

Investigations of Tensor Voting Modeling

Joanna Beltowska†‡
jb643@drexel.edu

Ken Museth‡
museth@acm.org

David Breen†
david@cs.drexel.edu

†Computer Science Department
Drexel University
Philadelphia, PA USA

‡Science and Technology Department
Linköping University
Norrköping, Sweden

ABSTRACT

Tensor voting (TV) is a method for inferring geometric structures from sparse, irregular and possibly noisy input. It was initially proposed by Guy and Medioni [Guy96] and has been applied to several computer vision applications. TV generates a dense output field in a domain by dispersing information associated with sparse input tokens. In 3-D this implies that a surface can be generated from a set of input data, giving tensor voting a potential application in surface modeling. We study the tensor voting methodology in a modeling context by implementing a simple 3-D modeling tool. The user creates a surface from a set of points and normals. The user may interact with these tokens in order to modify the surface. We describe the results of our investigation.

Keywords

implicit surface modeling, surface reconstruction, tensor voting

1. INTRODUCTION

As the use of implicit 3D representations gains popularity, the need for modeling software that provides implicit model editing capabilities increases. New surface editing approaches are continually being explored and much research is being conducted to find ways to interactively modify implicit models. Towards these ends we investigated tensor voting as a possible technology that could provide a new and interesting approach for editing implicit models.

Tensor voting is a method for grouping geometric features [Med00]. It can be used to generate surfaces from a sparse set of possibly noisy and irregular input data, and therefore may provide novel editing capabilities within a 3-D modeling context. The goal of our work was to investigate TV as a technique for interactive surface modeling. We were interested in determining if TV can be used to model simple objects, edit them interactively and control their shape via "input tokens." To achieve this, we developed a TV modeling system (TVMS) based on an already existing TV framework and conducted

experiments to evaluate TV as a 3-D modeling tool. Spheres were used to examine the parameters of TVMS. Two models were sculpted in order to evaluate the potential capabilities and limitations of the TV modeling process.

Related Work

We looked to other novel modeling systems for ideas and inspiration. In general, these are systems that model 3-D objects with geometry other than polygonal meshes. The Vector Field Based Shape Deformations of von Funck et al. [vF06, vF07] is an interactive method for shape modeling with volume preservation. Guo et al. [Guo04] use scalar-fields to drive a point-set surface editing process. The point-sample-based shape modeling presented by Pauly, Keiser, Kobbelt and Gross [Pau02] is a hybrid system which couples the benefits of implicit and parametric surface representations, using the implicit surface definition of moving least squares projection along with point-based visualization. Blinn [Bli82] presents a point-based modeling method where points emit a radial Gaussian field. An iso-surface away from the input points is then extracted from the cumulative field. Pointshop 3D is a framework developed by Zwicker et al. [Zwi02] that transfers the functionality of 2D image editing operations to 3D. Dewaele and Cani [Dew04] present a virtual clay modeling approach, which implements interactive shape modeling capabilities which may be based on haptic feedback. Museth et al. [Mus02, Mus05] developed interac-

Permission to make digital or hard copies of all or part 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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Copyright UNION Agency - Science Press, Plzen, Czech Republic.

tive editing operators for a relatively new type of implicit surface, level set models.

In our approach, as in many of the approaches described above, the object is represented using control points, which can be modified by the user in order to edit the surface. By using the TV methodology in a novel context, we have employed the ideas of local editing operators, point-based surface representation and a physical analogy for modeling surfaces.

2. TENSOR VOTING

This section provides an introduction to the Tensor Voting methodology. More details about the methodology can be found in [Guy97, Tan98, Med00, Tan01, Ton04].

TV has its background in early computer vision problems where the available data often is sparse and noisy, making it difficult to extract relevant information and structures. TV identifies local feature descriptions by spreading the information associated with shape-related input within a neighborhood while enforcing a smoothness constraint. This process refines the information and accentuates local features. By doing so, coherent, locally smooth, geometric features are defined and noise is discarded. Each data point communicates its information in a neighborhood through a voting process. The more information that is received at each data point, the stronger is the likelihood of a geometric feature being present at a certain location. This likelihood is expressed through a confidence measure, saliency, which is used in the feature extraction process.

