Volume Analysis Using Multimodal Surface Similarity

Martin Haidacher, Stefan Bruckner, *Member, IEEE Computer Society*,
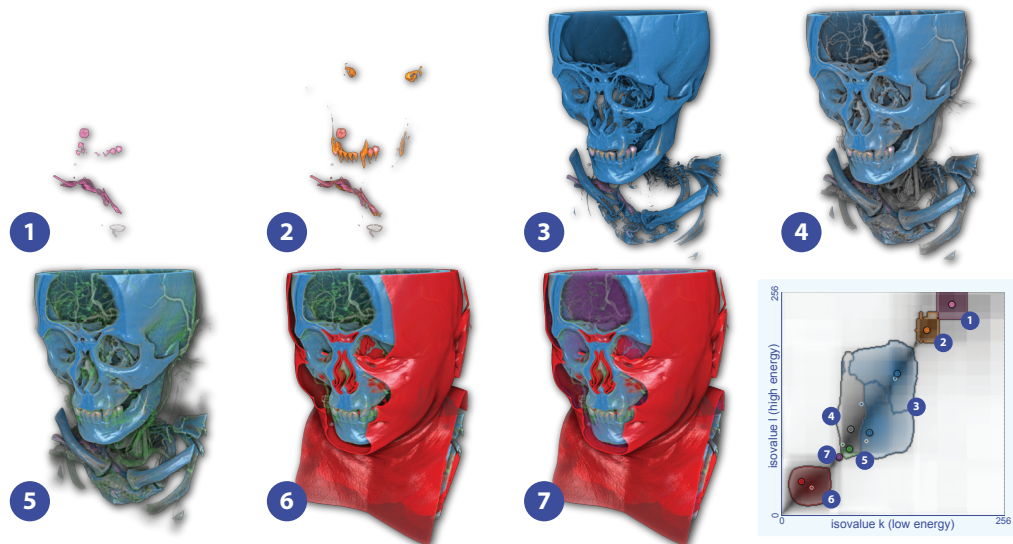and M. Eduard Gröller, *Member, IEEE Computer Society*

Fig. 1. Iterative control point specification for similarity-based classification of a dual energy CT (DECT) angiography data set. The individual steps are numbered from 1 to 7.

Abstract—The combination of volume data acquired by multiple modalities has been recognized as an important but challenging task. Modalities often differ in the structures they can delineate and their joint information can be used to extend the classification space. However, they frequently exhibit differing types of artifacts which makes the process of exploiting the additional information non-trivial. In this paper, we present a framework based on an information-theoretic measure of isosurface similarity between different modalities to overcome these problems. The resulting similarity space provides a concise overview of the differences between the two modalities, and also serves as the basis for an improved selection of features. Multimodal classification is expressed in terms of similarities and dissimilarities between the isosurfaces of individual modalities, instead of data value combinations. We demonstrate that our approach can be used to robustly extract features in applications such as dual energy computed tomography of parts in industrial manufacturing.

Index Terms—Multimodal data, volume visualization, surface similarity.

◆

---

1 INTRODUCTION

Imaging modalities have different advantages and disadvantages typically related to the physical principles they use to scan a specimen. They may suffer from different kinds of artifacts, can be differently affected by noise, may be able to distinguish different materials or tissues, and can have differences with respect to contrast and resolution. In order to gain insight into the phenomenon under investigation, it is essential to integrate this information effectively. The work presented in this paper focuses on the analysis and fusion of two registered volume data sets of the same specimen. While the data generated by each modality may be visualized separately, it is difficult to mentally integrate multiple three-dimensional sources, particularly if spatial relationships are important. Thus, the effective visual fusion of multiple volume data sets has long been an active area of research. As discussed

- *Martin Haidacher, Stefan Bruckner, and M. Eduard Gröller are with the Institute of Computer Graphics and Algorithms, Vienna University of Technology, E-mail: {haidacher|bruckner|groeller}@cg.tuwien.ac.at.*

by Cai and Sakas [4], this combination can occur at different stages. At the extreme ends of the spectrum, the two data sets are treated separately and are only blended at the image level, or, conversely the data values at each position are combined at the very beginning of the pipeline to form a single merged volume. Most commonly visual fusion is performed during the rendering phase which provides spatial integration and allows for a flexible mapping of data attributes to optical properties [24].

While straightforward blending can be an effective technique in 2D slice views, it has many disadvantages in 3D visualization. In particular, the projection of multiple volumetric data sets onto a single 2D image can quickly lead to visual clutter. Hence, it is important to provide the user additional guidance about the spatial similarities and differences between the individual modalities to enable goal-directed selection of features. Approaches which attempt to identify correspondences based only on the frequency of data values, however, suffer from the fact that data value ranges of corresponding structures of interest may differ significantly. In order to address this challenge, we propose *multimodal surface similarity* as a measure for identifying similarities and dissimilarities between two volumetric scalar fields. Instead of collecting statistics about the frequency of data values, we quantify spatial similarities between isosurfaces across two modali-
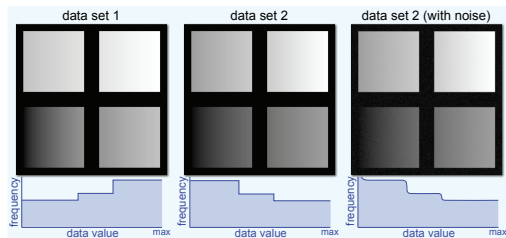
Fig. 2. Data sets containing the same structures with different value ranges (supplementary data). The histograms show the distributions of the data values.
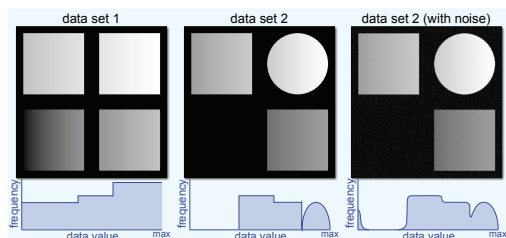

Fig. 3. Two synthetic data sets which represent complementary data. Data set 2 contains structures which are different from data set 1.

ties, i.e., how much does knowledge about one surface tell us about the others. By generating a multimodal similarity map, which encodes the similarity between all combinations of isosurfaces from two modalities, we can provide a concise overview of the differences between two scalar fields. This information can then be used to guide the identification of structures of interest. Based on this concept, we present a novel method for feature classification in similarity space which enables the user to easily take advantage of the complementary information provided by two modalities.

The remainder of the paper is structured as follows. In Section 2 we review related work on the visualization of multimodal volume data and other approaches connected to our work. Section 3 provides background information on the types of data we focus on. In Section 4, the general concept of multimodal surface similarity is introduced. In Section 5, we show how multimodal surface similarity can be used in the visualization process. Results obtained with our approach are presented in Section 6. Section 7 discusses implementation details. The implications of our approach as well as its limitations are discussed in Section 8. Finally, the paper is concluded in Section 9.

2 RELATED WORK

As mentioned in the introduction, two volumes can be fused in different stages of the visualization pipeline [13, 14]. The fusion in image space is covered by the field of image processing [37]. The drawback of the fusion in image space is the loss of 3D information. For the fusion in volume space, the spatial information of the data sets can be used to improve the fusion quality. The first methods for volume fusion were based on extracted surfaces. Levin et al. [27] generated a surface model from an MRI scan and mapped PET-derived measurements onto this surface. Evans et al. [10] generated an integrated volume visualization from the combination of MRI and PET. Noz et al. [29] introduced a framework for 3D registration and fusion of CT/MRI and SPECT data sets based on a polynomial warping technique. These works mainly focused on the combination of anatomical and functional modalities. A more general approach for the fusion of modalities was introduced by Zuiderveld and Viergever [42]. For this method an additional segmentation of the volumes is necessary to decide which one to show at a given sample point. Heinzl et al. [16] introduced a processing pipeline for surface extraction in dual energy CT.

