

Augmented Reality Interaction for Semiautomatic Volume Classification

A. del Río[†] J. Fischer[‡] M. Köbele D. Bartz and W. Straßer

WSI/GRIS-VCM, University of Tübingen, Germany

Abstract

In the visualization of 3D medical data, the appropriateness of the achieved result is highly dependent on the application. Therefore, an intuitive interaction with the user is of utter importance in order to determine the particular aim of the visualization. In this paper, we present a novel approach for the visualization of 3D medical data with volume rendering combined with AR-based user interaction. The utilization of augmented reality (AR), with the assistance of a set of simple tools, allows the direct manipulation in 3D of the rendered data. The proposed method takes into account regions of interest defined by the user and employs this information to automatically generate an adequate transfer function. Machine learning techniques are utilized for the automatic creation of transfer functions, which are to be used during the classification stage of the rendering pipeline. The validity of the proposed approach for medical applications is illustrated.

Categories and Subject Descriptors (according to ACM CCS): H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems — Artificial, augmented, and virtual realities; I.2.6 [Artificial Intelligence]: Learning — Knowledge acquisition; I.3.6 [Computer Graphics]: Methodologies and Techniques — Interaction techniques

1. Introduction

The application of visualization techniques to medical data has been a topic of intense research during the last decades. This interest has experienced an important impulse in recent years, due to the fast development and improvement of scanning devices such as computed tomography (CT) or magnetic resonance tomography (MRT) scanners. These devices provide a three-dimensional reproduction of the patient anatomy. Direct volume rendering provides a powerful alternative for displaying scanned volume data in a three-dimensional environment. Since all the data is directly processed and incorporated to the scene, it generates a valuable global overview of the whole dataset at once. However, even though such a general insight of the volume can be of great interest as a first approach to its visualization, usually only some regions within the dataset are of specific concern and must be highlighted with respect to the rest of the anatomy. Taking into account that this is highly dependent on the specific application, identifying and displaying these regions or features of interest can become a very challenging task. Classification is associated with direct volume ren-

dering and consists on the assignment of certain renderable optical properties (typically color and opacity) to different regions within the inspected dataset. This is normally performed through the application of a transfer function, which defines the correspondence between internal parameters of the analyzed data (i.e. intensity value) and the color and opacity values to be used during rendering.

An appropriate classification is mandatory in order to achieve a successful visualization. In medical applications, the aim of the visualization is to provide an informative insight of the scanned data. Since the information sought after varies on each specific case, user interaction is important to generate a proper classification. By translating the interaction between the user and the analyzed data into an augmented reality environment, a better and more direct manipulation of the volume is enabled. Moreover, given the inherent three-dimensional nature of the data, the definition of a key factor for its visualization, such as the transfer function, clearly benefits from a direct real-time interaction in 3D.

In this work, we propose a novel semi-automatic strategy for informative volume classification, based on user interaction in an AR environment. The remainder of this paper is structured as follows: In the following section, we provide a brief overview of related work. In Section 3, we introduce

[†] angel.delrio@gris.uni-tuebingen.de

[‡] fischer@gris.uni-tuebingen.de

ARGUS, our medical AR system. Following in Section 4, the transfer function definition process is described in detail. Results obtained with real clinical data are presented in Section 5. Finally, we summarize the most relevant aspects of our work in Section 6.

2. Related Work

Finding a good transfer function for volume classification is one of the key problems in direct volume rendering. Despite its crucial relevance, it has not been until recent years that a considerable research effort has been put on addressing this problem [PLB*01]. In the simplest case, a transfer function has a one dimensional domain (intensity), while its range can be characterized by one (opacity) or four dimensions (color and opacity). Even in this simple scenario, the standard approach of trial and error is a complex and time consuming task. This has resulted in several efforts being proposed to simplify this exploration process. He et al. [HHKP96] treat the search for a proper transfer function as a parameter optimization problem and employ heuristic search techniques, either based on user selection of intermediate thumbnail renderings, or automatically controlled by some objective image fitness function. Marks et al. [MAB*97] address the general problem of computer-assisted parameter setting with a Design Gallery interface, which generates a selection of images obtained with a broad selection of all possible transfer functions. A similar approach is presented in [KG01], where thumbnail renderings are also employed to assist the user, but decoupling the different search domains (data range, color, opacity) and benefiting from the use of the VolumePro technology [PHK*99] to provide real-time visual feedback.

