

Full length article

Nonstationarity Analysis of Materials Microstructures via Fisher Score Vectors

Kungang Zhang^a, Daniel W. Apley^{a,*}, Wei Chen^b^a Department of Industrial Engineering & Management Sciences, Northwestern University, Evanston, IL 60208, USA^b Department of Mechanical Engineering, Northwestern University, Evanston, IL 60208, USA

ARTICLE INFO

Article history:

Received 9 November 2020

Revised 1 March 2021

Accepted 15 March 2021

Available online 25 March 2021

MSC:

00-01

99-00

Keywords:

Microstructure

Fisher Score Vector

ABSTRACT

Microstructures are critical to the physical properties of materials. Stochastic microstructures are commonly observed in many kinds of materials (e.g., composite polymers, multiphase alloys, ceramics, etc.) and traditional descriptor-based image analysis of them can be challenging. In this paper, we introduce a powerful and versatile score-based framework for analyzing nonstationarity in stochastic materials microstructures. The framework involves training a parametric supervised learning model to predict a pixel value using neighboring pixels in images of microstructures (as known as micrographs), and this predictive model provides an implicit characterization of the stochastic nature of the microstructure. The basis for our approach is the Fisher score vector, defined as the gradient of the log-likelihood with respect to the parameters of the predictive model, at each micrograph pixel. A fundamental property of the score vector is that it is zero-mean if the predictive relationship in the vicinity of that pixel remains unchanged, which we equate with the local stochastic nature of the microstructure remaining unchanged. Conversely, if the local stochastic nature changes, then the mean of the score vector generally differs from zero. In light of this, our framework analyzes how the local mean of the score vector varies across one or more image samples to: (1) monitor for nonstationarity by indicating whether new samples are statistically different than reference samples and where they may differ and (2) diagnose nonstationarity by identifying the distinct types of stochastic microstructures that are present over a set of samples and labeling accordingly the corresponding regions of the samples. Unlike feature-based methods, our approach is almost completely general and requires no prior knowledge of the nature of the nonstationarities or the microstructure itself. Using a number of real and simulated micrographs, including polymer composites and multiphase alloys, we demonstrate the power and versatility of the approach.

© 2021 Acta Materialia Inc. Published by Elsevier Ltd. All rights reserved.

1. Motivation and Introduction

The physical properties of materials depend strongly on their microstructures. The discovery and design of materials with certain complex microstructures and superior or desirable properties are boosted by advances in fabrication and imaging techniques. In general, microstructures can be nonstationary in the sense that the nature of the microstructure varies across a single image or multiple image samples, due to variation in processing conditions, input materials, environmental conditions, etc. For example, state-of-the-art additive manufacturing (AM) gives unprecedented control of the microscopic phases of composite materials to customize the properties resulting in increasingly sophisticated multiphase microstructures [1]. After fabricating materials, advanced imaging

approaches can be used to non-intrusively and efficiently collect massive amount of 2D and 3D microscopic image data, examples of which include transmission electron microscopy (TEM) [2], scanning tunneling electron microscopy (STEM) [3], synchrotron-based tomography [4], magnetic resonance imaging [5], and confocal microscopy [6]. Recent literature shows that such advances have led to fast-growing image databases of complex microstructures of materials (also known as micrographs) [7–15], which in turn incentivizes further effort to design general, automated, and efficient workflows to process the micrographs.

An important problem for materials scientists and manufacturing/processing engineers is to automatically analyze nonstationarity of microstructure image samples for quality control purposes (e.g., to detect instability in the material processing that inadvertently results in changes in the microstructure and material properties) and, more generally, to more fully understand the nature of the material being produced [16–20]. By “nonstationarity”, we mean that the stochastic nature of the microstructure varies spa-

* Corresponding author.

E-mail address: apley@northwestern.edu (D.W. Apley).




Fig. 1. Examples of micrographs of microstructures. (a) A TEM image of silica particles in Polymethyl methacrylate (PMMA). (b) A SEM image of dual-phase steel [38]. (c)-(d) Two images of simulated microstructures.

tially, either within a single image sample or across multiple image samples. For example, Fig. 1b is an SEM image of steel showing two distinct phases, each of which corresponds to a distinct stochastic nature. And Fig. 1c and 1d show two simulated microstructure samples that, although stationary within each sample, are nonstationary across the two samples. In Sections 2 and 3, we more formally define nonstationarity in the context of a predictive supervised learning model fitted to the micrograph.

For microstructures whose behavior depends predominantly on clearly defined geometric features (e.g., particle size, volume fractions, particle distance, or simple inclusion shapes, etc.), one could directly apply a number of standard quality control monitoring methods [21] and/or profile monitoring methods [22–24]. However, such methods lack generality in the sense that they monitor only a specific set of predefined features and cannot detect more general changes in the microstructure, and they are often not applicable to common microstructures with more stochastic natures [25], for examples metals [26], polymer composites [27], ceramics [28,29], etc. Such stochastic microstructures do not have well-defined features that provide a complete characterization of the microstructure, as illustrated in Fig. 1.

Previous approaches to analyzing these types of stochastic microstructures, such as N -point correlation functions, lineal path functions, spectral density through Fourier transformation, and joint probability functions [30–35] use certain statistical features. These methods provide a richer representation of the stochastic nature than geometric and descriptor-based features, but have serious limitations [36]. First, the statistical features are only summary statistics for the full joint distribution of the stochastic microstructure and do not provide a sufficiently complete representation for some complex microstructures. For example, a 2-point correlation function is an incomplete statistical description and can also be computationally prohibitive for large images [37], and multi-point correlation functions are even more computationally prohibitive while still not providing a complete characterization. In general, there is a fairly severe tradeoff that must be balanced between generality and completeness of the statistical representation on the one hand, versus tractability and computational feasibility achieved by working with only a subset of the statistical features on the other hand; and balancing this tradeoff is complicated by the fact that the most appropriate subset of features is always material-dependent. Second, the most common statistical features are global features in the sense that they are computed over an entire image sample (or at least a sufficiently large subregion of the sample) and therefore cannot capture local behavior that varies on a finer scale.

In contrast, the score-based framework that we introduce in this paper is almost completely general and, in theory, entails a complete statistical representation of even the most complex stochastic microstructures. It also provides a pixel-by-pixel mea-

sure of nonstationarity and can monitor fine-scale local changes in the microstructure. Finally, because the score vectors are automatic by-products of the predictive model training, the computational expense is reasonable.

One component of our framework uses the approach of [39] and [16,17] to characterize the stochastic nature of microstructures. More specifically, the conditional distribution of each pixel value given its neighboring pixels is approximated by training a supervised learning model to predict the pixel value. In other words, the extremely high-dimensional joint distribution of the stochastic microstructure is implicitly represented by a trained supervised learning model for predicting an individual pixel value given its neighbors. For monitoring for global changes in the stochastic nature from sample to sample, [16] trained a separate supervised learning model on each image sample and then used a form of generalized likelihood ratio test to detect changes (i.e., nonstationarity) in the stochastic nature from sample to sample. However, their approach cannot monitor or diagnose nonstationarity within a sample, in addition to being computationally expensive due to the need to fit a separate supervised learning model for each sample. For detecting local anomalies or defects, [17] fitted a single supervised learning model to a reference image sample. They then computed various statistics on the residuals of the predictive model (each pixel has a predicted value and a residual error) to detect local defects like voids, tears, etc. However, their approach cannot monitor or diagnose general nonstationarity beyond the presence of certain types of local defects.

In this study, we develop a powerful and versatile framework based on Fisher score vector concepts to analyze general nonstationarity within and across samples. The Fisher score vector is defined as the gradient of the log-likelihood with respect to the parameters of the fitted predictive model, and each pixel is associated with its own score vector. In Section 3, we discuss a fundamental property of the score vectors that provides the basis for our approach. Namely, under fairly general conditions, the score vector is zero-mean if the predictive relationship in the vicinity of the pixel remains unchanged from the training sample on which the predictive model was fitted, and it is nonzero-mean otherwise. Because we equate changes in the stochastic nature of the microstructure with changes in the predictive model that implicitly represents the stochastic nature, our framework for analyzing nonstationarity reduces to analyzing the local mean of the score vector as it varies spatially across one or more image samples. We apply these concepts to two different but related contexts in Sections 4 and 5. In Section 4, we develop a nonstationarity monitoring (NM) method, the purpose of which is to indicate whether new samples are stationary and statistically equivalent to some reference samples and, if they are not, at which spatial locations they

may differ. In Section 5, we develop a nonstationarity diagnostic (ND) method, the purpose of which is to identify the distinct types of stochastic microstructure behavior present over a single sample or set of samples and label each region of the samples according to the microstructure type. Using real (e.g., Figs. 1a and 1b) and simulated (e.g., Figs. 1c and 1d) material examples, we illustrate use of the framework and demonstrate its power and versatility.

2. Background on Characterizing the Stochastic Nature of Microstructures via Supervised Learning

In this section, we review a general approach for modeling the stochastic nature of microstructures with micrographs of materials [39]. For stochastic microstructures, we can view each microstructure image sample as a realization of some underlying spatial random process. Let $\mathbf{X} = [X_1, X_2, \dots, X_m]^T$ denote the concatenated vector of pixel values of a micrograph (or set of micrographs), where m is the number of pixels in the micrograph(s). Let $P(\mathbf{X})$ denote the joint distribution of $\mathbf{X}$. If it were available, the joint distribution $P(\mathbf{X})$ completely characterizes the stochastic nature of the microstructure sample $\mathbf{X}$, and distinguishing two statistically different samples amounts to distinguishing their joint distributions.

