

Isosurface Similarity Maps

Stefan Bruckner and Torsten Möller

Graphics, Usability and Visualization (GrUVi) Lab
School of Computing Science, Simon Fraser University, Canada

Abstract

In this paper, we introduce the concept of isosurface similarity maps for the visualization of volume data. Isosurface similarity maps present structural information of a volume data set by depicting similarities between individual isosurfaces quantified by a robust information-theoretic measure. Unlike conventional histograms, they are not based on the frequency of isovalues and/or derivatives and therefore provide complementary information. We demonstrate that this new representation can be used to guide transfer function design and visualization parameter specification. Furthermore, we use isosurface similarity to develop an automatic parameter-free method for identifying representative isovalues. Using real-world data sets, we show that isosurface similarity maps can be a useful addition to conventional classification techniques.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms

1. Introduction

The field of volume visualization aims to provide insightful depictions of three-dimensional data. Modalities such as CT, MRI, or laser-scanning confocal microscopy allow physicians, scientists, and engineers to investigate the interior of complex objects. However, providing clear visualizations of the structures contained in a volume data set is a major challenge. One of the issues is the lack of explicit geometric information and limited semantics. A volume data set contains a large number of isosurfaces at different target scalar field values, while its structure is typically characterized by a finite number of feature isosurfaces that segment the data set into several important components. The data may be visualized directly by mapping the scalar values and/or derived attributes to optical properties, or a geometric surface representation may be extracted using techniques like the popular Marching Cubes algorithm [LC87]. Irrespective of the chosen visualization method, providing guidance in the identification of salient isovalues plays an important role in improving the exploration process. In this paper, we present a new approach for the visualization and detection of relevant isovalues which provides additional information compared to conventional histograms.

Histogram-based methods typically infer similarity from

the frequency at which isovalues occur. While there are cases where this assumption holds, in general only limited information can be deduced using this approach. Moreover, the presence of large homogenous regions, acquisition artifacts, and noise introduces additional problems. Instead of the frequency of individual data values and/or derived attributes, our method is based on the global notion of isosurface similarity. Using the information-theoretic measure of mutual information, we compare all combinations of isosurfaces to determine their degree of dependency. This process results in an isosurface similarity map which provides a compact overview of the similarities.

The main contribution of this paper is a new approach for quantifying and visualizing the similarity between isosurfaces in a scalar field. We demonstrate its applicability for simplifying isovalue selection and enhancing scalar-field visualization. However, we want to stress that we do not directly compete with the plethora of techniques which employ multi-dimensional transfer functions based on local voxel properties. We recognize that it is useful and, for many types of data, unavoidable to employ such classifiers in order to separate features which share the same value ranges. Indeed, as the isovalue is one axis in many multi-dimensional

transfer function domains, our approach complements these techniques.

2. Related Work

Due to the complex nature of volumetric data sets, techniques for providing a simplified view of the data are frequently used. Their purpose is to guide the user to interesting regions which can subsequently be investigated in a three-dimensional view. Histograms are one of the most common representations employed for this purpose. They visualize the data set by depicting the number of voxels for each data value. Carr et al. [CDD06], in a result later refined by Scheidegger et al. [SSD*08], demonstrated that this actually converges to the distribution of isosurface areas but is formally equivalent to the nearest neighbor interpolant. They proposed several practical measures which show better convergence behavior. Isosurface area, however, still conveys limited information about the nature and structure of a data set. For this reason, Bajaj et al. [BPS97] displayed a variety of additional isosurface statistics in their contour spectrum. Pekar et al. [PWH01] suggested the use of a Laplacian-weighted histogram to assist in the detection of significant isovalues.

Several approaches apply topological analysis to volume data. The contour tree [CSVDP10] is an abstraction of a scalar field that encodes the nesting relationships of isosurfaces. Takahashi et al. [TTFN06] employed a volume skeleton tree to identify isosurface embeddings in order to provide additional structural information. Hyper Reeb graphs, proposed by Fujishiro et al. [FTAT00], capture the topological skeleton of a volumetric data set and can serve as a reference structure for designing comprehensible volume visualizations. One problem of these methods is that they rely on geometric extraction processes which suffer from noise in real-world data.

