

Information-based Transfer Functions for Multimodal Visualization

Martin Haidacher¹, Stefan Bruckner¹, Armin Kanitsar², and M. Eduard Gröller¹

¹Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria

²AGFA HealthCare, Vienna, Austria

Abstract

Transfer functions are an essential part of volume visualization. In multimodal visualization at least two values exist at every sample point. Additionally, other parameters, such as gradient magnitude, are often retrieved for each sample point. To find a good transfer function for this high number of parameters is challenging because of the complexity of this task. In this paper we present a general information-based approach for transfer function design in multimodal visualization which is independent of the used modality types. Based on information theory, the complex multi-dimensional transfer function space is fused to allow utilization of a well-known 2D transfer function with a single value and gradient magnitude as parameters. Additionally, a quantity is introduced which enables better separation of regions with complementary information. The benefit of the new method in contrast to other techniques is a transfer function space which is easy to understand and which provides a better separation of different tissues. The usability of the new approach is shown on examples of different modalities.

Categories and Subject Descriptors (according to ACM CCS): I.4.10 [Image Processing and Computer Vision]: Volumetric, Multidimensional

1. Introduction

Volume visualization is a technique which enables physicians and scientists to gain insight into complex volumetric structures. Currently, the trend towards information acquisition using data sets from multiple modalities is increasing in order to facilitate better medical diagnosis. As different modalities frequently carry complementary information, our goal is to combine their strengths providing the user with a consistent interface.

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 overlap 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 paper, 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. This is done by using the information contained in the distribution of values in both modalities. Based on this information, the values of both modalities are fused. 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 tissues.

In Section 3 the new 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 4 briefly describes an efficient implementation of the new method. The usability of the new method is shown in Section 5 with some results. Conclusions and ideas for further work are given in Section 6. First an overview over related works on this topic is given in the following section.

2. Related Work

All different methods for multimodal visualization can be classified - as described by Cai and Sakas [CS99] - 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 [SBS*87, SZHV94].

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. [LHT*89] generated a surface model from an MRI scan and mapped the PET-derived measurement onto this surface. Evans et al. [EMT*91] 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 [ZV94]. 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. [HBKS05] 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. [KHGR02]. In this work the idea of multi-dimensional transfer functions to assign optical

properties to a combination of values was used. Akiba and Ma [AM07] 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. [KSW*04]. The classification is done on the basis of the dual histogram. Kim et al. [KEF07] 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 define 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 [Lev88]. 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 [KD98] introduced a semi-automatic approach for the visualization of boundaries between tissues. Pfister et al. [PBSK00] gave an overview on existing techniques to support the design task of transfer functions. The direct manipulation widgets introduced by Kniss et al. [KKH01] 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. [KPI*03] describe a way 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 paper, these additional properties are retrieved by methods from information theory founded by Shannon [Sha48]. 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. [WVA*96] 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 paper.

Rezk-Salama et al. [RSKK06] employed PCA 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. Information-based Transfer Functions for Multimodal Visualization

In this section we introduce a novel transfer function space for multimodal visualization. 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. The main contribution of the new approach is the use of methods from information theory for the design of this transfer function space. Figure 1 shows all necessary processing steps to classify a tuple of input values (f_1, f_2) in this new transfer function space with optical properties. The further sections describe these processing steps in detail.

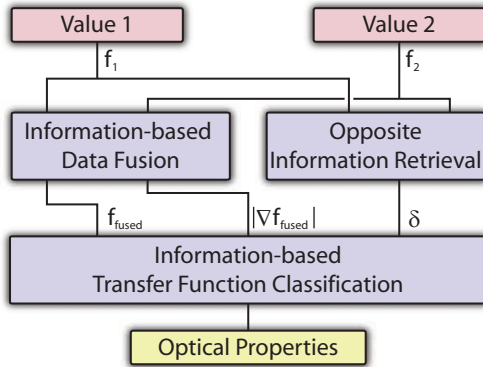


Figure 1: Processing pipeline for the classification of sample points in a multimodal visualization by an information-based transfer function.

