

Evaluating Shape Alignment via Ensemble Visualization

Mukund Raj, Mahsa Mirzargar, J. Samuel Preston, Robert M. Kirby, and Ross T. Whitaker ■ *University of Utah*

Visualizing variability in surfaces embedded in 3D provides a means of understanding the underlying distribution of a collection of surfaces. An expert-based evaluation of various ensemble visualization techniques demonstrates the efficacy of using a 3D contour boxplot ensemble visualization technique to analyze shape alignment and variability in atlas construction and analysis.

As computational tools for simulation and data analysis have matured, researchers, scientists, and analysts have become interested in understanding not only the deterministic

output of these tools, but also the uncertainty associated with their computations and data collection. Consequently, there is an increasing interest in uncertainty quantification as an integrated part of simulation and data science in various science and engineering disciplines. Uncertainty quantification views the simulation and data science pipelines as a random process containing possibly both epistemic (reducible) and aleatoric (by chance) uncertainty. Quantification efforts in this random process are divided into roughly two categories:

- efforts to understand the uncertainty and/or variability of the process by examining instances (samples) of the process and
- efforts to determine models (such as probability theory) that capture the nature of the process.

The first of these categories, and the focus of this study, utilizes an ensemble of solutions meant to capture the inherent variability or uncertainty in a computational or data-science pipeline. Although we assume that the variability seen in the ensemble can be attributed to some condition or property of the generating process, we do not assume that articulation of the process via a mathematical

model is straightforward, and hence we have only the ensemble members themselves from which to gain insight into the originating process.

Studying an ensemble in terms of the variability or dispersion between ensemble members can provide useful information and insight about the underlying distribution of possible outcomes. Correspondingly, ensemble visualization can be a powerful way to study this variability. However, a key challenge here is to convey the variability among ensemble members while preserving the main features they share. Preserving these features is particularly challenging in cases where the ensemble members are not fields over which statistical operations such as mean and variance are well-defined, but instead are derived or extracted features such as isosurfaces.

In this article, we examine the effectiveness of the *contour boxplot technique*,¹ a descriptive summary analysis and visualization methodology, in the context of a particular medical-data-science application: brain atlas construction and analysis. We conducted an expert-based evaluation of the visualization of ensembles generated through shape alignment using image deformation in the construction of atlases (or templates) for brain image analysis. To accomplish this evaluation, we constructed a prototype system for visualizing and interacting with ensembles of 3D isosurfaces through a combination of 3D rendering (isocontouring) and cut planes (slices through 3D volumetric fields). In addition, we generalized the algorithm¹ to three dimensions as a direct extension of their analysis of isocontours to isosurfaces—that is, from codimension-one objects embedded in 2D to codimension-one objects em-

bedded in 3D. This generalization lets us compare an ensemble's contour boxplot summaries to both the full enumeration of the ensemble as well as other traditional means of atlas evaluation (such as qualitative visual inspection of slices of the atlas image or individual volumetric images used to construct the atlas). In collaboration with domain experts, we used this system to explore the efficacy of using ensemble visualization techniques for evaluating 3D shape alignment of brain MRI images.

The purpose of this article is to study and evaluate the use of contour boxplots in a real-world data-science application: the alignment of 3D shapes or surfaces in a population-based ensemble. Our hypothesis is that the contour boxplot will let users summarize their data in a meaningful way that allows either better or more efficient (faster) assessment of the atlas construction as compared with explicit enumeration of the ensemble (that is, looking at each image individually) or through more coarse-grained characterizations, such as examining the average intensity image or label (segmentation) probability maps.

Our evaluation results show that the contour boxplot methodology has the potential to significantly benefit the application under study by providing a visualization of the ensemble's quantitative summaries. Although we have formulated our hypothesis in the context of a particular application, we believe that our evaluation may provide insight into other arenas where visualization and analysis of shape ensembles are desired.

Brain Atlas Construction

Constructing an anatomical atlas for a collection of brain images is an important problem in medical image analysis. The goal of various atlas construction schemes is to construct a statistical representative image and an associated set of coordinate transformations (deformations) from an ensemble of images.² Anatomical atlases provide a common coordinate system (atlas space) in which to define reference locations of brain structures.

As part of the atlas construction process, nonlinear registration techniques generate deformations that can map the anatomies in an individual image to the atlas space (see Figure 1). The atlas construction process jointly estimates a representative image defining the atlas space (the atlas image) and the deformations aligning individual images to this atlas image (that is, it maps the image individually to the atlas space). The atlas image generated by these techniques then represents the average (or normal) anatomy of this population. Such atlases help domain experts characterize the




Figure 1. An atlas construction scheme involves deforming and registering all ensemble members to the atlas. This deformation and registration process is called transformation to the atlas coordinate system or the atlas space.

expected anatomical structure and variability of a population and compare different populations in terms of their group atlases (for example, healthy and unhealthy groups). Differences in the atlas anatomy can be identified both qualitatively by inspecting unaligned structures (when mapped to the atlas space) and quantitatively by analyzing the deformations, quantifying the amount of change necessary to bring an individual ensemble member into alignment.

Atlas generation is an automated process, but it is not parameter-free, and the choice of parameters can greatly influence the quality of the result. In particular, nonlinear deformations computed for medical image registration are a trade-off between image matching and plausible deformations. For example, the deformation should not result in the elimination of anatomical features or noninvertible transformations. Hence, the deformation is often controlled by tuning parameters to find a compromise between the mismatch between images and the regularity (smoothness) of the transformation. Because of the regularization of the deformations and the inherent anatomical differences among ensemble members, not all features will be perfectly aligned. This imperfect alignment manifests as blurring in the atlas image where there is disagreement regarding voxel intensity among ensemble members when mapped to the atlas space.

