Optimal Mesh Signal Transforms

Hao Zhang

Hendrik C. Blok

*School of Computing Science
Simon Fraser University, Burnaby, BC Canada
Email: haoz, hcblok@cs.sfu.ca*

Abstract

We describe a simple stochastic model for the geometry of 3D meshes with a given connectivity based on linear prediction plus a Gaussian error term. We show that the resulting probabilistic distribution is a multivariate Gaussian, which may be singular. Furthermore, the eigenvectors of its covariance matrix coincide with the eigenvectors of the linear prediction operator, provided that the operator is symmetric positive semi-definite. This implies the optimality of a mesh signal transform, with respect to a specific class of mesh distributions and in the sense of basis restriction errors. We consider the problem of spectral compression of 3D meshes using the optimal transforms. The generality of our stochastic model allows us to analyze and improve their performance, e.g., on the handling of boundaries.

1. Introduction

A large body of work in discrete signal and image processing [14, 23] can be analyzed and carried out conveniently in the spectral domain through transformations such as the wavelet or discrete Fourier transform (DFT). With polygonal meshes quickly becoming the primary model for 3-D objects in computer graphics and visualization, spectral decomposition techniques for mesh geometry, especially *eigenvalue decomposition*, has received a great deal of attention [15, 21, 26, 27]. In this setting, the mesh geometry is represented by a 3D signal, i.e., the Cartesian (x, y, z) coordinates, defined over the vertices of the underlying graph. The mesh signal transform is given by a projection of the signal onto the eigenvectors of a suitably defined discrete, combinatorial Laplacian operator [27, 15].

Note that for meshes with regular or semi-regular connectivity, the classical DFT may still be applicable, as shown by Karni and Gotsman [16]. But a generalization has to be made for irregular meshes, in which the valence of a mesh vertex (number of outgoing edges) can be ar-

bitrary. This is where eigenvalue decomposition becomes necessary. The most frequently used decomposition adopts as its basis vectors the eigenvectors of the discrete uniform Laplacian operator [27], also known as the umbrella operator [17] or the (normalized) *Tutte Laplacian* [11].

Taubin [27] points out that the eigenvectors of the Tutte Laplacian represent the natural vibration modes of the mesh, while the corresponding eigenvalues capture its natural frequencies, much like the case for well-known image transforms such as the DFT. Geometrically, the intuition is that high-frequency eigenvectors of the Tutte Laplacian are used to capture the small-scale details over the mesh surface, while the low-frequency ones dictate its large-scale variations and overall shape. In general, these eigenvectors possess no analytical form and there are no fast methods, analogous to the Fast Fourier Transform, to compute the corresponding mesh signal transform.

1.1. Combinatorial operators for spectral analysis

In addition to the Tutte Laplacian, the *Kirchhoff operator* and the *normalized graph Laplacian*, both as its variants, have also been used for eigenvalue decomposition. Note that these operators are all *combinatorial*, as they depend on the connectivity of a mesh only. The Kirchhoff operator and the graph Laplacian have traditionally been used in spectral graph theory [7] and their graph-theoretic properties have been studied intensively [20]. Recently, the Tutte Laplacian and the Kirchhoff operator have appeared in numerous applications for geometry processing, including compression [15, 16, 26], smoothing [27, 31], parameterization [11], and watermarking [21, 22] of mesh geometry, as well as shape similarity [19, 30] and data mining [3].

Despite of these developments, we have found little work on analyzing the properties of these combinatorial Laplacian operators and the mesh signal transforms they induce, especially in the context of geometry processing, where the emphasis, such as on speed and the quality of smoothing or compression, are quite different from those in graph the-

ory. We have addressed some of these issues in our recent work [29], where two new symmetric operators for spectral mesh processing and analysis have been proposed and some experimental results demonstrating their advantages over existing operators are given.

In this paper, we take a more theoretical approach and wish to formulate a connection between mesh signal transforms and the well-known *Karhunen-Loeve* or *KL transform* [14]. The KL transform, also known as the method of principle components [6], is of fundamental importance in digital signal and image processing. It is optimal in many ways, e.g., in the context of spectral compression, as we explain in Section 2.3. We are interested in establishing the same notion for mesh geometry processing, which we call *optimal mesh signal transforms*. To achieve this, we need to come up with a stochastic model for mesh geometry.

1.2. “Optimality” of the Kirchhoff operator

Recent work by Ben-Chen and Gotsman [4] appears to be the first attempt at proving the optimality of spectral mesh compression in 3D. Based on a stochastic model of the barycentric coordinate matrix [11], they derive a probabilistic distribution for the geometry of meshes *with fixed boundaries* and having a given connectivity. It is shown to be a multivariate Gaussian distribution [2], but this requires the valences of the mesh vertices to be *sufficiently large*. That is, this multivariate distribution is only approximately Gaussian in general. A simple 2D example is given to illustrate that a valence of six seems to be large enough for the Gaussian approximation to be reasonable.

Ben-Chen and Gotsman argue for the optimality of spectral mesh compression using the Kirchhoff operator K . They prove that K is *essentially* the inverse of the covariance matrix of the mesh geometry distribution and its optimality follows from the well-known result of principle component analysis [14]. The approximate nature of the statement has something to do with the fact that K is *singular*. The proof offered is somewhat lengthy and it resorts to the use of conditional distributions [5] and the notion of order statistics [12] and uniform spacings [24].

1.3. Our results and contribution

In this paper, we are able to make precise statements in regards to the optimality of certain mesh signal transforms based a simple stochastic model for the geometry of 3D meshes (with or without fixed boundaries) having a common connectivity. This model is defined by the familiar notion of linear prediction whose combined error and noise term is assumed to realize a multivariate Gaussian distribution with zero mean [2]. It is straightforward to prove that the resulting mesh distribution is also a multivariate

Gaussian, where the only nontrivial part is the inclusion of the singular case. This is necessary since all the combinatorial Laplacian operators considered in this paper, e.g., the Kirchhoff and Tutte Laplacian operators, are singular.

