

Random Composites Characterization Using a Classifier Model

H. Liu, A.M.ASCE¹; S. R. Arwade, A.M.ASCE²; and T. Igusa, A.M.ASCE³

Abstract: A new method is introduced for characterizing and analyzing materials with random heterogeneous microstructure. The method begins with classifiers which process information from high-fidelity analyses of small-sized simulated microstructures. These classifiers are subsequently used in a multipass moving window to identify subregions of potentially critical microscale behavior such as strain concentrations. In the derivation of the method, it is shown how information theory-based concepts can be formulated in a Bayesian decision theory framework that addresses microstructural issues. Furthermore, it is shown how a sequence of classifiers can be constructed to refine the analysis of microstructure. While the method presented herein is general, a relatively simple example of a two-dimensional, two-phase composite is used to illustrate the analysis steps.

DOI: 10.1061/(ASCE)0733-9399(2007)133:2(129)

CE Database subject headings: Microstructures; Decision making; Damage; Fracture; Composite materials; Uncertainty principles; Statistics; Bayesian analysis.

Introduction

With recent advances in computational power, there has been growing interest in analyzing materials down to the microstructural level. This interest is primarily driven by the fact that damage initiation and growth as well as critical material behavior are highly dependent on microstructural properties. While methods such as those based on multiscale techniques have proven to be promising for many materials, the analysis of materials with random microstructure has unique, interesting challenges. The main challenge is that the behavior of random microstructures cannot be extrapolated from the behavior of a single subregion of the microstructure, as is the case of periodic composites, which can be represented using a unit cell.

Most approaches for characterizing and analyzing materials with random microstructure use representations at the mesostructural level based on microscale statistics. The essential idea is to condense information at the microstructural level, sampled at multiple or theoretically infinite locations, into a relatively small set of parameters or functional representations. The most basic statistical methods use spatial averages over local regions that are large in comparison with the microstructure. More complex

statistical methods use second- or higher-moment information, wherein spatial averaging is performed using a product of properties at multiple points. The classical examples are effective medium theories, where micromechanics is used to represent random composite materials by equivalent deterministic continua (Benveniste 1987; Chen et al. 1992; Buryachenco et al. 2003). More recent examples include approaches that use the orientation distribution function (ODF) to characterize polycrystalline materials (Beaudoin et al. 1995; Adams et al. 2001). The ODF is a statistical model of the histogram of the crystal orientation vector obtained by sampling micrographical observations of crystalline structure. While the ODF is already a reduced form of microstructural information, it is typically reduced further using basis function expansions. Harmonic polynomials (Adams et al. 1994) and principal components (Ganapathysubramanian and Zabarar 2004; Acharjee and Zabarar 2003) are two types of such basis function expansions that have been tailored for specific material design applications.

While mesostructural representations based on microscale statistics have been proven to be useful, there is still a need to analyze material behavior using more detailed, spatially variable information at the microstructural level. The reason is that many material characteristics, especially those related to strongly nonlinear phenomena such as fracture, are closely related to short-range interactions between microstructural entities. For instance, the yield strength and fracture behavior of materials with inclusions have been found to be governed by the nonuniformity of the spatial distribution of inclusions (Lewandowski et al. 1989; Corbin and Wilkinson 1994). The difficulty with any detailed, high-fidelity micromechanical simulations is that, even with large-scale computational effort, only relatively small-sized material specimens can be analyzed (Arwade and Grigoriu 2004).

The goal of this paper is to introduce a method that can use information from micromechanical simulations of very small-sized material samples to characterize and analyze random composites at the mesoscale. The theoretical framework for this method is based on Bayesian classifiers. An important difference from existing statistics or simulation-based methods (Torquato

¹SDR Engineering, 3370 Capital Circle NE, Ste. G, Tallahassee, FL 32308. E-mail: hliu@sdengineering.com

²Assistant Professor, Dept. of Civil and Environmental Engineering, Univ. of Massachusetts, 223 Marston Hall, 130 Natural Resources Rd., Amherst, MA 01003 (corresponding author). E-mail: arwade@ecs.umass.edu

³Professor, Dept. of Civil Engineering, Johns Hopkins Univ., 3400 N. Charles St., Baltimore, MD 21218. E-mail: tigusa@jhu.edu

Note. Associate Editor: Arvid Naess. Discussion open until July 1, 2007. Separate discussions must be submitted for individual papers. To extend the closing date by one month, a written request must be filed with the ASCE Managing Editor. The manuscript for this paper was submitted for review and possible publication on February 8, 2005; approved on July 7, 2006. This paper is part of the *Journal of Engineering Mechanics*, Vol. 133, No. 2, February 1, 2007. ©ASCE, ISSN 0733-9399/2007/2-129-140/\$25.00.

et al. 2003; Saganopan and Pitchumani 1998) is that the characterization and analysis is tied to a search for microstructural patterns, defined broadly and abstractly within the classifier context.

Classifiers, which are essentially algorithms that search for a class of objects within a large data set, have been used with broad success in computer science applications (McLachlan 1992; Breiman et al. 1984). It is only recently that they have been used in analysis and design problems in structural mechanics (Acharjee and Zabarar 2003; Ganapathysubramanian and Zabarar 2004). The work that is most relevant to the present context is by Zabarar and his colleagues. Beginning with a statistically reduced description of the microstructure in the form of the ODF, they used classifiers to define the relationships between the ODF and mesoscopic behavior. In the present paper, classifiers are used on direct representations of the microstructure. As noted later in this paper, the two methods use statistics and classifiers in a fundamentally different way. This has positive implications in that the methods are, in fact, complementary.

The paper begins with a localized representation of microstructure. The simplest representation is a set of geometric and material-related descriptors that provides information that can be used, for example, to create a high-fidelity finite-element mesh. Such a representation, however, is too detailed and cumbersome for classifier analysis. Hence, for the approach in this paper, a low-dimensional representation is developed in terms of basis vector expansions of randomly generated, small-sized materials samples. Statistically efficient basis vectors for the microstructure are used that are specifically associated with critical response behavior such as local strain concentrations. Next, a sequence of Bayesian classifiers are derived in a decision-theoretic framework. Finally it is shown how the classifiers can be used with multipass moving windows to search for the most likely locations of damage initiation in a material specimen under external stresses.

While the concepts and associated equations presented herein are applicable to general three-dimensional (3D), models of microstructure, the figures and example used to illustrate the concepts are restricted to a relatively simple two-dimensional (2D) problem: the prediction of damage initiation in a 2D cross section of a fiber-reinforced composite material. In the usual approaches to analyzing this problem, the microstructure is characterized by the location of each inclusion and is analyzed using the distribution of these locations using spatial statistics. An important point-based statistic is the shortest distance between neighboring points; tessellation-based methods (Everett and Chu 1993; Pyrz 1994; Li et al. 1999a,b) provide information on the densities of inclusion clusters. Correlation-based methods (Berryman 1985; Baxter and Graham 2000) and multiscale techniques (Spowart et al. 2001) have also been used to investigate such composites. The Bayesian classifier method developed herein provides a significantly different approach to the composites analysis problem. The results of the analysis of this simple example problem show both that mechanically meaningful patterns exist in the microstructure of fiber-reinforced composites, and that classification algorithms can make use of these patterns to identify damage initiation sites much more quickly than would be possible using high-fidelity finite-element analyses.

The example is used in the next section to clarify the development of the theory behind the classifier-based method. Details of the example results are presented in the succeeding section.

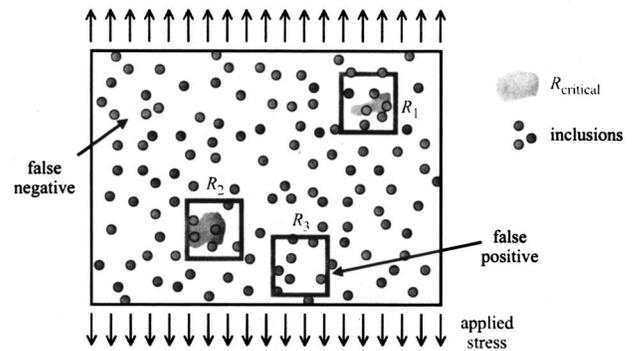


Fig. 1. Two-dimensional, two phase random composite with critical subregions $R_{critical}$, potentially critical subregions $\{R_j\}$, and examples of false positive and false negative errors

Development of Microstructural Classifiers

Basic Concepts

Consider a body of material with random heterogeneous microstructure under external stresses, such as the two-dimensional, two-phase composite shown in Fig. 1. Let R_{total} denote the total region of material under investigation and let $R_{critical} \subset R_{total}$ be the subregions that may experience critical micromechanical behavior. Critical behavior may include high stress or strain concentrations or other significant localized response such as damage initiation. This subregion, which may be disjoint, is indicated by the shaded areas in the figure. The goal is to derive an approach that can rapidly predict the location of $R_{critical}$ at a computational effort that is far less than that required for a full analysis of R_{total} . The tradeoff here is that such an approach would produce errors in identifying the critical region. Specifically, the method of interest would identify potentially critical subregions $\{R_j\} = \{R_j : j = 1, \dots, m\}$ whose union $\hat{R}_{critical} = R_1 \cup \dots \cup R_m$ contains most or all of the critical subregion $R_{critical}$. In Fig. 1, the three potentially critical subregions $\{R_1, R_2, R_3\}$ are indicated by the squares.

