Friction

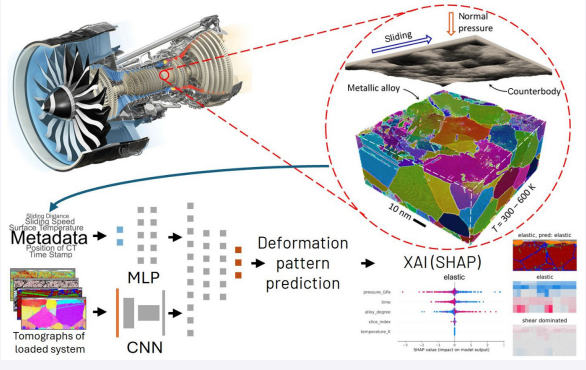
Research Article

Toward machine-learning-based interpretation of tribological deformation patterns in molecular dynamics data

Hendrik J. Ehrich[1,2,✉], Marvin C. May[3], Stefan J. Eder[1,2,✉]

Cite this article: Ehrich HJ, May MC, Eder SJ. *Friction* 2026, **14**: 9441269. https://doi.org/10.26599/FRICT.2026.9441269

ABSTRACT: Molecular dynamics (MD) simulations are widely used to study tribological deformation at the atomic scale, but translating the resulting high-dimensional data into interpretable deformation pattern maps remains a largely manual and resource-intensive task. In this work, we present a data-driven workflow that aims to automate this process using unsupervised and supervised machine learning techniques. Grain-orientation-colored tomographic images from CuNi alloy simulations were first compressed into a 32-dimensional representation using an autoencoder, demonstrating that microstructural features can be retained under strong dimensionality reduction. The predictive model was trained independently using a dual-branch CNN–MLP architecture that combines image data with simulation metadata to classify dominant deformation patterns (transient and final states). The model achieves > 95% accuracy on spatially independent validation data. To assess generalization, entire simulations were excluded from training, yielding 100% final-state accuracy, as these conditions lie distant from regime boundaries. Evaluation on an artificial pressure–composition grid results in 87% accuracy, with misclassifications concentrated near regime boundaries. For transient states, the model achieves 68% accuracy on unseen simulations, increasing to > 92% when accounting for physically plausible misclassifications between adjacent regimes. A comparison with an image-only ResNet-50 model (~70% accuracy on spatially independent data, ~50% on unseen simulations) highlights the importance of integrating metadata. These results demonstrate that tribological deformation mechanisms can be automatically identified from microstructural data. This work provides a proof of concept for data-driven construction of deformation maps and represents a step toward predictive frameworks that may reduce reliance on large-scale MD simulations.



KEYWORDS: molecular dynamics (MD); plastic deformation; dual-branch neural network; explainable AI

1 Introduction

Friction and wear are major sources of energy loss and material degradation across engineering systems, contributing to roughly 23% of global energy consumption being dissipated in tribological contacts [1]. Understanding how materials deform under tribological loading is therefore essential for improving durability and energy efficiency. A wide range of modeling approaches has been developed to address tribological phenomena across multiple length and time scales, ranging from continuum and contact-mechanics models to atomistic simulations [2]. At the atomic scale, molecular dynamics (MD) simulations can provide detailed insights into the microstructural evolution and deformation patterns under controlled variations in composition, pressure, temperature, and sliding conditions [3]. Such simulations have been used to study a broad range of metals and alloys, including Al [4, 5], Ni [6], Fe [7, 8], Mo [9], or Cu [7, 10], among others. Such works highlight the roles of dislocation activity, stacking faults, and twinning in plastic deformation [5, 11], as well as the strong influence of grain size and alloy composition on the tribological response [12–14].

Besides the considerable computational expense of MD simulations, the significant manual effort required for interpreting large volumes of structural data is its major limitation. This challenge persists even when simulations avoid nanocrystalline artifacts associated with the Hall–Petch breakdown by considering sufficiently large grains [3, 15]. Existing automated tools, such as twin identification frameworks [16], address isolated tasks but do not provide a general method for identifying dominant deformation patterns across diverse tribological conditions.

[1] Institute for Engineering Design and Product Development, TU Wien, Lehárgasse 6–Objekt 7, Vienna 1060, Austria. [2] AC2T research GmbH, Viktor-Kaplan-Straße 2/C, Wiener Neustadt 2700, Austria. [3] School of Mechanical & Aerospace Engineering, Nanyang Technological University, 50 Nanyang Avenue, 639798, Singapore.

✉ Corresponding authors. E-mail: H. J. Ehrich, hendrik.ehrich@tuwien.ac.at; S. J. Eder, stefan.j.eder@tuwien.ac.at

Received: February 27, 2026; Revised: April 24, 2026; Accepted: May 20, 2026

© The Author(s) 2026. This is an open access article under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0, http://creativecommons.org/licenses/by/4.0/).

In parallel, machine learning (ML) has transformed many areas of materials research by enabling largescale screening, structure–property prediction, and microstructure optimization [17–20]. ML approaches have been applied to predict mechanical properties of complex alloys [21], optimize microstructures [22], and analyze heterogeneous experimental tribological data [23, 24]. Convolutional neural networks (CNNs) in particular offer powerful tools for interpreting high-dimensional image data and have been used to classify atomic environments [25], extract grain boundary features [26], and interpret diffraction patterns [27]. Although approaches using CNNs have attempted to classify wear mechanisms from scanning electron microscopy (SEM) images [28, 29], similar links to tribological molecular dynamics simulations remain largely unexplored. Existing efforts primarily target isolated structural descriptors or mechanical property predictions [30, 31], and no available method performs automated, mechanism-level classification of deformation modes in large-scale polycrystalline MD data. A recent study by Chen and Xu [32] comes closest in scope but focuses on identifying structural criticalities rather than classifying deformation regimes.

In this work, we report a proof-of-concept study [33] assessing to what extent data-driven ML models can assist in the interpretation of structural images from MD simulations of tribologically loaded polycrystals. Grain-orientation-colored computational tomographs are compressed using a deep autoencoder to explore low-dimensional structure representations. In parallel, a dual-branch CNN–MLP model is trained on raw images together with simulation metadata to predict dominant deformation mechanisms. Rather than presenting a finalized predictive tool, the emphasis lies on identifying robust trends, sensitivity to modeling choices, and current limitations.

This paper is structured as follows: we first briefly introduce how computational materials tribology based on MD simulations has been carried out over the past years, motivating the necessity for more resource efficient methods. We then walk the reader through our labeling process, the handling of the large amount of historic MD data, and an unsupervised learning approach. At the heart of this contribution lies a transparent discussion of the design choices and constraints involved in constructing a supervised prediction model, including a suitable strategy for splitting between training and validation data as well as an approach to assess the interpretability of the model's output. We then apply the introduced methods to a consistent data set consisting of computational tomographs of tribologically loaded CuNi alloy surfaces plus accompanying metadata. We discuss the ML model's performance and interpret its output, thus conducting sanity checks on the predictions. Finally, a broader outlook includes possible avenues of further work to make such models a part of engineering tools of the future.

2 Background and motivation

2.1 Conventional data processing workflow

Conventional MD analysis, which has been used in previous data evaluation [3, 34, 35] and can serve labeling and validation purposes for ML model construction and preparation, employs a multi-tiered data distillation process to derive robust trends, see Fig. 1. In the left column of this figure, we show representative examples of computational tomographs through the 3D MD model, with the atoms colored by (a) grain orientation in electron backscatter diffraction (EBSD) standard, (b) lattice type, grain boundaries, and defects, (c) advection (drift) velocity to visualize shearing, and (d) local stresses. As a first step in the data distillation process, these 3D data that are stored for each atom are averaged across the lateral system dimensions, revealing depth-resolved, time-dependent quantities of interest, as visualized in the heat map at the top of the middle column (e). Further elimination of the sample depth and time dimensions leads to time-resolved global quantities (f) and contact pressure dependent trends (g), which can be fitted with characteristic pressures that mark the transition between deformation patterns (h). As an outlook to the utility of such highly distilled data, we propose their incorporation into Ashby-style charts, as schematically shown in Fig. 1(i), which link material properties with tribological properties. This conventional approach accommodates the complexities of polycrystalline materials under tribological loading conditions and is guided by the underlying physics, resulting in this time-

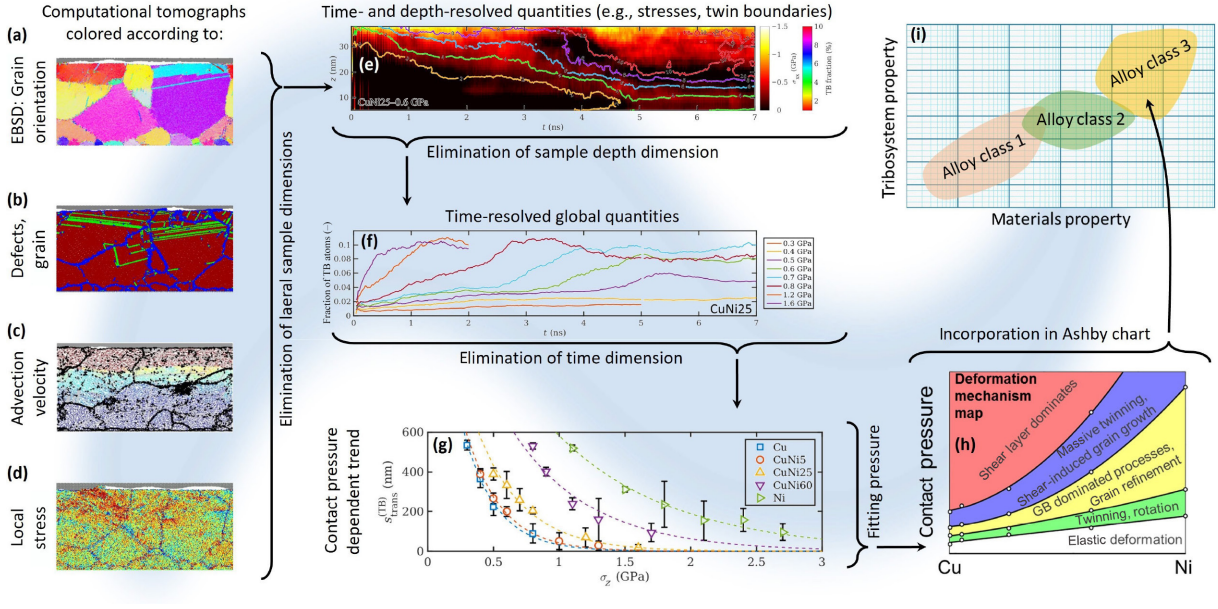


Fig. 1 Schematic of conventional MD data analysis. The data distillation workflow follows the tilde-shaped background from (a–d) to (i), based on the representative example of statistically evaluating the emergence of twin boundaries (TB) in the microstructure. A fully developed ML model could replace all the intermediate steps up to the deformation pattern map.

consuming procedure. This workflow also serves as an important reference point for assessing which aspects of physical insight can realistically be recovered by data-driven models. Thus, exploring whether parts of this workflow can be supported or accelerated by ML models is highly relevant. Figure 1 suggests that a fully developed ML model could eventually replace all intermediate steps up to the creation of a deformation pattern map, and the current study represents a foundational step toward this goal. At present, our methodology relies on a "conventionally obtained" deformation pattern map to provide necessary labels for training. Therefore, this work demonstrates automated classification based on an existing labeling framework rather than the autonomous generation of a new map for an unknown material system.