Furthermore, we show that the eigenvectors of the covariance matrix of the mesh distribution coincide with the eigenvectors of the prediction operator, if the operator is *symmetric positive semi-definite*. This implies the optimality of a mesh signal transform, with respect to a specific *class* of mesh distributions and in the sense of *basis restriction errors*, which we define in Section 2.3. We consider the problem of spectral compression of 3D meshes and compare the optimal mesh signal transforms. The generality of our stochastic model also allows us to analyze and improve their performance, e.g., on the handling of boundaries.

1.4. Notations

We denote matrices by capital letters, with A, T, D, U, C, K , and I reserved for special operators. For example, I_n is used for the $n \times n$ identity matrix. The trace, rank, and transpose of a matrix P are denoted by $tr(P)$, $rank(P)$, and P' , respectively, and the 2-norm of a vector by $\|\cdot\|$. Random vectors and variables are represented using calligraphic letters such as $\mathcal{X}, \mathcal{Y}, \mathcal{R}, \dots$, and ordinary vectors are denoted by $\mathbf{x}, \mathbf{y}, \mathbf{r}, \dots$, possibly with subscripts. The expectation, variance, and covariance matrix of a random variable or vector $\mathcal{X}$ is denoted by $\mathbb{E} \mathcal{X}$, $\sigma^2(\mathcal{X})$ and $\Sigma(\mathcal{X})$, respectively. Note that the matrix may be omitted when it is obvious from the context. If $\mathcal{X}$ has a Gaussian or normal distribution with mean μ and variance σ^2 , then we write $\mathcal{X} \sim N(\mu, \sigma^2)$. Other necessary notations will be defined as they enter our discussion.

2. Spectral compression and the KL transform

In this paper, we assume that all the meshes considered are connected. Meshes composed of disjoint parts can be handled in a component-wise manner. Consider a mesh M with n vertices whose vertex and edge sets are denoted by $V(M)$ and $E(M)$, respectively. Let us represent M by a pair $(A, \mathbf{x})$, where A is the $n \times n$ adjacency matrix of the underlying *mesh graph* $G = (V(M), E(M))$, specifying the connectivity among the mesh vertices. Namely, $A_{ij} = 1$ if $(i, j) \in E(M)$ and $A_{ij} = 0$ otherwise.

The geometric information about the mesh is specified by the real $n \times 3$ matrix $\mathbf{x}$, called the *coordinate vector*. Thus the mesh M can be viewed as a 3D signal defined over the vertices of the mesh graph. This is similar to the vector-space representation of discrete images [23], except that no particular order of the mesh vertices is assumed. For an image, either a column-wise or a row-wise raster scan order is used to generate a vector from the image.

2.1. ED-transforms and spectral compression

Consider a mesh $M = (A, \mathbf{x})$ with n vertices and an operator $P \in \mathbf{R}^{n \times n}$ whose eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ form a basis. Typically, P takes the form of a combinatorial Laplacian defined by connectivity and its eigenvalues capture the frequencies [15, 26, 27]. Unless otherwise specified, assume that the eigenvectors are ordered according to increasing eigenvalues. Let $E = [\mathbf{e}_1 | \mathbf{e}_2 | \dots | \mathbf{e}_n]$. The *eigenvalue decomposition* of $\mathbf{x}$ with respect to P is given by

$$\mathbf{x} = E\hat{\mathbf{x}} = \sum_{j=1}^n \mathbf{e}_j \hat{x}_j,$$

where $\hat{\mathbf{x}} \in \mathbf{R}^{n \times 3}$ is said to be the *ED-transform* of $\mathbf{x}$ with respect to P . The inverse ED-transform is $\hat{\mathbf{x}} = E^{-1}\mathbf{x}$. The projection of $\mathbf{x}$ onto the subspace spanned by the first m ($m \leq n$) eigenvectors of P is denoted by

$$\mathbf{x}_{(P,m)} = \sum_{j=1}^m \mathbf{e}_j \hat{x}_j.$$

This provides a spectral compression of the mesh $\mathbf{x}$ using the leading m low-frequency spectral coefficients. For any given probabilistic distribution of the coordinate vector $\mathbf{x}$, the expected value of the L_2 error

$$\mathfrak{E} \sqrt{\text{tr}[(\mathbf{x} - \mathbf{x}_{(P,m)})(\mathbf{x} - \mathbf{x}_{(P,m)})']}$$

is said to be a *basis restriction error* [14]. From now on, let us only deal with the x -component of the mesh, still denoted by $\mathbf{x}$; the same arguments can be made about y and z . We can then eliminate the $\text{tr}()$ from the definition of the basis restriction error, yielding $\mathfrak{E} \|\mathbf{x} - \mathbf{x}_{(P,m)}\|$.

2.2. Kirchhoff and Tutte Laplacian operators

Given a mesh $M = (A, \mathbf{x})$ with n vertices and associated mesh graph $G = (V(M), E(M))$, denote the valence of vertex $i \in V(G)$ by d_i . Let D be the diagonal matrix with d_i 's along the diagonal. The Kirchhoff operator $K \in \mathbf{R}^{n \times n}$ of M is given by $K = D - A$. That is

$$K_{ij} = \begin{cases} d_i & \text{if } i = j \\ -1 & \text{if } i \neq j \text{ and } (i, j) \in E(M) \\ 0 & \text{otherwise} \end{cases}$$

Following [11], let us call the matrix $T = D^{-1}K$ the (normalized) Tutte Laplacian. We have

$$T_{ij} = \begin{cases} 1 & \text{if } i = j \\ -1/d_i & \text{if } i \neq j \text{ and } (i, j) \in E(M) \\ 0 & \text{otherwise} \end{cases}$$

Figure 1. Spectral compression (100 out of 379 spectral coefficients) of a 4-8 mesh.

