Display of Merged Multimodality Brain Images Using Interleaved Pixels with Independent Color Scales

Kelly Rehm, Stephen C. Strother, Jon R. Anderson, Kirt A. Schaper and David A. Rottenberg

Health Computer Sciences and Departments of Radiology and Neurology, University of Minnesota; and PET Imaging Service, Veterans Administration Medical Center, Minneapolis, Minnesota

This article reviews common methods for two-dimensional display of registered multimodality brain images and describes a software package for presentation of merged MRI and PET images that runs on a workstation with an eight-bit color display. The software package displays merged brain images from multiple modalities in a way that is readily manipulated, visually pleasing and easy to interpret. The display method used, i.e., interleaving of alternate pixels with independent color scales, is effective in producing merged images with high contrast-detail for each volume. Interleaving images from different volumes creates unusual perceptual effects, one of which is the apparent camouflage of low-contrast signals by high values in the paired volume. **Methods:** The camouflage effect was thought to arise from perceptual merging of adjacent pixels. An observer experiment was conducted to investigate this tendency of high-activity PET data to obscure low-contrast detail in interleaved MRI data in spite of the digital independence of neighboring pixels. Four observers were presented with 20 combinations of signal plus background targets with uniform mask images, using a two-alternative forced-choice experimental design with 50 trials per combination. **Results:** The psychophysical evaluation of the ability of human observers to detect the simple test objects in an interleaved image presentation indicated a statistically significant camouflage effect of one volume on the other for some combinations of target and mask. The concept of perceptual merging of adjacent pixels was able to predict which combinations caused the greatest degradations in performance. **Conclusions:** The image interleaving approach to the display of two-dimensional slices from registered image volumes makes efficient use of an eight-bit color display. Contrast resolution of both individual volumes is high compared with that in other techniques and the volumes are presented in familiar color scales. However, the method yields an unexpected camouflage effect that tends to obscure low-contrast signals. The practical effect of such camouflage on the interpretation of clinical images remains to be investigated.

Key Words: multimodality displays; image interpretation; observer experiment; MRI; PET

J Nucl Med 1994; 35:1815–1821

Received July 28, 1993; revision accepted Apr. 4, 1994.

For correspondence or reprints contact: Kelly Rehm, Health Computer Sciences, University of Minnesota Medical School, Box 511, 420 Delaware Street S.E., Minneapolis, MN 55455.

The field of neuroimaging has been concerned for some time with the relationship between functional and structural information (*1*). Various techniques exist to register anatomic volumes obtained from MRI or CT scanners to functional volumes obtained from PET and SPECT cameras. Various methods that present the results of registration operations so that they can be easily interpreted are currently being investigated (*2–6*).

Techniques to display registered image volumes on 8- and 10- to 24-bit workstations have been previously discussed (*3–9*). Workstations with limited (eight-bit) color capabilities impose severe constraints on techniques that present composite (or "fused") volumes. This article reviews common methods to visualize registered structural MRI and PET (or other) functional volumes, describes a software package that has been developed at the Minneapolis Veterans Affairs Medical Center (VAMC) for two-dimensional display of registered volumes on an eight-bit display system and, finally, presents the results of a study of the VAMC visualization approach in regard to the detection of low-contrast signals in fused images.

TECHNIQUES FOR TWO-DIMENSIONAL PRESENTATION OF REGISTERED IMAGE VOLUMES

Parallel Display

In this approach, slices from two or more registered volumes are displayed side by side, with shared or independent color scales. For an eight-bit color display system, 256 colors are divided (evenly or unevenly) among the volumes. Spatial connections between matched slices are made by grid markings, simultaneous pointing or simultaneous display of regions of interest (ROIs) (*9–14*). An advantage of this approach is familiarity, i.e., slices from individual volumes can be presented in their conventional color scales. Users do not need to translate their mental atlas of anatomic and functional data to a new presentation space. A disadvantage is the user's difficulty in visually judging spatial registration. There are well-known visual illusions in which judgments of the relative sizes of circles (Titchner illusion) and the lengths of

lines (Müller-Lyer illusion) are affected by their surroundings in side-by-side presentation (*15*).