TV is based on two elements; a data representation, which is obtained by means of tensor calculus, and communication of data through linear voting, a process similar to linear convolution. The input elements, referred to as input tokens, are encoded into tensor form and communicate their information to their neighboring tokens via pre-calculated tensor voting fields. After this initial voting step, each token has its confidence and surface orientation encoded into a generic second order symmetric tensor. The tokens vote a second time to propagate their information throughout a neighborhood. The result is a dense tensor field which assigns a measure of confidence and saliency to each point in the domain. This dense map is decomposed into three dense maps, each representing a geometric feature (junctions, curves or surfaces), which are analyzed during feature extraction. TV can be generalized to N-D [Tan01]. The 3-D case, specifically surface voting, is sufficient for our needs and will therefore be the focus of this paper.

Tensor Representation

Diagonalizing a second order symmetric tensor, which can be represented by a 3×3 matrix, produces the associated characteristic equation. Solving this equation leads to a representation based on the eigenvalues $\lambda_1, \lambda_2, \lambda_3$ (in decreasing order) and the associated eigenvectors $\hat{e}_1, \hat{e}_2, \hat{e}_3$ of the tensor. A second order symmetric tensor may be graphically represented as an ellipse in 2-D or an ellipsoid in 3-D. The eigenvalues describe the general size and shape of the ellipsoid and the eigenvectors describe its principal directions. Because of the properties of second order symmetric tensors, the eigenvalues are real and positive, or zero, and the eigenvectors form an orthonormal basis. A tensor can be decomposed into three components defined by

$$\mathbf{T} = (\lambda_1 - \lambda_2)\hat{e}_1\hat{e}_1^T + (\lambda_2 - \lambda_3)(\hat{e}_1\hat{e}_1^T + \hat{e}_2\hat{e}_2^T) + \lambda_3(\hat{e}_1\hat{e}_1^T + \hat{e}_2\hat{e}_2^T + \hat{e}_3\hat{e}_3^T). \quad (1)$$

The first term corresponds to a *3-D stick tensor*, which implies a surface patch with normal \hat{e}_1 . The second term corresponds to a *3-D plate tensor* and implies a curve or surface intersection with tangent \hat{e}_3 that is perpendicular to the plane defined by \hat{e}_1 and \hat{e}_2 . The last term corresponds to a *3-D ball tensor* and implies a structure with no orientation preference.

Tensor Communication

The voting process is similar to convolution with the difference that convolution produces scalar values and tensor voting produces tensors. Voting kernels which encode certain constraints such as smoothness and proximity are used. These voting kernels are continuous tensor fields that assign a value to every point within the domain. Any voting kernel, regardless of dimension, can be derived from the 2-D stick tensor, which is therefore referred to as the fundamental 2-D stick voting field (VF) and is presented in Figure 1 (left). For its derivation, see [Med00] and [Tan01]. Given an input point and normal N at O , the most likely normal at P to a curve passing through O and P is defined by the osculating circle connecting O and P , because it keeps curvature constant. The 3-D VF is produced by rotating the 2-D VF around N .

The magnitude of the field is described by a decay function, which is expressed in spherical coordinates and is presented in Figure 1 (right),

$$DF(\gamma, \varphi, \sigma) = e^{-\left(\frac{\gamma^2 + c\varphi^2}{\sigma^2}\right)}. \quad (2)$$

γ is the arc length of the curve OP , φ the curvature, c the curvature scale factor and σ the scale of analysis. c