Correct tuning of the regularization parameters lets the deformations account for as much anatomical variability as possible by correctly aligning the corresponding anatomy and not simply matching similar intensities. This alignment of corresponding anatomy is essential for an atlas to be effective in later statistical analyses of the population. Optimization convergence can be easily checked, but the degree to which particular structures align is analyzed qualitatively by observing the amount of blurring in the atlas image and by checking each ensemble member's alignment (deformed to the atlas space) with the atlas image. The initial alignment is often unsatisfactory, which results in an iterative process of parameter tuning and rerunning the atlas-generation process.

In addition, because of problems with image scans, extreme variability among the ensemble members, or incorrect preprocessing, it may not be

Deformations computed for medical image registration are a trade-off between image matching and plausible deformations.

possible to achieve reasonable alignment with the atlas image for some set of outlier images. Identifying and removing such images is often another part of the atlas-generation procedure. Automated measures of global image alignment are available, but they do not give insight into why or in which spatial regions particular ensemble members have poor alignment. Depending on the proposed application of the atlas, these insights may be pertinent to the decision to prune or keep particular images (ensemble members).

This manual iteration of parameter tuning/pruning and atlas generation eventually yields the final atlas to be used in further analysis. We should note two important points about the final atlas image. First, this representative image/segmentation is not a member of the ensemble itself, but rather an image/segmentation generated through statistical operations on the deformation fields. That is, it is not a member of the population that best represents the population, but rather an attempt at statistically characterizing a representative image. Second, the iterative process does not guarantee that the resulting atlas image will be crisp, that there will be no blurry regions in the image. The image ensemble compared with the atlas image

scenario is similar in spirit to the feature-space averaging issue highlighted in earlier work¹—that is, the average field's isosurface (segmentation) is oftentimes not equivalent to a representative of a set chosen from isosurfaces of the individual fields. Thus, the avoidance of feature-space averaging is why we believe the contour boxplot methodology provides a useful way to summarize the type of ensemble data where analyzing feature sets and their representatives is important.¹ Because the manual, qualitative evaluation of shape alignment (as a result of image registration) is challenging, quantifying the variability of the shape alignment and visualizing this variability can facilitate the domain experts' ability to effectively validate the atlas-construction scheme.

Data Preprocessing for Atlases

The images analyzed in this article are 3D MRI images obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database.³ Each brain image in our ensemble was also provided with a corresponding label map volume with various anatomical structures segmented and marked, with each brain region having a unique integer value. To analyze a specific structure within the brain anatomy, we used the label assigned to that structure to select it and mask out the remaining region in all members of the ensemble.

The atlas construction scheme we used is the unbiased diffeomorphic atlas,² which was implemented as part of an open-source medical image atlas-construction package called AtlasWerks (www.sci.utah.edu/software/atlaswerks.html). We constructed atlases from MRI image ensembles using different parameter choices and/or different ensembles (subject groups). In each case, after constructing the atlas using the MRI images, the corresponding label map images were transformed to the common (atlas) coordinate space using deformation fields calculated during the atlas-construction process we described in the last section. These transformed label maps were then passed as input to the preprocessing pipeline (which we describe in the next section) for visualization.

For a well-constructed atlas, we can expect the anatomical structures in the brain to have a relatively small amount of variability after being transformed to the atlas space. We selected two anatomical structures in the brain expected to pose different levels of difficulty during atlas construction: the left ventricle and the cortex. The ventricle is often considered to be a distinct structure (with high contrast) in the brain image and, therefore, can be expected to exhibit good align-

ment among ensemble members in the atlas space (if all goes well). The cortex was selected as an example of an anatomical structure with a complex shape (see Figure 2), a significant challenge for registration and alignment.

Expert Evaluation Study Details

Domain experts use various open-source or commercial packages to visualize slices from individual volumetric images or simply from the average of the aligned images, but to the best of our knowledge, ours is the first attempt to study shape alignment in atlas construction using ensemble visualization techniques. For our evaluation study, we had access to a group of five domain experts who work with atlases regularly and who volunteered to participate in our expert evaluation study. This group included graduate students, staff researchers, and faculty who use atlases and medical image ensembles in their research projects.

We asked the participants to explain their current methodologies for evaluating the atlas construction scheme as well as the quality of the atlases in terms of being a representative of the ensemble. As mentioned earlier, we learned that this process is often performed qualitatively. A visual inspection is carried out to ascertain whether the shapes of the anatomical structures in the atlas space are realistic. Experts also mentioned that, in order for an atlas to be helpful for different medical imaging applications such as the segmentation of a specific brain structure, they need the atlas image and the anatomical structures therein to have sufficient contrast. For example, they expect to see a crisp boundary (in terms of the average combined image intensities) between gray and white matter in the brain. Therefore, the sharpness of the boundaries of the anatomical structures in the atlas image is another criterion examined qualitatively to evaluate the ensemble's alignment. These qualitative evaluations are often performed on a subset of the ensemble (in the atlas coordinate system) because visualizing the entire ensemble results in too much clutter and blurriness.

Figure 3 shows a snapshot of a slice of the brain atlas image used as a common (atlas) coordinate system to register individual label maps from the ensemble.

Visualization Pipeline

To describe our prototype system's visualization pipeline, we start with a brief summary of various ensemble visualization strategies that we considered and incorporated into our prototype system. We then provide an overview of the pipe-



Figure 2. Example anatomical structure: the cortex (green) and the ventricle (red). This image shows the segmentation provided by the label map volume for a typical ensemble member. The coarseness of the segmentation seen in this label map is mitigated by smoothing for the final visualization.




