

Surface Sculpting with Stochastic Deformable 3D Surfaces

Andrew Willis[†], Jasper Speicher[‡], David B. Cooper[†]
Brown University, Division of Engineering[†], Division of Arts[‡]

Abstract

This paper introduces a new stochastic surface model for deformable 3D surfaces and demonstrates its utility for the purpose of 3D sculpting. This is the problem of simple-to-use and intuitively interactive 3D free-form model building. A 3D surface is a sample of a Markov Random Field (MRF) defined on the vertices of a 3D mesh where MRF sites coincide with mesh vertices and the MRF cliques consist of subsets of sites. Each site has 3D coordinates (x,y,z) as random variables and is a member of one or more clique potentials which are functions of the vertices in a clique and describe stochastic dependencies among sites. Data, which is used to deform the surface can consist of, but is not limited to, an unorganized set of 3D points and is modeled by a conditional probability distribution given the 3D surface. A deformed surface is a MAP (Maximum A posteriori Probability) estimate of the joint distribution of the MRF surface model and the data. The generality and simplicity of the MRF model provides the ability to incorporate unlimited local and global deformation properties. Included in our development is the introduction of new data models, new anisotropic clique potentials, and cliques which involve sites that are spatially far apart. Other applications of these models are possible, e.g., stereo reconstruction.

1 Introduction

We define our shape model on an initial 3D mesh which consists of vertices, which are points lying on the original shape, and edges – connections between vertices. We then assign a MRF site to each vertex and define cliques which consist of one or more MRF sites. The MRF is specified as a Gibbs distribution characterized by cliques (subsets of sites) and clique potentials (clique energies), and these cliques and clique energies can be anisotropic, i.e., can be different in different directions and can be inhomogeneous, i.e., can vary over the surface. In the case of time varying models, the system state can in theory consist of the surface vertices at a succession of two or more instants (i.e., cycles). It is this fully 3D stochastic model for the representation and modification of surfaces and its ability to incorporate unlimited local and global deformation properties (either physically realizable or physically unrealizable) that is completely new, tremendously powerful and computationally fast (real time on a PC).

2 Mathematical Models

2.1 Surface Model

We assume that an initial shape provided is a triangular surface mesh with N vertices. Associated with vertex \mathbf{p}_i is a

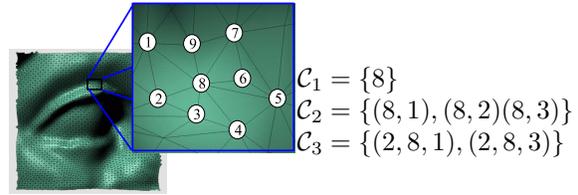


Figure 1: An example clique structure and 3 cliques.

site i . Hence, the MRF sites are labelled $1, 2, \dots, N$ where the i^{th} site has random variables $\mathbf{p}_i = (x_i, y_i, z_i)$. The vector $\mathbf{P} \equiv \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\}$ denotes all site variables. We denote the initial surface as the set of vertices at time $t = 0$ and indicate the time with a superscript, e.g., \mathbf{P}^0 denotes the vertices of the initial surface and $\mathbf{P}^t \equiv \{\mathbf{p}_0^t, \mathbf{p}_1^t, \dots, \mathbf{p}_N^t\}$ denotes the surface vertices at time t .

A clique is a subset of indices into the set of MRF sites, e.g., $\{1, 5, 3\}$, and has an associated clique potential function defined on the clique site variables $V_3(\mathbf{p}_1, \mathbf{p}_5, \mathbf{p}_3)$. \mathbf{P}_c denotes the set of site variables, i.e., vertices, in clique c . For our example, $\mathbf{P}_c = \{\mathbf{p}_1, \mathbf{p}_5, \mathbf{p}_3\}$. Cliques and their associated potential functions are distinguished by the cardinality of the set they define, i.e., single-vertex cliques, two-vertex cliques, three-vertex cliques, etc. We denote the sets of such cliques and clique potentials as $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \dots$, and V_1, V_2, V_3, \dots , respectively, and $\mathcal{C} = \bigcup_i \mathcal{C}_i$. For the simple mesh with indices shown in Figure (1), an example set of clique structures for site $i = 8$ and its neighbors are shown to illustrate the correspondence of clique definitions to their associated mesh vertices.