There are two primary types of error: *false positive*, where subregion R_j is completely noncritical and *false negative*, where some of the critical subregion lies outside the union $\hat{R}_{critical}$. These errors, which are illustrated in Fig. 1, are defined more precisely by the following:

$$\text{false positive in subregion } j \Leftrightarrow R_j \cap R_{critical} = \emptyset$$

$$\text{false negative} \Leftrightarrow \hat{R}_{critical} \cap R_{critical} \neq R_{critical} \quad (1)$$

One important application of such a method is in improving the efficiency of computational analysis of microstructured materials. Efficiency can be gained since high-fidelity analysis (a highly refined finite-element mesh, for example), need only be performed within $\hat{R}_{critical}$, whereas in the complementary region $\hat{R}'_{critical}$, a lower fidelity analysis (coarse finite-element mesh with homogenized properties) may suffice. In this application, when the subregions with false positive error form a relatively large volume, then the computational effort in the reduced analysis becomes unnecessarily high. If the false negative error is large, in which case a large portion of the critical subregion lies outside of $\hat{R}_{critical}$, then the accuracy of the reduced analysis is compromised. It can be seen that while both types of errors are important, the

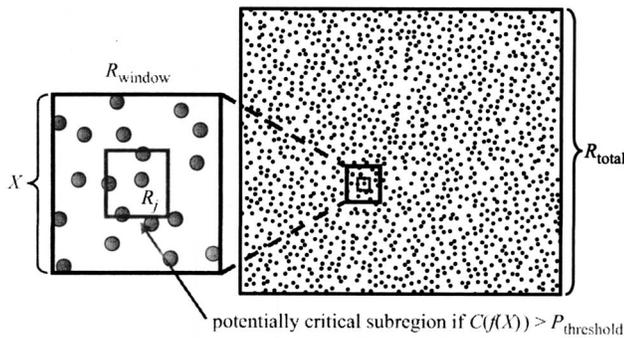


Fig. 2. Moving window R_{window} enclosing microstructure X with central subregion $R_{j=1}(R_{\text{window}})$ and associated inequality test for identifying potentially critical subregions

false negative error may be more detrimental in this computational analysis application. What is presented herein is a method which identifies the most likely critical subregions $\{R_i\}$, and provides an approximate ranking of these subregions in order of most critical to least critical. This ranking method can be subsequently used to investigate false positive and false negative errors using receiver operating characteristic (ROC) analysis (Provost and Fawcett 2001). A brief discussion of such a ROC analysis is given in numerical example.

A classifier would identify the subregions $\{R_i\}$ only through examination of the spatial configuration of the microstructure. As indicated in Fig. 2, the idea is to characterize the microstructure within a windowed region R_{window} that is at least as large as any given subregion R_j but considerably smaller than the entire specimen R_{total} . It is convenient to define the operator $\iota(R_{\text{window}}) = R_j$ for obtaining the central subregion within R_{window} as shown in Fig. 2. The classifier would use this characterization to provide an estimated likelihood of finding a critical subregion in the center of the window. As noted in the "Introduction," information from detailed microstructural analyses of a set of small samples would be needed to derive the classifier. Once this is done, the classifier can be used over the entire region R_{total} through a moving-window technique to rapidly identify and rank the potentially critical subregions $\{R_i\}$.

Some further notation is needed to explain the microstructural characterization process. Let X be an n_x -dimensional vector that fully describes the microstructure within R_{window} . In our two-phase composites example, the microstructure is divided into $n_p \times n_p$ square pixels. The microstructure vector X is of dimension n_p^2 with coordinates that are the elastic modulus values at the center of each of the pixels. The space of all possible values for the vector X is denoted as Ω_x . For a given set of boundary conditions with externally applied tractions over R_{window} , a stress analysis at the microstructural level is performed. The analysis results corresponding to microstructure X are collected in an n_y -dimensional vector Y ; hence $Y = Y(X)$ can be considered as a function of X . The set of all possible values for Y is denoted by Ω_y . For example, the response vector Y could have dimension $n_y = n_x$ with coordinates equal to the effective strain at the pixel centers. If a range of boundary conditions is of interest, then this information would have to be appended to the vector X ; to simplify notation and to focus on the classifier aspects of the method, only one set of boundary conditions is considered herein.

The simplest microstructural characterizations would be in terms of spatial averages of the field properties in X . From the classifier point of view, such characterizations can be obtained

using classifiers developed through *unsupervised* learning, where the term *unsupervised* indicates that the classifiers are developed without regard to the microstructural response Y . Besides spatial averages, spatial correlation functions, higher moments, and other statistical measures such as the shortest distance between inclusions also fall into this category of characterizations generated from *unsupervised* learning. The value of such characterizations is that they are often related to macrostructural response. Such relationships can be discovered through theoretical or simulation studies. It is noted that such relationships are obtained independently of the characterization process.

Herein, the emphasis is on characterizations of microstructure obtained from classifiers developed through supervised learning, where a microstructure configuration X_i is identified with some class or category of microstructural response $Y(X_i)$. In the following subsections, the mathematical form of this characterization process is presented and it is shown how this type of characterization can be used to identify the potentially critical subregions $\{R_i\}$ in Fig. 1.

Feature Vectors Obtained by Supervised Learning

Consider a set of N simulated random microstructures $X_i \in \Omega_x$ for regions with size and geometry given by R_{window} . In our example, these microstructures consist of circular inclusions placed according to a Poisson field on an underlying matrix to give a representation of a cross section through a fiber-reinforced composite. The raw microstructural data set is given by the set $D_x = \{X_i\}$. To perform supervised learning, the microstructural vectors X_i are augmented with the response vectors $Y_i = Y(X_i)$, yielding the set of $(n_x + n_y)$ -dimensional data vectors $D = \{(X_i, Y_i)\}$. Critical microstructural response would be defined by a subset of values $\Omega_y^* \subset \Omega_y$. This subset might be defined, for example, to be those microstructures in which the maximum strain in the central region of R_{window} is above some threshold. Microstructure X would then be defined to be critical if the associated response $Y(X)$ is critical. Hence, the critical microstructure set is

$$D_x^* = \{X_i; Y_i \in \Omega_y^*\} \quad (2)$$

While it is possible to extend these definitions to three or more classes of microstructure, there are computational complexities associated with such an extension. In "Classifiers for Microstructural Characterization," it is shown how a sequence of two-class analyses is sufficient to produce an approximate ranking of microstructure in order of most critical to least critical.

The dimension n_x of the microstructure X is, in general, too high for subsequent classifier analysis. For example, with a 40×40 pixel microstructural window for our composites problem, the vectors X are of dimension 1,600. Thus, it is necessary to provide low-dimensional representations that capture the most important properties of the microstructure in a compact form. Such representations are developed using *features*. In the most general sense, a microstructural feature is simply a function of X ; in the context of this work, features will be derived using supervised learning from the N simulated augmented data vectors $D = \{(X_i, Y_i)\}$. A collection of M features would define the coordinates of the M -dimensional feature space F as follows:

$$f(X) = [f_1(X), \dots, f_M(X)] \quad (3)$$

To be of practical use, the feature space must have a dimension M that is considerably smaller than the dimension n_x of the microstructure X . Vectors in the feature space F are also denoted as f ; it

will be clear from context whether f refers to the feature space vector or the feature function.

The most straightforward feature functions are defined using class-dependent basis vector expansions. Beginning with the critical microstructure set D_X^* , a set $\{e_j\}$ of n basis vectors would be obtained, where the expansion of any microstructural vector X would be of the form

$$X = \sum_{j=1}^n \alpha_j e_j + X_{\text{res}} \quad (4)$$

Here, α_j =basis vector coefficients and X_{res} =vector of residuals. The essential idea is that, for $X \in D_X^*$, some measure of the size of the residuals, such as the sum of the mean-squared components, would be minimized. Some popular methods to obtain such basis vectors include principal components analysis, where the coefficients α_j are uncorrelated and the basis vectors e_j are orthogonal, and factor rotation methods, where the statistical dependence between the coefficients α_j is minimized but orthogonality of the basis vectors e_j is not enforced. Once the basis vectors are obtained, then the basis vector coefficients can be used directly as feature functions, $f_j(X) = \alpha_j$ for $j=1, \dots, n$. As stated earlier, the number of basis vectors n must be much smaller than the dimension of the microstructure, n_X . In general, the feature vector f can also contain other components to describe the microstructure X .