Figure 3. Atlas image slice constructed using AtlasWerks: (a) Atlas image slice and (b) MRI image slice. The anatomical structures in the atlas image usually have lower contrast and fuzzier edges compared with the original MRI images. This fuzziness results from performing averaging while constructing the atlas.

line and our design choices to mitigate the challenge of visualizing and rendering an ensemble of 3D isosurfaces.

Ensemble Visualization Overview

Visualization is often data-driven, and therefore uncertainty visualization schemes are typically designed to deal with the type of data being visualized. For scientific data, users are often interested in visualizing derived features of their data, such as transition regions (or edges), critical points, and isosurfaces (of volumetric data), as well as the uncertainty associated with such feature sets.

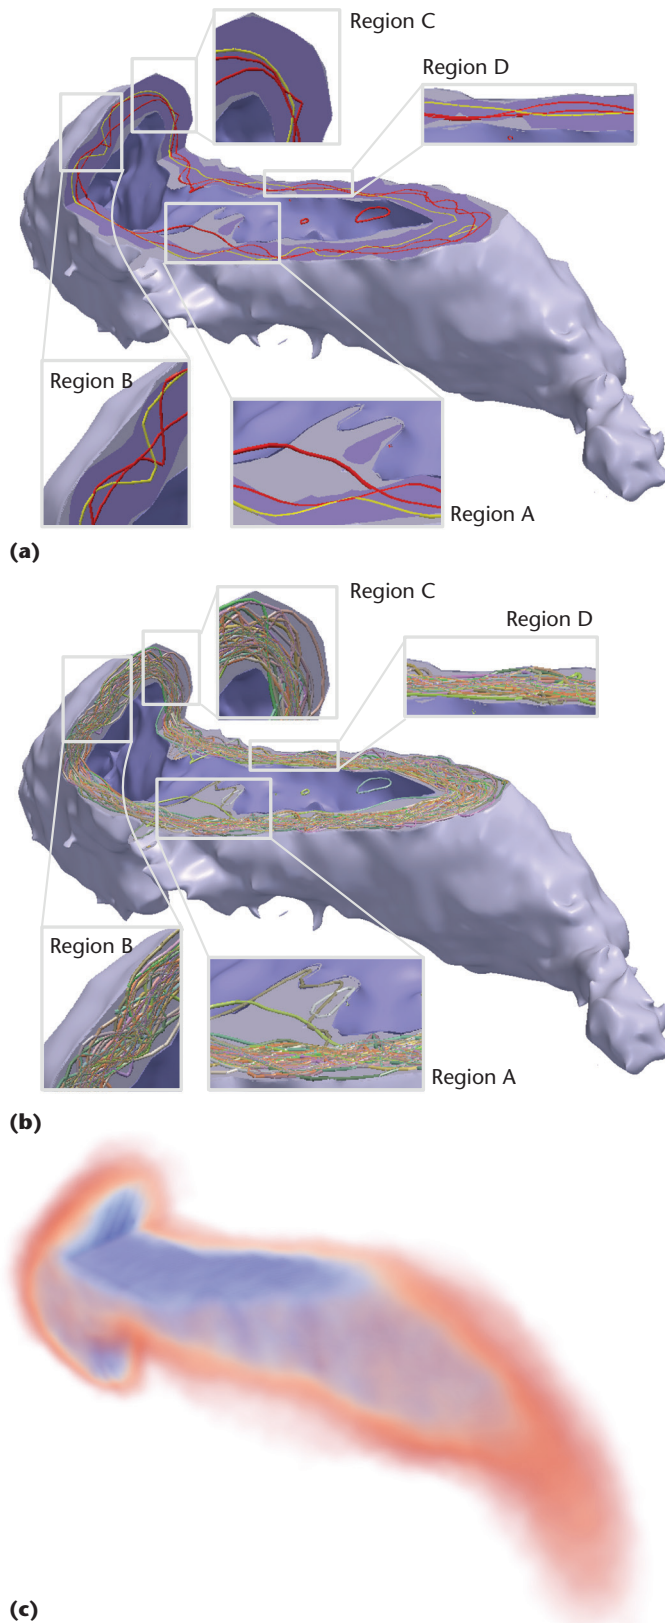


Figure 4. Three visualizations of ventricles from an ensemble containing 34 images from the Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset transformed to a common atlas space: (a) 3D contour boxplot visualization, where dark purple indicates the 50 percent volumetric band, the 100 percent band volume is in light purple, the median is in yellow, and outliers are in red (on the cutting plane); (b) direct visualization of the ensemble members (spaghetti plot); and (c) 3D average intensity image.

The focus here is the visualization of isosurfaces in the context of uncertain scalar fields, which has been studied somewhat extensively. Most relevant to the application under study (atlas construction) is the visualization of uncertain isosurface extracted from an ensemble of scalar fields.

Here we provide a brief summary of three classes of popular techniques for visualization of uncertain isosurfaces that are extracted from ensembles of scalar fields. The following three techniques were chosen to represent the range of strategies for representing an ensemble. To analyze the alignment, or lack thereof, of shapes in an ensemble, we incorporated representative members of these technique categories as part of our prototype system.

Enumeration. A widely used approach for ensemble visualization is the direct visualization of all ensemble members. Direct visualization has gained significant interest in applications such as weather forecasting and hurricane prediction.⁶ Ensemble-Vis is an example of a data analysis tool designed to visualize ensemble data.⁷ It uses multiple views of fields of interest to enhance the visual analysis of ensembles.

We incorporated direct visualization of 3D ensemble members (see Figure 4b) by rendering the curves formed by the region of intersection of each ensemble member's codimension one isosurface with a cut plane. As long as the isosurface embedded in 3D is closed, closed curves will be generated when the isosurface is sliced for visualization purposes. We refer to this visualization as a *spaghetti plot*.

