Volume Rendering Data with Uncertainty Information

Suzana Djurcilov[1], Kwansik Kim[1], Pierre F. J. Lermusiaux[2], and Alex Pang[1]

[1]Computer Science Department, UCSC
[2]Division of Engineering and Applied Sciences, Harvard University

Abstract. This paper explores two general methods for incorporating volumetric uncertainty information in direct volume rendering. The goal is to produce volume rendered images that depict regions of high (or low) uncertainty in the data. The first method involves incorporating the uncertainty information directly into the volume rendering equation. The second method involves post-processing information of volume rendered images to composite uncertainty information. We present some initial findings on what mappings provide qualitatively satisfactory results and what mappings do not. Results are considered satisfactory if the user can identify regions of high or low uncertainty in the rendered image. We also discuss the advantages and disadvantages of both approaches.

1 INTRODUCTION

Visualization is used for gaining an understanding of large amounts of data in a short period of time. Scientific datasets often have associated with them a measure of quality, reliability or uncertainty which also needs to be made a part of the visual output.

Uncertainty can be caused by many factors in the data collection and processing: from unreliable instrumentation and problems in transportation, to errors caused by the interpolation and modeling algorithms. While these errors can sometimes be ignored, it is important to alert the users to the trustworthiness of the image upon which they need to make a decision.

While the uncertainty is an essential part of the data, it has often been ignored while processing or displaying. This can be misleading to the user unaware that parts of the dataset contain unreliable information. For accurate interpretation it is important to display the original data together with its uncertainty. Uncertainty visualization techniques present data in such a manner that users are made aware of the locations and degree of uncertainties in their data so as to make more informed analyses.

In this paper we concern ourselves with uncertainty visualization using one particular rendering method, namely direct volume rendering. We present two general options: one which is calculated at the rendering time and presented as part of the volume rendering of the primary value, which we call inline processing; and one which combines the volume renderings of the primary value and of the uncertainty value as a post-processing method.

2 BACKGROUND

Visualizing uncertainty is a recognized challenge in the visualization community, and recently, more visualization research have focused on this area. For example, Cedilnik and Rheingans [1] looked at different ways of imparting uncertainty over 2D fields using procedural methods to distort overlaid grid lines, Interrante [3] discussed how one might use natural textures over a map to show uncertainty, Djurcilov and Pang [2] looked at different ways of incorporating uncertainty information in contour lines and isosurfaces of sparse data sets, Wittenbrink et al. [11] included uncertainty in direction and uncertainty in magnitude into glyph designs, and Pang et al. [7] described some general methods for incorporating uncertainty into visual displays.

The approaches above involve some modification of how the data is represented, and through this modification, impart the uncertainty information. The modifications are typically applied to geometric primitives and attributes such as grid lines, contour lines, glyphs, and textures. Unfortunately, volume rendering does not produce any intermediate geometric primitives that could be modified in order to represent uncertainty. Therefore, this paper explores different alternative techniques for including uncertainty information directly in volume renderings.

3 DATA WITH UNCERTAINTY

3.1 Ocean data and dynamical model

During July and August of 1996, ocean data were collected in the Middle Atlantic Bight (MAB) south of New England, as part of the “ONR Shelfbreak PRIMER Experiment” [6]. The dominant dynamical feature in the MAB consists of a temperature and salinity front, separating the shelf and slope water masses. This front is often located above the shelfbreak, near the 100 m isobath (see Figure 13). It is usually tilted, in the opposite direction of the bottom slope. The main objective was to study the influence of oceanographic variability on the propagation of sound between the shelf and slope regions. Intensive cruise surveys were carried out daily in a 45 km by 30 km domain between the 85 m and 500 m isobaths.

The physical variables or fields are the temperature, salinity, velocity and pressure. They are dynamically evolved by the numerical ocean model of the Harvard Ocean Prediction System [8]. Atmospheric fluxes based on buoy data are imposed in surface.