For our composites example, the basis vectors were obtained from principal components analysis. The number of basis vectors retained in the example is $n=20$ which is two orders of magnitude smaller than the original number of components used to describe the microstructure, $n_X=1,600$. It is noted that while the feature vector $f=(\alpha_1, \dots, \alpha_{20})$ contains only 20 scalars to describe the microstructure X , each of the 20 associated basis vectors e_j contains the full 1,600 components corresponding to the microstructure pixels.

Partition of Feature Space Using Classifiers

In this subsection, a standard tree-based classifier approach is briefly presented in the context of microstructural analysis. In the next subsection, this approach is generalized using decision-theoretic arguments. Of interest herein are classifiers that quantify the likelihood that a particular microstructure X , approximately represented by feature vector f , will have a critical response:

$$C[f] = \Pr[Y(X) \in \Omega_y^*; f(X) = f] \quad (5)$$

where $\Pr[\cdot]$ designates an estimated probability. The most straightforward classifiers $C[f]$ have a finite number of values and can be defined completely by a partition $\{F_j\}$ of F and a corresponding set of likelihood values $\{P_j\}$ as follows:

$$C[f] = P_j \quad \text{where } f \in F_j \text{ for some } 1 \leq j \leq L \quad (6)$$

To make the notation more transparent, it is useful to define an index function which assigns, for each microstructure X , a partition index j

$$j(X) = j[f(X)] = j \quad \text{where } f(X) \in F_j \quad (7)$$

We can also rewrite the expression for the classifier as simply

$$C(X) = C[f(X)] = P_{j(X)} \quad (8)$$

The approach used herein for determining the partition $\{F_j\}$ and associated likelihood values $\{P_j\}$ is based on Bayesian classification trees (Buntine 1992).

Briefly, a binary *classification tree*, such as that shown in

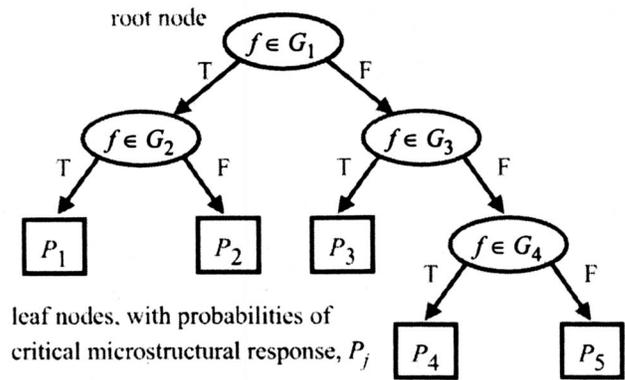


Fig. 3. Binary classification tree with binary expressions $f \in G_i$ at intermediate nodes; feature subset corresponding to leaf 4 is $F_4 = \{f: f \notin G_1, f \notin G_3, f \in G_4\}$

Fig. 3, consists of nodes and downward directed edges with a root node at the upper end and leaf nodes $j=1, \dots, L$ at the lower ends. Two edges are directed from each nonleaf node—they are associated with the binary relations $f \in G_i$ and $f \notin G_i$. For every leaf j , there is a unique sequence of edges from the root to the leaf. The binary expressions and logical values associated with this sequence define a subset F_j of the feature space F . For instance, in Fig. 3, the subset F_4 corresponding to leaf 4 is $F_4 = \{f: f \notin G_1, f \notin G_3, f \in G_4\}$. The collection of all such subsets $\{F_j\}$ form a partition of F . To complete the classification tree, probability values P_j are assigned to the leaves; these values are estimated using the microstructural descriptions and their associated response. For example, P_4 would be the probability that those values of X with features $f(X)$ lying with subset F_4 , that is, those microstructures the tree places at leaf node 4, would be critical, with $Y(X) \in \Omega_y^*$. The feature function f and the classification tree forms a classifier in that every microstructure X would be assigned a feature vector $f=f(X)$ and would subsequently be placed by the tree at a unique leaf node j . This leaf node corresponds to a unique subset F_j that contains f and a probability P_j of being critical. The functional relation between microstructure X and the leaf node j would define the index function $j(X)$ in Eq. (7).

To illustrate the manner in which the tree can classify microstructure, one of the trees from the composites example is briefly examined. A small portion of this tree, shown in Fig. 4, shows a sequence of edges, binary relations involving the feature vector f at the nonleaf nodes, and a probability of critical microstructure at leaf node 1. At the top node of this tree, all microstructures X are considered. Before proceeding through the tree, the feature vector $f(X)$ is rapidly evaluated without any mechanics-based stress analysis. As noted earlier, for the composites example, the components of f are simply the 20 coefficients of the basis function expansion in Eq. (4). Next, the tree uses f to classify the microstructure X according to the binary relations in the nonleaf nodes. For instance, only those microstructure whose first three components of f satisfy the three inequalities shown in Fig. 4 would be propagated along the leftmost branch of the figure. The terminal node at the bottom is leaf node 1, and the tree indicates that 17% of the microstructure X that fall into this node are critical with a maximum stress above a critical threshold. In other words, for any microstructure X falling in this node, the feature vector would be in the first partition, $f(X) \in F_1$, the index function would be

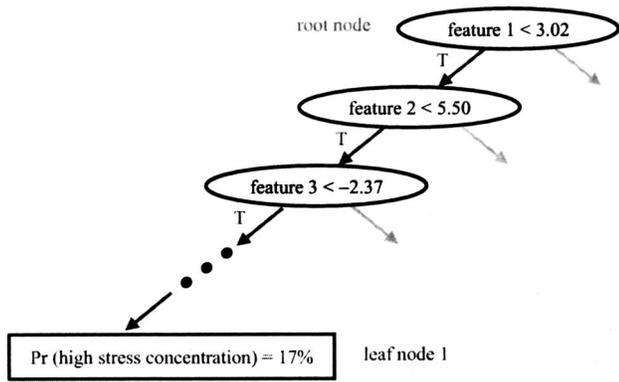


Fig. 4. Small portion of tree $T^{(1)}$ for class $k=1$ that identifies microstructures X with features satisfying $f^{(1)} < 3.02$, $f^{(2)} < 5.50$, $f^{(3)} < -2.37, \dots$, as having 17% likelihood of high stress concentration

$j(X)=1$, and the classifier result would be $C(X)=P_1=17\%$. While this tree can easily and rapidly classify microstructure X in this manner, the real work is in designing the tree.

To design the structure of the tree and the nature of the binary relations, standard classification tree algorithms use the (information) entropy for the heuristic. The entropy of leaf j of a tree is given by

$$I(P_j) = -P_j \log_2 P_j - (1 - P_j) \log_2 (1 - P_j) \quad (9)$$

which is minimized when $P_j=0$ or 1. These limiting values for P_j represent perfect information, or equivalently, zero uncertainty. For the present problem, this situation would occur when either all or no microstructure X assigned to leaf j has critical response. The weighted entropy I_T of the tree T is a weighted average of the leaf information entropies (Hastie et al. 2001)

$$I_T = \sum_{j=1}^L \pi_j I(P_j) \quad (10)$$

where π_j =proportion of the training microstructures assigned to leaf j . While the weighted entropy cannot be minimized in a computationally feasible manner, it has been successfully used in heuristics of several widely used tree construction methods (Breiman et al. 1984; Quinlan 1993). To prevent the size of the tree from growing too large and cumbersome, the heuristics in these methods include a penalty term that is a function of tree complexity.

Decision-Theoretic Classifiers

Standard information-theoretic classifiers, as reviewed in the preceding subsection, turn out to be ill suited for the microstructural analysis problem. Those classifiers must be generalized in a decision theory framework; the derivation of such generalized classifiers is presented in the following. The derivations are given here in outline form. Details of the calculations can be found in Liu and Igusa (2006), which is available for online download from the location given in the citation.

The Bayesian approach begins with utility functions; in the present application, the utility of interest is associated with the microstructure assigned to leaf j . This utility can be expressed as a function $q(P_j)$ of the probability P_j that a microstructure X assigned to leaf j is critical. The expected utility would then be given by

$$E[q|T, P_j] = \sum_{j=1}^L \pi_j q(P_j) \quad (11)$$

where the topology and associated binary expressions of the tree T is treated as a condition. A special case is where the utility of information is given by the negative of the entropy defined in Eq. (9)

$$q(P_j) = -I(P_j) \quad (12)$$

It can then be immediately shown that the expected utility in Eq. (11) is equal to the negative entropy $-I_T$ in Eq. (10). Thus, the weighted entropy heuristic is a special case of the Bayesian heuristic.

