Point-Based Probabilistic Surfaces to Show Surface Uncertainty

Gevorg Grigoryan and Penny Rheingans, *Member*, *IEEE Computer Society*

Abstract—Efficient and informative visualization of surfaces with uncertainties is an important topic with many applications in science and engineering. In these applications, the correct course of action may depend not only on the location of a boundary, but on the precision with which that location is known. Examples include environmental pollution borderline detection, oil basin edge characterization, or discrimination between cancerous and healthy tissue in medicine. This paper presents a method for producing visualizations of surfaces with uncertainties using points as display primitives. Our approach is to render the surface as a collection of points and to displace each point from its original location along the surface normal by an amount proportional to the uncertainty at that point. This approach can be used in combination with other techniques such as pseudocoloring to produce efficient and revealing visualizations. The basic approach is sufficiently flexible to allow natural extensions; we show incorporation of expressive modulation of opacity, change of the stroke primitive, and addition of an underlying polygonal model. The method is used to visualize real and simulated tumor formations with uncertainty of tumor boundaries. The point-based technique is compared to pseudocoloring for a position estimation task in a preliminary user study.

Index Terms—Uncertainty, visualizing surface uncertainty, point-based graphics.

1 Introduction

Everything we measure in science and engineering can have uncertainty associated with it. When the measured (or predicted) data are visualized, it is important to communicate information about the uncertainty along with the data itself, especially when the visualization is to be used in decision-making. For simple situations with modest amounts of data, there are standard methods for dealing with this problem, for example, representing uncertainty at points of a 2D graph with vertical bars or graphing probability density curves. In more complicated situations, the task becomes more challenging.

We have chosen to concentrate on the problem of visualizing surfaces with uncertainties. Approaches for visualizing the uncertainty of surfaces are particularly useful for studying the general problem of deducing and displaying border information from a volume data set. For example, in protein structure determination with X-ray crystallography, where atomic coordinates are deduced from electron density, properly estimating the uncertainties in atomic position resulting from thermal motion and crystallographic artifacts is as important as finding the coordinates themselves. Similar problems exist in medicine, where the task is to determine the boundary between two different types of tissue (healthy/sick or contaminated/noncontaminated) so that surgery or radiation therapy may be planned.

- *G. Grigoryan is with the Department of Biology, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139. E-mail: gevorg@mit.edu.*
- *P. Rheingans is with the Computer Science and Electrical Engineering Department, University of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250. E-mail: rheingans@cs.umbc.edu.*

Manuscript received 19 Feb. 2003; revised 27 Mar. 2003; accepted 20 Nov. 2003.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCGSI-0010-1102.

An ideal technique suitable for displaying surfaces with uncertainties should meet certain requirements. First, the method must give information about the geometry of the surface and the uncertainty in every region. The method should be intuitive so that, by simply looking at a visualization, one can easily surmise the desired information. In order to communicate sufficient information about a 3D surface in a still 2D picture, the method should be interactive to allow viewing from different positions. Whatever means of visualizing uncertainty is used, it should not distract one from the perception of likely surface location. And, ideally, one has to be able to visualize other variables along with uncertainty.

In this paper, we present an approach to visualizing surfaces with uncertainties which meets all of the above requirements. Our technique is based on using points as display primitives to render a surface. The main idea behind dealing with uncertainty is to follow the logical consequence of having an uncertainty at a point. Namely, if we are not certain about how the surface behaves in a particular region, then we are not certain about exactly where to put points representing the surface in that region. Hence, we should displace these points from their original location in a way that incorporates a certain amount of randomness proportional to the uncertainty in the region.

We apply our approach to the visualization of growing tumors. It is well-known that finding strict borders of a growing tumor is extremely important when it comes to making decisions on how to treat the tumor, but also very hard because areas of low tumor cell density can be the ones with the highest growth rates. We demonstrate our method using both modeled and diagnostic tumor formations.

2 Related Uncertainty Visualization Work

There is a rich and active body of research addressing the challenge of showing data values in the context of their

certainties in an organic and effective manner. These approaches may be grouped into three basic categories: sampling approaches, additional variable approaches, and geometry modifications.

Sampling approaches rely on showing uncertainty information only at discrete points of the space. Examples are uncertainty glyphs [1], [2], sonification [3], [4], [5], procedural annotations [6], and discrete probability distributions [7]. Uncertainty glyphs incorporate certainty information into the visual appearance of the glyph. For example, a glyph for wind direction might show an uncertain direction as a wedge, rather than a precise arrow. Sonification approaches map uncertainty at probed locations to sound. Using procedural annotations, uncertainty values may be used as parameters controlling the distortion or appearance of annotation marks. Finally, discrete probability distributions describe the regions of space spanned by a discrete number of possible surface configurations. These methods are useful as sampling devices; however, the fact that uncertainty is only shown at discrete points limits their utility in continuous domains.

Perhaps the simplest approach to representing uncertainty relies on simply treating uncertainty as an extra variable, effectively increasing the number of dimensions of the data set by one. This extra variable is then displayed using any of the standard methods for displaying multidimensional data sets. In particular, pseudocoloring [8] can be easily used for this task. Unlike the previous group, these approaches do convey information about the uncertainty continuously. However, by visualizing uncertainty as an independent variable, they fail to treat it as a quantity associated with another measure. Indeed, when we say uncertainty, we refer to uncertainty in some variable. If this variable is also shown in the visualization, then, ideally, the visual display mechanism of the uncertainty should correspond to that of the variable. For example, uncertainty in location is associated with coordinates. So, if we use pseudocolor to render uncertainty, we lose the connection between the uncertainty and the geometry.

The third widely used general method for visualizing uncertainties on surfaces modifies the geometry of the surface based on the uncertainty in each region. Approaches from this collection are fat surfaces [1], [9], displacements and other perturbations of geometry [1], [10], animation [11], and IFS fractal interpolation [12]. In the fat surfaces approach, several surfaces are rendered to show the range of possible locations of the data points. This technique is well-suited for situations when the uncertainty is given in the form of minimum/maximum value pairs. The animation approach is based on oscillating regions of the surface with amplitudes proportional to uncertainty. This approach is convenient, but has the disadvantage that it cannot be viewed on static media. Also, with the surface oscillating, it may be hard to display variables other than uncertainty. The IFS fractal interpolation method [12] produces surfaces with perturbed geometry, with larger perturbations in uncertain regions. This is similar to the technique employed here, but, with the IFS approach, the most salient feature of the visualization is the high frequency nature of uncertain regions. The distinction between certain and uncertain regions is clear, but the location of the most likely surface may be obscured. This paper builds upon a preliminary work of the authors [13], offering several enhancements and preliminary evaluation of the effectiveness of the approach.