All these techniques propose alternatives to facilitate the design of one-dimensional transfer functions. However it has become clear that the use of multi-dimensional transfer functions can drastically benefit the success of volume visualization [PLB*01]. Such approaches, first proposed by Levoy in 1988 [Lev88], incorporate additional parameters to the domain of transfer functions, hence increasing their capability to visually discern between different materials and structures. Kindlmann et al. [KD98] use first and second derivative information in the transfer function design in order to semi-automatically isolate structures within the volumetric dataset correlating with a material boundary model. An interactive approach is presented in [KKH01], where the gradient magnitude is also computed together with the Hessian matrix. This information is then incorporated into the transfer function design process by a set of manipulation widgets that the user can employ to select and highlight features within the dataset. A relevant aspect of this method is the concept of dual-domain, which connects the spatial and the transfer function domains. However, even for this interactive system proposed by Kniss et al. [KKH01] a considerable expertise from the user is necessary to achieve a meaningful visualization in a reasonable amount of time. In a medical scenario this would imply the necessity of a computer graph-

ics expert assisting the physician during the transfer function definition, which is not practicable in most situations.

The concept of direct interaction with the rendered result instead of with a representation of the transfer function domain has been taken one step further by Tzeng et al. [TLM03], which served us as inspiration for our approach. In this case, the transfer function space is kept completely hidden to the user, who only interacts with the volume itself by painting on sample slices of the dataset. The classification itself is performed by one multilayer perceptron (MLP) neural network for each predefined material class. In our interface proposal, we borrow the idea of using a multi-dimensional transfer function, while limiting the user interaction to the spatial domain. We also combine the result of the user interaction with machine learning methods in order to produce the final transfer function in an automated way. However, our approach clearly differentiates by employing an augmented reality paradigm on which a real 3D interaction with the volume is guaranteed, in contrast with a 2D slice-based solution. Moreover, our automatic classification process defines a standard multi-dimensional transfer function for the whole volume and does not require an extra neural network for each sought after material.

The method presented in this paper uses interaction in augmented reality for an easy-to-use semi-automatic volume classification procedure. Augmented reality denotes techniques which combine images of the real environment with three-dimensional computer-generated graphics. An overview of augmented reality is given by Azuma [Azu97]. The system described in this text was implemented using our medical augmented reality framework ARGUS [FNFB04]. The application of augmented reality in medical diagnostics and treatment has been in the focus of active research for many years. An early system for supporting ultrasound-guided needle biopsies was presented by State et al. [SLH*96]. Figl et al. have described a head-mounted operating microscope which is capable of overlaying additional graphical information over the conventional microscopic view [FBH*01]. A high performance video see-through augmented reality system for medical applications was presented by Vogt et al. [VKS*03]. Schmalstieg et al. developed the *Studierstube* project, a collaborative augmented reality system [SFH*02].

Augmented reality user interfaces have recently been used as a tool for defining transfer functions in a manual way [RZBB04]. The proposed paradigm corresponds, however, to a traditional one-dimensional transfer function, where the user manually combines a set of predefined functions in a trial and error manner, in order to determine a correct opacity transfer function for a gray-scale representation.

In this paper, we present a novel volume classification approach that combines the abilities of multi-dimensional transfer functions with a user friendly AR-based user interface. This type of interaction has the advantage of providing

a better understanding of the actual three-dimensional structure of the dataset, enabling a deeper integration of the user into the transfer function specification process. The actual transfer function design is kept transparent to the user by employing machine learning classifiers.