Color Compositing

Color compositing (also referred to as opacity-weighted compositing) is a computer graphics technique used to form composite images of overlapping objects of varying opacity (*16*). This method requires an additional level of encoding for data representation, i.e., each volume has an associated "matte" that defines the locations of its influence on the composite image. Matte values vary from zero (no coverage) to one (full coverage). A multiplicative factor, w, can explicitly address the opacity of a volume by scaling its matte values. This technique has been used in both two-dimensional and three-dimensional image presentation of composite volumes (*13,17–22*). Color compositing makes full use of the spatial correlation of two volumes and removes this burden from the viewer. Manipulation of color, matte and opacity allow the user to change the emphasis given each volume to the extremes of presenting each volume in its "pure" form. Compositing two volumes, each represented in 256 colors, requires 256^2 colors in the displayed image. For an eight-bit display system, the volumes must be coarsely quantized to use this method.

Red, Green, Blue (RGB) Color Encoding

This encoding uses the three primary colors (RGB) to represent each of a pair or trio of volumes. The color of the displayed pixel is an additive mixture of the two (or three) colors. The color scales of the individual volumes may be selected so that equal contributions from both volumes result in a colored pixel [R and G = yellow (*3,23,24*)] or in a gray pixel [R + 0.5 G and B + 0.5 G = gray (*8*)].

As with color compositing, the method requires coarse quantization of the data volumes. Spatial registration information is readily available to the user; however, the user must determine the contribution of pixels from the individual volumes by hue and lightness [in terms of the hue-lightness-saturation scheme of color description (*25*)]. Wahl et al. (*24*) uses this method in fusing MRI and PET body images; Alfano and Ney (*3,10*) discuss fusing different MRI sequence volumes. In such fusions, discernment of the relative contribution of pixels is done by judgment of a hue shift, with pixels that "reinforce" each other at midhue and increased brightness. Assessment of the absolute value of pixels in either volume requires recognition of the fact that a pixel represented by the maximum brightness of G in one volume can run the gamut from maximum G to maximum yellow in an RG composite image with a corresponding increase in brightness. This may make the method awkward when functional and anatomical images are combined. It is important to relate activity level to spatial location in this case.

Hue and Lightness Encoding

The method of hue and lightness encoding of volumes for presentation is similar to RGB encoding. However, hue, saturation and lightness are commonly assumed to be independent descriptors of color (*26,27*), which suggests that they are better suited than RGB to represent multivariate data. Techniques described in the literature have focused on hue and lightness variations, with saturation kept fixed (*28*). Weiss et al. (*7*) encoded paired MRI sequence volumes with 256 hues and 4 lightnesses. Wells et al. (*5*) encoded paired MRI sequences with 64 hues and 64 lightnesses, and Hawkes et al. (*9*) and Hemminger (*6*) encoded paired MRI and PET volumes with 32 lightnesses for the MRI and 8 hues for the PET data. Choice of the hue gamut has been fairly arbitrary, although Wells et al. (*5*) describe an experiment in which observers expressed a preference for the popular hot-body (or "hot-metal") color scale over a uniform chromaticity hue scale, and Crowe et al. (*4*) showed that hot-body and RGB scales compare favorably with the uniform chromaticity scale to convey sharpness and intensity differences.

As with RGB encoding, it may be difficult to link uniform areas cognitively in one volume when the area in the paired volume is nonuniform. Human observers have been shown to have difficulty recognizing identical hue values when a lightness difference is present (*26*). Variance in the lightness of a pixel modifies the perception of its hue, i.e., a phenomenon known as the Bezold-Brücke effect (*27,29*). This effect may be slight in a hue scale with large shifts between adjacent steps. However, a potentially more serious effect arises when anatomic and functional data are merged. When lightness is controlled by the anatomic image, the functional activity is partially constrained to follow anatomic boundaries, i.e., at very low lightnesses, a pixel is perceived as nearly black, regardless of the functional activity level. When the range of lightness or hue is compressed to a single value, each volume can be interpreted independently; however, a hue-encoded volume is less easily interpreted than when it is encoded in hue and lightness or lightness alone. It has been noted elsewhere that "the psychological similarity of colors is arranged around a circle, therefore differences in color are hard to interpret as increases in quantity" (*30*).

