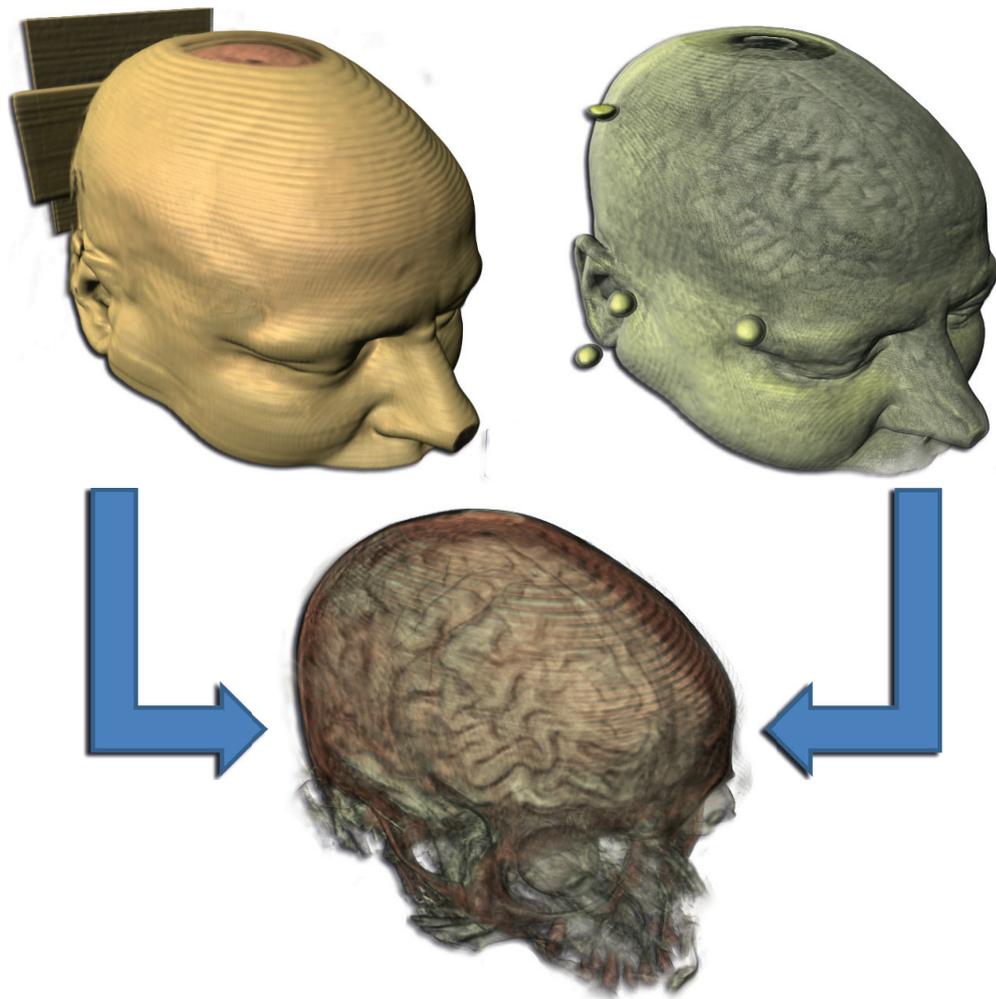


Information-based Feature Enhancement in Scientific Visualization

DISSERTATION



Martin Haidacher

Institute of Computer Graphics and Algorithms

Vienna University of Technology, Austria

martin.haidacher@cg.tuwien.ac.at

Abstract

Scientific visualization is a research area which gives insight into volumetric data acquired through measurement or simulation. The visualization allows a faster and more intuitive exploration of the data.

Due to the rapid development in hardware for the measurement and simulation of scientific data, the size and complexity of data is constantly increasing. This has the benefit that it is possible to get a more accurate insight into the measured or simulated phenomena. A drawback of the increasing data size and complexity is the problem of generating an expressive representation of the data.

Since only certain parts of the data are necessary to make a decision, it is possible to mask parts of the data along the visualization pipeline to enhance only those parts which are important in the visualization. For the masking various properties are extracted from the data which are used to classify a part as important or not. In general a transfer function is used for this classification process which has to be designed by the user.

In this thesis three novel approaches are presented which use methods from information theory and statistics to enhance features from the data in the classification process that are important for a certain task. With the tools of information theory and statistics it is possible to extract properties from the data which are able to classify different materials or tissues in the data better than comparable other approaches.

One approach adaptively extracts statistical properties, i.e. the mean value and the standard deviation, of the data values in the local neighborhood of each point in the data set. With these statistical properties it is possible to better distinguish between different materials in a data set even though the data is very noisy.

The other two approaches in this thesis employ methods from information theory to extract features from multimodal data sets. Thus it is possible to enhance features of the data which are either very similar or very dissimilar in both modalities. Through information theory the variations in the value ranges of both modalities do not influence the classification of these features.

All three approaches define novel transfer-function spaces which simplify the design process of a transfer function for the user. Different features of the data, such as different materials, can be clearly depicted in these spaces. Therefore, it is easier for a user to design a transfer function which enhances the features of importance for a certain task.

For each of the new approaches results and comparisons to other existing techniques are shown to highlight the usefulness of the proposed methods. Through the described research it is shown that information theory and statistics are tools which are able to extract expressive properties from the data.

In the introduction a broad overview over scientific visualization and the visualization pipeline is given. The classification process is described in more detail. Since information theory and statistics play an important role for all three approaches, a brief introduction to these concepts is given as well.

Kurzfassung

Wissenschaftliche Visualisierung ist ein Forschungsgebiet, das Einblick in gemessene oder simulierte volumetrischen Daten gibt. Die Visualisierung ermöglicht eine schnellere und intuitivere Erforschung der Daten.

Durch die rasante Entwicklung der Hardware für die Messung und Simulation von wissenschaftlichen Daten nimmt die Größe und Komplexität der Daten ständig zu. Dies hat den Vorteil, dass es möglich ist einen genaueren Einblick in die gemessenen oder simulierten Phänomene zu erhalten. Jedoch wird es zunehmend schwieriger, eine geeignete Darstellung für diese Daten zu finden.

Da in den meisten Fällen nur bestimmte Teile der Daten erforderlich sind, um eine Entscheidung zu treffen, können Teile der Daten verworfen werden, welche für einen bestimmten Anwendungsfall nicht erforderlich sind. Diese Klassifizierung in wichtige und unwichtige Teile der Daten erfolgt durch eine Transfer-Funktion. Die Transfer-Funktion beschreibt eine Abbildung von bestimmten Ausprägungen der Daten auf optische Eigenschaften.

In dieser Arbeit werden drei neue Ansätze vorgestellt, die Techniken aus der Informationstheorie und Statistik verwenden, um Eigenschaften aus den Daten für die Klassifizierung zu extrahieren. Durch die Informationstheorie und Statistik ist es möglich, Eigenschaften zu berechnen, welche verschiedene Materialien in den Daten besser unterscheidbar machen als bestehende Methoden.

Ein Ansatz extrahiert statistische Eigenschaften - wie den Mittelwert und die Standardabweichung - aus einer lokalen Umgebung um jeden Punkt in den Daten auf eine adaptive Weise. Durch die statistischen Eigenschaften ist es möglich, verschiedene Materialien besser zu unterscheiden, auch wenn die Daten sehr verrauscht sind.

Die beiden anderen Ansätze verwenden Methoden aus der Informationstheorie, um Merkmale aus multimodalen Daten zu extrahieren. Dadurch können Eigenschaften in den Daten hervorgehoben werden die entweder in beiden Modalitäten sehr ähnlich oder sehr unterschiedlich sind. Durch die Informationstheorie hat der Wertebereich der Daten in beiden Modalitäten keinen Einfluss auf die Klassifizierung dieser Merkmale.

Alle drei in dieser Arbeit vorgestellten Ansätze definieren Transfer-Funktions Räume die das Design einer Transfer-Funktion für den Benutzer erleichtern. Dies ist möglich, da unterschiedliche Eigenschaften der Daten - wie etwa verschiedene Materialien - in diesen Räumen klar unterscheidbar sind. Daher ist es für den Benutzer einfacher, bestimmte Teile der Daten hervorzuheben, welche für eine bestimmte Aufgabe benötigt werden.

In der Arbeit werden für jede neue Technik Ergebnisse und Vergleiche zu existierenden Methoden gezeigt, um den Nutzen der Techniken hervorzuheben. Durch die Forschung in diesem Bereich wurde bewiesen, dass die Informationstheorie und Statistik in der Lage ist, aussagekräftige Eigenschaften aus den Daten zu extrahieren.

Die Einleitung bietet einen Überblick über die wissenschaftliche Visualisierung und die Visualisierung-Pipeline. Die Klassifizierung wird im folgenden näher beschrieben. Da Informationstheorie und Statistik eine wichtige Rolle für alle drei Methoden spielen, wird die Einleitung mit einer kurzen Einführung in diese Theorie abgeschlossen.



Contents

Preface	vii
1 Introduction	1
1.1 Scientific Visualization	2
1.2 The Visualization Pipeline	3
1.3 Classification	6
1.4 Information Theory in Visualization	11
2 Volume Classification based on Statistical Transfer-Function Spaces	18
2.1 Related Work	19
2.2 Statistical Transfer-Function Space	20
2.3 Implementation	29
2.4 Results and Discussion	32
2.5 Conclusion	37
3 Information-based Transfer Functions for Multimodal Classification	38
3.1 Related Work	39
3.2 Information-based Transfer Functions for Multimodal Volume Classification . .	41
3.3 Implementation	46
3.4 Results	47
3.5 Conclusion and Discussion	50
4 Volume Classification Using Multimodal Surface Similarity	52
4.1 Related Work	53
4.2 Synthetic Multimodal Data Sets	54
4.3 Multimodal Surface Similarity	56
4.4 Similarity-Based Volume Fusion	59
4.5 Implementation	69
4.6 Discussion	69
4.7 Conclusion	71
5 Summary	75

Contents

vi

Bibliography

77

Nearly every man who develops an idea works it up to the point where it looks impossible, and then he gets discouraged. That's not the place to become discouraged.

— Thomas A. Edison



.....

Preface

THE start of every new invention is an idea. An idea is nothing which can be generated on command. It just appears every now and then. It is something which starts its existence in someone's mind. Unfortunately most of the ideas instantly stop existing since they cannot prove their usefulness in a brief mental review. But a few ideas make it to the next stage of their existence. In this stage an idea has to survive an internal stress test. In this stress test the idea has to face various scenarios and confirm its value. After this stage the idea is ready to be presented to other people and further investigated.

In the last three and half years while my PhD studies I had many ideas. A few of these ideas have evolved so far that they were worth of being investigated closer. Through these investigations some of the ideas turned out to be so good that they finally ended up in this thesis as novel approaches in scientific visualization.

Even though an idea arises in a single mind it is highly influenced by everybody and everything someone gets in contact. These influences are also important in a later stage of the idea when it is discussed with others or further investigated. For the ideas behind the approaches in this thesis I would have to thank pretty much everyone I ever met in my life. But there are a couple of people who deserve special credits since their influence had much more impact in the ideas and research activities which led to the presented results.

The person I have to thank most is my supervisor Meister Eduard Gröller. He was always able to show me new directions when I thought that I reached a dead end in my research. It is also due to him that he constantly caters for a nice environment in the visualization group. It is and was always fun to work there.

Due to different projects I was able to work together with different people. This gave me the possibility to gain insight into different operations and environments (Harvard University and King Abdullah University of Science and Technology). Therefore I want to thank Markus Hadwiger and Armin Kanitsar.

Another special thank goes to Stefan Bruckner. He helped me out several times no matter if I had a scientific question or just some technical problems. The Volumeshop which was developed and is maintained by Stefan made it much easier and faster to implement all methods.

Finally, I want to thank all the co-authors for the papers which I submitted during my PhD studies, all my colleagues, and also the students which I was supervising. With all of them I had many fruitful discussions which gave my new impulses for my research.

At the end of the preface I want to mention that after writing this thesis I am absolutely sure that I made the right decision when I decided to start my PhD here at the Institute of Computer Graphics and Algorithms. I hope my research during this period leaves at least a small impact on the visualization community. In my life this period will leave for sure a huge and positive impact.

The will to win, the desire to succeed, the urge to reach your full potential... these are the keys that will unlock the door to personal excellence.

— Confucius

CHAPTER



.....
Introduction

VISUALIZATION has the purpose to give an insight into data. Each data contains some information and through visualization this information should be presented to a viewer [70]. The visual representation of data has the benefit that it is easier to perceive by humans than raw data in digits or letters.

The development of hardware to measure or simulate some phenomena rapidly increases the amount of data which should be investigated by the user. Through the simultaneous development of graphics hardware for the visualization it would be possible to process this increasing amount of data. But the visualization is limited by human perception. A visualization which is able to represent the whole information embedded in the data at once might not be practical, since a user is not able to process all this information.

In most applications some parts of the data are not necessary to fulfill a certain task. These parts of the data should be masked in the visualization to reduce the amount of information which is presented to the user.

In this thesis methods are presented which are able to mask data which are not necessary for the user. As result a visualization can be generated which is not overloaded with details. Hence the user can focus on the most important parts of the data.

The control about the parts which are enhanced or masked is still exercised by the user. However, the methods which are presented in this thesis are able to support this user task and, hence, make it easier to emphasize the most relevant parts of the data. In contrast to a naive separation between more and less important parts in the data, information theory and statistics are used to guide the user in finding the most relevant parts.

The thesis is structured in the following way: In the remaining part of this section an introduction to scientific visualization is given. With the visualization pipeline all steps are explained which are necessary to generate a visualization based on some data. The methods in this thesis are tools to improve the classification step in the visualization pipeline. Therefore, the classification step is described in more detail. It is depicted in which level of classification the introduced methods are applied. The introduction is completed by a brief survey of common concepts in information theory.

In Chapters 2, 3, and 4 the different methods for an information-based classification are explained. These novel methods show ways to use the well-known information theory and statistics to improve the classification process. All three methods were designed to provide an intuitive user interface which supports the user in extracting the most relevant parts of the data.

Each aspect and all algorithms of the methods are explained in detail. Results in each section show the usefulness of each individual approach.

The method in Chapter 2 describes a technique which investigates the local neighborhood around each point. Statistics is used to extract some information about the properties of the data in this area. With this method it is easier to distinguish between different objects or materials, represented in the data.

Chapter 3 describes a technique which is able to enhance parts of the data with a high information content. For this reason the global distribution of data values is used as reference for the estimation of the information content at a single point inside the data.

The last method - which is described in Chapter 4 - extracts the structure of objects in the data and uses this information to find similarities between them. For the measurement of the similarity, information theory is used. With this method it is possible to enhance structures which represent stable object surfaces.

In Chapter 5 a summary and conclusion is given. The main contribution of the thesis is depicted as well as limitations with the introduced methods are described.

1.1 Scientific Visualization

Over the last decades visualization has become an important part of many domains. Whenever data has to be investigated, a visualization of the data is useful to understand the data faster and in a more intuitive way. This starts with simple graphs, such as election polls, and ends with highly complex visualizations of multidimensional data, such as a weather forecast.

Due to the variety of different visualization techniques based on different applications and data sources, the field of visualization was split into two branches: information visualization and scientific visualization [21]. The separation is defined by the characteristics of the underlying data which is visualized. Information visualization describes the field of visualizing high dimensional data with no inherent spatial reference. Examples for sources of such data are files, text, relationships in the internet, or polls [3].

Scientific visualization focuses on the visualization of data which has an inherent spatial reference. The name scientific originates from the fact that data with inherent spatial reference is usually acquired by scientific experiments or measurements. The data describes phenomena arising from medicine, biology, metrology, architecture, etc. In most cases the data is embedded in a three dimensional (3D) space. Each point in the 3D space represents one or more properties of the measurement or simulation at its spatial position.

An example for scientific visualization is *Computed Tomography (CT)* or *Magnetic Resonance Imaging (MRI)* in medicine. These are two techniques to get a non-invasive insight into the human body. The resulting data of CT and MRI represents physical properties at each 3D position inside a human body. A visualization of the data should be able to show all important parts of the body for a certain medical investigation.

The main difference in the visualization technique between information and scientific visualization is the inherent spatial reference. While the data in information visualization can be arranged in any useful way, the data in scientific visualization has to be arranged according to their inherent spatial reference. By maintaining the spatial reference in the visualization, the

result is easier to interpret by a user. The drawback with the inherent spatial reference is that scientific visualization has to deal with occlusions since it is possible that some areas of the data set are occluded by other parts.

The techniques which are described in this thesis are all used for scientific visualization. Therefore, the term visualization is used as a synonym for scientific visualization in the following. Nevertheless, information and scientific visualization share many techniques and also the methods in this thesis might be useful for some applications in information visualization but this is out of scope of this thesis.

1.2 The Visualization Pipeline

The starting point for every visualization is the data which has to be visualized. The result of a visualization is an image which represents the underlying data from a certain viewpoint and additional parameter settings. A set of images from different perspectives and with different parameter settings should provide enough insight into the data to result in a correct interpretation.

All the steps which are necessary to generate an image from the data can be described by the visualization pipeline of Haber and McNabb [25]. Figure 1.1 illustrates the pipeline. It starts with data acquisition on the left side. The acquisition process can be either a simulation or measurement. The result of the data acquisition is *raw data*. The raw data stores for each point in 3D space its properties. In many cases it is only a single value per position. This is the case, e.g., in CT where a single value at each point represents the physical density of the scanned material. If the underlying phenomena are time-varying then each point represents a value for each time step. It is also common that a simulation or measurement is done several times with different modalities or settings. In this case each point in space represents the different values for each simulation or measurement. Another possibility is that a measurement or simulation is represented by a vector at each position.

These differences in the type of the data can be formalized by the terms *scalar field* and *vector field*. A scalar field holds a scalar, i.e., a single value, for each point in the field. A scalar field f can also be expressed by:

$$f : \mathbb{R}^3 \rightarrow \mathbb{R} \quad (1.1)$$

In contrast to that a vector field v represents an N -dimensional vector at each position:

$$v : \mathbb{R}^3 \rightarrow \mathbb{R}^N \quad (1.2)$$

In the literature the term vector field is only used if each point in the data represents a real vector - such as the direction of a flow - and not just a set of different properties. The visualization of such data is often referred to as flow visualization. Multimodal data or time-varying data is usually represented by several scalar fields for each modality or time step. In this thesis we concentrate on the visualization of scalar fields from single or multiple modalities.

Depending on the acquisition process it might happen that the raw data for example is too noisy or too large for the further processing steps. In such a case the raw data is filtered. After the filtering the data is called *visualization data*. Typical filtering techniques which are used in this processing step are smoothing - to reduce the noise - or downsampling - to reduce the size

of the data. In Figure 1.1 the filtering can be seen as second processing step in the visualization pipeline.

The visualization data still represents the original properties of the data, such as physical density in the case of CT. These properties cannot be displayed directly on the screen since they do not match any display scheme in general. Therefore, the visualization data has to be converted to something which can be visually represented. This process is known as classification or mapping. The output of the classification is a *visual abstraction* of the data as depicted in Figure 1.1.

During the classification process it is possible to enhance or mask parts of the data. For most applications this is an important step since the data is usually very dense and many parts are occluded. After the classification the most important parts for a certain application should be visible. Usually the classification is based on a combination of automatically retrieved features from the data and user interaction.

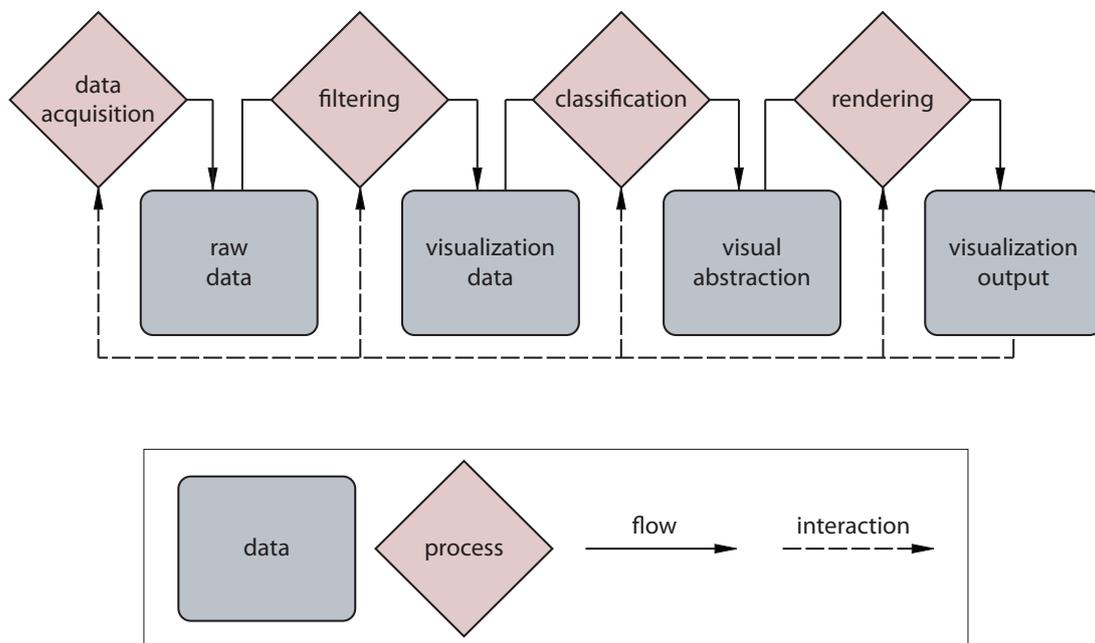


Figure 1.1: The visualization pipeline describes all steps from data acquisition until the final visualization output, i.e., an image.

The final step in the visualization pipeline is rendering. It results in a *visualization output*, i.e., an image which shows the visual abstraction from a certain perspective. Rendering techniques are divided into two groups [14]:

- Surface rendering
- Volume rendering
 - Object-order rendering

– Image-order rendering

For surface rendering a surface is extracted from the data before it is rendered. This extraction is done in the classification step and results in a surface model as visual abstraction. The surface model represents a surface of the object for given parameters, such as a certain density value in the case of CT. After the extraction of a surface the rendering is trivial since modern graphics hardware is optimized to render such surfaces. A well-known technique for surface rendering is called marching cubes [53]. It is a technique which generates triangles which represent the surface for a certain threshold.

The second rendering technique is volume rendering. For this rendering technique it is not necessary to generate an intermediate representation of the data, such as the surface model. The image is directly generated from the volumetric data. Therefore this technique is also referred to as *direct volume rendering* (DVR). All DVR techniques are classified as object-order or image-order techniques [79]. Figure 1.2 illustrates the two different rendering types for volume rendering. In object-order rendering on the left side each volume element (*voxel*) of the volumetric data is projected onto the image plane. The final color of a pixel on the screen is then calculated by a blending of all voxels which are projected to this pixel. A well-known representative of this technique is splatting introduced by Westover [79].

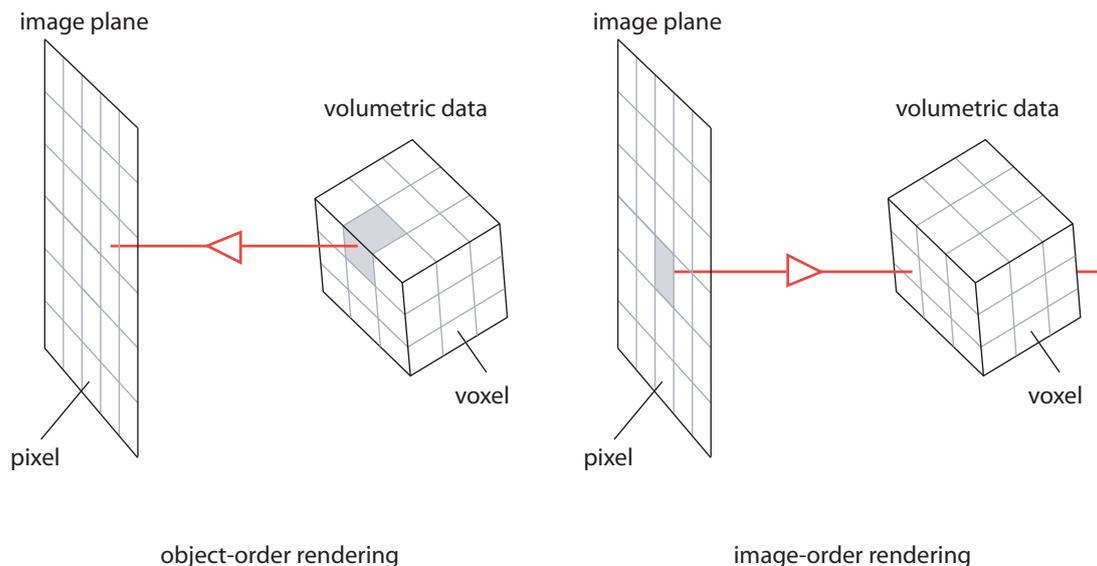


Figure 1.2: Different rendering orders for volume rendering.

The image-order techniques start from each pixel at the image plane. A ray is shot along the view direction into the 3D scene. On the right side of Figure 1.2 this process is illustrated. If the ray hits the object a color is calculated for the pixel by accumulating the colors of the voxels along the ray. This technique is also known as *volume raycasting* and was introduced by Levoy [49]. Due to the development of graphics hardware, raycasting can be implemented efficiently on the *graphics processing unit* (GPU). Therefore, raycasting became the state-of-art technique for many visualization applications.

The methods presented in this thesis are able to improve the classification process for certain applications. Raycasting is the technique of choice since it delivers better results than other rendering techniques while it is still fast on current graphics hardware. Therefore, a description of different rendering techniques is out of scope of this thesis. A detailed overview over different rendering techniques is given by Elvins [14].

After the rendering process the whole visualization pipeline has been traversed. The result of the visualization pipeline is an image which visually represents the raw data. In Figure 1.1 it can be seen that the user is able to interact with all processing steps of the pipeline. This is necessary because a single image generated with a set of different settings is in most cases not sufficient to understand the complexity of the underlying data.

The next section covers classification in more detail. It will be explained how the methods described in this thesis are integrated into the visualization pipeline. The goal of classification is to enhance parts of the data which are important for a certain task and mask all other parts. In Chapters 2, 3, and 4 novel techniques are described which are able to provide this classification goal.

1.3 Classification

In the classification step the visualization data is converted into a visual abstraction as shown in Figure 1.1. The classification has two purposes. One purpose is the mapping of the visualization data to optical properties which can be used in the rendering process. In most cases the optical properties are color and opacity. After the classification each voxel is represented by optical properties. This abstraction can then be used by the rendering algorithm, such as raycasting, to generate an image.

A second purpose of classification is the segmentation of the data. In most cases the data in scientific visualization is very dense. This means that many areas inside the volumetric data are occluded by other areas after the projection to an image in the rendering step. As an example, we can consider the volumetric data of a medical CT scan of a human head. In this case the bones are occluded by muscle tissue and skin. The bones would be occluded all the time no matter how the viewing direction is changed. The result on the left side of Figure 1.3 shows this case. If someone is interested in the bones the classification process can be adapted to mask the tissues which occlude the bones in the final projection. The result after such a classification is shown on the right side of Figure 1.3.

In the literature the classification is often formalized with the concept of a *transfer function* [15, 34]. A transfer function defines the mapping from features of the data to optical properties. By modifying the transfer function some parts of the data can be enhanced or masked. For the example in Figure 1.3 two different transfer functions were used to show the skin in one result image and the bones in the other.

The visualization pipeline in Figure 1.1 of the previous section indicates that the classification process is a separate step before rendering. This is only true when pre-classification is used. In this case each voxel of the visualization data is classified with optical properties before the rendering is applied. In the rendering process the optical properties of a sample point are calculated by interpolating between optical properties of the surrounding voxels. For the methods in this thesis

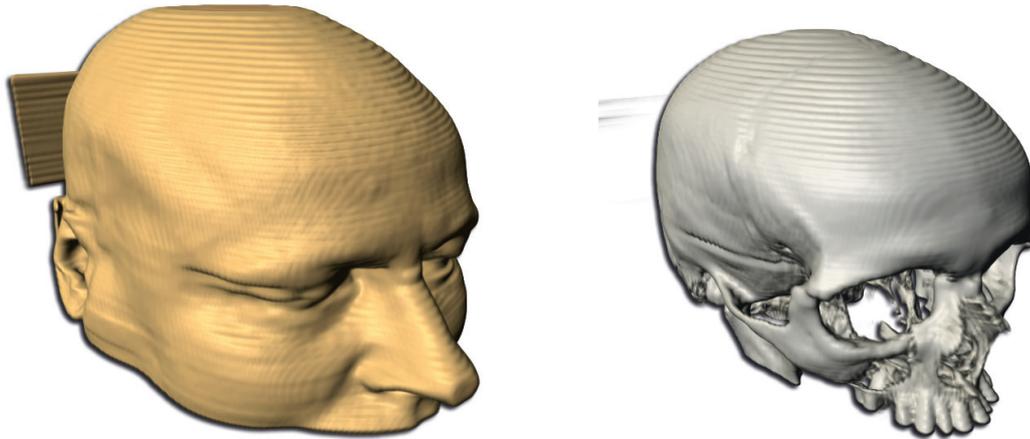


Figure 1.3: Influence of the classification on the result image of the visualization. On the left side the skin is occluding all inner tissues. In the visualization on the right side all tissues except the bones are masked.

post-classification is used. In post-classification the features at a sample point are calculated by interpolating the visualization data first. Based on this interpolated values optical properties are assigned to the sample point.

In contrast to pre-classification, post-classification needs more processing power during the rendering process but it results in smoother visualizations since less information is lost in the interpolation between the original data values. With modern graphics hardware the processing power is high enough to use post-classification without any limitations in the rendering speed.

Even though classification is often directly integrated into the rendering process, it can be seen as separate processing step in the visualization pipeline. Pre- and post-classification only differ in some implementation details. For the further thesis this difference is not important. Therefore we will only use the term classification to refer to pre- and post-classification.