3.2 Uncertainty forecasts

To dynamically evolve the physical uncertainty, an Error Subspace Statistical Estimation (ESSE) scheme [5] is employed. This scheme is based on a reduction of the evolving error statistics to their dominant components or subspace. Presently, statistics are measured based on a variance or least-squares criterion [9]: a subspace is then characterized by the dominant eigen decomposition of a covariance matrix. The objective is then to dynamically forecast the principal component decomposition of the uncertainty of the physical fields.

In the present MAB case, these error principal components are initialized combining data and dynamics. To account for nonlinearities, they are forecast by an ensemble of Monte-Carlo forecasts.

In the visualizations presented here, only temperature and salinity uncertainty forecasts are used. However, since physical fields are coupled, the effects of velocity errors are included in these forecasts and accurate estimates of temperature and salinity errors can thus be obtained.

As a first endeavor, we utilize the variances of the Monte-Carlo ensemble as a scalar representation for uncertainty at each point.

4 INLINE APPROACH

The classic volume rendering equation is:

$$C(a,b) = \int_a^b E(s)e^{-\int_a^s \delta(x)dx} ds \qquad (1)$$

where $C(a,b)$ is the color intensity contributions through a line from position a to b. E is the color emission function and δ the differential opacity function. Equation (1) calculates an integrated color for each pixel by summing up opacity weighted emittance values. One form or another of this equation is used to generate volume rendered images of 3D scalar fields. Because the data set generally consists of a single scalar field, the same scalar field is used to determine both the opacity and the material emittance values. This is typically achieved by transfer functions that map the scalar data value to both opacity and color.

Because our uncertainty is also represented as a 3D scalar field, we have the opportunity to map field values to color and uncertainty values to opacity, and experiment with different transfer functions. We refer to this approach as “inline” in the sense that the uncertainty information is directly incorporated into the rendering process. In this section, we describe two inline experiments.

4.1 1D Transfer Functions

In this experiment, we mapped salinity values to E and uncertainty values to δ. We then use a 1D transfer function to separately map the salinity and uncertainty values. We experimented with a transfer function that maps increasing uncertainty to increasing opacity, and composited the resulting images to a black background with white grid lines.

As a point of reference, Figure 1 is a traditional volume rendering of the mean salinity field. Figures 2 and 3 both show a volume rendering of the uncertainty in the salinity field. Figure 2 maps uncertainty values above 0.2 to high opacity values, while Figure 3 maps uncertainty values above 0.5 to high opacity values. High uncertainty regions show up as a bluish cloud. Dark regions have lower uncertainty. We note that most of the uncertainty lies along the salinity front on top of the shelfbreak.

One can map uncertainty to opacity in a number of ways. In this experiment, we mapped higher uncertainty values to higher opacity values. Field values such as salinity and temperature are mapped to color. A black background with white grid lines is used to accentuate the fact that more transparent regions have lower uncertainty. Note that

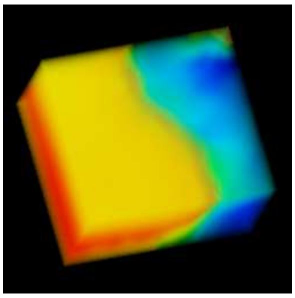

Fig. 1. Mean salinity.

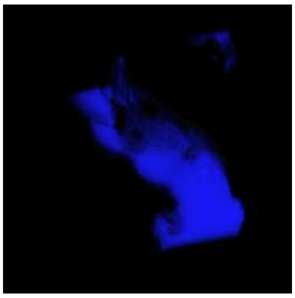

Fig. 2. Uncertainty ≥ 0.2.

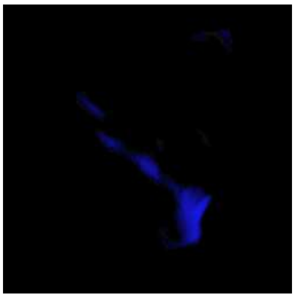

Fig. 3. Uncertainty ≥ 0.5.

regions with low uncertainty do not automatically produce more transparent regions because of the potential occlusion with accumulated opacities from different viewing angles.