To facilitate the interpretation of the individual ensemble members, we rendered each of these curves with distinct and random colors. There are a variety of options for rendering the enumeration of all 3D surfaces, including transparency, but clutter is a significant challenge.⁶ For this work, we present the surfaces of the innermost and outermost volumetric bands formed by all ensemble members. User studies have suggested the effectiveness of direct ensemble visualization techniques,⁶ but such approaches do not provide

(A thorough review of the rich literature on uncertainty visualization is beyond the scope of this article. Interested readers can consult these works for further details on recent advancements on this topic.^{4,5})

any quantitative information about the data uncertainty and rely solely on the user for interpreting data.

Parametric probabilistic summaries. Many uncertainty visualization schemes use probabilistic modeling to convey quantitative information regarding data uncertainty. These techniques often rely on a certain kind of statistical model such as multivariate normal distributions.

As a representative of such techniques, we chose to consider the concept of level-crossing probabilities (LCP).⁸ For visualization, we implemented the 3D probabilistic marching cubes algorithms (proposed based on LCP)⁹ as part of our initial visualization system. Probabilistic marching cubes rely on approximating and visualizing the probability map of the presence of the isosurface at each voxel location. However, the use of parametric modeling can limit the capability of this technique. Approximating the underlying distribution that gives rise to the ensemble and presenting the user with only aggregated quantities of the inferred distribution can be misleading in some applications. For instance, this approach can often hide or distort structures that are readily apparent in the ensemble.

Nonparametric descriptive summaries. An alternative strategy that relies on neither enumeration nor parametric modeling of the underlying distribution is to form descriptive statistics of an ensemble. Descriptive statistics offer an ensemble visualization paradigm for understanding or interpreting uncertainty from an ensemble's structure. The notion of centrality is a natural approach to understanding an ensemble's structure. Because an ensemble is an empirical description of its distribution, some instances from it are more central to the distribution and therefore are more typical within the distribution.

The notion of data depth provides a formalism for characterizing how central a sample is within an ensemble. Data depth provides a natural generalization of rank statistics to multivariate data.¹⁰ The univariate boxplot (or whisker plot) is a conventional approach to visualize order statistics. Boxplot visualizations provide a visual representation of an ensemble's main features, such as the most representative member (the median), quartile intervals, and potential outliers. The notation of data depth has been generalized for ensembles of isocontours.¹ Researchers have also proposed the contour boxplot technique to summarize robust and descriptive statistics of ensembles of 2D isocontours.¹

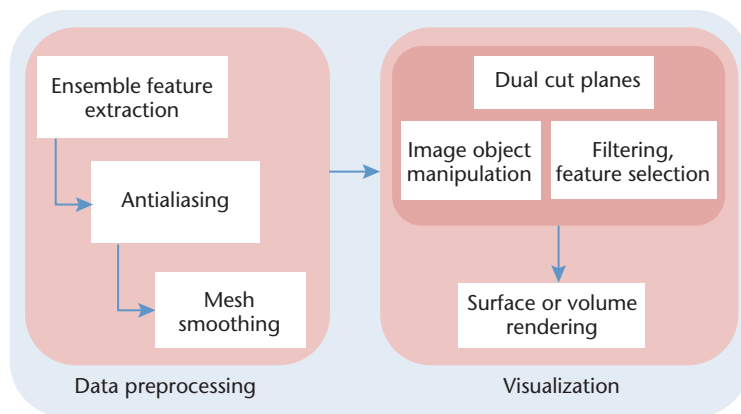


Figure 5. Overview of prototype system designed for shape alignment evaluation using ensemble visualization. The prototype system consists of data preprocessing and visualization stages.

In our system, we algorithmically extend and implement the contour boxplot analysis for iso-surfaces embedded in 3D (see Figure 4a) as an example of visualization techniques based on an ensemble's nonparametric descriptive statistical summaries.

Ensemble Visualization Prototype System

At a high level, our prototype system consists of two stages: data preprocessing and visualization (see Figure 5).

When visualizing isosurfaces of a binary 3D segmented image, it is often necessary to perform smoothing to reduce aliasing artifacts and facilitate 3D rendering and shading. First, we perform this smoothing in a two-step preprocessing stage. In the first step, the binary partitioned image is antialiased using an iterative relaxation process.¹¹ Next, a small amount of mesh smoothing is performed on the isosurface mesh generated from the antialiased binary image. All visualization preprocessing operations occur on the 3D volume (and corresponding codimension one isosurfaces) prior to cut-plane extraction.

The second stage includes some visualization strategies to facilitate the perception and navigation of the rendered 3D objects. To improve shape perception in our application, we include interactivity with renderings of 3D objects as part of the visualization system. In our settings, the user can rotate the object displayed on the screen using a standard trackball interaction mechanism. The system lets the user select cutting planes, which clip a portion of the volume displayed on the screen, to render cross-section views of surfaces embedded in 3D. The user can also interactively orient and translate the cutting plane. Additionally, the system provides the flexibility of having one or multiple cutting planes and interactively adjusting their position and orientation. The system

interface lets the user interactively select various features of interest for rendering in order to focus on particular features of interest. For example, the user can select specific ensemble members to be rendered individually.

In the case of 3D contour boxplots, we performed the analysis on the 3D binary segmented volumetric data (in the preprocessing stage) and rendered the results interactively. Although the analysis was performed on the volumetric data leading to volumetric 50 and 100 percent bands, we rendered the visualization of the statistical summaries only on chosen cut planes to deal with the issue of occlusion. For instance, in the absence of a cut plane, the 100 percent band entirely occludes the median shape and the 50 percent band.

The contour boxplot, as part of the atlas construction process, can help users tease apart the different aspects of variability.