Alternate Pixel Presentation

This technique differs from the methods discussed earlier. Each pixel displayed on the screen is derived entirely from one or the other volumes, without an algebraic combination. Each volume uses independent color scales, not necessarily RGB or hue-lightness scales. [This method may be familiar from computer graphics as the "screen-door transparency" approach (*25*).] Hawkes et al. (*9*) used this approach on an eight-bit color system for presentation of registered SPECT and MRI images. Although an eight-bit display system imposes limitations on this technique, as it does on the other compositing methods, the restrictions are not as severe. The 256 colors available represent N + M values, instead of N × M values; therefore, N and M can be larger for both volumes. An apparent advantage of this approach is that the individual volumes may each be presented with familiar color scales, as is done in a parallel

FIGURE 1. Interleaved presentation of registered pseudo-T1 MRI volume and FDG-PET volume from a normal volunteer.

display. Hawkes et al. described this method of presentation as "visually pleasing" but difficult to interpret because of color mixing. The software package described in this article uses alternate pixel presentation; however, in contrast to Hawkes et al.'s result, at this institution, the method is considered both visually pleasing and easy to interpret.

MATERIALS AND METHODS

Merged Display Package

The merged volume display package accepts as input two sets of volume data (e.g., MRI and PET) that have been previously registered to each other and the Talairach atlas by another program that provides a variety of alignment techniques (*31*). The package is written in PV-WAVE version 3.10 (Precision Visuals Inc., Boulder, CO) and runs on an eight-bit color SPARCstation (Sun Microsystems, Mountain View, CA) supporting X windows. Images are fused with the pixel "interleaving" method (alternate pixel presentation) on a slice-by-slice basis. Paired slices, which have previously been rescaled to pixel values of 0 to 255, are first recoded into odd and even pixel values, respectively. In effect, the data is reduced from eight to seven bits.

After recoding is done, the two registered slices are multiplied by "checkerboard" masks to select alternate pixels from each. Each square of the checkerboard is a single pixel, i.e., "black" squares select pixels from the odd-valued slice and "white" squares, from the even-valued slice. The resulting interleaved slice has the same pixel dimensions as the originals. In the color table of the display, RGB values assigned to odd-numbered entries show the odd-valued MRI pixels; even-numbered entries select the PET pixels. Figure 1 illustrates the image interleave presentation of a trio of registered MRI and PET slices of a normal volunteer. The MRI volume is a pseudo-T1 study (a linear combination of 20- and 80-msec scans), presented with a linear gray scale. The PET volume is a [^{18}F]fluorodeoxyglucose (FDG) study presented with a hot-metal color scale.

The package initially presents interleaved axial, coronal and sagittal slices at the center voxel of the brain volume. Menu buttons on the display (Fig. 1) allow navigation through the volumes, identification of structures and adjustment of the display characteristics. The user interface of the display package is entirely mouse driven. Navigation may be accomplished in three ways. First, soft buttons allow single-slice translations through the volume in the x, y and z directions. Second, the user may select a point of interest in any of the three orthogonal slices, with an automatic update of the other two slices. Third, the user may select an anatomic structure from a list of locations based on the Talairach atlas or a list of ROIs defined at this institution. The program then displays the trio of slices that correspond to that location. The "mark" and "what" soft buttons, respectively, flag the point of intersection in the three displayed slices and identify the Talairach structure nearest to the selected point. "Mirror" moves the point of intersection from one hemisphere to the other; slices are updated as required. The "zoom" option allows the user to select any enlargement factor up to 10.0. Both volumes are zoomed by bilinear interpolation (registration is maintained) and then interleaved for presentation. Slices can be zoomed about the point of intersection or about the center of each slice.

The "lut" soft button allows selection and manipulation of the color scale that represents each volume and the degree of emphasis given each volume. The user may select one of nine basic color scales to be applied to either volume. The current color scale for each volume is presented as a color stripe next to the menu. Color scales for each volume are manipulated completely independently from one another. Each color scale retains the adjustments made to it even when it is deselected, i.e., this allows a user to switch between maps without readjustment.

Brightness and contrast can be controlled by window and level adjustments and also by means of histogram-based modification of the color scale. As the value of each pixel is its key into the color table, histogram modification effectively assigns new colors to the pixel values. The pixel value histogram of the entire volume is

used to calculate a histogram equalization map (*32*). The user selects a weighting factor that controls the histogram equalization applied; the final mapping is a weighted combination of a linear and histogram equalization map. Histogram modification can create a nonlinear color scale that cannot be matched with window and level adjustments.