Alternatively, fusion can be performed without an intermediate fea-

ture extraction step. A straightforward method is fusion by linear intermixing of the data values. Such an approach is used for volumetric CSG construction where different volumetric parts are fused into a single object [36, 11, 8]. Hong et al. [17] describes how fusion techniques in volume space can be efficiently implemented using the graphics hardware. Eusemann et al. [9] have shown that this intermixing can be improved for dual energy CT by adapting the intermixing ratio to different tissues. A case study on visualization of multivariate data where multiple values are present at each sample point was presented by Kniss et al. [24]. In this work the idea of multi-dimensional transfer functions for assigning optical properties to a combination of values was used. Akiba and Ma [1] used parallel coordinates for the visualization of time-varying multivariate volume data. Multimodal visualization of medical data sets by using multi-dimensional transfer functions was discussed by Kniss et al. [25]. Kim et al. [22] presented a technique which simplifies transfer function design by letting the user define a separate transfer function for each modality. Their combination defines a two-dimensional transfer function. Haidacher et al. [15] defined a data fusion and transfer function space for multimodal visualization based on the information content of the individual modalities which aims to reduce the loss of information. In contrast to our method, this approach is only based on the global frequency distribution of values and not on structural similarities between the individual modalities.

In our approach, we use information theory [31] to measure similarities between the different modalities, but it has been applied to many aspects of visualization [35]. In flow visualization, for instance, Xu et al. [39] used information theory to select meaningful streamlines. Feixas et al. [12] presented an information-theoretic approach for optimal viewpoint selection. Chen and Jänicke [7] discussed a general information-theoretic framework for scientific visualization.

For many applications, such as industrial CT [16], surfaces are of particular interest. Surfaces can be used to represent the interfaces between different materials. In order to extract a stable isosurface, the selection of the isovalue is crucial. Khoury and Wenger [21] use the fractal dimension to measure how stable an isovalue is. The lower the dimension, the less noisy the corresponding isosurface is. The contour tree [6] is used to topologically analyze volume data. It is able to encode the nesting relationships of isosurfaces. Takahashi et al. [32] employed a volume skeleton tree to identify isosurface embeddings in order to provide additional structural information. Kindlmann and Durkin [23] introduced a transfer function space in which the gradient magnitude is used as additional classification dimension. Interfaces between materials show up as arches in this transfer function space. In LH histograms, introduced by Šereda et al. [34], the highest and lowest value along a local streamline in the gradient field are used for the classification. Sample points at interfaces between materials form clusters in this space, which represent stable surfaces.

Bruckner and Möller [3] introduced similarity maps which represent the similarity of isosurfaces for different isovalues. For the measurement of the similarity mutual information is used. In a similarity map clusters with high mutual information can be detected. In our approach we extend the idea of the similarity maps to multimodal data. The resulting multimodal similarity maps are used for analysis, fusion, and classification of multimodal data.

3 MULTIMODAL VOLUME DATA

There are different reasons for seeking to combine the information from multiple modalities. In medicine, for instance, it is frequently desired to simultaneously depict anatomical and functional data. Functional data contains information about physiological activities, such as metabolism or blood flow, within a certain tissue or organ. Anatomical imaging modalities, on the other hand, present structural information and typically provide higher resolution. In other fields, such as nondestructive testing, multiple industrial CT scans with different parameters are used for scanning an object. These parameters can affect the contrast and amount of artifacts in different regions. Thus, the goal in this case is to combine the advantages of different scans in order to obtain a better visualization of the object.
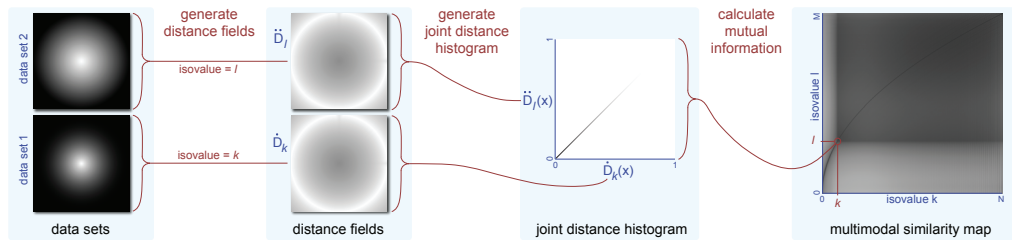
Fig. 4. Pipeline for the generation of a multimodal similarity map. The illustration shows which steps are necessary to calculate the similarity of isosurfaces for the isovalues k and l.

For the further description of our approach we will differentiate between two types of multimodal data which are depicted in this section. We will introduce synthetically generated data sets which represent the two different types of multimodal data sets. The data sets are all 3D data sets, where the slices are duplicates of the same image with a size of 512×512 pixels. To investigate the influence of noise on our method, we added Gaussian white noise with a standard deviation $\sigma = 1.5\%$ ($SNR = 17.6$) to one modality. In the subsequent sections the synthetic data sets are used to highlight the properties of multimodal similarity maps.

Supplementary Data

Multimodal data is often used to eliminate the drawbacks of a certain imaging modality. This is necessary when one modality contains undesirable noise or other artifacts in certain regions. In this case a second modality is used to compensate for these artifacts. In this paper we will refer to this type of data as supplementary data. Basically both modalities depict the same structures, but disadvantages of one modality may be compensated by the other and vice versa. An example for supplementary data is dual energy CT. It is used in medicine and industrial applications. The most common artifacts in CT scans in general are noise-induced streaks, beam hardening, partial volume effects, aliasing, and scattered radiation [2, 18]. Due to the fact that different energy levels have different attenuation characteristics, some of these artifacts appear prominently only in one energy level. Hence it is desired to reduce artifacts by the fusion of CT data sets of different energy levels.

We generated two synthetic data sets in order to simulate multimodal data with supplementary characteristics. In Figure 2 these two data sets and their histograms are depicted. Due to scaling reasons the frequency of background points with a value of zero is omitted in the histograms. On the right side of Figure 2, data set 2 with additional Gaussian white noise is shown. Both data sets contain four squares with gradually changing data values from left to right. Their value ranges are different in both data sets to simulate the effects of varying attenuation characteristics in modalities such as dual energy CT.

Complementary Data

In some cases, it is also advantageous to combine the information of modalities with more distinct characteristics. In such a scenario, a significant amount of information differs or is not represented in one of the modalities. We will refer to this type of multimodal data as complementary data. Complementary data is commonly encountered in medicine. Modalities such as CT and MRI measure different physical characteristics of the human body, and thus there are substantial differences between two such scans of the same patient. An even more pronounced example is the combination of anatomical and functional modalities, such as CT and PET. There is only a rough correspondence between the two modalities, as CT images contain no functional information at all.

We will use the synthetic data sets illustrated in Figure 3 to represent complementary data. Data set 1 contains four squares while data set 2 contains two squares and a circle. The missing square and the circle in data set 2 represent the complementary nature of the data. With the circle we want to show how differences in shape are depicted in the multimodal similarity map. The missing square is used to show the effect if one object is completely omitted from one modality. In the next section these synthetic data sets are used to explain multimodal surface similarity measurement.

4 MULTIMODAL SURFACE SIMILARITY

Isosurfaces are important features in volumetric data. An isosurface of a volumetric scalar field $f : \mathbb{R}^3 \to \mathbb{R}$ is the locus of all points in the scalar field at which f attains an *isovalue k*:

$$L_k = \left\{ x \in \mathbb{R}^3 : f(x) = k \right\} \tag{1}$$

In many cases, different material types correspond to different value ranges in the data set. For example, in medical CT data sets there are typically well-defined intensity ranges associated with soft tissue, fat, and bone. The important characteristic parameter is the intensity isovalue which defines an isosurface representing the boundary of a particular region. Isosurfaces, however, also exhibit a significant amount of redundancy and small variations caused by noise and partial volume effects will result in many similar isosurfaces. Histograms and other isosurface statistics [5, 30] can be used to obtain a better characterization of a data set by depicting distributions of isosurface properties over the range of data values. They are limited, however, in that they treat each isosurface in isolation and therefore cannot capture the spatial relationships between multiple structures.

As an alternative, the measure of isosurface similarity was introduced by Bruckner and Möller [3] to quantify how much information two isosurfaces have in common. They used a matrix of isosurface similarity for all combinations of isovalues within a single data set as the basis for identifying relevant isovalues. We will refer to this method as *self similarity maps* since the measurement of the similarity is between isosurfaces of a single data set. For multimodal data, in particular, it is difficult to investigate differences and similarities based on isosurface statistics as the order and range of corresponding data values may vary significantly. Thus, in order to characterize the correspondences between multiple modalities, we extend the idea to *multimodal similarity maps* which quantify the similarity between all combinations of isosurfaces from two scalar fields.