Evaluation

To demonstrate the efficacy of using ensemble visualization techniques to study the alignment of MRI brain images during brain atlas construction, we gathered feedback in an expert evaluation study of the proposed prototype system. We described the prototype system to our expert evaluators (whom we call participants) after a walk-through presentation of the different ensemble visualization techniques. The participants were able to interact with the system and switch through the various visualization. For our study, we solicited their feedback on the visualization of two anatomical structures: the left ventricle and the cortical surface. We paid particular attention to the participants' comments concerning the suitability of ensemble visualization for this application. As the participants interacted with our system, we gained useful insights into the atlas data. We describe three examples here.

Local Variability

In our first example, we focus on analyzing the variability within an ensemble of different regions of brain ventricles transformed to a common atlas space using the unbiased, diffeomorphic approach.² Ensemble visualization not only helps general users identify regions that are either well or poorly aligned, but also provides insight regard-

ing whether the variability is due to differences in shape, position, or both.

Figure 4 shows the three approaches to visualizing the aligned ventricles for an ensemble of 34 brains. In Figure 4a, one can immediately identify regions of high variability. For example, in region A, most of the variability is outside the 50 percent band, which means that less than half the ensemble members contributed to this variability.

Looking at the spaghetti plot in Figure 4b, we see there are in fact only two ensemble members that significantly differ from the other members in region A. These results show that the variability is due to overall position as well as shape in this region. In region B in Figure 4a, we notice that the variability can be attributed to significantly different shapes of the isocontours and that these shapes would not easily be aligned through the smooth transformations in this atlas and may require parameter tuning to achieve alignment. By observing region C in Figures 4a and 4b, we see that the variability comes mostly from the positions of the isocontours. Results in region C also show that no particular ensemble member is disproportionately responsible for the variability—the width of the 50 percent band is nearly that of the 100 percent band in this region, and outliers align well with the median contour.

Finally, region D in Figure 4a demonstrates an area of low variability across the ensemble and provides an example of good alignment of all the ventricles, which is confirmed by the spaghetti plot in Figure 4b. Figure 4c shows a volume-rendered 3D version of the average intensity image for comparison. The average intensity image is an essential part of the atlas, but it does not provide the same insights for debugging the atlas in a detailed way.

We also showed the participants volume renderings of LCP values, as suggested in earlier work.⁹ The participants noted that the LCP visualization shows almost the same information as the average intensity image in Figure 4c, which is already used extensively during atlas construction. They did not feel that further exploration of this form of ensemble uncertainty visualization for evaluating atlases would be useful, and therefore we did not include comprehensive results from LCP renderings in this study.

Overall Variability

The second example was chosen to evaluate whether ensemble visualization can also provide insight into the overall variability among the members of




Figure 6. Slices of average intensity atlases for ensembles of 30 brain images: (a) High value of regularization (transformation smoothing), $\lambda = 1.0$; (b) the same ensemble with low regularization value, $\lambda = 1/9$; and from a different ensemble (subject group) with the regularization/smoothing at $\lambda = 1/9$.

an ensemble of aligned shapes. An understanding of the overall variability, as opposed to local variability, is useful not only for understanding how well a particular atlas was constructed, but also for comparing different atlases.

For this example, we constructed three atlases, each with an ensemble size of 30. The first atlas was constructed with a high value of regularization (transformation smoothing), $\lambda = 1.0$; a second atlas was constructed for the same ensemble while using a low regularization value, $\lambda = 1/9$; and a third atlas was constructed from a different ensemble (subject group) with the regularization/smoothing at $\lambda = 1/9$.

Figures 6 and 7 shows slices of intensity atlases and contour boxplot visualizations for each of the three cases. Figure 6 presents a slice of the intensity image for each atlas, and Figure 7 demonstrates the 3D contour boxplot visualization of the cortical surfaces for atlases corresponding to the intensity images in Figure 6.

Using a high value for the regularization parameter enforces high smoothness of the deformation fields, which in turn makes it harder to arrive at a set of deformations that would perfectly align all the individual images. This lack of alignment leads to high variability between isosurfaces in the ensemble. Such high variability is easily visible

