Perception of color greatly enriches visual experience. Beyond its esthetic value, color vision is of great practical value for detection of patterns and objects that would be elusive in a world devoid of color.

P. Gouras, 1985

Chapter 5

Color Encoded Integrated Volume Visualization of Functional Image Data and Anatomical Surfaces

Abstract

This chapter evaluates the use of the HSV color model in a generic method, called 'Normal Fusion', for integrated visualization of functional data with anatomical surfaces. The first part of the method derives quantitative values from functional input data by sampling the latter along a path determined by the (inward) normal of an anatomical surface; 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 for a thorough assessment of the functional data in the integrated display. Two strategies for retrospective HSV manipulation are tested and several clinical examples are used to demonstrate the viability of the proposed approaches. Furthermore, the overall method for color manipulation is evaluated by five observers in a clinical study using 30 cases.

5.1 Introduction

3D imaging has become an essential procedure for a variety of clinical diagnostic and therapy planning procedures. Digital imaging modalities like MRI, CT, SPECT, and PET, generate huge amounts of volume data. Accordingly, there is a growing need for accurate volumetric rendering of the data so as to get a better understanding of anatomical structures and their interrelationships.

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. 1998). Examples of clinical applications that benefit from integrated 3D visualization are display of thickness of the skull from CT (Zuiderveld et al. 1995), display of electro-magnetic dipole data with CT and MR images (Van den Elsen and Viergever 1994, Van den Elsen et al. 1995), fusion of CT, MRI, and MRA 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 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 (see Chapters 3 and 4). Further examples of the fusion of functional and anatomical data can be found in (Gevins et al. 1995, 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 dataset (in this chapter we assume these features to be anatomical surfaces). In multimodality matching, the focus is shifting from techniques using frames, molds, 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 (Gerig et al. 1992, Kikinis et al. 1992, Shenton et al. 1992). 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, semi-automatic segmentation; these often simple methods are capable of segmenting a large volumetric dataset within a few minutes (Höhne and Hanson 1992, Robb and

Hanson 1996).

Assuming that the volume images have been matched and properly segmented, appropriate visualization techniques are required for presentation of the (usually intricate) information. The present chapter addresses this issue for the combined visualization of functional information with anatomical surfaces. Three problems have to be dealt with: *i)* To obtain the functional information that corresponds to a certain point on the anatomical surface; *ii)* To map the resulting functional values onto the anatomical surface; and *iii)* To optimally present the data to the clinician for interpretation.

In the previous chapter we applied the Normal Fusion technique to combine functional and anatomical information. This chapter describes the visualization strategy for Normal Fusion in more detail and uses the technique as a tool to discuss and evaluate the use of the HSV color model for integrated 3D visualization. This color model allows easy, rapid, and intuitive retrospective manipulation of the color encoding of the functional information, which will help interpretation. The aim of this chapter is to evaluate the HSV color model for the Normal Fusion method and test different strategies for the HSV color manipulation.

The chapter 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 anatomical surfaces. This method, called Normal Projection, traverses a volume (the functional dataset) along a secondary ray determined by the (inward) normal associated with the anatomical surface. A value is calculated from sample points along this trajectory and encoded onto the surface. Since 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 are presented and evaluated using several clinical examples. 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 chapter.

5.2 Methods and materials

5.2.1 3D visualization of anatomical surfaces

Realistic (computer) images from 3D datasets can be obtained with a process called volume visualization. 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 anatomical surfaces to be visualized have been identified during a preceding segmentation step. Since only visualization of surfaces is addressed in this chapter, this implies that the dataset which provides the anatomical context (usually a CT or MRI dataset) has been classified into surface and non-surface 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 Figure 5.1).

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, 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). However, for each point at the surface, the direction from point to observer (given by the unit vector E) has to be evaluated; this makes perspective projection computationally expensive. Furthermore, the image quality is only marginally improved compared to orthographic projection (Thaller et al. 1991), which assumes the vector E to be constant across the whole volume. Consequently, orthographic projection is still the preferred method for the majority of volume rendering applications.

