Lazy Wavelet Simplification using Scale-dependent Dense Geometric Variability Descriptors ⋆

Teodor Cioaca * **Bogdan Dumitrescu** * **Mihai-Sorin Stupariu** **

* *University Politehnica of Bucharest (e-mail: t.cioaca@gmail.com, bogdan.dumitrescu@acse.pub.ro)*
** *University of Bucharest (e-mail: stupariu@fmi.unibuc.ro)*

Abstract: Partitioning geometric data into two sets, one corresponding to high frequencies and the other to low frequencies, is a critical operation in the second generation wavelet multiresolution analysis. From a geometric point of view, a region with high variability within a vertex neighborhood at a certain scale indicates a correlation with a signal having a frequency that dominates at that scale. We thus prospect the abilities of several geometric variability descriptors to robustly identify features. We consider three descriptor families: based on principal component analysis, surface fitting and quadric error metrics. To assess the quality of each descriptor, we employ a *lazy wavelet* simplification of digitized 3D models since these usually contain noisy geometric structures from which multiple scales of resolutions can be inferred. The difference between a simplified model and the highest resolution representation is measured objectively using averaged local distance functions.

Keywords: Computer Graphics, Graphs, Differential Geometric Methods, Discriminators, Successive Approximations

1. INTRODUCTION

Modern object scanners produce very high density digitized models, but the raw data itself does not immediately lend itself to a qualitative interpretation, relevant to an end-user. Thus, the data density richness can also become a disadvantage when direct statistical or semantic processing is applied. In particular, hierarchical representations of the data at different resolution levels can greatly facilitate the identification of features at different scales. One of the most widely accepted techniques for constructing such representations is through the use of *wavelet multiresolution analysis* (WMRA), a term recognized by many authors as a synonym for this goal. With the advent of *second generation wavelet transform*, where the *lifting scheme* introduced by Sweldens (1996) plays a pivotal role, a comprehensive gamut of applications for geometry processing were proposed, with active research still being conducted.

The connection between WMRA and scale-dependent feature analysis is easier to establish by examining a series of relevant works on this topic. Even recently, Nader et al. (2014) proposed an adaptive, multi-scale point cloud editing algorithm that acts on the scale specific features of the geometric model similarly to the wavelet-based frameworks of Guskov et al. (1999) and Zorin et al. (1997). The key difference between feature-based model analysis and second generation wavelet analysis methods is the focus of the first category on the data variability instead on predict-update filter design. Nevertheless, this distinction is not conspicuous in the sense that the first category could be included into a broader class of a *lazy wavelet* decomposition, a term popularized by Sweldens (1996), where a geometric process is used to predict samples at different scales. Ranging from the influential work of Hoppe (1996) on progressive mesh representation to a linear predictive coding of 3D mesh sequences as described by Stefanoski et al. (2007), the implications of feature-driven geometry simplification, compression and filtering also motivate our present work.

WMRA approaches based on the lifting scheme exploit data redundancy and can operate directly on the spatial domain. Lazy wavelet decomposition, as an integral part of the lifting scheme, can be guided by a feature discrimination mechanism, as discussed by Cioaca et al. (2016). In this respect, geometric variability is a relevant redundancy indicator and can serve the purpose of partitioning the data samples. By convention, the samples deemed as disposable are assigned to a set called the *odd* subset, while the remaining samples are placed in a set called the *even* subset. The lifting mechanism implies that the odd samples be estimated from the even ones. Data compression then becomes possible by working on the set of *wavelet coefficients*, which are the differences between the odd samples and their predicted values.

In this work, we examine the performance of several geometric variability descriptors in conjunction with a lazy wavelet partitioning strategy. We achieve this by computing a *dense* scalar descriptor function (i.e. defined at each individual sample point) and then sorting the

⋆ This work was supported by the Swiss Enlargement Contribution in the framework of the Romanian-Swiss Research Programme, project WindLand, project code: IZERZO_142168/1 and 22 RO-CH/RSRP.

samples according to their descriptor values. For this goal, the following properties are desirable:

- scale dependence: the descriptor must reflect the variability of the data at a certain scale of interest,
- robustness to noise: the descriptor must be able to discriminate salient features even in the presence of noise.

1.1 Related work

The term "descriptor" has broad connotations depending on whether it represents single points or entire regions, or whether it takes on vector or scalar values. Some of the most common applications of descriptors involve shape matching for rigid or deformed objects. For the purpose of this work, where an importance value has to be assigned to all point samples, the discussion will be limited to scalar valued descriptors. One prominent choice involves computing robust principal curvatures estimators at multiple scales. To achieve this, Yang and Shen (2012) employed principal component analysis of ball and sphere vertex neighborhoods. Ho and Gibbins (2009) used the Casorati curvature to extract multi-scale features from surfaces and unstructured point clouds. The multi-scale shape index analysis developed by Bonde et al. (2013) generalized another curvature-based descriptor. Boyer et al. (2002) and Cipriano et al. (2009) resorted to local surface fitting for robust curvature estimation and saliency discrimination, while Sipiran and Bustos (2011) extended the Harris corner detector to 3D meshes by computing the Harris response in a vertex neighborhood of variable radius again by performing paraboloid fitting. Saliency and scale were further combined by Liu et al. (2007) in an algorithm for detecting critical points by averaging scalar vertex descriptors. A more complex automatic and robust scale selection mechanism based on curvature evolution according to the heat equation was proposed by Fadaifard and Wolberg (2011). For large point sets, Shtrom et al. (2013) have constructed a saliency map based on point feature histogram descriptors. Park et al. (2012) employed tensor voting and eigenanalysis to detect features and scales automatically for noisy point clouds. Maximo et al. (2011), Wu et al. (2013) and Tao et al. (2015) introduced a new descriptor class based on the Zernike coefficients of a patch sample, a method suited only for measuring geometric variability and akin to a spherical harmonics decomposition. Although robust to noise and scale independent, this type of descriptor requires the computation of 25 such coefficients on a regular, 16×16 grid around each vertex via heightmap fitting. Tangelder and Veltkamp (2008) and Tang and Godil (2012) performed a survey analysis of several existing shape descriptors in the context of 3D shape retrieval, while Creusot et al. (2013) described a machine learning approach aggregating several types of scalar descriptors for the purpose of identifying landmark features on face scans.

1.2 Main contributions

The contribution of this work can be summed up as follows:

- We compared the feature preserving capabilities of 3 families of scale-dependent descriptors.
- We extended the normal field variation of Hussain (2010) to variable size vertex neighborhoods.
- We proposed a principal component construction of a new quadric error metric that is more robust to noise, scale dependent and extensible to multiple dimensions.
- We improved the selectivity of the Harris 3D Sipiran and Bustos (2011) corner detector.

The rest of this article is structured as follows. In section 2, the notions of scale dependent neighborhoods and dense scalar descriptors are introduced, highlighting their role in driving a lazy wavelet partitioning algorithm. Two PCA-based descriptors are then presented in section 3. Three descriptors based on paraboloid fitting are examined in section 4. Two new descriptors based on the quadric error metric matrices are presented in section 5. The performance of all seven descriptors is discussed in section 6. Conclusions and future work directions are given in section 7.

2. SCALE FEATURES AND LAZY WAVELET SIMPLIFICATION

2.1 Scale dependent neighborhoods

As input, we consider a mesh, denoted by $M = (V, E)$, consisting of the initial set of points, $V \subset \mathbb{R}^3$, and the set of triangle edges, E. Working with a mesh representation has two interesting advantages: it facilitates local neighborhood analysis and allows operating local topological changes efficiently. Besides being the focus of local queries and processing, neighborhoods are also a means of defining the concept of scale, as exemplified in the works of Yang and Shen (2012), Rocca et al. (2011) and Ho and Gibbins (2009). When referring to vertices and their neighbors, one common concept is the so-called *n-ring* neighborhood (see Botsch et al. (2010)) of a vertex $\mathbf{v} \in V$, denoted by $\mathcal{N}_v^n(\mathbf{v})$. By definition, a vertex $\mathbf{v}_i \in V$ is an element of the $\mathcal{N}_v^n(\mathbf{v})$ set if the shortest topological path connecting it to vertex $\mathbf{v}$ consists of at most n edges. If one desires to include the interiors of the mesh triangles in the definition of the n-ring, the distinct notation $\mathcal{N}_f^n(\mathbf{v})$ is used to emphasize the inclusion of faces. Since most meshes are irregular in terms of sampling density and vertex degree (i.e. the number of edges incident at that vertex), an n-ring neighborhood alone may not accurately reflect a geometric scale. A more suitable concept is that of a local surface patch of radius r (see Mitra et al. (2006)), centered at vertex $\mathbf{v}$, denoted by $\mathcal{P}_v(\mathbf{v}, r)$. Another vertex, $\mathbf{v}_i$, is an element of this set if the distance between $\mathbf{v}$ and $\mathbf{v}_i$ is less than or equal to r. A common choice for measuring distances across surfaces is the geodesic distance, but, for small local neighborhoods, the Euclidean distance is also a suitable approximation to the otherwise more computationally complex geodesic distance. By combining the n-ring neighborhood and the local surface patch of a vertex, we can further define the notion of a *local window*,

$$\mathcal{W}_v^n(\mathbf{v}, r) = \mathcal{N}_v^n(\mathbf{v}) \cap \mathcal{P}_v(\mathbf{v}, r), \tag{1}$$

comprising all vertices within an r radius around $\mathbf{v}$, but no further than n rings from it (see figure 1).

