

Gaussian-Product Subdivision Surfaces

REINHOLD PREINER, Graz University of Technology, Austria
TAMY BOUBEKEUR, LTCI, Telecom ParisTech, Paris-Saclay University, France
MICHAEL WIMMER, TU Wien, Austria

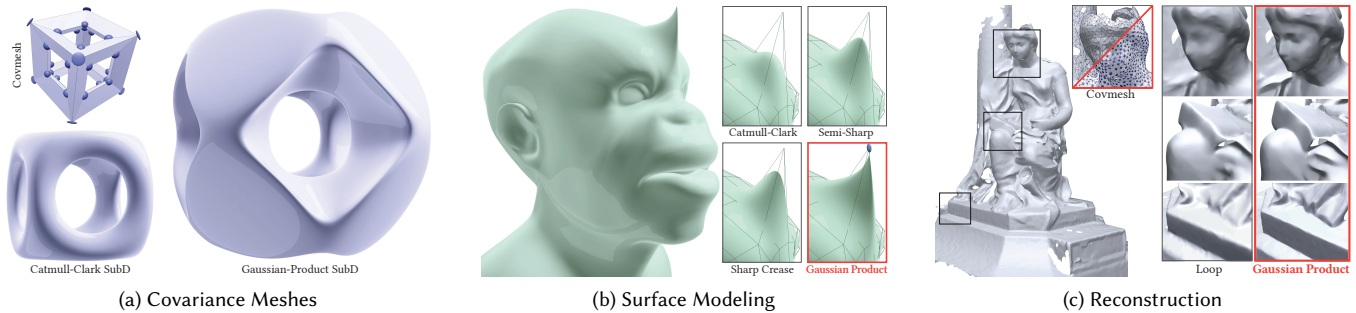


Fig. 1. We introduce a new nonlinear subdivision surface model, which is based on a control-mesh representation encoding Gaussian covariances in its vertices (a). Our surface definition relies on a refinement using Gaussian products, providing wider shape control on the smooth limit surface. This significantly widens the space of possible shapes stemming from a given control mesh, enabling better modeling features, (semi-)sharpness and concavities without changing the base connectivity (b). In addition, our representation naturally integrates into surface-reconstruction pipelines, recovering high- and mid-frequency structures even from a low-resolution control mesh (c).

Probabilistic distribution models like Gaussian mixtures have shown great potential for improving both the quality and speed of several geometric operators. This is largely due to their ability to model large fuzzy data using only a reduced set of atomic distributions, allowing for large compression rates at minimal information loss. We introduce a new surface model that utilizes these qualities of Gaussian mixtures for the definition and control of a parametric smooth surface. Our approach is based on an enriched mesh data structure, which describes the probability distribution of spatial surface locations around each vertex via a Gaussian covariance matrix. By incorporating this additional covariance information, we show how to define a smooth surface via a nonlinear probabilistic subdivision operator based on products of Gaussians, which is able to capture rich details at fixed control mesh resolution. This entails new applications in surface reconstruction, modeling, and geometric compression.

CCS Concepts: • **Computing methodologies** → **Parametric curve and surface models**; *Mixture models*; *Mesh geometry models*.

ACM Reference Format:

Reinhold Preiner, Tamy Boubekeur, and Michael Wimmer. 2019. Gaussian-Product Subdivision Surfaces. *ACM Trans. Graph.* 38, 4, Article 35 (July 2019), 11 pages. <https://doi.org/10.1145/3306346.3323026>

This work was partially supported by the EU framework programme 7, under grant 323567 Harvest4D, the French National Research Agency (ANR) under grant ANR 16-LCV2-0009-01 ALLEGORI, and by BPI France, under grant PAPAYA.

Authors' addresses: Reinhold Preiner, r.preiner@cgvtugraz.at, Graz University of Technology, CGV, Graz, Austria; Tamy Boubekeur, LTCI, Telecom ParisTech, Paris-Saclay University, IDS, Paris, France; Michael Wimmer, TU Wien, Vienna, Austria.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2019 Copyright held by the owner/author(s).

0730-0301/2019/7-ART35

<https://doi.org/10.1145/3306346.3323026>

1 INTRODUCTION

For the efficient processing of fuzzy geometric data like noisy point sets, probabilistic distribution models such as Gaussian mixtures have recently shown great potential for tasks like registration, filtering or resampling. This is largely due to their ability to model large fuzzy data using only a reduced set of atomic distributions, allowing for large compression rates at minimal information loss. Due to this compactness, it is desirable to be able to define a surface directly on such a sparse model, and avoid the need to expand to larger representations (e.g., meshes or point clouds) for further processing and rendering. In fact, such a sparse representation is also highly interesting for modeling applications, which aim at defining complex shapes using simple base representations. So far, there have been some attempts to define a probabilistic surface along the ridge contour of the probability density function (pdf) of a dense Gaussian mixture. However, this contour degenerates when the mixture is compressed to large anisotropic Gaussians, where discontinuities appear; also, the resulting surface definition is not amenable for further processing or modeling tasks.

In this paper, we introduce a new probabilistic surface representation that allows defining continuous, artifact-free surfaces even for a sparse set of Gaussians, while still closely resembling the ridge of their pdf. Our surface definition is based on a polygonization of the individual Gaussian components, resulting in a new, enriched mesh model that carries anisotropic covariance information at its vertices, called *covariance mesh*. The key idea to turn this representation into a continuous surface is to depart from the linear combination of individual Gaussians used in mixture models, and instead consider their *joint probability*, leading to a new interpolation method between Gaussians based on a product formulation. To apply this

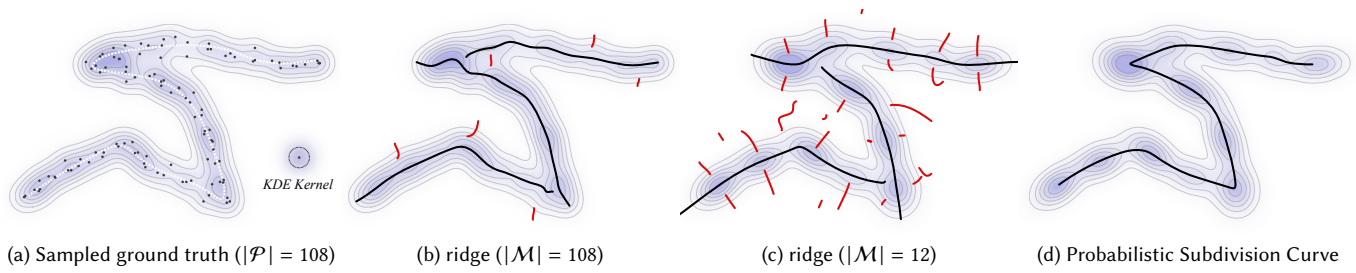


Fig. 2. (a) Surface contour (dashed white) sampled with normal distributed noise ($\sigma = 1.6\%$ of the shape diagonal). Convolving the Dirac distribution of the samples with a Gaussian kernel ($\sigma = 3.2\%$) results in a kernel density estimate of their pdf (blue height field). Note that the ground-truth contour runs closely along the ridge of this pdf, where the density is maximal. (b) The actual ridge contour contains spurious ridges (red), and discontinuities at highly curved features (black). (c) A maximum-likelihood simplification of the pdf to fewer anisotropic Gaussians generally smoothes the ridge contour, but reinforces discontinuities and the appearance of spurious ridges. (d) Gaussian-product subdivision curve interpolating the set of simplified Gaussians in (c).

idea to covariance meshes, we introduce a new family of subdivision operators, generalizing existing subdivision-surface algorithms by incorporating the anisotropic covariance information of the individual vertices (Fig. 1a). In particular, we define the subdivided vertex in a refined covariance mesh according to the new Gaussian interpolation method.

Despite representing a nonlinear subdivision scheme, the proposed subdivision operation based on Gaussians can actually be shown to be dual to traditional linear ones: We introduce a smooth map between covariance meshes and ordinary meshes in a dual space, such that all findings and tools developed for linear subdivision can be applied to our proposed Gaussian-product subdivision through these maps. Thus, one can use existing subdivision schemes like Loop or Catmull-Clark to define smooth probabilistic surfaces for covariance meshes of arbitrary complexity. In fact, our subdivision formulation is a generalization of traditional subdivision modeling methods: by modifying the covariances at the vertices of a covariance mesh, a user can achieve complex shape variations, including concavities and sharp creases, without increasing the complexity of the base mesh (Fig. 1b). Covariance meshes therefore provide a powerful new representation for modeling applications where subdivision surfaces have been used so far.

In order to compute the initial covariance mesh, we provide two different techniques: one based on a pure Gaussian mixture where we probabilistically triangulate its components – to be used to convert noisy point clouds into the new representation; and one based on an ordinary mesh, where covariances are inferred based on the locally surrounding geometry – usable in standard modeling scenarios where the user defines a coarse base mesh. We illustrate the power of covariance meshes and our new subdivision scheme with applications in interactive shape modeling (Fig. 1b) and surface reconstruction (Fig. 1c).