In Section 3.2, we describe how the input values can be fused to get just a single value for each pair of input values. Section 3.3 introduces an additional property which is used to refine the classification of different tissues through the transfer function. Finally, Section 3.4 describes how the fused values are used to define the new transfer function space and how the additional property is used to influence the classification. First of all we describe how the probabilities are estimated for the further approach.

3.1. Probabilities in Volume Data

To estimate the probabilities within the volume we first assume 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 occurrence 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 a histogram the count of a bin is increased if a value falls in 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 $P(f_1, f_2)$ for each tuple of values f_1 from modality 1 and f_2 from modality 2, also referred to as joint probability. Equally 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 $P(f_1, f_2)$ the probability of just a single value $P(f_1)$ 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. 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 a fusion 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 the 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 the information is based on the quantity, i.e. frequency, of the data. For this measurement the methods of information theory are used. 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:

$$I(f) = -\log_2(P(f)) \quad (1)$$

where $P(f)$ is the probability of occurrence for a certain value f . Through the \log_2 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 want to introduce the following equation:

$$\gamma(f_1, f_2) = \frac{I(f_2)}{I(f_1) + I(f_2)} \quad (2)$$

The γ value is 0 when the second modality has no information. It is 1 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 2 we get a number for each pair of values which can directly be used for the weighting in the fusion step. The fusion of two values, f_1 and f_2 , is simply done by the following equation:

$$f_{fused} = (1 - \gamma) * f_1 + \gamma * f_2 \quad (3)$$

The fused value f_{fused} is close to the value of one modality when this modality contains 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:

$$\nabla f_{fused} = (1 - \gamma) * \nabla f_1 + \gamma * \nabla f_2 \quad (4)$$

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.

3.3. 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 used 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 tool 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 pointwise mutual information (PMI) is a quantity for the mutual information for a certain combination of points. It is defined by the following equation:

$$PMI(f_1, f_2) = \log_2 \left(\frac{P(f_1, f_2)}{P(f_1) * P(f_2)} \right) \quad (5)$$

The PMI is 0 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 $P(f_1, f_2)$ is exactly the product of both marginal probabilities $P(f_1)$ and $P(f_2)$. If they occur together more frequently as one would expect by chance then the result of the calculation is greater than 0. Conversely, the value is lower than 0 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 0 because the occurrence is less frequent. For a joint probability $P(f_1, f_2)$ of 0 the PMI is by definition 0. For all other probabilities the PMI can be normalized to a value between 0 and 1 by subtracting the lower bound ($P(f_1) = 1$ and $P(f_2) = 1$) from the PMI and dividing it by the difference between the upper bound ($P(f_1) = P(f_1, f_2)$ and $P(f_2) = P(f_1, f_2)$) and the lower bound:

$$PMI_{norm}(f_1, f_2) = \frac{PMI(f_1, f_2) - \log_2(P(f_1, f_2))}{\log_2\left(\frac{1}{P(f_1, f_2)}\right) - \log_2(P(f_1, f_2))} \quad (6)$$

The value of PMI_{norm} approaches 0 if the information carried by the pair of values is high. Values close to 1 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} :

$$\delta(f_1, f_2) = 1 - PMI_{norm}(f_1, f_2) \quad (7)$$

Figure 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 ($P(f_1)$ and $P(f_2)$) 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 $P(f_1, f_2)$ 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




Figure 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.

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 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 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 in the classification process.

3.4. Information-based Transfer Function Classification

In the previous two sections we described how methods from information theory can be used to generate a fused value and 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.

