Integrated Volume Visualization of Functional Image Data and Anatomical Surfaces Using Normal Fusion

Rik Stokking,* Karel J. Zuiderveld, and Max A. Viergever

Image Sciences Institute, University Medical Center Utrecht, Utrecht, The Netherlands

Abstract: A generic method, called normal fusion, for integrated three-dimensional (3D) visualization of functional data with surfaces extracted from anatomical image data is described. The first part of the normal fusion method derives quantitative values from functional input data by sampling the latter along a path determined by the (inward) normal of a surface extracted from anatomical data; the functional information is thereby projected onto the anatomical surface independently of the viewpoint. Fusion of the anatomical and functional information is then performed with a color-encoding scheme based on the HSV model. This model is preferred over the RGB model to allow easy, rapid, and intuitive retrospective manipulation of the color encoding of the functional information in the integrated display, and two possible strategies for this manipulation are explained. The results first show several clinical examples that are used to demonstrate the viability of the normal fusion method. These same examples are then used to evaluate the two HSV color manipulation strategies. Furthermore, five nuclear medicine physicians used several other clinical cases to evaluate the overall approach for manipulation of the color encoded functional contribution to an integrated 3D visualization. The integrated display using the normal fusion technique combined with the added functionality provided by the retrospective color manipulation was highly appreciated by the clinicians and can be considered an important asset in the investigation of data from multiple modalities. *Hum. Brain Mapping* 12:203–218, 2001. © 2001 Wiley-Liss, Inc.

Key words: integrated visualization; multimodality imaging; brain imaging; volume visualization; image fusion; HSV color model

INTRODUCTION

Three-dimensional (3D) imaging has become essential for a variety of clinical diagnostic and therapy planning procedures. Digital imaging modalities such as magnetic resonance imaging (MRI), computed tomography (CT), single photon emission tomography (SPECT), and positron emission tomography (PET) generate huge

amounts of volume data. Accordingly, there is a growing need for accurate volumetric rendering of the data in order to get a better understanding of anatomical structures and their functional distributions.

In addition, there is a growing interest to combine data from multiple imaging modes (e.g., MR T1 and T2 images) and from multiple imaging modalities (e.g., SPECT and MRI). Here, even more than for single modality images, the problem of mentally reconstructing a 3D picture comprising the information provided by the various sources occurs.

Integrated 3D display depicting, simultaneously, aspects of anatomy and/or function is called for [Viergever et al., 1997]. Examples of clinical applications

*Correspondence to: R. Stokking, McConnell Brain Imaging Center, Webster 2B, Montreal Neurological Institute, 3801, rue University, Montreal, Quebec, Canada H3A 2B4. E-mail: rik@bic.mni.mcgill.ca
Received for publication 20 December 1999; accepted 5 October 2000

that benefit from integrated 3D visualization are display of thickness of the skull from CT [Zuiderveld et al., 1995], display of electromagnetic dipole data with CT and MR images [Van den Elsen and Viergever, 1994; Van den Elsen et al., 1995], fusion of CT, MRI, and MR angiography in skull base surgery [Hawkes et al., 1995], radiotherapy planning integrating renderings from CT and/or MRI visualizations with dose distributions [Bendl et al., 1995; Schlegel et al., 1996; Robb and Hanson, 1996], display of PET and MR images for surgical planning in epilepsy [Levin et al., 1989; Hu et al., 1990], and combined display of SPECT with MR images [Stokking et al., 1997, 1999]. Further examples of the fusion of functional and anatomical data can be found in Gevins et al. [1990] and Evans et al. [1996].

Important prerequisites for multiparameter and multimodality visualization are the availability of registration methods that reliably match separately acquired images and segmentation methods that identify and classify interesting features in the data set (in this article we assume these features to be surfaces extracted from anatomical image data). In multimodality matching, the focus is shifting from techniques using frames, moulds, or markers, i.e., extrinsic matching, to methods employing information intrinsically present in the acquisitions, i.e., intrinsic matching using features or voxel properties. Advantages of intrinsic matching over extrinsic matching are the absence of distress to patients and the possibility to apply these techniques retrospectively. For an extensive overview, see Maintz and Viergever [1998].

In the area of image segmentation, the time-consuming manual techniques for delineating structures in 2D images are replaced with (semi-) automated 3D techniques that reduce labor, time, and subjectivity. For MRI applications the abundance in segmentation approaches and the reported difficulties show that problems are still formidable: Robust, reproducible, fast, and fully automated segmentation of MR image data appears far away [for an overview see Clarke et al., 1995]. This observation has raised interest in user-guided, semiautomatic, simple segmentation methods that are capable of segmenting a large volumetric data set within a few minutes [Höhne and Hanson, 1992; Robb and Hanson, 1996] and fully automated segmentation methods that are aimed at a certain structure (typically the brain) [Brummer et al., 1993; Collins et al., 1995; Stokking, 1998; Atkins and Mackiewicz, 1998; Lemieux et al., 1999; Stokking et al., 2000].

Assuming that the volume images have been matched and properly segmented, appropriate visualization techniques are required for presentation of the

(usually intricate) information. This article addresses this issue for the combined volume visualization of functional information with surfaces extracted from anatomical image data. Three problems have to be dealt with: i) obtaining the functional information that corresponds to a certain point on the extracted anatomical surface; ii) mapping the resulting functional values onto the surface; and iii) optimally presenting the data to the clinician for interpretation.

We previously applied the normal fusion technique to combine functional and anatomical information, i.e., SPECT and MRI [Stokking et al., 1994, 1997; see also von Stockhausen, 1998]. However, this was primarily intended for a clinical audience and therefore the description of the technique was limited to a brief and general overview. Here, the visualization strategy for normal fusion is described in detail and the technique is extended by using the hue-saturation value (HSV) color model instead of the red-green-blue (RGB) color model. The HSV color model allows easy, rapid, and intuitive retrospective manipulation of the color encoding of the functional information, which will help interpretation ([see also Dimitrov, 1998] for the use of the HSV model for EEG/MRI). Furthermore, two different strategies for HSV color manipulation in the rendering results are evaluated.

This article is organized as follows. We start by providing some background on common shading techniques that are used for creating 3D renderings of volumetric data. Contemplating on a standard rendering technique, we describe a method to obtain functional measurements for points located on surfaces extracted from anatomical image data. This method, called normal projection, traverses a volume (the functional data set) along a secondary ray determined by the (inward) normal associated with the extracted anatomical surface. A value is calculated from sample points along this trajectory and encoded onto the surface. Because the visual system is very sensitive to color variations, the obvious approach is to modulate the surface color obtained from shading the anatomical data by the functional value. We provide some background on color models to support our choice for the HSV color model. Furthermore, different strategies for the color manipulation in the integrated functional-anatomical rendering results are presented and evaluated for different functional modalities using several clinical examples. Thereupon the overall approach for manipulation of the color encoded functional contribution to an integrated 3D visualization is evaluated in a clinical study. A discussion of the merits of the proposed approach concludes the article.

Figure 1.

3D visualization of anatomical surfaces. **(A):** Principle of rendering of surfaces as commonly done in medical imaging. An MRI data set has been classified into brain and skin voxels; the zoomed detail focuses on the (highest) grey values of the skin voxels. Shading calculations require the vector from the point at issue to the light source (light direction L), the surface normal at the point (outward

surface normal N), and the vector from point to the observer (view direction E); these variables are then used to evaluate a light model (see text). **(B):** A surface visualization of parts of the skin and brain from a healthy volunteer. The surface color as well as the light color were chosen white.