2 BACKGROUND AND RELATED WORK

2.1 Background

Gaussian Mixtures have been widely used in various scientific fields to describe the distribution of large complex data by a superposition of a sparse, tractable set of Gaussian components. In point-based

processing, they have been used for registration [Danelljan et al. 2016; Jian and Vemuri 2011], filtering [Calderon and Boubekeur 2014] and resampling [Preiner et al. 2014]. For a given set \mathcal{P} of discrete points sampled from a surface, Expectation Maximization (EM) [Dempster et al. 1977] or its faster hierarchical variants [Vasconcelos and Lippman 1999] can be used to compute a Gaussian mixture $\mathcal{M} = \{w_i, \Theta_i\}$ modeling their pdf

$$f_{\mathcal{M}}(\mathbf{x}) = \sum_i w_i f(\mathbf{x}|\Theta_i) \quad (1)$$

via a set of anisotropic Gaussians $\Theta_i = (\mu_i, \Sigma_i)$, where μ_i denotes the mean, Σ_i the covariance, w_i the prior weight, and

$$f(\mathbf{x}|\Theta_i) = |2\pi\Sigma_i|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{x}-\mu_i)^T \Sigma_i^{-1}(\mathbf{x}-\mu_i)} \quad (2)$$

the Gaussian pdf of the i -th mixture component.

Our principal aim is to define a *continuous, smooth surface* S that faithfully resembles the shape of the original surface solely based on this compact probabilistic representation \mathcal{M} . Intuitively, an optimal reconstruction with respect to the probability density of \mathcal{M} is one that places the surface S along the *ridge* of the pdf landscape $f_{\mathcal{M}}$, where the probability value is maximal. Various methods exist that extract this ridge by finding local maxima in the pdf along trajectories that are given by the smallest negative eigenvector of its Hessian [Li et al. 2010; Ozertem and Erdogmus 2011; Süßmuth and Greiner 2007]. However, general Gaussian mixture pdfs can exhibit discontinuities in the smallest Hessian eigenvector field, deteriorating this ridge contour. Figure 2a depicts a dense mixture pdf (blue iso-lines) of a noisy point set, produced by convolving each point with an isotropic Gaussian kernel. As shown in Figure 2b, this ridge formulation leads to discontinuous ridge contours (black lines), and the occurrence of secondary, *spurious* ridges (red lines), which have to be dealt with. Moreover, processing such dense mixtures is quite inefficient, and becomes infeasible for large input data. In contrast, we strive for a sparse surface model that exploits the compactness of compressed Gaussian mixtures. Figure 2c shows that a maximum likelihood compression of the dense mixture to only few anisotropic Gaussians still faithfully resembles the original pdf. However, while the resulting ridge contour is expectedly smoothed, the increased sparseness and anisotropy of the Gaussians

even increases the appearance of discontinuities and spurious ridges (Figure 2c).

Our proposed surface model follows an alternative approach that directly works on a compact, sparse set of anisotropic Gaussians: By explicitly defining the connectivity between the Gaussian components, we can apply a new probabilistic subdivision operator, producing limit contours that closely resemble the ridge of the Gaussians' pdf while maintaining its continuity (Figure 2d).

2.2 Related Work

Subdivision surfaces. Subdivision surfaces generalize spline surfaces by using any 2-manifold surface mesh as base control net. They are defined by a recursive process of applying a *subdivision mask* to a mesh and are often analyzed by the properties of its limit, i.e., the geometry defined by an infinity of subdivision steps. These properties include the continuity order, the approximation or interpolation level and the regularity/curvature distribution. Subdivision surfaces are often classified regarding their control mesh polygon type, triangles [Loop 1987] and quads [Catmull and Clark 1978] being the most common flavors. A complete overview of subdivision surfaces is given in the course by Zorin and Schröder [2000], and we refer the reader to the work of Brainerd et al. [2016] and Karčiauskas and Peters [2018] for a survey of recent evolutions.

To account for extra surface properties such as sharp creases and prescribed tangents, Hoppe et al. [1994] proposed alternative subdivision masks in the presence of tagged sharp edges and vertices for the Loop subdivision scheme [Loop 1987]. Later on, DeRose et al. [1998] introduced semi-sharp creases for Catmull-Clark subdivision surfaces, allowing modeling a continuum of features between smooth and sharp edges using a single scalar per edge. Other control mechanisms have been proposed, such as the method by Biermann et al. [2000] who prescribe the local tangent plane at the limit surface using a per-base-vertex normal vector. Introducing the idea of “guiding” the subdivision process to achieve certain properties, Levin [2006] proposed to use a polynomial function at extraordinary points to recover C^2 continuity, while Karčiauskas and Peters [2007] enabled the use of any piecewise smooth function as a guide.

Nonlinear subdivision schemes. In our work we build on a particular class of subdivisions that replace linear averaging rules by *nonlinear* refinement operators. It can be shown that such nonlinear refinements can be achieved by performing a linear averaging after certain locally invertible nonlinear mappings [Schaefer et al. 2008], both of which pass on their differentiability properties to the resulting nonlinear scheme. This constitutes a basic recipe for the construction of a large spectrum of both interpolatory and approximative nonlinear subdivision operators. For instance, Vaxman et al. [2018] recently constructed a nonlinear subdivision scheme by applying a linear subdivision after mapping the one-ring of each vertex to a canonical form using a Möbius transformation.

In general, the analysis of smoothness and convergence of such nonlinear schemes is much more elaborate than those of linear subdivision models. Wallner and Dyn [2005] and Grohs [2009; 2010] analyze these properties for nonlinear schemes on manifolds and in Lie groups based on their proximity relation to the linear schemes

from which they are derived. Other work by Wallner analyzed univariate interpolatory schemes generalized to Riemannian manifolds via a log/exp mapping [Wallner 2014], and investigated the convergence and smoothness properties of subdivision operations in symmetric spaces [Wallner et al. 2011]. The latter contains the space of positive definite symmetric matrices, such as the Gaussian covariance matrices, that is of particular interest for the subdivision of our covariance meshes. Subdivision schemes for this class of matrices have been studied by Itai and Sharon [2013]. The application of subdivision to the interpolation of such manifold-valued functions over smooth two-manifolds has been demonstrated by Weinmann [2010], who showed convergence and C_1 smoothness of nonlinear subdivision schemes on irregular meshes.

Instead of merely interpolating matrix-valued functions over manifolds, we refine covariance kernels located in 3d space to define the shape of an intrinsic two-manifold via subdivision. Using matrix weights for subdivision surfaces was pioneered by Yang [2016] and used to provide normal control. Our work expands on these methods and introduces a more powerful control mechanism, using Gaussian kernels to model sharp features, concavities, and more accurately control curvature without changing the control mesh topology. By applying concepts from nonlinear subdivision theory to contour the probability density of a Gaussian mixture, we thereby build a bridge between a probabilistic and a geometric view on the data.

Gaussian triangulation. Generating a covariance mesh from a sparse Gaussian mixture that is computed from an input point cloud requires establishing connectivity between the Gaussians. A straightforward strategy is to greedily and locally generate connectivity among samples. Lower-dimensional meshing techniques [Gopi et al. 2000; Linsen and Prautzsch 2003] or front-propagation methods [Cohen-Steiner and Da 2004] typically work well on dense enough data sets, are able to capture non-manifold regions, and are usually very fast to compute. Although they are based on heuristics and cannot guarantee manifoldness, their flexibility makes them good candidates for triangulating anisotropic Gaussian sets that model the potential space occupancy of a surface, particularly when the set is not guaranteed to be sampled under strict feature-size conditions. However, to the best of our knowledge, existing greedy meshing methods are solely based on point sets. In our case, each sample comes as an anisotropic Gaussian, for which a dedicated meshing method can exploit the richer information carried by the distribution, accounting for the underlying anisotropy to favor a specific local connectivity structure.

3 OVERVIEW

Our new surface representation is defined on a manifold mesh $\Pi = (\mathcal{V}, \mathcal{E}, \mathcal{F})$ that encodes individual Gaussians $\Theta_i = (\mu_i, \Sigma_i)$ in its vertices $\mathcal{V} = \{\Theta_i\}$, and defines their topological connectivity via edges \mathcal{E} and faces \mathcal{F} . Since this mesh definition extends the positional information μ_i stored in a vertex by a covariance matrix Σ_i , we call Π a *covariance mesh*, or *covmesh* (Fig. 3).