3 Tumor Growth Model Application

We apply our techniques to visualizing tumor formations simulated by a tumor growth model [14], [15]. One of the best known and simplest models of this type is the Gompertz model [14]. We apply it to build a tumor out of many small tumors initially dispersed in space. The growth of each of these small tumors is simulated using the same growth formula prescribed by the Gompertz model, but with different growth parameters corresponding to different conditions. In order to estimate uncertainty associated with the border of the tumor, we begin with the following observations: First, we observe that, no matter what the diagnostic detection technique, the uncertainty will be high in the areas where tumor cell density is low. Next, we observe that high cell density means limited nutrient supply, hence, a low growth rate and vice versa. From these two observations, we estimate uncertainty as the growth rate of the tumors, which can be calculated as the derivative of volume with respect to time.

As a modification of this computational model, we developed a method in which we incorporated metastasis. Metastasis is a well-known phenomenon in cancer development where cells detach from a tumor formation, travel along blood vessels, settle in other places in the organism, and start new tumors there. This behavior is modeled as follows: At each growing time interval, the sizes of all the tumors are increased according to the model. At this point, if the size of a tumor exceeds some threshold, then, with a certain probability, it is forced to metastasize. This means that a new tumor center is created a random distance away from the old one along the predefined blood vessel direction. This introduced a new parameter into our model —age, which we output along with uncertainty. In real life, the age of a tumor region would be hard to assess directly, however, the parameter is still very meaningful in understanding the growth of tumor formations.

Using a computational model to test our visualization method has several advantages. First of all, one can get more accurate certainty information from a computational model. Eventually, of course, a way of estimating uncertainty using medical images has to be developed, but with a working visualization technique this becomes a separate problem. One advantage of the approach proposed in this paper is the ability to add additional variables to the visualization already containing uncertainty data. A computational model allows us to develop meaningful supplementary variables on which to test the visualization approach.

A simple isosurface visualization of a data set generated using this model is presented in Fig. 1a. Fig. 1b shows the same model with pseudocolor indicating uncertainty. This representation conveys information about both surface shape as well as uncertainty. However, there are limitations to this type of a visualization approach. Pseudocolor only gives us information about relative uncertainty values throughout the surface. Even though one can introduce a colorbar into the visualization, one still will not have a clear spatial concept of how bad a particular uncertainty value is. In other words, there is no direct connection between uncertainty and geometry.

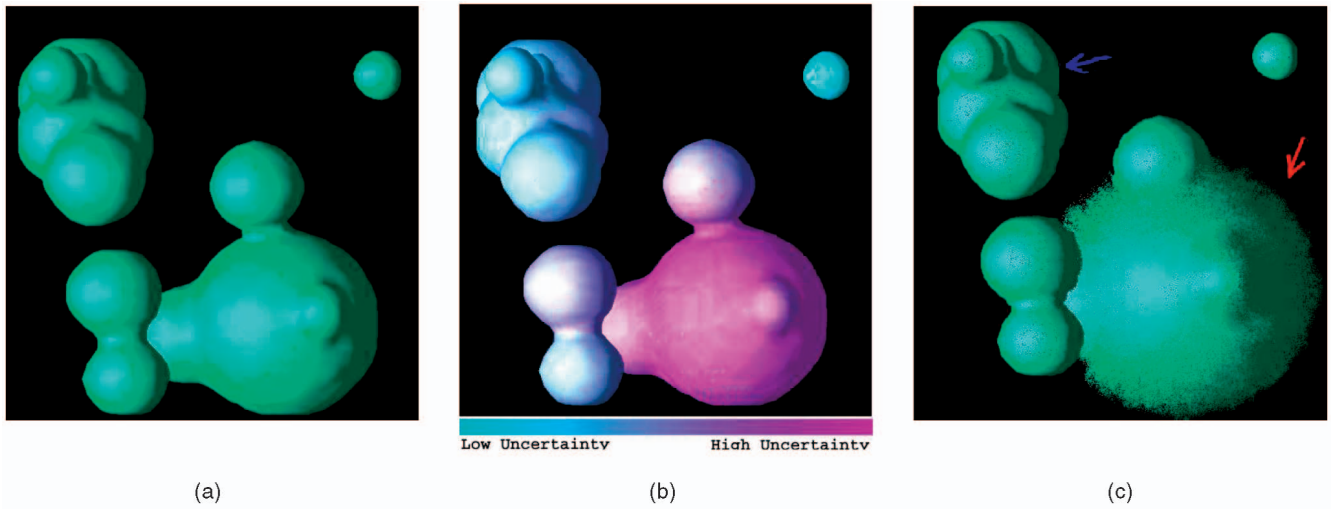

Fig. 1. Tumor formation. (a) Polygonal model only. Excellent surface geometry information, however, no uncertainty information. (b) Polygonal model with pseudocolor. Both surface geometry and relative uncertainty information are clearly conveyed. (c) Point-based model. Surface geometry as well as absolute uncertainty information (as it relates to coordinate uncertainty) are conveyed.

4 Approach

Our approach to visualizing surfaces with uncertainties corresponds directly to how we think of these uncertainties. A nonzero uncertainty value at a particular point on a surface indicates that we are not sure precisely where the point is. Specifically, that point in the actual object may be somewhere different from where it is on the rendering. The expected difference range is given by the size of the uncertainty value associated with this point. The heart of our approach is to displace each point on the surface along the surface normal at the point, with the displacement proportional to the uncertainty value at the point. Because the uncertainty value at each point shows the expected range of displacement, rather than a precise value, we show displacements randomly distributed in the range. Performed for many points, this will produce the expected displacement distribution in a region.

Triangles and other polygons have been historically used as display primitives in representing 3D objects because of the availability of hardware acceleration for polygon rendering. But, as scenes grow more complex, polygons shrink to sizes approaching that of a pixel. The possibility of using points as general display primitives was introduced by Levoy and Whitted [16]. It was shown that, as the complexity of the scene increases, choosing points as rendering primitives presents a great number of advantages in terms of making algorithms simple by enabling the rendering of surfaces without the necessity for underlying geometry. Hence, at some point, it becomes appropriate to use points instead of polygons [17], [18]. The advent of hardware support for point rendering as well as work done in the area of optimizing point-based rendering [19], [20] makes this representation even more attractive.

The idea of representing fuzzy phenomena as point clouds dates back to the particle systems [21] which were later applied to surface modeling [22]. Treating uncertain surfaces as fuzzy phenomena and rendering then with a point-based system seems a particularly natural choice. Note that, if we use a point-based representation of the surface, then, given an uncertainty value at a point, we are free to do whatever we want with that point without explicitly having to consider what happens to the rest of the surface. This inherent flexibility of point-based models has proven useful as the basis for a geometric modeling and editing package [23]. We use that flexibility here to make independent decisions about the position and appearance of each point making up a surface. Rendering time is unaffected by the specific uncertainty values. Hence, in our situations, it is convenient, natural, and efficient to render the scene point-by-point.

5 Basic Implementation

Our visualization program is implemented using the Open Graphics Library (OpenGL). It provides user parameters for the scale factor ($scale$) and falloff (a) for displacement. The basic algorithm is outlined below:

1. Read in surface and uncertainty information.
2. Create N random points inside each triangle.
3. Interpolate uncertainty and normals from vertices.
4. **for** *each point* P **do**
5. calculate the displacement:
   $disp = rand() * (uncert\ at\ P)^a * (scale)$,
6. Displace P in the direction of the normal at P.
7. **end for**
8. Display all the points.