In a decision-theoretic approach, it is first necessary to assign each microstructure X to either the critical or noncritical class. This is equivalent to defining a logical classifier C_{logical} that would be equal to the logical value of the statement $Y(X) \in \Omega_Y^*$. In terms of the probability classifier, $C(X)$ as defined in Eq. (8), the logical classifier would be

$$C_{\text{logical}}(X) = \begin{cases} \text{TRUE} & \text{if } P_{j(X)} > P^* \\ \text{FALSE} & \text{otherwise} \end{cases} \quad (13)$$

Next, costs or negative utilities must be assigned to false positive and false negative errors, defined in Eq. (1). Let the cost associated with a false positive result be 1 and that associated with a false negative result be c . Consider microstructure X that the tree assigns to leaf j . If $P_j > P^*$, then the logical classifier will be true, and X will always be considered to be critical. Hence, the only possible error is a false positive, with probability $1 - P_j$ and cost 1. Similarly, if $P_j < P^*$, then the only possible error is a false negative, with probability P_j and cost c . It immediately follows that the conditional expected utility $q(P_j)$ of microstructure placed in leaf j will be given by

$$q(P_j) = \begin{cases} P_j - 1 & \text{if } p_j > P^* \\ -cP_j & \text{otherwise} \end{cases} \quad (14)$$

It can be verified that $q(P_j)$ is maximized when

$$P^* = \frac{1}{c + 1}$$

For the special case where the costs of false negatives and false positives are all equal to 1, we have $P^*=1/2$ and $q(P_j)$ becomes symmetric about $P_j=0.5$, the point of zero information gain so that, for example, $q(0.4)=q(0.6)$.

The decision theory result for the expected utility of the tree would be given by the weighted sum in Eq. (11) with conditional expected utility $q(P_j)$ given by Eq. (14). While this result is a straightforward conclusion from decision theory, it has been noted (Hastie et al. 2001) that it does not necessarily lead to a good search heuristic. We can illustrate this with a simple example that is modified from the example in the aforementioned reference. Consider two trees, denoted as A and B, each with only two leaves, where half of the microstructure data X are placed in each leaf in Tree A, $\pi_1^A = \pi_2^A = 1/2$, and one fourth of the data are placed in leaf 1 of Tree B, $\pi_1^B = 1/4$ and $\pi_2^B = 3/4$. The proportions of critical microstructure in the two leaves of Tree A are $P_1^A = 1/3$, $P_2^A = 2/3$ and the corresponding proportions of Tree B are $P_1^B = 1/6$, $P_2^B = 7/18$. If the costs of false positives and false negatives are equal to 1, then the expected utility of the two trees would be equal, with $E[q|T^A, P_j^A] = E[q|T^B, P_j^B] = -1/3$, in which the decision theory result for the utility function q in Eq. (14) is

used. With the weighted entropy heuristic, Tree B is preferred over Tree A; this is because leaf 1 of Tree B is more pure than leaf 1 of Tree A which results in a significantly smaller entropy. The weighted entropy heuristic gives better guidance for constructing the tree because, after selecting Tree B, this heuristic would lead to tree growth in the direction of leaf 2.

This simple example provides an indication as to why the weighted entropy is desirable as a tree construction heuristic and, as noted above, is used in the most popular classification tree algorithms. Nevertheless, the decision theory approach is attractive in its treatment of unequal costs of false positives and false negatives. The strengths of the decision theory approach to general classification problems has been documented in many applications (Tan and Schlimmer 1990). What is done next in the present paper is to combine the concepts of weighted entropy and expected utilities to obtain a decision theory-based heuristic for constructing classification trees. Such a heuristic is derived by examining the *value of information* from the decision theory point of view.

The basic idea is to compare the preceding decision theory results with the limiting case where infinite information is available. To do this within the Bayesian framework, the probabilities $\{P_j\}$ must be modeled as beta-distributed random variables. If standard Bayesian conjugate prior and posterior distributions are used, the expected value of the information in the training data D_j associated with leaf j would be given by

$$E[Q_j|D_j] = \int_0^1 q(P_j)\phi_j(P_j|D_j)dP_j \quad (15)$$

Here $\phi_j(P_j|D_j)$ = posterior beta density function given the N_j data samples in D_j (DeGroot 1970). The difference in the expected information evaluated when P_j is treated as given information and when it is treated in the Bayesian sense as a random variable has the following asymptotic limit

$$\lim_{N_j \rightarrow \infty} \frac{\log_2(E[Q_j|P_j] - E[Q_j|D_j])}{N_j} = I(P_j) + 1 \quad (16)$$

where the total number of data N_j corresponding to leaf j increases to infinity. This gives the essential relationship between the expected value of information and the weighted entropy. If the weighted sum in Eq. (10) is used to combine the differences in expected information, then the following is obtained:

$$\lim_{N_j \rightarrow \infty} \frac{1}{N} \sum_{j=1}^L \log_2(E[Q_j|P_j] - E[Q_j|D_j]) = I_T + 1 \quad (17)$$

where we use $N_j = \pi_j N$ and note that the proportion π_j which would appear in the denominator of Eq. (16) would cancel with the weights in Eqs. (10). Since the weighted entropy I_T is used for the search heuristic, it is important to see how slight changes in I_T are related to changes in the expected information difference $E[Q_j|P_j] - E[Q_j|D_j]$. The result is approximately

$$\frac{1}{N \log 2} \sum_{j=1}^L \frac{\Delta(E[Q_j|P_j] - E[Q_j|D_j])}{E[Q_j|P_j] - E[Q_j|D_j]} \approx \Delta I_T \quad (18)$$

This shows that, in the entropy search heuristic, it is the relative changes in the expected information difference that is important. The weights π_j , which are proportional to the number of data in each leaf, are not used. As shown in Eq. (17), these weights are implicitly included in the expected information difference.

With the relationship between weighted entropy and Bayesian decision theory established, it is now possible to include the important difference of utilities associated with false negatives and false positives. A simple approximate approach to include this difference is to use a revised probability P'_i defined by

$$P'_i = \frac{cP_i}{1 + (c-1)P_i} \quad (19)$$

in place of the probability P_i in Eqs. (6) and (14). This revised probability can be interpreted as the probability of critical configuration at leaf j if the number of critical configurations is scaled by the relative cost of false negatives, c . It can be shown that the minimum of the entropy $I(P'_j)$ at leaf j evaluated using this revised probability is approximately at $P'_j = (c+1)^{-1}$; this corresponds to the maximum of the expected utility in Eq. (11).

The preceding results for a decision-theoretic classifier are new; it is shown in the following how these results can be combined with standard Bayesian techniques to reduce the complexity of the classifier. The essential idea is to use prior probabilities $\phi(T)$ for the tree T ; the prior acts essentially as a weight to influence tree design. For instance, if trees with a relatively small number of edges and nodes are preferred, then the prior $\phi(T)$ would be larger for such trees. The final heuristic for the design of Bayesian classification trees would be given by the product of the expected value of the information and the prior probability for the tree

$$T_{\text{optimal}} = \arg \max_T \{E[Q|T, \text{data}] \phi(T)\} \quad (20)$$

In practice, the optimization required in this equation is difficult; hence, heuristic algorithms have been developed. In the present context, a Bayesian approach to tree construction (Buntine 1992) is most relevant.

It is noted that any feature f_j that is not included in the binary expressions of the tree can always be removed from the feature vector f . Hence, an important benefit of the tree reduction process is in reducing the dimension of the feature vector space F .

Classifiers for Microstructural Characterization

In this subsection, it is shown how a sequence of classifiers can be developed for characterizing the microstructure in a specimen R_{total} larger than R_{window} . The first classifier $C^{(1)}$ and corresponding tree $T^{(1)}$, given by C and T in the preceding section, would be used to analyze the microstructure X within R_{window} . Specifically, the classifier would estimate the probability that the microstructure response is critical, as defined by the relation $Y(X) \in \Omega_Y^{(1)} = \Omega_Y^*$. If this probability exceeded some threshold value $P_{\text{threshold}}^{(1)}$, then the central subregion $R = \iota(R_{\text{window}})$ within R_{window} , as defined by the operator ι , would be marked as a potentially critical subregion. By moving the window throughout R_{total} , this classifier would identify a collection of potentially critical subregions $\{R_i\}$ as illustrated in Fig. 2. The union of these subregions would yield the potentially critical region, $\hat{R}_{\text{critical}}^{(1)}$. With only a single classifier, however, it is not possible to obtain $\hat{R}_{\text{critical}}^{(1)}$ with low likelihoods of false positives and false negatives. This is because, to keep the likelihood of false negative low, the threshold probability must be set so that $1 - P_{\text{threshold}}^{(1)}$ is nearly 1. This would make $\hat{R}_{\text{critical}}^{(1)}$ large, and would result in a high likelihood of false positives. To address this problem, two additional classifiers are defined that will reduce the size of the potentially critical region without increasing the likelihood of false negatives.