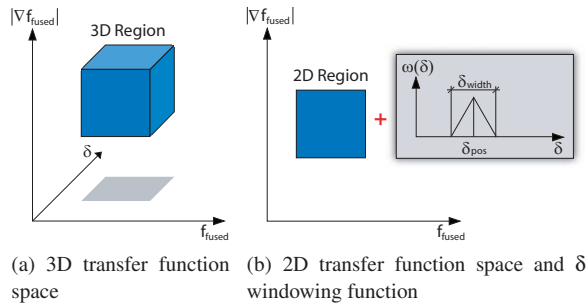


Figure 3: Transfer function space is converted from 3D (a) to 2D (b). Additionally, a simple windowing function for the δ value is used to modify the optical properties of each 2D region.

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(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 this task 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 perfect for this purpose. The windowing function can be expressed by the following equation:

$$\omega(\delta) = \max\left(\left|1 - \frac{\delta - \delta_{pos}}{0.5 * \delta_{width}}\right|, 0\right) \quad (8)$$

The parameters δ_{pos} and δ_{width} define the position and shape of the windowing function $\omega(\delta) \in [0, 1]$. 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(b) the separation in a 2D region and a corresponding windowing function is shown.

4. Implementation

For a fast and efficient volume rendering it is necessary to do as many calculations as possible in a pre-process. 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 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 round 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. With the fused value and the magnitude of the fused gradient a lookup in the lookup tables of the transfer function is done. One 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

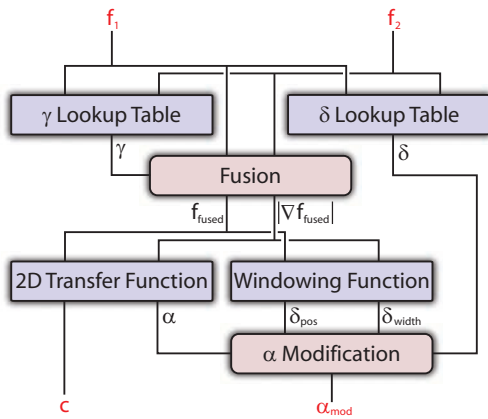


Figure 4: Overview over the processing steps for each sample point during the rendering process. Nodes with round corners are calculation steps and nodes with sharp corners are lookups.

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 no big issue because all processing steps can be executed quite fast on the graphics hardware to achieve real-time frame rates.

5. 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 structures. Soft tissues have a higher contrast in MRI. In Figure 5(a) a visualization of a CT scan is shown and in Figure 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 5(c) shows the rendering results 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 are visible 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 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.

Figures 5 (e) to (h) show the corresponding histograms for the visualizations in Figures 5 (a) to (d). 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 5(h), are highly related to the individual regions in the single modality visualizations, as shown in Figure 5(e) and Figure 5(f). The regions for the multimodal visualization, based on the dual histogram, are shown in Figure 5(g). The position and shape of the regions in this transfer function space are completely different in comparison to the regions for the single modality visualization. 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 3.4 the definition of a transfer function is done in two steps. In Figure 5(h) 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 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 used for Figure 5(d). It can be seen that the result contains a lot of artifacts. In comparison to that, Figure 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.

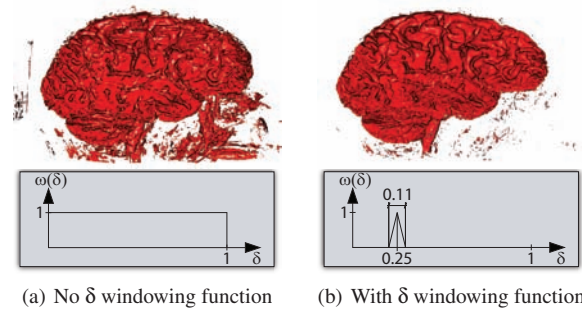


Figure 6: The two results show the effect of the usage of δ to modify the optical properties of a 2D region in the transfer function space.

