

Optimization Techniques for Approximation with Subdivision Surfaces

M. Marinov and L. Kobbelt

Computer Graphics Group, RWTH Aachen, Germany

Abstract

We present a method for scattered data approximation with subdivision surfaces which actually uses the true representation of the limit surface as a linear combination of smooth basis functions associated with the control vertices. This is unlike previous techniques which used only piecewise linear approximations of the limit surface. By this we can assign arbitrary parameterizations to the given sample points, including those generated by parameter correction. We present a robust and fast algorithm for exact closest point search on Loop surfaces by combining Newton iteration and non-linear minimization. Based on this we perform unconditionally convergent parameter correction to optimize the approximation with respect to the L^2 metric and thus we make a well-established scattered data fitting technique which has been available before only for B-spline surfaces, applicable to subdivision surfaces. Further we exploit the fact that the control mesh of a subdivision surface can have arbitrary connectivity to reduce the L^∞ error up to a certain user-defined tolerance by adaptively restructuring the control mesh. By employing iterative least squares solvers, we achieve acceptable running times even for large amounts of data and we obtain high quality approximations by surfaces with relatively low control mesh complexity compared to the number of sample points. Since we are using plain subdivision surfaces, there is no need for multiresolution detail coefficients and we do not have to deal with the additional overhead in data and computational complexity associated with them.

Categories and Subject Descriptors (according to ACM CCS):

I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations

1. Introduction

Scattered data approximation methods are a key technology for shape reconstruction and reverse engineering from measured geometry data. In a typical application scenario, raw data is generated, e.g., by some 3D scanning device and fitting a globally smooth surface to the set of sample points converts this data into a geometric representation of the original object that enables sophisticated downstream applications like, e.g., free-form shape editing. Most of the work in this area has been done based on classical tensor-product spline surfaces but with the availability of more flexible subdivision surfaces many ideas have been extended to this generalized setting during the last years. Instead of being constrained to rectangular patches, subdivision surfaces can represent globally smooth surfaces of arbitrary (manifold) topology by allowing for arbitrary irregular control meshes. We are using Loop subdivision surfaces [Loo87] in this paper, but the basic concepts could be transferred to other types of subdivision, e.g., Catmull-Clark subdivision surfaces [CC78].

The majority of the well-established scattered data approximation techniques focuses on the minimization of some form of the L^2 error. The main reason for this is that least squares problems are easy and efficiently handled by solving a simple linear system.

However, from the application point of view, L^∞ type errors are much more relevant since the user usually prescribes some maximum tolerance δ by which the fitted surface is allowed to deviate from the given data. Because the L^2 metric is computed by some integral over the surface, one often wastes effort (and degrees of freedom) when globally improving the approximation even if the maximum tolerance is only violated in some local region.

The motivation for the work presented in this paper is the observation that due to the flexibility of subdivision surfaces with respect to their control mesh structure, we can apply and iterate many different operations to progressively improve the approximation of the given data. If we just update the positions of the control vertices we can do least squares fitting just like for spline surfaces. However, in addition we can change the structure of the control mesh by locally inserting or removing control vertices. This allows us to efficiently reduce the L^∞ error by adaptively introducing new degrees of freedom (i.e. control vertices) in regions where the maximum tolerance is exceeded and by removing degrees of freedom in regions where the surface fitting problem is under-determined due to sparse sample data.

The specific contributions of this paper are that we present a complete scattered data fitting method for subdivision surfaces that




Figure 1: Original Iphigénie model (left): scanned data 351750 points, Loop surface approximation (center): maximum deviation is 0.05% of the bounding box diagonal, control mesh defining the surface (right): 15347 control points.

uses the true subdivision basis functions instead of some piecewise linear approximation. This is made possible by using Stam’s exact evaluation procedure [Sta98, Sta99] to set up the least squares system. Since we can also evaluate partial derivatives of a subdivision surface exactly, we propose a robust algorithm which finds the parameter value of the closest point on the approximating surface to a given sample. Thus we generalize the technique of parameter correction [Die95] from spline surfaces to subdivision surfaces and perform unconditionally convergent optimization of the approximation with respect to the L^2 error (Section 2).

To reduce the L^∞ error below a user prescribed error tolerance δ , we present an iterative procedure in Section 3 that adaptively refines the control mesh according to the local approximation error or coarsens it if the local sample density is insufficient. In combination with mesh connectivity regularization we are able to produce high quality approximations without having to add a fairness functional. Our technique is progressive and scalable in the sense that we can get a coarse fit after just a few seconds while we can further improve the approximation quality by letting the algorithm perform some more iterations.

1.1. Related work

The amount of work that has been done in the area of surface approximation is immense and a complete review is beyond the scope of this paper. We refer to [Sap94] as a standard reference and to [CMS03] for some more recent advances. Traditionally, tensor-product spline surfaces have been used for this task, but when it comes to the approximation of complex geometric objects, their rigid regular structure makes it necessary to fit several patches to parts of the data and then to stitch them together in a geometrically smooth fashion [EH96].

Another problem is that the regular structure of tensor-product patches prohibits the flexible adaption of the control mesh to the local shape complexity or sample density. As a consequence penalty functionals (a.k.a. fairing functionals) usually have to be added to the approximation problem in order to stabilize it [Die95, EH96]. All these difficulties compromise the flexibility and approximation power of spline surfaces for general approximation problems.