The *local window* set defined in formula (1) can be endowed with a mechanism capable of measuring the ge-

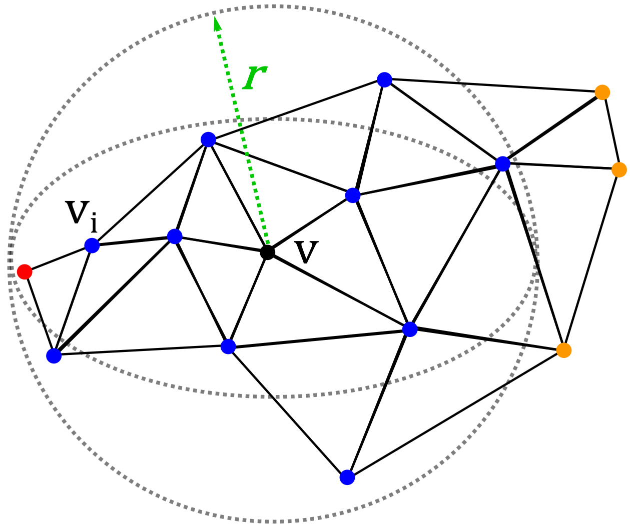

Fig. 1. A window neighborhood $\mathcal{W}_v^n(\mathbf{v}, r)$ with $n = 2$. The blue vertices are inside the window, while the orange vertices are outside the sphere of center $\mathbf{v}$ and radius r. The red vertex, although inside this sphere, is more than 2 topological rings away from $\mathbf{v}$.

ometric variability of the vertex patch contained within this window. Formally, this mechanism is represented by a dense descriptor function (see Bronstein et al. (2011)), $D(\mathbf{v}, n, r) : V \times \mathbb{N}^* \times (0, \infty) \to [0, \infty)$. At a certain scale identified through the number of rings, n, and the maximum radius, r, the set pair $(V, D(V, n, r))$ can be used to discriminate between and order the vertices of the mesh. In particular, by sorting the vertices in descending order according to the values of the descriptor function, the more salient features relevant to the (n, r) scale parameters can be identified.

2.2 Feature-driven lazy wavelet simplification

The scale and feature descriptor concepts defined earlier serve the purpose of constructing a multiresolution mesh representation. The resulting representation is a hierarchical chain of increasingly coarser approximations of the input mesh, $M = M^L \supset M^{L-1} \supset \ldots \supset M^0$. Any two consecutive elements in this chain are meshes whose sets of vertices are related through the following expression:

$$V^{i+1} = V_E^i \cup V_O^i, \tag{2}$$

where $V_E^i = V^i$ is called the *even* subset and $V_O^i = V^{i+1} \setminus V^i$ is called the *odd* subset. This terminology is adopted from wavelet analysis theory, where the *even-odd* partitioning of the samples of a signal is also referred to as *lazy wavelet decomposition*. The reason behind this categorization of samples is to exploit the local redundancy of the signal and express one subset of samples from the other. One immediate application is information *compression* by compressing the differences between the estimated samples and their local approximations derived from the samples of the coarser set. Usually, the *odd* samples are estimated from their *even* neighbors at each level i in the hierarchical decomposition. Besides compression, other applications of wavelet analysis include *multiresolution editing* and *filtering*. Multiresolution editing is a process where a coarser representation is edited and local details are added automatically from the precomputed difference vectors. Filtering implies working with the sets of difference vectors at each level i, treating them as frequency band equivalents. These vectors can be scaled, offering the possibility of either enhancing or suppressing certain scale-dependent features.

A feature weighing mechanism, which is offered by the $D(\mathbf{v}, n, r)$ function, can be used to guide the lazy wavelet partitioning of samples since it offers a means to measure redundancy. Ideally, at each scale, the odd samples should be selected in such a way that removing them altogether from the denser set would incur the lowest approximation error if the fine scale model were replaced with the coarser, even subset. By exploiting the mesh connectivity, a reliable means of constructing the even-odd partition is described in algorithm 1.

Algorithm 1 Descriptor-driven *lazy wavelet* partitioning

Require: $M = (V, E)$, $n \in \mathbb{N}^*$, $r > 0$
Ensure: $V = V_E \cup V_O$
  $V_E \leftarrow \varnothing$, $V_O \leftarrow \varnothing$
  $V' \leftarrow \mathtt{SORT}(V, \mathtt{key}(\mathbf{v}) = D(\mathbf{v}, n, r))$
  for all $\mathbf{v} \in V'$ **do**
    if $\mathbf{v} \notin V_E$ **then**
      $V_O \leftarrow V_O \cup \{\mathbf{v}\}$
      $V_E \leftarrow V_E \cup \mathcal{N}_v^1(\mathbf{v})$
    end if
  end for

Intuitively, the values of the descriptor function are used to greedily select those vertices with the lowest geometric saliency first and tag them as odd, while automatically including their one-ring neighbors in the even subset. In theory, other partitioning schemes are possible where an odd sample can be connected to other odd samples, but such strategies lead to poor estimations of the odd values from the even ones. Thus, assuming a semi-regular mesh connectivity, where each vertex is connected, on average, to six others, the odd-even cardinality percentage with respect to the total number of samples is, on average, 25%-75% (instead of the 50%-50% decimation ratio that classic, 1D wavelet analysis commonly is known for). Also, to reflect the change of scale after each odd subset is removed, the search radius r can be increased by a user selected factor (usually between 1.5 and 2). The number of rings does not need to be increased as it only acts as a barrier for the number of samples in regions with nonuniform density. The effect that both radius and ring number parameters have is also important when dealing with data corrupted by noise. In this sense, increasing these parameters strengthens the robustness to noise of the $D(\mathbf{v}, n, r)$ descriptor function. The trade-off is thus between distinguishing high frequency details and coherently estimating the geometric variability of the patch even in the presence of noise.

3. PRINCIPAL COMPONENT ANALYSIS AND SURFACE VARIATION

3.1 Surface variability index

Computing the principal components of the covariance matrix inside a $\mathcal{W}_v^n(\mathbf{v}, r)$ window constitutes a more robust alternative to evaluating discrete curvatures as a measure for surface variability. Pauly et al. (2002) proposed computing $D(\mathbf{v}, n, r)$ using a surface variation cost equivalent to the separability index of the singular value decomposition of the data matrix of the interest window. In practice, it suffices to compute the eigenvalues of the covariance matrix,

$$\mathbf{C}(\mathcal{W}_v^n(\mathbf{v}, r)) = \begin{bmatrix} (\mathbf{v}_{i_1} - \bar{\mathbf{v}})^\top \\ \cdots \\ (\mathbf{v}_{i_N} - \bar{\mathbf{v}})^\top \end{bmatrix}^\top \cdot \begin{bmatrix} (\mathbf{v}_{i_1} - \bar{\mathbf{v}})^\top \\ \cdots \\ (\mathbf{v}_{i_N} - \bar{\mathbf{v}})^\top \end{bmatrix}, \quad (3)$$

where $N = |\mathcal{W}_v^n(\mathbf{v}, r)|$, and $\bar{\mathbf{v}} = \frac{1}{N} \sum_{\mathbf{v}_{i_k} \in \mathcal{W}_v^n(\mathbf{v}, r)} \mathbf{v}_{i_k}$ is the centroid of the window. Then the eigenvalues can be indexed according to their increasing order $\lambda_0 \leq \lambda_1 \leq \lambda_2$. The associated separability index is expressed as

$$D(\mathbf{v}, n, r) = \frac{\lambda_0}{\lambda_0 + \lambda_1 + \lambda_2}. \quad (4)$$

The geometric reasoning behind this feature descriptor is based on the interpretation of the eigenvectors of $\mathbf{C}(\mathcal{W}_v^n(\mathbf{v}, r))$. The eigenvectors corresponding to λ_1 and λ_2 span a plane that offers a good approximation for the tangent plane to the mesh support surface at $\mathbf{v}$. Similarly, the eigenvector corresponding to the direction of least variation approximates the surface normal. In the presence of noise, these approximations are more robust than typical tangent plane and normal estimators that are based only on the one-ring elements around $\mathbf{v}$. Thus, the separability index defined in formula (4) attains large values for highly deformed, far from planar regions. It is important to mention that windows having $N \geq 3$ samples are considered, otherwise the descriptor evaluates to zero, corresponding to a rank deficient covariance matrix.

3.2 Normal field variation via normal voting

Hussain (2010) suggested measuring the one-ring variation of the normal field in order to iteratively simplify a mesh. Previously, Koibuchi and Yamada (2000) have used a similar measure, referred to as the *bending energy*, to study the first order phase transition of a mesh membrane. For the one-ring neighborhood of a vertex $\mathbf{v}$, this descriptor can be evaluated as:

$$NFV(\mathbf{v}) = \sum_{\mathbf{v}_i \in \mathcal{N}_v^1(\mathbf{v})} (1 - \hat{\mathbf{n}}^\intercal \hat{\mathbf{n}}_i), \quad (5)$$

where $\hat{\mathbf{n}}$ and $\hat{\mathbf{n}}_i$ are the normalized normals at vertices $\mathbf{v}$ and $\mathbf{v}_i$, respectively.