In the following, we first briefly revisit isosurface self similarity maps as presented by Bruckner and Möller [3] and then describe our extension to multimodal data.

4.1 Self Similarity Maps

A common measure for similarity is *mutual information*. Mutual information is a basic concept from information theory, measuring the statistical dependence between two random variables or the amount of information that one variable contains about the other. It is a particularly attractive measure because no assumptions are made regarding the nature of this dependence and because of its robustness against perturbations [38]. Therefore, mutual information has been applied in many areas including shape registration [19], multi-modality fusion [15], and viewpoint selection [33]. The mutual information of two discrete random variables X and Y can be defined as [40]:

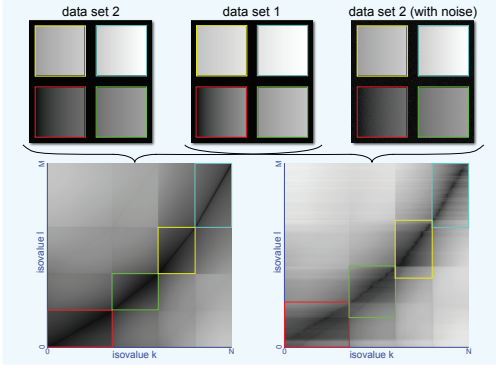$$I(X,Y) = H(X) + H(Y) - H(X,Y) \tag{2}$$

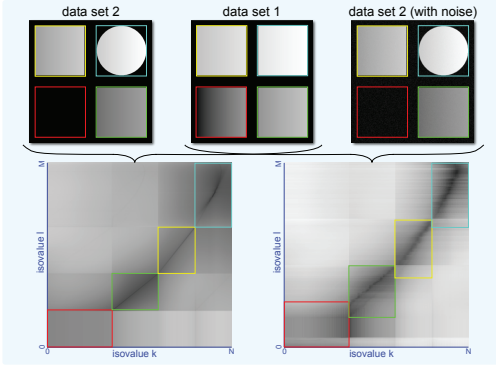Fig. 5. Multimodal similarity maps for supplementary data.

Fig. 6. Multimodal similarity maps for complementary data.

where $H(X,Y)$ is the joint entropy and $H(X)$ and $H(Y)$ are the marginal entropies of random variables X and Y. Since the mutual information is limited by the average marginal entropies, it can be normalized to a value range in $[0,1]$ by [26]:

$$\hat{I}(X,Y) = \frac{2I(X,Y)}{H(X) + H(Y)} \tag{3}$$

As a measure of isosurface similarity, Bruckner and Möller [3] proposed the normalized mutual information of the respective isosurface *distance fields*. For a given isovalue k and an isosurface L_k the distance field D_k can be defined as follows [20]:

$$D_k(x) = \min_{\forall y \in L_k} d(x,y) \tag{4}$$

where d is a distance measure between the points x and y. To measure the similarity between two isosurfaces L_i and L_j, the distance fields for both isosurfaces $D_i(x)$ and $D_j(x)$ are used as discrete random variables X and Y for the calculation of the mutual information based on Equation 3. This leads to a single quantity between 0 and 1 which expresses the similarity between isosurface L_i and isosurface L_j. Higher values mean that the isosurfaces are considered to be more similar.

If we consider N different isovalues $V = \{k_1, ..., k_N\}$ then the self similarity map can be defined as an $N \times N$ matrix $\text{SSM}(i,j)$. Each element (i,j) of the matrix represents the normalized mutual information for a combination of isovalues i and j.

4.2 Multimodal Similarity Maps

In this paper, we extend the concept of isosurface similarity maps to multimodal data. Instead of investigating the similarity of isosurfaces in a single data set, we explore the similarity of two different data sets representing the same object. The isosurfaces of both modalities are represented by:

$$\dot{L}_k = \left\{ x \in \mathbb{R}^3 : \dot{f}(x) = k \right\} \quad \ddot{L}_l = \left\{ x \in \mathbb{R}^3 : \ddot{f}(x) = l \right\} \tag{5}$$

where k and l are the two isovalues, and $\dot{f}$ and $\ddot{f}$ are the scalar-valued functions representing the two modalities. Based on the two isosurfaces, two distance fields $\dot{D}_k$ and $\ddot{D}_l$ can be generated:

$$\dot{D}_k(x) = \min_{\forall y \in \dot{L}_k} d(x,y) \quad \ddot{D}_l(x) = \min_{\forall y \in \ddot{L}_l} d(x,y) \tag{6}$$

Figure 4 illustrates how the mutual information for a combination of isovalues l and k is calculated. The first step is the generation of the distance fields $\dot{D}_k$ and $\ddot{D}_l$ for the isosurfaces $\dot{L}_k$ and $\ddot{L}_l$. In the next step the distances $\dot{D}_k(x)$ and $\ddot{D}_l(x)$ for each point x in the volume space are used to generate a *joint distance histogram*. The joint distance histogram represents the joint probability for a point x to have the distance $\dot{D}_k(x)$ to isosurface $\dot{L}_k$ and $\ddot{D}_l(x)$ to isosurface $\ddot{L}_l$. In Figure 4 an example of a joint distance histogram is shown for two identical isosurfaces. In this case, all points x in the volume space have the same distance to $\dot{L}_k$ and $\ddot{L}_l$.

Finally, the mutual information is calculated based on Equation 3. The joint and marginal probabilities for the calculation of the joint and marginal entropies can be directly retrieved from the joint distance histogram. For the example in Figure 4 the mutual information of isosurfaces for the isovalues k and l is maximal since the isosurfaces are identical.

To generate the entire multimodal similarity map, the steps in Figure 4 are repeated for every possible combination of isovalues in both modalities. If we assume that modality 1 has N different isovalues $\dot{V} = \{k_1, ...k_N\}$ and modality 2 has M different isovalues $\ddot{V} = \{l_1, ..., l_M\}$ then the multimodal similarity map can be defined as an $N \times M$ matrix $\text{MSM}(i,j)$. Each entry of the multimodal similarity map represents the similarity between the isosurface $\dot{L}_i$ of modality 1 with the corresponding isovalue i and the isosurface $\ddot{L}_j$ of modality 2 with the corresponding isovalue j. On the right side of Figure 4 the complete multimodal similarity map for the example data sets is shown. The dark line represents the combinations of isovalues i and j which result in identical isosurfaces $\dot{L}_i$ and $\ddot{L}_j$. For all other combinations of isovalues the similarity of their corresponding isosurfaces is lower.

Figures 5 and 6 show the multimodal similarity maps for the synthetic data sets introduced in Section 3. Dark regions denote a high similarity in these figures. For both types of multimodal data, the similarity maps for the combination of data sets without noise and with noise are shown.

For the supplementary data in Figure 5, both data sets contain four squares at the same locations. In the MSM each of the squares is represented by a rectangular area of higher similarity. In Figure 5 corresponding squares and rectangular regions are emphasized by colored frames. The band with the maximum similarity represents the combinations of isovalues k and l at which both data sets represent exactly the same isosurfaces. Due to different value ranges in both data sets this band does not follow the diagonal of the multimodal similarity map. In contrast to self similarity maps, multimodal similarity maps are not symmetrical along the main diagonal. If we investigate the influence of noise in the multimodal similarity map on the right side of Figure 5, it can be seen that the band with the higher similarity is expanded. The expansion of the band gets smaller the higher the data values are. This is due to a higher *SNR* and therefore a smaller impact of the noise on the similarity measurement for higher data values.

In Figure 6 the multimodal similarity maps for our complementary test data sets are shown. The regions in which both data sets contain contradictive information are clearly visible in the similarity map. In contrast to Figure 5, the lower left rectangular area (red frame) in Figure 6 has a considerably lower similarity. Furthermore the band with maximum similarity is missing since there are no isosurfaces for the corresponding isovalues in data set 2 as one square is completely omitted in data set 2. In the same area of the multimodal similarity map with the noisy data set we get a higher variation of similarity values. This is due to the similarity between the square in data set 1 and structures generated by the noise in the background areas of data set 2.

