimes they follow one-dimensional diffusion paths, exhibiting in addition longer range elastic interactions of a highly anisotropic character. The changes induced in the material by the production and diffusion of these defects, as observable at the macroscopic scale, are thus not only irradiation-enhanced, but often irradiation-induced.

Because it all starts at the atomic level, thorough and non-empirical physical models capable of describing the effects of irradiation up to the macroscopic level must necessarily entail a largely multiscale strategy. Fig. 1 depicts an overview of a modelling approach based on kinetic Monte Carlo (KMC) methods, which is one of the possible choices to bridge the gap from the electronic to coarse-grained level. Other possible approaches are mean field models, rate theory, cluster dynamics, etc. Interested readers are directed to Ref. [1,2] for a general overview. Addressing the macroscopic level requires higher-scale models such as dislocation dynamics and continuum mechanics calculations with adequate plastic flow laws: these are not shown in Fig. 1. In the lowest left corner, two fundamental sources of input data are: (a) Calculations based on first principle physics to address the electronic structure, most often using the density functional theory (DFT); (b) Any kind of experimental evidence, which can be of a wide

* Corresponding author.

E-mail address: nicolas.m.b.castin@gmail.com (N. Castin).

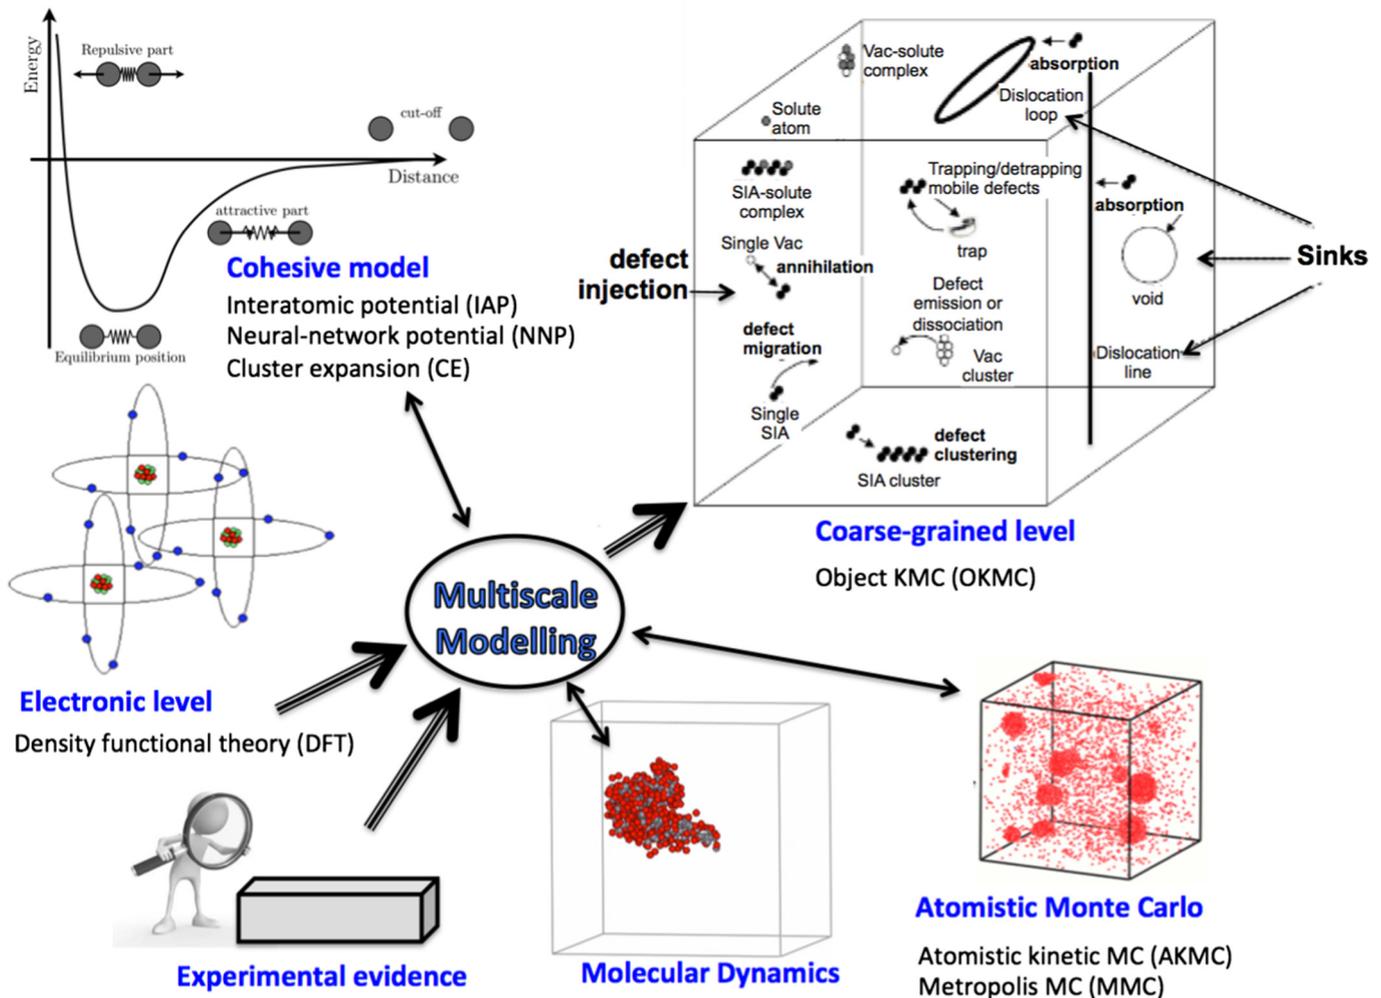


Fig. 1. Multiscale modelling approach for describing the effects of irradiation in structural materials.

variety, ranging from formation energies, diffusion coefficients, elastic constants, phase diagrams, etc. Unquestionably, experimental evidence remains the only true final judge on any model prediction. The story ends at the coarse-grained level, which is in this work implemented by the object kinetic Monte Carlo (OKMC) method. There, the processes undergone by the defects injected during irradiation are described as thermally-activated events that correspond, for the objects involved, generally clusters of point-defects, to specifically designed laws for mutual interactions, mobility and stability. All regular elements of the microstructure are also explicitly present in the model, for example grain boundaries, dislocations, solute or foreign interstitial impurities. A substantial number of reactions must be carefully parametrised in OKMC models, which is achieved summoning a variety of methods and models at lower scales, ranging from DFT, molecular dynamics, and atomistic kinetic Monte Carlo.

Playing the role of a reference metrics for the physical description of the studied material, vital pieces of this puzzle are, undoubtedly, cohesive models. Given a set of atoms described by their coordinates in space and their chemical species, their task is to provide a numerical value for the total energy of the system in any configuration, determining the driving forces for its evolution, including interatomic forces. In a quantum mechanism framework, *Ab initio* methods, e.g., DFT, are, for many good reasons, regarded as the most reliable choice to explore the energy landscape of the system. There are, however, limitations. Firstly, some

degree of empiricism remains. Unfortunately, a blind use of DFT may, in some cases, provide imperfect predictions (see e.g. [3,4]). Secondly, even with state-of-the-art computing facilities, using DFT still implies a huge cost in CPU resources which is often unaffordable and, sometimes, unwarranted. For these reasons, alternative cohesive models are still extensively used, empirical interatomic potentials (IAP) being a very popular example. On the one hand, in addition to the largely reduced computing cost, the main advantage of these potentials is the possibility to achieve a tunable compromise between various target properties. It can mix DFT-originated with experimental-originated data, as conveniently as required. For example, the FeCu potential proposed by Pasianot et al. in Ref. [5] was fitted to faithfully reproduce the experimentally observed Cu solubility limit in Fe, whereas DFT is known to underestimate it (see discussions in Ref. [4]). The affordable computing time allows large and complex systems, containing up to several million atoms, to be studied. On the other hand, the simplicity proper to IAP also implies an intrinsic limitation in their capacities of making accurate prediction of various set of properties at the same time.

Suitable low-computing-cost cohesive models allow the gap between the basic atomic level and the realm of longer time-scale Monte Carlo (MC) models to be bridged. Many more gaps remain to bridge, however. Even using IAP, the precise calculation of the energy barriers the system has to overcome through thermal activation, which are at the core of MC methods, pose severe

computing time limitations. Repetitive and heavy routine calculations are indeed required at key steps of a MC simulation, significantly limiting the range of applicability of the method. Similar limitations arise in connection with storage and integration of data/knowledge moving from one scale to a higher one. An illustrative example is the calculation of migration energies of single point-defects in a chemically changing environment. These need to be routinely calculated using a dedicated method in order to let a Monte Carlo simulation that describes diffusion-driven processes to progress, but they entail a prohibitive computational cost if the calculation is performed in a precise way each time. Bridging the gap between models thus here means finding a numerical solution for either speeding up the calculation of the migration energies, or, from another standpoint, extrapolate knowledge in the MC model from a limited amount of examples.

Machine learning techniques, specifically artificial neural networks (ANN), are very promising tools for taking this challenge. As reviewed in this paper, ANN-based tools have been progressively coupled to atomistic Monte Carlo models to improve their physical reliability, within affordable computing loads [6–14]. ANN are flexible and powerful regression techniques, capable to assimilate complex and inexpressible laws of interactions from a limited number of explicit examples. Once the ANN are properly designed, they provide the values required for the model at a very limited computing cost, thanks to their mathematical simplicity. The paper is organised as follows. We start, in Section 2, by providing a basic description of ANN techniques, and the underlying mathematical framework. Next, in Section 3, we explain how neural-network potentials (NNP) can be fitted directly from DFT. Later, in Section 4, we extensively review our enhanced atomistic kinetic Monte Carlo models, where ANN are used to evaluate migration energies on-the-fly. Finally, we provide our visions and recommendations for future perspectives in Section 5.

2. Artificial neural networks: a practical form of artificial intelligence

Artificial neural networks (ANN) are a concept of *weak artificial intelligence*, in a group of paradigms often denoted to as *machine learning*, or *computational intelligence*. In many applications, they can be seen as self-learning systems aimed at extracting hidden knowledge from their environment, in order to take relevant actions when entirely new situations are encountered. For the particular application to the design of numerical predicting tools, for instance, they are *surrogate models* sparing the user the (otherwise fastidious) need to explicitly formulate knowledge about the problem at hand.

A complete description of the theoretical and mathematical frameworks of ANN goes beyond the scope of this paper. For a complete and detailed reference textbook, interested readers are directed to the reference book by Bishop [15], or to a more concise overview in Ref. [16]. In Section 2.1, we start by providing a description of the specific implementation of ANN that is used in our work, namely, the *multilayer perceptron*. Next, in Section 2.2, we describe the fundamental aspects of training. Last, in Section 2.3 we briefly report a typical example of application of ANN as numerical regression, by predicting the amount of radiation-induced hardening in reactor pressure vessel steels.

2.1. The multi-layer perceptron: a universal approximation machine

Several types of ANN exist. The so-called *multilayer perceptron* depicted in Fig. 2, for instance, provides appropriate solutions for the design of general numerical regressions. Inspired from biological neural networks such as the human brain, the fundamental

idea is to create a network of simple processing units (either called neurons or nodes in the literature), thus constructing a sophisticated and complex response out of a set of simple individual rules. For the sake of simplicity, at least from the mathematical point of view, ANN are typically constituted of organised layers of nodes. On the left-hand side in Fig. 2, the input layer is the collection of the raw input signals for the whole network. On the right-hand side, the last layer is called the output layer, providing the answer of the network. In between, several intermediate layers may be introduced (there is only one in the figure). They are called *hidden layers*, for the reason that they are, in practice, invisible to the user. There is in theory no restriction about how nodes of different layers can be connected to each other. For simplicity, many networks do not allow backwards connections: a given node never receives as input output signals coming from nodes of either the same layer, or from the next layers (see e.g. Ref. [17]). Such a network is qualified as *feedforwards*. Last, the network shown in the figure is said to be *fully connected*, because all nodes in a given layer receive as input all output signals coming from the immediately preceding layer, and no *layer bypass* is allowed. The inputs for the first (and only) hidden layer are the raw input signals of the network.