The final adjustments available to the user control the apparent emphasis given each volume in the interleaved presentation. The user may instantly "hide" either volume, which sets all its color values temporarily to black. The user may also select a degree of emphasis for the color scale of each volume. When the emphasis of the color scale is decreased by one volume, its apparent transparency is increased (an emphasis of zero is the same as hiding the volume).

Signal Detection Experiment

A small two-alternative forced choice (2AFC) observer performance study was designed to study the "camouflage" effect of interleaving on the detection of test objects in images displayed on the workstation. The percentage of correct decisions in a 2AFC study is theoretically equivalent to the area under a receiver operating characteristic curve, although 2AFC designs usually require a greater number of trials to achieve the same error in the estimated area. However, because the observer is forced to choose an image, 2AFC methods allow measurement of the ability to detect objects at the threshold of visibility, without regard to the observer's willingness to "call" a low-visibility signal (*33*).

Four target images (signal plus background represented in gray scale) and four mask images (uniform field represented in hot-metal color scale) were tested in combination. The signal in the target images consisted of a light square superimposed on a darker background square, with a 1-pixel white frame on the background. The images were noise free. The set of target and mask combinations is shown in Figure 2. Each combination of target and mask images was presented to each of four observers in 50 trials on the same workstation under the same ambient lighting conditions. The duration of presentation of each trial was not restricted, nor were the observers constrained to remain at a fixed distance from the screen. All trials of a target/mask combination were presented in a single session, and the order of sessions was randomized across the observers.

RESULTS

The Weber ratio describes the minimum incremental change in luminance ($\Delta L/L$) at which a simple disk of luminance $\Delta L + L$ can be detected against a background of luminance, L (*29,34*). For photopic vision (luminances greater than about 1 cd/m^2), the Weber ratio is constant with respect to luminance but varies with the temporal and spatial frequency of the target. The smallest (most sensitive) value of this constant is approximately 0.01 (*27*). The luminance response of the SPARCstation monitor was measured for each color of the interleaved gray and hot-metal color table. The luminance ratios of adjacent entries in the gray and hot-metal color scales were greater than the maximum sensitivity threshold for all but the lowest values in each scale (where the luminances were so low, 0.02 cd/m^2, as to be indistinguishable by the photometer). These data are consistent with the results of other researchers (*35*) and indicate that the lowest contrast targets should be detectable. However, the luminance ratio threshold for normal working conditions is expected to be greater (i.e., less sensitive) than 0.01.

The luminance ratio partially describes the effect when a uniform mask image is interleaved with a target image. Over the area of the signal, the average luminance is $0.5 \cdot (\Delta L_{target} + L_{target}) + 0.5 \cdot L_{mask}$; over the background the luminance is $0.5 \cdot L_{target} + 0.5 \cdot L_{mask}$. The inclusion of L_{mask} in this way is mathematically equivalent to an increase in the background luminance without an increase in ΔL. This causes a reduction in the luminance ratio for the target. This "merged" luminance ratio was calculated for each of the 20 target/mask combinations. Table 1 illustrates the luminance ratio degradation caused when uniform masks are interleaved with the target images. The effect of interleaving the brighter masks with the darker targets resulted in luminance ratios near or below the maximum sensitivity threshold value of 0.01. For these combinations, a reduction in observer performance in target detection was predicted in spite of the physical independence of the target and mask pixels and the additional cue of color difference.

The average detection rates for the observers are presented in Table 2. The reductions in detection rate are roughly consistent with the pattern shown in the change in luminance ratio in Table 1. For the combinations with luminance ratios less than the minimum threshold, all observers were essentially guessing about which image contained the signal. Based on a binomial distribution with N = 50 and alpha = 0.05, detection rates of 92% and lower for an individual represent a statistically significant degradation in performance compared with that of noninterleaved images (100% detection rate) (*36*). The superscripts on the entries in Table 2 indicate the number of individuals that showed significant degradations.