Another interesting area in the multimodal similarity map of Figure 6 is the rectangle in the upper right corner (cyan frame). This rectangular area represents the similarity between the square in one
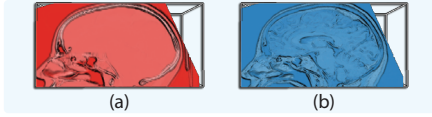
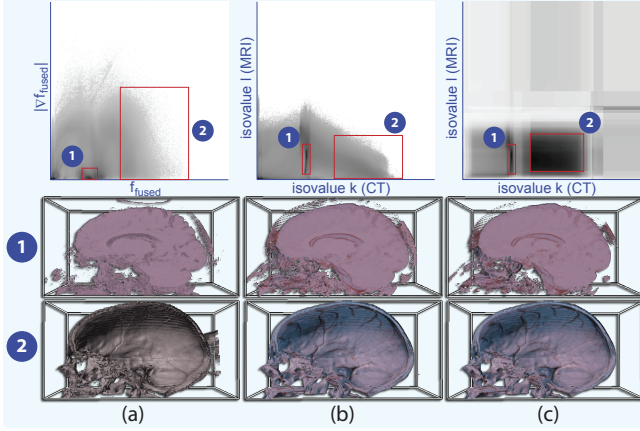Fig. 7. Clipping plane through a (a) CT and (b) MRI scan of a human brain.

Fig. 8. Comparison of (a) a fused transfer function space [15], (b) a dual histogram, and (c) a multimodal similarity map for selection guidance. Region 1 classifies brain tissue and region 2 classifies the cranial bone.

data set and the circle in the other data set. Because of the different shapes of the objects the isosurfaces are similar but not identical. In the similarity map this can be seen by the expanded band of maximum similarity.

5 SIMILARITY-BASED MULTIMODAL VOLUME VISUALIZATION

In this section, we discuss how the additional information provided by multimodal similarity maps can guide the process of exploring and analyzing multimodal volume data. It directs the user towards regions of high similarity or dissimilarity among the two modalities. We first show how salient regions in the similarity map can assist the user in identifying features. Next, we describe a simple approach for providing insight into the spatial differences in a multimodal data set by automatically identifying the most similar isosurfaces in two modalities. Finally, we present a novel approach for similarity-based classification of multimodal volume data.

5.1 Similarity-Based Exploration

Multimodal similarity maps can be used to enable better selection of features in multimodal volume data in a straightforward manner. As the coordinate system of the similarity map is defined by the data values in both modalities, a simple approach is to allow the user to select a region of interest (e.g., by specifying a rectangular selection) and to
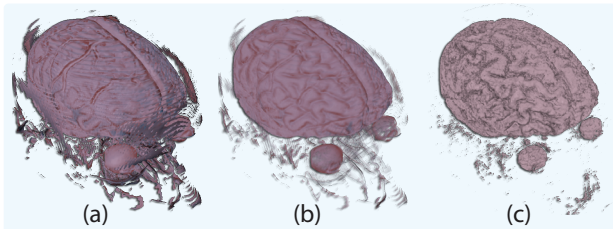

Fig. 9. Selection of brain tissue (a) without similarity weighting, (b) with similarity weighting, and (c) using the method of Haidacher et al. [15] after manually adjusting their δ weighting function to achieve the optimal result.

restrict the visualization to data values which lie within this range. The color and opacity maps are defined separately for each modality.

To demonstrate the advantages of similarity maps over other techniques, we choose a common example: the combination of CT and MRI data. While CT offers a standardized scale for identifying certain types of tissue, MRI provides significantly higher contrast in soft tissue regions. Figure 7 shows (a) a CT and (b) an MRI data set each rendered using a simple linear color map. MRI depicts more details in the brain tissue. But bone, due to its low water content, cannot be distinguished from air. In the CT scan, on the other hand, bone can be clearly identified.

We compare our method with the multimodal transfer function space presented by Haidacher et al. [15] as well as a simple dual histogram. The transfer function space of Haidacher et al. [15] presents a fused data value on the horizontal axis and a fused gradient magnitude on the vertical axis. The fusion is performed based on point-wise mutual information. However, in contrast to our approach this measure is not based on spatial information, but only on the estimated probability of occurrence of a data value combination. In the dual histogram, the frequency of each data value combination is represented by the intensity of the corresponding pixel with darker regions corresponding to higher frequencies. The resulting parameter spaces are depicted in the top row of Figure 8, where (a) shows the fused transfer function space, (b) depicts the dual histogram, and (c) presents the multimodal similarity map. It can be seen that all methods give a salient representation for regions corresponding to brain tissue. However, both the fused transfer function space and the dual histogram fail to give a clear indication of bone as they do not take into account spatial information. The data value ranges corresponding to bone are vastly different in CT and MRI data. Using the multimodal similarity map, on the other hand, a region corresponding to bone can be easily identified. The middle and bottom rows of Figure 8 show 3D visualizations of the data value ranges corresponding to the highlighted selection regions. In the case of the fused transfer function space and the dual histogram, the selection rectangle for bone had to be placed by trial-and-error.

The previous example employed a binary selection in the multimodal similarity map. In order to exploit the information provided by the similarity map we can further use the similarity directly to modulate the opacity of a sample within the selected region:

$$A'(x) = A(x) \frac{\text{MSM}(\dot{f}(x), \ddot{f}(x)) - \text{MSM}_{\text{min}}}{\text{MSM}_{\text{max}} - \text{MSM}_{\text{min}}} \tag{7}$$

where $A'(x)$ is the modulated opacity, $A(x)$ is the original opacity, and $\dot{f}(x)$ and $\ddot{f}(x)$ denote the data value of modality 1 and modality 2, respectively, at a sample position x. MSM_{min} and MSM_{max} are the minimum and maximum similarity values in the selected region. The result of this weighting is an enhancement of similar structures in both modalities while dissimilar structures are suppressed.

Figure 9 illustrates the effect of this weighting. In Figure 9 (a) no weighting is applied, while Figure 9 (b) shows the result obtained with weighting. For comparison, Figure 9 (c) depicts the method of Haidacher et al. [15] after manually adjusting their δ weighting function. The selected regions used for the fusion correspond to the regions for the brain in Figure 8. The results in Figure 9 show that similarity weighting produces comparable results to Haidacher et al. [15] without the necessary user interaction for the adjustment of the δ windowing function.

This example illustrates how multimodal similarity maps can be used to provide assistance in identifying features across multiple modalities. A main advantage over the method of Haidacher et al. [15] is that the original data values can be retained instead of combining them during preprocessing. Furthermore, we do not introduce a new transfer function space which may be unfamiliar to users and difficult to understand.

5.2 Maximum Similarity Isosurfaces

With the multimodal similarity map we gain information about the similarity of certain combinations of isovalues. The multimodal sim-
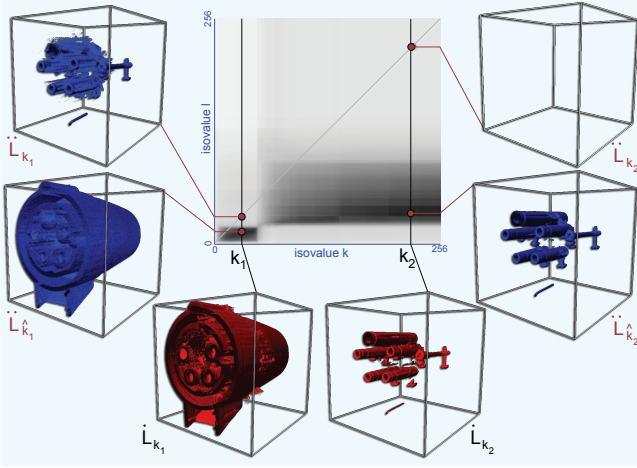
Fig. 10. Maximum similarity isosurface detection for two different iso-values k_1 and k_2. The results in the middle row show the most similar isosurfaces ($\ddot{L}_{\hat{k}_1}$, $\ddot{L}_{\hat{k}_2}$). The results in the top row show the isosurfaces for a naive selection of the isovalues, i.e., in both data sets the same isovalue is chosen.

ilarity maps for the examples in Section 4 have shown that combinations of isovalues for isosurfaces with a high similarity can be identified easily even if their ranges differ significantly. In many applications, such as industrial CT, users want to compare how well the object of interest is depicted in both modalities. Finding the isovalues which best represent the structure of interest in both scans, however, is difficult and requires time-consuming manual tuning.