Much work has focused on the collection of local properties such as first and second derivatives. Plotting these attributes against isovalues provides guidance for feature selection. Kindlmann and Durkin [KD98] demonstrated that a two-dimensional histogram of data value and gradient magnitude enables the identification of boundary regions which manifest themselves as arches. Kniss et al. [KKH02] extended this idea to the common notion of two-dimensional transfer functions and developed a direct manipulation interface for their specification. Lum et al. [LM04] used additional gradient-aligned samples depicted in a line-based histogram. Šereda et al. [ŠVSG06] extended this idea by searching for high and low values in paths that follow the gradient. Tenginkai et al. [TLM01] employed multi-dimensional histograms based on local higher-order moments to detect important data values. To better characterize the shape of local features, Sato et al. [SWB*00] used the matrix of second derivatives. The role of curvature was investigated by Kindlmann et al. [KWMT03] and

Hladůvka et al. [HKG00]. Roettger et al. [RBS05] extended histograms by incorporating spatial information. Lundström et al. [LLY06] proposed the use of local histograms to better represent the distribution of intensity values in a given neighborhood which improves tissue separation for the case of overlapping intensity ranges. Correa and Ma proposed classification approaches based on size [CM08], occlusion [CM09a], and visibility [CM09b].

Due to this large number of different classification criteria, several approaches proposed the use of dimensionality reduction techniques to identify regions based on high-dimensional voxel signatures. Tzeng et al. [TM04] used machine learning methods to generate classifications based on a simple painting interface. They also presented a cluster-space approach for interacting with multiple classification criteria [TLM05]. Šereda et al. [ŠVG06] employed hierarchical clustering of material boundaries. Pinto et al. [PF07] utilized self-organizing maps to reduce the dimensionality of the classification space.

Previous methods attempt to characterize volume data sets by analyzing global isosurface statistics, extracting topological relationships, or collecting local voxel signatures. Our approach is fundamentally different in that we measure similarities of isosurfaces as a whole based on a robust information-theoretic measure. We show that this notion can yield additional insights into the structure of the data and can serve as the basis for enhanced visualization.

3. Isosurface Similarity

Isosurfaces play a crucial role in visualizing and interpreting volumetric data. In many cases they represent important object and/or material boundaries. However, in practice it is difficult to identify salient isovalues. Traditionally, one- and two-dimensional histograms have been employed to assist the user in this process. These approaches depict the frequency of isovalues and other attributes (e.g. gradient magnitude). Peaks or clusters in these plots then guide the exploration and visualization process. Frequency, however, can be a problematic measure as large regions such as the background intensity tend to dominate.

We propose *isosurface similarity* as an alternative measure for identifying features and guiding visualization parameter specification. Instead of collecting statistics of individual isosurfaces and using their variation to obtain information about their significance, we are interested in investigating similarities between individual isosurfaces directly, i.e., how much does knowledge about one surface tell us about the others. In this section, we first introduce our new measure for isosurface similarity. We then apply this measure to the isosurfaces of a scalar field to provide an overview of the mutual similarities which complements traditional frequency-based depictions.

3.1. Similarity Measure

An *isosurface*, or level set, of a volumetric scalar-valued function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ is the locus of all points at which f attains the *iso*value h :

$$L_h = \left\{ x \in \mathbb{R}^3 : f(x) = h \right\} \quad (1)$$

A popular information-theoretic measure of similarity which has been applied in many areas including shape registration [HPM06], multi-modality fusion [HBKG08], and view-point selection [VFSG06] is *mutual information*. The mutual information of two discrete random variables X and Y can be defined as [Yao03]:

$$I(X, Y) = \sum_{x \in X} \sum_{y \in Y} p_{X,Y}(x, y) \log \left(\frac{p_{X,Y}(x, y)}{p_X(x) p_Y(y)} \right) \quad (2)$$

where $p_{X,Y}$ is the joint probability distribution function of X and Y , and p_X and p_Y are the marginal probability distribution functions of, respectively, X and Y . Mutual information quantifies the information that X and Y share by measuring how much knowing one of these variables reduces the uncertainty about the other. In other words, mutual information measures the dependence between the joint distribution of X and Y and what the joint distribution would be if X and Y were independent. The mutual information of X and Y is zero if and only if they are statistically independent. In the case of identity of X and Y , the mutual information is equal to the uncertainty associated with the random variable, i.e., its entropy. Mutual information can be equivalently expressed in terms of entropy as:

$$I(X, Y) = H(X) + H(Y) - H(X, Y) \quad (3)$$

with

$$H(X) = - \sum_{x \in X} p_X(x) \log(p_X(x)) \quad (4)$$

$$H(Y) = - \sum_{y \in Y} p_Y(y) \log(p_Y(y)) \quad (5)$$

$$H(X, Y) = - \sum_{x \in X} \sum_{y \in Y} p_{X,Y}(x, y) \log(p_{X,Y}(x, y)) \quad (6)$$

where $H(X)$ and $H(Y)$ denote the marginal entropies and $H(X, Y)$ is the joint entropy of X and Y . As it is convenient to work with values in $[0, 1]$, a normalized measure can be obtained by [Kvã87]:

$$\hat{I}(X, Y) = \frac{2I(X, Y)}{H(X) + H(Y)} \quad (7)$$

Volume data frequently exhibit an onion-peel-like structure and contain material inhomogeneities, partial volume effects, and noise. This results in several redundant isosurfaces, i.e., they do not represent substantial additional information. We would like to obtain a measure which classifies these kind of isosurfaces as similar. In the registration literature, shape representations based on implicit distance functions are commonly used as they have proven to be stable and robust to shape perturbations and noise [HPM06]. For

these reasons, we choose to represent individual isosurfaces using their *distance transform*. The distance transform D_h of an isosurface with isovalue h gives the minimum distance of a point x to the surface [JBŠ06]:

$$D_h(x) = \min_{y \in L_h} d(x, y) \quad (8)$$

where d is a distance measure (for the remainder of this paper we will assume the Euclidean distance). We can now consider the distances from any point to a pair of isosurfaces L_p and L_q as random variables X and Y . In order to compute the mutual information between the two isosurfaces, we need to estimate the joint distribution of X and Y . This can be accomplished using the joint histogram of D_p and D_q : for every voxel position x , we record the distances $D_p(x)$ and $D_q(x)$ in a two-dimensional histogram where each bin represents a certain range of distances. The marginal probability distributions of X and Y can be estimated by summing over the columns and rows, respectively, of the joint histogram. This allows us to directly compute $H(X, Y)$, $H(X)$, and $H(Y)$ to evaluate the normalized mutual information, as defined in Equation 7, of the two isosurfaces as a measure of their similarity.

3.2. Similarity Maps

In order to obtain information about the similarity relationships in a data set, we introduce the *isosurface similarity map*. Given the discrete set of N isovalues $V = \{h_1, \dots, h_N\}$ in a data set, we generate a $N \times N$ matrix $SM_V(i, j)$ containing the isosurface similarity, computed as described in the previous section, for each combination of isovalues h_i and h_j . We will also use the notation $SM_V(x, y)$ to denote the matrix element $SM_V(i, j)$ with $x = h_i$ and $y = h_j$ when convenient. Due to the properties of mutual information, the map is symmetric and one along the main diagonal. As it records the similarity between every pair of isosurfaces, it provides an overview of the similarity relationships in the data set. In contrast to histograms, which visualize the frequency of individual values, it instead depicts the similarity of isosurfaces measured by the mutual information of their distance fields.

By summing over the rows (or columns) of the isosurface similarity map and normalizing the result by the number of isovalues, we obtain the *isosurface similarity distribution* SD_V :

$$SD_V(i) = \frac{1}{|V|} \sum_{j=1}^{|V|} SM_V(i, j) \quad (9)$$

The isosurface similarity distribution describes the average similarity for each individual isosurface. Peaks in the similarity distribution correspond to those isosurfaces which are most similar to others while valleys indicate regions of rapid change. As will be shown in Section 4, it is frequently useful to investigate the similarity distribution for a specific subset




Figure 1: Example of our isosurface similarity measure. Row (1): column (a) shows an image of three concentric circles, in column (b) the innermost sphere is replaced by a square of approximately equal area, and columns (c) and (d) show the same two images with added noise. In row (2) the respective histograms are shown. Row (3) depicts the corresponding isosurface similarity maps. Row (4) shows the isosurface similarity distributions.

of isovalues. In practice, a summed-area table of the similarity map enables the efficient evaluation of similarity distribution values for arbitrary continuous subranges.

Figure 1 shows a simple example. The image in row (1), column (a) contains three concentric circles with different isovalues. In column (b) the innermost circle is replaced by a square of approximately equal area. Expectedly, the corresponding histograms shown in row (2) are essentially identical. The isosurface similarity maps depicted in row (3), however, show considerable differences. Similarity is linearly mapped to grayscale intensity where white means a similarity of zero and black corresponds to a similarity of one. While row (3), column (a) shows a high degree of mutual similarity between the three spheres, the square's presence is clearly indicated in row(3), column (b). This is also re-

flected in the corresponding similarity distributions depicted in row (4). Columns (c) and (d) of the figure demonstrate that the basic structure of similarity map and distribution remains the same even though noise has been added to the images.

4. Applications

In this section, we present applications of isosurface similarity maps and distributions for the visualization of volume data. We do not advocate replacement of well-proven methods such as histograms which are clearly useful for many purposes nor do we propose similarity as a sole classification criterion. Instead, we want to demonstrate that isosurface similarity provides additional information which can be exploited to build improved tools for volume visualization.