In classification the assignment of optical properties is dependent on features extracted from the data. A feature is, e.g., the data value at the voxel position which should be classified. In the example in Figure 1.3 the data value is used as feature. In this case it is possible to classify bones based on this feature alone. In many applications a single feature is not good enough for the classification because by itself it is not distinctive to enhance certain parts of the data. In such a case other features or additional features of the data are used for the classification.

In this thesis the features which are used for the classification are differentiated by the area which is involved in their extraction. Figure 1.4 shows an overview on different possibilities for the extraction of features. The red dot in the different levels of classification symbolizes the sample point which has to be classified.

If the classification is based only on a single location then just the original data value is used to assign optical properties to a voxel. A more sophisticated classification uses features extracted from the local neighborhood around the voxel. For some applications it is useful to calculate

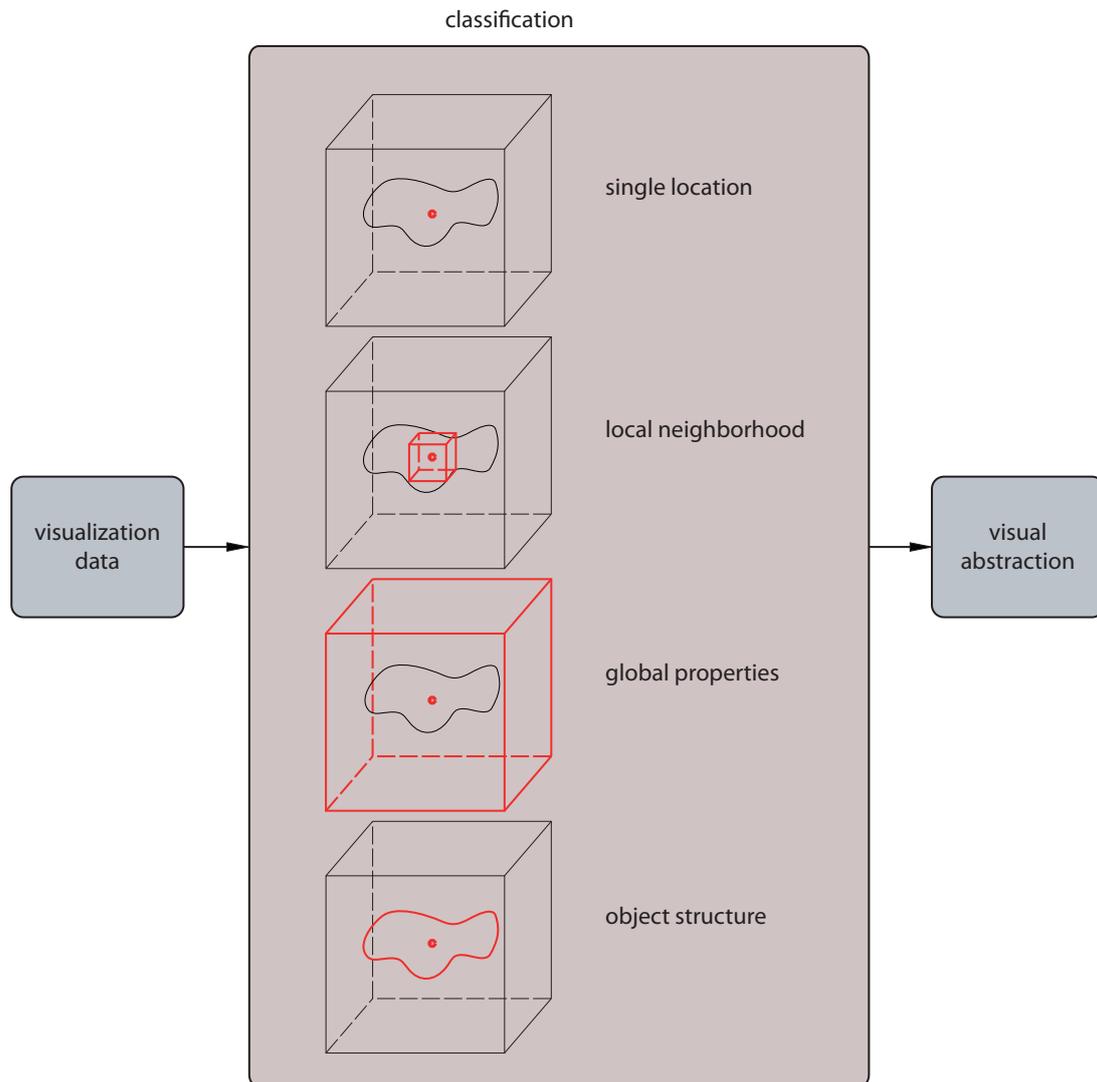


Figure 1.4: Different levels of classifying visualization data at a certain location (red dot). The complexity increases with the size and intricacy of the region which is considered for classification.

features based on global properties. The final possibility for extracting features is based on global structures of the object represented in the data.

Based on the features which are extracted by these different methods a transfer function can be designed which assigns optical properties to single sample points. In the methods described in this thesis we show how different features are extracted and how these features are used to improve the classification process for certain applications in scientific visualization. But first, the differences of the extraction of features at various levels are highlighted.

1.3.1 Single Location

The most trivial classification method is based only on the values at the position which has to be classified. In the simplest case this is a single value if only a single volume is represented by the data. In the case of multimodal data and time-varying data more values are involved in the classification.

The classification based only on the local value was first introduced by Levoy [49]. The transfer function for this mapping assigns optical properties to each value in the whole value range of the data. The same optical properties can be assigned to different data values. It is also possible to mask certain data values by assigning optical properties to them which are not visible in the final rendering. This can be achieved, e.g., by setting the opacity to zero.

The design of a transfer function is the process of defining and changing optical properties for different data values until the desired rendering result is generated. In the example with the bones from Figure 1.3 this would mean to set the opacity for all data values which do not represent bones to zero.

The design of a transfer function is done by the user of a visualization application. Usually the design is a trial-and-error process which needs experience. For every data set the transfer function has to be designed anew or at least adapted from a preset. Different hints about the characteristics of the data can support the design process. For example the frequency distribution of the data values can help to identify value ranges for soft tissues or bones in the case of CT. This is possible because more voxel belong to soft tissue or bones than to any other tissue.

Since the classification based on the value at a single location is simple, it is still widely used. It is best suited for data which contains materials with distinct data value ranges. In this case it is possible to assign different optical properties for each material based on the data value. If the data value ranges of different materials overlap, this method is not sufficient to distinctively classify different materials. Hence more features have to be extracted from the data to make a distinctive classification possible.

1.3.2 Local Neighborhood

If it is not possible to get a good classification based only on the data values, it is possible to extract additional features from the local neighborhood around the sample point which should be classified. This should result in a set of features which are distinctive for different materials and, hence, different optical properties can be assigned to each material.

Kniss et al. [40] employed the gradient magnitude as additional feature for the classification. To estimate the gradient magnitude the neighboring voxels are used. The gradient magnitude is larger in areas with high variations of the data values. Such areas are interfaces between materials where the values change rapidly. Therefore, the classification based on this additional feature is able to distinguish interfaces between materials from material interiors.

For many applications the interfaces between materials are of interest. Hence many methods were developed which are able to classify these interfaces. As additional features, e.g., curvature [29] or data values along the gradient direction [54, 72] are used. Through these additional features the transfer function gets more complicated since it has more input values. Transfer

functions which use additional features are also known as multi-dimensional transfer functions. The additional features define a two- or higher-dimensional transfer-function space.

If the interior of objects is of interest then it is possible to extract features which are able to depict salient regions of the same material. Lundstrom et al. [56] introduced a method which classifies different materials by their local histograms. Local histograms are more stable features to classify different materials than single data values. Caban and Rheingans [7] use textural properties of a local neighborhood to assign optical properties to different materials.

The extraction of additional features in a local neighborhood is more complex and needs additional processing time. The transfer-function design also gets more complicated since there are more degrees-of-freedom. On the other hand the classification of the data can be improved by the additional features. Different optical properties can be assigned to different materials even though this is not possible when the data value is used alone.

In Chapter 2 a method is introduced which extracts additional features in a local neighborhood. The transfer-function space is defined in a way to simplify the transfer-function design process.

1.3.3 Global Properties

In contrast to features which are extracted from the local neighborhood it is also possible to use global properties for the classification. With global properties it is possible to extract information about the general content in the data set. An example for a global property is the probability of the occurrence of a certain data value in the whole data set. Such information together with the local data value can be used to extract features for the classification.

Kniss et al. [43] introduced a classification method based on statistics. In this case the probability of the occurrence of a certain material in the data set was used as global property. Based on this estimation the membership of a value to one of the materials is calculated.

The benefit of global properties in contrast to features extracted from a local neighborhood is their robustness. Noise or other artifacts have only a low impact on global properties while they might have a high impact in a local neighborhood.

In Chapter 3 a technique is described which uses global properties for the classification of multimodal data. The global properties are used to estimate the information content for each combination of values. Through this it is possible to enhance parts with higher information content.

1.3.4 Object Structure

The classification based on features extracted from the object structure is the most sophisticated technique. This classification technique is useful if the structure of the object is of interest.

Hadwiger et al. [26] developed a region growing algorithm to detect structures of different sizes in the data. The size of the structure is the feature which is used for the classification with optical properties. Correa and Ma [12] use a multi-scale approach to detect features of different sizes.

Isosurfaces are another structural feature which can be extracted from the data. Tenginakai et al. [68, 69] introduced a method to extract salient isosurfaces. In their work statistical properties for isosurfaces are calculated and used as features for the classification. Carr et al. [9] use the

contour tree of isosurfaces to classify object structures. In the work of Khoury and Wender [35] the fractal dimension of isosurfaces is used as feature. Bruckner and Möller [6] compared different isosurfaces with each other to estimate their similarity. The similarity was then used as feature to classify stable isosurfaces.

In Chapter 4 a method is introduced which uses the similarity of isosurfaces in multimodal data to classify stable surfaces of both modalities. Thus it is possible to enhance or mask features which are very similar in both modalities.

With the different features extracted from the data it is possible to assign optical properties to different materials in the data. The classification technique which is used for the visualization depends on the data and the application.

Nevertheless how good a classification method is, some information gets lost in the classification process. In each step of the visualization pipeline some information is typically lost. This is in many cases unavoidable since it is rarely possible to visualize all parts of the raw data at once. The goal throughout the visualization pipeline is to discard only information which is not necessary to fulfill a certain task based on the data.

The naive approach for the classification to prevent loss of important information is to let the user control what will be discarded. For data sets which represent simple objects this is possible but for more complex data sets it is difficult for the user to modify the classification in a way to discard only information which is not necessary for a certain task.

In this thesis we introduce methods for the classification of visualization data which are based on statistical properties and information theory. Through this the user is supported in the task of modifying the classification to discard less important parts of the data and enhance the most important ones.

Since information theory is an important tool for the methods in this thesis, the next section gives a short introduction to it. It is also shown for which other aspects than classification information theory can be used in the scope of scientific visualization.

1.4 Information Theory in Visualization

Information theory was introduced in the late 1940s by Shannon [65]. Initially it was developed for the application of signal processing. Figure 1.5 shows the original setup which was used to formalize information theory. A sender sends a signal over a noisy channel to a receiver. With the introduction of information theory it was the first time possible to calculate the amount of information which can be transmitted without the loss of any information over a given channel.

Since then information theory was employed in many research fields such as electrical engineering, mathematics, computer science, physics, economics and arts [13]. Applications which are based on information theory are, e.g., data compression and image registration. In data compression the loss of information can be quantified by information theory for a certain compression rate. Hence, the compression rate can be chosen according to the maximum possible loss of information. In the field of image processing, image registration is one task which uses information theory. Two images are optimally registered to each other when one image contains the maximum amount of information about the other image. The field of image processing

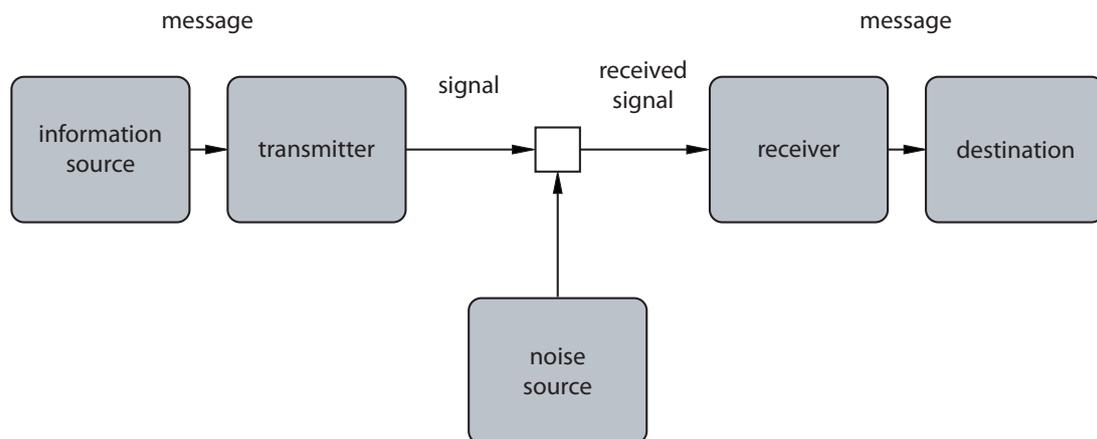


Figure 1.5: Information theory was originally introduced by Shannon [65] for signal processing.

inspired many applications in visualization which use information theory. Wang and Shen [74] depict the increasing influence of information theory in visualization.

The analogy between the utilization of information theory in visualization in comparison to usage in its original field of signal processing can be explained through the visualization pipeline. In Section 1.2 the visualization pipeline was explained. According to the illustration of signal processing in Figure 1.5 the raw data can be seen as sender and the visualization output as receiver. All the processing steps between are part of the transmission channel. In general the size and dimension of the raw data is too high to generate a visualization output which represents every detail of the data. Hence some data gets lost during the different processing steps along the visualization pipeline. With information theory it is possible to minimize the loss of information along the visualization pipeline by discarding data which is not necessary for a certain task [10].

The filtering step along the visualization pipeline is necessary especially when the raw data is noisy or contains some other artifacts. The goal is to preserve the shape and structure of the data as much as possible while the noise or artifacts are eliminated. In practice it is not possible to eliminate all noise without modifying the structure of the data. Information theory can be used in this case to measure the amount of information loss during the filtering. The filter can then be adjusted in a way to preserve the structure of the data up to a certain level. An example for a filtering method based on information theory is the work of Cheng et al. [11]. In this work the data is converted into a so called fuzzy domain. In this domain data can be discarded depending on its fuzziness.

The classification step is the processing step which is able to discard most of the data. Therefore, it is especially important to keep the information loss in this step as small as possible. In flow visualization, e.g., data is very dense and therefore it is often represented by streamlines. In the classification step the number and positions of streamlines are generated. Xu et al. [80] presented a method which employs information theory for the distribution of streamlines, based on the data which represents the flow field. Another classification task which uses information theory is the identification of representative isosurfaces for an object in the data. Bruckner and

Möller [6] introduced a measurement for the similarity of two isosurfaces based on information theory. The similarity is used to detect isosurfaces which best represent the structure of an object. Laidlaw et al. [45] apply Bayes' theorem on a small neighborhood around a sample point to classify mixed materials. This method can be used to detect interfaces between different objects.

In the rendering step some information gets lost through the mapping from 3D to 2D. Nevertheless the loss of information can be minimized with different techniques. The viewpoint is one criterion which highly influences the amount of information which can be perceived in the visualization output. In most visualization applications the user can interactively manipulate the viewpoint. Bordoloi and Shen [4] introduced a information-theoretic technique which automatically selects a viewpoint that preserves the maximum information in the rendering step. In some applications the data is represented in different resolutions. A higher resolution is able to represent more details in contrast to a lower resolution. Each region of the volume should be represented by a resolution which is sufficient enough to faithfully represent the data. Wang and Shen [73] developed a method which automatically selects the best resolution for each region of the volume. The automatic selection is based on information theory.

In this thesis three different techniques are described which employ information theory and its methods in the classification step to minimize the information loss (Chapters 2, 3, and 4). Since all three methods use techniques which are related to information theory the most important concepts are briefly described in the remaining part of this section.

1.4.1 Concepts of Information Theory

Information theory includes all methods which are used to quantify information. As mentioned earlier this quantification became important for many research fields after it was introduced by Shannon [65].

The basis for the quantification of information is the probability of the occurrence of a certain signal. If we consider for example the English language as source then each letter is a signal with a certain probability of occurrence. The letter e occurs more often than the letter x . Hence the quantity of information incorporated with the letter e is lower in comparison to the letter x . The reason is that a signal which appears less likely reveals more information. In the example with letters, an x appears less often in a word or text. Therefore, a better prediction about the whole message can be made by receiving an x instead of an e .

This simple example depicts the basic concept behind information theory. Since probabilities for the occurrence of signals is the basis for the quantification of information, statistics is a fundamental tool for all calculations of quantities in information theory. From the point of view of statistics each signal can be seen as a random variable X . In the discrete case X can have n different conditions x_i . Each of these conditions occur with a certain probability $p(x_i)$. If the random variable can be described by a continuous function then it can have infinite different conditions x with a probability of $p(x)$ for each condition x . On the left side of Figure 1.6 an example for the *probability density function* (pdf) $p(x)$ of a random variable X is shown.

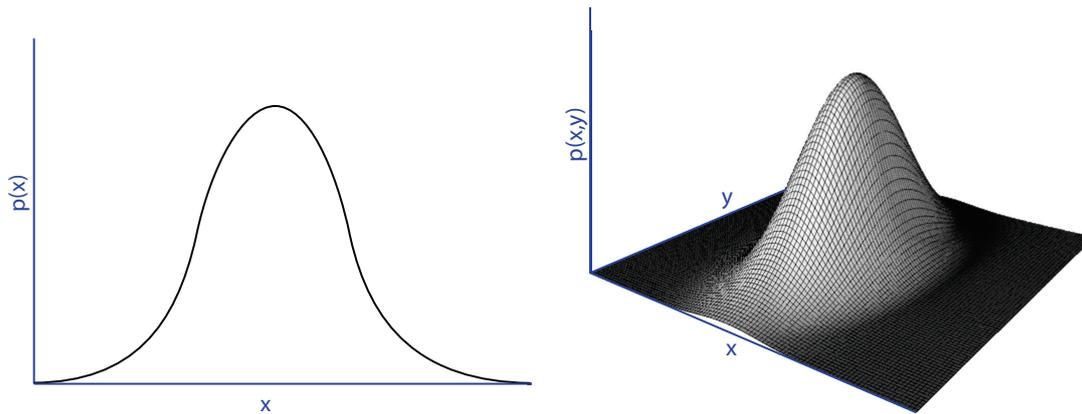


Figure 1.6: The figure shows the probability density function $p(x)$ for a random variable X on the left and the joint probability density function $p(x,y)$ for the co-occurrence of two random variables X and Y on the right.

The pdf has the property that the sum of probabilities for all conditions of x must be one in total:

$$\int_{\mathcal{X}} p(x) dx = 1 \quad (1.3)$$

In scientific visualization every feature of the data can be seen as a random variable. The data value itself, e.g., can be used as random variable. In general the pdf for a certain feature which is used as a random variable is not known. In this case the pdf can be estimated by a normalized histogram. For the generation of the histogram the value range of a feature is subdivided into a certain number of bins. Each bin is increased by one for a value within its borders. Finally the counts of each bin are divided by the total number of points. This results in a discrete probability function $p(x)$ for a feature X which can be used for further calculations.

The probability $p(x)$ of a random variable X can be used directly to calculate the self information $I(x)$:

$$I(x) = -\log(p(x)) \quad (1.4)$$

If the logarithm in the equation is to the base of two, the equation returns the self information measured in *bits*. The self information quantifies the storage which is necessary to encode the signal x . The lower the probability $p(x)$ the higher the self information $I(x)$ and the higher the storage which is necessary for encoding x .

In information theory many concepts rely not only on the information of a single signal itself but on the information which is included in the combination of two random variables X and Y . Such concepts are used when, e.g., the mutual information of two signals should be calculated. For these calculations it is necessary to know the joint pdf $p(x,y)$ of the co-occurrence of x and y . On the right side of Figure 1.6 an example of a joint pdf for two random variables X and Y is shown.

Two random variables are called independent when the joint probability $p(x,y)$ is the product of the individual probabilities $p(x) * p(y)$ for all x of X and y of Y . This is the case if the occurrence of one signal x has nothing to do with the occurrence of the other signal y .

The joint information $I(x,y)$ for the co-occurrence of two signals x and y can be calculated in the same way as the self information in Equation 1.4. Instead of $p(x)$ the joint probability $p(x,y)$ is used.

With these theoretical concepts of self information and joint information together with the probabilities of the occurrence of signals, several information-theoretic quantities can be calculated. The most important ones are briefly introduced in the following.

Entropy

The entropy $H(X)$ is a measure for the average information content of a random variable X . The entropy of a continuous random variable X can be calculated by the following equation:

$$H(X) = E(I(X)) = \int_X p(x)I(x) dx \quad (1.5)$$

E is the expected value and I is the information content. The entropy can equally be expressed by the integral of the probability-weighted information content.

In practice the entropy can be used to calculate the encoding length for the whole message. The range of the random variable X is the whole message in this case. If one particular signal x has the probability of 1 and all other signals have a probability of 0 then the entropy is 0. This case needs the lowest encoding length for the whole message. The other extreme is a random variable X where each signal x occurs with exactly the same probability. In this case the entropy is at a maximum and the whole message needs the highest encoding length.

Joint Entropy

The joint entropy $H(X,Y)$ expresses the average information content for the co-occurrence of two random variables X and Y . Similar to the entropy it can be formalized by the following equation:

$$H(X,Y) = - \int_Y \int_X p(x,y) \log(p(x,y)) dx dy \quad (1.6)$$

$p(x,y)$ is the joint probability for the co-occurrence of x and y . The two extrema for the joint entropy are also similar to the entropy. The joint entropy is maximal when all combinations of signals x and y occur with the same probability. The joint entropy is zero if only one combination of signals occurs.

Conditional Entropy

In information theory the conditional entropy $H(Y|X)$ expresses the remaining information, i.e. uncertainty, of a random variable Y while the other random variable X is already given. The conditional entropy can be expressed in the following way:

$$H(Y|X) = - \int_Y \int_X p(x,y) \log \frac{p(x,y)}{p(x)} dx dy \quad (1.7)$$

In other words the conditional entropy is low if the random variable Y contains similar information as the random variable X . In this case a lot of information is already known about the random variable Y by knowing the random variable X . If the conditional entropy is high the two random variables differ much more from each other. By knowing the random variable X less information about random variable Y is known. If both random variables are independent from each other then the conditional entropy $H(Y|X)$ is equal to the entropy $H(Y)$.

Mutual Information

The mutual information $I(X, Y)$ of two random variables X and Y is a quantity to measure the mutual dependency between the variables. It is the reduction in the uncertainty of one random variable due to the knowledge of the other one [13]. The mutual information can be formalized by:

$$I(X, Y) = \int_Y \int_X p(x, y) \log \left(\frac{p(x, y)}{p(x) p(y)} \right) dx dy \quad (1.8)$$

An extreme of the mutual information occurs when both random variables are independent from each other, i.e., $p(x, y) = p(x) * p(y)$. In this case the knowledge of one random variable does not give any information about the other random variable. Hence the mutual information is zero. The other extreme occurs when both random variables are identical. In this case all information about one random variable is known if the other random variable is given.

Instead of Equation 1.8 the mutual information can also be expressed by the entropies, joint entropy, and conditional entropies of random variable X and Y :

$$I(X, Y) = H(X) - H(X|Y) \quad (1.9)$$

$$= H(Y) - H(Y|X) \quad (1.10)$$

$$= H(X) + H(Y) - H(X, Y) \quad (1.11)$$

$$= H(X, Y) - H(X|Y) - H(Y|X) \quad (1.12)$$

Figure 1.7 illustrates the connection between the quantities of information theory. The two circles illustrate the individual entropies $H(X)$ and $H(Y)$. The intersection of both circles is the mutual information $I(X, Y)$. The combination of both circles represents the joint entropy $H(X, Y)$. The conditional entropies $H(Y|X)$ and $H(X|Y)$ are the individual entropies without the intersecting part from the other random variable.

The introduced concepts of information theory are used in the following chapters to calculate features for the classification of volumetric data. It is shown that information theory is a useful tool to enhance features of interest while masking parts of the data which are less important for the visualization. First of all a method is introduced in the next chapter which uses the statistical basis of information theory to classify different materials of a data set.

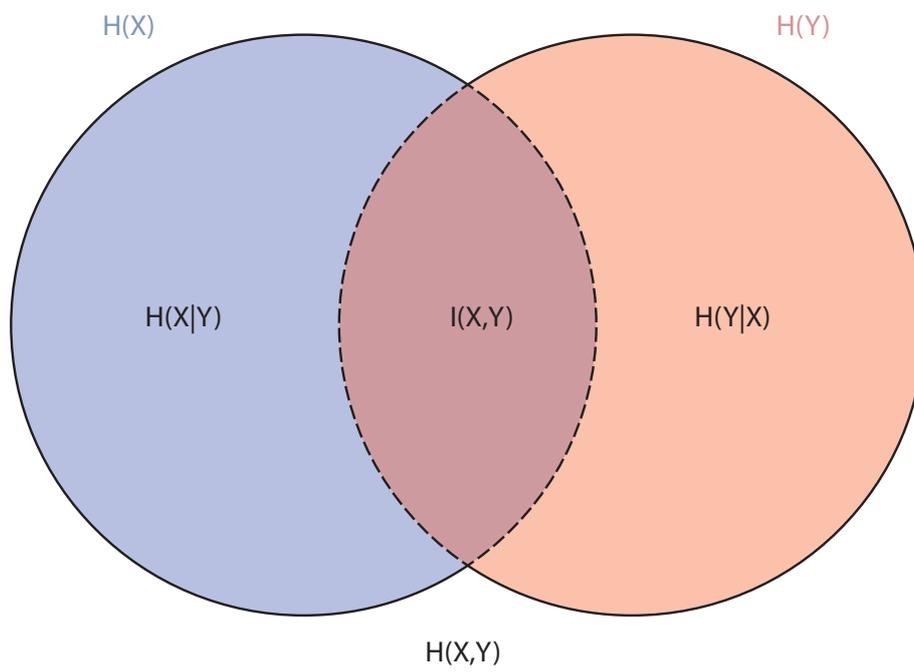


Figure 1.7: Correlation between different entropies H and the mutual information I .

CHAPTER

2

.....
**Volume Classification based on
Statistical Transfer-Function
Spaces**

Martin Haidacher, Daniel Patel, Stefan Bruckner, Armin Kanitsar and M. Eduard Gröller.
Volume Visualization based on Statistical Transfer-Function Spaces. In Proceedings of
IEEE Pacific Visualization Symposium 2010, Pages 17-24, 2010

As depicted in Chapter 1 classification is a major issue in volume visualization. The goal of classification is to enhance different materials or objects which are important for further analysis of the data. Transfer functions have been proven to be a powerful tool for classification. Nevertheless, in most cases it is a non-trivial task to find a good transfer function which is able to achieve the desired classification.

A transfer function (TF) is a general concept. Concrete implementations use one or more features, derived from the data, to define a transfer-function space. A transfer function is then designed in this space. How easily different materials can be distinguished by the transfer function, depends on the data set as well as on the used features in the definition of the transfer-function space.

Noise in the measured data is a typical problem, which complicates the classification process. The most frequently observed noise in measured data is Gaussian white noise. White noise has a mean value of zero and a symmetric variance. For different materials in a volume data set, the noise causes variations of the intensity values around an average value. Therefore, it is difficult to assign intensity values of sample points to a certain material, especially if average values of other materials are close by. This intermixing of materials in the intensity space cannot be resolved in the transfer-function space when only local features, such as the data values, are used to define the transfer-function space.

In this chapter a method is presented which considers a local neighborhood around each sample point to derive features for the transfer-function space. By this, we are able to estimate the distribution of noise around the average value of a material. The statistical properties of this distribution are used to describe the material at a certain sample position. Since different materials can be distinguished by their distributions of intensity values, we are able to separate them.

For the estimation of the statistical properties, we employ an adaptive growing approach at each sample point. The extent of the growing is dependent on the local neighborhood of a sample point. The estimated features are used to define the statistical transfer-function space.

Sample points from separate materials can be seen as separate clusters in this space. We introduce special transfer-function regions which are adapted for this space to design a transfer function. Furthermore, we demonstrate how the statistical properties can be used to steer visual properties such as shading. This results in higher quality visualizations, especially for noisy data.

We use the new statistical transfer-function space to generate images for medical MRI and industrial CT data. We show for these data sets, that our method classifies different materials better than other state-of-the-art transfer-function spaces.

2.1 Related Work

The method presented in this chapter spans several research fields. Since we are dealing with noise in the data, the analysis of noise in image processing is related to this approach. The growing of regions around each sample point is similar to the scale-space analysis where data is analyzed on different scales. In this chapter we introduce a new transfer-function space. Therefore, most of the related work is dedicated to other transfer-function spaces.

Image Processing. Noise in data is a well investigated field in image processing. For this approach we assume the noise in the data to be Gaussian white noise. This is specifically true for CT data sets [47, 75]. For other data sets, the Gaussian distribution is at least a good approximation of the noise distribution. In MRI, e.g., the real distribution is a Rician distribution, but for a low signal-to-noise ratio the difference to a Gaussian distribution is very small [24].

Scale-Space Analysis. Early works, such as Lindeberg [50], analyzed images on different scales. Over the years, different scale-spaces were investigated. The most common scale space is the linear scale space, which is generated by progressive Gaussian smoothing. In this scale-space Lindeberg [51] introduced a technique for feature detection and automatic scale selection. Due to the complexity of the scale-space generation for volume data, alternatives, such as Laplacian pyramids [23] or Wavelet transforms [58], were developed for an easier and faster representation of different scales. A method to improve the classification of features, based on a pyramid representation, was introduced by Lum et al. [55]. In contrast to scale-space analysis, our method uses different scales for each sample point because the growing is terminated depending on local features of the neighborhood region.

