

muView: A Visual Analysis System for Exploring Uncertainty in Myocardial Ischemia Simulations

Paul Rosen and Brett Burton and Kristin Potter and Chris R. Johnson

Abstract In this paper we describe the Myocardial Uncertainty Viewer (muView or μ View) system for exploring data stemming from the simulation of cardiac ischemia. The simulation uses a collection of conductivity values to understand how ischemic regions effect the undamaged anisotropic heart tissue. The data resulting from the simulation is multi-valued and volumetric, and thus, for every data point, we have a collection of samples describing cardiac electrical properties. μ View combines a suite of visual analysis methods to explore the area surrounding the ischemic zone and identify how perturbations of variables change the propagation of their effects. In addition to presenting a collection of visualization techniques, which individually highlight different aspects of the data, the coordinated view system forms a cohesive environment for exploring the simulations. We also discuss the findings of our study, which are helping to steer further development of the simulation and strengthening our collaboration with the biomedical engineers attempting to understand the phenomenon.

Paul Rosen

Scientific Computing and Imaging Institute, University of Utah, 72 S. Central Campus Dr., Salt Lake City, UT 84112, e-mail: prosen@sci.utah.edu

Brett Burton

Scientific Computing and Imaging Institute, University of Utah, 72 S. Central Campus Dr., Salt Lake City, UT 84112, e-mail: bburton@sci.utah.edu

Kristin Potter

University of Oregon, 460 Mckenzie Hall, 5246 University of Oregon, Eugene, OR 97403 e-mail: kpotter@cas.uoregon.edu

Chris R. Johnson

Scientific Computing and Imaging Institute, University of Utah, 72 S. Central Campus Dr., Salt Lake City, UT 84112, e-mail: crj@sci.utah.edu

1 Introduction

Myocardial ischemia is a disease that results from a metabolic imbalance caused by a restriction in blood flow in which the demand for oxygen and nutrients within the muscular tissue of the heart exceeds local supply. Left untreated, cardiac cells will gradually weaken and die; in many cases, leading to heart attack. These consequences make ischemic heart disease the leading cause of death for men and women in the U.S. and most industrialized countries [34]. Detection of cardiac ischemia often requires inspecting the results of an electrocardiogram (ECG) and looking for abnormalities, particularly within the ST segment of the ECG trace. However, the relationship between cardiac ischemia and abnormalities in the ST segment is still unclear [7, 53].

To better understand the underlying physiology of cardiac ischemia, mathematical models are created to study the effect of ischemic regions on cardiac electrical signals, such as the electrical potentials external to the cell. These models have input parameters whose values cannot be practically obtained, and thus a series of parameter perturbations are used to cover a range of possible values. The result of these multi-run simulations is an ensemble of realizations that explores numerous possibilities within the domain and range of the simulation, but also introduces uncertainty within the output. We are currently developing the Myocardial Uncertainty Viewer (muView or μ View) tool for visualizing the output of the cardiac model runs and, more specifically, the uncertainty present within the ensemble.

The goal of μ View is to both directly explore the simulation results, helping scientists design and troubleshoot experiments, and to help understand the relationship of conductivity uncertainties to size and shape estimates of the ischemic zone. The challenges to this goal stem mainly from the complexity of the data; we are given multiple simulation outputs for each voxel. The structure of the data is inherently difficult; the spatial domain of the data is 3D, so simply displaying the data causes occlusion and clutter. Indicating further attributes within the 3D context is a formidable challenge. To address this issue, we have created μ View to experiment with the collection of visualization techniques, including traditional two-dimensional and three-dimensional spatial displays, as well as the incorporation of information visualization approaches, to find a meaningful visual representation. The broader goal of this work is to develop visualization techniques that can concisely express the nature of the uncertainty within this type of complex data for domain scientists and health care professionals alike.

The main contributions of this work stem from the use of multiple visualization approaches to get a sense of the uncertainty in the data. Due to the complexity of the data and the domain, multiple views are employed to allow the user to explore different characteristics of the data. Each visual interface is designed to highlight specific aspects of the uncertainty within the data display, and interactions within a specific display are linked, as appropriate, to the other views. While each specific view is limited in its novelty for displaying this type of uncertainty information, the strength of our technique lies in the combination of our selected views to extract aspects of the uncertainty that is complimentary to the other views and appropriate

for the specific needs of our domain scientists. We feel that this approach contributes knowledge to the field by demonstrating a collection of visualization techniques appropriate for uncertainty information within a 3D spatial domain that addresses the needs of scientists working with this specific type of complex data.

2 Related Work

Interest in uncertainty visualization is steadily increasing [10, 18, 32, 35], and the topic has been identified as a top research problem [19]. Related to this work are techniques aimed at incorporating uncertainty information into volume rendering and isosurfaces, using linked multiple windows, the visual representation of probability distribution functions (PDFs), and displaying the results of parameter-space explorations.

Volume rendering and isosurfacing are techniques designed to convey spatial characteristics of volumetric scalar data. The challenge for these methods is that uncertainty often accompanies data as a multi-valued attribute and applying volume rendering and isosurfacing is not straightforward. Approaches to add uncertainty information to these displays include pseudo-coloring, overlay, transparency, glyphs and animation [4, 23, 30, 44]. Fout and Ma [8] propose a computational model that computes a posteriori bounds on uncertainty propagated through the entire volume rendering algorithm and developed an interactive tool to inspect the resulting uncertainty. Pfaffelmoser et al. derive a mathematical framework for computing confidence levels of gradients of uncertainty parameters to allow for analysis such as the stability of features in an uncertain scalar field [37].

Rather than using isosurfaces to directly convey uncertainty in data, they can be used to show shape and extent of clusters, for example, in the exploration of a supervised fuzzy classification of a 3D feature space plots [29]. Probabilistic formulations of marching cubes [40] and isocontours [39] allow for the display of positional uncertainty of isosurfaces colored by their distance from a mean [38].

While these three-dimensional representations are quite useful for conveying geometric structure and providing context, the complexity of the data often requires multiple presentation types to enable full understanding. For this reason multi-window linked-view systems are popular for addressing uncertainty – uniting a collection of visualization modes. Examples of such systems for understanding uncertainty have been used in magnetic resonance spectroscopy (MRS) [5], multi-dimensional cosmological particle data [11], and weather and climate modeling [43, 50].

Another way to look at uncertainty is to consider the multiple values as PDFs and to use statistical methods for characterizing them [42, 41]. Initial work in the area began by extending existing techniques to work with PDFs [31]. Clustering [2] and slice planes [26] can be used to reduce the dimensionality of the data for visualization, while colormaps, glyphs, and deformations have been used to express summaries and clusters [24, 25].

Finally, the type of data we are analyzing can be thought of in terms of parameter-space exploration in which the effect of perturbations of input parameters is related visually to outcomes through techniques such as parallel coordinates [1], preattentive highlighting [6], and high-dimensional boxplots [48].

3 Background

The electrocardiography forward problem aims to describe torso and cardiac electrical potentials that result from electrically active sources on or within the heart [16, 17, 15]. To this end, computational simulations have been developed to numerically approximate electrical outcomes like cardiac activation times, epicardial potentials and body surface potentials. However, several parameters are required for the simulations, many of which cannot be practically obtained through experimentation or other methods. The estimation of these parameters has the potential to introduce uncertainty into forward cardiac simulations, thereby reducing their potential for clinical applications.