The simplest version of our method is one in which the surface is represented as a collection of displaced points. An example of such a visualization is presented in Fig. 1c. Several advantages of this model are immediately apparent. First of all, it is very clear by looking at the visualization which regions of the surface have high uncertainties and which have low. Second, the spatial extent of the uncertainty is clear because one can see the region around the surface where points are likely to be found. So, if the size of this region is on the order of surface detail, we know that the uncertainty in that region is rather high and we have practically no idea about the location of the surface in that region. However, if the size of this region is rather small

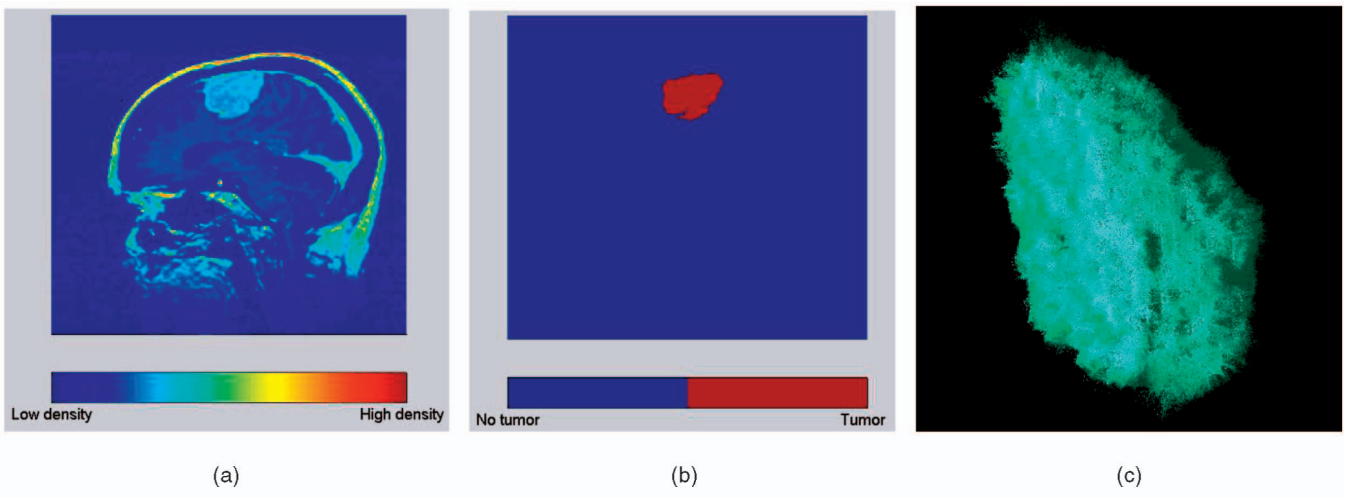

Fig. 2. Tumor in hand-segmented diagnostic data. (a) CT slice. (b) Tissue segments. (c) Tumor surface.

compared with surface detail, we can be fairly certain about how the surface behaves in the region. Looking at Fig. 1c, we can see that the region indicated by the red arrow is rather uncertain, as opposed to the region indicated by the blue arrow, where the surface behavior is clear.

Since our method displays 3D information on a 2D medium, it is essential for the display to be fast enough to be interactive. Our method provides real-time interaction and allows the user to rotate, move, zoom in and out, as well as allowing the control of representation parameters such as point density, uncertainty scale factor, and other optional enhancements.

6 Results on Diagnostic Data

In order to assess the wider applicability of this method to diagnostic data, we also used it to visualize both segmented and unsegmented medical volumes. These volumes differ from the growth model data in that neither tumor certainty nor age are directly available. Since only density information is given by the data, other quantities must be estimated. We use preliminary and simple methods for estimating tumor certainty here; more complex models could be easily substituted.

We first visualized the boundary surface of the tumor identified by hand-segmentation. Skilled hand-segmentations are still regarded as the gold standard for medical segmentation since they incorporate tacit knowledge about what structures should look like, allowing for the correction of imaging artifacts through expert subjective judgements. Segment boundaries follow the anatomist's best judgment of where the boundary is, rather than a strict contour of image intensity. Accordingly, the intensity is not necessarily constant across the surface. For such segmentations, it is particularly interesting to distinguish where the anatomist followed strong image boundaries and where more subjective judgments were made. For each section of this surface, tumor certainty was estimated by the inverse density gradient. Specifically, we expect that, where the density gradient at the surface is high, the image contains a sharp and obvious border, hence the uncertainty level regarding the location of the tumor boundary is low. On the contrary, if the density gradient is low, a definite boundary is hard to detect, therefore uncertainty is high (or, alternatively, the boundary in that section is rather subjective). Fig. 2a shows a slice of the original MRI data, while Fig. 2b shows the same slice after segmentation [24], [25]. Fig. 2c shows the presumed tumor boundary in the context of its certainty. The general shape of the tumor is apparent, as is the uncertainty about its precise extent.

We also used our method to visualize an unsegmented CT scan of human kidneys with tumor formations. In order to find the areas with tumor formations, we calculated the isosurface of the volume with an isovalue that was known to correspond to the tumor density. For the certainty measure, we once again used inverse density gradient at the surface. The visualization of the data is shown in Fig. 3. Again, areas of high and low certainty are apparent and this information does not interfere with the surface geometry information. We can clearly see that areas pointed to by red arrows are of particularly high uncertainty, while those pointed to by blue arrows are rather certain and clearly shaped.

7 Enhancements

Although the basic underlying concept for point-based probabilistic surfaces is extremely simple, it proves a powerful and flexible approach, in large part because it very naturally enables a wide range of enhancements. These enhancements include alternate appearance attributes for the points, alternate point distributions, alternate primitives, and the inclusion of additional variables in the display. This section discusses each of these enhancements in turn.

7.1 Transparent Points

Probabilistic surfaces intuitively convey the basic nature of uncertain boundaries by giving them a fuzzy appearance. This fuzzy appearance can be reinforced by also varying transparency with the uncertainty level in order to to see blurriness in the areas of low certainty and distinct features in areas of high certainty. Hence, areas with high uncertainty are given low alpha values and vice versa. This can be achieved by controlling the transparency of the individual points, making a point with a higher uncertainty value more transparent than a point with a lower uncertainty value. The regions having an aggregation of