Using the multimodal similarity map, we can automatically identify the isovalue for the isosurface in one modality which maximizes the similarity to a specific isosurface from another modality. If we assume that a user has specified an isovalue k for an isosurface in one modality, the isovalue $\hat{k}$ with the most similar isosurface in the second modality can be obtained by:

$$\hat{k} = \arg\max_{j} \text{MSM}(k, j) \tag{8}$$

Using this simple approach, it is possible to specify an arbitrary isovalue in either modality and instantly visualize the corresponding isosurfaces from both modalities. A typical setup may depict these isosurfaces side-by-side in linked views enabling the user to quickly identify the spatial differences between two volumetric data sets by browsing through the range of isovalues.

Figure 10 depicts an example for a dual energy CT scan. Due to the different attenuation characteristics for different energy levels, the value ranges in both data sets are different. This can be seen in the multimodal similarity map in the center of Figure 10. The images in the bottom row show isosurfaces for isovalues k_1 and k_2 in modality 1. The top row shows the isosurfaces for the same isovalues in modality 2. The middle row shows the isosurfaces in modality 2 for the isovalues $\hat{k}_1$ and $\hat{k}_2$ with the maximum similarity to k_1 and k_2. The isosurfaces for $\hat{k}_1$ and $\hat{k}_2$ match the isosurfaces in modality 1 much better than the isosurfaces for the naive selection of isovalues.

5.3 Similarity-Based Classification

Selection of simple regions in the multimodal similarity map, as described in the Section 5.1, allows quick exploration of multimodal data. This approach can be useful when only few specific features are of interest. For generating more complex visualizations, which depict multiple volumetric structures and take advantage of the additional information provided by multiple modalities, classification in the joint data space is necessary. The multimodal similarity map also opens up new avenues to assist in this process. Our idea is to use a nearest neighbor classifier in similarity space to determine the optical properties of

a sample. Intuitively, instead of trying to relate the two modalities in terms of their data values, we instead classify samples, i.e., combinations of data values from both modalities, according to their similarity to a set of user-specified isosurfaces from both modalities.

We assume two continuous three-dimensional scalar fields $\dot{f}, \ddot{f} : \mathbb{R}^3 \to \mathbb{R}$ which represent two co-registered input volumes. For multimodal volume visualization, we assign a color and opacity to every point $x \in \mathbb{R}^3$ in space based on the value of these functions. Our method takes as input a set of isovalue pairs $h_i = (\dot{h}_i, \ddot{h}_i)$ where $\dot{h}_i, \ddot{h}_i$ correspond to isovalues of $\dot{f}$ and $\ddot{f}$ respectively. Each pair of isovalues has an assigned color c_i, opacity α_i, and optional weight w_i.

For two data values $k \in \dot{f}$ and $l \in \ddot{f}$, we evaluate their multimodal similarity to the i-th isovalue pair in the following manner:

$$\dot{s}_i(k) = \text{MSM}(k, \ddot{h}_i) \quad \ddot{s}_i(l) = \text{MSM}(\dot{h}_i, l) \tag{9}$$

where MSM is the multimodal similarity map. This means that $\dot{s}_i$ is the similarity of the isosurface k of $\dot{f}$ and the isosurface $\ddot{h}_i$ of $\ddot{f}$ and $\ddot{s}_i$ is the similarity of the isosurface l of $\ddot{f}$ and the isosurface $\dot{h}_i$ of $\dot{f}$.

Based on the similarities $\dot{s}_i$ and $\ddot{s}_i$ we can now define a combined measure s_i of similarity between h_i and the two isovalues $k \in \dot{f}$ and $l \in \ddot{f}$ in multimodal similarity space:

$$s_i(k, l) = \dot{s}_i(k)\ddot{s}_i(l) \tag{10}$$

The rationale behind this choice is that we interpret the similarities $\dot{s}_i(k)$, $\ddot{s}_i(l)$ as independent probabilities of k being similar to $\ddot{h}_i$ and l being similar to $\dot{h}_i$. Thus, the joint probability of (k, l) being similar to h_i is the product $\dot{s}_i(k)\ddot{s}_i(l)$. Alternatively, we could consider $\dot{s}_i$ and $\ddot{s}_i$ as the membership functions of two fuzzy sets and s_i as the membership function of their intersection. In this case, another possible definition would be $s_i(k, l) = \min(\dot{s}_i(k), \ddot{s}_i(l))$ [41]. In our experiments, we found that both approaches lead to similar results.

Having defined a measure of closeness between two points in similarity space, we now let each pair of isovalues h_i to determine the optical properties of points that are closer to h_i than to any other isovalue pair h_j ($i \neq j$). This means a pair of data values (k, l) with $k \in \dot{f}$, $l \in \ddot{f}$ will assume the color and opacity of the isovalue pair $h_{m(k,l)}$ which maximizes $s_i(k, l)$:

$$m(k, l) = \arg\max_{i} s_i(k, l)w_i \tag{11}$$

where w_i is a weight which allows additional control over the influence of the isovalue pair h_i. During rendering, we can now evaluate this maximum for every sample location $x \in \mathbb{R}^3$ in space:

$$m_x = m(\dot{f}(x), \ddot{f}(x)) \tag{12}$$

Thus, m_x denotes the index of the isovalue pair which maximizes the similarity to the data value $\dot{f}(x)$, $\ddot{f}(x)$ at the sample location x. To visually encode the similarity of the sample to h_{m_x}, we additionally weight the sample opacity based on the similarity s_{m_x}. The color $C(x)$ and opacity $A(x)$ at the sample position x are then simply:

$$C(x) = c_{m_x} \quad A(x) = \alpha_{m_x} s_{m_x} \tag{13}$$

In practice, in order to obtain crisp boundaries, it is convenient to define an additional threshold t which specifies the minimum similarity of a sample with any of the isovalue pairs in order to be visible. If $s_{m_x} < t$, the sample is considered to be fully transparent.

In volume rendering, it is common to evaluate a local illumination model using the normalized gradient of the scalar field as the normal vector. To enable volume shading, we can combine the gradient information of both modalities using a similarity-based weighting:

$$g(x) = \frac{\dot{s}_{m_x}(\dot{f}(x))\nabla \dot{f}(x) + \ddot{s}_{m_x}(\ddot{f}(x))\nabla \ddot{f}(x)}{\dot{s}_{m_x}(\dot{f}(x)) + \ddot{s}_{m_x}(\ddot{f}(x))} \tag{14}$$
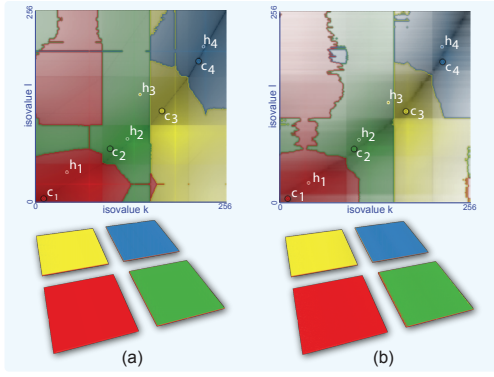
Fig. 11. Classification based on multimodal surface similarity for supplementary data (a) without noise and (b) with added noise.
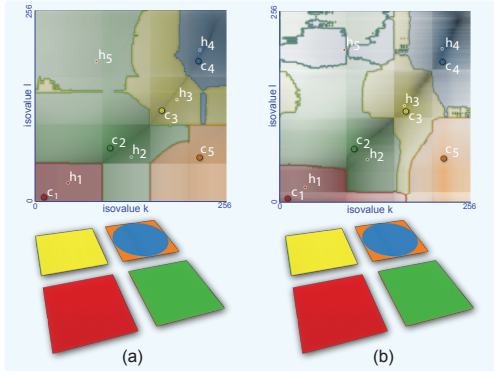

Fig. 12. Classification based on multimodal surface similarity for complementary data (a) without noise and (b) with added noise.