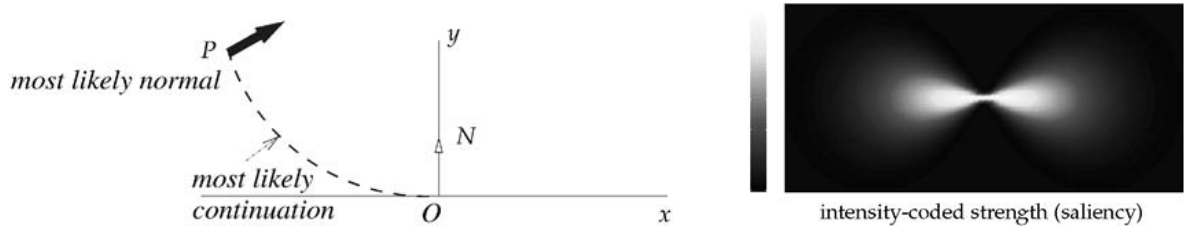


Figure 1: (left) The design of fundamental 2D stick voting field. (right) Magnitude (saliency) of the fundamental 2D stick voting field. From [Tan01].

provides additional control on the influence of curvature. σ determines the size of the voting neighborhood and is the only free parameter adjusted by a user. In practice the range of influence of a particular VF is demarcated once its magnitude drops below some value, e.g. 1%. Therefore the value of σ determines the extent of the voting field.

The field has an analogy with particle physics, motivating the choice of a Gaussian function for the energy decay function. The voting kernel shows the same behavior as an anisotropic energy field, where an emitting particle P_s radiates energy in all directions and a receiving particle P_r receives energy from all directions, but with a preference for a straight path. Therefore received energy is inversely proportional to distance traveled and curvature, producing a loss of energy along curved paths.

After voting, the saliency tensor $\mathbf{T}(i, j, k)$ at each location (i, j, k) is the tensor sum of the contributions from each voting token within range,

$$\mathbf{T}(i, j, k) = \sum_{m, n, o} V(\mathbf{S}(m, n, o), ((i, j, k) - (m, n, o))) \quad (m, n, o) \in N(i, j, k), \quad (3)$$

where \mathbf{S} is the tensor associated with the voting token at (m, n, o) , $N(i, j, k)$ returns the neighboring locations of (i, j, k) and V is a function that returns a tensor field corresponding to the ball, stick and surface voting kernel components of \mathbf{S} using the vector $(i, j, k) - (m, n, o)$, which is expressed in the $(\hat{e}_1, \hat{e}_2, \hat{e}_3)$ coordinate system of \mathbf{S} .

Feature Extraction

After the voting process, each point in the domain has been assigned a second order symmetric tensor that estimates the structure of the feature type(s) and the associated saliency. Thus, information is collected at each location in the domain, building a saliency map for each feature type. Salient features can be located by finding local extrema in this dense tensor map, which in 3-D is decomposed into three dense vector maps, expressing the occurrence of junctions, curves and surfaces in the domain. For

each feature type, the dense tensor map can be broken into a scalar field, expressing the saliency (strength) of the tensor field at any given point, and a vector field, expressing the direction of the feature corresponding to the saliency value. Thus, at any given location there is a scalar and a unit vector $\langle s, \hat{n} \rangle$ present. These dense saliency and vector maps are then used for vote interpretation.

Since we are only interested in the surface features produced by tensor voting, we only examine the Surface Map $\langle s, \hat{n} \rangle$ extracted from the tensor field [Tan01]. For this case,

$$s = (\lambda_1 - \lambda_2) \quad \hat{n} = \hat{e}_1, \quad (4)$$

where s is the surface saliency and \hat{n} is the normal to the most likely surface at each point in space.

The most likely surface produced by a set of input tokens (points and normals) is an *extremal surface* embedded in the Surface Map. A point is on the extremal surface when the saliency s is extremal in the normal direction \hat{n} at that point, i.e. $ds/d\hat{n} = 0$. Reformulating this equation, the surface may be found by identifying zeros of the the scalar field $\hat{n} \cdot \nabla s$, defined as the g field. For practical reasons the Surface Map is calculated on a regular grid and the zero crossings of the extremal surface are identified along grid lines.

3. TENSOR VOTING MODELING SYSTEM

The tensor voting modeling system (TVMS) is based on the TV3D framework for tensor voting developed by the USC Computer Vision Group¹, and uses QSplat² [Rus00] for point-based visualization. It has been developed on a MacBook Pro with a 2.2 GHz Intel Core 2 Duo processor. Functionality has been added to both frameworks and only surface voting is used from the TV3D framework. Since our goal is to extract a surface from a given input, we have omitted the point and curve voting functionality of TV3D.