3. Medical AR

Our new system for semi-automatic volume classification uses an augmented reality environment for displaying the volume datasets. Moreover, intuitive three-dimensional user interaction is provided by a specialized AR-based user interface. The system has been realized using our framework for medical augmented reality, ARGUS (Augmented Reality based on Image *GU*ided Surgery). Unlike many other experimental setups for medical augmented reality, ARGUS uses existing, commercially available medical equipment [FNFB04]. A VectorVision[®] intraoperative navigation device is the basis for our AR setup (see Fig. 1(a)). This device is equipped with a highly accurate infrared tracking system. We utilize the information delivered by these infrared cameras to track the video camera which delivers the augmented reality background images.

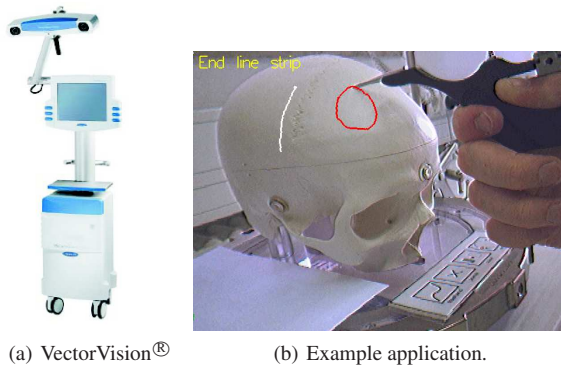


Figure 1: Medical AR. (a) IGS device produced by the BrainLAB company (Heimstetten, Germany). (b) 3D user interaction based on intraoperative navigation: operation plan drawings on a plastic skull.

The intraoperative navigation system is capable of tracking several objects simultaneously. A rigid configuration of two or three reflective spheres is attached to each object. Depending on the number of attached spheres, 5-DOF or 6-DOF position and orientation information can be obtained for the objects. Using these capabilities, we have designed and implemented a user interaction library based on our medical AR framework [FB05]. We use different pen-like and pointer-like tools as wireless interaction devices, which are tracked by the image guided surgery equipment. Our user interface system automatically detects different basic click gestures for the definition of points in 3D. Moreover, a full-fledged menu system with freely placeable menu items is provided. An example application of the interaction library is illustrated in Fig. 1(b).

4. Volume Classification

Our work proposes a new approach for generating a multi-dimensional transfer function that liberates the user from the internal complexity of the tedious design process. Following, the functional pipeline is described.

- Render volume in AR environment: The original dataset is rendered with a standard linear ramp transfer function for all color and opacity channels producing a gray-scale representation. The obtained result is displayed using 3D texture mapping.
- Inspect volume: The rendered volume can be directly examined in the AR environment using a clipping plane widget (see below).
- Select sample points: Combining the use of the clipping plane widget together with a pointer tool, sample points representing features of interest can be easily selected.
- Generate transfer function: The information corresponding to the selected points is processed and utilized to automatically generate an appropriate transfer function. This automatic process is achieved with the help of machine learning algorithms on the basis of a 2D histogram (voxel intensities and gradient magnitudes).
- Render classified volume: Once a satisfactory transfer function has been obtained, the volume is classified and a final rendering image is produced.

4.1. User Interaction




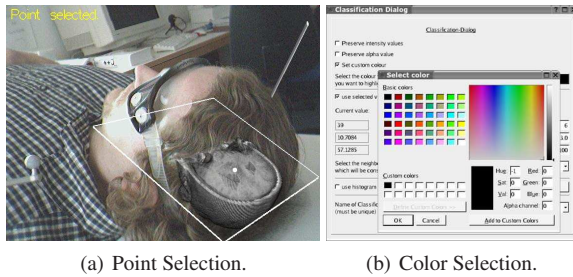
Figure 2: AR interaction tools: clipping plane (lower image) and pointer tool (upper image).

In our system, the user can directly manipulate and interact with the volume in an AR environment. The correct positioning of the scanned dataset (CT, MRI) is guaranteed by the registration process provided by the image-guided surgery (IGS) system [FNFB04]. Two simple tools are then employed by the user to inspect the volume and define regions of interest. Figure 2 shows both interaction tools together with their virtual representation in an AR environment: a clipping plane and a pointer tool.