The output y_j of a hidden node j reads:

$$y_j = \phi(v_j) \quad (1)$$

$$v_j = w_{j0} + \sum_{i=1}^N w_{ji} \cdot x_i \quad (2)$$

Here, v_j is the *internal activity* of the node. Function ϕ can, in principle, be any non-linear function. For convenience, to prevent large magnitudes of the signal, function ϕ is almost always taken to be bounded, e.g., a hyperbolic tangent. The internal activity is calculated as a weighted sum of the input signals to the node. The output O of the only node in the output layer is calculated in a similar way. In Eq. (2), the synaptic weights w_{j0} , w_{ji} monitor the strength of interaction between the nodes, and are therefore called *synaptic weights*, or also *synapses*.

The interest of the feed-forwards multilayer perceptron, in its simplest form as depicted in Fig. 2, is that it fulfils the *universal approximation theorem* [18]: For any continuous function $F(x_1, \dots, x_n)$, there exist an ANN with a finite number of hidden nodes that fulfils the following condition, for all set of inputs x_i and all ϵ :

$$|F(x_1, \dots, x_n) - O(x_1, \dots, x_n)| < \epsilon \quad (3)$$

Here, $O(x_1, \dots, x_n)$ is the output of the ANN. The multilayer perceptron can therefore, in theory, be regarded as a universal approximation machine capable of perfectly assimilating any problem with a numerical character. It is safe and reasonable to assume that such degree of idealness cannot be dreamed of in practice, unless perhaps for some academical examples. Concretely, it is the authors' opinion that the flexibility and generality of the ANN internal structure, avoiding the user to explicitly formulate knowledge about the problem at hand, comes at an unavoidable cost. Indeed, ANN might be very successful in understanding a given problem from a limited amount of examples, but expecting it to perform a truly general understanding is hazardous. In other words, it is safe to assume that ANN are inherently devoted (not to say limited) to interpolation problems, whereas extrapolation outside the domain of the input space covered during training is uncertain. Therefore, during any application of ANN, it is important to make sure that the domain of applicability can be determined. This can in fact be delicate in some cases, as discussed later.

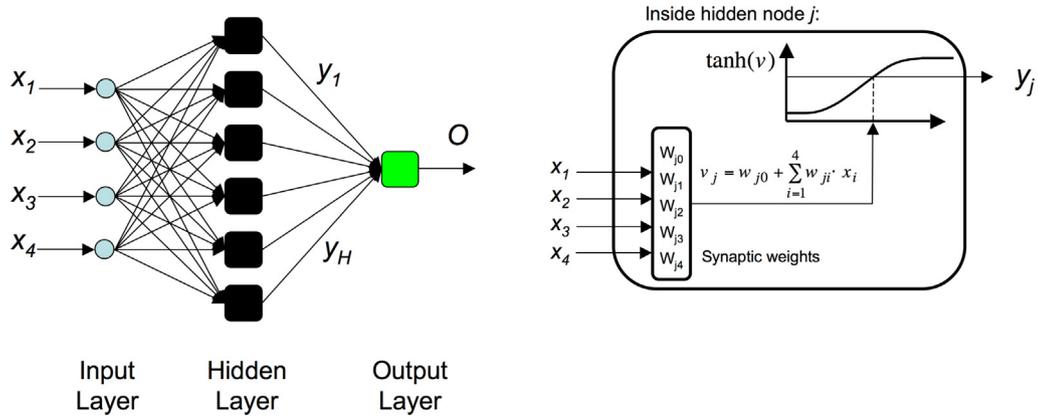


Fig. 2. Example of feed-forwards multilayer perceptron with 4 input variables (x_1, x_2, x_3 and x_4), one hidden layer with $H = 6$ nodes and one output (O). The right part of the figure shows the detail inside a node of the hidden layer. The signal is propagated from the left to the right of the network, in a layer-by-layer fashion.

2.2. Supervised training

In this work, ANN are designed to implement numerical regression tools. This is therefore a *supervised training problem*, and training can be regarded as an optimization problem that consists in the minimization of the following objective function:

$$f(N_H, \bar{w}) = \sum_{i=1}^{N_T} (o_i(N_H, \bar{w}) - d_i)^2 \quad (4)$$

Here N_T is the number of available examples of I/O for training, d_i is the desired output for the example i , and o_i is the corresponding prediction by the network. The latter is a function of the ANN architecture, i.e., the number N_H of nodes in the hidden layer, and, finally, the vector \bar{w} of synaptic weights. Assuming that N_H is fixed, the minimization of function f is thus the problem of determining the optimal numerical value of the synaptic weights \bar{w} : it can be undertaken by any classical method for non-linear optimization. In this work, we used the method proposed by Levenberg [19] and Marquardt [20] (LM).

As anticipated in the previous section, a major concern while training ANN is to guarantee that the predictions for new sets of inputs are equally accurate, compared to predictions on the available set of examples used for training. Indeed, without control, there is a risk that the ANN does not develop a general logic, but in reality rather memorizes the complete set of available examples, as illustrated in Fig. 3. We see that reasonable interpolation is achieved by the ANN if the latter is not too complex, i.e. if H is

small. Predictions are however not equally accurate for all training examples. If the number of hidden nodes is increased, the ANN manages more accurate predictions for all known points, but clearly loses generality. This pathology cannot be easily identified only on the basis of a limited amount of examples of I/O, especially if the dimensionality of the problem at hand is large. ANN training must therefore be *regularized*. The most common regularization approach is called *early stopping*, and is based on the idea that memorization of the provided examples, or more generally can we say *network overspecialisation*, develops only at a certain moment of training, i.e. after a certain number of training iterations, called *epochs* in ANN jargon. The most natural way to prevent it is therefore to divide the available table of examples of I/O into two different and non overlapping sets: (a) The training set is used to minimize function f in Eq. (4). Only these examples are thus used to calculate the gradients with the LM algorithm, and the synapses are updated during each epoch taking only them into account; (b) The reference set is used to measure, after every epoch, the average error of prediction on new cases. Fig. 3 shows the typical evolution of the average error of prediction on both sets during the training epochs. The error committed on the training set always decreases. The error on the reference set, however, ceases to decrease from a certain epoch, and then starts to increase, as a clear sign of the onset of overspecialisation. Training is therefore interrupted at that moment.

At this stage, the determination of the optimal network architecture is still an open question. According to our experience, acquired during the applications summarized in this paper, we invariably found optimal architectures using no more than one hidden layer in the ANN. Determining the optimal architecture is thus a mono-parametric study: networks with increasing N_H are trained separately, and the one committing the lowest error on the reference set is finally retained, as depicted in Fig. 3. Too small a N_H understandably leads to higher errors of prediction, because not enough degrees of freedom are available in the network. Oppositely, too high a N_H increases the risk of overspecialisation, and the error on the reference set increases as well. Other approaches exist for determining the optimal architecture. Interested readers can find an overview in Ref. [21–23]. Specifically for our work in ANN-based KMC models, we proposed a constructive method called GIACA in Ref. [9].

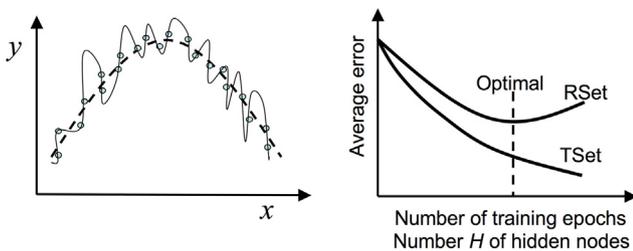


Fig. 3. (left) Schematic illustration of ANN over-specialization for a simple problem with one variable x and one output y . Blue points represent data in the training set. The dashed line shows predictions by an ANN with a small number H of hidden nodes, whereas the plain line shows predictions of another ANN where too many hidden nodes were introduced. (right) Typical evolution of the average error committed on the training set (TSet) and reference set (RSet) with increasing number of training epochs (N_H being fixed), or with increasing N_H (at the end of training). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.3. Example of application: radiation-induced reactor pressure vessel steel hardening

Reactor pressure vessel (RPV) steels are well known to harden and embrittle under neutron irradiation [24]. Hardening is custom-

arily measured as the increase of the yield stress, with tensile tests performed on samples of the RPV steel. Hardening is generally directly proportional to the increase in the ductile-brittle transition temperature, or conceptually equivalent transition temperature definition. Nuclear regulations impose safety margins on these temperatures, according to rules that change depending on the country, as safeguard against RPV failure in both service and accidental conditions. In the absence of a complete physical model, from the atomic to the macroscopic level, that can describe with accuracy the relevant processes taking place under irradiation, hardening and embrittlement are predicted by semi-empirical, or totally empirical models that are mostly based on numerical fittings using experimental data [25,26]. Although inadequate to cover all possible conditions, a large amount of data from surveillance capsules and from material test reactors does exist. One of the most important goals for utilities and other nuclear stakeholders is the development, based on "clever" interpolations and extrapolations of the available data, of reliable trend curves, providing estimates of steel embrittlement as a function of the most important influencing variables. ANN are therefore a potentially interesting candidate to achieve adequate predictions, because of their abilities to extract hidden knowledge from data, and also thanks to the fact that knowledge about the physical process must not be explicitly formulated.

In Ref. [27], ANN were trained to predict the increase $\Delta\sigma_Y$ of the yield stress. This application is very illustrative of the typical limitations and practical problems faced for regression problems from scarce and valuable (expensive) experimental data, because: (a) the amount of available examples of I/O is limited to the extreme. It imposes an upper boundary on the complexity of the ANN, because the number of synapses must, ideally, be lower than the number of training examples to avoid over-fitting; (b) the separation of the data in a training and reference set may be delicate, because of conflicting constrains and inhomogeneous coverage of the input space; (c) the choice of the most adequate input variables to the ANN is not obvious. In Ref. [27], four inputs were taken into account: the Cu content of the steel, the Ni content, the neutron fluence, and, finally the irradiation temperature. The RADAMO database [28] was used as set of 346 examples of I/O. Two different algorithms to define the training and reference sets were proposed, and compared. The achieved ANN quality of predictions on the reference set is shown in Fig. 4. Predictions on a separate set that corresponds to higher neutron fluences, thus evaluating the extrapolative capabilities, are also shown. The conclusions in Ref. [27] were that ANN can accurately predict embrittlement. Extrapolation skills, e.g. for higher fluences as in Fig. 4 or for never seen steel compositions, are possible, providing that the definition of the training and reference sets are thought in accordance with this objective.

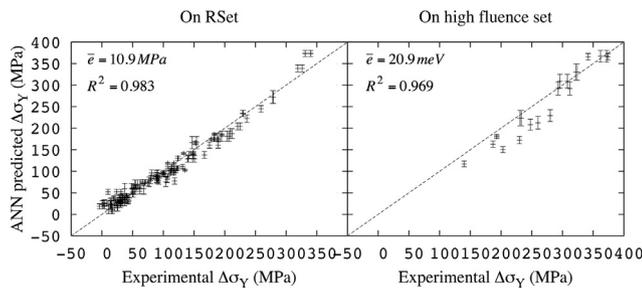


Fig. 4. ANN quality of prediction of the neutron irradiation induced increase $\Delta\sigma_Y$ in the yield stress of RPV steels. The number of input variables is 4, and the number of hidden nodes is 4. Error bars show the standard deviation of predictions made by 25 different ANN, trained with the same training and reference sets. (left) Prediction on the reference set. (right) Prediction on a separate set that corresponds to higher neutron fluences, that was not used for training.