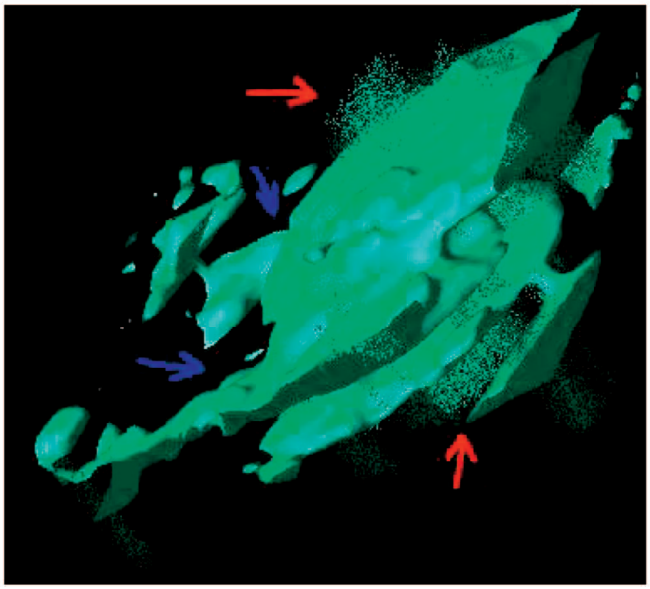

Fig. 3. Kidney tumors in unsegmented diagnostic data.

points with high uncertainty values will be a collection of highly displaced, blurred points and, hence, the shape of the surface in these regions will not be apparent. In this case, the transparency of the underlying point model is dependent on uncertainty values with the following relationship: $\alpha = 1.0 - err^c$, where err is the scaled uncertainty value (from 0 to 1) and c is a constant which controls how quickly the transparency increases with increasing uncertainty. The utility of this approach is apparent from Fig. 4a. Regions of high uncertainty become even more blurry than with the basic technique.

7.2 Alternative Point Distributions

The basic probabilistic surface algorithm uses a uniform distribution of displacements along the surface normal, up to a maximum displacement. This implies that the surface might be located anywhere in the window given by the maximum displacement with equal likelihood. In some situations, more may be known about the probability distribution of likely surface locations. For example, positions near the best estimate given by the surface position may be more likely than those farther away. The displacement of points may be modified in a straightforward manner to incorporate uncertainty of arbitrary distribution by adding a term to the calculation of displacement. Specifically, the displacement is given by:

$$disp = rand(PDF) * (uncert\ at\ P)^a * (scale), \qquad (1)$$

where a and $scale$ are controlled by the user and $rand(PDF)$ is a random number distributed according to the probability density function PDF. We used the rejection method [26] to generate random numbers with an arbitrary probability distribution.

Fig. 4b shows the result of using a Gaussian probability distribution function for displacement. Points are more likely to be displaced small distances than large, resulting in a somewhat more defined shape even in uncertain regions. We also tried using a probability distribution which ensures equal density of points at all distances from the surface (in a particular range of distances). However, for the particular surface geometries and displacements we considered, this did not seem to make much difference, so we used the simpler distribution. We expect that equal density displacement might produce a visible difference in cases of surfaces with highly varying curvature or when displacement distances are extremely large.

7.3 Line Primitives

Other types of stroke primitives can be used to produce similar effects in probabilistic surface visualization. We have implemented a variant using line primitives instead of points. Points on the surface are generated and displaced as before, but, instead of simply drawing the points, we draw a line segment linking the original and displaced point. In this way, uncertain sections of the surface grow furry with displacement lines. Fig. 4c shows the tumor growth model with furry uncertain regions. This technique may be thought of as a generalization of the error bars common in 2D graphs, showing both the estimated value and the variability of that quantity. Because of the larger screen extent of a line, compared to a point, fewer primitives are needed to create an effective visualization. As with points, the transparency of the line segment may be modulated to

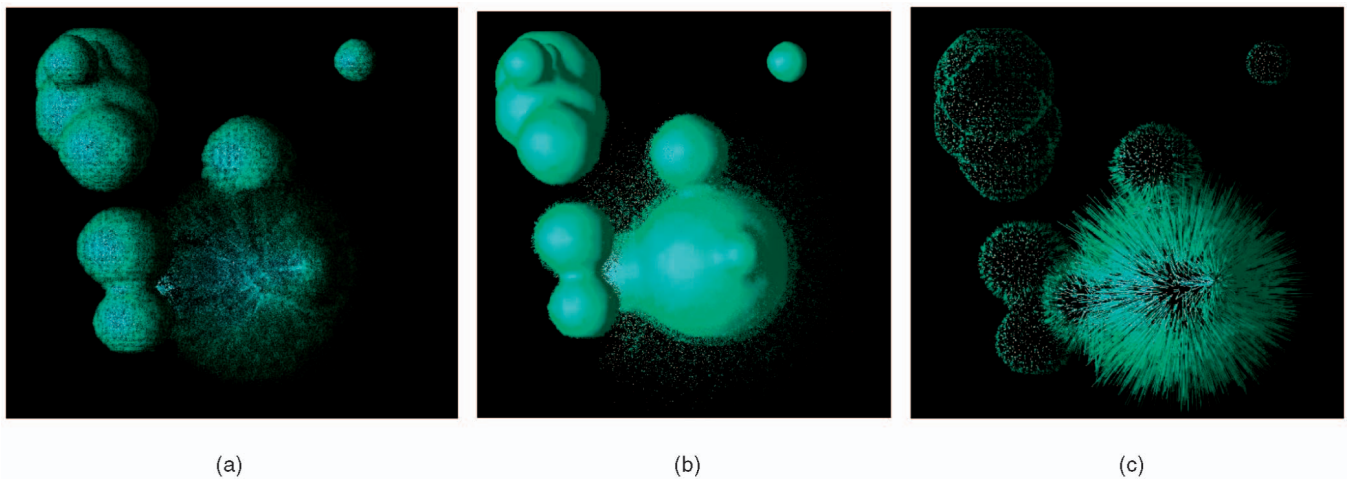

Fig. 4. Some enhancements to basic point-based surfaces. (a) Point opacity modulation. (b) Gaussian probability distribution. (c) Lines.

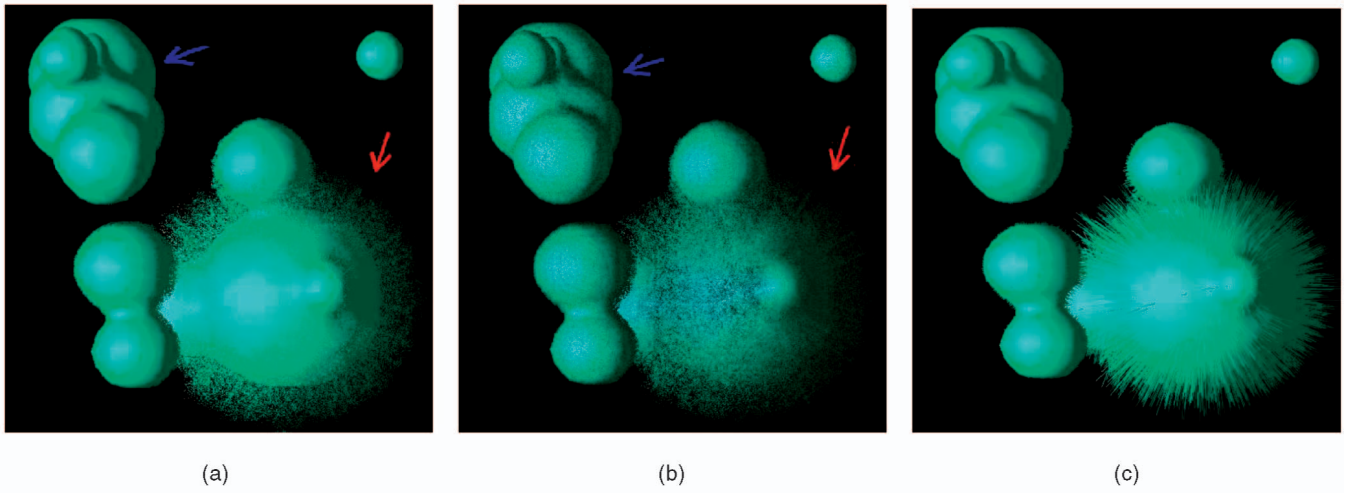

Fig. 5. Underlying polygonal model shown with (a) opaque points, (b) transparent points, (c) opaque lines. Polygonal model clarifies expected location of boundary.