DISCUSSION

The software package presented here, which uses the image interleave method to present registered two-dimensional brain slices, is a tool for navigation and interpretation of paired volume datasets, be they MRI with MRI,

TABLE 1
Luminance Ratio of Target-to-Background* as a Function of Mask Luminance

Luminance of target	No mask	Luminance of mask (cd/m^2) 0.02	2.63	20.30	52.10
1.34	0.136	0.133	0.042	0.007	0.003
5.49	0.070	0.070	0.046	0.014	0.006
33.88	0.031	0.031	0.029	0.018	0.012
69.43	0.023	0.023	0.022	0.018	0.013

*Target background is the next gray level less than the target. Mask values correspond to pixel values of 1, 85, 171 and 247 in a hot-body scale. Target values correspond to pixel values of 54, 86, 170 and 252 in a gray scale.

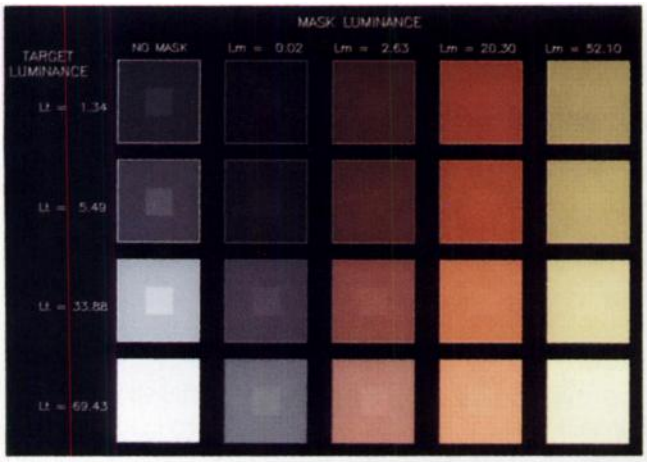

FIGURE 2. Signal plus background targets in gray scale interleaved with uniform masks in color scale. The central signal subtended approximately 1.1° of visual angle. For the purposes of illustration, the signal is stronger than that used in the 2AFC experiment.

PET with MRI or PET with PET. Compared with other techniques, the process makes good use of the limited capabilities of an eight-bit color display system by maintenance of contrast resolution and the use of familiar color scales. In Figure 3, the registered MRI and PET pair is presented with image interleaving, color compositing, RG encoding and hue-lightness encoding methods, as described previously.

There is an obvious contouring effect in the functional data values for all but the image interleaving approach. The level of functional detail available in the interleaved slices is markedly higher than that with the other methods. In the presentation of anatomic and functional data, it is perhaps acceptable to show only a few levels of activity, but when two anatomic datasets are merged, the artifact is more objectionable. Another, more subtle distinction exists among these presentation techniques. As predicted by the mechanism of hue and lightness encoding, PET activity levels in the hue/lightness slice are not apparent when the MRI values are at or near minimum. In the area of the sinuses, for example, activity shows well in the RG, composited and interleaved slices but shows poorly in the hue/lightness-encoded slice. The potential to impose anatomic