Our parameter of interest is the conductivity of cardiac tissue. Cardiac muscle fibers (known as myocytes) are anisotropic, electrically active cells. Currents pass preferentially along the longitudinal axis of these fibers, investigating action potentials that ultimately cause contraction of the heart muscle. Experimental approaches have been used in an effort to extract appropriate conductivity values, but there are large discrepancies in the literature with regard to these values, which can differ in magnitude from each other by as much as five times [3, 20, 21, 45, 46]. Table 1 shows the range of conductivity values found in the literature that were used for our simulation with $\bar{\sigma}_i$ and $\bar{\sigma}_e$ representing intracellular and extracellular conductivities, respectively. The significance of these conductivity values are explained below.

Table 1 Conductivity Ranges

	Longitudinal		Transmural	
	σ_i	σ_e	σ_i	σ_e
Min	0.00174	0.0012	0.000193	0.0008
Max	0.0034	0.00625	0.0006	0.00236
Ischemic Scaling	1/10	1/2	1/1000	1/4

3.1 Simulation

To solve the forward problem of electrocardiography, three components are linked together for a complete simulation. We defined an ischemic region with reduced transmembrane potential to mimic the ischemic condition, an anatomical cardiac

geometry with associated conductivity properties, and a representation of the electrical activity through the cardiac tissue.

The ischemic region was obtained through experimental methods. Ischemia was induced in a canine model by restricting blood flow to the left anterior descending artery (LAD) of the heart. Plunge, or needle, electrodes with a 1.6 mm transmural resolution were positioned within the LAD vascular bed. Recordings of electrical potentials through the thickness of the myocardial wall were recorded at 1 KHz and potential values during the ST segment of the cardiac cycle were extracted. As ischemia progressed, the ST segment potentials were depressed within the ischemic region. Values that were greater than one standard deviation from baseline values were extracted and used to define the ischemic region.

The same heart in which ischemia was induced was used to extract accurate, cardiac geometry and fiber structure by way of MRI and diffusion weighted tensor imaging (DTI). The canine heart was scanned in a 7 tesla MRI with 0.31 millimeter/pixel resolution. These images were segmented and meshed to represent the cardiac anatomy. DTI was used to determine the preferred diffusion direction within the cells, thereby defining the direction of the striated cardiac muscle.

Electrical activation of the heart was defined by the bidomain equations [12]. This model was adapted to generate cardiac potentials at a single time step, under the influence of ischemia, by reducing it to the passive current flow bidomain (Equation 1) [13]. The bidomain equations represent cardiac tissue by defining, on each node of a simulation, two continuous regions, or domains, that are coupled together by a membrane. The intracellular domain represents the region within cardiac myocytes while the extracellular region defines the extracellular space. Both of these regions are represented by respective conductivity tensors ($\bar{\sigma}_i$ and $\bar{\sigma}_e$). The extracellular and transmembrane potentials, V_e and V_m , are also represented as shown in the equation below. For the purposes of our study, V_m is assumed to have a constant potential difference of -35 mV.

$$\nabla \cdot (\bar{\sigma}_i + \bar{\sigma}_e) \nabla V_e = -\nabla \cdot \bar{\sigma}_i \nabla V_m \quad (1)$$

3.2 Conductivity Values

Given the variability in reported cardiac tissue conductivities [3, 20, 21, 45, 46], we selected and explored a range of conductivity values determined by the minimum and maximum values for longitudinal and transmural conductivities for both $\bar{\sigma}_i$ and $\bar{\sigma}_e$, as observed in the literature. However, when considering evaluation of high-dimensional, randomized parameters, it is necessary to consider both parameter distribution and computational complexity. To address both, we applied *generalized polynomial chaos with stochastic collocation* (gPC-SC) [54]. We used the gPC-SC method to reduce the amount of stochastic collocation points required to accurately compute statistical measures, thereby reducing computational complexity. We treated the conductivity ranges as uniformly distributed, stochastic process. Sub-

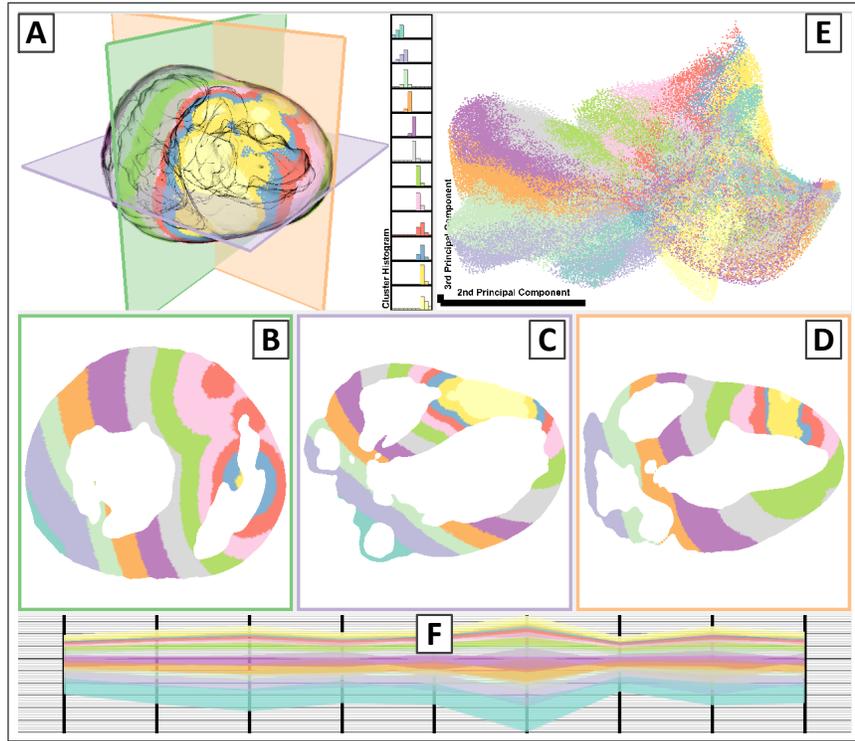


Fig. 1 *K9-2* dataset. Overview of the visualization system which combines 6-way linked views. A: Three-dimensional view, B-D: Two-dimensional views, E: Feature space view of the principal component analysis of the PDFs, F: Parallel coordinates. The data points are colored categorically using k-means clustering with L2-norm distance metric.

sets of values within the range were selected based on Smolyak construction [9, 51], a linear combination of tensor products that span the subspace, while considered the four randomized conductivity parameters and the desired level of representative points. We selected a first level (9 representative conductivity value combinations), second level (41 representative conductivity value combinations), and a third level (137 conductivity combinations) Smolyak representations with which to run the deterministic solutions of the above mentioned bidomain equations.

4 Visualization System

The data input into the visualization system consists of three required components and one optional component. The required data components are a set of vertices, a solid mesh of tetrahedral or hexahedral elements that connect the vertices and form the geometric anatomy of the heart, and finally, an ensemble of simulation

results attached to each vertex in the mesh. Optionally, a supply of tensor directional information can be added.