create more intuitive effects. One method would be to have lines fade from full opacity at the root (i.e., the end near the surface) to half opacity at the tip (i.e., the end displaced away from the surface). This mimics the structure of familiar furry objects with fiber diameter decreasing with distance from the skin. Like points, line primitives gracefully provide a mechanism for conveying both probable surface location and certainty by enabling independent decisions about line geometry and appearance properties to be made for each primitive.

7.4 Hybrid Point and Polygon Rendering

One disadvantage of a basic probabilistic surface visualization like the one in Fig. 1c is that there are clearly artifacts present due to the use of simple point primitives. Since points are chosen at random, it is hard to make the points dense enough to guarantee smoothness in the regions with low uncertainty. Specifically, one has to overshoot with the density in order not to miss any points in low uncertainty areas. This translates into wasted computational time. We offer an alternative method which combines point-based and polygon-based rendering to address this problem. Such a hybrid approach has been previously used in the context of optimizing the rendering of large scenes by combining the speed of point rendering for distant objects (low level of detail) and the quality of polygonal rendering for closer object (high level of detail) [27], [28]. Here, we use a similar system but for a different purpose. Along with our point-based model, we render an underlying polygonal model. In the regions where the uncertainty levels are low, we see a smooth surface regardless of the point density, while, in the regions with high uncertainty levels, we see both where the surface is predicted to be (the polygonal model) as well as where it may actually be (the point-based model). An example of such an illustration is presented in Fig. 5a.

The advantage of the dual-model approach is immediately apparent. Unlike with the simple point-based model, in this case, the regions with low uncertainty values are very well-defined. To see this, compare the region marked by blue arrows in Fig. 1c and Fig. 5a. The second model gives a smooth surface despite the fact that Fig. 1c was created with a point density five times larger than Fig. 5a. Hence, the hybrid approach is visually more appealing as well as more computationally efficient. The region indicated by the red arrow in Fig. 1c is also more informative than the corresponding region in Fig. 5a. Here, we see both the hypothetical surface as well as the level of uncertainty around that surface. Fig. 5b shows hybrid point and polygon rendering with transparent primitives. Fig. 5c shows line primitives with an underlying polygonal mesh.

7.5 Additional Variable Display

Ideally, a good method for showing surfaces with uncertainty should allow for additional information to be displayed along with surface geometry and uncertainty. By showing uncertainty with displacement, the color attributes of the surface are still available to show another variable using standard color mapping techniques. Fig. 6a shows tumor age for a modeled tumor formation with metastasis. The information about uncertainty conveyed by displacement causes little interference with information about tumor age. It is easy to see the distribution of tumor age throughout the tumor mass as well as the correlation (or its absence) between tumor age and uncertainty. For example, we can see that, in the region pointed to by the red arrow, there is a young tumor formation having a large uncertainty. This makes sense since young tumors generally have high growth rates and our model treats that as an indication of border uncertainty. However, in the region pointed to by the blue arrow, there is a young tumor region with a much lower uncertainty. In real tumors, this corresponds to the case where the growth conditions (such as nutrient supply) are not favorable. Also, in the region marked by the green arrow, tumor age is high as well as growth rate. In a real tumor, this would indicate favorable growing conditions and sufficient nutrient supply throughout the tumor. Hence, it is easy to see that the ages of two tumors need not strictly determine the relationship between their growth rates.

Fig. 6b shows the segmented surface of the brain tumor, with image intensity displayed at each location. This clearly shows how judgments about the tumor boundary do not strictly follow intensity isocontours.

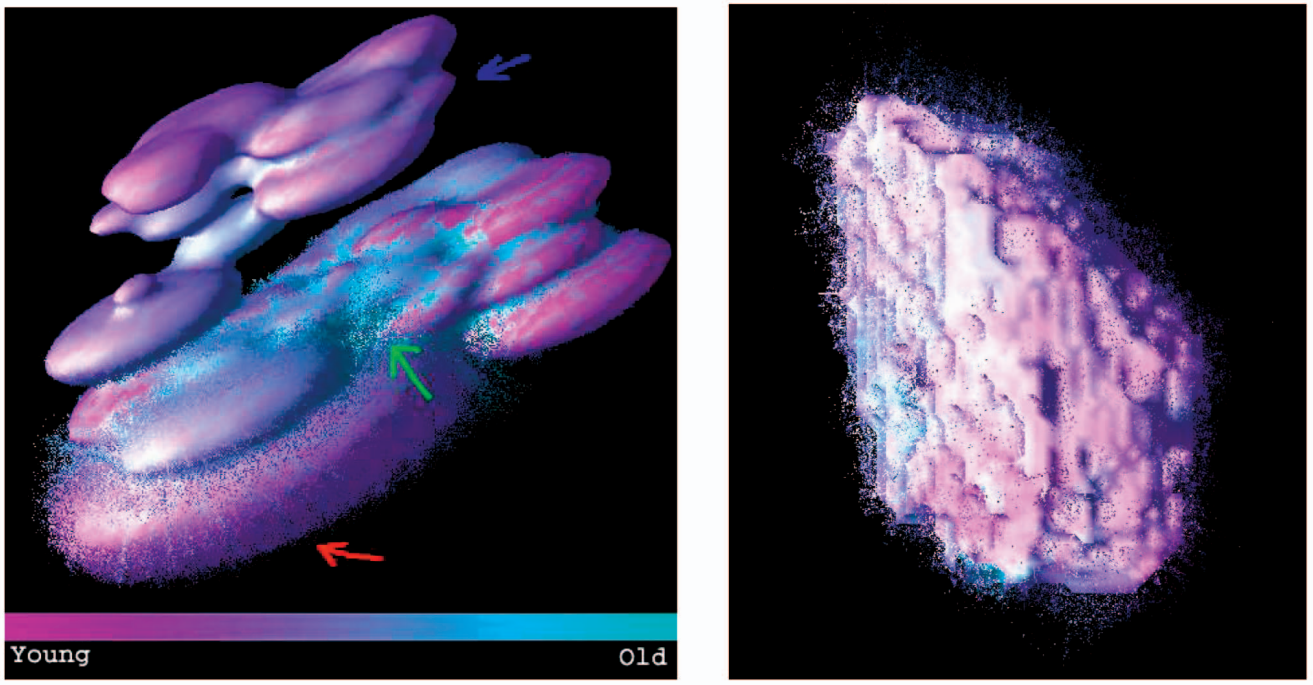

Fig. 6. Display of additional variable on point-based surface: (a) Age in growth model: Color represents age in tumor growth model. (b) Intensity in diagnostic data: Color shows intensity of along segmented boundary of MRI volume.