Directly modeling the distribution is of course computationally prohibitive. Two assumptions that reduce the problem to a more tractable one are the Markov locality property and stationarity [30,40] over the sample or at least over a subregion of the sample. The Markov locality property assumes that the conditional distribution of the i^{th} pixel value X_i , given some appropriate set of its neighboring pixels $\mathbf{N}(X_i)$, does not depend on the remaining pixels in the sample, i.e., that $P_i(X_i|\mathbf{X}_{-i}) \equiv P_i(X_i|\mathbf{N}(X_i))$, where $\mathbf{X}_{-i}$ denotes the vector $\mathbf{X}$ but excluding the entry X_i . Stationarity means that the conditional distribution is independent of the location of a pixel, i.e., $P_i(X_i = y|\mathbf{N}(X_i) = \mathbf{x}) \equiv P(X_i = y|\mathbf{N}(X_i) = \mathbf{x})$ is the same function of the tuple $(y, \mathbf{x})$ for all locations i in the region over which the image is stationary. Then, we can think of a stationary sample $\mathbf{X}$ as being a realization generated by the conditional distribution $P(X|\mathbf{N}(X))$ via a mechanism analogous to Gibbs sampling. In other words, the compact predictive model $P(X|\mathbf{N}(X))$ implicitly represents the stochastic nature of the micrograph $\mathbf{X}$ and can be treated as the “fingerprint” of a microstructure. Modeling this conditional distribution by fitting some supervised learning model to predict X given $\mathbf{N}(X)$ is a tractable problem, and if we can effectively learn this model through image sample data, we can detect and analyze changes in the stochastic nature of the microstructure via detecting and analyzing changes in the conditional distribution $P(X|\mathbf{N}(X))$. The remainder of the paper develops our approach for accomplishing this.

3. Fundamental Theory for Score-Based Nonstationarity Monitoring and Diagnostics

In this section, we first relate microstructure nonstationarity to stationarity of the parametric supervised learning model that represents $P(X|\mathbf{N}(X))$ in Section 3.1. Then, the main concepts behind our score-based framework for analyzing nonstationarity are discussed in Section 3.2. We discuss a spatial smoothing technique for estimating the local mean of the score vectors in Section 3.3. Throughout this section, we illustrate the statistical concepts behind our method with a simple and transparent simulation example with micrographs generated similarly to those depicted in Fig. 1c and 1d. The images were generated according to the following spatial autoregressive (AR) model to generate stochastic microstructures. The model is (for $i \in \{1, 2, \dots, m\}$):

$$U_{r_i, c_i} = c_0 + \sum_{r=0}^{l_g} \sum_{c=0}^{l_g} \phi_{r,c} U_{r_i-r, c_i-c} + \epsilon_{r_i, c_i}$$

$$\begin{aligned} \epsilon_{r_i, c_i} &\sim NID(0, \sigma_{AR}^2) \\ X_i &= h(U_{r_i, c_i}) \end{aligned} \quad (1)$$

where (r_i, c_i) are the row and column coordinates of the i^{th} pixel; c_0 is an intercept parameter; U_{r_i, c_i} 's are latent variables; l_g is the length scale of data generation window; $\{\phi_{r,c}\}_{r,c=0}^{l_g}$ with $\phi_{0,0} = 0$ are the AR coefficients; σ_{AR}^2 is the variance of i.i.d Gaussian noise random variable ϵ_{r_i, c_i} ; and the function $h(\cdot)$ takes a latent variable U_{r_i, c_i} and outputs a pixel value X_i . The simulated micrographs are generated with some randomly initialized edges and later those edges are cut off so that only stationary pixels remain. Note that the two micrographs shown in Figs. 1c and 1d are generated by two sets of coefficients and a nonlinear $h(\cdot)$ function and they appear similar to some real stochastic microstructures of materials, e.g., the sandstone micrographs in [41] and silica-filled rubber matrix micrographs in [39]. Denote by $\phi^{(CL)}$ and $\phi^{(M)}$ the 2D AR coefficients for generating micrographs for the CL-selection data (like Fig. 1c) and the monitoring data (like Fig. 1d) (which will be formally defined in Section 4), respectively. The parameters and configurations are: $c_0 = 0.1$; $\sigma_{AR} = 1$; the row-by-row concatenated AR coefficient vectors $[\phi_{0,0}^{(CL)}, \phi_{0,1}^{(CL)}, \phi_{0,2}^{(CL)}, \phi_{1,0}^{(CL)}, \phi_{1,1}^{(CL)}, \phi_{1,2}^{(CL)}, \phi_{2,0}^{(CL)}, \phi_{2,1}^{(CL)}, \phi_{2,2}^{(CL)}] = [0, 3.59e-01, 1.07e-01, 9.98e-03, -1.82e-03, 1.72e-05, 3.51e-01, 4.21e-02, 1.76e-03]$ and $[\phi_{0,0}^{(M)}, \phi_{0,1}^{(M)}, \phi_{0,2}^{(M)}, \phi_{1,0}^{(M)}, \phi_{1,1}^{(M)}, \phi_{1,2}^{(M)}, \phi_{2,0}^{(M)}, \phi_{2,1}^{(M)}, \phi_{2,2}^{(M)}] = [0, 3.59e-01, 1.07e-01, 9.98e-03, -1.82e-03, 1.72e-05, 3.12e-01, 4.21e-02, 1.76e-03]$; and $h(u) = u$. This data set will also be used in Fig. 5b and 6b in Section 6.1.2.

3.1. Representing Microstructure Nonstationarity via a Parametric Supervised Learning Model

As discussed in Section 2, the conditional distribution $P(X|\mathbf{N}(X))$ can be treated as a “fingerprint” of the stochastic nature of the microstructure. To model $P(X|\mathbf{N}(X))$, we fit a parametric supervised learning model to the data $\{(X_i, \mathbf{N}(X_i))\}_{i=1}^m$ from one or more training micrographs, which we will temporarily treat as stationary. The approximated (learned) conditional distribution is denoted as $P(y|\mathbf{x}; \theta)$, where θ is the vector of parameters of the supervised learning model (e.g., the set of weights for all nodes in a neural network), y is the value of a target pixel, and $\mathbf{x}$ are the values of its neighboring pixels. Consider grayscale images, which are a common form of microstructure image data. We model the conditional distribution as normal with mean $g(\mathbf{x}; \theta)$ and variance σ^2 , where $g(\mathbf{x}; \theta)$ is the parametric supervised learning model for predicting the mean of the pixel, and σ^2 is the variance of the residual errors of the predictions.

The normality assumption here only means that the conditional distribution of pixel X , given its neighbors $\mathbf{N}(X)$, is normally distributed. Intuitively, this means that the prediction errors $X - g(\mathbf{N}(X); \theta)$ are assumed normally distributed. Note that it does not imply that the marginal distribution of each pixel is normally distributed, or that the joint distribution of all pixels is multivariate normal, which are assumptions that would often be strongly violated in practice. Because of the nonlinearity of the supervised learning model $g(\mathbf{x}; \theta)$ as a function of $\mathbf{x}$, the joint and marginal distributions of all pixels can be far from normal even though their conditional distributions are normal. For our assumption of a normal conditional distribution, even if it is violated in practice, we believe the resulting nonstationarity monitoring statistics and algorithm still remain meaningful and valid. In this sense, we view the conditional normality assumption as only a tractable and convenient means to an end.

Under certain identifiability assumptions, it is reasonable to treat the parameter vector θ of a given parametric model as the

“fingerprint” of the microstructure. If we define $\theta_i, i \in \{0, 1\}$, as the “true” values of parameters for two micrographs or two regions within one micrograph, nonstationarity (stationarity) across the two regions can be represented as $\theta_0 \neq \theta_1$ ($\theta_0 = \theta_1$). There are a number of desirable aspects of this modeling procedure. First, the model provides a concise representation of the microstructure regardless of the size of the training data and can be easily applied to new image samples. Second, it can represent general stochastic microstructures and does not require or involve any microstructure-specific knowledge or features. Third, it provides a well-defined mathematical representation of nonstationarity of stochastic microstructures that is flexible in the sense that it can be used in conjunction with any parametric supervised learning model that suitably models the microstructure. Lastly, the score vectors (defined in Section 3.2) of each observation can be easily computed and are by-products of model training or fine-tuning via the popular stochastic gradient descent (SGD) or related algorithms, which results in our framework having reasonable computational expense. We elaborate on these aspects in subsequent sections. One should really view the functional form of $g(\mathbf{x}; \theta)$, along with θ , as the fingerprint. This is especially true if different microstructure phases in different regions are best represented with different functional forms of $g(\mathbf{x}; \theta)$. However, it would be difficult to extend our score-based approaches to allow multiple functional forms. Hence, we recommend choosing a functional form for $g(\mathbf{x}; \theta)$ that is sufficiently rich and flexible (e.g., a neural network) to represent the stochastic nature of all microstructure phases that are present.

Returning to the illustrative example, suppose one fit the linear regression supervised learning model

$$g(\mathbf{x}; \theta) = \theta^T \mathbf{x}, \quad (2)$$

to a set of micrograph training data (details and guidelines for fitting the model are discussed in Section 3.3). Although one would typically fit a more complex model like a neural network for $g(\mathbf{x}; \theta)$, it is helpful to illustrate the concepts with a linear model, since the subsequent computations become more transparent. Note also that for these AR simulation data generated via (1) with $h(u) = u$, the conditional distribution $P(X|\mathbf{N}(X); \theta)$ actually is Gaussian with conditional mean $g(\mathbf{x}; \theta) = \theta^T \mathbf{x}$. In the examples of Section 6, we consider more complex microstructures and nonlinear models for $g(\mathbf{x}; \theta)$, including real data examples and also a version of (1) with nonlinear $h(u)$.

In this study, we demonstrate our framework with grayscale microstructure images, for which predicting pixel values is a regression problem. The framework can be easily extended to classification problems which correspond to micrographs for which each pixel has been converted to categorical values that indicate to which phase the pixel belongs. In this classification setting, the supervised learning model directly produces $P(X|\mathbf{N}(X))$, which are the multinomial probabilities of the pixel being in each phase.