METHODS AND MATERIALS

3D visualization of anatomical surfaces

Realistic (computer) images from 3D data sets can be obtained with a process called volume rendering. This process relies on shading techniques that model light absorption, reflection, and transmission along surfaces. Adequate rendering speeds can be achieved using simple techniques, because photorealism is usually not required.

We assume that the surfaces to be visualized have been identified during a preceding segmentation step. As only visualization of surfaces is addressed in this article, this implies that the data set, which provides the anatomical context (usually a CT or MRI data set), has been classified into surface and nonsurface voxels, or structure and nonstructure voxels. Given a point on the surface, calculation of the light reflected from it requires the direction and intensity of the light that hits the surface, the direction of the observer, the surface direction, and a light model that models the reflection properties of the surface (see Fig. 1).

Most rendering algorithms assume a single light source at an infinite distance, while shadowing is usually ignored. The light intensity I_l as well as its direction, given by the unit vector L (labeled in Fig. 1A), is therefore assumed to be constant across the entire volume to be visualized, which greatly simplifies the shading algorithms and thus reduces computational requirements.

Perspective projection should be used to produce visually realistic 3D images; with this projection, the ray density decreases as one moves away from the observer [Hagen, 1991]. For each point at the surface, the direction from point to observer (given by the unit vector E, labeled in Fig. 1A) has to be evaluated; this makes perspective projection computationally expensive. Consequently, orthographic projection, which assumes the vector E to be constant across the whole volume, is still the preferred method for the majority of volume rendering applications.

Most modern volume rendering techniques calculate the surface direction (unit vector N, labeled in Fig. 1A) from the original grey-value data, e.g., using the normalized grey level gradient [Höhne et al., 1990]. The gradient is calculated either from the grey levels of its six first order neighbors or from the grey-level data in a second order ($3 \times 3 \times 3$) neighborhood of the point of interest. Normalization of the resulting gradient value then yields the surface normal.

Calculation of the light reflected from the surface is straightforward given the vectors N, L, and E. Because photorealism is not required, a simple light reflection model is adequate for most visualization purposes. The most used light model is the Phong model [Phong, 1975], which separates the reflected light into three components: i) an ambient (k_a), ii) a diffuse (k_d), and iii) a specular (k_s) component. The Phong model was later modified by an approximation of the spec-

Figure 2.

Principle of normal projection. Quantitative information at a pre-defined anatomical surface is derived by performing calculations along a trajectory defined by the (inward) surface normal. On the left, (outward) surface normals for the brain are obtained using the

gradient of an MR image. In the corresponding SPECT image on the right, samples on the trajectory along the (inward) surface normals are used to derive quantitative information from the SPECT data.

ular component [Schlick, 1994], thereby significantly improving the rendering speed.

Given the monochromatic light source I_l , the reflection I from the surface can be estimated by Schlick's modified Phong light model:

$$I = I_a k_a + I_l \left(k_d (L \cdot N) + k_s \frac{t}{n - nt + t} \right)$$

where I_a is the intensity of ambient light, n $[1, \infty]$ a parameter that controls the size of the specular highlight, while t is obtained by $t = N \cdot H$ where H is the "half-angle vector" calculated by $H = (L + E)/|L + E|$ [Blinn, 1977].

Figure 1B contains a typical rendering of a surface that was obtained using a $3 \times 3 \times 3$ neighborhood gradient calculation, orthographic projection, and Schlick's modified Phong light model. This figure depicts anatomical information, i.e., the surface of the skin and brain from MRI data. 3D visualizations of the brain are increasingly used by clinicians because it allows easier and more rapid appreciation of the gyral and sulcal pattern [Levin et al., 1989; Kikinis et al., 1992]. When not only anatomical information but also functional information of the patient is acquired with, e.g., SPECT, PET, or functional MRI (fMRI), techniques for integrated visualization are called for to convey the information to the clinician. In this article we focus on the specific problem of combining func-

tional data with the surface of the brain extracted from anatomical image data. The first step toward integrated visualization of functional images and the brain surface is to project relevant functional information to points on this surface. This is accomplished by a technique called normal projection, which is discussed next.

Normal projection

We previously introduced a technique that uses the surface normal for an anatomically accurate and view-point independent mapping of quantitative values from SPECT onto the surface of the brain from MRI [Stokking et al., 1994, 1997]. Figure 2 shows the principle of this normal projection technique. Surface normals are obtained by applying gradient operators to each point at the surface that should be visualized; a $3 \times 3 \times 3$ gradient operator provides an accurate and relative noise-free estimation of the normal direction with our MR data sets. These normals are then used for evaluation of the light model as discussed in the previous section but can also be used to derive quantitative data from the functional data set(s). For the application depicted in Figure 2, the SPECT activity is evaluated on a trajectory along the inward surface normal where the depth and sampling rate are adjustable.

Normal projection defines a path along which quantitative values corresponding to each point at the surface can be obtained. The best method to quantify the functional information depends on the application. Feasible options are the maximum value along the path, the mean value, or a weighted average in which the weight factor depends on the distance to the surface.

Integrated display or “fusion” of the quantitative information and the surface extracted from anatomical image data is done by color encoding the calculated quantities onto the surface. The next section discusses the selection of appropriate color assignments.

Color models

The complex characteristics of human color perception make the selection of color assignments for graphical display purposes far from trivial. It is beyond the scope of this article to give an extensive overview of color perception; good introductions can be found in the chapters by Foley et al. [1990] and Gouras [1991], and in a series of papers [Murch, 1984a, 1984b, 1984c]. Instead, we present a brief overview of the two color models that we used in our work on integrated visualization.

The RGB color model is most widely used as it has a convenient mapping to hardware. Display hardware allows for independent control of the contribution of each of the RGB colors. However, the RGB model lacks intuitive appeal. Given a color, it proves hard to estimate its correct RGB values, which indicates that the RGB color description system does not match well with perceptual properties. A more intuitive interface for color selection was proposed by Smith [1978]; the color model is a relatively simple nonlinear transformation of the RGB cube [see Foley et al., 1990; Watt, 1993 for pseudocode]. Smith’s HSV color model maps better on the visual sensations caused by colored light [Murch, 1984a; Foley et al., 1990; Lutz et al., 1991]. Here, hue refers to the wavelength that enables us to distinguish one color from another; saturation refers to the purity of the color, while value refers to the perceived intensity. The HSV model uses a cylindrical coordinate system and is usually represented by an inverted cone, as illustrated in Figure 3A. Changes in hue rotate around the axis of the cone while the saturation is greatest at the outer edge of the cone. Finally, dark colors are those close to the apex of the cone, while light colors are located at the cone’s base. The HSV model can be interpreted in terms of an achromatic (grey) part, i.e., the value component, and a chromatic component, i.e., the hue and saturation

components (see also the General Discussion section on separation of achromatic and chromatic information in the human visual system).

Although often considered as a “perceptual” color model, the HSV model is not perceptually linear; for example, maximum intensity yellow has a higher perceived brightness than maximum intensity blue [Keller and Keller, 1992]. Color models as CIELUV and TekHVC [Taylor et al., 1989] overcome this problem; they represent perceptually uniform color spaces in which measured and perceived distances are approximately equal. However, use of these models requires measurement of the colorimetric performance of the used display device; unfortunately, changing lighting conditions as well as manipulation of monitor contrast/brightness makes this calibration cumbersome.

Color encoding for integrated visualization

Several authors have used color for integrated 2D visualization of medical image data from multiple sources. Their techniques can be roughly divided into four categories: i) alternate pixel display, ii) RGB integration, iii) HSV integration, and iv) color compositing. We note that the distinction between the last three categories can be difficult, as they overlap. Furthermore, literature is not always clear whether a “change of color” is a change of hue or also a change in saturation and value [see also Christ, 1975].