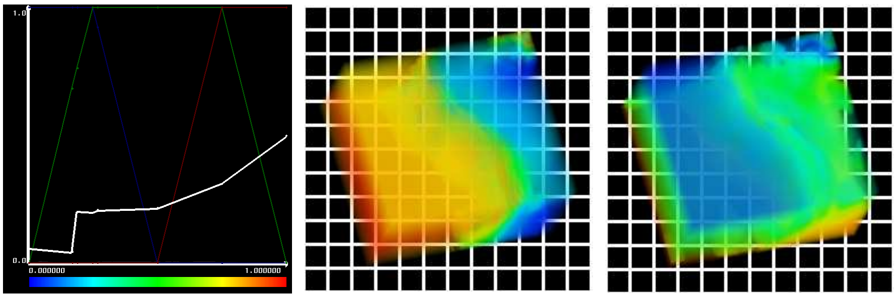

(A) Transfer function

(B) Salinity

(C) Temperature

Fig. 4. (A) Shows the transfer function for both the field values and uncertainty values. All values have been normalized to lie between 0 and 1. The increasing white curve maps higher uncertainty to higher opacity. (B) Volume rendering of the mean salinity field mapped to color and uncertainty in salinity mapped to opacity. (C) Volume rendering of the mean temperature field mapped to color and uncertainty in temperature mapped to opacity.

One can also experiment with an increasing uncertainty to decreasing opacity mapping so that the regions of uncertainty show up as transparent regions rather than opaque regions. The choice of mapping increasing/decreasing opacity seems to depend on the volume data to be studied. Looking at the uncertainty of the temperature field alone confirms that, indeed, the regions of high uncertainty in the right columns of Figures 4 and 10 (color plate) are in the greenish opaque regions. On the other hand, the fine structural details in the uncertainty field are washed out and lost in the resulting rendering. In addition, there is some ambiguity in interpreting the image. The ambiguities can be attributed to a number of factors including varying thickness of the volume from a given viewpoint, the depth within a volume of a region of high uncertainty, interaction of the color and opacity compositing. The image in Figure 10 is similar to Figure

4 except that we used a transfer function which produces more contrast between high and low uncertainty regions. We also removed the white grid lines to see if it is better without them or not.

4.2 2D Transfer Functions

In this experiment, we use 2D transfer functions similar to those used by Kindlmann and Durkin [4]. However, instead of looking at the first and second derivatives of the data, we look at data versus uncertainty values. Figure 5 is a 2D scatter plot showing the distribution of mean salinity versus uncertainty in salinity. We use this 2D scatter plot as the basis for our transfer function, mapping different regions of the scatter plot to different color values. Figures 11 and 12 show different 2D transfer functions and the corresponding volume rendered images of the combined salinity and uncertainty fields.

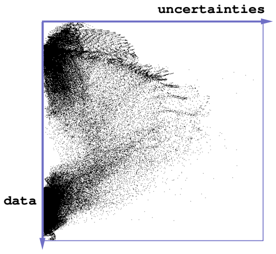

Fig. 5. Scatter plot of mean salinity (Y-axis) versus uncertainty (X-axis). Mean salinity values increase towards the bottom, while uncertainty values increases towards the right.

Unlike 1D transfer functions where we mapped uncertainty to opacity, 2D transfer functions primarily use color to show regions with varying uncertainty. For example, the middle images of Figures 11 and 12 (color plate) use a constant opacity regardless of uncertainty. However, opacity can be used to also emphasize or de-emphasize uncertainty. For example, the right images of Figures 11 and 12 use a step function that maps low uncertainty data to an almost transparent value, and high uncertainty data (greater than 0.2, as in Figure 2) to high opacity. The result is a volume rendering of the salinity data, but with obvious structural features showing the location of the high uncertainty regions. In Figure 12 blue and cyan regions have higher uncertainty. Middle and right images use the same uncertainty to opacity mapping as the corresponding images in Figure 11.