3.2. The Score Vector and Its Zero-Mean Property Under Stationarity

Our framework for microstructure nonstationarity analysis is inspired by recent work on score-based concept drift monitoring [42] for detecting temporal changes in predictive relationships with data collected over time. For a given model $P(X|\mathbf{N}(X))$, the score function/vector associated with an individual pixel X (having value y and neighborhood values $\mathbf{x}$) is formally defined as the gradient of the log-likelihood:

$$\mathbf{s}(\theta; y, \mathbf{x}) = \nabla_{\theta} \log P(X = y | \mathbf{N}(X) = \mathbf{x}; \theta) \quad (3)$$

where $P(X = y | \mathbf{N}(X) = \mathbf{x}; \theta)$ is our parametric conditional likelihood for an individual observation $(y, \mathbf{x})$ and ∇_{θ} is the gradient

operator with respect to the parameters θ . According to a fundamental property of score functions (Proposition 3.4.4 from [43] if certain regularity and identifiability conditions are met and if the parametric likelihood is the correct model with a true parameter vector θ , then the expectation of the score function evaluated at the true parameter is zero, i.e.,

$$E_{\theta}[\mathbf{s}(\theta; X, \mathbf{N}(X)) | \mathbf{N}(X)] := \int \mathbf{s}(\theta; y, \mathbf{N}(X)) P(y | \mathbf{N}(X); \theta) dy = \mathbf{0}. \quad (4)$$

In other words, assuming that the reference micrographs are realizations of the same stationary spatial random process whose conditional distribution is correctly represented by $P(y | \mathbf{x}; \theta)$, the expectation of the score vector for each pixel is zero. Note that the expectation in Eq. (4) is conditioned on a specific set of neighbor pixels $\mathbf{N}(X)$, and Eq. (4) holds for any such $\mathbf{N}(X)$. Consequently the unconditional expectation is also zero, i.e.,

$$E_{\theta}[\mathbf{s}(\theta; X, \mathbf{N}(X))] = \mathbf{0} \quad (5)$$

In real data sets, the expectation in Eq. (5) is replaced by the empirical mean, which should also be 0:

$$\begin{aligned} \hat{E}_{\theta}[\mathbf{s}(\hat{\theta}; X, \mathbf{N}(X))] &:= \frac{1}{m} \sum_{i=1}^m \mathbf{s}(\hat{\theta}; y_i, \mathbf{x}_i) = \mathbf{0}, \text{ where} \\ \hat{\theta} &:= \operatorname{argmax}_{\theta} \hat{E}_{\theta}[\log P(X | \mathbf{N}(X); \theta)] \\ &:= \operatorname{argmax}_{\theta} \frac{1}{m} \sum_{i=1}^m \log P(X = y_i | \mathbf{N}(X) = \mathbf{x}_i; \theta), \end{aligned} \quad (6)$$

and the operator $\hat{E}_{\theta}$ denotes a sample average over the training data $\{(y_i, \mathbf{x}_i)\}_{i=1}^m$ (here, y_i is the observed value of the i^{th} pixel, and $\mathbf{x}_i$ is the vector of observed values of all pixels in the neighborhood of the i^{th} pixel, and m is the number of pixels in the reference training micrographs that are enough far away from the image boundary to have full neighborhoods) generated under the true parameter vector θ , and $\hat{\theta}$ is the maximum-likelihood estimator (MLE) of θ for the training data.

Eq. (6) generally holds for training micrographs no matter whether the parametric conditional distribution (or model) is correct. That is because the estimated parameter vector $\hat{\theta}$ is the optimum solution when the empirical mean of the log-likelihood is maximized, the gradient of which is the empirical mean of the score vectors at $\hat{\theta}$, which will be the zero vector since $\hat{\theta}$ is the maximizer. When micrographs or micrograph regions are statistically different from the training/reference micrographs, the parameters $\hat{\theta}$ estimated for the reference micrograph generally no longer provide the best fit for the statistically different regions. In such regions the score vector, which is also collinear to the gradient vector of the log-likelihood for individual observations, is generally not zero-mean. This underscores the generality of the score-based framework described below for monitoring nonstationarity by monitoring for changes in the local mean of the score vectors, because when characterizing stochastic microstructures, we can choose from among a wide range of parametric models for the mean function $g(\mathbf{x}; \theta)$ that are flexible and convenient to work with, without requiring that the model exactly represents the true distribution. In our study in Section 6, we found that for complex materials, linear models often provide nonstationarity analyses that are nearly as effective as those for nonlinear models like neural networks, but at a much cheaper training cost.

Based on the rationale discussed in Section 2, define and represent nonstationarity as the change in the parameters of the conditional distribution $P(X | \mathbf{N}(X); \theta)$. This and Eq. (5) or (6) imply a general method for analyzing nonstationarity through monitoring the mean behavior of the score vectors defined in Eq. (3). More

specifically, as shown in [42] for monitoring for temporal nonstationarity, under fairly general conditions, when the true parameter vector changes from θ to a different vector θ' , the expected score vector in Eq. (5) or (6) differs from zero, i.e., $E_{\theta'}[\mathbf{s}(\theta; X, \mathbf{N}(X))] \neq \mathbf{0}$ or $\hat{E}_{\theta'}[\mathbf{s}(\hat{\theta}; X, \mathbf{N}(X))] \neq \mathbf{0}$. Since our formulation equates $\theta' \neq \theta$ in the vicinity of a pixel with the microstructure in the vicinity of that pixel having a different distribution than that of the training/reference microstructure, we can monitor and visualize the mean of the score vector (details for which are provided in Sections 4 and 5) to signal when statistically different microstructures are encountered and specifically where those differences are.

Returning again to the illustrative example with micrographs generated via (1) and fitted linear regression supervised learning model (2), the mathematical expressions for the conditional distribution and the score vectors are:

$$P(X = y | \mathbf{N}(X) = \mathbf{x}; \theta) = \frac{1}{\sqrt{2\pi}\sigma^2} \exp \left\{ -\frac{\|y - \theta^T \mathbf{x}\|_2^2}{2\sigma^2} \right\}$$

$$\mathbf{s}(\theta; y, \mathbf{x}) = (y - \theta^T \mathbf{x}) \mathbf{x}. \quad (7)$$

If the fitted model $g(\mathbf{x}; \theta) = \theta^T \mathbf{x}$ captures the conditional mean (i.e., $y = \theta^T \mathbf{x} + \epsilon$ with zero-mean ϵ), then after applying the model to newly collected micrographs that have the same stochastic behavior (i.e., generated by the same model with the same parameters $\theta' = \theta$), the expectation of the score vectors in Eq. (7) is $E_{\theta'}[\mathbf{s}(\theta; y, \mathbf{x})] = E_{\theta'}[(\theta'^T \mathbf{x} + \epsilon - \theta^T \mathbf{x}) \mathbf{x}] = E[\mathbf{x} \mathbf{x}^T](\theta' - \theta) = \mathbf{0}$ (a vector of zeros of the same length as θ). In contrast, if the newly collected micrographs have different stochastic behavior (e.g., follow the same conditional mean model $y = \theta^T \mathbf{x} + \epsilon$ but with different parameters $\theta' \neq \theta$), then the expectation of the score vector will differ from the zero vector.

3.3. Training the Model and Estimating the Local Mean of the Score Vector

To train a supervised learning model $g(\mathbf{x}; \theta)$ to represent microstructures, we first need to choose the neighborhood $\mathbf{N}(X)$ for each pixel X . As shown in Figs. 2a and 2b, there are two basic choices: a causal or a non-causal neighborhood window (excluding the target pixel colored as red). For stationary microstructures with the Markov locality property, either a causal or a non-causal neighborhood window can serve to implicitly characterize the joint distribution $P(X)$ of the pixels in the micrograph (the former via the decomposition $P(X) = P(X_1)P(X_2|X_1)P(X_3|X_1, X_2) \cdots P(X_m|X_1, \dots, X_{m-1})$, and the latter via a mechanism akin to Gibbs sampling). In general, the choice depends on the purpose of using the trained model. For computational reasons, in certain applications the causal neighborhood window is usually chosen if the goal is to generate new samples of microstructures with the trained model [39]. Also for some specific applications, a causal neighborhood window may result in better performance than a non-causal neighborhood window [16,17]. For our purposes within our score-based monitoring framework, we choose a non-causal neighborhood window, because it obviously provides more accurate prediction of the pixel values, and it resulted in better monitoring performance with roughly the same computational cost for our usage. Intuitively, a pixel X should depend on its neighbors in all directions, which suggests a non-causal neighborhood window will lead to more accurate prediction of X .

After choosing a proper neighborhood for each pixel, training a model is simply fitting the corresponding model to the data set $\{(X_i, \mathbf{N}(X_i))\}_{i=1}^m$. Regarding the size of the neighborhood window, this can be chosen empirically via cross-validation. Using the trained parametric likelihood (based on the trained supervised learning model $g(\mathbf{x}; \theta)$), the score vectors $\mathbf{s}(\hat{\theta}; X, \mathbf{N}(X))$ can be calculated for any set of image pixels. In SGD algorithms, the score

vectors are by-products produced during training or predicting (e.g., for training data, $\mathbf{s}(\hat{\theta}; y_i, \mathbf{x}_i) = \nabla_{\theta} \log(P(X_i = y_i | \mathbf{N}(X_i) = \mathbf{x}_i; \hat{\theta}))$ and are automatically computed as the gradient for $(y_i, \mathbf{x}_i)$), which is a main reason why our score-based approach is computationally reasonable.

After obtaining the score vectors, directly monitoring the individual score vectors would be an ineffective way to monitor mean changes, because the individual vectors are noisy. To handle the noise issue and estimate the local mean of the score vectors as it varies spatially across the image sample, we employ ideas similar to multivariate exponentially weighted moving average (EWMA) control chart concepts. A multivariate EWMA is of the most effective methods for monitoring for changes in the mean of general random vectors [21] over temporal or spatial domains. Monitoring for spatial nonstationarity (within or across samples) reduces to monitoring for changes in the mean of score vectors over regions of one sample or across multiple samples.

Similar to the temporal EWMA, we calculate the spatial weighted moving average (WMA) of the score vectors $\mathbf{s}(\hat{\theta}; y, \mathbf{x})$ to smooth out noise and estimate the local mean, where $\hat{\theta}$ is the MLE of the parameters computed over the entire training data. Specifically, we calculate the weighted average of the score vectors of pixels in a WMA window that is shown as the square with blue edges in Fig. 2c, with the weight function being a truncated 2D normal distribution centered at the target pixel over the WMA window (including the target pixel colored as red) as illustrated in Fig. 2d. The standard deviation σ_w (defined below) of this truncated normal distribution governs how fast this dependency decays away from a target pixel in different directions. Notice that to distinguish the two windows, one for training a model for conditional distribution (neighborhood window), $\mathbf{N}(X)$, and the other for WMA smoothing, we use different colors, shadings, and notations in Fig. 2.