Alternate pixel display [Hawkes et al., 1990; Rehm et al., 1994] presents information from two input images in an alternating fashion by using the ‘even’ pixels of the first image and the ‘odd’ pixels of the second image. Although not primarily aimed at color integration, the technique has been applied for integration of a color image with a grey or other color image. The contributions of the two images remain separate and color may be independently changed. Both Hawkes et al. [1990] and Rehm et al. [1994] consider this display visually pleasing and report perceptual interactions between neighboring pixels [see also Livingstone and Hubel, 1988; Murch, 1984b], but they disagree on the ease of interpretation. Hawkes et al. [1990] find the display difficult to interpret owing to “color smearing,” Rehm et al. [1994] consider the display easy to interpret despite camouflaging effects. Their difference of findings may be contributed to the perceptually simpler hot-metal color scale used by Rehm et al. [1994] compared to the saturated rainbow scale used by Hawkes et al. [1990].

With RGB integration, images to be combined are each individually assigned to the primary colors red,

Figure 3.

green, and blue. Three sources of information can be integrated, e.g., multiple PET tracer images [Freiherr, 1988], multiparameter MR images [Kamman et al., 1989; Alfano et al., 1995], or two SPECT tracer images and CT [Ricard et al., 1993]. Integration of two images, e.g., PET and MRI [Wahl et al., 1993], leaves one of the RGB components that can be used to make images more appealing, e.g., to assess registration accuracy for CT to CT registration [Van Herk and Kooy, 1994]. For integration of comparable images, e.g., multiparameter MRI or CT with CT, RGB integration is a natural choice, but for integration of functional with anatomical images RGB encoding appears nonintuitive [see also Brown et al., 1991; Stokking, 1998].

With HSV integration, usually two source images separately encode the hue and value parameters, and the saturation parameter is kept fixed, e.g., for multiparameter MR display [Weiss et al., 1987] and for multimodality colorwash display [Pelizzari et al., 1989; Levin et al., 1989]. However, one of the sources can also be assigned to the saturation instead of the hue component, for instance, to present an overlay of different saturation levels of a hue onto grey values.

Color compositing originates from the classic work on volume rendering by Porter [Porter and Duff, 1984] where a transparency value is assigned to each pixel (the so-called α -value). This value determines the contribution of the pixel content to the final image. The Montreal group has adopted this technique for integrated display of PET and MR images and denotes it "opacity weighted display" [Evans et al., 1991, 1996]. Brown et al. [1991] describe the "multichannel color composite" method to integrate MR parameter images. With this technique, each of the input images contributes to the red, green, and blue components through multiplication with independent constants to obtain more appealing images for the integrated MRI data.

Overall, while integrated 2D visualization using color encoding presents some unexpected perceptual effects and limitations with 8-bit displays [Hawkes et al., 1990; Brown et al., 1991; Rehm et al., 1994], color

has proven a powerful cue for simultaneous display of functional and anatomical information.

A variety of techniques have been reported for integrated 3D visualization of functional and anatomical images. The two image types can be independently rendered and the resulting (2D) images can be combined with any of the previously mentioned techniques for integrated 2D visualization, e.g., painting a color onto a grey surface [Levin et al., 1989; Hu et al., 1990], or using color compositing [Evans et al., 1996]. Integration of information can also be performed by texture mapping functional information onto a surface, e.g., for the brain [Payne and Toga, 1990], or by first mapping functional information onto the anatomical volume followed by rendering of the combined volume for the heart [Heffernan and Robb, 1984] or for the brain [Valentino et al., 1991]. In the technique presented here—normal fusion—local functional information is color encoded onto a surface extracted from anatomical image data [Stokking et al., 1994, 1997].

Normal fusion with the HSV color model

In previous work we used the normal fusion technique to project SPECT information onto the surface of the brain volume rendered from MRI. A clinical evaluation was conducted, where several of the observers reported the desire to manipulate the color-encoding scale of the functional information in the rendering results to improve understanding of the data [Stokking et al., 1997]. However, as already mentioned, the RGB model does not offer an intuitive and simple approach for color manipulation. The HSV model appeared more appropriate to color encode quantitative data onto surfaces and separate the anatomical and functional information into an achromatic and a chromatic component. This separation is highly intuitive in that it appears to exploit the different pathways in the visual system quite efficiently (see Discussion for more details). It also allows easy, rapid, and retrospective manipulation of the color encoding in the render-

Figure 3.

(A): HSV color model, after [Foley et al., 1990]. (B): Lookup table used in this article. (A) shows on the left a schematic representation of the cylindrical coordinate system used by the HSV color model, and on the right the colors at the base of the inverted color cone. The superimposed contour shows the trajectory and control points of the lookup table presented in (B). (C): The strategies for HSV color manipulation, where separate storage is a subset of the recalculation scheme. The quantitative and anatomical information is either separately stored or calculated from an integrated image and the original lookup table (defined by several control points). A novel integrated visualization can be obtained by color encoding the quantitative information onto the anatomical information using a new lookup table. Manipulation of the control points of the lookup table readily changes the color encoding of the quantitative information in the integrated visualization. The lookup table presented in this figure was applied in the clinical evaluation to signal both cold and hot spots.

ing results without the need for a new volume rendering.

Given a white surface and a monochromatic light source, the intensity of the light reflected from the surface can be readily calculated using Schlick's modified Phong light model. This yields the value component of the HSV model. Color display is often used for clinical evaluation of functional and/or quantitative information, which makes the choice to use the hue and saturation components for the quantitative value obtained from the functional data by the normal projection technique described earlier, quite straightforward.

Several authors indicate the usefulness of color encoding for interpretation of functional images, e.g., for improving detection [Arnstein et al., 1990; Stapleton et al., 1994], but, to our knowledge, little work has been done to standardize or validate the use of different color scales for interpretation of functional images. Also, most observers have their own preferences. These perceptual "problems" are usually tackled by offering a range of lookup tables to the observer with some means to manipulate the chosen lookup table [see also Foley et al., 1990]. We have adopted this strategy in our research by allowing manipulation of the lookup table.

The lookup table and manipulation

Human perception is such an intricate process that there appears to be no ubiquitous "best" strategy for hue/saturation assignment. In general, the important functional information is probably best presented with red for hot spots and blue for cold spots. This is highly intuitive and at short (blue) and long (red) wavelengths there are more distinguishable steps of saturation for each hue than for the midspectral region (green) [Foley et al., 1990; Gouras, 1991]. On the negative side, the human visual system seems to be biased against blue as we have relatively few cones sensitive to blue [Foley et al., 1990] and the lens absorbs almost twice as much light in the blue region as in the yellow and red region (which also increases as we get older) [Murch, 1984c]. Other considerations are that red appears closer to the observer and blue seems to be more distant [Foley et al., 1990; Murch, 1984c], and red and blue must be of a much greater intensity than a green or yellow to be perceived [Murch, 1984c]. Whenever hues are required but should not interfere with the anatomical information, e.g., for the transition from insignificant to highly significant functional data, we suggest using the orange-yellow hue as humans are maximally sensitive to luminance changes for these

hues [Levkowitz and Herman, 1992]. The description of the applied lookup table is preceded by an explanation of the strategies for color manipulation as these impose some constraints on the lookup table.

We have investigated two strategies for manipulation of the color encoding in the rendering results, namely i) separate storage, and ii) recalculation (see Fig. 3C). The first strategy requires separate storage of the quantitative and anatomical information, which are combined into an output image. With the second strategy, the quantitative information at a surface voxel, is reconstructed from the hue and saturation components stored in the output file of the volumetric rendering and the lookup table that was used for the color encoding of the quantitative information. The latter technique requires a one-to-one mapping of the calculated quantitative information to the hue and saturation components.