It follows that $T = I_n - D^{-1}A = I_n - C$, where C is referred to as the *centroid matrix* of M . By applying C to $\mathbf{x}$, each vertex is transformed into the centroid of its immediate, i.e., first-order or *one-ring*, neighbors; this is the well-known Laplacian smoothing [10]. In general, T and C are not symmetric and their eigenvectors cannot be made orthonormal [29]. We also consider the operator

$$U = T'T,$$

proposed in [29], for eigenvalue decomposition. Clearly, U is a symmetric second-order (since the nonzero entries of U extends to the second-order neighbors of a vertex) operator. For convenience, let us refer to it as the *symmetric second-order Tutte Laplacian*, or *SSTL*, for short.

Our experiments show that the Kirchhoff operator often generates undesirable artifacts around vertices with low valences, e.g., 3 or 4, when applied to spectral compression. This is especially problematic for the class of so-called *semi-regular 4-8 meshes* [28], as shown in Figure 1. The SSTL, on the other hand, always seems to produce better visual quality at the same compression rate.

2.3. The KL transform

For a given distribution of random signals specified by its mean μ and covariance matrix Σ , it is well-known [14]

that the optimal transform, one which minimizes the basis restriction error for this distribution, is given by the eigenvalue decomposition with respect to the *autocorrelation matrix* $\Phi = \Sigma + \mu\mu'$. That is, for any P and m ,

$$\mathfrak{E} \|\mathbf{x} - \mathbf{x}_{(\Phi, m)}\| \leq \mathfrak{E} \|\mathbf{x} - \mathbf{x}_{(P, m)}\|.$$

The ED-transform of $\mathbf{x}$ with respect to Φ is referred to as the KL transform. If the random signals have zero mean, then $\Phi = \Sigma$. In general, the transforms of Σ and Φ need not be the same. To prove the optimality of a mesh signal transform with respect to P , for a given distribution, it suffices to show that the eigenvectors of P coincide with the eigenvectors of the autocorrelation matrix of that distribution.

3. A stochastic model for mesh geometry

A stochastic model describes a digital signal, e.g., an image, as a member of an ensemble. This permits development of algorithms that are useful for an entire class of signals rather than a single one. The ensemble is referred to as a *discrete random field*, and it is often sufficient, e.g., for our study of optimal basis for spectral mesh compression, to characterize it by its mean and covariance matrix [14]. In this section, we define a stochastic model for all meshes with n vertices having a common connectivity.

3.1. The observation model and linear prediction

Let us denote the random mesh by $\mathcal{X}$ and start with the classical observation model [14, 18]

$$\mathcal{X} = \bar{\mathbf{x}} + \mathcal{R}_m,$$

where $\mathcal{X}$ is the observed mesh signal subject to some additive measurement noise $\mathcal{R}_m$ and $\bar{\mathbf{x}}$ represents the true shape of the mesh. We can provide a prediction $\tilde{\mathcal{X}}$ for the true shape based on some observation $\mathcal{Y}$; this could simply be $\mathcal{X}$. For linear predictions, we have

$$\tilde{\mathcal{X}} = P\mathcal{Y} \quad \text{with} \quad \tilde{\mathcal{X}} = \bar{\mathbf{x}} - \mathcal{R}_p,$$

where $P \in \mathbf{R}^{n \times k}$ is a linear prediction operator, $\mathcal{Y} \in \mathbf{R}^k$ is used for the prediction, and $\mathcal{R}_p$ is the prediction error term. Thus we have

$$\mathcal{X} = P\mathcal{Y} + (\mathcal{R}_m + \mathcal{R}_p) = P\mathcal{Y} + \mathcal{R}. \quad (1)$$

This is the familiar *Wiener filtering equation* [14], for which the goal is to find P that minimizes the error $\mathcal{R}_p$. Alexa [1] has used this model for Wiener filtering of meshes, where prediction is based on the observed signal $\mathcal{X}$.

In the context of signal or image restoration, the prediction operator P is regarded as a *blurring* operator [23]. Thus

the image acquisition process is assumed to be subject to blur and additive noise, given by $\mathcal{R}$. One often requires *a priori* statistical knowledge of the noise/error term. A natural and frequently used model is that of a Gaussian white noise or a more general multivariate Gaussian distribution.

3.2. Multivariate Gaussian distributions

Let $\mathcal{X} = [\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_n]'$ be a discrete random field, from which a sample is denoted by $\mathbf{x} = [x_1, x_2, \dots, x_n]'$. Let Σ be an $n \times n$ real *symmetric positive definite* matrix and let the vector $\mu = [\mu_1, \mu_2, \dots, \mu_n]^T$, where each μ_i is a real constant. It can be shown that the nonnegative function

$$f(\mathbf{x}) = f(x_1, x_2, \dots, x_n) = \frac{e^{-\frac{(\mathbf{x}-\mu)^T \Sigma (\mathbf{x}-\mu)}{2}}}{\sqrt{(2\pi)^n \det(\Sigma^{-1})}} \quad (2)$$

is a joint probability density function of the n random variables $\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_n$. We say that $\mathcal{X}$ has a *nonsingular multivariate Gaussian distribution*, and $\mathcal{X} \sim N(\mu, \Sigma^{-1})$, where μ and Σ^{-1} are the mean and covariance matrix of the distribution, respectively [12]. A Gaussian white noise is a discrete random field in which each component $\mathcal{X}_i$ is a zero-mean Gaussian and any two different components are mutually uncorrelated. Thus it is a multivariate Gaussian $N(\mathbf{0}, \Sigma)$, where Σ is a nonnegative diagonal matrix.

One of the primary ways of defining multivariate Gaussian distributions is through linear transformations [2]. Specifically, if $\mathcal{Y} \sim N(\mu, \Sigma)$ and $\mathcal{X} = P\mathcal{Y}$ with P a nonsingular matrix, then $\mathcal{X} \sim N(P\mu, P\Sigma P')$. The nonsingularity of the distribution is related to the requirement that Σ be nonsingular, as its determinant appears in the denominator in (2). However, it is possible to generalize this and include the singular case, where the density would not exist and the distribution is characterized by other means [12].