A natural choice for the second class of microstructural response is simply the complement of the first, $\Omega_Y^{(2)} = (\Omega_Y^*)'$, representing uncritical cases. For the composites example, such uncritical cases correspond to microstructure response with small strain in the central pixel. Using the data subset $D_X^{(2)}$ in Eq. (2), the feature vector $f^{(2)}$ can be derived, as explained after that equation. The classifier $C^{(2)}$ and tree $T^{(2)}$ for this second class would, in general, not be a simple complement of the classifier and tree associated with the first class. With this second classifier, a second pass of the moving window R_{window} would be performed. Since the probability $P_j^{(2)}$ of the second classifier $C^{(2)}$ is the estimated probability that a region is not critical, a second upper threshold value $P_{\text{threshold}}^{(2)}$ would be used with $C^{(2)}$ to eliminate some of the regions identified in the first pass. In this manner, the second classifier reduces the size of the potentially critical region without significantly affecting the probability of false negatives. This process is explained explicitly in the following:

1. Using a moving window R_{window} over the total region R_{total} the first pass approximation for the critical region would be obtained using the first classifier, $C^{(1)}$

$$\hat{R}_{\text{critical}}^{(1)} = \bigcup_{R_{\text{window}} \subset R_{\text{total}}} \{ \iota(R_{\text{window}}) : C^{(1)}[f^{(1)}(R_{\text{window}})] > P_{\text{threshold}}^{(1)} \} \quad (21)$$

Here the union is a collection of subregions, $\hat{R}_{\text{critical}}^{(1)} = \{R_{\text{window},1}, R_{\text{window},2}, \dots\}$, and $f^{(1)}(R_{\text{window}})$ = feature vector associated with the microstructure in R_{window} . The operator $\iota(R_{\text{window}})$ which extracts the central portion of the window, is applied since the goal of the classification is to identify the actual site of critical behavior, and the trees have been developed in terms of the responses within this central portion. Thus, the classifiers actually provide predictions about the criticality of the central portion of the window based on the microstructural contents of the entire window.

2. The second pass would produce a refinement using the second classifier, $C^{(2)}$

$$\hat{R}_{\text{critical}}^{(2)} = \bigcup_{R_{\text{window}} \in \hat{R}_{\text{critical}}^{(1)}} \{ \iota(R_{\text{window}}) : C^{(2)}[f^{(2)}(R_{\text{window}})] < P_{\text{threshold}}^{(2)} \} \quad (22)$$

where the union remains as a collection of subregions.

In theory, it is possible to replace the aforementioned two classifiers with a single classifier to do essentially the same classification problem. Such a single classifier would have at least three classes, which predicts the likelihood that a microstructure X within R_{window} is a member of the critical class, noncritical class, and one or more intermediate classes. This approach was not pursued herein because of complexity issues. As noted earlier, the optimization of a classification tree is computationally infeasible, and suboptimal search algorithms must be used in practice. The reason for proposing a separation of the classification problem into two separate problems, one for critical microstructure and the other for noncritical microstructure, is to reduce computational complexity. It was found that using a single classification tree, there were convergence difficulties and the tree construction algorithm produced larger, yet less accurate trees.

While the use of two classifiers provides a smaller set of potentially critical locations than the use of a single classifier, the result for the set of potentially critical subregions $\hat{R}_{\text{critical}}^{(2)}$ is still too large to be practical. The essential source of the

difficulty is in the binary nature of the classifiers, with assessments restricted to relations of the form $Y(X) \in \Omega_Y^*$ and $Y(X) \notin \Omega_Y^*$. To address this problem, a third classifier is used that, while still within the binary framework, provides a more detailed assessment of microstructural response. The third classifier is based on a relational operator $<$ that is used to compare microstructural responses, where $Y(X_A) < Y(X_B)$ implies that the response of microstructure X_B is more critical or severe than that of microstructure X_A . If the response function $Y(X)$ is scalar, then the operator $<$ can simply be replaced by the inequality $<$; if $Y(X)$ is a vector or tensor, then a more general form of response comparison is needed. In the composites example, a scalar response function was used, given by the strain in the center of the microstructural window. Once this operator is defined, then the third classifier would be given by a generalization of Eq. (5)

$$C^{(3)}[f] = \Pr[Y(X_A) < Y(X_B) : f(X_A, X_B) = f] \quad (23)$$

Here, the feature function is defined for the $2n_X$ -dimensional vector (X_A, X_B) . The data subset associated with this third classifier would be defined by a generalization of Eq. (2)

$$D_{(X_A, X_B)}^{(3)} = \{ (X_A, X_B) : Y(X_A) < Y(X_B) \} \quad (24)$$

There are two ways to define the feature vector $f^{(3)} \in F^{(3)}$. The first would be to use a basis vector expansion based on an analysis of the subset $D_{(X_A, X_B)}^{(3)}$ similar to that described after Eq. (2); the components of the feature vector would be given by $\alpha_j^{(3)}$ of Eq. (4) with X replaced by (X_A, X_B) . A simpler way to define feature vector $f^{(3)}$ is to use the product space $F^{(3)} = F^{(1)} \times F^{(2)} \times F^{(1)} \times F^{(2)}$ of feature vectors for the first two classes resulting in a $2(n^{(1)} + n^{(2)})$ -dimensional vector

$$f^{(3)}(X_A, X_B) = [f^{(1)}(X_A), f^{(2)}(X_A), f^{(1)}(X_B), f^{(2)}(X_B)] \quad (25)$$

In either case, a third classification tree $T^{(3)}$ would be used with partitions $\{F_j^{(3)}\}$ of the feature space $F^{(3)}$. At the j th leaf of this tree, $P_j^{(3)}$ would be the probability that microstructures X_A and X_B with feature value $f(X_A, X_B) \in F_j^{(3)}$ has responses satisfying $Y(X_A) < Y(X_B)$.

The third classifier $C^{(3)}$ would be used to rank order the set of regions $\hat{R}_{\text{critical}}^{(2)}$ identified by the first two classifiers in the following manner, in which a higher ranked region has a more critical response:

3. The probability P_i that a region $R_{\text{window},i} \in \hat{R}_{\text{critical}}^{(2)}$ is more critical than any other region in $\hat{R}_{\text{critical}}^{(2)}$ is approximated by the following product, based on the third classifier, $C^{(3)}$:

$$P_i = \prod_{j \neq i} C^{(3)}[f^{(3)}(R_{\text{window},j}, R_{\text{window},i})] \quad (26)$$

If the indices of the m largest values for the probabilities $\{P_i\}$ are given by $i(1), i(2), \dots, i(m)$, then

$$\hat{R}_{\text{critical}} = \bigcup_{j=1}^m \iota(R_{\text{window},i(j)}) \quad (27)$$

would be the final estimate for the critical region. Here, the indices $i(j)$ would be rank ordered such that $P_{i(1)} > \dots > P_{i(m)}$; hence $\iota(R_{\text{window},i(1)})$ is the estimate of the most likely critical subregion.

Implicit in the approximation for the probabilities P_i is the assumption that the events $Y(X_A) < Y(X_B)$ and $Y(X_C) < Y(X_B)$ are independent for distinct regions X_A, X_B , and X_C . This assumption

would be valid provided that the regions are sufficiently separated in the spatial sense. Even if the independence assumption was not accurate, the preceding approximation for P_i would still produce a useful estimate for the indices $i(1), \dots, i(m)$ and the associated rank order of probabilities. If the regions X_i overlap or are otherwise spatially aggregated, clustering techniques can be used to form groups of regions. The central region would be used to represent each group before performing the rank order analysis.

Application to 2D Composites

General Concepts

This section contains a detailed analysis of the microstructural example that was described in the "Introduction," and subsequently used to illustrate the concepts used in the proposed method. The microstructure considered is that of a fiber reinforced composite with parallel fibers. The problem is rendered two dimensional by considering only a cross section of the material taken perpendicular to the fiber direction. The matrix and inclusions are modeled as homogeneous, isotropic, elastic-brittle materials with deterministic properties; the only uncertainty is the position of the inclusions. The Young's moduli, Poisson's ratio, and ultimate strains for the two phases are given by E_J , ν_J , and ϵ_J with index J =matrix or inclusion. The fibers, or inclusions, are cylindrical with diameter d , and the applied load σ^∞ is uniaxial, perpendicular to the inclusions as shown in Fig. 1. In the remainder of this subsection, an overview is presented to show how the classifier-based method can be used to identify the most likely locations for damage initiation, without performing a finite element, or other computationally intensive analysis. The next subsections provides specific details for a numerical example.