7.6 Additional Uncertain Variable Display

The independent nature of individual point primitives also very naturally supports the inclusion of uncertainty information about the additional variable mapped to color. Just as an additional variable could have a distribution of values in a region of the surface, so can the colors given to points representing that region be given a distribution of values. This assignment can follow any known distribution function for that variable. We have implemented a method which randomly colors some primitives in the variable mapped to color. Fig. 7 shows the effect when tumor age is mapped to color. Regions where tumor age is believed to be known precisely are colored clearly to indicate that age. Regions where age is less certain contain an increased number of gray primitives, producing the visual effect of decreased saturation. This gives an intuitive indication of tumor age uncertainty.

8 Enhanced Implementation

Incorporating the various enhancements into the algorithm yields the more general algorithm, with user controlled parameters for displacement falloff (a), displacement scale factor ($scale$), transparency falloff (c), and pseudocolor falloff (β). Also, the random function ($rand(PDF)$) now generates a random number according to a given probability function (PDF). This results in the following enhanced algorithm:

1. Read in surface and uncertainty information.
2. Create N random points inside each triangle.
3. Interpolate uncertainty and normals from vertices.
4. **for** *each point* P **do**
5. calculate the displacement: $disp = rand(PDF) * (uncert\ at\ P)^a * (scale)$,
6. Displace P along the normal at P resulting in P'.
7. **if** transparency used **then**
8. Calculate the transparency (alpha value): $\alpha = 1 - (uncert\ at\ P)^c$,
9. **end if**
10. **if** pseudocolor used **then**
11. **if** $(rand() < (uncert\ in\ variable\ at\ P)^\beta)$ **then**
12. $color = neutral\ color$.

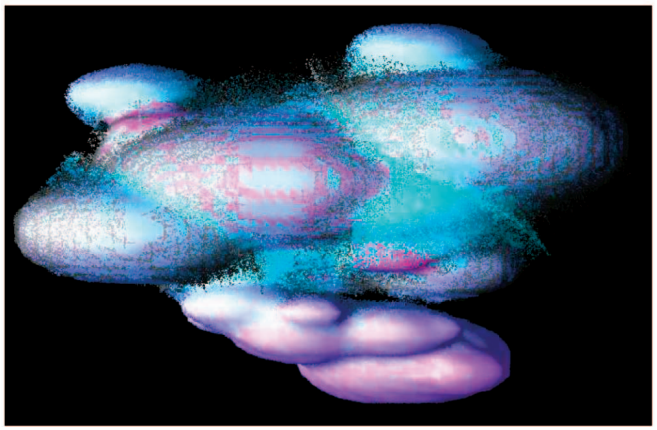

Fig. 7. Display of an uncertain additional variable on a probabilistic surface. Meaning of color here is the same as in Fig. 6a. Areas of low variable uncertainty are dominated by primitives indicating variable value, while areas with high variable uncertainty contain many primitives colored a neutral gray.

13. **else**
14. $color = lookup(variable\ at\ P)$.
15. **end if**
16. **else**
17. $color = default$.
18. **end if**
19. **if** line primitives are wanted **then**
20. Generate line from P to P'
21. **else**
22. Generate P'
23. **end if**
24. **end for**
25. Display all the points or lines.
26. **if** underlying polygonal model is used **then**
27. Display all the polygons.
28. **end if**

9 Performance

One of the notable advantages of our method is the ability to use it interactively. Table 1 summarizes the performance of our implementation with the data and parameters used in some of the figures. It shows the data complexity for each of the cases in terms of the number of polygons in the polygonal model shown in column Polygons (N/A when a polygonal model absent), the number of points per polygon and the total number of points in the point-based model in columns Density and Points (N/A when a point-based model is absent). The running time complexity is demonstrated in terms of the time required for generating all the primitives and the time required for displaying the model. These are shown in columns Build and Display, respectively (time is in seconds). The tests were performed on a computer running RedHat Linux 6.1, with an Intel 1 GHz processor, 256 Mb of memory, and NVIDIA GeForce3 graphics card.

Fig. 1c corresponds to our basic method—where only a point-based model is used. From Table 1, we can see that, since we have to use 100 points per polygon in order to guarantee smoothness in the regions of low uncertainty, a total of 1,853,200 points are present in the scene. When, instead, we employ the point-based and polygonal hybrid model, we obtain better smoothness in low uncertainty areas (see Fig. 5a) using fewer points per polygon—only 20, which translates into five times fewer points in total—370,640. An additional advantage of the hybrid model is that we see both where the surface is thought to be as well as an idea of how certain it is. Comparing the running times for the two models, we see that the hybrid model is about five times faster in both displaying as well as building the primitives. Both models are interactive with time to redisplay below one second, however, the hybrid model has a frame rate of 6.25 frames per second while the basic model only rund at 1.64 frames per second.

Additionally, we see that adding transparency to the hybrid model has practically no effect on the running time (compare entries for Fig. 5a and Fig. 5b). The data set used for visualizing tumor age as an extra variable (Fig. 6a) is slightly more complex, having about 1.5 times more points in total than the data set for the previous figures. Hence, we see a 1.5-fold increase in running time both for displaying as well as building the model. However, here again the model is still interactive, having a frame rate of 4.0 frames per second. As we move to the more complex data set corresponding to the real kidney tumor data (Fig. 3), running time increases proportionally with the number of points.

TABLE 1
Running Time (in Seconds) for the Visualization Method on the Data Sets from the Figures in This Paper

Fig	Polygons	Density	Points	Display	Build
1(a)	18,532	N/A	N/A	0.034	0.034
1(b)	18,532	N/A	N/A	0.034	0.033
1(c)	N/A	100	1,853,200	0.61	5.9
3	26,486	20	529,720	0.32	2.78
5(a)	18,532	20	370,640	0.16	1.33
5(b)	18,532	20	370,640	0.16	1.33
6(a)	28,088	20	561,176	0.25	2.02

As is common in visualization systems, some aspects of image quality are at odds with the need for interactive perfomance. In some cases, we included a parameter to let the user select the desired balance point. The optional hybrid rendering and controllable number of points per polygon are examples of this. In other cases, design choices reflect the need for interactivity. Examples of this include the use of hardware-supported transparent points and the constant number of points per polygon. In the first case, sorting of points by depth would be required to ensure correct compositing. In the second case, a scaling of the number of points by the size of the polygon would result in a more even distribution of points. In both cases, we chose a simpler solution in the interests of interactivity.