To delineate the boundaries between dominant deformation patterns in the generated deformation pattern map (Fig. 1(h)), equilibrium data from large-scale MD simulations were distilled into distinct regimes based on quantitative microstructural and kinematic thresholds, as described in detail in Ref. [3]. The boundary for *elastic deformation* is defined by the composition-dependent critical twinning stress ($\tau_{crit}^{(TB)}$). The transition to *grain refinement* was established by setting a grain boundary (GB) fraction threshold of 11%, representing non-fcc ordered atoms not part of a twin boundary, which effectively distinguishes refined grains (diameter < 18 nm) from the bulk aggregate. For the *grain growth* pattern, an atomic advection velocity threshold of 10 m/s was utilized to define the shear layer thickness (d_{shear}). The pattern transition is drawn where d_{shear} exceeds 2 nm, a value chosen to filter out localized influences from surface asperities. Finally, a transition to *shear layer domination* is identified when the shear layer thickness surpasses the depth of the previously formed refined grain layer ($d_{refined}$). These discrete transition points were interconnected using first- and second-order polynomial fits to produce the final continuous boundaries. At elevated temperatures [34] or extreme sliding velocities [35], these boundaries shift as thermal softening and phase transitions alter the energy barriers for dislocation activity and shear localization. The same quantities used to define these deformation regimes are also employed for time-resolved state classification, as summarized in Table 1.

2.2 Computational cost

Each MD simulation run of 7 ns of sliding for a 25M atom system typically produces some 200 GB of raw compressed data and requires 160k CPU-h for both computation and postprocessing. There are approximately 60 data points available for 5 CuNi alloys at two temperatures and various contact pressures. In addition, there are approximately 20 more points available for 4 CuNi alloys at various sliding speeds. This leads to some 16 TB of raw data that took around 13M CPU-h to calculate, which corresponds to a wall time of approximately one year at a job size of 1,440 CPU cores. It therefore becomes clear that the computational expense for a reasonable parametric sweep for one single alloy class is considerable. Another bottleneck is the interpretative step, which requires human resources to evaluate the visualizations and guide the data distillation. This motivates the search for automated methods that can on one hand reduce the volume of necessary calculations and on the other hand dramatically cut corners during the laborious interpretation.

3 Methodology overview

3.1 MD data and system description

The dataset used in this work consists of CuNi alloy simulation snapshots (Fig. 2) with varying Ni content (0%, 5%, 25%, 60%, and 100%), normal pressure (ranging from 0.1 to 2.7 GPa depending on composition), see Fig. 3(a), with constant temperature (300 K) and sliding speed (80 m/s). The initial microstructures were generated using a Voronoi construction where grain nodes are randomized within a simulation box and then padded with 26 identical copies to create a 3×3×3 superbox, ensuring perfect 3D periodicity [36].

These resulting polyhedra are filled with the specific crystal lattices (fcc Cu with random Ni replacement according to the

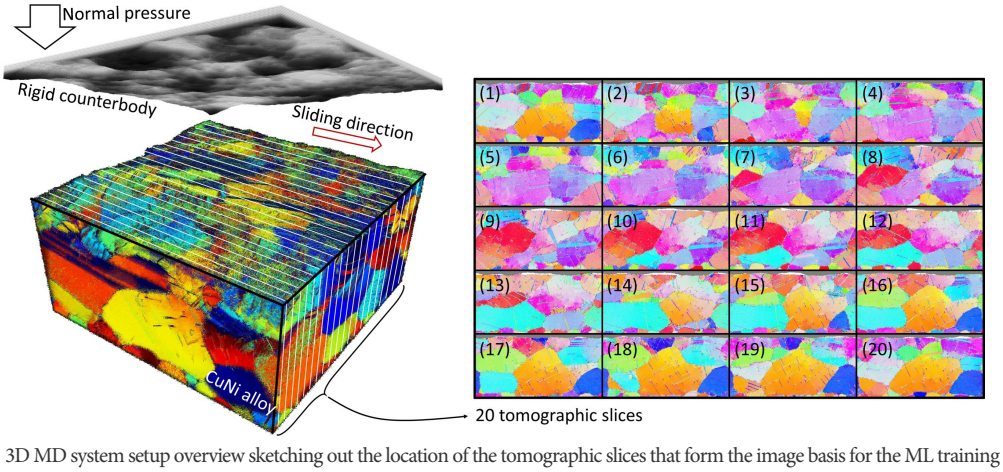


Fig. 2 3D MD system setup overview sketching out the location of the tomographic slices that form the image basis for the ML training data.

Table 1 Thresholds of calculated system quantities used for transient state label assignment. As described in the text, triggering thresholds may overlap and require further refinement of the labeling procedure. To avoid unphysical temporal sequences (e.g., shear-dominated states preceding twinning), trigger times were manually defined and, where necessary, adjusted to prevent overlap. This may lead to deviations from the strict thresholds reported in this table

System quantity	Elastic	Twinning	GB dominated	Shear dominated
Fraction of twin boundary atoms	< 2% (constant)	> 6% (distinct increase)	—	—
Fraction of grain boundary atoms	< 10% (constant)	< 10% (constant)	> 12% (distinct increase)	—
Shear layer thickness	< 10 nm (constant)	< 10 nm (constant)	< 10 nm (constant)	> 10 nm (distinct increase)

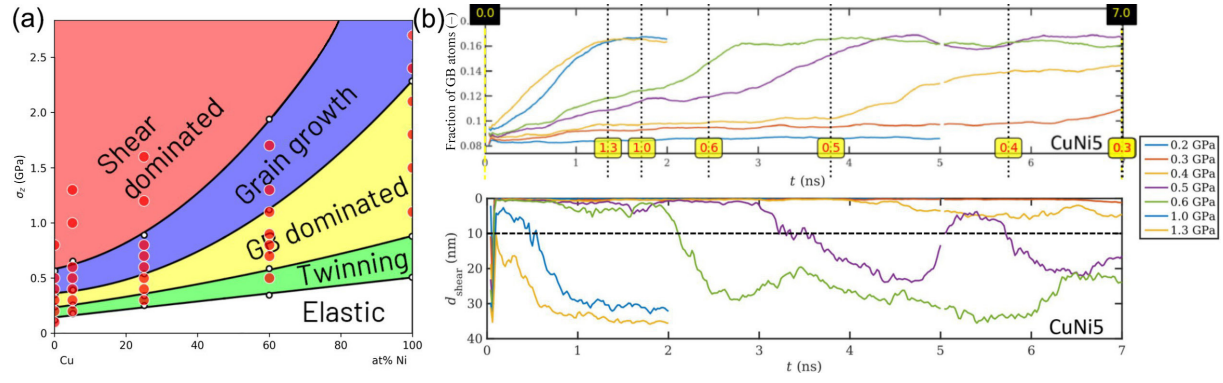


Fig. 3 (a) Deformation pattern map as a function of Ni fraction and normal pressure, with the deformation regimes that serve as labels for the final states and the parameter combinations of the MD simulations marked as red circles. (b) Exemplary intermediate plots of relevant time-dependent quantities (top: grain boundary fraction, transition times labeled; bottom: shear layer depth, no distinct transition times) used to assign labels to the transient states. Reproduced with permission from Ref. [3] for (a), © American Chemical Society 2020.

desired Ni content) that are assigned random orientations based on Euler angle triplets [37]. Specific care was taken during the construction process to avoid an excessive occurrence of low-angle grain boundaries, ensuring a representative distribution of high-angle interfaces. For the specific microstructure relevant here, two separate Voronoi tessellations were employed: one for coarse surface grains (mean grain diameter 40 nm) and the other for a "sacrificial base layer" of fine grains (mean grain diameter 10 nm), which are joined by a rough fractal interface with an RMS roughness of 1.5 nm to prevent the formation of an artificial sliding interface between the two layers [15]. While this method provided the historic data currently being analyzed, it is worth noting that nowadays, we typically use more advanced tools such as Dream.3D [38] to generate even more sophisticated and realistic microstructures. Each system is then represented as a series of computational tomographs sliced normal to the y-direction (with x being the direction of sliding and the sliding surface being normal to the z direction), each carrying associated metadata (Ni content, normal pressure, temperature, time step, and slice position).

3.2 Labels

In complex tribological interfaces, multiple deformation patterns often coexist. However, the conventional workflow, and thus, our labeling framework focuses on identifying the dominant deformation pattern observed in the MD runs. This focus is specifically motivated by tribological engineering, in that critical operation points to be avoided in practical applications can be determined by identifying when the most destructive dominant mechanisms occur.

Initially, labels for each simulation run were derived directly from the system conditions, specifically alloy composition and applied normal pressure, which define the axes of the deformation pattern map previously established through manual analysis [3]. In this approach, each simulation corresponded to a single point in the deformation pattern map, such as shear-dominated or elastic, resulting in a straightforward labeling procedure of the final outcome of a single simulation run (see Fig. 3(a)). However, this method underutilized the available data, as each label encompassed multiple simulations with thousands of time-resolved samples, many of which corresponded to early stages of the simulation in which the microstructure did not yet resemble the final deformation pattern. Consequently, the image data associated with each label were highly heterogeneous, and first models trained with these labels tended to rely predominantly on metadata rather than visual information.

To address this limitation, transient state labels derived from interim results of the manual workflow were incorporated. These intermediate metrics, such as the evolution of the shear layer depth or the twin boundary and stacking fault fraction over time, provide insight into the current structural state of the simulation rather than only the final outcome. Transient labels were assigned based on thresholds in these metrics (see Fig. 3(b)(top)). The available data supported reliable labeling only for the elastic, twinning, grain-boundary-dominated, and shear-dominated regimes, which is why the transient-state classifier was restricted to these four classes. This approach enables the model to associate images with the structural state they depict, ensuring that visual information is meaningfully incorporated into training and improving the model's capacity to capture the temporal evolution of deformation patterns. To provide a transparent overview of these classifications, Table 1 summarizes the specific physical quantities used for assignment (such as grain boundary fraction or shear layer depth), and the corresponding thresholds employed in the manual labeling workflow. While final state labels are based on several distillation steps in the conventional workflow [3], as summarized in Section 2.1, transient state labels are directly derived from intermediate quantities of the systems like fraction of twin boundary atoms.