To begin the analysis, the size of the window region, R_{window} , must be determined. This window size is used in generating the microstructural samples X_i in the data set $D=\{(X_i, Y_i)\}$ which is used in supervised learning to determine the feature set to be used in classification. The tradeoff in selecting window size is between retaining more of the mechanically relevant microstructure with a large window, and reducing the dimension n_X of X , the vector describing the microstructural geometry, with a small window. A large window results in inefficiencies associated with the high dimensionality of X_i and with the computation cost in evaluating $Y_i=Y(X_i)$ for each sample i . On the other hand, a small window may not allow the supervised learning algorithm to capture the relevant, physically meaningful features of the microstructure. The concept of a representative volume element (RVE) with periodic boundary conditions (Hill 1963) is widely used for obtaining an appropriate moving window size. For the present application, the size of the RVE should be large enough such that its averaged elastic modulus is not sensitive to the positions of the inclusions (Drugan and Willis 1996). To help find an appropriate RVE, Drugan and Willis (1996) provided a micromechanics-based constitutive equation relating the averaged stress and strain for a wide range of random linear elastic composites. By comparing the results of this constitutive equation with those obtained from a constant overall modulus tensor, it is found that the minimum RVE size is typically on the order of just several inclusion diameters. Once the RVE is chosen, each microstructural sample X_i can be generated with randomly placed inclusions using a spatial Poisson point process with standard techniques to handle periodicity and inclusion overlap. To represent the possible vari-

ability in local volume fraction in the total material specimen, R_{total} , samples X_i with different numbers of inclusions are needed.

Damage would occur in the central region ($u(R_{\text{window}})$) of the window when the local strain exceeds the fracture strain. Hence a natural choice for the microstructural response Y is the ratio of the local and fracture strain in the material phase in the center of the window. The set of critical responses would be defined in terms of a lower limit: $\Omega_Y^{(1)}=\{Y \geq Y^{(1)}\}$. Due to the periodicity in the analysis of X_i , it is always possible to shift the microstructure within the window so that the largest value of the strain ratio is at the center; the shifted microstructure is denoted as $X_i^{(1)}$. Thus, for the critical microstructures data set $D_X^{(1)}=\{X_i^{(1)}: Y_i \in \Omega_Y^{(1)}\}$, the most likely damage initiation location of each sample $X_i^{(1)}$ is at the center. Similarly, for the uncritical microstructures $D_X^{(2)}=\{X_i^{(2)}: Y_i \in \Omega_Y^{(2)}\}$, the set of critical responses would be defined in terms of an upper limit: $\Omega_Y^{(2)}=\{Y \leq Y^{(2)}\}$. The microstructure $X_i^{(2)}$ is obtained from X_i by shifting the microstructure so that the largest value of the strain ratio is at the corners. In the way, it is highly unlikely that damage initiation would occur in the centers of the microstructures in the uncritical set $D_X^{(2)}$. It can be seen that the computationally intensive analysis $Y(X_i)$ is needed once for each sample X_i to determine the two data sets $D_X^{(k)}$ for classes $k=1,2$ which correspond to critical and non-critical microstructural response. After obtaining the basis vector coefficients $\alpha_j^{(k)}$, which define the feature subvector $f^{(k)}$ for $j=1, \dots, n^{(k)}$, three classifiers $C^{(1)}, C^{(2)}$, and $C^{(3)}$ are derived, following the general procedure described above.

Numerical Example

In the specific numerical example analyzed herein, the stiffness of the inclusion is three times that of the matrix, $E_{\text{inclusion}}=3E_{\text{matrix}}$, and the Poisson's ratio for both phases is $\nu=1/3$. The fracture strain for the inclusion is five times as large as that of the matrix, $\epsilon_{\text{inclusion}}=5\epsilon_{\text{matrix}}$, so that damage initiation always occurs in the matrix. The volume fraction of inclusions is set to be 12%. The composite is subjected to stresses σ_x^∞ in the x direction under plane-strain conditions. Using the aforementioned RVE concept, an appropriate window size was determined to be $5d$ where d is the inclusion diameter. A triangular spring lattice model (Ostoja-Starzewski et al. 1996) is used for the stress analysis. The spring length is chosen to be $d/8$ to accurately capture the stress concentrations between inclusions. With this spring length and window size, a 40×40 lattice is needed to model each microstructure X_i .

In the Poisson point process generation of microstructures X_i , the minimum allowable distance between inclusion centers was set at $3d/2$. Three volume fractions, 19, 25 and 31% were used, corresponding to 3, 4, and 5 inclusions/sample. A total of $N=600$ samples were generated and analyzed. For each of the resulting two data sets, $D_X^{(1)}$ and $D_X^{(2)}$, $n^{(1)}=n^{(2)}=10$ feature coordinates were retained. As noted in "Feature Vectors Obtained by Supervised Learning," the basis vector coefficients were obtained by principal components analysis. With these feature coordinates, the original 1,600-dimensional vectors X_i are reduced to 20 dimensions. Fig. 5 shows three of the feature basis vectors derived from the critical training samples, denoted, as $e_j^{(1)}$ in Eq. (4), and Fig. 6 shows three derived from the noncritical training samples, denoted as $e_j^{(2)}$ in Eq. (4). The darker regions represent regions of higher stiffness, which correspond to high likelihood of inclusions; the lighter regions represent locations of low likelihood of inclusions. It is noted that these basis vectors do not necessarily

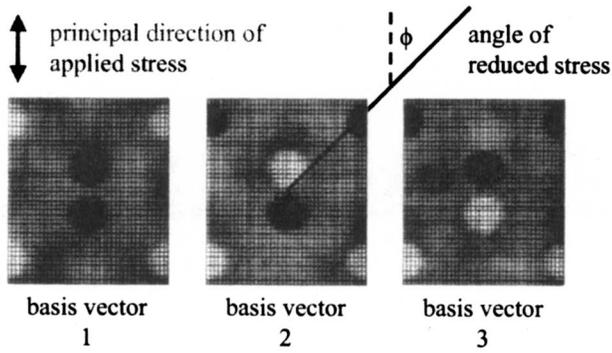


Fig. 5. First three basis vectors $e_1^{(1)}$, $e_2^{(1)}$, and $e_3^{(1)}$, for critical response, showing high (dark) and low (light) likelihood of inclusions; angle ϕ corresponds to direction of long-range stress reduction

represent critical microstructures: the classifiers $C^{(k)}$, used as described in “General Concepts,” are needed to identify such microstructures. The feature basis vectors indicate *sensitivity* of the local stress at the center to inclusion patterns, but as noted earlier, it is the classifiers $C^{(k)}$, that *identify* critical inclusion patterns most likely to induce high local stresses.

The basis vectors can partially be interpreted based on the Eshelby solution to the problem of an inclusion in a homogeneous, infinite matrix subject to remote tension (Liu 2003). This solution describes stress and strain concentrations caused by a stiff, circular inclusion. Critical basis vector $e_1^{(1)}$ shows two inclusions aligned with the loading axis, for which configuration superposition of the Eshelby solution would predict extreme strain concentration. Noncritical basis vector $e_1^{(2)}$, on the other hand, shows inclusions (dark regions) far away from the center, and a light color, indicating a low likelihood of inclusion presence, near the center of the window. Such a configuration would be expected to result in low strain near the center of the window. The other basis functions shown are more difficult to interpret mechanically. The mechanical meaning of the basis functions is a topic of active research by the writers.

With the data sets, $D_X^{(k)}$, for classes $k=1,2$ and the associated feature vectors $f^{(k)}$, the classifiers $C^{(k)}$ are obtained next. The corresponding Bayesian classification trees $T^{(k)}$ were constructed so that the prior density for the tree in Eq. (20) favored relatively small trees: $T^{(1)}$ had 19 nodes and eight leaves, $T^{(2)}$ had only 11 nodes and four leaves. As discussed earlier, Fig. 4 shows a small portion of the tree $T^{(1)}$ for class $k=1$. This portion of the tree shows that if microstructure X within a window R_{window} has features satisfying $f_1^{(1)} < 3.02, f_2^{(1)} < 5.50, f_3^{(1)} < -2.37, \dots$, then

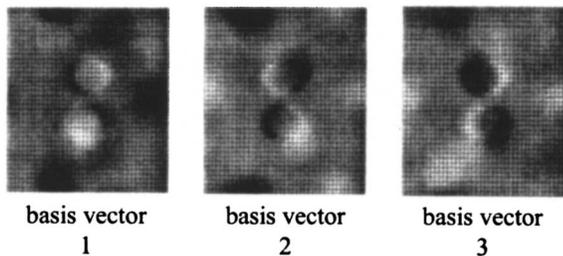


Fig. 6. First three basis vectors $e_1^{(2)}$, $e_2^{(2)}$, and $e_3^{(2)}$, for noncritical response, showing high (dark) and low (light) likelihood of inclusions

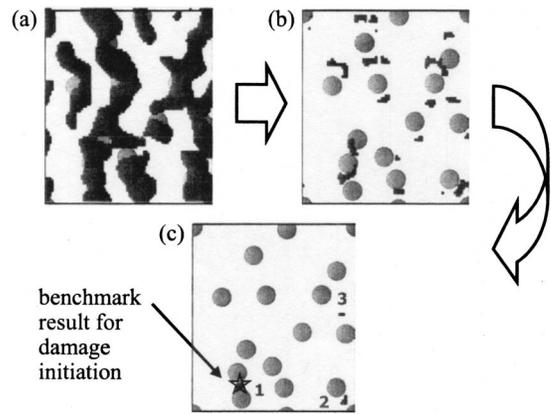


Fig. 7. Evolution of estimate for likely subregions of damage initiation, $\hat{R}_{\text{critical}}$: (a) first pass approximation obtained from classifier $C^{(1)}$; (b) second pass reduction using $C^{(2)}$; and (c) third pass reduction and comparative analysis using $C^{(3)}$, with numerals indicating estimates for three most likely locations

there is a 17% likelihood that the central portion of the window is characterized by high stress concentrations. Following “General Concepts,” the comparative classifier $C^{(3)}$ is also obtained using the simplified form for the feature vector in Eq. (25). This classifier, which operates on the product of feature spaces, is considerably more complex than the first two classifiers: the corresponding tree $T^{(3)}$ had 115 nodes and 58 leaves.