3. Neural-network potentials (NNP) fitted from DFT

As previously introduced in Section 1, cohesive models are a central element in a multiscale modelling strategy. Density functional theory (DFT) is, in this respect, often an ideal and thus preferred choice, as it provides a reliable model, based on little approximation from the first principles of physics. Cluster expansion models, e.g., the FeCrW by Bonny et al. [29], were originally proposed to extrapolate DFT data for rigid-lattice configurations [30]. In this case, the total energy is decomposed in contributions from clusters of atoms, according to predefined patterns: these are, e.g., pairs of atoms found in first nearest-neighbour position, triangles, quadruplets, etc. While this approach has been especially taken for addressing simulations in bulk materials driven by vacancy defects [31], it has also been used to consider configurations with self-interstitials [32]. It is worth noting that the deployment of a CE in the latter case is far more delicate, mainly because of the increasing complexity in the geometry of the explored configurations and the ensuing reduction of exploitable symmetries. As a consequence, the required numbers of clusters in the expansion for a proper general description rapidly explodes.

Given these limitations, neural-network potentials (NNP) are a very promising alternative method. They naturally benefit from the generality and portability proper of ANN, therefore not making any pre-assumption on the kind of interactions between the chemical species in the target alloy. Such qualities are, with little doubt, predicated to surpass the abilities of EAM-like potentials or any CE, provided that enough data are available for training. As in the case of the CE, NNP are meant to learn from DFT directly, and no experimental data can be directly incorporated. Differently from CE, however, NNP are not limited to a rigid lattice formalism, nor by the geometrical complexity of the configurations described.

In Ref. [14], we proposed a method for designing NNP for binary Fe-based alloys, as briefly summarized in the following of this section. Consistently with the prerequisites for high-dimensional potentials, the total energy of a given atomic configuration is decomposed as proposed by Behler and Parrinello [33]:

$$E_{\text{ANN}}^{\text{(Tot)}} = \sum_{a=1}^N E_{\text{ANN}}^{X(a)}(\rho^{(a)}) \quad (5)$$

Here, ρ is the local atomic density. Superscript (a) refers to a particular atom within the N constituting the studied configuration, and $X(a)$ denotes the chemical species for atom (a) , i.e. $X(a) = \text{Fe, Cu, Cr, } \dots$. Functions $E_{\text{ANN}}^{X(a)}$ are atomic energy functions (AEF), providing an estimation of the energy assigned to every atom of the corresponding chemical species; subscript "ANN" refers to the fact that each AEF is implemented by an individual ANN. Their input variables are a description of the local atomic density ($\rho^{(a)}$ in Eq. (5)), using symmetry functions defined as follows:

$$Q_{nl}^{(a)} \triangleq \sum_{m=-l}^l |C_{nlm}^{(a)}|^2 \quad (6)$$

With

$$C_{nlm}^{(a)} = \frac{1}{N_{\text{RC}}^{(a)}} \sum_{i=1}^{N_{\text{RC}}^{(a)}} R_n(r_i) Y_{lm}(\theta_i, \varphi_i) \quad (7)$$

Here, n and m are integers defining the complexity for an expansion in series. The summation in Eq. (7) is performed over all neighbouring atoms i found within the prescribed cut-off R_C , located in space using their relative $(r_i, \theta_i, \varphi_i)$ spherical coordinates. Functions $R_n(r_i)$ are series of orthogonal radial functions, and Y_{lm} are the Laplace spherical harmonics. In a system with one chemical species, the vector of ANN input variables is defined as:

$$\bar{Q}^{(a)} \triangleq \left[Q_{nl}^{(a)} \right] \quad (8)$$

with $1 \leq n \leq N_{\text{Max}}$ and $0 \leq l \leq L_{\text{Max}}$. For alloys, information about the chemical species is included with the approach taken by Behler et al. [34]. Considering the example of a binary FeX alloy, three distinct \bar{Q} vectors are combined, each including the contribution from chemical sub-ensembles in the neighbouring atoms, i.e.:

$$\bar{Q}^{(a)} \triangleq \left[Q_{\text{Fe}}^{(a)} Q_X^{(a)} Q_{\text{FeX}}^{(a)} \right] \quad (9)$$

Here, $Q_{\text{Fe}}^{(a)}$, $Q_X^{(a)}$ and $Q_{\text{FeX}}^{(a)}$ are identical as defined in Eq. (8), but only consider either Fe atoms, X atoms, or both, respectively. Given N_{Max} and L_{Max} , the number N_Q of symmetry functions (corresponding to the ANN input variables) is thus given by:

$$N_Q = 3N_{\text{Max}}(L_{\text{Max}} + 1) \quad (10)$$

In the work described in Ref. [14], NNP were fitted for both the bcc FeCu and the FeCr system. Primarily aiming to incorporate DFT-based energies in rigid-lattice Monte Carlo models, training configurations were extracted from DFT calculation of single point-defects migration energies (both single vacancy and single self-interstitial), using the nudged elastic band (NEB) method [35,36], as illustrated in Fig. 5. To fix the ideas, each NNP was fitted from 2000 to 5000 NEB calculation (requiring 10–30 million CPU hours), providing a total number of 20,000–50,000 atomic configurations. The vector of input variables was defined using $N_{\text{Max}} = 5$ and $L_{\text{Max}} = 10$, thus leading to $N_Q = 165$. In each case, the number H of nodes in the ANN hidden layer was 3; the total number of ANN synapses was thus 502. The accuracy of prediction, after training, is summarized in Fig. 6. Similar results were later on obtained for the FeNi system, as also shown in the figure.

In addition to refine our predictions for thermal annealing experiments (see later in Section 4.1.1), the potentials were used to evaluate the phase diagram, using a Metropolis Monte Carlo method. As shown in Fig. 7, we see that DFT predicts a Cr solubility limit that is very consistent with an extensive review of experimental data previously performed in Ref. [37,38]; alloying Cr atoms should remain fully solute in the ferritic matrix at all temperatures, up to a concentration near 9at%Cr from which α' phases form.

4. Enhanced atomistic kinetic Monte Carlo models

Kinetic Monte Carlo (KMC) methods [39–42] are widespread simulation tools dedicated to describe diffusion-controlled phenomena at the atomic level. They are suitable to study a wide variety of materials up to experimentally relevant length and time scales, shedding light on the resulting microstructural and microchemical evolution during operational conditions, e.g. under irradiation [43–45]. Generally based on a rigid-lattice approach, they

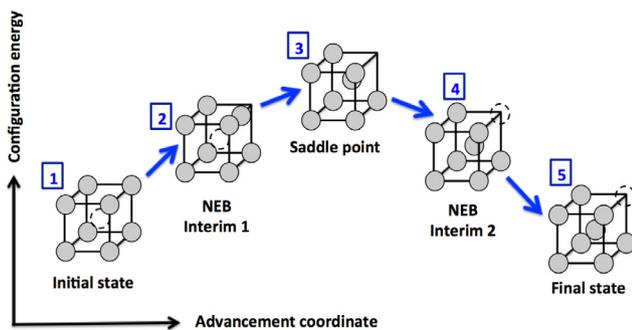


Fig. 5. Atomic configurations extracted from NEB calculations (vacancy migration event) with DFT, used to design an ANN-based potential.

feature an explicit spatial characterization of the diffusion of lattice defects and atoms, enabling a detailed investigation of the kinetics of formation of fine microstructural features. KMC methods have been widely employed to simulate metallic alloys under irradiation [41,42,40], in particular the formation of embrittling solute-defect clusters [46,32,47]. While these methods are in principle well suited for the investigation of the underlying atomic-scale mechanisms, the task is extremely challenging due to the chemical complexity of the reference alloys. Specifically, in atomistic KMC (AKMC), the evolution of the alloy proceeds through migration events of single defects (vacancies and/or interstitials) [48,39,49], which are stochastically selected at each step based on their transition rates:

$$\Gamma = \Gamma_0 \cdot \exp\left(\frac{-E_m}{k_B T}\right) \quad (11)$$

Here, k_B is Boltzmann's constant, T the absolute temperature, Γ_0 the attempt frequency and E_m the migration energy, generally evaluated in static atomic-level calculations. The accuracy of the latter parameter is thus crucial to ensure the physical reliability of the model, as it embodies both thermodynamics and kinetics properties of the system being studied.

Migration rates associated with single point defects are traditionally computed with several approaches (see, e.g., Ref. [50,51] for extensive reviews). Some are based on first-principles methods such as density functional theory (DFT) [52,53], while others rely on system-specific interatomic potentials (IAP) [40,42]. In any case, suitable mathematical expressions are necessary to predict the transition energies associated with each atomic configuration, as well as the frequencies associated with each possible transition (often migration events). These are usually based on cohesive models constructed on pair-interaction [40,54] or more sophisticated cluster-expansion methods [53], supported by limited datasets of experimental and *ab initio* properties. Their range of applicability is thus limited by their intrinsic rigid-lattice approach and by their poor transferability to new kinds of configurations beyond their original intended scope. IAP enable the description of any stable or metastable configuration, allowing the portability to lattice-free MC models [55–58]. However, their direct “on-the-fly” use as cohesive models to calculate energy barriers corresponding to all possible transitions at a given time, to be able to decide the following event in otherwise classical rigid-lattice KMC, though possible [55], is impractical, because the exact saddle-point configurations in each transition event are unknown and must be sought with time-consuming procedures.

In past work we have proposed and optimized a concrete solution to overcome this technical limitation. ANN were trained to predict the migration energy of single point-defects, otherwise obtained using the NEB method and a given cohesive model. The proposed concept is schematically depicted in Fig. 8, and can be summarized as follows. The AKMC module is found on the left-hand part of the figure, where an example of migration event for a single vacancy is indicated by the black arrow. A similar setup is used for single SIA migration events. The local atomic configuration (LAC) is defined by the species located in the closest neighbour lattice nodes to both initial and final vacancy site, denoted as A to M and encircled by a blue line in the figure. This LAC is described by a vector of numerical signals, aimed at communicating with other modules of the simulation code, and serves two purposes: (a) it is used to construct an atomic supercell suitable for lattice-free static calculations, thereby populating a database of migration energies associated to different LACs (right-hand side of the figure); (b) once the database is large enough, an ANN is trained to replace the migration energy calculation in the atomic supercell (left-hand side of the figure). The migration energies for each

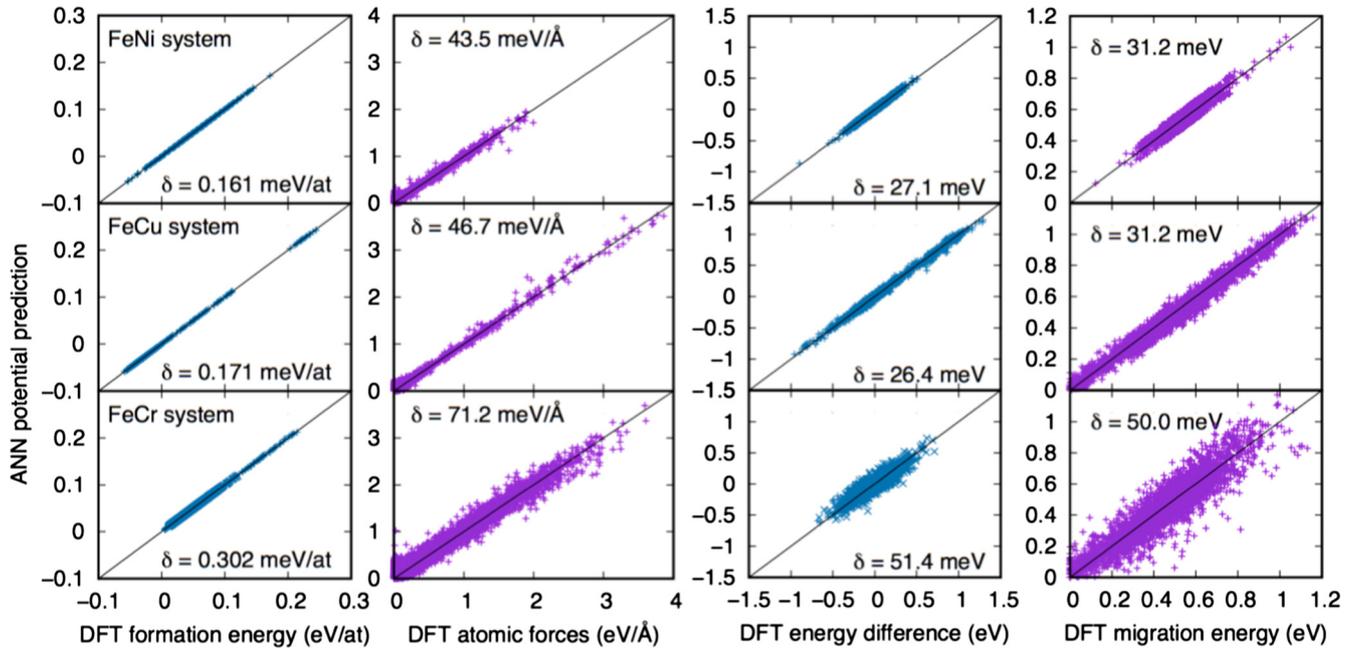


Fig. 6. Accuracy of prediction achieved by the ANN potential while reproducing NEB calculations for evaluating the saddle point during the migration of single point-defects (either a single vacancy or a single SIA) in binary bcc alloys. δ is the mean absolute error of prediction.