¹ <http://www.cs.ust.hk/~cstws/research/TensorVoting3D>

² <http://graphics.stanford.edu/software/qsplat>

TV3D uses an input of oriented or unoriented points, curvels and surfels. In TVMS, the input is restricted to points with associated normals. While a surface may be constructed from unoriented points only, we have chosen to include direction information in the input dataset. The reason for this is that it gives more accurate results, allows us to omit the first voting step, which estimates and assigns a direction to all unoriented input tokens, and finally, it provides the user with an additional tool for manipulating the resulting surface.

These points can be interactively translated and their corresponding normals rotated, as well as added to and removed from the scene. Translation is done by first rotating the normal of a selected point and then translating the point along the normal. This is a simple yet intuitive interaction that meets our needs. The user can either create an entirely new model on his or her own, or read an existing model from a file. A new token is inserted by clicking between two existing tokens. The new token is inserted at the position halfway between the two tokens closest to the user's mouse click, with the new normal calculated as the average of the normals of the two neighboring tokens. The user can update the surface of the model that (s)he is working on at anytime. Each input token has a position, a normal, a weight and a local σ , which can be modified by the user.

Once the user is satisfied with the initial configuration of input tokens, (s)he selects an initial global sigma value and runs the tensor voting process to extract a surface. The surface is quickly rendered as points by default, but given more time, the user can also render the surface as a mesh.

Modeling Parameters

The original TV framework has only one free parameter, namely σ , which is a coefficient in the saliency decay function (Eq. 2), and establishes the voting field scale of analysis. σ is a positive scalar which determines the size of the voting fields used in the voting process and is set with an initial global value. Having a single global σ value will create a model with one level of detail. In TVMS, we name this parameter σ_G , and it is used to initially give all input tokens the same scale of analysis. We introduce a second σ , namely σ_L , which is used to set the scale of analysis locally. The user may assign σ values to the individual tokens. Both σ_G and σ_L are used in Equation 2 in the same way. The only difference between them is that σ_L may be used to modify the voting field of a token, while σ_G defines a global, default voting field. Initially σ_L is set to -1 for each token, so that the system can distinguish modified tokens from unmodified ones. σ_G

is set to 10 by default and used for all tokens whose σ_L value has not been changed from the initial value of -1. Instead of a single global voting field, up to n , where n is the number of input tokens, individual voting fields can be used to vote for a surface. This allows for the level of detail of the surface to be varied by using a smaller σ_L in those regions where the user desires finer details.

A weight has been added as a third modeling parameter. The weight is a scalar value associated with each input token. The token's normal is multiplied by the weight before voting and surface extraction. Increasing the weight value strengthens the influence of the plane defined by the token; thus flattening the resulting surface near the token. The default weight value is set to 1.0.

4. RESULTS

In order to investigate the geometric modeling potential of tensor voting we conducted a number of experiments. Our experiments were performed in two parts. In the first part, we looked at the local influence of the individual modeling parameters on a simple shape. In the second part, we examined the overall effectiveness of the TV modeling process when attempting to model two specific shapes.

Influence of Modeling Parameters

The scale of analysis parameter σ affects the overall computational extent of the model and the cost of calculating a surface. Because we wish to extract the surfaces as quickly as possible, the size of the voting fields, and thus the size of the models, have to be kept small. We found that the radius of our models, in units of voxels, should be approximately the same as σ . Having a value of σ between 8 and 10 produced results in a reasonable amount of time, which implied a model radius of slightly more than 10 units. Using larger σ values would require bigger models, thus a higher resolution computational grid, and would require significantly more time in the surface extraction step. Also, σ should not be set too large, since this causes votes to cancel each other out [Med00] and produces surface holes, further motivating the use of small σ values.

The model used in our parameter influence experiments is a sphere with a radius of 12.5 units, that was defined by 18 input tokens. See Figure 3 (left). TV3D performs tensor voting and interpretation on a 25 x 25 x 25 grid, and usually requires 3 to 5 cpu-seconds to produce a result.