The clipping plane tool (see lower part of Figure 2) is tracked with the help of three non-aligned reflecting spheres (6-DOF) in such a way that its position and orientation can

be retrieved from the IGS system. With this information a plane containing the three reflecting spheres can be easily modelled. An intersection test between this plane and the volume bounding box provides the information necessary to cull off all voxels *above* the defined plane. This way the user can browse through the volume rendered in the AR environment by just moving the clipping plane tool in any direction, providing a direct insight of the data.

The second interaction tool is a pointer. As can be seen in the upper part of Figure 2, the pointer tool is tracked using two spheres (5-DOF). The position of these two points determine a vector that can be used to represent the direction of a straight line. A simple plane-line intersection test between this straight line and the plane defined by the clipping plane tool can be used to indicate a position within the dataset. In order to complete the type of interaction needed by our system we need to be able to generate an activation *click* event informing that the current position corresponds to a region of interest and should be used for the transfer function design process. We solved this problem with our AR user interaction library (see Section 3). By holding the pointer still during a predefined time (~ 2 seconds), the user generates an event informing the system that the point selection process must be activated.



(a) Point Selection.

(b) Color Selection.

Figure 3: Sample point selection process.

Once a point has been selected (see Figure 3(a)) a dialog allows the user to select color and opacity values (see Figure 3(b)), as well as an identifying name for the class corresponding to the current material. This point selection process can be repeated either to select new points corresponding to already defined material classes, or to introduce new materials to be taken into account during the classification of the volume.

As mentioned above, the points selected by the user constitute the input to the automatic transfer function generation. In order to minimize the number of points that must be defined by the user and to make the point selection more robust against noise in the dataset and inaccuracies during the point selection, a larger set of sample points is generated out of the initial subset defined by the user. For each point selected with the interaction tools, a small surrounding region around the point is determined. The user has the opportunity to choose between two different specifications for this surrounding region: A 3D neighborhood and

a planar neighborhood. The 3D neighborhood is formed by the $3 \times 3 \times 3$ first order neighbors along the three orthogonal directions (X,Y,Z). The planar neighborhood is a two-dimensional 3×3 set of points on the intersection polygon determined by the clipping plane tool. In either case, the computed set of points is taken and the average values for both intensity and gradient magnitude are calculated. This way, even if noise is present in the dataset, the user can rely on the fact that the selected point does correspond to a representative value of the spatial region which has been pointed to.

Next, for each of these average points representing the region selected by the user, a small two-dimensional interval in the transfer function domain is taken around the given values, and the cartesian product of these two subsets is employed to generate the set of points that are passed to the actual classifier.

$$[I_i - \Delta I, |\nabla I_i| - \Delta |\nabla I|] \times [I_i + \Delta I, |\nabla I_i| + \Delta |\nabla I|] \quad (1)$$

Equation 1 illustrates this set of points, where I_i and $|\nabla I_i|$ denote, respectively, the voxel intensity and the gradient magnitude of the average point representing one region selected by the user.

After several tests, we have found that, for 8 bits/voxel datasets, an interval width (ΔI) of 6 HU (Hounsfield units) in intensity and 3 in terms of gradient magnitude ($\Delta |\nabla I|$) are enough to obtain a proper representation of the selected points. Hence, 72 sample points are generated for each user selection, ensuring a sufficient amount of input data for the automatic classification process. Due to this procedure, the spatially guided selection performed by the user is combined with a data centric generation of similar points in the transfer function domain.

4.2. Machine Learning Classifiers

The actual transfer function design process is performed automatically using machine learning methods. In this work, we have tested two different approaches in order to generate a 2D transfer function. The first is based on the utilization of an artificial neural network, more specifically a *Multi-Layer Perceptron* (MLP), while the second employs a *k-Nearest Neighbors* classifier.