5.4 Classification Specification

The described classification is equivalent to a generalized Voronoi decomposition of similarity space, i.e., using non-Euclidean distances defined by our similarity measure. Every sample, which is a pair of values from the two modalities, is assigned to the most similar isovalue pair which determines its color and opacity. We also visualize this classification on the similarity map itself by simply evaluating Equation 11 for each location, i.e., each combination of data values, in the similarity map and coloring the corresponding pixel accordingly. When depicted on the two-dimensional similarity map, where the coordinate system is defined by data values, these regions may be disconnected and non-convex (see, for example, Figures 11 and 12 which are discussed in detail below). Furthermore, based on the structure of the similarity map, the site, i.e., the isovalue pair that defines a region may not be contained within this region. While this may initially sound counter-intuitive, the following situation exemplifies such a case: Assume two isosurfaces for $\dot{h}_i$ and $\ddot{h}_i$ which are highly dissimilar. There will likely be other isosurfaces they are more similar to than to each other.

To provide an additional means for manipulating the classification regions instead of directly modifying the isovalues themselves, we define a user-specified control point $c_i = (\dot{c}_i, \ddot{c}_i)$ for each isovalue pair h_i, which can be freely moved. h_i is initialized with c_i and is then used to compute the similarity-weighted centroid of the region it defines. The isovalue pair h_i is then moved to the position of the centroid:

$$h_i = \frac{\sum\limits_{(k,l) \in R(c_i)} (k,l) s_{m(k,l)}(k,l)}{\sum\limits_{(k,l) \in R(c_i)} s_{m(k,l)}(k,l)} \qquad (15)$$

where $R(c_i) = \{(k,l) | m(k,l) = i\}$ is the similarity-space region as-

signed to c_i. This essentially corresponds to one iteration of Lloyd's algorithm [28]. Note, however, that we do not perform the full relaxation since our goal is not to perform a full centroidal decomposition of the similarity space. Instead, our aim is for regions to follow their control points.

Based on this approach, we developed a simple user interface for similarity-based classification of multimodal volume data. The user is presented with the multimodal similarity map and can interactively add and remove control points, move them on the similarity map, and change their colors and opacities. When moving control points they behave similar to well-known "magic wand"-type selection tools – the regions they define snap to clusters in the similarity map. Slightly modifying a control point will, in accordance with the structure of the similarity map, not cause major changes of the classification result.

Figures 11 and 12 show examples of our classification approach using the previously introduced synthetic data sets. The colored regions encode the nearest neighbors, in similarity space, of each of the white-outlined points for isovalue pairs h_i in the corresponding color. This means a point on the similarity map is assigned to a region if it is more similar to the corresponding isovalue pair than to any other one. This also means that during volume rendering a sample with the corresponding value combination will be assigned the respective color and opacity (see Equation 13). The control points c_i are depicted as the slightly larger points with dark outlines. It can be seen that in regions of high similarity the control points c_i will be close to the corresponding isovalue pairs h_i, but in other areas this is not necessarily the case. Figure 11 (a) illustrates that our approach is successful in identifying the correspondences between both data sets. By placing control points along the band of maximum similarity, the resulting regions will subdivide the map such that all rectangles in the data are assigned the user-specified color even though their data value ranges vary. Small perturbations in the placement of the control points leave the resulting classification unaffected. Adding noise to one of the data sets has little effect on the visual result, as shown in Figure 11 (b). Figure 12 shows that this approach makes it possible to easily exploit complementary information from both data sets. The red square, which is only present in one data set, as well as the blue circle and orange square can all be separated. This enables the generation of a combined visualization which contains all these features. Using a conventional approach, where the user defines a 1D transfer function for each modality and the results are blended, it is much more difficult to separate the features, since their corresponding data value ranges significantly vary. The additional weights w_i used in Equation 11 therefore allow to control the sizes of the respective regions and can be interactively modified. In Figures 11 and 12 all weights w_i are set to one.

6 RESULTS

An application for which our similarity-based classification approach is particularly suitable is the study of industrial parts where the goal is to detect manufacturing defects. Dual energy CT is of interest in such scenarios, since different materials can cause scanning artifacts at certain energy levels. The low energy scan typically has high precision but is affected by severe artifacts, while the high energy scan is nearly artifact-free but suffers from reduced precision and noise. It is desirable to combine the advantages of both energy levels, i.e., to generate a visualization which uses the global structure from the high energy scan to remove the artifacts from the low energy scan while preserving subtle details. Heinzl et al. [16] presented a processing pipeline for extracting surfaces from dual energy CT scans. With our similarity-based classification approach it is now possible to directly visualize structures which exhibit high surface similarity between both modalities. An example is given in Figure 13. Figure 13 (a) shows the low energy scan of a 400 Volt power connector rendered using a conventional 1D transfer function. It is not possible to find an opacity setting which suppresses all artifacts but leaves the surface intact. In Figure 13 (b) the corresponding high energy scan, also rendered using a 1D transfer function, is shown. This result gives a better impression of the actual surface, but is noisy and lacks details. Using our similarity-based classification approach, as shown in Figure 13 (c),
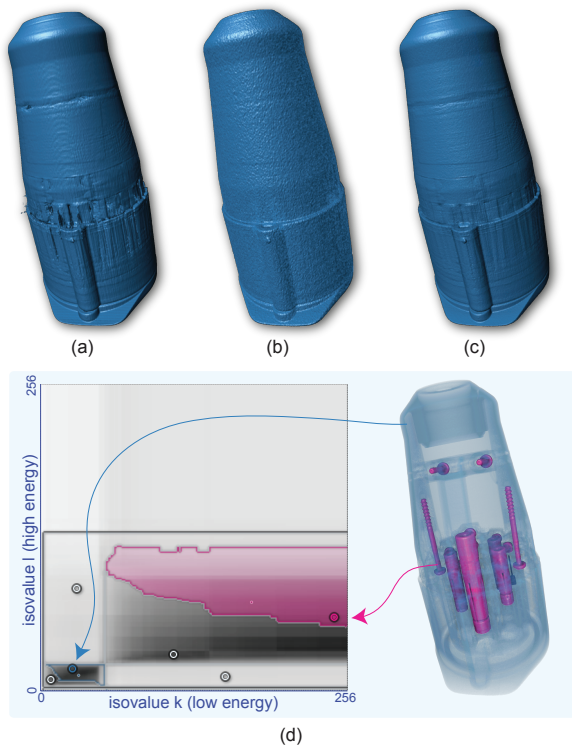
Fig. 13. Similarity-based fusion of a dual energy CT scan of a power connector. The low-energy scan (a) and the high-energy scan (b) provide supplementary information which can be used to remove most of their respective drawbacks, as shown the similarity-based classification (c). The corresponding similarity map (d) shows the placement of the control points. In the image to its right the opacity of the outer surface has been reduced to reveal the interior parts of the connector.

we can remove the artifacts by choosing control points which select regions of high dissimilarity and setting their opacity to zero. The feature of interest, the outer surface of the connector, is similar in both data sets. Since the opacity of a sample is based on the global surface similarity of its data values to the isosurface pair, holes and artifacts in the low energy scan can be remedied using information from the high energy scan. Figure 13 (d) shows the similarity map together with the specified control points. The image to the right of the similarity map uses the same control points as Figure 13 (c), but the opacity of the outer surface has been lowered to reveal the interior parts of the connector.

Furthermore, our approach can be used to assist the classification of ambiguous structures. One example is CT angiography, where it is desired to clearly separate contrast-enhanced blood vessels from bone. In a single modality scan this is typically not possible as the data values of the contrast agent partially overlap with lower-density bone regions and cartilage. This is illustrated in Figure 14 (a), where it was attempted to specify different colors for bones and vessels using a 1D transfer function on a low energy CT scan of the lower extremities. While it is also not possible to achieve this separation using a high-energy scan, as shown in Figure 14 (b), it can be seen that the classified structures are slightly different in both modalities. Using similarity-based classification, we can therefore achieve a better separation, as depicted in Figure 14 (c). The corresponding similarity map is shown in Figure 14 (d). In order to illustrate how regions in the similarity map correspond to structures of different ossification levels, they have been assigned different colors in the rightmost image. Vessels are orange and different bone structures are white, blue, and green.