To demonstrate the effect of translating input tokens, we translate one single token on the sphere, namely the right-most one, one unit at a time. The results are presented in Figure 2, which consists of spheres of radius




Figure 2: Translating the right-most token one unit at a time. Left: $\sigma = 7$. Right: $\sigma = 10$. The token locations are $x = 12.5, 13.5, 14.5, 15.5, 16.5, 17.5$.

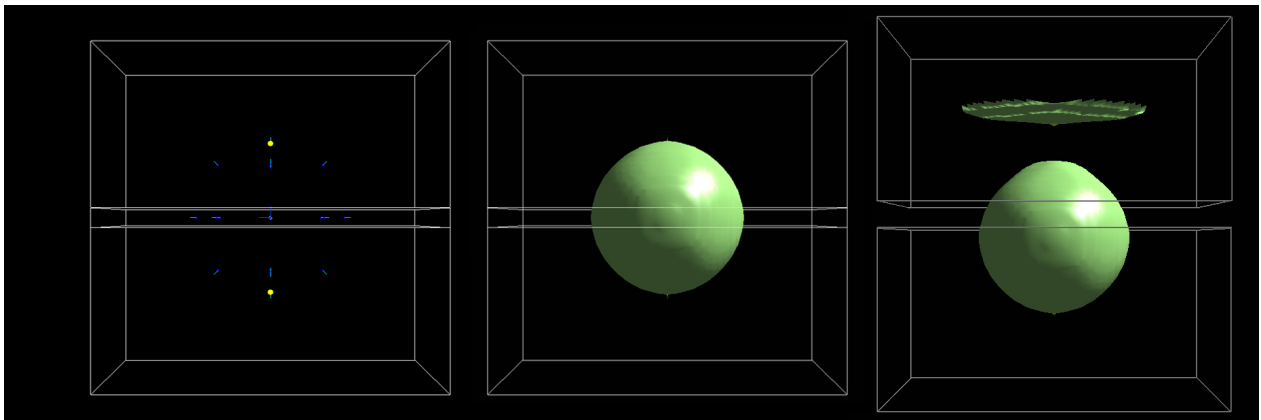


Figure 3: Left: The 18 input tokens needed to define a sphere. The top and bottom tokens have been highlighted and the extent of their voting fields ($\sigma_G = 10$) have been displayed. (center) The resulting sphere, produced by the 18 input tokens. (right) The extent of the voting fields for the top and bottom tokens, after token translation. The two fields no longer overlap and a disjoint surface has been formed.

12.5 produced with $\sigma = 7$ and $\sigma = 10$. The token is translated by 1.0 additional unit for each example. It can be seen that the surface does begin to bulge out to fit to the surface patch implied by the translated token. But if the token is translated too far the surface breaks up.

Insight into this unwanted artifact is provided in Figure 3. Here, the extents of the voting fields for two of the sphere's 18 tokens are displayed. As the top token is translated away from the others the two displayed extents no longer overlap. In this situation two disjoint surfaces are produced. The TV process interprets the translated token as an isolated point and tries to infer a separate plane from it. This clearly shows that in order to produce a closed mesh, one needs to ensure that there is significant overlap between the voting fields of neighboring tokens.

Figure 2 also highlights the effect of σ on the surface. On the left, with $\sigma = 7$, the resulting surface has flatter, somewhat sharper, features. When σ is increased to 10, more voting fields overlap because of their increased size.

The increased number of contributing VF's at any location creates a smoothing effect on the surface. This is most evident in the top-right and bottom-left surfaces in each of the examples. As the right-most token is translated, the resulting bump has a smoother transition to the sphere when σ is greater.

To demonstrate the effect of rotating normals, Figure 4 presents the changes to the sphere produced by rotating the normal of the right-most token. It can be seen that applying small rotations (angles less than 45°) to a token can be easily accommodated by the tensor modeling system. The surface smoothly blends in the disoriented surface patch implied by the rotated token. As the rotation angles increases past 45° the surface breaks up. Using a higher σ gives slightly better results, but this still cannot resolve the conflicting input information that implies that a piece of surface is directed into the sphere.