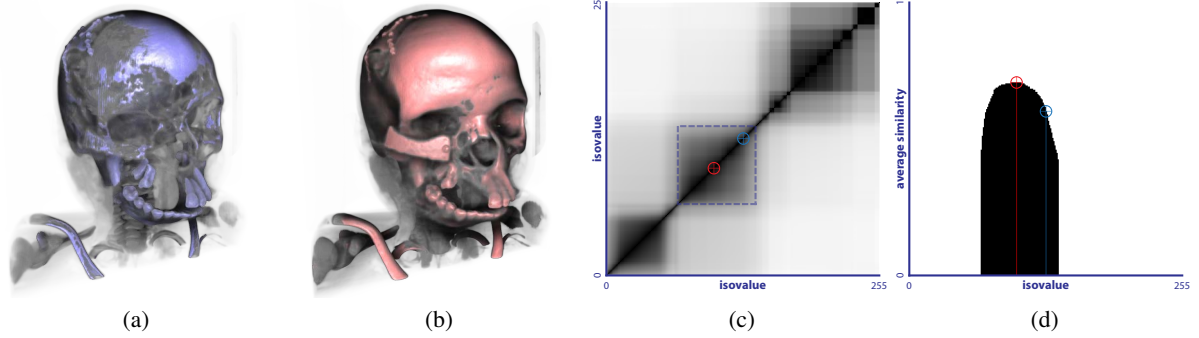


Figure 2: Result images for similarity-enhanced isosurface visualization are shown in (a) and (b), (c) depicts the isosurface similarity map of the data set, and (d) shows the isosurface similarity distribution for the highlighted region in the similarity map. The isovalues depicted in (a) and (b) are marked with corresponding colors in (c) and (d).

4.1. Similarity-Enhanced Isosurface Visualization

A common problem in volume visualization is that even minor changes in the selected isovalues can have dramatic impact on the depicted features. The importance of providing the user with information on this kind of uncertainty was demonstrated by Lundström et al. [LLPY07] in the context of stenosis assessment. In a similar spirit, isosurface similarity allows us to indicate the stability of an isosurface by visually encoding the similarity of a sample point with respect to a specified isoalue.

We want to depict the isosurface at an user-selected isoalue h_u in a focus+context manner using direct volume rendering [HMBG01]. In addition to the isosurface itself (focus), information about similar regions (context) should be conveyed to the user. To prevent visual overload, the contextual regions should be depicted in a sparse way. We use the following importance function $\gamma(x)$ which determines the degree-of-interest in a sample at location x :

$$\gamma(x) = \prod_{y \in C(x)} SM_V(h_u, f(y)) \quad (10)$$

where f is the scalar field and $C(x)$ denotes a local neighborhood of samples around x . In practice, we reuse the sample values required for gradient estimation using central differences as the sample neighborhood. The effect of this function is that, due to the product of similarities of neighboring voxels, only contextual regions with high local similarity will tend to exhibit high importance. For visualizing the isosurface h_u we specify the opacity of a sample $\alpha(x)$ as:

$$\alpha(x) = \begin{cases} 1 & \text{if } f(x) \geq h_u \\ \gamma(x) & \text{otherwise} \end{cases} \quad (11)$$

To clearly distinguish between focus and context, γ is also used to control the color and the degree of surface shading. Additionally, the directional occlusion model of Schott et al. [SPH*09] is used uniformly for all samples.

An example is shown in Figure 2. The isoalue selected

in Figure 2 (a) is very unstable – the extent of the cloud surrounding the surface indicates that it only captures part of the structure of interest. Figure 2 (b) shows more stability. Note that the similarity cloud has the same shape in both images – this means that both isosurfaces are part of a cluster of high mutual similarity. This is confirmed by locating the isovalues in the similarity map shown in Figure 2 (c) – both lie within one large cluster. Figure 2 (d) shows the similarity distribution for the indicated square region. It can be seen that the isoalue used in Figure 2 (a) has low average similarity, while the value from Figure 2 (b) is located at a peak, i.e., it represents the region well.

4.2. Similarity-Based Isoalue Remapping

Isovalues are typically selected and modified using user interface widgets such as sliders or by linearly mapping them to mouse movement. This, however, can be quite non-intuitive: if a subrange of isovalues corresponds to very similar isosurfaces, large changes of the value will have almost no effect on the depicted structures. Conversely, in regions of high dissimilarity even a minor modification can completely alter the appearance. Ideally, the effects in the visualization should correspond to the magnitude of change in the corresponding user interface component. Thus, instead of directly translating changes of the controlling element to changes in the isoalue, we use a nonlinear mapping based on isosurface similarity. Let ML_V be the conventional direct mapping function which maps values $[0, 1]$ linearly to the range of isovalues $[h_{min}, h_{max}]$ in the set V :

$$ML_V(x) = h_{min} + x(h_{max} - h_{min}) \quad (12)$$

The idea is to use a monotonous function $MS_V(x)$ with $MS_V(0) = ML_V(0)$ and $MS_V(1) = ML_V(1)$ whose derivative is controlled by the similarity of neighboring isovalues. For this purpose, we define the *cumulative similarity* of an




Figure 3: Transition using linear mapping $ML_V(x)$ (top row) and similarity-based mapping $MS_V(x)$ (bottom row) for $x \in [0.22, 0.32]$.

ordered set of isovalues V :

$$SC_V(i) = \sum_{j=1}^i SM_V(j-1, j) \quad (13)$$

where $SM_V(0, 1) = 0$. Our similarity-based mapping function can then be written as:

$$MS_V(x) = ML_V \left(\frac{SC_V(x(|V| - 1) + 1)}{SC_V(|V|)} \right) \quad (14)$$

The fact that this function is piecewise constant does not matter in practice, since we are typically only interested in discrete isovalues. However, one can simply use any interpolant for SC_V to remedy this.