Subdivision surfaces are globally smooth (mostly even piecewise polynomial) surfaces that do not suffer from these limitations. Complex shapes can be represented with one control mesh and local adaption of the control vertex density is straightforward [ZSS*00]. This is why several papers [HKD93, HDD*94, STK99, TSK*00, MZ00, LLS01a] have addressed the scattered data fitting problem by using subdivision surface representations. However, since subdivision surfaces have no obvious explicit parametrization, modifications and simplifications of the general setting have been used for the sake of efficiency.

One issue is that, assuming the canonical parametrization, subdivision surfaces are much easier to evaluate at dyadic barycentric parameter values than at arbitrary parameter values. This is why special uniform parameterizations of the given sample data have been preferred by many authors [STK99, LMH00, LLS01a, LLS01b, JK02, MMTP02]. Although this leads to well conditioned least squares systems and extremely simple quasi-interpolation operators [LLS01a, LLS01b], the major drawback of using uniform parameterizations is that the evaluated approximation errors differ significantly from the actual geometric deviation. In [LLS01b], Litke et al. use uniform parameterization in combination with re-sampling the target surface. While this leads to a geometrically meaningful error metric, it may affect the local sample density potentially leading to under-sampling in regions where the surface normals of the fitting surface and the target surface strongly differ or where the fitting surface has high curvature (Fig. 2).

Parameter correction is a technique that does exactly the opposite. For each given sample point, the closest point on the approximating surface is found which implies an obvious geometric interpretation of the approximation error (Fig. 2). However, whenever parameter correction has been used for subdivision surfaces in the literature, the correction has been computed with respect to a piecewise linear approximation instead of the true limit surface [HDD*94]. In practice this often leads to numerical instabilities - mostly due to the fact that the piecewise linear approximation does not have a continuous surface normal field.

From the conceptual point of view our work is closest to [HDD*94]. In this paper Hoppe et al. describe a procedure that does least squares fitting of a subdivision surface to scattered data,




Figure 2: Uniform parametrization (top) vs. parameter correction (bottom). The dots represent the approximated samples and the position of the squares indicate the corresponding parameter values. In the uniform setting we have to find the closest data sample for uniformly distributed “sensor” points on the approximating surface. In the parameter correction setting we find the closest point on the approximating surface for each sample leading to a more reliable geometric distance measure.

however the parametrization is computed and evaluated only on a piecewise linear approximation of the approximating surface. This procedure is interleaved with a “random descent” mesh optimization scheme, which iteratively attempts to reduce the L^2 error. Minimization of the L^∞ error is not discussed.

A quite different approach to subdivision surface fitting is described in [BKZ01, JK02, LMH00, LLS01a, LLS01b]. Here the so-called *multiresolution* subdivision surfaces [Zor97] are used which assign an additional displacement vector to every control vertex of an adaptively refined control mesh. While this representation provides a natural hierarchy that distinguishes different levels of detail, the mathematical representations becomes extremely redundant as can be seen from the number of detail coefficients that are necessary to closely approximate complicated objects [LLS01a].

1.2. Notation

With $S = \{s_1 \dots s_M\}$ we denote the set of given data samples and with D the subdivision surface which is fitted to S . Its control mesh, the base mesh $C_0(P_0, T_0)$, is composed of two sets - vertices and triangles. We denote $N = \#(P_0)$ as the number of degrees of freedom (control vertices) and we usually assume $M \gg N$. The uniform refinement levels of C_0 are $\{C_1 \dots C_k\}$ and the meshes $\hat{C}_j(\hat{P}_j, T_j)$ are obtained by applying the limit surface projection operator to the control vertices P_j of C_j .

2. Optimization of the L^2 metric

We examine the problem of finding the best L^2 approximation of a given set of samples $S = \{s_i\}$ by a Loop subdivision surface D with fixed number of control points and connectivity. Since S is a discrete set, the L^2 error is expressed by the following term:

$$L^2(\{s_i\}, D) = \left(\sum_{i=1}^M \|s_i - D(t_i)\|_2^2 \right)^{\frac{1}{2}}, \quad (1)$$

where $\{t_i\}$ are the parameter values assigned to the samples $\{s_i\}$ with respect to some parametrization of the surface D . The most common way is to use barycentric coordinates with respect to the triangles of the base mesh C_0 , i.e., $t_i = \langle f_i, (v_i, w_i) \rangle$, where $f_i \in T_0$ indicates the patch to which the sample s_i is mapped and $(1 - v_i - w_i, v_i, w_i)$ define the barycentric coordinates of t_i within this triangle.