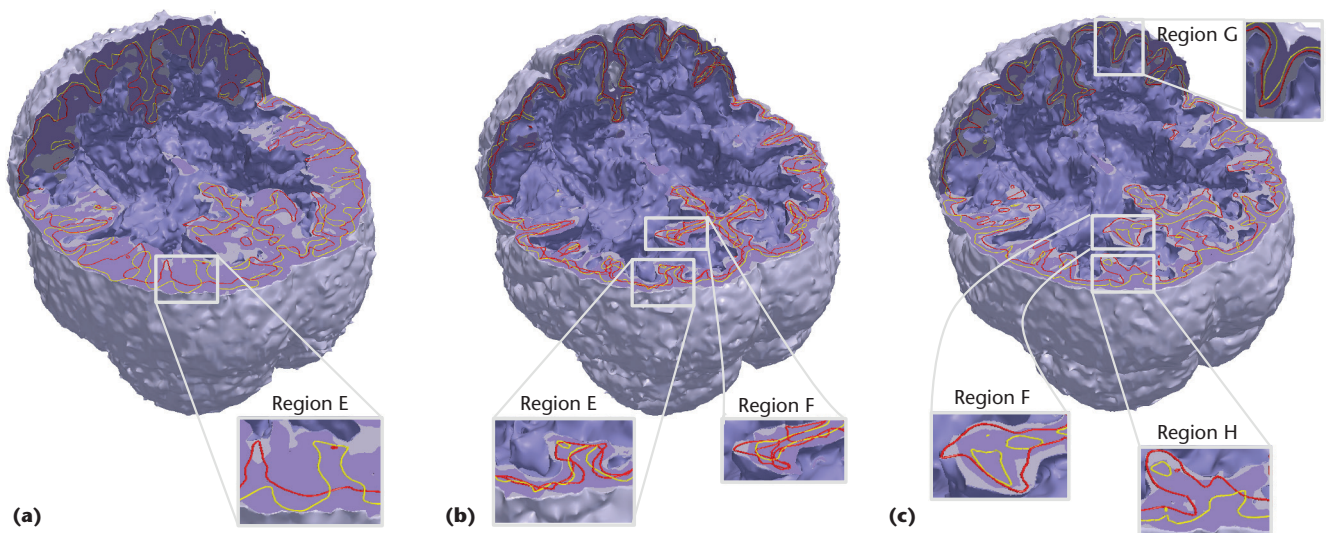


Figure 7. Associated contour boxplot visualizations for cortical surfaces in Figure 6: (a) Atlas constructed with high regularization of deformation, (b) atlas constructed with low regularization, and (c) atlas with low regularization using a different ensemble than in the other columns.

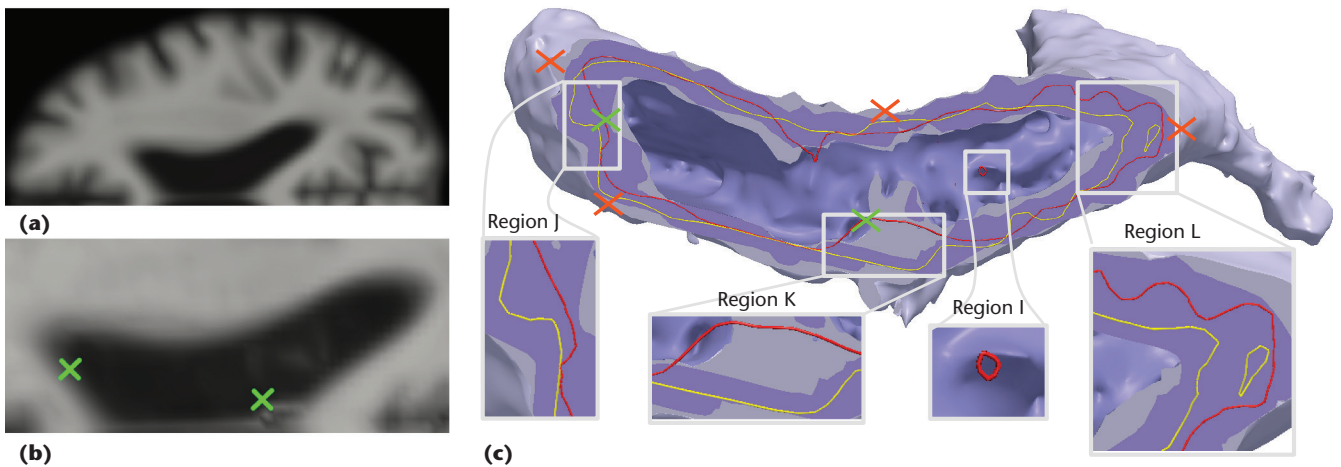


Figure 8. Visualizations of left ventricles. Crosses mark the correspondence between the images. (a) Left ventricle slice from an intensity image of the atlas. (b) Left ventricle slice of an ensemble member identified as an outlier by data depth analysis. (c) Contour boxplot visualization of an ensemble of 34 ventricles in the atlas space.

by looking at region E in Figure 7a, where the 50 and 100 percent bands are wider than in the corresponding region of the atlas with low regularization (see Figure 7b). Better image alignment when the atlas is constructed with low regularization is also evident in region E by comparing contours of the median and outlier shapes rendered on the cut plane in Figures 7a and b. We see that the median and the outlier shapes are poorly aligned for images aligned with an atlas constructed with high regularization (see Figure 7a), whereas the alignment is much better when the atlas is constructed with low regularization.

Finally, the third atlas in this example (see Figures 6c and 7c) demonstrates the effect of inherent variability among the ensemble members (brain images) on the atlas construction process. We see that in many regions of Figure 7c, for instance in region F, the 100 percent band is significantly wider than the 50 percent band, indicating a significant spread in the distribution of surfaces, which differs from the variability seen in the corresponding region in Figure 7b, where both bands nearly overlap. Furthermore, in the third atlas we see that the outlier is well aligned with the median in some regions (see region G), but it is poorly aligned in others (see region H). This example demonstrates that shape and surface variability in atlases depends, in addition to construction parameters, on the inherent variability of shapes in the ensemble. Thus, the contour boxplot, as part of the atlas construction process, can help users tease apart these different aspects of variability.

Member Alignment

In addition to aiding in the understanding of the general alignment of shapes in an ensemble, the contour boxplot is also useful in conveying to the

general user how well a particular shape is aligned with respect to the rest of the ensemble. Such knowledge is particularly useful in the case of outlier shapes. Atlas construction is often an iterative process, and identifying outlier images that do not align sufficiently with the atlas is an important intermediate step in the process.

In the contour boxplot in Figure 8c, we see a single outlier shape and its alignment relative to the ensemble. By comparing this visualization with an average intensity image of the left ventricle region (Figure 8a), we see that an anomaly in region I (Figure 8c) shows as a barely perceivable increase in intensity in Figure 8a. A similar observation can be made from the intensity image slice of the outlier member shown in Figure 8b. However, the anomaly shows up clearly in the contour boxplot, and because it is outside the 100 percent band, we know that the degree of misalignment of this shape is rare within the ensemble of ventricles.

Region I also demonstrates the challenges of assessing geometry in 3D because distances between surfaces can be exaggerated when viewing them on a single cut. However, interacting with the visualization by moving and rotating the cut plane can help verify the 3D shapes of rank statistics and the surface geometries and separation distances.

