

This is the accepted manuscript made available via CHORUS. The article has been published as:

Spatial strain correlations, machine learning, and deformation history in crystal plasticity

Stefanos Papanikolaou, Michail Tzimas, Andrew C. E. Reid, and Stephen A. Langer

Phys. Rev. E **99**, 053003 — Published 16 May 2019

DOI: [10.1103/PhysRevE.99.053003](https://doi.org/10.1103/PhysRevE.99.053003)

Spatial strain correlations, machine learning and deformation history in crystal plasticity

Stefanos Papanikolaou,^{1,2} Michail Tzimas,¹ Andrew C.E. Reid,³ and Stephen A. Langer⁴

¹*The West Virginia University, Department of Mechanical & Aerospace Engineering, Morgantown, WV, 26505*

²*The West Virginia University, Department of Physics, Morgantown, WV, 26505*

³*Material Measurement Laboratory, National Institute of Standards and Technology, Gaithersburg, MD 20899*

⁴*Information Technology Laboratory, National Institute of Standards and Technology, Gaithersburg, MD 20899*

Systems far from equilibrium respond to probes in a history-dependent manner. The prediction of the system response depends on either knowing the details of that history or being able to characterize all the current system properties. In crystal plasticity, various processing routes contribute to a history dependence that may manifest itself through complex microstructural deformation features with large *strain gradients*. However, the complete spatial strain correlations may provide further predictive information. In this paper, we demonstrate an explicit example where spatial strain correlations can be used in a statistical manner to infer and classify prior deformation history at various strain levels. The statistical inference is provided by machine-learning (ML) techniques. As source data, we consider uniaxially compressed crystalline thin films generated by two dimensional discrete dislocation plasticity simulations, after prior compression at various levels. Crystalline thin films at the nanoscale demonstrate yield-strength size effects with very noisy mechanical responses that produce a serious challenge to learning techniques. We discuss the influence of size effects and structural uncertainty to the ability of our approach to distinguish different plasticity regimes.

I. INTRODUCTION

The term “far from equilibrium” describes statistical ensembles that are not exploring the available microstates in an ergodic manner. Thus, the system’s response to probes (*e.g.* mechanical stress, electric/magnetic fields) is highly dependent on the particular initial condition and its history. In metallurgy, far from equilibrium microstructures are easily created by plastic deformation but their history is typically hidden (*e.g.* various processing routes used, such as extrusion, forging *etc.*). It is common that a sample that had been processed in various ways, then polished, is further mechanically tested as a part of a component [1–6]. In such cases, characterizing the mechanical response of the sample requires either the precise history of the processing routes taken, or the precise knowledge of the particular state realized and its state variables. In plasticity, the main observable, the strain tensor of the freshly deformed sample, is usually not capable of uniquely and reliably characterizing the mechanical response [7]. Nevertheless, spatially resolved strain correlations do reflect the full spatial structure of the microstructure [8–11], and are formally equivalent [8] to capturing microstructural strain gradients [12–14]; strain gradients have been shown to classify plastically deformed microstructures in various cases [15, 16].

In this paper, we present a direct and simple example where we can statistically infer the prior deformation history of a crystalline thin film sample through a small-strain mechanical test by only using spatial strain correlations. We utilize discrete dislocation dynamics simulations to produce test samples with well-controlled mechanical processing histories. We demonstrate, using unsupervised machine learning, that spatial correlations encoded in the strain images of the reloaded samples, contain adequate information to produce a full classification of dislocation-driven deformation history at multiple scales and perform reliable predictions.

It is an axiom of materials science [17] that microstructure controls properties, and of course microstructure is the result of processing. In metallurgy, standard metal preparation techniques involve a variety of processing steps [18] with associated process-structure-property linkages; consequently, components of the same net shape and composition, but different prior processing technique(s), may exhibit different mechanical properties. In some fields, these processing-structure-property relationships are empirically well-established, and in some others, only general qualitative knowledge is common. However, in theory [8] as well as in practice [18], this type of linkage may be questioned when memory-dependent, hysteretic phenomena [19] such as friction, plastic-

ity or/and fracture are involved, but only partial sample history information exists.

Nevertheless, in the context of crystal plasticity, there has been strong evidence that the formal absence of the complete sample deformation history information, may be remedied to a large degree by accounting for developed local gradients in the plastic/elastic distortion [20, 21]. These gradients naturally represent a “truncated truth” of what happened to the material through its deformation, but remain at least a set of observables that one may check by direct mechanical tests and also, display a sense of clarity by connecting to the microstructural feature of strain localization [7]. Strain gradients in a microstructure can be directly captured through measures of spatial strain correlations [8]. The fact that spatial correlations of an observable can capture and classify its gradient, has been the principal reason that spatial correlations have been traditionally used as a direct way to assess phase transitions in statistical mechanics [22], driven by interface phenomena of appropriately defined order parameters [23]. In phase transitions, the spatial correlations of the order parameter develop a long-distance expectation value which can be a direct phase signature. However, crystal plasticity’s complexity may allow for a large variety of possible spatial correlation features (depending on type of plasticity activated), and thus, it is natural to avoid a direct –and possibly constraining– assessment of particular short-distance features.

Spatial correlation features of crystal plasticity may contain various plastic deformation signatures that may include dislocation-driven motions and processes, diffusional creep, mechanical twinning or grain boundary sliding [24, 25]. Despite the multitude of origins, the success of strain gradients in capturing deformation history of plastically deformed materials, is based on the main spatial signatures of strain localization and shear banding [7, 26]. Shear bands have been identified as possible indicators of prior deformation, given that the creation of slip bands during various small-load mechanical tests of polished samples, shows strong dependence on the prior deformation history [1, 2, 4]. However, the presence/absence of shear banding may not suffice for characterizing prior processing, and a more complete characterization of crystal plasticity history should require the classification of the full spatial correlations of stress and strain tensorial fields [21].




FIG. 1. Schematic of experimental non invasive test in modeling and sequence of sample loading: A schematic of the sequence of events when loading, unloading and reloading a sample is shown, with the corresponding stress and strain graphs vs the time step of simulation. A sample is obtained from 2D-DDD simulation and it’s loaded to a specific strain value – “Prior” Loaded stage (Stage L). Then, the sample is unloaded to zero stress and the remaining plastic strain can be calculated – “Prior” Unloaded stage (Stage U). Finally, the sample is reloaded to a testing deformation – “Test” Reloaded stage (Stage T). Even though a sample has been plastically deformed (Stage L), the samples obtained from experiments can be polished, thus the surface of a sample is not able to provide information about deformation (Sample at stage U can be seen in the figure as having smooth surface). Such techniques are applicable in experiments such as digital image correlation, where randomly placed tracking nanoparticles are detected optically and contribute to correlation statistics are applied to the sample. Then, as the sample is reloaded, the permanent deformation can be observed, since there are changes in the distances between tracked nanoparticles.

In this paper, we only focus on crystal plasticity signatures during uniaxial compression of thin films, obtained through a two dimensional discrete dislocation dynamics model (2D-DDD). Given that dislocation movements are the principal instigators of plastic deformation on materials, this study captures a large number of experimentally relevant cases and can provide a transparent application framework [27–30]. The simulated systems begin in a state with no pre-existing mobile dislocations, but with a set of dislocation sources and obstacles present (see also Fig. 1). The systems are then run through a compression and release, with the amplitude of this compression being

the “prior deformation”. After release, the systems are “polished” (by erasing any pre-existing memory), and then they are subjected to a second, low-amplitude compression, the amplitude of which is called the “testing deformation”. The strain developed during the testing deformation is imaged, and information derived from this imaging process is the input to a machine-learning (ML) approach, whose goal is to accurately determine the amplitude of the prior deformation. Because we use a discrete set of prior deformation amplitudes, our ML task is one of classification – for a given data set obtained from the testing deformation, the ML set-up should be able to say which discrete prior deformation the sample had undergone. While we measure the strain developed during the testing deformation directly through simulation, the process is meant to mimic a readily-accessible, non-destructive experimental technique, namely, digital image correlation (DIC) [31–38].

