Multiresolution Analysis

Georges-Pierre Bonneau[1], Gershon Elber[2], Stefanie Hahmann[3], and Basile Sauvage[3]

[1] Université Joseph Fourier, Grenoble, France,
Georges-Pierre.Bonneau@imag.fr
[2] Technion, Israel Institute of Technology, Israel,
gershon@cs.technion.ac.il
[3] Institut National Polytechnique de Grenoble, France,
Stefanie.Hahmann@imag.fr, basile.sauvage@imag.fr

Summary. Multiresolution analysis has received considerable attention in recent years by researchers in the fields of computer graphics, geometric modeling and visualization. They are now considered a powerful tool for efficiently representing functions at multiple levels-of-detail with many inherent advantages, including compression, Level-Of-Details (LOD) display, progressive transmission and LOD editing.

This survey chapter attempts to provide an overview of the recent results on the topic of multiresolution, with special focus on the work of researchers who are participating in the AIM@SHAPE European Networks of Excellence [4].

1 Introduction

Multiresolution analysis has received considerable attention in recent years by researchers in the fields of computer graphics, geometric modeling and visualization [83]. Its attraction is its utility as a powerful tool for efficiently representing functions at multiple levels-of-detail with many inherent advantages, including compression, Level-Of-Details (LOD) display, progressive transmission and LOD editing. A plethora of publications can be easily found on these topics.

This survey chapter attempts to provide an overview of the recent results on the topic of multiresolution, with special focus on the work of researchers who are participating in the AIM@SHAPE European Networks of Excellence.

In Section 2, hierarchical freeform representations are introduced and discussed. Multiresolution methods for freeform spline spaces are discussed in Section 3, including linear and non-linear constraints, and intrinsic multiresolution decomposition. Multiresolution representation of piecewise linear and triangular irregular meshes are considered in Section 4. Finally, we conclude this chapter, in Section 5.

[4] AIM@SHAPE Project, http://www.aimatshape.net

2 Hierarchical Freeform Representations

Forsey and Bartels pioneered the idea of hierarchical B-splines [25]. B-splines can be locally refined using overlays. Based on this model, these researchers created a complex surface such as a dragon's head from a rectangular domain with a hierarchical edition. However, since this model was established over tensor product splines, it is restricted to tensor product mesh and topology.

Localized-hierarchical surface splines [32] extended the hierarchical spline paradigm to surfaces of arbitrary topology. These are defined locally on a hierarchy of meshes using the "reference plus offset" model of Forsey and Bartels for encoding the details. Since they are based on C^1-surface-splines [72], the surface is defined explicitly by low-degree triangular and quadrangular Bézier patches, while requiring the structure to satisfy a particular regularity property through all levels of the hierarchy.

Hierarchical triangular splines [90] are the most recent method for hierarchical modeling of smooth surfaces of arbitrary topology. Based on the previously developed triangular interpolating scheme [45], this method enables LOD construction and surface editing by interpolating the vertices of a hierarchy of locally refined meshes. The initial mesh, referred to as the base mesh, can be any triangular two-manifold mesh. Given a base mesh, a polynomial interpolating surface is computed. It is smooth in the sense that it is overall tangent-plane continuous. LOD is then added by iterative local refinement and editing of the surface. Each local surface refinement replaces a set of coarse surface patches by a set of finer surface ones, while maintaining both the overall tangent-plane continuity and the shape. The user can add detail by editing the refined surface patches. A hierarchical editing tool is also provided thanks to a "reference plus offset" representation. The main features include:

- Any *triangular mesh* can be handled, which means there are *no restrictions* on topology, geometry, genus or boundaries. It only has to be a two-manifold mesh.
- The surface *interpolates* a hierarchy of meshes, thus offering direct control for surface modifications, in different resolutions.
- *Uniform surface model*: The same interpolant is applied to both the initial surface and its different refinement steps, in order to locally recompute the new surface part. Thus, only a *few geometric quantities* are stored for each surface patch in order to completely evaluate the surface.

Further properties that the hierarchical surface inherits from the underlying surface model include:

- Overall tangent-plane continuity.
- Each surface patch is represented as a parametric, polynomial triangular Bézier patch of degree five.
- The surface has local control, i.e., the modification of a mesh vertex modifies only the surrounding surface patches, leaving the surface unchanged outside this region.

Figure 1 presents a hierarchical editing process for a surface composed of triangular patches. The different colors, from white through yellow to red, denote the different levels of detail (see Figure CP-1 in Appendix C).

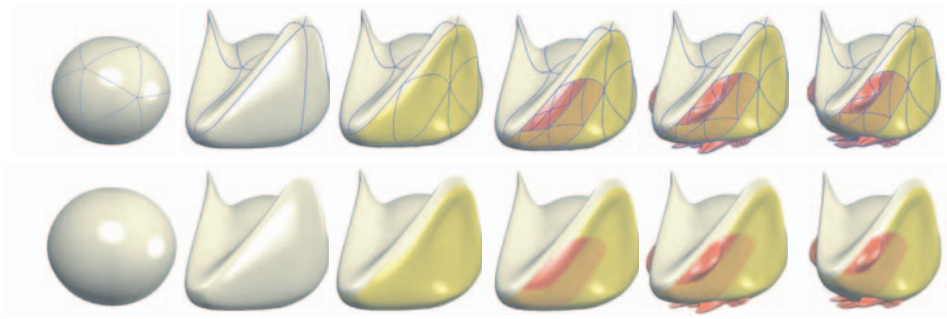

Fig. 1. Hierarchical surface representation and editing: From left to right, a hierarchical editing of an object is shown. Colors correspond to different levels of detail. Starting with an initial surface at the coarsest level, local refinements and local editings are gradually introduced. Finally, the surface is edited at a vertex of the coarsest level, thus naturally deforming all finer details depending hierarchically on this vertex.

3 Multiresolution Methods for Freeform Representations

In the literature, the term multiresolution (MR) is employed in different contexts, including MR-based wavelets, subdivision and hierarchies or multigrids. Multiresolution representations based on wavelets have been developed for parametric curves [14, 22, 62], and can be generalized to tensor-product surfaces [22, 52], to surfaces of arbitrary topological type [61], to spherical data [80], and to volume data [15]. Wavelets provide a rigorous unified framework. Herein, a complex function is decomposed into "coarser" low-resolution parts, together with a collection of detail coefficients and different resolution levels, necessary to recover the original function. Other multiresolution representations exist for data defined for tensor-product surfaces, known as hierarchical B-splines [25], and for volumetric data sets represented using tri-variate functions [75].

In the context of geometric modeling, LOD editing is an attractive MR application because it allows modification of the overall shape of a geometric model at any scale while automatically preserving all fine details. In contrast to classical control-point-based editing methods where complex detail-preserving deformations need to manipulate a lot of control points, MR methods can achieve the same effect by manipulating only a few control points of some low-resolution representation; see [22, 83]. However, there are application areas, including Computer Aided Geometric Design (CAGD) and computer animation, where deformations under constraints

are required. It is obvious that constraints offer additional and finer controls over the deformations applied to curves and surfaces.

The remainder of this section surveys recent results on MR methods in the context of freeform spline geometry, with and without constraints. In Section 3.1 we briefly describe B-wavelets, or wavelets of B-spline functions. In Section 3.2, direct manipulation of freeform curves and surfaces are presented whereas in Sections 3.3 and 3.4 linear and non-linear constraints are discussed, respectively. In Section 3.5, intrinsic MR decomposition of freeform geometry is considered, employing curvature signatures of the shapes. The application of MR to metamorphosis is considered in Section 3.6 and finally, variational design that aims at optimizing and/or fairing the shape is discussed in the context of MR representations, in Section 3.7.

3.1 Wavelet Decomposition of B-spline Functions

Multiresolution manipulation of geometry draws from the ability to project geometry G_i in space $\mathcal{S}_i$ onto another subspace $\mathcal{S}_{i+1} \subset \mathcal{S}_i$. Spline spaces are solely defined by the knot sequences $\tau_{\mathbf{i}}$ (and the orders o_i). In [14, 62], wavelet decomposition of spline spaces, both uniform and non-uniform, were presented. Subspaces are typically selected by removing every second knot, preserving the uniformity of the knot sequence or possibly by weighing the importance of the knots, as is done, for example, in knot removal algorithms [63].

Consider a curve $C(t) \in \mathcal{S}_i$ with a uniform knot sequence $\tau_j = j$. The removal of a single knot, τ_k, creates a sub space $\mathcal{S}_{i+1}$ in which no discontinuity can be present at parameter value τ_k. The B-spline wavelet (also known as B-wavelet) Ψ_k that corresponds to knot τ_k spans the complementary subspace of $\mathcal{S}_i - \mathcal{S}_{i+1}$. While many ways exist to define the function that spans the complementary space, seeking a unique orthogonal representation to Ψ_k, we constrain Ψ_k to be orthogonal to all the B-spline basis functions in $\mathcal{S}_{i+1}$. Since $\Psi_k \in \mathcal{S}_i$, Ψ_k has one additional degree of freedom, which is typically used to normalize Ψ_k, for example with the constraint of $\langle \Psi_k, \Psi_k \rangle = 1$. By using only uniform knot sequences, all B-wavelets are just translations (and scales) of each other. Yet, nothing in the above prevents one from using non-uniform knot sequences with the cost of no possible precomputations. All B-wavelets must now be reevaluated for every new knot sequence. Figures 2 and 3 show several B-wavelet functions for the quadratic and cubic cases, respectively. Both uniform (computed once up to translation and scale!) and non-uniform B-wavelets are shown.

A B-spline curve $C(t)$ is typically decomposed into a low-resolution curve $C_0(t)$ and a sequence of detail curves $D_i(t)$ at different resolutions so that

$$C(t) = C_0(t) + \sum_{i=1}^{n} D_i(t). \tag{1}$$

$C_0(t)$ is the lowest or coarsest resolution and typically contains no interior knots in its subspace. Every additional detail curve $D_i(t)$ contains additional knots all the way to $D_n(t)$. These knots are all shared by the original space of $C(t)$. The vector

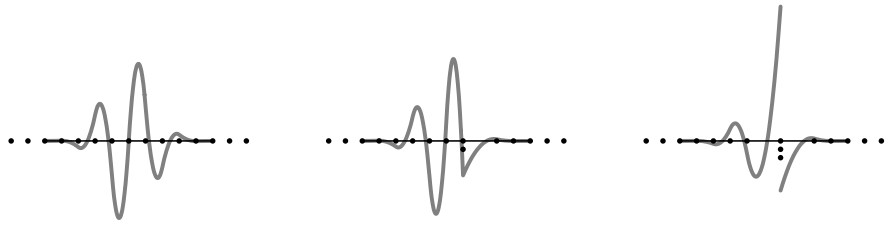

Fig. 2. Quadratic B-wavelets for the uniform case (left) and multiple knots (middle and right). Note a triple knot renders the quadratic B-wavelet discontinuous.