To mathematically define the WMA, let $\mathbf{z}_{r_i, c_i}$ denote the score vector at a pixel X_i smoothed by the WMA window with row and column coordinates of X_i denoted by (r_i, c_i) ; and use (r, c) as the row and column coordinates of neighboring pixels over which we calculate $\mathbf{z}_{r_i, c_i}$. For notational simplicity, also define $\mathbf{s}_{r_i, c_i} := \mathbf{s}(\hat{\theta}; X_i, \mathbf{N}(X_i))$, and denote the truncated bivariate Gaussian density function by $p(r, c; (r_i, c_i, \sigma_w))$, which is viewed as a function of (r, c) with the distribution centered at (r_i, c_i) and having bivariate covariance matrix $\sigma_w^2 \mathbf{I}$ with $\mathbf{I}$ the 2×2 identity matrix. Our score vector WMA is defined as

$$\mathbf{z}_{r_i, c_i} = \sum_{r=r_i-l_w}^{r_i+l_w} \sum_{c=c_i-l_w}^{c_i+l_w} p(r, c; (r_i, c_i, \sigma_w)) \mathbf{s}_{r, c} \quad (8)$$

where the window size is $2l_w + 1$ as shown in Fig. 2c. Note that the weights $p(r, c; (r_i, c_i, \sigma_w))$ sum to one over the WMA window, by definition of the truncated distribution.

In all of our examples, we use $l_w = \sigma_w$. Regarding choice of the hyper-parameter σ_w , we recommend choosing it to be larger than the size of normal microstructure features that are present within stationary phases of the microstructure, but not so large that neighboring regions that have different microstructure characteristics (and that should be viewed as two distinct microstructure phases) will be smeared together. In general, using a larger σ_w will smooth out more noise and false alarms that result from the microstructure transitioning between features that are part of normal stationary single-phase behavior. This allows better detection of small-magnitude microstructure differences that are sustained over larger spatial regions; but it will also tend to smooth out localized variations that are best viewed as nonstationarities and that one would like to detect. In contrast, using a smaller σ_w will be more sensitive to detecting variation in localized characteristics that are truly indicative of nonstationarities; but it will also result




Fig. 2. Illustration of various windows used in computing and monitoring the score vectors. (a) Causal neighborhood window (excluding the red target pixel) for modeling the mean of $P(X|N(X))$, i.e., $g(\mathbf{x}; \theta)$ and (b) Non-causal neighborhood window (excluding the red target pixel) with the side length as $2l_s + 1$ pixels for modeling the mean of $P(X|N(X))$, i.e., $g(\mathbf{x}; \theta)$ (patches with brown dashed edges) for modeling the conditional distribution of X_i (the red pixel) given the neighboring pixels in those windows. (c) WMA window (including the red target pixel) with the side length as $2l_w + 1$ pixels for spatially smoothing score vectors. WMA window for spatially smoothing the score vectors to estimate their local mean. To differentiate it from the neighborhood window, we use blue solid lines as the edges here. (d) Truncated Gaussian distribution over the WMA window. A 2D Gaussian distribution centered (at the red pixel) and truncated over the WMA window. The height is proportional to the density value of the truncated distribution.

in a noisier $\mathbf{z}_{r_i, c_i}$ and will therefore less accurately detect small-magnitude but sustained differences. Proper choice of σ_w will inevitably involve some user subjectivity. Fortunately, it is computationally inexpensive to try different σ_w and visually compare the results of characterizing nonstationarity, which can be done after fitting and tuning the supervised learning model and computing the score vectors over the entire micrograph. We recommend that users do this and plot the monitoring statistics, which will be discussed in details in Section 4, on top of the micrograph image (e.g., as in Figs. 7 and 8) to visually assess the most appropriate σ_w . Complicating any purely quantitative choice of σ_w is the fact that the notion of microstructure stationarity is inextricably linked to the scale of interest. At a sufficiently fine scale (e.g., atomic or subatomic levels) and hypothetically small micrograph sizes, virtually all microstructures can be viewed as nonstationary; and at a sufficiently coarse scale (e.g., product level) and hypothetically large micrograph sizes, most can be viewed as stationary. For additional discussion on the relationship between nonstationarity, feature size, and window size/scale, see, e.g., [30,39,44–47].

For the same illustrative example with linear supervised learning model (2), training the model means simply fitting a linear regression model to estimate the regression coefficients θ . In this case, the score vectors for each pixel are calculated via Eq. (7), and then the mean of the score vector in Eq. (6) at any spatial location can be approximated by averaging the score vectors throughout the WMA window in Fig. 2d.

4. Nonstationarity Monitoring (NM)

The goal of this monitoring is to determine whether a given set of micrograph samples are statistically equivalent to a given reference sample or samples, where, for example, the reference

sample(s) represent normal conditions; and to potentially provide some interpretation if nonstationarity is detected. One practical problem falling into this category in materials fabrication or manufacturing is when we have available image samples from batches of materials that were produced under normal or well-calibrated conditions, and the goal is to determine whether subsequently produced batches of materials are statistically equivalent to the reference batch(es), based on image samples collected periodically from the subsequent batches.

In order to monitor for local changes in the mean of the score vectors via a control chart, we use a Hotelling T^2 [48,49] statistic for $\mathbf{z}_{r_i, c_i}$:

$$T_i^2 = (\mathbf{z}_{r_i, c_i} - \bar{\mathbf{s}})^T \hat{\Sigma}^{-1} (\mathbf{z}_{r_i, c_i} - \bar{\mathbf{s}}) \quad (9)$$

where $\bar{\mathbf{s}}$ and $\hat{\Sigma}$ are the empirical mean vector and covariance matrix of the training score vectors $\{\mathbf{s}_{r_i, c_i}\}_{i=1}^m$ over the entire training data, respectively. Notice that we first apply the WMA smoothing window and then calculate the Hotelling T^2 statistic for $\mathbf{z}_{r_i, c_i}$ as opposed to the Hotelling T^2 statistic for the individual score vector $\mathbf{s}_{r_i, c_i}$. This reduces the adverse effects of noise and gives a more accurate estimate of the local score vector mean, so that, with a proper choice of moving window size, the control limits (defined below) are tighter and the control chart is more sensitive to the mean change in score vectors.

In the supervised learning model $g(\mathbf{x}; \theta)$, if we view the residual standard deviation parameter σ as an additional parameter along with θ , then changes in θ and/or σ indicate nonstationarity of the microstructure. Consequently, we compute and monitor the score vector with respect to both θ and σ . Since changes in θ versus changes in σ represent different types of changes in the microstructure, we have found it more effective to compute a score vector for θ and a score vector for σ and treat them jointly but

distinctly, as described below. This is relative to computing the T^2 statistic for the score vector for (θ, σ) together.

As mentioned above, we approximate the conditional distribution of the i^{th} pixel value X_i given its neighboring pixels $\mathbf{N}(X_i)$ as a normal distribution:

$$P(X_i = y | \mathbf{N}(X_i) = \mathbf{x}; \theta, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left\{ -\frac{[y - g(\mathbf{x}; \theta)]^2}{2\sigma^2} \right\}$$

$$l(\theta, \sigma; y, \mathbf{x}) := \log(P(y | \mathbf{x}; \theta, \sigma^2))$$

$$= -\log(\sigma) - \frac{[y - g(\mathbf{x}; \theta)]^2}{2\sigma^2} + \text{Const.} \quad (10)$$

The score vectors for θ and σ , evaluated at the MLEs $\hat{\theta}$ and $\hat{\sigma}$ computed from the entire training data $\{y_i, \mathbf{x}_i\}_{i=1}^m$, are defined as

$$\mathbf{s}_{\theta,i} := \frac{\partial l(\theta, \sigma; y_i, \mathbf{x}_i)}{\partial \theta} \bigg|_{(\theta, \sigma) = (\hat{\theta}, \hat{\sigma})} = \frac{1}{\hat{\sigma}^2} [y_i - g(\mathbf{x}_i; \hat{\theta})] \nabla_{\theta} g(\mathbf{x}_i; \hat{\theta}) \quad (11)$$

$$s_{\sigma,i} := \frac{\partial l(\theta, \sigma; y_i, \mathbf{x}_i)}{\partial \sigma} \bigg|_{(\theta, \sigma) = (\hat{\theta}, \hat{\sigma})} = -\frac{1}{\hat{\sigma}} + \frac{[y_i - g(\mathbf{x}_i; \hat{\theta})]^2}{\hat{\sigma}^3}. \quad (12)$$

Since a fixed $\hat{\sigma}$ in Eqs. (11) and (12) drops out of the T^2 statistic in Eq. (9), it can be ignored, in which case monitoring for changes in the means of $\mathbf{s}_{\theta}$ and s_{σ} reduces to monitoring for changes in the mean of $[y - g(\mathbf{x}; \theta)] \nabla_{\theta} g(\mathbf{x}; \theta)$ and changes in the mean of $[y - g(\mathbf{x}; \theta)]^2$.

The goal is to detect local changes in the mean of $\mathbf{s}_{\theta,i}$ and/or the mean of $s_{\sigma,i}$. Analogous to Eqs. (8) and (9), we define the WMAs

$$\mathbf{z}_{\theta,r_i,c_i} = \sum_{r=r_i-l_w}^{r_i+l_w} \sum_{c=c_i-l_w}^{c_i+l_w} p(r, c; (r_i, c_i, \sigma_w)) \mathbf{s}_{\theta,r,c} \quad (13)$$

$$z_{\sigma,r_i,c_i} = \sum_{r=r_i-l_w}^{r_i+l_w} \sum_{c=c_i-l_w}^{c_i+l_w} p(r, c; (r_i, c_i, \sigma_w)) s_{\sigma,r,c} \quad (14)$$

and the T^2 statistics

$$T_{\theta,i}^2 = (\mathbf{z}_{\theta,r_i,c_i} - \bar{\mathbf{s}}_{\theta})^T \hat{\Sigma}_{\theta}^{-1} (\mathbf{z}_{\theta,r_i,c_i} - \bar{\mathbf{s}}_{\theta}), \quad (15)$$

where $\bar{\mathbf{s}}_{\theta}$ and $\hat{\Sigma}_{\theta}$ are the empirical mean vector and covariance matrix of the training score vectors $\{\mathbf{s}_{\theta,r_i,c_i}\}_{i=1}^m$. Because $\mathbf{s}_{\theta,i}$ is a vector and $T_{\theta,i}^2$ is intended to detect a mean change in $\mathbf{s}_{\theta,i}$ away from $\bar{\mathbf{s}}_{\theta}$ in any direction in the θ space, the control chart for $T_{\theta,i}^2$ has only an upper control limit (UCL_{θ}), the purpose of which is to signal a change if $T_{\theta,i}^2 > UCL_{\theta}$. Since $s_{\sigma,i}$ is a scalar, its WMA chart detects a change in its mean if either $z_{\sigma,r_i,c_i} < LCL_{\sigma}$ or $z_{\sigma,r_i,c_i} > UCL_{\sigma}$, where LCL_{σ} and UCL_{σ} denote its lower and upper control limits, respectively.