Most modern rendering techniques calculate the surface direction (unit vector N) 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 ($3\times3\times3$) 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. Since photorealism is not required, a simple light reflection model is adequate for most visualization purposes. The most used light model is that of Phong (1975) which separates the reflected light into three components, *viz.*; *i)* an ambient (k_a), *ii)* a diffuse (k_d), and *iii)* a specular (k_s) component. Schlick (1994) modified the 'Phong' light model by an approximation of the specular component, 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:

$$I = I_a k_a + I_l \left(k_d (L \cdot N) + k_s \frac{t}{n - nt + t} \right) \tag{5.1}$$

where I_a is the intensity of ambient light, $n \in [1, \infty]$ 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

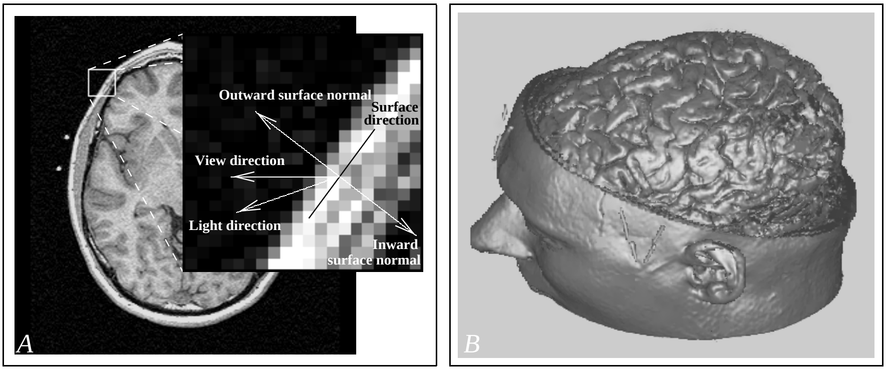

Figure 5.1 *3D visualization of anatomical surfaces. Frame (A): Principle of rendering of surfaces as commonly done in medical imaging. An MRI dataset has been classified into brain and skin voxels; the zoomed detail focuses on the (highest) grey values of the skin voxels. Shading calculations require the vector from the point at issue to the light source ("Light direction"), the surface normal at the point, and the vector from point to the observer ("View direction"); these variables are then used to evaluate a light model (see text). Frame (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.*

vector" calculated by $H = (L+E)/|L+E|$ (Blinn 1977).

Figure 5.1*B* contains a typical rendering of surfaces that was obtained using a $3\times3\times3$ neigborhood 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 fMRI, techniques for integrated visualization are called for to convey the information to the clinician. In this chapter we focus on the specific problem of combining functional data with the cortical surface of the brain. The first step towards integrated visualization of functional images and cortical anatomy is to project relevant functional information to points on the cortex. This is accomplished by a technique called Normal Projection, which is discussed next.

5.2.2 Normal Projection

In Chapter 4 we have introduced a technique that uses the surface normal for an anatomically accurate and viewpoint independent mapping of quantitative values from SPECT onto the surface of the brain from MRI. Figure 5.2 shows the principle of this Normal Projection technique. Surface normals are obtained by applying

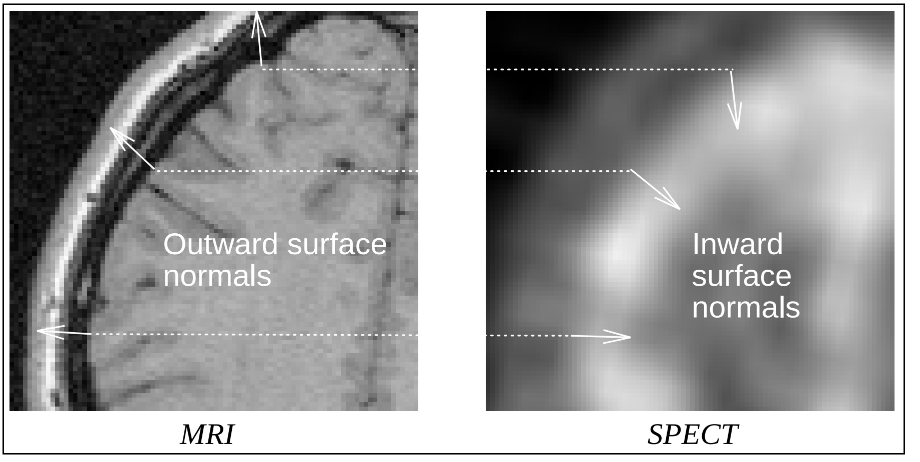

Figure 5.2 *Principle of Normal Projection. Quantitative information at a predefined anatomical surface is derived by performing calculations along a trajectory defined by the (inward) surface normal. On the left, (outward) surface normals for the brain are obtained using the gradient of an MR image. 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.*

gradient operators to each surface point that should be visualized; these normals are 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 dataset(s). For the application depicted in Figure 5.2, 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 of the quantitative information and the anatomical surface is done by color encoding the calculated quantities onto the surface. The next section discusses the selection of appropriate color assignments.