5 POST-PROCESS APPROACH

Due to the use of transparency, images produced by volume rendering algorithms have a soft and smooth quality to them. This aspect lends itself into exploring the use of discontinuity as a means of representing uncertainty. We use discontinuity in several ways by introducing speckles, noise and texture as options used in post-processing of an image to highlight areas where data is uncertain.

5.1 Inserting Speckles/Holes

This task is accomplished in several steps:

1. Produce a standard volume rendering of the field values (see Figure 14).
2. Produce a gray scale volume rendering of uncertainty values from the same viewpoint (see Figure 6). Note that converting a color volume rendering of the uncertainty field to gray scale will not produce the same desired effect.
3. Dither the gray scale rendering into a black and white bitmap image with inverted values (see Figure 7). The purpose of this step is to create a rendering in which each black dot will be a representation of uncertainty in that neighborhood. The dithering itself makes sure that the dots are evenly distributed and visually pleasing.
4. Generate a composite image by multiplying the color volume rendering with the bitmap image pixel by pixel (see Figure 15).

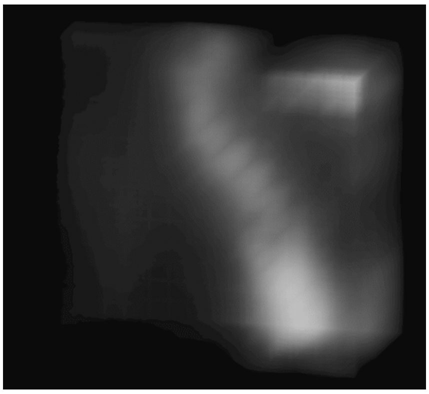

Fig. 6. Gray-scale rendering of uncertainty.

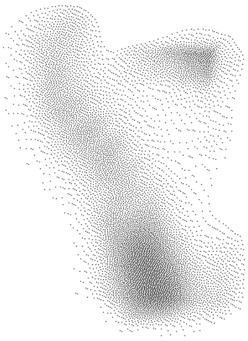

Fig. 7. Inverted bitmap rendering.

Figure 15 shows the outcome of the operations 1 to 4 - an image in which the volume rendering of the primary data value is modified to show pixel-sized holes in areas of high uncertainty. The user is still able to grasp the overall structure of the primary value throughout the dataset, and yet has an understanding of where the data is unreliable.

One possible pitfall of this method is that at a distance the small holes may blend into the image and cause the volume rendering to appear darker in regions with high error. This may be undesirable and can be improved by increasing the size of the holes, thus making it more apparent that the disturbance is not a coloring artifact, but indeed an intended feature of the image. We show one such example in Figure 16 where the holes are increased four-fold in order to emphasize the uncertainty. This image was produced by first sub-sampling the volume rendering of the uncertainty values (output of step 2) by a factor of four, then proceeding with the dithering, after which the image is brought back to its original size and multiplied as in step 4. The end-result is an image where the holes are four pixels large.

We would like to point out that in these examples we have used black to color the holes and match the background color. It would be up to the user to decide the choice of color for the speckles, but we recommend black as an intuitive choice for representing holes.

5.2 Adding Noise

Noise seems a natural option for conveying uncertainty - our minds easily accept the idea that a picture containing noise is less reliable than a clear one. Noise also has the convenient property that it can be introduced into an image without worry for side-effects, as its random nature eliminates any possibility for artifacts appearing as regular patterns.

We apply this idea to the volume rendered images by selectively disturbing the images in the area if high uncertainty. The output (see Figure 17) shows how randomized color dots can be added to an area, thus causing it to appear uneven and fuzzy.

The pseudo code for the algorithm is quite simple.

For each color pixel of the original volume rendering:

1. Find the matching pixel in the uncertainty rendering.
2. Rescale the uncertainty gray shade value to between 0 .. P. P is the probability that the original color will be changed.
3. Replace the original color pixel with a random color with probability P.

