Visualization Viewpoints

Editor:
Theresa-Marie Rhyne

Visualization of Uncertainty without a Mean

Kristin Potter, Samuel Gerber, and Erik W. Anderson
University of Utah

Uncertainty is a tricky business. The term encompasses all sorts of unknowns, including error, deviations, missing information, or confidence levels. Numerous methods exist for quantifying and expressing uncertainty, and its existence is persistent and accumulative throughout the visualization pipeline.[1] For visualization researchers, this term is particularly complex owing to the visual display medium's limitations.

Consequently, aggregation is a common technique for summarizing uncertainty for visualization purposes. Two particular summary statistics stand out as the de facto characterization: mean and standard deviation. These statistics reduce uncertainty to an expected value and variation from that value and are particularly effective in expressing normally distributed data. However, they aren't always appropriate or even feasible, particularly when we can't describe the uncertainty in a dataset as a probability distribution function.

Consider a problem from medical-image processing: segmenting a brain volume into specific tissue types using fuzzy classification. Each voxel has 11 probability distributions, one for each possible tissue type. Because these probabilities are assigned over *categorical data* (in which a variable value is one of a limited number of nominal categories), traditional uncertainty measures aren't well defined. In this setting, we can use entropy, in the information-theoretic sense, to quantify the spread of probabilities over the categories and encode uncertainty for visualization. For example, in Figure 1, a grayscale color map encodes each voxel's entropy. The white areas indicate regions where the assignment of a voxel to a particular tissue type is inherently uncertain.

Uncertainty Visualization

Since the first calls for research in uncertainty visualization,[2,3] much work has been done in response.[4,5] This goes hand in hand with computing-technology advances leading to larger, more complex datasets along with greater availability of uncertainty information for use in visualizations. In most cases, visualization researchers are ultimately presented with a probability distribution of possible instances of the data. The challenge here is to augment a visual representation of a single instance to incorporate the uncertainty represented in the probability distribution. This additional information poses challenges related to visual encoding, including visual clutter, cognitive overload, and data obfuscation, which can defeat uncertainty visualization's main goals. In most cases, a comprehensive visualization of uncertainty isn't feasible; a summary of the probability distribution is required.

The typical framework for summarizing is to pick the "most likely" representation of the data, such as the mean, and add information about the variability through color, spatial distortion, and so on. So, uncertainty visualization's major challenge is simply the perceptually efficient encoding of the dataset's appropriate summary statistics. This article focuses on such statistics.

Although summary statistics such as the mean, the standard deviation, and quantiles are typical expressions of uncertainty,[6] their use is limited to data classes in which they're well defined. One data class for which they don't work is categorical data. For example, in medical images, a pixel (or voxel) might be one of several tissue types. In remote sensing, satellite image data might depict a particular ground cover, such as vegetation, urban space, or water. (Although such data might actually have an ordering provided by the imaging modality, the classification into nominal categories often relies on both the scalar value and spatial location. So, the existent ordering is insufficient.)

Categorical data is inherently discrete; we can encode uncertainty through a discrete distribution

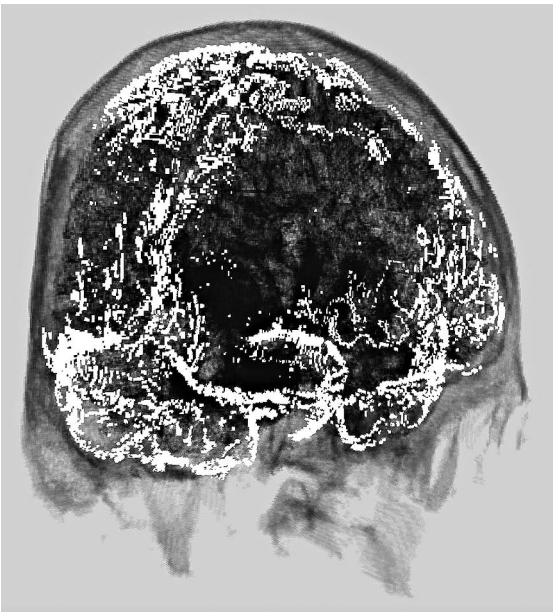

Figure 1. This volume rendering of magnetic resonance imaging (MRI) brain data uses entropy to show areas in which the type of brain tissue is uncertain. The high-entropy regions, in white, highlight tissue boundaries where MRI couldn't distinguish between the defined tissue types.