Given such a covmesh, we introduce a family of nonlinear subdivision schemes that uses the covariance information for a refinement based on the *product* of vertex Gaussians, leading to a continuous limit contour that closely approximates the ridge of their associated

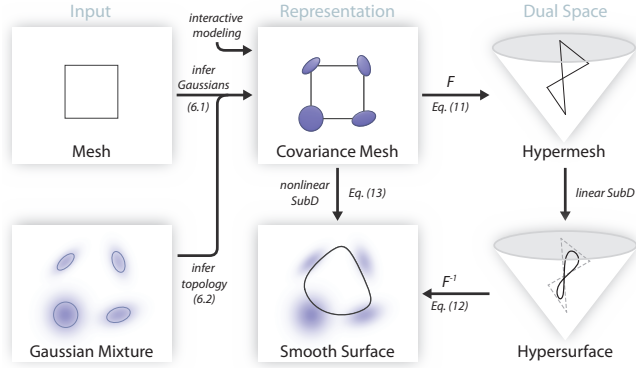


Fig. 3. Overview of our surface model. Its central representation is a *covariance mesh* (center), a mesh structure that extends vertices by Gaussian covariances and can be generated from a given mesh, or from a Gaussian mixture (left). They define smooth surfaces via a nonlinear subdivision operator, which corresponds to linear subdivision in a dual space (right).

probability density (Section 4). This new *probabilistic* subdivision scheme exhibits several interesting properties (Section 5): it corresponds to given linear schemes in a dual space via a smooth map, is thus easy to implement and integrate into existing subdivision tools, and produces limit surfaces exhibiting the same continuity order as their corresponding linear schemes, while modeling more complex features at the same base-mesh complexity.

We will discuss ways to generate a covmesh both from given classical meshes, by inferring the required covariance information, and from Gaussian mixtures, via a heuristic triangulation method that is guided by the covariances (Section 6). Finally, we demonstrate the merits of our new surface definition for applications such as interactive modeling of sharp features using sparse control meshes (Section 7.1), and highly efficient surface reconstruction from points via Gaussian mixtures (Section 7.2).

4 PROBABILISTIC SUBDIVISION SURFACES

In this section we develop a smooth surface definition for sparse covariance meshes by introducing a probabilistic subdivision operator that acts on the individual Gaussian distributions stored at its vertices. This results in a new family of subdivision surfaces, which extend existing linear schemes like Loop or Catmull-Clark subdivision in a simple and elegant way.

Probabilistic Refinement. We first assess the basic example of two topologically connected Gaussians Θ_i and Θ_j shown in black in Figure 4a. Our aim is to define a smooth contour between their means describing a continuous path of highest possible probability density. While the ridge of their common pdf (blue contour) would have maximum density along its path, it has problems to meet the continuity requirement, especially close to where the Gaussian kernels overlap. Therefore, in contrast to a ridge-based curve formulation based on the *mixture* distribution, our strategy is to recursively refine the mesh by inserting a new Gaussian that models the probabilistic overlap of the original Gaussians. To this end, we employ the Gaussians' *joint* distribution, which is given by the (normalized)

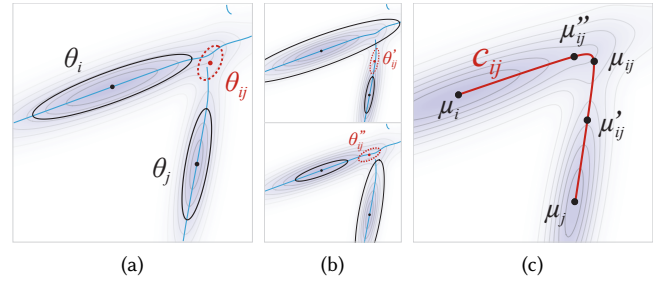


Fig. 4. (a) Two Gaussians Θ_i and Θ_j and their common product Gaussian Θ_{ij} . Dots indicate means, ellipses the unit-variance isocontours of their kernels. Blue lines represent the ridge of their common mixture pdf, exhibiting a discontinuity near their pdf overlap. (b) Movement of the product at double (resp. half) the kernel of Θ_i (resp. Θ_j) (upper image) and vice versa (lower image). (c) Continuous contour c_{ij} interpolating the means of the product Gaussians shown in (a) and (b).

product of their individual pdfs,

$$f(\mathbf{x}|\Theta_{ij}) = \omega^{-1} f(\mathbf{x}|\Theta_i) f(\mathbf{x}|\Theta_j), \quad (3)$$

and results in the pdf of a new, subdividing Gaussian Θ_{ij} (red dashed). Here, the term $\omega = \int_{\mathbb{R}^d} f(\mathbf{x}|\Theta_i) f(\mathbf{x}|\Theta_j) d\mathbf{x}$ only accounts for the fact that this product is generally not a pdf that integrates to 1, but does not influence the location or shape of the resulting product Gaussian. In the following, we will always use a term ω^{-1} to indicate respective pdf normalizations.

Figure 4a shows that the mean $\mu_{ij} = E[\mathbf{x}|\Theta_{ij}]$ of this product Gaussian gives a suitable location for a refinement point of high joint probability density, even in a region where the continuity of the analytic pdf ridge tends to break down. By further subdividing the new edge between Θ_i and Θ_j , we get

$$f(\mathbf{x}|\Theta_{iij}) = \omega^{-1} f(\mathbf{x}|\Theta_i) f(\mathbf{x}|\Theta_{ij}) = \omega^{-1} f(\mathbf{x}|\Theta_i)^2 f(\mathbf{x}|\Theta_j).$$

By repeating this process, any resulting Gaussian can be expressed as a weighted product of the two base Gaussians, where the weights modify the *power* of their pdfs:

$$f(\mathbf{x}|\Theta_{ij}) = \omega^{-1} f(\mathbf{x}|\Theta_i)^{\alpha_i} f(\mathbf{x}|\Theta_j)^{\alpha_j}. \quad (4)$$

Increasing the power of a Gaussian pdf results in a downscaling of its covariance kernel Σ_i by α_i^{-1} . As shown in Fig. 4b, this also causes Θ_{ij} to shift towards the Gaussian with higher power, or similarly, smaller kernel. The set of all weighted product means μ_{ij} , for weights $\alpha_i \geq 0$,

$$\mu_{ij} = \left(\alpha_i \Sigma_i^{-1} + \alpha_j \Sigma_j^{-1} \right)^{-1} \left(\alpha_i \Sigma_i^{-1} \mu_i + \alpha_j \Sigma_j^{-1} \mu_j \right), \quad (5)$$

describes a curve c_{ij} connecting the Gaussians' means, as the weights are invariant to any scaling and can thus be expressed using a single parameter $t = \frac{\alpha_i}{\alpha_i + \alpha_j}$. This curve closely resembles the ridge of the common pdf and thus gives a robust definition of a probabilistic surface contour (Fig. 4c). To extend this refinement strategy to covariance meshes embedded in \mathbb{R}^3 , we generalize Eq. (3) to a weighted product of $n > 2$ topologically connected Gaussians Θ_i ,

which is given by their joint pdf

$$f(\mathbf{x}|\Theta_J) = \omega^{-1} \prod_{i \in J} f(\mathbf{x}|\Theta_i)^{\alpha_i} \quad (6)$$

with weights $\alpha_i \geq 0$. Eq. (6) can now be used to guide the insertion of new (*odd*) face and edge vertices as well as the update of existing (*even*) vertices in a covariance mesh. What remains to be ascertained is a choice of weights α_i that ensures that the limit mesh of the resulting probabilistic subdivision is in fact a smooth surface.

Choice of Weights. For linear subdivision schemes, analyzing their corresponding subdivision matrices makes it comparatively easy to determine suitable weights for the linear combination of adjacent vertices in a subdivision computation, such that the limit surface exhibits a particular smoothness. In general, however, the smoothness analysis of nonlinear subdivision schemes, as the one introduced above, is much more elaborate, making it difficult to directly determine similarly suitable weights α_i for our probabilistic refinement scheme. In the following, however, we will show that by using a particular mapping of the Gaussians to a higher-dimensional space, we can reformulate the nonlinear subdivision operator from Eq. (6) to an ordinary linear subdivision scheme, from which we can directly adopt both its definition of weights α_i and its corresponding smoothness properties of the resulting limit surface.

To this end, we reformulate the pdf of a Gaussian Θ_i in Eq. (2), by expressing the quadratic form in its exponent by a quadratic polynomial basis $\mathbf{b}(\mathbf{x})$ and a corresponding coefficient vector q_i ,

$$f(\mathbf{x}|\Theta_i) = c \cdot e^{-\frac{1}{2} \mathbf{b}(\mathbf{x})^T q_i}, \quad (7)$$

where we let c collect all non-exponential factors. This way, we can identify every Gaussian Θ_i with a point q_i encoding the elements of μ_i and Σ_i^{-1} via a bijective map $F : \Theta_i \mapsto q_i$ to a space Q of quadratic functions with basis \mathbf{b} . By substituting Eq. (7) into Eq. (6) and simplifying, the pdf of the product Gaussian Θ_J can be defined as a single exponential,

$$f(\mathbf{x}|\Theta_J) = c \cdot e^{-\frac{1}{2} \mathbf{b}(\mathbf{x})^T \sum_i \alpha_i q_i} = c \cdot e^{-\frac{1}{2} \mathbf{b}(\mathbf{x})^T q_J}. \quad (8)$$

The right side of Eq. (8) now makes apparent that in Q , the point q_J corresponding to the product Θ_J is defined by a linear combination

$$q_J = \sum_{i \in J} \alpha_i q_i \quad (9)$$

of the individual Gaussian images q_i . This indicates that if we map a covariance mesh $\Pi = (\{\Theta_i\}, \mathcal{E}, \mathcal{F})$ to a corresponding hypermesh $M^* = (\{F(\Theta_i)\}, \mathcal{E}, \mathcal{F})$, then applying a linear subdivision with given weight coefficients α_i to M^* corresponds to a probabilistic subdivision of Π with identical weights. Thus, for our probabilistic subdivision scheme defined in Eq. (6), we can adopt the weights α_i provided by any given linear scheme L that suits the topology of Π .