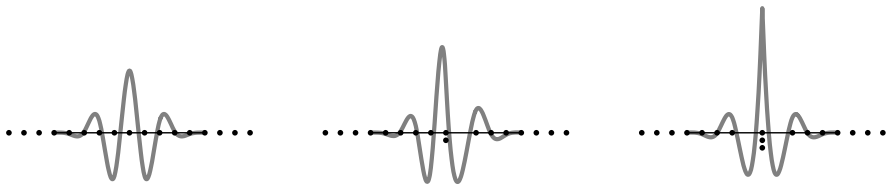

Fig. 3. Cubic B-wavelets for the uniform case (left) and multiple knots (middle and right). Note a triple knot renders the cubic B-wavelet C^0 continuous.

function addition in Equation (1) is always possible since both $C_0(t)$ and $D_i(t)$ remain in the subspace of the original space. In other words, by refinement, one can always elevate $C_0(t)$ and $D_i(t)$ to the original space, where the sum presented in Equation (1) reduces to adding the respective control points of the curves. Figure 4 presents one example of a multiresolution decomposition of a freeform curve.

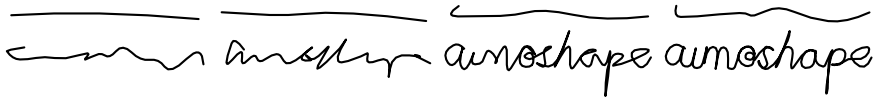

Fig. 4. A decomposition of a B-spline curve into various resolutions. The original quadratic curve is shown at the bottom right and contains over a hundred control points.

By modifying a single control point in $C_0(t)$, the entire shape of $C(t)$ is affected. By modifying the $D_i(t)$ vector functions, one is able to create modifications in different resolutions, from a coarse resolution for $D_1(t)$ all the way to fine details in $D_n(t)$. Figure 5 presents an example of manipulating a freeform curve at different MR levels.

One typical application for MR analysis of spline geometry could be found in the direct manipulation of a freeform shape (see also Section 3.2 below). The local support of the B-spline representation is also the weakest point of the representation. Global modifications are no longer possible in a highly refined B-spline curve. Recognizing this deficiency, in [36, 22], wavelet decomposition was proposed for

Fig. 5. Modification of a B-spline curve at various resolutions. A vertical select-and-drag operation at the top of the 's' character at four different resolutions. The original curve is presented in gray.

uniform B-spline curves toward MR editing control of the shape. When the user wishes to add small details to the shape, a fine subspace is used during the manipulation whereas when global changes are necessary, a coarse resolution is employed.

One clear advantage of using uniform knot sequences is that it allows wavelet decomposition to be performed a-priori, as the decomposition depends solely on the subspaces of the splines and is completely independent of the control points of the shapes. Yet, in reality, many curves and surfaces that are created using contemporary geometric modeling tools possess non-uniform knot sequences. Further, in order to preserve the uniformity of the knots, in a given curve with a uniform knot sequence, every subspace must present half the number of knots of its immediate containing space. That is, $\tau_{\mathbf{i+1}}$ of $\mathcal{S}_{i+1}$ will consist of half the knots in $\tau_{\mathbf{i}}$, with every second knot in $\tau_{\mathbf{i}}$ being removed, preserving the uniformity in the knot spacing.

The work of [36, 22] was extended to non-uniform knot sequences for curves and surfaces, in [52]. Direct manipulation of non-uniform B-spline curves and surfaces is presented in [52] with the aid of a B-wavelet decomposition [62]. Figure 6 shows an example of MR interactive editing, in different resolutions, of a freeform tensor-product B-spline surface in the shape of a chess knight (see also Figure CP-2 in Appendix C).

3.2 Direct Freeform Curve and Surface Manipulation

As already stated, direct manipulation of freeform shapes is a crucial and vital tool in any modern geometric modeling environment. Being able to sculpt the geometry allows novice users to intuitively and interactively manipulate the shape.

Direct manipulation of freeform surfaces is not new and, for example, in [25], a hierarchical representation of B-spline surfaces is presented that allows local and focused manipulation of freeform geometry. Adding degrees of freedom to a freeform surface is usually translated into the insertion of new knots into the shape—an action that affects a whole row or column in the mesh of the surface, and hence is not

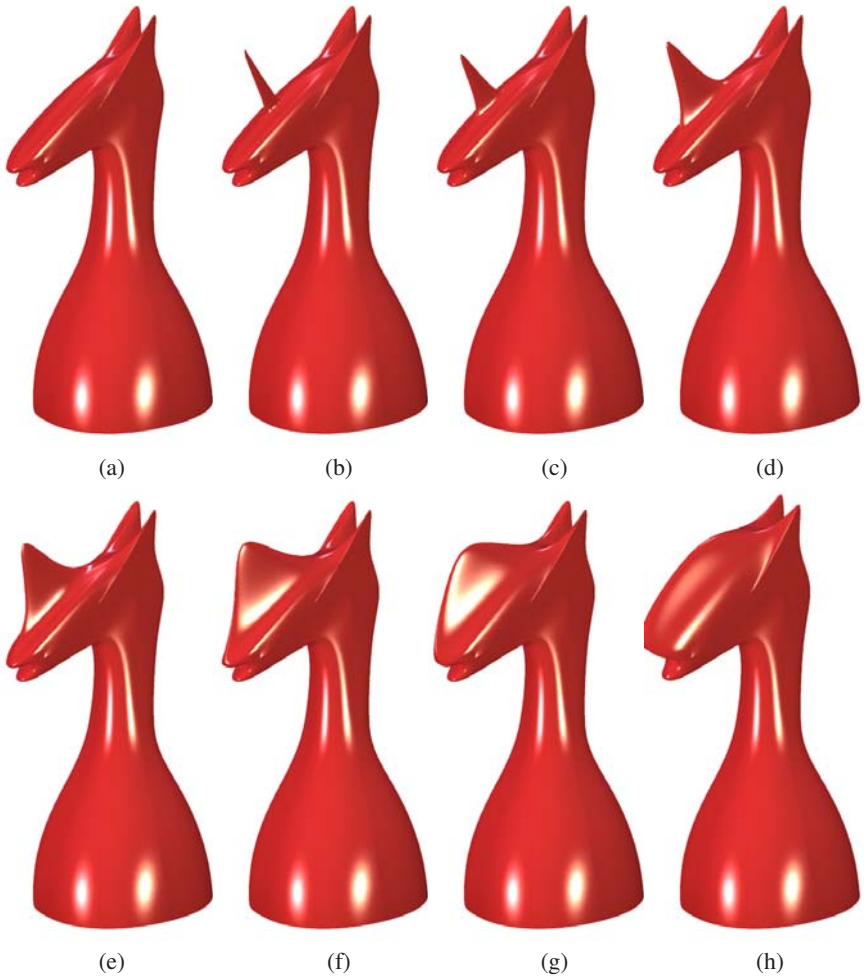

Fig. 6. An example of multiresolution editing of a tensor-product B-spline surface in the shape of a chess knight. A forehead location is selected and dragged upwards in several different resolutions. (a) shows the original surface, while (b) to (h) present the results of the select-and-drag operations in the different resolutions from the highest (b) to to lowest (h).

really local. In contrast, in [25], a hierarchy of partially independent surfaces is imposed that allows the end user to locally affect only a small region in a given surface, by applying a small patch with the new detail at the desired location. This occurs while fixing the outermost rows and columns of the new small patch to preserve the proper continuity. A related scheme for volumetric representations was offered in [75]. Here, a hierarchy of tri-variate functions of different resolution is used to define the sculpted surface. The (iso) surface itself is defined as the zero set of the sum of these tri-variates.

Rather than manipulating control points, Bartels and Beatty showed, in [2], how to select any point on a B-spline curve and change its location, i.e., the curve is constrained to pass through a user-specified location. The new curve shape is computed by minimizing the control points' offset. In [28] Fowler and Bartels controlled the shape of a B-spline curve by enforcing prescribed geometric constraints, such as the position of a curve point, tangent direction and magnitude, or curvature magnitude. An extension to tensor-product B-spline surfaces is given in [26]. This satisfies the user-defined position of surface points, normal direction, tangent-plane rotation (twisting effect), and the first partial derivative's magnitude (tension effect). Borel and Rappoport [9] deformed B-spline surfaces by determining the displacement and radius of influence for each constrained surface point. Hsu et al. [50] proposed points selection for freeform deformations. Curve constraints, i.e., enforcing the surface to contain a given curve or to model a character line, were considered in [12, 38, 71]. Direct shape manipulation techniques are closely related to variational design, where the objective of obtaining fair and graceful shapes is achieved by minimizing some energy; see Section 3.7. In general, a freeform shape has many more degrees of freedom than constraints to satisfy. In order to compute a new shape, the remaining degrees of freedom are stipulated by minimizing some energy functional, such as bending. For example, Welch et al. [89] maintained the imposed constraints while calculating a surface that is as smooth as possible. Celniker and Welch [12] derived interactive sculpting techniques for B-spline surfaces based on energy minimization, keeping some linear geometric surface-constrained features unchanged. Celniker and Gossard [11] enforced linear geometric constraints for shape design of finite elements governed by some surface energy. While energy minimization affects the surface globally, finite element methods allow for local control. Forsey and Bartels [25] later used the technique of hierarchical B-splines in an attempt to overcome this drawback for B-spline surfaces.

In the context of MR, [36, 22] offered direct multiresolution manipulation of uniform B-spline curves and surfaces. While no constraint support was offered in these publications, they demonstrated, for the first time, the hidden power in MR editing and direct manipulation of freeform curves and surfaces. Exact B-spline wavelet (B-wavelets) decomposition was used to perform the MR analysis. In [52] and using the results of [62], the approach of using precise B-wavelet decomposition in direct curve and surface manipulation was extended to non-uniform B-spline space. Also demonstrated in [36, 22] were abilities to add details of different shapes to an existing curve—another modeling feature of high interest.

The work of [22, 36, 52] computed the exact orthogonal projections of the freeform geometry into lower dimensional spaces, employing the B-wavelet decomposition of uniform and non-uniform B-spline representations. While fairly simple to compute in the case of uniform knot sequences, this decomposition, in the non-uniform case, must be recomputed for each newly defined space and is computationally intensive. Fortunately, one can recognize that the explicit orthogonal decomposition is not really necessary [35], alleviating these computational difficulties. In [20], an MR curve editor that is based on a non-orthogonal decomposition was also presented. The major deficiency of this non-orthogonal decomposition lies in its ambiguous representation, by offering the user, for example, the option of conducting many fine high-resolution operations, which can, in fact, be represented as a few low-resolution operations. The (approximated) projection of a curve to a low-dimensional space is simple, and for direct manipulation purposes, it might be sufficient.

3.3 Linear Constraints

In [27, 29, 89], surface editing schemes that satisfy zero-dimensional constraints such as positions, tangents and normals, were presented. The constraints, being linear, are efficiently solved, allowing for the interactive manipulation of the freeform geometry. [89] also considered *transfinite constraints* where the constraints might have a non zero dimensionality. While some cases might be of a finite dimension, such as the containment of a polynomial curve in a polynomial surface when posed as a composition, other cases might necessitate an approximation. The composition of the polynomial curve $\gamma(t) = (u(t), v(t))$ and polynomial surface $S(u, v)$ yields $S(t) = S(u(t), v(t))$, a curve over S, which is a polynomial as well. The degree of $S(t)$ equals the product of the sum of the degrees of S and the degree of $\gamma(t)$. Hence, m linear constraints, where m is the order of $S(t)$, fully prescribe a polynomial curve over a polynomial surface. This result also extends to rationals.