over the possible classification by defining a probability for each nominal category. These categories have no meaningful ordering; that is, we can't say a particular category is closer to another. So, we can't define a metric across the space. Moreover, we can't combine categories such that a variable exists partially as multiple categories. Thus, we can't meaningfully compute measures such as the mean or standard deviation because this would require averaging across different categorical types.

As an example, we use data from the BrainWeb project (www.bic.mni.mcgill.ca/brainweb), which seeks to provide a ground truth for medical-image-processing techniques through a database of simulated magnetic resonance imaging (MRI) brain scans. BrainWeb provides a probabilistic anatomical model used for simulation. The anatomical model is represented through 11 volumetric scalar fields, each corresponding to one tissue type identified in the MRI scan. A single volume represents the probability of the corresponding tissue at each voxel. So, across all volumes, the 11 probabilities at each voxel add to 1. Formally, we have a set of random variables $X_{(i,j)}$ at each voxel location (i, j), with a discrete distribution $p_{(i,j)}$ over the 11 tissue types.

An important question in this context is, what's each voxel's tissue type? To get a complete answer, we must look at the probability distribution at each voxel and encode it into a visual representation. Our approach follows the framework we mentioned earlier in this section. We encode the most likely representation as the maximum-probability tissue type—that is, the maximum statistic at each voxel, $\arg\max_x \in X_{(i,j)} p_{(i,j)}(x)$. We quantify the uncertainty through the distribution's entropy, $H[X_{(i,j)}]$.

Entropy as a Measure of Uncertainty

Uncertainty visualization approaches have focused almost exclusively on applications that derive the uncertainty's magnitude from a continuous distribution. Uncertainty stemming from a discrete distribution over (unordered) categorical variables presents unique challenges. In the continuous case, we can characterize the uncertainty by a measure of the distribution's *spread*, most commonly described by the distribution's variance or quantiles. For categorical data, such a measure is possible only if we can meaningfully order the variables. For categorical data with no ordering, we believe that entropy is a more appropriate measure.

Let X be a random variable with the probability density p: $\Omega \rightarrow \mathbb{R}^+$ on the discrete sample space $\Omega = \{x_1, \ldots, x_n\}$. The mean $E[X] = \sum_{i=1}^{n} x_i p(x_i)$ and variance $\mathrm{var}[X] = \sum_{i=1}^{n} \left(E[X] - x_i\right)^2$ are meaningful only if the addition of the random variables x_i is sensible. Consider the example of different tissue types in a brain MRI scan. Adding, for example, white matter and cerebral spinal fluid (CSF) clearly isn't sensible; tissue can't be both white matter and CSF. So, we reasonably can't define the sample mean and variance. The Shannon entropy $H[X] = -\sum_{i=1}^{n} p(x_i) \log\left(p(x_i)\right)$ relies on only the probabilities of tissue types rather than the value of x_i. So, it's sensible for categorical variables such as these tissue types.

Informally, entropy, like variance, measures how spread out a distribution is. A discrete distribution's entropy ranges from 0 to $-\log(1/n)$; 0 corresponds to a distribution with a probability of 1 for a single outcome, and $-\log(1/n)$ is a discrete uniform distribution. In the tissue type example, a voxel with maximal entropy would indicate that the brain MRI scan is inconclusive regarding tissue type. Conversely, voxels with minimal entropy would indicate, with a probability of 1, that the tissue is of a specific type.

However, unlike variance, entropy provides no information about the random variable's value. Entropy only measures randomness, whereas the combination of mean and variance indicates a range of most likely values. The formal definition of entropy reflects this by observing that only the probability density influences the entropy, not the random variable's actual values.