This algorithm ensures that the areas with high uncertainty (lighter gray shades) on the uncertainty image are translated into regions with higher numbers of disturbed pixels in the original rendering. The use of probability allows a portion of the pixels to retain their original color even in areas of high uncertainty, so that the overall color context is not lost. In our experiments, we found that setting P to 20 produced a desirable effect. This will preserve at least 80% of the original color pixel values, and yet introduce enough noise in high uncertainty areas. An alternative scheme is to use the uncertainty value as an amount (rather than as a probability) to change the original color value in color space.

5.3 Adding Texture

Similarly to the previous option, we explore the use of textures in the post-processing context. We use 2D grainy, gray scale textures with varying intensity or contrast levels to represent different levels of uncertainty. Low contrast represents low uncertainty, while high contrast represents high uncertainty. We then use the texture brightness (value in HSV space) to alter the brightness of the original color image (value in HSV space). Naturally, in areas of very low or no uncertainty we do not apply any modifications.

The algorithm for adding textures to represent uncertainty in a volume rendered image is also carried out on a per-pixel basis. The difference from the previous method is that the different levels of texture contrast have to be created first. Each texture is tiled so that they are at least as large as the volume rendered image.

For each color pixel of the original volume rendering:

1. Find the matching pixel in the uncertainty rendering.
2. Bin the uncertainty value to one of 5 contrast levels, i.
3. Find the corresponding pixel from texture map i.
4. Adjust the brightness of the original pixel to that obtained from the texture map.

The example in Figure 18 uses a sandstone texture to alter the original volume rendering. Figures 8 and 9 show the sandstone texture at 2 of the 5 different uncertainty levels. In our experiments, we found that 5 levels of contrast to represent different uncertainty levels was sufficient. Beyond 5 levels, it was difficult to distinguish additional levels of uncertainty.

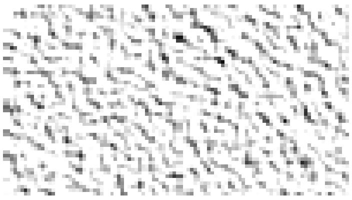

Fig. 8. Low uncertainty texture.

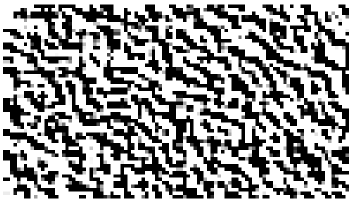

Fig. 9. High uncertainty texture.

6 DISCUSSION AND CONCLUSION

We have described some experiments on how one might include volumetric uncertainty information in a volume rendering. They can be classified as either inline or post-process. Of course, one can also use a pre-process approach where the two volumes are first combined to produce a single scalar volume. Different strategies may be employed to combine the two volumes. For example, one can perform a point wise multiplication of the two fields and volume render the result. In this case, low values would indicate either low data value, low uncertainty value, or both. Converse is true for high values. We did not experiment with this approach because it would be difficult to distinguish between data and uncertainty values in the resulting images.

One can argue which is the better approach: inline or post-process ? The inline method has the advantage that the uncertainty information is integrated into the volume rendering calculation, taking into account their 3D positions within the volume, and hence the results are more faithful. On the other hand, more research is needed to design transfer functions that will unambiguously show the uncertainty information together with the data values. The post-process approach has the advantage of producing images that intuitively show the locations and extent of uncertainty in the volume renderings. However, it is not as faithful to the data in the sense that the uncertainty presentations are really just image embellishments on the volume rendering of the data. For example, if there is a region of high uncertainty embedded within the volume, the post-process approach does not accurately capture the interaction of this region of uncertainty with the corresponding embedded data values.