5.2.3 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 chapter to give an extensive overview of color perception; good introductions can be found in the chapters (Foley et al. 1990) and (Gouras 1991), and in a series of papers (Murch 1984[a-c]). 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 since 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 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 non-linear transformation of the RGB cube. Smith's HSV color model maps better on the visual sensations caused by colored light (Foley et al. 1990, Murch 1984a, Lutz et al. 1991). Here, hue refers to the wavelength which 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 5.3*A* (see page 81). 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.

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.

Since conversion between HSV and RGB is simple (see (Foley et al. 1990) and (Watt 1993) for pseudo-code) and HSV seemed an adequate basis for our purposes, we decided to use the HSV model for color encoding quantitative information onto anatomical surfaces.

5.2.4 Color encoding of quantitative information

Several authors have used color for integrated visualization. Their techniques can be roughly divided into four categories, *viz.*; *i)* alternate pixel display, *ii)* RGB integration, *iii)* HSV integration, and *iv)* color compositing. We note that the distinction between the last 3 categories can be difficult as these 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 (see Figure 3.1*B*). 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. and Rehm et al. consider this display 'visually pleasing' and report perceptual interactions between neighboring pixels (see also (Livingstone and Hubel 1988) and (Murch 1984b)), but they disagree on the ease of interpretation. Hawkes et al. find the display difficult to interpret owing to 'color smearing', Rehm et al. 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. compared to the saturated rainbow scale used by Hawkes et al.

With RGB integration, images to be combined are assigned to the primary colors red, 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 which 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 non-intuitive (see also (Brown et al. 1991) and Chapter 3).

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 MRI 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, *e.g.*, to present an overlay of different saturation levels of a hue onto grey values. It is noteworthy that with integrated HSV visualization low grey values for the value component cancel out (darken) the hue and saturation component (see also (Levkowitz and Herman 1988, Van Herk and Kooy 1994)), but with interpretation of brain images this appears insignificant.

Color compositing originates from the classic work on volume rendering by 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, Evans et al. 1996). Brown et al. (1991) describe the "multichannel color composite" method to integrate MR parameter images. With this method, 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 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 integrated with any of the previously mentioned 2D techniques, *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). The technique Normal Fusion, color encodes local functional information onto an anatomical surface (Chapter 4).

5.2.5 Normal Fusion with the HSV color model

We used the Normal Fusion technique to project SPECT information onto the brain surface 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 to improve understanding of the data (see Chapter 4). However, as indicated in the Section 5.2.3, 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* keep the information from the shading and the functional information separate. This is not only intuitive, but it also allows easy and retrospective manipulation of the color encoding without the need for a new rendering.

Given a white surface and a monochromatic light source, the intensity of the light reflected from the anatomical 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.

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

$\overrightarrow{\textbf{Figure}}$ **5.3** *Frame (A): HSV color model, after (Foley et al. 1990). Frame (B): Lookup table used for the clinical examples. Frame (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 Frame (B).*

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).

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, Levkowitz and Herman 1992, Encarnação et al. 1994)). We have adopted this strategy in our research by allowing manipulation of the lookup table.

5.2.6 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

$\overrightarrow{\textbf{Figure}}$ **5.4** *The strategies for HSV color manipulation, where separate storage is a subset of the recalculation scheme. The quantitative and anatomical information are 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.*