We use a multi-chart to monitor for mean changes in either $\mathbf{s}_{\theta,i}$ or $s_{\sigma,i}$. A change in either in the vicinity of pixel X_i indicates that the stochastic nature of the microstructure has changed in the vicinity of this pixel, relative to the reference microstructure samples. The multi-chart is defined as signaling either if $T_{\theta,i}^2 > UCL_{\theta}$ or if $z_{\sigma,r_i,c_i} < LCL_{\sigma}$ or $z_{\sigma,r_i,c_i} > UCL_{\sigma}$. To display the control chart results in a single 3D plot (e.g., as in Fig. 5, later), we define the plotted scaled statistics from the two component charts as

$$C_{\theta,i} = \frac{T_{\theta,i}^2}{UCL_{\theta}} - 1 \quad (16)$$

$$C_{\sigma,i} = \frac{z_{\sigma,r_i,c_i} - \frac{UCL_{\sigma} + LCL_{\sigma}}{2}}{\frac{UCL_{\sigma} - LCL_{\sigma}}{2}}, \quad (17)$$

which would individually signal if either $C_{\theta,i} > 1$ or $|C_{\sigma,i}| > 1$. The plotted statistic in the multi-chart statistic is then defined as

$$C_{M,i} = \text{sign}(C_{\theta,i} + C_{\sigma,i}) \times \max(|C_{\theta,i}|, |C_{\sigma,i}|), \quad (18)$$

and the multi-chart signals nonstationarity in the vicinity of pixel X_i if $C_{M,i}$ falls outside the range $[-1, 1]$.

The three control limits (UCL_{θ} , LCL_{σ} , and UCL_{σ}) are determined empirically, as follows. We divide the reference data $\{(y_i, \mathbf{x}_i)\}_{i=1}^M$ into two sets $\{(y_i, \mathbf{x}_i)\}_{i=1}^m$ (which we refer to as the training data) and $\{(y_i, \mathbf{x}_i)\}_{i=m+1}^M$ (which we refer to as the CL-selection data). We fit the predictive model $g(\mathbf{x}; \theta)$ to the training data, and then for each pixel X_i in the CL-selection data we compute the chart statistics $\{T_{\theta,i}^2, z_{\sigma,r_i,c_i}\}_{i=m+1}^M$. We then compute the empirical cdfs of $\{T_{\theta,i}^2\}_{i=m+1}^M$ and $\{z_{\sigma,r_i,c_i}\}_{i=m+1}^M$. The UCL_{θ} , LCL_{σ} , and UCL_{σ} are determined such that (1) the two component charts for $C_{\theta,i}$ and $C_{\sigma,i}$ each have the same empirical false alarm rate over the CL-selection data, and (2) the multi-chart has an empirical false alarm rate over the CL-selection data that is equal to some user-specified desired false alarm rate. The three control limit values can be efficiently found via searching for the required empirical false alarm rate (denoted by $\alpha_{\theta,\sigma}$) over the CL-selection data for the two component charts (i.e., for $\mathbf{s}_{\theta}$ and s_{σ} individually) so that the empirical false alarm rate (denoted by α_M) of the multi-chart over the CL-selection data is the desired value. This can be accomplished via binary search algorithm by noting that α_M is monotonically non-decreasing with $\alpha_{\theta,\sigma}$. In the sequel, we denote these two component charts and the multi-chart by as SWMA- θ , SWMA- σ , and SWMA-M, where the “S” stands for “score-based”.

Selecting a benchmark method to which to compare our score-based approach is difficult, because there are very few existing methods that have been developed to monitor for nonstationarity on a pixel-by-pixel basis. To the best of our knowledge, the closest existing method is the residual-based method in [17]. Although it was more intended to detect local defects in the images, it can easily be adapted to monitor for nonstationarity on a pixel-by-pixel basis. To adapt the method for this purpose, we monitor the residuals (i.e., prediction errors) $r_i = X_i - g(\mathcal{N}(X_i); \hat{\theta})$ for each pixel. The control charts for residuals (which we refer to as the RWMA) have both a LCL and an UCL , which we choose based on the empirical distribution of the residuals over the same set of CL-selection data, similarly to how we determine the control limits for the score-based charts.

The following summarizes the steps of the NM approach.

- **Step 1 (Training):** The data $\{(y_i, \mathbf{x}_i)\}_{i=1}^M$ obtained from reference micrograph/micrographs is split into the two subsets $\{(y_i, \mathbf{x}_i)\}_{i=1}^m$ and $\{(y_i, \mathbf{x}_i)\}_{i=m+1}^M$. The first subset $\{(y_i, \mathbf{x}_i)\}_{i=1}^m$ is used to train the parametric supervised learning model $g(\mathbf{x}; \theta)$, mean of the conditional likelihood $P(y | \mathbf{x}; \hat{\theta})$ of the individual pixels $(y, \mathbf{x})$, which implicitly represents the underlying joint distribution of the micrograph pixels. During training, we compute the (regularized) MLE $\hat{\theta}$ by minimizing the cost function, $-\sum_{i=1}^m l(\theta, \sigma; y_i, \mathbf{x}_i) + J(\theta)$, where $J(\theta)$ is a regularization term (in this study we use $J(\theta) = \lambda \|\theta\|_{L_2}^2$) with all hyper-parameters (including the neighborhood window size l_s) chosen by cross-validation. As discussed in [42], including regularization in the training loss function does not affect the salient point that the mean of the score functions changes if and only if the predictive relationship is nonstationary.
- **Step 2 (CL-selection):** The supervised learning model from Step 1 is applied to the second subset $\{(y_i, \mathbf{x}_i)\}_{i=m+1}^M$ to obtain the score vectors $\{\mathbf{s}_{\theta,i}, s_{\sigma,i}\}_{i=m+1}^M$ (and the residuals $\{r_i\}_{i=m+1}^M$ for the residual-based benchmark) and to select the control limits as described above to provide an empirical false alarm rate that

is equal to some user-specified desired false alarm rate α (e.g., 0.01 or 0.001).

- **Step 3 (Monitoring):** New micrograph samples are collected and, converted to a monitoring data set $\{(y_i, \mathbf{x}_i)\}_{i \geq M}$, for each pixel of which the same score vectors, residuals, and monitoring statistics $\{C_{\theta,i}, C_{\sigma,i}, C_{M,i}, r_i\}$ are computed and compared to their respective control limits. For a particular chart, if a significant portion of the charted statistics fall beyond its control limit, we conclude that the new sample is nonstationary in the sense that its microstructure differs stochastically from the reference sample.

Before fitting the predictive model $g(\mathbf{x}; \theta)$ to the training data and applying it to the CL-selection or monitoring image data to compute the score vectors or residuals, we always scale each micrograph to have zero mean and unit variance, because a change in the mean or variance of the pixel grayscale values across different micrographs could be due to the different light exposure or contrast levels, which should not be treated as nonstationarity. If one suspects that such mean or variance changes could be the net result of actual nonstationarity of the microstructure, then one could supplement our score-based multi-chart with additional component charts that monitor a local WMA estimate of the pixel mean and variance. For the monitoring micrograph(s) (i.e., the new set of micrographs to be monitored and compared to the reference micrographs), we define the *power* (i.e., probability of correctly detecting nonstationarity) for a control chart as the percentage of pixels in the monitoring samples for which the control chart signaled, i.e., for which the control chart statistic fell outside its control limit(s). The larger the power, the more clearly the control chart correctly indicated nonstationarity, when the monitoring micrographs are truly nonstationary. In Section 6.1, we apply and demonstrate this NM approach on a real and simulated materials data.

5. Nonstationarity Diagnostics (ND)

The NM approach in Section 4 is intended to (1) indicate whether the samples are nonstationary in the sense that a monitoring sample or part of a monitoring sample is different than the CL-selection reference sample(s); and (2) if nonstationary, highlight regions that are most likely to be different. If there are multiple phases in a nonstationary sample (or samples), a related important objective is to identify and demark the regions of the sample that correspond to the different phases (i.e., different distinct types of microstructure stochastic behavior). We refer to this objective as nonstationarity diagnostics (ND) to distinguish it from the NM objective. In this section, we develop an approach for this that can be used as a follow-up to NM, e.g., if the NM indicates a sample is nonstationary, to determine how many material phases there are and, more generally, what is the nature of the nonstationarity. Alternatively, our ND approach can be used as a standalone approach in which we are given a single sample or multiple samples and want to know the nature of the nonstationarity. For example, suppose a materials scientist has just created a new sample of material in the laboratory using some new processing technique or settings, and one goal is to understand the nature of the new material and whether there are multiple material phases mixed together. Or suppose that engineers in a commercial-scale process have just implemented a new processing method and want to know the same.

The intuition behind our score-based method for ND is as follows. If a micrograph is a realization of a stationary random process, then the same value of θ will represent the microstructure behavior everywhere in the micrograph. In this case, in addition to Eq. (6) holding, the sample score vectors $\{\mathbf{s}(\hat{\theta}; y_i, \mathbf{x}_i)\}_{i=1}^m$ will have

local empirical mean close to zero over every local region, providing the region is large enough that the noise in the score vectors averages out. In contrast, if the micrograph has multiple phases in it as in Figs. 1a and 1b, there will be multiple θ 's, say $\{\theta^{(l)}\}_{l=1}^k$, representing the stochastic nature of the k different phases. If we hypothetically had *a priori* pixel-wise phase labels to serve as ground truth for all k phases in the micrograph/micrographs, we would train k different supervised learning models to represent the k different conditional distributions $\{P(y|\mathbf{x}; \theta^{(l)})\}_{l=1}^k$ for the k different phases. For each phase, as in Eq. (5) or (6), we would have $E_{\theta^{(l)}}[\mathbf{s}(\theta^{(l)}; X, \mathbf{N}(X))] = \mathbf{0}$ or $\hat{E}_{\theta^{(l)}}[\mathbf{s}(\hat{\theta}^{(l)}; X, \mathbf{N}(X))] = \mathbf{0}$, where $\hat{\theta}^{(l)}$ is the MLE of $\theta^{(l)}$ over its corresponding region. In reality, we do not have such labels *a priori*, and even though we fit a single model to the entire training data with multiple phases present, there is no single parameter vector θ that can represent the multiphase microstructure. However, we can still leverage the preceding concepts that distinguish from the case of a stationary microstructure to diagnose the nonstationarity in the micrographs. More specifically, we can train a single model for the entire set of training data, in which case the empirical (sample) mean of the score vectors over the entire training data is zero, i.e., $\sum_{i=1}^m \mathbf{s}(\hat{\theta}; y_i, \mathbf{x}_i) \cong \mathbf{0}$, where $\hat{\theta}$ is the MLE over the entire training data. But locally, within regions that fall inside any of the k single phases, the score vectors will have nonzero mean. This is because each individual phase has different stochastic behavior than the mixture of all k phases, and the latter is what $\hat{\theta}$ represents.