For our simulations and chosen tests, we use the exemplary case of thin *Al* films ($\leq 2\mu\text{m}$) under uniaxial compression. The films may have undergone prior deformation of 0.1, 1 or 10% deformation in terms of prior uniaxial compression, before unloading to reach mechanical equilibrium (*cf.* Fig. 2). Over the last decade, it has been shown that experimental [39] and simulation tests [28, 29] on such thin-film geometries, display very noisy characteristics that may develop unpredictable features, such as avalanches [40], especially at very small widths $w < 500\mu\text{m}$. (see for example the deformation features in Fig.A.1 in Appendix A of the Supplementary Information (SI) [41]). While the overall behavior of the films may be traceable to the behavior of the film’s mean dislocation density, in this paper we explore how well we could infer prior deformation if neither dislocations nor prior deformations are observable but instead, only the total strain in a small-strain compressive thin-film test. We believe that our study can become an exemplary test problem for “black-box” and *physics-agnostic* machine-learning methods in Data Science and Informatics, before these methods become applicable to devices and geometries that undergo mechanical deformation with elasticity-plasticity.

The remainder of this paper is organized as follows: In section II we describe our 2D-DDD model and our approach for quantifying strain correlation patterns. In section III we discuss the ML approaches we use for processing our data samples,




FIG. 2. **Schematic of the various loading histories & testing deformation on samples:** A material has an assumed prior deformation history (Stage L, red circles). We unload the sample and obtain the Stage U (blue circles). How does the strain field, which characterizes Stage T (green circles), reflect the prior history? The testing deformation $\epsilon_T - \epsilon_U$ is constant in all T-U cases.

and present results based on system sizes, reloading strain, slip systems and alternate processing techniques. In section IV we make several remarks and present a summary of our work. Appendix A in the SI [41] contains more details of the numerical simulations, while in Appendix B in the SI we discuss the loading/unloading results from a physical consistency perspective. In the Appendix C of the SI, we describe the calculation of the statistical correlations. In the Appendix D of the SI, we discuss principal component analysis (PCA), a statistical approach that is heavily used in this work, while in the Appendix E of the SI we present a more complete set of parameters for our results. Finally, in the Appendix F of the SI we show that spatial resolution does not affect the classification of prior deformation history in crystal plasticity.

II. THE MODEL

A. Discrete dislocation dynamics

We use a 2D-DDD model that is a simplified version of crystal plasticity. It can capture the most important crystal plasticity mechanisms, namely dislocation gliding, nucleation and mutual interactions. Even though there are 3D-DDD mod-




FIG. 3. **The 2D discrete dislocation plasticity model of uniaxial compression of thin films:** Slip planes (lines) span the sample, equally spaced at $d = 10b$, but planes close to corners are deactivated to maintain a smooth loading boundary. Surface and bulk dislocation sources (red dots) and forest obstacles (blue dots) are spread homogeneously across the active slip planes. Initially the sample is stress and dislocation free.

els that include more detailed aspects of dislocation microstructures [15], this 2D model suffices to demonstrate the feasibility of our approach while increasing statistical accuracy. We consider a 2D plane stress problem, considering only infinitely long straight dislocation lines along the third dimension. While minimal, this model has been used extensively in the past for thin-film modeling [27–30], and the choice of parameters in the model (dislocation sources, obstacles, slip spacing) are based on previous studies of nanopillar compression, where realistic dislocation densities were obtained [29].

In this model, plastic flow occurs by the nucleation and glide of edge dislocations, on single or double slip systems. Our primary focus will be on double slip system samples (see Fig. 3), and we will compare the performance with single slip in Sec. III C. Samples are modeled [29] by a rectangular profile of width w , height h and aspect ratio α ($\alpha = h/w$). We study sample widths ranging in powers of 2 from $w = 0.125$ (or w_0) to $2 \mu\text{m}$ with $\alpha = 4$. 2D-DDD samples are discretized on a finite element mesh of 320×80 square elements, independently of the width. The top and bottom edges ($x = 0, w$) are traction free, allowing dislocations to exit the sample. Loading is taken to be ideally displacement controlled, by prescribing the y -displacement at the lateral edges ($y = 0, h$). The applied strain rate (for both loading and unloading regimes), $\dot{h}/h = 10^4 \text{ s}^{-1}$, is held constant across all our simulations, similar to experimental practice. Plastic deformation of the crystalline samples is described using the discrete dislocation framework for small strains [27]. Each dislocation

is treated as a singularity in a linear elastic background solid with Young’s modulus E and Poisson ratio ν , whose analytic solution is known at any position. We assume that the Burgers vector $b = 0.25nm$.

In the model, slip planes are separated by $10b$ and oriented at $\pm 30^\circ$ from the loading direction (Fig. 3). In the single slip model, planes are also separated by $10b$ but are oriented in just one direction (30° from the loading direction). Bulk sources are randomly distributed over slip planes and locations, and their strength is selected randomly from a Gaussian distribution with mean value $\bar{\tau}_{\text{nuc}} = 50 \text{ MPa}$ and 10 % standard deviation. Forest dislocation obstacles with strength τ_{obs} are also distributed on the samples. Their strength is Gaussian distributed with mean 300 MPa and 20 % standard deviation (see Appendix A of the SI [41]).

At the beginning of the calculation, the crystal is stress free and there are no mobile dislocations. We only consider glide of dislocations, neglecting the possibility of climb. The motion of dislocations is determined by the Peach-Koehler force in the slip direction. Once nucleated, dislocations can either exit the sample through the traction-free sides, annihilate with a dislocation of opposite sign when their mutual distance is less than $6b$, or become pinned at an obstacle. Our simple obstacle model is that a dislocation stays pinned until its Peach-Koehler force exceeds the obstacle-dependent value $\tau_{\text{obs}}b$. If dislocations approach the physical boundary of the sample then a geometric step is created on the surface along the slip direction (see Fig. 4) Our simulations are carried out for material parameters that are reminiscent of aluminum: $E = 70 \text{ GPa}$, $\nu = 0.33$. The effective Young’s modulus for plane stress problems is $E_{\text{eff}} = E/(1 - \nu^2) = 78.55 \text{ GPa}$. In the case of double-slip systems, we consider 50 random realizations of sources and obstacles in each parameter case (loading of 0.1, 1, 10 %) for a total of $n = 150$ samples. For single-slip systems, we consider 9 random realizations for each parameter case for a total of $n = 27$ samples.

The simulation is carried out incrementally, using a time step that is a factor 20 smaller than the nucleation time $t_{\text{nuc}} = 10 \text{ ns}$. At the beginning of every time increment, nucleation, annihilation, pinning at and release from obstacle sites are evaluated. After updating the dislocation structure, the new stress field in the sample is determined, using the finite element method to solve for the




FIG. 4. **Strain profiles captured from the 2D-DDD simulation at stage L and stage T – $w = 2 \mu\text{m}$ – Double slip system:** The full sample is shown, in both deformed (a) or undeformed (b) coordinates. A sample is loaded to 10 % strain, unloaded to zero stress and reloaded to testing deformation of 0.1 %. (a) Strain profile for Stage L at 10 % strain. Plastic steps are allowed to emerge on the film surface [29]. (b) A strain profile at the Stage T, after subtracting the residual plastic deformation at Stage U. Such strain profiles are analogous to typical DIC experimental strain profiles. The strain maps are unitless.

image fields [27].

The test that we wish to imitate would measure the strain field in the sample after it has been strained and relaxed, as described above, and then subjected to a subsequent “testing” deformation. We consider a testing reload regime that is governed mainly by the degree of invasiveness we introduce to the data set. All tests have been carried out for prior deformation (see Sec. I in three different amplitudes (0.1 %, 1 %, 10 %) of total strain. Figs. 1,2 show a schematic of the way we create our data set: “As annealed” samples (see Fig 1) are loaded to 3 different amplitudes (L stages). For each stage L, we unload (at 0 applied stress) to obtain U stages. In stage U, the samples are stress free, but there is some remaining strain due to plasticity. We then reload the samples to a specific testing deformation (stage T).