A further example is shown in Figure 1. In this case, a dual energy CT data set of a human head is used. The similarity map, shown on the bottom right, provides good guidance for iteratively selecting the

individual tissues numbered from 1 to 7. The information provided by the two energy levels is sufficient to allow differentiation between bone (selected in step 3), major vessels (step 4), and minor vessels (step 5).

7 IMPLEMENTATION

The calculation of the multimodal similarity map is a preprocessing step which is implemented in C++ and runs on the CPU. It has to be performed only once for a single multimodal data set. After the preprocessing step the multimodal similarity map is simply represented as a two-dimensional image. During rendering, the similarity of a combination of isovalues from the two modalities can be retrieved by a single lookup in a 2D texture.

The user interface for our similarity-based classification approach was implemented using the Qt toolkit. The user-interface widget generates a set of isovalue pairs, colors, and weights, which are passed to a GPU-based volume renderer implemented in GLSL. In the shader, the similarity between the data values at the current sample point and each isovalue pair is determined using two texture lookups (see Equations 9 and 10) and the maximum is computed. While this is more expensive than a conventional transfer function lookup, the additional costs are limited due to the fact than only few control points will be required in many applications. In our implementation, for a typical number of five control points, the average render time increases by a factor of approximately 1.4 compared to a single conventional transfer-function lookup. The color and opacity of the maximally similar isovalue pair then determines the color and opacity of the current sample, as described in Section 5.3. The gradient vectors of both modalities are computed by central differences, combined with Equation 14, and used to evaluate a local illumination model if shading is enabled.

8 DISCUSSION

As shown in our examples, multimodal surface similarity can provide a useful tool for the visual analysis of multimodal volume data. However, isosurface similarity as a measure is only useful in cases where there is some correspondence between features and isosurfaces. For example, in data where textures or patterns are of central importance, isosurface similarity will likely fail to provide valuable insights. While this is a clear limitation of our approach, we want to emphasize that also the lack of distinct structures in a multimodal similarity map provides additional information to the user. As our approach deliberately avoids to position itself as a new technique central to the visualization process, the lack of distinct features (like the lack of distinct features in a histogram) simply means that little additional guidance can be provided for the particular data set. However, in our experiments we found that even for challenging data combinations, such as CT and PET, which exhibit little correspondence, multimodal surface similarity is still able to assist in finding joint data value ranges which correspond to joint structures of interest.

The computation time for the multimodal similarity map of two data sets is approximately twice the computation time of a self similarity map for a data set of the same size. This is due to the lack of symmetry. As reported by Bruckner and Möller [3], a feasible strategy to limit the duration of this pre-processing step is to use downsampled versions of the distance transforms (which are computed at the original data set resolution) for the mutual information computation. The computation times for all data sets used in this paper are given in Table 1. The second column in the table lists the downsampling rate for the respective volume which is automatically chosen to limit the computation time to approximately one minute. Even though downsampling is performed quite aggressively, a distance field is a rather redundant representation and the downsampled version essentially acts as a shape descriptor and is not used for precise spatial measurements. To the results of Bruckner and Möller [3] we can also add information about additional experiments on the effects of quantization in the value domain. We found that for real-world data a quantization to 8 bits results in practically no structural differences in the similarity map, as exemplified in Figure 15.
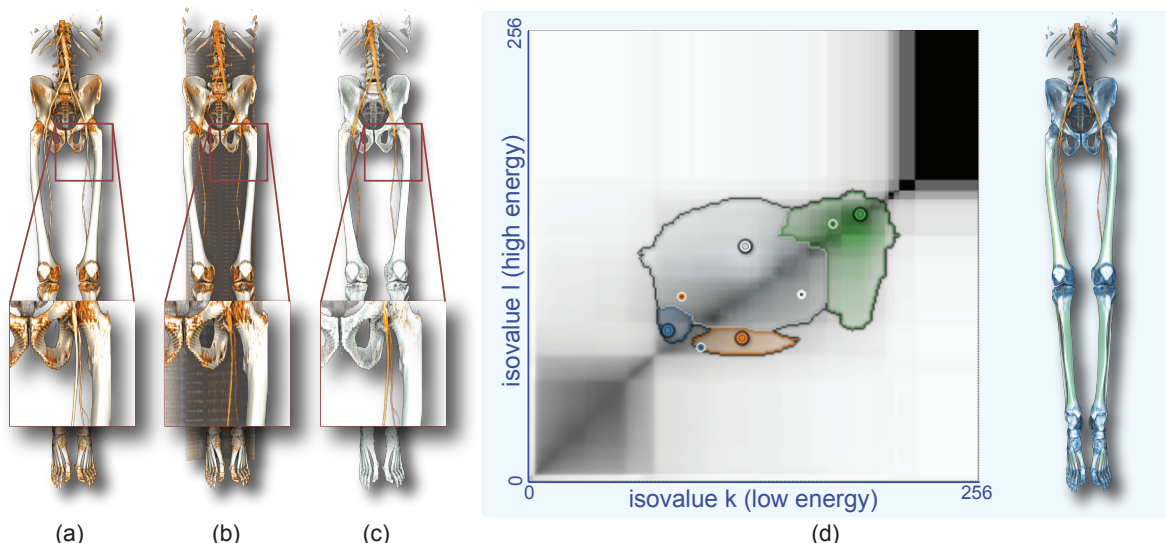
Fig. 14. Similarity-based classification of blood vessels in a dual energy CT angiography scan of the lower extremities. When using only the information from the low energy scan (a) or the high energy scan (b), it is not possible to separate blood vessels and bones. Using multimodal similarity (c) this can be achieved. The corresponding control points are shown on the similarity map (d) where different colors have been assigned to vessels and bones of different densities.

Table 1. Computation times for the multimodal similarity maps shown in the paper as measured in an Intel Core i7 950 CPU with a clock rate of 3.07 GHz and 12 GB RAM. The first column reports the data type and size. The second column gives the downsampling rate for the distance fields. The last column gives the total computation time for the similarity maps which is the sum of the computation times for all distance transforms (third column) and the mutual information of all isovalue combinations (fourth column).

Data Set	Downsample Rate	Distance Transform	Mutual Information	Total
Supplementary $512 \times 512 \times 6$	2	28.57s	32.97s	61.54s
Complementary $512 \times 512 \times 6$	2	21.17s	32.41s	53.58s
CT-MRI $256 \times 256 \times 128$	4	4.64s	32.86s	37.50s
Industrial DECT $425 \times 551 \times 895$	16	2.12s	12.08s	14.20s
Head DECT $512 \times 512 \times 575$	16	2.62s	11.30s	13.92s
Extremities DECT $512 \times 512 \times 855$	16	2.86s	14.55s	17.41s

Another limitation of our work is that the described approach only considers data sets consisting of two modalities. While this applies to many application scenarios, a solution for a larger number of modalities would be desirable. A multi-dimensional similarity map of similarities between all isovalue combinations of the respective data sets, however, would be computationally infeasible. A potential solution could be to only consider pair-wise similarities between the individual modalities resulting in a matrix of multimodal similarity maps. The investigation of whether such an approach is effective is an interesting topic for future research. Furthermore, our technique could also be applied to investigate time-dependent data by generating a set of similarity maps between subsequent time steps. Temporal similarity maps could help to identify stable features and to pinpoint discontinuities.

A further limitation of similarity maps in general is that they do not contain frequency information, i.e., small structures which exhibit high similarity receive the same prominence as very large regions with a similar degree of similarity. This can be regarded as an advantage
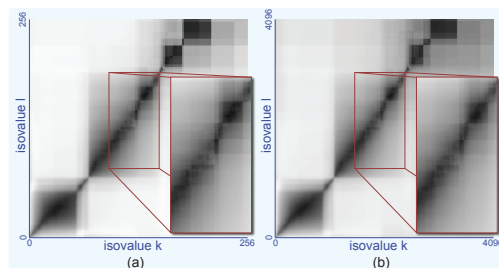

Fig. 15. A comparison between the multimodal similarity map with a isovalue precision of 8 bits (a) and 12 bits (b) for the data set shown in Figure 1.

with respect to histograms, where large regions tend to dominate and logarithmic scaling is typically required. It can also be a drawback since data value combinations which do not occur at all are not clearly indicated. Ideally, a combination of both types of information would be desired, but identifying a good visual encoding for this purpose is not straightforward and remains an area of future research.

