

Decomposition and Visualization of Fourth-Order Elastic-Plastic Tensors

Alisa G. Neeman¹ Rebecca Brannon² Boris Jeremić³ Allen Van Gelder¹ and Alex Pang¹

¹Computer Science Department, UC Santa Cruz

²Department of Mechanical Engineering, University of Utah

³Department of Civil and Environmental Engineering, UC Davis

Abstract

Visualization of fourth-order tensors from solid mechanics has not been explored in depth previously. Challenges include the large number of components ($3 \times 3 \times 3 \times 3$ for 3D), loss of major symmetry and loss of positive definiteness (with possibly zero or negative eigenvalues). This paper presents a decomposition of fourth-order tensors that facilitates their visualization and understanding. Fourth-order tensors are used to represent a solid's stiffness. The stiffness tensor represents the relationship between increments of stress and increments of strain. Visualizing stiffness is important to understand the changing state of solids during plastification and failure. In this work, we present a method to reduce the number of stiffness components to second-order 3×3 tensors for visualization. The reduction is based on polar decomposition, followed by eigen-decomposition on the polar "stretch". If any resulting eigenvalue is significantly lower than the others, the material has softened in that eigen-direction. The associated second-order eigentensor represents the mode of stress (such as compression, tension, shear, or some combination of these) to which the material becomes vulnerable. Thus we can visualize the physical meaning of plastification with techniques for visualizing second-order symmetric tensors.

Keywords: Stiffness tensor, tensor decomposition, Reynolds glyph.

1. Introduction

Modeling and simulations of static and dynamic behavior of solids and structures made up of various materials (soils, concrete, wood, steel, etc.) is a focus of current research in civil, mechanical and other branches of engineering. In such modeling and simulations, a pivotal role is played by the fourth-order 3-D stiffness tensor, which can be derived for any point inside the domain of interest. Visualization of fourth-order tensors representing stiffness of solid materials has not been explored in depth previously. Until now, much of the visualization community was unaware of the problem, and the engineering community has had limited visualization tools to study these and other higher order tensors. One of the purposes of this paper is to introduce this problem domain to the visualization community. We present the terminology and describe some of the challenges. Then we describe and evaluate one method that makes it possible to

filter for the most critical changes in stiffness and visualize them as second-order tensors.

In early works on physical modeling for computer graphics, it was popular to use gimbal-jointed spring meshes, with stiffness assigned to springs according to various heuristics. Very early engineering studies used similar models for analysis of truss bridges and similar structures. In these oversimplified models, stress and strain (roughly, force and deformation) could be represented as vectors [TF88, VG98]. Stiffness is the derivative of stress with respect to strain. When stress and strain are represented as vectors, stiffness is a second order tensor, or matrix. Informally, vectors are first-order tensors and matrices are second-order tensors.

The above models are inadequate for representing the elastic properties of solids because they cannot account for shear. To account for shear effects, both stress and strain require second-order tensors. Then stiffness, the derivative of stress with respect to strain, becomes a fourth-order tensor. Because of the *symmetries* present in the stress and strain

tensors of natural solid materials, the fourth-order stiffness tensor has up to 36 independent quantities (rather than 81).

Visualization of 36 dimensional quantities, whether cast as matrices or tensors, is largely unexplored territory. One thesis of this paper is that meaningful decompositions are essential to reduce the complexity of a visualization to a level that allows human comprehension, while preserving the important information in the simplified representations.

This paper's approach to visualization is based on analysis of the behavior of stiffness tensors and identification of properties that indicate susceptibility to failure and other important changes of material behavior. This reduces the complexity of the visualization task and addresses the specific need to better understand the physical effects of a changing stiffness tensor. Readers are encouraged to consult the technical appendix and animation submitted to this workshop as supplemental material.

One novelty in the proposed technique is *polar decomposition* followed by eigen-decomposition (also called spectral decomposition) producing eigenvalues and eigentensors or eigenvectors. Polar decomposition is known in the plasticity and mechanics literature, but our use of it to analyze the stiffness tensor is new, as far as we know. Nielsen and Schreyer [NS93] showed that eigen-decomposition on the *symmetric part* of the stiffness tensor produces information regarding the mode of a material's vulnerability to stress. We are able to obtain other details with our approach. Like Nielsen and Schreyer, when a significantly low eigenvalue occurs, we can detect loss of stiffness and the associated eigentensor that represents the stress mode that will produce the greatest deformation. The polar decomposition also appears to tell us about the changing flow direction of permanently deforming material.

Our method is applicable to general elasto-plastic models. It can also be applied to any material behavior where the stiffness tensor changes as the loading progresses (pressure sensitive elasticity, damage, etc.).

2. Related Work

Related work falls into two categories; visualization and solid mechanics methods. For visualization we look at some current approaches for visualizing second-order and fourth-order tensors.

For solid mechanics we look at the motivation for various methods to decompose the fourth order stiffness tensor, particularly when major symmetry is lost.

2.1. Related Visualization Techniques

A substantial part of our visualization involves the use of Reynolds glyphs and extensions. Here we review previous work along these lines.

Ellipsoidal glyphs have long been used to represent second-order symmetric tensors with applications in engineering. Hashash *et al.* [HYW03] provides a good overview of the glyphs used in solid mechanics with a list of advantages and shortcomings of each. They use color to show the magnitude of the vector from the origin to the surface points. We extend the color scheme to cover both negative and positive values of vector length since shape alone is ambiguous.

In HOT-lines, Hlawitschka and Scheuermann [HS05] used spherical harmonics to compute Reynolds glyphs more efficiently. Their method could be applied to *fully symmetric* tensors of even order on a regular grid, and were applied to fourth-order diffusion tensors from Magnetic Resonance Imaging (DT-MRI) data. The spherical harmonics also helped detect global maxima which represented direction of major eigenvectors for fiber tracking in DT-MRI [HS05].