Samples of different widths (w) undergo the same unload-reload protocol to create our data set. We have the option to select at which strain the unload process begins, as well as the testing deformation level we want to introduce. We perform tests at two different reloading strains (*small-reload* data set: 0.1 % and *large-reload* data set: 1.0 %). Reloading strain is the strain difference between stages T and U ($\epsilon_T - \epsilon_U$). Fig. 5 shows two typical stress-strain curves like the one shown schematically in Fig. 1. Stages L (triangles), as well as stages U (squares) and T (circles) are shown. Note that in Figure 5 (a), the inset figure shows that the slope of the stress-strain




FIG. 5. **$w = 2 \mu\text{m}$ – Small vs large reload stress strain curves:** Samples are loaded to 3 different deformations levels (Stage L- \blacktriangledown). For each level, the sample is then unloaded (to zero applied stress) creating Stage U (\blacksquare). The predeformed sample is then reloaded to a testing deformation (Stage T - \bullet). Each color represents a particular class of prior deformed states. (a) Testing deformation = 0.1 %. The inset figure represents a zoom-in of the circled unloading-reloading region (Stages U/T at 10 % strain.). The dashed line represents the ideal elastic response by just considering the parent material’s elastic modulus. The solid lines correspond to the unloading and reloading curves. The reloading curves slightly deviate from the ideal elastic response even in small stresses, due to the presence of inelastic precipitates and also, due to the ideal form of applied uniaxial loading. (b) Testing deformation = 1 %.

curve, in the case of 10 % stage L amplitude, exceeds the elastic modulus when the testing load is applied. Dislocation motion relieves stress and reduces the effective modulus, but obstacles and boundary conditions (the top and bottom boundaries are constrained to be straight) impede relaxation and stiffen the system. The excess modulus arises from the work required to maintain straight boundaries as the system deforms. For more information on the unloading-reloading procedures, see Appendix B of the SI [41]. Figure 5 (b) shows a stress strain curve obtained through reloading to larger testing deformation, 1 %. The main difference is at the reload points, which show further deformation of the sample, in contrast to samples reloaded to smaller strain (0.1 %). For the small-reload data set case, the reload strain value is small enough that does not introduce further plastic deformation.

B. Extracting machine learning input data from local strain information

By running 2D-DDD simulations, we acquire strain information at the L,U and T stages. In order to remove the prior memory, we form the quantity $T-U = \epsilon_T(\vec{r}) - \epsilon_U(\vec{r})$, or ϵ_{TU} , which is the testing deformation in Fig. 2. This process is similar to polishing a sample, applying speckles and tracking their movement as the sample is further deformed, which is naturally similar to DIC [31–38].

After removing the prior strain, it is not straightforward to characterize the plastic behavior of the samples without prior knowledge (i.e. the degree of plasticity incurred from L and T stages). For example, in Fig. 6, without prior knowledge we would not know that the samples in (a),(b) are loaded to 10 % strain while in (c),(d) the samples are loaded to 1 % strain. The figures appear to be quite different, and the similarity of their histories is not recognizable by eye. However, ML’s trained eye is able to detect the initial deformation history. Indeed, in later sections we will show how ML algorithms can show that the figures are quantifiably and fundamentally different. With the help of ML we are able to find the initial deformation history of various samples, as long as the testing deformation does not overwrite it.

We compute the strain ϵ_{TU} on a grid of 2000×500 points overlaid on the 320×80 finite element mesh, using finite element interpolation to compute the strain at each point. At each point we then construct a scalar, the determinant of the deviatoric total strain, $\phi \equiv \frac{1}{2}(\epsilon_{xx}^2 + \epsilon_{yy}^2) - \epsilon_{xx}\epsilon_{yy} + 2\epsilon_{xy}^2$, where the tensor ϵ is the total strain. The choice of the total strain was based on making contact with what is easily accessible in experimental settings; nevertheless, other choices such as the plastic strain (*cf.* Appendix E [41], Fig.E9) do not appear to qualitatively influence the conclusions of this work. The scalar ϕ can be interpreted as a color to form a strain profile image, such as those shown in Figs. 4, 6. These strain profiles are then put through a correlation algorithm, following the scheme of the Materials Knowledge System (MKS) [42]. It is worth noting that the choice of the deviatoric ϕ -invariant satisfies local rotation invariance but is not unique in any way; as we show in Sec. III F and also in Appendix E [41] (*cf.* Fig. E9), our results do not depend on the par-




FIG. 6. **Variety of strain profiles in 2D-DDD simulations for smaller and larger testing strain:** For the spatial scale of the figures see Fig. 4. A sample is loaded to a high deformation of strain (which could be either 1 % or 10 %) and then unloaded. This predeformed sample is then reloaded to a testing strain. In (a),(b) the samples are loaded to 10 % strain, while in (c),(d) the samples are loaded to 1 % strain. (a) Small testing deformation (0.1 %). $w=2 \mu\text{m}$. Double slip system simulation. (b) Large testing deformation (1 %). $w=2 \mu\text{m}$. Double slip system simulation. (c) Small testing deformation (0.1 %). $w=1 \mu\text{m}$. Single slip system simulation. (d) Large testing deformation (1 %). $w=1 \mu\text{m}$. Single slip system simulation. For description of color map see Fig. 4 (b).

ticular choices made.

In the general MKS scheme [43], one selects spatially-varying quantity(-ies) which characterize the microstructure. The space of all possible values of these quantities is called the local state space, \mathcal{H} , and a point in this space is denoted n . Some care with the vocabulary is required, since physically speaking, these quantities are simply the values of fields of interest to us, and may or may not correspond to thermodynamic state variables. In this study, our quantity is the determinant of the deviatoric total strain invariant.

In the MKS method, one further considers a “microstructure function”, defined on the product space of the microstructure state variables \mathcal{H} , and physical space x , $m(n, x)$. In general use, this function may be thought of a probability density on these spaces. In our case, where we have a succession of particular microstructural instances, the microstructure function corresponding to each instance is a delta-function in n at each point in space.

In order to obtain data suitable for constructing two-point correlations, it is necessary to bin the state variables. We make use of the PyMKS software [44] which offers tools to accomplish this. The most basic n -axis discretization scheme is




FIG. 7. $w = 2 \mu\text{m}$ – 2D projection of PCA results for thin films – Double slip system: n_0, n_0 auto-correlation. (a) Projection of data set on first two principal components. Red blobs denote samples with 0.1 % strain (stage L), blue triangles samples with 1 % strain (stage L) and green squares denote samples with 10 % strain (stage L), respectively. (b) First principal component of PCA, shown in sample coordinates (Fig. 3, Sec. III B). (c) Second principal component of PCA, shown in sample coordinates (Fig. 3, Sec. III B). The colormaps are unitless, showing the intensity the PCA-transformed correlations.

the so-called “primitive basis” scheme, in which one selects some number P of evenly-spaced levels, n_0, n_1, \dots, n_P , and, at a point in space where the state variable has value n , selects amplitudes ω_i for these levels such that $\sum_i \omega_i n_i = n$, with the additional restriction that only the n_i ’s which are directly below and directly above the local value n are nonzero, and $\sum_i \omega_i = 1$. The entire system is thus described by a set of values $\{\omega_i\}$ in each spatial point x . In our simulations we discretize the state space into 3 different bins, corresponding to 3 local states n_0 , n_1 , and n_2 at low, intermediate, and high local strains. 6 possible correlations are observed for $(i, j) \in Z$, where space Z is defined by the values of $(i, j) = ((0, 0), (0, 1), (0, 2), (1, 1), (1, 2), (2, 2))$ for the 3 local states. The 2-point correlations are:

$$C^{[k]}[r|l, l'] = \frac{1}{S} \sum_s m[s, l] m[s + r, l'], \quad (1)$$

where $C^{[k]}[r|l, l']$ is the conditional probability of finding the local states l and l' at a distance and orientation away from each other defined by the vector r , for the k^{th} sample. S is the total number of spatial cells in the microstructure and s is a specific spatial cell. When the 2 local states are the same $l = l'$, the correlation is called an auto-correlation. If the 2 local states are not the same, it is a cross-correlation.

III. PRE-PROCESSING, CLUSTERING AND CLASSIFICATION FOR STRAIN PROFILES OF CRYSTALLINE THIN FILMS

In the model, signatures of plastic deformation are concentrated in collective features of the 2-




FIG. 8. Efficiency of ML compared to visual inspection of results – $w = 2 \mu\text{m}$ – Double slip system: n_0, n_0 auto-correlation. The colors follow the definition of Fig. 7. (a) Strain profiles obtained through our model. The top figure corresponds to strain profiles for a sample with stage L at 0.1 % strain. The middle figure is a sample strain profile obtained from stage L at 1 % strain. Finally, the bottom figure is a sample strain profile from stage L at 10 % strain. All strain profiles show the strain localizations formed from the quantity ϵ_{TV} (see Sec. II B). For the spatial scale of the figures see Fig. 4 (b) PCA map for samples that have similar strain profiles as (a). 3 distinct clusters are formed. The projection is upon the first and second principal components. For description of color map see Fig. 4 (b).