The weight parameter provides another way to modify the influence of a token. The weight parameter may either




Figure 4: Rotating the normal of the right-most token. Starting at the top left, the normal is rotated 0° , 15° , 30° , 45° , 60° and 75° . (left) Voting with $\sigma_G = 7$. (right) Voting with $\sigma_G = 10$.

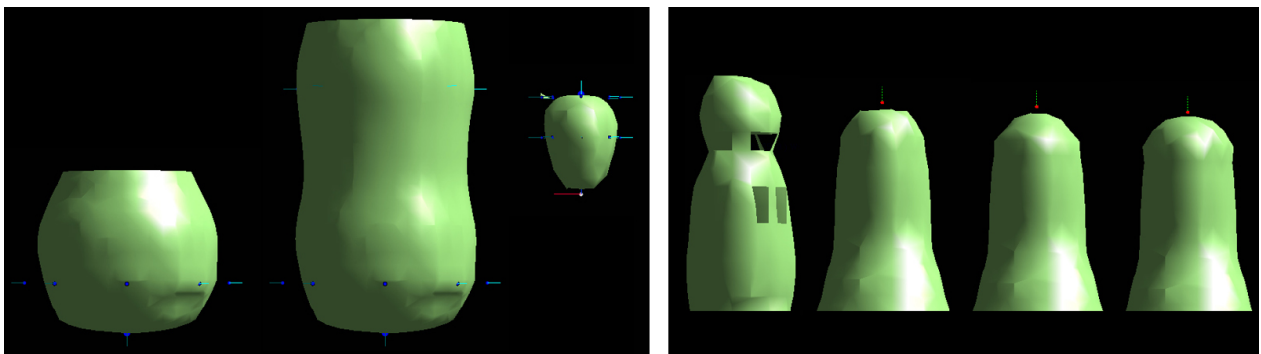


Figure 5: (left) Modeling the bottle. Three rings of tokens were added to define the body. The top was modeled separately and connected to the rest of the model by adding three smaller circles of tokens to create the neck. (right) The weight parameter was used to change the shape of the bottle's top. Weight values, from left to right are 1.3, 2.0 and 2.8. All the tokens on the circle below the topmost token have a weight value of 1.2.

elongate or shorten the length of the normal \hat{n} associated with a token after the voting process. Given a large weight the plane associated with the token's normal has greater influence on the surface; thus locally attracting the overall surface to the plane, as seen in Figure 5 (right). Initially the bottle's top token is overwhelmed by the information from the side tokens. As the top token's weight is increased, the most likely curved surface emanating from it increasingly blends with the rest of the bottle surface.

Creating Complete Models

To explore the overall tensor voting modeling process, two complete models were created, a bottle and a car. Both are fairly simple objects, but have geometric properties, such as varying curvature and level of detail, that potentially could reveal any weaknesses or problems with our TV modeling system. Both models were created by initially defining the tokens needed to produce the basic shape. The tokens were then modified to add details or fix problems. To change the curvature of the surfaces, a

number of normals were rotated. To flatten regions of the surfaces, weights of nearby tokens were increased. To fill holes, new tokens were added or σ_L of the neighboring tokens were increased.

For the car, the body was first created. The four wheels were then added as details later. The bottle model, defined with 56 tokens, is presented in Figures 5 and 6. As seen in Figure 5, the bottom of the bottle was created using two circles of tokens and one bottom token. Then a third circle of tokens was added to extend the body. Finally, the top was modeled separately and connected to the rest of the model by adding three smaller circles of tokens, thus creating the neck. The weight of the top token was increased until the desired blending and curvature was achieved. The TV3D system evaluated the bottle's Surface Map on a $14 \times 38 \times 14$ grid in 11.5 cpu-seconds. The resulting surface may be displayed with points. These are the zero-crossings of the computational grid used to calculate the associated g field. The Marching Cubes algorithm [Lor87] may also