4.2.1. Multi-Layer Perceptron Classifier

Artificial neural networks are composed by simple processing elements (artificial neural cells) organized in architectures characterized by a high degree of interconnection inspired by the parallel architecture of animal brains. They result specially attractive due to their ability to approximate functions based on sparse data through a training process and to apply this to solve new problems of similar nature. This training process adapts the weights modulating the value across each connection until the network implements a desired function.




Figure 4: Structure of a three-layer perceptron.

For our purposes, we have chosen a three-layer perceptron topology, using the supervised training method known as *Feed-Forward Back-Propagation* algorithm. Figure 4 shows an illustration of such a network. Our network is composed by one input layer of two cells, one for each value in the domain of the transfer function (I_i , $|\nabla I_i|$), one output layer with as many cells as material classes have been defined by the user, and one hidden layer. After several tests, we have determined that a hidden layer of 15 cells is able to properly discern among the user defined classes (typically 4 – 8 materials).

For the implementation of the neural network, a third party library (LTI-Lib from the University of Aachen (Germany)) was selected. The functioning principle is simple: The set of sample points created out of the user selected positions is utilized as a training set for the neural network. Once the network has been properly trained, each voxel intensity and gradient magnitude are fed to the network and an output value is produced at each cell of the output layer. As mentioned above, the output layer has one cell for each material class. The value generated at each of these cells indicates the probability that the input voxel belongs to the material associated to the output cell. This way, a probability threshold can be set so that each voxel producing a probability over the threshold be classified with the color and opacity of the corresponding material, while the remaining voxels with probabilities below the threshold can be either ignored (*opacity* = 0), or classified with the initial standard transfer function.

4.2.2. k-Nearest Neighbors Classifier

A k-nearest neighbors (kNN) analysis provides a different alternative for the automatic classification of new objects out of a number of known examples. In our approach we use a 1-nearest neighbor classifier ($k = 1$). The operation routine is very simple, though effective. The known examples constitute classes defined by the position of a *prototype*. The initial position of each prototype is given by the first sample point corresponding to a material class being passed to the input of the classifier. Each subsequent sample point is assigned to the class whose prototype is at the lowest Euclidean distance computed on the transfer function domain. Once the assignment has been completed, the prototype is recomputed as

the average position of all the points belonging to the class it represents. When the next sample point arrives, the same process is repeated until all the points generated as a result of the direct user interaction have been classified and the corresponding prototypes have been respectively repositioned to their final location. Then the kNN classifier can be utilized to perform the classification of the volume, voxel by voxel. Figure 5 illustrates a simple example of this process. The





Figure 5: Structure of a kNN Classifier ($k=1$). The red point illustrates an example of a voxel belonging to class B.

labels A, B, C, D, E show the prototype position corresponding to five different material classes selected by the user with the AR-based interaction tools. In this scenario, each voxel is classified as belonging to the material class whose prototype is the closest in terms of intensity and gradient magnitude. The red point in Figure 5 represents a voxel being classified. Since its nearest neighbor ($k = 1$) is prototype B, the voxel will be classified with the color and opacity associated to this material.

Even though the kNN classifier does not provide a direct measure of the probability for each voxel to belong to a specified material class, the distance to the closest prototype can be used for this purpose. This way it is possible to set a threshold again in order to decide whether all voxels should be classified (assigned to one of the defined classes), or those which are not close enough to any prototype should be ignored (*opacity* = 0) or classified with the initial gray-scale linear ramp transfer function.

5. Results



(a) Unclassified Volume.

(b) Insight View.

Figure 6: Rendering of unclassified MRI head.

A first prototype implementation of our proposed method

has been realized in a master project. Based on this prototypical implementation, we have performed tests with several scanned medical datasets (CT, MRI). Due to space restrictions, only one representative example is analyzed in detail here (see also video). Results for other datasets are graphically presented to demonstrate the effectivity of our method. The example dataset is an MRI scan of a human head with 8 bits/voxel of resolution and size $256 \times 256 \times 124$ voxels. Our test system is a PC with an Intel[®] Xeon[™] processor running at 2.66 GHz and a graphics card based on an NVidia[®] GeForce[™]FX 6800 chipset.