FIG. 9. $w = 0.5 \mu\text{m}$ – 2D projection of PCA results for thin films – Double slip system: n_0, n_0 auto-correlations. The colors follow the definition of Fig. 7. (a) Projection of data set on first two principal components with a clustering algorithm applied to the data set, demonstrating a failure in clustering the various deformation levels. (b) Projection of data set on first two principal components without a clustering algorithm applied to the data set, justifying (a). (c) First principal component of PCA, shown in sample coordinates (Fig. 3, Sec. IIIB). (d) Second principal component of PCA, shown in sample coordinates (Fig. 3, Sec. IIIB). For description of colormaps, see Fig. 7

point correlation function (see Eqn. 1, Sec. IIB). In order to identify and classify these collective features, a statistical approach needs to be implemented in a multitude of training samples. Similar work has been examined in materials science in the past [43, 45–47]. We use Principal Component Analysis (PCA, [48]) as a dimension-reduction scheme to pick out relevant axes in a high-dimensional space, and then we do classification on the points projected along these important axes. PCA takes a cloud of points in a high dimensional space and computes the orthogonal directions (“principal components”) in that space in which the cloud has the largest variance. Projecting the points onto the subspace defined by the principal components allows the points to be distinguished from one another succinctly.

The inputs to our ML algorithm are the correlation functions of the MKS local states discussed in Sec. IIB. Having computed the correlation functions, we now wish to see if ML can extract prior histories from them. Our ML workflow will: (1)

convert correlation functions to vectors, (2) find the significant features of the vectors by PCA, and (3) apply a clustering algorithm to identify samples with similar histories.

The correlation functions are evaluated at displacements $r = (r_x, r_y)$ in a box around $(0, 0)$. The range of the correlation function can be limited by choosing the size of the box. We usually use a 40×40 square. By assigning integers v to each r , we can convert the 2-dimensional set of points to a list, and thus convert the correlation function to a high dimensional vector:

$$\mathbf{d}_{ij} = (C^{[k]}[r_1|n_i n_j], C^{[k]}[r_2|n_i n_j], \dots, C^{[k]}[r_q|n_i n_j]) \quad (2)$$

Here, q is the total number of points in the box, (i, j) label the MKS local states from which the correlation function was computed, and k labels the samples (*i.e.* the simulation run). Thus the Stage T strain from each simulation run has been mapped to a point in a high dimensional space. Our goal is to see if different Stage L strains show up as clusters in this space.

Alternatively, we can remove some of the high-resolution information, by selecting small 5×5 , or bigger 20×20 , squares on the overlaid interpolated mesh (Sec. IIB). In these squares, we average the strain information, thus “blurring” the images. By averaging this information, we also limit the spatial resolution of the images and we can examine longer range correlations (see Appendix F of the SI [41]) instead of the 40×40 short range correlations.

At the end we have a matrix \mathbf{D} in the form $n \times m$, with n rows, where n is the number of statistical samples. Each row contains the vector \mathbf{d}_{ij} which may or may not be truncated. The matrix \mathbf{D} has m columns, where m is the number of spatial correlation instances:

$$\mathbf{D} = \begin{bmatrix} C^{[1]}[r_1|n_i n_j] \dots C^{[1]}[r_m|n_i n_j] \\ \vdots \\ C^{[n]}[r_1|n_i n_j] \dots C^{[n]}[r_m|n_i n_j] \end{bmatrix} \quad (3)$$

The rows of \mathbf{D} are the data vectors on which PCA operates, and the resulting principal components are linear combinations of the basis vectors of this set ([48], also see Appendix D of the SI [41]).




FIG. 10. $w = 1 \mu\text{m}$, The choice of the correlation domain and how it impacts the PCA maps – **Double slip system** : n_0, n_0 auto-correlation. For description of colormaps and colors of PCA maps, see Fig. 7. Projection of data set on first two principal components. (a) 40×40 domain of correlation matrix. Highly smooth in the center and towards the boundaries of the domain. (b) 100×100 domain of the correlation matrix. A highly focused area near the center of the domain is shown, where the phenomena are focused. The smoothness present in (a) is slowly removed from this domain. (c) 200×200 domain of correlation matrix. We have rich phenomenology present towards the center of the correlation matrix and at the boundaries. (d) PCA maps for 40×40 domain. (e) PCA map for 100×100 domain. The variance of the data has changed and the projections have shifted. The information provided by (b) does not change the cluster formations, but introduces unnecessary information that has shifted the results along the PC1 and PC2 axes. (f). PCA map for 200×200 domain. The variance of the data has changed even more compared to (e). The distances between the blue and green clusters have increased an order of magnitude compared to (e) and 2 orders of magnitude compared to (d). The information provided by (c) does not affect the clusters that are formed from our algorithm. For description of correlation domains, see Sec. III B

A. Clustering and classification

We use the Continuous k-Nearest Neighbors (CkNN) algorithm [49] to classify samples after running PCA on the data set. The CkNN algorithm is a clustering algorithm, with the advantage that the number of clusters is not arbitrarily defined by the user, as in K-Means clustering [50], but is calculated through a distance based approach. In particular, CkNN recognizes samples that are close to each other and calculates the most probable number of clusters for the data set. After the number of clusters has been found, the algorithm classifies the samples similar to the K-Means approach. The algorithm is an unsupervised method that detects natural clusters within

a data set, and our interest in it is the degree to which the natural clusters correspond to the prior deformation (stage L). The input to this algorithm must consist of a set of points, which in our case is the projections of the correlation matrix on the 3 principal components. As an output, the algorithm produces the classified samples, based on the cluster to which they belong.

The size of the data set (see Sec. II A), as in most classification algorithms, imposes a limitation on the algorithm. The algorithm groups data samples with similar PCA vectors into one cluster. We find that the algorithm works better for larger data sets. This introduces a limitation on classification, especially for single slip systems. The method is successful if the samples with different prior load-




FIG. 11. $w = 1 \mu\text{m}$, 3D projection of PCA results for thin films – Double slip system : n_0, n_0 auto-correlation. The colors follow the definition of Fig. 7. 3 different clusters are shown like in Fig. E3. Introducing the 3rd component into the PCA map, does not affect the results.

ing are grouped into different clusters. Note that the clustering is done in three dimensions using all three principal components, but most of our plots are two dimensional, which can sometimes hide the degree of clustering.

In Figure 11 we show a 3D PCA map for material samples of $w = 1 \mu\text{m}$. It is obvious that the clustering isn't affected by the 3rd dimension, and in this case the information provided by PC3 is irrelevant to our results.

The results shown on this paper, except for results shown in section IIID, are extracted by applying PCA and the CkNN algorithm to the whole data set. The same PCA and CkNN steps are applied to all simulations. The remainder of this paper discusses how well the clustering algorithm works in various situations.

Fig. 8 (b) and Fig. 12 (b)




FIG. 12. $w = 2 \mu\text{m}$, – Strain profiles and 2D projection of PCA results for thin films – Single slip system : n_0, n_0 auto-correlation. The colors follow the definition of Fig. 7. Projection of data set on first two principal components. (a) Different strain profiles are seen for a single slip system of $w = 2 \mu\text{m}$. The top figure is a sample's strain profile with stage L at 0.1 % strain. The middle figure is a sample's strain profile obtained from stage L at 1 % strain. Finally, the bottom figure is a sample's strain profile from stage L at 10 % strain. Strain localization are formed from the quantity ϵ_{TV} (see Sec. IIB). For the spatial scale of the figures see Fig. 4 (b) PCA projection of of our results for samples of $w = 2 \mu\text{m}$. The similarity between strain profiles at 1 and 10 % strain does not affect the formation of separate clusters for samples that were initially loaded at these strains. For description of color map see Fig. 4 (b).