Transfer Function Spaces. In an early work, Levoy [49] used the data value alone to define a transfer function space. Kniss et al. [40] employed the data value and the gradient magnitude for the classification of different materials and borders between them. Since they only consider single data values and a very small neighborhood for the gradient magnitude, this technique is not well suited for the classification of noisy data. Hladůvka et al. [29] proposed curvature as an additional property for the classification. With this method special features, like ridges and valleys, could be extracted. An extension to multi-dimensional transfer functions was introduced by Roettger et al. [63]. The method includes spatial information in the transfer-function space. They simplified the transfer-function design-process by using the spatial information to color the transfer-function space. However, for noisy data different materials overlap in this transfer-function space.

In the work of Lum and Ma [54], a larger region is considered for the definition of the transfer-function space. Besides the data value at a sample point, a data value along the gradient direction is used as well. In data sets with sharp transitions, such as CT data, this technique can be

used to highlight border areas. An extension to this method was introduced by Šereda et al. [72] named LH histograms. This method looks for homogeneous regions along both directions of the gradient streamline. The detected low and high data values are used to define the transfer-function space. This method provides good results for data sets with little noise. For noisy data sets, values in homogeneous regions have a high variance. Therefore, the clusters representing homogeneous regions are getting larger and overlap each other in the LH histogram space.

A method which also uses a larger neighborhood for the classification was presented by Hadwiger et al. [26]. They use region growing to detect features of different sizes in industrial CT data. In a 3D transfer-function space these different features can be classified. In the work of Correa and Ma [12], a multi-scale approach is used to detect the size of features in a data set. The feature size is then used as an additional parameter for the definition of a transfer function. In both approaches the shape of a feature in the data set is the main criterion for the classification. Instead, in our method the statistical properties of materials are used for the classification. These properties are independent of object shapes.

Lundström et al. [56] introduced a method to classify different tissues by the local histograms in the neighborhood around a sample point. Caban and Rheingans [7] used textural properties to differentiate between materials, possibly with similar data values. These methods are able to separate materials but they use a neighborhood with a fixed size for the extraction. Thus, these approaches do not differentiate between homogeneous and inhomogeneous regions.

Laidlaw et al. [45] use Bayes' theorem on a small neighborhood of a voxel to classify mixed materials. Tenginakai et al. [68, 69] introduced a method to extract salient iso-surfaces based on statistical methods. A different classification based on statistics was introduced by Kniss et al. [43]. For the estimation of the statistical characteristics certain features of the different materials have to be known. For our approach no prior knowledge of material properties is necessary. Lundström et al. [57] used the variance in a neighborhood of a voxel to separate materials. In comparison to our method, they used a fixed neighborhood size to estimate the variance. In a previous publication (Patel et al. [59]) we used statistical properties to manually classify materials for differently sized neighborhood regions. In this approach we extract the statistical properties for the best suited neighborhood size semi-automatically. Furthermore, we use these statistical properties to define a transfer-function space and to enhance the visual appearance of the resulting rendering.

2.2 Statistical Transfer-Function Space

The idea behind the statistical transfer-function space is that materials are distinguishable according to their statistical properties. Since the data is not segmented, we are not able to calculate the statistical properties for different materials in general. Therefore, we introduce a technique which extracts statistical properties for the neighborhood of each sample point individually. We expect that sample points from the same material get similar statistical properties. In the new transfer-function space this leads to clusters for different materials, which makes it possible to design meaningful transfer functions. In this section we describe all steps which are necessary to generate the statistical transfer-function space.

materials have high standard deviations ($\sigma_1 = 0.14$, $\sigma_2 = 0.09$, and $\sigma_3 = 0.11$), consequently the distributions considerably overlap each other. The blue line gives the frequency distribution of all three materials together.

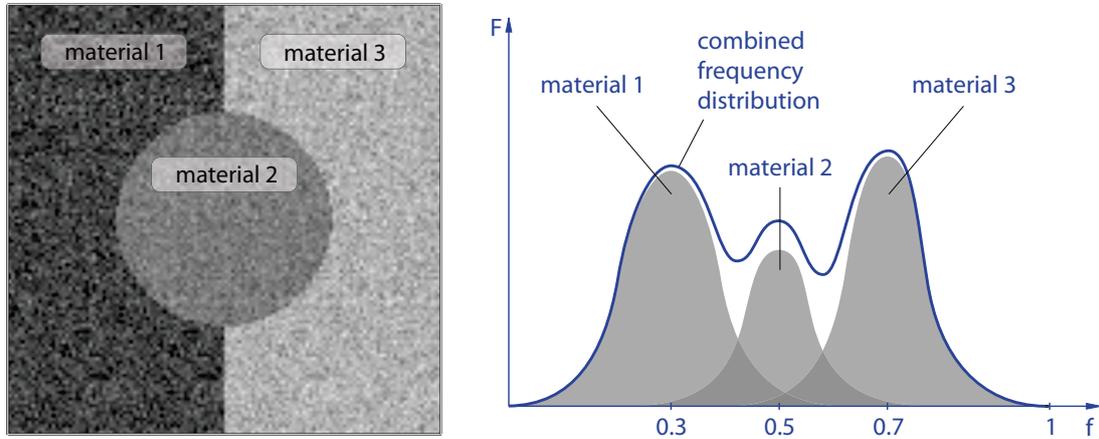


Figure 2.2: The figure depicts a slice through the synthetic data on the left side and the frequency distribution of the data values on the right side.

The synthetic data set exemplifies noisy data. The overlapping of the distributions for different materials is a problem which often occurs in real-world data sets [72]. In the further explanation of the method, the synthetic data set is used to show the effects of the different processing steps.

2.2.1 Statistical Properties Extraction

The extraction of the statistical properties is essential for our statistical transfer-function space. For sample points within a certain material, the extracted statistical properties should be close to the statistical properties of the entire material. We achieve this by investigating a neighborhood region around each sample point. To keep the neighborhood within the same material, we introduce an adaptive growing which is dependent on the local features.

The distribution of data values of a single material in real-world data sets, such as MRI or CT, can be approximated very well by the Gaussian white noise model. Therefore, we consider only the Gaussian distribution as basis for the calculation of statistical properties. A Gaussian distribution is described by its mean value and standard deviation. Hence, we use these two parameters as our statistical properties.

For the extraction of the statistical properties, we iteratively grow a spherical neighborhood by increasing the radius by one voxel in each step. We compare for each growing step if the newly grown hull still belongs to the same material. Figure 2.3 shows a cross section of such a neighborhood. In the following explanation we use two different notations for the statistical properties. The mean value μ_r and the standard deviation σ_r for a certain radius r , are the estimations for the statistical properties of all points within a sphere of radius r . $\hat{\mu}_r$ and $\hat{\sigma}_r$ are the statistical properties of the points in the outer hull of the sphere (see Figure 2.3).

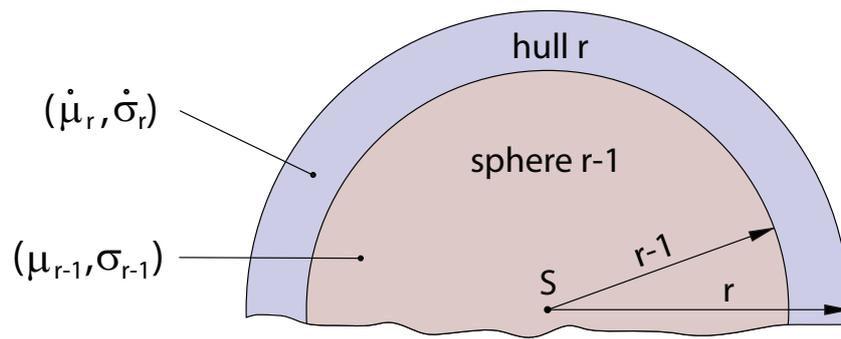


Figure 2.3: The standard deviation and mean for the sphere and the outer hull.

In each growing step several calculations are done to decide if the growing should be terminated or not. Figure 2.4 shows the processing steps for each growing step. The goal of these steps is to detect whether the sphere grows into another material. In such a case the loop is terminated.

As initial parameters for the loop, the mean value and the standard deviation of a sphere with a radius of one are used. The data values of points in the sphere may not be normally distributed. This can happen when the sphere intersects two materials. Such a situation should be detected in an early stage so growing can be terminated. Therefore, we apply a normal-distribution test, as described later in this section, to check if the points in the initial sphere are normally distributed. Only if this test is passed, the loop is started with the initial statistical properties. Otherwise μ_1 and σ_1 of the initial sphere are used as statistical properties for the actual sample point.

In the loop depicted in Figure 2.4, the first step is the calculation of the statistical properties $\hat{\mu}_r$ and $\hat{\sigma}_r$ for points in the hull at a radius r . In the next step it is tested if the values in the hull are normally distributed. If this is the case, the properties are compared with μ_{r-1} and σ_{r-1} . In the case the statistical properties are similar, the statistical properties of the hull are merged with the statistical properties of the sphere $r-1$. If a sample point lies in the center of a large homogeneous area, the loop is terminated when the maximum radius r_{max} is reached.

For the extraction of the statistical properties, a confidence level ω has to be set. This confidence level expresses the general confidence in the distribution of data values in a data set with respect to the general noise level. It can be set differently to adapt the model for various data types such as MRI or CT. In the following part of this section, the usage of ω and the other processing steps are described in more detail.

Properties Estimation

In each successive cycle of the extraction loop, the statistical properties of a larger region are considered. Since the statistical properties for sphere $r-1$ are already known, we are interested in the statistical properties $\hat{\mu}_r$ and $\hat{\sigma}_r$ of the additional points in the hull r (see Figure 2.3).

As the distribution is considered to be Gaussian, we estimate the mean value and the standard deviation for the points in the hull r [19]. The mean value $\hat{\mu}_r$ is the average and $\hat{\sigma}_r$ is the biased

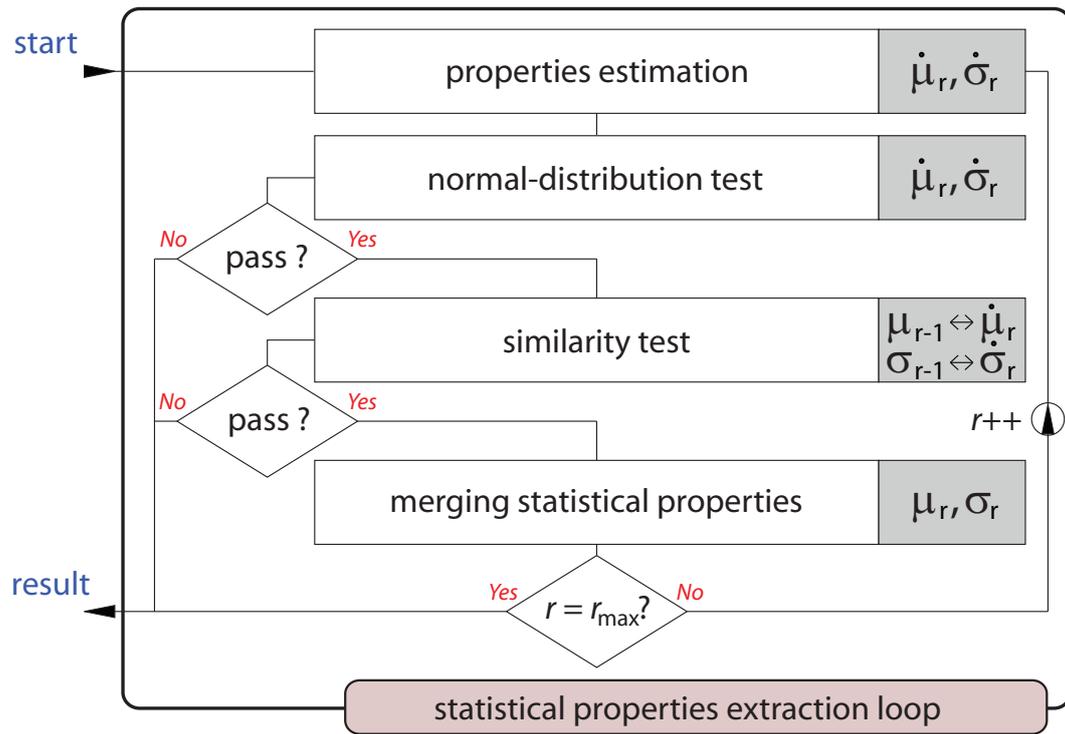


Figure 2.4: Calculation loop for the extraction of the statistical properties.

standard deviation of all points:

$$\dot{\mu}_r = \frac{1}{\dot{N}_r} \sum_{i=1}^{\dot{N}_r} f_i, \quad (2.1)$$

$$\dot{\sigma}_r = \sqrt{\frac{1}{\dot{N}_r} \sum_{i=1}^{\dot{N}_r} (f_i - \dot{\mu}_r)^2} \quad (2.2)$$

\dot{N}_r is the number of points included in the hull r . f_i denotes the data value of a sample point in the hull. With these estimations for the mean value and standard deviation we expect to get values which are close to the real statistical properties of the material in the hull.

Normal-Distribution Test

Before we apply a similarity test with the derived properties of the hull and the statistical properties of the inner sphere, it must be ensured that the distribution is normally distributed. This is necessary because the similarity test is based on normal distributions. When two materials are intersecting the hull, the distribution is not normally distributed. In such a case the distribution would have two peaks. With the normal-distribution test we want to detect such situations and terminate the loop.

In statistics, several normality tests exist. We chose the Jarque-Bera test (JB). This test uses the third-order moment, i.e., the skewness \hat{S}_r , and the fourth-order moment, i.e., the kurtosis \hat{K}_r , of the points in the hull at radius r for the calculation of JB :

$$JB = \frac{\hat{N}_r}{6} \left(\hat{S}_r^2 + \frac{(\hat{K}_r - 3)^2}{4} \right) \quad (2.3)$$

The benefit of using this test is that it is not necessary to sort the point values of the distribution. Therefore, it can be efficiently implemented on the graphics hardware.

The parameter JB is used to test for the null hypothesis with a test level $1 - \alpha$ of 99.9%. This results in a test value of 13.82 according to statistical lookup-tables:

$$H_0 : JB < 13.82_{(1-\alpha=0.999)} \quad (2.4)$$

With the high test level, only distributions with a high divergence to a normal distribution are declined by the null hypothesis. This is necessary, due to the low number of samples which are used for the test. Only distributions which are very different from a normal distribution will fail the test.

If the null hypothesis is declined, the loop is terminated and μ_{r-1} and σ_{r-1} of sphere $r - 1$ are taken as statistical properties for the sample point S . If the test passes the null hypothesis, we continue with the similarity test.

Similarity Test

In this processing step, we measure the similarity between the statistical properties in the hull r and the statistical properties of the sphere $r - 1$. The goal is to detect whether the hull is still part of the same material as the sphere $r - 1$.

Through the properties-estimation step, we get the statistical properties of the hull. We have calculated the statistical properties of the sphere $r - 1$ in the preceding loop cycle. Since all the values are estimations and the numbers of points which are involved in the estimation is rather low, we use a variant of the student's t-test for the similarity test. This test is best suited for Gaussian-distributed populations with small sample sizes [19].

In our case, we have two independent samples which were used for the estimations. Since the mean values as well as the standard deviations of both estimations can vary, we use a generalized form of the student's t-test also known as the Welch's t-test [77]. As primary parameter for the similarity test, a t_r parameter is calculated:

$$t_r = \frac{\mu_{r-1} - \hat{\mu}_r}{\sqrt{\frac{\sigma_{r-1}^2}{N_{r-1}-1} + \frac{\hat{\sigma}_r^2}{\hat{N}_r-1}}} \quad (2.5)$$

The t_r parameter is dependent on the mean values, the standard deviations, and sample sizes of both distributions. Additionally, a degree-of-freedom δ_r has to be calculated:

$$\delta_r = \frac{\left(\frac{\sigma_{r-1}^2}{N_{r-1}-1} + \frac{\hat{\sigma}_r^2}{\hat{N}_r-1} \right)^2}{\frac{\sigma_{r-1}^4}{(N_{r-1}-1)^3} + \frac{\hat{\sigma}_r^4}{(\hat{N}_r-1)^3}} \quad (2.6)$$

The degree-of-freedom δ_r is only dependent on the standard deviations and the sample sizes, but not on the mean values. The δ_r value together with the confidence level ω are used as parameters to get a reference $t_\omega(\delta_r)$ value in a t-test lookup-table. This value is used to test for the null hypothesis H_0 :

$$H_0 : |t_r| < t_\omega(\delta_r) \quad (2.7)$$

If the null hypothesis is true, it is assumed that both Gaussian distributions are the same. If the null hypothesis is declined then both distributions are expected to be different with a probability of $1 - \omega$. Therefore, a small confidence level ω results in a high probability that both distributions are not similar if the null hypothesis is declined. On the other hand, the reference $t_\omega(\delta_r)$ value for a small ω is high which makes the similarity test less selective.

As in the step earlier μ_{r-1} and σ_{r-1} of sphere $r - 1$ are taken as statistical properties if the test is failed. Otherwise we continue with the next growing step.

Merging Statistical Properties

If the statistical properties have passed the normal-distribution test and the similarity test, we assume that the material in the outer hull still is the same as in the sphere $r - 1$. Therefore, the statistical properties of both areas can be merged together.

This step results in a new μ_r and σ_r . These statistical properties represent the distribution of all points in the sphere r . The merged statistical properties are used in the successive cycle of the loop to do the similarity test with the next larger hull with radius $r + 1$.

The loop is terminated when the normal-distribution test or the similarity test fails or when the maximum radius is reached. In the first two cases we store μ_{r-1} and σ_{r-1} as statistical properties μ and σ for the sample point S . In the third case we take the statistical properties μ_r and σ_r after the merging step.

Additionally, we store the radius r_{break} at which the loop is terminated. The closer r_{break} is to r_{max} the more significant the statistical properties are at this point, because the population of points for the estimation is larger. In the next section, r_{break} is used to highlight statistical properties with a higher significance.

Figure 2.5 shows the statistical properties μ and σ for the synthetic data set at different confidence levels ω . For a low ω of 0.1%, the similarity test is more easily passed. Therefore, the sphere grows larger and results in more consistent values for μ and σ . For the transfer-function space this means that the clusters of materials are smaller. If the confidence level ω is high, e.g., 30%, the loop has a higher probability of being terminated. Since in this case a smaller area is used to estimate the statistical properties, the results for μ and σ are less smooth compared to a low ω but details, such as borders, are preserved better.

Congruent to the characteristic in this example, ω should be chosen according to the type of data. If a modality is very noisy and the distribution of points does not exactly follow a normal distribution, the confidence value should be chosen rather low. Therefore, some details get lost but the clusters for different materials in the transfer-function space are smaller. MRI is an example of such data. In such a case, ω should be set to a low value, such as 0.1%. For less noisy data types with a Gaussian-like distribution of the data values, e.g., CT, ω can be set to a higher value,

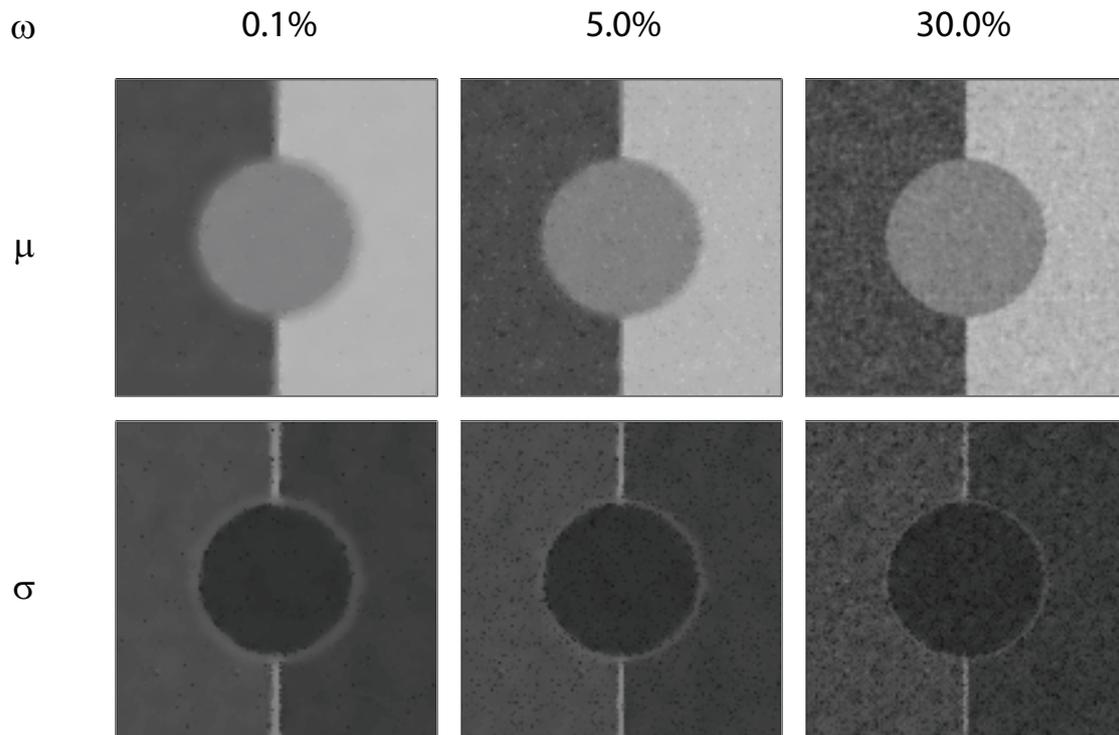


Figure 2.5: Statistical properties (μ and σ) of the synthetic data for three different confidence levels $\omega = \{0.1\%, 5.0\%, 30.0\%\}$. The brighter a point is, the higher is its value.

such as 20%. In this case details are better preserved and the clusters are small enough due to the low noise level. The synthetic data set is rather noisy but the distribution of the points follows exactly a normal distribution. Therefore, a confidence level of 5% was chosen. The values for the confidence levels were found through experiments, where we tried to achieve the best balance between noise reduction and detail preservation. This has to be done only once for a certain data type. In Section 2.4 we demonstrate that the selected ω -values work well for different MRI and CT data sets.

2.2.2 Definition of Statistical Transfer-Function Space

By extracting the statistical properties we get a mean value μ and a standard deviation σ for each sample point. In the next step we use this information for the design of a transfer function. First we have to define a transfer-function space, which is used for the design of a transfer function. The intent of this statistical transfer-function space is to separate different materials in the presence of noise.

For the transfer-function space we use the original data value f of each sample point together with the mean value μ and the standard deviation σ . We follow the convention of using the horizontal axis of a transfer-function space for the data value. The horizontal axis is, however, also used to depict the mean μ . In the new transfer-function space the data value f on the horizontal

axis is considered as starting point of a line segment. The statistical properties μ and σ for each sample point define a second point, where the standard deviation is given on the vertical axis.

In Figure 2.6 the transfer-function space is shown. On the left side the features of a sample point S are drawn as a line segment in this space.

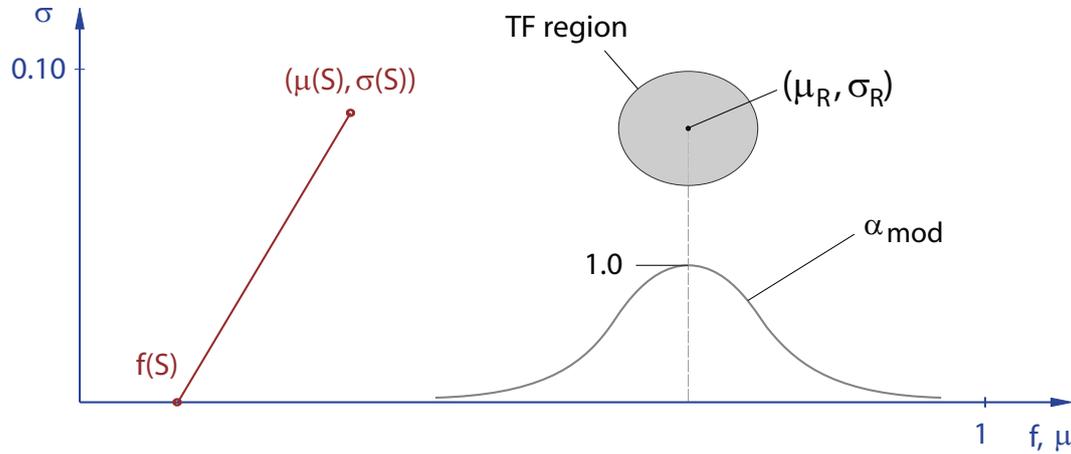


Figure 2.6: Schematic representation of the statistical transfer-function space. On the left side the features of a sample point are shown as a line segment. On the right side a transfer-function region in the new space can be seen.

Since the clusters for each material are expected to show up around a certain mean value and standard deviation, an elliptical area is used to define a transfer-function region. On the right side of Figure 2.6 an example of such a region is shown. The center of the region is at (μ_R, σ_R) . A sample point is only classified by this region when its statistical properties (μ, σ) lie within the elliptical area.

The optical properties for each region are represented with a color c_R and an opacity α_R . The position of the transfer-function region (μ_R, σ_R) also defines a corresponding Gaussian distribution of data values along the horizontal axis. The opacity of sample points is manipulated according to this distribution:

$$\alpha_{mod} = \alpha_R * e^{-\frac{1}{2} \left(\frac{f - \mu_R}{\sigma_R} \right)^2} \quad (2.8)$$

α_R is the opacity defined for the respective transfer-function region. f is the data value of the sample point and (μ_R, σ_R) is the center point of the region. Equation 2.8 is represented as a curve along the horizontal axis in Figure 2.6, below the transfer-function region. The opacity for data values f close to the center μ_R of the region are modified less in contrast to the opacity of data values f , which are more different to μ_R . By this modification, we reduce the influence of sample points which are less likely part of the desired distribution.

Figure 2.7 shows the sample points of the synthetic data set in the statistical transfer-function space. In Figure 2.7(a) each sample point is represented as a line segment. Since the line segments of points with high σ values might occlude points with low σ values, an alternative representation is shown in Figure 2.7(b). In this representation only a single dot at (μ, σ) is drawn for each

sample point. In both representations three large clusters for the different materials and three smaller clusters for the borders can be seen. The border between material 1 and 3 actually consists of two small clusters. This results from the fact that the mean values of material 1 and 3 are very different. Therefore, it makes a difference if the original sample point of a neighborhood-region lies in material 1 or 3, because more points from one or the other material are then used for the estimation of the mean value.

If the difference between the mean values of two materials is smaller, e.g., between material 1 and 2, then only one cluster shows up in the transfer-function space.

To enhance points with higher significance, we weight the opacity of line segments or dots by the termination radius r_{break} of the calculation loop. Thus, statistically well defined interior regions are emphasized.

2.2.3 Statistical Properties in Visualization

In addition to the definition of a transfer function we use the statistical properties as input for the shading process. In volume visualization gradient-based techniques are most common for shading. They are computationally less expensive and faster as gradient-free shading techniques. For noisy data the gradient-based techniques have the disadvantage that the noise in the data deteriorates the estimation of the gradient and, therefore the shading [5]. Especially in homogeneous regions, where the gradient magnitude is rather low, noise has a high impact on the estimation of the gradient direction.

To reduce the influence of noise on the gradient, we use the mean values of the sample points to estimate the gradient direction. The mean values are smoother than the original data values and, therefore, the gradient direction in nearly homogeneous regions is estimated better in comparison to the widely used central-difference method. With this gradient we calculate a color c_{shaded} based on Phong shading [61].

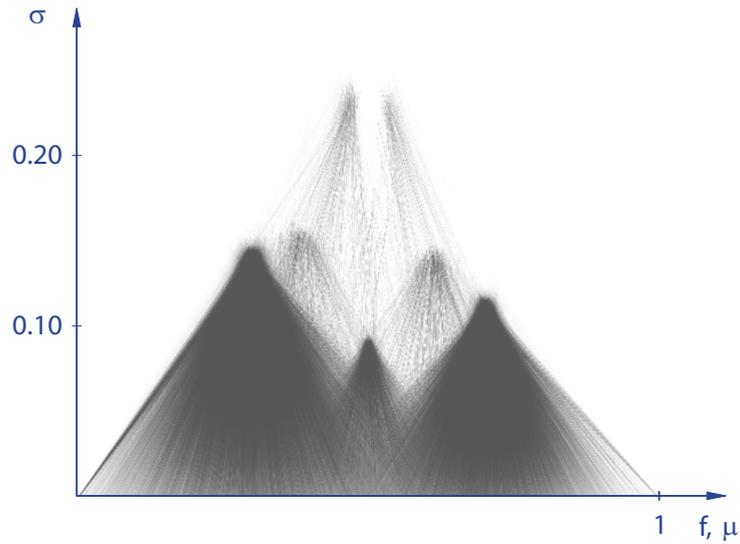
For the observer, shading is important for the perception of surfaces. To avoid visual clutter resulting from shading all parts of the volume, we apply shading only for border regions. For this purpose the standard deviation σ can be used. σ is higher in border regions and it is less affected by noise than the gradient magnitude. Therefore, σ is used to interpolate between the shaded color c_{shaded} and the unshaded color $c_{unshaded}$:

$$c = (1 - \sigma)c_{unshaded} + \sigma c_{shaded} \quad (2.9)$$

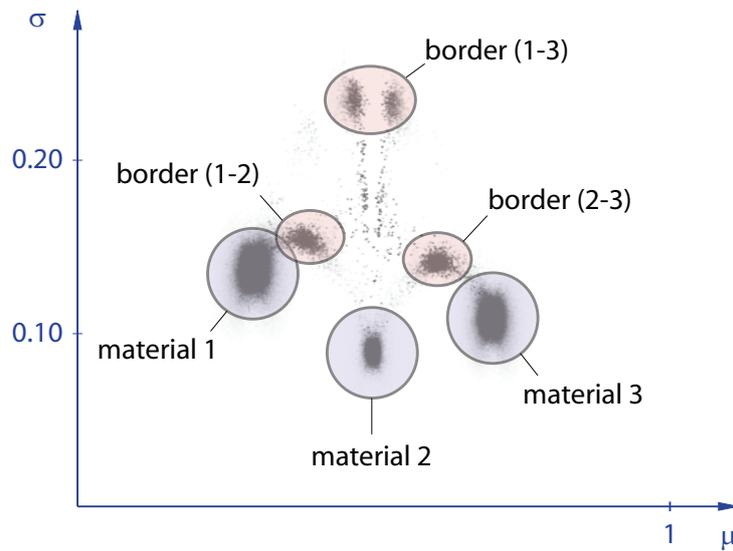
The lower σ is, the less shading is applied. This leads to a visualization where the border areas are shaded more as opposed to the rest. The synthetic data with traditional shading in Figure 2.8(a) is shown in contrast to the shading based on statistical properties in Figure 2.8(b). With statistical shading the influence of noise is clearly reduced without modifying the data itself by filtering or other techniques.