Another perspective comes from information theory results for constructing optimal (minimal length) codes. We can translate those results into finding the minimal expected number of yes or no (binary) questions,[7] #*Q*, required to determine the value of an observation of *X*. This relates directly to the random variable's entropy (in $\log_2$): $H[X] \leq E[\#Q] < H[X] + 1$. This entropy intuitively describes the uncertainty by encoding how close to deterministic a random variable acts.

Generally, the logarithm's base only scales the entropy value and doesn't affect the visualization's qualitative aspects. In the following, we use a base two logarithm, which lets us interpret the entropy in bit units—that is, the expected number of binary questions required to determine the category.

As a concrete example, consider the following discrete distributions over four categories and the resulting entropies:

$p_1 = \{1, 0, 0, 0\} \Rightarrow H[X] = 0.00,$
$p_2 = \{0.85, 0.15, 0, 0\} \Rightarrow H[X] = 0.61,$
$p_3 = \{0.85, 0.05, 0.05, 0.05\} \Rightarrow \mathrm{H}[\mathrm{X}] = 0.85,$
$p_4 = \{0.25, 0.25, 0.25, 0.25\} \Rightarrow H[X] = 2.00.$

The first example, p_1, is completely deterministic. Determining the outcome requires zero yes or no questions, corresponding to zero entropy. Examples p_2 and p_3 demonstrate increased entropy because the probabilities are distributed over more categories. Finally, p_4 is the maximal-entropy case for four categories and results in the expected two binary questions required to determine an observation from p_4.

Visualization Methodology

Using entropy to encode each voxel's uncertainty lets us visualize the uncertainty across the entire volumetric dataset. This allows for a more contextual understanding of the data, including the individual tissues' shape as well as their juxtaposition in the volume.

Figure 2 shows a slice through the volume of entropy values in which the distribution's entropy at each voxel is color-mapped from low (blue) to high (white). White indicates where multiple tissue types are possible; dark blue indicates a single tissue type. This is most readily apparent in the head's outer edges, where a voxel might be classified as skull, skin, muscle, CSF, fat, or connective tissue. Although this visualization appropriately expresses the regions of higher entropy (and thus uncertainty) around tissue boundaries, it's missing the structural information from the original data.

To reincorporate structural and contextual information, we look at a particular tissue's most likely location, which we represent as the maximum-probability tissue type at each voxel. A natural visual encoding of this information is to label each tissue type with a color and assign to each voxel the color of the tissue with the maximum probability at that location. Such an encoding is effective only when the colors are easily discernable. So, we must choose colors carefully, and the number of distinguishable color labels is generally limited to a dozen or so. Figure 3 shows a slice through this tagged dataset. Each of the 11 tissues is assigned a color, and each voxel in the slice represents a single tissue.

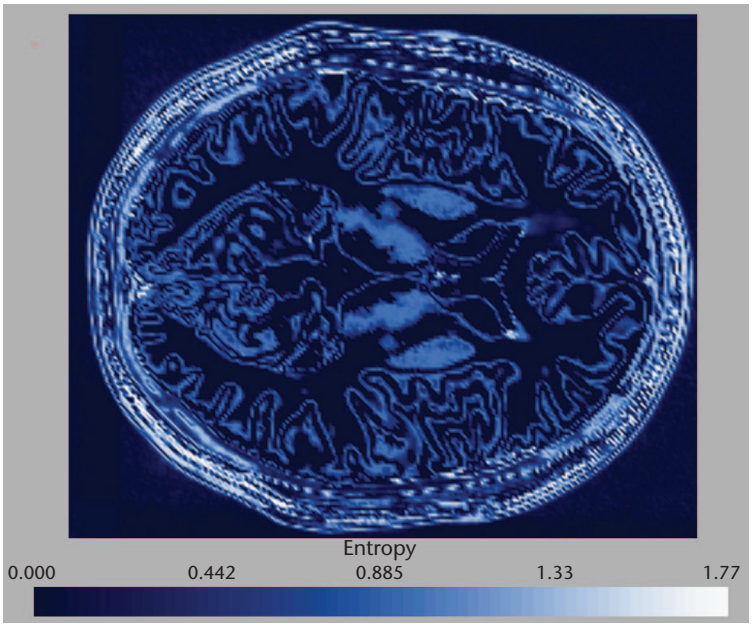

Figure 2. A slice through the entropy volume. Low to high entropy is color-mapped blue to white. Low entropy indicates where a voxel's category is known, such as the volume's four corners, where "background" is the only possible category. High entropy indicates where multiple categories can exist, such as around tissue boundaries.