As an alternative interpretation, score vectors are the updating vectors in a SGD algorithm to fit the model to maximize the log-likelihood function, i.e., $\hat{\theta}_{t+1} = \hat{\theta}_t + \eta \mathbf{s}(\hat{\theta}_t; y_i, \mathbf{x}_i)$, at iteration t of the algorithm. The score vectors from within any particular stationary phase from among the k phases, say the l^{th} phase, will tend to steer the current parameter values $\hat{\theta}$ towards $\theta^{(l)}$. If the k phases are sufficiently different, then $\{\theta^{(l)}\}_{l=1}^k$ will be sufficiently different in the high-dimensional space, and so the mean vectors of the k score vectors, $\mathbf{s}(\hat{\theta}; y^{(l)}, \mathbf{x}^{(l)})$, $l \in \{1, 2, \dots, k\}$ should generally be different and non-zero, where $(y^{(l)}, \mathbf{x}^{(l)})$ represents the values of a pixel and its neighborhood from within the l^{th} phase. Based on this intuition, our score-based ND approach uses *k-means* clustering on the score vectors $\{\mathbf{s}(\hat{\theta}; y_i, \mathbf{x}_i)\}_{i=1}^m$ (or some transformed version thereof) to diagnose the spatial nonstationarity of microstructures.

To reduce the effects of noise in the score vectors and improve the clustering performance by taking into account spatial proximity information, we conduct clustering on $\{\mathbf{z}_{r_i, c_i}\}_{i=1}^m$ in Eq. (8) instead of original score vectors, as follows. First, as a visualization tool to help approximately estimate the number k of phases in the sample, we developed a 3D plot in which the two horizontal axes represent the 2D spatial coordinates of the micrograph, and the vertical axis represents the magnitude of the vectors $\{\mathbf{z}_{r_i, c_i}\}_{i=1}^m$. Moreover, the red-green-blue color of this plotted surface represents the first 3 principal components analysis (PCA) scores. See Fig. 8a, later, as an example of this plot. In this plot, the pixels falling into different phases can be distinguished based on having different vertical axes height and/or different color, because the smoothed score vectors from different phases have different magnitudes and directions. From this plot, the estimated number $\hat{k}$ of phases is taken, roughly, to be the number of patches with different heights and different colors. This $\hat{k}$ is then used in setting the number of centroids in the *k-means* clustering algorithm. A 3D scatter plot of the top-3 PCA scores of $\{\mathbf{z}_{r_i, c_i}\}_{i=1}^m$ along with their clustering labels can be constructed, and those labels can be also overlaid on the initial micrograph to show regions of different phases. In Section 6.2, we demonstrate this ND approach on real and simulated materials data.




Fig. 3. Example of (a) reference and (b) subsequent nonstationary micrographs of silica particles dispersed in PMMA.

6. Experimental Results

In this section, we present the results of our score-based framework using real and simulated micrograph data for NM (Section 6.1) and ND (Section 6.2).

6.1. Results for the NM Approach

We consider two data sets: a PMMA data set and a simulated 2D AR data set shown in Fig. 3 and 1c-1d, respectively. The former is a real data set that consists of TEM images of silica particles in PMMA with octyl functional modification. How densely those particles are dispersed in the matrix can be controlled by various processing conditions, and the dispersion density can affect the physical properties (e.g., the breakdown stress or dielectric constant of the material). We select two sets of samples with different dispersion densities and treat one as the reference data and the other as the monitoring data to be monitored for nonstationarity. The simulation data are generated via a 2D spatial autoregressive (AR) model with various choice of AR coefficients. The 2D AR model allows us to generate many Monte Carlo replicates of data sets to investigate and compare the false alarm rate and the power of the various control charts. We show that our score-based charts are far more effective than the residual-based chart, and its performance is further enhanced by the multi-chart.

6.1.1. PMMA Data Analysis

We trained two different supervised learning models $g(\mathbf{x}; \theta)$, a linear regression model and a fully-connected neural network model with one hidden-layer having 10 nodes, to serve as the mean function of the conditional distribution. The L_2 regularization parameter, $\lambda = 0.01$, was chosen via cross-validation. The results for the PMMA micrographs in Figs. 3 are shown in Fig. 4 for the neural network model. Although the linear model results in much less accurate prediction of each pixel value, it provided similar monitoring results as the neural network and so is omitted here.

Figs. 4a and 4b show 3D control charts and 2D heat-maps, respectively, of the results. In the 3D control charts shown in Fig. 4a, we plot the UCL for the SWMA- θ chart and the LCL and UCL for the RWMA chart as horizontal planes, calculated from CL-selection micrographs. For comparison, the first monitoring micrograph is a reference micrograph, and the other two monitoring micrographs are from nonstationary micrographs which are statistically different from reference ones (so we would like the chart to signal for these micrographs).

From the control charts in Fig. 4a, we can see that the SWMA- θ chart out-performs the RWMA chart in terms of the power (i.e., the out-of-control percentage of signals for the two nonstation-

ary images). For the heat-maps in Fig. 4b, we used the following convention: (1) for all heat-maps of each monitoring statistic (each row), we apply the same color scheme, meaning the color bars correspond to the same range of numbers, the minimum and maximum values of which correspond to the two extremes of the monitoring statistic values of all the heat-maps in that row; and (2) the color bars of different rows share the same colors, but correspond to different ranges of numbers. In this way, within each row, we can map the relative differences in metric values to the differences of colors in heat-maps. And then we can compare the relative differences across rows. We observe the SWMA- θ chart has the best performance in terms of best distinguishing the two monitoring images from the reference image. Besides the global nonstationarity, the score-based method seems to be also effective in detecting local nonstationarity, based on the peaks in the heat-maps coinciding with what appear to be somewhat unusual local agglomerations. Here, we only consider the SWMA- θ chart but not the SWMA- σ chart, because the latter generally requires the CL-selection data to be of larger size, as in the example in Section 6.1.2. This is perhaps because the scores for the SWMA- σ chart are the squares of the residuals, which typically have higher variability than the scores for the SWMA- θ chart.

6.1.2. 2D AR Data Analysis

In order to access the power of the score-based methods on different kinds of nonstationarity, we continue to use the 2D AR model (1) to generate stochastic microstructures. For the microstructures in Figs. 5a and 6a, the parameters and configurations are: $c_0 = 0.01$; $\sigma_{AR} = 1$; the row-by-row concatenated AR coefficient vectors $[\phi_{0,0}^{(CL)}, \phi_{0,1}^{(CL)}, \phi_{0,2}^{(CL)}, \phi_{1,0}^{(CL)}, \phi_{1,1}^{(CL)}, \phi_{1,2}^{(CL)}, \phi_{2,0}^{(CL)}, \phi_{2,1}^{(CL)}, \phi_{2,2}^{(CL)}] = [0, 3.59e-01, 1.07e-02, 3.90e-01, 4.21e-02, 1.76e-03, 9.98e-02, -1.82e-03, 1.72e-05]$ and $[\phi_{0,0}^{(M)}, \phi_{0,1}^{(M)}, \phi_{0,2}^{(M)}, \phi_{1,0}^{(M)}, \phi_{1,1}^{(M)}, \phi_{1,2}^{(M)}, \phi_{2,0}^{(M)}, \phi_{2,1}^{(M)}, \phi_{2,2}^{(M)}] = [0, 2.74e-1, 2.93e-2, 4.31e-1, 1.50e-1, -1.17e-2, 4.52e-2, -2.96e-2, -2.41e-1]$; and $h(u) = \min(5, \max(0.05, \exp(u)))$. For the microstructures in Fig. 5b and 6b, the parameters and configurations are as specified in Section 3.

The control charts in Fig. 5 show the results of the NM approach. We again fit a neural network model for $g(\mathbf{x}; \theta)$, using a single hidden-layer with 10 nodes. Because here we can generate an arbitrary number of samples, we increase the number of CL-selection micrographs to 4 with each micrograph of size 256×256 , which will allow us to more accurately control the false alarm rate when the CL-selection and monitoring micrographs are statistically equivalent (we only show a single CL-selection micrograph and a single monitoring micrograph in Fig. 5 for brevity). We see that in terms of the power, the SWMA- θ chart performs better than the SWMA- σ chart in Fig. 5a, and the SWMA- σ chart performs better than the SWMA- θ chart in Fig. 5b. In both figures, the SWMA-M chart achieves the best performance of the SWMA- θ chart and the SWMA- σ chart and is much better than the baseline RWMA chart.

To more comprehensively investigate both the power and the false alarm rate of our score-based approach, we gradually increase the difference between the CL-selection and monitoring micrographs. More specifically, we denote the two sets of 2D AR coefficients used to generate any pair of CL-selection and monitoring micrographs in Fig. 5 as $\phi^{(p)} = \{\phi_{r,c}^{(p)}\}_{r,c=0}^2$ with $\phi_{0,0}^{(p)} = 0$ and $p \in \{0, 1\}$, where $p = 0$ and $p = 1$ denote reference and nonstationary micrographs, respectively. Then, we define a parameter γ , which governs how different the CL-selection micrographs are from the monitoring micrographs via:

$$\begin{aligned} \phi^{(CL)} &= \phi^{(0)} \\ \phi^{(M)} &= (1 - \gamma)\phi^{(0)} + \gamma\phi^{(1)} \end{aligned} \quad (19)$$




Fig. 4. Control charts of SWMA- θ and RWMA on the PMMA data set from the neural network model. The window length scale for training is $I_s = 5$, and the WMA window length scale $I_w = 30$.