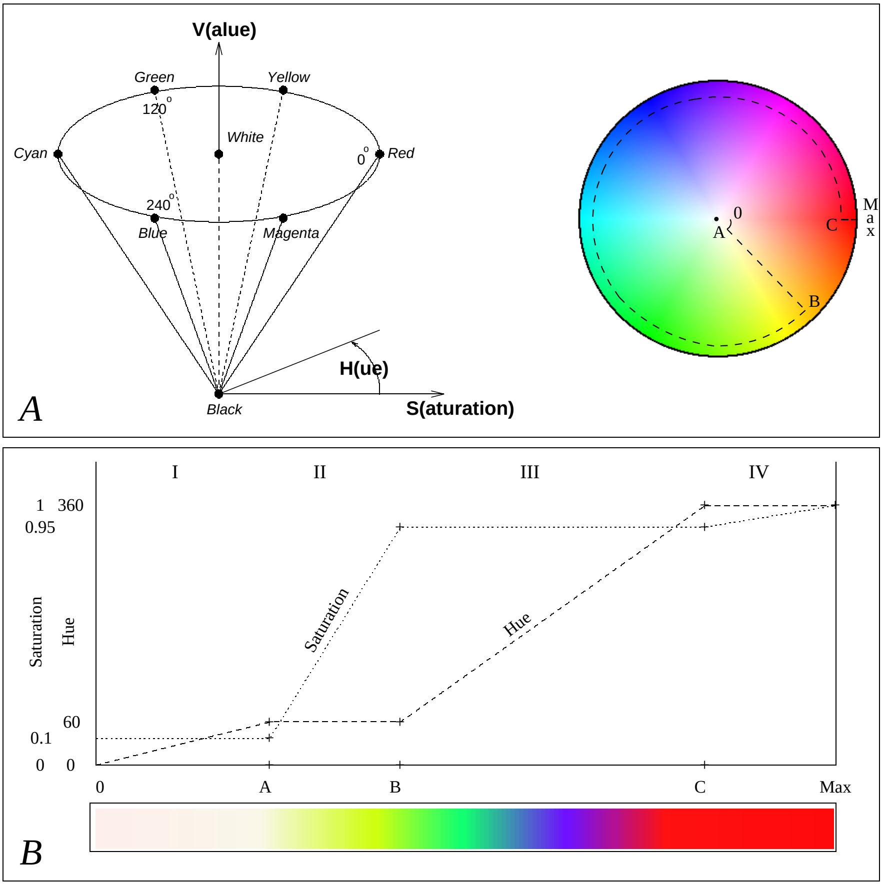

Figure 5.3 *See page 80.*

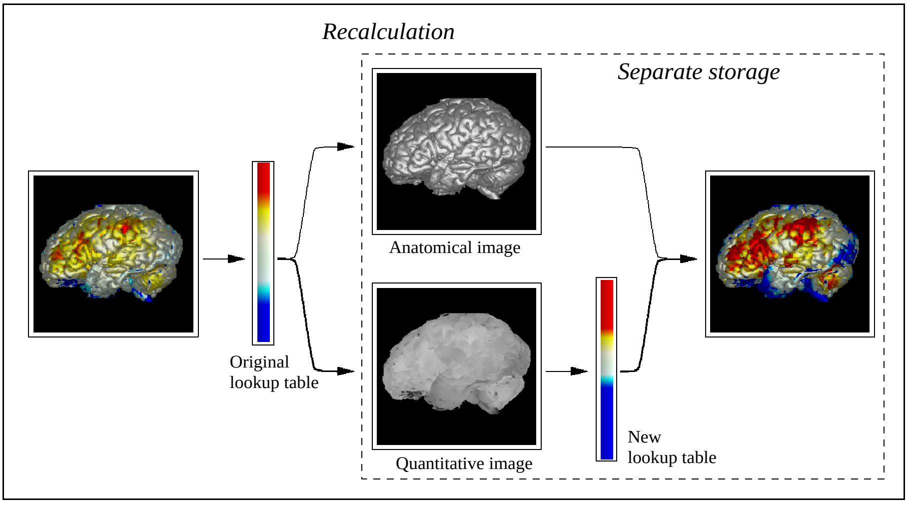

Figure 5.4 *See page 80.*

(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 to use the orange-yellow hue as humans are maximum 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 developed software to evaluate two strategies for manipulation of the color encoding, *viz.*; *i)* separate storage, and *ii)* recalculation (see Figure 5.4). The first strategy requires separate storage of the quantitative and anatomical information which are then combined into an output image. With the second strategy, the quantitative information at a surface voxel is reconstructed from the hue and saturation information 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 (see Section 5.3) and a simple HSV color scale aimed at signaling hot-spots (see Figure 5.3*B*). The range is divided into four parts and in each part of the range only the hue or the saturation is gradually increased to obtain the necessary one-to-one relationship. The quantitative information of interest is represented by a hue gradually increasing from 60° (yellow) to 360° (red) (area III of the graph in Figure 5.3*B*). The area of quantitative information of minor interest (area I) is represented by low saturations to obtain a non-colored 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 quite straightforward; we therefore used manipulation of the location of the control points, *i.e.*, the numbers corresponding to locations A, B, and C in Figure 5.3*B* for manipulating the color assignment schemes.