B. Distinguishing plasticity regimes for small testing deformation (0.1 %)

We ran a multitude of tests for different w . For large w ($> 0.5 \mu\text{m}$), our algorithm correctly clusters and classifies data into 3 different groups, one for each of the prior strain values, which was the main objective of our work. Figure 7 (a) shows that clustering is easily observed for $w = 2 \mu\text{m}$, where 3 distinct clusters appear in the PCA of the n_0, n_0 autocorrelation. It is clear that there is enough cluster separation to reliably classify plastically deformed metals into heavily deformed and less deformed categories. For these larger sized systems the CkNN algorithm has 100 % accuracy, but for smaller sized systems with $w \leq 0.5 \mu\text{m}$ the clustering algorithm fails to cluster data points according to their deformation state. That is evident in Fig. 9 (b), where one can see what a correct clustering and classification would look like for specimens of $w = 0.5 \mu\text{m}$. In figure 9 (a) one can observe




FIG. 13. $w = 2 \mu\text{m}$ – Comparison of principal components among double and single slip systems: Components shown correspond to the analysis of Figs. 8 (b), 12 (b). (a) First principal component of double slip system. (b) First principal component of single slip system. (c) Second principal component of double slip system. (d) Second principal component of single slip system. For description of colormaps, see Fig. 7.

the results after the CkNN algorithm is applied to the data set. Other figures in Appendix E.1 of the SI [41] show how specimens of various sizes are classified using the CkNN algorithm. The plastic noise fluctuations in the system, as well as the finite size of the system, interferes with the classification of smaller sized data samples, while for larger w the samples are classified correctly.

Figures 7 (b) and (c) show the representation of the first two principal components of the data matrix \mathbf{D} for samples of $w = 2 \mu\text{m}$, shown in their natural sample coordinates (*i.e.*, the PCA vectors have been converted back to the 2D grid representation of a correlation function, see Sec. III). If two correlation functions were randomly chosen from the data set, the difference between them would most likely look like 7 (b) (with some scaling) mixed with a smaller amount of 7 (c). Note that the first principal component is roughly isotropic, while the second is strongly anisotropic. Figures 7, 9 show the progression of our ML work-flow as sample width decreases. We can observe that the first PCA component at larger w is relatively isotropic.

While in Fig. 7 (b) we notice a concrete isotropy of the first principal component of the analysis, it gradually becomes anisotropic as the sample width

decreases (Fig. 9 (b)). This change is correlated with the onset of stochastic fluctuations at small scales and mechanical annealing [51] that promotes concrete slip bands even at small testing strains. While both principal components for $w = 2 \mu\text{m}$ (Fig. 7 (b,c)) are smooth, they gradually become less structured as w decreases (Fig. 9 (b,c)), naturally an effect of stochastic fluctuations at small length scales. For $w = 2 \mu\text{m}$ there is a distinct difference between the first and second principal components, related to a spatial symmetry breaking. This distinction disappears as w decreases. For smaller w , due to the emerging crystal plasticity size effects [29], the data set is not as distinguishable as we would have wanted with our clustering technique, because of the noise associated with strengthening. (Fig. 5 (a)).

The area of the correlations, with respect to the sample area can be calculated by: $A_{corr} = N_x \times N_y \times (w^2/500^2) \mu\text{m}^2$, where N_x, N_y is the number of nodes in x,y directions respectively. For example, for short range (40×40) correlations: $A_{corr} = 40 \times 40 \times (w^2/500^2) = (4/625) \times w^2 \mu\text{m}^2$. Figure 10 shows how A_{corr} can influence the results. Differences can be observed in Figs. 10 (d), (e), (f), only with respect to the variance of the projected points. The principal components in (a), (b), (c) have small differences, mostly on their intensity. We deduce that our results do not depend on the examined area of correlations, and in order to reduce computational resources and time, we examine short (40×40) range correlations, from the center of the sample.

One deficiency of our ML work-flow emerged as we examined the results: as w decreases, the distance between the PCA-transformed samples also decreases. It is known that classification algorithms have an inherent limitation: when the distance between points in one cluster is similar to the distance separating two clusters, then the algorithm has difficulty distinguishing the clusters. In particular, Fig. 7 shows that the cluster distances in the PC1 direction are of order of magnitude 10^{-2} to 10^{-1} . For $w \leq 0.5 \mu\text{m}$ (see Fig. 9 a) the cluster between PCA-transformed samples is on the order of 10^{-3} to 10^{-2} , similar to the distance between the samples itself, and the data samples cannot be classified correctly. For smaller systems, it is evident that samples with stage $L = 0.1 \%$ or 1% strain (red circles and blue triangles, respectively) are so close to each other that the classifier regards them as belonging to the same cluster.

As mentioned in sections I, we were inspired by experimental techniques. For this reason, we have considered strong statistical variations in the initial microstructures. Frank-Read sources (see Sec. II A) are distributed randomly with a random nucleation stress. Obstacles (mimicking precipitates) are also distributed randomly, with a random resistance stress. This variability causes strong noise and limited spatio-temporal resolution (as can be seen for samples of $w < 0.5 \mu\text{m}$, Appendix B of the SI [41]). Furthermore, this noise propagates into PCA maps where the variance for samples loaded to 1 or 10 % strain is very high. However, these variations do not affect the successful application of the ML work-flow, and this is one of the main findings of this work. Nevertheless, in order for our work to be comparable to experimental data we need to limit the resolution of the examined strain profiles as well (since in experiments, typical image resolution can reach $\approx 1 \mu\text{m}$). Our generated profiles originally have the nanoscale resolution of the 2D-DDD grid. In Appendix F of the SI [41], we show that through averaging out the slip band information, and moving to more realistic resolutions (factor of 10-20 lower), we can still conclude that strain correlations can reveal the deformation history.

C. Distinguishing plasticity regimes for single slip samples

As mentioned in Sec. II A, we model single and double slip systems. So far, we have shown how emergent shear bands can be observed in our simulations for both of these systems (Fig. 6), as well as PCA results for double slip (Figs. 7, 9). PCA results for single slip are consistent with double slip, as shown in Fig. 12. Specifically, Fig. 12 shows results of single slip system simulations for $w = 2 \mu\text{m}$. The clustering properties for these larger sized systems are similar to the properties observed for similar systems for double-slip simulations. A comparison between the results of single slip and double slip systems for samples of $w = 2 \mu\text{m}$ can be made by observing Fig. 8 (b) and Fig. 12 (b). The PCA results contain distinctly separated clusters.

Figure 13 compares the principal components for single and double slip systems.




FIG. 14. $w = 2 \mu\text{m}$ – 2D projection of PCA results for thin films – Double slip system – Validation: n_0, n_0 auto-correlation. Red blobs denote samples with 0.1 % strain (stage L), blue triangles samples with 1 % strain (stage L) and green squares denote samples with 10 % strain (stage L), respectively. Red stars depict *testing* samples of 0.1 % strain (stage L), blue stars *testing* samples of 1 % strain (stage L) and green stars *testing* samples of 10 % strain (stage L). Validated-split data set. Projection on first two principal components.

D. Validation and accuracy of the algorithm

An ML algorithm, in order to be considered successful, should be validated with “unknown” data sets (testing data) which have the same features as the data set the algorithm was designed for (training data). In many cases, testing data sets are hard to find, so the whole data set is split into two parts (not necessarily a half and half split), and the ML algorithm can be trained on part of the data set and its effectiveness tested on the rest. Other than this subsection, the results shown in the paper are an application of our ML work-flow on the whole data set (for a given w), and cannot be used to determine the validity of the classifier.

For validation purposes, we “trained” the algorithm by computing the PCA transformation from a randomly chosen half of the $w = 2 \mu\text{m}$ samples and applying the CkNN algorithm. Then we applied the PCA transformation to the remaining half of the samples and examined whether or not they were projected into the correct clusters. The results are shown in Fig. 14. It is evident that the testing data perfectly matches the training set. Similarly “training” the algorithm to samples of

various sizes (i.e., half of the samples instead of all the samples), follows the results of section II B. For samples with $w \geq 1 \mu\text{m}$ the “testing” data set is projected to the 3 classified clusters that have formed. In contrast, for smaller systems, the training data set is misclassified (as happens when examining the whole data set) and the testing data set falls within the misclassified results.

We can quantify the degradation of the clustering process using some of the tools provided in the scikit-learn metrics module [52]. In particular, we examine the accuracy score of the algorithm, as well as the F_β score. Accuracy is the fraction of samples that were classified correctly. We apply the CkNN algorithm and generate clusters. Because we know the prior strain for each sample, we can immediately check whether the clusters correspond to the strain levels. Perfect clustering is when each cluster contains only samples with identical prior strains. The results are summarized in Fig. 15. For $w \geq 1 \mu\text{m}$ the accuracy score is 1 as seen in Fig. 15 (a); that is, all the samples are correctly classified. For smaller samples $w \leq 0.5 \mu\text{m}$ (or $w/w_0 \leq 2^2$ as in the figure), we have a 0.33 accuracy score, because only the samples of one cluster are correctly classified. The accuracy score is not affected by the wrongly classified samples, and cannot provide a measure for the correct classification of individual clusters.