In this general sense, a random vector $\mathcal{X}$ of n components with $\mathfrak{E} \mathcal{X} = \mu$ and $\mathfrak{E} (\mathcal{X} - \mu)(\mathcal{X} - \mu)' = \Sigma$ is said to have a multivariate Gaussian distribution $N(\mu, \Sigma)$ if there exists a linear transformation $\mathcal{X} = P\mathcal{Y} + \mathcal{Z}$, where $P \in \mathbf{R}^{n \times k}$ and k is the rank of Σ , and $\mathcal{Y}$ has a nonsingular multivariate Gaussian distribution [2]. The distribution for $\mathcal{X}$ is singular if $k < n$. Furthermore, if $\mathcal{Y} \sim N(\mu, \Sigma)$, then $\mathcal{X} = P\mathcal{Y}$ is distributed according to $N(P\mu, P\Sigma P')$, where P is allowed to be a non-square matrix without a full rank.

4. Optimal mesh signal transforms

Let us first consider the case for meshes without fixed boundaries. Consider the stochastic model (1) where linear prediction is based on $\mathcal{X}$ itself, as in Wiener filtering of meshes [1]. We have

$$(I_n - P)\mathcal{X} = Q\mathcal{X} = \mathcal{R}, \quad \text{where } \mathcal{X}, \mathcal{R} \in \mathbf{R}^n$$

Let $q = \text{rank}(Q)$ be the rank of Q , $q \leq n$. If $q = n$, then Q is nonsingular. Let us assume that the error term $\mathcal{R}$ is a Gaussian white noise and that all the components of $\mathcal{R}$ have the same variance σ^2 . Thus $\mathcal{R} \sim N(\mathbf{0}, \sigma^2 I_n)$. It is not hard to see that $\mathcal{X} \sim N(\mathbf{0}, \sigma^2 (Q'Q)^{-1})$. So the ED-transform with respect to the operator $Q'Q$ is optimal.

4.1. The singular distribution case

The more difficult case is when $\text{rank}(Q) = q < n$, in which case Q is singular. Let us suppose that Q is symmetric positive semi-definite however. We first observe that then the error term $\mathcal{R}$ *cannot* be a white noise, since the subspace for $\mathcal{R}$ is only q -dimensional. In fact, this subspace is the *image space* of Q , $\text{Im } Q = \{Q\mathbf{x} \mid \mathbf{x} \in \mathbf{R}^n\}$.

Inevitably, the mapping from the space of random mesh signals $\mathcal{X}$ to the error space would be many-to-one. To make this clear, let us first fix some notations. Denote by E the $n \times n$ matrix whose columns are the orthonormal eigenvectors of Q , i.e., $E' = E^{-1}$. Then

$$QE = E\Lambda, \text{ where } \Lambda = \begin{bmatrix} 0 & 0 \\ 0 & \Lambda_1 \end{bmatrix},$$

and $\Lambda_1 \in \mathbf{R}^{q \times q}$ is a diagonal matrix with q positive eigenvalues of Q along the diagonal. We shall partition $E, \hat{\mathcal{X}}$, and $\hat{\mathcal{R}}$ as follows.

$$E = [E_0 \ E_1], \ \hat{\mathcal{X}} = \begin{bmatrix} \hat{\mathcal{X}}_0 \\ \hat{\mathcal{X}}_1 \end{bmatrix}, \ \hat{\mathcal{R}} = \begin{bmatrix} \hat{\mathcal{R}}_0 \\ \hat{\mathcal{R}}_1 \end{bmatrix},$$

where $E_1 \in \mathbf{R}^{n \times q}$ and $\hat{\mathcal{X}}_1, \hat{\mathcal{R}}_1 \in \mathbf{R}^q$. Now recall our stochastic model $Q\mathcal{X} = \mathcal{R}$. It follows that

$$E\Lambda E' \mathcal{X} = \mathcal{R} \Rightarrow \Lambda \hat{\mathcal{X}} = \Lambda E' \mathcal{X} = \hat{\mathcal{R}} = E' \mathcal{R},$$

where $\hat{\mathcal{X}} = E' \mathcal{X}$ is the ED-transform of $\mathcal{X}$ with respect to Q . We must have $\hat{\mathcal{R}}_0 = 0$ and $\hat{\mathcal{R}}_1 = \Lambda_1 \hat{\mathcal{X}}_1$. Observe that no matter how we choose $\hat{\mathcal{X}}_0$, as long as $\hat{\mathcal{X}}_1 = \Lambda_1^{-1} \hat{\mathcal{R}}_1$, we would have $Q\mathcal{X} = E\Lambda \hat{\mathcal{X}} = EE' \mathcal{R} = \mathcal{R}$. Thus to have a unique mesh for a given error $\mathcal{R}$, we must impose a constraint on $\hat{\mathcal{X}}_0$. In our model, we let $\hat{\mathcal{X}}_0 = 0$. Therefore,

$$\mathcal{X} = E\hat{\mathcal{X}} = E_1 \hat{\mathcal{X}}_1.$$

A **key assumption** we make for our stochastic model is that the *spectral-domain error term* $\hat{\mathcal{R}}_1$ is a Gaussian white noise.¹ That is,

$$\hat{\mathcal{R}}_1 \sim N(\mathbf{0}, \Omega), \quad (3)$$

where Ω is a diagonal matrix with a positive diagonal. Since $\hat{\mathcal{X}}_1 = \Lambda_1^{-1} \hat{\mathcal{R}}_1$, we have

$$\hat{\mathcal{X}}_1 \sim N(\mathbf{0}, \Lambda_1^{-1} \Omega (\Lambda_1^{-1})') \equiv N(\mathbf{0}, \Omega \Lambda_1^{-2}),$$

¹Note that in the nonsingular case, the spatial-domain error $\mathcal{R}$ is a Gaussian white noise with constant variance σ^2 if and only if the spectral-domain error $\hat{\mathcal{R}}$ is a Gaussian white noise with the same variance, assuming Q is symmetric. As we will see, this is not so in the singular case.