μ View is an interactive 6-way linked visual analysis system consisting of displays designed to highlight various aspects of the data in two main contexts. The first context is the anatomic or spatial context. In this context, the first view is a three-dimensional visualization of the data (Figure 1 A), while three additional views contain orthogonal two-dimensional slices through the volume (Figure 1 B-D). The second context is the PDF or uncertainty context. To explore this context, a feature space view that depicts the result of principal component analysis (PCA) applied to the PDFs (Figure 1 E), and a parallel coordinates view of the PDFs (Figure 1 F) shows each dimension of the data. Data point colors are held consistent across all visualizations for visually cross-correlating data between views and contexts (i.e. points receive the same coloring as their counterparts in other views). The interfaces are manipulated through mouse interactions and a small menu system (not shown).

4.1 Visual Interfaces

Three-Dimensional View (Figure 1 A): Anatomic context is important in many medical applications, and thus we provide a three-dimensional rendering of the geometry. We extract a surface mesh from the volume by selecting tetrahedral faces that are not shared. Using the surface mesh, we darken triangle faces that are near perpendicular to the view direction to help contextualize the data within the three-dimensional shape, as seen in Figure 3 left. This method has the advantage of showing both the surface and cavities of the heart. However, it can be difficult to determine the orientation and shape of the heart without interacting with it. Other more illustrative rendering methods [28] could be inserted in its place.

Within this interface, the data can be visualized either through a series of iso-surfaces as discussed in Section 4.2 or by rendering the data points colored via a transfer function as discussed in Section 4.3. The three colored planes serve as a cross-reference mechanism to the two-dimensional views and can be switched on or off as desired.

Two-Dimensional View (Figure 1 B-D): For a view of the data without perspective distortion, two-dimensional visualization slices of the volume are extracted by intersecting the solid elements with a plane and linearly interpolating vertices, triangles, and PDFs. The three slice planes, axial, coronal, and sagittal, are a more natural way for health care professionals to view the data, and are displayed using transfer functions to color the mesh with isolines to help highlight the variations in value. Isolines are extracted from the data with their frequency controlled by the user (i.e. every 0.5 step in data value). The orientation planes, optionally visible in the three-dimensional visualization, are color coordinated with the borders surrounding the views. These assist users in identify the three-dimensional spatial location of the slices and cross-referencing phenomena of interest.

Feature Space View (Figure 1 E): A wide-variety of dimensionality reduction methods [27, 36, 49] allow for the conversion from a difficult to visualize and interpret high-dimensional space, to an easier to understand low-dimensional space. We have employed this approach in the feature space view with a goal of preserving as many features as possible.

We treat our PDFs as high-dimensional points and reduce them down to two-dimensional points. Principle component analysis (PCA) [36] is used as our dimensionality reduction approach. PCA works by extracting a vector representing the strongest component of the dataset. It then finds the next strongest vector which is orthogonal to the first. This process can continue to any number of dimensions. Our interface allows selecting any pair of principle components (not just the first two, see Figure 2) to generate the display. This allows exploring the high-dimensional space for features which might not otherwise be visible in the spatial domain.

The coordinate system is placed in the lower left, with the length of the axes indicating the scaling of that principal component. For example, in Figure 2 right, the second principal component (horizontal component) is significantly larger than the third principal component (vertical component). Therefore, the vertical axis is stretched, indicating that the natural dimensionality of the data may be reached.

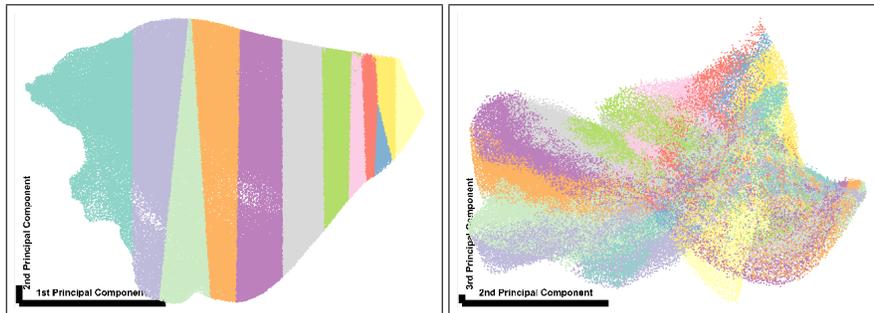


Fig. 2 *K9-2* dataset. Visualization of PCA using the first and second principal components (left) and the second and third principal components (right). The data points are colored categorically using k-means clustering with L2-norm distance metric.

Parallel Coordinate View (Figure 1 F): Parallel coordinates are an alternative method for exploring the full high-dimensional space of the data. We supply an interactive parallel coordinates interface where each dimension represents a single simulation. Again, the data points (represented as lines in parallel coordinates) are colored using the same transfer function as in the other linked-views for easy cross-correlation. For example, in Figure 1 it can be seen in the parallel coordinates that the yellow colored clusters are likely enclosing the ischemic zone. Similar observations are true of the other clusters as well, though some are difficult to see because of their limited size in the parallel coordinate view. We also enable interactive selection of the data points and reordering of axes within the parallel coordinates view for further exploration.

4.2 Isosurfaces over the PDFs

The range of values for an individual data point makes isosurface location unclear [14]. Each dimension may maintain its own isosurface for a given isovalue, meaning the isosurface for the PDF could exist anywhere within a range of locations. To account for this, we reduce the PDF to a single dimension by applying an operator, such as minimum, maximum, or mean, to the data samples at each point. Isosurfaces are then extracted from the single dimensional field. Figure 3, left, shows an example where the minimum isosurface is blue and the maximum isosurface is yellow implying an envelope containing the range of possible surfaces.

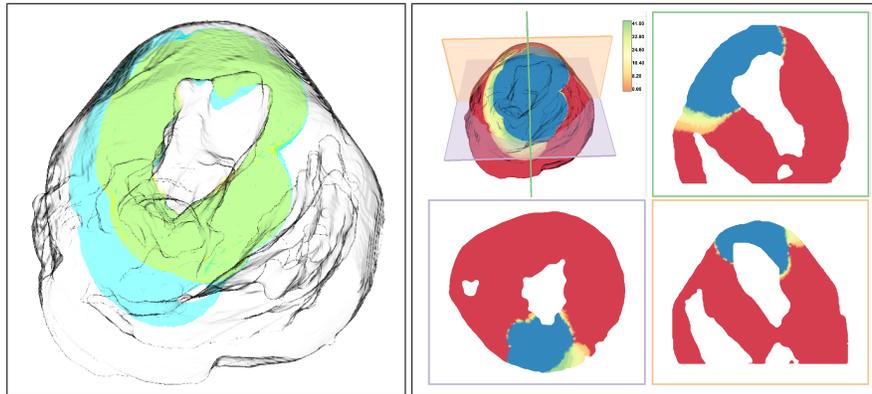


Fig. 3 *K9-1 dataset*. Left: Isovalue visualization is used to describe the range of potential isosurfaces of a given value with the minimum operator (in blue) and maximum operator (in yellow), combined with silhouettes (in black) for anatomic context. Right: Isovalue color mapping shows regions above (blue), below (red), and in-between (yellow) an isovalue.

4.3 Transfer Functions over the PDFs

Because of the complexity of the data, we have adopted a number of transfer functions to color the data, each designed to aid understanding in a unique way.