Many of the metrics used to derive transient states, such as twin boundary and grain boundary fractions, are often correlated, meaning that a single time threshold may be insufficient to capture the true evolution of the deformation pattern. Furthermore, the labeling process does not account for boundary states, where the microstructure exhibits characteristics of multiple regimes simultaneously (see Fig. 3(b)). To address these issues, transient labels could be subdivided based on the evolution of key quantities, such as shear layer depth, and assigned throughout the simulation in a time-resolved manner. This assignment should also take the final state into account, ensuring physically consistent labeling (e.g., an ultimately elastic simulation would not include a shear-dominated transient state at any point). This approach enables a more accurate and consistent representation of the evolving structural state in the images, and is pursued in ongoing work.

3.3 Dataset curation

The available simulation data provided nearly 200,000 tomograph–metadata samples. However, due to imbalances originating from both the simulation design and the structure of the labeling scheme, only a balanced subset of approximately 12,000 samples can ultimately be used for model training and

evaluation (e.g., only one simulation was explicitly situated in the elastic regime, so a lot less data were available for this scenario). This balancing ensures that all transient–final label combinations were represented sufficiently often, which is essential because the goal of the model is to learn characteristic microstructural signatures, not the arbitrary frequency of deformation patterns in the simulation campaign.

Balancing joint labels in this way, however, introduces a secondary effect. When the transient and final labels are considered independently, their individual distributions are no longer balanced. Some classes naturally occur in many more transient–final combinations than others, so the balancing by joint labels implicitly oversamples classes with fewer combinations and undersamples those with many.

The practical relevance of this imbalance is limited, but worth noting. Since the learning target is the class-specific structure, not class frequency, the model does not use marginal frequency as a cue for prediction, and balancing does not distort the physical interpretation of the classes. The model is thus exposed to all classes sufficiently often, which is the intended effect. The main consequence is confined to evaluation rather than training, in that per-class accuracies for transient and final labels may not reflect the class proportions in the full simulation dataset. Instead, they reflect the intentionally equalized joint-label sampling strategy.

The simulation campaign itself did not sample all compositions under loading conditions that would lead to every deformation pattern at least once. For example, an elastic final state was only obtained for simulations of pure copper (see Fig. 3) for a visualization of the simulation conditions, whereas other alloy compositions consistently evolved toward shear-dominated or grain-boundary-dominated behavior. This imbalance does not affect the image input branch, because atomic species are not visually distinguishable in the tomographic representations. It does, however, introduce bias in the metadata branch, as certain metadata combinations, such as pure copper under low pressure, serve as strong indicators of specific final states, in this case the elastic regime. Transient-state labeling is less influenced by these simulation-design constraints, since system evolution is tracked meticulously, and sufficient transient samples were available for each class.

3.4 Unsupervised representation learning

Autoencoders are ML models trained to reconstruct the input as accurately as possible by learning an encoder–decoder mapping [39]. Concretely, an autoencoder consists of an encoder $f_\theta(x) = y$ that maps high dimensional, original inputs x to a typically lower dimensional latent output y and a decoder $g_\delta(y) = \hat{x}$. Given the model structures f and g, both parameter sets θ and δ are adapted to reduce the reconstruction error $L = ||x-\hat{x}||$. Specifically, in an undercomplete autoencoder, this mapping f compresses the input into a latent representation with a much lower dimensionality than the original data, in short $|x| >> |y|$ [40]. Thus, in this case in order to investigate the compressibility of high-dimensional MD tomographic images and assess whether essential structural information could be retained, we employed an undercomplete autoencoder with a 32-dimensional bottleneck layer using the mean squared error (MSE) as the loss function, ensuring that omissions of small but relevant defects were penalized during training. Unlike most experimental machine learning workflows for deformation pattern classification, where defects such as twins or stacking faults are manually annotated [29], the present approach aims for the model to discover and utilize relevant structural features autonomously.

This approach aimed at reducing the images to a compact latent representation without requiring manual annotation of features such as twin boundaries, stacking faults, or phase transformations, as is often necessary in experimental transmission electron microscopy (TEM) or scanning electron microscopy (SEM) datasets. In preliminary experiments, deeper architectures were evaluated, but produced visibly blurrier reconstructions, leading to the loss of fine structural features. As a result, we adopted a shallower architecture as visualized in Fig. 4. Although this design led to a higher parameter count and a more abrupt transition from spatial representations to the flattened bottleneck, it preserved important local features more effectively.

Ultimately, the autoencoder was not employed as a preprocessing step for the subsequent architectures, as it has not yet been fully validated and its suitability for generating robust latent representations remains uncertain. This outcome underscores that representation learning from atomistic images is highly sensitive to architectural choices and loss functions, and should not be assumed transferable without careful validation.

3.5 Supervised prediction model

A ResNet-50 [41] model was trained on the computational tomographs only to predict both final and transient states, to explore the performance of well-established architectures in the domain at hand, see Fig. 5(a). The pretrained model's weights (from ImageNet) were frozen, and the top layer was replaced with a global average pooling layer followed by a dense layer and a dropout leading into the two categorization end layers of final and transient state.

Independent of previously explored architectures, we trained a dual-branch neural network [42], consisting of two branches for image and metadata, also using paired transient and final state labels. The architecture separately processes image data and the

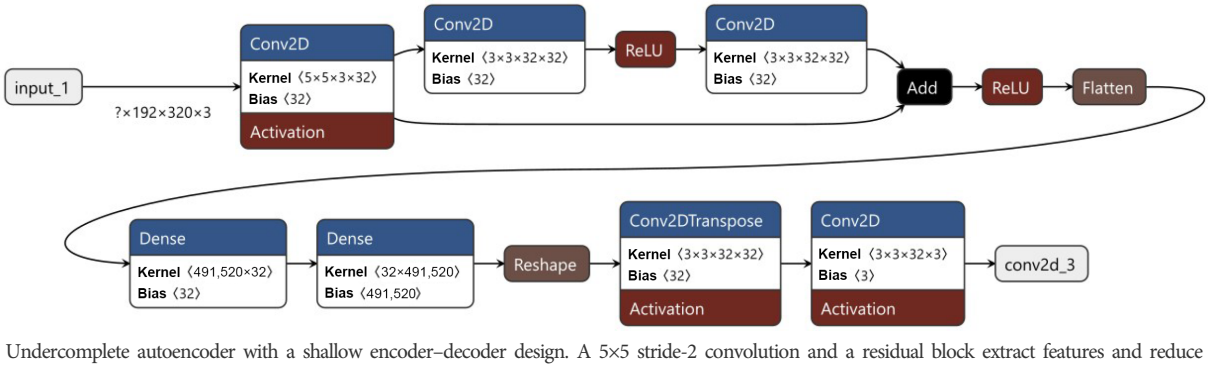


Fig. 4 Undercomplete autoencoder with a shallow encoder–decoder design. A 5×5 stride-2 convolution and a residual block extract features and reduce spatial resolution before flattening into a 32-dimensional latent space. The decoder expands the latent vector via a dense projection, reshaping, and a stride-2 transposed convolution, followed by a final 3×3 sigmoid layer to reconstruct the image.

清華大学出版社 Tsinghua University Press | SciOpen

simulation metadata, in numerical form, before combining both modalities in a shared latent space as visualized in Fig. 5(b).

More specifically, the concrete architecture was selected following the state-of-the-art approach of Ref. [43] as dual-branch neural network architecture optimization is notoriously hard to optimize with automated search architectures or AutoML [44], especially if different modalities are regarded. The resulting architecture is visualized in Fig. 6.

The image branch consists of three convolutional blocks with 32, 64, and 128 filters, each using 3×3 kernels and ReLU activations, followed by max-pooling layers. To reduce computational cost while preserving structural information, the input images were downscaled by a factor of five, keeping the original aspect ratio, resulting in an input size of 294 pixels × 146 pixels. A global average pooling layer preserves spatial information while drastically reducing the number of parameters, producing a feature vector that passes through a 128-unit dense layer with ReLU activation, dropout (0.3), and layer normalization.

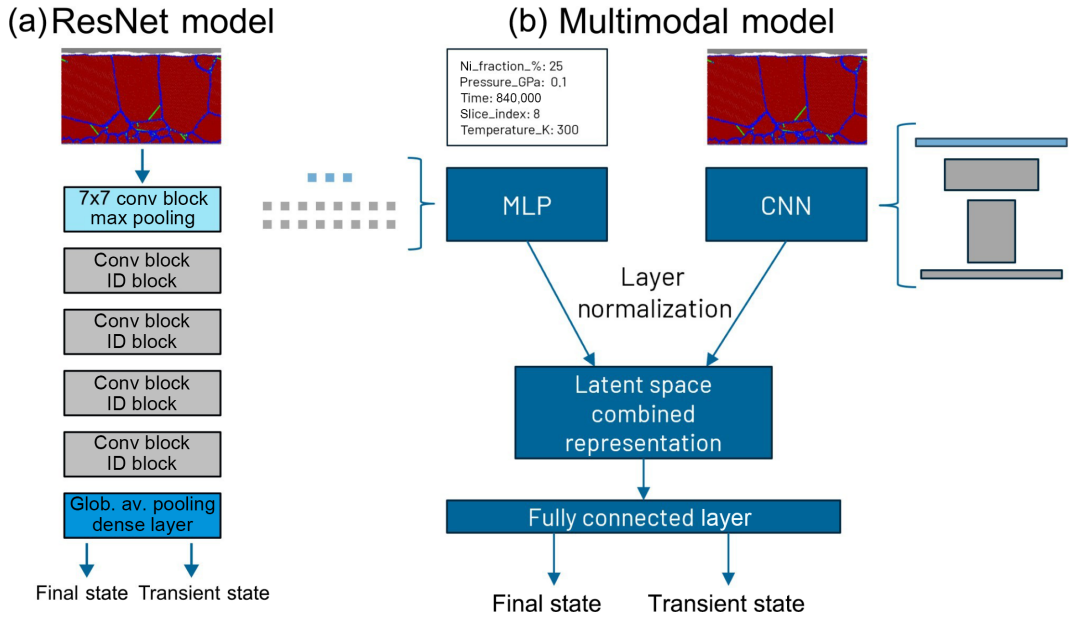


Fig. 5 Architectures used for supervised training. Both models predict two labels, the transient and final state (4-class and 5-class). (a) The ResNet-50 model is a standard backbone for computer vision tasks. (b) Our own two-branch neural network architecture combines numeric and categorical metadata and image data to predict the classes. Image data are processed through a convolutional neural network, while metadata are handled via a multi-layer perceptron (MLP). The outputs of both branches are concatenated and passed through additional fully connected layers to produce probabilistic predictions for both labels.