To quantify the performance of the classification process, we also use the F_β score [53, 54] which is computed separately for each cluster:

$$F_\beta = (1 + \beta^2) \cdot \frac{p \cdot r}{(\beta^2 \cdot p) + r} \quad (4)$$

where precision p is the number of correctly classified samples in the cluster divided by the number of all classified samples in the same cluster, and recall r is the number of correctly classified samples in the cluster divided by the number of samples that should have been in that cluster. The β number changes the weight of recall *vs* precision. For $\beta > 1$ recall is weighted more than precision, while for $\beta < 1$ precision is weighted more than recall. For $\beta = 1$, we have the F_1 -score, with precision and recall having the same weight in the equation. Fig. 15 (b,c,d) shows the F_1 , F_2 and $F_{0.5}$ scores for our results.

For samples with $w \geq 1 \mu\text{m}$ (or $w/w_0 \geq 2^3$) we have value of 1 on all scores and all clusters, but for smaller w we observe that the line with the

squares, which corresponds to samples with 10 % initial compressive loading returns non-zero values, varying as the β value changes. For samples that are classified in the cluster, we do not obtain the highest possible result, because the number of correctly classified samples is smaller than the number of samples in the cluster (i.e., the precision is small). The line with the circles, which corresponds to samples with 0.1 % initial strain loading, has value 0 for $w \leq 0.5 \mu\text{m}$ because no samples have been classified as belonging to that cluster. The last line, with the triangles corresponding to samples with 1 % initial loading has non-trivial values because in some cases there are some samples that are classified correctly (the recall and precision are very small). In summation: For the “square” cluster we have low precision but high recall, since we classify the samples that actually belong to that cluster correctly, but we also classify samples from other clusters; for the “triangle cluster” we have low recall and low precision, since we classify a small number of samples into that cluster.

We also tested the response of the algorithm with respect to accuracy and F_1 -score while changing the number of tested samples. Fig. 16 shows the algorithm’s reduced effectiveness when the number of samples is less than 20 % of our maximum. Fig. 16 (b) shows the average F_1 -score across the three clusters instead of the score for each cluster individually.

E. Distinguishing plasticity regimes for large testing (1 %) total strain

The results from our “large reload” data set, with 1 % testing strain, show that delicate handling is required to obtain the desired cluster separation. 1 % testing strain does not produce the clear separation obtained with 0.1 % strain. As the testing deformation increases so do strain localization features and shear band sizes. With a shear band spanning the whole specimen, we expect that the statistical correlations differ significantly from the statistical correlations of the “small-reload” data set. That is due to the overall effect of localization, from a structural correlation viewpoint. High loads lead to strain localization in the form of shear bands, which are inhomogeneous and anisotropic, unlike the low strain plastic response. Our methods pick up the transition between the two responses. Indeed, even in the case




FIG. 15. Measures of success for classification of samples - 0.1 % testing strain: n_0, n_0 auto-correlations. (a) Accuracy score for the samples. Maximum value 1 means that all the samples have been correctly classified. (b) F_1 -score of our 3 clusters that are formed. The line with the squares represents the cluster with samples at stage $L = 10$ % strain, while the line with the triangles is for the cluster with samples at stage $L = 1$ % strain. Finally, the line with the circles is for the cluster with samples at stage $L = 0.1$ % strain. For smaller sized systems we have observed that most of the samples are classified as belonging in the “square” cluster, hence the scored value for that cluster only. Since the algorithm correctly classifies the samples that were initially loaded to 10 % strain, but also classifies more samples as belonging to that cluster, then the score does not have the maximum value of 1 but lower. (c) F_2 -score of our 3 cluster that have formed. The definition of the colored lines follows (b). Since for F_2 -score we have increased weight of the recall, the 0.7 maximum value is expected for the “square” cluster. (d) $F_{0.5}$ -score of our 3 clusters. The colors definitions follow (b). Since we have reduced weight of the precision, for lower sample widths it is expected to have lower score than F_1 for the “square” cluster.

of low reload strain, the distance between clusters is small and in smaller systems (Fig. 9) the samples are unclassifiable.

Fig. 17 compares the small and large reload testing regimes. Fig. 17 (b) and (d) show the results of PCA with a clustering algorithm applied to the data set. From Fig. 17 (d) we can see that higher testing deformation renders samples indistinguishable in PCA coordinates. The separation that was present for the low testing deformation (0.1 %)




FIG. 16. Measures of success for classification of samples - 0.1 % testing strain: n_0, n_0 auto-correlations. The x-axis of each graph is percentage of samples tested for classification. (a) Accuracy score for samples of $w = 2 \mu\text{m}$ (stars) and $w = 1 \mu\text{m}$ (disks). Maximum value 1 means that all the samples have been correctly classified. (b) Averaged F_1 -score across the 3 clusters that have formed for samples of $w = 2 \mu\text{m}$ (stars) and $w = 1 \mu\text{m}$ (disks). It is obvious that we have good agreement for the classified samples even when we test less than 30 % of the total number of samples.

is missing for higher values. Figs. 17 (a,c) show the strain profiles captured when the sample is reloaded to low (a) and high (c) testing strains. It is obvious that for higher testing deformation there is much more mixing of the samples, thus the classification algorithm fails. Fig. 18 shows another difference between the two testing regimes. For large reload strain the first principal component (a) is highly anisotropic, while it becomes nearly isotropic at small reload strains (see Fig.9 (c)). This observation extends to other components (e.g., 2nd, Figs. 18 (b) and 9 (d)) and is correlated to the emergent anisotropy of strain localization. A more comprehensive comparison for these regimes can be found in Appendix E.3 of the SI [41].

F. Dependence of unsupervised learning capacity on pre-processing aspects.

As discussed in Sec. IIB, the discretization scheme defines the form and dimensions of the correlation functions to which we apply a PCA transformation. We can choose to examine correlations between different local states h . We can categorize samples based on their deformation history either for n_0, n_0 auto-correlations or n_1, n_1 auto-correlations. We find that cross-correlations aren’t helpful for classifying samples according to their




FIG. 17. **Large-reload vs. small-reload testing – Example of PCA projection results for thin films of $w = 0.5 \mu\text{m}$:** n_0, n_0 auto-correlations. The colors follow the definition of Fig. 7. Strain profiles are created from the quantity ϵ_{TU} (a) Sample with stage L = 10 % strain is unloaded to zero stress and then reloaded to small testing deformation (Stage T = 0.1 %). (b) Stage T at small testing deformation (0.1 %), without a clustering algorithm applied to data set. Projection on two principal components. Actual representation of the data set, with some mixing of the samples. The clusters have shifted closer to one another but not indistinguishable. (c) Sample with stage L = 10 % strain is unloaded to zero stress and then reloaded to large testing deformation (Stage T = 1 %). (d) Stage T at large testing deformation (1 %), without a clustering algorithm applied to data set. Actual representation of the data set. For the higher testing deformation of 1 %, we can see that there is much more mixing of the samples. Reloading to higher strain values, adds plastic memory to the samples, rendering our process inapplicable for these cases. For description of color map see Fig. 4 (b).

deformation levels. Fig. 19 shows results obtained from various correlation functions: in general, as w decreases, we observe that distances between each cluster are also decreasing. In particular, Fig. 19 (a) shows that the distances in each cluster are measured in an order of magnitude 10^{-4} to 10^{-3} while in Fig. 19 (b) the order of magnitude is 10^{-2} to 10^{-1} , similar to the one in Fig. 8 (b). This difference in Fig. 19 (a),(b) is enough for the clustering algorithm to find the different deformation levels and classify our data set with 100 % accuracy.

Another choice we can make is the quantity that characterizes the microstructure. Until now, we




FIG. 18. **First and second principal component of PCA application on thin films of $w = 0.5 \mu\text{m}$ shown in sample coordinates:** (a) First principal component - Stage T = 1 % (b) Second principal component- Stage T = 1 %. For description of colormaps, see Fig. 7.

considered an isotropic measure of the total deformation strain in the sample. Our classification scheme produces similar results if we use the more common 2nd invariant of the strain deformation tensor, $J_2 = \epsilon_{ik}\epsilon_{ki}$. Fig. 20 shows the results for different microstructural measure calculations. For larger systems ($w = 1, w = 2$) the only notable difference is the overall variance of the data in PCA coordinates.