Value-based Coloring: The first transfer function simply assigns a single value to each data point and applies an intensity-based sequential color map. The values can be related to individual dimensions, or derived values such as the mean (Figure 4 left) or standard deviation (Figure 4 right).

Coloring by Isovalue: We have also explored coloring points by isovalue. This method takes each PDF and counts the number of dimensions above and below the isovalue. In this scheme, we choose solid colors to represent data points where all

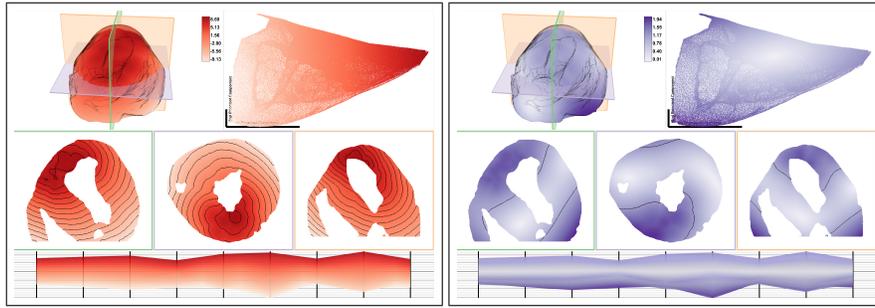


Fig. 4 *K9-1* dataset. Single value, sequential color mapping of mean value (left) and standard deviation (right).

dimensions were above (red) or below (blue) the isovalue. The remaining points are colored using a sequential color map (orange to blue-green) which partially indicates how many dimensions fall above or below the isovalue. An example is seen in Figure 3, right.

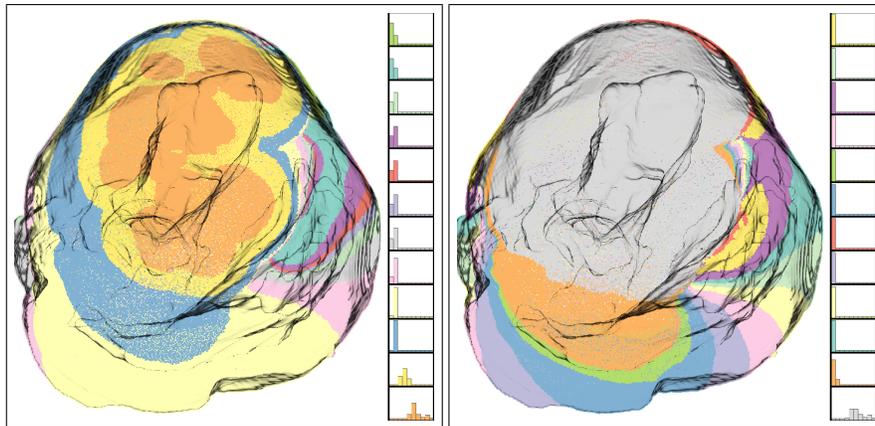


Fig. 5 *K9-1* dataset. K-means clustering applied to the data using the L2-norm (left) and Pearson correlation coefficient (right) as distance metrics.

Coloring by Clustering: Clustering can reduce the set of data under investigation by grouping similar data together, such as points that respond similarly to variations in initial conditions. As points are placed into clusters, they are colored using a categorical color map. A collection of histograms showing the mean of each cluster is placed to the right.

We use k-means clustering [33] to exploring this space. We employ multiple distance metrics for comparing the underlying PDFs. The L2-norm (Figure 5 left) groups points that are similar in a Euclidean sense and is defined by $d(X, Y) =$

$\sqrt{\sum(X_i - Y_i)^2}$, where X and Y each contain the ensemble of simulation dimensions of the data points.

Pearson correlation coefficient [47] (Figure 5 right) clusters points that respond similarly to changes in input and is defined as $d(X, Y) = 1 - \frac{\sum(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum(X_i - \bar{X})^2} \sqrt{\sum(Y_i - \bar{Y})^2}}$, where \bar{X} and \bar{Y} are the means of the sets.

Selecting the number of clusters is a well known challenge to the k-means approach. Initially, the number of clusters is selected using an information-theoretic approach [52] where $k = \log(n/2)$. The user is then given the opportunity to adjust the number of clusters based upon their intuition about the data.

Painting Over Feature Space: We give the user the opportunity to interact indirectly with the high-dimensional data through a painting interface included in the feature space view. The interface allows users to select a paint color, brush over a region of interest, and see the resulting color change in both the feature and other domain visualizations. An example is seen in Figure 6, where an interesting structure in feature space has been selected in various shades of purple. In this case, it turns out that the structures are related to the ischemic zone of the heart.

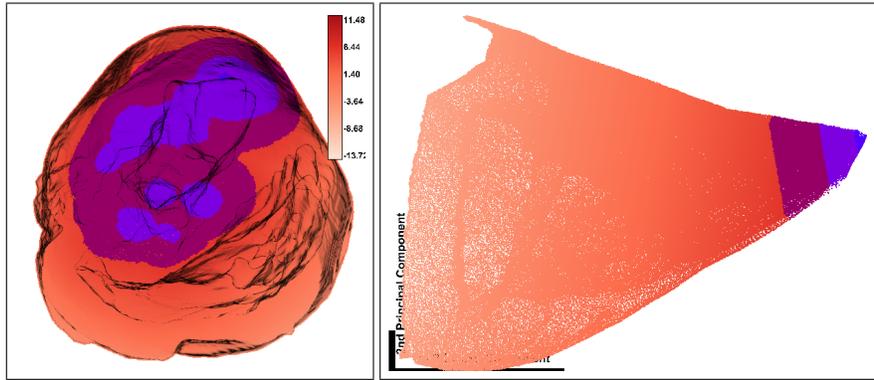


Fig. 6 *K9-1 dataset*. The result of painting the feature space (right) correlates a structure to its spatial counterpart (left). In this case, the ischemic zones are identified in the first principal component.

Coloring by Fiber Direction: To better understand the impact of input versus output, we have included a common method of visualizing input fiber directions. In this visualization, red, green, and blue are assigned to x, y, and z, respectively. This gives the opportunity to correlate fiber directions with structures in either the feature or spatial domains. Figure 7 shows an example of this visualization.

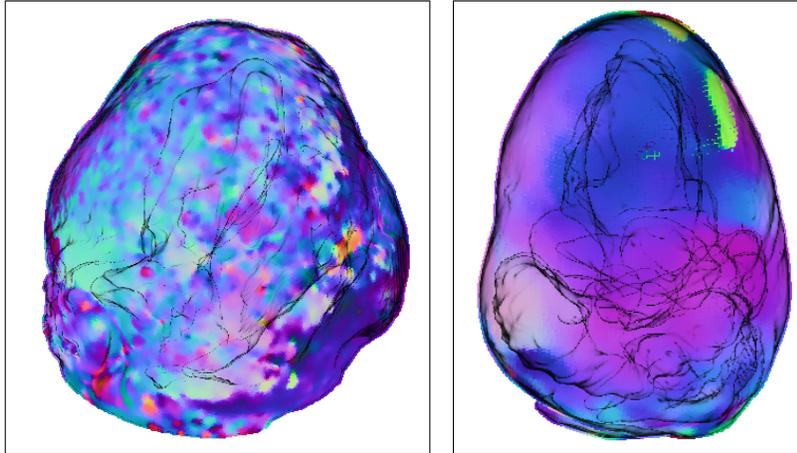


Fig. 7 Coloring by fiber directions in the three-dimensional view shows noisy MR diffusion tensor data used as input for the *K9-1 dataset* (left). A more anatomically realistic model would be expected to vary more smoothly across the heart such as that of *K9-2 dataset* (right).