Kriz *et al.* also used a fourth-order Reynolds glyph representation for materials [KGM95]. The representation came from Christoffel's equation for waves propagating in anisotropic media. The shape of the glyph is mapped to wave velocity magnitude or the eigenvalues, while the color on the glyph is mapped to vibration direction or eigenvectors.

In a recent paper, Basser and Pajevic [BP07] applied eigen-decomposition to a fourth-order symmetric covariance tensor for use in DT-MRI. The resulting eigentensors and eigenvalues were used to represent and visualize variability. The authors also noted that the mathematical underpinnings of these eigentensors were from Lord Kelvin's work to identify modes of deformation in linear elastic material and classify symmetries in anisotropic materials. Basser and Pajevic visualized the second-order eigentensors and fourth-order covariance tensors as Reynolds glyphs. While the decomposition was applied to fully symmetric positive definite covariance tensors, it motivated our own research with respect to more general fourth-order stiffness tensors.

2.2. Related Decomposition Methods

Nielsen and Schreyer were the first to use eigen-decomposition to determine deformation modes associated with materials breaking (bifurcation) [NS93]. Their criterion for bifurcation was loss of positive definiteness. They performed eigen-decomposition on $\text{sym}(\mathbf{E})$, the *symmetric part* of stiffness tensor \mathbf{E} , since a \mathbf{E} is positive definite if and only if $\text{sym}(\mathbf{E})$ is positive definite. When positive definiteness was lost, they used the eigentensor of the zero eigenvalue of $\text{sym}(\mathbf{E})$ as the mode of bifurcation. This occurs *before* \mathbf{E} becomes singular, which was the previous criterion. However, the eigentensor of $\text{sym}(\mathbf{E})$ is not necessarily an eigentensor of \mathbf{E} itself.

This paper uses polar decomposition to obtain a different symmetric component of \mathbf{E} for eigen-decomposition. The technical appendix shows that when the polar symmetric

component has a zero eigenvalue, the associated eigentensor is also an eigentensor of the stiffness tensor itself. This provides an alternate representation of the failure mode.

3. Modeling Elasticity and Plasticity of Solids

Mechanical engineers model the elastic and plastic properties of solids with tensors of various orders. This section reviews tensor notation, conventions, and operations. Then the basics of modeling solids are stated, and modeling issues are discussed.

3.1. Tensor Notation and Operations

Tensors are linear operators that can be represented as multi-dimensional arrays of coefficients. For 3-D solids, a fourth-order tensor is a $3 \times 3 \times 3 \times 3$ array, a second-order tensor is a 3×3 array, etc. The order of a tensor is the same as the number of subscripts needed to write a typical element. Thus, if \mathbf{E} is a fourth-order tensor, a typical element is denoted by E_{ijkl} . Scalars, vectors, and matrices represent tensors of orders zero, one, and two, respectively.

Operations using tensors are usually denoted using the *Einstein convention* that repeated indices in different tensors are implicitly summed; e.g., matrix multiplication is denoted as $C_{ij} = A_{ik}B_{kj}$, rather than the explicit equation,

$$C_{ij} = \sum_k A_{ik}B_{kj}. \quad (1)$$

This operation is called a *single contraction* in tensor terminology.

In elasticity (and many other physical processes) the *double contraction* operator is important. It is denoted by “:” as an infix symbol, and involves summing over two indices, e.g.,

$$C_{ij} = \mathbf{A} : \mathbf{B} = A_{ijkl}B_{kl} = \sum_k \sum_\ell A_{ijkl}B_{kl}. \quad (2)$$

Double contraction can also be applied to two fourth-order tensors, yielding a new fourth-order tensor.

Where many operations on first and second order tensors use single summation, their generalizations to second and fourth order tensors use double summation. The above example of double contraction is thus the generalization of multiplying a matrix by a vector. Two important cases are the scalar *inner product*, $A_{ij}B_{ij}$, and the *dyad* or *outer product*, $A_{ij}B_{kl}$ that results in a fourth-order tensor.

3.2. Stress, Strain and Stiffness

In 3-D, *stress* is force per unit area, *strain* is fractional change in length (e.g., a strain of .001 means a length L increased to $1.001L$), and *stiffness* is the ratio of *stress* to *strain*. Since strain is dimensionless, stiffness has the same units as stress.

Table 1: How a stiffness tensor can change with increasing stress.

Material Behavior	Stiffness Tensor Properties	
	elastic	symmetric
elastic-plastic	asymmetry	eigenvalue reduction, non-singular if hardening
failure, localized	asymmetry	non-positive definite, possibly singular

A solid object such as a steel spring stretches when a tension force is applied, and returns to its original length when the force is removed. This means that upon unloading (removing the force) all the deformation (except possibly rigid motion) is recovered; i.e., the solid returns to its original shape. Hooke’s Law [Lov44] describes this simple case. More generally, the material’s stress-strain relationship is described by the *constitutive equation*:

$$\sigma_{ij} = E_{ijkl} \epsilon_{kl} \quad (3)$$

where σ_{ij} is a second-order symmetric stress tensor, ϵ_{kl} is a second-order symmetric strain tensor, and $\mathbf{E} \equiv E_{ijkl}$ is the fourth-order stiffness tensor. Indices i, j, k, ℓ range over 1, 2, 3 and represent three orthogonal spatial axes, x , y , and z .

3.3. Stiffness Changes

Solid materials undergo a sequence of changes under increasing loads. In general, increasing loads beyond what is called the *elastic limit* induce changes in the stiffness tensor described in table 1.