Given a fixed correspondence $s_i \leftrightarrow t_i$, the problem of minimizing (1) is solved in the *least squares* sense by finding that solution P_0 which minimizes the L^2 residuum of the over-determined linear system

$$AP_0 = S. \quad (2)$$

In order to compute the matrix $A = [\phi_j(t_i)]_{M \times N}$ for arbitrary parametrizations $\{t_i\}$, we have to evaluate the basis functions $\{\phi_1 \dots \phi_N\}$ which define D at $\{t_1 \dots t_M\}$ (see 2.1). Solving the system (2) gives us the optimal position of the control vertices P_0 of D . The sparsity of the matrix A depends on the support of the basis functions ϕ_j . In the case of Loop subdivision surfaces, each patch (corresponding to one triangle of the base mesh) is affected by 12 control vertices on average. Hence the matrix A has about 12 non-zero coefficients per row. There are many different ways to efficiently solve (2). In our implementation we use an iterative method such as CGLS [Elf78].

Another way to minimize (1) is by performing parameter correction, i.e., by finding for every sample point s_i the parameter t_i of the closest point on D . Iterating least square fitting and parameter correction generates a sequence of solutions D^0, D^1, D^2, \dots , each of which has a smaller L^2 error, until the approximation quality cannot be improved any more, i.e.

$$L^2(\{s_i\}, D^k) - L^2(\{s_i\}, D^{k+1}) < \varepsilon,$$

or some other criterion is met, for example a maximum number of steps is performed.

2.1. Exact evaluation of the subdivision basis functions

A significant improvement of our method is due to the fact that the fitting procedure depends on the *exact* subdivision surface D instead of depending on some piecewise-linear approximation of it, as in [HDD*94]. To achieve this, we use the representation of D as a linear combination of one smooth basis function for each control vertex, i.e.

$$D(t) = \sum_{j=0}^N \phi_j(t) p_{0,j},$$

where the parameter domain of D is modeled as the faces of the base mesh C_0 . Unlike uniform tensor-product spline surfaces, the subdivision basis functions are not mere translates of each other.




Figure 3: From left to right: a) original model with 30696 sample points, b) initial approximation with 461 control points, L^2 error 82.9, c) after 25 parameter correction and re-fitting steps, L^2 error 42.6, d) and e) show color coded errors before and after the 25 optimization steps. Notice the concentration of the error to local “hot spots” which makes it easier to detect regions where the insertion of new control vertices effectively reduces the L^∞ error. f) convergence behavior (error vs. number of iterations) plot.

Instead, each ϕ_j depends on the valence of the corresponding control vertex $p_{0,j}$ and on the valences of its direct neighbors in C_0 . Hence, the easiest way to evaluate a basis function ϕ_j is to assign an additional scalar attribute σ to each control vertex and set $\sigma = 1$ for $p_{0,j}$ and $\sigma = 0$ for all other vertices in P_0 [BS02], and then to apply Stam’s evaluation procedure on the so-defined scalar-valued mesh.

2.2. Exact closest point search and parameter correction

The next ingredient of our L^2 error minimization procedure is an algorithm for finding the *exact* closest point $D(t_i)$ on a Loop subdivision surface D given an arbitrary sample point s_i . This is done by performing stabilized Newton iteration. The starting value for the iteration can be either a previously assigned parameter value t_i or a value obtained by searching for the closest point on a discrete piecewise linear approximation \widehat{C}_k . In the second case we use a variant of the MESH-framework for evaluating Hausdorff distances between surfaces [ASCE02], which employs a spatial data structure in order to minimize “closest point on triangle” evaluations. We denote the initial solution as $t_{i,0}$.

In the j -th step of the Newton iteration we linearize the surface D at the current approximate solution $t_{i,j}$ by computing the tangent plane $T_{i,j}$, which is given by the Jacobian $\nabla D(t_{i,j}) \in \mathbb{R}^{3 \times 2}$. In order to find an update vector $q_{i,j} \in \mathbb{R}^2$ in the parameter domain towards an improved estimate for the closest point $D(t_i)$, we determine the orthogonal projection of s_i onto $T_{i,j}$ by solving the following 2×2 linear system for $q_{i,j}$:

$$\begin{aligned} \{s_i - (D(t_{i,j}) + \nabla D(t_{i,j}) \cdot q_{i,j})\} \cdot \frac{\partial D}{\partial v}(t_{i,j}) &= 0 \\ \{s_i - (D(t_{i,j}) + \nabla D(t_{i,j}) \cdot q_{i,j})\} \cdot \frac{\partial D}{\partial w}(t_{i,j}) &= 0 \end{aligned}$$

Special care has to be taken when actually updating the parameter value $t_{i,j} = \langle f_{i,j}, (v_{i,j}, w_{i,j}) \rangle$ since the parameter domain of the surface D is split into disjoint triangles corresponding to the faces of the base mesh C_0 . In order to avoid an excessive number of special cases, we simply consider the following three cases:

1. When $(v_{i,j}, w_{i,j}) + q_{i,j}$ still lies in the same triangle $f_{i,j}$ then $t_{i,j+1} := \langle f_{i,j}, (v_{i,j}, w_{i,j}) + q_{i,j} \rangle$.
2. When $(v_{i,j}, w_{i,j}) + q_{i,j}$ lies outside $f_{i,j}$, i.e., the update moves into a neighboring patch, then we scale $q_{i,j}$ down by a factor $0 < \lambda < 1$ such that the updated parameter value $t_{i,j+1} := \langle f_{i,j}, (v_{i,j}, w_{i,j}) + \lambda q_{i,j} \rangle$ lies exactly on the boundary of the patch $f_{i,j}$. By this we avoid the re-parametrization that would be necessary to compute the proper barycentric coordinates of the vector $q_{i,j}$ in the next parameter triangle.
3. If $t_{i,j}$ lies already on an edge of $f_{i,j}$ and $q_{i,j}$ is pointing outside then we switch to the neighbor face $f_{i,j+1}$ into which $q_{i,j}$ points, i.e., $t_{i,j+1} := \langle f_{i,j+1}, (v_{i,j+1}, w_{i,j+1}) \rangle$ where $(v_{i,j+1}, w_{i,j+1})$ are the barycentric coordinates of the same common boundary point with respect to the new triangle $f_{i,j+1}$. The actual parameter update will be executed in the next iteration.