2.3 Implementation

The steps which are described in Section 2.2 can be divided into two parts of computation. The first part is the extraction of statistical properties. This is a pre-processing step which has to be



(a) Line representation



(b) Dot representation

Figure 2.7: Line and dot representation in the statistical transfer-function space of the synthetic data.

recalculated only when the confidence level ω is changed. The other part is the visualization with the usage of the statistical transfer function. This part is done in real-time.

Since the extraction of the statistical properties is a highly parallel process, we use graphics hardware. Nevertheless, the extraction is expensive because many lookups have to be performed to estimate the statistical properties.

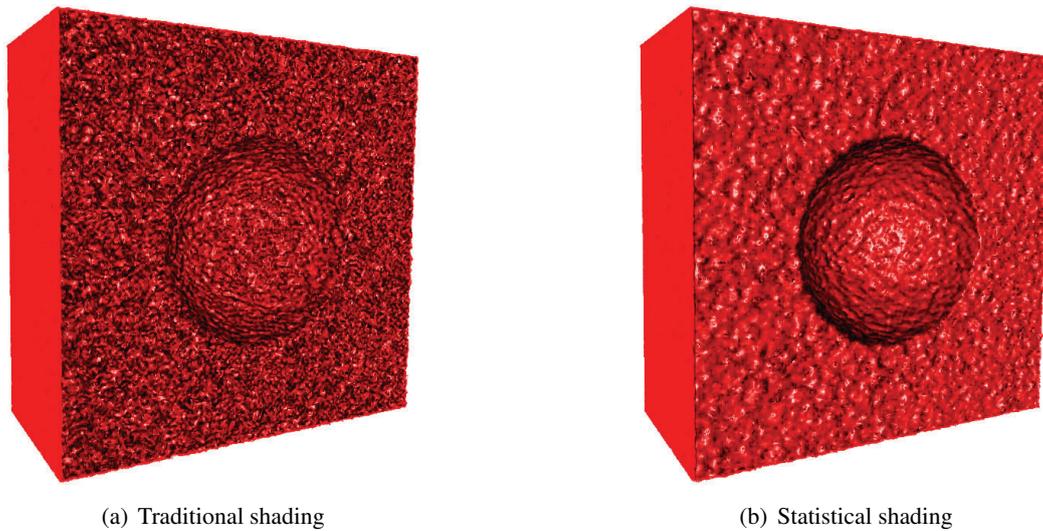


Figure 2.8: Visual difference between traditional shading and statistical shading. With statistical shading the influence of noise is reduced.

The maximum number of lookups per sample point is dependent on the maximum radius r_{max} of the calculation loop. We decided to set the maximum radius to a value of 6. With this radius, a maximum of 925 lookups are done per sample point. This leads to the following 99% confidence limits:

$$\mu = \bar{\mu} \pm (0.0424\sigma) \quad (2.10)$$

This means that the estimated μ is, with a confidence of 99%, not farther away from the real mean value $\bar{\mu}$ than 0.0424 times σ . Even for a large σ the interval for μ is small enough for our purpose of estimating the statistical properties.

In addition to the maximum radius r_{max} , the confidence level ω influences the speed of the calculation. The lower ω , the less likely it is that the calculation loop is terminated. Therefore, more lookups have to be done for the estimation of the statistical properties. In Figure 2.9 the termination level at different radii is shown. The graph shows the percentage of all sample points for which the calculation loop is already terminated at a certain radius. For the synthetic data set different confidence levels are used. It can be seen that the termination level at the maximum radius of 6 is much lower for an ω of 0.1% in comparison to an ω of 30%. For the low ω the loop is only terminated for points close to the border. A high ω causes terminations also for large variations in homogeneous regions. As shown in Figure 2.5, the visual difference of various ω values can be seen in the smoothness.

Apart from the synthetic data set, Figure 2.9 also shows the termination level for two real-world data sets with according confidence levels for each type of data. In the case of CT data, we get a high termination level even at a low radius. This results from the fact that areas with zero variation, e.g., air, do not pass the initial normal-distribution test. For MRI data, the curve is more linear, due to the fact that all parts contain at least some noise.

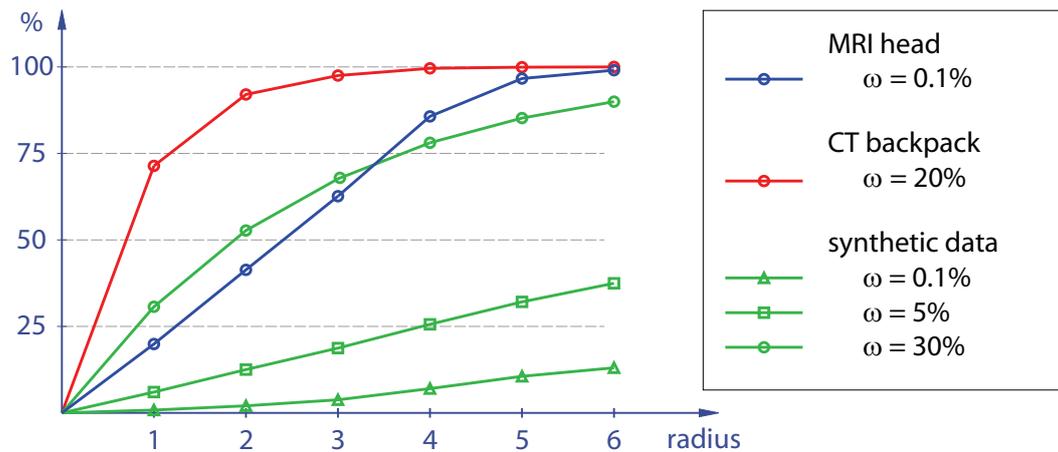


Figure 2.9: Termination level of the calculation loop for different data sets and confidence levels ω .

Table 2.1 shows the time measurements for the estimation of the statistical properties for the data sets of Figure 2.9 on a GeForce GTX 260. The different settings for the level ω have only a small influence on the calculation time. This results from the fact that most time is used to initialize the graphics hardware for the calculation.

MRI Head	CT Backpack	Synthetic Data		
256×256×128	256×256×186	128×128×128		
$\omega=0.1\%$	$\omega=20\%$	$\omega=0.1\%$	$\omega=5\%$	$\omega=30\%$
2.386 _{sec}	2.330 _{sec}	0.654 _{sec}	0.649 _{sec}	0.646 _{sec}

Table 2.1: Time measurement for the estimation of statistical properties for different data sets.

After the extraction of the statistical properties, we use raycasting for the visualization. The statistical properties μ and σ are stored in additional channels of the volume. To classify a sample point with the designed transfer function, the parameters of all transfer-function regions are handed over to the graphics card. There it is tested if the statistical properties of a sample point lie within a transfer-function region. If so, the color and opacity is assigned to the sample point as described in Section 2.2.2. This classification can be implemented efficiently on the graphics card. We get interactive rendering rates for data sets of size $256 \times 256 \times 256$ on a GeForce 8800 GT graphics card.

2.4 Results and Discussion

In this section we show some results generated with our new method and compare them with other techniques.

2.4.1 Synthetic Data Set

For the explanation of the method we have introduced a synthetic data set, as shown in Figure 2.2. The three different materials in this data set are rather noisy. Therefore, it is difficult to separate the materials in common transfer-function spaces. In Figure 2.10(a) the 2D transfer-function space with axis f and $|f'|$ and in Figure 2.10(b) the LH histogram-space were used to classify the different materials. With the 2D transfer function we were not able to classify all points correctly, because of the density overlapping. Especially at the border between material 1 and material 3 (blue and yellow) points are classified as material 2 (red). In the LH histogram-space it is easier to separate the different materials and the border but transitions are very ragged.

In Figure 2.10(c) the result of our method is shown for an ω of 5%. Since the different materials have different statistical properties they can be clearly seen as clusters in the statistical transfer-function space. For the synthetic data, smoothing techniques would be able to reduce the cluster sizes in the 2D transfer-function space and in the LH histogram-space. However, the smoothing only clusters the data values. Our approach uses the standard deviation for the classification as well. Therefore, we are better able to classify different materials especially if they differ by their standard deviation, such as at border regions.

2.4.2 Real-World Data Sets

In the real world, noise is typically present in measured data sets. The amount of noise varies between acquisition techniques. In MRI data sets, the noise level is rather high. Therefore, it is especially difficult to classify different materials in such data sets. A common problem is the visualization of the brain in an MRI scan of the head. Figure 2.11 shows different results of this task for different classification techniques. Below each rendering result, the settings of the transfer function for each of these spaces are shown. In comparison to the 1D (Figure 2.11(a)) and 2D transfer function (Figure 2.11(b)) as well as the LH histogram-space (Figure 2.11(c)), we can better separate the brain from other tissues with our method (Figure 2.11(d)). In the statistical transfer-function space it is also easier to design a transfer function because the brain tissue has different statistical properties than other tissues in this data set and is more tightly clustered. As can be seen in Figure 2.11(d) the cluster is rather large due to the different matters in the brain but it is distinguishable from other clusters of other tissues. In comparison to this, in the 2D transfer-function space and in the LH histogram-space no cluster for the brain tissue can be seen. Additionally, Figure 2.11 shows that shading based on statistical properties (Figure 2.11(d)) results in a smoother surface in comparison to normal gradient-based shading (Figure 2.11(a)-(c)).

The ability of the new method to separate different materials by their statistical properties can be used for many applications. An example is the detection of certain materials in CT scans for security checks. Figure 2.12 shows the scan of a backpack containing three different fluids. With a 2D transfer function, as shown in Figure 2.12(a), the different fluids cannot be classified without classifying also other parts of the data set. It is also hard to detect the fluids in the transfer-function space because they do not show up as clusters. In comparison, our method can clearly classify the fluids, as shown in Figure 2.12(b). It can also be seen that the fluids show up as clusters with very low standard deviation in the transfer-function space. This makes it

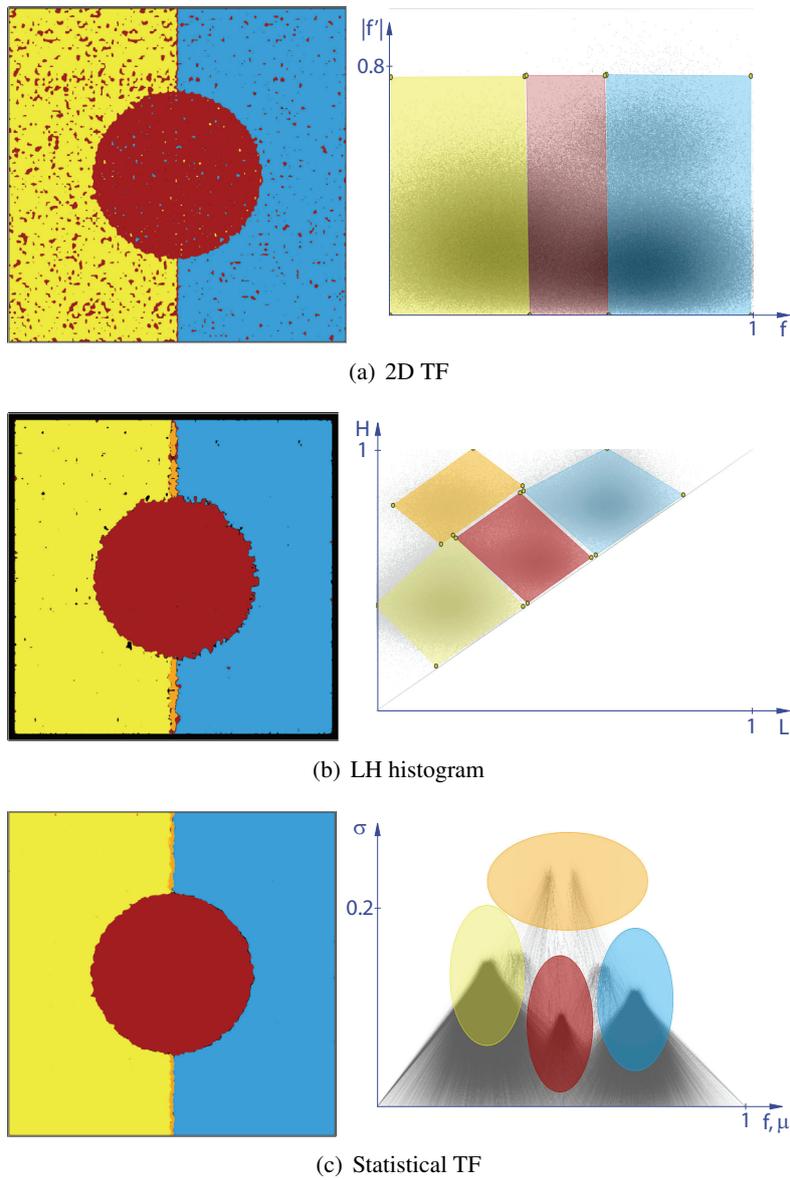


Figure 2.10: Classification results of the materials in the synthetic data set with different methods.

much easier to define a transfer function. Furthermore, the shading with our method is slightly smoother.

Figure 2.13 shows a result of an MRI scan. The data set contains a tumor inside the brain. The statistical properties of the tumor are actually different from the rest of the brain which is captured in our transfer-function space. This can be seen in Figure 2.13. Since the tumor is rather small, only a few sample points show up in the area of the classification region of the tumor (red

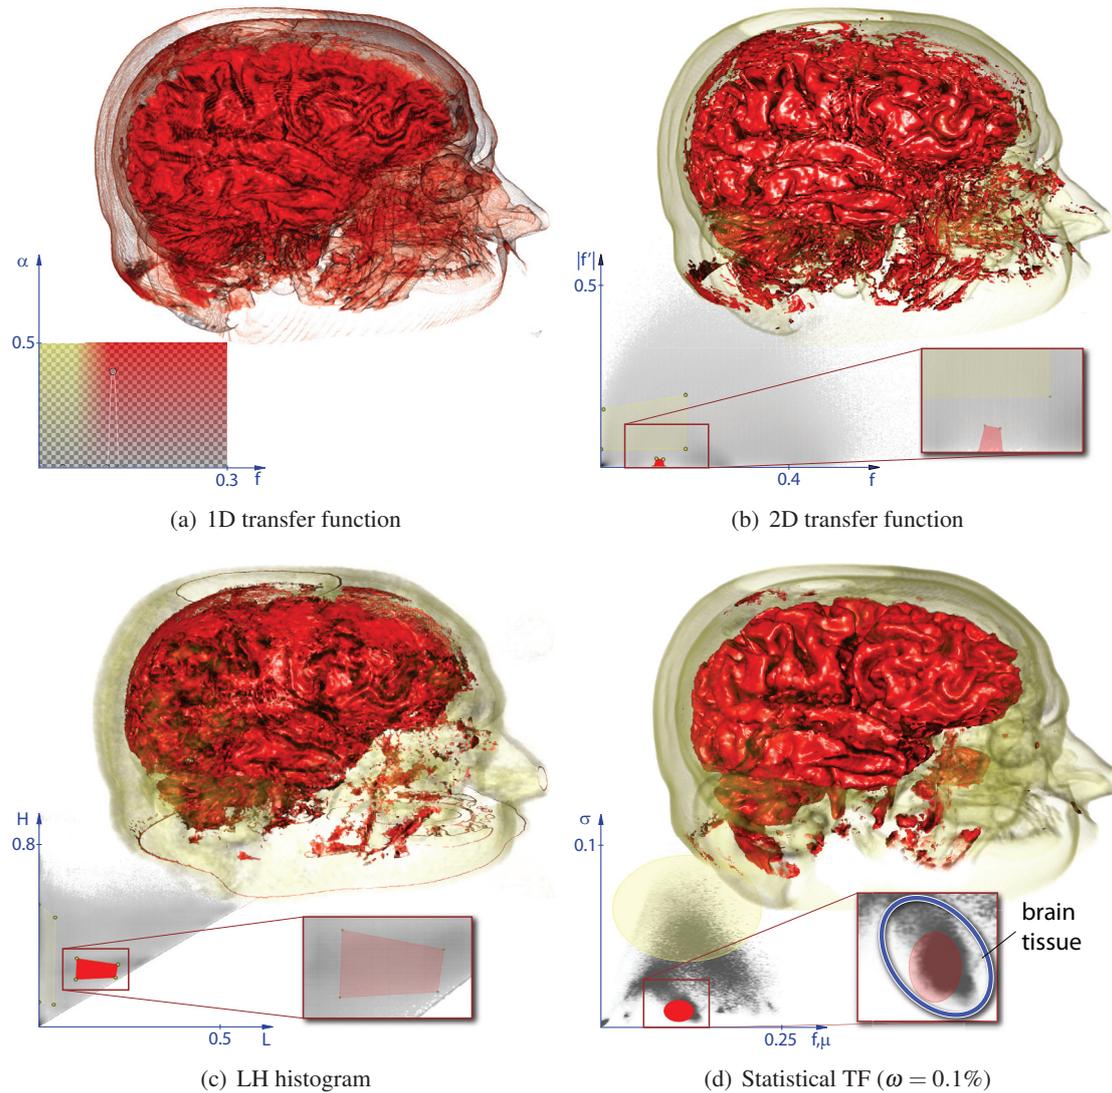


Figure 2.11: Comparison of the statistical transfer-function space with the 1D and 2D transfer function space as well as the LH histogram-space. The task for the generation of the results was to classify the brain in the different spaces.

region). However, with other methods, such as 1D, 2D, and LH transfer functions, we were not able to clearly separate the tumor from the brain.

The results show that the new method can be used for various data sets and different tasks. The main reason for this is the confidence level ω , which can be set according to the type of data. For MRI, e.g., where the material distributions slightly differ from a Gaussian distribution, we set ω to a low level in comparison to CT data. We are able to classify different materials even if the data type is different. This is not as easily possible in other common transfer-function spaces.

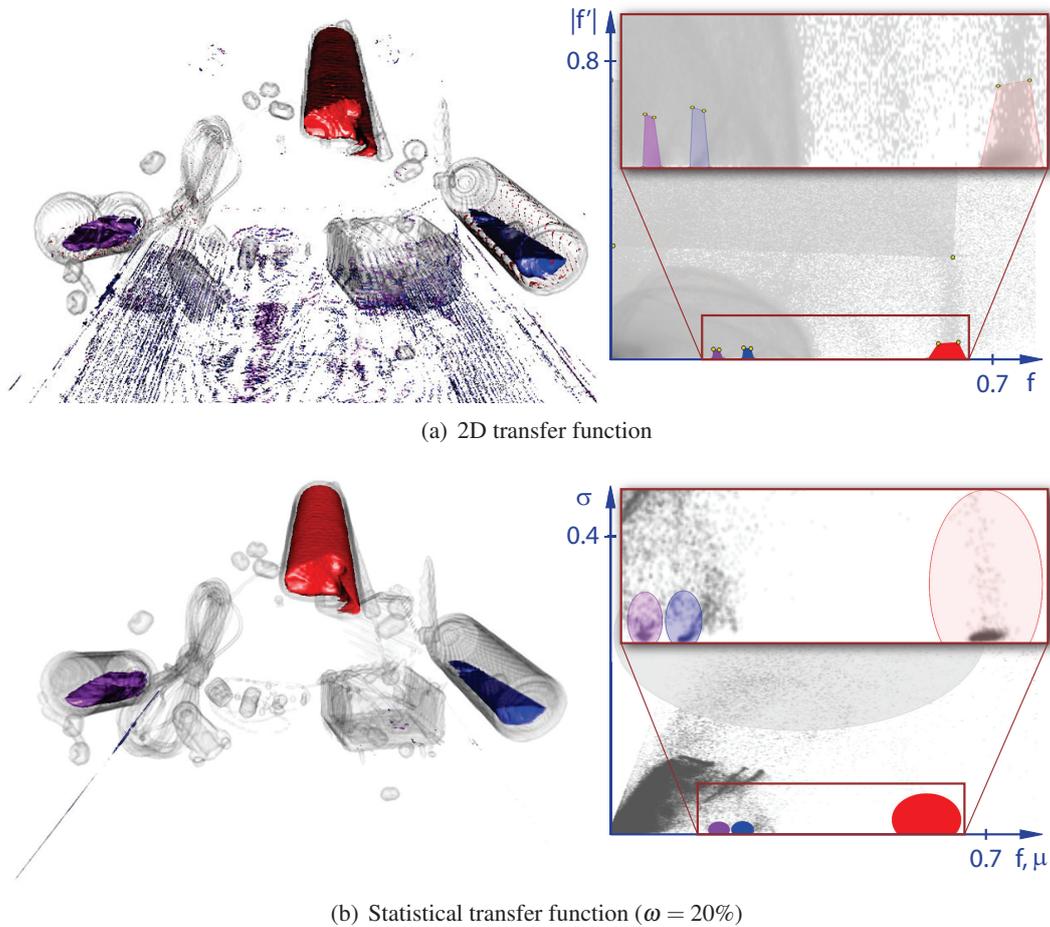


Figure 2.12: Detection of different fluids in a CT scan of a backpack.

There are some limitations of the new technique. One drawback is the rather high memory consumption because for each sample point two different statistical properties have to be stored together with the data value. Thus the data size is tripled. For large data sets this could exceed the memory of a graphics card. Another penalty can occur for noise distributions very different from Gaussian white noise. In such cases the test methods have to be adapted to the given frequency distribution in the data sources. For this approach we concentrated on measured data, where the distributions of data values are similar to a Gaussian distribution.

Although the confidence level is the only parameter which has to be set by the user, this parameter should be defined automatically. This should be possible by using a termination level. With the termination level at different radii it is probably possible to detect if ω is either too high or too low for a given data set. Furthermore, an automatic detection of clusters in the statistical transfer-function space can be implemented in a future work. By using the exit radius r_{break} of the calculation loop, it should be possible to automatically find significant clusters for different

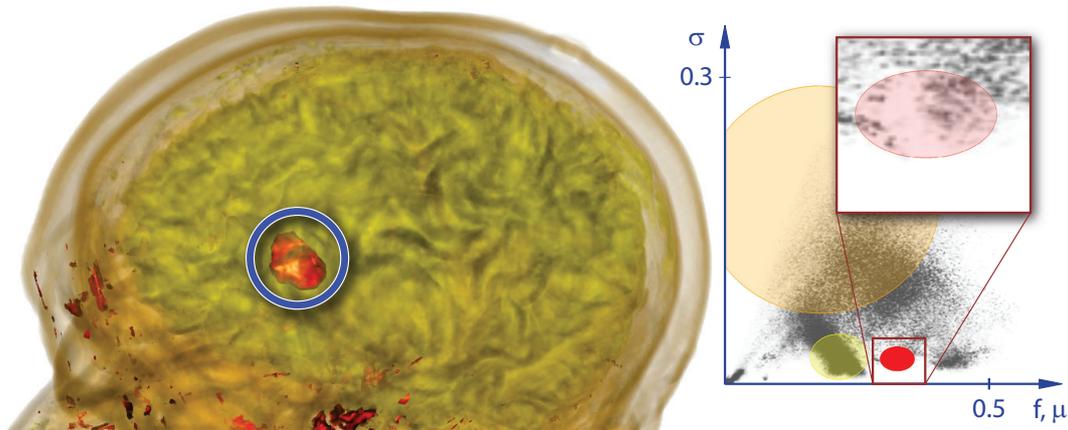


Figure 2.13: Detection of a brain tumor. For the extraction of the statistical properties ω was set to 0.1%.

materials. This additional step should accelerate the design process for transfer functions because an initial setting can be provided.

2.5 Conclusion

In this chapter a method was presented which uses the local neighborhood around each sample point to extract some features based on statistics. With the employment of statistics it is possible to reduce the influence of small variations in the data, such as noise or other small artifacts, in the classification process. The statistics is able to reveal some information about the materials represented in the data which cannot be detected by considering only local features, such as the data value alone.

The statistical properties were used to define a novel transfer-function space. In this transfer-function space it is possible to better separate different materials in comparison to other widely used transfer-function techniques. Even for very noisy data sets, such as MRI data, the materials are still distinguishable by their statistical properties.

Additionally, the statistical properties are used in the shading calculations. Through this approach, the influence of noise on the shading is reduced as well.

In our experiments, the novel transfer-function space has proven to produce better results for different tasks than other common transfer-function techniques. Therefore, we believe the statistical transfer-function space can be used for classifying different materials in a volume data set. Through the user-specified confidence level it can be employed for data sets from various modalities.

In the next chapter a method is described which uses the value distribution in the entire data set to extract features for the classification. Through this it is possible to extract different information about the data in comparison to features which are extracted in the local neighborhood.

No man was ever so completely skilled in the conduct of life, as not to receive new information from age and experience.

— Jonathan Swift

CHAPTER

3

Information-based Transfer Functions for Multimodal Classification

Martin Haidacher, Stefan Bruckner, Armin Kanitsar, and M. Eduard Gröller. *Information-based Transfer Functions for Multimodal Visualization*. In VCBM 2008, Pages 101-108, 2008

THE classification of data should enhance parts of the data which contain most information for a certain task. In the previous chapter a method was introduced which is able to classify different materials in a data set. Through this it is possible to assign different optical properties to different materials in a novel transfer-function space. The features for the transfer-function space were extracted in the local neighborhood around each sample point. The method is adapted for the classification of a data set from a single modality.

Currently, the trend towards data acquisition using data sets from multiple modalities is increasing in order to facilitate, e.g., better medical diagnosis. As different modalities frequently carry complementary information, the goal of the method in this chapter is to provide the user with a consistent interface for the enhancement of complementary information in the different modalities through the classification process.

Normally a side-by-side view is provided in medical applications for the inspection of the different modalities. A physician can simultaneously scroll through both registered modalities. This practice has two main drawbacks. One is the missing direct visual combination of the data. A physician has to mentally overlay the two images to get the corresponding points of one modality in the other one. A second drawback is the restriction to a 2D visualization. These drawbacks can be eliminated by the fused display of both data sets together in a 3D multimodal visualization. The challenge for such a visualization is the density of information in space. For each sample point at least two values from the different modalities are present. To reduce the density a transfer function can be used which defines optical properties, such as color and opacity, for certain values. The transfer function can be controlled by the user to change the appearance of the result image. The more input values are taken to classify a sample point and assign optical properties to it, the harder it is for the user to find a good transfer function. This is the main problem of multimodal visualization because there are at least two values involved.

In this chapter, we introduce a novel concept for defining transfer functions in multimodal volume visualization. Our method aims to reduce the complexity of finding a good transfer function. A new transfer-function space is provided which can be controlled by the user in an intuitive and familiar way. The information-theoretic features which are used to define the

transfer-function space are extracted from the global value distributions of all voxels in the data set.

Since the design of a transfer function is a non-trivial task, we designed the transfer-function space in a way which is similar to well-known transfer-function spaces. Hence the design process of a transfer function is easier in comparison to a completely new transfer-function space. To achieve this, the values of both modalities are fused based on the features which are extracted through information theory. This results in a fused transfer-function space with a single value and a single gradient magnitude as parameters. A measure for the complimentary information of both modalities is used as additional parameter for more user control and a better separation of different materials in the classification process.

In Section 3.2 the approach is described in detail. We show how the retrieved information of the value distribution can be used to generate the transfer-function space. Section 3.3 briefly describes an efficient implementation of the new method. The usability of the new method is shown in Section 3.4 with some results. Conclusions and ideas for further work are given in Section 3.5. First an overview of related works on this topic is given in the following section.

3.1 Related Work

All different methods for multimodal visualization can be classified - as described by Cai and Sakas [8] - according to the level in the rendering pipeline in which they are applied. In the illumination-model-level intermixing optical properties are assigned to a combination of values from the different modalities. The accumulation-level intermixing fuses the values after optical properties are assigned to each modality individually. In the image-level intermixing the fusion is done after the 2D images have been rendered.

The image-level intermixing is the simplest way for the fusion of two modalities, but it has the disadvantage that the 3D information is lost. Therefore this fusion technique is typically just applied on single slices of the volume. Several techniques have been developed for this purpose, such as alternate pixel display or linked cursor [64, 66]. These techniques are covered by the field of image processing [76].

Due to the increasing speed of computers and graphics hardware volume rendering became more popular and, therefore, also the multimodal fusion could be done in the volume space. The first methods were based on surface models. Levin et al. [48] generated a surface model from an MRI scan and mapped the *Positron Emission Tomography* (PET)-derived measurement onto this surface. Evans et al. [18] generated an integrated volume visualization from the combination of MRI and PET. These works are mainly focused on the combination of anatomical and functional images. A more general approach for the fusion of modalities was introduced by Zuiderveld and Viergever [83]. For this method an additional segmentation of the volumes is necessary to decide which one to show at a given sample point. A more recent work by Hong et al. [30] describes how fusion techniques in this intermixing level can be efficiently implemented using the graphics hardware.

More sophisticated but more complex methods for multimodal visualization are directly applied in the illumination-model-level. The intermixing in this level directly generates optical properties from the combination of the values and additional properties of the two volumes

at a single sample point. A case study for the rendering of multivariate data where multiple values are present at each sample point was done by Kniss et al. [39]. 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 shown by Kniss et al. [42]. The classification is done on the basis of a dual histogram which depicts the co-occurrence of value pairs from both modalities. Kim et al. [36] presented a technique which simplifies the transfer-function design by letting the user define a separate transfer function for each modality. The combination of them defines the two-dimensional transfer function. The problem with this technique is the loss of information by reducing the multi-dimensional transfer function to two 1D transfer functions.