We propose extending this descriptor to a window patch by computing the Gaussian weighted sum of normal variations over the window:

$$D(\mathbf{v}, n, r) = \frac{1}{S} \sum_{\mathbf{v}_i \in \mathcal{W}_v^n(\mathbf{v}, r)} w(\mathbf{v}_i)(1 - \hat{\mathbf{n}}^\intercal \hat{\mathbf{n}}_i), \quad (6)$$

where $S = \sum_{\mathbf{v}_i \in \mathcal{W}_v^n(\mathbf{v}, r)} w(\mathbf{v}_i)$ is the sum of the Gaussian weights $w(\mathbf{v}_i) = \exp\left(-\frac{\|\mathbf{v}_i - \mathbf{v}\|^2}{(r/3)^2}\right)$. The evaluation of vertex-wise normals is sensitive in the presence of noise and corrupts the discriminative abilities of the window descriptor. To counteract this disadvantage, we suggest smoothing out the normal field through the normal voting method of Page et al. (2002). Their algorithm does not modify the positions of the vertices and it requires gathering normal votes from the same window from formula (6). Individual normal votes are collected from the vertices in this window, each vote consisting of an orientation estimation. More explicitly, the vote that a vertex $\mathbf{v}_i$ casts implies transporting the normal vector at $\mathbf{v}_i$ along a circular arc connecting this vertex to $\mathbf{v}$. Since the $(\hat{\mathbf{n}}_i, \mathbf{v}_i)$ pair completely defines a plane, the circular arc then lies on the sphere which is tangent to this plane and passes through $\mathbf{v}$. Referring to figure 2, the expression for the normal vote from $\mathbf{v}_i$ at $\mathbf{v}$ is:

$$\tilde{\mathbf{n}}_i = \hat{\mathbf{n}}_i - 2\frac{\hat{\mathbf{n}}_i \cdot (\mathbf{v}_i - \mathbf{v})}{\|\mathbf{v}_i - \mathbf{v}\|^2}(\mathbf{v}_i - \mathbf{v}). \quad (7)$$

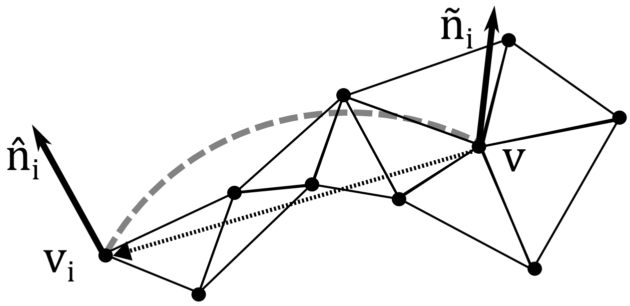

Fig. 2. Normal vote cast from $\mathbf{v}$ to $\mathbf{v}_i$.

To actually recover the robust estimate of the normal vector at $\mathbf{v}$, a covariance matrix of the normal votes must be constructed:

$$\mathbf{C}_{normal}(\mathbf{v}) = \sum_{\mathbf{v}_i \in \mathcal{W}_v^n(\mathbf{v}, r)} \frac{w(\mathbf{v}_i)}{\|\tilde{\mathbf{n}}_i\|^2} \tilde{\mathbf{n}}_i \tilde{\mathbf{n}}_i^\intercal, \quad (8)$$

where the $w(\mathbf{v}_i)$ weights attenuate the contribution of those vertices that are not close to $\mathbf{v}$. The eigenvector corresponding to the largest eigenvalue of the covariance matrix is the sought-after normal estimate. We denote this normalized normal estimate by $\bar{\mathbf{n}}$ and rewrite formula (6) as:

$$D(\mathbf{v}, n, r) = \frac{1}{S} \sum_{\mathbf{v}_i \in \mathcal{W}_v^n(\mathbf{v}, r)} w(\mathbf{v}_i)(1 - \bar{\mathbf{n}}^\intercal \bar{\mathbf{n}}_i). \quad (9)$$

4. PARABOLOID FITTING

4.1 Curvature estimation via paraboloid fitting

The problem of measuring discrete curvatures on meshes or point clouds has been well studied and we refer the interested reader to the work of Brentzen et al. (2012). The majority of these methods rely on discretizations of Differential Geometry concepts. Often, only the one-ring neighborhood of vertices is used in the computations, leading to noise sensitivity. One may argue that noisy data sets could be subjected to a denoising treatment, but this kind of approach is not always desirable since it can smooth out certain features or it may be undesirable in scenarios where the original information is not allowed to be changed. Since computing curvatures on polynomial surfaces of the type $z = f(x, y)$ is trivial, one common solution is to fit paraboloid surfaces in each window. This approach is inherently more robust to noise than the one-ring alternatives, the main trade-off being the requirement to solve a least squares problem. In case the input set contains outliers, one may replace the least squares solver with a nonlinear RANSAC (see Kang et al. (2011)) to diminish their impact. The vertices of $\mathcal{W}_v^n(\mathbf{v}, r)$ can be seen as samples on a Monge patch, each having its coordinates written as $(x, y, f(x, y))$. Thus, in the local window of a vertex $\mathbf{v}$, assumed to be the origin, the height function is approximated by a paraboloid:

$$f(x, y) = ax^2 + bxy + cy^2 + dx + ey. \tag{10}$$

The fitting process inside $\mathcal{W}_v^n(\mathbf{v}, r)$ consists of the following steps:

(1) compute a pseudonormal at $\mathbf{v}$, e.g. $\bar{\mathbf{n}} = \sum_{\mathbf{v}_i \in \mathcal{W}_v^n(\mathbf{v},r)} \hat{\mathbf{n}}_i$.
(2) eliminate normal outliers by replacing the search window with $\mathcal{W}_v^n(\mathbf{v}, r)' = \{\mathbf{v}_i \in \mathcal{W}_v^n(\mathbf{v}, r) : \hat{\mathbf{n}}_i \cdot \bar{\mathbf{n}} \geq 0\}$.
(3) translate the window so that $\mathbf{v}$ corresponds to the origin of a local coordinate system, i.e. $\mathcal{W}_v^n(\mathbf{v}, r)'' = \{\mathbf{v}_i - \mathbf{v} : \mathbf{v}_i \in \mathcal{W}_v^n(\mathbf{v}, r)'\}$.
(4) construct $\mathcal{W}_v^n(\mathbf{v}, r)'''$ by rotating the $\mathcal{W}_v^n(\mathbf{v}, r)''$ set such that $\bar{\mathbf{n}}$ becomes aligned with the $\hat{\mathbf{z}} = (0, 0, 1)^\intercal$ axis. One may argue that this rotation should cancel the first degree terms in formula (10), but such a situation would occur only if the fitted paraboloid had its origin normal perfectly aligned with $\bar{\mathbf{n}}$.
(5) minimize $\sum\limits_{\mathbf{p} \in \mathcal{W}_v^n(\mathbf{v},r)'''} \|f(x_p, y_p) - z_p\|^2$ in the least squares sense, where $\mathbf{p} = (x_p, y_p, z_p)^\intercal$.

Given the above transformations, it suffices to compute the principal, mean and Gaussian curvatures at $f(0, 0)$ since these values correspond to $\mathbf{v}$. Based on the principal curvatures, $\kappa_1(\mathbf{v}) \leq \kappa_2(\mathbf{v})$, two feature descriptors can be formulated:

$$\kappa_C(\mathbf{v}) = \sqrt{\frac{\kappa_1^2(\mathbf{v}) + \kappa_2^2(\mathbf{v})}{2}}, \tag{11}$$

i.e. the Casorati curvature at $\mathbf{v}$, and

$$SI(\mathbf{v}) = \frac{2}{\pi} \left| \arctan \left(\frac{\kappa_2(\mathbf{v}) + \kappa_1(\mathbf{v})}{\kappa_2(\mathbf{v}) - \kappa_1(\mathbf{v})} \right) \right|, \tag{12}$$

usually known as the *shape index*. Both of these descriptors are versatile enough to describe how the local surface is bending. For more information on these descriptors, we refer the interested reader to the work of Koenderink and van Doorn (1992).

4.2 3D Harris response

The Harris corner detector, introduced by Harris and Stephens (1988), is an efficient feature identification mechanism for grayscale images. As its name implies, the features that are easily recognized correspond to regions where the gradient of the image has a high variation in both horizontal and vertical directions. Thus, if $I(x, y)$ represents the image value at pixel coordinates (x, y), and if $W_{(x,y)}$ represents a window centered at this pixel, the corner features correspond to those pixels where the following autocorrelation function attains local maxima:

$$E(\Delta x, \Delta y) = \sum_{(x_i,y_i) \in W_{(x,y)}} w(x_i, y_i) \cdot \cdot \left[I(x_i, y_i) - I(x_i + \Delta x, y_i + \Delta y) \right]^2, \tag{13}$$

where $w(x_i, y_i)$ is the (Gaussian) weight of the (x_i, y_i) pixel and $(\Delta x, \Delta y)$ is a 2D shift vector.

Using a first order Taylor expansion of the image function, the autocorrelation function in expression (13) becomes:

$$E(\Delta x, \Delta y) \approx \mathbf{S} \sum_{(x_i,y_i) \in W_{(x,y)}} w(x_i, y_i) \begin{bmatrix} I_x(x_i, y_i)^2 & I_x(x_i, y_i) I_y(x_i, y_i) \\ I_x(x_i, y_i) I_y(x_i, y_i) & I_y(x_i, y_i)^2 \end{bmatrix} \mathbf{S}^\intercal, \tag{14}$$

where $\mathbf{S} = [\,\Delta x\ \Delta y\,]$. By denoting the weighted matrix sum in formula (14) with $\mathbf{M}(x, y)$, we can rewrite this expression in a condensed form:

$$E(\Delta x, \Delta y) = \mathbf{S}\mathbf{M}(x, y)\mathbf{S}^\intercal. \tag{15}$$

The Harris response for the (x, y) pixel then takes the following expression:

$$h(x, y) = det(\mathbf{M}(x, y)) - kTr(\mathbf{M}(x, y))^2, \tag{16}$$

where k is a threshold parameter empyrically chosen in the $[0.04, 0.07]$ range.

Sipiran and Bustos (2011) have extended the Harris corner detector to triangular meshes by using a paraboloid fitting to substitute the $I(\cdot, \cdot)$ image function. The fitting process is performed just as described in subsection 4.1, but the pseudonormal $\bar{\mathbf{n}}$ is estimated by computing the PCA of the $\mathcal{W}_v^n(\mathbf{v}, r)$ window, i.e. being the eigenvector corresponding to the smallest eigenvalue. The authors justify this normal choice due to its increased insensitivity to noise.