3.3.1. Elastic Strain Increment

As an example of elastic deformation, imagine a paper clip: it will (elastically) change shape to hold a small stack of papers, suffering virtually no permanent deformation after the papers are removed. Theories of plasticity typically presume that there exists some contiguous domain of stress levels for which permanent (non-recoverable) deformation is negligible. However, it is not assumed that the stress/strain relationship is linear within this elastic limit. Therefore the constitutive equation is generalized to express the relationship incrementally with the *tangent* stiffness tensor:

$$\dot{\sigma}_{ij} = E_{ijkl} \dot{\epsilon}_{kl} \quad (4)$$

where $\dot{\sigma}_{ij}$ is a second-order symmetric stress *increment* tensor, $\dot{\epsilon}_{kl}$ is a second-order symmetric strain *increment* tensor, and $\mathbf{E} \equiv E_{ijkl}$ is the fourth-order *tangent* stiffness tensor. As before, indices i, j, k, ℓ range over 1, 2, 3 and represent three orthogonal spatial axes, x , y , and z .

3.3.2. Elastic-Plastic Strain Increment

When a paper clip is deformed by a large amount, only a portion of the deformation is recovered upon removal of the

load. As discussed by Brannon [Bra07], this familiar observation is generalized in plasticity theory by additively decomposing the strain increment into two parts: a recoverable elastic part plus a permanent plastic part.

3.3.3. Yielding and Plastic Strain Increment

Materials modeled as elastic-plastic are characterized by a *yield function*, $f(\sigma_{ij})$, a function of stress. The yield criterion, which delineates points at which stress causes elastic-plastic deformation versus pure elastic deformation is defined at

$$f(\sigma_{ij}) = 0 \quad (5)$$

The yield surface is the *zero isosurface* of the yield function. It is a convex surface in six-dimensional stress space. Stress states on the interior of the yield surface correspond to elasticity, while stress states outside the yield surface are unattainable except via plastic loading.

Elastic-plastic materials are also often characterized with a plastic flow function. Like the yield function, this is a function of stress, and its gradient describes the direction of plastic flow, in stress space. In *associated* (or normal) elasto-plasticity, the flow direction (i.e., the direction of the plastic-strain increment) is parallel to the yield surface normal while *non-associated* plastic flow can deviate from the normal.

3.3.4. Hardening and Softening

Classical equations of plasticity theory [Bra07] model the stress increment as a fourth-order elastic-plastic tangent stiffness tensor acting on the total strain increment. In the absence of elastic-plastic coupling, associativity corresponds to symmetry of the elastic-plastic tangent stiffness. For a non-hardening material, the yield surface is *fixed* and, therefore, the stress increment during plastic loading is constrained to move tangent to the yield surface. In this case, the elastic-plastic tangent stiffness tensor is a positive-semidefinite rank-1 projection operator that essentially removes the part of the elastic stress increment that points outside the yield surface.

For a *hardening* material, the yield surface can expand outward away from the stress so that increasingly high levels of stress are required to continue plastic loading. In this case, the elastic-plastic tangent stiffness tensor is positive definite and invertible.

For a *softening* material, the yield surface contracts in the neighborhood of the stress, so diminishing stresses are required to continue plastic loading. In this case, the elastic-plastic tangent stiffness tensor can become *non-positive definite* because there exist strain increment modes that decrease the stress level required for continued plastic loading. Depending on the nature of the loading, softening can lead to an instability in the material response that is characterized by intense localization of the deformation into narrow bands.

Mathematically, the onset of such behavior is determined by examining properties of the fourth-order elastic-plastic tangent stiffness tensor.

3.3.5. Localized Failure

With enough bending force, a paper clip will deform permanently and eventually break. According to early literature a solid (in this case a clip) is said to fail *locally* when the stiffness tensor is singular. Here, locally means that the material breaks or liquefies at a set of points or region inside the domain [MH79, NS93, RMW96]. Singularity is identified by at least one zero eigenvalue. Later literature identified localized failure at the point of loss of positive definiteness [NS93]. A primary goal of this visualization work is to filter the field of evolving tangent stiffness tensors to identify those that are near or beyond the localization threshold and to then visualize the corresponding modes of susceptibility to failure.

4. Approach

Our approach is to filter the tensor field from a single time step for the plastic part of deformation and then visualize second-order tensors representing the modes of change in stiffness (hardening or softening elasto-plasticity). These second-order tensors are calculated via the following steps

1. Unroll the 3 tensor into a 6×6 matrix.
2. Perform a polar decomposition on the 6×6 matrix, producing two 6×6 matrices, called the *rotation part* and the *stretch part*, which is symmetric.
3. Perform an eigen-decomposition on the *stretch part* of the polar decomposition, yielding 6 eigentensors and 6 real eigenvalues.
4. Select a single eigentensor and compose it from a 6-D vector into a symmetric 3×3 tensor.
5. Visualize the second-order eigentensor with a glyph that shows its structure and also reflects its eigenvalue.

4.1. Unrolling Tensors

Numerical methods for eigen-decomposition and polar decomposition only exist for matrices (rather than fourth-order tensors). The fourth-order stiffness tensor E_{ijkl} (representing both elastic and plastic components of stiffness) is transformed into a 6×6 matrix \mathbf{E} , called the *Mandel components*. Other 6×6 matrices and 6-vectors computed during the processing can be converted back into fourth-order tensors and symmetric second-order tensors, respectively.

The straightforward representation of a 3×3 tensor would be as a 9-vector with one component for each tensor element. As applied first by Mandel, and later rigorously justified by others [Man62, Its00, Hel01, Tar06, Bra07], due to the symmetry of the 3×3 tensor space of interest, an orthonormal change of basis can force the last three components of the 9-vector to be zero. This orthonormal change of basis in 9-D

simply consists of 45° 2-D degree rotations on the three pairs of vector components that correspond to symmetric pairs of *off-diagonal* tensor elements. Similarly, if the set of transformations is restricted to those that produce symmetric results, after applying the orthonormal change of basis, a 6×6 matrix suffices. The equations are given in the technical appendix submitted as supplemental material.

4.2. Application Constraints

Our proposed method can be applied to stiffness tensors for a constrained set of solid mechanical simulations. It is applicable to general elasto-plastic models, and it can also be applied to any material behavior where the stiffness tensor changes as the loading progresses (pressure sensitive elasticity, damage, etc.).