and furthermore, since $\mathcal{X} = E_1 \hat{\mathcal{X}}_1$,

$$\mathcal{X} \sim N(\mathbf{0}, E_1 \Omega \Lambda_1^{-2} E_1').$$

As $\hat{\mathcal{X}}_1$ has a nonsingular multivariate Gaussian distribution and $E_1 \in \mathbf{R}^{n \times q}$ with $q < n$, the random mesh $\mathcal{X}$ has a *singular* multivariate Gaussian distribution by definition. The mean and covariance matrix of $\mathcal{X}$ are $\mathbf{0}$ and $\Sigma = E_1 \Omega \Lambda_1^{-2} E_1'$, respectively. We can easily verify that the eigenvectors of Σ coincide with the eigenvectors of the linear prediction operator Q , since

$$\begin{aligned} \Sigma E &= E_1 \Omega \Lambda_1^{-2} E_1' E = E_1 \Omega \Lambda_1^{-2} \begin{bmatrix} 0 & I_q \end{bmatrix} \\ &= E_1 \begin{bmatrix} 0 & \Omega \Lambda_1^{-2} \end{bmatrix} = E \begin{bmatrix} 0 & 0 \\ 0 & \Omega \Lambda_1^{-2} \end{bmatrix}. \end{aligned}$$

Therefore, the optimal transform for our mesh distribution for $\mathcal{X}$ is given by the eigenvalue decomposition with respect to Q . Note that the eigenvalues of Q are in *reverse* order against the eigenvalues of Σ . It is also worth noting that results we have derived so far would not hold if the prediction operator Q does not have an orthonormal set of eigenvectors, as is the case for the Tutte Laplacian.

Now let us consider the spatial-domain error term $\mathcal{R}$. Recall that $\mathcal{R} = E\hat{\mathcal{R}} = E_1 \hat{\mathcal{R}}_1$. Since $\hat{\mathcal{R}}_1$ is a nonsingular multivariate Gaussian, $\mathcal{R}$ is a *singular* Gaussian distributed according to $N(\mathbf{0}, E_1 \Omega E_1')$, by definition. Note that the rank of the covariance matrix of $\mathcal{R}$ and $\mathcal{X}$ are both q . We summarize our results in the following theorem.

Theorem: Given a stochastic model $Q\mathcal{X} = \mathcal{R}$. If Q is symmetric positive definite and $\mathcal{R} \sim (\mathbf{0}, E_1 \Omega E_1')$ with some positive diagonal matrix $\Omega \in \mathbf{R}^{q \times q}$, $q = \text{rank}(Q)$, and E_1 is the matrix whose columns are the eigenvectors of Q corresponding to its nonzero eigenvalues (let $QE_1 = E_1 \Lambda_1$), then $\mathcal{X}$ is a multivariate Gaussian $N(\mathbf{0}, E_1 \Omega \Lambda_1^{-2} E_1')$ and the optimal transform for the distribution of $\mathcal{X}$ is given by the ED-transform with respect to Q . $\square$

In fact, we are able to make a stronger statement about the optimality of the ED-transforms given by Q . To see this, let $0 < \lambda_1 \leq \dots \leq \lambda_q$ be the values along the diagonal of Λ_1^{-2} . Observe that if we change the values of the λ_i 's but still maintain the order $\lambda_1 \leq \dots \leq \lambda_q$, then the Gaussian distribution for $\mathcal{X}$ changes — it has a different covariance matrix now — but the eigenvectors of the covariance matrix does not change and neither does the optimality of the ED-transform given by Q . Therefore, this ED-transform is really optimal for a *class* of mesh distributions. For example, some members of this class include the distributions of $\mathcal{X}$ for which $f(Q)\mathcal{X} = \mathcal{R}$, where f is a *monotone* polynomial. That is, $f(a) > f(b)$ whenever $a > b$.

Optimality of the Kirchhoff operator and the SSTL

Recall that the Kirchhoff operator $K = DT$ and the symmetric second-order Tutte Laplacian, or SSTL, $U =$

$T'T$, where T is the Tutte Laplacian operator. It is not hard to see that for a connected mesh, the *null space* of T , $\text{Nul } T = \{\mathbf{x} \mid T\mathbf{x} = \mathbf{0}\}$, has rank 1. In fact, $\text{Nul } T = \{a\mathbf{1} \mid a \in \mathbf{R} \text{ and } \mathbf{1} = [1 \dots 1]'\}$, which is the set of all meshes that degenerate to a point. Also, it can be shown that $\text{Nul } K = \text{Nul } U = \text{Nul } T$. Thus $\text{rank}(K) = \text{rank}(U) = \text{rank}(T) = n - 1$ and these operators are all singular.

As we have shown, to ensure the uniqueness of a mesh for a given error (drawn from the image space of the prediction operator), we need to impose the constraint that $\hat{\mathcal{X}}_0 = 0$. In the case for K and U , it is equivalent to requiring that the centroid of all the vertices, i.e., the DC value corresponding to eigenvalue 0, be invariantly at the origin. Since both K and U are symmetric positive definite [29], the results from the previous section do apply. That is,

- K is optimal for the singular multivariate Gaussian distribution $\mathcal{X}$ with $K\mathcal{X} = \mathcal{R}$, where the prediction operator $P = I_n - K$.
- U is optimal for the distribution $\mathcal{X}$ with $U\mathcal{X} = \mathcal{R}$, where the prediction operator $P = I_n - U$.