When $\gamma = 0$, micrographs in the CL-selection and monitoring data are generated by the same 2D AR model so that they are statistically equivalent, in which case the power is the false alarm rate α . As γ increases from 0 to 1, the images become increasingly statistically different. The results in Fig. 5 correspond to $\gamma = 1.0$.

In practice, if the sizes (in terms of total number of pixels) of the CL-selection data and the monitoring data are both sufficiently large, then the false alarm rate during monitoring will be approximately the desired value of α specified when selecting the control limits using the CL-selection data. This would give a common basis for comparison of the power of the different control charts. For all




Fig. 5. Spatial WMA control chart results for the 2D AR data. Top row shows two pairs of microstructures. The left micrograph in each pair (panel (a) or panel (b)) is representative of the reference images, and the right micrograph is statistically different and representative of the monitoring images. The training and WMA window length scales are $l_s = 5$ and $l_w = 30$.

of our simulation experiments, we achieved a common false alarm rate differently, by bypassing the CL-selection data and selecting the control limits to directly control (at least approximately) the false alarm rate over a large set of monitoring data. From Fig. 6, which plots the power of the different control charts as a function of γ , we can see that between the two component score-based charts, sometimes the SWMA- θ chart has higher power than the SWMA- σ chart, and sometimes vice-versa. And the SWMA-M chart always performs similarly to the best of the two component score-based control charts. In contrast, for this example the RWMA chart

is completely ineffective at detecting the change and has a power that is not much higher than the false alarm rate even for the larger γ values. Note that when $\gamma \cong 0.2$ in Fig. 6, the difference of microstructures are difficult to discern with the human eye (we omit the micrographs for brevity), but our score-based method can still detect the differences with reasonable power, which further demonstrates its effectiveness.

We also trained and used linear models for $g(\mathbf{x}; \boldsymbol{\theta})$ and found that the score-based method performed almost as well (in terms of power at detecting nonstationarity) as the neural network mod-




Fig. 6. Comparison of the statistical power of the four control charts in detecting nonstationarity of micrographs corresponding to the Fig. 5 example, as γ varies from 0 to 1. The power for $\gamma = 0$ is the false alarm rate. Panels (a) and (b) are for the AR models corresponding to Figs. 5a and 5b, respectively. For each value of γ , the box plots are the power values across ten Monte Carlo replicates.




Fig. 7. Results of our score-based method applied to an SEM image of a dual-phase steel sample. We used a neural network model with a single hidden-layer having 10 neurons and parameters $l_s = 5$ and $l_w = 30$. The color-coded predicted labels (top-right panel) are overlapped (bottom-left panel) with the mask showing the “true” phases provided by [38]. The bottom-right panel is the original SEM micrograph overlaid with the labels from our ND approach. The top-left plot are the first three PCA score for the set of score vectors, color-coded by ND phase labeling, which is useful for understanding the distribution of PCA components of score vectors in a high-dimensional space and debugging the data preprocessing and calculations of score vectors. The electronic version has color images of higher resolution.

els, and so we omit the results for brevity. This is an interesting observation, because even though a linear model is not a correct model structure for $g(\mathbf{x}; \theta)$ for this example (because of the non-linear transformation $h(\cdot)$), our score-based framework still works reasonably well. Based on this observation, one potential strategy is to initially use a simpler model to take advantage of the lower computational expense of fitting the model and computing the score vectors, and then switch to a more complex model to get a better performance and interpretation if the simpler model signals a change.


6.2. Results for the ND Approach

We apply the ND approach to examples involving two different sets of micrograph data. The first data set consists of SEM images of dual-phase steel [38], in which there are some martensite islands in ferrite matrix. Diagnosing/segmenting such multiphase (thus nonstationary) images is important for quality control and for understanding properties of steel samples. The second data set consists of TEM images of silica particles dispersed in PMMA with octyl functional modification. We choose images with different dispersion density and paste them together to form some artificial nonstationary microstructures and represent the practical problem of segmenting micrographs with multiple phases in each sample. Automatically segmenting regions with different microstructure characteristics like particle density is of interest to materials researchers, because the microstructure affects physical properties of the materials. We note that the ND approach can be applied to analyze and segment the different microstructures in a single multiphase micrograph sample or in a collection of nonstationary micrograph samples.


6.2.1. Dual-Phase Steel Data Set

Fig. 1b shows an SEM image of dual-phase steel consisting of a ferrite matrix with martensite in the form of islands. This micrograph is especially challenging to analyze because some parts of the martensite regions are similar to parts of the ferrite regions and because the martensite areas are not connected. A human attempt to draw boundaries around the phase regions is very time-consuming, tedious, and error-prone. In contrast, the results of our score-based method shown in Fig. 7 are very effective at identifying the multiple phases that are present and distinguishing them, in addition to being fully automated. In particular, the boundaries obtained from our score-based method are smooth and closely aligned with the true boundaries between the phases. Note that the true boundaries in Fig. 7 were determined in [38] via various thresholding, smoothing, and post-processing methods that required a substantial amount of manual processing and human expertise. In addition, the non-connected regions of martensite are successfully detected. The phase labels from our ND method are consistent with the mask, which can be treated as the ground truth, except that the boundaries of our labeled phase regions are smoothed to some extent. Similarly, some of the small islands of martensite have been smoothed out. The smoothing is the inevitable consequence of our WMA window and the neighborhood window having size larger than a single pixel. We emphasize that our score-based method is highly automated and uses minimal human intervention, while the method of [38] requires many steps and heavy human-involved pre-processing, tuning, and post-processing.

Note that the width of the unlabeled region around the borders of the bottom images in Fig. 7 is equal to size of the neighborhood window plus WMA window ($l_s + l_w$). In order to produce phase labels for each pixel in the micrograph, we need to substantially modify/extend the current approach, which we are currently investigating.



(a) Visualization of the score vectors of a linear model from side (left panel) and top views (right panel).



(b) The results of the our score-based ND approach with a linear model.

(c) The results of the our score-based ND approach with a neural network model with one hidden-layer having 10 nodes.

Fig. 8. Results of our score-based method applied to a TEM image with nonstationarity generated by pasting together silica-PMMA materials samples with different particle dispersion densities. Here, we use parameters $l_s = 5$ and $l_w = 20$. The electronic version has color images of higher resolution.

6.2.2. PMMA Data Set

The second data set consists of TEM images of silica particles dispersed in PMMA, some of which were shown in Fig. 1a and analyzed in Section 6.1. We created nonstationary micrographs by pasting a micrograph with low particle dispersion density to the upper left quadrant and three micrographs with higher particle dispersion density to the other three quadrants. As shown in Fig. 8a, the result of our visualization method with 3D height and color mapping described in Section 5 implies there are mainly two kinds of microstructures in the sample. With this information, we applied our ND score clustering method with two clusters. Fig. 8b and 8c show that our score-based ND approach with either the linear model or the neural network as the supervised learner $g(\mathbf{x}; \theta)$ can effectively diagnose the nonstationarity and accurately label the phases within the micrograph. The 3D clustering figures also help visualize the distribution of PCA components of score vectors in a high-dimensional space to help understand how score vectors are different for different material phases and also help in debugging the data preprocessing and calculations of score vectors.

In our investigation, we found that compared with the score-based ND approach, the score-based NM approach in Section 4 tends to indicate differences between different phases

with higher power. This makes sense, because when the model is trained on a stationary micrograph(s) of a single phase, the variance of the score vectors will typically be smaller than when the model is trained on nonstationary micrograph(s) with multiple phases, and hence it becomes easier to detect differences between a new phase and the reference phase. On the other hand, the NM approach requires more data in the sense that one must have one or more reference micrographs available. Moreover, the ND approach provides more diagnostic information than the NM approach and explicitly identifies the micrograph regions corresponding to the multiple phases.

7. Conclusions

In this study, we have developed a powerful and versatile score-based framework for nonstationarity analysis of stochastic microstructures of materials. This problem is of increasing importance due to the increasing availability of complex multiphase micrograph data and the lack of effectiveness of traditional methods. Modeling the stochastic nature through parametric supervised learning models and analyzing nonstationarity through our score-based framework have a solid theoretical foundation and, as we

demonstrated with a number of examples, are effective and efficient at monitoring and diagnosing nonstationarity. We have developed two components in our framework: nonstationarity monitoring and nonstationarity diagnostics, which are intended for use in two common related but different practical scenarios. The consistently good performance across our examples and the higher statistical power that we demonstrated through Monte Carlo simulations are evidences of the effectiveness of our score-based approach and its advantages over the residual-based benchmark existing approach. The framework has substantial potential for automating and improving image analysis of materials microstructures and can be combined with other state-of-the-art machine learning and deep learning techniques, e.g., classification and segmentation deep learning networks, which we are currently investigating as extensions.

Another potential extension of the approach is to aid in determining the appropriate representative volume element (RVE) size for subsequent numerical simulation (e.g., finite element) of material properties. We anticipate that the smallest WMA window size for which our averaged score vector is stable (i.e., only varies “negligibly” across the microstructure) constitutes an upper-bound on the RVE size. We are currently investigating this as future work.

Our approach is a parametric one in that the predictive supervised learning model $g(\mathbf{x}; \theta)$ has a set of estimated parameters θ . Many parametric supervised learning models (like the neural networks we have used in our examples) can have so many parameters that one may view them as virtually nonparametric, in terms of the flexibility they provide in capturing even the most complex nonlinear behavior. We are currently investigating extending our approach to the case of extremely high-dimensional θ and also to truly nonparametric models. Related to this, in other areas of microstructure analysis such as microstructure synthesis, nonparametric methods such as Markov random fields have shown much success (e.g., [50]). We view our nonstationarity monitoring approach not as a competitor to, but as something that could be used in conjunction with, nonparametric microstructure synthesis methods. One would first use our approach to identify whether the microstructure is stationary and, if not, to identify micrograph sub-regions within which the microstructure is stationary. One could then use existing microstructure reconstruction algorithms intended for stationary microstructures within each sub-region. Integration of the approaches is a topic that warrants further research.