However, the method is based on *small deformation theory*, and is not applicable to simulations outside this theory. “Small deformation” refers to problems with only small displacement gradients. This means that the stiffness tensors are cast with respect to the reference configuration such that any changes in stiffness associated with rigid material rotation are eliminated from consideration. Another limitation is that the current visualization relies on the second-order tensors being symmetric, which constrains the fourth-order stiffness tensor to exhibit minor symmetry.

5. Polar Decomposition of Fourth-Order Tensors

Polar decomposition is a technique to separate a matrix into two component matrices, $\mathbf{E} = \mathbf{Q}\mathbf{S}$, where \mathbf{S} is a pure stretch (symmetric positive-semidefinite matrix), and \mathbf{Q} is a pure rotation (orthonormal matrix) [Hig86, Gan90, HS90, HP94, ZZ95, GVL96]. As described in Section 4, we compute on tensors by transforming them into matrices and vectors. We believe that polar decomposition *in this space* is new. This section describes how we use the results of the polar decomposition for analysis and visualization of solid mechanics simulations.

The technical appendix, submitted as supplemental material, reviews technical details on a new computational method for the polar decomposition in the presence of ill-conditioning or singularity [VG08]. As described below, this case is central to our methodology, and we did not find previously published methods to be satisfactory.

5.1. Polar Stretch and Mode of Vulnerability

Eigen-decomposition is performed on \mathbf{S} , the stretch matrix component of the polar decomposition. This can be done using the LAPACK library [ABD*90]. We analyze the eigen-decomposition of \mathbf{S} to find the eigenvector associated with its smallest eigenvalue. Converting this eigenvector to a symmetric second-order tensor (the reverse of unrolling), we interpret this tensor as a mode of strain. Being an eigentensor,

after multiplying by the eigenvalue, it is the corresponding mode of stress. This is expected to be the mode to which the material is most vulnerable to failure.

The initial eigenvalues of stiffness are based on the bulk modulus and shear modulus of the material, and the initial stiffness matrix is a pure stretch; i.e., its rotation component is the identity.

Upon inelastic loading, the stiffness may change in a manner such that resistance to certain stress increments is considerably diminished. A polar decomposition of \mathbf{E} quantifies this degradation. Eigenvalues of \mathbf{S} (the “stretches”) are larger than their initial values if the material has stiffened. They are smaller if it has become more compliant. The corresponding eigentensors of \mathbf{S} represent loading modes. Suppose, for example, the stretch eigenvector associated with the smallest eigenvalue is $[0, 0, 0, 1, 0, 0]^T$ (vectors are column vectors). The corresponding normalized stretch eigentensor is:

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & \sqrt{1/2} \\ 0 & \sqrt{1/2} & 0 \end{bmatrix}$$

This is a pure shear. If the corresponding eigenvalue is 0.75 of its initial value, then the material is 25% less resistant (stiff) to that particular loading mode than it was initially. Selecting the lowest stretch eigenvalue can therefore reveal which regions are suffering significant plastic loading even in a multi-material simulation. As the stiffness matrix approaches singularity, having a robust computation procedure, as described in the technical appendix, becomes crucial. An eigenvalue approaching zero is considered to be a failure mode for the material.

If \mathbf{E} is nonsingular but with negative eigenvalues, the physical meaning of the polar decomposition is unclear. In this case, a symmetric-antisymmetric decomposition could still be used, but its meaning would also be unclear. The occurrence of a negative eigenvalue in the stiffness matrix corresponds to dynamic softening instability, corresponding to a change in type for the differential equations of motion. The stiffness at a point becomes no longer physically meaningful because the process becomes scale dependent. Therefore, any points in the domain having a negative stiffness should be filtered out of the visualization on physical grounds unless the strain increment at that point belongs to the eigenspace of the stiffness associated with positive stiffness eigenvalues.

5.2. Polar Rotation and Elasto-Plastic Materials

The rotation part of the polar decomposition also carries interesting information. As we have mentioned in Section 3.3, elastic-plastic materials are characterized by a yield function and a plastic flow “direction”. In the case of *associated* elasto-plasticity, the flow “direction” is perpendicular to the “surface” on which the yield function is zero. In this case there is no rotation component (\mathbf{Q} mentioned at the beginning of this section is the identity). We have put “direction”

and “surface” in quotation marks because we are operating in stress space, which is either the space of symmetric second-order tensors or their unrolled versions, which are 6-vectors. Thus “direction” and “surface” do not have their usual 3-D meanings.

We conjecture that for *non-associated* materials, the polar rotation (in 6-D) quantifies the degree of misalignment of the yield surface normal and the flow direction. In other words, we conjecture that applying the polar rotation to the yield surface normal produces the flow direction to a good approximation. This conjecture has been confirmed anecdotally, but it requires further empirical and theoretical evaluation.

6. Eigentensor Glyphs and Physical Meaning

Reynolds glyphs visually depict characteristic modes of stress or strain [BP07]. The formula for a point on the surface of the second-order Reynolds glyph is (tensor notation is reviewed in Section 3.1):

$$\sigma_N = n_i \sigma_{ij} n_j \quad (6)$$

where i, j ranges over 1, 2, 3 and 1, 2, 3 are orthogonal axes. n is a vector of the direction cosines perpendicular to the plane of interest for a ray of unit length.

Stress is characterized by five basic modes, and a stress tensor may represent a pure mode or some combination of them. The modes are:

1. spherical tension (three equal positive eigenvalues)
2. spherical compression (three equal negative eigenvalues)
3. triaxial compression (two equal eigenvalues with the distinct eigenvalue more compressive than the repeated eigenvalues)
4. triaxial tension (two equal eigenvalues with the distinct eigenvalue less compressive than the repeated eigenvalues)
5. pure shear (one zero eigenvalue, with the others equal and opposite in sign)

An isotropic stress (with three equal eigenvalues) corresponds to a spherical Reynolds glyph (figure 1a). This case is similar to isotropy for DT-MRI [WPG*97] but we can represent either compression or tension, distinguished by color. Triaxial tension or triaxial compression produce an axisymmetric stress (symmetric about one axis) and is characterized by two equal eigenvalues. In this case, formula 6 produces an axisymmetric Reynolds glyph with the axis of symmetry aligned with the eigenvector of stress corresponding to the distinct eigenvalue (figure 1b). The disk shape in figure 1b is similar to planar anisotropy for DT-MRI second-order tensors, also characterized by two equal eigenvalues [WPG*97].