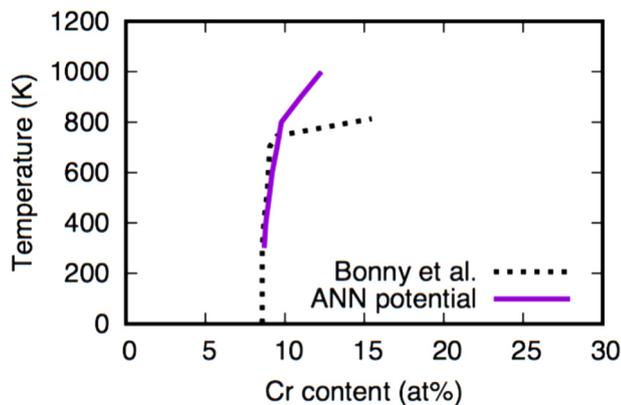


Fig. 7. Phase diagram for the FeCr binary alloy, as predicted with Metropolis MC based on our rigid-lattice NNP. The dashed line is a review from experimental data in Ref. [37,38]

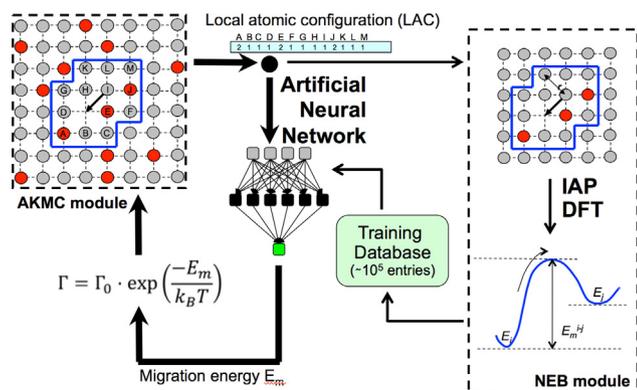


Fig. 8. AKMC model aided with ANN for fast and accurate on-the-fly estimations of the activation energies associated to the migration of single point-defects.

LAC, with defined initial and final states, are obtained by NEB calculations, using a suitable cohesive model to evaluation of total energy and atomic forces. These calculations show that, if the migration events are adequately defined, there is a unique minimum energy path between the initial and final states, thus leading to the definition of a single energy barrier. This may be directly returned to the AKMC module (on-the-fly mode), but most importantly it is also stored in a database. From a numerical standpoint, the database contains LAC vectors, with each of which a single numerical value of migration energy is associated. This database is therefore used to provide examples of connections between LAC and migration energy on which a suitable ANN is trained. The ANN is obviously expected to accurately return the migration energies corresponding to the same configurations as previously added in the database, but most importantly it is aimed at making faithful predictions for new (never previously calculated) configurations, i.e., for any atomic configuration that may be encountered during the AKMC simulation. This prediction is very fast produced by any computer, thereby gaining orders of magnitude in terms of computing time.

Clearly, the quality of the proposed simulation scheme entirely relies on the predictive abilities of the ANN for on-the-fly estimation of the energy barriers. It is worth noting that in Fig. 8 no feedback of any kind is provided to improve the ANN predictions. It was initially proposed [6] to use an additional module to perform a feedback on the ANN predictions, implemented with a fuzzy-logic-based set of rules. Experience suggested, however, that it is more practical to train ANN that are, by design, expected to be accurate during the whole simulation, making sure that both training and reference sets contain a large enough number of representative configurations [9] for any state that the simulated system can be found in.

This ANN-based AKMC model was applied to a variety of different problems, as described in the following sections. Table 1, summarizes the features and accuracies of the different ANN trained over the years for specific systems, while the goodness of the predictions as compared to the reference NEB values is given in Fig. 9.

Table 1

Summary of ANN training and accuracies. The abbreviation PU stands for previously unpublished, N_{LAC} is the number of entries in the LAC vectors, W is the number of synapses in the ANN, and δ is the mean absolute error of prediction after training. Note that the mean error of prediction is always negligible in magnitude, making the ANN predictions unbiased. All alloys are in bcc structure, except when otherwise indicated.

Alloy	Cohesive model	N_{LAC}	Num. NEB calculations	W	δ (meV)	Ref.	Case in Fig. 9
Single vacancy migration in otherwise perfect lattice							
FeCu	IAP [5]	223 (11nn)	30000	2750	14.5	[9]	A
FeCr	IAP [59]	223 (11nn)	30000	3000	20.1	[9]	B
FeCu	DFT [4]	77 (5nn)	2000	625	23.5	[4]	C
FeCr	DFT [14]	77 (5nn)	2000	340	40.6	[14]	D
FeCuNi	IAP [60]	173 (9nn)	60000	1250	43.2	[8,61]	E
FeNiCr (fcc)	IAP [62]	113 (6nn)	60000	2500	43.7	[62]	F
FeCuNiMn	IAP [63]	77 (5nn)	100000	2000	60.4	PU	G
Single SIA migration in otherwise perfect lattice							
FeCr	IAP [59]	173 (9nn)	50000	900	29.3	[64]	H
FeCr	NNP [14]	83 (6nn)	35000	880	39.1	[14]	I
Single vacancy migration when found in cluster							
Fe	IAP [65]	173 (9nn)	50000	2800	22.3	[61]	J
W	IAP [66]	83 (6nn)	20000	421	26.1	[67]	K
FeCu	IAP [5]	77 (5nn)	75000	2150	30.3	[12]	L
FeCuNi	IAP [60]	77 (5nn)	75000	1100	27.8	[61]	M
FeCrW	IAP [68]	313 (13nn)	50000	2760	49.3	[68]	N
Single SIA migration when found in cluster							
Fe	IAP [65]	77 (5nn)	90000	4100	48.6	[61]	O
Lattice free AKMC							
FeCr (near grain boundary)	IAP [69]	$R_c = 2.75a_0$ (10nn)	31000	5000	123	[13]	P

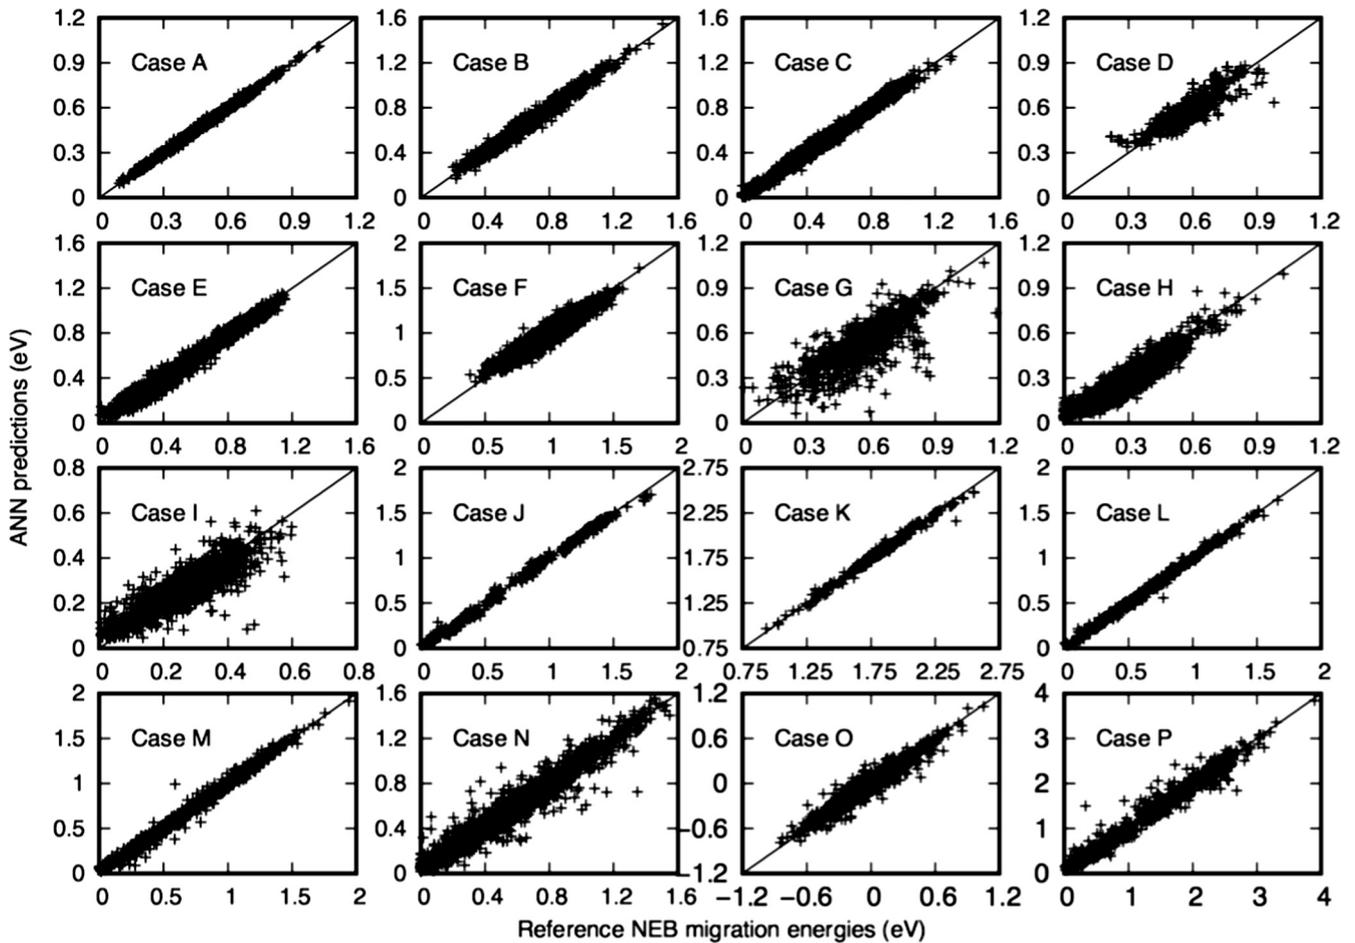


Fig. 9. ANN qualities of predictions for all systems listed in Table 1.