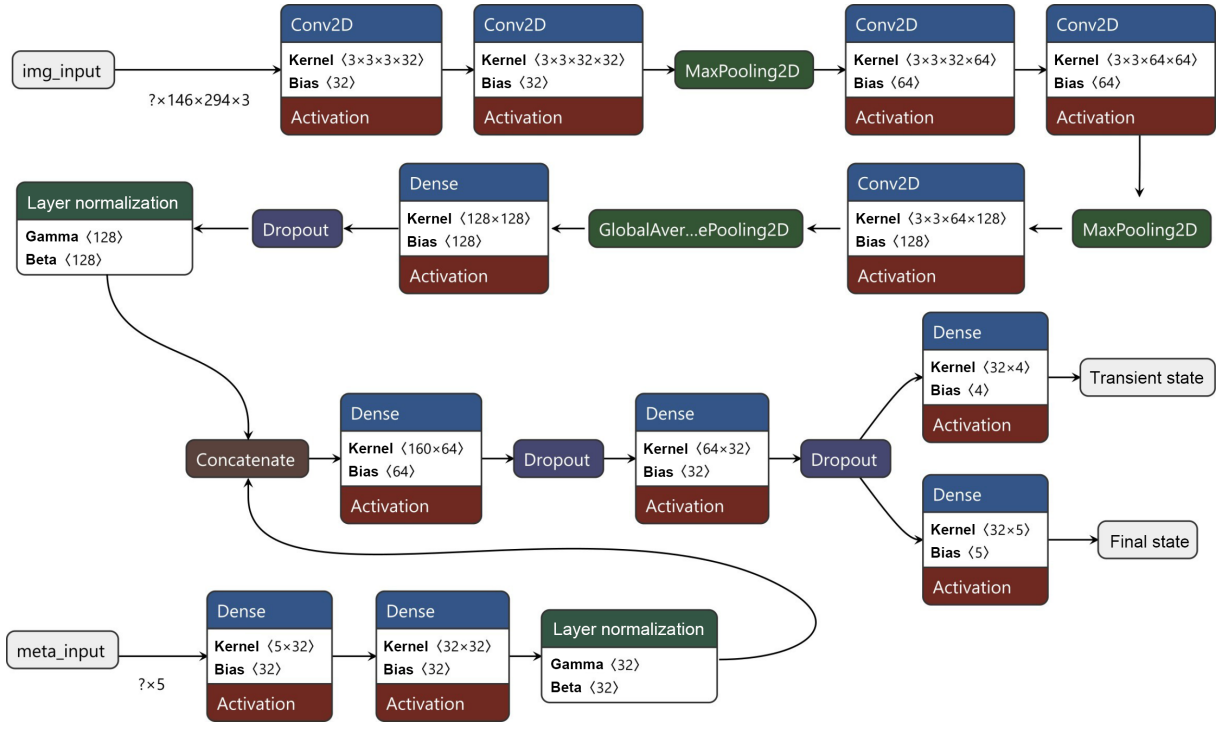


Fig. 6 Detailed dual-branch neural network architecture. The image branch (top) applies successive convolutional and dropout layers before flattening and projecting the features to a dimensionality comparable to that of the metadata branch (bottom), which consists of fully connected layers. The two feature streams are concatenated into a shared representation that is processed by additional fully connected layers to jointly predict the transient and final deformation labels.

On the other hand, the metadata branch comprises two fully connected layers of 32 units each, also followed by layer normalization. At this stage, metadata are sparse as only a single simulation configuration is available, varying pressure and alloy composition (Ni fraction), but with constant temperature and sliding speed. As a result the corresponding weights are nilled, the overall architecture however shall be applicable to various simulation configurations, each with varying temperature and sliding speed. Thus, the metadata inputs consist of Ni fraction, normal pressure, timestep, temperature, and slice index (analogous to spatial (y-)position within the system).

Outputs from both branches are concatenated and passed through two additional dense layers (64 and 32 units, ReLU activations) with dropout regularization following best practice [43]. The network produces two outputs, one, via a five-class softmax for the final deformation state and the other one via a four-class softmax for the transient state. Softmax is applied due to superior classification capabilities in earlier studies [45]. The resulting model contains approximately 170k parameters, largely due to the parameter reduction achieved by the global average pooling layer. Training was conducted for 60 epochs with a batch size of 32, using the Adam optimizer [46] and categorical cross-entropy as the loss function [47].

3.6 Revised validation strategy

Neighboring tomographs exhibited strong microstructural similarity due to their spatial proximity within the simulated volume. Consequently, a random train–validation split would result in both sets containing highly similar microstructures, leading to artificially inflated evaluation metrics. In such a case, the model would effectively have been exposed to each microstructural configuration at least once, even if at a different simulation time, thereby testing recognition rather than generalization. To more rigorously assess the model's ability to generalize, contiguous tomographic regions were excluded from training. This strategy mimics evaluation on unseen microstructures, similar to testing on newly generated simulations with independently initialized microstructures. Strong performance under these conditions indicates that the model learned generalizable representations of microstructural features rather than memorizing specific image patterns [48].

For this purpose, the five most distinct tomographs from the dataset were selected for training, with one tomograph on either side reserved as validation data (see Fig. 7(b)) to provide a controlled but structurally distinct test set. To stress-test the model, we also performed an extreme experiment in which the model was trained on only three tomographs (two maximally different positions plus one additional slice adjacent to one of those positions, see Fig. 7(c)) and validated on tomographs maximally different from the training data.

To address potential information-leakage concerns where the model might learn "shortcut mappings" between specific metadata and labels, we further extended the validation strategy. In addition to the spatial splits, we implemented a simulation-level exclusion strategy where entire pressure–composition combinations were withheld from the training phase (see Fig. S5 in the Electronic Supplementary Material (ESM)). Furthermore, we evaluated the model against artificially generated final states to test the boundaries of its learned physical representations (see Fig. S4 in the ESM).

3.7 Model interpretation

Model interpretability is achieved using SHAP (Shapley additive explanations) [49], which attributes a quantitative contribution to each input feature for a given prediction. SHAP is based on Shapley values from cooperative game theory, where each feature is treated as a "player" contributing information to the model's output.

For a target sample, SHAP evaluates the model repeatedly on combinations of feature subsets constructed from representative background data. Each feature's true value is paired with background values for the remaining features, and the resulting change in model prediction is compared to the case where that feature is replaced by background values. Averaging these prediction differences across many feature subsets yields the SHAP value for that feature, representing its average marginal contribution to the prediction relative to the dataset baseline [49].

The model output for each sample can thus be expressed as

$$f(x) = E[f(X)] + \sum_j \phi_j(x)$$

where $E[f(X)]$ is the expected model output over the dataset, and $\phi_j(x)$ is the SHAP value quantifying the contribution of feature j. Positive SHAP values indicate that a feature increases the model's evidence toward a specific class, while negative values indicate the opposite. When computed on the model's logit outputs, these contributions are additive and directly comparable across classes [50].

4 Results

4.1 Feature representation and reconstruction

The autoencoder architecture introduced in Section 3.4 was

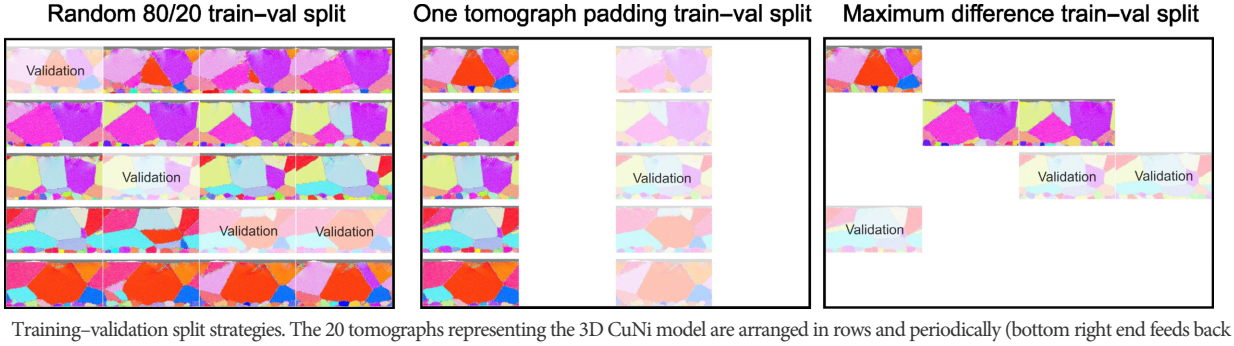


Fig. 7 Training–validation split strategies. The 20 tomographs representing the 3D CuNi model are arranged in rows and periodically (bottom right end feeds back into the top left). (a) Classic random splitting of the data into 80% for training and 20% for validation. (b) "One tomograph padding" strategy, where 50% of the data are selectively discarded to reduce self-similarity, and the validation dataset is interleaved between the training set with only little overlap by similarity. (c) Maximum difference split, where 70% of data were discarded and tomographs for training and validation were selected manually so that there is the least possible overlap (worst-case scenario).

清華大学出版社 Tsinghua University Press | SciOpen

trained on two types of tomographic images: EBSD-colored tomographs encoding grain orientation on the one hand, and common neighbor analysis (CNA)-colored tomographs, which encode local structural classifications (fcc, hcp, bcc, undefined, and areas with advection velocity in x direction of over 10 m/s), on the other. These representations emphasize microstructural features differently and therefore probe different aspects of compressibility.

In the CNA-colored images, the dominant features are grain boundaries and dislocations (in blue), and extended defects such as stacking faults and twins (in green), which appear as sharp, localized deviations within otherwise uniformly colored crystalline grains (in red). The autoencoder reconstructed grain boundaries reliably as evidenced in Fig. 8, indicating that the global morphology of the microstructure was well captured in the low-dimensional latent space. However, isolated defects such as single stacking faults were frequently omitted in the reconstructions. This behavior was reflected directly in the training curves as the MSE loss plateaued after a moderate number of epochs, consistent with the pixel-wise loss function penalizing small, spatially sparse defects that the network could not reproduce faithfully. Because the loss is dominated by the large, homogeneous grain interiors, these small defects had negligible influence on optimization and remained unreconstructed.

By contrast, the EBSD-colored tomographs, despite visually containing more complex color variations and similarly fine details, were reconstructed with much higher accuracy. Even localized defects, including individual stacking faults, were retained in the reconstructed images. This difference indicates that the orientation-based representation distributes microstructural information more smoothly across the image domain, enabling the autoencoder to encode subtle rotational variations and defect signatures more effectively within the limited available latent dimensionality.