Figure 3 depicts a transition using $ML_V(x)$ (top row) and $MS_V(x)$ (bottom row) with x varying from 0.22 to 0.32 in increments of 0.02. Even though the differences are subtle, the images generated using the similarity-based mapping function show a more uniform progression. The isosurface similarity map for the data set is shown in Figure 4 (a) – the highlighted area indicates the range of isovalues of the transition. Note that the dissimilar nature of the isosurfaces in the interval is clearly visible. The graph in Figure 4 (b) depicts the function $ML_V(x) - MS_V(x)$ for the chosen interval, i.e., the difference in isovalue resulting from using the similarity-based mapping function instead of the linear one.

4.3. Representative Isovalue Selection

An important problem in volume visualization has been the identification of relevant isovalues. Many approaches combine different isosurface statistics to infer salient isovalues which characterize the function well. The isosurface similarity map can be used to provide guidance in finding representative isovalues. By selecting rectangular regions in the map and investigating their similarity distribution, insight on




Figure 4: Similarity-based isovalue remapping. (a) Isosurface similarity map for the data set shown in Figure 3 – the highlighted area indicates the interval of isovalues in the transition. (b) Plot of the difference between the linear mapping function $ML_V(x)$ and the similarity-based mapping function $MS_V(x)$ for $x \in [0.22, 0.32]$.

the relationships of isovalues and the corresponding structures can be gained. While manual analysis is useful and unavoidable for many exploratory tasks as the intent of the user is not known, the highly structured nature of the isosurface similarity map also provides us with means to automatically identify relevant isovalues. Regions of high similarity manifest themselves as distinct squares in the isosurface similarity map. We developed a simple algorithm which allows us to automatically identify these regions and, based on their similarity distributions, select the most representative isovalues for each of them:

Step 1 – Our aim is to identify representative isovalues, i.e., values with high similarity to many other values. Initially, the value with the highest average similarity to all others is identified. Then, the maxima of the similarity distribution for only the values below and above this value are cho-



Figure 5: Representative isovalue selection algorithm applied to a CT data set. The isosurface similarity map is shown on the left and the six most representative isovalues are marked. The corresponding isosurfaces are depicted in the middle section numbered from one to six with decreasing relevance. The image on the right shows a cutaway view of the data set classified according to maximum similarity with the six isovalues.

sen, one so on. Thus, our strategy recursively partitions the set of isovalues V by selecting the maximum of the similarity distribution for the current subset. The chosen isovalue m is inserted into a priority queue Q based on its similarity distribution value weighted by the number of isovalues in the current subset. The following procedure **prioritize**(Q, V) summarizes these operations:

$$\begin{aligned}
 m &\leftarrow \arg \max_{h_i \in V} SD_V(i) \\
 p &\leftarrow |V| SD_V(m) \\
 &\text{enqueue}(Q, m, p) \\
 V_1 &\leftarrow \{h_i \in V : 1 \leq i < m\} \\
 &\text{prioritize}(Q, V_1) \\
 V_2 &\leftarrow \{h_i \in V : m < i \leq |V|\} \\
 &\text{prioritize}(Q, V_2)
 \end{aligned} \tag{15}$$

Step 2 – Next, we remove the isovalue h_m with maximum priority from the queue. To prevent similar values from being chosen, all remaining entries are penalized based on their similarity with the selected value in the following manner:

$$p_i \leftarrow \frac{p_i}{1 + SM_V(h_m, h_i)} \tag{16}$$

where p_i is the priority of the i -th item in the queue and h_m is the selected value with the highest priority. This process repeats until no more items remain in the queue.

This simple algorithm results in a reordering of the isovalues. Early values in the resulting order have high similarity with many other isovalues, i.e., they represent a certain range of isovalues well, but low mutual similarity meaning that they are likely to correspond to distinct structures. One major advantage of this approach is that it does not require any kind of threshold or parameter. The user can simply examine the isosurfaces in the order generated by the algorithm. After the first few isovalues corresponding to distinct features of the data set, subsequent values will typically be less




Figure 6: Automatically classified CT data set using the eight most representative isovalues.

representative values for the same structures as no further dissimilar values can be found.