In some cases, aligned shapes can differ in size from the rest of the ensemble. For instance, Figure 8c shows that the outlier ventricle is noticeably smaller than the median ventricle in regions J and K, which is not the case in region L. This observation is not possible in the corresponding intensity images. These size differences occur for several reasons. In this example, for instance, the outlier ventricle may have been different and irregular to begin with. Another reason could be mislabeling of the ventricular region during the

segmentation process to generate that image's labels. Finally, the process of generating deformations during the atlas construction might fail, leading to irregularities for an ensemble member when mapped onto the atlas space. The contour boxplot can provide information that can help the user decide whether or not any particular outliers need to be removed from the ensemble or if further investigation is necessary to identify causes of possible misalignment.

Technique Comparison

At the conclusion of our study, we asked the participants to comment on their experience with the system, including the applicability of such a system if it were integrated into an atlas construction software. We also asked them to compare the ensemble visualizations to the evaluation techniques they currently use. As we mentioned earlier, the two main techniques currently used for atlas evaluation are inspecting unaligned structures (when mapped to the atlas space) or analyzing the deformations, quantifying the amount of change necessary to bring an individual ensemble member into alignment.

We collected the following observations from the participants in this study:

- Being able to visualize the extent of the variation among the ensemble of aligned shapes in terms of quantitative percentile information using the contour boxplot visualization was helpful for comparing various atlas-construction schemes (or comparing atlases that were constructed from different ensembles or parameter settings). The participants also mentioned that the contour boxplot has the potential to help reduce the time needed to gain insights regarding the quality of the atlas.
- State-of-the-art techniques for evaluation and visualization of atlases provided limited information about the variability that remained within an ensemble after transforming it to the atlas space. Deformation and image match energies (quantities that are optimized during image registration in atlas construction) are not able to provide insight into the geometric discrepancies that are crucial to understanding atlas quality.
- The contour boxplot's ability to effectively locate and characterize different types of variability was valuable in atlas construction.
- An automated and statistically robust way of identifying and visualizing outliers in an ensemble can play a major role in constructing an atlas.

- The spaghetti plot helped the participants view the contours of specific ensemble members other than the median or outliers.
- The contour boxplot and the spaghetti plot were able to convey important details pertaining to the variability in an ensemble, whereas the average intensities had limited utility because of their general fuzziness.

The goal of the application we describe here is to evaluate the alignment of 3D shapes, in particular the alignment of 3D MRI images that have been transformed to a common atlas space using various ensemble visualization methods. We found that the ensemble visualization methods are helpful in characterizing the shape alignment

The ensemble visualization methods are helpful in characterizing the shape alignment and provide insights that are useful in understanding the variability in alignment.

and, furthermore, provide insights that are useful in understanding the variability in alignment. An understanding of the type or location of the variability can be helpful in tuning parameters used in atlas construction and/or removal of outliers to achieve better alignment.

We also observed that the contour boxplot emerged as a clear favorite of our participants. One of the contour boxplot's salient features that makes it distinct from the other ensemble visualization approaches is its ability to convey an aggregated result from the analysis of all shape regions in the ensemble on any arbitrary cut plane. For example, visualizing a slice of the intensity image, or contours on a cut plane using the spaghetti plot, conveys the variability for only the region intersecting the cut plane, whereas a contour boxplot visualization using the same cut plane also provides information about the median and outlier contours that are calculated from a global analysis of contours.

The contour boxplot, however, has a drawback: it does not give the user much information about specific ensemble members, other than the median or the outliers. For such cases, the spaghetti plot with interactivity that allows highlighting of specific ensemble members can augment the





Figure 9. Contour boxplot visualizations. (a) Visualization for an ensemble of 100 simulated HIV protein. Here, we see the median contour in yellow and the outlier contours in red. (b) Visualization of the isosurface of a pressure field in a fluid flow. The pressure is considered a function of depth to generate a 3D pressure volume. The median contour is in yellow, and the outlier contours are in red.

contour boxplot by providing more detail if the general user wishes to focus on specific anatomical areas or members of the ensemble.

Future work for our system in the context of the current application includes refining the system in order to address the study participants' suggestions, such as viewing the specific structures in the context of the whole brain and more interaction options. Furthermore, the ensemble visualization approaches we discussed here can be integrated into an atlas construction package to provide users with the ability to interactively inspect the shape alignments and the variability among ensemble members after atlas construction. Motivated by the feedback from the participants, a more comprehensive study is required to examine the applicability of ensemble visualization to compare different atlas-construction schemes.

Studying shape variability has applications in various branches of science. In molecular dynamics, researchers study different types of molecular structures and the shapes of their potential fields in solutions (which vary stochastically) to understand, for instance, their biochemical properties.¹² Scientists are also interested in the evolution of the shape of molecules. For example, the surfaces of 3D molecular chains are of significant interest for comparison of various types of protein structures.¹² Figure 9a shows the contour boxplot visualization of the surface of an ensemble of simulated HIV molecules. The ensemble members underwent a Procrustes alignment (translation, rotation, scale) using the positions of the underlying molecules. The potential fields that form these contours are inherently smooth, and thus there was no need to preprocess this volume data.

Another application where the study of shape variability and alignment is of significant interest is fluid mechanics. In fluid mechanics, when developing models of vortex behavior, scientists often study the variability in the shapes of vortex structures among different simulations (for example, using slightly different parameter settings or boundary conditions) to confirm that their observations are repeatable,¹³ rather than a numerical artifact of a particular simulation. The center of an eddy corresponds to low pressure values in the flow, and hence studying the pressure field of a fluid flow can help detect the position of the eddies and regions of high vortices. For this case, we used the 2D incompressible Navier-Stokes solver as part of the open source package Nektar++ (www.nektar.info) to generate an ensemble of 28 fluid flow simulation runs. These simulations were designed for a steady fluid flowing past a cylindrical obstacle. For each of the ensemble members, we randomly perturbed the initial conditions such as inlet velocity and Reynolds number. For this example, the pressure dependence in the third dimension was computed analytically. Figure 9b shows the contour boxplot visualization of the isosurfaces of the pressure volume.