The $\mathbf{M}(x, y)$ matrix is still a 2×2 matrix:

$$\mathbf{M}(x, y) = \begin{pmatrix} A(x, y) & C(x, y) \\ C(x, y) & B(x, y) \end{pmatrix}, \tag{17}$$

where the $A(x, y), B(x, y), C(x, y)$ elements are computed as follows:

$$A(x, y) = \frac{1}{\sqrt{2\pi}r} \int_{\mathbb{R}^2} e^{\left(\frac{-(u^2+v^2)}{2r^2}\right)} f_x(u, v)^2 du dv, \tag{18}$$

$$B(x, y) = \frac{1}{\sqrt{2\pi}r} \int_{\mathbb{R}^2} e^{\left(\frac{-(u^2+v^2)}{2r^2}\right)} f_y(u, v)^2 du dv, \tag{19}$$

$$C(x, y) = \frac{1}{\sqrt{2\pi}r} \cdot \int_{\mathbb{R}^2} e^{\left(\frac{-(u^2+v^2)}{2r^2}\right)} f_x(u, v) f_y(u, v) du dv. \tag{20}$$

Through experimentation, we have found that the formulation of the Harris 3D feature response, as expressed in formula (16), has very poor discriminative properties. Instead we propose the following expression for the Harris response descriptor:

$$D(\mathbf{v}, n, r) = \frac{2 \det(\mathbf{M}(x, y))}{\varepsilon + Tr(\mathbf{M}(x, y))}, \tag{21}$$

where ε is a small design parameter, meant to avoid division by zero in flat regions, set in our experiments to $\varepsilon = 10^{-5}$.

5. QUADRIC ERROR METRICS

5.1 Smoothed quadric error metric

Initially conceived for incremental mesh simplification, the quadric error metric (Garland and Heckbert (1998)) was extended to cope with meshes with vertex-wise attributes such as color, texture coordinates or any other data. To better understand how this mechanism can be adapted to a multiscale scenario, we briefly recall its local definition. When dealing with meshes, one is interested in efficiently computing the distance between a vertex $\mathbf{v}$ and the support plane of a triangular face f with unit normal $\mathbf{n}_f$ which contains a point $\mathbf{p}$. The square of this distance is:

$$\begin{aligned} d(\mathbf{v}, f)^2 &= (\mathbf{n}_f \cdot (\mathbf{v} - \mathbf{p}))^2 = \\ &= \mathbf{v}^\intercal(\mathbf{n}_f\mathbf{n}_f^\intercal)\mathbf{v} - 2(\mathbf{p}^\intercal\mathbf{n}_f)\mathbf{n}_f^\intercal\mathbf{v} + (\mathbf{n}_f^\intercal\mathbf{p})^2. \end{aligned} \quad (22)$$

It can be observed that formula (22) can be expressed in matrix form as:

$$d(\mathbf{v}, f)^2 = \begin{pmatrix} \mathbf{v} \\ 1 \end{pmatrix}^\intercal \mathbf{Q}_f \begin{pmatrix} \mathbf{v} \\ 1 \end{pmatrix}, \quad (23)$$

where $\mathbf{Q}_f$ is the following block matrix:

$$\mathbf{Q}_f = \begin{pmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{b}^\intercal & c \end{pmatrix}, \quad (24)$$

with $\mathbf{A} = \mathbf{n}_f\mathbf{n}_f^\intercal$, $\mathbf{b} = (\mathbf{p}^\intercal\mathbf{n}_f)\mathbf{n}_f$, $c = (\mathbf{n}_f^\intercal\mathbf{p})^2$. A fundamental property of the $\mathbf{Q}_f$ matrices is that the sum of squared distances from $\mathbf{v}$ to two different faces, f_i and f_j, can be obtained by employing formula (23), where $\mathbf{Q}_f = \mathbf{Q}_{f_i} + \mathbf{Q}_{f_j}$. This property is extremely useful because the geometric variation in the one-ring neighborhood of a vertex can be stored in a matrix form as $\mathbf{Q}_\mathbf{v} = \sum_{f \in \mathcal{N}_f^1(\mathbf{v})} \mathbf{Q}_f$ at each vertex $\mathbf{v}$. Moreover, since the basic incremental simplification operation is the edge contraction where the end-point vertices $\mathbf{v}_i, \mathbf{v}_j$ are replaced by a single vertex $\bar{\mathbf{v}}$, it is easy to incorporate the variation contribution by setting the quadric matrix of $\bar{\mathbf{v}}$ to be $\mathbf{Q}_{\bar{\mathbf{v}}} = \mathbf{Q}_{\mathbf{v}_i} + \mathbf{Q}_{\mathbf{v}_j}$.

In terms of feature scale, the quadric error metric is capable of measuring only the variation within the one-ring of a vertex. A natural extension is to include all the vertices from the $\mathcal{W}_v^n(\mathbf{v}, r)$ window using Gaussian weights to derive a matrix that reflects the variation in the entire window:

$$\mathbf{Q}_{\mathcal{W}_v^n(\mathbf{v},r)} = \frac{1}{S} \cdot \sum_{\mathbf{v}_i \in \mathcal{W}_v^n(\mathbf{v},r)} w(\mathbf{v}_i) Q_{\mathbf{v}_i}. \quad (25)$$

In consequence, the descriptor value using these averaged matrices can be expressed similarly as:

$$D(\mathbf{v}, n, r) = \frac{1}{S} \cdot \sum_{\mathbf{v}_i \in \mathcal{W}_v^n(\mathbf{v},r)} w(\mathbf{v}_i) \begin{pmatrix} \mathbf{v}_i \\ 1 \end{pmatrix}^\intercal \mathbf{Q}_{\mathcal{W}_v^n(\mathbf{v},r)} \begin{pmatrix} \mathbf{v}_i \\ 1 \end{pmatrix}. \quad (26)$$

5.2 PCA-based quadric error metric

In the presence of noise, the $\mathbf{Q}_{\mathbf{v}_i}$ matrices no longer accurately reflect the local geometric variation, and computing their smoothed window average will not solve the issue. Both quadric error metric and principal component analysis can visually represent this variability as isosurfaces that are quadrics. Bearing this in mind and given the properties of the eigenvectors of a covariance matrix, we can also construct a PCA-based definition of a quadric metric matrix. If λ_i is an eigenvalue corresponding to the $\mathbf{e}_i$ eigenvector, we can scale $\mathbf{e}_i$ by $\alpha_i = \frac{1}{1+\lambda_i}$. This scaling helps to construct a non-Euclidean distance function that attains larger values for any vector quantities aligned with the axes of highest variation, just like the Mahalanobis distance does. To proceed with the constructions, we arrange the eigenvectors into an orthonormal matrix:

$$\mathbf{R} = (\, \mathbf{e}_1 \; \mathbf{e}_2 \; \dots \; \mathbf{e}_d \,), \quad (27)$$

where d is the number of dimensions of the space into which the mesh is embedded (in our case, $d = 3$, but the same logic is valid in the general case). If $\mathbf{S}$ is the scaling matrix with $s_{i,i} = \alpha_i^{-1}$ and $s_{i,j} = 0$ for $i \neq j$, then the $\mathbf{RS}$ product represents a change of basis transforming a vector $\mathbf{v}_u$ written in the scaled local principal component space into another vector $\mathbf{v}_e$ written in the canonical Euclidean space. Thus, the reciprocal transform becomes:

$$\mathbf{v}_u = \mathbf{S}^{-1}\mathbf{R}^{-1}\mathbf{v}_e, \quad (28)$$

where $\mathbf{R}^{-1} = \mathbf{R}^\intercal$ considering the definition of $\mathbf{R}$.

The squared norm of the $\mathbf{v}_u$ vector can be easily computed from formula (28):

$$\mathbf{v}_u^\intercal\mathbf{v}_u = \mathbf{v}_e^\intercal\mathbf{M}\mathbf{v}_e, \quad (29)$$

where $\mathbf{M} = \mathbf{R}(\mathbf{S}^{-1})^2\mathbf{R}^\intercal$. If $\mathbf{p}$ is a vertex of the mesh for which the $\mathbf{R}$ and $\mathbf{S}$ matrices have been computed, the squared pseudo-Mahalanobis distance from $\mathbf{p}$ to an arbitrary vector $\mathbf{x}$ can be written in matrix form as:

$$d_{pM}(\mathbf{p}, \mathbf{x})^2 = (\mathbf{x} - \mathbf{p})^\intercal\mathbf{M}(\mathbf{x} - \mathbf{p}). \quad (30)$$

Through basic algebraic manipulation, we can also derive a block matrix expression for this distance such that:

$$d_{pM}(\mathbf{p}, \mathbf{x})^2 = \mathbf{x}^\intercal\mathbf{M}\mathbf{x} - 2\mathbf{b}^\intercal\mathbf{x} + c, \quad (31)$$

where $\mathbf{A}$ is a $d \times d$ matrix with elements

$$a_{i,k} = \sum_{j=\overline{1,d}} \mathbf{e}_{\mathbf{j}_i} \cdot \alpha_j^2 \mathbf{e}_{\mathbf{j}_k}, \quad (32)$$

and $\mathbf{b} = \mathbf{M}\mathbf{p}$ is a $d \times 1$ vector, and $c = \mathbf{p}^\intercal\mathbf{M}\mathbf{p}$ is a scalar term. Using these $\mathbf{M}$, $\mathbf{b}$ and c block elements, we can now define the PCA-based quadric error matrix for a vertex $\mathbf{v}$ as:

$$\mathbf{Q}_\mathbf{v}^{PCA} = \begin{pmatrix} \mathbf{M} & \mathbf{b} \\ \mathbf{b}^\intercal & c \end{pmatrix}. \quad (33)$$

The descriptor value at vertex $\mathbf{v}$ is then computed as described in formula (26) by replacing $\mathbf{Q}_{\mathcal{W}_v^n(\mathbf{v},r)}$ with $\mathbf{Q}_{\mathcal{W}_v^n(\mathbf{v}_i,r)}^{PCA}$. By construction, this PCA-based quadric error metric design is not as sensitive to the presence of noise, although an objective measure needs to be applied to better confirm this quality.