However the distributions induced by K and U are quite different. To get a sense of that, consider the simplest case where the covariance matrix of the spectral-domain error $\hat{\mathcal{R}}_1$ is $\Omega = I_q$. Thus the covariance matrix of the error $\mathcal{R}$ is given by $\Sigma_r = E_1 E_1' = I_n - E_0 E_0'$. For K (or U), E_0 is simply the eigenvector of K (or U) corresponding to the (unique) zero eigenvalue. The normalized version of this eigenvector is simply $E_0 = [1/\sqrt{n} \dots 1/\sqrt{n}]'$. Thus,

$$\Sigma_r = \begin{bmatrix} (n-1)/n & -1/n & \dots & -1/n \\ -1/n & (n-1)/n & \dots & -1/n \\ \dots & \dots & \dots & \dots \\ -1/n & \dots & -1/n & (n-1)/n \end{bmatrix}.$$

Consider the distribution $\mathcal{X}$ given by $K\mathcal{X} = DT\mathcal{X} = \mathcal{R}$. We have $T\mathcal{X} = D^{-1}\mathcal{R}$ and the covariance matrix of $T\mathcal{X}$ is $D^{-1}\Sigma_r D^{-1}$. Thus at a vertex i with valence d_i , the variance of the i -th component of $T\mathcal{X}$, which can be seen as a rough, discrete measure of curvature at vertex i [27], is $(n-1)/(nd_i^2)$. So this variance is sensitive to the valence and the smaller d_i is, the larger the variance.

Random mesh generation and spectral compression

An implication of the above observation is that the distribution induced by K tends to generate meshes that have “spikes” at vertices with low-valences, as shown in Figure 2(a)-2(c), where some of the first a few random meshes generated from K ’s distribution are shown. It turns out that the eigenvectors of K , when visualized as a scalar field defined over the vertices of the mesh graph, also tend to have “spikes” at vertices with a low-valence, e.g., 3 or 4 [29].

Figure 2. Some random meshes generated with the same semi-regular 4-8 connectivity of the sphere in Figure 1. (a)-(c): Using K ’s distribution. (d)-(f): Using U ’s distribution.

The way we generate these random meshes is as follows. We use the `RANDN()` function from Matlab to generate the spectral-domain error term $\hat{\mathbf{r}}_1 \in \mathbf{R}^{n \times 3}$ according to $N(0, I)$. Then we generate the mesh $\mathbf{x} = E_1 \Lambda_1^{-2} \hat{\mathbf{r}}_1$, as explained before. We see that the diagonal values of Λ_1^{-2} , which are the square reciprocals of the eigenvalues of K (or U), have the effect of attenuating the high-frequency content of error term. In general, the characteristics of the random meshes generated may be examined from the eigenvalues and eigenvectors of the prediction operator and the nature of the prediction it provides.

With the same error statistics, the distribution induced by the SSTL operator U appears to favor much smoother meshes without any effects from the vertex valences, as shown in Figure 2(d)-2(f). As far as linear prediction is concerned, the operator for U ’s distribution

$$P = I_n - U = C' + C - C'C = C + C'T \quad (4)$$

is suitable for smooth meshes. It predicts the location of a vertex i as the centroid of its neighbors plus a vector displacement. This displacement is a weighted average of the discrete Tutte Laplacian measures, given by T , at the immediate neighbors of i , where the weight used at neighbor j is $1/d_j$, d_j being the valence.

Now consider applying K and U to spectral compression of a mesh M . If M has a sufficient number of “spikes” around low-valence vertices, e.g., boundary vertices tend to fall into this category, then it is more likely to be generated from K ’s distribution. So we would expect K to perform better in terms of approximation quality (in the mean square error sense) for compression, as it gives the (statistically)

optimal transform for the distribution. On the other hand, if M is highly smooth, then it is more likely to be generated by U 's distribution and then U should provide better compression quality. In Section 5, we show some experimental results to support these arguments.

To summarize, suppose that we wish to find an operator Q which performs well in spectral compression for a particular class of meshes. We would then want this operator to be a good predictor for this class of meshes, so that the error term would tend to be small and close to its statistical mean $\mathbf{0}$. In this case, this class of meshes have a higher probability of being generated from the mesh distribution defined by Q . Due to Q 's optimality for this distribution, we would expect spectral compression of meshes from the class to have smaller basis restriction errors. We will use this approach to improve the performance of the SSTL in compressing meshes with boundaries in Section 5.

Comparison with Ben-Chen and Gotsman's model

A main difference between the stochastic model of Ben-Chen and Gotsman [4] and ours is that they focus on generating only *valid* meshes, i.e., meshes with no inverted triangles. This is hard to achieve in 3D, especially when mesh boundaries are free to move. In fact, Ben-Chen and Gotsman do not guarantee mesh validity in their stochastic model for 3D meshes with fixed boundaries.

Our model may generate invalid geometries, as shown in Figure 2(a), although all valid geometries can also be generated according to our distribution. Moreover, our linear prediction operator, which plays a similar role as the barycentric coordinate matrix in their model, is fixed and not chosen stochastically. It is unclear whether and to what extent does this make our model more restrictive.

Ben-Chen and Gotsman do not handle the singular case. We provide treatment of this, as well as meshes with fixed boundaries or any set of anchor points [26], in the next section. We again utilize the linear prediction model and arrive at a different distribution and a different optimality result. Finally, our stochastic model and proofs are given in a general setting and our optimality statements do not depend on any assumptions on vertex valences and they are applicable to a class of mesh distributions.

4.2. Meshes with fixed boundaries or anchors

Let us use $\mathcal{X}$ again to represent the random mesh vector and without loss of generality, assume that all interior vertices precede all boundary vertices in the vertex ordering. Let $\mathcal{X}_I$ and $\mathbf{x}_B$ denote the subvectors of interior and boundary vertices, respectively. Suppose that the mesh boundary $\mathbf{x}_B$ is fixed. We use the linear prediction model (1) again, but this time for the interior vertices $\mathcal{X}_I$ only,

$$\mathcal{X}_I = P\mathcal{X} + \mathcal{R} = P[\mathcal{X}'_I \mid \mathbf{x}'_B]' + \mathcal{R}, \quad (5)$$

where we note that the prediction is based on the interior, as well as the fixed boundary vertices. Note also that $\mathbf{x}_B$ does not have to be just the mesh boundary; it may be any set of “anchor” points which are fixed in space.