Finally, we may use the plastic strain determinant as the microstructural deformation state variable, which effectively corresponds to examining the unloaded stage T dislocation ensembles. Instead of computing correlation functions and clustering on Stage T, we can unload the testing strain to create a new Stage S, and look for clustering there. Fig. 21 shows that classification still works and there is an observable difference of the data variance in PCA coordinates. For more figures on the differences in pre-processing aspects and the role of the plastic strain, please see Appendix E in




FIG. 19. **Auto-correlations vs. cross-correlations for pre-processing – Example of PCA projection maps for $w = 2 \mu\text{m}$:** The colors follow the definition of Fig. 7. $w = 2 \mu\text{m}$. (a) n_0, n_1 cross-correlations (b) n_1, n_1 auto-correlations.




FIG. 20. **Effect of strain invariant type for pre-processing – Examples of PCA projection maps:** The colors follow the definition of Fig. 7. (a) n_0, n_0 auto-correlations. $w = 2 \mu\text{m}$. Deviatoric ϕ -invariant. (b) n_0, n_0 auto-correlations. $w = 2 \mu\text{m}$. J_2 -invariant. (c) n_1, n_1 auto-correlations. $w = 1 \mu\text{m}$. Deviatoric ϕ -invariant. (d) n_1, n_1 auto-correlations. $w = 1 \mu\text{m}$. J_2 -invariant.

the SI [41].

G. Independence from the choice of discretization schemes and dimension reduction methods

We find that our protocol is not sensitive to reasonable variations of the microstructural binning of the local strain variable. As a test, we discretize the microstructure into $L = 2, 3, 4$ and 5 parts. We are able to distinguish the initial de-




FIG. 21. **Total vs. Residual/Plastic strain for pre-processing: Examples of PCA projection maps:** The colors follow the definition of Fig. 7. (a) Plastic strain. n_0, n_0 auto-correlations. $w = 2 \mu\text{m}$ (b) Total strain. n_0, n_0 auto-correlations. $w = 2 \mu\text{m}$




FIG. 22. **Effect of discretization schemes on pre-processing: Examples of PCA projection maps for $w = 1 \mu\text{m}$:** The colors follow the definition of Fig. 7. n_0, n_0 auto-correlations. (a) 2 local states (b) 3 local states (c) 4 local states (d) 5 local states

formation history of all the samples when calculating the n_0, n_0 auto-correlations and the n_1, n_1 auto-correlations. These results are independent of the discretization scheme (i.e. the number of local states used). Fig. 22 shows the results for data samples of $w = 1 \mu\text{m}$, as the number of local states L increases. Clustering and classification is possible, and the clustering algorithm has 100 % accuracy independently of the number of local states, but the overall noise of the data increases with the number of local states.

The noise is due to the use of a fixed number of DDD simulations for each prior strain level. The signal strength in each correlation function increases with system size and the number of dislocations, but decreases as the data is distributed into more bins L . This effect is more pronounced for cross-correlations because they decrease for short distances and our correlation function range is truncated. Hence we do not obtain classifiable results for any cross correlations.

While PCA is one of the most common and useful tools for dimensionality reduction, some data sets could be so large that it is impractical. With that in mind, we compared our PCA results with other common algorithms, such as Incremental Principal Component Analysis (IPCA) and the Truncated Singular Value Decomposition (TSVD). IPCA uses a different form of processing a data set that allows for partial computations which in

most cases match the results of PCA. Incremental PCA stores estimates of component variances and updates the variance ratio of a component incrementally. It is faster and uses memory more efficiently than PCA. TSVD on the other hand, implements a variant of singular value decomposition (SVD) that only computes the largest singular values. Given that PCA works on the basis of the singular value decomposition, we expect little to no difference with this method.

No significant differences are seen when applying some of these variations of PCA to our data sets. The results are shown on Fig. 23. Note that no additional parameters, other than the initialization of the different methods, have been modified; in particular, the same clustering algorithm is used as with the PCA methods. The TSVD results do not display any differences from regular PCA, besides slight changes in data variance and data cluster positions. The IPCA results, on the other hand, are mirrored from the PCA results in both the PC1 and PC2 axes (negative values). If we calculated the absolute values we would see just minor differences in data variance and cluster positions as in the TSVD results.

IV. REMARKS AND CONCLUSIONS

Our results could be generalized in a number of ways. Our work is applicable to thin films [28–30], but may also apply to more general families of materials. Firstly, the material parameters can change to correspond to composites and/or polymers, and their associated mechanical behavior during testing. For composites it is important to model and study ductile fracture while for polymers creep phenomena might be of interest. Secondly, an expansion can be made to the experimental protocol. Instead of examining uniaxial compression of thin films and their spatially resolved strain correlations, we could have included multi-cycle loading-unloading tests, multiaxial compression or nanoindentation. Thirdly, the data matrix \mathbf{D} (see Eq. 3, Sec. III) can be defined in different ways. While the protocol would have remained the same within the algorithm, we could have used geometrically necessary dislocations or local misorientations to calculate spatial correlations. In this particular case, the required data for the correlations would have been obtained, for example, by EBSD. In future studies, we will

examine data from theoretical solutions and aim to compare them with experimental data sets for dislocation-density related problems. A natural next step in our approach is a development of a regression method which can provide a continuous assessment of clustering and classification, and naturally provide error bars. Instead of using only 3 values of the applied strain at Stage L, we can use a continuous set of values, and apply regression based methods (eg. decision trees [55]) to identify features at each load.

In addition, there are some caveats of the approach that one has to be careful with: When samples used for ML have either been reloaded to high strain (1 %) or exhibit large noise due to their nano size ($w \leq 0.5 \mu\text{m}$), our classification method does not work. There are many possible reasons that the algorithm occasionally fails to identify these samples. For example, in the case of smaller w , short-range correlations may not be enough to distinguish the deformation history. Moreover, we use a simple ML work-flow, that may not distinguish features of the data matrix \mathbf{D} (see Sec. III, III A). Advanced ML protocols such as neural networks [56] or deep learning algorithms [57] could capture more information than correlation based approaches. The occasional failure of our methodology to distinguish prior deformation could also stem from basic aspects of the physical phenomenon of crystal plasticity at small scales: The data shows a substantial amount of noise at smaller widths [28–30] making classification occasionally unsuccessful (see Sec. III B, and Appendix E in the SI [41]), and at larger reload strain (see Sec. III E) the prior deformation history is overwritten [7] and becomes undetectable by the algorithm.

In summary, we examined the applicability of spatial correlations to practical and relatively inexpensive experimental methods for the detection of the degree of *prior* plastic deformation of thin films, especially when they display significant plasticity size effects. Our overall conclusion is that ML algorithms can achieve our objective with varying levels of success. Through mimicking experimental protocols with two dimensional discrete dislocation plasticity simulations, we identified realistic cases (single and double slip thin films with widths larger than $1 \mu\text{m}$) where data clustering and classification is possible, based on the degree of prior plastic deformation. When size effects come into play, we found that clustering and classifica-




FIG. 23. **Comparison of different dimensionality reduction methods for $w = 2 \mu\text{m}$:** The colors follow the definition of Fig. 7. n_1, n_1 auto- correlations. (a) PCA (b) TSVD (c) IPCA

tion becomes gradually more difficult, since the intrinsic, plasticity-induced crackling noise causes large variance in smaller systems. In general, for the success of our methodology for thin films, the physical size of the samples should exceed 500 nm in the lateral direction (see Sec. III A, III B), while the data set should consist of more than 50 samples (See Sec. III D). Furthermore, we uncovered a crucial parameter for the applicability of our methods, namely the testing total strain during reloading. The stage T reload strain should be small enough that it does not overwrite the prior deformation history of the samples; Reload strains less than 0.4% could be applicable for detecting deformation history. While for a small-reload level of 0.1% (half of the commonly defined engineering yield stress, found at at engineering strain 0.2%), our methods are highly successful (see Sec. III B), they are clearly not successful one order of mag-

nitude higher, at 1% (see Sec. III E). Another output of our calculations was that smoothing the short-range correlations and keeping features that refer to larger distances improved the capacity of our learning apparatus to reliably recognize material history (see Appendix F in the SI [41]).

V. ACKNOWLEDGEMENTS

We would like to thank D. Wheeler and S. Kalidindi for illuminating discussions. We would like to thank Marilyn Y. Vasquez Landrove for sharing the CkNN clustering code [49], and Erik Van Der Giessen for insightful comments on our work. We also acknowledge funding from Department of Commerce under award No. 1007294R (SP).