6. RESULTS

To distinguish between the qualities of the descriptors, we have conducted a series of experiments using four different mesh sets and compiled the results of these experiments according to two criteria: colormap comparison and average local distance.

We now examine the approximation quality and robustness of the descriptors presented in the previous sections. To generate the approximations, we have repeatedly used algorithm 1 and produced a chain of 8 fine-to-coarse approximations of 4 datasets. Two of the sets are closed, genus 0 meshes (available at `liris.cnrs.fr/meshbenchmark/`, the Dragon model, consisting of 50,000 points, and the Ramesses model, totaling 820,000 points). The other two sets were processed from LiDAR scans: one scan fragment of the Great Smoky Mountains (available through the `http://opentopo.sdsc.edu/` portal), and another custom scanned fragment of the Romanian Carpathian Mountains. The first LiDAR

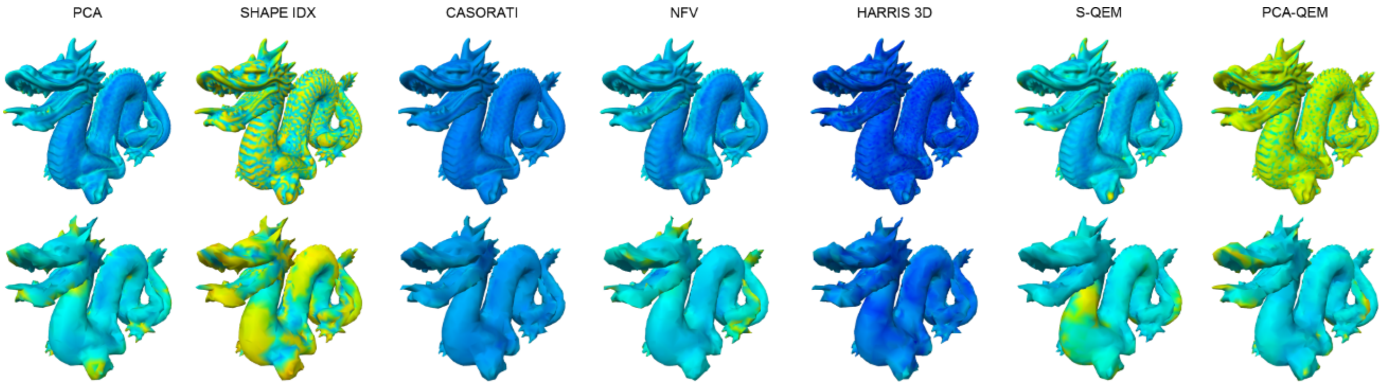

Fig. 3. Dragon mesh. The top row depicts the original set, while the bottom row shows the mesh after 8 simplification passes.

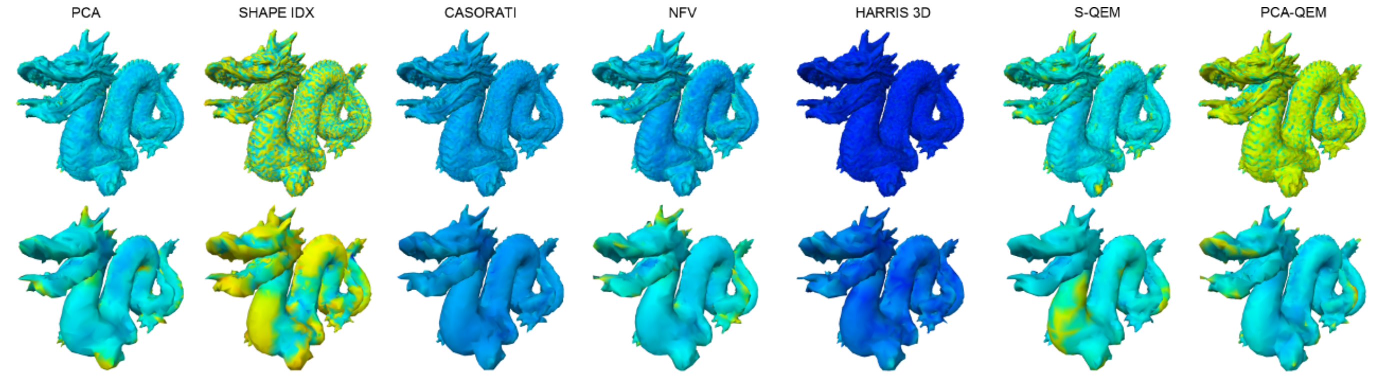

Fig. 4. Noise affected Dragon set. The top row depicts the unsimplified set, while the bottom row shows the mesh after 8 approximation passes.

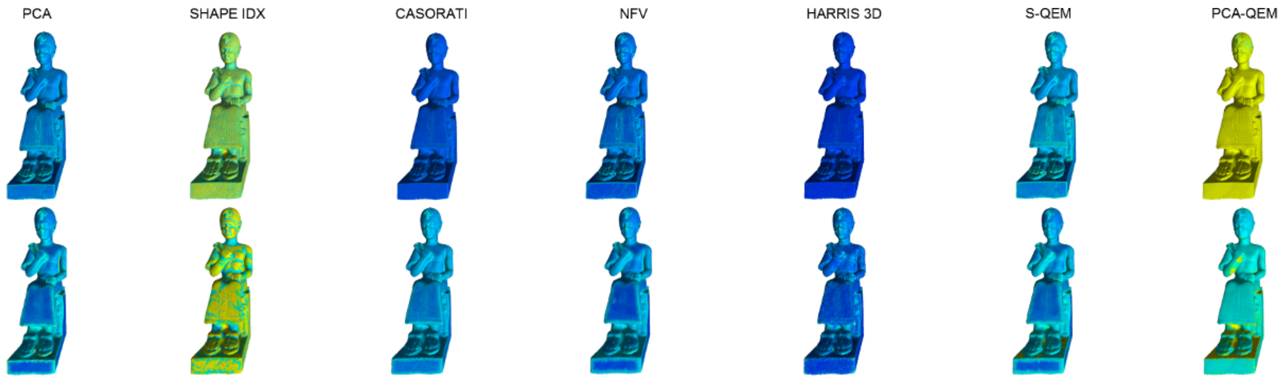

Fig. 5. Ramesses set. The top row depicts the original set, while the bottom row shows the mesh after 8 simplification passes.

set, representing a scan fragment of the Great Smoky Mountains, contains 270,000 points at a density of 2.2 points per square meter. In case of the Carpathian Mountains fragment, the point count is approximately 9 million points and the average density is of 20 points per square meter. To assess the robustness against noise, the points of the datasets were perturbed by adding artificial Gaussian noise of variance magnitude equal to one third of the average edge length.

All models were subjected to a sequence of 8 applications of algorithm 1. For these simplification experiments, the window parameters were initialized as follows: a limit of $n = 4$ topological rings and the window radius r being set to 3 times the average edge length. After each full iteration of algorithm 1, the output set of even vertices, V_E, becomes the input for the next stage. Afterwards, the V_O odd subset is extracted, the value of r is increased by 50% to compensate for the decrease in the point density caused by the change of scale. The average vertex count of the window has been experimentally determined to be 40 vertices for the highest resolution model. For the coarser resolutions, the average vertex count increases significantly, but does not exceed 70. For a completely regular mesh (i.e. each vertex having degree 6) with no radius constraints, the size of the window for $n = 4$

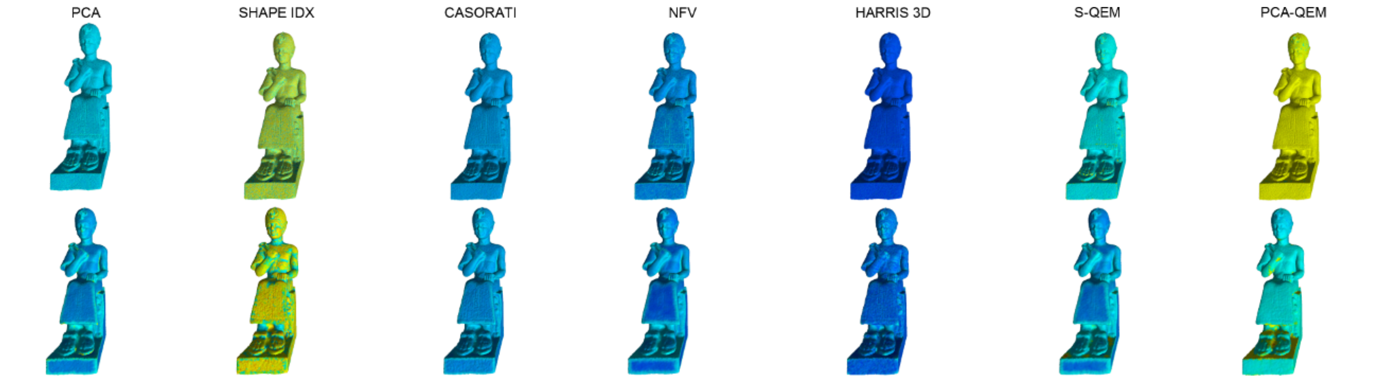

Fig. 6. Noise affected Ramesses set. The top row depicts the unsimplified set, while the bottom row shows the mesh after 8 approximation passes.