As mentioned before, the assignment of optical properties in multimodal visualization is dependent on more than one value. If the whole information space is used then a multi-dimensional transfer function is needed. In general it is a non-trivial task to design a multi-dimensional transfer function because of its complexity. Nevertheless, multi-dimensional transfer functions are commonly used for volume visualization. 2D transfer functions were first introduced by Levoy [49]. In addition to the data value the gradient magnitude was used as second dimension to classify a sample point. Due to the fact that the design of a 2D transfer function is non-trivial, methods were developed, to support this task. Kindlmann and Durkin [37] introduced a semi-automatic approach for the visualization of boundaries between tissues. Pfister et al. [60] gave an overview on existing techniques to support the design task of transfer functions. The direct manipulation widgets introduced by Kniss et al. [38] can be used to find regions of interest in the multi-dimensional transfer-function space in an intuitive and convenient way. In other work, Kniss et al. [41] describe an approach to efficiently represent multi-dimensional transfer functions by Gaussian functions instead of storing a multi-dimensional lookup table.

For the definition of the multi-dimensional transfer functions, in addition to the values from the two volumes, further properties can be used to better distinguish between tissues. In this chapter, these additional properties are retrieved by methods from information theory founded by Shannon [65]. He described how the probability of occurrence of a signal can be used to define the information content of the signal. In imaging, information theory is used in different areas. Image registration is one of these areas. Wells et al. [78] maximized the mutual information to find a good registration position for two images or volumes. This idea is the basis for the information-based part of the new approach in this chapter.

Rezk-Salama et al. [62] employed principal component analysis to assist the generation of more effective transfer functions based on semantics. Our approach provides additional derived quantities for evaluating the joint information of multiple modalities. In future work, a combination of both methods could lead to even more intuitive user control for multimodal volume visualization.

3.2 Information-based Transfer Functions for Multimodal Volume Classification

In this section we introduce a novel transfer-function space for multimodal volume classification. The aim of all steps described here is the design of a transfer function space which is as simple as possible but still able to separate different tissues or materials. The main contribution of the new approach is the use of methods from information theory for the design of this transfer-function space. Figure 3.1 shows all necessary processing steps to classify a tuple of input values (\dot{f}, \ddot{f}) in this new transfer-function space with optical properties. The further sections describe these processing steps in detail.

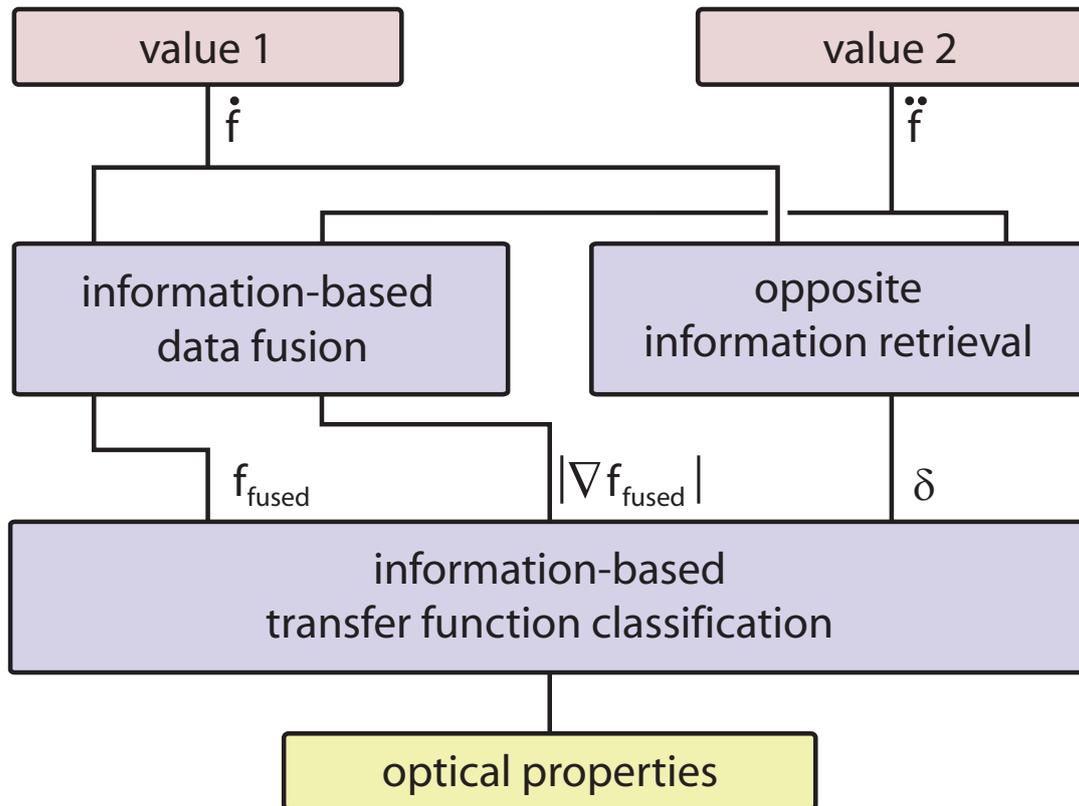


Figure 3.1: Processing pipeline for the classification of sample points in a multimodal visualization by an information-based transfer function.

In Section 3.2.2, we describe how the input values can be fused to get just a single value for each pair of input values. Section 3.2.3 introduces an additional property δ which is employed to refine the classification of different tissues through the transfer function. Finally, Section 3.2.4 describes how the fused values are used to define the new transfer-function space and how the additional property δ is used to influence the classification. In addition to the introduction to

information theory in Chapter 1, we will first describe in more detail in Section 3.2.1 how the probabilities for the further calculations are estimated.

3.2.1 Probabilities in Volume Data

To estimate the probabilities within a volume we first assume that the volume is given as a set of regularly arranged grid points. The simplest way to estimate the probability of a certain value in such a volume is done by counting its occurrences in the whole data set and by dividing this number by the total number of points in the volume. To do this for all values a histogram is generated. In the histogram the count of a bin is increased if a value falls into the range of this bin. When the counted numbers for all bins are divided by the total number of points in the volume, we get a probability distribution $p(f)$ which returns a probability of occurrence for each value f .

For retrieving the information content of the joint occurrence of two values from two modalities another probability distribution is needed. It returns a probability $p(\hat{f}, \check{f})$ for each tuple of values \hat{f} from modality 1 and \check{f} from modality 2, also referred to as joint probability. Analog to the probability for the occurrence of only one value this probability distribution can also be estimated by a histogram. Due to the dependency of two values, the histogram is defined in 2D. This histogram is often referred to as dual histogram.

In the context of the joint probability $p(\hat{f}, \check{f})$ the probability of just a single value $p(\hat{f})$ is referred to as marginal probability. These two types of probabilities are further used in the following sections to generate a new transfer-function space based on the methods of information theory.

3.2.2 Information-based Data Fusion

At some point in a multimodal visualization pipeline the information from both data sets has to be combined, as each sample point can only have one color and opacity. The idea behind the information-based data fusion is to have an approach which loses as little as possible information. Information can be measured based on the quality or the quantity of the data. To be measured by quality, user interaction would be necessary to decide which region is important in which modality. This would be a good measurement but it is a time-consuming process and has to be repeated for each new data set.

A second way to measure information is based on quantity, i.e., frequency, of the data. For this measurement the methods of information theory are employed. The idea behind this measurement is that values which occur very often have less information than values which occur not so often. For medical data sets this can be interpreted that larger regions with the same value, such as the background, contain less information than smaller regions, such as border areas or small tissues. The information content can be expressed by the following equation:

$$I(f) = -\log(p(f)) \quad (3.1)$$

where $p(f)$ is the probability of occurrence for a certain value f . Through the \log function the information $I(f)$ is high for values with a low probability. The fusion should then be done in a way to weight the value with more information content higher than the value with less information

content. To formalize this weighting we introduce the following equation:

$$\gamma(\dot{f}, \ddot{f}) = \frac{I(\ddot{f})}{I(\dot{f}) + I(\ddot{f})} \quad (3.2)$$

The γ value is zero when the second modality has no information. It is one if the first modality has no information. For a value of 0.5 both modalities contain the same amount of information for a given pair of values.

With Equation 3.2 we get a quantity γ for each pair of values which can directly be used for the weighting in the fusion step. The fusion of two values, \dot{f} and \ddot{f} , is simply done by the following equation:

$$f_{fused} = (1 - \gamma) * \dot{f} + \gamma * \ddot{f} \quad (3.3)$$

The fused value f_{fused} is close to the value of one modality when this modality contains much more information than the other modality. Therefore, points with high information content in just one modality are only slightly modified in contrast to their original value. This property makes it easier to find such points in the new transfer-function space because they have almost the same value as they would have in volume visualization of this modality alone. For points with a γ around 0.5 the fused value is a mixture of both values and, therefore, is distinguishable from points with high information content in one modality.

The gradients of both modalities are fused in the same manner as the values to get an appropriate fused gradient according to the values:

$$\nabla f_{fused} = (1 - \gamma) * \nabla \dot{f} + \gamma * \nabla \ddot{f} \quad (3.4)$$

The fusion of the gradients is needed for the shading calculation as well as for classification by the transfer function based on gradient magnitude. The result of the fusion is a single value for each sample point like for the visualization of a single volume. This fused value together with the magnitude of the fused gradient can be used for the classification by a transfer function. Unfortunately some tissues are overlapping in this fused transfer-function space. Therefore an additional parameter is introduced in the following section which supports the transfer-function design for a better separation of different tissues.

3.2.3 Opposite Information Retrieval

In the previous section a quantity was calculated which indicates which of the two values has more information. In this section we will define a quantity which indicates the information contained in the joint occurrence of two values rather than the information contained in the occurrence of a single value. This new quantity will be employed as another attribute for the classification of a sample point. It allows for a better separation of different tissues.

For image and volume registration the maximization of the mutual information is a common approach to find a good registration position. In this context the best registration position is found when the mutual information is at a maximum. This means that in this position both data sets contain the lowest possible opposite information. The mutual information is a quantity for the

whole data set. In contrast the point-wise mutual information (PMI) is a quantity for the mutual information of a certain combination of values. It is defined by the following equation:

$$PMI(\dot{f}, \ddot{f}) = \log \left(\frac{p(\dot{f}, \ddot{f})}{p(\dot{f}) * p(\ddot{f})} \right) \quad (3.5)$$

The PMI is zero when a pair of values occurs exactly as frequently as one would expect by chance. This is the case when both values are statistically independent from each other and the joint probability $p(\dot{f}, \ddot{f})$ is exactly the product of both marginal probabilities $p(\dot{f})$ and $p(\ddot{f})$. If the two values \dot{f} and \ddot{f} occur together more frequently as one would expect by chance then the result of the calculation is greater than zero. Conversely, the value is lower than zero if a pair of values occurs less frequently as one would expect by chance. By the definition of Shannon this case contains more information than a result value greater than zero because the occurrence is less frequent. For a joint probability $p(\dot{f}, \ddot{f})$ of zero the PMI is by definition zero. For all other probabilities the PMI can be normalized to a value between zero and one by the lower bound ($p(\dot{f}) = 1$ and $p(\ddot{f}) = 1$) and upper bound ($p(\dot{f}) = p(\dot{f}, \ddot{f})$ and $p(\ddot{f}) = p(\dot{f}, \ddot{f})$) of the PMI:

$$PMI_{norm}(\dot{f}, \ddot{f}) = \frac{PMI(\dot{f}, \ddot{f}) - \log(p(\dot{f}, \ddot{f}))}{\log(\frac{1}{p(\dot{f}, \ddot{f})}) - \log(p(\dot{f}, \ddot{f}))} \quad (3.6)$$

This can be further reduced to:

$$PMI_{norm}(\dot{f}, \ddot{f}) = \frac{1}{2} - \frac{PMI(\dot{f}, \ddot{f})}{2 \log(p(\dot{f}, \ddot{f}))} \quad (3.7)$$

The value of PMI_{norm} approaches zero if the information carried by the pair of values is high. Values close to one represent low information content. To get a high value for pairs of values with high information content we define a new quantity δ as an inversion of PMI_{norm} :

$$\delta(\dot{f}, \ddot{f}) = 1 - PMI_{norm}(\dot{f}, \ddot{f}) \quad (3.8)$$

Figure 3.2 illustrates the behavior of δ . The different regions, labeled with capital letters, have different colors to symbolize regions of different values in both modalities. The red crosses are sample points for which the δ value should be calculated. For the sample point S_1 the involved marginal probabilities ($p(\dot{f})$ and $p(\ddot{f})$) are rather low because only a small area (C_1 and C_2) has the same value in both modalities. For the sample point S_2 the marginal probability in the second modality is higher because the sample point lies in a larger area B_2 . The joint probability $p(\dot{f}, \ddot{f})$ is the same for both sample points because the combination of C_1 and C_2 occurs exactly as often as the combination of D_1 and B_2 . By calculating the δ values with these probabilities we, however, get a smaller value for the sample point S_1 than for the sample point S_2 .

This example can be interpreted in a way that for sample point S_1 both modalities contain correlated information whereas for S_2 modality 1 complements the information of modality 2 because the region D_1 is only represented in modality 1. This means that the δ value responds with a high value for regions with high opposite information content. So this value can be used to separate tissues which only show up in one modality from tissues which are present in both

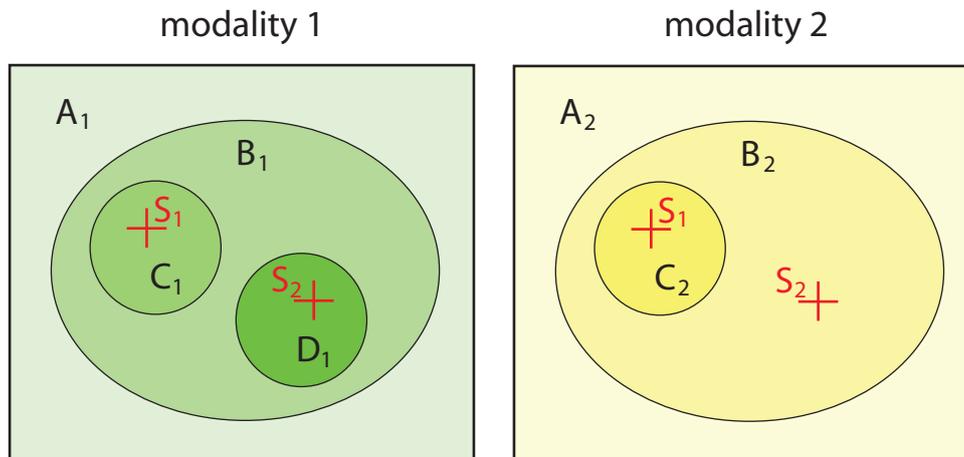


Figure 3.2: Example of slices of two different modalities to explain how the δ value is affected by the value distribution. S_1 and S_2 are sample points for which the δ value is calculated.

modalities. It can be seen as a quantity which indicates the difference of information content in both modalities at each sample point. Noise in the data sets does not influence the δ value. It flattens the probability distribution function of a certain material but the relation between the probabilities does not change and, therefore, the δ value is not affected. The following section describes how this property can be integrated into the classification process.

3.2.4 Information-based Transfer-Function Classification

In the previous two sections we described how methods from information theory can be taken to generate a fused value and a fused gradient as well as an additional property δ which indicates the opposite information. These values together will be used now for the assignment of optical properties.

Due to the existence of three values (f_{fused} , $|\nabla f_{fused}|$, δ) for each sample point the classification could be done in a 3D space. For every triple of values optical properties would be assigned. This approach is shown in Figure 3.3(a). The problem with this approach is the complexity of the transfer-function design and, therefore, it is hard to find a good transfer function. To avoid this we reduce the degree of freedom by defining a region only in the 2D transfer-function space (f_{fused} , $|\nabla f_{fused}|$). The design task in this space is easier because the 2D space is already well-known from volume visualization of only one volume. Additionally, for each region a simple windowing function is defined for the δ value. The selection of a windowing function for the extraction of δ results from the fact that the δ values for points of one tissue in anatomical modalities or a level of activity in functional modalities are in a certain value range. To extract such parts only points with a δ value in this range should be selected. A windowing function is easy to adjust to a certain value range and, therefore, is well suited for this purpose. The windowing function can

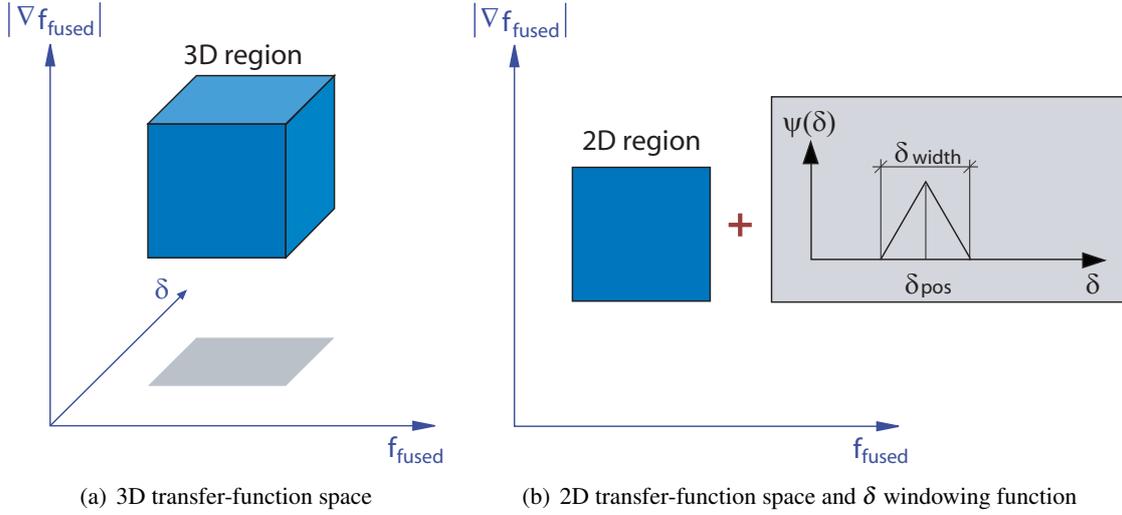


Figure 3.3: The transfer-function space is converted from 3D (a) to 2D (b). Additionally, a simple windowing function for the δ value is employed to modify the optical properties of each 2D region.

be expressed by the following equation:

$$\psi(\delta) = \max\left(\left|1 - \frac{\delta - \delta_{pos}}{0.5 * \delta_{width}}\right|, 0\right) \quad (3.9)$$

The parameters δ_{pos} and δ_{width} define the position and shape of the windowing function $\psi(\delta) \in [0, 1]$. The original opacity α , assigned according to a 2D region in the transfer-function space, is multiplied with this value to fade out points with a low value of this windowing function. In Figure 3.3(b) the separation into a 2D region and a corresponding windowing function is shown.

3.3 Implementation

For a fast and efficient volume rendering it is necessary to do as many calculations as possible in a pre-processing step. The most time-consuming part of the whole process is the generation of the dual histogram and the two individual histograms of both modalities for the estimation of the probabilities. This can be done before the rendering because the histograms are static for two given volume data sets and do not change during the rendering process. The histograms are used to calculate the γ and δ values as described in the previous section. Each of these values can be stored in a 2D lookup table. They also do not change for two given volume data sets.

Figure 3.4 shows the processing steps for each sample point during the rendering process. The processing steps with sharp corners are lookups and the processing steps with rounded corners are calculations. As first step lookups in the a priori generated γ and δ lookup tables are done. The γ value is used to fuse the two input values as described in Section 3.2.2. With the fused value and the magnitude of the fused gradient a lookup in the tables of the transfer function is done. One

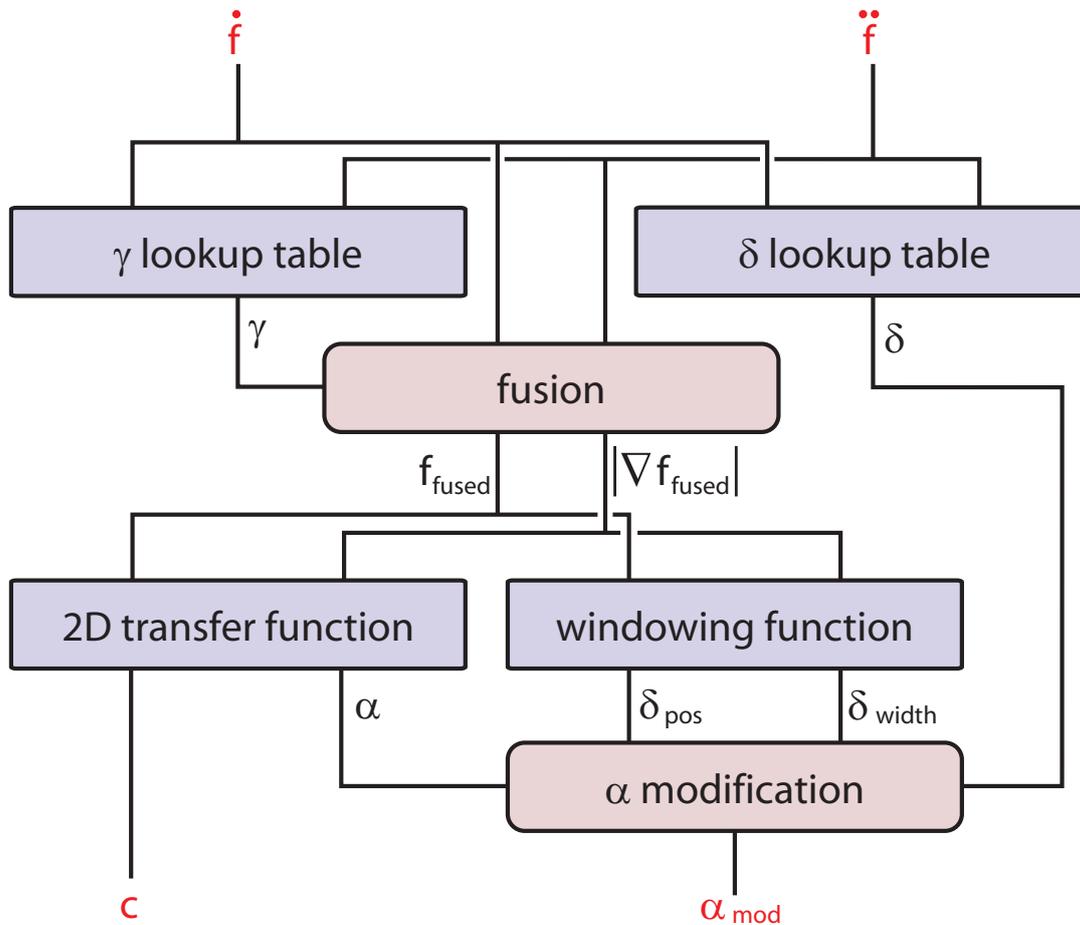


Figure 3.4: Overview over the processing steps for each sample point during the rendering process. Nodes with rounded corners are calculation steps and nodes with sharp corners are lookups.

lookup table stores the color c and opacity α for each point in the transfer-function space. The second lookup table stores the parameters δ_{pos} and δ_{width} of the windowing function. The color c of the 2D transfer function is directly used for further processing steps, such as shading. The opacity α is modified by the windowing function according to the parameters δ_{pos} and δ_{width} as well as the δ value. As output of this calculation step we get a modified opacity α_{mod} which is further used in the rendering process. The speed of the implementation is not a big issue because all processing steps can be executed quite fast on the graphics hardware to achieve real-time frame rates.

3.4 Results

Modalities can be generally classified into two groups: functional and anatomical modalities. The most common anatomical modalities are CT and MRI. CT is typically used to show bone

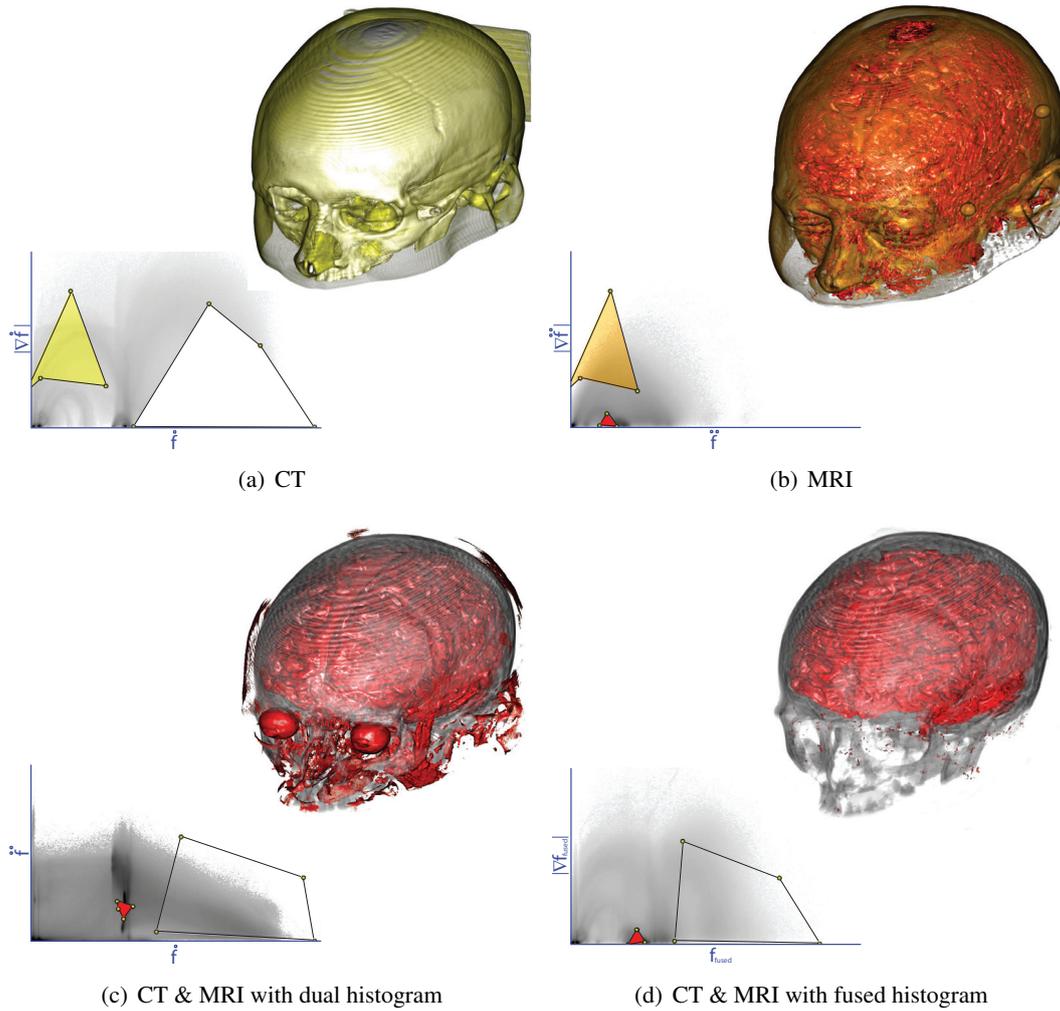


Figure 3.5: The images show single volume visualizations of CT data (a) and MRI data (b) in contrast to multimodal visualizations by using a dual transfer-function space (c) (as described in [42]) and the fused transfer-function space (d). The corresponding histograms with colored regions for the assignment of optical properties are shown below the results.

structures. Soft tissues have a higher contrast in MRI. In Figure 3.5(a) the visualization of a CT scan is shown and in Figure 3.5(b) the visualization of an MRI scan. Both visualizations can be useful for special examinations but it can also be seen that both data sets contain some joint information. Furthermore some regions with less information, such as the tissue around the brain in the MRI scan, are hiding regions with more information, such as the brain itself.

The goal of a multimodal visualization is to combine relevant tissues from both modalities and show them together to provide additional context. The relevance of a tissue is dependent on the kind of examination. In a combination of CT and MRI of a head the brain could be the relevant

part of the MRI scan and the bones could be the relevant parts of the CT scan. Figure 3.5(c) shows the rendering result of a multimodal visualization based on the dual histogram. Both relevant tissues, the brain and the bones, are visible but also a lot of artifacts show up in the result. This follows from the fact that the brain cannot be better separated in the transfer-function space based on the dual histogram. Figure 3.5(d) shows the result generated by the new method. In comparison to the result generated with the traditional multimodal visualization technique the brain is clearly separated from other tissues and only a few artifacts are visible.

Below the results in Figures 3.5(a) to (d) the corresponding histograms are shown. The regions which were used to classify sample points with optical properties, such as color and opacity, are also shown on top of these histograms. It can be seen that the regions for classifying the brain tissue and the bones in the new fused transfer-function space, as shown in Figure 3.5(d), are highly related to the individual regions in the single modality visualizations, as shown in Figure 3.5(a) and Figure 3.5(b). The regions for the multimodal visualization, based on the dual histogram, are shown in Figure 3.5(c). The position and shape of the regions in this transfer-function space are completely different in comparison to the regions for the single modality visualizations. This makes it much harder for the user to define regions for the transfer function because the knowledge from the single modality visualization cannot be used.

As described in Section 3.2.4 the definition of a transfer function is done in two steps. In Figure 3.5(d) only the regions are shown which assign a color and non-zero opacity to sample points. Furthermore for each of these regions a windowing function for the δ value is defined. This function is used to refine the separation by the transfer function. In Figure 3.6(a) the rendering result is shown which is generated without the usage of a windowing function for δ . The region which is used to assign optical properties to the brain is the same as for Figure 3.5(d). It can be seen that the result contains a lot of artifacts. In comparison to that, Figure 3.6(b) shows a result which is generated by the additional usage of a windowing function for δ to modify the opacity. Through the refinement of the classification with the windowing function most of the artifacts are gone and the brain is clearly separated.

Besides the reduction of artifacts the strength of the additional δ value is the ability to find regions with high differences in both data sets. This can be very helpful for several applications, such as the finding of a tissue which only shows up in one modality. Due to the properties of δ as described in Section 3.2.3 regions with opposite information in both data sets have a high δ value. Figure 3.7 shows the response of the δ value for the combination of two example data sets. In Figure 3.7(a) and Figure 3.7(b) two data sets are shown which only differ at one region where in modality 1 a sphere exists and in modality 2 not. Figure 3.7(c) shows the corresponding distribution of δ values for the two modalities. In the region where the sphere is represented in only one modality the δ value is the highest due to complementary information.