In order to evaluate both strategies we used three clinical cases of different modality combinations and a simple HSV color scale aimed at signaling hot spots (see Fig. 3B). The range of the color scale is divided into four parts. In each part of the range, only the hue or the saturation is gradually increased to obtain the required one-to-one relationship for the recalculation strategy. The quantitative information of interest is represented by a hue gradually increasing from 60° (yellow) to 360° (red) (area III of the graph in Fig. 3B). The area of quantitative information of minor interest (area I) is represented by low saturations to obtain a noncolored surface. In area II of the graph there is a gradual increase in saturation from area I to area III of the graph, which means a gradual transition from white to yellow. Area IV was constructed to allow manipulation of information higher than point C. This lookup table was aimed at the presentation of hot spots using color encoding "increasing" from white over yellow, green, blue to red.

The lookup table can be characterized by a few control points with their respective hue and saturation entries. This makes manipulation of the color encoding for the functional information in the rendering results quite straightforward; we therefore used manipulation of the location of the control points, i.e., the numbers corresponding to locations A, B, C, and Max in Figure 3B for manipulating the color assignment schemes. Changing the lookup table for the functional information does not require new volume renderings, only the recalculation of the color contributions in the rendering results. This dramatically reduces the number of required computations allowing very rapid manipulation of the color encoding.

Processing and visualization

Registration of the data sets was done either by using the surface-matching facilities in the ANALYZE™ [Robb and Hanson, 1996] software package or the mutual information technique [Maes et al., 1997]. We chose to match to the anatomical data to avoid degradation of the visualizations. For the segmentation of the brain from the anatomical data we used a supervised segmentation procedure incorporated into ANALYZE™. The method was originally proposed by Höhne [Höhne and Hanson, 1992] and is based on region growing and morphological operations (erosion and geodesic dilation). For visualization, we used the software package VROOM [Zuiderveld, 1995], developed at our department; it is essentially a collection of C++ classes, which was specifically designed for exploration of novel strategies for integrated volumetric visualization.

The software for color manipulation was written in C++ and the graphical user interface was developed using the Tcl/Tk toolkit [Oosterhout, 1994], which greatly simplifies the building of user interfaces. OpenGL [Neider et al., 1993] was used for the display of the colored images to allow for real-time interaction and manipulation needed for clinical usage.

RESULTS

Clinical cases

We selected three cases of combined functional and anatomical brain images to evaluate our approach: i) PET/MRI and fMRI/MRI of a volunteer for monitoring of brain activity with a finger opposition stimulus, ii) PET/MRI of an epileptic patient, iii) SPECT/MRI of a patient with the Gilles de la Tourette syndrome. The first case is to illustrate the normal fusion technique, and the latter two cases are to illustrate more specifically the benefits of the HSV color-encoding scheme. For all cases, the surface of the brain was indicated by the segmented data and visualized from the corresponding original MRI-T1 data using a $3 \times 3 \times 3$ neighborhood gradient calculation, orthographic projection, and Schlick's modified Phong light model. This resulted in the normal used for shading and normal projection. Thereupon the functional data set was sampled each mm along the inward surface normal until a depth of 10 mm was reached (see Fig. 2). The maximum value of these samples was calculated and then color encoded onto the surface using the lookup table of Figure 3B. Applied in this way, the normal fusion image depicts the maximum functional

activity below the visualized brain surface until a depth of 10 mm. The results are shown in Figures 4 and 5.

PET/MRI and fMRI/MRI monitoring of finger opposition

Functional MRI (fMRI) has emerged over the past few years as a promising technique to image brain function. In an experiment to improve understanding of fMRI, $H_2^{15}O$ PET was also used to measure regional cerebral blood flow for cross-validation [Ramsey et al., 1996]. Both modalities were used to monitor the brain of a subject with or without stimulation of the primary sensory motor (PSM) cortex area.

The stimulation was evoked by a simple finger opposition task with the right hand by a right-handed subject. The task entailed repeatedly and sequentially touching the thumb once with each one of the digits. A total of eight PET and eight MRI data sets (four stimulus and four nonstimulus sets) were acquired and processed similarly to properly compare both functional modalities. Statistical analysis utilized standardized normal variates based on repeated measurements within a single subject, referred to as a z_t -map [for details we refer to Van Gelderen et al., 1995]. Both an fMRI and a PET z_t -map were calculated to denote the difference in respectively fMRI and PET activity between activated and nonactivated data. Furthermore, the PET data was resampled to the fMRI resolution, and the lookup tables of the integrated visualizations of both PET and fMRI had to be identical because both z_t -maps were statistically processed in order to have identical levels of activation.

MRI scans were obtained with a clinical 1.5 Tesla scanner (SIGNA, General Electric). The fMRI data was acquired with a 3D PRESTO sequence [Van Gelderen et al., 1995] (TE = 35 ms, TR = 24 ms, flip angle 11 degrees, slab thickness 65 mm with 90-mm field of view (FOV), data matrix $64 \times 50 \times 24$, 6-sec scan time, pixels of $3.75 \times 3.75 \times 3.75$ mm). For registration purposes an IR sequence (TI/TR 800/3,000 ms, slice thickness 2.75 mm, 1 mm gap, 24 slices spanning 90 mm, FOV 240 mm, 256×128 , 7-min scan time) was acquired, which matched the fMRI data in both location and orientation. Anatomical images of the whole brain volume were acquired with a spoiled GRASS sequence (TE/TR 5.1/20 ms, 124 contiguous slices of 1.2 mm thickness, FOV 300 mm, 256×256) and used for localization and for registration purposes of all functional data. The $H_2^{15}O$ PET images were obtained with a Scanditronix brain tomograph (15 contiguous

Figure 4.

Surface color encoding of PET subtraction activity (top row) and fMRI data (bottom row) for a finger opposition task. The brain cortex was extracted from MRI data. The presented 3D normal fusion renderings are stereo images (cross-fusion) of the left

hemisphere. The maximum functional value over a depth range of 0–10 mm was used to color encode the corresponding surface voxel. The lookup table used for all renderings is shown on the right.

slices, 6–6.5 mm in-plane and axial resolution after reconstruction).

The results for PET/MRI and fMRI/MRI are shown in Figure 4. The hot spots visible in both the PET/MRI and fMRI/MRI visualizations bear a high resemblance in size and activity. Furthermore, the hot spots are both located in the section of the PSM area corresponding with the finger opposition stimulus.

On the 2D image slices, the hot spots could be easily recognized for this case, but the anatomical localization requires mental integration of the functional with the anatomical data. Also, the pattern of gyri and sulci is hard to follow when using 2D slices only. The normal fusion visualizations assist in establishing the relationship of the functional information to the anatomy.

PET/MRI brain images of epilepsy

This case reports on a patient diagnosed with an epileptic focus in the right hippocampus. FluoroDe-

oxy Glucose (FDG) PET was used as a metabolic tracer to acquire functional information. The normal fusion procedure was applied to display the metabolic effects of the epileptic focus on the cortex.

The FDG-PET data was acquired in the interictal state with a 951 CTI/Siemens tomograph. A total of 31 contiguous transaxial planes parallel to the long axis of the temporal lobe were acquired, 5 mm in thickness, with a 3.4 mm slice separation (center to center) simultaneously covering 11 cm of axial FOV (top of brain was not included). A 3D T1-weighted gradient-echo MR image was acquired with 140 1.2 mm contiguous axial slices, TR = 30 ms, TE = 13 ms, 256 × 256 matrix, and 230 mm FOV of the head with a whole-body Philips Gyroscan 0.5 Tesla.