The distinction between case (2) and (3) is necessary because it is difficult to predict if the update vector $q_{i,j+1}$ in the next Newton iteration will point into the same direction as $q_{i,j}$ or in the opposite direction.

If the starting value $t_{i,0}$ is not sufficiently close to the exact solution, it might happen that the Newton iteration suggests a parameter value where the L^2 distance actually increases, i.e. $\|s_i - D(t_{i,j+1})\|_2 \geq \|s_i - D(t_{i,j})\|_2$. This usually means that the length of the update step $q_{i,j}$ is incorrect, which is a common behavior of any root-finding algorithm for multivariate functions. To handle such situations in a robust manner, we switch to a reliable univariate optimization technique like Brent minimization [Bre73] to find the minimum \hat{h} of the function $g(h) = \|s_i - D((v_{i,j}, w_{i,j}) + hq_{i,j})\|_2$, $h \in (0, 1)$. Finally we set $t_{i,j+1} = \langle f_{i,j}, (v_{i,j}, w_{i,j}) + \hat{h}q_{i,j} \rangle$.

We stop the closest point search whenever $\|q_{i,j}\|_2 < \epsilon$ or $j > n$. In our test cases we always used $\epsilon = 10^{-6}$ and $n = 500$. Because of the robust minimization-based backtracking, we observed even for very complicated models with more than 200000 sample points less than 0.01% failures to converge with respect to the tolerance ϵ in less than n update steps. If such a failure occurs, we compare the newly found solution at $t_{i,n}$ with the old solution at $t_{i,0}$ and

keep the better one. This guarantees the unconditionally stable and monotonic convergence of the parameter correction procedure. A case where $t_{i,0}$ is actually better than $t_{i,n}$ occurs extremely rarely. In most of our experiments the Newton iteration converges on average in less than 6 update steps.

It is important to notice the difference of our parameter correction scheme to previous approaches [HDD*94], where a piecewise linear approximation of the limit surface was used for the closest point search. An update step that reduces the distance between a sample point s_i and a piecewise linear approximation \hat{C}_k of the limit surface D does not necessarily reduce the distance between s_i and D . To illustrate this we repeated the experiment of Fig. 3 searching for the closest point on \hat{C}_2 , as proposed in [HDD*94]. The L^2 error after the 25 optimization steps was larger by 26.24%. Since the test control mesh was relatively coarse, we could repeat the experiment again, this time using \hat{C}_4 as a piecewise linear approximation of D . The exact solution was still better by 11.4%. Note that this approach is not applicable for the optimization of large control meshes (without the use of a sophisticated adaptive subdivision algorithm), since it requires us to subdivide the control mesh at every step to a 64 times larger mesh.

3. Optimization of the L^∞ metric

While the L^2 error is a good measure for globally fitting a surface to sample data, the L^∞ error is much closer to the intuitive notion of approximation tolerance. Hence we present a technique to effectively reduce the L^∞ error by changing the structure of the control mesh. The L^∞ error for discrete sample data $\{s_i\}$ is defined as:

$$L^\infty(\{s_i\}, D) = \max_{1 \leq i \leq M} \|s_i - D(t_i)\|_2. \quad (3)$$

In engineering applications this maximum tolerance $\delta \geq 0$ is usually set by the user. In [LMH00, BKZ01, LLS01a, LLS01b] *multiresolution* subdivision surfaces are used to satisfy such user defined tolerances. Faces for which the maximum error is exceeded are subdivided adaptively and displacement vectors are added to the newly inserted vertices. This approach is robust and leads to quite good results, has the advantage of a genuine multiresolution semantics and is convenient for applications such as progressive transmission and compression. However, it also has several drawbacks:

1. Looking at the results, e.g., in [LLS01a], it seems that the representation is highly redundant in terms of the number of coefficients that we need to faithfully represent complex objects. The representation is no longer unique, especially if multiple detail coefficients are assigned to the same control vertex on different refinement levels.
2. The conceptual simplicity of the subdivision surface is lost since we have to deal with a combination of basis functions from different refinement levels. The evaluation can become tricky in regions where the adaptive refinement level changes.

To compensate for these difficulties, we present an iterative method for optimizing the control mesh of a plain subdivision surface D such that the approximation satisfies a given L^∞ error tolerance $\delta \geq 0$. We derive different heuristics to improve both the approximation and the surface quality. The method is based on a set of fundamental operations:

1. Adaptive insertion of control vertices where the tolerance is not met (Section 3.1).
2. Removal of control vertices in under-sampled regions (Section 3.2).
3. Re-establishing the parameterization $\{t_i\}$ of the samples $\{s_i\}$ after the control mesh (and hence the surface D) has changed. This is done by re-running the parameter correction algorithm of Section 2.2.
4. Connectivity regularization (Section 3.3).