10 Evaluation of Effectiveness

We conducted a preliminary user study to compare the effectiveness of our point-based method and the pseudocolor approach for visualizing positional uncertainty. Eighteen volunteers used the two techniques to judge whether a small teapot-shaped marker was inside or outside the error margin of a surface. Marker locations were generated randomly near the edge of the error margin. Fig. 8a shows the display screen for a pseudocolor trial. The color of each part of the surface shows the thickness of the error margin at that point. A legend for the colors and the associated distance was provided below the display. Fig. 8b shows the display screen for a point-based trial. In this display, the error margin was shown by the spread of points out from the surfaces.

After a brief introduction to the representations and the task, each subject completed eight trials with one representation, followed by eight trials with the other representation. Half of the subjects used the pseudocolor representation first, while the other half used the point-based representation first. Subjects were randomly assigned to a trial order.

Analysis of the results showed an average increase in accuracy of 20 percent using the point-based representation, a statistically significant difference ($p < 0.01$). This level of significance indicates a less than 1 percent likelihood that the differences observed resulted from random chance. Table 2 shows summary statistics over all subjects: mean

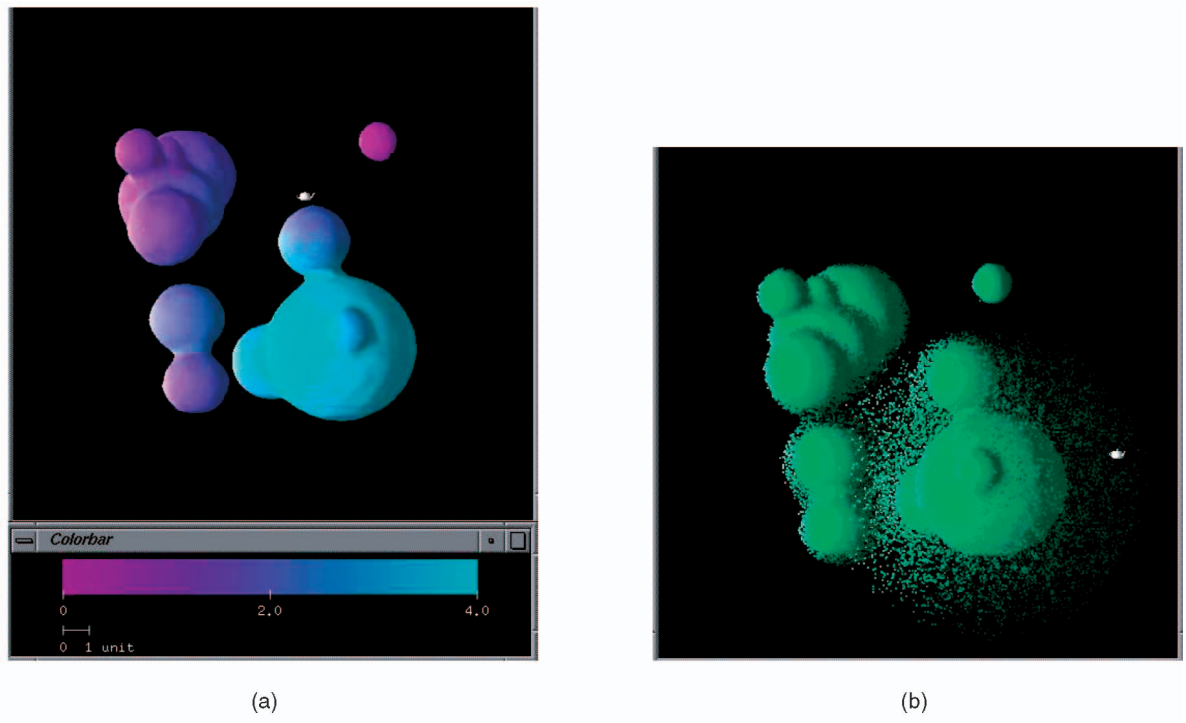

Fig. 8. Sample display screens for user study trials. (a) Pseudocolor trial. (b) Point-based trial.

correct for each condition and their difference, 95 percent confidence interval for each condition interval and their difference, F-value, and p-value.

Subjects made their judgments slightly faster using the point-based representation, an average of three seconds faster per trial, but this difference was not quite statistically significant ($0.05 < p < 0.10$). This represents a likelihood of less than 10 percent that the differences observed resulted from random variation, but, by convention, this is not regarded as statistically significant.

Subject ratings of the ease of use, general satisfaction, and confidence in their answers with the two techniques also show higher ratings for the point-based model, with all differences statistically significant ($p < 0.01$). Table 3 shows mean and 95 percent confidence interval for ease, confidence, and satisfaction with each of the representations and the difference between representations, as well as F-value and p-value.

Comments from subjects also indicated a clear preference for the point-based technique for proximity judgments. Several subjects wished they could combine the two techniques, showing a pseudocolored point distribution. One subject pointed out that pseudocoloring was better for the interpretation of numeric error values, while the point-based technique was better for proximity judgments. A few subjects thought the pseudocolored representations were prettier because they were less cluttered.

11 Conclusions and Future Work

In this paper, we have proposed a technique for visualizing surfaces along with uncertainties associated with regions on them. The basic method relies on point-based rendering of surfaces and displacing individual points according to uncertainty values at the points. Several enhancements to the method are presented. Arbitrary point displacement functions allow for the use of representative probability distributions where available. Lines are provided as an alternative stroke primitive, giving an appearance of something like 3D error bars for surface position. The introduction of an underlying polygonal mesh increases running time efficiency by guaranteeing smoothness in low uncertainty areas without the need to use a large density of points. The addition of transparency as one of the parameters controlled in accordance to uncertainty values makes the illustrations more intuitive by making areas of low and high uncertainty more apparent.

TABLE 2
Analysis of Accuracy with Pseudocolor and Point-Based Representations

	c right	p right	diff
mean	3.9	5.6	1.61
CI	1.0	0.7	1.17
F			9.32
p			<0.01

TABLE 3
Mean Ratings of Eace, Confidence, and Satisfaction on a Scale from -4 to 4

	Ease			Confidence			Satisfaction		
	c	p	D	c	p	D	c	p	D
mean	1.0	2.4	1.7	0.6	2.1	1.5	1.1	2.5	1.4
CI	1.0	0.7	1.0	1.1	1.0	1.0	1.0	0.7	7.2
F			14.0			11.0			8.6
p			<0.01			<0.01			<0.01

Among the advantages of our approach is that it represents uncertainty in a nondistracting manner. We showed that uncertainty did not interfere with surface geometry information and, moreover, it was also possible to effectively visualize at least one additional variable with optional uncertainty. This makes for a six-dimensional data set, which can be effectively displayed using our visualization technique. Another highlight of our method is the fact that uncertainty is represented in an intuitive way. Unlike other possible approaches, ours explicitly connects uncertainty with the variable with which the uncertainty is associated. In the data presented here, uncertainty was associated with location; however, it can potentially be associated with any variable in the visualization as long as this variable is displayed as one of the coordinates of the points. Finally, our method is fast enough to be interactive. This is important for any technique that displays data with more than two dimensions on a 2D medium.