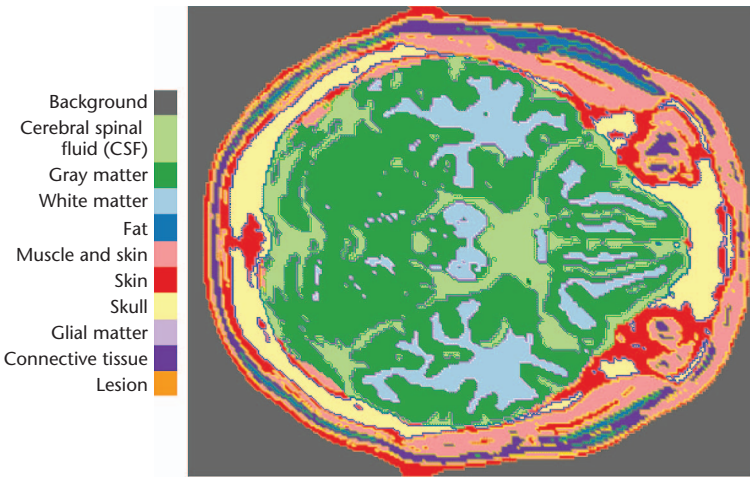

Figure 3. A slice through the tagged data, with each voxel having the color of the corresponding maximum-probability tissue.

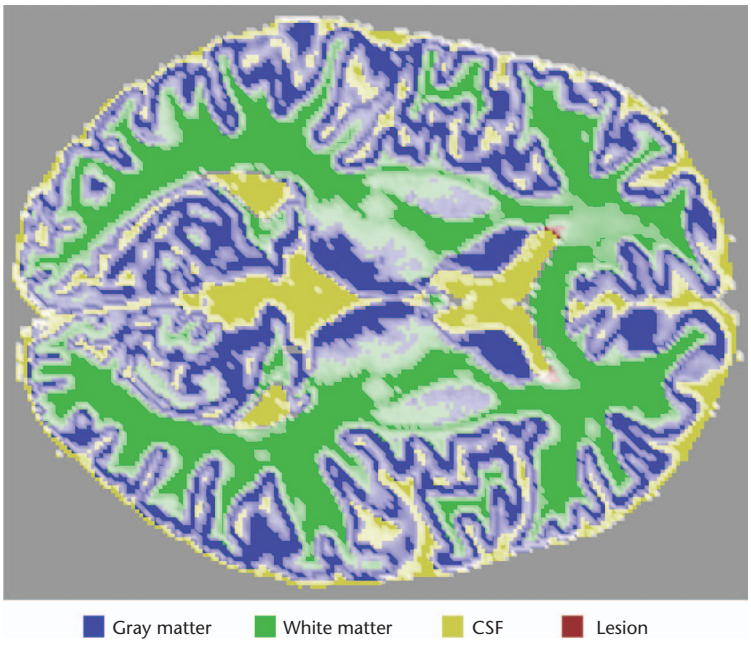

Figure 4. A slice through the combined entropy volume and tagged voxel information. We color-mapped each voxel on the basis of its maximum-probability tissue and the amount of entropy. As a voxel tends toward white, the higher the entropy and the less certainty of a particular tissue type.