Although the autoencoder confirmed that the tomographic data can be compressed to a low-dimensional representation, the simplicity of the final architecture imposes constraints on the generality of the learned latent space. The encoder employs a shallow architecture with few convolutional operations before projection into the latent space. This design performs an aggressive dimensionality reduction after only minimal hierarchical feature extraction. Consequently, the latent variables predominantly reflect the specific filtering characteristics of this shallow encoder rather than a broadly structured feature hierarchy. As previously mentioned, deeper encoder variants with additional convolutional stages were also evaluated, but these models did not achieve comparable reconstruction performance. Their training usually failed to converge towards a reconstruction quality that preserved the relevant microstructural detail. This outcome indicates that the apparent architectural narrowness of the final model is not merely a design choice, but a practical constraint imposed by the data characteristics, the limited variability in the training set, and the pixel-wise MSE objective, which harshly penalizes even small misalignments.

As a result, the learned latent space should be interpreted primarily as a task-specific compression tailored to this particular architecture and reconstruction objective, rather than as a general-purpose representation suitable for downstream transfer. Because the flattening step collapses spatial organization abruptly, the model has no mechanism to form intermediate, physically interpretable feature hierarchies that would support cross-task reuse.

4.2 Prediction performance

Both the ResNet-based and the dual-branch-based deformation pattern classification models were trained using the three training–validation splits introduced in Section 3.6.

During training, the ResNet model performs better on the 1 padding split and maxdiff. This can be attributed to the reduced diversity compared to the 80/20 split where the model sees all various tomograph positions. On validation, however, reduced amount of training data and increased heterogeneity between training and validation set lead to lower accuracies for both state predictions (see Fig. 9(c)).

The observations on the multimodal model training results summarized below hold for all splits, with differences only in degree. Performance variations arising from the specific dataset configurations are examined in the following section. The trained dual-branch model demonstrated strong generalization performance with no indications of overfitting. Prediction accuracy for final deformation states exceeded that for transient states. This is expected, as the labels (deformation pattern defined by the Ni fraction and pressure) are constant for each simulation, so the combination of final microstructure and consistent metadata appears frequently and is therefore easier for the model to learn. By contrast, transient states exhibit greater structural variability and weaker coupling to the metadata, making them inherently more difficult to classify.

Validation accuracy occasionally exceeded training accuracy

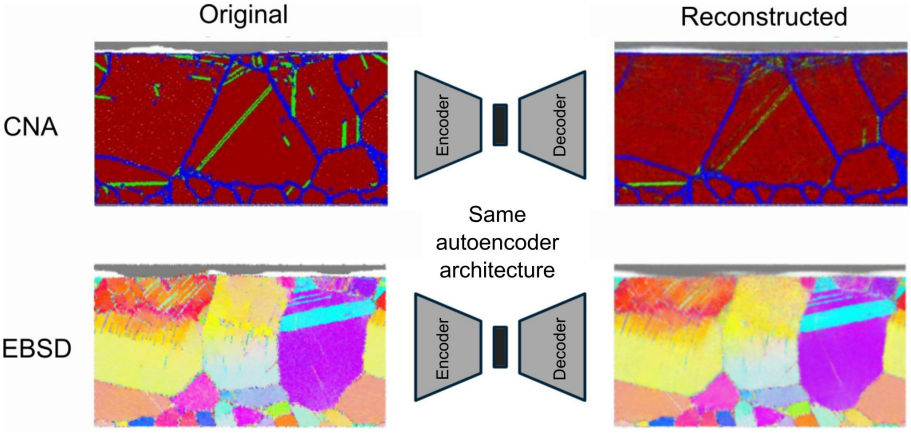


Fig. 8 Autoencoder reconstructions of MD tomographs highlighting microstructure (CNA) and grain orientation (EBSD). The rigid counterbody is shown in gray. Separate models were trained on each representation to compare reconstruction performance. For CNA representations, grain boundaries are accurately preserved, while small-scale defects like stacking faults are often lost during reconstruction. Dense defect regions show reduced fidelity. EBSD images retain fine structural details.

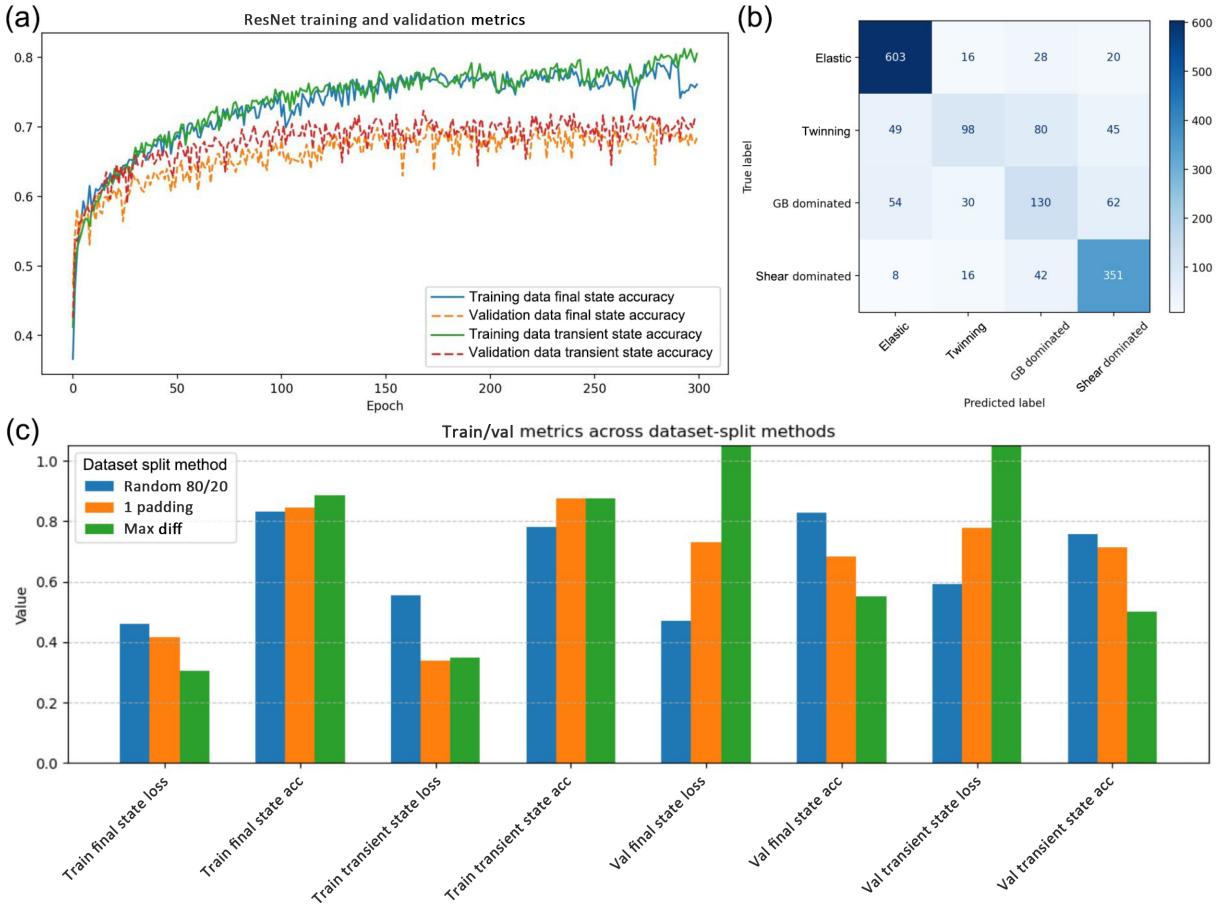


Fig. 9 ResNet training and validation metrics. (a) Training and validation metrics over 300 epochs of training, stagnancy of validation metrics is observable after 150 epochs with the continuing trend of training stats suggesting possible overfitting. (b) Transient state prediction confusion matrix of the "1 padding" split. (c) Comparison of losses and accuracies between the three different dataset split methods.

due to dropout being active only during training. When the model was re-evaluated on both training and validation datasets under identical inference conditions (dropout disabled), their performance was nearly identical, confirming that the network is not overfitting and that the apparent discrepancy arises solely from the regularization behavior of dropout (see Fig. 10) [51].

4.3 Generalization across spatial regions

Using the spatially independent validation strategy introduced in Section 3.6, the model's ability to generalize beyond microstructures it had been exposed to during training was tested. Because of the grain sizes and periodic boundaries in our simulations, it is nearly impossible to find tomographs that do not share grains with others. However, shifting roughly five slices left or right typically produces a visibly different scene with distinct grain morphologies satisfying the suggestions for human in the loop for image-based machine learning [48]. Although this approach naturally reduces the apparent accuracy compared to a random train–validation split, since large contiguous regions of similar tomographs are withheld, it provides a more realistic measure of true generalization. Even under these stricter conditions, the model achieved ~99% accuracy for final-state prediction and > 95% for transient-state prediction, demonstrating that it learned robust, transferable representations rather than relying on microstructural similarity between neighboring slices. Even in the highly constrained setting of only 3 images for training ("Max diff"), the model retained strong predictive power, achieving ~95% accuracy for final-state predictions and ~85% accuracy for transient-state predictions on the validation set. These results indicate that the network learns transferable microstructural representations and can generalize to structurally distinct regions even under reduced training data. Because excluding contiguous regions substantially reduces the amount of available training data, the model is effectively evaluated under information-limited conditions, yet still performs comparably to random-split experiments.

The extended validation of final state prediction accuracy on artificial unseen pressure and composition combinations yields a final-state prediction accuracy of 87%. Misclassifications occurred primarily at regime boundaries and were frequently accompanied by increased prediction uncertainty (compare Fig. S4 in the ESM). A cluster of mispredictions in the region of high Ni content and low pressure can be attributed to the absence of training data in this regime.