In the future, we would like to experiment by adding more parameters under the control of uncertainty such as specular coefficient or refractive index to see if this gives a more intuitive look to the areas of high and low uncertainty. Additionally, we would like to explore point distribution functions which are more sensitive to surface shape characteristics.

ACKNOWLEDGMENTS

This work was supported by grants from the US National Science Foundation (#121288 and #9996043). The kidney tumor data set is courtesy of Dr. Ron Summers of the US National Institutes of Health. The brain tumor data comes from the Brigham and Women's Hospital Surgical Planning Lab and NSG Brain Tumor Database courtesy of Drs. Simon Warfield, Michael Kaus, Ron Kikinis, Peter Black, and Ferenc Jolesz.

REFERENCES

[1] A. Pang, C. Wittenbrink, and S. Lodha, "Approaches to Uncertainty Visualization," *The Visual Computer,* vol. 13, no. 8, pp. 370-390, 1997.
[2] C. Wittenbrink, A. Pang, and S. Lodha, "Glyphs for Visualizing Uncertainty in Vector Fields," *IEEE Trans. Visualization and Computer Graphics,* vol. 2, no. 3, pp. 226-279, Sept. 1996.
[3] G. Kramer, *Auditory Display, Sonification, Audification, and Auditory Interfaces,* pp. 1-78. Addison-Wesley, 1994.
[4] R. Minghim and A.R. Forrest, "An Illustrated Analysis of Sonification for Scientific Visualization," *Proc. Visualization '95,* pp. 110-117, 1995.
[5] S. Lodha, C. Wilson, and R. Sheehan, "LISTEN: Sounding Uncertainty Visualization," *Proc. Visualization '96,* pp. 189-195, 1996.
[6] A. Cedilnik and P. Rheingans, "Procedural Annotation of Uncertain Information," *Proc. IEEE Visualization '00,* pp. 77-84, 2000.
[7] P. Rheingans and S. Joshi, "Visualization of Molecules with Positional Uncertainty," *Proc. Data Visualization '99,* E. Gröller, H. Löffelmann, and W. Ribarsky, eds., pp. 299-306, 1999.
[8] C. Wittenbrink, A. Pang, and S. Lodha, "Verity Visualization: Visual Mappings," Technical Report UCSC-CRL-95-48, Univ. of California Santa Cruz, 1995.
[9] R. Barnhill, K. Opitz, and H. Pottmann, "Fat Surfaces: A Trivariate Approach to Triangle-Based Interpolation on Surfaces," *Computer Aided Geometric Design,* vol. 9, no. 5, pp. 365-378, 1992.
[10] S. Lodha, R. Sheehan, A. Pang, and C. Wittenbrink, "Visualizing Geometric Uncertainty of Surface Interpolants," *Proc. Graphics Interface,* pp. 238-245, May 1996.
[11] C. Ehlschlaeger, A. Shortridge, and M. Goodchild, "Visualizing Spatial Data Uncertainty Using Animation," *Computers in GeoSciences,* vol. 23, no. 4, pp. 387-395, 1997.
[12] C.M. Wittenbrink, "IFS Fractal Interpolation for 2D and 3D Visualization," *Proc. Visualization '95,* pp. 77-84, 1995.
[13] G. Grigoryan and P. Rheingans, "Probabilistic Surfaces: Point Based Primitives to Show Surface Uncertainty," *Proc. IEEE Visualization 2002,* 2002.
[14] G. Steel, *Growth Kinetics of Tumors.* Oxford: Clarendon Press, 1977.
[15] A. Kansal, S. Torquato, G. Harsh, E. A. Chiocca, and T. Deisboeck, "Simulated Brain Tumor Growth Dynamics Using a Three-Dimensional Cellular Automaton," *J. Theoretical Biology,* vol. 203, pp. 367-382, 2000.
[16] M. Levoy and T. Whitted, "The Use of Points as a Display Primitive," Technical Report 85-022, Univ. of North Carolina at Chapel Hill, Jan. 1985.
[17] M. Gross, "Are Points the Better Graphics Primitives?" *Computer Graphics Forum,* vol. 20, no. 3, 2001.
[18] S. Rusinkiewicz and M. Levoy, "Qsplat: A Multiresolution Point Rendering System for Large Meshes," *Proc. SIGGRAPH 2000,* pp. 343-352, July 2000.
[19] H. Pfister and J. van Baar, "Surfels: Surface Elements as Rendering Primitives," *Proc. SIGGRAPH 2000,* K. Akeley, ed., pp. 335-342, July 2000.
[20] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. Silva, "Point Set Surfaces," *Proc. IEEE Visualization '01,* pp. 21-28, 2001.
[21] W. Reeves, "Particle Systems: A Technique for Modeling a Class of Fuzzy Objects," *ACM Trans. Graphics,* vol. 2, no. 2, pp. 91-108, 1983.
[22] R. Szeliski and D. Tonnesen, "Surface Modeling with Oriented Particle Systems," *Proc. SIGGRAPH 1992,* pp. 185-194, July 1992.
[23] M. Zwicker, M. Pauly, O. Knoll, and M. Gross, "Pointshop 3D: An Interactive System for Point-Based Surface Editing," *ACM Trans. Graphics,* vol. 21, no. 3, pp. 322-329, 2002.
[24] M. Kaus, S.K. Warfield, A. Nabavi, P.M. Black, F.A. Jolesz, and R. Kikinis, "Automated Segmentation of MRI of Brain Tumors," *Radiology,* vol. 218, no. 3, pp. 586-591, 2001.
[25] S.K. Warfield, M. Kaus, F.A. Jolesz, and R. Kikinis, "Adaptive, Template Moderated, Spatially Varying Statistical Classification," *Medical Image Analysis,* vol. 4, no. 1, pp. 43-55, 2000.
[26] W.H. Press, *Numerical Recipes in C: The Art of Scientific Computing,* second ed., p. 290. Jan. 1993.
[27] B. Chen and M.X. Nguyen, "POP: A Hybrid Point and Polygon Rendering System for Large Data," *Proc. Visualization 2001,* pp. 45-52, 2001.
[28] J. Cohen, D. Aliaga, and W. Zhang, "Hybrid Simplification: Combining Multi-Resolution Polygon and Point Rendering," *Proc. IEEE Visualization '01,* pp. 37-44, 2001.

Gevorg Grigoryan received the BS degree in computer science and the BS degree in biochemistry from the University of Maryland, Baltimore County, in 2002, where he worked as an undergraduate research assistant. Currently, he is a PhD student at the Massachusetts Institute of Technology, Department of Biology. His current research interests lie in computational biology and include protein design, de novo structure prediction, electrostatics in protein-protein interactions, and energetics of protein folding.

Penny Rheingans received the PhD degree in computer science from the University of North Carolina, Chapel Hill, and the AB degree in computer science from Harvard University. She is an associate professor of computer science at the University of Maryland, Baltimore County. Her current research interests include uncertainty in visualization, multivariate visualization, volume visualization, information visualization, perceptual and illustration issues in visualization, dynamic and interactive representations and interfaces, and the experimental validation of visualization techniques. She is a member of the IEEE Computer Society