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

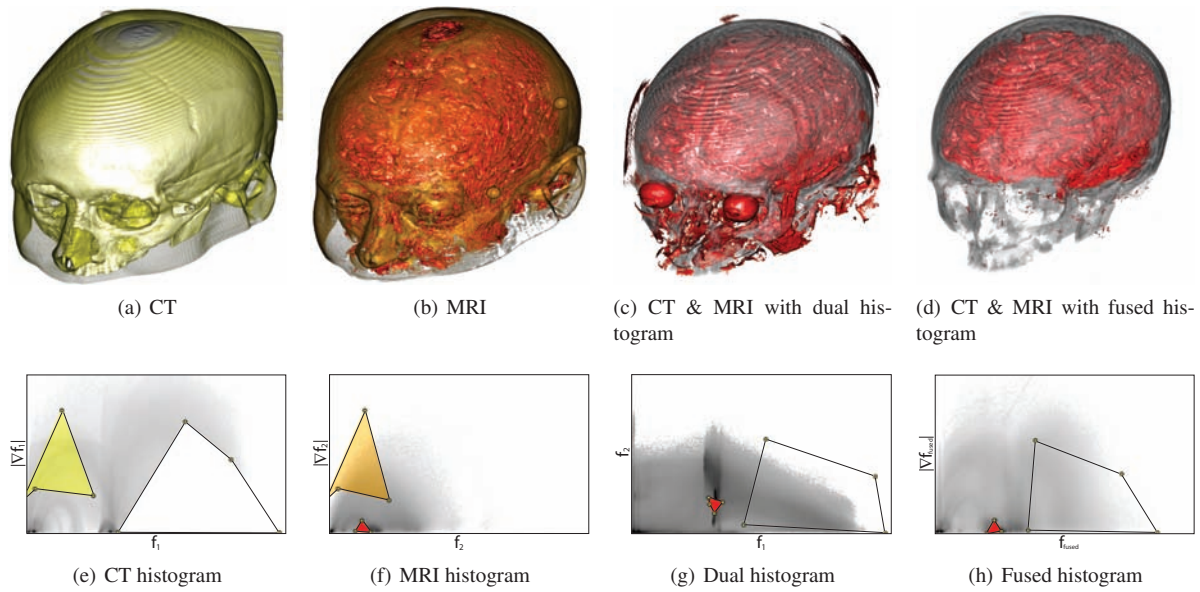


Figure 5: The images show single volume visualizations of CT data (a) and MRI data (b) in contrast to multimodal visualizations by using the dual transfer function space (c) and the fused transfer function space (d). Histograms (e)-(h) with the colored 2D regions for the assignment of optical properties correspond with the above visualizations.

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 3.3 regions with opposite information in both data sets have a high δ value. Figure 7 shows the response of the δ value for the combination of two example data sets. In Figure 7(a) and Figure 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 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 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.

Figure 8 shows a 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 proves that the method also works with the combination of anatomical and functional modalities and, furthermore, with different spatial resolutions.




Figure 8: Multimodal visualization of a CT and PET scan. The more opaque regions symbolizes regions of high activity such as in the brain and in the tumor on the neck.

6. Conclusion and Future Work

In this paper we introduced a novel approach for the definition of transfer functions for multimodal visualization. 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 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. An additional δ value, which describes the complementary information contained in a pair of values, is used for a better separation of different tissues. In the result section we have shown how the new transfer function space can be used to show relevant parts of both modalities together.

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. But anyway, the user can control this by defining a transfer function which has low opacity for such unimportant parts.

An idea for future work is the extension of the method to more than two modalities. For this reason the approach can be modified to generate a single fused value and fused gradient as a combination of all values and gradients from all modalities. The same modification can be done with the calculation of the opposite information as well. With these modifications we still get a transfer function space with the same dimensionality as for two modalities. The question is if different tissues are still separable in this transfer function space. The answer to that question will be part of the new research.