4.1. Rigid-lattice model

In rigid-lattice models, atoms are assumed to occupy always the nodes of a perfect lattice, e.g., bcc. This hypothesis is reasonable if single vacancies are the only defect present in the system. Defects that create a strain field, even single SIA, question the validity of this assumption. However, rigid lattice models can include more complex defects, such as vacancy clusters, single SIA or even small SIA clusters, if somehow the effect of the presence of a strain field on the mutual defect interaction is implicitly included in the characteristic energies handled by the model. This is possible without drastic modifications if the strain field only distorts slightly the lattice, without changing radically the connectivity between lattice nodes. If this is true, then the LAC vector in Fig. 8 may remain limited to a minimal amount of information: since the atomic coordinates are deducible from the crystallographic structure, which by assumption remains constant throughout the simulation, only the chemical nature of each lattice site needs to be specified. The LAC is thus unequivocally defined with a vector of integers: each of them takes a predefined value that varies with the chemistry (e.g., value 1 stands for Fe, 2 for Cu, 3 for Ni, etc.) for the atoms sitting at sites A to M in the figure. Clearly, a convention for the lattice sites ordering in the LAC vector must be defined (e.g., what relative position from the migrating vacancy is tagged as the A site in Fig. 8, etc.), and respected through the whole simulation. The symmetries in the crystallographic structure allow the migration energy calculation to be limited to one of the several equivalent LAC vectors. Considering migration events towards 1nn positions in bcc structures, for instance, each vector may be transformed into 6 equivalent ones, applying rotations along the $\langle 111 \rangle$ direction of migration, and projections along the perpendicular $[110]$ plane. See e.g. Ref. [61].

In the absence of a formal feedback on the ANN predictions, their reliability for new atomic configurations must be maximized with a proper choice of the examples in both the training and the reference set. The inherent capability of extrapolation for the numerical regression as implemented by the ANN (predict a real number, the migration energy, from a vector of integers) may be delicate to appraise. Mathematically speaking, a new configuration is a new combination of integers in the LAC vector. Since each of these integers are seen at all positions many times in the training database (in other words, many cases are included with either an Fe or an X atom sitting in a given lattice site), new configurations are actually never in an extrapolative area in the input space. However, a new configuration can be considered as extrapolative from a physical point of view, if it describes a case governed by kinds of interactions that were never included in the training set. An example is predicting the migration energy of a single vacancy in a concentrated alloy, while the ANN was trained using configurations that correspond to dilute concentrations only.

4.1.1. Thermal annealing experiments

The methodology described above was applied to simulate thermal annealing experiments in Fe-based model alloys. In a cubic simulation box with periodic boundary conditions, the prescribed content in solutes is initially introduced in random positions, as illustrated on the left-hand side in Fig. 10, together with a single vacancy. The simulation proceeds by computing its possible migration events towards each 1nn position (eight in bcc structures). The individual migration energies are evaluated by the procedure illustrated in Fig. 8, which entirely pilots the evolution of the system. Homogeneous precipitation of solutes eventually takes place, as depicted in Fig. 10, if so dictates the thermodynamics embedded in the underlying cohesive model, reflected by the ANN if the error of prediction is low enough.

For a binary FeX alloy, the integers in the LAC vectors may be any (e.g. 1 for Fe and 2 for X) without, in principle, affecting ANN training [61]. For more complex alloys, experience showed that keeping a binary description for the LAC leads to optimal ANN predictions, in spite of the increased number of input variables. For example, a ternary FeXY alloy is described using 2 bits for each lattice site: Fe atoms are coded by 00, X atoms by 10 and Y atoms by 01. Quaternary FeXYZ alloys are dealt with equally: 000 for Fe, 100 for X, 010 for Y and, finally, 001 for Z.

In Ref. [6–9], a rigorous study was performed to assess the efficiency of ANN training, and the consequent quality of the ANN predictions, as function of various factors here summarized.

Choice of configurations for training: The ANN capabilities to make relevant predictions for never seen configurations was maximized with an adequate choice of the atomic configurations added to both the training and reference set. Special care was taken that an equal proportion of examples represent each of the different stages of the solute precipitation, as depicted in Fig. 10: from configurations in a random solid solution (left), to those with small clusters formed (middle), completed with cases where the vacancy migrates near big clusters (right). It was thus reasonably assumed that the vacancy could be found in three different kinds of LAC, thereby covering the full range of possible configurations for the studied system; any new configuration either corresponds exactly to one of the three kinds, or is in an intermediate state. To our experience, a fully random and unguided choice includes totally unrealistic configurations that may mislead or even jeopardize the learning process of the ANN.

The number of neighbouring sites included in the LAC: It corresponds to the number of ANN input variables. Using an IAP, the most accurate predictions are obtained when the LAC is defined up to approximately 1.5 times the cut-off distance, in such a way that not only the direct chemical interactions (within the cut-off distance), but also the long-range chemical interactions (beyond the cut-off distance), are taken into account. Since the cut-off is at the 5th nearest-neighbours (5nn) distance for most potentials used in our works, the typical number of entries in the LAC vector is round 223 (11nn). Defining the LAC up to shorter distances removes useful information for the ANN, whereas including more shadows it by increasing the complexity of the training problem (too many synapses to fit by the LM algorithm). For ANN trained from DFT data (there is thus no formal cut-off in the cohesive model), it is the limited number of examples that restrains the possibility to successfully connect far away neighbours to the network. We found an optimal number of ANN input variables including the 5nn in the LAC.

Optimal architecture for the ANN: As already mentioned, a specific constructive algorithm (called GIACA) was proposed in Ref. [9]. In short, this training procedure connects the successive layers of close neighbours progressively, and gradually add nodes in the ANN hidden layer, until no more progress is achieved. The merit of the method is to seek for the lightest ANN possible, being economical for both the input and the number of synaptic weights.

The ANN trained for thermal annealing problems are listed in Table 1 and their performance shown in Fig. 9. They are denoted as single vacancy migration in otherwise perfect lattice, because the migrating single vacancy is the only defect present in the system, expect for solutes atoms. Using an IAP as cohesive model, the trained ANN for the FeCu and FeCr systems are tagged as A and B, respectively. We see that they are the most accurate of all. In these cases, it is reasonable to assume that the residual error of prediction almost qualifies the ANN as an undistinguishable substitute to NEB in the AKMC simulation. Later, similar ANN were trained,

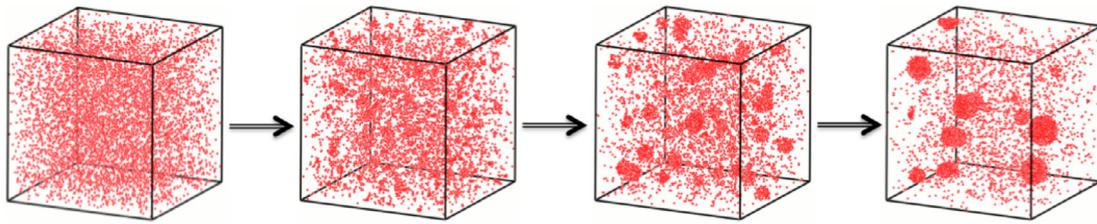


Fig. 10. Snapshots from an AKMC simulation of a thermal annealing experiment in a Fe-1.34%Cu at 500 °C.

using DFT instead of an IAP. They are indicated as C and D, respectively. Remarkably, we see that the quality of the prediction is also excellent, in spite of the (relatively) limited amount of examples available for training. Here, the LAC only includes up to the 5nn. The number W of synapses in the ANN is thus significantly smaller. More complex alloys were also considered, tagged as E, F and G in Table 1. We see that the quality of the prediction remains high, except perhaps the quaternary FeCuNiMn alloy, where the residual error of prediction is significantly larger than in the other cases. Remarkably, these ANN were trained using an amount of examples that was only 2–3 times larger than for the binary alloys, which highlights the strength of ANN regressions and their robustness with respect to the increase of the mathematical complexity.

In Ref. [9], the model predictions based on IAP were compared with experimental data for an Fe-20%Cr alloy annealed at 500 °C, as shown in Fig. 11. The comparison with the prediction obtained using a simple Kang-Weinberg (KW) [70] decomposition for the migration energy (based on a rigid-lattice calculation of the energy difference associated to the vacancy migration), demonstrates the benefits of relying on NEB-calculated barriers: AKMC predictions based on the KW overestimate the Cr clusters density. Finally, in Ref. [14], the results obtained with the ANN trained from DFT resulted in even better agreement with experimental evidence, as both the predicted clusters density and the average size coincide.

In Ref. [11,4], the model predictions based on IAP or DFT were compared with experimental data for Fe-1.34%Cu, Fe-1.1%Cu and Fe-0.6%Cu alloys, annealed at 500–700 °C. The simulations for these systems are more demanding than in the FeCr case, because of the high binding energy between the vacancy and the solute Cu atoms that slows down the simulation. A hybrid model was therefore proposed, including features of object KMC models for what concerns Cu clusters once they are formed. The AKMC model

as described in Fig. 8 was thus applied for the description of clusters seeding (the vacancy is not found in a Cu cluster), while the mobility and stability of the formed Cu clusters was managed with a coarser-grained description. For this purpose, a large series of separate AKMC simulations was conducted, in order to parametrize the diffusion coefficients and lifetimes of VacCu complexes of any sizes, up to the upper limit of the model when Cu clusters are no longer coherent with the bcc matrix (approx. 5000 atoms), see Ref. [11,61,4] for more details. In these works, the mechanism for the coherent stages of Cu precipitation in Fe was clearly highlighted, further stressing the importance of Cu clusters mobility [52]: classical Ostwald ripening is not sufficient to explain the rapid kinetics of Cu precipitation. This was rendered possible thanks to the model hybridization, and the accurate parametrization enhanced by the ANN. Even if the predictions obtained from an IAP resulted very satisfactory (see right panels in Fig. 11), the DFT-based parametrization maximized the agreement with experimental evidence. As argued in [4], the solubility limit as predicted by DFT did not match the experimental one, with a consequent overestimation of the clusters density. Nevertheless, the model significantly improved from the point of view of time rescaling, necessary to convert the time in the MC simulation into a physical time comparable to the experiment, that resulted more consistent.

In complement to the above-mentioned binary systems, the simulation of thermal annealing experiments were reported for FeCuNi alloys in Ref. [8]. Next, diffusion coefficients for single vacancies in a function of the alloy composition were reported in Ref. [62] in FeNiCr alloys.

4.1.2. Computation of diffusion coefficients for point-defects clusters

An important use of AKMC models for the parametrization of higher-scale models such as OKMC is the calculation of the diffusion coefficients and lifetime of vacancy-solute clusters. These cannot be obtained by molecular dynamics because of the too long timescale involved. In Ref. [12], the methodology initially proposed for single point-defects (shown in Fig. 8) was generalised to vacancy clusters: ANN were designed to predict the migration energy towards a 1nn position of a single vacancy, taking into account the presence of other vacancies nearby. Here too it is assumed that the LAC can be described by a vector of integers. In other words, single vacancies found in the LAC are described as an additional chemical species. ANN were trained using a very similar method, as described earlier in Section 4.1.1. The ANN quality of prediction is very satisfactory, in spite of the increased complexity: the presence of other vacancies in the LAC does not require a much larger number of NEB-calculated examples for training.