To test the classifier approach, ten material specimens are generated, each with region R_{total} that is four times as large as the window-sized regions R_{window} used for the samples X_i . To maintain the volume fraction of 12%, each specimen includes 15 randomly located inclusions. The benchmark analysis is performed using an 80×80 spring network loaded until damage initiates by failure of a spring element. The classifier-based procedure described in “General Concepts” is used to determine the most likely locations of damage initiation. The three main steps of the analysis, illustrated in Fig. 7, are as follows: (1) the first pass approximation for the critical region is obtained using the first classifier, $C^{(1)}$, resulting in a rather large region $\hat{R}_{\text{critical}}^{(1)}$ for the possible locations of damage initiation; (2) the second pass using the second classifier, $C^{(2)}$, eliminates most of $\hat{R}_{\text{critical}}^{(1)}$, resulting in a substantially smaller region for consideration; and (3) the comparative analysis using the third classifier, $C^{(3)}$, provides the final estimates for the three most likely locations for damage initiation, indicated by the numerals in the figure. A simple clustering algorithm, based on the average linkage method, was used to group the closely spaced regions. This method identifies clusters simply based on the average Euclidean distance between all points in each cluster. This simple clustering algorithm is implemented in the statistics toolbox of the MATLAB scientific computing program, and this implementation was used to perform the clustering in this study. The benchmark analysis result for damage initiation, indicated by a star in Fig. 7(c), coincides with the most likely site from the classifier analysis for this specimen. In Fig. 8, four additional representative specimens are shown with comparisons between benchmark and classifier results. The benchmark result coincides with the most likely location identified by the classifier in Figs. 8(a and b), the second most likely location in Fig. 8(c), and with none of the top four likely locations identified by the classifier in Fig. 8(d). While the result in Fig. 8(d) indicates that the classifier has room for improvement, it is noted that the

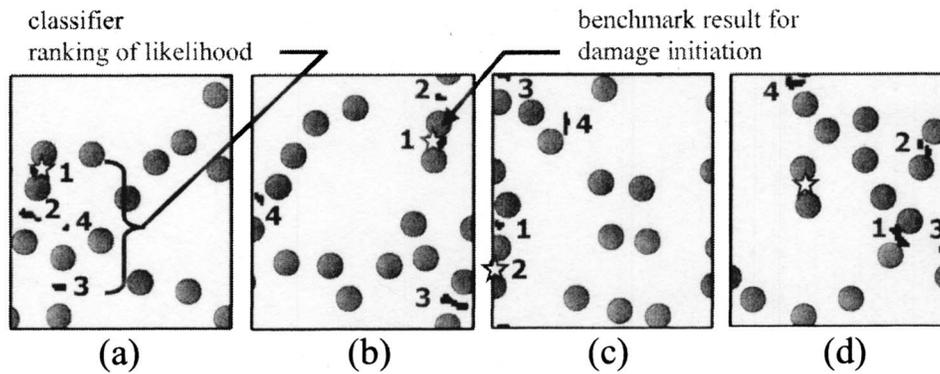


Fig. 8. Estimates for most likely locations of damage initiation compared with benchmark results for four representative specimens

classifier results are based on a relatively small training data set and that its accuracy can be enhanced by increasing the size of the training data set. Overall, in 80% of the comparisons, the benchmark result was among the four most likely locations and in half of the comparisons, the benchmark result coincided with the most likely location identified by the classifier. These results are remarkably accurate, considering that the classifier analysis is several orders of magnitude faster than the benchmark analysis for identifying the locations of damage initiation.

It is noted that the concepts of false positives and false negatives with nonequal costs were used in developing a decision-theory based heuristic that incorporates the desirable properties of the entropy-based heuristic. This heuristic was needed to drive the construction of the tree-based classifiers. It was not possible to perform a meta comparison of several classifiers, simply because standard classifier-based methods, such as using off-the-shelf classifier algorithms to construct a single classification tree for finding the most likely damage initiation point, performed so poorly. Moreover, the intent herein is not to criticize such widely used algorithms, but to explain how minor modifications to the underlying heuristic can result in a classifier particularly suited to the microstructure characterization problem. It is noted that if the notions of false positives and false negatives, which were vital in defining this heuristic, are also applied to the final classifier in the example problem then the probability of false positive is high and the probability of false negatives would be very close to zero. Specifically, if m is the total number of locations within each specimen, n is the final number of locations with each specimen that is retained as possible locations for the highest stress level, and p is the probability that the highest stress level in a specimen lies within these n grid points, then the probability of false positives is $1-p/n$ and the probability of false negatives is p/m . Here, each location refers to a cluster of points, so that in Fig. 8, there are $n=4$ locations per specimen identified out of a total of approximately $m=1,000$ possible locations. For $n=4$, the probabilities of false positive and false negative are $1-0.8/4=80\%$ and $0.8/1,000=0.08\%$, whereas for $n=1$, these probabilities are $1-0.5/1=50\%$ and $0.5/1,000=0.05\%$. Hence, if these points were plotted on a ROC curve, they would lie nearly horizontally at the top of ROC coordinates. To make use of this comparative information, a measure of cost of false positive and false negative errors is needed for this problem of finding the most highly stressed location in a specimen. This is beyond the present study, and is different from the costs of false positive and false negative errors used in the classifier heuristic, which is the focus of this paper.

Conclusions

In the foreseeable future, high-fidelity simulations of microstructural response are largely limited to very small material sample sizes. The research work presented herein was conducted with this important limitation in mind. Specifically, an approach was developed that can use such simulation results to characterize and analyze materials with random microstructure at the mesoscale. Bayesian classifiers are used in a manner that is considerably different from standard classifier applications in pattern recognition and other data-mining problems. What is noteworthy herein is: (1) the generalization of the classifier heuristic from the standard information-theoretic to the broader decision-theoretic setting that can be tied to microstructural analysis goals; and (2) the manner in which a sequence of classifiers are used to identify the critical region of a material specimen. A relatively simple example is used to illustrate the method.

Outstanding issues include the estimation of the errors inherent in the procedure, the characterization of microstructure with random phase geometry, the consideration of multiple sets of boundary conditions, and the analysis of more complex examples. In the paper, there was only a brief discussion of the basis vector expansion that was used in an intermediate step; the compatibility of Bayesian classifiers to various expansion methods, particularly approaches using factor rotation and principal surfaces and manifolds, could be explored. While more complex classifiers using flexible discriminants or support vector machines can also be examined, there is always the fundamental statistical issue of bias-variance tradeoff where more complex methods can fit data more faithfully, but end up with poorer predictive accuracy (Vapnik 1996; Cherkassky and Mulier 1998).

While the classifiers can be used directly in problems such as the identification of damage initiation sites, they can also provide insight into macromechanical behavior of composites. Furthermore, classifiers can guide the construction of computational models. A main challenge in analyzing the response of material microstructures is the degree of discretization necessary to obtain reliable computational results. By identifying regions of a microstructure that are likely to participate in damage initiation, for example, classification has the potential to guide multiscale modeling and discretization of microstructures.

It was noted earlier that there has been some seminal work on materials with random microstructure by Zabaras and his colleagues that also use statistics and classifiers. The two approaches examine the materials problem from different viewpoints, how-

ever, and work at different aspects of the materials problem. While the methods in the aforementioned work do not require detailed microstructural simulations, they also are not developed to extract information from such simulated data. Thus, the two approaches are complementary. The combination of these approaches would allow for the analysis of information from both experimentally sampled microstructural data, such as captured by the orientation distribution function, and small-scale mechanical simulations of microstructure. While this is not addressed herein, it would be a natural next step in this research topic.

Acknowledgments

The writers wish to acknowledge the support of this work by NSF through grant no. DMI-0423582.