Other finite linear constraints are treated with ease. A positional constraint, following [27, 29, 89, 28, 19], could be prescribed as, for curves,

$$P = C(t_p) = \sum_j Q_j B_{i,n}(t_p),$$

and for surfaces,

$$P = S(u_p, v_p) = \sum_{jk} Q_{jk} B_{j,n}(u_p) B_{k,m}(v_p).$$

Similarly, a normal constraint could be written as

$$0 = \langle N, C'(t_n) \rangle = \sum_j \langle N, Q_j \rangle \, B'_{i,n}(t_n),$$

for curves and the normal or tangent-plane constraint yields

$$0 = \left\langle N, \frac{\partial S(u_n, v_n)}{\partial u} \right\rangle = \sum_{jk} \langle N, Q_{jk} \rangle \, B'_{j,n}(u_n) B_{k,m}(v_n),$$

$$0 = \left\langle N, \frac{\partial S(u_n, v_n)}{\partial v} \right\rangle = \sum_{jk} \langle N, Q_{jk} \rangle \, B_{j,n}(u_n) B'_{k,m}(v_n).$$

For surfaces, the normal constraint is related to tangency constraints. The two partials of S, which span the tangent-plane if S is regular, also uniquely determine the orthogonal, normal space, of S. That is, $\frac{\partial S(u_n,v_n)}{\partial u} \times \frac{\partial S(u_n,v_n)}{\partial v} \neq 0$. Hence, the normal constraints as listed above could be similarly written as $C'(t_n) = T$ with one important difference. By coercing $C'(t_n) = T$, the length of the tangent field is exactly fixed, achieving C^1 continuity at this point. By posing the constraint as $\langle N, C'(t_n) \rangle$, G^1 continuity is gained, necessitating fewer degrees of freedom.

3.4 Bi-Linear and Non-Linear Constraints

The advantage of having linear constraints is obvious. The solution is much simplified and is typically more robust to compute. Several types of non-linear constraints could also be expanded and solved with little effort. The preservation of the area enclosed by a closed planar curve is important, for example, when one designs a cross-section of an airplane's fuselage that is assumed to hold a fixed volume. This area (and volume in R^3) constraint could be represented as a bi- (tri-) linear constraint [19, 47].

Consider again $C(t) = (x(t), y(t))$ being a regular, closed planar parametric curve. Employing Green's theorem, the (signed) area, $\mathcal{A}$, enclosed by $C(t)$, equals (see, for example [16, 31]),

$$\mathcal{A} = \frac{1}{2} \oint -x'(t)y(t) + x(t)y'(t)dt = \frac{1}{2} \oint |C(t) \times C'(t)|dt. \tag{2}$$

Equation (2) is clearly quadratic in t. Yet, in [19, 47], it is recognized that Equation (2) could be decomposed into a bi-linear form in t as $\mathcal{A} = x(t)My(t)$. This decomposition eases the incorporation of an area constraint into an MR framework. In [19], $x(t)$ and $y(t)$ are interleavingly fixed while solving the remainder of the linear constraint in $y(t)$ and $x(t)$, respectively. In Figure 7, a nose in an outline of a face is pulled without constraints, and then with positional, and positional and area constraints. This comparison shows how positional constraints could anchor the shape at certain points, and how the fixed area constraints have a global effect on the shape even for local changes. A local nose expansion automatically reacts by shrinking the entire shape, in order to keep the area constant.

In [47] another area preserving MR editing method for uniform B-splines was introduced. Herein, a wavelet-based MR analysis similar to [22] has been used. It enables in particular the derivation of a multiresolution representation of the area functional for the curve at any level of resolution. Let us briefly introduce this MR framework here, since it is different from the non-uniform multiresolution setting of Section 3.1. In this setting we are given some functional space E and some nested

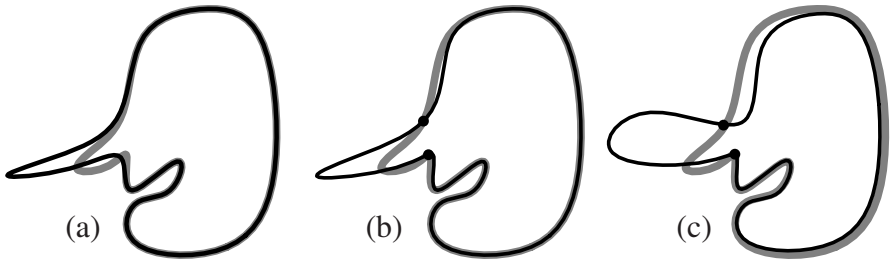

Fig. 7. Multiresolution editing with linear and bi-linear area constraints, before (wide gray) and after (thin black) the editing operation. In (a), the nose is interactively pulled to the left with no additional constraints. In (b), two positional constraints are placed at the base of the nose, while in (c), the area is fixed.

linear *approximation spaces* $S^j \subset E$ with $S^0 \subset S^1 \subset \cdots \subset S^n$. Since we are dealing with closed curves, these spaces have finite dimension. Let S^j be spanned by a set of basis functions $(\varphi_i^j)_i$, called *scaling functions*. A space W^j being the complement of S^j in S^{j+1} is called the *detail space*. Its basis functions $(\psi_i^j)_i$ are such that together with φ^j they form a basis of S^{j+1}. The functions ψ_i^j are called *wavelets*. The space S^n can, therefore, be decomposed as follows:

$$S^n = S^{n-1} \oplus W^{n-1} = S^{n-2} \oplus \bigoplus_{j=n-2}^{n-1} W^j = \cdots = S^0 \oplus \bigoplus_{j=0}^{n-1} W^j. \tag{3}$$

Condition (3) implies that the scaling functions are refinable; that is, for all $j \in \{0, \ldots, n\}$ there must exist some matrices P^j and Q^j such that the following refinement equations hold:

$$\begin{aligned} \varphi^{j-1} &= (P^j)^T \varphi^j, \\ \psi^{j-1} &= (Q^j)^T \varphi^j. \end{aligned} \tag{4}$$

On the other hand, the "fine" scaling functions φ^j can be constructed from the coarser scaling functions and wavelets with the aid of some matrices A^j and B^j:

$$\varphi^j = (A^j)^T \varphi^{j-1} + (B^j)^T \psi^{j-1}. \tag{5}$$

Note that $[P^j \mid Q^j]$ and $\begin{bmatrix} A^j \\ B^j \end{bmatrix}$ are both square matrices, and that

$$\begin{bmatrix} P^j \; Q^j \end{bmatrix} \begin{bmatrix} A^j \\ \\ B^j \end{bmatrix} = I. \tag{6}$$

The choice of the scaling functions determines the structure of the matrices P^j, Q^j, A^j, and B^j . Sparse matrices are desirable for most of the applications.

Based on this framework let an *MR curve* $C(t)$, $t \in I$ be defined as an element of some functional space S^n. Let $D2^n$ be the dimension of the space S^n, with $D = dim(S^0)$. $C(t)$ can be written as:

$$C(t) = \sum_{i=0}^{D2^n-1} x_i^n \varphi_i^n = (\mathbf{x}^n)^T(\varphi^n), \tag{7}$$

where $\mathbf{x}^n$ is a column vector of $D2^n$ scaling coefficients $x_0^n, \ldots, x_{D2^n-1}^n \in \mathbb{R}^2$. In a more general setting, the vector of scaling coefficients $\mathbf{x}^n$ can be thought of as a discrete signal with $D2^n$ samples.

The relations (4) and (5) now allow us to create a low-resolution signal $\mathbf{x}^{n-1}$ with less samples by using the low pass filter matrix A^n:

$$\mathbf{x}^{n-1} = A^n \mathbf{x}^n,$$

where A^n is of size $(D2^{n-1} \times D2^n)$. The details, which have been lost in this filtering process, can be captured as another signal $\mathbf{d}^{n-1}$ using the high-pass filter matrix B^n:

$$\mathbf{d}^{n-1} = B^n \mathbf{x}^n,$$

where B^n is also a $(D2^{n-1} \times D2^n)$ matrix. The process of splitting a signal $\mathbf{x}^n$ into a coarser signal $\mathbf{x}^{n-1}$ with details $\mathbf{d}^{n-1}$ can now be repeated recursively with the new signal $\mathbf{x}^{n-1}$. Finally, the original signal will be decomposed into a low-resolution signal $\mathbf{x}^0$ and details $\mathbf{d}^0, \ldots, \mathbf{d}^{n-1}$. This recursive process is known as a **filter bank** [64]; see Figure 8. At any intermediate level of resolution, L, the original signal is decomposed into a coarser signal $\mathbf{x}^L$ and details $\mathbf{d}^L, \ldots, \mathbf{d}^{n-1}$. Note that the vector $\mathbf{d}^j$ is of size $D2^j$, $j = 0, \ldots, n-1$.

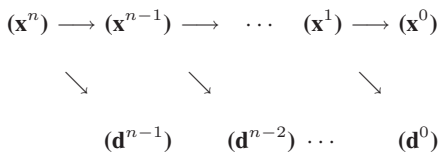

Fig. 8. A filter bank

Using the matrices P^j and Q^j the original signal $\mathbf{x}^n$ can be reconstructed recursively with

$$\mathbf{x}^j = P^j \mathbf{x}^{j-1} + Q^j \mathbf{d}^{j-1}, \qquad \text{for } j = 1, \ldots, n.$$

Since the filter bank also applies to the scaling coefficients of an MR curve (7), such a curve can be represented at any *level of resolution*, $L \in \{0, \ldots, n\}$, by some coarse coefficients $\mathbf{x}^L$ that form approximations of the initial coefficients and by the detail coefficients $\mathbf{d}^L, \ldots, \mathbf{d}^{n-1}$ as follows:

$$C(t) = (\mathbf{x}^L)^T(\varphi^L) + (\mathbf{d}^L)^T(\psi^L) + \cdots + (\mathbf{d}^{n-1})^T(\psi^{n-1}), \quad L = 0, \ldots, n. \tag{8}$$

In this wavelet-based MR framework, the area functional (2) of an MR curve (8) can now be evaluated at any level of resolution L. This leads to the bi-linear equation

$$2\mathcal{A} = (\mathbf{X}^L) \left[M^L \right] (\mathbf{Y}^L)^T, \quad \forall L \in \{0, \ldots, n\}, \tag{9}$$

where X^L and Y^L are the line vectors of the x- and y-coordinates, respectively, of all $D2^n$ coefficients (coarse and wavelet coefficients) of the MR representation of the curve, i.e.,

$$\begin{pmatrix} \mathbf{X}^L \\ \mathbf{Y}^L \end{pmatrix} = (\mathbf{x}^L, \mathbf{d}^L, \mathbf{d}^{L+1}, \ldots, \mathbf{d}^{n-1}),$$

and

$$M^L = \begin{bmatrix} I(\varphi^L, \varphi^L) & I(\varphi^L, \psi^l)_{l=L}^{n-1} \\ I(\psi^k, \varphi^L)_{k=L}^{n-1} & I(\psi^k, \psi^l)_{k,l=L}^{n-1} \end{bmatrix} = \begin{bmatrix} A & B \\ -B^T & C \end{bmatrix}. \tag{10}$$

The MR area constraint is then linearized during the optimization process in order to locally or globally deform a curve at any level of resolution while preserving the enclosed area; see Figure 9.