In Ref. [78,12], vacancy-Cu clusters were studied in an otherwise pure Fe matrix. Example of results are shown in Fig. 12. The mechanism of migration and dissolution of the clusters was carefully analysed. We concluded that the addition of vacancies does not enhance the mobility and dissolution of the Cu clusters, contrarily to what one could have anticipated, so long as the vacancies are surrounded by Cu atoms. Instead, as shown in the figure, vacancies tend to interact with each other inside of the Cu clusters,

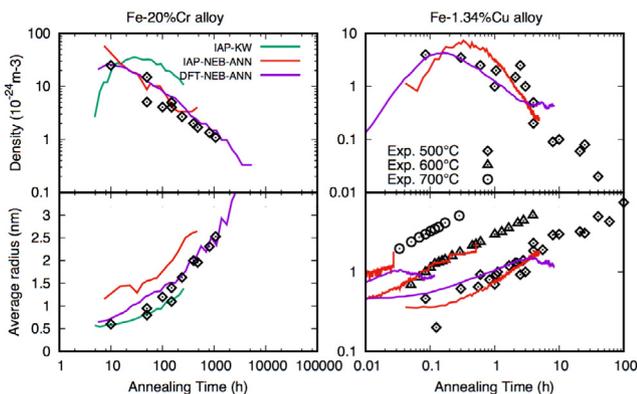


Fig. 11. Predicted densities and average sizes of solute clusters by the AKMC model, for the simulation of thermal annealing experiments. Three different methodologies to evaluate the vacancy migration energies were used. The series “IAP-KW” refers to the use of a Kang-Weinberg formula, based on an IAP, “IAP-NEB-ANN” refers to the ANN trained from NEB calculations using an IAP, whereas “DFT-NEB-ANN” is the same method using DFT instead of the IAP. (left) in Fe-20%Cr at 500 °C. Experimental data is taken from Ref. [71–73]; (right) in Fe-1.34%Cu at 500 °C, 600 °C and 700 °C. Experimental data is taken from Ref. [74–77].

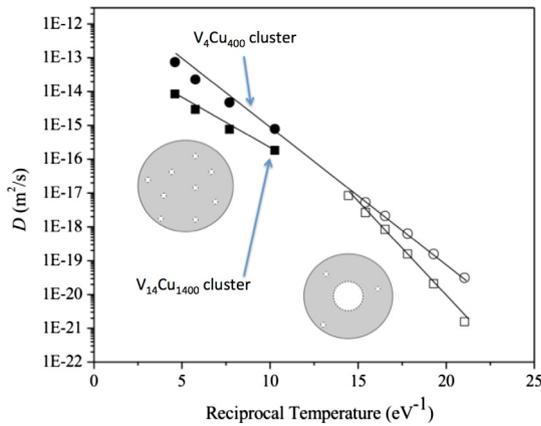


Fig. 12. Example of Arrhenius plot computed with the ANN-based AKMC for the diffusion coefficient of Vac-Cu clusters.

thus reducing the effective migration. Once again, only the use of a reliable IAP and the NEB-ANN method could give results not accessible to molecular dynamics. An interesting exercise was performed in Ref. [78]: for small clusters, the geometry of the problem is reduced to such an extent that it is actually feasible to catalogue all possible vacancy migration events. Therefore, a direct comparison could be made between, on the one hand, the calculated diffusion coefficient while using the IAP-NEB only, and on the other hand, the coefficients obtained while using the trained ANN to replace the NEB. We could see that the effect of the ANN error of prediction was negligible.

A similar methodology was followed for describing small SIA clusters. The ANN was trained to predict the energy difference associate to SIA migration events in pure Fe, taking into account the presence of other SIA nearby. It is tagged as O in Table 1 and in Fig. 9.

4.1.3. Simulation of resistivity recovery experiments

In a metallic sample, resistivity is directly related to the amount and type of defects it contains. The diffusion of the latter is activated at different temperatures. Hence, with a gradual temperature increase and a monitoring of sudden changes in resistivity, it is possible to experimentally determine the activation energy of each defect type. Such experiments can be closely reproduced with AKMC, and thus constitute a precious benchmark for AKMC model development. In Ref. [64], such experiments were simulated in Fe-Cr alloys. In addition to the ANN already mentioned in Section 4.1.1, dedicated to the prediction of the single vacancy migration energy, another one was trained for dealing with single SIA. Using first an IAP by Olsson et al. [59], the obtained ANN is tagged as case H in Table 1 and in Fig. 9. A similar network, tagged as case I, was later on fitted in Ref. [14], relying on DFT as cohesive model. It is worth noting the NNP shown in Fig. 6 was in fact employed in an intermediate step.

The work reported in Ref. [64] highlighted how the prediction by the AKMC model are improved thanks to a proper criterion in vacancy-SIA recombination. Instead of using a fixed distance criterion, relative stable or unstable configurations were systematically evaluated with static relaxation. Also, the employment of the above-mentioned ANN, thus predicting migration energies as function of the exact LAC, lead to a realistic prediction for the progressive suppression of stage I_E with increasing Cr content.

4.2. Lattice-free model

The ANN-based AKMC algorithm as described in the previous section is not directly applicable to systems where the rigid lattice

assumption is no longer valid, e.g., near free surfaces, grain boundaries, in the presence of dislocations, or in the presence of nanostructural features such as dislocation loops. By removing the rigid lattice assumption, the definition of the transitions and especially of the LAC becomes much more complex. Thus a different formulation is needed.

Fully lattice-free AKMC models do exist. They constantly explore the local curvature of the potential energy surface, and find transitions to nearby basins looking for saddle-points, as specified by the transition-state theory [49]. For example, in Ref. [55], Henkelmans and Jonsson used the Dimer method [79], which in theory searches for all possible transitions without the need to make assumptions. An elegant alternative to this method was proposed in Ref. [80], searching for the same saddle points but using a single image of the studied system, thus alleviating the computing cost. Other authors developed different schemes, though accomplishing the same finality, e.g., the ART method [81] in Ref. [82]. In these, the definition of the migration events, and the calculation of their corresponding migration barriers, is made on the fly. Clearly, the advantage is the flexibility with respect to the simulated system. However, any saddle-point searching algorithm must be adequately parametrized to guarantee that most transitions are found, which can be delicate for some systems. Most importantly, the required computing time is prohibitive for performing long simulations, inherently limiting the practical application of the method to no more than a few thousands of events. This is undoubtedly insufficient to study long-term and slow processes such as precipitation or depletion of solutes at interfaces.

In Ref. [13], we developed similar ideas than our ANN-based model described in the previous sections, for proposing a compromise between the afore-discussed extremes, as briefly summarized as follows:

- The search for any possible transition events with a fully general method such as the Monomer is equivalent to test migration vectors in the 3N space, starting from the present state of the system. This allows the definition of a large number of migration events, which can become unmanageable. Additionally, many events found are likely to be minor (e.g., slight displacement of a single atom to a very nearby stable position), thus not letting the system advance significantly in time. Instead, we defined a generic procedure for defining transitions, assuming the most likely events. In Ref. [13], the studied system was a grain boundary in FeCr alloys. Making legitimate assumptions regarding its specificities, migration events were defined looking for sites with enough open volume (thus playing the role of a pseudo-vacancy), using a geometrical criterion based on Voronoi vertexes, as shown in Fig. 13a. In the figure, seven atoms are found within a prescribed maximal distance from an eligible site, thus leading to the definition of seven transition events (the migration of any of them towards the site). A vector of migration is thus

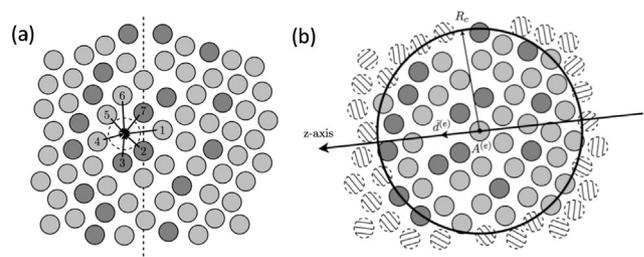


Fig. 13. Application of the ANN-based AKMC model to a lattice-free problem. (a) Definition of possible transition events with a generic procedure. (b) Definition of a migration vector (denoted as $d^{(e)}$) and a LAC (atom within a given cut-off radius R_c) for a given event.

naturally defined for each transition, as illustrated in Fig. 13b. The direction of evolution for the system is thus reduced to a vector in three dimensions (indicated as $\vec{d}^{(e)}$ in Fig. 13b), expressing the relative translation for the migrating atom.

- The activation energy associated with each events was defined using a semi-rigid-lattice procedure: the migrating atom is first rigidly translated to the destination site, and later on full static relaxation using conjugate gradients [83] is applied. The obtained configuration is defined as the final state assigned to the event, and NEB is applied to evaluate the migration energy.
- Similarly to the idea illustrated in Fig. 8, many examples of transition events are gathered in a database, until enough data are available for designing an ANN, aimed at fully replacing the NEB calculations. Similarly to the rigid-lattice AKMC, its purpose is strictly to provide a numerical estimate of the migration energy. The required inputs are, again, a description of the initial state (calculated from the relaxed atomic coordinates in the 3N space). The direction of migration is, also, implicitly provided, because the z-axis of the attached referential system of spherical coordinates is, by convention, aligned with the migration vector $\vec{d}^{(e)}$. It is worth noting that the final state of the system, i.e., after a given migration event is applied, is not known. Differently from the rigid lattice case, it cannot be fully deduced without the application of static relaxation.
- Following a basic Monte Carlo algorithm, a list of events is established at every step of the simulation using the generic procedure, and the ANN is used to estimate the associated migration energies. One event is thus stochastically selected amongst the others. From this point, the underlying method for finding transition events (Monomer, ART, etc) must be fully applied. In the case of our simplified approach in Ref. [13], only static relaxation is necessary, after the migration atom is rigidly placed at its destination. The application of NEB is then optional, and may be reserved for occasional feedback on the ANN predictions.

The major advantage of our proposed lattice-free AKMC model is that the CPU cost is reduced to the minimum necessary. Instead of applying a general transition search method many times at each step of the simulation (for which it is in principle necessary to define a list of events), this operation is performed only once, i.e., for the selected event only. Once again, this algorithm entirely relies on the capacity of an ANN to make faithful predictions of the events migration energies, given as input a description of the initial state, and of the migration vector $\vec{d}^{(e)}$. We demonstrated in Ref. [13] the feasibility, successfully training an ANN for the case of a single vacancy migrating near a grain boundary in FeCr alloys. It is tagged as P in Table 1 and in Fig. 9. Interested readers are directed to Ref. [13] for a full description of the method we followed to train it, and in particular how the LAC was described in ANN input variables. Concisely, similar ideas as developed in Section 3 were followed. Taking advantage of the symmetry of the system, the ANN input variables were taken to be the moduli of the C_{nm} coefficients in Eq. (7). They are, by construction, invariant with respect to rigid rotations around the migration vector $\vec{d}^{(e)}$ in Fig. 13.

We see in Table 1 and in Fig. 9 that the quality of prediction of the obtained ANN is very high, and that the required number of NEB-calculated examples was not exceedingly high. Nevertheless, open issues for an efficient portability of the model to more complicated system do remain. First, the question of the ANN prediction reliability for never seen cases is delicate to appreciate in practice. Differently from the rigid-lattice case, new configurations may leave the domain in the input space that was covered during training, putting thus the ANN in extrapolating modes. This is in fact not an issue, however, because a constant feedback is auto-

matically performed on the ANN predictions, as discussed above. The ANN may thus be retrained as necessary, while the AKMC simulation proceeds. Secondly, the method may work in a less efficient way when the transition events imply large distortions from the initial state, as highlighted in Ref. [13].

5. Future perspectives: proper modelling of radiation-induced hardening in ferritic steels

Artificial neural networks and machine learning schemes in general may find manifold applications in a multiscale modelling framework. The examples provided in the present review are essentially all based on the idea of predicting how chemically complex local atomic configurations, and/or specific strain fields, influence the activation energy of selected thermally activated events, short-cutting computationally heavy calculation methods. Appropriately trained ANN are then used to inform more or less standard atomistic KMC models. Similar schemes can be used to calculate other quantities than activation energies and to transfer atomistic details to non-atomistic models, thereby helping in the effort of bridging between scales. Here a few examples are discussed.