A natural choice for P is the centroid matrix C restricted to the interior vertices. That is, if the centroid matrix is

$$C = \begin{bmatrix} C_I & C_B \\ C_0 & C_1 \end{bmatrix}, \quad \text{then the predictor } P = \begin{bmatrix} C_I & C_B \end{bmatrix}.$$

Let $T_I = I - C_I$. Then we have

$$T_I \mathcal{X}_I = C_B \mathbf{x}_B + \mathcal{R}.$$

If there are at least one boundary vertex and the mesh is connected, it can be shown that T_I is a irreducibly diagonally dominant matrix [25] and that it is invertible. Therefore,

$$\mathcal{X}_I = T_I^{-1} C_B \mathbf{x}_B + T_I^{-1} \mathcal{R}.$$

We shall assume that the error term $\mathcal{R}$ is a Gaussian white noise distributed according to $N(\mathbf{0}, \sigma^2 I)$. It follows that $\mathcal{X}_I$ has a nonsingular multivariate Gaussian distribution specified by $N(\mu, \Sigma) \equiv N(T_I^{-1} C_B \mathbf{x}_B, \sigma^2 (T_I' T_I)^{-1})$. The optimal transform for this distribution is given by the eigenvalue decomposition with respect to the autocorrelation matrix $\Sigma + \mu\mu'$. Unlike K and U , this matrix depends on the spatial information about the mesh, i.e., $\mathbf{x}_B$.

Figure 3. (a) A mesh patch. (b) The mean mesh interpolating the fixed mesh boundary.

Note that the *mean mesh* $\mu = T_I^{-1} C_B \mathbf{x}_B$ has zero total energy at the interior vertices, where the energy is defined by the partial Tutte Laplacian $[T_I \mid -C_B]$, which is the full Tutte Laplacian restricted to the interior vertices only. Thus our stochastic model favors smooth and regular meshes interpolating the given boundary, as shown in Figure 3.

Finally, if we choose another linear predictor, e.g., $P = I - U$, in place of the centroid matrix, we would arrive at a different optimal mesh signal transform for a different distribution, but the derivations will be the same.

5. Experimental results

In practice, we have found that in terms of L_2 approximation errors, the Kirchhoff operator K and the SSTL U have very similar performances for spectral compression. However, meshes compressed using K consistently exhibit little “spikes” around low-valence vertices. Results shown in Figure 1 exemplify this situation. The SSTL operator, on the other hand, always seems to produce compressed meshes with comparable or better visual quality (this is especially true for relatively smooth meshes with small or no boundaries) and without any artifacts caused by vertices with unusual valences.

Visual quality of compression may be measured roughly using the *visual error* proposed by Karni and Gotsman [15] and later studied by Sorkine et al. [26]. A slightly modified version of the visual error is used in our experiments. We compute the sum of: (1) squared differences between the *scale-dependent Laplacian*² measures (as vectors) at corresponding vertices and (2) squared vertex displacements projected along the direction of such a vector (as an estimate of the normal) at the vertex. In Figure 4(c)-4(f), we show some sample compression results for a coarse version of the bunny and a cube. The plots of the mean square approximation errors and visual errors are given in Figure 5.

For mesh patches with boundaries, the Kirchhoff operator tends to achieve a smaller approximation error mostly because it is able to preserve the boundaries better, as shown in Figure 6(b). The SSTL U , on the other hand, smooths the mesh boundary more aggressively, as shown in Figure 6(c). However, we can improve the performance of U in these situations by modifying the centroid matrix C in the predictor (4) so as to do a better job in preserving the boundary vertices. This can be achieved by assigning a positive weight, instead of zero as in the centroid matrix, to a boundary vertex and reducing the weights at its neighbors, keeping the sum of weights at 1. In Figure 6(d), we see slight improvement at the boundary achieved by assigning a weight of 1/2 to the boundary vertices in the prediction. The corresponding mean square approximation errors, which are shown in Figure 7, should be more telling.

Note that precise preservation of the mesh boundaries is possible if we take the improvement described above to the limit. That is, we assign a weight of 1 to a boundary vertex and zero to its neighbors in the prediction. The corresponding results are similar to those obtained using the optimal transform described in Section 4.2, except for slightly larger errors around interior vertices close to the boundary. A comparison by the Metro tool [8] reveals that the approx-

Figure 4. (a) Original bunny with 557 vertices. (b) Compressed with fixed boundaries using the transform of Section 4.2 ($m = 80$ spectral coefficients). (c) Compressed with K ($m = 120$). (d) Compressed with U ($m = 120$). (e) Compressed with K ($m = 120$). (f) Compressed with U ($m = 120$).

imation errors are slightly larger for the former approach. Note that to apply the optimal boundary-preserving transform, we need to set the error variance σ^2 to obtain a reasonable distribution. The choice of this value, which depends on the fixed boundary $\mathbf{x}_B$, may be optimized during encoding and it is then transmitted for decompression.

If one were to apply this optimal transform to compress a large mesh composed of many patches, so as to fix the boundary shrinkage problem pointed out by Karni and Gotsman [15], the patch boundaries need to be encoded separately, e.g., using delta-encoding. The decoded connectiv-

²This is a variant of the Tutte Laplacian proposed by Desbrun et al. [9]. It depends on the geometry, namely, the edge lengths, of the mesh. Compared with the combinatorial Tutte Laplacian, it provides a more accurate discrete curvature measure at the mesh vertices.

Figure 5. Mean square approximation error and visual error plots for the bunny (top) and the cube (bottom).

Figure 7. Approximation error improvement with a better predictor for boundary vertices.

Figure 6. (a) Original patch with 503 vertices. (b) Compressed with K ($m = 80$). (c) Compressed with original U ($m = 80$). (d) Compressed with modified U ($m = 80$).

ity and the error variances transmitted will be used to construct the operators for spectral decompression. Note that there will be a different operator for each of the x , y , and z component for each patch.

6. Conclusion and future work

The main contribution of this paper is an answer to the question: what would be the equivalents of KL transforms in the spectral analysis of meshes with irregular connectivity? We show that if one assumes a Gaussian white noise model for the error *spectrum* of the linear prediction model (1), then any symmetric positive definite operator, e.g., the SSTL U and the Kirchhoff operator K , is optimal, in the sense of minimization of basis restriction errors, for the probabilistic distribution induced by this operator.