Whether these weights actually result in a probabilistic limit surface $S \subset \mathbb{R}^d$ with similar smoothness properties than the limit hypersurface $S^* \subset Q$ resulting from L depends on the properties of the inverse map F^{-1} : If S^* is C^n continuous in its points q , then S is C^n continuous in the means of their image Gaussians $F^{-1}(q)$ provided the inverse mapping function F^{-1} is smooth over the domain of S^* . Next, we will show that for $\alpha_i \geq 0$ this is the case.

Mapping Function. The definition of the mapping $F : (\mu, \Sigma) \mapsto q$, satisfying $\mathbf{b}(\mathbf{x})^T q + c_q = (\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu)$, depends on the choice of the polynomial basis \mathbf{b} , which, to our convenience, we define as

$$\mathbf{b}(\mathbf{x})^T = (\text{vech}(2 \mathbf{x} \mathbf{x}^T - \text{diag}(\mathbf{x})^2)^T, -2\mathbf{x}^T), \quad (10)$$

where the first and second part represent the bases of the quadratic and linear coefficients, respectively. Here, *vech* is the half-vectorization operator linearizing the lower triangular part of a symmetric matrix. Using the above basis, the map $F(\mu, \Sigma)$ is given by the vector $q = (\check{q}^T, \bar{q}^T)$, consisting of quadratic and linear sub-coefficient vectors:

$$\check{q} = \text{vech}(\Sigma^{-1}) \quad \bar{q} = \Sigma^{-1} \mu \quad (11)$$

Note that it is not necessary for this basis to model the constant term $c_q = \mu^T \Sigma^{-1} \mu$ of the quadratic polynomial, since being in the exponent of the pdf in Eq. (7), it only affects the scale c of its amplitude. For covmeshes defined over \mathbb{R}^d , their image space Q thus has dimensionality $k = \binom{d+2}{d} - 1$ (number of linear and quadratic polynomials in \mathbb{R}^d). To perform the inverse mapping $F^{-1} : q \mapsto (\mu, \Sigma)$, we first restore its covariance, which is then used to restore its mean:

$$\Sigma = F_{\Sigma}^{-1}(q) = [\check{q}]^{-1} \quad \mu = F_{\mu}^{-1}(q) = \Sigma \bar{q} \quad (12)$$

where we use $[\cdot]$ to denote the inverse half-vectorization operator vech^{-1} restoring a symmetric matrix. Note that using Eq. (9), (11) and (12), we can also directly express the refined Gaussian parameters by the Gaussians from the coarser level:

$$\begin{aligned} \Sigma_J &= [\check{q}_J]^{-1} = \left[\sum_i \alpha_i \check{q}_i \right]^{-1} = \left(\sum_i \alpha_i \Sigma_i^{-1} \right)^{-1}, \\ \mu_J &= \Sigma_J \left(\sum_i \alpha_i \bar{q}_i \right) = \left(\sum_i \alpha_i \Sigma_i^{-1} \right)^{-1} \left(\sum_i \alpha_i \Sigma_i^{-1} \mu_i \right), \end{aligned} \quad (13)$$

which extends the product definition in Eq. (5) to n Gaussians.

5 PROPERTIES

This section discusses the smoothness and the shape control provided by Gaussian-product subdivision. Differential geometric properties like tangents and curvature are analyzed in Appendix B.

Smoothness. By definition, the Gaussian covariances Σ_i stored in a covmesh Π are positive-definite matrices, and so are their inverses Σ_i^{-1} . Within the subspace \check{Q} of quadratic coefficients spanned by the first part of the basis \mathbf{b} defined in Eq. (10), their images \check{q}_i therefore all map to a conical region $\check{Q}_{pd} \subset \check{Q}$ known as the *positive-definite cone* [Hill and Waters 1987]. The boundary $\partial \check{Q}_{pd}$ of this cone represents a set of singular, i.e., non-invertible symmetric matrices, for which both the map F in Eq. (11) and its inverse F^{-1} in Eq. (12) are undefined. Discontinuities in the probabilistic limit subdivision surface S will thus occur at the intersection $S^* \cup \partial \check{Q}_{pd}$ of its corresponding limit hypersurface with this singular cone boundary. However, such intersections can be avoided if the linearly refined hyperpoints q_J resulting from Eq. (9) are ensured to stay within the convex hull of the input hypervertices $q_i \in \check{Q}_{pd}$, which is the case for any *convex* set of weights α_i . Besides affinity ($\sum_i \alpha_i = 1$), which is a requirement for non-diverging subdivision schemes anyway,



Fig. 5. Reflection lines for the standard Catmull-Clark (left) and Gaussian-product Catmull-Clark limit surface (right) shown in Fig. 1a.

this implies the usage of non-negative weights $\alpha_i \geq 0$, as provided by approximative schemes like Loop or Catmull-Clark subdivision.

For a probabilistic subdivision of a given covmesh Π , we can therefore adopt the weights $\alpha_i \geq 0$ of such an approximative scheme L and obtain a continuous limit surface $S = \{F_\mu^{-1}(q) : q \in S^*\}$. Since within Q_{pd} the inverse map F_μ^{-1} given by Eq. (12) is well defined and C^∞ differentiable, the resulting probabilistic surface adopts the smoothness properties of S^* , that is, the continuity orders at regular and extraordinary points defined by the associated linear scheme L . This is demonstrated in Figure 5 by comparing reflection lines between linear and Gaussian-product Catmull-Clark subdivision.

Covariance Shape Control. Given our new probabilistic mesh representation and subdivision-surface definition, we are interested in the effect of individual Gaussian covariance configurations on the resulting limit subdivision surface. In general, the anisotropy of the covariance kernels attached to the control vertices, that is, the ratio of their eigenvalues, directly affects the smoothness (i.e., the curvature) of the local subdivision contour. Figure 6a illustrates this feature by successively reducing the anisotropic extent of the Gaussians from Figure 4a (dashed). With decreasing anisotropy, the contour c'_{ij} becomes smoother (red), until they are both isotropic, where the resulting curve c''_{ij} reduces to a straight line (blue).

This behavior also lets the subdivision contour intuitively reflect the level of uncertainty encoded in the Gaussians. In a sparse mixture representing a noisy point set, stronger noise will cause larger surface-orthogonal variance, resulting in reduced anisotropy and smoother contours (Fig. 6b). In contrast, a low level of noise will create thinner, more anisotropic Gaussians, allowing for a more precise definition of sharp features (Fig. 6c). This suggests that we can control the smoothness of the contour by reducing the anisotropy of the covariances, as can be achieved by convolving the mixture with an isotropic Gaussian smoothing kernel (blurring). Eq. (13) reveals two further properties of the influence of the covariance kernels on the surface shape. First, applying a global scale factor s to the covariances Σ_i does not change the shape of the curve, as this scale factor cancels out in the definition of μ_J . Second, using a globally constant covariance $\Sigma_i = \Sigma_c$ and convex weights α_i , the subdivision Gaussian Θ_J reduces to

$$\Sigma_J = \Sigma_c, \quad \mu_J = \left(\sum_i \alpha_i \Sigma_c^{-1} \right)^{-1} \left(\sum_i \alpha_i \Sigma_c^{-1} \mu_i \right) = \sum_i \alpha_i \mu_i, \quad (14)$$

which equals a linear subdivision of the vertices μ_i , with stationary covariance Σ_c . Thus, our probabilistic subdivision is in fact a true generalization of ordinary linear subdivision schemes. An example

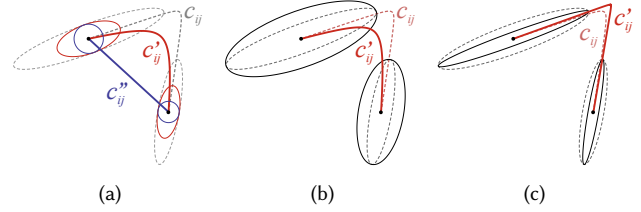


Fig. 6. Influence of the shape of two Gaussian kernels on their common subdivision contour c_{ij} . (a) Successively reducing their tangential variance straightens out the curve up to the point where c_{ij} is a straight line. (b) Increasing the surface-orthogonal variance produces smoother contours as well, while (c) reducing the variance results in sharper features.

is given by the two isotropic kernels in Fig. 6a (blue), leading to a straight line c''_{ij} .

6 COVMESH GENERATION

In this section, we discuss different ways for generating a covariance mesh depending on the type of input and considering two complementary scenarios. In the first one, we start from an existing surface mesh, and enrich its vertices with anisotropic Gaussian covariance information; in the second one, we start with a set of anisotropic 3d Gaussians, and infer topological connectivity among them to form a covmesh. While the former strategy addresses modeling scenarios, the latter allows for probabilistic point-based reconstruction.