Declaration of Competing Interest

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

Acknowledgement

This work was funded in part by the Air Force Office of Scientific Research Grant # FA9550-18-1-0381, which we gratefully acknowledge. The micrographs of silica particles in PMMA are courtesy of Prof. Linda Schadler (Linda.Schadler@uvm.edu) and Prof. Catherine Brinson (cate.brinson@duke.edu). This work used the Extreme Science and Engineering Discovery Environment (XSEDE) [51], which is supported by National Science Foundation grant number ACI-1548562, and the Quest high performance computing facility at Northwestern University, which is jointly supported by the Office of the Provost, the Office for Research, and Northwestern University Information Technology. We also thank two anonymous reviewers for a number of suggestions that have improved the paper.

References

- [1] A. Bandyopadhyay, B. Heer, Additive manufacturing of multi-material structures, *Materials Science and Engineering: R: Reports* 129 (2018) 1–16.
- [2] F. Stenley, *Scanning and transmission electron microscopy*, 1993.
- [3] J.A. Strosio, W.J. Kaiser, *Scanning tunneling microscopy*, 27, Academic press, 1993.
- [4] J.H. Kinney, M.C. Nichols, X-ray tomographic microscopy (xtm) using synchrotron radiation, *Annual review of materials science* 22 (1) (1992) 121–152.
- [5] L. Howle, R. Behringer, J. Georgiadis, Visualization of convective fluid flow in a porous medium, *Nature* 362 (6417) (1993) 230–232.
- [6] J. Fredrich, B. Menendez, T.-F. Wong, Imaging the pore structure of geomaterials, *Science* 268 (5208) (1995) 276–279.
- [7] D. Gagliardi, Material data matter—standard data format for engineering materials, *Technological Forecasting and Social Change* 101 (2015) 357–365.
- [8] K. Takahashi, Y. Tanaka, *Materials informatics: a journey towards material design and synthesis*, *Dalton Transactions* 45 (26) (2016) 10497–10499.
- [9] B. Blaiszik, K. Chard, J. Pruyne, R. Ananthakrishnan, S. Tuecke, I. Foster, The materials data facility: Data services to advance materials science research, *JOM* 68 (8) (2016) 2045–2052.
- [10] B. Puchala, G. Tarcea, E.A. Marquis, M. Hedstrom, H. Jagadish, J.E. Allison, The materials commons: a collaboration platform and information repository for the global materials community, *JOM* 68 (8) (2016) 2035–2044.
- [11] A. Jain, G. Hautier, S.P. Ong, K. Persson, New opportunities for materials informatics: resources and data mining techniques for uncovering hidden relationships, *Journal of Materials Research* 31 (8) (2016) 977–994.
- [12] A. Jain, K.A. Persson, G. Ceder, Research update: The materials genome initiative: Data sharing and the impact of collaborative ab initio databases, *APL Materials* 4 (5) (2016) 053102.
- [13] C. Kim, G. Piliaria, R. Ramprasad, From organized high-throughput data to phenomenological theory using machine learning: the example of dielectric breakdown, *Chemistry of Materials* 28 (5) (2016) 1304–1311.
- [14] F. Rose, C. Toher, E. Gossett, C. Oses, M.B. Nardelli, M. Fornari, S. Curtarolo, Aflux: The lux materials search api for the aflow data repositories, *Computational Materials Science* 137 (2017) 362–370.
- [15] AFRL, Aflr (2018) air force research laboratory (afrl) additive manufacturing (am) modeling challenge series., 2018.
- [16] A.T. Bui, D.W. Apley, Monitoring for changes in the nature of stochastic textured surfaces, *Journal of Quality Technology* 50 (4) (2018) 363–378.
- [17] A.T. Bui, D.W. Apley, A monitoring and diagnostic approach for stochastic textured surfaces, *Technometrics* 60 (1) (2018) 1–13.
- [18] J.J. Liu, J.F. MacGregor, Estimation and monitoring of product aesthetics: application to manufacturing of “engineered stone” countertops, *Machine Vision and Applications* 16 (6) (2006) 374.
- [19] H.-D. Lin, Computer-aided visual inspection of surface defects in ceramic capacitor chips, *Journal of Materials Processing Technology* 189 (1–3) (2007) 19–25.
- [20] M. Bharati, J. MacGregor, W. Tropper, Softwood lumber grading through on-line multivariate image analysis techniques, *Industrial & Engineering Chemistry Research* 42 (21) (2003) 5345–5353.
- [21] D.C. Montgomery, *Introduction to statistical quality control*, John Wiley & Sons, 2007.
- [22] K. Paynabar, C. Zou, P. Qiu, A change-point approach for phase-i analysis in multivariate profile monitoring and diagnosis, *Technometrics* 58 (2) (2016) 191–204.
- [23] M. Grasso, A. Menafoglio, B.M. Colosimo, P. Secchi, Using curve-registration information for profile monitoring, *Journal of Quality Technology* 48 (2) (2016) 99–127.
- [24] R. Viveros-Aguilera, S.H. Steiner, R.J. MacKay, Monitoring product size and edging from bivariate profile data, *Journal of Quality Technology* 46 (3) (2014) 199–215.
- [25] S. Torquato, Statistical description of microstructures, *Annual review of materials research* 32 (1) (2002) 77–111.
- [26] J.J. Lewandowski, M. Seifi, Metal additive manufacturing: a review of mechanical properties, *Annual Review of Materials Research* 46 (2016) 151–186.
- [27] S.C. Ligon, R. Liska, J. Stampfl, M. Gurr, R. Mülhaupt, Polymers for 3d printing and customized additive manufacturing, *Chemical reviews* 117 (15) (2017) 10212–10290.
- [28] Z. Chen, Z. Li, J. Li, C. Liu, C. Lao, Y. Fu, C. Liu, Y. Li, P. Wang, Y. He, 3d printing of ceramics: A review, *Journal of the European Ceramic Society* 39 (4) (2019) 661–687.
- [29] T. Moritz, S. Maleksaeedi, Additive manufacturing of ceramic components, in: *Additive Manufacturing*, Elsevier, 2018, pp. 105–161.
- [30] A.A. Efros, T.K. Leung, Texture synthesis by non-parametric sampling, in: *Proceedings of the seventh IEEE international conference on computer vision*, 2, IEEE, 1999, pp. 1033–1038.
- [31] B. Lu, S. Torquato, Lineal-path function for random heterogeneous materials, *Physical Review A* 45 (2) (1992) 922.
- [32] P.-E. Øren, S. Bakke, Reconstruction of berea sandstone and pore-scale modelling of wettability effects, *Journal of petroleum science and engineering* 39 (3–4) (2003) 177–199.
- [33] S. Torquato, Necessary conditions on realizable two-point correlation functions of random media, *Industrial & engineering chemistry research* 45 (21) (2006) 6923–6928.

- [34] Y. Jiao, F. Stillinger, S. Torquato, Modeling heterogeneous materials via two-point correlation functions. ii. algorithmic details and applications, *Physical Review E* 77 (3) (2008) 031135.
- [35] Y. Liu, M.S. Greene, W. Chen, D.A. Dikin, W.K. Liu, Computational microstructure characterization and reconstruction for stochastic multiscale material design, *Computer-Aided Design* 45 (1) (2013) 65–76.
- [36] R. Cang, H. Li, H. Yao, Y. Jiao, Y. Ren, Improving direct physical properties prediction of heterogeneous materials from imaging data via convolutional neural network and a morphology-aware generative model, *Computational Materials Science* 150 (2018) 212–221.
- [37] B.L. DeCost, E.A. Holm, A computer vision approach for automated analysis and classification of microstructural image data, *Computational materials science* 110 (2015) 126–133.
- [38] S. Banerjee, S.K. Ghosh, S. Datta, S.K. Saha, Segmentation of dual phase steel micrograph: An automated approach, *Measurement* 46 (8) (2013) 2435–2440.
- [39] R. Bostanabad, A.T. Bui, W. Xie, D.W. Apley, W. Chen, Stochastic microstructure characterization and reconstruction via supervised learning, *Acta Materialia* 103 (2016) 89–102.
- [40] E. Levina, P.J. Bickel, et al., Texture synthesis and nonparametric resampling of random fields, *The Annals of Statistics* 34 (4) (2006) 1751–1773.
- [41] X. Li, Y. Zhang, H. Zhao, C. Burkhart, L.C. Brinson, W. Chen, A transfer learning approach for microstructure reconstruction and structure-property predictions, *Scientific reports* 8 (1) (2018) 1–13.
- [42] K. Zhang, A.T. Bui, D.W. Apley, Concept drift monitoring and diagnostics of supervised learning models via score vectors, 2020. Submitted
- [43] P.J. Bickel, K.A. Doksum, *Mathematical statistics: basic ideas and selected topics*, volume I, 117, CRC Press, 2015.
- [44] X. Liu, V. Shapiro, Random heterogeneous materials via texture synthesis, *Computational Materials Science* 99 (2015) 177–189.
- [45] D.M. Turner, S.R. Kalidindi, Statistical construction of 3-d microstructures from 2-d exemplars collected on oblique sections, *Acta Materialia* 102 (2016) 136–148.
- [46] I. Javaheri, V. Sundararaghavan, Polycrystalline microstructure reconstruction using markov random fields and histogram matching, *Computer-Aided Design* 120 (2020) 102806.
- [47] P. Acar, V. Sundararaghavan, A markov random field approach for modeling spatio-temporal evolution of microstructures, *Modelling and Simulation in Materials Science and Engineering* 24 (7) (2016) 075005.
- [48] C.A. Lowry, W.H. Woodall, C.W. Champ, S.E. Rigdon, A multivariate exponentially weighted moving average control chart, *Technometrics* 34 (1) (1992) 46–53.
- [49] H. Hotelling, *Multivariate quality control. techniques of statistical analysis*, McGraw-Hill, New York (1947).
- [50] A. Kumar, L. Nguyen, M. DeGraef, V. Sundararaghavan, A markov random field approach for microstructure synthesis, *Modelling and Simulation in Materials Science and Engineering* 24 (3) (2016) 035015.
- [51] J. Towns, T. Cockerill, M. Dahan, I. Foster, K. Gaither, A. Grimshaw, V. Hazelwood, S. Lathrop, D. Lifka, G.D. Peterson, et al., Xsede: accelerating scientific discovery, *Computing in science & engineering* 16 (5) (2014) 62–74.