Figure 6 shows the starting point for our approach. The volume is first rendered in an AR environment using a standard linear ramp one-dimensional transfer function for all color and opacity (alpha) channels. The volume rendering is performed with 3D texture mapping using preclassification (classification previous to interpolation) to apply the transfer functions. As can be seen on the images (see Figures 6(a) and 6(b)), the initial rendering does not provide a sufficient insight of the structures present in the scan volume and a new classification must be performed in order to reveal the information it contains. Using our direct interaction tools (clipping plane and pointer tool), we define five different materials as described in Table 1.

Material	Color	Opacity	# Points
air	black	0.0	3
brain	blue	0.8	3
ventricle	red	0.5	2
skull	dark yellow	0.4	5
edges	white	0.6	6


Table 1: Material classes defined for the MRI head dataset.

For each material class, a color and opacity value has been defined. The selected sample points are reproduced in Figure 7(c) over the 2D histogram of the dataset. Each colored rectangle illustrates the size of the intervals in terms of intensity and gradient magnitude employed for the generation of the input to the automatic classifiers. Since these labelled regions define a color and opacity for those voxels represented below the respective areas on the histogram, a preliminary manual classification can be performed. The Figures 7(a) and 7(b) show the volume after classifying these voxels with the selected color and opacity values. Obviously, the result is still not satisfactory enough, but gives an idea of the type of information the automatic classifiers receive.

Using these 19 sample points, a training set for the MLP neural network is generated as described above. For this example, we have configured the training process to a maximum length of 400 epochs (iterations of the back-propagation algorithm). Under these conditions, the network was trained after approximately 2 minutes. A threshold of 70% was set as the minimum probability to consider a voxel as belonging to one of the defined material classes. Voxels with a probability inferior to 70% are rendered using




(a) Pre-Classified Volume. (b) Insight View.



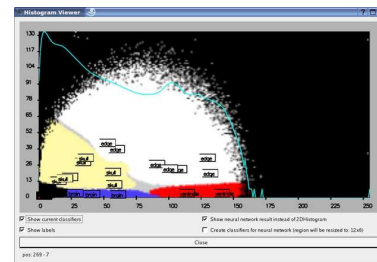
(c) Histogram with user defined sample points.

Figure 7: Result of manual classification with user defined sample points (MRI head).

the initial linear ramp transfer function. Figure 8 shows the obtained result after applying the transfer function generated by the MLP neural network. The transfer function itself is represented in Figure 8(c) over the 2D histogram of the dataset. Each point on this histogram corresponds to one



(a) Classified Volume. (b) Insight View.




(c) Histogram showing transfer function.

Figure 8: Result of automatic classification with MLP neural network (MRI head).


voxel in the volume, while the colors indicate the distribu-

tion of the classes produced by the automatic classification process. As can be seen in Figures 8(a) and 8(b), the inner structure of those materials selected by the user during the direct manipulation stage has been successfully highlighted and the dark cloud produced by the air surrounding the head has also been effectively removed.



(a) Classified Volume.

(b) Insight View.



(c) Histogram showing transfer function.

Figure 9: Result of automatic classification with kNN classifier (MRI head).

In order to be able to perform a fair comparison of both automatic classifiers, we use the same set of sample points (see Figure 7(c)) for the initialization of the kNN classifier. This setup consists in the creation and proper location of the class prototypes. Given the reduced amount of sample points (19) that must be processed and the limited number of material classes (5), this initialization step is accomplished almost instantaneously. In this case, no threshold was set, letting the classifier act over the whole volume. The obtained result is illustrated in Figure 9. Figures 9(a) and 9(b) clearly show the ability of the kNN classifier to emphasize those features explicitly selected by the user in the AR environment. Specifically, the ventricle system has been clearly highlighted and the external circumvolutions of the brain have been made visible too.

Two additional illustrations of results obtained with our semiautomatic volume classification approach are shown in Figure 10 for another MRI scan of a human head and a CT acquisition of a human thorax. The setup utilized in both cases is analogous to that described above for the first MRI head example dataset.