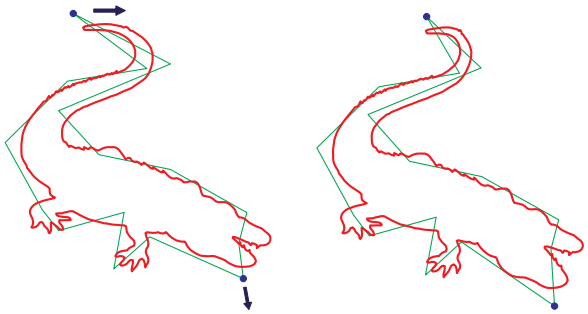

Fig. 9. Multiresolution editing of a coarse level of resolution with preservation of the enclosed area and a positional constraint.

Some works that preserve **volume** while manipulating the shape are also available. In [74], a cuboid volume was manipulated while preserving its volume, handling the problem as a non-linear optimization problem. In [19], it was also shown that the volume constraint, which is cubic in general, could also be posed as a trilinear constraint. Volume-preserving editing of MR surfaces represented by wavelets for uniform tensor-product B-splines following the MR setting described above has been developed in [77]. An example of volume-preserving MR editing is shown in Figure 10.

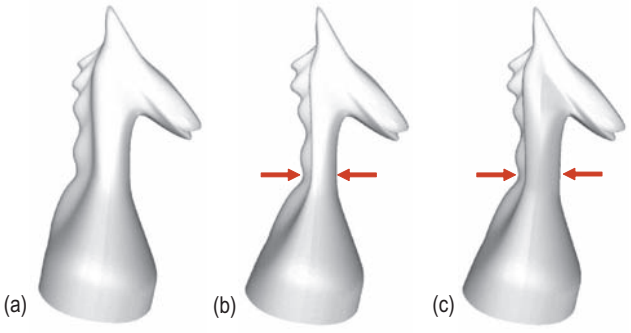

Fig. 10. An example of multiresolution editing with volume-preservation. (a) shows the original tensor-product B-spline surface. In (b), a deformation is applied at a coarse level of resolution without volume-preservation. In (c), the same deformation is applied but the volume of the original surface is preserved.

Other non-linear constraints that are commonly considered are second-order differential constraints such as convexity [51], and first- and second-order fairing constraints, typically in the form of strain and stress surface shape optimization functionals [89]. Another non-linear constraint of high interest is the preservation of the arc-length of the curve. In [78], the arc-length of a curve was presented as a non-linear constraint that is preserved during the curve's manipulation. Herein the constraint is integrated into an MR editing system that allows intuitive control of the deformation's extent and aspect. In [79] this length-constrained MR deformation has been integrated in a wrinkling tool for soft tissue modeling.

The exploitation of first and second differential order constraints, in real-time, is also highly intensive computationally. In [73], an interactive surface editing system that supports real-time surface manipulation with convexity/developability constraints was reported. It exploits a careful symbolic pre-computation of the curvature fields.

3.5 Intrinsic Multiresolution Decomposition of Freeform Shapes

The fundamental problem of MR decomposition is that the decomposition is typically not intrinsic. A curve or a surface could be arbitrary closely approximated using different knot sequences and even different control points. Likewise, two similarly looking objects could be represented using completely different polygonal meshes, as it is evident by the vast remeshing results that have been published in recent years.

It is, therefore, plausible to try and execute this MR decomposition in a way that is independent of the representation underneath, taking into account only the intrinsic geometry, and ignoring, for example, the parameterization.

One such possibility with regard to a planar C^2 freeform curve is to represent the shape by its curvature signature, $\kappa(t)$:

$$\kappa(t) = \frac{x'(t)y''(t) - x''(t)y'(t)}{\left(x'^2(t) + y'^2(t)\right)^{3/2}}, \tag{11}$$

assuming $C(t)$ is regular or $||C'(t)|| \neq 0$.

Assume $C(s)$ is an arc-length parameterized curve. Then, $\kappa(s) = x'(s)y''(s) - x''(s)y'(s)$. Further,

$$C'(s) = T(s), \quad C''(s) = T'(s) = \kappa(s)N(s),$$

where $T(s)$ and $N(s)$ are the unit tangent and normal fields of $C(s)$. $T(s) = (x'(s), y'(s))$ is a unit size vector and hence is always on the unit circle. Let θ be the angle between $T(s)$ and the x-axis,

$$\theta(s) = tan^{-1}\left(\frac{y'(s)}{x'(s)}\right),$$

and consider $\theta'(s)$,

$$\begin{aligned}
\theta'(s) &= \left(tan^{-1}\left(\frac{y'(s)}{x'(s)}\right)\right)' \\
&= \frac{1}{1+\left(\frac{y'(s)}{x'(s)}\right)^2}\left(\frac{y'(s)}{x'(s)}\right)' \\
&= \frac{x'^2(s)}{x'^2(s)+y'^2(s)}\frac{x'(s)y''(s)-x''(s)y'(s)}{x'^2(s)} \\
&= \frac{x'(s)y''(s)-x''(s)y'(s)}{x'^2(s)+y'^2(s)} \\
&= x'(s)y''(s)-x''(s)y'(s).
\end{aligned}$$

In other words, $\theta'(s) = \kappa(s)$ or a curve $C(s)$ could be reconstructed from $\kappa(s)$ by (see also [10])

$$C(s) = \int_0^s T(\bar{s})d\bar{s} = \int_0^s Circ\left(\int_0^{\bar{s}} \kappa(\tilde{s})d\tilde{s}\right)d\bar{s},$$

up to a rigid-motion, where $Circ(\cdot)$ is an arc-length parameterized unit circle. We are now able to switch back and forth between a regular parametric form of a planar curve $C(t)$ and its curvature signature $\kappa(t)$, up to rigid-motion.

While polynomial parametric curves are not arc-length, in general, one can approximate a given polynomial parametric curve as an arc-length polynomial parametric curve to an arbitrary precision; see, for example [18]. Figure 11 shows one example of a curvature signature computed to an approximation of an arc-length polynomial curve.

Multiresolution decomposition could now be applied to $\kappa(s)$ instead of $C(s)$. Alternatively, details could be added to low-resolution shapes by modulating the base $\kappa(s)$ signature and reconstructing the curve. Practical attempts of this procedure turned out to be quite slow and a large number of $\kappa(s)$ samples were necessary to achieve a reasonable reconstruction.

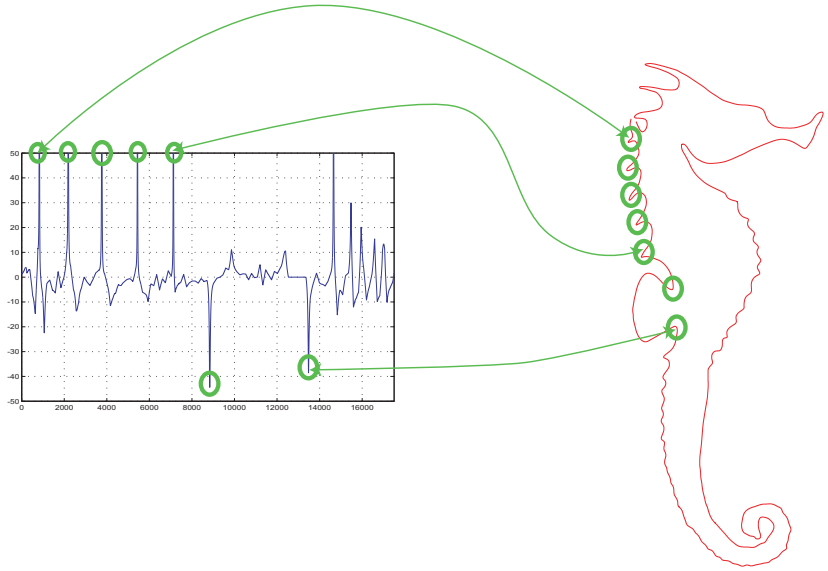

Fig. 11. A (portion of a) curvature signature (left) of an approximately arc-length parameterized curve (right).

An actual intrinsic MR decomposition of a freeform curve using its curvature signature is presented in Figure 12. A multiresolution analysis of $\kappa(t)$ was performed using Haar wavelets. New curvature functions κ_{small}, κ_{mean}, κ_{large} were computed by partially reconstructing the wavelet decomposition using only detail coefficients greater than a given (small, mean, large) threshold. The curves were then obtained by the integration of the new curvature functions.

Further research in this direction of intrinsic MR decomposition of freeform geometry is in order. One such research direction should seek an ability to preserve the continuity of closed, periodic curves throughout the intrinsic MR process.

3.6 Multiresolution Morphing

Morphing (or metamorphosis) is known as the smooth and progressive transformation of one shape into another. The shape can be an image or a planar curve in 2D space, or it can be a surface or a volume in 3D space. The problem is to create an aesthetic and intuitive transition between two shapes. The intermediate shapes should preserve the appearance and the properties of the input shapes. A trivial linear interpolation is often not appropriate, since the intermediate shapes tend to vary a lot in their volume or they lose the proportions of their shape features. Another negative effect is that the geometric details can disappear and re-appear later during the transition. Good results are generally achieved not by interpolating the positions of the boundary representation but by interpolating elements of alternative representations. In the case of 2D polygonal shapes, Sederberg et al. [81] represented polygons by a

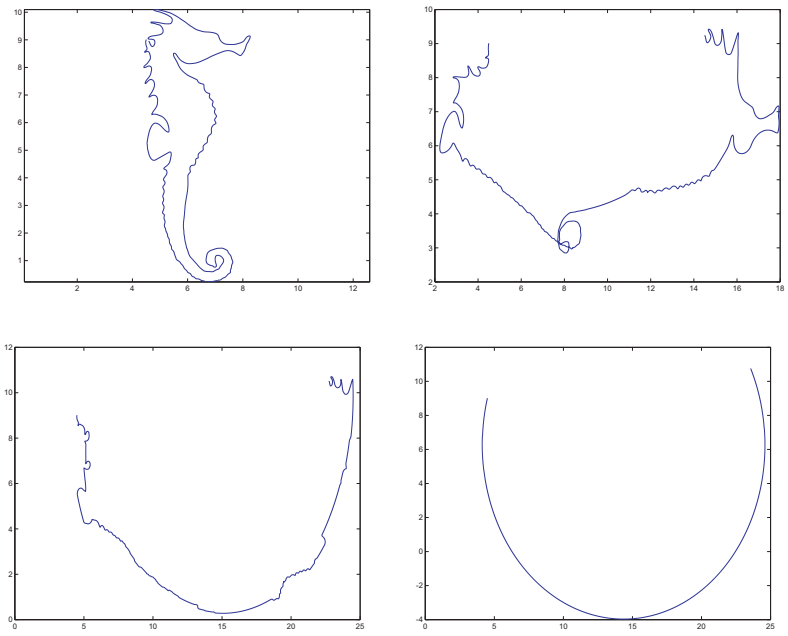

Fig. 12. Multiresolution analysis and partial synthesis of a seahorse curve based on an intrinsic curvature signature.

set of lengths and angles. Shapira and Rappoport [82] used a star-skeleton representation. Goldstein and Gotsman [30] used an MR representation based on curve evolution. Alexa et al. [1] morphed compatible triangulations by locally least-distorting maps. There is also the MR mesh morphing technique by Lee et al. [59]. The key to a successful method thus seems to be the use of a representation based on intrinsic properties of the object geometry such that interpolation of its elements achieves automatically pleasing morphs.