References

- [AM07] AKIBA H., MA K.-L.: A tri-space visualization interface for analyzing time-varying multivariate volume data. In *Proceedings of Eurographics/IEEE VGTC Symposium on Visualization* (2007), pp. 115–122.
- [CS99] CAI W., SAKAS G.: Data intermixing and multi-volume rendering. In *Computer Graphics Forum* (1999), vol. 18, pp. 359–368.
- [EMT*91] EVANS A., MARRETT S., TORRESCORZO J., KU S., COLLINS L.: MRI-PET correlation in three dimensions using a volume-of-interest (VOI) atlas. *Journal of Cerebral Blood Flow and Metabolism* 11, 2 (1991), A69–A78.
- [HBKS05] HONG H., BAE J., KYE H., SHIN Y.-G.: Efficient multimodality volume fusion using graphics hardware. In *International Conference on Computational Science (3)* (2005), pp. 842–845.
- [KD98] KINDLMANN G., DURKIN J. W.: Semi-automatic generation of transfer functions for direct volume rendering. In *VVS '98: Proceedings of the 1998 IEEE Symposium on Volume Visualization* (1998), pp. 79–86.
- [KEF07] KIM J., EBERL S., FENG D.: Visualizing dual-modality rendered volumes using a dual-lookup table transfer function. *Computing in Science and Engineering* 9, 1 (2007), 20–25.
- [KHGR02] KNISS J., HANSEN C., GRENIER M., ROBINSON T.: Volume rendering multivariate data to visualize meteorological simulations: a case study. In *VISSYM '02: Proceedings of the symposium on Data Visualisation 2002* (2002), pp. 189–195.
- [KKH01] KNISS J., KINDLMANN G., HANSEN C.: Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. In *VIS '01: Proceedings of the 12th IEEE Visualization 2001* (2001), pp. 255–262.
- [KPI*03] KNISS J., PREMOZE S., IKITS M., LEFOHN A., HANSEN C., PRAUN E.: Gaussian transfer functions for multi-field volume visualization. In *VIS '03: Proceedings of the 14th IEEE Visualization 2003* (2003), pp. 65–72.
- [KSW*04] KNISS J., SCHULZE J. P., WÖSSNER U., WINKLER P., LANG U., HANSEN C.: Medical applications of multi-field volume rendering and VR techniques. In *Proceedings of Eurographics/IEEE VGTC Symposium on Visualization* (2004), pp. 249–254.
- [Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications* 8, 3 (1988), 29–37.
- [LHT*89] LEVIN D., HU X., TAN K., GALHOTRA S., PELIZZARI C., CHEN G., BECK R., CHEN C., COOPER M., MULLAN J.: The brain: integrated three-dimensional display of MR and PET images. *Radiology* 172 (1989), 783–789.
- [PBSK00] PFISTER H., BAJAJ C., SCHROEDER W., KINDLMANN G.: The transfer function bake-off. *VIS '00: Proceedings of the 11th IEEE Visualization 2000* (2000), 523–526.
- [RSKK06] REZK-SALAMA C., KELLER M., KOHLMANN P.: High-level user interfaces for transfer function design with semantics. In *VIS '06: Proceedings of the 17th IEEE Visualization 2001* (2006), pp. 1021–1028.
- [SBS*87] SCHAD L., BOESECKE R., SCHLEGEL W., HARTMANN G., STURM V., STRAUSS L., LORENZ W.: Three dimensional image correlation of CT, MR, and PET studies in radiotherapy treatment planning of brain tumors. *Journal of Computer Assisted Tomography* 11, 6 (1987), 948–954.
- [Sha48] SHANNON C. E.: A mathematical theory of communication. *Bell System Technical Journal* 27 (1948), 379–423, 623–656.
- [SZHV94] STOKKING R., ZUIDERVELD K. J., HULSHOFF POL H. E., VIERGEVER M. A.: SPECT/MRI visualization for frontal-lobe-damaged regions. *Visualization in Biomedical Computing 1994* 2359, 1 (1994), 282–290.
- [WVA*96] WELLS III W. M., VIOLA P., ATSUMI H., NAKAJIMA S., KIKINIS R.: Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis* 1 (1996), 35–51.
- [ZV94] ZUIDERVELD K. J., VIERGEVER M. A.: Multi-modal volume visualization using object-oriented methods. In *VVS '94: Proceedings of the 1994 IEEE Symposium on Volume Visualization* (1994), pp. 59–66.