A state of pure shear stress (where one eigenvalue is zero and the other two are equal and opposite in sign) produces a




Figure 1: Reynolds Glyph representations of modes of stress. Cool colors (blue and cyan) represent negative values, hot colors (red and yellow) represent positive. X, Y, Z axes are red, green, blue.

glyph such as the one in figure 1c. The distinct butterfly appearance of pure shear is caused by equal magnitude and opposite values across orthogonal directions on a single plane.

Compared to the second-order glyphs for DT-MRI applications [WPG*97, Kin04], we see a larger number of distinct shapes because stretching constants can take both negative and positive values. Stress, strain and eigentensors glyphs need not be simple ellipsoids. Negative values serve to push the vector n from equation 6 in the reverse direction relative to the center of the sphere. However, like DT-MRI glyphs, they will be symmetric since the tensors they represent are symmetric. We also use a color mapping that covers both negative and positive values, to distinguish between tension and compression. This is unnecessary for diffusion tensor data.

7. Results

We tested our technique on two simple experiments from geomechanics. Both experiments consisted of 224 finite elements, each with 8 integration points on an irregular grid (a total of 1792 points for the volume). The experiments were run on OpenSees, a finite-element simulator for performance of structural and geotechnical systems subjected to earthquakes [Ope05]. Data for the visualization was gathered during the runs, which took 10.5–12 minutes on an Intel 6600 2.4 Ghz dual core desktop with 2.0 gigabytes of RAM. The visualization runs at 1–4 frames per second, depending on window size. Excerpts are shown in the figures. The software and experiment scripts are available from the authors and will soon be available through the VEES project website [Nee08].

The experiments started with self-weight analysis of the soil, compressing along $-Z$, for 25 time steps. This is a common practice when studying soil behavior. The initial compression from the soil’s own weight induces deformation and may also change the stiffness, depending on the model. The second part of the experiment was two point loads on the top of the volume. The loads included a small component




Figure 2: Color-mapped displacements, time step 124, Drucker-Prager material. Color map is clamped at 0.08. X, Y, Z axes are red, green, blue.




Figure 3: Top view (Z from above) of eigentensors for Drucker-Prager material, time step 124, colored by minimum "stretch" eigenvalue.




Figure 4: Eigentensors for Dafalias-Manzari material, time step 124. Black cubes denote negative eigenvalues. Color map clamped at 60623.

along $-Z$ (0.9659 kiloNewtons) and a large lateral component (1294 kiloNewtons). This part of the experiment ran for another 100 steps. Figure 2 shows the volume colored by displacement (Euclidean distance from initial position) on the last time step. The black arrows point to where the loads were applied along $+X$.

In Figs. 3 and 4, we drew glyphs for each integration point in the volume and the glyph size was scaled by the size of the encapsulating finite element. These eigentensor glyphs were colored by the lowest eigenvalue and scaled across all time steps to give an overall idea of softening or hardening. We used a special glyph for the case of a negative eigenvalue in the stiffness tensor (a black cube) so the engineer could quickly identify singularity and the inability of the simulator to produce the mode of the instability.

The first soil material was Drucker-Prager. The material simply fails under tension, but remains unchanged under compression (a non-hardening material). There were no changes in the material during the self-weight steps. The material experienced compression in front of the point loads and the reverse (tension) behind the point loads. There is a trail of points that have undergone singularity behind the point loads. The green and cyan glyphs indicate zero and very slightly negative lowest eigenvalues, respectively (see Fig. 3). These glyphs show a clear change in orientation indicating the mode of failure.

The second simulated material was Dafalias and Manzari's plasticity sand model [DM04]. The model's stiffness is pressure-dependent and non-associated. In order to make the experiment repeatable we used the material constants given in the article, with the exception of the initial void ratio (e_0). This material softens when it is denser (with smaller void ratios). Since we wanted to test our method's ability to detect softening, we reduced the e_0 to 0.65. In this experiment, self-weight induced hardening. The greatest hardening occurred at the bottom of the volume, where the finite-element nodes

were fixed and experiencing the greatest load from above. At first, stiffness increased with the depth of the volume, as expected with self-weight loading. With continued self-weight softening occurred in the volume we saw some integration points with negative eigenvalues. As the experiment continued, softening occurred along the side boundaries of the volume, but there was frequent change between elastic and elastic-plastic states. It was difficult to correlate the stress to the behavior because of the changing yield surface and nonassociativity. Another issue was that isotropic hardening and softening were represented by the eigentensor associated with the smallest eigenvalue, but with isotropy there are five equal eigenvalues so its orientation was random. However, we did notice softening on the top surface near the applied point loads late in the experiment (Fig. 4).

8. Conclusions and Future Work

A meaningful decomposition is an essential first step toward visualizing fourth-order tensor fields. In this work, we presented a decomposition method robust to loss of major symmetry and approaching singularity. As a result, engineers have the ability to visualize the critical mode of stress and reduction in stiffness across a volume. Many challenges remain. In some cases, a material may have more than one mode in which it softens, and there is as yet no technique to visualize a field that has multiple second order tensors at each location. Continuous visualization techniques tensor fields representing non-symmetric and non-positive definite tensors have not yet appeared, and interpolation techniques for irregular grids of finite elements could be improved. In many cases it is critical to respect internal element boundaries; shared faces may represent the border between two discrete materials, such as layers of sand and clay. Finally, we still have not attempted dealing with Large Deformation Theory. We plan to attack some of these problems in the near future and hope that this early work will open the door to these new challenges.