The overall optimization procedure is described by the following pseudo-code:

```

Adaptively insert control vertices
do
    Regularize connectivity
    Re-establish the correspondence
    Remove control vertices in
        under-sampled regions
while removed_vertices_number > 0
Fit the new control mesh to the samples

```

The presented technique does not guarantee the achievement of the criterion $L^\infty(\{s_i\}, D) < \delta$ in one single iteration, it is only a heuristic which identifies regions of C_0 which should be optimized with respect to the L^∞ error and the current parametrization $\{t_i\}$. In practice, one usually needs several iterations of the above procedure to satisfy the criterion. In Section 4 we interleave this iteration with the technique described in Section 2 to bound the growth of the number of control vertices.

3.1. Adaptive insertion of control vertices

With $Sf_i = \{s_i | t_i \in f_i\}$ we denote the set of samples mapped to $f_i \in T_0$. For every $f_i \in T_0$ we define

$$L^\infty(f_i) = \max_{s_k \in Sf_i} \|s_k - D(t_k)\|_2.$$

If $L^\infty(f_i) > \delta$, we have to split the face f_i and so locally add new degrees of freedom. We denote the set of all to-be-split faces by $G = \{f_i | L^\infty(f_i) > \delta\}$.

There are several common ways to split a face, e.g., longest edge-split, 1-to-3 split, or 1-to-4 split with crack-fixing. We empirically found that the best way to adaptively refine the mesh in terms of surface quality and approximation is to 1-to-4 split every face from G and then to fix the resulting cracks by bisecting neighboring faces. This way of adaptive refinement least affects the regularity of the mesh since all newly inserted vertices have valence 5 or 6 and only the crack-fixing changes the valence of some existing vertices by one. Other adaptive refinement operators tend to produce much more irregular, i.e., non-valence-6, vertices which has a negative effect on the quality of the resulting limit surface.

Let Q be the set of control vertices affected by the adaptive refinement of G (including the crack-fixing). These vertices have a natural one-to-one correspondence to certain control vertices from the mesh C_1 obtained by applying the uniform Loop subdivision operator to the given mesh C_0 . In the adaptively refined mesh, we assign to all control vertices from Q the vertex positions of the corresponding vertex in C_1 while the other control vertices keep their




Figure 4: Adaptive insertion of control vertices. The red squares correspond to the newly inserted control vertices when splitting the central face. Blue dots correspond to the original control vertices shifted to their position in C_1 . Dashed lines denote crack-fixing edges.

position from C_0 (Fig. 4). By this we minimize the instant modification of the surface D since the locally refined resolution of the control mesh is compensated by shifting the control vertices to their corresponding position on the next refinement level. The resulting change of the limit surfaces is significantly smaller compared to the adaptive refinement operator that simply inserts the new vertices at the midpoints of the edges. Notice, however, that we still treat the resulting control mesh as a single resolution plain subdivision surface.

3.2. Removal of control vertices in under-sampled regions

During the iteration of parameter correction, least squares fitting, and the adaptive insertion of new control vertices it can happen that some control vertices in C_0 become under-determined if too few samples (or more precisely, their associated parameter values) lie close enough to the center of the corresponding basis function's support. This leads to very unpleasant artifacts like ripples and bumps. The standard answer to these kind of instabilities, especially in the spline world, is to add a penalizing term (a fairing functional) to the minimization problem which, however, might affect the approximation quality in the properly sampled regions. Again by exploiting the flexibility of subdivision surfaces with respect to the connectivity of their control meshes, we propose an alternative approach that is based on detecting under-sampled vertices and removing them from the mesh.

From the mathematical point of view, under-determined control vertices degrade the condition number of the matrix A in the least squares equation (2) which makes the solution less robust. An appropriate value to rate the degree of being under-determined for a given vertex is the sum of the absolute values of the coefficients in the corresponding column of the matrix A . In the case of Loop subdivision, the basis functions and hence the matrix coefficients are all positive anyway.

The major drawback of this stability measure is that we need to evaluate it during the L^∞ optimization phase when no valid matrix A is available and hence we have to compute each of these coefficients by Stam's evaluation procedure. For efficiency reasons we therefore check if a control vertex is under-determined with a simplified criterion. The idea is to simply check if there are samples

present at all in some region around the center of a basis function's support.

Since we observed that the stability of the least squares system is less critical if the approximating surface already fits very well to the given sample data, we actually define two criteria. One to be used when the L^∞ error is above some threshold τ and one to be used when it is below.

For every control vertex p_i we define the *Voronoi region* V_i as the union of all Voronoi sectors corresponding to the adjacent (parameter) triangles [MDSB03] and the *one-ring region* W_i as the union of all adjacent (parameter) triangles. Both regions cover some inner part of the basis function ϕ_i 's support. Further we denote the set of samples that are associated with some parameter value in V_i or W_i by SV_i or SW_i respectively. Based on these definitions, we decide that a vertex p_i is under-determined iff the current L^∞ error is above the threshold τ and the set SV_i is empty or the current L^∞ error is below the threshold τ and the set SW_i is empty (Fig. 5).