In this paper, we applied different ideas of incorporating uncertainty into volume rendering using the data set from ocean modeling. Of course, the techniques are applicable to data sets from other domains as well. Some of the questions seeking further research include: How many levels of uncertainty are necessary and can one perceive? What transfer function best combines data and uncertainty, and perhaps their derivatives? And if one has a probability distribution function at each voxel, such as the Monte-Carlo ensemble, how does one go about visualizing such a data set? Finally, while volume rendering does not produce any geometry to be rendered, it does produce derived data in its rendering pipeline. These derived data, when combined with the uncertainty information, can also be used to depict uncertainty information [10]. This approach should also be investigated further.

7 ACKNOWLEDGEMENTS

We would like to thank Craig Wittenbrink for discussions on transfer functions for volume rendering, as well as Dr. Pat Haley and the members of the HOPS group led by Prof. A. R. Robinson for guidance on the HOPS software. We would also like to thank the members of the Advanced Visualization and Interactive Systems laboratory at Santa Cruz for their feedback. We are grateful to the Office of Naval Research for their support during the Uncertainty Pilot Working Group initiative, under grant N00014-00-1-0764 and N00014-00-1-0771. This project is supported in part by NASA grant NCC2-5281, LLNL Agreement No. B347879 under DOE Contract No. W-7405-ENG-48, NSF ACI-9908881, and DARPA grant N66001-97-8900.

References

1. Andrej Cedilnik and Penny Rheingans. Procedural annotation of uncertain information. In *Proceedings of Visualization 00*, pages 77–84. IEEE Computer Society Press, 2000.
2. Suzana Djurcilov and Alex Pang. Visualizing sparse gridded datasets. *IEEE Computer Graphics and Applications*, 20(5):52–57, September 2000.
3. Victoria Interrante. Harnessing natural textures for multivariate visualization. *IEEE Computer Graphics and Applications*, 20(6):6–11, November/December 2000.
4. G. Kindlmann and J.W. Durkin. Semi-automatic generation of transfer functions for direct volume rendering. In *IEEE Symposium on Volume Visualization*, pages 79–86, 170. IEEE, 1998.
5. P.F.J. Lermusiaux. Data assimilation via error subspace statistical estimation, Part ii: Middle Atlantic Bight shelfbreak front simulations and ESSE validation. *Monthly Weather Review*, 127(7):1408–1432, 1999.
6. E. Levy, G. Gawarkiewicz, and F. Bahr. The ONR shelfbreak PRIMER experiment: shelfbreak frontal dynamics in the Middle Atlantic Bight. URL: http://matisse.whoi.edu/primer_cd, 1999.

7. A. Pang, C.M. Wittenbrink, and S. K. Lodha. Approaches to uncertainty visualization. *The Visual Computer*, 13(8):370–390, 1997.

8. A.R. Robinson. Physical processes, field estimation and an approach to interdisciplinary ocean modeling. *Earth-Science Review*, 40:3–54, 1996.

9. A. Tarantola. *Inverse Problem Theory. Methods for Data Fitting and Model Parameter Estimation.* Elsevier Science Publishers, 1987.

10. Craig M. Wittenbrink. IFS fractal interpolation for 2D and 3D visualization. In *IEEE Visualization '95*, pages 77–84, Atlanta, GA, November 1995. IEEE.

11. Craig M. Wittenbrink, Alex T. Pang, and Suresh K. Lodha. Glyphs for visualizing uncertainty in vector fields. *IEEE Transactions on Visualization and Computer Graphics*, 2(3):266–279, September 1996. Short version in SPIE Proceeding on Visual Data Exploration and Analysis, pages 87-100, 1995.

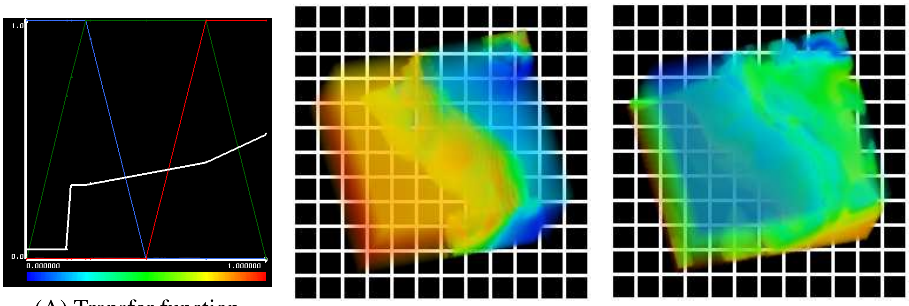

(A) Transfer function (B) Salinity (C) Temperature

Fig. 10. High contrast transfer function.