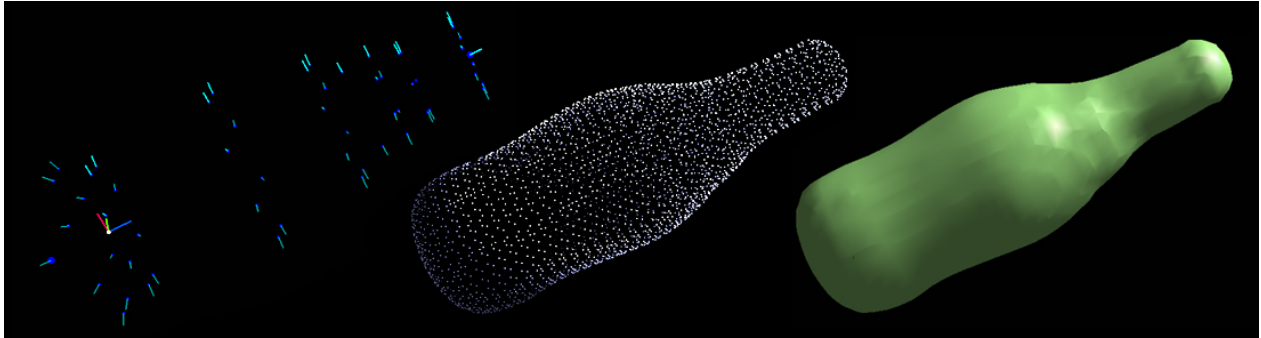


Figure 6: The bottle model created from 56 tokens. (left) Input tokens. (center) Resulting TV surface displayed as points. (right) TV surface displayed as a mesh.




Figure 7: The car model created from 151 tokens. (left) Input tokens. (center) Resulting TV surface displayed as points. (right) TV surface displayed as a mesh.

be applied to the g field grid to produce a mesh. The car model, defined with 151 tokens, is presented in Figure 7. The TV3D system evaluated its Surface Map on a $36 \times 22 \times 18$ grid in 11 cpu-seconds. The resulting surface is displayed as points and a mesh.

5. DISCUSSION

During the modeling process, a significant problem quickly emerged. Small, densely sampled areas may produce asymmetric results, which we believe is due to the TV3D surface extraction implementation. The algorithm begins at the most salient input token and grows a surface from it until the saliency of the neighborhood drops below a certain threshold. It then goes to the second most salient input token and so on. For the bottle, the asymmetry can be seen on the neck, which is not perfectly symmetric. When modeling the car, the asymmetry problem imposed serious problems shaping the more finely detailed areas. After creating the body, one wheel at a time was added. When the second wheel was added, its tokens influenced the tokens of the first wheel, changing its shape. We changed our strategy and modeled a full half of the car in one go, which was then

deduplicated, mirrored and merged with the first half. The resulting mesh was not closed and after repairing one hole, new holes emerged on other parts of the model. The surface extraction of TV3D appears to propagate the influence of the tokens forward as they are processed; thus the order in which tokens are processed affects the final extracted surface. The current TV3D implementation is elegant, but may not be the correct computational engine for an interactive modeling application.

The local influence of an individual input token on the final surface is not as strong nor with the type of properties as had been desired. Changing the parameters of a particular input token will in most cases render a change in the extracted surface, but only if the parameter values remain within a certain range. For instance, if a token is translated far away from the other tokens, it will lose its influence on its neighbors and will produce a surface separated from the rest of the model. Also, the influence of a single token decreases exponentially with its distance. As a token is moved away from the remaining tokens, its influence on the overall surface diminishes and its neighbors' influence increases. The neighbors then smooth out the contribution of the moved token. Determining these

acceptable parameter ranges and including them as constraints will be a necessary part of a future TVMS.

Surface extraction and voting is costly. Each time the model is modified and the surface must be updated, TVMS recalculates the entire surface instead of just the part that will be affected by the change. If the system could re-vote only the modified input tokens and update just the affected areas of the surface, surface regeneration could possibly be done in real-time. This will require a redesign of the surface extraction algorithm.

6. CONCLUSION

This work has investigated the tensor voting method, a technique for geometric feature grouping, in a 3D modeling context. The new system, which employs the TV3D framework for tensor voting and the point-rendering functionality of QSplat for visualization, extracts surfaces from sparse input data, consisting of points and normals. These points and normals can be interactively edited by the user to change the resulting surface. The initial goal of modeling simple objects using TV has been achieved and two new modeling parameters, a local σ , σ_L , and a weight w , have been introduced and implemented.