An example is shown in Figure 5. The six most representative isovalues of a CT data set determined using the described algorithm are marked in the isosurface similarity map and the corresponding isosurfaces are shown numbered from 1 to 6 with decreasing relevance. While the first three values correspond to distinct structures, the remaining ones only partially segment these features. The righthand side of the figure depicts a cutaway view where each voxel is classified according to its maximum similarity with any of the six isovalues – as the first three isovalues exhibit more similarity, the three less relevant isosurfaces are not visible. A further result for the classification of CT data using the most representative isovalues is shown in Figure 6. In CT data sets, there is a clear correspondence between isovalues and different tissue types. Even though other types of volume data do not exhibit the same kind of relationships and are



Figure 7: Representative isovalue selection algorithm applied to an MRI angiography data set. The isosurface similarity map is shown on the left and the eight most representative isovalues are marked. The corresponding isosurfaces are depicted in the middle section numbered from one to eight with decreasing relevance. The image on the right shows a cutaway view of the data set classified according to maximum similarity with the eight isovalues.

therefore difficult to classify based on isovalues alone, our approach can still identify salient structures in these cases. Figure 7 shows an MRI data set and its eight most representative isovalues as identified by our method. In all depicted examples the only manual interventions were specification of the viewpoint and clipping.

5. Implementation

Our tool for the generation of isosurface similarity maps was implemented in C++. The computation process involves two steps. First, the distance transforms for all isovalues are computed. As this can, depending on the resolution, require a substantial amount of space they are immediately written to disk in compressed form. In the second step, we compute the mutual information of the distance transforms for each pair of isovalues. Since mutual information is a symmetric measure, only half of the combinations need to be evaluated. The computation is performed by generating the joint histogram of the two distance transforms which allows estimation of the joint and marginal entropies as described in Section 3.1 using the CUDA-based implementation of Shams and Barnes [SB07].

6. Discussion

In our experiments we found that isosurface similarity maps provide a concise overview of a data set. Distinct features manifest themselves as squares and their size informs the user about the corresponding value range. Transitional regions can be detected through the nesting structure of these squares. In contrast to histograms, large regions do not tend to dominate the depiction. Due to the choice of mutual information as a similarity measure, uncorrelated noise has little influence. These properties indicate that the presented concept has the potential to improve the visualization and analysis of volume data even beyond the examples shown in this paper.

One obvious disadvantage of our approach is the considerable cost of generating the isosurface similarity map. Our implementation can require several hours of processing time. Even though this is a one-time preprocess, we consider this fact a serious limitation of our current method. The most performance-critical component is the construction of the joint histogram of two distance transforms since it has to be performed for each pair of isovalues. In order to reduce the computation time, we performed experiments with down-sampled distance transforms. Interestingly, it seems that the resolution can be considerably reduced without substantial changes in the resulting similarity map. Figure 8 compares similarity maps computed from distance transforms at several resolutions. Note that the distance transforms are generated at the original resolution of the data set and then down-sampled as opposed to computing them from a downsampled version of the volume. We believe that the reason for this stability is that the distance transform captures the unique characteristics of an isosurface even at reduced resolutions. Since it is not used to perform accurate spatial comparisons but rather as a shape descriptor, we consider downsampling a viable strategy. Table 1 lists the distance transform resolutions and similarity map computation times for all data sets used in the paper. On a typical notebook, generation of the isosurface similarity map using a distance transform resolution of approximately $64 \times 64 \times 64$ takes about 25 minutes for standard data sets. Throughout the paper, we used a fixed number of 128×128 bins in the computation of the joint histogram.

Even though lowering the resolution of the distance transform dramatically reduces the computation time to a level we consider acceptable, our generation method is still a brute-force approach. For a more fundamental improvement, one potential direction is the use of a different method for joint probability density estimation. While joint histograms are widely used, other methods are gaining increasing recognition. It may even be possible to use an alternative ap-




Figure 8: Comparison of isosurface similarity maps computed with different distance transform resolutions for the data set shown in Figure 5. The total computation times (in minutes) were (a) 569.1, (b) 35.8, (c) 22.1, and (d) 20.6.

proach which does not require explicit computation of the distance transform. Interesting recent work by Rajwade et al [RBR09] points in this direction and remains to be explored. Additionally, an adaptive strategy for approximating the full similarity map could also be employed. While our current method for generating isosurface similarity maps can be considered a proof-of-concept, we have shown that they have useful applications for the visualization of volume data. There are several other areas, however, where the proposed concept may be of interest. In closing, we would like to briefly list some avenues which could be promising directions for further exploration:

Volume quantization and compression. The notion of similarity may be useful in developing new approaches for quantizing and/or compressing volume data with higher fidelity. Isosurfaces which exhibit a high degree of redundancy could be grouped together while still preserving essential features in the data set.