Although a closer look at the class of mesh distributions for which a prediction operator is optimal is still needed, of practical importance is our study of ED-transforms defined by K , U , and their variants. We believe the SSTL operator is the preferable choice for spectral compression of meshes. We have also proposed two methods to perform boundary-preserving compressions, which are useful to remedy the boundary shrinkage problem for spectral compression of meshes composed of many patches [15]. Admittedly, the major obstacle for using spectral compression is still its computational cost.

Our stochastic model based on linear prediction is quite natural and its generality allows us to plug in various prediction operators to study their corresponding probabilistic distributions. It would be interesting to experiment with predictors defined with other weights, e.g., geometry-driven Laplacian operators [9], and examine the characteristic of the random meshes generated. In some sense, these random meshes reveal the “shape” of the operator, analogous to the notion of *connectivity shapes* studied by Isenburg et al. [13].

References

- [1] M. Alexa, "Wiener Filtering of Meshes," *Proceedings of Shape Modelling International*, pp. 51-59, 2002.
- [2] T. W. Anderson, *Introduction to Multivariate Statistical Analysis*, John Wiley and Sons, 1984.
- [3] M. Belkin and P. Niyogi, *Semi-Supervised Learning on Manifolds*, Technical Report TR-2002-12, Dept. of Computer Science, University of Chicago, 2002.
- [4] M. Ben-Chen and C. Gotsman, "On the Optimality of Spectral Compression of Meshes," *preprint*, 2003.
- [5] J. Besag and C. Kooperberg, "On Conditional and Intrinsic Autoregressions," *Biometrika*, Vol. 84, No. 4, pp. 733-746, 1995.
- [6] M. Bilodeau and D. Brenner, *Theory of Multivariate Statistics*, Springer, 1999.
- [7] F. R. K. Chung, *Spectral Graph Theory*, CBMS, American Math Society, 1997.
- [8] P. Cignoni, C. Rochini, and R. Scopigno, "Metro: Measuring Error on Simplified Surfaces," *Computer Graphics Forum*, Vol. 17, No. 2, pp. 167-174, 1998.
- [9] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr, "Implicit Fairing of Irregular Meshes using Diffusion and Curvature Flow," *SIGGRAPH 1999*, pp. 317-324.
- [10] D. A. Field, "Laplacian Smoothing and Delaunay Triangulations," *Communications in Applied Numerical Methods*, Vol. 4, pp. 709-712, 1988.
- [11] C. Gotsman, X. Gu, and A. Sheffer, "Fundamentals of Spherical Parameterization for 3D Meshes," *ACM Transaction on Graphics*, Vol. 22, No. 3, pp. 358-363, 2003.
- [12] R. V. Hogg and A. T. Craig, *Introduction to Mathematical Statistics*, Fourth Edition, Macmillan Publishing, 1978.
- [13] M. Isenburg, S. Gumhold, and C. Gotsman, "Connectivity Shapes," *IEEE Visualization 2001*.
- [14] A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice Hall, 1989.
- [15] Z. Karni and C. Gotsman, "Spectral Coding of Mesh Geometry," *SIGGRAPH 2000*, pp. 279-286.
- [16] Z. Karni and C. Gotsman, "3D Mesh Compression Using Fixed Spectral Basis," *Proceeding of Graphics Interface*, pp. 1-8, 2001.
- [17] L. Kobbelt, S. Campagna, J. Vorsatz, and H. P. Seidel, "Interactive Multi-Resolution Modeling on Arbitrary Meshes," *SIGGRAPH 1998*, pp. 105-115.
- [18] S. Mallat, *A Wavelet Tour of Signal Processing*, Second Edition, Academic Press, 1999.
- [19] D. McWherter, M. Peabody, W. C. Regli, and A. Shokoufandeh, "Transformation Invariant Shape Similarity Comparison of Solid Models," *Proceeding of ASME Design Engineering Technical Conference*, Pittsburgh, 2002.
- [20] B. Mohar, "Some Applications of Laplacian Eigenvalues of Graphs," *Graph Symmetry: Algebraic Methods Applications*, G. Hahn and G. Sabidussi (Ed.), pp. 225 - 275, 1997.
- [21] R. Ohbuchi, S. Takahashi, T. Miyazawa, and A. Mukaiyama, "Watermarking 3-D Polygonal Meshes in the Spectral Domain," *Proceedings of Graphics Interface*, pp. 9-18, 2001.
- [22] R. Ohbuchi, H. Ueda, and S. Endoh, "Watermarking 2D Vector Maps in the Mesh-Spectral Domain," *Proceedings of Shape Modelling International*, pp. 216-225, 2003.
- [23] W. K. Pratt, *Digital Image Processing*, Third Edition, Wiley-Interscience, 2001.
- [24] R. Pyke, "Spacings," *Journal of the Royal Statistical Society*, pp. 395-439, 1965.
- [25] Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Company, 1996.
- [26] O. Sorkine, D. Cohen-Or, and S. Toledo, "High-Pass Quantization for Mesh Encoding," *Eurographics Symposium on Geometry Processing*, pp. 42-51, 2003.
- [27] G. Taubin, "A Signal Processing Approach to Fair Surface Design," *SIGGRAPH 1995*, pp. 351-358.
- [28] L. Velho, "Using Semi-Regular 4-8 Meshes for Subdivision Surfaces," *Journal of Graphics Tools*, Vol. 5, No. 3, pp. 35-47, 2001.
- [29] H. Zhang and E. Fiume, "Laplacian Operators for Spectral Mesh Processing," *preprint*, March 2003.
- [30] H. Zhang and E. Fiume, "Shape Matching of 3D Contours using Normalized Fourier Descriptors," *Proceedings of Shape Modelling International*, pp. 261-268, 2002.
- [31] H. Zhang and E. Fiume, "Butterworth Filtering and Implicit Fairing of Irregular Meshes," *Pacific Graphics 2003*, to appear.