Figure 3.8 shows the result of a multimodal visualization for the combination of a CT scan and a PET scan generated by the new approach. The regions of high activity inside the brain and in the tumor on the neck are shown more opaque. This example shows that the method also works with the combination of anatomical and functional modalities and, furthermore, with different spatial resolutions.

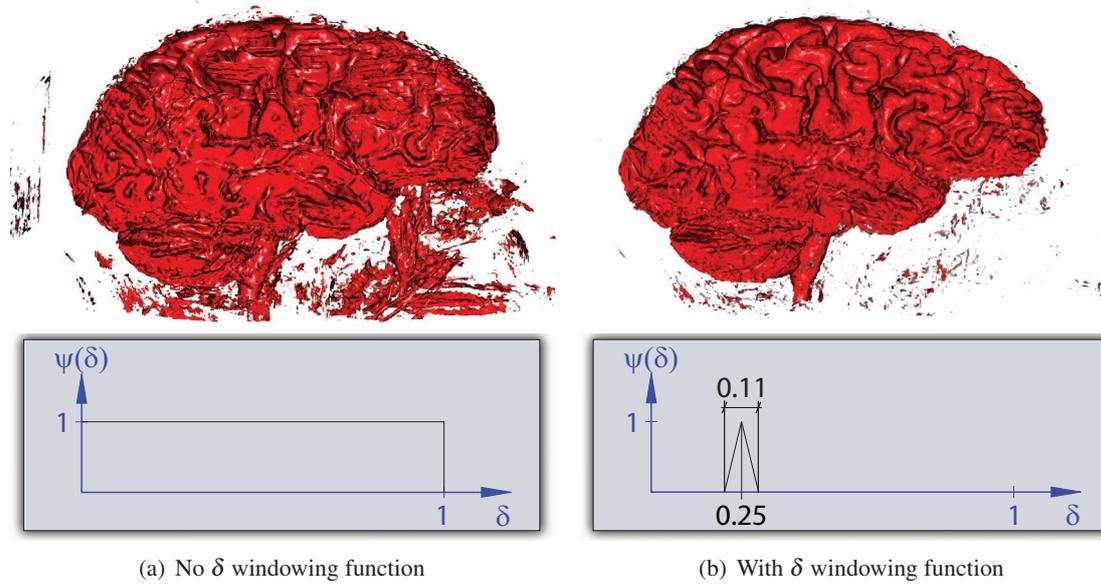


Figure 3.6: The two results show the effect of the usage of δ to modify the optical properties of a 2D region in the transfer-function space.

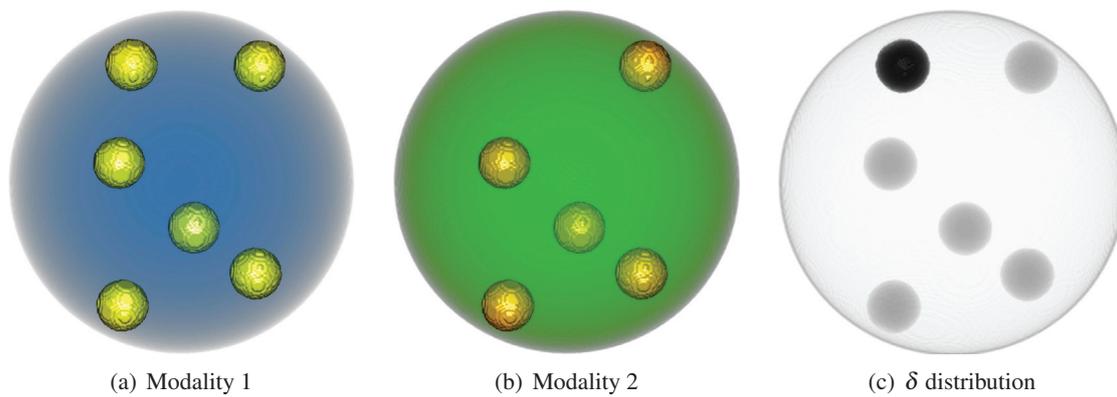


Figure 3.7: The image in (c) shows the distribution of δ in volume space. It is highest in regions with the largest difference. In this case the largest difference occurs where in modality 1 (a) a sphere exists and in modality 2 (b) not.

3.5 Conclusion and Discussion

In this chapter a novel approach for the definition of transfer functions for multimodal visualization was introduced. The initial idea was to define a user-friendly transfer-function space, which makes it easy to find an expressive transfer function in order to visualize certain tissues of both modalities. Through the fusion of the data values, based on the information content, a

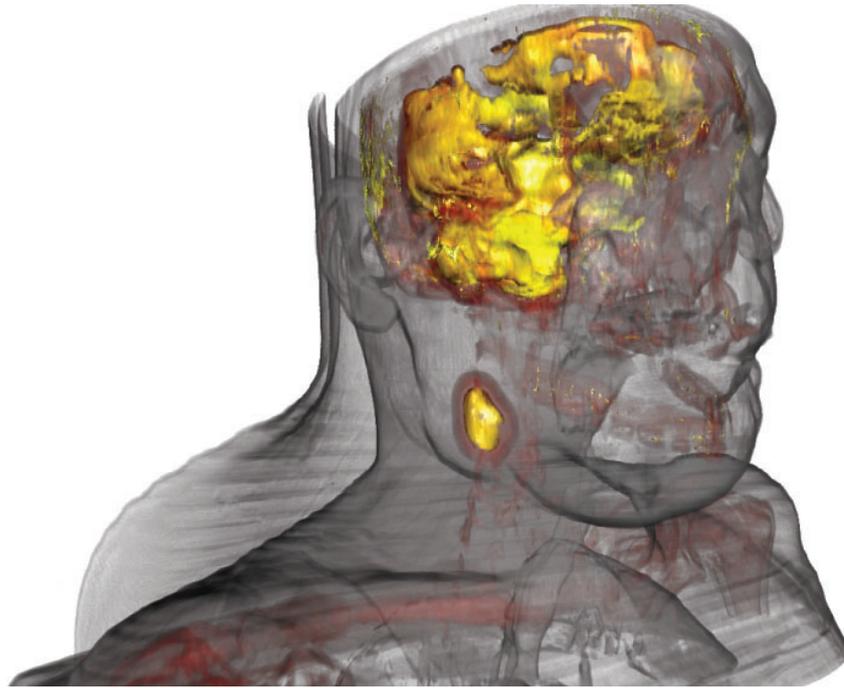


Figure 3.8: Multimodal visualization of a CT and PET scan. The more opaque regions indicate regions of high activity such as in the brain and in the tumor on the neck.

2D transfer-function space is defined which is similar to the well-known 2D transfer-function space of single volume visualization with value and gradient magnitude as the two dimensions. Therefore, the distribution of points in this transfer-function space is easier to understand by the user in comparison to other multimodal transfer-function spaces [42]. An additional δ value, which describes the complementary information contained in a pair of values, is used for a better separation of different materials. In the result section we have shown how the new transfer-function space can be used to classify the most relevant parts of both modalities in a single visualization.

In comparison to other approaches, which are used for multimodal visualization, the benefit of the new approach is the conversion of the classification problem to a problem which is already known from classification in single volume rendering. A penalty of the new method is that more information does not always mean more importance. So it can happen that, e.g., artifacts can have high information content while other, more important parts have lower information content. The user can control this by defining a transfer function which has low opacity for such less important parts.

In the method of this chapter the global distribution of data values was used as feature to extract the information content. This is a good feature if small parts of the data sets or differences in both modalities are of interest. In the next chapter a method is introduced which uses the structure of objects in multimodal data to extract information-theoretic features. Through this it is possible to enhance parts of the data based on their structure and not data values.

Though no one can go back and make a brand new start, anyone can start from now and make a brand new ending.

— Ivern Ball

CHAPTER

4

Volume Classification Using Multimodal Surface Similarity

Martin Haidacher, Stefan Bruckner, and M. Eduard Gröller. *Volume Classification Using Multimodal Surface Similarity*, Technical Report, Institute of Computer Graphics and Algorithms, Vienna University of Technology, April, 2011

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. While the data generated by each modality may be visualized separately, it is difficult to mentally integrate multiple sources, particularly if spatial relationships are important. Thus, the effective combination, or fusion, of multiple data sets in order to improve visualization, and, eventually provide insight into the phenomenon under investigation has long been an active area of research.

The method presented in the previous chapter is able to enhance parts in multimodal data sets which contain more information according to their data values. The approach is sensitive to small variations in both modalities. Such variations can be caused by artifacts in one data set. In this chapter we introduce a different approach which extracts features from the structure of the objects. In this case we investigate the surface of the objects represented in the multimodal data. Small variations, such as artifacts, have only a low impact in this classification since they are only a small part of the whole object surface.

In the previous chapter the features which were extracted for the classification were used to define a new transfer-function space. For the method described in this chapter, we chose a different approach. Instead of generating a new transfer-function space which may be non-intuitive and difficult to understand, we start by analyzing the joint information provided by both modalities. We introduce the notion of multimodal surface similarity as a novel concept to understand the similarities and dissimilarities between two scalar fields. By generating an isosurface similarity matrix for each combination of data values from the two modalities, we obtain a concise structural overview of the joint data space. We show that the information obtained in this way can be used to assist multi-volume fusion in a non-invasive manner – the similarity is used as a visualization and interaction tool for guiding the exploration process. Furthermore, we introduce a classification approach which uses multimodal surface similarity to assign optical properties to regions in space based on a distance measure in similarity space.

The remainder of the chapter is structured as follows. In Section 4.1, we review related work on the visualization of multimodal volume data and other approaches connected to our

method. Section 4.2 introduces synthetic data sets for the further explanation of our multimodal classification approach. In Section 4.3, the general concept of multimodal surface similarity is introduced. In Section 4.4, we show how multimodal surface similarity can be used in the visualization process. Implementation details are presented in Section 4.5. The implications of our approach as well as its limitations are discussed in Section 4.6. Finally, the chapter is concluded in Section 4.7.

4.1 Related Work

As mentioned in the related work of Chapter 3, two volumes can be fused in different stages of the rendering pipeline [22].

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. [48] and Evans et al. [18] introduced methods which extract surfaces from MRI and visualize the activity level of a PET scan on these surfaces. Heinzl et al. [28] introduced a fusion workflow specialized for dual-energy CT. With this method it is possible to fuse the most prominent features from both data sets.

The mentioned methods first extract surfaces in the different modalities before the fusion. A more sophisticated but more complex approach is the fusion without the intermediate step of a surface extraction. A straight-forward method is the fusion by linear intermixing of the data values. Eusemann et al. [17] has shown that this intermixing can be improved for dual-energy CT by adapting the intermixing ratio to different tissues. Kniss et al. [39, 42] and Akiba and Ma [1] introduced multi-dimensional transfer-function spaces based on features in a local neighborhood around each sample point to classify multimodal data. Kim et al. [36] presented a technique which simplifies the transfer-function design by letting the user define a separate transfer function for each modality. The combination of them defines a two-dimensional transfer function. The fusion of different modalities involves loss of information. In Chapter 3 a data fusion and transfer-function space for multimodal visualization based on the information content was defined. With this approach the loss of information defined by information theory [65] is minimized. In contrast to the method in this chapter the information is just retrieved by the global frequency distribution of values and not by structural similarities between the different modalities.

Apart from multimodal visualization, information theory is used for other visualization tasks [74]. In flow visualization, e.g., Xu et al. [80] uses information theory to select meaningful streamlines. Feixas et al. [20] compute the best viewpoint based on information theory. The viewpoint should provide the best view on the data with respect to information loss. Chen and Jänicke [10] introduce an information-theoretic framework for scientific visualization. For our approach we are using the information theory to measure similarities between the different modalities.

For many applications, such as industrial CT [28], 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 Wender [35] 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 [9] is used to topologically analyze the volume

data. It is able to encode the nesting relationships of isosurfaces. Takahashi et al. [67] employed a volume skeleton tree to identify isosurface embeddings in order to provide additional structural information. Stable surfaces can also be classified by additional features in the transfer-function space. Kindlmann and Durkin [37] introduced a transfer-function space in which the gradient magnitude is used as additional classification feature. Interfaces between materials show up as arches in this transfer-function space. In the LH histogram of Šereda et al. [72], the highest and lowest value along a 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 [6] 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. These clusters are indicators of stable isosurfaces. 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.

4.2 Synthetic Multimodal Data Sets

The usage of multiple data sets of the same object may offer better insight and understanding of the function or structure of the object. There are different reasons why multimodal data sets are used. In medicine, e.g., functional data is acquired together with anatomical data. The functional data contains information about the function of inner organs in low resolution. The anatomical data augments this information with the anatomy of the body in high resolution. In other fields - such as industrial CT - the various data sets differ by the parameters which were used for the scan.

For the further description of our algorithm 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. In the subsequent sections these data sets are used to explain multimodal similarity maps and highlight the usefulness of our method.

Supplementary Data

Multimodal data is often used to supplement one data set with additional features. This is necessary when a data set contains noise or other artifacts in some regions. In this case a second data set is used to compensate for these artifacts but it introduces other disadvantages. In this chapter we will refer to these data sets as supplementary data types. Basically both data sets contain the same information. The disadvantage of one data set is compensated by the other one and vice versa.

An example for supplementary data is dual energy CT. It is used in medicine and industrial CT. The most common artifacts in CT scans in general are noise-induced streaks, beam hardening, partial volume effects, aliasing, and scattered radiation [2, 31]. Due to the fact that different energy levels have different attenuation characteristics, some of these artifacts appear prominently

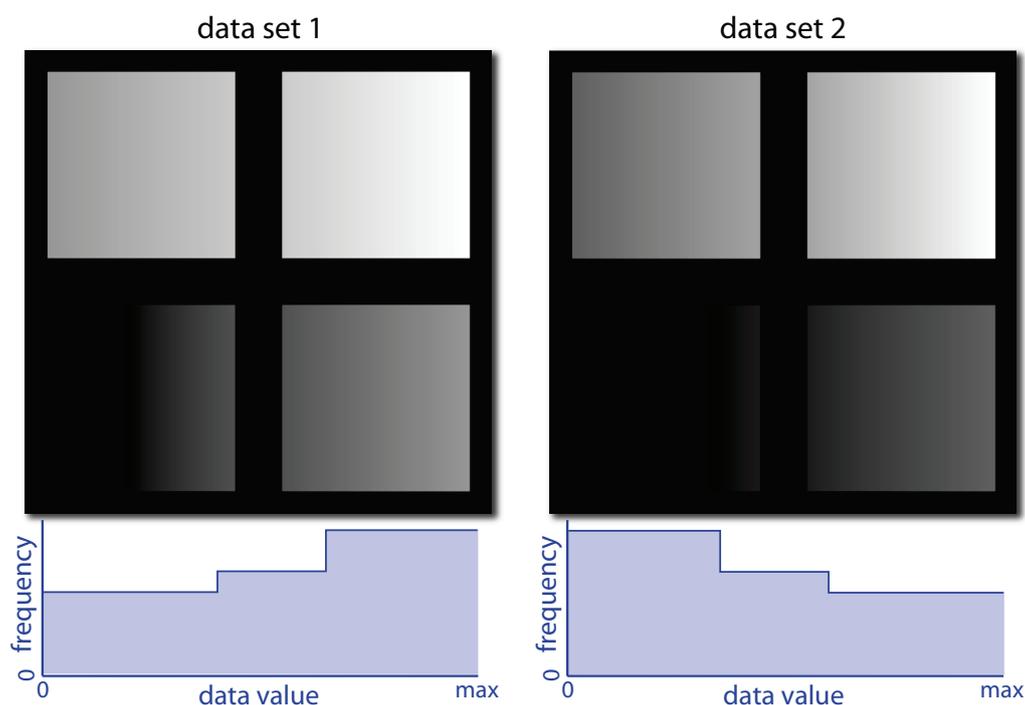


Figure 4.1: Data sets containing the same information but with different value ranges. The histograms show the distributions of the data values.

only in one energy level. Hence it is possible to reduce artifacts by the fusion of CT data sets of different energy levels.

For the description of our algorithms we generated two synthetic data sets which simulate the behavior of supplementary data. In Figure 4.1 these two data sets and their frequency distributions are illustrated. Both modalities contain the same information, i.e., four squares with gradually changing data values. The value ranges for the squares are in both data sets slightly different. This addresses the different attenuation characteristics in dual energy CT.

Complementary Data

Multimodal data sets are also used to combine information from different modalities. In this case, a significant amount of information differs, or is not represented in one of the data sets. Fusion aims to combine this information. We will refer to these multimodal data sets as complementary data.

A typical example for complementary data is medicine. Modalities such as CT and MRI measure different physical characteristics of the human body, and thus there are significant 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

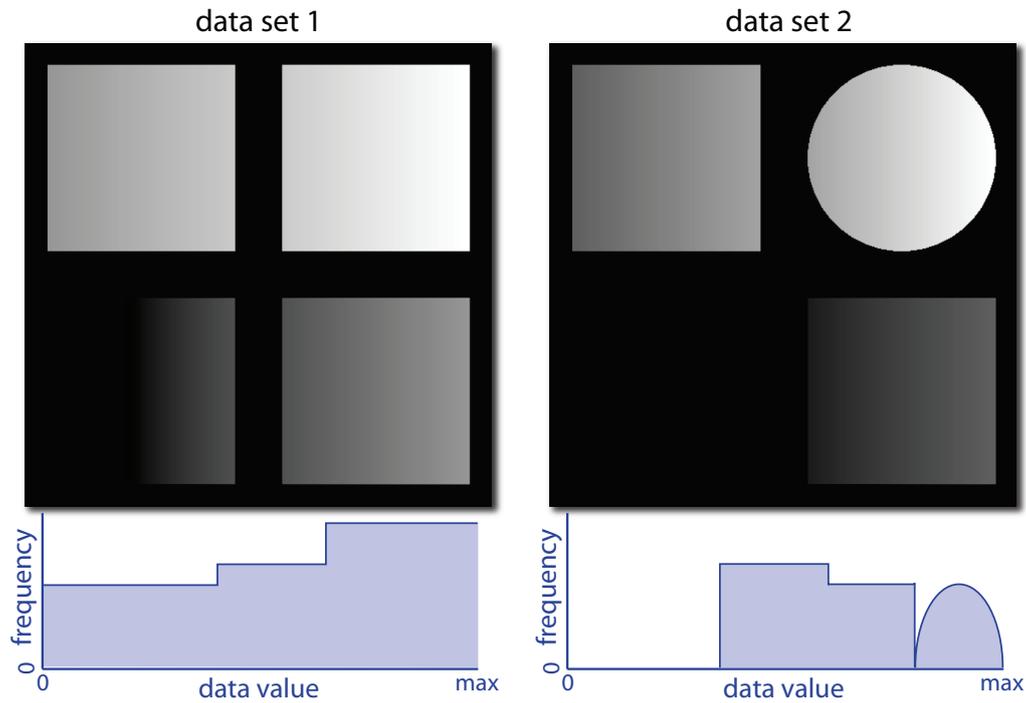


Figure 4.2: Two synthetic data sets which represent complementary data. Data set 2 contains information which is different from data set 1.

correspondence between the two modalities, as CT images contain no functional information at all.

We will use the synthetic data sets illustrated in Figure 4.2 to demonstrate our algorithms with 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 different information given in complementary data. In the next section these synthetic data sets are used to explain the multimodal surface similarity measurement.

4.3 Multimodal Surface Similarity

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

$$L_k = \{x \in \mathbb{R}^3 : f(x) = k\} \quad (4.1)$$

The measure of isosurface similarity was introduced by Bruckner and Möller [6] as a means of quantifying 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 two isosurfaces of a single data set. In this section, we first briefly review the original approach and then introduce the concept of *multimodal similarity maps* to represent the similarity between isosurfaces of different modalities.

4.3.1 Self Similarity Maps

A robust method to measure the similarity between isosurfaces is *mutual information*. It has been applied in many areas including shape registration [32] and viewpoint selection [71]. The mutual information of two discrete random variables X and Y can be defined as [81]:

$$I(X, Y) = H(X) + H(Y) - H(X, Y) \quad (4.2)$$

where $H(X, Y)$ is the joint entropy and $H(X)$ and $H(Y)$ are the marginal entropies of random variables X and Y (see also Section 1.4). Since the mutual information is limited by the average marginal entropies, it can be normalized to a value range of $[0, 1]$ by [44]:

$$\hat{I}(X, Y) = \frac{2I(X, Y)}{H(X) + H(Y)} \quad (4.3)$$

As a measure of isosurface similarity, Bruckner and Möller [6] proposed to compute 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 [33]:

$$D_k(x) = \min_{\forall y \in L_k} d(x, y) \quad (4.4)$$

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

4.3.2 Multimodal Similarity Maps

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

$$\dot{L}_k = \{x \in \mathbb{R}^3 : \dot{f}(x) = k\} \quad (4.5)$$

$$\ddot{L}_l = \{x \in \mathbb{R}^3 : \ddot{f}(x) = l\} \quad (4.6)$$

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

$$\dot{D}_k(x) = \min_{\forall y \in \dot{L}_k} d(x, y) \quad (4.7)$$

$$\ddot{D}_l(x) = \min_{\forall y \in \ddot{L}_l} d(x, y) \quad (4.8)$$

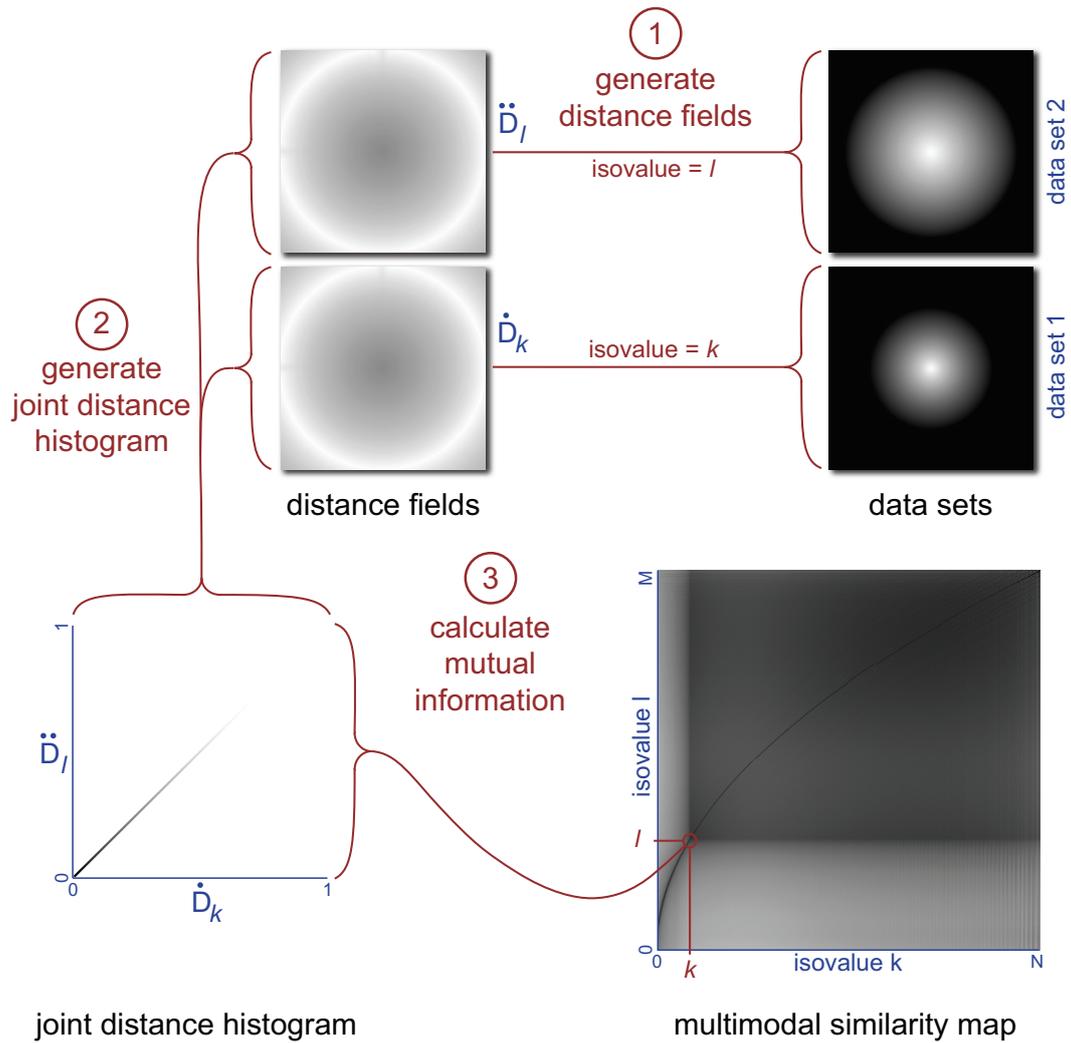


Figure 4.3: Pipeline for the generation of a multimodal similarity map. The illustration shows which steps are necessary to calculate the similarity of isosurfaces for the isovalues k and l .

Figure 4.3 illustrates how the mutual information for a combination of isovalues l and k is calculated. The first step is the generation of distance fields \dot{D}_k and \dot{D}_l for the isosurfaces \dot{L}_k and \dot{L}_l . In the next step the distances \dot{D}_k and \dot{D}_l 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 to isosurface \dot{L}_k and \dot{D}_l to isosurface \dot{L}_l . In Figure 4.3 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 \dot{L}_l .

Finally, the mutual information is calculated based on Equation 4.2. The joint and marginal probabilities for the calculation of the joint and marginal entropies can be directly retrieved from the joint distance histogram.

If we assume that modality 1 has N different isovalues $\dot{V} = \{k_1, \dots, k_N\}$ and modality 2 has M different isovalues $\dot{V} = \{l_1, \dots, l_M\}$ 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 isosurface \dot{L}_j of modality 2 with the corresponding isovalue j .

Figures 4.4 and 4.5 show the multimodal similarity maps for the synthetic data sets introduced in Section 4.2. Dark regions denote a high similarity in these figures. For the supplementary data types in Figure 4.4 both data sets contain four squares at the same location. In the MSM each of the squares is represented by a rectangular area of higher similarity. In Figure 4.4 corresponding squares and rectangular regions are emphasized by colored frames. The band with the maximum similarity represents the combination of isovalues k and l at which both data sets represent exactly the same isosurface. 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.

In Figure 4.5 the multimodal similarity map for our complementary test data set is shown. The regions in which both data sets contain contradictive information is clearly visible in the similarity map. In contrast to Figure 4.4, the lower left rectangular area (red frame) in Figure 4.5 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.

Another interesting area in the multimodal similarity map of Figure 4.5 is the rectangle in the upper right corner (cyan frame). This rectangular area represents the similarity between the square in one 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.

4.4 Similarity-Based Volume Fusion

The conventional approach for visualizing volume data in 3D is to define a transfer function which specifies the color and opacity of each sample based on one or several data attributes. Many different transfer-function spaces have been proposed, each with its own set of advantages and disadvantages. One dimensional transfer functions based on the original data values, as proposed by Levoy [49], however, still enjoy considerable popularity due to their relative simplicity. In particular in the context of multimodal data, additional transfer-function dimensions can result in complex user interfaces. In this section, we describe how multimodal similarity maps can be used as a tool to explore multimodal data. In this section, we discuss how the additional information provided by the multimodal similarity map can guide and assist the process of exploring multimodal data. We specifically do not want to introduce a new transfer-function space. Instead, we aim to explore how multimodal similarity maps can be used as a visualization and interaction tool for conventional volume visualization techniques.

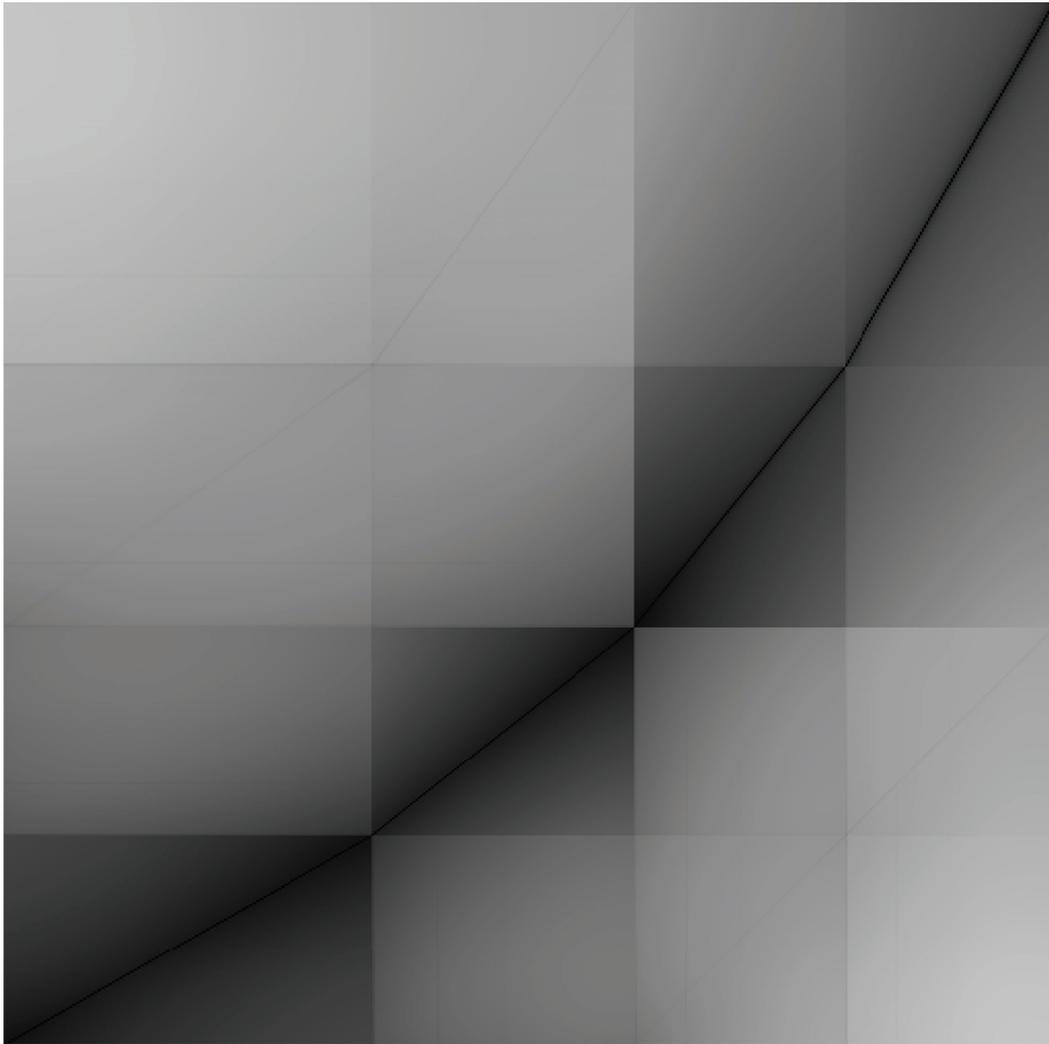


Figure 4.4: Multimodal similarity map for supplementary data types.