Comparing the results obtained with both automatic classification methods, the outcome produced is surprisingly similar for both the MLP artificial neural network and the

kNN classifier, specially considering their completely different internal structure. Both approaches have been proven to be effective removing undesired elements of the dataset (e.g. surrounding air) and also highlighting materials and features of interest (e.g. ventricular system) over uninteresting background regions. However it is also possible to recognize some characteristic differences between them. The artificial neural network, probably due to its more complex structure, produces a slightly more accurate classification of the data than the k-nearest neighbor classifier. This is particularly true for fine structures like the circumvolutions of the brain in the MRI head datasets or the bronchi in the CT thorax dataset (see Figure 10). On the other hand, the simplicity of the kNN classifier makes it also attractive given its shorter runtime and similar visual performance. In any case, further testing should be performed before a definitive evaluation of both classifiers can be proclaimed, since at the current state both alternatives have proven their adequateness.

6. Conclusions and Future Work

An appropriate transfer function is crucial for an effective visualization of a volume with direct volume rendering. Specially in medical routine, the aim of the visualization is highly dependent on the application. Thus, in order to achieve a successful volume classification in a reasonable period of time, part of the knowledge of an expert user must be introduced in the transfer function design process. This makes user interaction a first order priority for any useful volume classification method. Augmented reality can be employed to provide intuitive manipulation tools, allowing the user to directly interact in 3D with the volume to be visualized. This way, an intrinsic volumetric problem like volume visualization can be addressed in a three-dimensional environment. In this paper, we have presented a novel approach for semi-automatic transfer function design which is based on direct interaction between the user and the rendered volume in an AR environment. Based on a set of sample points defined by the user directly in the volume, an automatic volume classification process is carried out using machine learning techniques. Specifically, an artificial neural network (multi-layer perceptron) and a k-nearest neighbor classifier have been implemented and tested. Both alternatives have proven to be appropriate for the transfer function design process. As future work, we would like to perform a deeper comparative analysis of the presented and probably other machine learning solutions (e.g. RBF neural networks) in order to determine which techniques are better suited for which particular applications. We also plan to improve the volume renderer implementation by adopting a post-classification scheme, hence reducing interpolation artifacts and producing more appealing images. Even though this does not have a direct effect on the proposed methodology, we believe that a more accurate rendering quality can benefit the proper evaluation of the different classification methods.

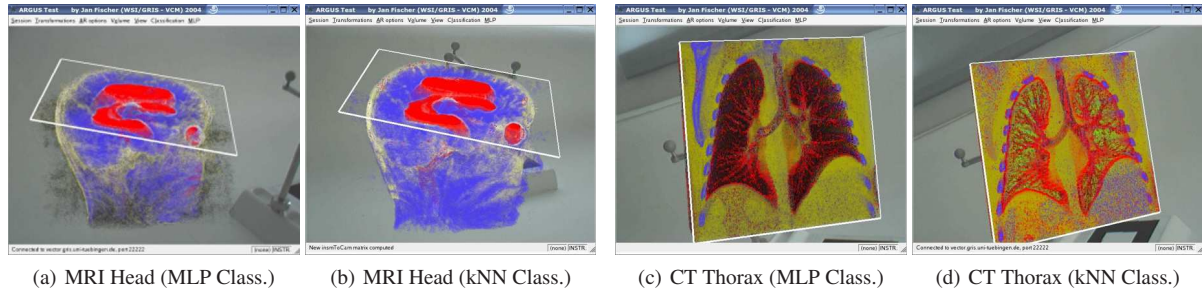


Figure 10: Results of automatic classification with MLP ((a), (c)) and kNN ((b), (d)) classifiers for a second MRI scan of a human head ((a), (b)) and a CT scan of a human thorax ((c), (d)).

References

[Azu97] AZUMA R.: A Survey of Augmented Reality. *Presence: Teleoperators and Virtual Environments* 6, 4 (1997), 355–385.

[FB05] FISCHER J., BARTZ D.: *Utilizing Image Guided Surgery for User Interaction in Medical Augmented Reality*. Technical Report WSI-2005-04, WSI/GRIS, University of Tübingen, March 2005.