The results for PET-FDG (for examples, see Fig. 5A) were suggestive of i) an atrophic right temporal lobe (especially noticeable in stereo images and a movie sequence of this case) and ii) decreased PET activity in the right frontotemporo-parietal area. From the example images it is not clear whether the right temporal

Figure 5.

3D normal fusion renderings of the right and left hemisphere using the maximum value of a functional data set over a depth range of 0–10 mm to color encode the corresponding surface voxel extracted from MRI data (in each frame the lookup table used for the renderings is shown on the right). **(A)**: Surface color encoding of PET-FDG activity for an epileptic patient. **(B)** is the result of color

manipulation of **(A)** via the separate storage scheme, **(C)** is the result with the recalculation scheme. **(D)**: Surface color encoding of SPECT-HMPAO for a patient with the Gilles de la Tourette syndrome. **(E)** is the result via the separate storage scheme, **(F)** with the recalculation scheme.

lobe is hypometabolic compared to the left temporal lobe.

SPECT/MRI brain images of the Gilles de la Tourette syndrome

The image data presented in this subsection concern a 7-year-old right-handed patient diagnosed with the Gilles de la Tourette syndrome.

Information on brain anatomy was acquired from a 3D T1-weighted gradient-echo MR image (127 1.3 mm contiguous axial slices with TR = 30 ms, TE = 13 ms, 256×256 matrix, and 230 mm FOV of the head with a whole-body Philips Gyroscan 0.5 Tesla, the top of the brain was not included in the FOV). Information on cerebral blood perfusion was obtained from a SPECT HMPAO scan acquired with a Picker PRISM® three-detector gamma camera and reconstructed to 44 slices with a 64×64 matrix, a slice thickness of approximately 7.1 mm, and a plane resolution of 7.5 mm FWHM.

The integrated display shown in Figure 5D is the result of the normal fusion procedure applied to the SPECT/MRI data sets. Several differences can be noted when comparing the left and right hemisphere: i) a hot spot in the right lateral fronto-orbital region, ii)

increased activity in the left dorsal parietal lobe, and iii) increased activity in the left dorsal cerebellum, with a normal right cerebellum. It is not clear from this image what the level of activity is in the right lateral fronto-orbital hot spot compared to the rest of the visualized cortical activity.

Retrospective color manipulation

Strategies for retrospective color manipulation

The two cases depicted in Figure 5 illustrate the effects of interactive color manipulation. Initial examination of the PET/MRI normal fusion results suggested an atrophic right temporal lobe. This was most obvious as an enlarged superior temporal sulcus when viewing stereo images and/or movie sequences of visualizations of the right temporal lobe for this case. Whether the metabolic activity of this temporal lobe was abnormal could not be discerned from this image. Manipulation of the color encoding using the HSV scheme (see Fig. 5B, C) allowed rapid appreciation of the functional information; indeed the right temporal lobe had a lower metabolic activity when compared to other cortical regions. With the SPECT/MRI Tourette case, it proved of interest to refer the activity of the

right lateral fronto-orbital hot spot to activity in other cortical regions. Manipulation of the color encoding (see Fig. 5E and F) rapidly showed that the fronto-orbital hot spot was the most prominent of the visualized cortical regions.

The two strategies that were tested have both advantages and disadvantages. The first strategy, separate storage (see Fig. 5B and E) requires only a simple tool, no limitations need to be set for the lookup table, and the presentation is precise; but unfortunately, dedicated software is required for presentation of the image(s). The second strategy, recalculation (see Fig. 5C and F), introduces artifacts when using interpolating rendering algorithms. These artifacts are best noted in the border pixels of the object in Figure 5C. Changing the background color during rendering to a color not present in the HSV lookup table, e.g., black, solves the artifacts at the border pixels (see Fig. 5F). Furthermore, recalculation requires a one-to-one mapping of the calculated quantitative information to the hue and saturation components. It also requires a more intricate tool than with separate storage because the functional data has to be calculated from the H and S color components and the lookup table. On the other hand, recalculation has the major advantage that all the required information can be stored and presented in any image format that supports (lossless) 24-bit color images such as TIFF, PNM, and PNG.

Clinical evaluation of the color manipulation

Following the evaluation of the two strategies, we applied the separate storage strategy in a preliminary clinical evaluation using five SPECT/MRI cases (randomly selected from a total of 30) and five nuclear medicine physicians. The observers were already performing another validation study using the 30 SPECT/MRI cases [Stokking et al., 1999]. We extended the latter study to assess the opinion of clinicians on the color manipulation strategy with the normal fusion technique. An application was written to present six (previously) rendered orthogonal views in one window. This same window also contained an image of the corresponding lookup table used for the integrated visualizations and superimposed lines depicting the location of the control points in the lookup table. The position of these lines could be manipulated with the mouse to rapidly change the lookup table and consequently the color encoding of the functional information in the rendering results. For the color manipulation we decided to use separate storage of the anatomical and functional information instead of recalculation because we wanted to avoid the restric-

tions imposed by the required one-to-one mapping with the recalculation strategy. As both cold and hot spots had to be signaled, the lookup table was modified such that it was aimed at signaling cold spots with the color blue and hot spots with red. Consequently, the lookup table was more intricate than the one initially used. This required more control points (six) and thus more interaction from the observers. A simple solution involved coupling an additional function to a mouse button (i.e., changes in the upper or lower control point also affected the intermediate control points). For the background we could not use a neutral grey [this is suggested by Foley et al., 1990] or red/blue, as this interfered with the rendering of the brain surface; a low saturated green was therefore used instead.

It was our intention to perform a full clinical evaluation following this preliminary evaluation. However, observations were so obviously in favor of the color manipulation technique, as opposed to presenting the 3D visualizations with a fixed lookup table, that we decided to skip the second (and more thorough) validation altogether. The nuclear medicine physicians immediately considered color manipulation an asset because interpretation of the functional-anatomical images, and functional images in general, tends to rely on several steps with each an optimal color encoding. This calls for multiple color-encoded images or color manipulation. For instance, correlation of functional information to a reference area like the cerebellum requires a different color encoding than localization of a hot spot or detection of patterns of functional activity. Also, high activity present in extracranial tissue, e.g., the salivary glands, or due to markers may seriously interfere with the determination of a proper lookup table.

The experiments with these 5 SPECT/MRI cases revealed that the proposed lookup table and the user interface for color manipulation were highly intuitive and (thus) required little or no training.

DISCUSSION

Both the proposed normal projection technique and the color-encoding scheme are simple to implement. The use of the local grey level gradient is customary for shaded volume rendering, which makes implementation of the normal projection technique straightforward. With our data sets, the standard $3 \times 3 \times 3$ grey-level gradient was sufficiently accurate and noise-free to provide excellent results; for very noisy data sets, usage of Gaussian-filtered gradients might be appropriate. The HSV color model offered an intu-

itive and simple approach for retrospective color manipulation of the functional data in the rendering results. The color-mapping tables can be characterized by just a few control points since piece-wise linear hue and saturation assignment schemes give good results and a graphical user interface for the color manipulation is straightforward both in VROOM and as a stand-alone program. Also, the histogram of the functional information in the brain can be calculated for an initial guess of the control points of the graph. For the examples presented in this article we used percentages of the area under the graph, i.e., 80%, 90%, 95%, and 100% as initial guess for control points A, B, C, and Max. Overall, the approach does not require a new rendering each time the color encoding is changed. This not only allows rapid manipulation, as the image manipulation is basically 2D processing, but also requires little storage especially when compared to the original data sets.