6.1 Gaussian Inference

For a given sparse input mesh $\Pi = (\mathcal{V}, \mathcal{E}, \mathcal{F})$, we would like to have means for complementing its vertices with expressive anisotropic covariance information that produces a plausible subdivision surface shape while being able to model inherent features encoded in the mesh. Different approaches are conceivable for automatically inferring Gaussian covariances Σ_i at its vertices $v_i \in \mathcal{V}$. An intuitive and straightforward way is to employ a *vertex-based* estimate, given by the sample covariance of the local mesh umbrella vertices,

$$\Sigma_i^v = \text{cov}(v_i \cup \mathcal{N}(i)) + \sigma_0^2 I, \quad (15)$$

where the right term biases the diagonal with some variance σ_0^2 to prohibit rank-deficient matrices in flat areas. However, we propose a more feature-preserving, *face-based* variant that makes use of the same Gaussian product formulation that is central to our subdivision-surface definition. Our idea is to model the area of each neighboring face Δ_j of a vertex v_i by a ‘face Gaussian’ Θ_{Δ_j} and define the inferred tensor as the covariance of their product

$$\Sigma_i^f = \text{cov} \left(\prod_j f(\mathbf{x} | \Theta_{\Delta_j}) \right) = \left(\sum_j \Sigma_{\Delta_j}^{-1} \right)^{-1}. \quad (16)$$

Here, the face covariance is given by the sample covariance of the face vertices, biased by some bandwidth as before:

$$\Sigma_{\Delta_j} = \text{cov}(v \in \Delta_j) + \sigma_0^2 I. \quad (17)$$

The probabilistic limit subdivision surface of the covmesh resulting from such an automatic Gaussian inference allows modeling inherent surface features in the input mesh much more faithfully than

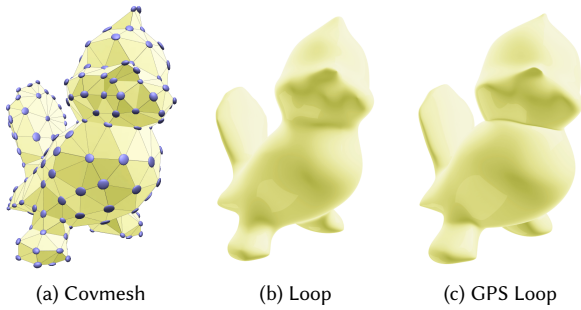


Fig. 7. Loop and Gaussian-product Loop subdivision surface on a covmesh.

standard linear subdivision schemes. Figure 7 shows an example of a coarse triangle mesh of the Tweety model, where Gaussians were inferred using Eq. (16). In comparison to its standard Loop subdivision surface, the resulting GPS surface creates a more characteristic neck line and beak.

6.2 Topological Inference

For tasks like surface reconstruction, covariance meshes can be used to directly define a reconstructed probabilistic surface based on a sparse Gaussian mixture that has been computed from a potentially very large input point cloud for efficient processing. In contrast to the previous section, such applications require us to solve the inverse problem, that is, for a given mixture, finding a manifold polygon mesh that connects its individual components to a topological surface. We propose a simple yet efficient way to obtain a triangulation of such a set of Gaussians, thus resulting in a triangular covariance mesh.

An ordinary Delaunay-based triangulation, as used for simple point sets, is obviously unsuitable for our purpose, as the anisotropic extent of the individual Gaussian kernels calls for different metrics for the assessment of local distance relations. One way is to approach this problem as general Delaunay triangulation for a Riemannian manifold, whose local metric is stretched according to the Gaussians' covariance tensors [Budninskiy et al. 2016; Rouxel-Labbé et al. 2016]. However, such strategies lead to volume meshes, from which the extraction of a surface structure relies on input sampling guarantees we cannot provide in practice, and ends up being even harder than the triangulation problem itself [Amenta et al. 2001]. Therefore, instead of resorting solely to such *geometric* interpretations of the mixture model, we use a triangulation approach that makes use of the *probabilistic* information encoded in the Gaussians.

Greedy triangulation. Our method is based on a simple greedy front-growing triangulation of point sets [Cohen-Steiner and Da 2004], which starts with an initial seed triangle and then iteratively advances its edge front by adding the next connected candidate triangle Δ_{ijk} that is optimal with respect to a certain plausibility grade $P(\Delta_{ijk})$. For point sets, Cohen-Steiner and Da [2004] measure this grade of a candidate triangle geometrically by the reciprocal radius of the smallest empty circumsphere. To avoid the appearance of

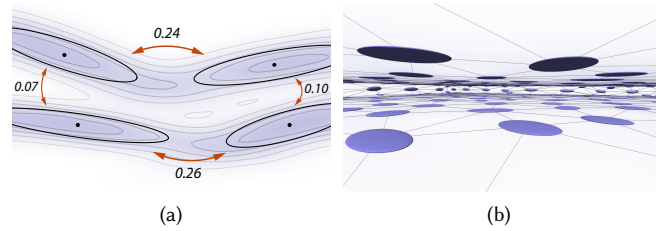


Fig. 8. (a) Thin sheets modeled by 4 anisotropic Gaussians. Simple distance-based criteria fail to express their topological connectivity. We therefore measure the overlap of their probability extent (numbers over red arrows). (b) Faithful triangulation of a cuboid mixture, where the distribution of Gaussian means does not meet the sampling criterion.

slivers, where a newly added triangle creates a fold, candidate triangles Δ_{ijk} are discarded if their dihedral angle β_{ijk} with a connected existing triangle falls below a certain threshold β_{min} .

Probabilistic plausibility grading. For triangulating a set of anisotropic Gaussians describing the probability distribution of surface points, such a Delaunay-based criterion is unsuitable, as the mere Euclidean distance between Gaussians is inconclusive about their topological connectivity. Figure 8a demonstrates this on the example of two twin density sheets, represented by 4 Gaussians. While two horizontally aligned Gaussians can have a larger distance between their means, they are still more plausible to be connected, since their individual probability distributions partially model the same region of the surface, i.e., they *overlap* to a considerable degree. Our intuition therefore is to use a grading that assesses exactly this overlap of their probability distributions. To this end, we use the Bhattacharyya coefficient [Kailath 1967]

$$BC_{ij} = \int_{\mathbb{R}^3} \sqrt{f(\mathbf{x}|\Theta_i) f(\mathbf{x}|\Theta_j)} dx \quad (18)$$

between two Gaussians Θ_i and Θ_j as a basic measure of connectivity. This coefficient quantifies the amount of overlap between two statistical populations, has a closed-form expression for Gaussians,

$$BC_{ij} = |\tilde{\Sigma}|^{-\frac{1}{2}} |\Sigma_i \Sigma_j|^{\frac{1}{4}} e^{-\frac{1}{8}(\mu_i - \mu_j)^T \tilde{\Sigma}^{-1}(\mu_i - \mu_j)}, \quad (19)$$

with $\tilde{\Sigma} = (\Sigma_i + \Sigma_j)/2$, and has a tractable range $0 < BC \leq 1$, where 1 indicates maximal overlap in case of coinciding distributions. The Bhattacharyya coefficients shown for pairs of Gaussians in Figure 8a indicate that they provide a robust measure of topological closeness even for difficult configurations like these thin sheets.

To grade the plausibility of a triangle connecting three Gaussians Θ_i , Θ_j and Θ_k , we request each respective pair of Gaussians to provide sufficient mutual overlap. We thus measure their probabilistic plausibility by

$$P_{prob}(\Delta_{ijk}) = BC_{ij} \cdot BC_{jk} \cdot BC_{ik} \quad (20)$$

and consider a candidate triangle valid only if all its three pairwise coefficients exceed a minimal threshold BC_{min} . A higher value for BC_{min} is more restrictive and skips triangles in areas where the common density overlap of its Gaussians is very low, indicating a higher probability for an actual hole in the surface, but might also leave more undesired holes in the triangulation, while smaller values

are more prone to produce incorrect triangulations. In practice, we found values around $BC_{min} = 0.01$ to produce best results.

In addition to the probabilistic grade, our intuition is that at similar probability overlaps, the candidate triangle with the larger dihedral angle is more plausible to provide a good triangulation. Therefore, we add a geometric weight

$$P_{\beta}(\Delta_{ijk}) = \beta_{ijk} / \pi \quad (21)$$

based on the smallest dihedral angle β_{ijk} between Δ_{ijk} and its neighboring triangles. Our final plausibility grade is thus given by

$$P(\Delta_{ijk}) = P_{prob}(\Delta_{ijk}) \cdot P_{\beta}(\Delta_{ijk}). \quad (22)$$

To initialize the triangulation, we use the seed triangle with the highest probabilistic plausibility P_{prob} amongst all candidates. We experienced a single seed triangle to be sufficient in our examples, however, several seeds can be used to provide increased robustness for larger, more complex models, especially if comprising several distinct parts. To avoid the appearance of slivers, we adopt the dihedral angle threshold β_{min} from Cohen-Steiner and Da [2004], and only consider triangles with $\beta_{ijk} > \beta_{min}$. As suggested in their original paper, we typically use $\beta_{min} = \pi/6$ for our models. The algorithm stops if no open boundary edge is left to advance, in which case a watertight mesh was produced, or if no valid candidate triangle satisfying the β_{min} and BC_{min} thresholds is left for triangulation.