Figure 5: Left: the Voronoi cell V_i of a control vertex p_i is defined as the union of all Voronoi sectors of p_i . Right: the one-ring region W_i is defined as the union of the control faces adjacent to p_i . The dots represent the parametrization of the samples in the domains defined by the control faces. Samples are associated to the set SV_i (resp. SW_i) if their parameter value is inside V_i (resp. W_i). Depending on the L^∞ error, a vertex p_i is classified as under-sampled if SV_i or SW_i is empty.

The threshold τ could be set by the user. However, in our experiments we found that the choice of τ is not very critical so we simply set it to 0.1% of the bounding box diagonal of S for all the models that we tested.

Once a vertex is classified as under-determined we remove it by collapsing that half-edge connected to it which minimizes the connectivity metric (4) of the control mesh after the collapse.

3.3. Connectivity regularization

We use the following common metric [SG03] for rating the regularity of the connectivity of a mesh M :

$$R(M) = \sum_{v \in M} (d(v) - d_{opt}(v))^2, \quad (4)$$

where $d(v)$ is the valence of the vertex v and $d_{opt}(v) = 4$ if v is a boundary vertex or $d_{opt}(v) = 6$ if v is a non-boundary vertex. An edge flip is called *good* if it decreases $R(M)$. We build a candidate set H of all good edge flips and perform a greedy optimization by




Figure 6: a) Original Cyberware scan of a male head with 320k triangles, $M = 160k$. From left to right: approximations produced by our method with relative tolerance 0.03% (as in [LLS01a]) for different values of the parameter J , \widehat{C}_0 is shown: b) $N = 16792$, $J = 1$, 8min, c) $N = 14365$, $J = 3$, 15min, d) $N = 13642$, $J = 5$, 22min, e) shows (d) shaded. The initial approximating surface has 1600 vertices.

choosing the best flip $h_i \in H$, i.e., the one that maximizes the improvement $R(C_0) - R(C'_0)$, in every step. After flipping we remove h_i and all edge flips affected by it from H and continue the greedy selection. Once H is empty we build a new set of candidates and check if there exist more good flips. The procedure stops once there are no more good edge flips in C_0 . Although this approach is not as sophisticated as the one in [SG03] it usually converges quickly to a local minimum of $R(C_0)$ and successfully prevents the occurrence of high or low valence vertices in the control mesh.

4. Overall approximation procedure

Finding a good balance between the optimization of the L^2 error (Section 2) and the optimization of the L^∞ error (Section 3) is one of the key issues for achieving a high-quality approximation and in our implementation this balance is determined by the user who has to select two parameters - K and J . Here, K is the maximum number of optimization steps, and every J -th step we perform optimization with respect to the L^∞ error, i.e., change the structure of C_0 . The user also prescribes the L^∞ error tolerance δ . The following pseudo-code implements the main approximation loop:

```

while k < K and  $L_k^\infty > \delta$ 
  if ((k+1 mod J) == 0)
    Optimize the  $L^\infty$  error
  else
    Optimize the  $L^2$  error
  end
  k=k+1
end

```

As a rule of thumb, using relative large $J = 5, 6, \dots, 10$ is a good idea and often leads to control meshes with smaller complexity (Fig. 6) since it is generally worth investing effort in finding the best approximating surface with the current number of degrees of freedom, before trying to optimize the fit by adaptively inserting control vertices in the high error regions. One can also look at J as a parameter controlling the trade-off between mesh complexity and running time for a given tolerance.

4.1. Initial fitting surface

As in [TSK*00, MMTP02], given a polygonal mesh S we find the initial approximating subdivision surface D by decimating S using QEM-based mesh simplification [GH97] until the number of degrees of freedom, i.e., the number of control vertices of C_0 , reaches some predefined number. A necessary condition for C_0 is to have the same topology as S . The final quality of the approximation might vary depending on the number of C_0 control vertices, but values from 1% to 5% of M provide always very good results for dense S . One might expect that using relatively large initial N over 5% of M will produce better results, but this often leads only to unjustified waste of degrees of freedom which do not contribute to the quality of the approximation and the overall minimization of the error.

The connectivity information of S is not used at any other place throughout the approximation procedure, therefore one could use any other method for determining the initial surface. In the future we intend to examine alternative approaches to construct the initial approximating surface with the same topology as the sample set, which will allow us to perform approximation also of non-triangulated point sets.

5. Results

We tested our approximation method on several models (Table 1). The goal was to achieve high-quality approximation with L^∞ error not larger than 0.05% of the bounding box diagonal of the corresponding model. We also illustrate that we are able to quickly produce a relatively coarse fit with tolerance less than 1% and progressively improve the approximation by investing more time into the fitting procedure (Fig. 7).