Volume segmentation. Isosurface similarity could also be employed as an alternate metric for segmentation algorithms such as region growing. These methods typically use similarity criteria based on the difference between isovalues, so our measure may help to increase robustness.

Multi-dimensional classification. While there is nothing

Figure(s)	Orig. Resolution	DT Resolution	Time
2	$512 \times 512 \times 361$	$64 \times 64 \times 45$	23.9
3, 4	$512 \times 512 \times 333$	$64 \times 64 \times 41$	21.8
5	$256 \times 256 \times 230$	$64 \times 64 \times 57$	22.1
6	$512 \times 512 \times 361$	$64 \times 64 \times 45$	23.7
7	$512 \times 512 \times 125$	$64 \times 64 \times 15$	21.6

Table 1: Depicting figure, original data set resolution, resolution of the downsampled distance transform, and total computation time (in minutes) for the isosurface similarity maps used in the paper are listed. System configuration: Intel Core 2 Duo 2.53 GHz CPU, 4 GB RAM, NVidia GeForce 9600M GT GPU.

ing in our approach that prevents combination with multi-dimensional classification approaches using gradient magnitude [KKH02], curvature [KWTM03], occlusion [CM09a], or other measures proposed in the literature, we did not investigate this area. Indeed, as the isovalue typically represents one axis in multi-dimensional transfer function schemes, similarity could help to better separate features and to decrease the influence of noise.

7. Conclusions

In this paper, we introduced the notion of isosurface similarity for the visualization of volume data. This new measure quantifies the similarity of two isosurfaces as the normalized mutual information of their respective distance transforms. The resulting isosurface similarity map provides a visualization of these similarities and gives an overview of a data set which complements traditional depictions. Additionally, the similarity map can be used to improve rendering and parameter specification. Its structured nature enables automatic detection of representative isovalues to assist the exploration process. The presented concept opens up several interesting directions for future investigation.

Acknowledgements

The work presented in the paper was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC). The data sets used in Figures 2, 5, and 6 are courtesy of the OsiriX Foundation (<http://www.osirix-viewer.com>).

References

- [BPS97] BAJAJ C., PASCUCCI V., SCHIKORE D.: The contour spectrum. In *Proceedings of IEEE Visualization 1997* (1997), pp. 167–173. 2
- [CDD06] CARR H., DUFFY B., DENBY B.: On histograms and isosurface statistics. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006), 1259–1265. 2
- [CM08] CORREA C., MA K.-L.: Size-based transfer functions: a new volume exploration technique. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1380–1387. 2