Cliques and clique potentials determine the forces among MRF site variables, i.e., the surface material behavior. The total energy of the MRF is described in terms of the sum of the defined clique potentials over all cliques (see §2.1.1).

$$U(\mathbf{P}|\mathbf{P}^0) = \sum_{c \in \mathcal{C}} V_c(\mathbf{P}_c|\mathbf{P}^0) \quad (1)$$

In general, potential functions are tailored to applications.

The surface probability distribution is modeled as a Gibbs Distribution which is a representation for an MRF. The Gibbs pdf $p_S(\mathbf{P}|\alpha)$ is specified as (2) (see §2.1.1).

$$p_S(\mathbf{P}|\alpha, \mathbf{P}^0) = \frac{1}{Z} \exp \{-U(\mathbf{P}|\alpha, \mathbf{P}^0)\} \quad (2)$$

2.1.1 Surface Model Example

A very simple surface model has the following cliques containing the i^{th} site: $\mathcal{C}_1 = \{i\}$ and $\mathcal{C}_2 = \{(i, j_1), (i, j_2), (i, j_3), \dots\}$ where j_k denotes the label of a site in a clique with site i , and the j_k sites are mesh neighbors of site i , i.e. are connected by a mesh edge.

The site clique potentials are $V_1(\mathbf{p}_i) = \alpha_1 \|\mathbf{p}_i - \mathbf{p}_i^0\|^2$ and $V_2(\mathbf{p}_i, \mathbf{p}_j) = \alpha_2 \|\mathbf{p}_i - \mathbf{p}_i^0 - \mathbf{p}_j + \mathbf{p}_j^0\|^2$. Here the parameters are the same for all sites. The potential V_1 is the squared distance from the deformed vertex \mathbf{p}_i to the initial vertex \mathbf{p}_i^0 and is a function of this local deformation from the surface. V_2 is the squared distance between the surface deformations at sites i and j and is a function of the relative positions of vertices at sites i and j before and after the deformation. α_1 and α_2 are weighting coefficients controlling the contributions of V_1 and V_2 to the total MRF energy. We denote the vector of clique potential coefficients as $\boldsymbol{\alpha}$, here $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)$.

2.2 Data Model

The data is modeled by (3).

$$p_{\mathcal{D}|\mathcal{S}}(\mathbf{D}|\mathbf{P}, \boldsymbol{\beta}) \quad (3)$$

This is a probability density function for the input data vector \mathbf{D} given the surface shape \mathbf{P} and parameters $\boldsymbol{\beta}$. Here, the components of $\mathbf{D} \equiv \{\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_M\}$ are the data points and $\boldsymbol{\beta}$ is a vector of parameters having values specified by the user.

2.2.1 Data Model Example

There are a number of data distribution models that come to mind as appropriate. Ideally, the data should lie on the surface we wish to represent. However, because of noise in the data generation process, each data point is a noisy perturbation of a point on the desired surface: but a perturbation of which surface point? One approach we and others have taken [2] is to assume an a-priori distribution for N' points on the surface from which the data point could have occurred. If the perturbation distribution is isotropic Gaussian having 3x3 covariance matrix $\sigma^2 \mathbf{I}$, then the conditional pdf of data-point \mathbf{d}_m given the surface \mathcal{S} is equation (4).

$$p(\mathbf{d}_m|\mathbf{P}) = \frac{1}{N'} \sum_{i \in N'} \frac{1}{(2\pi)^{\frac{3}{2}} |\sigma^2 \mathbf{I}|^{\frac{1}{2}}} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{d}_m - \mathbf{p}_i\|^2\right) \quad (4)$$

Note that $p(\mathbf{d}_m|\mathbf{P})$ is then a mixture of N' Gaussians, but it is only the Gaussians located at the vertices that are near \mathbf{d}_m that contribute significantly to this conditional probability. If it is assumed that the M data points are generated as independent random vectors, then the conditional probability $p(\mathbf{D}|\mathbf{P})$ is the product of the conditional probabilities for the individual data points, hence, a product of mixtures of Gaussians. Note, the pdf for data point \mathbf{d}_i is characterized by the two parameters N' and σ^2 which control the weight to be associated with a point and the extent of the region on the surface that is attracted by the data point. Hence, in our general framework from §2.2, $\boldsymbol{\beta} = (N', \sigma^2)$.