Figure 8b shows the inside of a triangulated mixture representing a thin cuboid, where the average distance between neighboring Gaussians lies above the distance between its upper and lower faces.

7 EXPERIMENTS AND COMPARISONS

7.1 Surface Modeling

We will now analyze the capabilities of our surface representation for shape modeling, and investigate the effect of different covariance configurations in a covmesh on the resulting Gaussian-product subdivision surface (in the following denoted *GPS* for short).

Shape variability. Our surface definition allows modeling complex shapes using sparse control meshes. Figure 9 shows different shapes




Fig. 9. Shape variations of the limit surface of a cubic covariance mesh under different covariance combinations (blue ellipsoids) vs. standard Catmull-Clark limit surface (upper left).

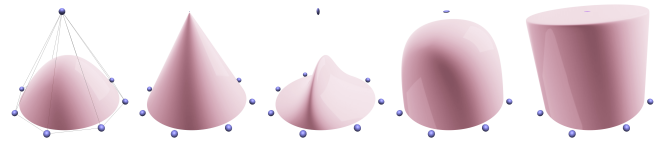


Fig. 10. GPS shape variations achieved by manipulating the covariance at the apex of a cone (f.l.t.r): ordinary Loop (isotropic), pointy apex (singular), concave ridge (vertically anisotropic), smooth plateau (horizontal anisotropic) and sharp plateau (horizontal flat).

produced by a cubic control covmesh with fixed vertex positions and varying covariances. The first cube at the top left shows the standard Catmull-Clark limit surface on an ordinary mesh (blue frame). The second cube represents a covmesh with uniform isotropic covariances at all vertices. As discussed in Eq. (14), this reduces to the linear case and thus produces the exact same surface. Further shape variations are achieved via different orientations, scales, and levels of anisotropy of the Gaussians, leading to complex shapes exhibiting high-curvature creases and various concavities.

For modeling subdivision surfaces using covmeshes, the fundamental difference to standard subdivision is the interpretation of a control vertex location. While in ordinary subdivision surfaces, cusps and creases are modeled via special constructs like semi-sharp creases, covariance meshes intrinsically encode such features through the shape of the associated Gaussian. Figure 1b shows a pointed tip at the end of a concave ridge modeled through a single covmesh vertex, a combination that cannot be achieved in such a compact way by standard subdivision methods without resorting to both explicit crease modification and control-mesh refinement. Figure 10 shows the wide range of shape variations that are possible by manipulating just the covariance at the apex of a cone. Moreover, covmeshes not only help to explore a wider space of subdivision surfaces, but can also be used to significantly reduce typical artifacts stemming from traditional subdivision schemes, such as the many bumps appearing in the helix in Figure 11.




Fig. 11. Thin quads in (a) create crease artifacts in subdivision surfaces (b), while in covmeshes, surface geometry is mostly encoded in the covariances, making its surface less susceptible to the mesh geometry (c).

Interactive modeling. The extra degrees of freedom induced by covariances offer great flexibility when designing local shape structures, making it possible to quickly create stylized shapes without adding additional control vertices. In Figure 12, we inferred Gaussians to a sparse control mesh of an axe and sharpened its blade by interactively flattening the two covariances at its cusps along

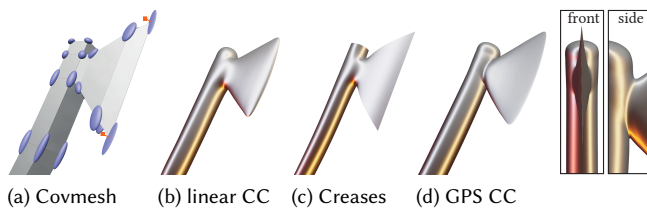


Fig. 12. Interactive modeling using linear and GP subdivision surfaces.

their smallest eigendirections (red handles). Supported by automatic inference, this process allows quickly generating interesting shapes, comparable to ordinary subdivision surfaces using crease edges.

7.2 Reconstruction

Building on a sparse probabilistic description of geometry, our representation offers new applications in efficient surface reconstruction.

For ordinary subdivision surfaces, the geometric quality depends solely on the quality of the control mesh, which is mostly the result of a careful, topology-aware modeling process. They are thus rarely applied to meshes resulting from a reconstruction process, unless their topology is suitably consolidated in a postprocess. Covmeshes, however, encode most of the local geometric information around a vertex in its covariance, allowing for more flexible subdivision-based surface modeling even in cases of problematic topology. A reconstruction process that preserves this covariance information along with the topological information of the mesh makes it thus feasible to apply our nonlinear subdivision even for meshes of coarse and irregular topology. Fig. 1c shows a sparse triangular covmesh reconstructed from a set of points by first converting them to a Gaussian mixture (see Appendix A) and then triangulating the Gaussians as described in Section 6.2. As shown in the closeups, the resulting GPS is able to model crisp features and rich detail, whereas an ordinary Loop subdivision surface defined on the same sparse triangle structure suffers from a significant loss of detail.

Reconstruction performance. To assess the quality and efficiency of a reconstructed covmesh for surface reconstruction, we compare its accuracy, speed and memory requirements against screened

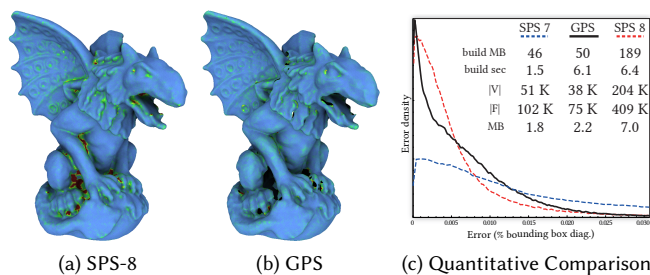


Fig. 13. Reconstruction error of screened Poisson (a) compared to the Gaussian-product surface of a reconstructed covmesh (b), along with error histograms and performance measures for comparable reconstructions (c).

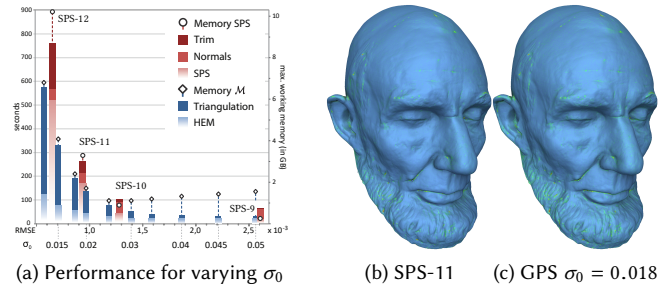


Fig. 14. (a) Timings for screened Poisson reconstruction (red) vs. covmesh computation (blue) for surfaces of different accuracy, controlled by octree depth and kernel size σ_0 , respectively. Dashed flags indicate the used working memory. (b) and (c) Surface error of two comparable reconstructions, where GPS exhibits superior performance values.

Poisson reconstruction [Kazhdan and Hoppe 2013]. To this end, we use a point scan of the Gargoyle model (303K points), which was virtually scanned to exhibit realistic scanner artifacts, holes and registration errors [Berger et al. 2013] and consolidated via a robust projection operator [Preiner et al. 2014], producing an unorganized point set exhibiting holes and realistic residual imperfections. We have produced two reconstructions of comparable quality and reconstruction time, shown in Fig. 13. After normal estimation for the input points, the screened Poisson surface (SPS) was reconstructed at octree depth 8, using code version 10.05. The resulting mesh is shown in Fig. 13a. For our covmesh reconstruction, we converted the point cloud to a mixture of 38K Gaussians, which was again triangulated to a covmesh. Note that this method does not require properly oriented surface normals. Fig. 13b shows the resulting Gaussian-product limit surface (GPS).

To measure the reconstruction error for each method, we computed the distance of each input point to the resulting surfaces. Surface colors depict the error values of the closest point. Fig. 13c plots the histogram of point errors for our GPS surface as well as SPS at depth 7 and 8. The inset table gives the used reconstruction time and working memory (*sec* and *build MB*) as well as the vertex and face counts and the memory footprint of the resulting mesh representation. Note that the additional covariance information requires another 6 floats per vertex. The graph shows that our reconstructed GPS surface (black) allows for a much more accurate representation than SPS at depth 7, while exhibiting only a slightly higher memory footprint. At the same time, it almost matches the quality of SPS at depth 8, while using a much more efficient representation, which overall requires only less than a third of the memory.

Fig. 14 compares the reconstruction performance of both SPS and GPS at the example of the Lincoln mask model, analyzing the trade-offs between quality and processing time/memory requirements. For Gaussian-product surfaces, the quality can be influenced by varying the kernel size σ_0 , which controls the bandwidth used for initializing the Gaussian mixture based on the input points (see Appendix A). A larger kernel produces mixtures of higher variance, resulting in a stronger mixture compression and thus a smoother surface. As before, at comparable RMSE values, GPS is faster, and for lower

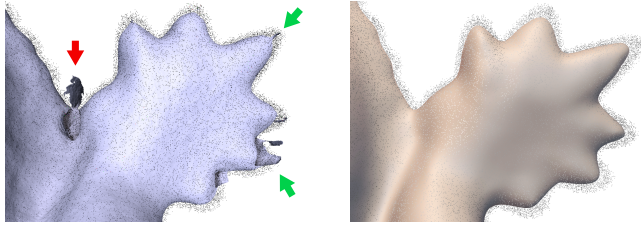


Fig. 15. Left: defects of ridge surfaces in high curvature regions, where a tip collapses to one sheet (green) or a spurious ridge merges with the main one (red). Right: Robust continuous surface provided by our surface model.