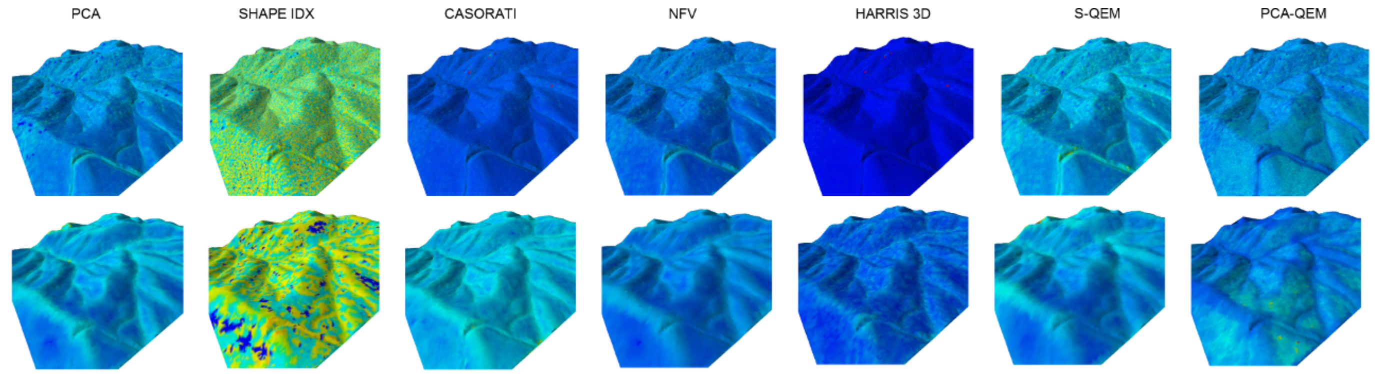

Fig. 7. Smoky data set. The top row depicts the original set, while the bottom row shows the mesh after 8 approximation passes.

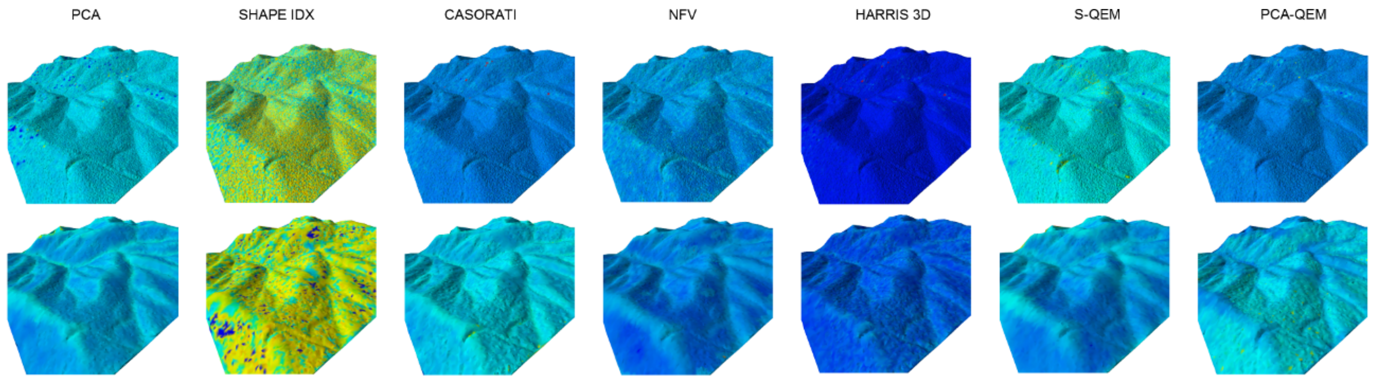

Fig. 8. Noise affected Smoky set. The top row depicts the original set, while the bottom row shows the mesh after 8 approximation passes.

is 60, thus the observed window sizes are acceptable in comparison to the ideal reglar mesh case.

For brevity, we have denoted the descriptors as follows: the PCA separability index (PCA), the Casorati curvature (CASORATI), the shape index (SHAPE IDX), the normal field variation (NFV), the Harris 3D corner detection response (HARRIS3D), the smoothed quadric error metric (S-QEM) and the PCA-based quadric error metric (PCA-QEM).

For visualization purposes, the models were rendered using a heat colormap obtained by normalizing the descriptor function values to the $[0, 1]$ interval. Besides the original sets, their noise-perturbed counterparts were also rendered. The results are presented in figures 3-10. These renderings allow for a preliminary interpretation that is data independent. The CASORATI and HARRIS3D descriptors achieve poor visual feature separation, at least for the highest resolution sets, where several isolated local maxima dominate the profile. The other descriptors, on the other hand, present higher levels of contrast and the distribution of the descriptor values better highlights the presence of scale features. In terms of feature highlighting abilities, the SHAPE IDX and PCA-QEM descriptors exhibit an exaggerated feature differentiation, at least for the Dragon (figure 3) and Ramesses (figure 5) closed surfaces.

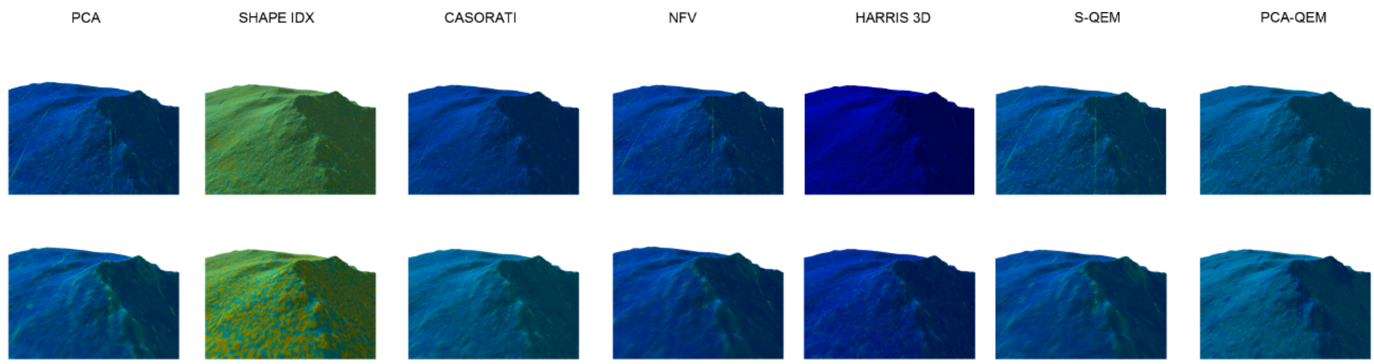

Fig. 9. Carpathian data set. The top row depicts the original set, while the bottom row shows the mesh after 8 approximation passes.

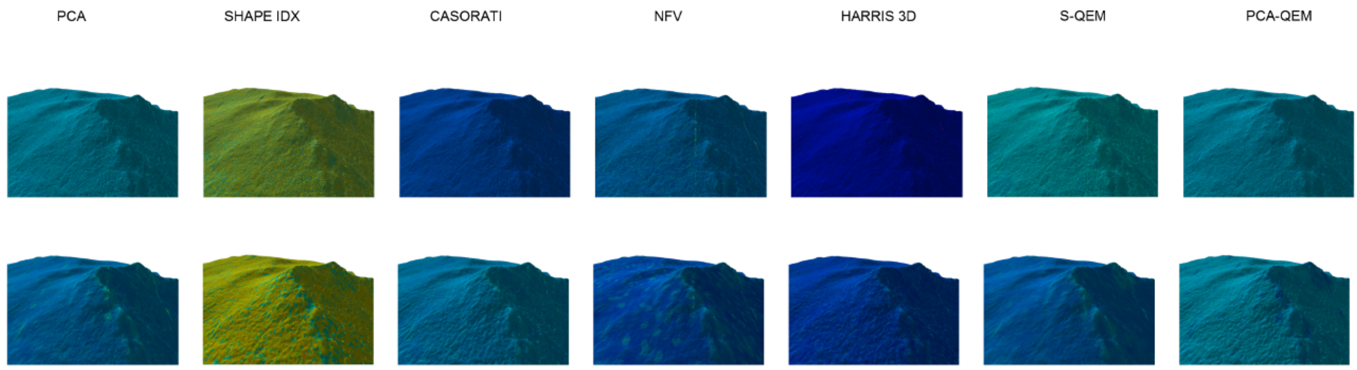

Fig. 10. Noise affected Carpathian set. The top row depicts the original set, while the bottom row shows the mesh after 8 approximation passes.

The PCA, NFV and S-QEM descriptors produce very similar colormap profiles for both the highest and lowest resolution models. Due to its construction, the PCA-QEM descriptor assigns a higher cost to the vertices closer to ridge or valley features, but does not accurately associate local extrema in the colormap profile with their geometric counterparts. When the data is perturbed artificially, the colormap contrast decreases for most descriptors, as can be observed in figures 4, 6, 8 and 10. In this sense, the PCA and S-QEM exhibit a dramatic decrease in contrast, especially where the higher frequency features are measured, i.e. at the highest level of resolution.

Besides a visual performance assessment, an objective approximation quality measure is required. Since the goal of our method is to be coupled with a lifting scheme, one option is to track the average local error that occurs after the even-odd partitioning. To define this error measure, we use the distance to the local one-ring best fit plane, first proposed by Schroeder et al. (1992) in the context of mesh decimation, and formulate the following distance operator:

$$\mathrm{apd}(V_O, V_E) = \frac{1}{\bar{l}|V_O|} \sum_{\mathbf{v} \in V_O} d(\mathbf{v}, \pi(\mathcal{N}_v^1(\mathbf{v}))), \qquad (34)$$

where $\bar{l}$ is the average edge length of the initial mesh, $\pi(\mathcal{N}_v^1(\mathbf{v}))$ is the best fit plane of $\mathcal{N}_v^1(\mathbf{v})$, and $d(\mathbf{v}, \pi(\mathcal{N}_v^1(\mathbf{v})))$ is the distance from $\mathbf{v}$ to this plane. Although this distance formulation does not directly depend on how the mesh is triangulated locally, it offers a good indication of the redundancy present in the neighborhood of a vertex, this being of immediate relevance for the application of wavelet analysis. Since the triangulation process can be performed locally inside the one-rings of the removed vertices, any valid triangulation algorithm can be employed. In our experiments we have opted for a modified version of the split plane strategy proposed by Schroeder et al. (1992). Instead of recursively splitting the one-ring hole that results after a vertex is removed, we collapse the edge emanating from $\mathbf{v}$ whose length is minimum and whose collapse does not create non-manifold artifacts. To prevent illegal topological changes or any non-manifold artifacts, the same mechanisms proposed by Schroeder et al. (1992) were adopted.