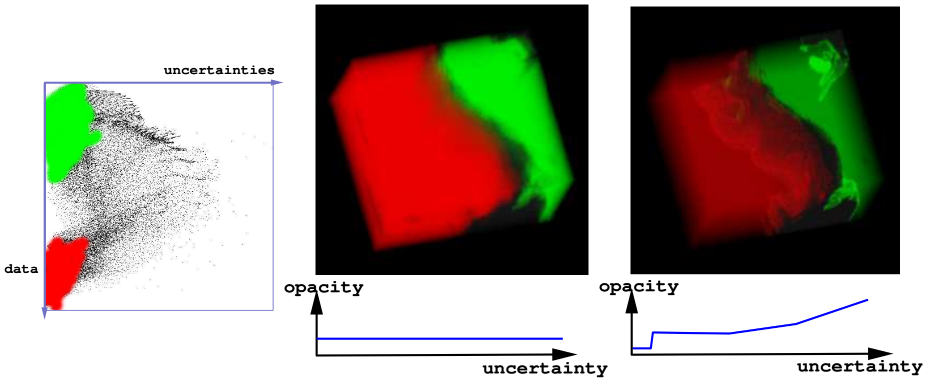

Fig. 11. The scatter plot in Figure 5 is used as a 2D transfer function. Good (low uncertainty) data with low values are mapped to green, while good data with high values are mapped to red. Rest are mapped to gray.

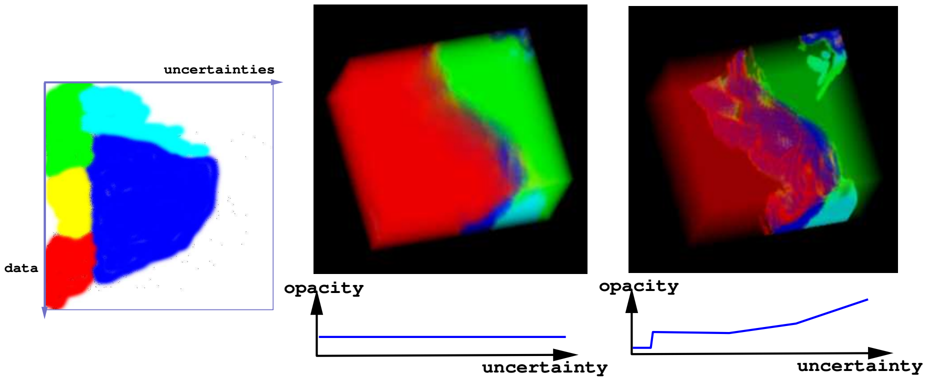

Fig. 12. The 2D transfer function identifies 5 regions instead of just 2.

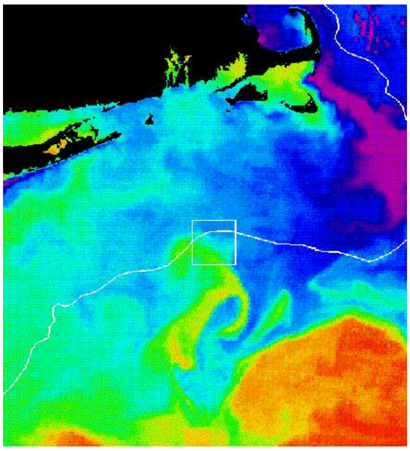

Fig. 13. Surface temperature.