Retraining the model using a dataset split with a 1-padding tomograph filter and excluding three simulations entirely for validation resulted in 100% final state accuracy, as the excluded simulations correspond to conditions further away from regime boundaries (see Fig. S5 in the ESM). Applying this strategy to transient state prediction initially led to a notable drop in accuracy (~55%), indicating that the model had partially relied on correlations between pressure, composition, and simulation time. To address this, the model was retrained without including time as a metadata input. When evaluated on unseen simulations, the model showed improved generalization (~68%), with most misclassifications occurring between adjacent, physically related regimes. To better capture this behavior, two additional metrics

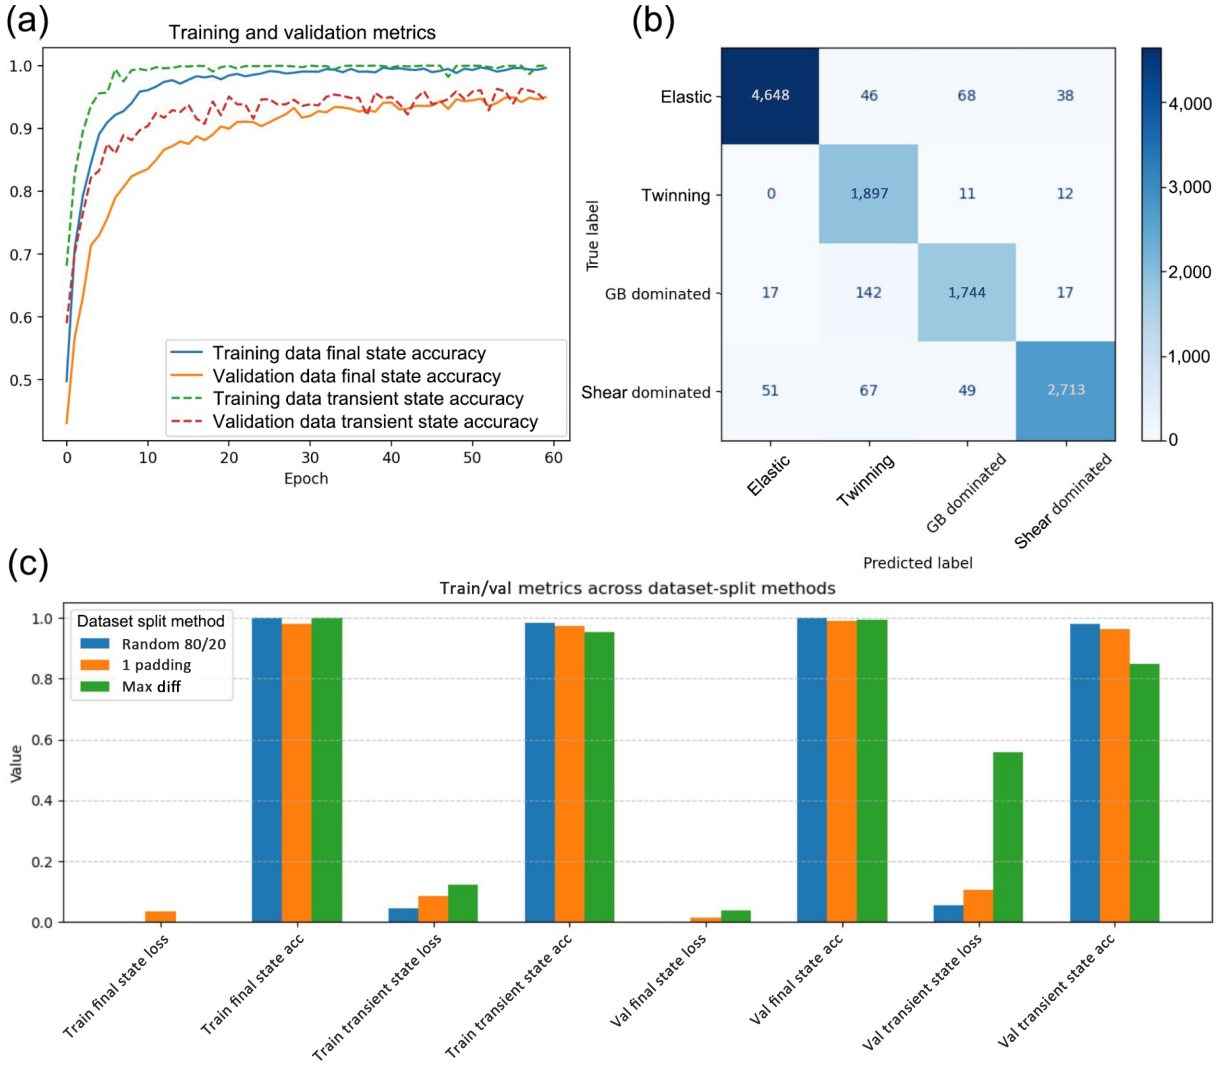


Fig. 10 Dual-branch model training and performance metrics. (a) Accuracies of training and validation data for final and transient state predictions during training with the "1 padding" train–val split. (b) Confusion matrix for true and predicted labels for the transient states ("1 padding" train–val split). (c) Comparison of losses and accuracies between the three different train–val split approaches.

were introduced, a top-2 categorical accuracy and a custom "off-by-one" accuracy. The latter considers a prediction correct if it falls within one adjacent class of the true label, while the former counts a prediction as correct if the true class corresponds to the second-highest predicted probability. These metrics reach ~92% (off-by-one) and ~98% (top-2), respectively, indicating that misclassifications predominantly occur between neighboring regimes and that the correct class is typically among the top predicted probabilities, consistent with boundary cases and underscoring the model's understanding of the relationships between classes and the underlying structure of the data.

Validating the ResNet model on excluded simulations results in a substantial drop in prediction accuracy to approximately 75% for final states and 50% for transient states (see Fig. S6 in the ESM).

5 Discussion

Comparing the performance of the classical ResNet architecture with the proposed multimodal model highlights the importance of combining image data with metadata to provide additional context. The image-only model performs consistently worse across all dataset splits. This is expected for final state prediction, as these labels are directly derived from metadata. For transient state prediction, the advantage of the multimodal approach becomes particularly evident, yielding consistently higher accuracies by leveraging complementary information from both input sources. Nevertheless, the results suggest that integrating a ResNet as the image-processing branch within a dual-branch architecture could further improve performance and should be explored in future work.

The relative importance of image and metadata inputs was evaluated through an ablation study. For final-state predictions, accuracy decreased sharply when only image data were used, whereas removal of the image branch had minimal effect, highlighting the dominant role of metadata for final state prediction. For transient-state predictions, removal of either images or metadata resulted in an approximately 40% decrease in accuracy, indicating comparable contributions from both input types. These results indicate that predictions of transient deformation evolution rely on both structural information presented in the form of images and contextual information from metadata.

When interpreting these behaviors, it is important to note again that the model learns from *paired* transient–final labels, since both outputs are predicted jointly. This pairing inherently introduces

uneven sampling across the combined label space. Rare combinations, such as a grain-growth final state paired with a transient twinning state (a configuration that may occur only briefly during evolution), became strongly underrepresented. Ideally, the MD simulation matrix would be expanded to ensure full coverage and balanced occurrence of all deformation patterns across compositions and loading conditions. However, achieving such balance is the subject of future work, as the required simulations are computationally intensive and beyond the scope of the present study.

Results indicate that final-state predictions are heavily dominated by metadata, especially pressure and alloy composition. From an engineering perspective, this is consistent with the fact that the final dominant pattern is physically governed by these boundary conditions. However, the image branch remains essential for the transient case, where structural evolution, rather than just initial metadata, is the primary indicator of the active deformation mechanism.

The observation that misclassifications are concentrated at regime boundaries is consistent with the inherent physical ambiguity within the system and the manual effort of defining sharp transitions in deformation pattern evolution. To illustrate this, we provide a comparison in Fig. S7 in the ESM between the fitted regime boundaries in the underlying study by Eder et al. [3] and critical twinning stress values derived from density functional theory (DFT) [52, 53]. This comparison reveals that the physical quantities used to derive these boundaries exhibit error bars significantly larger than the delineating lines themselves. Consequently, it is physically plausible, and even expected, that a model predicting the instantaneous state from image and metadata will occasionally assign a sample to a neighboring class at these transition points. In fact this is similar to human evaluators and given their preference and uncertainty being part of the training data, in fact it is expected that model inherits this uncertainty in allocation.

Metadata SHAP values for each input parameter and deformation class were analyzed using class-specific beeswarm plots. For clarity, the main manuscript focuses on two representative transient-state classes, which illustrate the characteristic and physically interpretable metadata trends (Fig. 11). The corresponding beeswarm plots for the remaining transient and all final-state classes are provided in Fig. S1 in the ESM and are discussed below. For the transient classes, the extreme deformation states, purely elastic and fully shear-dominated, suggested clear, interpretable trends. Low pressure and high Ni fraction increased the probability of the elastic transient state, while progression in time decreased it as other deformation patterns emerged. Conversely, high pressure, low Ni fraction, and later time stamps consistently increased the probability of the shear-dominated transient state. These trends align with the labeling system used to assign the ground-truth labels. The intermediate transient states, twinning and grain-boundary dominated, suggested substantially weaker and less consistent metadata influence. Here, none of the metadata parameters exhibited a systematic monotonic effect, and overall SHAP magnitudes were smaller than in the elastic or shear-dominated cases.

Counterintuitively, the final-state predictions exhibited even less clear metadata influence in the SHAP analysis, despite ablation studies indicating that the model relied almost exclusively on metadata for these classifications. This apparent contradiction can be explained by the labeling procedure. Multiple combinations of pressure and Ni fraction can lead to the same final state, so no single metadata trend dominates within a class. Hence, fractional changes as in the SHAP analysis have little effect. For the elastic final state, low pressure generally increased class probability, while Ni-fraction effects were mixed but included clear low-Ni cases with higher elastic likelihood, which results from the condition imbalance mentioned in Section 3.3.

For the twinning final state, metadata SHAP values are scattered, with both high and low pressure or Ni fraction producing positive or negative contributions. GB-dominated and grain-growth states exhibited similar ambiguity, though broader trends matched the widening of their respective regions in the deformation pattern map. Higher Ni fraction and lower pressure favored GB-dominated outcomes, whereas grain-growth was promoted by a lower Ni fraction and higher pressure. The shear-dominated final state showed the clearest metadata trend, with strong preference for low Ni fraction and high pressure.

The positional encoding (slice index) contributed negligibly across all classes. Additionally, many samples exhibited SHAP values clustered near zero for individual metadata entries, suggesting that for a substantial subset of cases no single metadata parameter had a decisive marginal effect on the model output. This reflects the multivariate nature of the pattern map, where classifications may depend on combinations rather than isolated metadata variables.

For the image branch, SHAP was applied using partition-based explanations, which decompose an image into smaller regions and estimate each region's contribution to the class probability. The resulting overlays color-code positive and negative influence, allowing spatial interpretation of where the model "looks" when

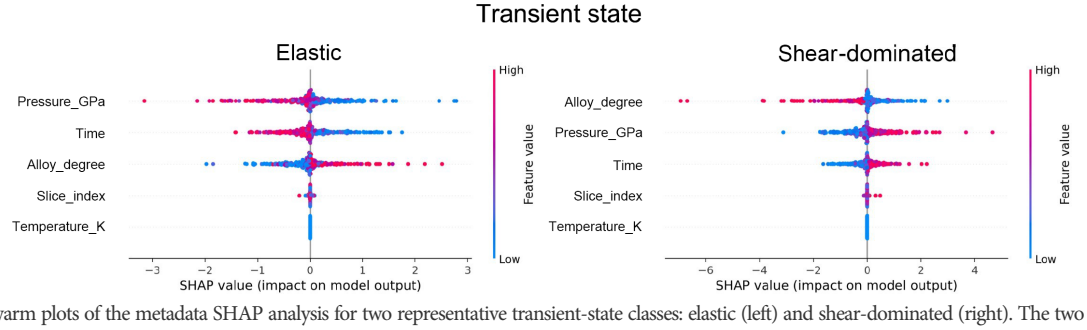


Fig. 11 Beeswarm plots of the metadata SHAP analysis for two representative transient-state classes: elastic (left) and shear-dominated (right). The two classes exhibit opposing and physically consistent metadata contributions. High normal pressure and advanced simulation time increase the likelihood of the shear-dominated state while reducing the probability of the elastic state, reflecting the increasing energy input that must be accommodated by plastic deformation, ultimately, near-surface shearing. The influence of Ni fraction is likewise opposite for the two classes. Higher Ni contents favor the elastic state, indicating that Ni-rich alloys withstand severe plastic deformation longer than Cu-rich compositions, which transition earlier to shear-dominated behavior. The slice index shows negligible influence, and temperature does not contribute, as it was constant across the dataset.