Acknowledgments

Work was partially supported by a GAANN fellowship and the UCSC/Los Alamos Institute for Scalable Scientific Data Management (ISSDM). We thank the anonymous reviewers for their feedback and helpful suggestions.

References

- [ABD*90] ANDERSON E., BAI Z., DONGARRA J., ET AL.: LAPACK: A portable linear algebra library for high-performance computers. In *Supercomputing '90* (Washington, DC, 1990), IEEE, pp. 2–11. 5
- [BP07] BASSER P. J., PAJEVIC S.: Spectral decomposition of a fourth-order covariance tensor: Applications to diffusion tensor MRI. *Signal Processing* 87, 2 (2007), 220–236. 2, 6
- [Bra07] BRANNON R. M.: Elements of phenomenological plasticity: Geometrical insight, computational algorithms, and applications in shock physics. In *Shock Wave Science and Technology Reference Library: Solids I*. Springer, 2007, ch. 6. 4
- [DM04] DAFALIAS Y. F., MANZARI M. T.: Simple plasticity sand model accounting for fabric change effects. *J. of Engineering Mechanics* 130, 6 (2004), 622–634. 7
- [Gan90] GANDER W.: Algorithms for the polar decomposition. *SIAM Journal on Scientific and Statistical Computing* 11 (1990), 1102–1115. 5
- [GVL96] GOLUB G. H., VAN LOAN C. F.: *Matrix Computations*, third ed. Johns Hopkins, 1996. 5
- [Hel01] HELNWEIN P.: Some remarks on the compressed matrix representation of symmetric second-order and fourth-order tensors. *Computer Methods in Applied Mechanics and Engineering* 190 (16 February 2001), 2753–2770(18). 4
- [Hig86] HIGHAM N. J.: Computing the polar decomposition – with applications. *SIAM Journal on Scientific and Statistical Computing* 7 (1986), 1160–1174. 5
- [HP94] HIGHAM N. J., PAPADIMITRIOU P.: A parallel algorithm for computing the polar decomposition. *Parallel Computing* 20 (1994), 1161–1173. 5
- [HS90] HIGHAM N. J., SCHREIBER R. S.: Fast polar decomposition of an arbitrary matrix. *SIAM Journal on Scientific and Statistical Computing* 11 (1990), 648–655. 5
- [HS05] HLAWITSCHKA M., SCHEUERMANN G.: HOT-lines: Tracking lines in higher order tensor fields. In *Visualization '05* (Minneapolis, 2005), pp. 27–34. 2
- [HYW03] HASHASH Y. M., YAO J. I.-C., WOTRING D. C.: Glyph and hyperstreamline representation of stress and strain tensors and material constitutive response. *International Journal for Numerical and Analytical Methods In Geomechanics* 27, 7 (2003), 603–626. 2
- [Its00] ITSKOV M.: On the theory of fourth-order tensors and their applications in computational mechanics. *Computer Methods in Applied Mechanics and Engineering* 189, 2 (1 September 2000), 419–438(20). 4
- [KGM95] KRIZ R., GLAESSGEN E., MACRAE J.: Eigenvalue-eigenvector glyphs: Visualizing zeroth, second, fourth and higher order tensors in a continuum. In *Proceedings of the NCSA-NIST Workshop on Modeling the Development of Residual Stresses During Thermoset Composites Curing* (September 1995). 2
- [Kin04] KINDLMANN G.: Superquadric tensor glyph. In *Vissym'04* (2004), pp. 147–154. 6
- [Lov44] LOVE A. E. H.: *A Treatise of the Mathematical Theory of Elasticity*. Dover Publications, Inc., 1944. 3
- [Man62] MANDEL J.: Plastic waves in an infinite three dimensional medium. *J. de Mécanique* 1 (1962), 3–30. 4
- [MH79] MAIER G., HUECKEL T.: Nonassociated and coupled flow rules of elastoplasticity for rock-like materials. *International Journal of Rock Mechanics and Mining Science* 16 (1979), 77–92. 4
- [Nee08] NEEMAN A. G.: VEES open source project home page. neesforge.nees.org/projects/vees/, 2008. 6
- [NS93] NEILSEN M. K., SCHREYER H. L.: Bifurcations in elastic-plastic materials. *International Journal of Solids and Structures* 30, 4 (1993), 521 – 544. 2, 4
- [Ope05] OPENSEES DEVELOPMENT TEAM (OPEN SOURCE PROJECT): OpenSees: open system for earthquake engineering simulations. 1998-2005. 6
- [RMW96] RIZZI E., MAIER G., WILLAM K.: On failure indicators in multi-dissipative materials. *Int'l J. of Solids and Structures* 33, 20-22 (1996), 3187–3214. 4
- [Tar06] TARANTOLA A.: *Elements for Physics: Quantities, Qualities, and Intrinsic Theories*. Springer, 2006. 4
- [TF88] TERZOPOULOS D., FLEISCHER K.: Modeling inelastic deformation: Viscoelasticity, plasticity, fracture. In *ACM SIGGRAPH'88* (Atlanta, 1988), pp. 269–278. 1
- [VG98] VAN GELDER A.: Approximate simulation of elastic membranes by triangulated spring meshes. *Journal of Graphics Tools* 3, 2 (1998), 21–42. 1
- [VG08] VAN GELDER A.: Relaxed jordan canonical form for computer animation and visualization. (submitted for publication), July 2008. 5
- [WPG*97] WESTIN C.-F., PELED S., GUDBJARTSSON H., KIKINIS R., JOLESZ F. A.: Geometrical diffusion measures for MRI from tensor basis analysis. In *ISMRM '97* (Vancouver Canada, April 1997), p. 1742. 6
- [ZZ95] ZIELIŃSKI P., ZIĘTAK K.: The polar decomposition — properties, applications and algorithms. *Applied Mathematics, Annals of the Polish Mathematical Society* 38 (1995), 23–49. 5

Technical Appendix for “Decomposition and Visualization of Fourth-Order Elastic-Plastic Tensors”

These appendix sections provide additional technical details that could not be included in the proceedings paper, due to space limitations. In the interest of self-containment, we first review some tensor basics.