Object kinetic Monte Carlo (OKMC) models have been widely used to describe the microstructural evolution in materials under irradiation, in terms of radiation defects, i.e. vacancies, SIA and their clusters. For example, a model of this type was successfully developed in Ref. [84] to simulate irradiation processes in Fe-C systems, as reference system for models addressing steels, i.e. with a more complex composition, and in Ref. [67] for W-C systems. In both cases, however, the only objects explicitly treated were point-defects and their clusters. The parameters describing their migration and dissociation were specific for Fe or W, and the effect of C was effectively introduced in terms of traps for mobile clusters, without explicitly introducing C atoms. Likewise, models of this type have been extended to model chemically more complex systems, e.g. Fe-Mn-Ni as representative of RPV steels [45] or Fe-Cr as representative of ferritic-martensitic steels [85]. The effect of the presence of solutes was implicitly introduced in terms of changes of parameters due to the presence of solutes, without explicitly introducing them in the simulation box. This was achieved in a simplified way, assuming that the solute atoms are always uniformly and randomly distributed in the simulation volume, so that their effect is inherently independent of local fluctuations of composition.

A future development of these models would consist in making the parameters of migration and dissociation of point-defect clusters sensitive to the local composition. As a matter of fact, phenomena of radiation enhanced or induced heterogeneous precipitation are expected to create significant fluctuations in the local composition. Thus, vacancies will form complexes with solute atoms, mobile SIA clusters and dislocation loops will be repelled or attracted by regions rich in a given solute, while immobilized defects of the same type will become decorated by solutes, especially when these are dragged by point-defects. In order to describe these processes, it becomes necessary to know, for example, how the formation, trapping, and migration energy of point-defect clusters changes as a function of the local composition. Such a type of information cannot be calculated on-the-fly. First of all, such calculation would offset the advantage of non-atomistic modelling tools, i.e. the fact of simulating relatively long timescale. Secondly, concentration-dependent quantities are inherently averages or else randomly selected values for different configurations with the same local solute concentrations. ANN trained to calculate these quantities as functions of the local atomic configurations would allow the introduction of the chemical and atomistic detail in models that in fact do not include atoms but handle only, in the best case, local concentrations.

Therefore, a more realistic description of solute dragging by single point-defects could be provided, beyond the assumption of infinite dilution limit, incorporating the effect of local chemical configurations. For instance, the dragging of a solute by a point-defect may be reduced by the presence of other solutes, or the carrying defect may more favourably bind to a different solute encountered along the way. Given an adequate cohesive model as input, our ANN-based technique described in Section 4.1 could be integrated in the OKMC model to provide on-the-fly appropriate parameters at negligible computational cost.

Another example concerns the behaviour of SIA loops decorated by solutes. By DFT some data on the relevant binding energy can be obtained in presence of a limited number of solutes at a time and it is impossible to explore all possible ways in which solutes will aggregate around a loop. A valuable improvement in the model would be to train ANN to accurately predict the binding energy between SIA loops and decorating solutes, ideally on DFT data. Given as input a loop size in number of SIA and a content in solutes (either as an explicit set of spatial coordinates, or as a composition in regions of space), ANN could assimilate how the effective binding energy changes.

Similarly, and without mutual exclusion, ANN could be used to predict how the migration energy of mobile defects is influenced by the presence, in non-atomistic models such as those of OKMC type, of strain fields, such as those created by dislocation loops and lines, grain boundaries, etc. It is of course possible to describe this effect by associating with each extended defect the strain field that it generates, as calculated in an elasticity theory framework, treating for example single point-defect as elastic dipoles and parametrising their migration energy on DFT as function of the dipole/strain field interaction. However, this is an approximation that slows down significantly the calculation and is limited to relatively simple strain fields, the description of which becomes complicated when for example the materials are anisotropic, or when strain field superposition occurs. In contrast, ANN could be trained in this sense directly on atomistic models, to provide, say, the migration energy of a single point-defect as function of the local strain field. Moving to different scales and processes, when applying dislocation dynamics models to describe plastic deformation in irradiated materials, it becomes necessary to define local rules of interactions between specific types of dislocations and obstacles. These obstacles can be solute clusters or precipitates of different type and composition, loops orientated in different ways, decorated or not by solutes, etc. By molecular dynamics it is possible to explore only a limited number of configurations, temperature being also a key variable to define the result of the dislocation/obstacle reaction. Once again, ANN trained to predict the key parameters that govern the interaction between dislocations and defect, depending on the specific features of both, would enable dislocation dynamics models to become more local and take into account variables, such as chemical composition of obstacles, that are currently very difficult or impossible to include in these models.

All the examples given concern quite strictly machine learning applied to varying atomistic configurations. In a more general and futuristic way, partially out of the scope of the present review, machine learning may one day shortcut and replace numerical simulation models that require long computing times to provide an answer. These models, e.g. suite of codes that, through different scales, provide the increase of yield strength as a consequence of irradiation for a given temperature, flux, fluence, materials composition, etc, may be used to provide examples of results for given initial and operating conditions, on which suitable machine learning schemes could be trained, possibly also completed by experimental examples. These machine learning schemes would then eventually be able to provide equivalent results, at negligible

computational cost. Such a scheme of use would represent a step forward, in terms of physical reliability, with respect to the example provided in Section 2.3 in which ANN were used as a regression tool trained directly (and physically blindly) on experimental measurements of yield strength increase, as a function of operation and initial material conditions.

6. Conclusive remarks

In this paper, we have reviewed in detail our work aimed at achieving a more physically accurate parametrization of atomic-scale modelling, more specifically kinetic Monte Carlo models devoted to the description of irradiation-induced microstructure changes in metals and alloys. However, the value of the concepts and examples here presented goes beyond the field of materials modelling, as there exist many more research fields where ANN, and machine-learning techniques in general, can provide substantial contributions and help building more advanced and accurate models. It is in fact common in many modelling activities to face the need for more powerful and flexible regression techniques, which may appreciatively enhance the quality and applicability of the model. Always beyond the specific application to irradiated materials, another valuable take-home message that can be applied in plenty of other fields is to always treat ANN with a critical eye. Our experience tells that ANN, expectedly or not, lead to accurate and trustworthy predictions, providing that at least two important conditions are fulfilled.

First of all, the provided set of training reference data must be strictly self-consistent. Often, this condition might seem obvious and fulfilled by default, but this is not always the case. For instance, in the case of a single vacancy migrating in an Fe-based alloy, all events treated by the ANN are unquestionably of the same kind, thus self-consistency concerns essentially the numerical parametrization. In that respect, it is necessary to handle the large amount of required NEB calculations with some automatic scripts, which might be more challenging than expected considering that it is no longer possible to carefully examine and verify each individual case. Many times, we have realized that something was not optimal in the NEB setup, only after having trained the ANN and analysed its predictions (e.g., similar plots to those in Figs. 6 and 9). While outliers clearly mark specific (pathological) cases that are easily addressable, it is surely not straight-forward to identify and solve the issue in a blurry cloud of points. Sometimes it was necessary to recompute whole NEB batches with more stringent convergence parameters, or to improve the convergence of the end-state relaxations, to obtain a more satisfactory ANN quality of prediction. The task was most challenging in the lattice-free AKMC models, given the variety and complexity of many encountered migration paths, and the possibility that even a specific type of transition (e.g., a vacancy migrating or not) can unexpectedly change during the NEB relaxation.

Secondly, a particular attention must be paid to the risk of ending up with hidden correlations in the training database. In addition, some clumsy mistakes, such as the accidental mixing of unrelated but not incompatible data, can be equally dangerous. In a way, this may be regarded as the worst, and at the same time most natural enemy of a black-box approach, that can transform their best quality into their most serious shortcoming. If the training set is sufficiently self-consistent, the ANN will undoubtedly manage to perform “accurate” predictions of the reference set, even when there is something terribly wrong with the former. Hidden correlations and mistakes in the data can thus be hard to spot, because of the ANN capability of assimilating and integrating them in its inner logic, without showing evident symptoms to the blind user. We provide here three illustrative examples:

- In the hardening prediction based on surveillance data, reported in Section 2.3, it was easy to accidentally compile a training set referring to samples with correlated chemical compositions. In the surveillance dataset, the Ni content was in a linear correlation with the Cu content, which is not surprising, as this is dictated by a certain justified technological logic. If this dataset is used as such for training, the obtained ANN is unavoidably bound to that logic, and is not treating the two parameters Ni and Cu content as independent variables, as opposed to the original scope.
- Hidden correlations were also found in the application of the lattice-free AKMC in Section 4.2. With an incautious approach, it is easy to put together a database of migration events that are not sufficiently diversified, or whose associated end states are highly correlated. For instance, the migration vector $\vec{d}^{(e)}$, which should be in principle independent from the initial state, was found instead to be statistically deducible from it, due to an inadequate sampling of the possible migration events, or to the fact that the training configurations were too similar to bulk bcc. As a consequence, the ANN most likely exploited this accidental statistical correlation, thus learning a logic that is only partially applicable to all events that can be encountered during the simulation.
- Finally, we witnessed incredible cases where the ANN was able to assimilate to a high degree of accuracy training data that was completely and unquestionably wrong, as a consequence of (human) mistakes that can be likely when handling large amounts of data with automatic scripts. For example, while assembling the training dataset to design the FeNi NNP described in Section 3, we accidentally mixed two batches of incompatible data, namely referring to two different definitions of the cohesive energy of each chemical species. In spite of the evident inconsistency, the obtained NNP was (apparently) satisfactory, at least for the cases shown in Fig. 6! During the atomistic simulations, nothing wrong emerged as long as the system remained close to regular bcc; however, as soon as it departed from it (e.g. during static relaxation, or at the saddle point of a migration event), the energy landscape became unstable and inconsistent, which made us realize that the NNP training went, in fact, completely wrong.

To conclude, ANNs are clearly very promising tools, but they must be handled with care. With their black-box approach, they often provide a high-quality parametrization that can be comfortably and fruitfully exploited, but that can hide completely unphysical results. In other words, this shows that artificial intelligence and machine learning do not free us human scientists from the duty of critical thinking. They rather allow us to build up more advanced models, as well as to divert our thinking from lower-level repetitive tasks (such as looking for mathematical expressions to describe a migration energy as a function of the LAC) to more valuable and meaningful aspects: is the problem well formulated from a mathematical and physical standpoint? Is the provided data relevant or adequate? Are all physical aspects taken into account?

Acknowledgements

This work was performed, and partially funded, in the framework of various projects by the European Commission: (a) The FP6 PERFECT project under contract FI60-CT-2003-508840; (b) The Euratom FP7 (2007–2011), grant agreement number 212175 (GetMat project); (c) Partial funding from Euratom's 6th Framework Programme integrated project PERFECT, Contract No. FI60-CT-2003-5088-40; (d) From the European Atomic Energy Commu-

nity's 7th Framework Program (FP7/2007–2011) (Grant No. 232612) (Perform60 project); (e) the H2020 European project SOTERIA (No. 661913); (f) the FP7-Euratom-Fission European project MATISSE (No. 661913); (g) Contributions to the Joint Programme on Nuclear Materials of the European Energy Research Alliance; (h) H2020 project M4F (No. 755039).

M.I.Pascuet acknowledges support from the Belgian Federal Scientific Policy Office (post-doctoral Fellowship). N. Castin acknowledges PIP-CONICET 804/10 project (post-doctoral Fellowship).

The high-performance computing resources were provided by Électricité de France (R&D) and by the Swedish National Infrastructure for Computing (SNIC).