The morphing method we report on in this section is based on a new intrinsic MR representation. It decomposes the source and the target shapes into a coarse approximation and a set of detail coefficients. It computes separately the sequence of coarse intermediate shapes and details, and then reconstructs all intermediate shapes. The choice of the MR representation is crucial for the quality of the resulting shapes. For example, a wavelet-based MR analysis would not preserve the orientation of the details during deformation. In fact, the details here are encoded in a global coordinate system. An MR representation that encodes the details using local frames similar to [22] solves this problem, but the coarse representation of the curve in [22] is still not intrinsic.

In [46] a curvature-based MR representation for 2D polygonal curves was introduced. This MR representation is based on an intrinsic parameterization of both the coarse shape and the detail coefficients. All coefficients will be represented intrinsi-

cally by lengths and angles. Similar to local frames, the MR representation preserves the orientation of the details during deformation.

Let $P_i = (x_i, y_i)$, $i = 0, \ldots, N-1$ denote the vertices of a polygon to be MR-analyzed. The initial polygon needs to be transformed from (x, y)-coordinates into so-called (θ, l)-coordinates, where $\theta_i = \angle(\overline{P_{i-1}P_i}, \overline{P_iP_{i+1}})$ is the counterclockwise angle of the two consecutive polygon segments at P_i and $l_i = \|\overline{P_iP_{i+1}}\|$, $i = 0, \ldots, N-2$. The (x, y)-coordinates of the control points P_i can be recovered directly using, for example, P_0 as an anchor point and $\overline{P_0P_1}$ as an anchor line (determining the translation and rotation, rigid-motion, degrees of freedom). Note the (θ, l)-coordinates are rigid-motion invariant.

Following the filter bank algorithm presented in Figure 8, an *angle-length MR representation* can be computed as follows. From a polygon with 2^{n+1} segments, one analysis step creates a polygon with 2^n segments and 2^n detail coefficients, which are represented by two-dimensional vectors of the form:

$$\begin{array}{ccc} (\theta^{n+1}, \mathbf{l}^{n+1}) & \rightarrow & (\theta^n, \mathbf{l}^n) \\ & \searrow & \\ & & (\alpha^n, \beta^n), \end{array}$$

where $\theta^n = (\theta_0^n, \ldots, \theta_{2^n-1}^n)$ and is analogous for $\mathbf{l}^n, \alpha^n, \beta^n$.

The MR analysis is the recursive procedure of splitting the vector of coefficients of a polygon (θ^{n+1}, l^{n+1}) into a vector of coarse coefficients of a lower resolution polygon (θ^n, l^n) and a vector of detail coefficients; see Figure 13. Let the upper index n denote the level of resolution. Both coarse shape and detail coefficients of level n must be computed directly from (θ^{n+1}, l^{n+1}) and vice-versa. The coarse shape and detail coefficients are computed using the basic cosine trigonometric rule for triangles (also known as the Al-Kashi formula for triangles); see [46] for more details.

Given two polygons with the same number of corresponding control points P_S and P_T, called source and target polygons, the MR morphing algorithm constructs in-between polygons P_t that gradually change P_S into P_T for $t \in [0, 1]$, where $P_S = P_0$ and $P_T = P_1$. For both polygons one disposes of two sets of coarse coefficients and of two sets of detail coefficients. In principle, the in-between morphs are now generated by interpolating the coefficients of this intrinsic MR representation. However, the coarse coefficients are interpolated using the locally least distorting interpolation [1] in order to get "as-rigid-as-possible" intermediate morphs. Figure 14 shows two examples of curves with a complex shape and with a lot of fine details that are difficult to interpolate with standard morphing techniques.

3.7 Variational Multiresolution Methods for Freeform Surface

The variational modeling paradigm is used in order to find the "best" curve or surface amongst all solutions that meet some constraints. The constraints may result from the particular modeling technique used, for example, sample point approximation, or direct curve manipulation (see Section 3.2), or they may be one of the constraints

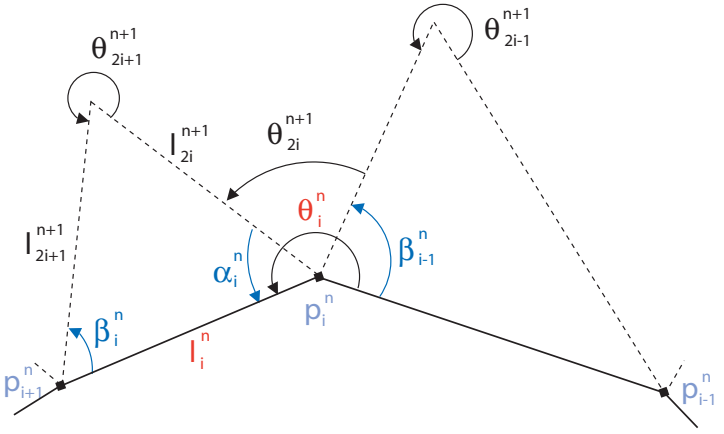

Fig. 13. The analysis. The dotted polygon belongs to resolution level $n+1$, the fat polygon belongs to level n.

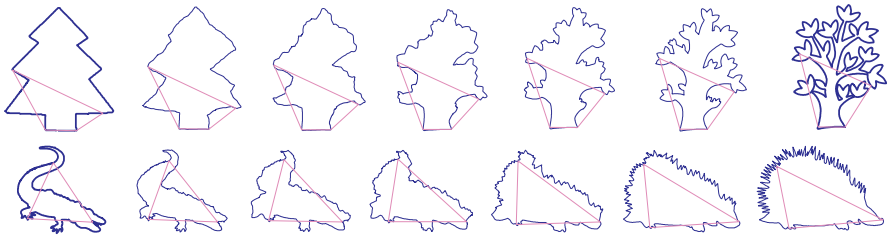

Fig. 14. Two examples of intrinsic multiresolution morphing. The tree curves have 2048 control points. The animal curves have 1536 control points. The algorithm applies a multiresolution analysis to the source (left) and target (right) polygons. Then, the coarse shape and detail coefficients are independently interpolated, and the intermediate curves are reconstructed. The resulting curves are shown alongside the coarse polygons.

described in Section 3.3. In the context of smooth curve and surface design, the notion of "best" is formulated by minimizing some energy functional.

Variational Shape Design

Although it is difficult to define exactly, in mathematical terms, what *fairness* of a curve or surface is, it is commonly accepted that smooth and graceful shapes are obtained by minimizing the amount of energy stored in the surface. The energy functionals originating from elasticity theory, such as the bending energy for curves $\int \kappa^2(t)dt$ or the thin-plate energy for surfaces $\int \kappa_1^2+\kappa_2^2 dA$, are in general non-linear. These and other higher order, non-linear, energy functionals were used in [67, 37].

In order to accelerate computations, linearized versions of these energy functionals are generally used; see, for example, [11, 12, 89, 33]

$$\mathcal{E} = \int_{\sigma} (\alpha \text{ stretch} + \beta \text{ bend}) d\sigma,$$

where α and β are weights on the stretching and bending energies. These produce a surface that tends to minimize its area to avoid folding and to distribute curvature over large regions in order to result in fair shapes. The stretch-and-bend functionals are typically approximated via the following quadratic terms: $\alpha_{11}X_u^2 + \alpha_{12}X_uX_v + \alpha_{22}X_v^2$ and $\beta_{11}X_{uu}^2 + \beta_{12}X_{uv}^2 + \beta_{22}X_{vv}^2$, respectively, only to be linearized in the optimization process.

Historically, the use of such energy functionals goes back to early spline and CAGD literature [65, 76] and has today led to a research area called variational design (of smooth curves and surfaces) [21, 43, 42, 7, 44].

Variational Multiresolution Modeling

Gortler and Cohen [33] showed how the variational constraint, which generalizes least-squares, can be solved through an MR formulation of a planar curve. A wavelet-based MR curve satisfying some linear constraints and minimizing a linearized bending energy functional may be found by solving the following linear system [89]:

$$\begin{bmatrix} \bar{H} & \bar{A}^T \\ \bar{A} & 0 \end{bmatrix} \begin{bmatrix} \bar{\mathbf{x}} \\ \lambda \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{b} \end{bmatrix},$$

where $\bar{A}$ is the constraint matrix, $\bar{H}$ is the Hessian matrix of the basis functions, and λ is the vector of Lagrange multipliers. The bars signify that the variables are wavelet coefficients. The wavelets allow acceleration of the iterative conjugate gradient-solving of the variational problem.

Variational subdivision is another modeling technique where constraints are combined with classical subdivision. Instead of applying explicit rules for the new vertices, Kobbelt's variational subdivision scheme [56] computes the new vertices such that a fairness functional is minimized. At each step a linear system has to be solved. The resulting curves have minimal total curvature. Furthermore, [58] showed how wavelets can be constructed by using the Lifting Scheme [85], which is appropriate for variational subdivision curves. Weimer and Warren [86, 87, 88] developed variational subdivision schemes that satisfy partial differential equations for, for instance, fluid or thin-plate equations.

4 Multiresolution Analysis for Irregular Mesh-based Representations

A lot of work has been done in the past ten years on MR analysis of models based on a decomposition of the shape into triangles. This section will focus on two types, scalar data defined on triangulations and mesh-based freeform surfaces.

Only certain types of data sets can be analyzed by wavelet MR analysis. One principal restriction is that the grid on which the data is defined has to be obtained

by successive subdivisions of a coarse grid. These subdivisions define a sequence of grids such that the cells of one grid are subdivided by the cells of the next grid. Such a sequence of grids is deemed "connected by subdivision". This restriction is due to the fact that wavelet analysis needs a *nested* sequence of approximating spaces; see Section 3.1. In the case of quadrilateral or triangular grids, the regular four-way split, as illustrated in Figure 15, is generally used to create a grid with subdivision connectivity, since the grids will not tend to degenerate after several subdivisions.

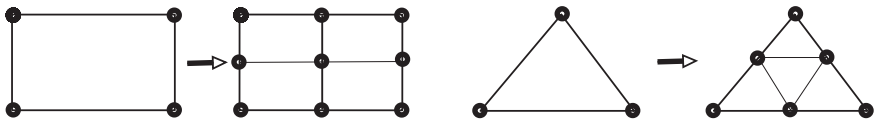

Fig. 15. Regular four-way split for a quad mesh and a triangle mesh.

However, data defined on triangulations as well as freeform surface meshes are generally of more complex structure due to acquisition techniques such as observation and laser range scanning. Thus, classical wavelet theory cannot be adapted directly to these so-called *irregular meshes*, since it is impossible to associate a sequence of grids with subdivision connectivity to this data. In the case of quadrilateral grids with subdivision connectivity, the one-dimensional wavelet-based MR analysis applies directly by tensor-product [84, 66]. If the data is defined on a regular triangular grid, classical wavelet theory can, and has also been adapted as well [68, 80]. Nevertheless, the case of freeform surfaces is more complicated since surfaces of arbitrary topology cannot be parameterized on regular quad or triangle meshes.

The aim of the present section is to focus on MR analysis for irregular mesh-based representations. Section 4.1 addresses the simplification of numerical data attached to an irregular mesh. Section 4.2 covers the simplification of surface meshes.