We first compare our algorithm with the B-spline approximation method in [EH96]. The best approximation of the bunny model presented in that paper has a relative maximum deviation of 1.44%. The approximating surface consists of 153 patches. Taking the inter-patch G^1 smoothness conditions into account, we count on average 4 dofs (degrees of freedom) per bi-cubic patch. Note that the actual patches are defined by more control vertices, however most of them are used up to satisfy the C^0 continuity and the G^1




Figure 7: From left to right: a) the original Stanford bunny model. Different approximations (N , relative $L^\infty(S, D)$, time in min:sec): b) 612, 0.63%, 0:32, c) 913, 0.30%, 1:54 d) 4680, 0.12%, 2:36, e) 8440, 0.049%, 4:39.

Model	M	initial N	final N	$L^\infty(\%)$	time h:m
Fig. 1	352K	3518	15347	0.049	0:30
Fig. 6	160K	1600	13642	0.029	0:22
Fig. 7	37K	612	8440	0.048	0:05
Fig. 8	40K	804	4494	0.036	0:03
Fig. 9	546K	4093	17995	0.049	1:04
Fig. 10	31K	307	4698	0.048	0:05
Fig. 11	51K	2028	4733	0.049	0:06

Table 1: Results obtained by the procedure described in Section 4. The L^∞ error is given as a percentage of the bounding box diagonal. $J = 5$ in all of the experiments and the algorithm converged in less than the maximum allowed ($K = 100$) optimization steps. Timings are taken on 2.8GHz Pentium IV with 2Gb RAM.

smoothness constraints across the patch boundaries. The estimated complexity corresponds to $153 \times 4 = 612$ dofs in the approximating Loop surface, where each dof corresponds to one control vertex. Using our procedure with an initial surface obtained by decimating the bunny model down to 612 vertices and performing 5 parameter correction steps, gives a relative maximum deviation of 0.63% and takes 32s to compute on 2.8GHz Pentium IV including the decimation (Fig. 7).

Next we compare our results to the multiresolution subdivision surface fitting technique proposed in [LLS01a] since this is, to our knowledge, the only work where subdivision surfaces have been used to produce high-quality approximations of complex objects comparable to ours. The algorithm presented in that paper is very efficient due to the quasi-interpolation fitting and the multiresolution hierarchy. However, as we show in Fig. 6, the number of degrees of freedom required in [LLS01a] for obtaining the same precision is significantly larger (8 times for this example) than in our method.

6. Future work

The parameter correction procedure, which we use to establish the correspondence between the samples and the approximating sur-

face, does *not* guarantee one-to-one mapping in all cases. Nevertheless, in practice we observed flipping only when the initial fitting control mesh was extremely coarse and could always be avoided by allowing enough degrees of freedom from the beginning. The connectivity regularization procedure could be improved using [SG03]. Full support of piecewise-smooth subdivision surfaces [ZK02] and a lot of performance optimizations are still pending in our implementation. We also investigate the possibility to use an anisotropic remeshing [ACSD*03] of S as an initial control mesh for the fitting procedure, which could reduce the final control mesh complexity.

7. Acknowledgment

This work was partially supported by the European Union Research Training Network "Multiresolution in Geometric modeling (MINGLE)" under grant HPRN-CT-1999-00117. The Buddha model is courtesy of Igor Guskov and Zoë Wood and is produced with the procedure described in [GW01]. We would like to thank Sergey Alekseev and Mario Botsch for their help to convert the original models on Fig. 1 and Fig. 6 and the anonymous reviewers for their constructive comments.

References

- [ACSD*03] ALLIEZ P., COHEN-STEINER D., DEVILLERS O., LEVY B., DESBRUN M.: Anisotropic polygonal remeshing. *ACM Transactions on Graphics. Special issue for SIGGRAPH conference* (2003), 485–493. 8
- [ASCE02] ASPERT N., SANTA-CRUZ D., EBRAHIMI T.: MESH: Measuring errors between surfaces using the Hausdorff distance. In *Proceedings of the IEEE International Conference on Multimedia and Expo* (2002), vol. I, pp. 705 – 708. <http://mesh.epfl.ch>. 4
- [BKZ01] BIERMANN H., KRISTJANSSON D., ZORIN D.: Approximate Boolean operations on free-form solids. In *SIGGRAPH'2001* (2001), pp. 185–194. 3, 5
- [Bre73] BRENT R. P.: *Algorithms for Minimization without Derivatives*. Prentice-Hall, Englewood Cliffs, NJ, 1973. 4
- [BS02] BOLZ J., SCHRÖDER P.: Rapid evaluation of Catmull-Clark subdivision surfaces. In *Proceeding of the seventh international conference on 3D Web technology* (2002), ACM Press, pp. 11–17. 4




Figure 8: From left to right: Cyberware scan of a machine detail, initial D , final D and base mesh C_0 overlaid on the final D .




Figure 9: Left: an approximation D of the Buddha model. Right: top - close-up views of the original data, bottom - close-up views of D .




Figure 10: Left: an approximation D of the Bust model. Right: solid and wire-frame close-up views: left - original data, middle - final D and \hat{C}_2 , right - C_0 overlaid on D and \hat{C}_0 .

Figure 11: Top: an approximation of the Horse model. Bottom: close-up wire-frame views from behind on the left back leg: left - the original mesh, right - \hat{C}_1 . Note the non-uniform density of the approximating mesh corresponding to the varying geometrical complexity of the model.