It can be argued that integrated 3D visualizations similar to the normal fusion visualizations can be obtained using approaches applied in packages such as ANALYZE/AVW, medX, and SPM. However, these packages typically produce integrated 3D visualizations where the functional (or statistical) information of interest completely replaces the corresponding depth information from the rendering of the anatomical structure. For small areas of functional information, the loss of depth information will hardly be noticed, but for bigger, more widespread, and/or multiple areas of functional (or statistical) information of interest, this loss severely degrades the appreciation of the anatomical framework provided by the rendering [e.g., see Price et al., 1999; Kajimura et al., 1999]. Furthermore, the aforementioned packages typically employ the viewing direction for the integration of the functional data, which results in visualizations where the size, color encoding, and location of the functional information of interest depends on the viewing direction. With normal fusion this is not the case because the local surface direction is used. For a more extensive discussion and example images we refer to Stokking et al. [1997].

When we compared the 3D visualization results obtained with the normal direction and techniques using the viewing direction [see also Levin et al., 1989; Stokking et al., 1997], we observed that color and shape both contribute to the delineation of structures like gyri and sulci [see also Livingstone and Hubel, 1988]. However, color can also seriously camouflage the anatomical information of the brain surface as was most noticeable with the use of the viewing direction where colors are painted over sulci (see also Christ,

1975, on subject's accuracy in identifying achromatic target features when adding colors to the display). Integration of information along the inward normal direction presented images where color and shape strengthened one another. The additional manipulation of the color encoding effectively counteracts any remaining camouflaging effects.

Our approach for color manipulation was highly appreciated by the clinicians because, in practice, observers tend to have different preferences with respect to color encoding and manipulation. This is caused by personal characteristics such as experience, training, and perceptual (dis)abilities but also by environmental factors such as the monitor and lighting conditions. The variation between humans in their color discrimination capability is considerable [Murch, 1984c]. Color display is one of those fields where one observer's dream can be another observer's nightmare [cf. Murch, 1984b]. A fair fraction (9%) of the male population is color-blind [Gouras, 1991]. However, only a fraction of these are monochromatic or truly color-blind [Murch, 1984c]. The dichromats confuse only a small fraction of possible color pairs and can differentiate most color pairs at all relative brightnesses [Livingstone and Hubel, 1988]. This implies that careful choice of the lookup table and subsequent manipulation may well prove to circumvent the impoverished color vision. This assumption was substantiated by work done with a dichromatic nuclear medicine physician who did not report any problems in performing the required tasks.

Color has been reported to give excellent results in a variety of tasks; it is a powerful tool, but considerable caveats must be considered. We agree wholeheartedly with the statements: "Color should be used conservatively" and "In general, it is good to minimize the number of different colors being used (except for shading of realistic images)" [Foley et al., 1990]. One has to be very careful to paint a rather complex surface like the brain, as color may have undesired perceptual effects and may cause serious problems in the interpretation of data. For example, perception of depth and size is influenced by the color [Gershon, 1990], and the choice of background color influences the size of objects [Gershon, 1994]. The perceived color of an area is affected by the color of the surrounding area, although this effect is minimized when the surrounding areas are some shade of grey or are relatively unsaturated colors [Foley et al., 1990]. Use of large areas of saturated colors is undesirable because an afterimage of the large area will appear, which is disconcerting and causes eye strain [Foley et al., 1990]. It is best to use black, white, and grey for fine detail,

and reserve chromatic color as a means of attracting attention [Murch, 1984b].

The use of chromatic information for attention and achromatic information for the surface of the brain apparently exploits the different pathways of the visual system quite efficiently. The human visual system appears to process grey and color independently and they only combine at a high level in the perceptual hierarchy. Entirely separate channels are thought to handle distinct parts of the visual information where three parallel pathways process information for motion and depth, form, and color separately. The parvocellular interblob system is thought to specialize in high-resolution form perception, the parvocellular blob system specializes in color, and the magnocellular system is specialized for motion and spatial relationships. The automatic interpretation of a 2D image into 3D information seems to be performed only in the achromatic magno system, not in the parvo system. For an extensive overview we refer to Livingstone and Hubel [1988], Gouras [1991], and Kandel [1991]. One important aspect is not used in our technique; that is, the magnocellular system is also specialized in motion and this may provide a natural extension to use as a next item for the improvement of interpretation of the integrated medical images.

We have demonstrated that our approach is an important asset in the investigation of data from multiple modalities. The clinician can be supplied with several renderings of the multimodal volume data from well-chosen viewpoints over different depths [see also Stokking et al., 1997] with dedicated software for manipulating the color encoding. This offers the clinician a powerful tool not only for investigation of the functional data of the peripheral cortex in relation to the anatomy but also for discussing the findings with others without the need for the original data and the rendering software.

With highly convoluted surfaces and low-resolution functional data it is conceivable that the normal fusion technique color encodes functional data onto a neighboring gyrus instead of the correct gyrus. This is why we always advise to verify the reliability of the normal fusion results using the (original and registered) 2D images and not to use an integration depth of more than 15 mm.

The examples we have used in this article all focus on the surface of the brain extracted from anatomical image data. This surface is only part of the anatomical brain surface or cortical surface since about 60% is not visible from the outside. However, this limitation is caused by the choice of our examples and consequently the applied segmentation. We chose to seg-

ment, and thus visualize, only the outside or “visible” part of the anatomical surface, but nothing prevents us from segmenting other parts of the anatomy from the image data. For example, separating the brain into left and right hemisphere allows the visualization of the mesial aspects of both hemispheres. The same applies to the removal of the cerebellum or for the segmentation of a specific gyrus through its corresponding sulcus/sulci.

In this article we only presented examples for the brain, but the normal fusion technique extends to applications with other surfaces extracted from image data, e.g., heart, liver, and bone. An example can be found in Zuiderveld et al. [1995], where the technique is applied to calculate and visualize the thickness of the skull. The HSV color-encoding approach also extends to other 3D multimodality visualization techniques such as the multimodal cutplane [Payne and Toga, 1990; Stokking et al., 1994] or for display of information from time series such as EEG [Dimitrov, 1998; Gevins et al., 1999], MEG, or perfusion CT or MRI. A final remark is that we have only used normal fusion and HSV manipulation with volume rendering. However, this technique is also suitable for graphics pipelines that use surface rendering. In this case, a polygonal mesh would model the anatomical surface whereas the corresponding functional information can be stored into a texture that is then mapped onto the polygon mesh.

CONCLUSIONS

We have applied the normal fusion technique with the HSV color model for integrated 3D visualization of anatomical surfaces and functional data. The normal fusion technique is independent of the viewing direction and accurate in anatomical localization because it follows the curvature of the surface (here: brain) to calculate the regional quantitative information (here: activity of cortical cells). Fusion of the calculated activity with the information from the rendering is based on the HSV color model; hue and saturation are used for the functional information, value for the rendering information. This allows easy, rapid, and intuitive retrospective manipulation of the color encoding of the integrated visualization. Experimental evidence is presented that the integrated display enhances appreciation of functional data of the peripheral cortex within an anatomical frame of reference from MRI.

ACKNOWLEDGMENTS

We are indebted to our colleagues W.F.C. Baaré, Dr. R. Debets, Dr. J.K. Buitelaar, Dr. B. Sadzot, Dr. N.F.

Ramsey, Dr. H.E. Hulshoff Pol, J.H. de Groot, A.W.L.C. Huisjes, J.W. van Isselt, R. Jonk, Dr. J.M.H. de Klerk, Dr. J.B.A. Maintz, Dr. L.C. Meiners, Dr. I.J.R. Mertens, M. Metselaar, Dr. P.P. van Rijk, G.R. Timmens, and Dr. T. van Walsum for their contributions.