Tensor Notation and Operations

Tensors are linear operators that can be represented as multi-dimensional arrays of coefficients. For 3-D solids, a fourth-order tensor is a $3 \times 3 \times 3 \times 3$ array, a second-order tensor is a 3×3 array, etc. The order of a tensor is the same as the number of subscripts needed to write a typical element. Thus, if \mathbf{E} is a fourth-order tensor, a typical element is denoted by E_{ijkl} . Scalars, vectors, and matrices represent tensors of orders zero, one, and two, respectively.

Operations using tensors are usually denoted using the *Einstein convention* that repeated indices in different tensors are implicitly summed; e.g., matrix multiplication is denoted as $C_{ij} = A_{ik} B_{kj}$, rather than the explicit equation,

$$C_{ij} = \sum_k A_{ik} B_{kj}. \quad (7)$$

This operation is called a *single contraction* in tensor terminology, and is often denoted by “ \cdot ” as an infix symbol.

In elasticity (and many other physical processes) the *double contraction* operator is important. It is denoted by “ $\cdot\cdot$ ” as an infix symbol, and involves summing over two indices, e.g.,

$$C_{ij} = \mathbf{A} : \mathbf{B} = A_{ijkl} B_{kl} = \sum_k \sum_\ell A_{ijkl} B_{k\ell} \quad (8)$$

Double contraction can also be applied to two fourth-order tensors, yielding a new fourth-order tensor.

Where many operations on first and second order tensors use single summation, their generalizations to second and fourth order tensors use double summation. The above example of double contraction is thus the generalization of multiplying a matrix by a vector. Two important cases are the scalar *inner product*, $A_{ij} B_{ij}$, and the *dyad* or *outer product*, $A_{ij} B_{kl}$ that results in a fourth-order tensor.

Appendix A: Unrolling Plasticity Tensors to Matrices and Vectors

Operations on fourth-order 3-D tensors with minor symmetries are more conveniently computed and analyzed by a transformation to 6×6 matrices. Recall that the minor symmetries are $E_{ijkl} = E_{jikl} = E_{ijlk} = E_{jilk}$. Symmetric second-order tensors are transformed into 6-vectors. As explained below, under this transformation, the usual linear-algebra vector and matrix operations correspond to the tensor operations involving double contraction (eq. 8); single contraction and dyad formation (“zero” contraction) also cor-

respond to 6-D vector operations. This transformation is informally called *unrolling*.

Numerical methods for eigen-decomposition and polar decomposition only exist for matrices (rather than fourth-order tensors). Therefore it is computationally advantageous to represent second-order tensors as vectors and *linear transformations* on second-order tensors as matrices. (The natural representation of a linear transformation from second-order tensors to second-order tensors is a fourth-order tensor.)

The straightforward representation of a 3×3 tensor would be as a 9-vector with one component for each tensor element. As applied first by Jean Mandel, (“Ondes plastiques dans un milieu indéfini à trois dimensions,” *Journal de Mécanique*, Vol. 1 (1962), pp. 3–30), and later rigorously justified by others, (see main paper for citations), due to the symmetry of the 3×3 tensor space of interest, an orthonormal change of basis can force the last three components of the 9-vector to be zero. This orthonormal change of basis in 9-D simply consists of 45° 2-D degree rotations on the three pairs of vector components that correspond to symmetric pairs of *off-diagonal* tensor elements.

Similarly, the straightforward representation of a *linear transformation* on 3×3 tensors would be a 9×9 matrix, but if the set of transformations is restricted to those that produce symmetric results, after applying the orthonormal change of basis, a 6×6 matrix suffices.

In summary, as long as the *physical quantities* of interest have the structure of symmetric 3×3 tensors, the corresponding vectors can be 6-D instead of 9-D, and linear transformations of such tensors can be represented with 6×6 matrices. Because the tensors used in the models we visualize always enjoy the minor symmetries, the transformation into 6-D suffices, and is described here.

The first part of the unrolling involves a mapping from single indices in the range $1, \dots, 6$ into pairs of indices in the range $1, 2, 3$.

$$\begin{array}{c|cccccc} k & 1 & 2 & 3 & 4 & 5 & 6 \\ \hline \mu(k) & (1,1) & (2,2) & (3,3) & (1,2) & (2,3) & (1,3) \end{array} \quad (9)$$

Other orders of the last three pairs are acceptable, but one order must be used consistently. Note that μ^{-1} is well defined and maps pairs of indices into single indices. Minor symmetries dictate values for tensor elements whose index pairs do not appear in the table.

Let \mathbf{I}_3 and $\mathbf{0}_3$ denote the 3×3 identity matrix and zero matrix. With the above notation we define the 6×6 matrix \mathbf{E} that represents the unrolling of the fourth-order tensor E_{ijkl} :

$$\mathbf{E} = \begin{bmatrix} \mathbf{I}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & \sqrt{2}\mathbf{I}_3 \end{bmatrix} \begin{bmatrix} E_{\mu(1),\mu(1)} & \cdots & E_{\mu(1),\mu(6)} \\ \vdots & & \vdots \\ E_{\mu(6),\mu(1)} & \cdots & E_{\mu(6),\mu(6)} \end{bmatrix} \begin{bmatrix} \mathbf{I}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & \sqrt{2}\mathbf{I}_3 \end{bmatrix} \quad (10)$$

Similarly, the 6-D column vector \mathbf{s} that represents the un-

rolling of the second-order tensor S_{ij} is given by:

$$\mathbf{s} = \begin{bmatrix} \mathbf{I}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & \sqrt{2}\mathbf{I}_3 \end{bmatrix} \begin{bmatrix} S_{\mu(1)} \\ \vdots \\ S_{\mu(6)} \end{bmatrix} \quad (11)$$

To recover tensors from matrices and vectors, simply use the inverse of the scaling matrix and the inverse of μ .