4.1 Irregular Triangulations

Wavelet methods assume that the mesh on which the data is defined can be reached by recursive subdivision of a basic mesh. Thus, every wavelet-based scheme is associated with hierarchies that have a tree structure (where every parent node is subdivided into a set of child nodes). Wavelet volume visualization [39] is related to Octree structures. Wavelet radiosity [34] and wavelets over triangulated domains [61, 68, 80] are based on Quadtree structures.

On the other hand, irregular triangular meshes cannot be reached by subdivision rules, therefore hierarchical structures that have been developed to handle them are more complicated than trees. These include, for example, hierarchical Delaunay triangulations [55, 23], or progressive meshes [48, 49]. These data structures are appropriate for LOD models, see [24], but not for MR analysis as described in the present chapter.

The approach introduced in [3] fills the gap between wavelet methods (on subdivision hierarchies) and hierarchical structures on irregular triangular meshes for a special type of data set, i.e., for piecewise constant data defined on irregular planar or spherical triangulations.

The basic idea is to relax the restrictions imposed by classical wavelet-based MR analysis, while preserving good properties such as constant memory requirements, linear computational time, and the ability to accurately approximate data with few detail coefficients. The relaxed restrictions are related to the approximation spaces associated with the MR analysis. These spaces are the functional spaces that correspond to each level of resolution where the original function is successively approximated during the analysis. These spaces have to be nested, i.e., the space corresponding to one resolution has to be a subspace of all spaces corresponding to finer resolutions. This property of nested spaces is the reason why the grids of data have be to connected by subdivision.

The generalized framework of MR analysis for irregular triangulation introduced in [3] does not require the nested property. The latter is replaced by a weaker condition that is related to the growth of the approximation spaces. If the data is defined on irregular triangulations, it becomes possible to associate them with a sequence of approximation spaces corresponding to coarser irregular triangular grids. There exist numerous algorithms for reducing the number of triangles in a mesh—independently of the data that is defined on this mesh. Delaunay-removal can be applied to planar or convex triangular meshes, and edge-collapse to general triangular meshes. If the mesh comes from the recursive four-way split of some triangles in a base mesh, then the obvious way to simplify it, is to replace each group of four sub-triangles by their parent triangle. The common setting of these decimation algorithms is that a set of n triangles is replaced by a set of m triangles covering the same domain, with $m < n$, as is shown in Figure 16.

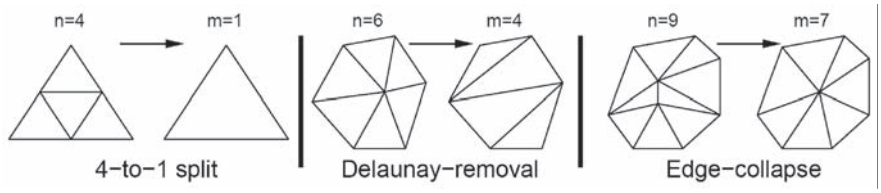

Fig. 16. Local triangle decimation.

A sequence of triangulations obtained by successive decimation is generally not nested because it is not connected by subdivision. However, the generalized framework allows the development of MR analysis algorithms that generate a coarse approximation of the original data and a set of detail coefficients. Therefore, the same types of applications that are mentioned in Section 3.1 are possible by selecting a subset of the detail coefficients and by synthesizing the data set using only the se-

lected coefficients. This idea has been used for different types of data sets defined on irregular triangulations in [8, 4, 5, 6]. Below, we describe the basic principle.

Let T denote a triangle of the domain and s a data value (scalar) defined on a triangle. The superscript f (fine) denotes quantities before the local decimation algorithm, and c (coarse), after the decimation. Bold letters denote vectors. The pair $(\mathbf{T}, \mathbf{s})$ denotes the piecewise constant function equal to s_i on the triangles T_i. If $\mathbf{Q}$ is a matrix, then $\mathbf{Q}_k$ denotes the k-th column vector of $\mathbf{Q}$.

Let us focus here on the following setting: we are given a piecewise constant function $(\mathbf{T}^f, \mathbf{s}^f)$ on n triangles, and a set of m triangles $\mathbf{T}^c$ covering the same domain on the surface, with $m < n$.

The essence of an MR analysis is the filter bank algorithm [64] explained in Section 3.4. It consists of the local decomposition and reconstruction algorithm. In the present irregular setting this algorithm is the same as in wavelet theory: a function living in a fine space (in our case, the piecewise constant function on the finer triangulation) is decomposed into a coarser approximating function (piecewise constant on the coarser triangulation) and error functions (piecewise constant on the finer triangulation). These error functions have two properties: they can be used to recover the original data exactly, and their norm is a measure of the error between the input function and the approximation.

Intuitively, we have to define one smoothing operator that maps the input function onto its approximation, and one error operator that captures the difference between the input function and its approximation. These two operators are defined by two rectangular matrices $\mathbf{A}$ and $\mathbf{B}$ of size $m \times n$ and $(n - m) \times n$, respectively:

$$\mathbf{s}^c = \mathbf{A}\mathbf{s}^f \tag{12}$$

$$\mathbf{d} = \mathbf{B}\mathbf{s}^f \tag{13}$$

The smoothing operator (12) computes the coarser coefficients $\mathbf{s}^c$ from the finer coefficients $\mathbf{s}^f$, and the error operator (13) computes the detail coefficients $\mathbf{d}$. The actual computation of the so-called analysis matrices $\mathbf{A}$ and $\mathbf{B}$ is detailed in [3]. One step of the filter bank algorithm can thus be illustrated as follows:

$$\begin{array}{ccc} \mathbf{s}^f & \longrightarrow & \mathbf{s}^c \\ & \searrow & \\ & & \mathbf{d} \end{array}$$

In order to keep a constant memory size for the data values, the original coefficients $\mathbf{s}^f$ are cleared from memory after the decimation, and replaced by the coarse and detailed coefficients $\mathbf{s}^c$ and $\mathbf{d}$. Of course, the sum of the sizes of $\mathbf{s}^c$ and $\mathbf{d}$ equals the size of $\mathbf{s}^f$. Since $\mathbf{s}^f$ is cleared from memory, the decomposition formulas (12) and (13) have to be invertible, in order to be able to recover the original data values. This is the purpose of the reconstruction formula:

$$\mathbf{s}^f = \mathbf{P}\mathbf{s}^c + \mathbf{Q}\mathbf{d}. \tag{14}$$

The so-called synthesis matrices $\mathbf{P}$ and $\mathbf{Q}$ are of sizes $n \times m$ and $n \times (n - m)$, respectively. Intuitively, the operator $\mathbf{P}$ is the inverse of the smoothing operator $\mathbf{A}$:

$\mathbf{P}$ acts as a subdivision operator, although subdivision is not possible if the triangular domains are non-nested. The operator $\mathbf{Q}$ adds the details $\mathbf{d}$ to the oversampled data $\mathbf{P}\mathbf{s}^c$, in order to recover the original data $\mathbf{s}^f$. The matrices $\mathbf{P}$ and $\mathbf{Q}$ can be computed from $\mathbf{A}$ and $\mathbf{B}$ by:

$$(\mathbf{P}\ \mathbf{Q}) = \begin{pmatrix} \mathbf{A} \\ \mathbf{B} \end{pmatrix}^{-1}.$$

To be more precise about the properties of our decomposition, let us rewrite the reconstruction formula (14) with a functional point-of-view instead of a coefficient point-of-view:

$$(\mathbf{T}^f, \mathbf{s}^f) = (\mathbf{T}^f, \mathbf{P}\mathbf{s}^c) + \sum_{k=1}^{n-m} d_k(\mathbf{T}^f, \mathbf{Q}_k), \tag{15}$$

where $\mathbf{Q}_k$ denotes the k-th column vector of $\mathbf{Q}$, and d_k denotes the k-th detail coefficient of $\mathbf{d}$.

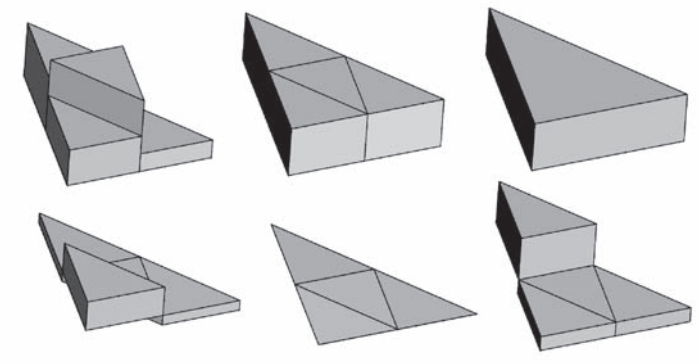

Fig. 17. Local decomposition by 4-to-1 split: finer, intermediate and coarser approximations on top, detail coefficients times wavelet functions on bottom. In this case, the coarse and fine triangular domains are nested, and therefore, the intermediate approximation equals the coarser approximation. The relative high magnitudes of the detail coefficients (bottom part) show the large L^2 error between the fine and coarse approximations.

Figs. 17 and 18 illustrate the local decomposition on two examples. In both figures, the top part shows, from left to right, the finer function $(\mathbf{T}^f, \mathbf{s}^f)$, the intermediate function $(\mathbf{T}^f, \mathbf{P}\mathbf{s}^c)$, and the final coarser function $(\mathbf{T}^c, \mathbf{s}^c)$. The bottom part shows the detail coefficients times the wavelet functions: $\mathbf{d}(\mathbf{T}^f, \mathbf{Q}^f)$. Figure 17 shows the local decomposition on a 4-to-1 split example. This leads to a traditional Haar wavelet decomposition for irregular triangular meshes. Note in this case that the intermediate function $(\mathbf{T}^f, \mathbf{P}\mathbf{s}^c)$ (top-middle), although defined over a finer mesh, equals the coarser function $(\mathbf{T}^c, \mathbf{s}^c)$ (top-right). In Figure 18, the block results from the removal of one interior vertex. Therefore, two detail coefficients are computed (bottom part).

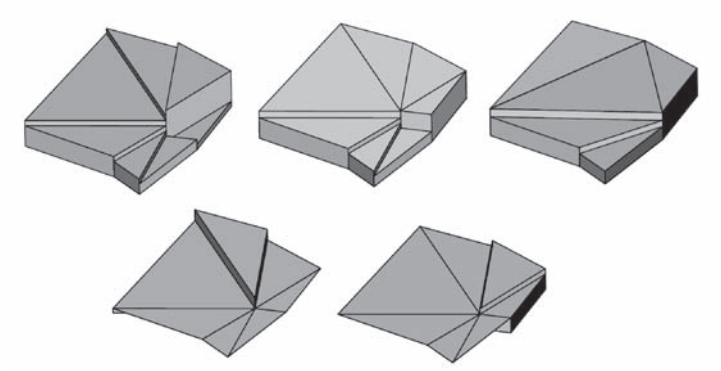

Fig. 18. Local decomposition by Delaunay-removal: finer, intermediate and coarser approximations on top, detail coefficients times wavelet functions on bottom. Since the triangular domains are non-nested, the intermediate approximation differs from the coarse approximation. The relative low magnitude of the detail coefficients (bottom part) shows the small L^2 error between the fine and coarse approximations.