In addition to those we have showcased here, there are many possible applications that could benefit from the contour boxplot summary and visualization technique. 

Acknowledgments

We thank Avantika Vardhan, Clement Vachet, Ben Galvin, and Sarang Joshi for discussions on atlas construction and help in processing the brain data we use here. We also thank Lee Makowski and Jaydeep Bardhan for providing the HIV molecule data and Shireen Elhajian for help in processing the HIV molecule data.

We received valuable feedback from Theresa-Marie Rhyne on improving the visualizations. This work was supported by National Science Foundation (NSF) grant IIS-1212806 and by the National Institute of General Medical Sciences of the National Institutes of Health under grant P41 GM103545-17.

References

1. R.T. Whitaker, M. Mirzargar, and R.M. Kirby, "Contour Boxplots: A Method for Characterizing Uncertainty in Feature Sets from Simulation Ensembles," *IEEE Trans. Visualization and Computer Graphics*, vol. 19, no. 12, 2013, pp. 2713–2722.
2. S. Joshi et al., "Unbiased Diffeomorphic Atlas Construction for Computational Anatomy," *Neuroimage*, vol. 23, no. 1, 2004, pp. 151–160.
3. C.R. Jack et al., "The Alzheimer's Disease Neuroimaging Initiative (ADNI): MRI Methods," *J. Magnetic Resonance Imaging*, vol. 27, no. 4, 2008, pp. 685–691.
4. K. Brodlić, R.A. Osorio, and A. Lopes, "A Review of Uncertainty in Data Visualization," *Expanding the Frontiers of Visual Analytics and Visualization*, Springer Verlag, 2012, pp. 81–109.
5. K. Potter, P. Rosen, and C.R. Johnson, "From Quantification to Visualization: A Taxonomy of Uncertainty Visualization Approaches," *Uncertainty Quantification in Scientific Computing*, Springer, 2012, pp. 226–249.
6. J. Cox, D. House, and M. Lindell, "Visualizing Uncertainty in Predicted Hurricane Tracks," *Int'l J. Uncertainty Quantification*, vol. 3, no. 2, 2013, pp. 143–156.
7. K. Potter et al., "Ensemble-Vis: A Framework for the Statistical Visualization of Ensemble Data," *Proc. IEEE Int'l Conf. Data Mining Workshops (ICDMW)*, 2009, pp. 233–240.
8. T. Pfaffelmoser, M. Reitingner, and R. Westermann, "Visualizing the Positional and Geometrical Variability of Isosurfaces in Uncertain Scalar Fields," *Computer Graphics Forum*, vol. 30, no. 3, 2011, pp. 951–960.
9. K. Pöthkow, B. Weber, and H.-C. Hege, "Probabilistic Marching Cubes," *Computer Graphics Forum*, vol. 30, no. 3, 2011, pp. 931–940.
10. S. López-Pintado and J. Romo, "On the Concept of Depth for Functional Data," *J. Am. Statistical Assoc.*, vol. 104, no. 486, 2009, pp. 718–734.
11. R.T. Whitaker, "Reducing Aliasing Artifacts in Iso-surfaces of Binary Volumes," *Proc. IEEE Symp. Volume Visualization (VVS)*, 2000, pp. 23–32.
12. X. Zhang et al., "Application of New Multiresolution Methods for the Comparison of Biomolecular Electrostatic Properties in the Absence of Global Structural Similarity," *Multiscale Modeling and Simulation*, vol. 5, no. 4, 2006, pp. 1196–1213.

13. C.H.K. Williamson, "Vortex Dynamics in the Cylinder Wake," *Ann. Rev. Fluid Mechanics*, vol. 28, no. 1, 1996, pp. 477–539.

Mukund Raj is a PhD student in computing and is affiliated with the School of Computing and the Scientific Computing and Imaging institute at the University of Utah. His research interests include scientific visualization, graph analytics, and statistical analysis. Raj has an MS in computing from the University of Utah. He is a member of SIAM. Contact him at mrj@cs.utah.edu.

Mahsa Mirzargar is a postdoctoral research associate at the Scientific Computing and Imaging Institute at the University of Utah. Her research interests include scientific visualization, statistical analysis, uncertainty quantification, and multidimensional signal processing. Mirzargar has a PhD in computer engineering from the University of Florida. Contact her at mirzargar@sci.utah.edu.

J. Samuel Preston is a PhD student in computing in the School of Computing and the Scientific Computing and Imaging institute and a software engineer at the Scientific Computing and Imaging Institute at the University of Utah. His research interests include scientific visualization and image analysis. Preston has an MS in computing from the University of Utah. He is a member of SIAM. Contact him at jsam@sci.utah.edu.

Robert M. Kirby is a professor of computer science in the School of Computing and the Scientific Computing and Imaging Institute at the University of Utah. His research interests include high-order methods, scientific visualization, concurrent programming, and high-performance computing. Kirby has a PhD in applied mathematics from Brown University. Contact him at kirby@cs.utah.edu.

Ross T. Whitaker is the director in the School of Computing and a professor in the School of Computing and the Scientific Computing and Imaging Institute at the University of Utah. His research interests include image analysis, geometry processing, and scientific computing. Whitaker has a PhD in computer science from the University of North Carolina at Chapel Hill. He is an IEEE fellow. Contact him at whitaker@cs.utah.edu.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.