2.3 Deformation

Surface deformations occur by moving the mesh vertices \mathbf{P}^0 to a modified vertex vector \mathbf{P}^1 where \mathbf{P}^1 is:



Figure 2: Surface interpolations (see §2.3 for explanation).

$$\mathbf{P}^1 = \arg \max_{\mathbf{P}} \ln (p_{\mathcal{S}}(\mathbf{P}|\boldsymbol{\alpha}, \mathbf{P}^0) p_{\mathcal{D}|\mathcal{S}}(\mathbf{D}|\mathbf{P}, \boldsymbol{\beta})) \quad (5)$$

Note, $p_{\mathcal{S}}(\mathbf{P}|\boldsymbol{\alpha}, \mathbf{P}^0) p_{\mathcal{D}|\mathcal{S}}(\mathbf{D}|\mathbf{P}, \boldsymbol{\beta})$ is the joint pdf of \mathbf{D} and \mathbf{P} and we denote it

$$p_{\mathcal{D}, \mathcal{S}}(\mathbf{D}, \mathbf{P}|\boldsymbol{\alpha}, \mathbf{P}^0, \boldsymbol{\beta}) \quad (6)$$

\mathbf{P}^1 , the \mathbf{P} that maximizes (6), is the so call MAP (maximum a posteriori probability) estimate of \mathbf{P} given the data \mathbf{D} . Figure (2a) shows a surface mesh obtained from a laser scanner. The mesh structure is irregular and contains large holes and tears. Figures (2b) and (2c) show our low resolution and fine resolution MRF mesh interpolations where the mesh vertices from Fig. (2a) are data and a plane is the initial MRF shape, \mathbf{P}^0 . The meshes produced do not have holes and are more regular. Our interpolation uses the surface model from §2.1.1 and the data model from §2.2.1.

2.3.1 Optimization and an Example

The joint pdf (6) has an energy (1) defined for all possible positions of the mesh vertices, i.e. values which the site (x, y, z) random variables may assume. Optimization corresponds to maximizing the joint pdf (6) by applying any standard non-linear maximization method such as gradient ascent.

For optimization, we compute the gradient of (5) which is $\mathbf{f}_g(\mathbf{P}) = \nabla_{\mathbf{P}} \ln(p_{\mathcal{S}}(\mathbf{P}|\boldsymbol{\alpha}, \mathbf{P}^0)) + \nabla_{\mathbf{P}} \ln(p_{\mathcal{D}|\mathcal{S}}(\mathbf{D}|\mathbf{P}, \boldsymbol{\beta}))$ where $\nabla_{\mathbf{P}}$ indicates that we are differentiating with respect to the vector of site variables \mathbf{P} . Hence, the clique potentials and data pdf combine to generate a motion vector field for the MRF sites, \mathbf{f}_g ; *this may be interpreted as a generalized force*. \mathbf{f}_g is the resultant of two distinct forces : (1) $\nabla_{\mathbf{P}} \ln(p_{\mathcal{S}}(\mathbf{P}|\boldsymbol{\alpha}, \mathbf{P}^0))$ is the force between MRF sites, which models the intrinsic surface behavior (e.g., flexibility); (2) $\nabla_{\mathbf{P}} \ln(p_{\mathcal{D}|\mathcal{S}}(\mathbf{D}|\mathbf{P}, \boldsymbol{\beta}))$ is the force exerted on the MRF sites by the data. Hence, the cliques and their associated clique potentials exert virtual *force fields* on the MRF sites.