4.4.1 Similarity Operations

With the multimodal similarity map we gain information about the similarity of certain combinations of isovalues. Based on this information we can think of different operations which can be used to improve the multimodal visualization.

One of the simplest ways to visualize multimodal data is to perform no fusion at all, but rather view both data sets side-by-side or blended on top of each other. This approach can be quite effective in slice views – the user can simply synchronously browse the two volumes and study the data. In a 3D visualization, however, this approach is typically less useful, as occlusion severely limits the number of isosurfaces which can be depicted simultaneously. An operation which is instantly feasible with the multimodal similarity maps is the detection of maximally

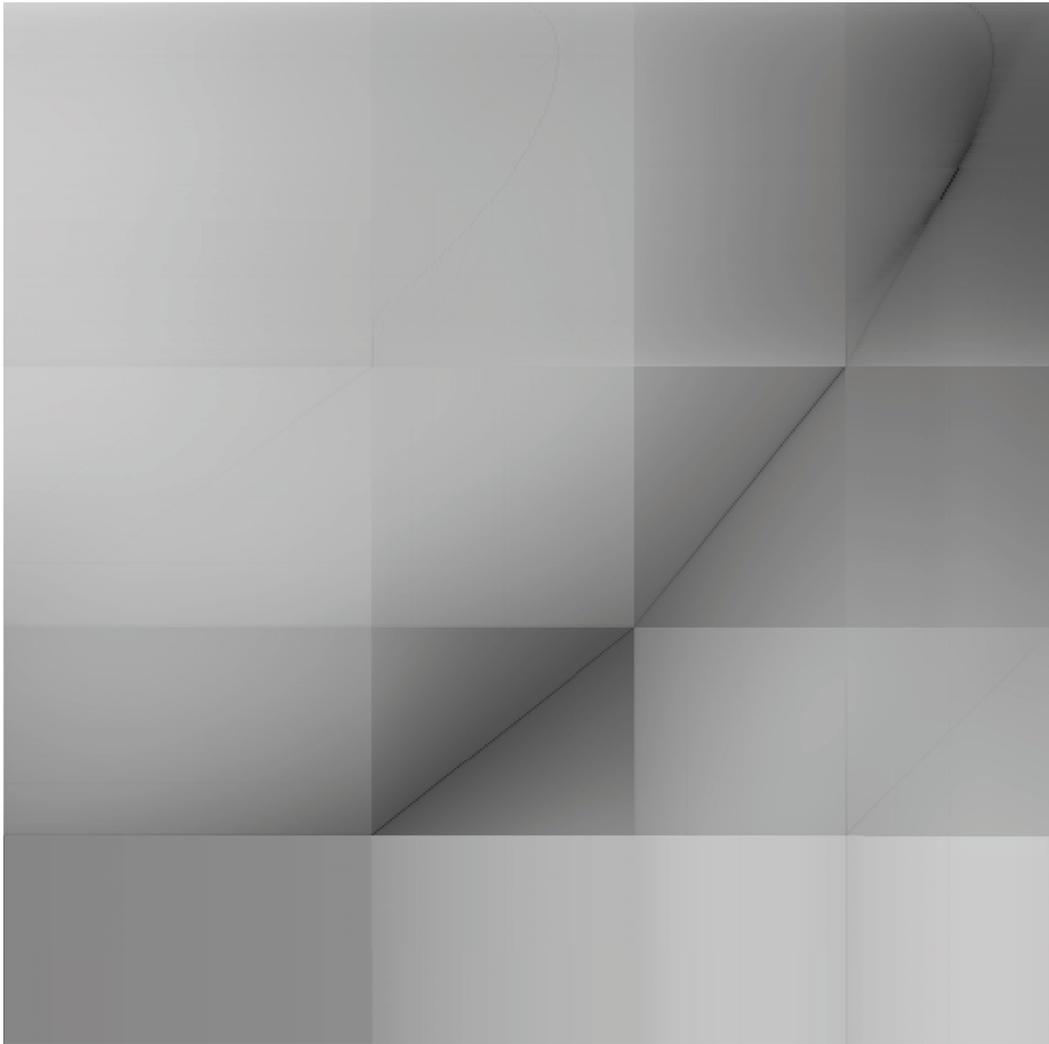


Figure 4.5: Multimodal similarity map for contradictory complementary data.

similar isosurfaces. If we assume that a user is choosing 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:

$$\hat{k} = \arg \max_j \text{MSM}(k, j) \quad (4.9)$$

The result in Figure 4.6 shows the effect of the maximum similarity selection. The data sets for this result originate from a dual-energy CT. 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 4.6. The results 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. In contrast to that, the middle row shows the isosurfaces in modality 2

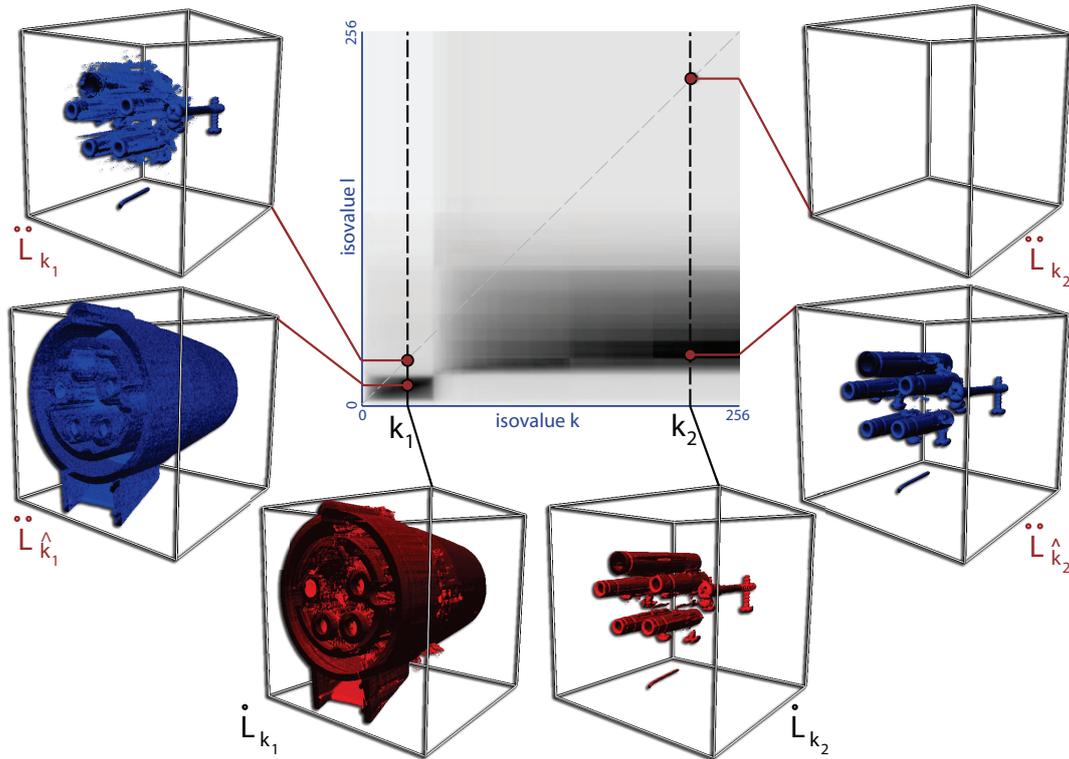


Figure 4.6: Robust isosurface selection for two different isovalues k_1 and k_2 . The results in the middle row show the most similar isosurfaces (\hat{L}_{k_1} , \hat{L}_{k_2}). The results in the top row show the isovalues for a naive selection of the isovalues, i.e., in both data sets the same isovalue is chosen.

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.

A second common approach for fusing two volumes is to blend the colors and opacities, defined by two separate transfer functions, at each sample point during direct volume rendering. The drawback of this method is that it requires the user to ensure, by tuning both transfer functions, if the desired parts of each modality are visible. As the multimodal similarity map provides information about the global similarity between the isosurfaces passing through a sample point, it can be used to guide visibility. The multimodal similarity map is used in this fusion to weight combinations of values based on their isosurface similarity.

Figure 4.7 shows the effect of the similarity weighting on the combination of CT and PET. PET contains functional information in a low resolution. The CT is used as complementary anatomical information. The combination of PET and CT is used to detect and localize cancer. A common fusion technique for CT and PET simply weights the CT values with the activity value in the PET [46]. With this method, regions of high activity are emphasized without considering the anatomical structure. Thus the anatomical structure, such as tissue borders, is faded away.

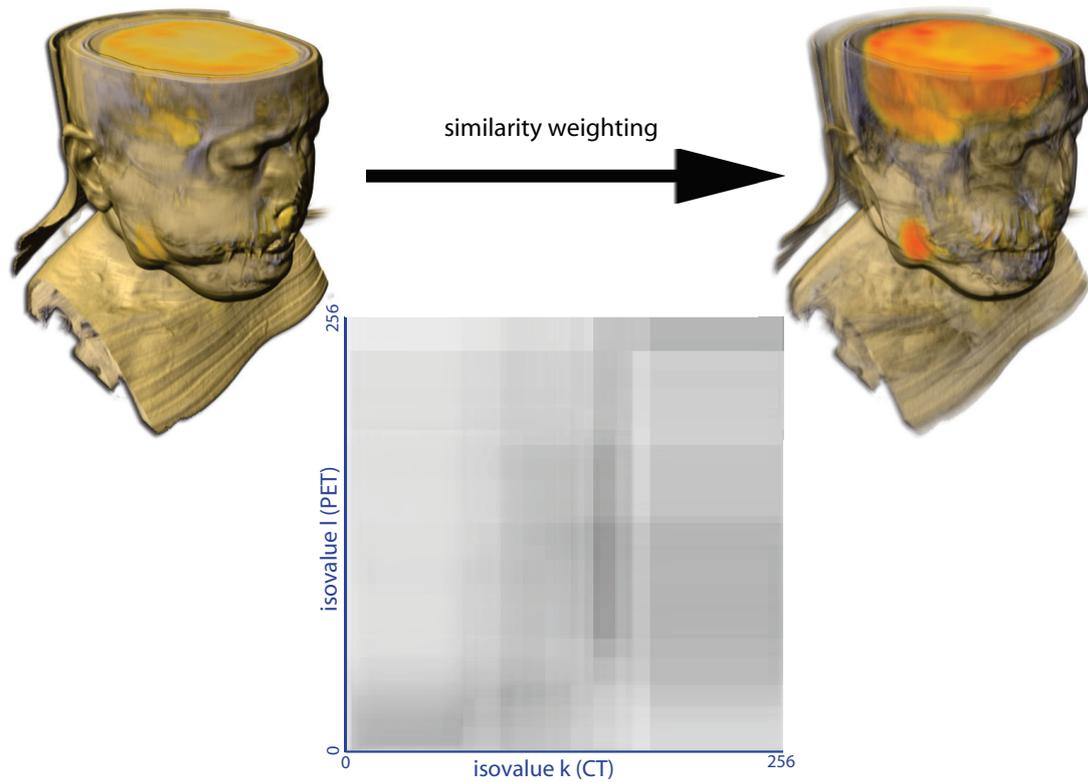


Figure 4.7: Effect of similarity weighting shown on a combination of a PET and a CT data set.

The similarity weighting preserves areas which represent similar structures in both modalities. In the case of CT and PET these are tissue borders surrounding an area of high activity in the PET scan. Therefore anatomical structures for regions of high activity can be preserved, like the shape of the brain in Figure 4.7. In this example the similarity weighting is used to modify the opacity α of the fused optical properties.

The similarity weighting is applied to all isovalues of both data sets. It is also possible to narrow down the selection of isovalues which are used for the fusion. For this operation the multimodal similarity map is not only used for the weighting but also as a guidance map to set the selection boundaries. The ranges $[k_{low}, k_{high}]$ and $[l_{low}, l_{high}]$ serve as selection ranges in the modalities. Each isovalue k and l is classified to be inside or outside this range:

$$\begin{aligned} \dot{a} &= \begin{cases} true & k_{low} < k < k_{high} \\ false & else \end{cases} \\ \ddot{a} &= \begin{cases} true & l_{low} < l < l_{high} \\ false & else \end{cases} \end{aligned} \quad (4.10)$$

This results in two Boolean variables \dot{a} and \ddot{a} . These variables can be combined by logical operations. Figure 4.8 shows the result of the application of AND, OR, and XOR on a combination of CT and MRI data sets. In the multimodal similarity map in Figure 4.8 the ranges for the

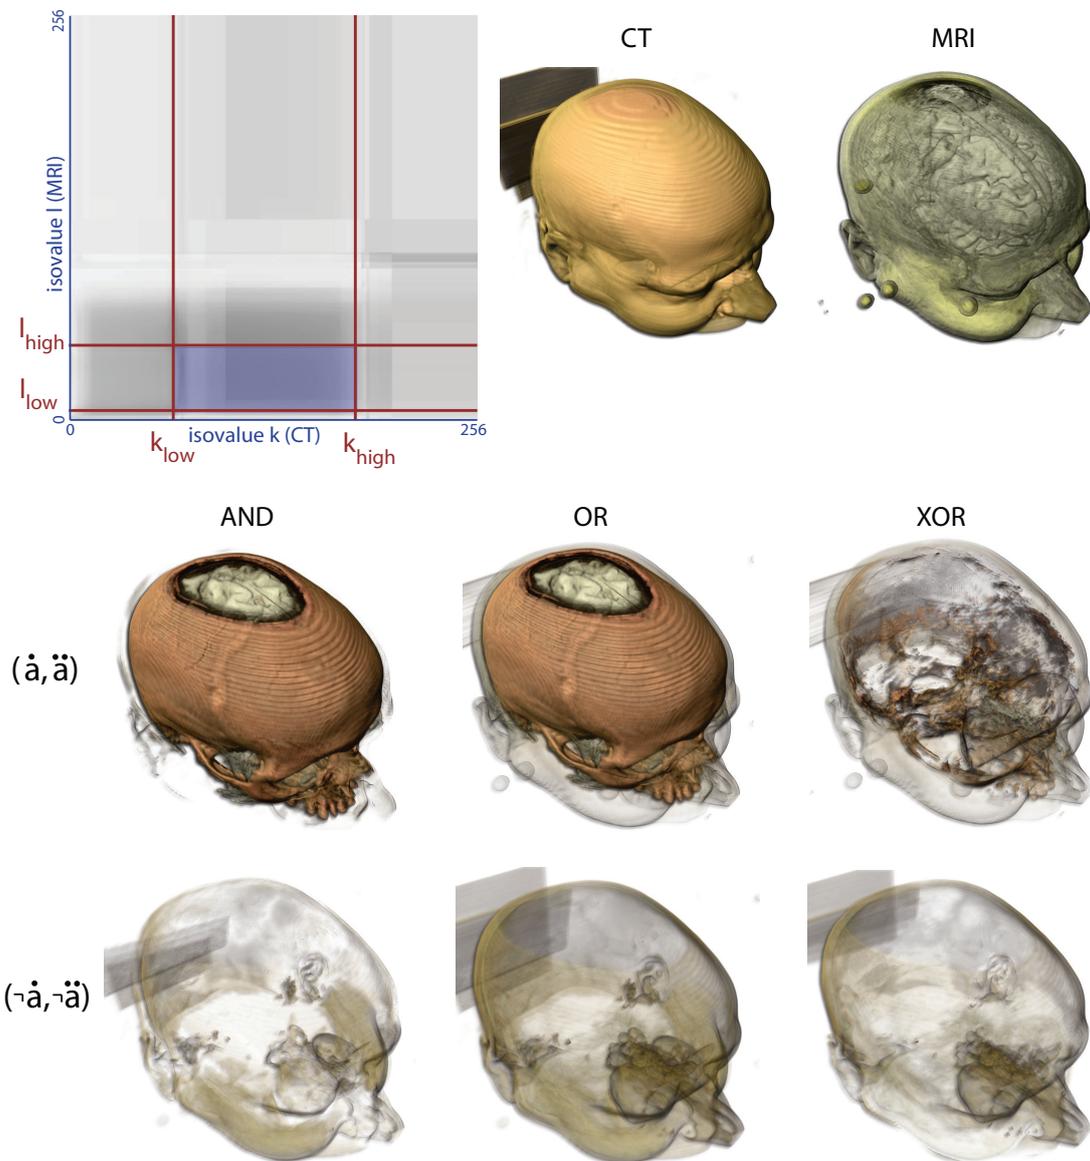


Figure 4.8: The similarity weighting can be narrowed to certain ranges. The logical operations AND, OR, and XOR are used to combine the classification results of both ranges.

selections are depicted. On the right side of the similarity map the two individual modalities are visualized as reference with their corresponding transfer functions.

For the result in Figure 4.8 the goal was to visualize the bone structure from CT together with the brain from MRI. This has been achieved by setting the selection range for CT to isovalues which represent bones and the selection range for MRI to isovalues which represent the brain.

By combining the Boolean variables \dot{a} and \ddot{a} with AND only the bones together with the brain remain visible. As transfer functions for the fused result the unmodified transfer functions of both modalities were used. With the method introduced in Chapter 3 it is possible to achieve similar results for the combination of CT and MRI but it is less intuitive, since the selection is done in a special high dimensional transfer-function space.

With the different logical operations this classification can be used in various modes. It is also possible to depict artifacts in the modalities by the selection range. By using a combination of the inverted Boolean variables $\neg\dot{a}$ and $\neg\ddot{a}$ these artifacts can be eliminated.

4.4.2 Similarity Classifier

The simple operations discussed in the previous section allow 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 want to perform classification based on the similarity of the actual isosurfaces these data values describe. We assume two continuous three-dimensional scalar fields $\dot{f}, \ddot{f} : \mathbb{R}^3 \rightarrow \mathbb{R}$ which represent two co-registered input volumes. For multimodal volume visualization, we want to 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, respectively, \dot{f} and \ddot{f} . 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:

$$\begin{aligned} \dot{s}_i(k) &= \text{MSM}(k, \ddot{h}_i) \\ \ddot{s}_i(l) &= \text{MSM}(\dot{h}_i, l) \end{aligned} \quad (4.11)$$

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:

$$s_i(k, l) = \dot{s}_i(k)\ddot{s}_i(l) \quad (4.12)$$

The rationale behind this choice is that we interpret the similarities $\dot{s}_i(k), \ddot{s}_i(l)$ as independent probabilities of, respectively, 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 $s_i(k, l) = \min(\dot{s}_i(k), \ddot{s}_i(l))$ [82]. 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 want 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 $k \in \hat{f}$, $l \in \check{f}$ will assume the color and opacity of the isovalue pair $h_{m(k,l)}$ which maximizes $s_i(k, l)$:

$$m(k, l) = \arg \max_i s_i(k, l) w_i \quad (4.13)$$

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:

$$m_x = m(\hat{f}(x), \check{f}(x)) \quad (4.14)$$

To visually encode the actual 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 $\alpha(x)$ at the sample position x are then simply:

$$\begin{aligned} c(x) &= c_{m_x} \\ \alpha(x) &= \alpha_{m_x} s_{m_x} \end{aligned} \quad (4.15)$$

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:

$$g(x) = \frac{\hat{s}_i(\hat{f}(x)) \nabla \hat{f}(x) + \check{s}_i(\check{f}(x)) \nabla \check{f}(x)}{\hat{s}_i(\hat{f}(x)) + \check{s}_i(\check{f}(x))} \quad (4.16)$$

4.4.3 Classification Specification

The described classification is equivalent to a generalized - i.e., using non-Euclidean distances defined by our similarity measure - Voronoi decomposition of similarity space. In the spatial domain, the depicted volumetric structures correspond to the Voronoi cells of the chosen isovalue pairs. We can now also visualize this classification on the similarity map itself by simply evaluating Equation 4.13 for each location - i.e., each combination of data values - in the similarity map and coloring the corresponding pixel accordingly.

Figures 4.9 and 4.10 show examples of our classification approach using the previously introduced synthetic data sets. The colored regions signify the nearest neighbors, in similarity space, of each of the white-outlined points in the corresponding color. The slightly larger points with dark outlines have a different meaning which will be explained in the next paragraph. The colored regions therefore constitute generalized Voronoi cells in similarity space. When depicted on the two-dimensional similarity map, where the coordinate system is defined by isovalues, these cells may be disconnected non-convex regions. Furthermore, based on the structure of the similarity map, the site - i.e., the isovalue pair - which defines a region may not be contained

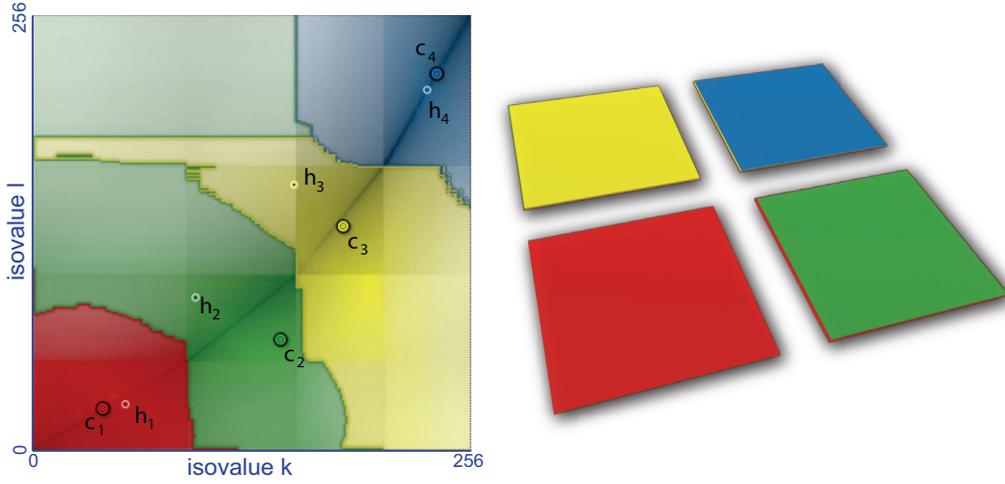


Figure 4.9: Classification based on multimodal similarity for supplementary data.

within this region. While this may initially sound counter-intuitive, the following situation exemplifies such a case. For two isosurfaces with high dissimilarity, there will likely be other isosurfaces they are more similar to than to each other. An example of this can be seen in Figure 4.10 where the isovalue pair which defines the orange region is located on the opposite side of the similarity map.

To provide a means for manipulating the classification regions instead of the isovalues themselves, we define a user-specified control point $c_i = (\hat{c}_i, \check{c}_i)$ for each isovalue pair h_i , which can be freely moved. At the beginning h_i is initialized with the values of c_i . When the control point is modified, we compute the similarity-weighted centroid of the region corresponding to c_i to obtain the new value of h_i :

$$h_i = \frac{\sum_{(k,l) \in R(c_i)} (k,l) s_{m(k,l)}(k,l)}{\sum_{(k,l) \in R(c_i)} s_{m(k,l)}(k,l)} \quad (4.17)$$

where $R(c_i) = \{(k,l) | m(k,l) = i\}$ is the similarity-space region assigned to c_i . This corresponds to one iteration of Lloyd's algorithm [52]. Note, however, that we do not perform the full relaxation as we want users to employ the control points as handles rather than obtaining a centroidal decomposition of the similarity space. Instead, we only want the regions to follow their control points. In Figures 4.9 and 4.10, the control points c_i are depicted as the slightly larger points with dark outlines, while the corresponding isovalue pairs are shown with white 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 , while in regions of low similarity they are essentially mirrored along the axis of maximum similarity.

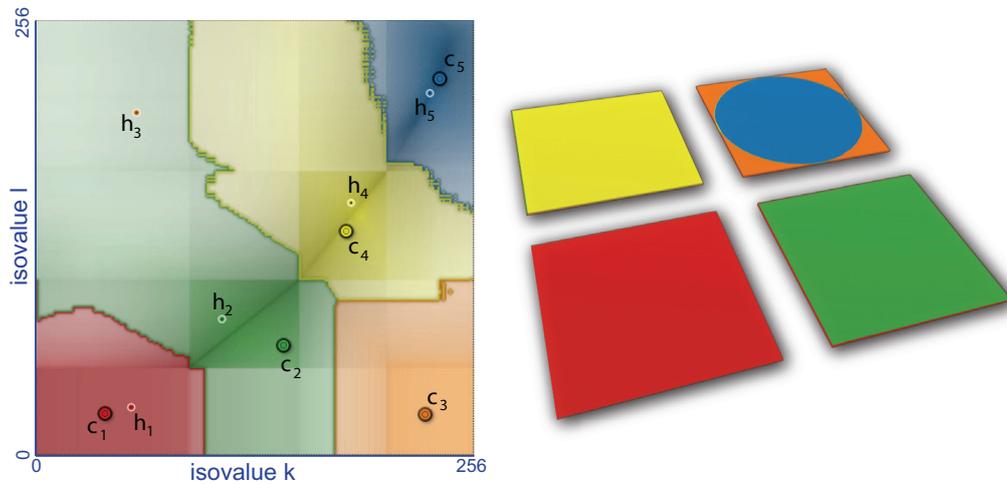


Figure 4.10: Classification based on multimodal similarity for complementary data.

Based on this approach, we can now provide a simple interface for interactively classifying multimodal data sets which takes the place of conventional transfer-function widgets. Users can interactively add and remove control points, move them on the similarity map, and change their colors and opacities. As shown in Figures 4.9 and 4.10, the corresponding classification regions are overlaid over the similarity map. When moving control points they behave similar to well-known "magic wand"-type selection tools – they snap to clusters in the similarity map. Moving a control point will, in accordance with the structure of the similarity map, not cause major changes of the classification result. This stands in stark contrast to conventional transfer functions, where a slight change may cause substantial changes in the resulting image. As mentioned in the previous section, we use a similarity threshold t for discarding samples with low similarity in the 3D visualization. The same threshold is also used in the 2D widget. The additional weights w_i used in Equation 4.13 therefore allow to control the size of the respective region.

An application for which our similarity-based classification approach is particularly suitable is the study of industrial parts using dual energy CT. In such scenarios, 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. An example is shown in Figure 4.11. An isosurface of the low energy scan of a 400 Volt power connector is shown on the top left. It is not possible to find an isovalue which suppresses all artifacts but leaves the surface intact. The center image shows the high-energy scan which gives a better impression of the actual surface, but is very noisy and lacks details. We can remove the artifacts by choosing control points which select regions of high dissimilarity and setting their opacity to zero. The blue control point corresponds to a region with high similarity, so sample points take into account the information of both energy levels. The result is shown on the top right of the figure – we achieve results

comparable to those of Heinzl et al. [28] without any additional filtering. The bottom right image shows that we can select the interior parts of the connector in an analogous manner.

A further example is shown in Figure 4.12. In this case, a dual energy CT angiography 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).

4.5 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 image. For the fusion the similarity of the combination of two isovalues is also retrieved by a single lookup in a 2D texture.

The user interface for our similarity-based classification approach was implemented using the Qt user interface 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 as uniform arrays. Alternatively, these values could also be stored in textures, but we found that common numbers of isovalue pairs are sufficiently low to make the use of uniform parameters feasible.

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 4.11 and 4.12) and the maximum is computed. The color and opacity of the maximally similar isovalue pair then determines the color and opacity of the current sample, as described in Section 4.4.2. The gradient vectors of both modalities, which are stored in a single 2-component 3D texture, are computed by central differences, combined with Equation 4.16, and used to evaluate a local illumination model if shading is enabled.

4.6 Discussion

As shown in our examples, multimodal surface similarity can provide a useful tool for 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.

Data Set	Red. Factor	DF	MI	Total
Supplementary 512×512×6	2	28.57 _{sec}	32.97 _{sec}	61.54 _{sec}
Complementary 512×512×6	2	21.17 _{sec}	32.41 _{sec}	53.58 _{sec}
CT-MRI 256×256×128	4	4.64 _{sec}	32.86 _{sec}	37.50 _{sec}
CT-PET 512×512×128	8	3.31 _{sec}	17.06 _{sec}	20.38 _{sec}
Industrial DECT 425×551×895	16	2.12 _{sec}	12.08 _{sec}	14.20 _{sec}
Medical DECT 512×512×575	16	2.62 _{sec}	11.30 _{sec}	13.92 _{sec}

Table 4.1: Computation times for the multimodal similarity map as measured in an Intel Core i7 950 CPU with a clock rate of 3.07 GHz and 12 GB RAM. DF stands for the calculation of both distance fields. MI stands for the calculation of the mutual information. The reduction factor indicates by which factor all dimensions (x, y, z) were reduced.

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 [6], 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 chapter are given in Table 4.1. The second column in the table lists the downsampling factor for the respective volume which is automatically chosen to limit the computation time to approximately one minute. Even though the reduction factors are chosen 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 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 4.13.

One limitation of our approach 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 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 remains to be explored in future work. Furthermore,

our technique could also be applied to investigate time-dependent data by generating a set of similarity maps between subsequent time steps. Temporal isosurface similarity maps could help to identify stable features and to pinpoint discontinuities.

4.7 Conclusion

In this chapter, 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 similarity information can assist the classification process without requiring the introduction of an entirely new multimodal transfer-function space. Instead, the main contribution of this work is a new approach guiding the visual analysis of multimodal data. Furthermore, we contributed a novel way of exploiting similarity information for interactive classification and manipulation.