References

- [1] C. Becquart, A. Barbu, J. Bocquet, M. Caturia, C. Domain, C.-C. Fu, S. Golubov, M. Hou, L. Malerba, C. Ortiz, A. Souidi, R. Stoller, J. Nucl. Mater. 406 (2010) 39.
- [2] C.S. Becquart, C. Domain, Met. Mat. Trans. A 42 (2011) 852.
- [3] P. Olsson, T. Klaver, C. Domain, Phys. Rev. B 81 (2010) 054102.
- [4] L. Messina, N. Castin, C. Domain, P. Olsson, Phys. Rev. B 95 (2017) 214117.
- [5] R. Pasianot, L. Malerba, J. Nucl. Mater. 360 (2007) 118.
- [6] N. Castin, R.P. Domingos, Int. J. Comput. Intell. Syst. 1 (2008) 4.
- [7] N. Castin, L. Malerba, Nucl. Instrum. Methods Phys. Res. B 267 (2009) 3148.
- [8] N. Castin, L. Malerba, G. Bonny, M. Pascuet, M. Hou, Nucl. Instrum. Methods Phys. Res. B 267 (2009) 3002.
- [9] N. Castin, L. Malerba, J. Chem. Phys. 132 (2010) 074507.
- [10] N. Castin, G. Bonny, D. Terentyev, M. Lavrentiev, D. Nguyen-Manh, J. Nucl. Mater. 417 (2011) 1086.
- [11] N. Castin, M.I. Pascuet, L. Malerba, J. Chem. Phys. 135 (2011) 064502.
- [12] N. Castin, M.I. Pascuet, L. Malerba, J. Nucl. Mater. 429 (2012) 315.
- [13] N. Castin, J. Fernandez, R. Pasianot, Comput. Mater. Sci. 84 (2014) 217.
- [14] N. Castin, L. Messina, C. Domain, R. Pasianot, P. Olsson, Phys. Rev. B 95 (2017) 214117.
- [15] C.M. Bishop, Neural Networks for Pattern Recognition, Clarendon, Oxford, 1995.
- [16] C. Bishop, Rev. Sci. Instrum. 65 (1994) 1803.
- [17] S. Fahlman, The recurrent Cascade-correlation architecture, in: Richard P. Lippmann, John E. Moody, David S. Touretzky (Eds.), Advances in Neural Information Processing Systems, 1991.
- [18] M.S.K. Hornik, H. White, Neural Netw. 2 (1989) 359.
- [19] K. Levenberg, Q. Appl. Math. 2 (1944) 164.
- [20] D. Marquardt, SIAM J. Appl. Math. 11 (1963) 431.
- [21] T.-Y. Kwok, D.-Y. Yeung, IEEE Trans. Neural Netw. 8 (1997) 630.
- [22] C. Campbell, Neural Network Systems Technologies and Applications, Constructive Learning Techniques for Designing Neural Network Systems, C. T. Leondes Academic, San Diego, 1997.
- [23] R. Parekh, J. Yang, V. Honavar, IEEE Trans. Neural Netw. 11 (2000) 2.
- [24] IAEA Nuclear Energy Series, Integrity of Reactor Pressure Vessels in Nuclear Power Plants: Assessment of Irradiation Embrittlement Effects in Reactor Pressure Vessel Steels, No. NP-T-3.11, 2009.
- [25] P. Petrequin, A Review of Formulas for Predicting Irradiation Embrittlement of Reactors Vessel Materials, AMES Report No. 6, November 1996.
- [26] A.-S. Bogaerts, R. Gerard, R. Chaouadi, Comparison of embrittlement trend curves to high fluence RPV surveillance results, in: International Symposium on Contribution of Materials Investigations to Improve the Safety and Performance of LWRs (Fontevraud 7) Avignon, France.
- [27] N. Castin, L. Malerba, R. Chaouadi, J. Nucl. Mater. 408 (2011) 30.
- [28] R. Chaouadi, R. Gerard, J. Nucl. Mater. 345 (2005) 65.
- [29] G. Bonny, N. Castin, C. Domain, D. Terentyev, Philos. Mag. 97 (2016) 1.
- [30] A. van de Walle, G. Ceder, J. Phase Equilibria 23 (2002) 348.
- [31] A.V. der Ven, G. Ceder, Phys. Rev. B 71 (2005) 054102.
- [32] R. Ngayam-Happy, C. Becquart, C. Domain, L. Malerba, J. Nucl. Mater. 426 (2012) 198.
- [33] J. Behler, M. Parrinello, Phys. Rev. Lett. 98 (2007) 146401.
- [34] W.J. Szlachta, A. Bartók, G. Csányi, Phys. Rev. B 90 (2014) 104108.
- [35] G. Henkelman, H. Jónsson, J. Chem. Phys. 113 (2000) 9901.
- [36] G. Henkelman, H. Jónsson, J. Chem. Phys. 113 (2000) 9978.
- [37] G. Bonny, D. Terentyev, L. Malerba, Scr. Mater. 59 (2008) 1193.
- [38] G. Bonny, D. Terentyev, L. Malerba, J. Phase Equilibria Diffus. 31 (2010) 439.
- [39] A.B. Bortz, M.H. Kalos, J.L. Lebowitz, J. Comput. Phys. 17 (1975) 10.
- [40] Y.L. Bouar, F. Soisson, Phys. Rev. B 65 (2002) 094103.
- [41] P. Krasnochtchekov, R.S. Averback, P. Bellon, Phys. Rev. B 75 (2007) 144107.
- [42] G. Bonny, D. Terentyev, L. Malerba, D.V. Neck, Phys. Rev. B 79 (2009) 104207.
- [43] C.-C. Fu, J.D. Torre, F. Willaime, J.-L. Bocquet, A. Barbu, Nat. Mater. 4 (2005) 68.
- [44] Y. Fan, A. Kushima, S. Yip, B. Yildiz, Mechanism of void nucleation and growth in bcc Fe: atomistic simulations at experimental time scales, Phys. Rev. Lett. 106 (2011) 125501.
- [45] M. Chiapetto, L. Malerba, C. Becquart, J. Nucl. Mater. 462 (2015) 91, and references therein.
- [46] E. Vincent, C.S. Becquart, C. Domain, J. Nucl. Mater. 382 (2008) 154.

- [47] Y. Zhang, P. Chakraborty, S.B. Biner, Modeling of Late Blooming Phases and Precipitation Kinetics in Aging Reactor Pressure Vessel (RPV) Steels, Tech. rep., Idaho National Laboratory (INL), 2013.
- [48] W.M. Young, E.W. Elcock, Proc. Phys. Soc. Lond. 89 (1966) 735.
- [49] K.A. Fichthorn, W.H. Weinberg, J. Chem. Phys. 95 (1991) 1090.
- [50] A. Chatterjee, D. Vlachos, An overview of spatial microscopic and accelerated kinetic monte carlo methods, J Comput.-Aided Mater. Des. 14 (2007) 253.
- [51] C.S. Becquart, C. Domain, Phys. Status Solidi B 247 (2010) 9.
- [52] F. Soisson, C.-C. Fu, Cu-precipitation kinetics in α -Fe from atomistic simulations: vacancy-trapping effects and Cu-cluster mobility, Phys. Rev. B 76 (2007) 214102.
- [53] M.Y. Lavrentiev, R. Drautz, D. Nguyen-Manh, T.P.C. Klaver, S.L. Dudarev, Monte carlo study of thermodynamic properties and clustering in the bcc Fe-Cr system, Phys. Rev. B 75 (2007) 014208.
- [54] S. Schmauder, P. Binkele, Comput. Mater. Sci. 24 (2002) 42.
- [55] G. Henkelman, H. Jónsson, J. Chem. Phys. 115 (2001) 9657.
- [56] F. El-Mellouhi, N. Mousseau, L.J. Lewis, Phys. Rev. B 78 (2008) 153202.
- [57] D. Konwar, V.J. Bhute, A. Chatterjee, J. Chem. Phys. 135 (2011) 174103.
- [58] H. Xu, R. Stoller, L.K. Béland, Y.N. Osetsky, Comput. Mater. Sci. 100 (2015) 135.
- [59] P. Olsson, J. Wallenius, C. Domain, K. Nordlund, L. Malerba, Phys. Rev. B 72 (2005) 214119.
- [60] G. Bonny, R. Pasianot, N. Castin, L. Malerba, Philos. Mag. 89 (2009) 3531.
- [61] N. Castin, Atomistic Kinetic Monte Carlo Simulation of Precipitation and Segregation in Metals for Nuclear Applications, Using a Novel Methodology Based on Artificial Neural Networks (Ph.D. thesis), Université Libre de Bruxelles (ULB), Belgium, 2011.
- [62] G. Bonny, N. Castin, D. Terentyev, Model. Simul. Mater. Sci. Eng. 21 (2013) 5004.
- [63] G. Bonny, D. Terentyev, A. Bakaev, E. Zhurkin, M. Hou, D.V. Neck, L. Malerba, J. Nucl. Mater. 442 (2013) 282.
- [64] D. Terentyev, N. Castin, C. Ortiz, J. Phys.: Condens. Matter. 24 (2012) 475404.
- [65] M. Mendelev, S. Han, D. Srolovitz, G. Ackland, D. Sun, M. Asta, Phil. Mag. 83 (2003) 3977.
- [66] M.-C. Marinica, L. Ventelon, M. Gilbert, L. Proville, S. Dudarev, J. Marian, G. Bencteux, F. Willaime, J. Phys.: Condens. Matter 25 (2013) 395502.
- [67] N. Castin, A. Bakaev, G. Bonny, A. Sand, L. Malerba, D. Terentyev, J. Nucl. Mater. 493 (2017) 280.
- [68] G. Bonny, N. Castin, J. Bullens, A. Bakaev, T. Klaver, D. Terentyev, J. Phys.: Condens. Matter. 25 (2013) 315401.
- [69] G. Bonny, R. Pasianot, D. Terentyev, L. Malerba, Philos. Mag. 91 (2011) 1724.
- [70] H.C. Kang, W.H. Weinberg, J. Chem. Phys. 90 (1989) 2824.
- [71] S. Novy, P. Pareige, C. Pareige, J. Nucl. Mater. 384 (2009) 96.
- [72] S. Novy, Ph.D. thesis, University of Rouen, France, 2009.
- [73] F. Bley, Acta Metall. Mater. 40 (1992) 1505.
- [74] S.R. Goodman, S.S. Brenner, J.R. Low, Metall. Trans. 4 (1973) 2363.
- [75] R. Kampmann, R. Wagner, in: C. Janot, W. Petry, D. Richter, T. Springer (Eds.), Atomic Transport and Defects in Metals by Neutron Scatterings, Springer, New York, 1986.
- [76] M.H. Mathon, A. Barbu, F. Dunstetter, F. Maury, N. Lorenzelli, C.H. de Novion, J. Nucl. Mater. 245 (1997) 224.
- [77] M. Perez, F. Perrard, V. Massardier, X. Kebler, A. Deschamps, H. de Monestrol, P. Pareige, G. Covarel, Philos. Mag. 85 (2005) 2197.
- [78] M. Pascuet, N. Castin, C. Becquart, L. Malerba, J. Nucl. Mater. 412 (2011) 106.
- [79] G. Henkelman, H. Jónsson, J. Chem. Phys. 111 (1999) 7010.
- [80] V. Ramunni, M. Alurralde, R. Pasianot, Phys. Rev. B 74 (2006) 054113.
- [81] G. Barkema, N. Mousseau, Comput. Mater. Sci. 20 (2001) 285.
- [82] L.K. Béland, P. Brommer, F. El-Mellouhi, J.-F. Joly, N. Mousseau, Phys. Rev. E 84 (2011) 046704.
- [83] M.C. Payne, M.P. Teter, D.C. Allan, T.A. Arias, J.D. Joannopoulos, Rev. Mod. Phys. 64 (1992) 1045.
- [84] V. Jansson, L. Malerba, J. Nucl. Mater. 443 (2013) 274.
- [85] M. Chiapetto, L. Malerba, C. Becquart, J. Nucl. Mater. 465 (2015) 326.