Notation

The following symbols are used in this paper:

- $C[f]$ = classifier;
- C_{logical} = logical classifier;
- $C(X)$ = alternative classifier notation;
- c = cost of false negative error;
- D = a set of microstructure and response vector pairs $\{X_i, Y_i\}$;
- D_X = a set of multiple microstructure vectors;
- $D^{(k)}$ = training data set for response class k ;
- D_X^* = subset of D^* corresponding to critical response;
- d = fiber diameter;
- E = elastic modulus;
- $E[\cdot]$ = expectation;
- e_j = basis vectors defining coordinates of feature space;
- F = feature space;
- $f(X)$ = feature space vector;
- $f_i(X)$ = feature function;
- G_i = feature space partition;
- I = entropy of leaf node;
- I_T = weighted entropy of classification tree;
- $j(X)$ = partition index;
- M = dimension of feature space;
- N = number of microstructures used in classification training;
- N = total number of data used in tree construction;
- N_j = number of data used in estimation of leaf node likelihoods;
- n = number of basis functions used to represent microstructure;
- n_p = the number of pixels along a side of R_{window} ;
- n_X = dimension of X ;
- n_Y = dimension of microstructural response vector;
- P'_i = revised likelihoods;
- P_j = likelihood values of the classifier;
- p^* = threshold likelihood;
- $q(P_j)$ = utility function;
- R_{critical} = microstructural subdomain in which critical behavior occurs;
- $\hat{R}_{\text{critical}}$ = total classifier identified critical subregion, union of R_i ;
- R_i = classifier identified potentially critical subregions;

- R_{total} = microstructural domain;
- R_{window} = microstructural subdomain used in moving window classification;
- T = classification tree;
- T_{optimal} = optimal classification tree;
- X = microstructure vector;
- X_i = the i th element of a set of microstructure vectors D_X ;
- X_{res} = residual from microstructure expansion;
- Y_i = i th element of a set of microstructure response vectors $Y_i = Y(X_i)$;
- $Y(X)$ = microstructural response function;
- α_j = microstructure expansion coefficients;
- ε = ultimate strain;
- $\iota(R_{\text{window}})$ = central pixel of a microstructural subdomain;
- ν = Poisson's ratio;
- π_j = entropy weight;
- σ^∞ = remote applied stress;
- $\phi(T)$ = prior probabilities for classification tree;
- Ω_X = space of all possible values of X ;
- Ω_Y = space of all possible values of Y ;
- Ω_Y^* = set of critical response vectors; and
- $<$ = relational operator for comparing microstructure criticality.

References

- Acharjee, S., and Zabarar, N. (2003). "A proper orthogonal decomposition approach to microstructure model reduction in rodrigues space with applications to optimal control of microstructure-sensitive properties." *Acta Mater.*, 51, 5627–5646.
- Adams, B. L., Etinghof, P. I., and Sam, D. D. (1994). "Coordinate free tensorial representation of n-point correlation functions for microstructure by harmonic polynomials." *Proc., 10th Int. Conf. on Textures and Materials*, Clausthal, Germany, 387–294.
- Adams, B. L., Henrie, A., Henrie, B., Lyon, M., Kalidindi, S. R., and Garmestani, H. (2001). "Microstructure-sensitive design of a compliant beam." *J. Mech. Phys. Solids*, 49, 1639–1663.
- Arwade, S. R., and Grigoriu, M. (2004). "Characterization and modeling of random polycrystalline microstructures with application to intergranular fracture." *J. Eng. Mech.*, 130(9), 997–1006.
- Baxter, S. C., and Graham, L. L. (2000). "Characterization of random composites using moving window technique." *J. Eng. Mech.*, 126(4), 389–397.
- Beaudoin, A. J., Dawson, P. R., Mathur, K. K., and Kocks, U. F. (1995). "A hybrid finite element formulation for polycrystal plasticity with consideration of macrostructural and microstructural linking." *Int. J. Plast.*, 11, 501–521.
- Benveniste, Y. (1987). "New approach to the application of Mori-Tanaka's theory in composite materials." *Mech. Mater.*, 6, 147–157.
- Berryman, J. G. (1985). "Measurement of spatial correlation functions using image processing techniques." *J. Appl. Phys.*, 57, 2374–2384.
- Breiman, L., Friedman, R. A., Olson, J. H., and Stone, C. J. (1984). "Statistics probability." *Classification and regression trees*, Wadsworth, Belmont, U.K.
- Buntine, W. (1992). "Learning classification tree." *Stat. Comput.*, 2, 63–73.
- Buryachenco, V., Pagano, N. J., Kim, R. Y., and Spowart, J. E. (2003). "Quantitative description and numerical simulation of random microstructures of composites and their elastic moduli." *Int. J. Solids Struct.*, 40, 47–72.
- Chen, T., Dvorak, G. J., and Benveniste, Y. (1992). "Mori-Tanaka estimates of the overall elastic moduli of certain composite materials." *J. Appl. Mech.*, 59, 539–546.

- Cherkassky, V., and Mulier, F. (1998). *Learning from data*. Wiley, New York.
- Corbin, S. F., and Wilkinson, D. S. (1994). "Influence of particle distribution on the mechanical response of a particulate metal matrix composite." *Acta Metall. Mater.*, 42, 1311–1318.
- DeGroot, M. H. (1970). *Optimal statistical decisions*, McGraw-Hill, New York.
- Drugan, W. J., and Willis, J. R. (1996). "Micromechanics based nonlocal constitutive equation and estimates of representative volume element size for elastic properties." *J. Mech. Phys. Solids*, 44, 497–524.
- Everett, R. K., and Chu, J. H. (1993). "Modeling of nonuniform composite microstructures." *J. Compos. Mater.*, 27, 1128–1144.
- Ganapathysubramanian, S., and Zabarbas, N. (2004). "Design across length scales: A reduced order model of polycrystal plasticity for the control of microstructure-sensitive properties." *preprint* (personal communication).
- Hastie, T., Tibshirani, R., and Freidman, J. H. (2001). *The elements of statistical learning*, Springer, New York.
- Hill, R. (1963). "Elastic properties of reinforced solids: Some theoretical principles." *J. Mech. Phys. Solids*, 11, 357–372.
- Lewandowski, J. J., Liu, C., and Hunt, W. H., Jr. (1989). "Effects of matrix microstructure and particle distribution on fracture of an aluminum metal matrix composite." *Mater. Sci. Eng., A*, 107, 241–255.
- Li, M., Ghosh, S., Rochmond, O., Weiland, H., and Rouns, T. N. (1999a). "Three dimensional characterization and modeling of particle reinforces metal matrix composites. Part II: Damage characterization." *Mater. Sci. Eng., A*, 265, 221–240.
- Li, M., Ghosh, S., Rochmond, O., Weiland, H., and Rouns, T. N. (1999b). "Three dimensional characterization and modeling of particle reinforces metal matrix composites. Part I: Quantitative description of microstructural morphology." *Mater. Sci. Eng., A*, 265, 153–173.
- Liu, H. (2003). "Bayesian classifiers for uncertainty modeling with applications to global optimization and solids mechanics problems." Ph.D. thesis, Johns Hopkins Univ., Baltimore (<http://www.ce.jhu.edu/abstracts/#Liu2>).
- Liu, H., and Igusa, T. (2006). "Feature-based classifiers for design optimization." *Res. Eng. Des.*, <http://dx.doi.org/10.1007/s00163-006-0024-4> (published electronically).
- McLachlan, G. J. (1992). *Discriminant analysis and statistical pattern recognition*, Wiley, New York.
- Ostoja-Starzewski, M., Sheng, P. Y., and Alzebdeh, K. (1996). "Spring network models in elasticity and fracture of composites and polycrystals." *Comput. Mater. Sci.*, 7, 82–93.
- Provost, F., and Fawcett, T. (2001). "Robust classification for imprecise environments." *Mach. Learn.*, 42, 203–231.
- Pyrz, R. (1994). "Quantitative description of the microstructure of composites. Part I: Morphology of unidirectional composite systems." *Compos. Sci. Technol.*, 50, 197–208.
- Quinlan, J. R. (1993). *Programs for machine learning*, Morgan Kaufman.
- Saganopan, D., and Pitchumani, R. (1998). "Application of genetic algorithms to optimal tailoring of composite materials." *Compos. Sci. Technol.*, 58, 571–589.
- Spowart, J. E., Maruyama, B., and Miracle, D. B. (2001). "Multiscale characterization of spatially heterogeneous systems: Implications for discontinuously reinforced metal matrix composite microstructures." *Mater. Sci. Eng., A*, 307, 51–66.
- Tan, M., and Schlimmer, J. C. (1990). "Two case studies in cost-sensitive concept acquisition." *Proc., AIAA-90*, 854–860.
- Torquato, S., Hyun, S., and Donev, A. (2003). "Optimal design of manufacturable three-dimensional composites with multifunctional characteristics." *J. Appl. Phys.*, 94(9), 5748–5755.
- Vapnik, V. (1996). *The nature of statistical learning theory*, Springer, New York.