Following the evaluation of the two strategies, we applied one of the strategies in a clinical evaluation using 30 SPECT/MRI cases and five nuclear medicine physicians. These observers were already performing the validation study described in Chapter 6 using the same 30 SPECT/MRI cases. We extended the study to assess the opinion of clinicians on the color manipulation strategy with the Normal Fusion technique. For the color manipulation we decided to use separate storage of the anatomical and functional information instead of recalculation. We wanted to avoid the restrictions imposed by the required one-to-one mapping with the recalculation strategy. As both cold and hot-spots had to be signaled, we applied a more intricate lookup table than the one initially used. Consequently, more control points (six) and thus more interaction was required from the observers. A simple solution was to couple 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.

5.2.7 Processing and visualization

Registration of the datasets was done either by using external arrow-shaped skin markers and in-house developed software (Van den Elsen 1993), or by the surface-matching facilities in the ANALYZE^{TM} (Robb and Hanson 1996) software package when markers were not used in the acquisition process. We chose to match to the anatomical data to avoid degradation of the visualizations and we used ANALYZE^{TM} for the segmentation of the datasets. 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.

5.3 Results I: Clinical examples

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 (TS). The first case is to illustrate the Normal Projection technique, and the latter two cases are to illustrate more specifically the benefits of the HSV color encoding scheme. For all cases, MRI data were used to calculate the normal for shading and Normal Projection, whereupon the functional dataset was sampled each mm along the inward surface normal until a depth of 10 mm was reached (see Figure 5.2). The maximum value of these samples was calculated and then color encoded onto the surface using the lookup table of Figure 5.3*B*. Applied in this way, the Normal Fusion image depicts the maximum functional activity below the brain surface until a depth of 10 mm. The results are shown in Figures 5.5 and 5.6.

5.3.1 PET/MRI and fMRI/MRI monitoring of finger opposition

Over the last few years fMRI has emerged 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.

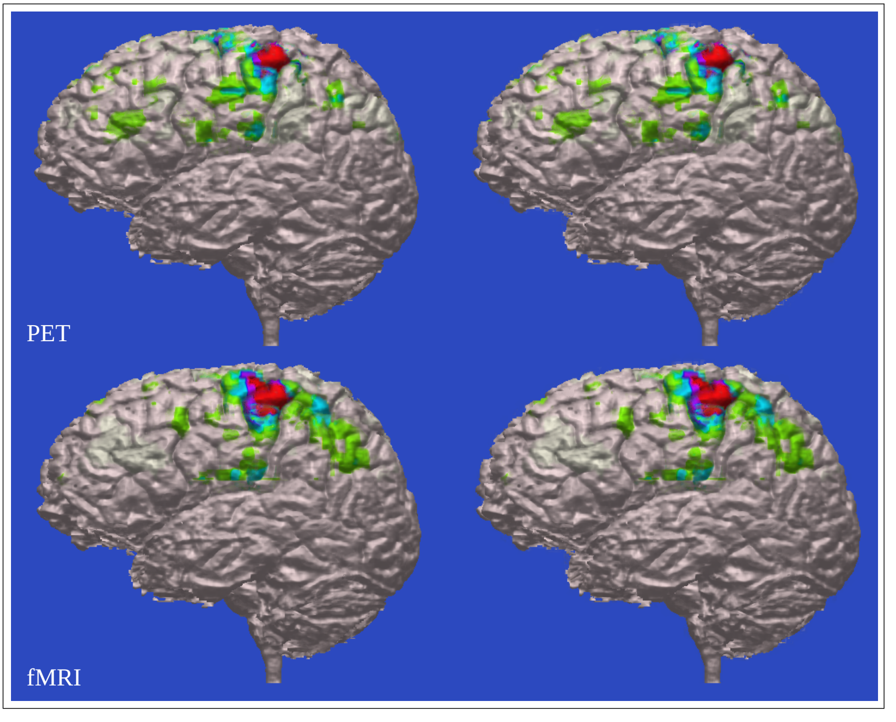

Figure 5.5 *Surface color encoding of* PET *subtraction activity (top row) and* fMRI *data (bottom row) for a finger opposition task. The brain cortex was extracted from MRI data. The presented 3D renderings are stereo images (cross-fusion) of the left hemisphere with a depth range of 0–10 mm. The maximum value over the depth range was used to color encode the corresponding surface voxel. The lookup table is shown in Figure 5.3B.*