9 CONCLUSION

In this paper, we introduced multimodal surface similarity maps as a tool for the investigation of multimodal volume data sets. The multimodal similarity map provides an overview of the differences and similarities between the isosurfaces of two modalities in a compact manner. The analysis of parameter spaces is an increasingly important topic for knowledge discovery in scientific data. Our approach showed that spatial similarity information can assist the visualization process by guiding the selection of features. By exploiting similarity information, we introduced a novel way for the interactive classification and visualization of multimodal volume data.

REFERENCES

[1] H. Akiba and K.-L. Ma. A tri-space visualization interface for analyzing time-varying multivariate volume data. In *Proceedings of EuroVis 2007*, pages 115–122, 2007.

[2] J. F. Barrett and K. Nicholas. Artifacts in CT: Recognition and avoidance. *Radiographics*, 24(6):1679–1691, 2004.

[3] S. Bruckner and T. Möller. Isosurface similarity maps. *Computer Graphics Forum*, 29(3):773–782, 2010.

[4] W. Cai and G. Sakas. Data intermixing and multi-volume rendering. In *Computer Graphics Forum*, volume 18, pages 359–368, 1999.

[5] H. Carr, B. Duffy, and B. Denby. On histograms and isosurface statistics. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1259–1265, 2006.

[6] H. Carr, J. Snoeyink, and M. van de Panne. Flexible isosurfaces: Simplifying and displaying scalar topology using the contour tree. *Computational Geometry: Theory and Applications*, 43(1):42–58, 2010.

[7] M. Chen and H. Jänicke. An information-theoretic framework for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1206–1215, 2010.

[8] M. Chen and J. V. Tucker. Constructive volume geometry. *Computer Graphics Forum*, 19:281–293, 2000.

[9] C. Eusemann, D. R. Holmes III, B. Schmidt, T. G. Flohr, R. Robb, C. McCollough, D. M. Hough, J. E. Huprich, M. Wittmer, H. Siddiki, and J. G. Fletcher. Dual energy CT: How to best blend both energies in one fused image? In *Proceedings of SPIE Medical Imaging 2008*, pages 1–8, 2008.

[10] A. Evans, S. Marrett, J. Torrescorzo, S. Ku, and L. Collins. MRI-PET correlation in three dimensions using a volume-of-interest (VOI) atlas. *Journal of Cerebral Blood Flow and Metabolism*, 11(2):A69–A78, 1991.

[11] S. Fang and R. Srinivasan. Volumetric-CSG - a model-based volume visualization approach. In *Proceedings of the 6th International Conference in Central Europe on Computer Graphics and Visualization*, pages 88–95, 1998.

[12] M. Feixas, M. Sbert, and F. González. A unified information-theoretic framework for viewpoint selection and mesh saliency. *ACM Transactions on Graphics*, 6:1–23, 2009.

[13] M. Ferre, A. Puig, and D. Tost. A framework for fusion methods and rendering techniques of multimodal volume data: Research articles. *Computer Animation and Virtual Worlds*, 15:63–77, 2004.

[14] R. Fuchs and H. Hauser. Visualization of multi-variate scientific data. *Computer Graphics Forum*, 28(6):1670–1690, 2009.

[15] M. Haidacher, S. Bruckner, A. Kanitsar, and M. E. Gröller. Information-based transfer functions for multimodal visualization. In *Proceedings of Visual Computing for Biomedicine 2008*, pages 101–108, 2008.

[16] C. Heinzl, J. Kastner, and M. E. Gröller. Surface extraction from multi-material components for metrology using dual energy CT. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1520–1527, 2007.

[17] H. Hong, J. Bae, H. Kye, and Y.-G. Shin. Efficient multimodality volume fusion using graphics hardware. In *Proceedings of the International Conference on Computational Science 2005*, pages 842–845, 2005.

[18] J. Hsieh. *Computed Tomography: Principles, Design, Artifacts and Recent Advances*. SPIE Press, 2003.

[19] X. Huang, N. Paragios, and D. Metaxas. Shape registration in implicit spaces using information theory and free form deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28:1303–1318, 2006.

[20] M. Jones, J. Baerentzen, and M. Šrámek. 3D distance fields: a survey of techniques and applications. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):581–599, 2006.

[21] M. Khoury and R. Wenger. On the fractal dimension of isosurfaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1198–1205, 2010.

[22] J. Kim, S. Eberl, and D. Feng. Visualizing dual-modality rendered volumes using a dual-lookup table transfer function. *Computing in Science and Engineering*, 9(1):20–25, 2007.

[23] G. Kindlmann and J. W. Durkin. Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of the IEEE Symposium on Volume Visualization 1998*, pages 79–86, 1998.

[24] J. Kniss, C. Hansen, M. Grenier, and T. Robinson. Volume rendering multivariate data to visualize meteorological simulations: a case study. In *Proceedings of VisSym 2002*, pages 189–195, 2002.

[25] J. Kniss, J. P. Schulze, U. Wössner, P. Winkler, U. Lang, and C. Hansen. Medical applications of multi-field volume rendering and VR techniques. In *Proceedings of VisSym 2004*, pages 249–254, 2004.

[26] T. Kvålseth. Entropy and correlation: Some comments. *IEEE Transactions on Systems, Man, and Cybernetics*, 17:517–519, 1987.

[27] D. Levin, X. Hu, K. Tan, S. Galhotra, C. Pelizzari, G. Chen, R. Beck, C. Chen, M. Cooper, and J. Mullan. The brain: integrated three-dimensional display of MR and PET images. *Radiology*, 172:783–789, 1989.

[28] S. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.

[29] M. E. Noz, G. Q. Maguire, M. P. Zeleznik, E. L. Kramer, F. Mahmoud, and J. Crafoord. A versatile functional/anatomic image fusion method for volume data sets. *Journal of Medical Systems*, 25(5):297–307, 2001.

[30] C. Scheidegger, J. Schreiner, B. Duffy, H. Carr, and C. Silva. Revisiting histograms and isosurface statistics. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1659–1666, 2008.

[31] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423,623–656, 1948.

[32] S. Takahashi, Y. Takeshima, I. Fujishiro, and G. Nielson. Emphasizing isosurface embeddings in direct volume rendering. In G. Bonneau, T. Ertl, and G. Nielson, editors, *Scientific Visualization: The Visual Extraction of Knowledge from Data*, pages 185–206. Springer, 2006.

[33] I. Viola, M. Feixas, M. Sbert, and M. E. Gröller. Importance-driven focus of attention. *IEEE Transactions on Visualization and Computer Graphics*, 12:933–940, 2006.

[34] P. Šereda, A. V. Bartrolí, I. W. Serlie, and F. A. Gerritsen. Visualization of boundaries in volumetric data sets using LH histograms. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):208–218, 2006.

[35] C. Wang and H.-W. Shen. Information theory in scientific visualization. *Entropy*, 13(1):254–273, 2011.

[36] S. W. Wang and A. E. Kaufman. Volume-sampled 3D modeling. *IEEE Computer Graphics and Applications*, 14(5):26–32, 1994.

[37] Z. Wang, D. Ziou, C. Armenakis, D. Li, , and Q. Li. A comparative analysis of image fusion methods. *IEEE Transactions on Geoscience and Remote Sensing*, 43:1391–1402, 2005.

[38] W. Wells III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical image analysis*, 1(1):35–51, 1996.

[39] L. Xu, T.-Y. Lee, and H.-W. Shen. An information-theoretic framework for flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1216–1224, 2010.

[40] Y. Yao. Information-theoretic measures for knowledge discovery and data mining. In Karmeshu, editor, *Entropy Measures, Maximum Entropy Principle and Emerging Application*, pages 115–136. Springer, 2003.

[41] L. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1995.

[42] K. J. Zuiderveld and M. A. Viergever. Multi-modal volume visualization using object-oriented methods. In *Proceedings of the IEEE Symposium on Volume Visualization 1994*, pages 59–66, 1994.