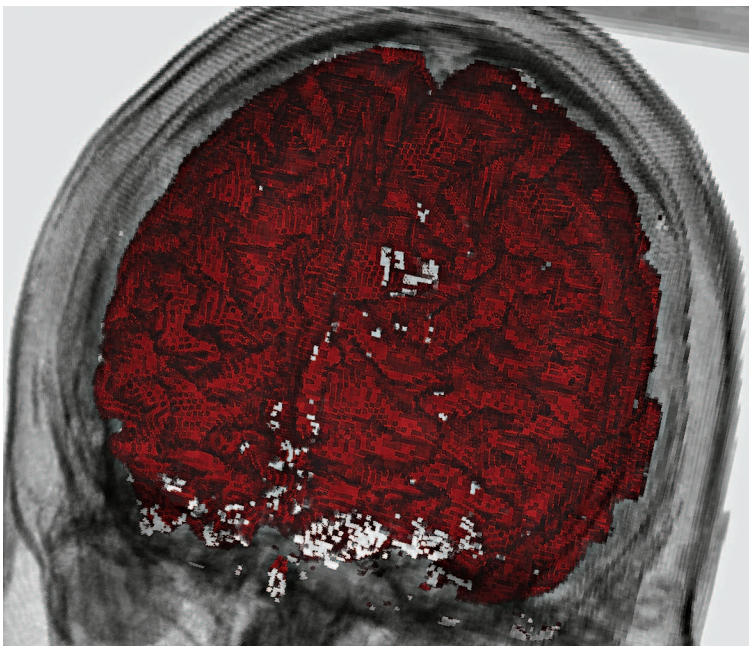

Figure 5. A volume rendering of the combined entropy volume and tagged voxel information in which red indicates white matter and white indicates high entropy. For more context, the skull and skin appear in shadow.

Combining the entropy volume and tagged voxel information provides a holistic view of the data. Figure 4 demonstrates this visualization, restricting the number of tissues we're interested in to four: gray matter, white matter, CSF, and lesion, and using two-step color mapping. For each voxel of the slice, we find the maximum-probability tissue type and take its color tag. We then blend this color through white, on the basis of that location's entropy. This encoding of entropy and maximum-probability tissue type is perceptually meaningful.

A maximal-entropy voxel, with each tissue type equally likely, will appear white and not indicate a particular tissue type. For a minimal-entropy voxel, the color will clearly identify its tissue type. Between these two extremes, the entropy depends on both the magnitude of the probability mass on the likely tissue type and the distribution of the remaining probability mass over the other tissues, as illustrated in the numerical example we gave earlier. In Figure 4, this approach highlights the uncertainty of the lesion's size and position.

Whereas slice-based rendering can be advantageous in understanding tissue interactions restricted to small regions of the data, volume rendering can give insights to contextual information in the 3D spatial domain. Figure 5 shows a volume rendering using a coloring technique similar to that of the slice-based method. We've selected a single tissue type—white matter. All maximum-probability voxels of that tissue type are red if the voxel's entropy is below a threshold or white if the entropy is above the threshold. In this visualization, viewers can see the white matter's relationship to the rest of the brain volume. They can also gain an understanding of the location, shape, and magnitude of the uncertainty associated with assigning the voxel a certain tissue type.

Such approaches to employing entropy might not be suitable in other applications or fulfill other specific visualization goals. However, we foresee that future visualizations will leverage entropy to aggregate complex information and expose the uncertainty's location and magnitude.

This use of entropy isn't novel; geophysics research has employed it extensively. (For a look at this and other uses of entropy in visualization, see the sidebar.) However, we believe it's important to exemplify entropy as an uncertainty measure for visualization. With the increased popularity of understanding uncertainty through visualization, visualization experts need a variety of uncertainty measures, and those measures should be commonplace in the visualization vocabulary. Moving beyond mean and standard deviation allows for more expressive control and greater understanding of the nuances of uncertainty quantification and visualization. For this reason, we expect entropy

Entropy in Visualization

Information theory and, more specifically, entropy have increasingly been leveraged as a way to make quantitative statements about various properties broadly throughout the visualization framework.[1,2] For example, researchers have applied entropy to measure the amount of information in the volume rendering of 3D datasets to determine the most informative camera viewpoint or a minimal set of representative views of a 3D scene.[3] Likewise, for large-scale time-varying data, researchers have used entropy to pick important time steps or subregions with maximal information to enable visualization with limited resources or time.[4] Entropy can also serve to measure quality for evaluation and comparison of level-of-detail algorithms for multiresolution volume rendering.[5] Finally, in flow field and vector visualization, researchers have employed entropy both to quantify the information present to evaluate a visualization's effectiveness[6] and to generate noise reflecting the amount of information in texture-based methods.[7]