- [CM09a] CORREA C., MA K.-L.: The occlusion spectrum for volume classification and visualization. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1465–1472. 2, 9
- [CM09b] CORREA C., MA K.-L.: Visibility-driven transfer functions. In *Proceedings of PacificVis 2009* (2009), pp. 177–184. 2
- [CSVDP10] CARR H., SNOEYINK J., VAN DE PANNE M.: Flexible isosurfaces: Simplifying and displaying scalar topology using the contour tree. *Computational Geometry: Theory and Applications* 43, 1 (2010), 42–58. 2
- [FTAT00] FUJISHIRO I., TAKESHIMA Y., AZUMA T., TAKAHASHI S.: Volume data mining using 3D field topology analysis. *IEEE Computer Graphics and Applications* 20, 5 (2000), 46–51. 2
- [HBKG08] HAIDACHER M., BRUCKNER S., KANITSAR A., GRÖLLER M.: Information-based transfer functions for multimodal visualization. In *Proceedings of Visual Computing for Biomedicine 2008* (2008), pp. 101–108. 3
- [HKG00] HLADŮVKA J., KÖNIG A., GRÖLLER M.: Curvature-based transfer functions for direct volume rendering. In *Proceedings of the Spring Conference on Computer Graphics* (2000), pp. 58–65. 2
- [HMBG01] HAUSER H., MROZ L., BISCHI G., GRÖLLER M.: Two-level volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 7 (2001), 242–252. 5
- [HPM06] HUANG X., PARAGIOS N., METAXAS D.: Shape registration in implicit spaces using information theory and free form deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28 (2006), 1303–1318. 3
- [JBŠ06] JONES M., BAERENTZEN J., ŠRÁMEK M.: 3D distance fields: a survey of techniques and applications. *IEEE Transactions on Visualization and Computer Graphics* 12 (2006), 581–599. 3
- [KD98] KINDLMANN G., DURKIN J.: Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of the IEEE Symposium on Volume Visualization 1998* (1998), pp. 79–86. 2
- [KKH02] KNISS J., KINDLMANN G., HANSEN C.: Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 8 (2002), 270–285. 2, 9
- [Kvã87] KVALSETH T.: Entropy and correlation: Some comments. *IEEE Transactions on Systems, Man, and Cybernetics* 17 (1987), 517–519. 3
- [KWTM03] KINDLMANN G., WHITAKER R., TASDIZEN T., MÖLLER T.: Curvature-based transfer functions for direct volume rendering: Methods and applications. In *Proceedings of IEEE Visualization 2003* (2003), pp. 513–520. 2, 9
- [LC87] LORENSEN W., CLINE H.: Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics* 21 (1987), 163–169. 1
- [LLPY07] LUNDSTRÖM C., LJUNG P., PERSSON A., YNNERMAN A.: Uncertainty visualization in medical volume rendering using probabilistic animation. *IEEE Transactions on Visualization and Computer Graphics* 13 (2007), 1648–1655. 5
- [LLY06] LUNDSTRÖM C., LJUNG P., YNNERMAN A.: Local histograms for design of transfer functions in direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 12, 6 (2006), 1570–1579. 2
- [LM04] LUM E., MA K.-L.: Lighting transfer functions using gradient aligned sampling. *Proceedings of IEEE Visualization 2004* (2004), 289–296. 2
- [PF07] PINTO F., FREITAS C.: Design of multi-dimensional transfer functions using dimensional reduction. In *Proceedings of EuroVis 2008* (2007), pp. 130–137. 2
- [PWH01] PEKAR V., WIEMKER R., HEMPEL D.: Fast detection of meaningful isosurfaces for volume data visualization. *Proceedings of IEEE Visualization 2001* (2001), 223–230. 2
- [RBR09] RAJWADE A., BANERJEE A., RANGARAJAN A.: Probability density estimation using isocontours and isosurfaces: applications to information-theoretic image registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31 (2009), 475–491. 9
- [RBS05] ROETTGER S., BAUER M., STAMMINGER M.: Spatialized transfer functions. In *Proceedings of EuroVis 2005* (2005), pp. 271–278. 2
- [SB07] SHAMS R., BARNES N.: Speeding up mutual information computation using NVIDIA CUDA hardware. In *Proceedings of Digital Image Computing: Techniques and Applications 2007* (2007), pp. 555–560. 8
- [SPH*09] SCHOTT M., PEGORARO V., HANSEN C., BOULANGER K., BOUATOUCH K.: A directional occlusion shading model for interactive direct volume rendering. *Computer Graphics Forum* 28 (2009), 855–862. 5
- [SSD*08] SCHEIDEGGER C., SCHREINER J., DUFFY B., CARR H., SILVA C.: Revisiting histograms and isosurface statistics. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1659–1666. 2
- [ŠVG06] ŠEREDA P., VILANOVA A., GERRITSEN F.: Automating transfer function design for volume rendering using hierarchical clustering of material boundaries. In *Proceedings of EuroVis 2006* (2006), pp. 243–350. 2
- [ŠVSG06] ŠEREDA P., VILANOVA A., SERLIE I., GERRITSEN F.: Visualization of boundaries in volumetric data sets using lh histograms. *IEEE Transactions on Visualization and Computer Graphics* 12 (2006), 208–218. 2
- [SWB*00] SATO Y., WESTIN C., BHALERAO A., NAKAJIMA S., SHIRAGA N., TAMURA S., KIKINIS R.: Tissue classification based on 3D local intensity structures for volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 6, 2 (2000), 160–180. 2
- [TLM01] TENGINAKAI S., LEE J., MACHIRAJU R.: Salient isosurface detection with model-independent statistical signatures. In *Proceedings of IEEE Visualization 2001* (2001), pp. 231–238. 2
- [TLM05] TZENG F.-Y., LUM E., MA K.-L.: An intelligent system approach to higher-dimensional classification of volume data. *IEEE Transactions on Visualization and Computer Graphics* 11, 3 (2005), 273–284. 2
- [TM04] TZENG F.-Y., MA K.-L.: A cluster-space visual interface for arbitrary dimensional classification of volume data. *Proceedings of VisSym 2004* (2004), 17–24. 2
- [TTFN06] TAKAHASHI S., TAKESHIMA Y., FUJISHIRO I., NIELSON G.: Emphasizing isosurface embeddings in direct volume rendering. In *Scientific Visualization: The Visual Extraction of Knowledge from Data*, Bonneau G., Ertl T., Nielson G., (Eds.). Springer, 2006, pp. 185–206. 2
- [VFSG06] VIOLA I., FEIXAS M., SBERT M., GRÖLLER M.: Importance-driven focus of attention. *IEEE Transactions on Visualization and Computer Graphics* 12 (2006), 933–940. 3
- [Yao03] YAO Y.: Information-theoretic measures for knowledge discovery and data mining. In *Entropy Measures, Maximum Entropy Principle and Emerging Application*, Karmeshu, (Ed.). Springer, 2003, pp. 115–136. 3