清華大学出版社 Tsinghua University Press | SciOpen

assigning a label. Unlike metadata beeswarm plots, these explanations cannot be aggregated into a meaningful class-level average, as they represent SHAP values for a single sample only, with a different dimensionality. The computational cost is substantially higher than for metadata when perturbed versions of each image are evaluated against the background dataset. Consequently, metadata must be held fixed during image SHAP computation, as allowing both image regions and metadata values to vary jointly would be prohibitive in dual-branch cases [50].

In practice, interpreting the image-based SHAP results is challenging. For most samples, overlays intuitively highlight microstructural features, e.g., regions with strong deformation contributing negatively to the elastic class probability (Fig. 12). However, in many cases the SHAP-attributed influential regions do not align with the model's actual predicted class. Even when evaluating metadata SHAP values and image SHAP values for the same sample, it is often difficult to derive a coherent explanation. Additional effort may therefore be required to better understand and interpret these outputs, especially given the high-dimensional, multimodal nature of the model.

6 Conclusions and outlook

The workflow developed in this study provides a first practical assessment of ML-assisted deformation pattern classification based on large-scale molecular dynamics simulations using machine learning. By demonstrating that deformation patterns can be inferred directly from structural tomographs using learned feature representations, we provide a proof of concept for data-driven mechanism identification. The dual-branch CNN–MLP model achieves high classification accuracy even on previously unseen configurations, and the latent-space analysis confirms that essential microstructural characteristics can be captured in a condensed, compact form. Together, these results highlight the feasibility of automating deformation regime labeling and of building predictive tools that can augment, accelerate, or partially reduce reliance on costly MD simulations and subsequent steps.

Neural network architectures offer a large search space for optimizing their performance, including architecture, convolutional depth, fusion strategies, activation functions, regularization schemes, and loss formulations as well as hyperparameters. A systematic exploration of these requires AutoML approaches and significant time and effort and, thus, is suggested for future work. A deliberately conservative architecture, based on existing state-of-the-art dual-branch neural networks architectures, was therefore chosen to establish feasibility and to analyze generalization behavior. The strong performance obtained nevertheless indicates the potential of marrying machine learning, MD simulations, and tribological studies. In particular, additional architectural tuning and hyperparameter optimization could further enhance robustness and accuracy.

Future developments could extend the present workflow beyond static deformation-state classification towards the analysis of multiple stages of tribosystem evolution. A natural next step is to assess whether microstructural evolution itself can be learned, which would require incorporating temporal information and exploring richer data representations. Promising approaches include voxel-based encodings that retain the spatial relationship of the 3-dimensional system, point-cloud descriptions of defects, grain boundaries or orientations, and graph-based representations in which grains form nodes that are connected by grain boundary edges. Such representations would allow the application of advanced architectures like transformers [54], physics informed neural networks (PINNs) [55], graph neural networks (GNNs), spatio-temporal GNNs [56, 57], and physics-embedded graph networks for learning microstructure evolution on topological graphs [58]. Such architectures outline possible directions for moving from state classification to predictive modeling, in which the model iteratively predicts future microstructural states and behavior.

Besides algorithmic enhancement and developments, the present study highlights that broader and more diverse MD datasets are likely to be a limiting factor for further progress. Expanding simulation coverage across alloy classes, compositions, pressures, temperatures, and loading conditions, including highly loaded systems that traverse multiple deformation modes within short time intervals, will be essential. Such highly loaded configurations serve not only to enrich the range of observable mechanisms but also provide means to evaluate whether shortened simulations containing compressed sequences of deformation events are sufficient for learning material behavior. If feasible, this could allow the model to infer long-term or steady-state responses from truncated trajectories, thereby reducing the reliance on fully resolved and computationally expensive MD runs. In a similar vein, hybrid approaches that capitalize on execution in machine learning and are coupled with certain MD simulation aspects are worth exploring.

In the long term, the perspective extends beyond deformation pattern identification from static tomographs. A realistic objective is the gradual development of predictive frameworks that infer aspects of material behavior from limited input, such as initial microstructure and loading conditions, while relying on MD

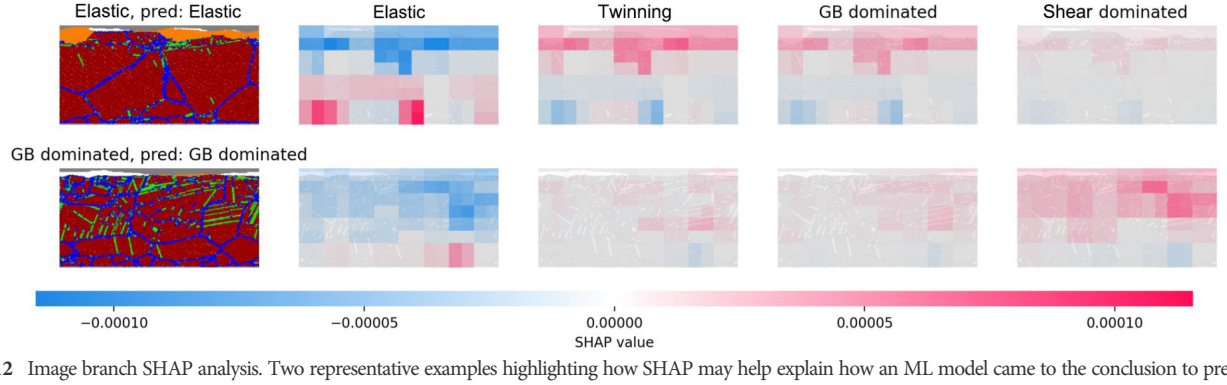


Fig. 12 Image branch SHAP analysis. Two representative examples highlighting how SHAP may help explain how an ML model came to the conclusion to predict a given label. Image overlays show regions of identified positive (red) and negative (blue) class prediction influence. Top row: While the top end of the image (with the sliding interface) would favor a prediction of either twinning or GB dominated deformation patterns, the lower end favors the elastic regime. Bottom row: The visible grain refinement and defect density in the image on the left correctly precludes a labeling as elastic and suggests shear-dominated behavior, but consideration of the metadata leads to labeling as GB-dominated.

simulations primarily for calibration and validation. Achieving this would move computational tribology in the direction of data-driven, mechanism-aware representations analogous to tribological Ashby charts and toward AI-assisted simulation pipelines that formalize and reuse the mechanistic knowledge embedded in large MD datasets.

Author contributions

H. J. Ehrich: Conceptualization, data curation, methodology, software, visualization, and writing—original draft. M. C. May: Validation, writing—review & editing. S. J. Eder: Conceptualization, supervision, writing—review & editing, and data curation.

Acknowledgements

The authors acknowledge the TU Wien Bibliothek for financial support through its Open Access Funding Programme. Part of the presented results were realized as part of the COMET Centre InTribology (FFG No. 906860), a project of the "Excellence Centre for Tribology" (AC2T research GmbH). COMET is managed by the Austrian Research Promotion Agency (FFG). The computational results presented here were obtained using the Vienna Scientific Cluster (VSC). Part of the results are supported by the Nanyang Technological University Start-up grant (No. SUG025861-00001).

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article.

Data availability statement

The data that support the findings of this study are available from the corresponding authors upon request.

Electronic Supplementary Material

Supplementary material is available in the online version of this article at https://doi.org/10.26599/FRICT.2026.9441269.