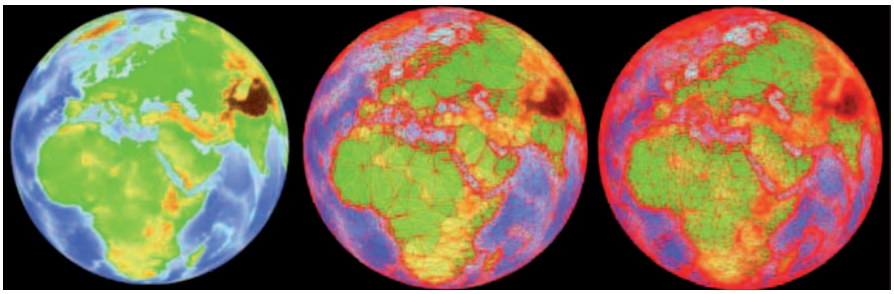

Fig. 19. Original data set with 1.3 M faces. Middle: partial reconstruction with 60000 faces. Right: partial reconstruction with 150000 faces.

Combining successive removal of sets of non-adjacent vertices with the filter bank algorithm leads to the construction of a hierarchy of triangulations that corresponds to an MR analysis of the initial data set. Threshold reconstruction is one of the possible applications of such a wavelet-based MR analysis. It consists of applying the reconstruction formula (15) only to the blocks of the hierarchy whose wavelet coefficients are greater than a fixed threshold. The visual effect is illustrated in Figure 19. The original data set (ETOP05 data set) in Figure 19 (left) has been fully analyzed with the preservation of the coastlines. Figures 19 (middle and right) show two partial threshold reconstructions. See also Figures CP-3 in Appendix C.

4.2 Surface Meshes

Developing MR modeling methods for large manifold surface meshes has been the subject of a great many research papers. Motivated by the ever-increasing size of polygonal meshes resulting from laser range scanners, researchers have tried to generalize the theory of wavelet MR analysis, successfully applied to 2D images, in order to compress, efficiently render, transmit, and *edit* such large meshes. One innovative application of MR methods for surface meshes is the ability to perform editing operations at different resolutions, as illustrated in Section 3.1 for curves. However, these surface meshes are generally irregular meshes, i.e., they do not have subdivision connectivity. Different approaches exist in the literature to define an MR analysis for manifold surface meshes.

For example, one of the early papers [17] first computes an approximation of the original data, a two-dimensional manifold triangular mesh, using new data defined on a grid with subdivision connectivity. The new data can then be analyzed with a wavelet analysis.

The groundbreaking work in the area of MR mesh representation was done by [61]. This paper also proposed a theory close to wavelet MR analysis. Its MR surface mesh model is closely related to subdivision surfaces. In general, in every MR analysis/synthesis scheme, the synthesis (or reconstruction) process can be seen as the combination of a subdivision step with a correction step. Based on this observation, [61] built an MR analysis/synthesis scheme on top of well known surface subdivision schemes, including Loop and Butterfly subdivision schemes (see the Chapter on Subdivision surfaces and applications in this volume). It turns out that only interpolating subdivision schemes lead to a linear time analysis process, while the synthesis process can always be performed in linear time. In order to build an MR analysis/synthesis scheme on top of a surface subdivision scheme, Lounsberry et al. [61] introduced a scalar product for functionals defined on the surface domain, and used this scalar product in order to define wavelet functions with good approximating properties. Based on this MR scheme, a fine mesh with subdivision connectivity can be represented on a wavelet basis. Thus, compression of the mesh can be performed by neglecting small wavelet coefficients. Progressive transmission is efficiently implemented by sorting the wavelet coefficients and transmitting them, starting with the most significant. The progressive transmission and its application to MR mesh viewing is the topic of [13].

This pioneering work was followed by [91] and [57]. These two papers did not rely on a genuine wavelet decomposition of the meshes. Rather, they mimicked the analysis process of the wavelet MR representation, by using a smoothing procedure to convert fine meshes into coarser meshes, and by encoding the error occurring during this smoothing procedure. [91] introduced highly adaptive procedures, with the aim of being able to edit large meshes in real-time. While [91] was still restricted to meshes with subdivision connectivity, [57] proposed a generalization to arbitrary meshes.

The idea of remeshing the irregular surface mesh into a semi-regular mesh (see Figure 20) with subdivision connectivity [60, 41] before computing wavelet-

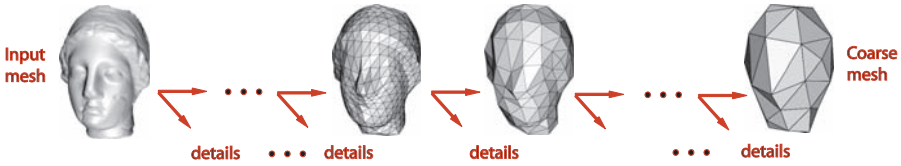

Fig. 20. Multiresolution analysis of a semi-regular mesh.

based MR analysis has also been used for signal processing applications, such as coding and compression of surface meshes. The compression allows a compact storage or a fast transmission of these surface data in a bandwidth-limited application. Wavelets are now frequently exploited to perform efficient compression. Based on MR analysis, wavelet coders do not only achieve better compression rates [61, 54, 40, 53, 69, 70] than methods based on signal quantization, but also make the progressive transmission and adaptive display easier.

5 Conclusions and Open Issues

Multiresolution representations play an increasing role in geometric design. Their use for both polygonal meshes and freeform polynomial and rational representations is expected to increase as their usefulness is further recognized. Nevertheless, many issues are still open and need to be resolved before the full power of this representation can be revealed.

Starting with MR representation for B-spline curves, both uniform and non-uniform knot spacings are, by now, fully supported and understood. Yet, the computation of the B-wavelet basis functions for spaces with non-uniform knot sequences is expensive and methods should be sought to reduce this cost.

Another related problem is the question of the inherent imprecisions of interactive MR editing of freeform curves. Being imprecise, it is difficult to employ in precise engineering design. In order to improve the precision, linear constraints are already embedded with interactive MR editing as well as a few non-linear constraints such as area and arc-length. The efficiency in solving linear constraints makes them attractive but also limited. Other families of non-linear constraints should be embedded with MR as well. Examples include curvature prescriptions, fairing requirements, higher order moments, etc.

Features are, many times, viewed as high frequency details of the geometry. A smooth shape of some animation could be combined with high frequency details of hair, thorns, or just scales. However, the simple algebraic sum of the two shapes would yield a result that is not necessarily the most appealing one. This is due to the fact that the details, when added algebraically, are not oriented along the smooth shape's geometry. Different MR decomposition schemes, which are intrinsically geometric and not algebraic, might be able to resolve such problems. In this sense,

intrinsic MR decomposition schemes, such as the presented curvature signatures, should be further explored.

References

1. M. Alexa, D. Cohen-Or, and D. Levin. As-rigid-as-possible shape interpolation. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 157–164, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.
2. R. Bartels and J. Beatty. A technique for the direct manipulation of spline curves. *Graphics Interface '89*, pages 33–39, 1989.
3. G.-P. Bonneau. Multiresolution analysis with non-nested spaces. *Computing Supplementum*, 13:51–66, 1996.
4. G.-P. Bonneau. Multiresolution analysis on irregular surface meshes. *IEEE Transactions on Graphics and Visualization*, 4(4):365–378, 1998.
5. G.-P. Bonneau and A. Gerussi. Hierarchical decomposition of datasets on irregular surface meshes. In *Proceedings of CGI'98,*, pages 59–63, Hannover, Germany, June 1998.
6. G.-P. Bonneau and A. Gerussi. Level of detail visualization of scalar data sets on irregular surface meshes. In *Proceedings Visualization'98,*, pages 73–77. IEEE, 1998.
7. G. P. Bonneau and H. Hagen. Variational design of rational bézier curves and surfaces. In L. Laurent and L. Schumaker, editors, *Curves and Surfaces*, volume II, pages 51–58. AK Peters, 1994.
8. G.-P. Bonneau, S. Hahmann, and G.M. Nielson. Blac-wavelets: a multiresolution analysis with non-nested spaces. In *Proceedings Visualization'96,*, pages 43–48. IEEE, 1996.
9. P. Borel and A. Rappoport. Simple constrained deformations for geometric modeling and interactive design. *ACM Transactions on Graphics*, 13(2):137–155, 1994.
10. M. P. Do Carmo. *Differential Geometry of Curves and Surfaces.* Cambridge University Press, 1976.
11. G. Celniker and D. Gossard. Deformable curve and surface finite-elements for free-form shape design. In *ACM SIGGRAPH Conference Proceedings*, pages 257–266. ACM, 1991.
12. G. Celniker and W. Welch. Linear constraints for deformable b-spline surfaces. In *Symposium on Interactive 3D Graphics*, pages 165–170, 1992.
13. A. Certain, J. Popovic, T. DeRose, T. Duchamp, D. Salesin, and W. Stuetzle. Interactive multiresolution surface viewing. *Computer Graphics*, 30(Annual Conference Series):91–98, 1996.
14. C. K. Chui and E. G. Quak. Wavelets on a bounded interval. In D. Braess and L. Schumaker, editors, *Numerical Methods of Approximation Theory, Volume 9*, pages 53–75. Birkhäuser Verlag, Basel, 1992.
15. P. Cignoni, C. Montani, E. Puppo, and R. Scopigno. Multiresolution representation and visualization of volume data. *IEEE Trans. on Visualization and Comp. Graph*, 3(4):352–369, 1997.
16. D. Eberly and J. Lancaster. On gray scale image measurements: I. arc length and area. *CVGIP: Graphical Models and Image Processing*, 53(6):538–549, 1991.
17. M. Eck, T. DeRose, T. Duchamp, H. Hoppe, T. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. In *Computer Graphics Proceedings (SIGGRAPH 95)*, pages 173–182, 1995.
18. G. Elber. Symbolic and numeric computation in curve interrogation. *Computer Graphics forum*, 14(1):25–34, March 1995.