5 Results

Simulations were performed using two separate data sets. The first, “*K9-1*”, was created by MRI segmentations of an excised canine heart with atrial tissue removed. The mesh consists of 1.4M tetrahedra and 350K vertices. In like manner, the second, “*K9-2*”, was created by using the same imaging techniques from a different canine dataset. This mesh contains 2.5M tetrahedra connecting 435K vertices. Ischemic regions in both meshes were determined by thresholding measured potential values observed within the heart under experimental, ischemic conditions. Thresholds were determined to be within 1 standard deviation of control, baseline ECG values observed in the ST segment. As part of our experiments, each mesh was used to calculate first, second, and third order gPC-SC which produced 9, 41, and 137 realizations, respectively.

Performance is reported on Windows-based system with a 3.2GHz Intel Core i7 CPU, 8 GB of RAM, and an NVIDIA GeForce GTX 580. It is important to note, we also able to run all of these experiments on a MacBook Air. The performance was slightly slower but still respectable. Our software takes a few moments to load data (15-30 seconds for smaller data and a few minutes for large data) and a few seconds to apply some operators (1-60 seconds; clustering being the slowest). OpenMP is used where possible to take advantage of the multicore environment. The individual OpenGL contexts can run at 50-100 frames per second, though there is a slowdown when all 6 contexts are rendered simultaneously. Nevertheless, the system remains interactive.

We developed this tool as part of a collaboration between visualization and biomedical researchers to better understand the physiology of cardiac ischemia.

μ View is being actively developed simultaneously with the development of the simulation model, allowing: results from the simulation to be explored within μ View; insights gleaned from μ View to be incorporated back into the simulation; and μ View revised and refined based on feedback from collaborators.

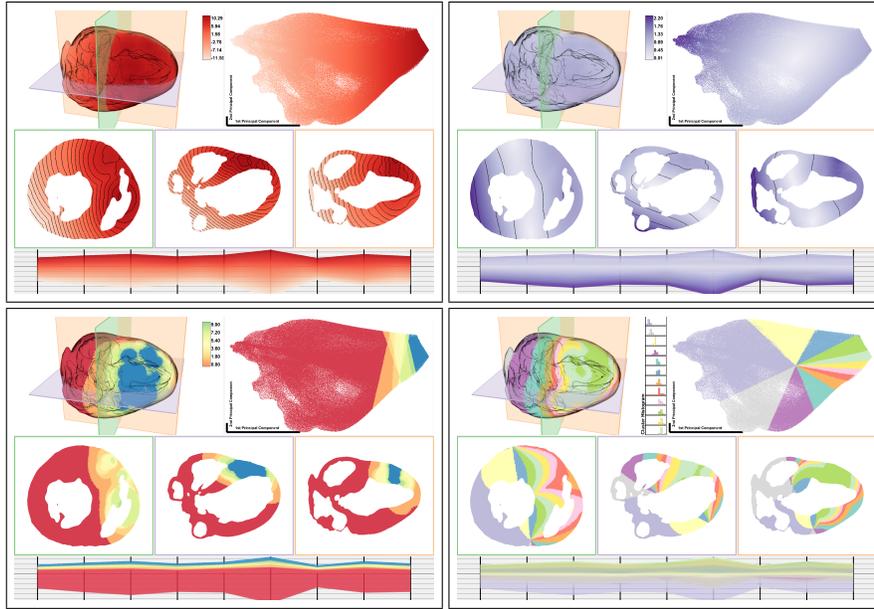


Fig. 8 *K9-2* dataset. The mean (top left) and standard deviation (top right) visualizations often challenge the user searching for ischemic zones, while the isovalue (bottom left) and L2-norm clustering (bottom right) can enable quick identification.

5.1 Identifying Ischemic Zones

Identifying the ischemic zones is one of the first tasks we engaged in when exploring the datasets. Many of the visualizations do not immediately enable access to this type of information. For example, in Figure 8 top, the mean and standard deviation visualizations do not clearly delineate the ischemic zones. However, by exploring the isovalues of the data we can better understand the location of diseased tissue. Thresholding is a common approach used by biomedical engineers for capturing the shape of the ischemic zone. By selecting an isovalue, this visualization performs a pseudo thresholding while providing additional uncertainty information (Figure 8 bottom left). It turns out that the L2-norm is also very useful in identifying ischemic zones. Figure 8 bottom right shows an example where the light green cluster represents the ischemic zones. The clustering approach also has the advantage of being

almost entirely automatic. Finally, feature space painting can be used to potentially identify the ischemic region. As seen in Figure 6, the first principal component primarily encodes a value that when explored, enables selecting the ischemic region. Our hypothesized ischemic regions can then be confirmed using the parallel coordinates view to note that the cluster or region has the highest voltage value, a known indication of the ischemic zone.

5.2 Detecting Noisy Fiber Orientations

The conductivities of the heart are highly dependent on the fiber directions across the tissue. The fiber direction data can be created any number of ways, such as rule-based methods or, as in our case, using DTI. Part of our study is understanding the impact of conductivity to fiber direction.

As we began our study, we noticed a bulge in the *K9-1* dataset for many isovalues (most prominently visible as the yellow area in Figure 3 right) that we could not easily explain. We then dove into the data by directly visualizing some of the input, such as fiber direction. This led us to discover that the fiber directions from our DTI imaging were noisy and poorly aligned for the first few millimeters of the heart surface (see Figure 7 left).

This finding drove us to obtain a secondary model, the *K9-2* dataset, which has much smoother fiber directional data (see Figure 7 right). In this new model, no such irregular bulge is visible, leading us to believe that the feature was caused by the noisy tensors in the original data.

5.3 Comparing Higher-Order Methods

For our experiments, we ran first- through third-order gPC-SC with the goal of determining if there were significant advantages in higher-order methods over lower-order ones. Figure 9 show the results for the *K9-1* dataset, and Figure 10 shows the results for the *K9-2* dataset. Both compare first- and third-order gPC-SC.

Both the first- and third-order versions of the data appear virtually identical. In *K9-1* (Figure 9), there is a mirroring in the second principal component in the feature space visualization. However, the shape of the feature space is virtually identical. Otherwise, only the slightest of differences can be seen with close inspection.

In the *K9-2* dataset (Figure 10), the only easily identifiable difference lies in the L2-norm clustering (row 3). However, this difference may be attributed as an artifact of initial mean selection in the k-means clustering algorithm.

The results of both of these experiments convinced our team that for conductivity-based uncertainty, higher-order gPC-SC does not provide a significant variation or advantage over the lower-order method. This assertion might not hold for other ischemia related uncertainties, such as positional uncertainty of the ischemic zone.