The evolution of the average distance in formula (34) was tracked at each level for each individual set. The results are charted in figures 11-18. The PCA descriptor achieved mediocre approximation quality regardless of the data set. The descriptors based on paraboloid fitting, SHAPE IDX, CASORATI, HARRIS 3D, outperformed, on average, all other proposed alternatives. However, the SHAPE IDX delivered better approximations for the Ramesses (figure 13) and Smoky (figure 15) sets, while the CASORATI descriptor seemed a better choice for the Dragon (figure 11) and Carpathians (figure 17) sets. The HARRIS 3D descriptor exhibited a much more consistent behavior, achieving, on average, the best results, regardless of the input set. Contrary to the intuition offered by the colormap profiles, the proposed NFV and S-QEM descriptors did not help achieve good quality approximations, yielding errors

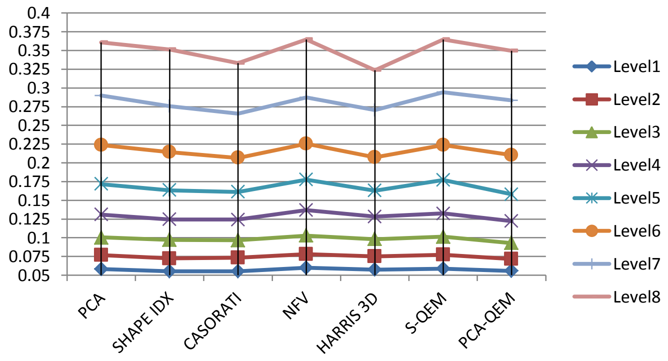

Fig. 11. Average plane distance evolution for the Dragon data set. The horizontal axis corresponds to the used descriptor, while the vertical axis corresponds to the averaged local distance. All intermediate errors are plotted, yielding evolution curves for a better method comparison.

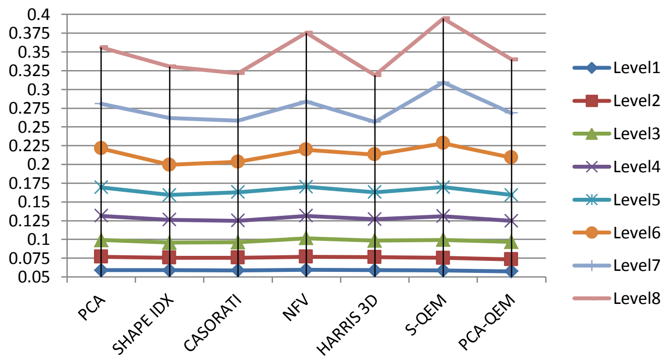

Fig. 12. Average plane distance evolution for the noise affected Dragon data set. The horizontal and vertical axes have the same meaning as in figure 11.

almost 10% higher than those offered by the paraboloid fitting descriptors. The PCA-QEM did outperform the pure PCA descriptor and, for the terrain data sets (figures 15 and 17), even surpassed the paraboloid fitting alternatives.

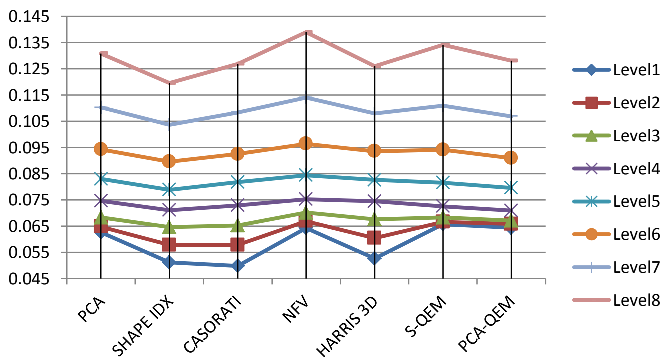

Fig. 13. Average plane distance evolution for the Ramesses data set. The meaning of each axis is the same as in figure 11.

The addition of artificial noise did not seem to affect the approximation quality. The plots shown in figures 12, 14, 16 and 18 were obtained using formula (34)

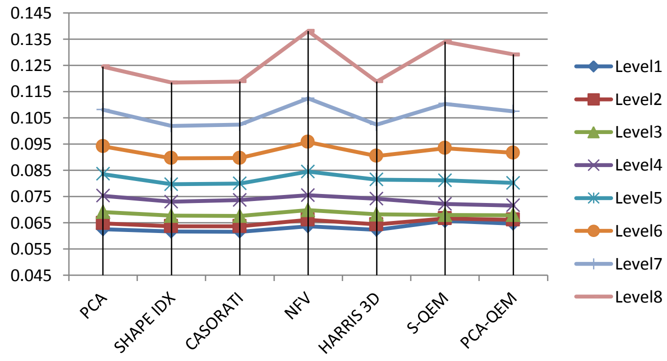

Fig. 14. Average local distance evolution for the noise affected Ramesses data set. The horizontal and vertical axes have the same meaning as in figure 11.

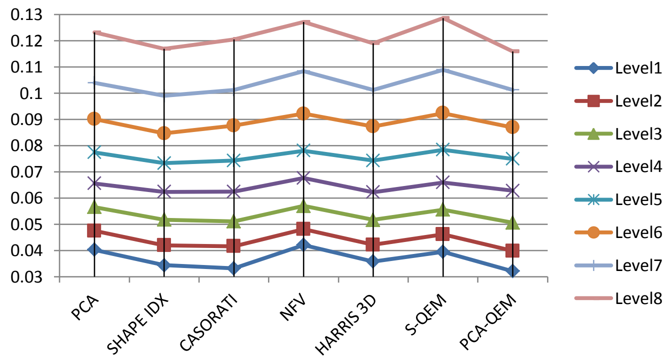

Fig. 15. Average local distance evolution for the Smoky data set. The horizontal and vertical axes have the same meaning as in figure 11.

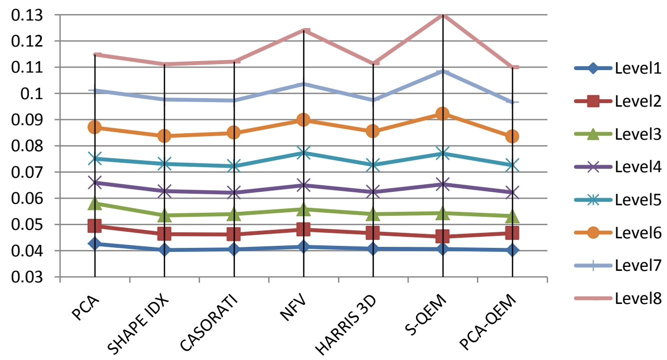

Fig. 16. Average local distance evolution for the noise affected Smoky data set. The horizontal and vertical axes have the same meaning as in figure 11.

where the initial vertex coordinates were used, the noise affected ones being considered only when computing the descriptor values. The simplification experiments reveal a mild increase in the approximation error for the fine resolution levels. For the coarser approximations, the error actually decreases on average, but the relative behavior of a descriptor with respect to its competitors does not change from the one observed when using the unperturbed sets. In conclusion, these charts indicate that, for the chosen noise intensity, the influence of noise on the error is not considerable, all descriptors being reasonably robust.

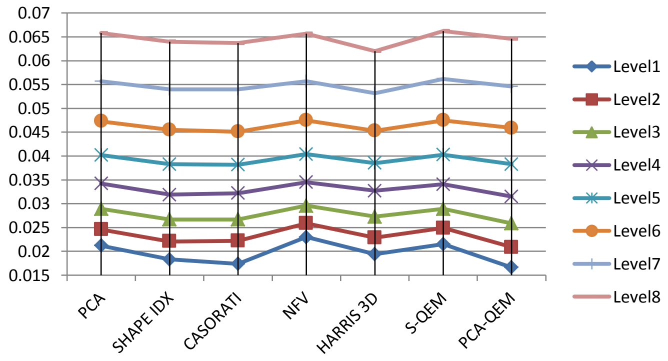

Fig. 17. Average local distance evolution for the Carpathians set. The horizontal and vertical axes have the same meaning as in figure 11.

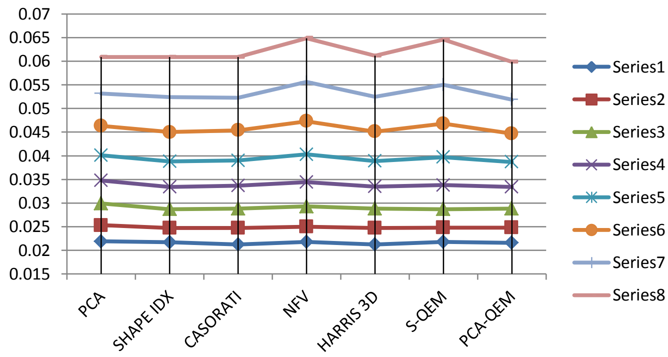

Fig. 18. Average local distance evolution for the noise affected Carpathians set. The horizontal and vertical axes have the same meaning as in figure 11.

7. CONCLUSION

In this work, we have experimentally analyzed the discriminative properties of three different families of 3D feature descriptors: PCA-based (PCA, NFV), paraboloid fitting methods (CASORATI, SHAPE IDX, HARRIS3D), and quadric error metric-based (S-QEM, PCA-QEM). A lazy wavelet partitioning mechanism, coupled with a fast local triangulation strategy was employed to simplify closed meshes and terrain meshes with a boundary. To assess the quality of the resulting approximations, a local redundancy measure based on the distance from a vertex to its average one-ring plane was used. The PCA-QEM was found to produce better approximations than the S-QEM and NFV descriptors. Since PCA-QEM has the same algebraic properties as the S-QEM, it is suitable for scenarios where both geometry and attribute data are present. The modified Harris 3D descriptor delivered the best approximations, but this descriptor is amenable to geometric analysis only and has the highest computational overhead. If such an overhead is relevant, the results shown in our work advocate for the use of either the PCA descriptor (for 3D data only) or the PCA-QEM one, since these deliver intermediate quality results. Although the other descriptors we have proposed, the NFV and S-QEM, managed to visually highlight features, their measured accuracy did not recommend them as suitable alternatives.