19. G. Elber. Multiresolution curve editing with linear constraints. *The Journal of Computing & Information Science in Engineering*, 1(4):347–355, December 2001.
20. G. Elber and C. Gotsman. Multiresolution control for nonuniform bspline curve editing. In *The Third Pacific Graphics Conference on Computer Graphics and Applications, Seoul, Korea*, pages 267–278, [August 1995.
21. G. Farin, G. Rein, N. Sapidis, and A. J. Worsey. Fairing cubic b-spline curves. *Computer Aided Geometric Design*, 4:91–103, 1987.
22. A. Finkelstein and D. H. Salesin. Multiresolution curves. *Computer Graphics Proceedings (SIGGRAPH 94)*, pages 261–268, 1994.
23. L. De Floriani. A pyramidal data structure for triangle-based surface description. *IEEE Computer Graphics and Applications*, 9(2):67–78, 1989.
24. L. De Floriani, M. Alexa, Marie-Paule Cani, Paolo Cignoni, Emanuele Danovaro, Thomas Di Giacomo, HyungSeok Kim, Nadia Magnenat-Thalmann, and Enrico Puppo. Level-of-detail shape modeling. In *chapter 5 of this book*. Springer, 2005.
25. D. Forsey and R. Bartels. Hierarchical b-spline refinement. *Proceedings of SIGGRAPH'88, ACM New York*, pages 205–212, 1988.
26. B. Fowler. Geometric manipulation of tensor product surfaces. In *1992 Symposium on Interactive 3D Graphics*, pages 101–108, 1992.
27. B. Fowler. Geometric manipulation of tensor product surfaces. In *Proceedings of the 1992 symposium on Interactive 3D graphics*, pages 101–108. ACM Press, 1992.
28. B. Fowler and R. Bartels. Constraint-based curve manipulation. *IEEE Computer Graphics and Applications*, 13(5):43–49, 1993.
29. M. Gleicher. Integrating constraints and direct manipulation. In *Proceedings of the 1992 symposium on Interactive 3D graphics*, pages 171–174. ACM Press, 1992.
30. E. Goldstein and Craig Gotsman. Polygon morphing using a multiresolution representation. In *Graphics Interface '95*, pages 247–254. Canadian Inf. Process. Soc., 1995.
31. C. Gonzalez-Ochoa, S. Mccammon, and J. Peters. Computing moments of objects enclosed by piecewise polynomial surfaces. *ACM Transaction on Graphics*, 17(3):143–157, July 1998.
32. C. Gonzalez-Ochoa and J. Peters. Localized-hierarchy surface splines (less). In *Proceedings of the 1999 symposium on Interactive 3D graphics*, pages 7–15. ACM Press, 1999.
33. S. Gortler and M. Cohen. Hierarchical and variational geometric modeling with wavelets. In *1995 Symposium on 3D Interactive Graphics*, pages 35–41, 1995.
34. S. Gortler, P. Schröder, M. Cohen, and P. Hanrahan. Wavelet radiosity. *Computer Graphics Proceedings (SIGGRAPH 93)*, pages 221–230, 1993.
35. S. J. Gortler. Private communications.
36. S. J. Gortler. Wavelet methods in computer graphics. *PhD thesis, Department of Computer Science, Princeton*, 1994.
37. G. Greiner. Variational design and fairing of spline surfaces. In *Proc. Eurographics 1994*, pages 143–154, 1994.
38. G. Greiner and J. Loos. Data dependent thin plate energy and its use in interactive surface modeling. *Eurographics '96 (1996)*, 15:176–185, 1996.
39. M. Gross, L. Lippert, R. Dietrich, and S. Häring. Two methods or wavelet-based volume rendering. *Computers & Graphics*, 21(2):237–252, 1997.
40. I. Guskov, A. Khodakovsky, P. Schröder, and W. Sweldens. Hybrid meshes: Multiresolution using regular and irregular refinement. In *Proceedings of SoCG 2002*, 2000.
41. I. Guskov, K. Vidimce, W. Sweldens, and P. Schröder. Normal meshes. In Kurt Akeley, editor, *Siggraph 2000, Computer Graphics Proceedings*, pages 95–102. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.

42. H. Hagen and P. Santarelli. Variational design of smooth b-spline surfaces. In H. Hagen, editor, *Topics in Geometric Modeling*, pages 85–94. SIAM Philadelphia, 1992.
43. H. Hagen and G. Schulze. Automatics smoothing with geometric surface patches. *Computer Aided Geometric Design*, pages 231–236, 1987.
44. S. Hahmann. Shape improvement of surfaces. *Computing Suppl.*, 13:135–152, 1998.
45. S. Hahmann and G.-P. Bonneau. Polynomial surfaces interpolating arbitrary triangulations. *IEEE Transactions on Visualization and Computer Graphics*, 9(1):99–109, 2003.
46. S. Hahmann, G.-P. Bonneau, B. Caramiaux, and M. Cornillac. Multiresolution morphing of planar curves. *Computing*, 2007. to appear.
47. S. Hahmann, B. Sauvage, and G.-P. Bonneau. Area preserving deformation of multiresolution curves. *Computer Aided Geometric Design*, 22(4):349–367, 2005.
48. H. Hoppe. Progressive meshes. *Computer Graphics Proceedings (SIGGRAPH 96)*, pages 99–108, 1996.
49. H. Hoppe. View-dependent refinement of progressive meshes. *Computer Graphics Proceedings (SIGGRAPH 97)*, pages 189–198, 1997.
50. W. M. Hsu, J. F. Hughes, and H. Kaufman. Direct manipulation of free-form deformations. In *Computer Graphics (SIGGRAPH 92 Proceedings)*, pages 177–184, July 1992.
51. P. D. Kaklis and N. S. Sapidis. Convexity-preserving interpolatory parametric splines of nonuniform polynomial degree. *Comput. Aided Geom. Des.*, 12(1):1–26, 1995.
52. R. Kazinnik and G. Elber. Orthogonal decomposition of non-uniform bspline spaces using wavelets. *Computer Graphics forum*, 16(3):27–38, September 1997.
53. A. Khodakovsky and I. Guskov. Compression of normal meshes, 2003.
54. A. Khodakovsky, P. Schröder, and W. Sweldens. Progressive geometry compression. In Kurt Akeley, editor, *Siggraph 2000, Computer Graphics Proceedings*, pages 271–278. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.
55. D. Kirkpatrick. Optimal search in planar subdivisions. *SIAM Journal on Computing*, 12(1):28–35, 1983.
56. L. Kobbelt. A variational approach to subdivision. *Computer Aided Geometric Design*, 13:743–761, 1996.
57. L. Kobbelt, S. Campagna, J. Vorsatz, and HP. Seidel HP. Interactive multiresolution modeling on arbitrary meshes. In *Computer Graphics Proceedings (SIGGRAPH 98)*, pages 105–114, 1998.
58. L. Kobbelt and P. Schröder. A multiresolution framework for variational subdivision. *ACM Trans. on Graph.*, 17(4):209–237, 1998.
59. A. W. F. Lee, D. Dobkin, W. Sweldens, and P. Schröder. Multiresolution mesh morphing. *Computer Graphics Proceedings (SIGGRAPH 99)*, pages 343–350, 1999.
60. Aaron W. F. Lee, Wim Sweldens, Peter Schröder, Lawrence Cowsar, and David Dobkin. MAPS: Multiresolution adaptive parameterization of surfaces. *Computer Graphics*, 32(Annual Conference Series):95–104, 1998.
61. M. Lounsbery, T. De Rose, and J. Warren. Multiresolution analysis for surfaces of arbitrary topological type. *ACM Transaction on Graphics*, 16(1):34–73, 1997.
62. T. Lyche and K. Morken. Spline wavelets of minimal support. In D. Braess and L. Schumaker, editors, *Numerical Methods of Approximation Theory*, pages 177–194. Birkhäuser Verlag, Basel, 1992.
63. Tom Lyche and Knut Morken. Knot removal for parametric b-spline curves and surfaces. *Comput. Aided Geom. Des.*, 4(3):217–230, 1987.
64. S. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:674–693, 1989.
65. E. Mehlum. Non-linear spline. In R. E. Barnhill and R. F.R iesenfeld, editors, *Computer Aided Geometric Design*, pages 173–208. Academic Press, 1974.

66. M.Gross, O. Staadt, and R. Gatti. Efficient triangular surface approximations using wavelets and quadtree data structures. *IEEE Transactions on Visualization and Computer Graphics*, 2(2):130–143, 1996.
67. H. P. Moreton and C. H. Séquin. Functional optimisation for fair surface design. *Computer Graphics*, 26(2):167–176, 1992.
68. G. Nielson, IH. Jung, and J. Sung. Haar-wavelets over triangular domains with applications to multiresolution models for flow over a sphere. In *IEEE Visualization'97*, pages 143–150, november 1997.
69. F. Payan and M. Antonini. An efficient bit allocation for compressing normal meshes with an error-driven quantization. *Computer Aided Geometric Design*, 22:466–486, July 2005.
70. F. Payan and M. Antonini. Mean square error approximation for wavelet-based semiregular mesh compression. *IEEE Transactions on on Visualization and Computer Graphics (TVCG)*, 2006. to appear.
71. J. P. Pernot, S. Guillet, J. C. Leon, F. Giannini, B. Falcidieno B., and E. Catalano. A shape deformation tool to model character lines in the early design phases. In *Proceedings Shape Modeling International 2002, Banff, Canada*, 2002.
72. J. Peters. C1-surface splines. *SIAM J. Numer. Anal.*, 32(2):645–666, 1995.
73. M. Plavnik and G. Elber. urface design using global constraints on total curvature. In *The VIII IMA Conference on Mathematics of Surfaces*, September 1998.
74. A. Rappoport, A. Sheffer, and M. Bercovier. Volume-preserving free-form solids. In *Proceedings of Solid Modeling 95*, pages 361–372, May 1995.
75. A. Raviv and G. Elber. Three dimensional freeform sculpting via zero sets of scalar trivariate functions. *CAD*, 32(8/9):513–526, July/August 2000.
76. C. H. Reinsch. Smoothing by spline functions ii. *Num. Math.*, 16:451–454, 1967.
77. B. Sauvage. *Déformation de courbes et surfaces multirésolution sous contraintes*. Phd, Institut National Polytechnique de Grenoble (INPG), December 2005.
78. B. Sauvage, S. Hahmann, and G.-P. Bonneau. Length preserving multiresolution editing of curves. *Computing*, 72:161–170, 2004.
79. B. Sauvage, S. Hahmann, and G.-P. Bonneau. Length constrained multiresolution deformation for surface wrinkling. In *International Conference on Shape Modeling and Applications, SMI'06*, pages 113–121, Matsushima, June 2006. IEEE Computer Society Press.
80. P. Schröder and W. Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. *Computer Graphics Proceedings (SIGGRAPH 95)*, pages 161–172, 1995.
81. T.W. Sederberg, P. Gao, G. Wang, and H. Mu. 2-d shape blending: An intrinsic solution to the vertex path problem. *Computer Graphics,(SIGGRAPH 93 Proceedings)*, 27:15–18, 1993.
82. M. Shapira and A. Rappoport. Shape blending using the star-skeleton representation. *IEEE Comput. Graph. Appl.*, 15(2):44–50, 1995.
83. E. Stollnitz, T. DeRose, and D. Salesin. *Wavelets for Computer Graphics: Theory and Applications*. Morgan-Kaufmann, 1996.
84. Eric J. Stollnitz, Tony D. DeRose, and David H. Salesin. Wavelets for computer graphics: A primer, part 2. *IEEE Computer Graphics and Applications*, 15(4):75–85, 1995.
85. W. Sweldens. The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.*, 29(2):511–546, 1997.
86. J. Warren and H. Weimer. Variational subdivision for natural cubic splines. *Approximation Theory IX*, 2:345–352, 1998.
87. H. Weimer and J. Warren. Subdivision schemes for thin plate splines. *Computer Graphics Forum (Proceedings of Eurographics 98)*, pages 303–313, 1998.

88. H. Weimer and J. Warren. Subdivision schemes for fluid flow. *Computer Graphics (SIGGRAPH 99 Conference Proceedings)*, pages 111–120, August 1999.
89. W. Welch and A. Witkin. Variational surface modeling. *Computer Graphics (SIGGRAPH '92 proceedings)*, 26:157–166, July 1992.
90. A. Yvart, S. Hahmann, and G.-P. Bonneau. Hierarchical triangular splines. *ACM Transactions on Graphics*, 24(4):1374–1391, 2005.
91. D. Zorin, P. Schröder, and W. Sweldens. Interactive multiresolution mesh editing. *Computer Graphics Proceedings (SIGGRAPH 97)*, pages 259–268, 1997.