- [CC78] CATMULL E., CLARK J.: Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer Aided Geometry Design* 10, 6 (Sep 1978), 350–355. [1](#)
- [CMS03] COHEN A., MERRIEN J.-L., SCHUMAKER L. L. (Eds.): *Curve and Surface Fitting: Saint-Malo 2002*. Nashboro Press, Brentwood, TN, 2003. [2](#)
- [Die95] DIETZ U.: *Erzeugung glatter Flächen aus Messpunkten*. Technical Report 1717, Department of Mathematics, University of Darmstadt, Germany, 1995. [2](#)
- [EH96] ECK M., HOPPE H.: Automatic reconstruction of B-spline surfaces of arbitrary topological type. *Computer Graphics* 30, Annual Conference Series (1996), 325–334. [2, 7](#)
- [Elf78] ELFVING T.: *On the Conjugate Gradient Method for Solving Linear Least Squares Problems*. Tech. Report LiTH-MAT-R-78-3, Linköping University, Sweden, 1978. [3](#)
- [GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. *Computer Graphics* 31, Annual Conference Series (1997), 209–216. [7](#)
- [GW01] GUSKOV I., WOOD Z.: Topological noise removal. In *Proceedings of Graphics Interface* (2001), Watson B., Buchanan J. W., (Eds.), pp. 19–26. [8](#)
- [HDD*94] HOPPE H., DEROSE T., DUCHAMP T., HALSTEAD M., JIN H., McDONALD J., SCHWEITZER J., STUETZLE W.: Piecewise smooth surface reconstruction. In *SIGGRAPH'94* (July 1994), pp. 295–302. [2, 3, 5](#)
- [HKD93] HALSTEAD M., KASS M., DEROSE T.: Efficient, fair interpolation using Catmull-Clark surfaces. In *Computer Graphics* (1993), ACM Press, pp. 35–44. [2](#)
- [JK02] JEONG W.-K., KIM C.-H.: Direct reconstruction of a displaced subdivision surface from unorganized points. *Graph. Models* 64, 2 (2002), 78–93. [2, 3](#)
- [LLS01a] LITKE N., LEVIN A., SCHRÖDER P.: Fitting subdivision surfaces. In *IEEE Visualization 2001* (October 2001), pp. 319–324. [2, 3, 5, 7, 8](#)
- [LLS01b] LITKE N., LEVIN A., SCHRÖDER P.: Trimming for subdivision surfaces. *Computer Aided Geometry Design* 18, 5 (2001), 463–481. [2, 3, 5](#)
- [LMH00] LEE A., MORETON H., HOPPE H.: Displaced subdivision surfaces. In *SIGGRAPH'2000* (2000), Akeley K., (Ed.), ACM Press / ACM SIGGRAPH / Addison Wesley Longman, pp. 85–94. [2, 3, 5](#)
- [Loo87] LOOP C.: *Smooth subdivision surfaces based on triangles*. Master's thesis, Department of Mathematics, University of Utah, 1987. [1](#)
- [MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and Mathematics III*, Hege H.-C., Polthier K., (Eds.). Springer-Verlag, Heidelberg, 2003, pp. 35–57. [6](#)
- [MMTP02] MA W., MA X., TSO S.-K., PAN Z.: Subdivision surface fitting from a dense triangle mesh. In *Geometric Modeling and Processing* (2002), pp. 94–103. [2, 7](#)
- [MZ00] MA W., ZHAO N.: Catmull-Clark surface fitting for reverse engineering. In *Geometric Modeling and Processing* (2000), pp. 274–284. [2](#)
- [Sap94] SAPIDIS N. S. (Ed.): *Designing fair curves and surfaces: shape quality in geometric modeling and computer-aided design*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1994. [2](#)
- [SG03] SURAZHISKY V., GOTSMAN C.: Explicit surface remeshing. In *Proceedings of Eurographics Symposium on Geometry Processing* (June 2003), pp. 17–28. [6, 7, 8](#)
- [Sta98] STAM J.: Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values. In *SIGGRAPH'98* (1998), vol. 32, pp. 395–404. [2](#)
- [Sta99] STAM J.: Evaluation of Loop subdivision surfaces. In *SIGGRAPH'99 Course notes* (1999). [2](#)
- [STK99] SUZUKI H., TAKEUCHI S., KANAI T.: Subdivision surface fitting to a range of points. In *Pacific Graphics* (1999), pp. 158–167. [2](#)
- [TSK*00] TAKEUCHI S., SUZUKI H., KIMURA F., KANAI T., SHIMADA K.: Subdivision surface fitting using QEM-based mesh simplification and reconstruction of approximated B-spline surfaces. In *Pacific Graphics* (2000). [2, 7](#)
- [ZK02] ZORIN D., KRISTJANSSON D.: Evaluation of piecewise smooth subdivision surfaces. *The Visual Computer* 18, 5-6 (August 2002), 299–315. [8](#)
- [Zor97] ZORIN D.: *Subdivision and Multiresolution Surface Representations*. PhD thesis, Caltech, Pasadena, California, 1997. [3](#)
- [ZSS*00] ZORIN D., SCHRÖDER P., SWELDENS W., DEROSE T., KOBBELT L., LEVIN A.: Subdivision for modeling and animation. In *SIGGRAPH'2000 Course notes* (2000). [2](#)