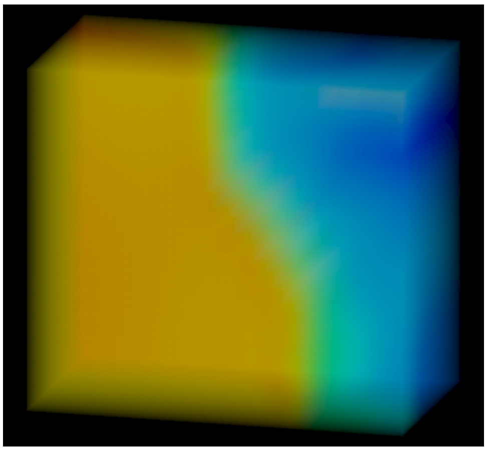

Fig. 14. Rendering of mean salinity.

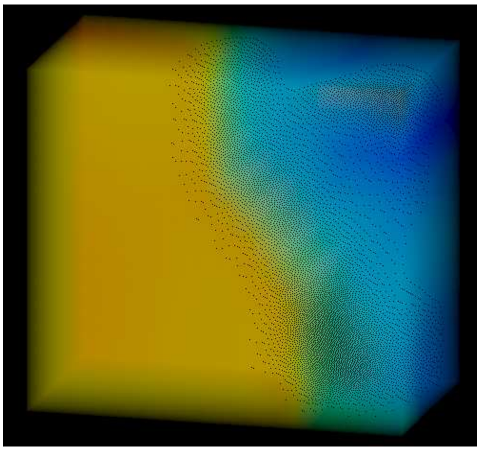

Fig. 15. Rendering with speckles.

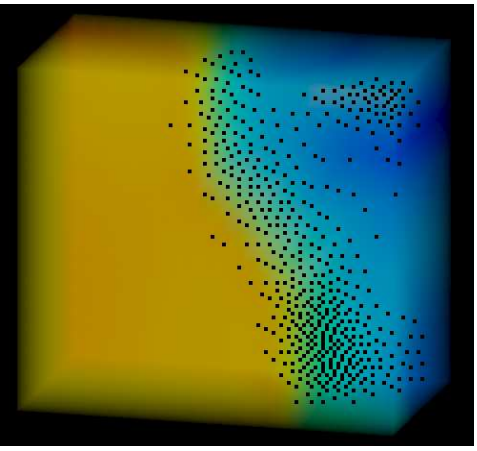

Fig. 16. Larger speckles emphasize holes.

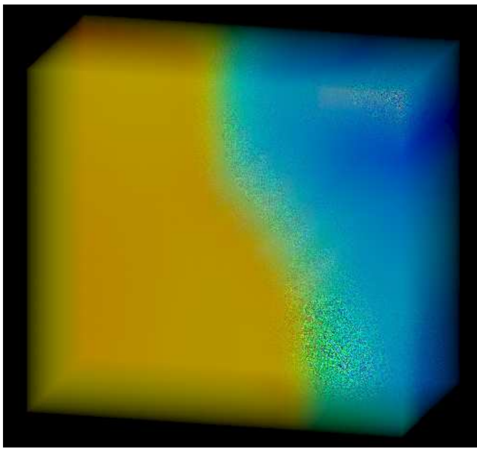

Fig. 17. Noise in high uncertainty areas.

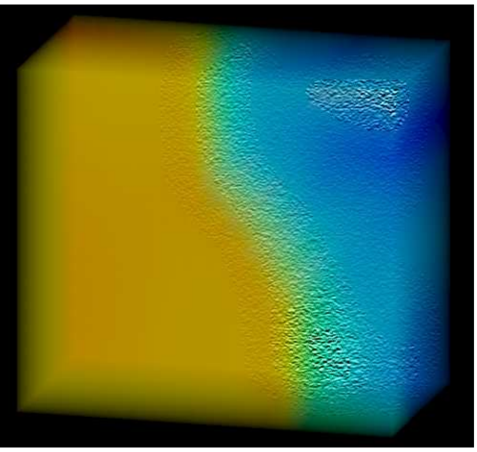

Fig. 18. Texture is another option.