3 Time Varying Surface Behaviors

In §2 the time index for surfaces could assume only two values $t = (0, 1)$, here we simply allow the index t to vary $t \in [0, \infty)$ for the site random variables, \mathbf{P}^t , the potential weighting coefficients, $\boldsymbol{\alpha}^t$, and the data, \mathbf{D}^t . In this case, (7)

denotes a time-varying pdf for 3D surfaces, or, alternatively, can be thought of as a pdf having indexing parameters (i, t) for site location and time.

$$p_S(\mathbf{P}|\boldsymbol{\alpha}^t, \mathbf{P}^t) = \frac{1}{Z} \exp \{-U(\mathbf{P}|\boldsymbol{\alpha}^t, \mathbf{P}^t)\} \quad (7)$$

3.1 Elasticity & Plasticity

Assume that the surface vertices from the previous time step, \mathbf{P}^{t-1} , remain fixed at the initial surface i.e., $\mathbf{P}^{t-1} = \mathbf{P}^0 \forall t$. If the clique potentials for the i^{th} MRF site are all convex quadratic functions symmetric about $\mathbf{p}_{i,0}$, then $p_S(\mathbf{P}|\boldsymbol{\alpha}^{t-1}, \mathbf{P}^{t-1})$ defines a Gaussian distribution on the MRF sites $\mathcal{N} \sim (\mathbf{P}^0, \Sigma)$ where Σ depends on the weighting coefficients $\boldsymbol{\alpha}^t$. The maximum is located at the mean, hence, $\arg \max_{\mathbf{P}} p_S(\mathbf{P}|\boldsymbol{\alpha}^t, \mathbf{P}^{t-1}) = \mathbf{P}^0$. Note that the clique potentials from §2.1.1 are a special case of this model where $\boldsymbol{\alpha}^t = \{\alpha_1^0, \alpha_2^0\} \forall t$. In this case, the surface will return to its original vertex positions in the absence of external forces from data. This behavior for the MRF is analogous to that of an elastic surface deformation for physically-based models.

Using the same Gaussian MRF site distributions we consider a new data set from time t . We optimize (8) to obtain the deformed MAP surface \mathbf{P}^t .

$$\mathbf{P}^t = \arg \max_{\mathbf{P}} \ln (p_S(\mathbf{P}|\boldsymbol{\alpha}^{t-1}, \mathbf{P}^{t-1})p_{\mathcal{D}|\mathcal{S}}(\mathbf{D}^t|\mathbf{P}, \boldsymbol{\beta})) \quad (8)$$

To make the deformation plastic, we update the mean of the MRF site distributions at the end of the time step i.e., $\mathcal{N} \sim (\mathbf{P}^t, \Sigma)$, and remove the data set \mathbf{D}^t . Now our MLE surface in the absence of data is the deformed surface \mathbf{P}^t . This behavior for the MRF is analogous to that of a plastic deformation for physically-based models.

4 Example : 3D Sculpting

Interactive sculpting of free-form 3D shapes has been an active area of research in the computer graphics community for quite some time. Existing systems provide examples of many different surface representations and work continues to develop their utility for *intuitive* sculpting [3, 1, 6, 7]. We present here our interactive sculpting system and explain how the surface model is represented to a virtual sculptor and provides a heretofore unexplored analogy to facilitate controlled surface deformation.

Our basic sculpting system uses three clique structures : $\mathcal{C}_1 = \{i\}$, $\mathcal{C}_2 = \{(i, j_1), (i, j_2), \dots, (i, j_K)\}$ and $\mathcal{C}_{K+1} = \{(i, j_1, j_2, \dots, j_K)\}$ where j_k denotes a mesh neighbor (see §2.1.1 for clarification) and K denotes the number of sites that are mesh neighbors of site i . The clique potentials and their associated cliques are provided in Table (1).