The most popular applications of entropy arise in geoscience. For example, uncertainties might arise when classifying remotely sensed data into vegetation types or modeling geological structures and data assimilation. In such cases, entropy encompasses uncertainties in these systems and can serve as an axis for parallel coordinates[8] or a color-mapping method.[9,10]

References

1. M. Chen and H. Jänicke, "An Information-Theoretic Framework for Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 16, no. 6, 2010, pp. 1206–1215.
2. C. Wang and H.-W. Shen, "Information Theory in Scientific Visualization," *Entropy*, vol. 13, no. 1, 2011, pp. 254–273.
3. U.D. Bordoloi and H.-W. Shen, "View Selection for Volume Rendering," *Proc. IEEE Visualization 2005* (Vis 05), IEEE, 2005, pp. 487–494.
4. G. Wu, Y. Cao, and J. Yin, "Entropy Based Information Visualization of Scientific Data," *Int'l J. Virtual Reality*, vol. 10, no. 3, 2010, pp. 65–72.
5. C. Wang and H.-W. Shen, "LOD Map—a Visual Interface for Navigating Multiresolution Volume Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, 2006, pp. 1029–1036.
6. L. Xu, T.-Y. Lee, and H.-W. Shen, "An Information-Theoretic Framework for Flow Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 16, no. 6, 2010, pp. 1216–1224.
7. H. Wang et al., "Fuzzy Feature Visualization of 3D Vector Field by Information-Entropy-Based Texture Adaptation," *Int'l J. Virtual Reality*, vol. 10, no. 4, 2011, pp. 37–43.
8. Y. Ge et al., "Exploring Uncertainty in Remotely Sensed Data with Parallel Coordinate Plots," *Int'l J. Applied Earth Observation and Geoinformation*, vol. 11, no. 6, 2009, pp. 413–422.
9. M. Goodchild, L. Chih-Chang, and Y. Leung, "Visualizing Fuzzy Maps," *Visualization in Geographical Information Systems*, H.M. Hearnshaw and D. Unwin, eds., John Wiley & Sons, 1994, pp. 158–167.
10. J.F. Wellman and K. Regenauer-Lieb, "Uncertainties Have a Meaning: Information Entropy as a Quality Measure for 3-D Geological Models," *Tectonophysics*, vols. 526–529, 2011, pp. 207–216.

to become one of a collection of mainstream measures of uncertainty for visualization.

References

1. A. Pang, C. Wittenbrink, and S. Lodha, "Approaches to Uncertainty Visualization," *The Visual Computer*, vol. 13, no. 8, 1997, pp. 370–390.
2. C.R. Johnson, "Top Scientific Visualization Research Problems," *IEEE Computer Graphics and Applications*, vol. 24, no. 4, 2004, pp. 13–17.
3. C.R. Johnson and A.R. Sanderson, "A Next Step: Visualizing Errors and Uncertainty," *IEEE Computer Graphics and Applications*, vol. 23, no. 5, 2003, pp. 6–10.
4. K. Brodlie, R.A. Osorio, and A. Lopes, "A Review of Uncertainty in Data Visualization," *Expanding the Frontiers of Visual Analytics and Visualization*, part 2, J. Dill et al., eds., Springer, 2011, 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, 2011, pp. 226–249.
6. B.N. Taylor and C.E. Kuyatt, "Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results," tech. note 1297, US Nat'l Inst. of Standards and Technology, 1994.
7. T.M. Cover and J.A. Thomas, *Elements of Information Theory*, John Wiley & Sons, 1991.

Kristin Potter *is a research scientist at the University of Utah's Scientific Computing and Imaging Institute. Contact her at kpotter@sci.utah.edu.*

Samuel Gerber *is a visiting assistant professor in Duke University's Mathematics Department. Contact him at sgerber@math.duke.edu.*

Erik W. Anderson *is a postdoctoral research associate at the University of Utah's Scientific Computing and Imaging Institute. Contact him at eranders@sci.utah.edu.*

Contact department editor Theresa-Marie Rhyne at theresamarierhyne@gmail.com.