Though the approach is promising, several limitations have been found in our current system. There is an asymmetry in the extracted surfaces, probably due to the surface extraction implementation in TV3D. Modeling parameter values may be easily set to produce unstable and unwanted results. Voting and surface extraction is slow for large models and large scales of analysis. The curve and junction capabilities of tensor voting should be explored for modeling sharp features and fine details. Each of these limitations should be addressed in future versions of a TV-based modeling system.

Acknowledgements We would like to thank Wai-Shun Tong who provided us with the tensor voting library and much technical assistance, and Christoffer Westberg for programming our initial point rendering software.

REFERENCES

- [Bli82] Blinn, J. A generalization of algebraic surface drawing. In *ACM Transactions on Graphics*, 1, No. 3, pp. 235–246, 1982
- [Dew04] Dewaele, G. and Cani, M.P. Interactive global and local deformations for virtual clay. In *Graphical Models*, 66, No. 6, pp. 352–369, 2004
- [Guo04] Guo, X., Hua, J., and Qin, H. Scalar-function-driven editing on point set surfaces. In *IEEE Comput. Graph. Appl.*, 24, No. 4, pp. 43–52, 2004
- [Guy96] Guy, G. and Medioni, G. Inferring global perceptual contours from local features. In *Int. J. Comput. Vision*, 20, No. 1-2, pp. 113–133, 1996
- [Guy97] Guy, G. and Medioni, G. Inference of surfaces, 3D curves, and junctions from sparse, noisy 3D data. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19, No. 11, pp. 1265–1277, 1997
- [Lor87] Lorensen, W. and Cline, H. Marching Cubes: A high resolution 3D surface construction algorithm. In *Proc. SIGGRAPH '87*, pp. 163–169, 1987
- [Med00] Medioni, G., Lee, M.S., and Tang, C.K. *Computational Framework for Segmentation and Grouping*. Elsevier Science Inc., New York, NY, USA, 2000
- [Mus02] Museth, K., Breen, D.E., Whitaker, R.T., and Barr, A.H. Level set surface editing operators. In *Proc. SIGGRAPH '02*, pp. 330–338, 2002
- [Mus05] Museth, K., Breen, D., Whitaker, R., Mauch, S., and Johnson, D. Algorithms for interactive editing of level set models. In *Computer Graphics Forum*, 24, No. 4, pp. 821–841, 2005
- [Pau02] Pauly, M., Keiser, R., Kobbelt, L., and Gross, M. Shape modeling with point-sampled geometry. In *Proc. SIGGRAPH '02*, pp. 322–329, 2002
- [Rus00] Rusinkiewicz, S. and Levoy, M. QSplat: a multiresolution point rendering system for large meshes. In *Proc. SIGGRAPH '00*, pp. 343–352, 2000
- [Tan98] Tang, C.K. and Medioni, G. Inference of integrated surface, curve, and junction descriptions from sparse 3D data. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20, No. 11, pp. 1206–1223, 1998
- [Tan01] Tang, C.K., Medioni, G., and S.Lee, M. N-dimensional tensor voting and application to epipolar geometry estimation. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23, No. 8, pp. 829–844, 2001
- [Ton04] Tong, W.S., Tang, C.K., Mordohai, P., and Medioni, G. First order augmentation to tensor voting for boundary inference and multiscale analysis in 3D. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26, No. 5, pp. 594–611, 2004
- [vF06] von Funck, W., Theisel, H., and Seidel, H.P. Vector field based shape deformations. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 25, No. 3, pp. 1118–1125, 2006
- [vF07] von Funck, W., Theisel, H., and Seidel, H.P. Explicit control of vector field based shape deformations. In *Proc. Pacific Graphics*, pp. 291–300, 2007
- [Zwi02] Zwicker, M., Pauly, M., Knoll, O., and Gross, M. Pointshop 3D: an interactive system for point-based surface editing. In *Proc. SIGGRAPH '02*, pp. 322–329, 2002