[FBH*01] FIGL M., BIRKPELLNER W., HUMMEL J., HANEL R., HOMOLKA P., WATZINGER F., WANSCHITZ F., EWERS R., BERGMANN H.: Current Status of the Varioscope AR, a Head-Mounted Operating Microscope for Computer-Aided Surgery. In *Proc. of IEEE International Symposium on Augmented Reality (ISAR)* (October 2001), pp. 20–29.

[FNFB04] FISCHER J., NEFF M., FREUDENSTEIN D., BARTZ D.: Medical Augmented Reality based on Commercial Image Guided Surgery. In *Proc. of Eurographics Symposium on Virtual Environments (EGVE)* (June 2004), pp. 83–86.

[HHKP96] HE T., HONG L., KAUFMAN A., PFISTER H.: Generation of Transfer Functions with Stochastic Search Techniques. In *Proc. of IEEE Visualization* (1996), pp. 227–237.

[KD98] KINDLMANN G., DURKIN J. W.: Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering. In *Proc. of IEEE Symposium on Volume Visualization* (1998), pp. 79–86.

[KG01] KÖNIG A., GRÖLLER E.: Mastering Transfer Function Specification by Using VolumePro Technology. In *Proc. of Spring Conference on Computer Graphics (SCCG)* (April 2001), pp. 279–286.

[KKH01] KNISS J., KINDLMANN G., HANSEN C.: Interactive Volume Rendering Using Multi-Dimensional Transfer Functions and Direct Manipulation Widgets. In *Proc. of IEEE Visualization* (2001), pp. 255–262.

[Lev88] LEVOY M.: Display of Surfaces from Volume Data. *IEEE Computer Graphics and Applications* 8, 3 (1988), 29–37.

[MAB*97] MARKS J., ANDALMAN B., BEARDSLEY P. A., FREEMAN W., GIBSON S., HODGINS J., KANG T., MIRTICH B., PFISTER H., RUMMLER W., RYALL K., SEIMS J., SHIEBER S.: Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Animation. In *Proc. of ACM SIGGRAPH* (1997), pp. 389–400.

[PHK*99] PFISTER H., HARDENBERGH J., KNITTEL J., LAUER H., SEILER L.: The VolumePro Real-Time Ray-Casting System. In *Proc. of ACM SIGGRAPH* (1999), pp. 251–260.

[PLB*01] PFISTER H., LORENSEN B., BAJAJ C., KINDLMANN G., SCHROEDER W., AVILA L. S., MARTIN K., MACHIRAJU R., LEE J.: The Transfer Function Bake-Off. *IEEE Computer Graphics and Applications* 21, 3 (2001), 16–22.

[RZBB04] REITINGER B., ZACH C., BORNIK A., BICHEL R.: User-Centric Transfer Function Specification in Augmented Reality. In *Proc. of WSCG* (Plzen, Czech Republic, February 2004), vol. 12.

[SFH*02] SCHMALSTIEG D., FUHRMANN A., HESINA G., SZALAVÁRI Z., ENCARNÇÃO L. M., GERVAUTZ M., PURGATHOFER W.: The Studierstube Augmented Reality Project. *Presence: Teleoper. Virtual Environ.* 11, 1 (2002), 33–54.

[SLH*96] STATE A., LIVINGSTON M., HIROTA G., GARRETT W., WHITTON M., FUCHS H., PISANO E.: Technologies for Augmented-Reality Systems: Realizing Ultrasound-Guided Needle Biopsies. In *Proc. of ACM SIGGRAPH* (August 1996), pp. 439–446.

[TLM03] TZENG F.-Y., LUM E. B., MA K.-L.: A Novel Interface for Higher-Dimensional Classification of Volume Data. In *Proc. of IEEE Visualization* (2003).

[VKS*03] VOGT S., KHAMENE A., SAUER F., KEIL A., NIEMANN H.: A High Performance AR System for Medical Applications. In *Proc. of IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (October 2003), pp. 270–271.