Each of the cliques and associated clique potentials in Table (1) is assigned a descriptive phrase symbolizing how the clique potential influences interpolations of sculpted

| Description | c | $V_c(\mathbf{P}_c)$ |
|------------------------|---------|---|
| Mean Field | 1 | $\ \mathbf{p}_i - \mathbf{p}_i^{t-1}\ ^2$ |
| Mean Smoothing | 2 | $\ \mathbf{p}_i - \mathbf{p}_i^{t-1} - \mathbf{p}_{j_k} + \mathbf{p}_{j_k}^{t-1}\ ^2$ |
| Edge Preserving | 2 | $\ \mathbf{p}_i - \mathbf{p}_{j_k}\ - \ \mathbf{p}_i^{t-1} - \mathbf{p}_{j_k}^{t-1}\ $ |
| Surface Flattening* | $K + 1$ | $\ \kappa_{\nabla} \ ^2$ |
| Curvature Preserving* | $K + 1$ | $(\kappa_{\nabla} - \frac{1}{K} \sum_{k=1}^K \kappa_{\nabla}(j_k))^2$ |
| Symmetry (see Fig. 4a) | 2 | $\ \mathbf{p}_i \cdot \mathbf{n}_p\ ^2 - \ \mathbf{A}\mathbf{p}_j \cdot \mathbf{n}_p\ ^2 + \ \mathbf{p}_i \times \mathbf{n}_p\ ^2 - \ \mathbf{A}\mathbf{p}_j \times \mathbf{n}_p\ ^2$ |

Table 1: Sculpting Clique Potentials.

*For potentials involving curvature, i.e., κ , we have applied the method from [5] to obtain principle curvatures and principle directions: $(\bar{\kappa}_1, \bar{\kappa}_2)$. In general, the notation κ_{∇} denotes the curvature in the direction of the vector ∇ .

data sets. For example, the *Mean Field* description is assigned the clique potential $\|\mathbf{p}_i - \mathbf{p}_i^{t-1}\|^2$ which increases as a quadratic function of the displacement of the vertex. Hence, this potential encourages points to remain at their original or *mean* position, i.e., \mathbf{p}_i^{t-1} for the i^{th} vertex. Hence, we can think of each clique and its associated clique potential as specifying a *specific type of force field* for the MRF sites which is identified by our descriptive phrase. For example, we can refer to the mean force field as those forces which move vertices from their current (x, y, z) position to their original (x, y, z) position.

The relative strength of each force field type causes the surface to interpolate the sculpted data differently. Hence, using this force-field analogy, the sculptor can modify force field strengths, i.e., clique potential function coefficients, to intuitively control the surface. The user controls the 6 potential function coefficients, $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6)$, for each of the rows from Table (1) which, combined with the 2 data model parameters $\boldsymbol{\beta} = (N', \sigma^2)$ from §2.2.1, makes 8 user-controlled parameters in total. Their geometric meanings are given in §4.1.

The sketched data is provided at different time instances and is consequently a situation where the time varying MRF models of §3 are appropriate. The sculptor is empowered to deform the shape via a single custom built device which can interactively (add/remove) data, control the data and material parameters, and define primitive cliques. The sculpting process is then described in terms of cycles. At the beginning of each cycle the sculptor may do any of the following functions:

- Introduce a data set, \mathbf{D}^t , to the system for surface deformation with varying data types:

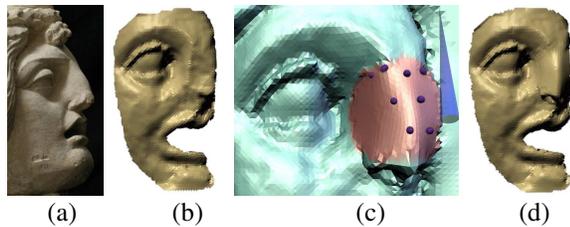


Figure 3: Figures (a-d) show the reconstruction of a nose on a damaged archaeological sculpture discovered in Petra, Jordan. (a) an image of the archetype nose used for reconstruction (b) an interpolation of the damaged mask data (c) interactive modification for rebuilding the missing nose using anisotropic clique potentials and only a few data points shown in blue (d) the reconstructed mask.

- 3D space curves, i.e., 3D points obtained from 3D tracker integrated into a sculpting pen,
- Directional feature enhancements, i.e., indicate areas to apply anisotropic diffusion,
- A cloud of 3D data points from a template.