REFERENCES

- Alfano B, Brunetti A, Arpaia M, Ciarmiello A, Covelli EM, Salvatore M (1995): Multiparametric display of spin-echo data from MR studies of brain. *J Magn Reson Imaging* 5:217–225.
- Arnstein ND, Chen DC, Siegel ME (1990): Interpretation of bone scans using a video display. A necessary step towards a filmless nuclear medicine department. *Clin Nucl Med* 15:418–423.
- Atkins MS, Mackiewicz BT (1998): Fully automatic segmentation of the brain in MRI. *IEEE T Med Imaging* 17:98–107.
- Bendl R, Hoess A, Schlegel W (1995): Advanced tools for 3D radiotherapy planning. In: Lemke HU, Inamura K, Jaffe C, Vannier M, editors. *Computer assisted radiology '95*. Berlin: Springer-Verlag. p 1094–1099.
- Blinn J (1977): Models of light reflection for computer synthesized pictures. *Comput Graph* 11:192–198.
- Brown HK, Hazelton TR, Silbiger ML (1991): Generation of color composites for enhanced tissue differentiation in magnetic resonance imaging of the brain. *Am J Anat* 192:23–34.
- Brummer M, Mersereau RM, Eisner RL, Lewine RRJ (1993): Automatic detection of brain contours in MRI data sets. *IEEE Trans Med Imaging* 12:153–166.
- Christ RE (1975): Review and analysis of color coding research for visual displays. *Hum Factors* 17:542–570.
- Clarke LP, Velthuizen RP, Camacho MA, Heine JJ, Vaidyanathan M, Hall LO, Thatcher RW, Silbiger ML (1995): MRI segmentation: methods and applications. *Magn Reson Imaging* 13:343–368.
- Collins DL, Holmes CJ, Peters TM, Evans AC (1995): Automatic 3-D model-based neuroanatomical segmentation. *Hum Brain Mapp* 3:190–208.
- Dimitrov LI (1998): Texturing 3D-reconstructions of the human brain with EEG-activity maps. *Hum Brain Mapp* 6:189–202.
- Evans AC, Collins DL, Neelin P, Marrett TS (1996): Computer-integrated surgery. In: Taylor RH, Lavallée S, Burdea GC, Mösges R, editors. *Correlative analysis of three-dimensional brain images*. Cambridge, MA: MIT Press. p 99–114.
- Evans AC, Marrett TS, Torrescorzo J, Ku S, Collins DL (1991): MRI-PET correlation in three dimensions using a volume-of-interest (VOI) atlas. *J Cereb Blood Flow Metab* 11(Suppl 1):A69–A78.
- Foley JD, van Dam A, Feiner SK, Hughes JF (1990): *Computer graphics—principles and practice* (2d ed.). Reading, MA: Addison-Wesley.
- Freiherr G (1988): PET scanning leans toward practical side of medicine. *Diagn Imaging* June 146–154.
- Gershon ND (1990): Visualization and three-dimensional image processing of positron emission tomography (PET) brain images. Bellingham, WA: SPIE Press. p 144–149.
- Gershon ND (1994): From perception to visualization. In: Rosenblum L, Earnshaw R, Encarnação, J, Hagen H, Kaufman A, Klimenko S, Nielson G, Post F, Thalmann D, editors. *Scientific visualization—advances and challenges*. London: Academic Press. p 129–139.
- Gevens A, Brickett P, Costales B, Le J, Reutter B (1990): Beyond topographic mapping: towards functional—anatomical imaging with 124-channel EEGs and 3-D MRIs. *Brain Topogr* 3:53–64.
- Gevens A, Le J, McEvoy LK, Smith ME (1999): Deblurring. *J Clin Neurophysiol* 18:204–213.
- Gouras P (1991): Color vision. In: Kandel ER, Schwartz JR, Jessell TM, editors. *Principles of neural sciences*. London: Prentice-Hall International, Inc. p 467–480.
- Hagen MA (1991): How to make a visually realistic 3D display. *Comput Graph* 25:76–81.
- Hawkes DJ, Hill DLG, Lehmann ED, Robinson GP, Maisey MN, Colchester ACF (1990): Preliminary work on the interpretation of SPECT images with the aid of registered MR images and an MR derived 3D neuro-anatomical atlas. In: Höhne KH, Fuchs H, Pizer S, editors. *3D imaging in medicine*. Berlin: Springer-Verlag. p 242–251.
- Hawkes DJ, Ruff CF, Hill DLG, Studholme C, Edwards PJ, Wong WL (1995): Medical imaging: analysis of multimodality 2D/3D images. In: Beolchi L, Kuhn MH, editors. *3D multimodal imaging in image guided interventions*. Volume 19 of *Studies in health, technology and informatics*. Amsterdam: IOS Press. p 83–100.
- Heffernan PB, Robb RA (1984): A new procedure for combined display of 3-D cardiac anatomic surfaces and regional functions. *Comput Cardiol* 1111–1114.
- Höhne KH, Bomans M, Pommert A, Riemer M, Schiers C, Tiede U, Wiebecke G (1990): 3D visualization of tomographic volume data using the generalized voxel model. *Vis Comput* 6:28–36.
- Höhne KH, Hanson WH (1992): Interactive 3D segmentation of MRI and CT volumes using morphological operations. *J Comput Assist Tomogr* 16:285–294.
- Hu X, Tan KK, Levin DN, Pelizzari CA, Chen GTY (1990): A volume-rendering technique for integrated three-dimensional display of MR and PET data. In: Höhne KH, Fuchs H, Pizer SM, editors. *3D imaging in medicine*. Berlin: Springer-Verlag. p 379–397.
- Kajimura N, Uchiyama M, Takayama Y, Uchida S, Uema T, Kato M, Sekimoto M, Watanabe T, Nakajima T, Horikoshi S, Ogawa K, Nishikawa M, Hiroki M, Kudo Y, Matsuda H, Okawa M, Takahashi K (1999): Activity of midbrain reticular formation and neocortex during the progression of human non-rapid eye movement sleep. *J Neurosci* 19:10065–10073.
- Kamman RL, Stomp GP, Berendsen HJC (1989): Unified multiple-feature color display for MR images. *Magn Reson Med* 9:240–253.
- Kandel ER (1991): Perception of motion, depth, and form. In: Kandel ER, Schwartz JR, Jessell TM, editors. *Principles of neural sciences*. London: Prentice-Hall. p 441–466.
- Keller PR, Keller MM (1992): *Visual cues—practical data visualization*. Los Alamitos, CA: IEEE Computer Society Press.
- Kikinis R, Shenton ME, Gerig G, Martin J, Anderson M, Metcalf D, Guttmann CRG, McCarley RW, Lorensen WE, Cline HE, Jolesz FA (1992): Routine quantitative analysis of brain and cerebrospinal fluid spaces with MR imaging. *J Magn Reson Imaging* 2:619–629.
- Lemieux L, Hagemann G, Krakow K, Woermann FG (1999): Fast, accurate, and reproducible automatic segmentation of the brain in T1-weighted volume MRI data. *Magn Reson Med* 42:127–135.
- Levin DN, Hu X, Tan KK, Galhotra S, Pelizzari CA, Chen GTY, Beck RN, Chen C-T, Cooper MD, Mullan JF, Hekmatpanah J, Spire J-P (1989): The brain: integrated three-dimensional display of MR and PET images. *Radiology* 172:783–789.
- Levkowitz H, Herman GT (1992): Color scales for image data. *IEEE Comput Graph* 12:72–80.