The methodology employed in our analysis followed the requirements of a lazy wavelet partitioning setup where the odd data samples are removed in order to produce the coarser approximation. The limitation imposed by this partitioning strategy is inherent to the methodology, and its influence on the experiments could be explored in a future work. In this respect, the NFV and S-QEM descriptors could be re-examined to better understand why their discriminative properties did not aid in lowering the approximation error. Principal component analysis, required for computing the PCA, NFV, HARRIS 3D and QEM-PCA descriptor values, is not fully impervious to noise. While robustness could be increased by expanding the local window parameters (radius and number of rings), the ability of identifying fine scale features would decrease. Another factor that could affect the reliability of feature discrimination is a highly irregular sampling density. An interesting question is hence how to improve the descriptor computation in order to adjust to irregularly sampled sets.

REFERENCES

Bonde, U., Badrinarayanan, V., and Cipolla, R. (2013). Multi scale shape index for 3d object recognition. In A. Kuijper, K. Bredies, T. Pock, and H. Bischof (eds.), *Scale Space and Variational Methods in Computer Vision*, volume 7893 of *Lecture Notes in Computer Science*, 306–318. Springer Berlin Heidelberg. doi:10.1007/978-3-642-38267-3_26.

Botsch, M., Kobbelt, L., Pauly, M., Alliez, P., and uno Levy, B. (2010). *Polygon Mesh Processing*. AK Peters.

Boyer, K.L., Srikantiah, R., and Flynn, P.J. (2002). Saliency sequential surface organization for free-form object recognition. *Comput. Vis. Image Underst.*, 88(3), 152–188. doi:10.1006/cviu.2002.0973.

Brentzen, J.A., Gravesen, J., Anton, F., and Aans, H. (2012). *Guide to Computational Geometry Processing: Foundations, Algorithms, and Methods*. Springer Publishing Company, Incorporated.

Bronstein, A.M., Bronstein, M.M., Guibas, L.J., and Ovsjanikov, M. (2011). Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Trans. Graph.*, 30(1), 1:1–1:20. doi:10.1145/1899404.1899405.

Cioaca, T., Dumitrescu, B., and Stupariu, M. (2016). Graph-based wavelet representation of multi-variate terrain data. *Comput. Graph. Forum*, 35(1), 44–58. URL `http://onlinelibrary.wiley.com/doi/10.1111/cgf.12670/full`.

Cipriano, G., Phillips, G., and Gleicher, M. (2009). Multiscale surface descriptors. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6), 1201–1208. doi:10.1109/TVCG.2009.168.

Creusot, C., Pears, N., and Austin, J. (2013). A machine-learning approach to keypoint detection and landmarking on 3d meshes. *Int. J. Comput. Vision*, 102(1-3), 146–179. doi:10.1007/s11263-012-0605-9.

Fadaifard, H. and Wolberg, G. (2011). Multiscale 3d feature extraction and matching. In *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2011 International Conference on*, 228–235. doi:10.1109/3DIMPVT.2011.36.

Garland, M. and Heckbert, P.S. (1998). Simplifying surfaces with color and texture using quadric error metrics.

In *Proceedings of the Conference on Visualization '98*, VIS '98, 263–269. doi:10.1145/288216.288280.

Guskov, I., Sweldens, W., and Schröder, P. (1999). Multiresolution signal processing for meshes. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, 325–334. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA. doi:10.1145/311535.311577.

Harris, C. and Stephens, M. (1988). A combined corner and edge detector. In *In Proc. of Fourth Alvey Vision Conference*, 147–151.

Ho, H. and Gibbins, D. (2009). Curvature-based approach for multi-scale feature extraction from 3d meshes and unstructured point clouds. *Computer Vision, IET*, 3(4), 201–212. doi:10.1049/iet-cvi.2009.0044.

Hoppe, H. (1996). Progressive meshes. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '96, 99–108. ACM, New York, NY, USA. doi:10.1145/237170.237216.

Hussain, M. (2010). Fast and reliable decimation of polygonal models based on volume and normal field. In *Proceedings of the 6th International Conference on Advances in Visual Computing - Volume Part I*, ISVC'10, 64–73. Springer-Verlag, Berlin, Heidelberg.

Kang, M., Gao, J., and Tang, L. (2011). Nonlinear ransac optimization for parameter estimation with applications to phagocyte transmigration. In *Proceedings of the 2011 10th International Conference on Machine Learning and Applications and Workshops - Volume 01*, ICMLA '11, 501–504. IEEE Computer Society, Washington, DC, USA. doi:10.1109/ICMLA.2011.104.

Koenderink, J.J. and van Doorn, A.J. (1992). Surface shape and curvature scales. *Image Vision Comput.*, 10(8), 557–565. doi:10.1016/0262-8856(92)90076-F.

Koibuchi, H. and Yamada, M. (2000). Phase transition of a model of crystalline membrane. *International Journal of Modern Physics C*, 11(08), 1509–1518. doi:10.1142/S0129183100001462.

Liu, Y.S., Liu, M., Kihara, D., and Ramani, K. (2007). Salient critical points for meshes. In *Proceedings of the 2007 ACM Symposium on Solid and Physical Modeling*, SPM '07, 277–282. ACM, New York, NY, USA. doi:10.1145/1236246.1236285.

Maximo, A., Patro, R., Varshney, A., and Farias, R. (2011). A robust and rotationally invariant local surface descriptor with applications to non-local mesh processing. *Graph. Models*, 73(5), 231–242. doi:10.1016/j.gmod.2011.05.002.

Mitra, N.J., Guibas, L., Giesen, J., and Pauly, M. (2006). Probabilistic fingerprints for shapes. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing*, SGP '06, 121–130. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland. doi:10.1145/1281957.1281973.

Nader, G., Guennebaud, G., and Mellado, N. (2014). Adaptive multi-scale analysis for point-based surface editing. *Comput. Graph. Forum*, 33(7), 171–179. doi:10.1111/cgf.12485.

Page, D.L., Sun, Y., Koschan, A.F., Paik, J., and Abidi, M.A. (2002). Normal vector voting: Crease detection and curvature estimation on large, noisy meshes. *Graph. Models*, 64(3-4), 199–229. doi:10.1006/gmod.2002.0574.

Park, M.K., Lee, S.J., and Lee, K.H. (2012). Multi-scale tensor voting for feature extraction from unstructured point clouds. *Graph. Models*, 74(4), 197–208. doi:10.1016/j.gmod.2012.04.008.

Pauly, M., Gross, M., and Kobbelt, L.P. (2002). Efficient simplification of point-sampled surfaces. In *Proceedings of the Conference on Visualization '02*, VIS '02, 163–170. IEEE Computer Society, Washington, DC, USA. doi:10.1109/VISUAL.2002.1183771.

Rocca, L., Giorgis, N.D., Panozzo, D., and Puppo, E. (2011). Fast neighborhood search on polygonal meshes. In A.F. Abate, M. Nappi, and G. Tortora (eds.), *Eurographics Italian Chapter Conference*, 15–21. Eurographics.

Schroeder, W.J., Zarge, J.A., and Lorensen, W.E. (1992). Decimation of triangle meshes. *SIGGRAPH Comput. Graph.*, 26(2), 65–70. doi:10.1145/142920.134010.

Shtrom, E., Leifman, G., and Tal, A. (2013). Saliency detection in large point sets. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, 3591–3598. doi:10.1109/ICCV.2013.446.

Sipiran, I. and Bustos, B. (2011). Harris 3d: A robust extension of the harris operator for interest point detection on 3d meshes. *Vis. Comput.*, 27(11), 963–976. doi:10.1007/s00371-011-0610-y.

Stefanoski, N., Liu, X., Klie, P., and Ostermann, J. (2007). Scalable linear predictive coding of time-consistent 3d mesh sequences. In *3DTV Conference, 2007*, 1–4. doi:10.1109/3DTV.2007.4379461.

Sweldens, W. (1996). The lifting scheme: A custom-design construction of biorthogonal wavelets. *Applied and Computational Harmonic Analysis*, 3(2), 186 – 200. doi:http://dx.doi.org/10.1006/acha.1996.0015.

Tang, S. and Godil, A. (2012). An evaluation of local shape descriptors for 3d shape retrieval. In *Three-Dimensional Image Processing (3DIP) and Applications II, Burlingame, California, USA, January 22, 2012*, 82900N. doi:10.1117/12.912153. URL `http://dx.doi.org/10.1117/12.912153`.

Tangelder, J.W. and Veltkamp, R.C. (2008). A survey of content based 3d shape retrieval methods. *Multimedia Tools Appl.*, 39(3), 441–471. doi:10.1007/s11042-007-0181-0.

Tao, P., Cao, J., Li, S., Liu, X., and Liu, L. (2015). Mesh saliency via ranking unsalient patches in a descriptor space. *Computers & Graphics*, 46, 264–274. doi:10.1016/j.cag.2014.09.023.

Wu, J., Shen, X., Zhu, W., and Liu, L. (2013). Mesh saliency with global rarity. *Graph. Models*, 75(5), 255–264. doi:10.1016/j.gmod.2013.05.002.

Yang, Y.L. and Shen, C.H. (2012). Multi-scale salient feature extraction on mesh models. In *Proceedings of the First International Conference on Computational Visual Media*, CVM'12, 122–129. Springer-Verlag, Berlin, Heidelberg. doi:10.1007/978-3-642-34263-9_16.

Zorin, D., Schröder, P., and Sweldens, W. (1997). Interactive multiresolution mesh editing. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '97, 259–268. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA. doi:10.1145/258734.258863.