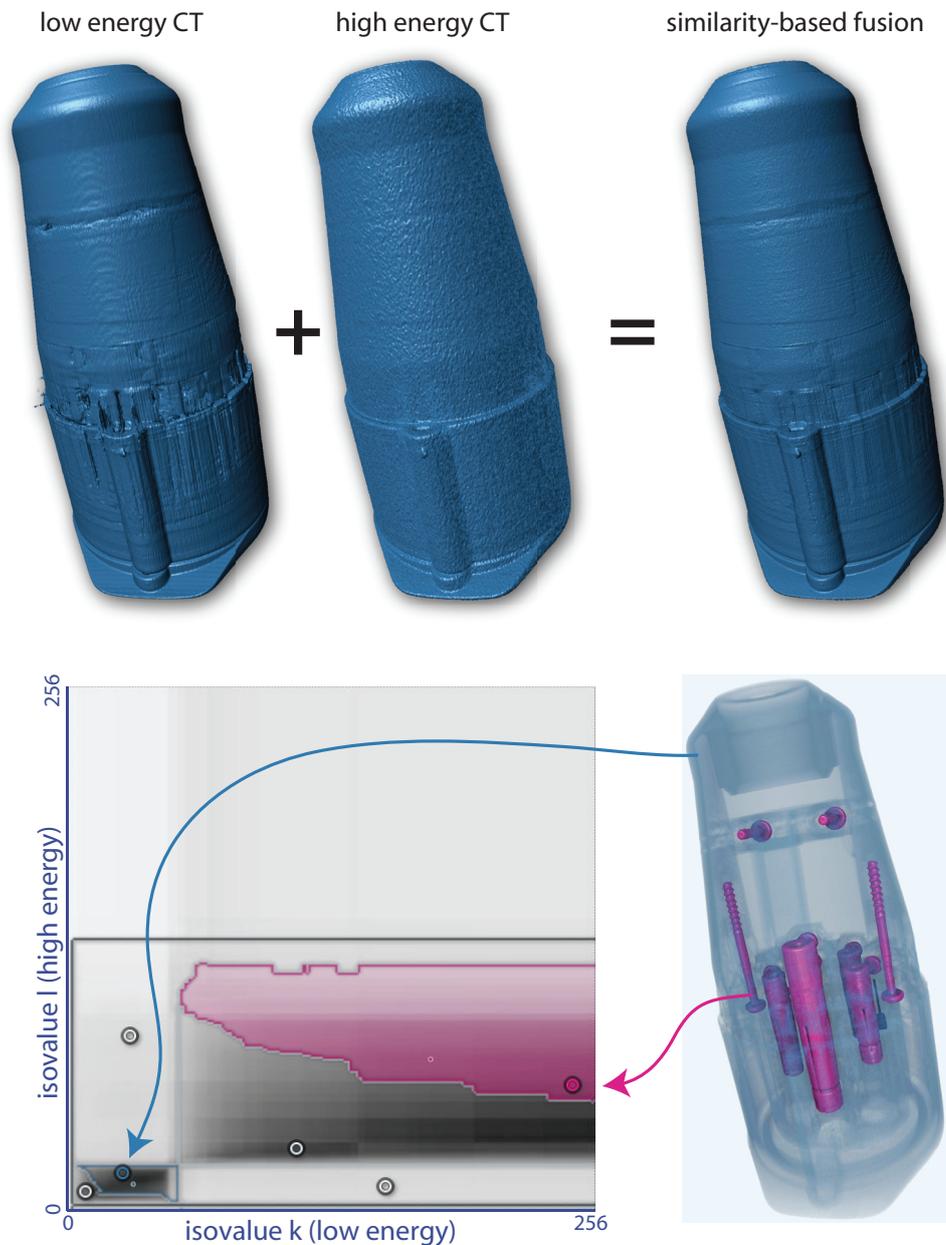


Figure 4.11: Similarity-based fusion of a dual energy CT scan of a power connector. The low-energy scan (top left) and the high-energy scan (top center) provide supplementary information which can be used to remove most of their respective drawbacks (top right). The corresponding similarity map is shown on the bottom left. The bottom right image depicts a different opacity setting which reveals the interior parts of the object.

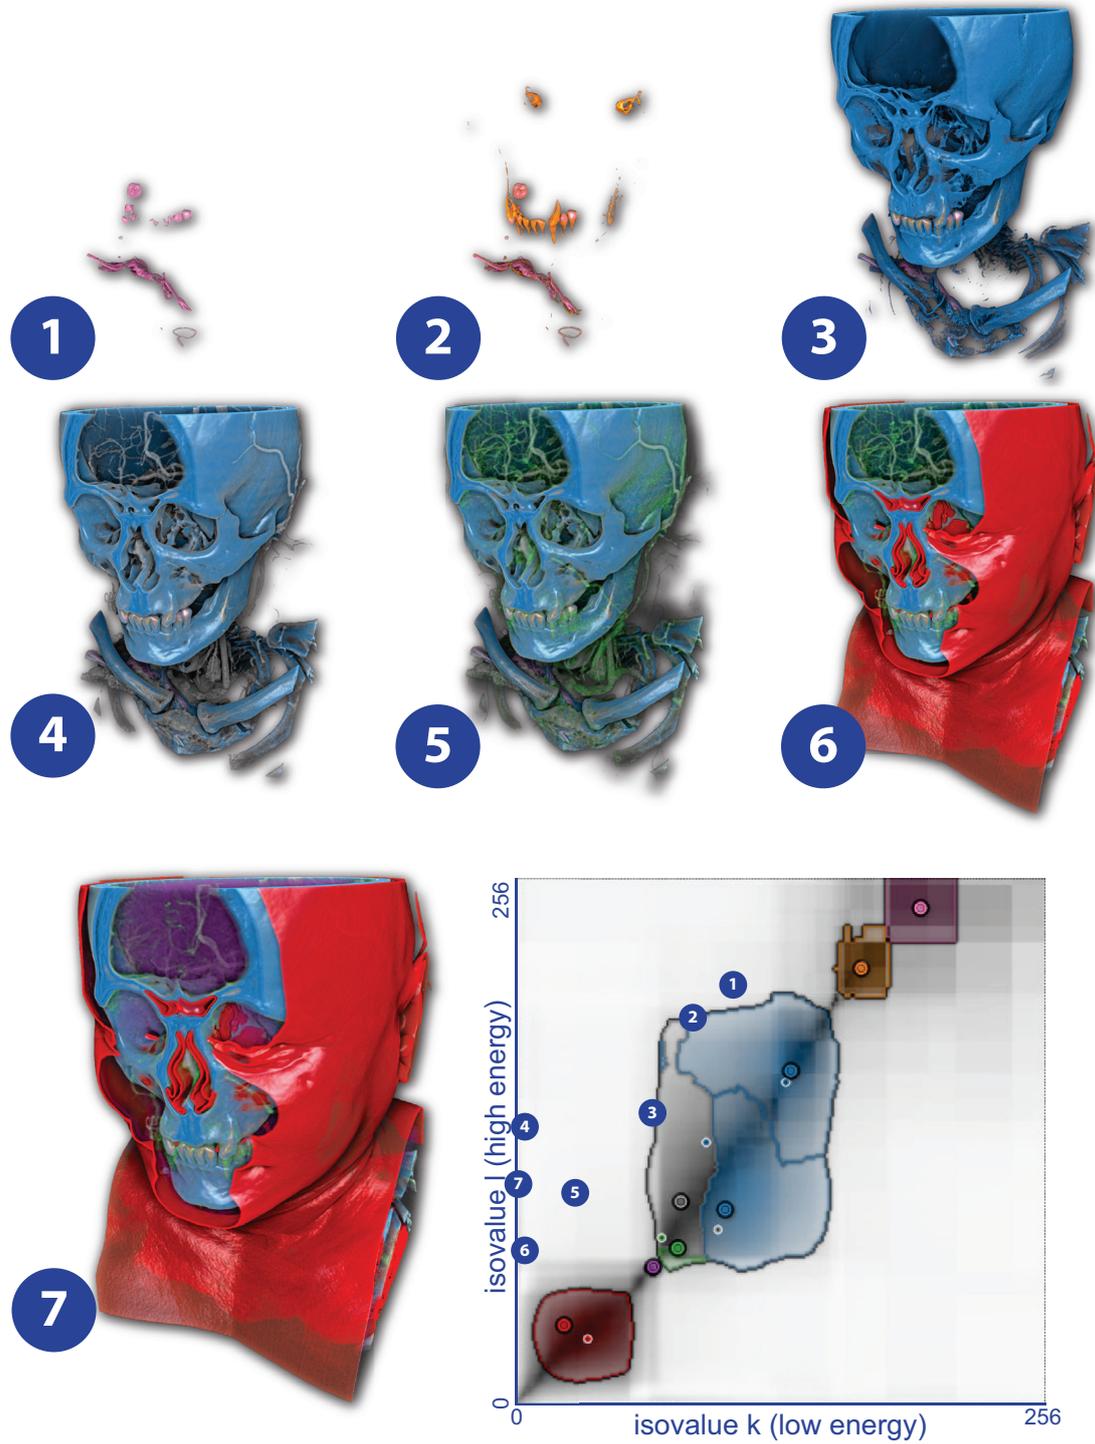


Figure 4.12: Iterative control point specification for similarity-based classification of a dual-energy CT angiography data set. The individual steps are numbered form 1 to 7.

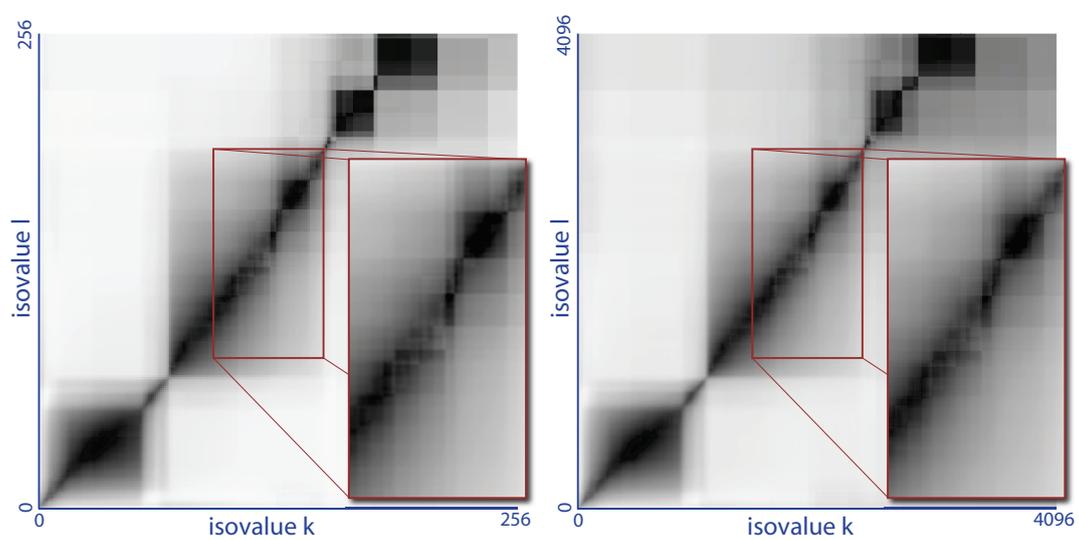


Figure 4.13: A comparison between the multimodal similarity map with a isovalue precision of 8 bits and 12 bits.

I may not have gone where I intended to go, but I think I have ended up where I needed to be.

— Douglas Adams

CHAPTER



.....
Summary

IN the research field of visualization the extraction of information from raw data is a major goal. This can be achieved through the different processing steps of the visualization pipeline by masking parts of the data which contain no information for a certain task. Through this, the remaining parts are enhanced and are easier to perceive by a user of a visualization system.

Information is a measurement which is hard to quantify. Which parts of the whole data set contain more information is depending on the task which has to be performed by the user. Through the complexity and size of data it became difficult for a user to manually interact with all processing steps of the visualization pipeline to enhance the parts with most information. Hence, techniques were developed which support this task by masking parts of the data with less information automatically or semi-automatically.

For the development of such a technique it is necessary to quantify the information content or importance of a certain part of the data set. In recent years information theory and its tools from statistics were discovered by the visualization community for this purpose [74]. Information theory is used in many different research fields - such as signal processing [13], data compression [27], or image processing [16] - for decades and with a high impact. In visualization it is also likely that information theory will become more important over the next years.

In this thesis three different approaches were introduced which already employ information theory and its statistical tools for the classification process along the visualization pipeline. The features which are employed for the classification were chosen in a way to enhance parts of the data which contain more information for a certain task.

Furthermore, the focus in the development of all three methods was on a user-friendly interaction. Many transfer-function spaces lack simple interaction. Either the transfer-function space is too complex by using too many features for the definition of the space or the employed features for the classification are not intuitive for the user. In both cases a user needs a lot of experience and patience to design a good transfer function.

For the methods in this thesis, transfer-function spaces were developed which simplify the design process of a transfer function. This has been achieved by using only a few but very expressive features and a sophisticated design of the transfer-function space. So it is possible to provide transfer-function spaces which are simple and similar to already well-known transfer-function spaces, even though novel features were used for the definition.

The method described in Chapter 2 discusses features for the classification with statistical methods. The features are extracted from an adaptive local neighborhood around each sample

point. With the statistical information about a local neighborhood it is possible to accurately classify different materials in the data set, even though the data is very noisy in general.

A transfer-function space is defined based on the extracted features, i.e., the mean value and standard deviation, of the local neighborhood. Due to the stability of the statistical properties in respect to noise, it is possible to separately classify different materials in this transfer-function space. In other comparable transfer-function spaces this is not possible since the extracted features for different materials overlap each other.

The methods introduced in Chapter 3 and 4 were developed for the classification of multimodal data. This kind of data has become more important in recent years since the acquisition techniques got faster and less expensive.

In Chapter 3 a technique was described which employs information theory to extract features for the classification. These features are dependent on the global distribution of the data values in both modalities. In an intermediate step the values of both modalities are fused according to their information content. By this approach it is possible to enhance the parts of both modalities which contain most of the information.

For the final classification a well-known transfer-function space is used. This makes it more intuitive for a user to define an appropriate transfer function. An additional parameter was introduced to enhance or mask parts of the volume with high mutual information.

The approach in Chapter 4 introduced another classification technique for multimodal data. In this approach the structures of the objects, i.e., isosurfaces, represented in the data were used to extract features for the classification. In a multimodal similarity map the similarities of isosurfaces of both modalities are depicted.

Based on this map a classification technique is presented which enables the extraction of structures which are similar in both modalities. Hence it is, e.g., possible to extract salient surfaces or suppress noise or artifacts. For this purpose the multimodal similarity map is directly used as transfer-function space for the classification. This makes it intuitive and practical for the user since it is not necessary to get used to a new transfer-function space.

All three approaches present new ways of classifying volumetric data based on information theory and statistics. The results in the chapters show that the methods can perform certain visualization tasks better in comparison to other existing approaches. Nevertheless, none of the introduced methods is able to solve all classification issues in scientific visualization. For certain applications and data types information theory and statistics are useful tools to enhance parts of the data which are important for a certain task in the classification process. It might be a good idea for future research to employ information theory for other visualization tasks which are beyond the methods introduced in this thesis. This could lead to a lower information loss along the information pipeline and, hence, to better conclusions about data based on visualizations.



Bibliography

- [1] Hiroshi Akiba and Kwan-Liu Ma. A tri-space visualization interface for analyzing time-varying multivariate volume data. In *Proceedings of EuroVis 2007*, pages 115–122, 2007.
- [2] Julia F. Barrett and Keat Nicholas. Artifacts in CT: Recognition and avoidance. *Radio-graphics*, 24(6):1679–1691, 2004.
- [3] Benjamin B. Bederson and Ben Shneiderman. *The Craft of Information Visualization: Readings and Reflections*. Morgan Kaufmann, 2003.
- [4] Udepta D. Bordoloi and Han-Wei Shen. View selection for volume rendering. In *VIS '05: Proceedings of the IEEE Visualization 2005*, pages 487–494, 2005.
- [5] Stefan Bruckner and M. Eduard Gröller. Instant volume visualization using maximum intensity difference accumulation. *Computer Graphics Forum*, 28(3):775–782, 2009.
- [6] Stefan Bruckner and Torsten Möller. Isosurface similarity maps. *Computer Graphics Forum*, 29(3):773–782, 2010.
- [7] Jesus J. Caban and Penny Rheingans. Texture-based transfer functions for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1364–1371, 2008.
- [8] Wenli Cai and Georgios Sakas. Data intermixing and multi-volume rendering. *Computer Graphics Forum*, 18(3):359–368, 1999.
- [9] Hamish Carr, Jack Snoeyink, and Michiel van de Panne. Flexible isosurfaces: Simplifying and displaying scalar topology using the contour tree. *Computational Geometry: Theory and Applications*, 43(1):42–58, 2010.
- [10] Min Chen and Heike Jänicke. An information-theoretic framework for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1206–1215, 2010.
- [11] Heng Da Da Cheng, Yen-Hung Chen, and Ying Sun. A novel fuzzy entropy approach to image enhancement and thresholding. *Signal Processing*, 75:277–301, 1999.

- [12] Carlos Correa and Kwan-Liu Ma. Size-based transfer functions: A new volume exploration technique. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1380–1387, 2008.
- [13] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2006.
- [14] T. Todd Elvins. A survey of algorithms for volume visualization. *SIGGRAPH Computer Graphics*, 26:194–201, 1992.
- [15] Klaus Engel, Markus Hadwiger, Joe Kniss, Christof Rezk-Salama, and Daniel Weiskopf. *Real-Time Volume Graphics*. A K Peters, Ltd., 2006.
- [16] Francisco Escolano, Pablo Suau, and Boyán Bonev. *Information Theory in Computer Vision and Pattern Recognition*. Springer-Verlag GmbH, 2009.
- [17] Christian Eusemann, David R. Holmes III, Bernhard Schmidt, Thomas G. Flohr, Richard Robb, Cynthia McCollough, David M. Hough, James E. Huprich, Michael Wittmer, Hasan Siddiki, and Joel 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.
- [18] Alan C. Evans, Sean Marrett, Jaime Torrescorzo, Shyan Ku, and Louis 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.
- [19] Brian S. Everitt. *The Cambridge Dictionary of Statistics*. Cambridge University Press, 1998.
- [20] Miquel Feixas, Mateu Sbert, and Francisco González. A unified information-theoretic framework for viewpoint selection and mesh saliency. *ACM Transactions on Graphics*, 6: 1–23, 2009.
- [21] Michael Friendly. Milestones in the history of thematic cartography, statistical graphics, and data visualization. *Engineering*, 9(2):2008, 2008.
- [22] Raphael Fuchs and Helwig Hauser. Visualization of multi-variate scientific data. *Computer Graphics Forum*, 28(6):1670–1690, 2009.
- [23] Mohammad H. Ghavamnia and Xue D. Yang. Direct rendering of Laplacian pyramid compressed volume data. In *VIS '95: Proceedings of the IEEE Visualization 1995*, pages 192–199, 1995.
- [24] Hákon Gudbjartsson and Samuel Patz. The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 34(6):910–914, 1995.
- [25] Robert B. Haber and David A. McNabb. Visualization idioms: A conceptual model for scientific visualization systems. *IEEE Visualization in Scientific Computing*, pages 74–93, 1990.

- [26] Markus Hadwiger, Laura Fritz, Christof Rezk-Salama, Thomas Höllt, Georg Geier, and Thomas Pabel. Interactive volume exploration for feature detection and quantification in industrial CT data. *IEEE Transactions on Visualization and Computer Graphics*, 14(6): 1507–1514, 2008.
- [27] Darrel Hankerson, Greg A. Harris, and Peter D. Johnson Jr. *Introduction to Information Theory and Data Compression*. CRC Press, 1997.
- [28] Christoph Heinzl, Johann Kastner, and M. Eduard 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.
- [29] Jiří Hladůvka, Andreas König, and M. Eduard Gröller. Curvature-based transfer functions for direct volume rendering. In *Spring Conference on Computer Graphics 2000 (SCCG 2000)*, volume 16, pages 58–65, 2000.
- [30] Helen Hong, Juhee Bae, Heewon Kye, and Yeong-Gil Shin. Efficient multimodality volume fusion using graphics hardware. In *Proceedings of the International Conference on Computational Science 2005*, pages 842–845, 2005.
- [31] Jiang Hsieh. *Computed Tomography: Principles, Design, Artifacts and Recent Advances*. SPIE Press, 2003.
- [32] Xiaolei Huang, Nikos Paragios, and Dimitris 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.
- [33] Mark W. Jones, J. Andreas Baerentzen, and Milos Šrámek. 3D distance fields: a survey of techniques and applications. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):581–599, 2006.
- [34] Arie Kaufman and Klaus Mueller. Overview of volume rendering. In Christopher Johnson and Charles Hansen, editors, *The Visualization Handbook*, pages 127–174. Academic Press, Inc., 2004.
- [35] Marc Khoury and Rephael Wenger. On the fractal dimension of isosurfaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1198–1205, 2010.
- [36] Jinman Kim, Stefan Eberl, and Dagan Feng. Visualizing dual-modality rendered volumes using a dual-lookup table transfer function. *Computing in Science and Engineering*, 9(1): 20–25, 2007.
- [37] Gordon Kindlmann and James W. Durkin. Semi-automatic generation of transfer functions for direct volume rendering. In *VVS '98: Proceedings of the IEEE Symposium on Volume Visualization 1998*, pages 79–86, 1998.

- [38] Joe Kniss, Gordon Kindlmann, and Charles Hansen. Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. In *VIS '01: Proceedings of the IEEE Visualization 2001*, pages 255–262, 2001.
- [39] Joe Kniss, Charles Hansen, Michel Grenier, and Tom Robinson. Volume rendering multivariate data to visualize meteorological simulations: a case study. In *VisSym '02: Proceedings of the Symposium on Data Visualisation 2002*, pages 189–195, 2002.
- [40] Joe Kniss, Gordon Kindlmann, and Charles Hansen. Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):270–285, 2002.
- [41] Joe Kniss, Simon Premoze, Milan Ikits, Aaron Lefohn, Charles Hansen, and Emil Praun. Gaussian transfer functions for multi-field volume visualization. In *VIS '03: Proceedings of the IEEE Visualization 2003*, pages 65–72, 2003.
- [42] Joe Kniss, Jürgen P. Schulze, Uwe Wössner, Peter Winkler, Ulrich Lang, and Charles Hansen. Medical applications of multi-field volume rendering and VR techniques. In *VisSym '04: Proceedings of the Symposium on Data Visualisation 2004*, pages 249–254, 2004.
- [43] Joe Kniss, Robert Van Uitert, Abraham Stephens, Guo-Shi Li, Tolga Tasdizen, and Charles Hansen. Statistically quantitative volume visualization. In *VIS '05: Proceedings of the IEEE Visualization 2005*, pages 287–294, 2005.
- [44] Tarald O. Kväseth. Entropy and correlation: Some comments. *IEEE Transactions on Systems, Man, and Cybernetics*, 17:517–519, 1987.
- [45] David H. Laidlaw, Kurt W. Fleischer, and Alan H. Barr. Partial-volume Bayesian classification of material mixtures in MR volume data using voxel histograms. *IEEE Transactions on Medical Imaging*, 17(1):74–86, 1998.
- [46] Christopher Lau, Sayan D. Pathak, Lixin Gong, Paul Kinahan, Phillip Cheng, and Lydia Ng. Advanced PET/CT fusion workstation for oncology imaging. In *Proceedings of SPIE Medical Imaging 2005*, pages 670–676, 2005.
- [47] Tianhu Lei and Wilfred Sewchand. Statistical approach to X-ray CT imaging and its applications in image analysis - Part I: Statistical analysis of X-ray CT imaging. *IEEE Transactions on Medical Imaging*, 11(2):53–61, 1992.
- [48] David N. Levin, Xiaoping Hu, Kim K. Tan, Simranjit Galhotra, Charles A. Pelizzari, George T. Chen, Robert N. Beck, Chin-Tu Chen, Malcom Cooper, and John Mullan. The brain: integrated three-dimensional display of MR and PET images. *Radiology*, 172: 783–789, 1989.
- [49] Marc Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8(3):29–37, 1988.

- [50] Tony Lindeberg. Scale-space for discrete signals. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:234–254, 1990.
- [51] Tony Lindeberg. Feature detection with automatic scale selection. *International Journal of Computer Vision*, 30:79–116, 1998.
- [52] Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [53] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *SIGGRAPH Computer Graphics*, 21:163–169, 1987.
- [54] Eric B. Lum and Kwan-Liu Ma. Lighting transfer functions using gradient aligned sampling. In *VIS '04: Proceedings of the IEEE Visualization 2004*, pages 289–296, 2004.
- [55] Eric B. Lum, James Shearer, and Kwan-Liu Ma. Interactive multi-scale exploration for volume classification. *The Visual Computer*, 22(9–11):622–630, 2006.
- [56] Claes Lundström, Patric Ljung, and Anders Ynnerman. Local histograms for design of transfer functions in direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 12(6):1570–1579, 2006.
- [57] Claes Lundström, Patric Ljung, and Anders Ynnerman. Multi-dimensional transfer function design using sorted histograms. In *Proceedings Eurographics/IEEE Workshop on Volume Graphics 2006*, pages 1–8, 2006.
- [58] Shigeru Muraki. Volume data and wavelet transforms. *IEEE Computer Graphics and Applications*, 13(4):50–56, 1993.
- [59] Daniel Patel, Martin Haidacher, Jean-Paul Balabanian, and Meister Eduard Gröller. Moment curves. In *Proceedings of the IEEE Pacific Visualization Symposium 2009*, pages 201–208, 2009.
- [60] Hanspeter Pfister, Chandrajit Bajaj, Will Schroeder, and Gordon Kindlmann. The transfer function bake-off. *VIS '00: Proceedings of the IEEE Visualization 2000*, pages 523–526, 2000.
- [61] Bui Tuong Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.
- [62] Christof Rezk-Salama, Maik Keller, and Peter Kohlmann. High-level user interfaces for transfer function design with semantics. In *VIS '06: Proceedings of the IEEE Visualization 2001*, pages 1021–1028, 2006.
- [63] Stefan Roettger, Michael Bauer, and Marc Stamminger. Spatialized transfer functions. In *Proceedings of EuroVis 2005*, pages 271–278, 2005.

- [64] Lothar R. Schad, Robert Boesecke, Wolfgang Schlegel, Günther H. Hartmann, Volker Sturm, Ludwig G. Strauss, and Walter J. Lorenz. Three dimensional image correlation of CT, MR, and PET studies in radiotherapy treatment planning of brain tumors. *Journal of Computer Assisted Tomography*, 11(6):948–954, 1987.
- [65] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423,623–656, 1948.
- [66] Rik Stokking, Karel J. Zuiderveld, Hilleke E. Hulshoff Pol, and Max A. Viergever. SPECT/MRI visualization for frontal-lobe-damaged regions. *Visualization in Biomedical Computing 1994*, 2359(1):282–290, 1994.
- [67] Shigeo Takahashi, Yuriko Takeshima, Issei Fujishiro, and Gregory M. Nielson. Emphasizing isosurface embeddings in direct volume rendering. In Georges-Pierre Bonneau, Thomas Ertl, and Gregory M. Nielson, editors, *Scientific Visualization: The Visual Extraction of Knowledge from Data*, pages 185–206. Springer, 2006.
- [68] Shivaraj Tenginakai and Raghu Machiraju. Statistical computation of salient iso-values. In *VisSym '02: Proceedings of the Symposium on Data Visualisation 2002*, pages 19–24, 2002.
- [69] Shivaraj Tenginakai, Jinho Lee, and Raghu Machiraju. Salient iso-surface detection with model-independent statistical signatures. In *VIS '01: Proceedings of the IEEE Visualization 2001*, pages 231–238, 2001.
- [70] Edward R. Tufte. *Visual Explanations*. Graphics Press, 1990.
- [71] Ivan Viola, Miquel Feixas, Mateu Sbert, and M. Eduard Gröller. Importance-driven focus of attention. *IEEE Transactions on Visualization and Computer Graphics*, 12:933–940, 2006.
- [72] Petr Šereda, Anna Vilanova Bartrolí, Iwo W.O. Serlie, and Frans A. Gerritsen. Visualization of boundaries in volumetric data sets using LH histograms. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):208–218, 2006.
- [73] Chaoli Wang and Han-Wei Shen. LOD map - a visual interface for navigating multiresolution volume visualization. *IEEE Transaction on Visualization and Computer Graphics*, 12:1029–1036, 2006.
- [74] Chaoli Wang and Han-Wei Shen. Information theory in scientific visualization. *Entropy*, 13(1):254–273, 2011.
- [75] Jing Wang, Tianfang Li, Hongbing Lu, and Zhengrong Liang. Penalized weighted least-squares approach to sinogram noise reduction and image reconstruction for low-dose X-ray computed tomography. *IEEE Transactions on Medical Imaging*, 25(10):1272–1283, 2006.
- [76] Zhijun Wang, Djemel Ziou, Costas Armenakis, Deren Li, and Qingquan Li. A comparative analysis of image fusion methods. *IEEE Transactions on Geoscience and Remote Sensing*, 43:1391–1402, 2005.

- [77] B. L. Welch. The generalization of "student's" problem when several different population variances are involved. *Biometrika*, 34(1-2):28-35, 1947.
- [78] William M. Wells III, Paul Viola, Hideki Atsumi, Shin Nakajima, and Ron Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1:35-51, 1996.
- [79] Lee Westover. Footprint evaluation for volume rendering. *SIGGRAPH Computer Graphics*, 24:367-376, 1990.
- [80] Lijie Xu, Teng-Yok Lee, and Han-Wei Shen. An information-theoretic framework for flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6): 1216-1224, 2010.
- [81] Yiyu 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.
- [82] Lotfi A. Zadeh. Fuzzy sets. *Information and Control*, 8:338-353, 1995.
- [83] Karel J. Zuiderveld and Max A. Viergever. Multi-modal volume visualization using object-oriented methods. In *VVS '94: Proceedings of the IEEE Symposium on Volume Visualization 1994*, pages 59-66, 1994.