- Livingstone M, Hubel D (1988): Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science* 240:740–749.
- Lutz R, Pun T, Pelligrini C (1991): Colour displays and lookup tables: real time modification of digital images. *Comput Med Imaging Graph* 15:73–84.
- Maes F, Collignon A, Vandermeulen D, Marchal G, Suetens P (1997): Multimodality image registration by maximization of mutual information. *IEEE T Med Imaging* 16:187–198.
- Maintz JBA, Viergever MA (1998): A survey of medical image registration. *Med Imaging Anal* 2:1–36.
- Murch G (1984a): The effective use of color: cognitive principles. *Tekniques* 8:25–31. (Issued by Tektronix, Inc., Beaverton, OR.)
- Murch G (1984b): The effective use of color: perceptual principles. *Tekniques* 8:4–9. (Issued by Tektronix, Inc., Beaverton, OR.)
- Murch G (1984c): The effective use of color: physiological principles. *Tekniques* 7:13–16. (Issued by Tektronix, Inc., Beaverton, OR.)
- Neider J, Davis T, Woo M (1993): *Open GL programming guide: the official guide to learning OpenGL*. Reading, MA: Addison-Wesley.
- Oosterhout JK (1994): *Tcl and the Tk toolkit*. Reading, MA: Addison-Wesley.
- Payne BA, Toga AW (1990): Surface mapping brain function on 3D models. *IEEE Comput Graph* 10:33–41.
- Pelizzari CA, Chen GTY, Spelbring DR, Weichselbaum RR, Chen C-T (1989): Accurate three-dimensional registration of CT, PET, and/or MR images of the brain. *J Comput Assist Tomogr* 13:20–26.
- Phong BT (1975): Illumination for computer generated pictures. *Commun ACM* 18:311–317.
- Porter T, Duff T (1984): Compositing digital images. *Comput Graph* 18:253–259.
- Price CJ, Veltman DJ, Ashburner J, Josephs O, Friston KJ (1999): The critical relationship between the timing of stimulus presentation and data acquisition in blocked designs with fMRI. *Neuroimage* 10:36–44.
- Ramsey NF, Kirkby BS, van Gelderen P, Berman KF, Duyn JH, Frank JA, Mattay VS, van Horn JD, Esposito G, Moonen CTW, Weinberger DR (1996): Functional mapping of sensorimotor cortex with 3D BOLD fMRI correlates highly with H₂¹⁵O PET rCBF. *J Cereb Blood Flow Metab* 16:755–764.
- Rehm K, Strother SC, Anderson JR, Schaper KA, Rottenberg DA (1994): Display of merged multimodality brain images using interleaved pixels with independent color scales. *J Nucl Med* 35:1815–1821.
- Ricard M, Tenenbaum F, Schlumberger M, Travagli J-P, Lumbroso J, Revillon Y, Parmentier C (1993): Intraoperative detection of pheochromocytoma with iodine-125 labeled meta-iodobenzylguanidine: a feasibility study. *Eur J Nucl Med* 20:426–430.
- Robb RA, Hanson DP (1996): The ANALYZE software system for visualization and analysis in surgery stimulation. In: Taylor RH, Lavallée S, Burdea GC, Mösges R, editors. *Computer-integrated surgery*. Cambridge, MA: MIT Press. p 175–189.
- Schlegel W (1996): Requirements in computer-assisted radiotherapy. In: Taylor RH, Lavallée S, Burdea GC, Mösges R, editors. *Computer-integrated surgery*. Cambridge, MA: MIT Press. p 681–691.
- Schlick C (1994): A fast alternative to Phong's specular model. In: Heckbert P, editor. *Graphics gems IV*. Boston: Academic Press. p 385–387.
- Smith AR (1978): Color gamut transformation pairs. *Comput Graph* 12:12–19.
- Stapleton SJ, Caldwell CB, Leonhardt CL, Ehrlich LE, Black SE, Yaffe MJ (1994): Determination of thresholds for detection of cerebellar flow deficits in brain SPECT images. *J Nucl Med* 35:1547–1555.
- Stokking R (1998): *Integrated visualization of functional and anatomical brain images*. Ph.D thesis, Utrecht University, The Netherlands.
- Stokking R, van Isselt JW, van Rijk PP, de Klerk JMH, Huiskens AWLC, Mertens IJR, Buskens E, Viergever MA (1999): Integrated visualization of functional and anatomical brain data: a validation study. *J Nucl Med* 40:311–316.
- Stokking R, Vincken KL, Viergever MA (2000): Automatic morphology-based brain segmentation (MBRASE) from MRI-T1 data. *Neuroimage* 12:726–738.
- Stokking R, Zuiderveld KJ, Hulshoff Pol HE, van Rijk PP, Viergever MA (1997): Normal fusion for three-dimensional integrated visualization of SPECT and magnetic resonance brain images. *J Nucl Med* 38:624–629.
- Stokking R, Zuiderveld KJ, Hulshoff Pol HE, Viergever MA (1994): Integrated visualization of SPECT and MR images for frontal lobe damaged regions. In: Robb RA, editor. *Visualization in biomedical computing 1994*. Proceedings SPIE, volume 2359. Bellingham, WA: SPIE Press. p 282–290.
- Taylor JM, Murch GM, McManus PA (1989): TekHVC: a uniform perceptual color system for display users. *Proc SID* 30:15–21.
- Valentino DJ, Mazziotta JC, Huang H (1991): Volume rendering of multimodal images: application to MRI and PET imaging of the human brain. *IEEE T Med Imaging* 10:554–562.
- Van den Elsen PA, Maintz JBA, Pol E-JD, Viergever MA (1995): Automatic registration of CT and MR brain images using correlation of geometrical features. *IEEE T Med Imaging* 14:384–396.
- Van den Elsen PA, Viergever MA (1994): Marker guided multimodality matching of the brain. *Eur Radiol* 4:45–51.
- Van Gelderen P, Ramsey NF, Liu G, Duyn JH, Frank JA, Weinberger DR, Moonen CTW (1995): Three dimensional functional MRI of human brain on a clinical 1.5 T scanner. *Proc Natl Acad Sci USA* 92:6906–6910.
- Van Herk M, Kooy HM (1994): Automatic three-dimensional correlation of CT-CT, CT-MRI, and CT-SPECT using chamfer matching. *Med Phys* 7:1163–1178.
- Viergever MA, Maintz JBA, Stokking R (1997): Integration of functional and anatomical brain images. *Biophys Chem* 68:207–219.
- von Stockhausen H-M (1998): *3D-visualisierung der Funktion und der Morphologie des menschlichen Gehirns aus tomographischen Daten*. Ph.D thesis, University of Cologne, Germany.
- Wahl RL, Quint LE, Cieslak RD, Aisen AM, Koepp RA, Meyer CR (1993): "Anatomometabolic" tumor imaging: fusion of FDG PET with CT or MRI to localize foci of increased activity. *J Nucl Med* 34:1190–1197.
- Watt A (1993): *3D computer graphics* (2d ed.). Reading, MA: Addison-Wesley.
- Weiss KL, Stiving SO, Herderick EE, Cornhill JF, Chakeres DW (1987): Hybrid color MR imaging display. *Am J Roentgenol* 149:825–829.
- Zuiderveld KJ (1995): *Visualization of multimodality medical volume data using object-oriented methods*. Ph.D thesis, Utrecht University, The Netherlands.
- Zuiderveld KJ, Stokking R, Viergever MA (1995): Integrated visualization of quantitative information with anatomical surfaces. In: Lemke HU, Inamura K, Jaffe C, Vannier M, editors. *Computer assisted radiology '95*. Berlin: Springer-Verlag. p 195–200.