References

[1] Holmberg K, Erdemir A. Influence of tribology on global energy consumption, costs and emissions. *Friction* **5**(3): 263–284 (2017)
[2] Pastewka L, Vakis A I, Eder S J, Aghababaei R, Almqvist A, Carbone G, Chandross M, Dini D, Ehrich H J, Ewen J P, et al. Modeling in tribology: Recent advances, applications, and open questions. *Tribol Int* **218**: 111326 (2026)
[3] Eder S J, Rodríguez Ripoll M, Cihak-Bayr U, Dini D, Gachot C. Unraveling and mapping the mechanisms for near-surface microstructure evolution in CuNi alloys under sliding. *ACS Appl Mater Interfaces* **12**(28): 32197–32208 (2020)
[4] Dan T, Saitoh K I. Microstructure evolution in polycrystalline metal under severe plastic deformation by strain-controlled molecular dynamics. *JMMP* **6**(1): 48–60 (2012)
[5] Shi Y Q, Szlufarska I. Effect of growth twins on strength and microstructural evolution of nanocrystalline aluminum. *J Mater Sci* **56**(26): 14587–14597 (2021)
[6] Dmitriev A I, Nikonov A Y, Österle W. Molecular dynamics sliding simulations of amorphous Ni, Ni-P and nanocrystalline Ni films. *Comput Mater Sci* **129**: 231–238 (2017)
[7] Romero P A, Järvi T T, Beckmann N, Mrovec M, Moseler M. Coarse graining and localized plasticity between sliding nanocrystalline metals. *Phys Rev Lett* **113**(3): 036101 (2014)
[8] Eder S J, Cihak-Bayr U, Bianchi D, Feldbauer G, Betz G. Thermostat influence on the structural development and material removal during abrasion of nanocrystalline ferrite. *ACS Appl Mater Interfaces* **9**(15): 13713–13725 (2017)
[9] Hu Y Q, Xu J F, Su L, Zhang Y H, Ding S H, Shen Y H, Xia R. Tensile responses of polycrystalline Mo via molecular dynamics simulation: Grain size and temperature effects. *Mater Chem Phys* **296**: 127270 (2023)
[10] Li A, Szlufarska I. How grain size controls friction and wear in nanocrystalline metals. *Phys Rev B* **92**(7): 075418 (2015)
[11] Dong C S, Jiang W. Molecular dynamics study on the effect of grain size on the plastic deformation of Fe–Cr–Ni alloy nanopolycrystals. *Mater Today Commun* **38**: 107770 (2024)
[12] Hansen N. Hall–Petch relation and boundary strengthening. *Scr Mater* **51**(8): 801–806 (2004)
[13] Schuh C A, Nieh T G. Hardness and abrasion resistance of nanocrystalline nickel alloys near the hall-petch breakdown regime. *MRS Online Proc Libr* **740**(1): 18 (2003)
[14] Argibay N, Chandross M, Cheng S, Michael J R. Linking microstructural evolution and macro-scale friction behavior in metals. *J Mater Sci* **52**(5): 2780–2799 (2017)
[15] Eder S J, Cihak-Bayr U, Gachot C, Rodríguez Ripoll M. Interfacial microstructure evolution due to strain path changes in sliding contacts. *ACS Appl Mater Interfaces* **10**(28): 24288–24301 (2018)
[16] Ehrich H J, Dollmann A, Grützmacher P G, Gachot C, Eder S J. Automated identification and tracking of deformation twin structures in molecular dynamics simulations. *Comput Mater Sci* **236**: 112878 (2024)
[17] Zhou T, Song Z, Sundmacher K. Big data creates new opportunities for materials research: A review on methods and applications of machine learning for materials design. *Engineering* **5**(6): 1017–1026 (2019)
[18] Louie S G, Chan Y H, da Jornada F H, Li Z L, Qiu D Y. Discovering and understanding materials through computation. *Nat Mater* **20**(6): 728–735 (2021)
[19] Bishara D, Xie Y X, Liu W K, Li S F. A state-of-the-art review on machine learning-based multiscale modeling, simulation, homogenization and design of materials. *Arch Comput Meth Eng* **30**(1): 191–222 (2023)
[20] Yin N, Yang P F, Liu S K, Pan S H, Zhang Z N. AI for tribology: Present and future. *Friction* **12**(6): 1060–1097 (2024)
[21] Zhang L, Qian K, Huang J, Liu M, Shibuta Y. Molecular dynamics simulation and machine learning of mechanical response in non-equiatomic FeCrNiCoMn high-entropy alloy. *J Mater Res Technol* **13**: 2043–2054 (2021)
[22] Lin J R, Hasan M, Acar P, Blanchet J, Tarokh V. Neural network accelerated process design of polycrystalline microstructures. *Mater Today Commun* **36**: 106884 (2023)
[23] Yin N, Xing Z G, He K, Zhang Z N. Tribo-informatics approaches in tribology research: A review. *Friction* **11**(1): 596 (2023)
[24] Paturi U M R, Palakurthy S T, Reddy N S. The role of machine learning in tribology: A systematic review. *Arch Computat Methods Eng* **30**(2): 1345–1397 (2023)
[25] Fukuya T, Shibuta Y. Machine learning approach to automated analysis of atomic configuration of molecular dynamics simulation. *Comput Mater Sci* **184**: 109880 (2020)
[26] Gomberg J A, Medford A J, Kalidindi S R. Extracting knowledge from molecular mechanics simulations of grain boundaries using machine learning. *Acta Mater* **133**: 100–108 (2017)
[27] Kaufmann K, Zhu C Y, Rosengarten A S, Maryanovsky D, Wang H R, Vecchio K S. Phase mapping in EBSD using convolutional neural networks. *Microsc Microanal* **26**(3): 458–468 (2020)
[28] Sieberg P M, Kurtulan D, Hanke S. Wear mechanism classification using artificial intelligence. *Materials* **15**(7): 2358 (2022)
[29] Sieberg P M, Hanke S. Challenges and potentials in the classification of wear mechanisms by artificial intelligence. *Wear* **522**: 204725 (2023)
[30] Jin W Y, Pei J Y, Xie P, Chen J C, Zhao H Y. Machine learning-based prediction of mechanical properties and performance of nickel–graphene nanocomposites using molecular dynamics simulation data. *ACS Appl Nano Mater* **6**(13): 12190–12199 (2023)
[31] Mińkowski M, Laurson L. Predicting elastic and plastic properties of small iron polycrystals by machine learning. *Sci Rep* **13**: 13977 (2023)

[32] Chen S H, Xu N. Detecting microstructural criticality/degeneracy through hybrid learning strategies trained by molecular dynamics simulations. *ACS Appl Mater Interfaces* **15**(7): 10193–10202 (2023)
[33] Ehrich H J, May M C, Eder S J. Machine-learning-enabled interpretation of tribological deformation patterns in large-scale MD data. *arXiv preprint* arXiv: 2512.05818 (2025)
[34] Eder S J, Grützmacher P G, Rodríguez Ripoll M, Dini D, Gachot C. Effect of temperature on the deformation behavior of copper nickel alloys under sliding. *Materials* **14**(1): 60 (2021)
[35] Eder S J, Grützmacher P G, Ripoll M R, Gachot C, Dini D. Does speed kill or make friction better?—Designing materials for high velocity sliding. *Appl Mater Today* **29**: 101588 (2022)
[36] Eder S J, Bianchi D, Cihak-Bayr U, Gkagkas K. Methods for atomistic abrasion simulations of laterally periodic polycrystalline substrates with fractal surfaces. *Comput Phys Commun* **212**: 100–112 (2017)
[37] Eder S J, Cihak-Bayr U, Bianchi D. Large-scale molecular dynamics simulations of nanomachining. In: *Advanced Machining Processes*. Markopoulos A P, Davim J P, Eds. Florida (USA): CRC Press, 2017: 141–177.
[38] Groeber M A, Jackson M A. DREAM. 3D: A digital representation environment for the analysis of microstructure in 3D. *Integr Mater Manuf Innov* **3**(1): 56–72 (2014)
[39] Yeh C K, Wu W C, Ko W J, Wang Y F. Learning deep latent space for multi-label classification. In: Proceedings of the 31st AAAI Conference on Artificial Intelligence, San Francisco, California, USA, 2017, **31**(1): 2838–2844.
[40] Goodfellow I, Bengio Y, Courville A. *Deep Learning*. Cambridge (USA): MIT Press, 2016.
[41] He K M, Zhang X Y, Ren S Q, Sun J. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 2016: 770–778.
[42] Gao F, Huang T, Wang J, Sun J P, Hussain A, Yang E F. Dual-branch deep convolution neural network for polarimetric SAR image classification. *Appl Sci* **7**(5): 447 (2017)
[43] Hunter D, Yu H, Pukish M S, Kolbusz J, Wilamowski B M. Selection of proper neural network sizes and architectures: A comparative study. *IEEE Trans Ind Inf* **8**(2): 228–240 (2012)
[44] Luo R Q, Tian F, Qin T, Chen E H, Liu T Y. Neural architecture optimization. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems, Red Hook, NY, USA, 2018, **31**: 7827–7838.
[45] Liang X Z, Wang X B, Lei Z, Liao S C, Li S Z. Soft-margin softmax for deep classification. In: *Neural Information Processing*. Liu D, et al. Eds. Springer Cham, 2017: 413–421.
[46] Kingma D P, Ba J. Adam: A method for stochastic optimization. *arXiv preprint* arXiv: 1412.6980 (2014)
[47] Feng L, Shu S L, Lin Z Y, Lv F M, Li L, An B. Can cross entropy loss be robust to label noise? In: Proceedings of the 29th International Joint Conference on Artificial Intelligence, Yokohama, Japan, 2020: 2206–2212.
[48] Yu F, Seff A, Zhang Y D, Song S R, Funkhouser T, Xiao J X. LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint* arXiv: 1506.03365 (2015)
[49] Lundberg S M, Lee S I. A unified approach to interpreting model predictions. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, Red Hook, NY, USA, 2017: 4768–4777.
[50] Wang H J, Liang Q X, Hancock J T, Khoshgoftaar T M. Feature selection strategies: A comparative analysis of SHAP-value and importance-based methods. *J Big Data* **11**(1): 44 (2024)
[51] Eelbode T, Sinonquel P, Maes F, Bisschops R. Pitfalls in training and validation of deep learning systems. *Best Pract Res Clin Gastroenterol* **52–53**: 101712 (2021)
[52] Bernstein N, Tadmor E. Tight-binding calculations of stacking energies and twinnability in FCC metals. *Phys Rev B* **69**(9): 094116 (2004)
[53] Kibey S, Liu J B, Johnson D D, Sehitoglu H. Predicting twinning stress in FCC metals: Linking twin-energy pathways to twin nucleation. *Acta Mater* **55**(20): 6843–6851 (2007)
[54] Gao Z H, Zhu C S, Shu Y F, Wang C L. Exploring time-series transformers for spatio-temporal prediction of microstructural evolution of polycrystalline grain. *Mater Today Commun* **40**: 109477 (2024)
[55] Hasan M, Ender Eger Z, Senthilnathan A, Acar P. Microstructure-sensitive deformation modeling and materials design with physics-informed neural networks. *AIAA J* **62**(5): 1864–1874 (2024)
[56] Qin Y G, DeWitt S, Radhakrishnan B, Biros G. GrainGNN: A dynamic graph neural network for predicting 3D grain microstructure. *J Comput Phys* **510**: 113061 (2024)
[57] Fan S X, Hitt A L, Tang M, Sadigh B, Zhou F. Accelerate microstructure evolution simulation using graph neural networks with adaptive spatiotemporal resolution. *Mach Learn Sci Technol* **5**(2): 025027 (2024)
[58] Xue T J, Gan Z T, Liao S H, Cao J. Physics-embedded graph network for accelerating phase-field simulation of microstructure evolution in additive manufacturing. *npj Comput Mater* **8**: 201 (2022)



Hendrik J. Ehrich is a Ph.D. student in mechanical engineering at TU Wien and a junior researcher at AC2T research GmbH. His work focuses on molecular dynamics simulations, tribology, and data-driven analysis methods for materials science applications. He has experience in developing analysis workflows for atomistic simulations and applying machine learning techniques to engineering problems, including automated defect detection and microstructural analysis. He collaborates with researchers at Karlsruhe Institute of Technology and has presented his work at international academic events, including at Nanjing University of Aeronautics and Astronautics. He is currently involved in the EU project SiToLub and research on machine learning methods for reactive molecular dynamics simulations.



Marvin C. May is a professor and researcher in industrial AI with expertise in manufacturing systems, robotics, and data-driven engineering. He earned multiple degrees with distinction, including a Dr.-Ing. degree in mechanical engineering, from Karlsruhe Institute of Technology and previously held research positions at Karlsruhe Institute of Technology and Massachusetts Institute of Technology before joining Nanyang Technological University. His work focuses on applying AI to complex industrial environments with sparse and domain-specific data, particularly in advanced manufacturing. He is also affiliated with CIRP and serves as General Chair of the CIRPe 2026 Conference.



Stefan J. Eder leads nanoscopic modeling activities in the tribology group at TU Wien and is a principal scientist at AC2T research GmbH. His research focuses on computational tribology, particularly microstructural evolution in polycrystalline alloys under tribological loading and tribochemical simulations. He received his Ph.D. degree in physics from TU Wien in 2012 and completed his habilitation in computational materials tribology in 2024. Dr. Eder has collaborated internationally with researchers at University of California, Merced, Imperial College London, and Karlsruhe Institute of Technology. He has authored around 70 peer-reviewed publications, contributed extensively to international conferences, and is active in scientific editorial and research network activities, including COST Action CA23155 Ocean Tribology.