Appendix B: Polar Decomposition of Plasticity Matrices

Once we have a tensor in matrix form we can perform polar decomposition. The method described here is designed to be robust in the presence of a very small or zero eigenvalue. Experience has shown that another published method performs poorly in these situations. The method we use follows a recent paper (A. Van Gelder, "Relaxed Jordan Canonical Form for Computer Animation and Visualization," submitted for publication July 2008, available from the author), and is described here for self-containment. For our application, we assume that the matrix is square, the determinant is nonnegative, and there is at most one eigenvalue that is zero. In particular, we are not aware of any interpretation of the polar decomposition in this application when $\det \mathbf{E} < 0$, and do not perform polar decomposition in this case.

The polar decomposition on the $n \times n$ square matrix \mathbf{E} is defined as

$$\mathbf{E} = \mathbf{Q}\mathbf{S} \quad (12)$$

where \mathbf{Q} is an orthogonal matrix and \mathbf{S} is a symmetric positive semidefinite matrix. (Conventionally, the term "orthogonal" in this context includes the requirement that rows and columns be unit-length, besides being pairwise orthogonal.) If the determinant of \mathbf{E} is nonnegative, then the determinant of \mathbf{Q} is $+1$. The main paper cites previous methods in the literature, which are either more complicated or more restricted than the method we adopt, described below.

It is well known that the decomposition is unique for $\det \mathbf{E} > 0$. If \mathbf{E} has one eigenvalue of 0, the decomposition is still unique with the specification that $\det \mathbf{Q} = +1$. (Proofs of this and other claims in the appendix are available from the authors in manuscript.) In all cases the \mathbf{S} part of the decomposition is unique.

The steps are summarized in the following equations, where \mathbf{Q} and \mathbf{S} are unknown until they appear on the left side of an equation, and " \equiv " introduces a definition of an unknown. A single subscript on a matrix denotes a *column* of that matrix.

$$\mathbf{M} = \mathbf{E}^T \mathbf{E} = \mathbf{S}^T \mathbf{S} = \mathbf{S}^2$$

$$\mathbf{M} = \mathbf{T}\mathbf{J}\mathbf{T}^T \quad \text{where } \mathbf{J} \text{ is diagonal, ascending order}$$

$$\mathbf{S} = \mathbf{T}\sqrt{\mathbf{J}}\mathbf{T}^T \quad \text{where } \sqrt{\mathbf{J}} \text{ is nonnegative}$$

$$\mathbf{C} \equiv \mathbf{Q}\mathbf{T}$$

$$\mathbf{B} = \mathbf{E}\mathbf{T} = \mathbf{Q}\mathbf{S}\mathbf{T} = \mathbf{C}\sqrt{\mathbf{J}}$$

$$\mathbf{C}_j = \mathbf{B}_j / \sqrt{\mathbf{J}_{jj}} \quad \text{for } j = 2, \dots, n$$

$$\mathbf{C}_1 = \text{Gram-Schmidt completion of 6-D orthonormal basis}$$

$$\mathbf{Q} = \mathbf{C}\mathbf{T}^T$$

$$\mathbf{S} = \text{sym}(\mathbf{Q}^{-1}\mathbf{E})$$

Higham recommends the same computation as the last line, except using \mathbf{Q}^T in place of \mathbf{Q}^{-1} , and in theory they are equal. We obtain slightly more accuracy with \mathbf{Q}^{-1} . In the Gram-Schmidt completion on the next to last line, replace the column \mathbf{C}_1 with $-\mathbf{C}_1$ if $\det \mathbf{C}$ and $\det \mathbf{T}$ have opposite signs. The Jacobi method is very robust and accurate for the computation of eigenvalues and eigenvectors on the second line.

The correctness of the procedure is shown in the cited paper and follows from well known linear algebra properties of real symmetric matrices; in particular, \mathbf{M} is positive semidefinite and \mathbf{T} can be chosen to be orthogonal, so that $\mathbf{T}^T = \mathbf{T}^{-1}$.

As applied in this paper, $n = 6$ and the 6×6 matrix being decomposed is usually the plastic stiffness matrix, which results from unrolling the plastic stiffness tensor (see Appendix A). If there is an eigenvalue of zero for \mathbf{S} , its eigenvector is found in column one of \mathbf{T} ; in this case, that column is also an eigenvector for the zero eigenvalue of the stiffness matrix and is of special interest. In addition, the matrices \mathbf{T} and $\sqrt{\mathbf{J}}$, which are by-products of the decomposition procedure, are useful for various simulations.

Appendix C: Isotropic Stiffness Matrix

Many materials exhibit *isotropic* elasticity properties. For such materials the stiffness tensor can be expressed in terms of two parameters, K , the *bulk modulus*, and G , the *shear modulus*. The stiffness matrix (unrolled stiffness tensor, see Appendix A) for isotropic materials is given by

$$\mathbf{E} = \begin{bmatrix} K + \frac{4}{3}G & K - \frac{2}{3}G & K - \frac{2}{3}G & 0 & 0 & 0 \\ K - \frac{2}{3}G & K + \frac{4}{3}G & K - \frac{2}{3}G & 0 & 0 & 0 \\ K - \frac{2}{3}G & K - \frac{2}{3}G & K + \frac{4}{3}G & 0 & 0 & 0 \\ 0 & 0 & 0 & 2G & 0 & 0 \\ 0 & 0 & 0 & 0 & 2G & 0 \\ 0 & 0 & 0 & 0 & 0 & 2G \end{bmatrix} \quad (13)$$

The eigen-decomposition for isotropic stiffness plays an important role in the visualization. One eigenvalue is $3K$ and the other five are $2G$. The eigenvector for $3K$ is $[1, 1, 1, 0, 0, 0]^T$. All 6-vectors orthogonal to this vector are eigenvectors for $2G$; they span a 5-D subspace.