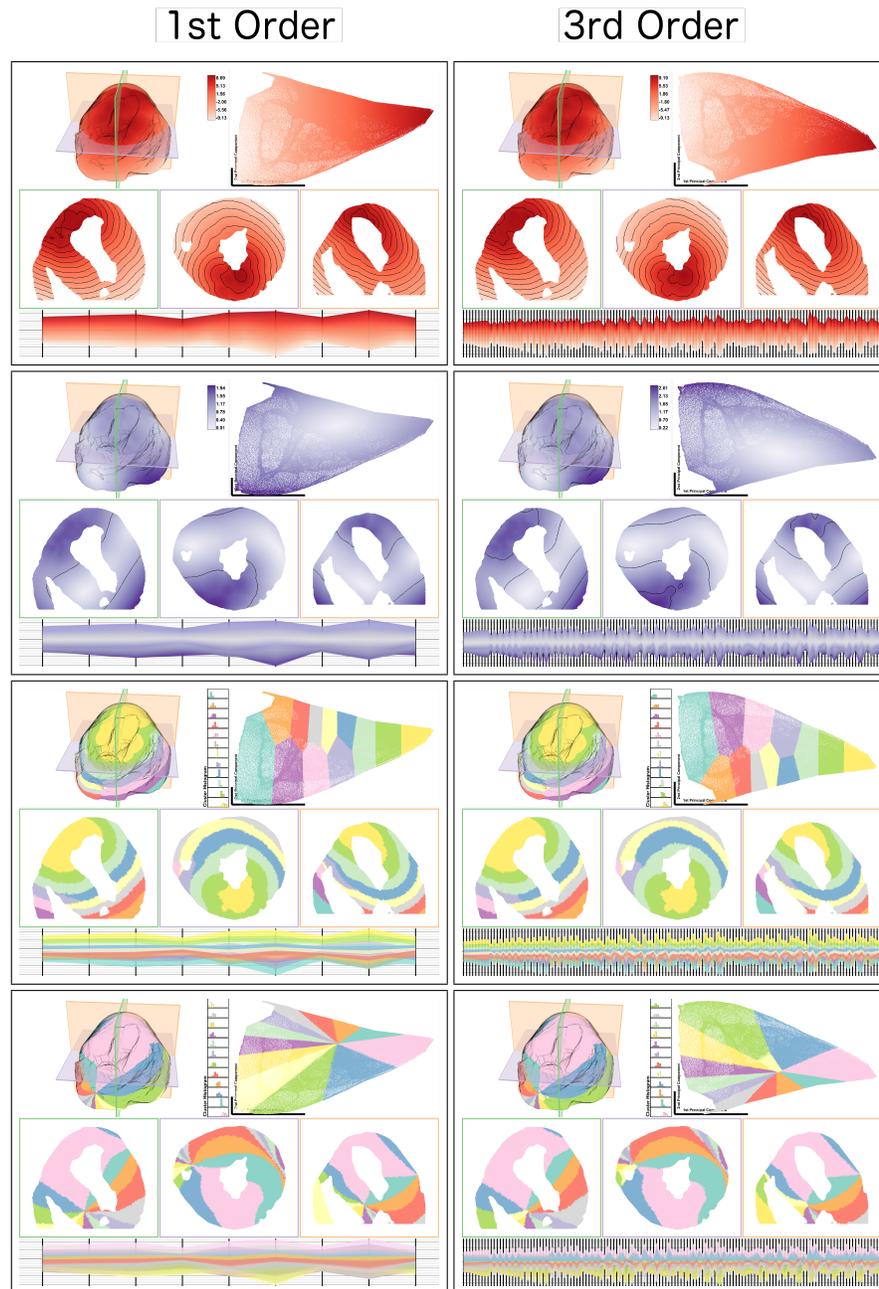


Fig. 9 *K9-1* dataset. Comparison of the first- (left) and third-order (right) gPC-SC through mean (row 1), standard deviation (row 2), L2-norm clustering (row 3), and Pearson correlation clustering (row 4) visualizations.

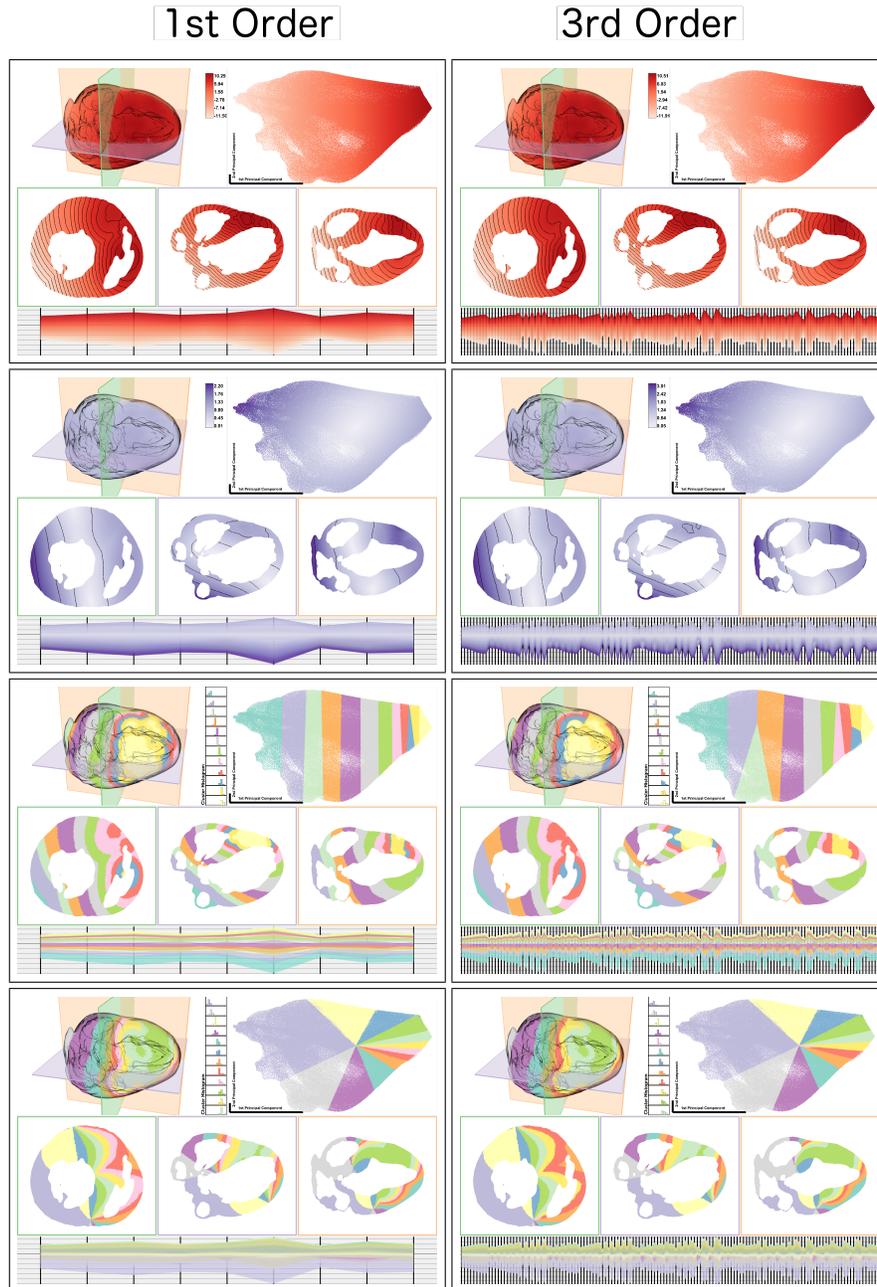


Fig. 10 *K9-2 dataset*. Comparison of the first- (left) and third-order (right) gPC-SC through mean (row 1), standard deviation (row 2), L2-norm clustering (row 3), and Pearson correlation clustering (row 4) visualizations.