-
- [1] J. Diehl and F. Hinzner, *physica status solidi (b)* **7**, 121 (1964).
 - [2] V. Novák, S. Kadečková, B. Šesták, and N. Zárubová, *Crystal Research and Technology* **19**, 781 (1984).
 - [3] J. Richter, *physica status solidi (b)* **46**, 203 (1971).
 - [4] V. Novák, B. Šesták, and N. Zárubová, *Crystal Research and Technology* **19**, 793 (1984).
 - [5] L. Kaun, A. Luft, J. Richter, and D. Schulze, *physica status solidi (b)* **26**, 485 (1968).
 - [6] D. Veselý, *physica status solidi (b)* **29**, 685 (1968).
 - [7] R. Asaro and V. Lubarda, *Mechanics of solids and materials* (Cambridge University Press, 2006).
 - [8] P. M. Chaikin, T. C. Lubensky, and T. A. Wit-
ten, *Principles of condensed matter physics*, Vol. 1
(Cambridge university press Cambridge, 1995).
 - [9] S. Papanikolaou, E. Luijten, and E. Fradkin, *Physical Review B* **76**, 134514 (2007).
 - [10] K. Raman, E. Fradkin, R. Moessner, S. Papanikolaou, and S. Sondhi, in *Quantum Magnetism* (Springer, 2008) pp. 139–150.
 - [11] S. Papanikolaou, C. S. O’Hern, and M. D. Shattuck, *Physical review letters* **110**, 198002 (2013).
 - [12] E. C. Aifantis, *International Journal of Plasticity* **3**, 211 (1987).
 - [13] M. E. Gurtin, *Journal of the Mechanics and Physics of Solids* **48**, 989 (2000).
 - [14] N. Fleck, G. Muller, M. Ashby, and J. Hutchinson, *Acta Metallurgica et Materialia* **42**, 475 (1994).

- [15] V. Bulatov and W. Cai, *Computer simulations of dislocations*, Vol. 3 (Oxford University Press on Demand, 2006).
- [16] H. Cleveringa, E. V. Giessen, and A. Needleman, *Philosophical Magazine A* **79**, 893 (1999).
- [17] B. S. Bokstein, M. I. Mendelev, and D. J. Srolovitz, *Thermodynamics and kinetics in materials science: a short course*, Vol. 1 (Oxford University Press, 2005).
- [18] A. Cottrell, *An introduction to metallurgy* (Universities Press, 1990).
- [19] J. P. Sethna, K. Dahmen, S. Kartha, J. A. Krumhansl, B. W. Roberts, and J. D. Shore, *Physical Review Letters* **70**, 3347 (1993).
- [20] E. C. Aifantis, *Journal of Engineering Materials and technology* **106**, 326 (1984).
- [21] A. G. Evans and J. W. Hutchinson, *Acta Materialia* **57**, 1675 (2009).
- [22] M. E. Fisher (1969).
- [23] L. P. Kadanoff, *Journal of Statistical Physics* **137**, 777 (2009).
- [24] M. Meyers and K. Chawla, *Mechanical Behavior of Materials* (Cambridge University Press, 2009).
- [25] M. Kassner and T. Hayes, *International Journal of Plasticity* **19**, 1715 (2003).
- [26] D. Bigoni and T. Hueckel, *International Journal of Solids and structures* **28**, 197 (1991).
- [27] E. V. der Giessen and A. Needleman, *Modelling and Simulation in Materials Science and Engineering* **3**, 689 (1995).
- [28] L. Nicola, Y. Xiang, J. J. Vlassak, E. Van der Giessen, and A. Needleman, *Journal of the Mechanics and Physics of Solids* **54**, 2089 (2006).
- [29] S. Papanikolaou, H. Song, and E. Van der Giessen, *Journal of the Mechanics and Physics of Solids* **102**, 17 (2017).
- [30] H. Song, V. Deshpande, and E. Van der Giessen, *Proc. R. Soc. A* **472**, 20150877 (2016).
- [31] W. Peters and W. Ranson, *Optical engineering* **21**, 213427 (1982).
- [32] M. Sutton, W. Wolters, W. Peters, W. Ranson, and S. McNeill, *Image and vision computing* **1**, 133 (1983).
- [33] W. Peters, W. Ranson, M. Sutton, T. Chu, and J. Anderson, *Optical Engineering* **22**, 226738 (1983).
- [34] M. Sutton, C. Mingqi, W. Peters, Y. Chao, and S. McNeill, *Image and Vision Computing* **4**, 143 (1986).
- [35] H. Schreier, J.-J. Orteu, and M. A. Sutton, *Image correlation for shape, motion and deformation measurements* (Springer US, 2009).
- [36] A. D. Kammers and S. Daly, *Experimental Mechanics* **53**, 1743 (2013).
- [37] M. A. Sutton, N. Li, D. Joy, A. P. Reynolds, and X. Li, *Experimental mechanics* **47**, 775 (2007).
- [38] A. D. Kammers and S. Daly, *Experimental Mechanics* **53**, 1333 (2013).
- [39] Xiang, Y. and Vlassak, J.J., 2006, Bauschinger and size effects in thin-film plasticity, *Acta Materialia*, 54(20), 5449-5460.
- [40] Papanikolaou, S., Cui, Y. and Ghoniem, N., 2017. Avalanches and plastic flow in crystal plasticity: an overview. *Modelling and Simulation in Materials Science and Engineering*, 26(1), p.013001.
- [41] See Supplemental Material at [URL will be inserted by publisher] for further details of the numerical simulations.
- [42] T. Fast and S. R. Kalidindi, *Acta Materialia* **59**, 4595 (2011).
- [43] P. Steinmetz, Y. C. Yabansu, J. Hötzer, M. Jainta, B. Nestler, and S. R. Kalidindi, *Acta Materialia* **103**, 192 (2016).
- [44] D. Wheeler, D. Brough, T. Fast, S. Kalidindi, and A. Reid, *PyMKS: Materials Knowledge System in Python* (2014).
- [45] A. Choudhury, Y. C. Yabansu, S. R. Kalidindi, and A. Dennstedt, *Acta Materialia* **110**, 131 (2016).
- [46] Y. C. Yabansu, P. Steinmetz, J. Hötzer, S. R. Kalidindi, and B. Nestler, *Acta Materialia* **124**, 182 (2017).
- [47] A. Khosravani, A. Cecen, and S. R. Kalidindi, *Acta Materialia* **123**, 55 (2017).
- [48] J. Shlens, “A tutorial on principal component analysis: Derivation, discussion and singular value decomposition.(2003),” (2003).
- [49] T. Berry and T. Sauer, *arXiv preprint arXiv:1606.02353* (2016).
- [50] J. A. Hartigan and M. A. Wong, *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **28**, 100 (1979).
- [51] Z.-W. Shan, R. K. Mishra, S. A. Syed Asif, O. L. Warren, and A. M. Minor, *Mechanical Annealing and Source-Limited Deformation in Submicrometre-Diameter Ni Crystals*, Vol. 7 (2008) pp. 115–9.
- [52] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, *et al.*, *Journal of machine learning research* **12**, 2825 (2011).
- [53] D. M. Powers, (2011).
- [54] R. Baeza-Yates and B. Ribeiro-Neto, “Modern information retrieval: The concepts and technology behind search (acm press books),” (2011).
- [55] J. R. Quinlan, *Machine learning* **1**, 81 (1986).
- [56] C. Bishop, C. M. Bishop, *et al.*, *Neural networks for pattern recognition* (Oxford university press, 1995).
- [57] Y. LeCun, Y. Bengio, and G. Hinton, *nature* **521**, 436 (2015).

- [58] J. P. Hirth and J. Lothe, (1982).
- [59] S. R. Kalidindi, ISRN Materials Science **2012**, 13 (2012).
- [60] S. R. Niezgoda, D. T. Fullwood, and S. R. Kalidindi, Acta Materialia **56**, 5285 (2008).
- [61] D. T. Fullwood, S. R. Kalidindi, and B. L. Adams, in *Electron backscatter diffraction in materials science* (Springer, 2009) pp. 177–188.
- [62] S. R. Niezgoda, D. M. Turner, D. T. Fullwood, and S. R. Kalidindi, Acta Materialia **58**, 4432 (2010).
- [63] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, New York, NY, USA, 2007).