RSME values also requires less working memory. For higher kernel sizes, GPS working memory increases due to larger neighborhood queries, requiring larger spatial search structure caches. Note that RMSE values are only rough indicators of apparent visual quality, since they are not robust against strong noise and outliers. Finally, Fig. 15 demonstrates the defects of ridge-based reconstruction [Süßmuth and Greiner 2007] as discussed in Section 2.1, and the ridge approximation quality of our proposed surface model.

8 DISCUSSION

Limitations and future work. Our experiments on reconstruction do not overcome state-of-the-art techniques, but can be the basis of more advanced topological inference. For instance, a covariance mesh topology could be pre-defined using a low-resolution Poisson reconstruction, and its vertices then enriched with Gaussians computed from the original data. For interactive modeling, directly manipulating Gaussian covariances may at times be counter-intuitive to adjust to a particular aim. Sophisticated mappings of these parameters to intuitive interaction paradigms could be investigated, for instance, solving an inverse problem under user-prescribed features (e.g., sketch) to optimize the covariance mesh. Furthermore, the mapping we introduced in Section 4 allows for directly using the vast directory of tools already developed for linear subdivision, such as GPU real-time rendering [Nießner et al. 2012] or exact parametric limit evaluation [Stam 1998]. Finally, we plan to study how to simplify high-resolution meshes to low-resolution covariance meshes, reproducing features in real-time Gaussian-product subdivision.

Conclusion. Gaussians are naturally capable of modeling uncertainty in data and compactly representing large surface regions. Our key idea is a new interpolation between Gaussians based on a product formulation, which we use to define a family of novel nonlinear subdivision operators for meshes extended with per-vertex Gaussians, i.e., covariance meshes. We derive a mapping to a dual space where each of our nonlinear operators corresponds to a linear one, from which it inherits properties like smoothness. Using a covariance mesh, subdivision modeling can be performed both with vertex positions and covariances. Concavities, creases and curvature can be controlled using simple modifications of the covariances without requiring a denser control mesh, widening the space of limit surfaces. When reconstructed from points, our surfaces compete with state-of-the-art, e.g., screened Poisson surface reconstruction, while

being independent of orientable normal information and reducing the reconstruction time thanks to its sparsity.

ACKNOWLEDGMENTS

The Lincoln mask model was provided by the Smithsonian Institution’s 3D model repository. We thank Johannes Wallner for his valuable input on smoothness characteristics of subdivision surfaces, Jochen Süßmuth for providing his ridge reconstruction tool, and the anonymous reviewers for their helpful suggestions.

REFERENCES

- Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. 2001. The Power Crust. In *Proceedings of the Sixth ACM Symposium on Solid Modeling and Applications (SMA ’01)*. ACM, New York, NY, USA, 249–266.
- Matthew Berger, Joshua A. Levine, Luis Gustavo Nonato, Gabriel Taubin, and Claudio T. Silva. 2013. A Benchmark for Surface Reconstruction. *ACM Trans. Graph.* 32, 2, Article 20 (April 2013), 17 pages.
- Henning Biermann, Adi Levin, and Denis Zorin. 2000. Piecewise Smooth Subdivision Surfaces with Normal Control. In *Proc. SIGGRAPH*. 113–120.
- Wade Brainerd, Tim Foley, Manuel Kraemer, Henry Moreton, and Matthias Nießner. 2016. Efficient GPU Rendering of Subdivision Surfaces using Adaptive Quadtrees. *ACM Transactions on Graphics (TOG)* (2016).
- Max Budninskiy, Beibei Liu, Fernando de Goes, Yiyi Tong, Pierre Alliez, and Mathieu Desbrun. 2016. Optimal Voronoi Tessellations with Hessian-based Anisotropy. *ACM Trans. Graph.* 35, 6, Article 242 (Nov. 2016), 12 pages.
- Stephane Calderon and Tamy Boubekeur. 2014. Point Morphology. *ACM Transactions on Graphics (Proc. SIGGRAPH 2014)* (2014).
- Edwin Catmull and James Clark. 1978. Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design* 10, 6 (1978), 350 – 355.
- David Cohen-Steiner and Frank Da. 2004. A Greedy Delaunay-based Surface Reconstruction Algorithm. *Vis. Comput.* 20, 1 (April 2004), 4–16.
- Martin Danelljan, Giulia Meneghetti, Fahad S. Khan, and Michael Felsberg. 2016. A Probabilistic Framework for Color-Based Point Set Registration. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1818–1826.
- Arthur P. Dempster, Natalie Laird, and Donald Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Stat. Soc. Ser. B* 39, 1 (1977), 1–38.
- Tony DeRose, Michael Kass, and Tien Truong. 1998. Subdivision Surfaces in Character Animation. In *Proc. SIGGRAPH*. 85–94.
- Meenakshisundaram Gopi, Shankar Krishnan, and Cláudio T. Silva. 2000. Surface Reconstruction based on Lower Dimensional Localized Delaunay Triangulation. *Computer Graphics Forum* 19, 3 (2000), 467–478.
- Philipp Grohs. 2009. Smoothness of interpolatory multivariate subdivision in Lie groups. *IMA J. Numer. Anal.* 29 (07 2009).
- Philipp Grohs. 2010. A General Proximity Analysis of Nonlinear Subdivision Schemes. *SIAM J. Math. Analysis* 42 (01 2010), 729–750.
- Richard D. Hill and Steven R. Waters. 1987. On the cone of positive semidefinite matrices. *Linear Algebra Appl.* 90 (1987), 81 – 88.
- Hugues Hoppe, Tony DeRose, Tom Duchamp, Mark Halstead, Hubert Jin, John McDonald, Jean Schweitzer, and Werner Stuetzle. 1994. Piecewise Smooth Surface Reconstruction. In *Proc. SIGGRAPH*. 295–302.
- Uri Itai and Nir Sharon. 2013. Subdivision Schemes for Positive Definite Matrices. *Foundations of Computational Mathematics* 13, 3 (01 Jun 2013), 347–369.
- Bing Jian and Baba C. Vemuri. 2011. Robust Point Set Registration Using Gaussian Mixture Models. *IEEE Trans. Pattern Anal. Mach. Intell.* 33, 8 (Aug. 2011), 1633–1645.
- Thomas Kailath. 1967. The Divergence and Bhattacharyya Distance Measures in Signal Selection. *IEEE Trans. on Communication Technology* 15, 1 (February 1967), 52–60.
- Kęstutis Karčiauskas and Jörg Peters. 2007. Concentric Tessellation Maps and Curvature Continuous Guided Surfaces. *Comput. Aided Geom. Des.* 24, 2 (2007), 99–111.
- Kęstutis Karčiauskas and Jörg Peters. 2018. A New Class of Guided C2 Subdivision Surfaces Combining Good Shape with Nested Refinement. *Computer Graphics Forum* 37, 6 (2018), 84–95.
- Michael Kazhdan and Hugues Hoppe. 2013. Screened Poisson Surface Reconstruction. *ACM Trans. Graph.* 32, 3, Article 29 (July 2013), 13 pages.
- Adi Levin. 2006. Modified Subdivision Surfaces with Continuous Curvature. *ACM Trans. Graph.* 25, 3 (2006), 1035–1040.
- Ruosi Li, Lu Liu, Ly Phan, Sasakthi Abeysinghe, Cindy Grimm, and Tao Ju. 2010. Polygonizing Extremal Surfaces with Manifold Guarantees. In *Proc. of the 14th ACM Symposium on Solid and Physical Modeling (SPM ’10)*. ACM, New York, 189–194.
- Lars Linsen and Hartmut Prutzsch. 2003. Fan Clouds - An Alternative To Meshes. In *Geometry, Morphology, and Computational Imaging*, Tetsuo Asano, Reinhard Klette, and Christian Ronse (Eds.). Lecture Notes in Computer Science, Vol. 2616. Springer, Berlin, Germany, 39–57.