5.4 *Uncertainty from Conductivity*

The final issue we addressed was whether the uncertainty is in some way meaningful. Our conclusion after exploring the data using many visualization methods over a significant period of time is that, no, the uncertainty in this data is not particularly meaningful. We believe that the results can be explained primarily through scale and bias. This is most apparent in the parallel coordinates view of Figure 1. The fact that the clustered colors maintain constant tracks is a clear indicator of the scale and bias effect. The other indication is that in principal component analysis, the first principal component almost entirely summarizes the voltage. Additionally, the secondary principal components are of significantly smaller scale, indicating that the underlying dimensionality of the data is probably just one or two (see Figure 2).

6 Discussion and Future Work

The ultimate clinical goal of this work is to be able to assess the uncertainty in determining the ischemic zone from an inverse simulation linked to ST segment waveform changes. If we better understand how the uncertainty in the conductivities affect the size and shape of the ischemic zone, it will help determine what levels of uncertainty will be of consequence clinically and how much confidence can be assigned to the understanding of the ischemic zone size, shape, and location.

From a scientific point of view, these studies can also give us a better understanding of the relationship of conductivity uncertainty to both forward and inverse simulations of cardiac ischemia. In the future, we aim to provide confidence criteria of the simulation results as a function of both conductivity uncertainty and the problem we are trying to solve. It may be that for some problems, the level of uncertainty will not greatly effect the results, while the uncertainty may invalidate other applications. It may also indicate that uncertainty levels in the conductivities would have to be reduced in order to use such a method for a particular application, which could then spark research into generating better conductivity values.

6.1 *Sources of Uncertainty*

We are able to identify a wide variety features associated with the location and magnitude of uncertainty. However, the source of the uncertainty remains largely a mystery (i.e. the classical correlation vs. causation problem). In Section 5.2, we suspected that the DTI noise was the source of the error, but were only able to find indirect evidence to validate our suspicion. Such a discovery encourages further investigation of the importance of fiber direction in our simulations, which we plan do through techniques such as those proposed by Jones [22] for visualizing uncertainty in fiber orientation. More generally, visualizations that indicate correlation and cau-

sation of uncertainty features will greatly improve the efficiency of these types of studies.

6.2 Repeated Features

We have presented a number of visualization approaches, many of which highlight the same features. Each visualization has its own advantages and disadvantages. Clustering has the advantage of highlighting multiple features simultaneously. However, it requires significant effort in visual search to wade through less important features. Using the isovalue visualization limits the number of features visible, making concentration easier, but requiring additional interaction. It is increasingly important to find visualizations which balance these modes of operation and identify which types of visualizations are most efficient from the perspectives of speed, accuracy, and cognitive load. Providing users with choice in visualization is valuable, but too much choice will overwhelm.

6.3 When have we found all features?

The last issue arising from this work that often plagues visualizations is identifying when all of the data features have been located. We have presented visualizations that find a wide variety of features, but for a long time we lacked the confidence to claim the search had been exhaustive and all interesting features had been found. Making such claims demands further validation of which types of features are and are not identifiable by each visualization method.

7 Summary

This work is an initial exploration of uncertainty data obtained from the forward simulation of cardiac ischemia. We believe our close collaboration with the biomedical simulation scientists will greatly guide the choices we make regarding visualization, particularly in light of how our system has, to date, helped improve our understanding of the simulation. Both clinical problems and scientific exploration provide opportunities for improvement in uncertainty visualization techniques, and we look forward to extending μ View to have greater research and clinical impact.

Acknowledgements This project was supported by grants from the National Center for Research Resources (5P41RR012553-14), National Institutes of Health's National Institute of General Medical Sciences (8 P41 GM103545-14), DOE NETL, and King Abdullah University of Science and Technology (KUS-C1-016-04).

References

1. Berger, W., Piringer, H., Filzmoser, P., Gröller, E.: Uncertainty-aware exploration of continuous parameter spaces using multivariate prediction. *Comp. Graph. Forum* **30**(3), 911 – 920 (2011)
2. Bordoloi, U.D., Kao, D.L., Shen, H.W.: Visualization techniques for spatial probability density function data. *Data Science Journal* **3**, 153–162 (2004)
3. Clerc, L.: Directional differences of impulse spread in trabecular muscle from mammalian heart. *J. Physiol.* **255**, 335–346 (1976)
4. Djurcilov, S., Kim, K., Lermusiaux, P., Pang, A.: Visualizing scalar volumetric data with uncertainty. *Computers and Graphics* **26**, 239–248 (2002)
5. Feng, D., Kwock, L., Lee, Y., II, R.M.T.: Linked exploratory visualizations for uncertain mr spectroscopy data. *SPIE Vis. and Data Analysis* **7530**(4), 1–12 (2010)
6. Feng, D., Kwock, L., Lee, Y., II, R.M.T.: Matching visual saliency to confidence in plots of uncertain data. *IEEE TVCG* **16**(6), 980–989 (2010)
7. Fleischmann, K.E., Zgre-Hemsey, J., Drew, B.J.: The new universal definition of myocardial infarction criteria improves electrocardiographic diagnosis of acute coronary syndrome. *J Electrocardiol.* **44**, 69–73 (2011)
8. Fout, N., Ma, K.L.: Fuzzy volume rendering. *IEEE TVCG* **18**(12), 2335–2344 (2012)
9. Geneser, S., Hinkle, J., Kirby, R., Wang, B., Salter, B., Joshi, S.: Quantifying variability in radiation dose due to respiratory-induced tumor motion. *Medical Image Analysis* **15**(4), 640–649 (2011)
10. Griethe, H., Schumann, H.: Visualization of uncertain data: Methods and problems. In: *SimVis*, pp. 143–156 (2006)
11. Haroz, S., Ma, K.L., Heitmann, K.: Multiple uncertainties in time-variant cosmological particle data. In: *IEEE PacificVis*, pp. 207–214 (2008)
12. Henriquez, C.: Simulating the electrical behaviour of cardiac tissue using the bidomain model. *Crit. Rev. Biomed. Eng* **21**(1), 1–77 (1993)
13. Hopenfeld, B., Stinstra, J., Macleod, R.: Mechanism for st depression associated with contiguous subendocardial ischemia. *Journals of Cardiovascular Electrophysiology* **15**(10), 1200–1206 (2004)
14. Jiao, F., Phillips, J., Gur, Y., Johnson, C.: Uncertainty visualization in HARDI based on ensembles of ODFs. In: *IEEE PacificVis*, pp. 193–200 (2012)
15. Johnson, C.: Numerical methods for bioelectric field problems. In: J. Bronzino (ed.) *The Biomedical Engineering Handbook*, pp. 161–188. CRC Press, Boca Ratan (1995)
16. Johnson, C., MacLeod, R., Matheson, M.: Computer simulations reveal complexity of electrical activity in the human thorax. *Computers in Physics* **6**, 230–237 (1992)
17. Johnson, C., MacLeod, R., Matheson, M.: Computational medicine: Bioelectric field problems. *IEEE Computer* **26**(26), 59–67 (1993)
18. Johnson, C., Sanderson, A.: Next step: Visualizing errors and uncertainty. *IEEE CG&A* **23**(5), 6–10 (2003)
19. Johnson, C.R.: Top scientific visualization research problems. *IEEE CG&A* **24**(4), 13–17 (2004)
20. Johnston, P.R.: A cylindrical model for studying subendocardial ischaemia in the left ventricle. *Mathematical Biosciences* **186**, 43–61 (2003)
21. Johnston, P.R., Kilpatrick, D.: The effect of conductivity values on st segment shift in subendocardial ischaemia. *IEEE Trans. Biomedical Eng.* **50**, 150–158 (2003)
22. Jones, D.K.: Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor mri. *Magnetic Resonance in Medicine* **49**, 7–12 (2003)
23. Josphe, A.J., Lodha, S.K., Renteria, J.C., Pang, A.: Uisurf: Visualizing uncertainty in isosurfaces. In: *Computer Graphics and Imaging*, pp. 184–191 (1999)
24. Kao, D., Dungan, J.L., Pang, A.: Visualizing 2d probability distributions from eos satellite image-derived data sets: a case study. In: *IEEE Vis*, pp. 457–561 (2001)