- Make the current deformation plastic (see §3.1).
- Remove previous datasets.
- Change deformation parameters of the system.
- Change the clique structure of the system.
- Remesh regions of the surface using the method in [4].

At the end of the cycle we compute a solution to the optimization problem (5) as described in §2.3.1. The system is implemented in Java, and performs in real time for deformations involving in excess of 5000 vertices.

4.1 Sculpting Potentials

For each of the potentials listed in Table (1) the time index $t - 1$ refers to the time of the last plastic deformation specified by the user, see §3.1. The following list provides a geometric interpretation of the potentials listed in Table (1) (see §4 for Mean Field).

Mean Smoothing Strengthening this force will encourage neighboring MRF sites to have similar displacements from their original position which makes the surface smoothly interpolate the data.

Edge Preserving Strengthening this force will encourage the mesh edges to remain fixed which makes the surface material rigid.

Surface Flattening Strengthening this force encourages the surface region to be linear in the direction of $\vec{\nabla}$.

Curvature Preserving Strengthening this force encourages the surface region to preserve curvature, i.e., mesh neighbors are encouraged to have the same curvature in the direction of $\vec{\nabla}^\perp$.

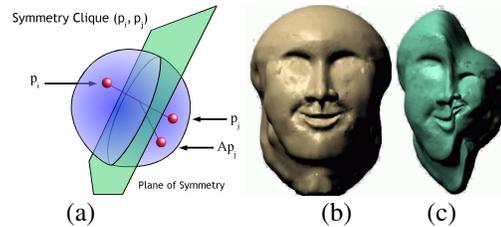


Figure 4: Symmetry Cliques : Spatially distant sites are associated by their height and (x, y) positions relative to a user-defined plane as shown in (a). In (b), we use the symmetry clique potential from Table (1) setting $\mathbf{A} = \mathbf{I}$ to create an original sculpting of a head from an initial capsule shape. In (c), we can change the symmetry clique potential to be non-isometric scaling : $diag(\mathbf{A}) = (1.25, 1.25, 1)$.

Symmetry Strengthening this force encourages the surface to be symmetric about a predefined plane as shown in Fig. (4).

5 Conclusions

We have presented a novel deformable surface model and demonstrated a sculpting system which intuitively uses this model and interpreted it as sculpting via virtual force-fields. Simple clique forces provide good surface interpolations as shown in Fig. (2) and more complex forces allow for *easier sculpting of useful features* as shown in Figs. (3) and (4). The power available due to the generality, versatility, and computational simplicity of the model makes it perfect for building arbitrary free-form shapes where the user provides data to the computer which is interpolated by the MRF model at an interactive rate. This material is based upon work supported by the National Science Foundation Grant No. 0205477.

References

- [1] J. Baerentzen et al. Volume Sculpting Using the Level-Set Method. In *Proc. of SMI*, pages 175–182, 2002.
- [2] R. Bolle and D. B. Cooper. On optimally combining pieces of information, with application to estimating 3-D complex-object position from range data. *PAMI*, 8(5):619–638, 1986.
- [3] H. Du and H. Qin. Direct Manipulation and Interactive Sculpting of PDE Surfaces. *Computer Graphics Forum*, 19(3), 2000.
- [4] L. Kobbelt, T. Bareuther, and H. Seidel. Multiresolution Shape Deformations for Meshes with Dynamic Vertex Connectivity. In *Proc. of EUROGRAPHICS*, pages 249–260, 2000.
- [5] G. Taubin. Estimating the Tensor Of Curvature Of A Surface From A Polyhedral Approximation. In *Proc. of ICCV*, pages 902–907, 1995.
- [6] W. Welch and A. Witkin. Free Form Shape Design Using Triangulated Surfaces. *SIGGRAPH*, 28:247–256, 1994.
- [7] M. Zwicker, M. Pauly, O. Knoll, and M. Gross. Pointshop 3D: An Interactive System for Point-Based Surface Editing. *SIGGRAPH*, pages 322–329, 2002.