- Charles Loop. 1987. *Smooth Subdivision Surfaces Based on Triangles*. Ph.D. Dissertation.
- Matthias Niefner, Charles Loop, Mark Meyer, and Tony DeRose. 2012. Feature-adaptive GPU rendering of Catmull-Clark subdivision surfaces. *ACM Trans. Graph. (TOG)* 31, 1 (2012), 6.
- Umut Ozertem and Deniz Erdogmus. 2011. Locally Defined Principal Curves and Surfaces. *Journal of Machine Learning Research* 12 (2011), 1249–1286.
- Reinhold Preiner, Oliver Mattausch, Murat Arikan, Renato Pajarola, and Michael Wimmer. 2014. Continuous Projection for Fast L1 Reconstruction. *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH 2014)* 33, 4 (Aug. 2014), 47:1–47:13.
- Mael Rouxel-Labbé, Mathijs Wintraecken, and Jean-Daniel Boissonnat. 2016. Discretized Riemannian Delaunay triangulations. In *Proceedings 25th International Meshing Roundtable (IMR25)*. Elsevier, Washington DC, United States.
- Scott Schaefer, Etienne Vouga, and Ron Goldman. 2008. Nonlinear subdivision through nonlinear averaging. *Computer Aided Geometric Design* 25, 3 (2008), 162 – 180.
- Jos Stam. 1998. Exact Evaluation of Catmull-Clark Subdivision Surfaces at Arbitrary Parameter Values. In *Proc. of the 25th Annual Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. ACM, New York, NY, USA, 395–404.
- Jos Stam. 1999. Evaluation of Loop Subdivision Surfaces. In *SIGGRAPH '99 Course Notes*.
- Jochen Süßmuth and Günther Greiner. 2007. Ridge Based Curve and Surface Reconstruction. In *Proc. of the Fifth Eurographics Symposium on Geometry Processing (SGP '07)*. Eurographics Association, Aire-la-Ville, Switzerland, 243–251.
- Nuno Vasconcelos and Andrew Lippman. 1999. Learning Mixture Hierarchies. In *Advances in Neural Information Processing Systems 11*, M. J. Kearns, S. A. Solla, and D. A. Cohn (Eds.). MIT Press, 606–612.
- Amir Vaxman, Christian Müller, and Ofir Weber. 2018. Canonical Möbius Subdivision. *ACM Trans. Graph.* 37, 6, Article 227 (Dec. 2018), 15 pages.
- Johannes Wallner. 2014. On Convergent Interpolatory Subdivision Schemes in Riemannian Geometry. *Constructive Approximation* 40 (12 2014).
- Johannes Wallner and Nira Dyn. 2005. Convergence and C1 analysis of subdivision schemes on manifolds by proximity. *Computer Aided Geometric Design* 22, 7 (2005), 593 – 622.
- Johannes Wallner, Esfandiar Nava Yazdani, and Andreas Weinmann. 2011. Convergence and smoothness analysis of subdivision rules in Riemannian and symmetric spaces. *Advances in Computational Mathematics* 34, 2 (01 Feb 2011), 201–218.
- Andreas Weinmann. 2010. Nonlinear Subdivision Schemes on Irregular Meshes. *Constructive Approximation* 31, 3 (01 Jun 2010), 395–415.
- Xunnian Yang. 2016. Matrix Weighted Rational Curves and Surfaces. *Comput. Aided Geom. Des.* 42, C (2016), 40–53.
- Denis Zorin and Peter Schroder. 2000. Subdivision for Modeling and Animation. In *ACM SIGGRAPH Course*.

A FROM POINTS TO GAUSSIANS

We review the process of converting a given input point set $\mathcal{P} = \{p_i\}$ to a geometry-preserving sparse mixture of anisotropic Gaussians $\mathcal{M} = \{\Theta_s\}$ using a geometrically regularized variant of *hierarchical expectation maximization (HEM)* [Preiner et al. 2014]. First, each input point p_i is converted to a small Gaussian $\Theta_i = (w_i, \mu_i, \Sigma_i)$, either by computing anisotropic covariances via local PCA, or by using isotropic covariances $\sigma_0^2 I$, with σ_0 chosen according to the input point density. These Gaussians constitute an initial level-0 mixture, which is then hierarchically reduced by applying alternating EM-steps optimizing an objective log-likelihood function

$$\mathcal{L}_{\log}(\mathcal{M}'|\mathcal{M}) = \sum_i \log \sum_s w_s \mathcal{L}(\Theta'_s|\Theta_i).$$

Here, the likelihood of a component Θ'_s in the next level representing a component Θ_i in the current level is given by

$$\mathcal{L}(\Theta'_s|\Theta_i) = \left[f(\mu_i|\Theta'_s) \exp\left(-\frac{1}{2} \text{tr}(\Sigma_s^{-1} \Sigma_i)\right) \right]^{w_i|\mathcal{P}|}.$$

In the E-step, the convex responsibilities of each Θ'_s for the current Θ_i are defined as

$$r_{is} = w_s \mathcal{L}(\Theta'_s|\Theta_i) \Big/ \sum_{s'} w_{s'} \mathcal{L}(\Theta'_{s'}|\Theta_i).$$

In the M-step, these responsibilities are used to distribute the mass of the current Gaussians among the Gaussians of the next level \mathcal{M}' :

$$(w'_s, \mu'_s) = \sum_i (r_{is} w_i, \omega_{is} \mu_i), \quad \Sigma'_s = \sum_i \omega_{is} (\Sigma_i + \mu_i \mu_i^T) - \mu'_s \mu'_s{}^T$$

with weights $\omega_{is} = r_{is} w_i / \sum_{i'} r_{i's} w_{i'}$. To ensure convergence and avoid too strong degeneration of geometric features modeled by the mixture, responsibilities between components Θ_i and Θ'_s are clamped to zero if their Kullback-Leibler divergence $D_{KL}(\Theta_i || \Theta'_s)$ exceeds a certain threshold $\alpha^2/2$. A larger α allows merging more distant and dissimilarly oriented Gaussians and thus achieving stronger compressions. Values for α between 1.5 and 2.5 have proven to yield good compression rates without too strong geometric degeneration.

B DIFFERENTIAL GEOMETRIC ANALYSIS

We are interested in the differential geometric properties of our subdivision surface S , in particular, the analytic definition of the surface normal and curvature in any surface point $\mu \in S$. We derive these quantities from the geometric properties of the limit hypersurface S^* resulting from an associated linear subdivision scheme L . Tangent vectors for S^* are typically extracted through an eigenanalysis of the local subdivision matrix associated with L . Reformulating S^* locally as a parametric surface $q(u, v)$ [Stam 1998, 1999] also allows deriving second-order derivatives used for curvature analysis.

Tangent plane. Let $q(u, v) = (\check{q}, \bar{q})$ be a parametric hyperpoint on S^* , $q_u = (\check{q}_u, \bar{q}_u) = \frac{\partial}{\partial u} q$ a known tangent to q , and $(\mu, \Sigma) = F^{-1}(q)$ the surface point and associated covariance of S corresponding to q . Then the tangent μ_u in μ , mapping back the tangent direction of q_u to the probabilistic surface S , is given by

$$\begin{aligned} \mu_u &= \frac{\partial}{\partial u} F_{\mu}^{-1}(q) = \frac{\partial}{\partial u} ([\check{q}]^{-1} \bar{q}) \\ &= -[\check{q}]^{-1} [\check{q}_u] [\check{q}]^{-1} \bar{q} + [\check{q}]^{-1} \bar{q}_u = -\Sigma [\check{q}_u] \Sigma \bar{q} + \Sigma \bar{q}_u \\ &= \Sigma (\bar{q}_u - [\check{q}_u] \mu). \end{aligned}$$

Any pair of linearly independent hypertangents q_u, q_v can thus be mapped to surface tangents μ_u, μ_v , yielding an analytic expression of the surface normal in μ solely via known quantities in Q .

Curvature. Both the mean and the Gaussian curvature are entirely defined via the first and second fundamental forms of S in μ . While the former is determined by the tangents μ_u and μ_v , the latter also requires knowledge of the second-order derivatives in μ . Denote by $(q_{uu}, q_{uv}, q_{vv}) = \left(\frac{\partial^2}{\partial u^2}, \frac{\partial^2}{\partial u \partial v}, \frac{\partial^2}{\partial v^2} \right) q$ the known second-order derivatives of q . Then we have

$$\begin{aligned} \mu_{uu} &= \frac{\partial^2}{\partial u^2} F_{\mu}^{-1}(q) = \frac{\partial}{\partial u} \mu_u \\ &= \frac{\partial}{\partial u} ([\check{q}]^{-1}) (\bar{q}_u - [\check{q}_u] \mu) + \Sigma \frac{\partial}{\partial u} (\bar{q}_u - [\check{q}_u] \mu) \\ &= -\Sigma [\check{q}_u] \Sigma (\bar{q}_u - [\check{q}_u] \mu) + \Sigma (\bar{q}_{uu} - [\check{q}_{uu}] \mu - [\check{q}_u] \mu_u) \\ &= \Sigma (\bar{q}_{uu} - [\check{q}_{uu}] \mu - 2 [\check{q}_u] \mu_u), \end{aligned}$$

μ_{vv} is defined accordingly, and

$$\begin{aligned} \mu_{uv} &= \frac{\partial^2}{\partial u \partial v} F_{\mu}^{-1}(q) = \frac{\partial}{\partial v} \mu_u \\ &= -\Sigma [\check{q}_v] \Sigma (\bar{q}_u - [\check{q}_u] \mu) + \Sigma (\bar{q}_{uv} - [\check{q}_{uv}] \mu - [\check{q}_u] \mu_v) \\ &= \Sigma (\bar{q}_{uv} - [\check{q}_{uv}] \mu - [\check{q}_u] \mu_v - [\check{q}_v] \mu_u). \end{aligned}$$