25. Kao, D., Kramer, M., Luo, A., Dungan, J., Pang, A.: Visualizing distributions from multi-return lidar data to understand forest structure. *Cartogr. J.* **42**(1), 35–47 (2005)
26. Kao, D., Luo, A., Dungan, J.L., Pang, A.: Visualizing spatially varying distribution data. In: *Information Visualisation*, pp. 219–225 (2002)
27. Kruskal, J., Wish, M.: *Multidimensional scaling*. Sage University Paper series on Quantitative Application in the Social Sciences **07-011** (1978)
28. Lawonn, K., Moench, T., Preim, B.: Streamlines for illustrative real-time rendering. *Computer Graphics Forum* **32**(3pt3), 321–330 (2013)
29. Lucieer, A.: Visualization for exploration of uncertainty related to fuzzy classification. In: *IEEE Intl Conf on Geoscience and Remote Sensing*, pp. 903–906 (2006)
30. Lundström, C., Ljung, P., Persson, A., Ynnerman, A.: Uncertainty visualization in medical volume rendering using probabilistic animation. *IEEE TVCG* **13**(6), 1648–1655 (2007)
31. Luo, A., Kao, D., Pang, A.: Visualizing spatial distribution data sets. In: *Symp. on Data Vis.*, pp. 29–38 (2003)
32. MacEachren, A.M., Robinson, A., Hopper, S., Gardner, S., Murray, R., Gahegan, M., Hetzler, E.: Visualizing geospatial information uncertainty: What we know and what we need to know. *Cartography and Geographic Information Science* **32**(3), 139–160 (2005)
33. MacQueen, J., et al.: Some methods for classification and analysis of multivariate observations. In: *Berkeley Symp. on Math. Stat. and Prob.*, vol. 1, p. 14 (1967)
34. Mathers, C.D., Fat, D.M., Boerma, J.T.: *The global burden of disease: 2004 update*. Tech. rep., World Health Organization (2004)
35. Pang, A., Wittenbrink, C., Lodha, S.: Approaches to uncertainty visualization. *Vis. Comp.* **13**(8), 370–390 (1997)
36. Pearson, K.: On lines and planes of closest fit to systems of points in space. *Philosophical Magazine* **2**(6), 559–572 (1901)
37. Pfaffelmoser, T., Mihai, M., Westermann, R.: Visualizing the variability of gradients in uncertain 2d scalar fields. *IEEE Transactions on Visualization and Computer Graphics* **19**(11), 1948–1961 (2013)
38. Pfaffelmoser, T., Reitingner, M., Westermann, R.: Visualizing the positional and geometrical variability of isosurfaces in uncertain scalar fields. *Comp. Graph. Forum* **30**(3), 951–960 (2011)
39. Pöthkow, K., Heg, H.C.: Positional uncertainty of isocontours: Condition analysis and probabilistic measures. *IEEE TVCG* **PP**(99), 1–15 (2010)
40. Pöthkow, K., Weber, B., Hege, H.C.: Probabilistic marching cubes. *Comp. Graph. Forum* **30**(3), 931–940 (2011)
41. Potter, K., Kirby, R., Xiu, D., Johnson, C.: Interactive visualization of probability and cumulative density functions. *International Journal of Uncertainty Quantification* **2**(4), 397–412 (2012)
42. Potter, K., Kniss, J., Riesenfeld, R., Johnson, C.: Visualizing summary statistics and uncertainty. In: *Computer Graphics Forum (Proceedings of Eurovis 2010)*, vol. 29, pp. 823–831 (2010)
43. Potter, K., Wilson, A., Bremer, P.T., Williams, D., Doutriaux, C., Pascucci, V., Johnson, C.R.: Ensemble-vis: A framework for the statistical visualization of ensemble data. In: *IEEE Workshop on Knowledge Discovery from Climate Data: Prediction, Extremes.*, pp. 233–240 (2009)
44. Rhodes, P.J., Laramée, R.S., Bergeron, R.D., Sparr, T.M.: Uncertainty visualization methods in isosurface rendering. In: *Eurographics Short Papers*, pp. 83–88 (2003)
45. Roberts, D.E., Hersh, L.T., Scher, A.M.: Influence of cardiac fiber orientation on wavefront voltage, conduction velocity and tissue resistivity in the dog. *Circ. Res* **44**, 701–712 (1979)
46. Roberts, D.E., Scher, A.M.: Effects of tissue anisotropy on extracellular potential fields in canine myocardium in situ. *Circ. Res* **50**, 342–351 (1982)
47. Rodgers, J.L., Nicewander, W.A.: Thirteen ways to look at the correlation coefficient. *The American Statistician* **42**(1), 59–66 (1988)
48. Ross T. Whitaker, M.M., Kirby, R.M.: Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles. In: *IEEE Visualization*, p. to appear (2013)

49. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**(5500), 2323–2326 (2000)
50. Sanyal, J., Zhang, S., Dyer, J., Mercer, A., Amburn, P., Moorhead, R.J.: Noodles: A tool for visualization of numerical weather model ensemble uncertainty. *IEEE TVCG* **16**(6), 1421 – 1430 (2010)
51. Smolyak, S.: Quadrature and interpolation formulas for tensor products of certain classes of functions. *Soviet Mathematics, Doklady* **4**, 240–243 (1963)
52. Sugar, C.A., Gareth, James, M.: Finding the number of clusters in a data set: An information theoretic approach. *Journal of the American Statistical Association* **98**, 750–763 (2003)
53. Toyoshima, H., Ekmekci, A., Flamm, E., Mizuno, Y., Nagaya, T., Nakayama, R., Yamada, K., Prinzmetal, M.: Angina pectoris vii. the nature of st depression in acute myocardial ischaemia. *Am. J. Cardiology* **13**, 498–509 (1964)
54. Xiu, D.: Efficient collocational approach for parametric uncertainty analysis. *Comm. Comp. Phys.* **2**, 293–309 (2007)