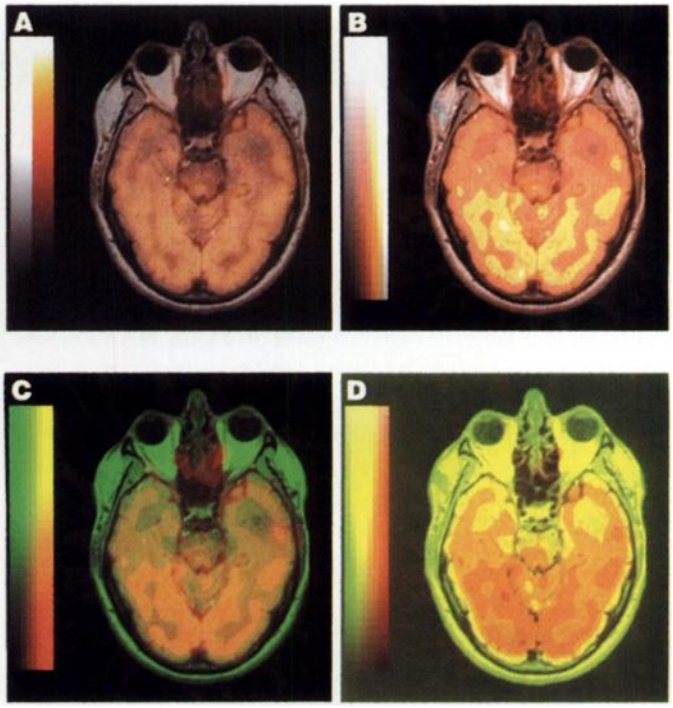

FIGURE 3. Axial MRI and PET slices presented with alternate methods. (A) Image interleave, i.e., MRI in gray and PET in hot-body scale. (B) Color compositing, i.e., MRI in 30 gray steps and PET in 8-step hot-body scale with a uniform matte. (C) RG encoding, i.e., MRI coded in linear red and PET coded in linear green. (D) Hue-lightness encoding, i.e., MRI in 30 lightness steps and PET in 8 hue steps (minimum lightness value is nonzero to allow hue changes to be discerned). For B to D, the color stripe represents MRI encoding in the vertical dimension and PET encoding in the horizontal dimension.

TABLE 2
Percentage of Targets Detected in Two-Alternative Forced Choice Trials as a Function of Mask Value

Luminance of target	No mask	Luminance of mask 0.02	2.63	20.30	52.10
1.34	100%	92^{1*}	96^{1}	65^{4}	54^{4}
5.49	100	100	100	90^{3}	52^{4}
33.88	100	98^{1}	96^{1}	100	92^{2}
69.43	100	99	95^{2}	89^{2}	73^{3}

*Average signal detection rate for four observers. Superscript indicates no. of observers exhibiting a significant reduction in detection rate.

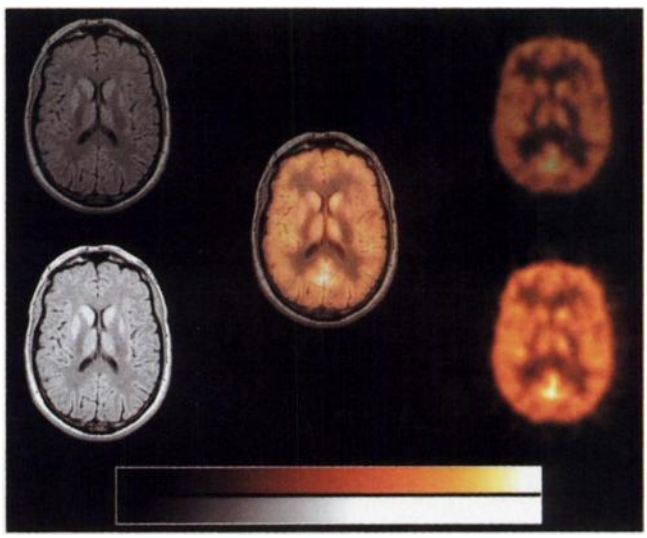

FIGURE 4. Interleaved, pure and hidden volumes. (Center) Interleaved MRI and PET slice. (Upper left) Interleaved MRI slice, with PET slice set to black. (Lower left) Noninterleaved MRI slice (odd-value encoding). (Upper right) Interleaved PET slice, with MRI slice set to black. (Lower right) Noninterleaved PET slice (even-value encoding).

boundaries unwittingly on activity data is a drawback to this method.

Although the interleaved pixel presentation package has been favorably received by users at the VAMC, a question remains as to the effectiveness of merged versus individual presentation of two-dimensional slices. There are several observations to be made about the characteristics of interleaved images, as illustrated in Figure 4. The interleaving of two differently colored volumes gives the appearance of one partially transparent volume overlaying another. At a comfortable viewing distance of approximately 16 in., users have commented that they tend not to perceive individual elements of the checkerboard unless the contrast of adjacent pixels is high. As mentioned previously, the ability to hide either slice allows independent interpretation of its pair. However, as can be seen in the upper left and upper right slices of Figure 4, when one of the interleaved slices is hidden, the remaining slice retains spatial detail but appears dark compared with a noninterleaved slice. Another effect is that colored slices are more vivid when displayed alone than when interleaved with gray slices. A potentially more serious effect observable with interleaved images is that low-contrast structures in the MRI volume seem more difficult to perceive when interleaved with the colored PET volume.

These effects may be explained by the fact that the SPARCstation monitors used have pixel matrices of 1152 × 900 on a field of 11.5 × 9.0 in. At a viewing distance of 16 in., the maximum spatial frequency of the interleaved image is approximately 14 cycles/degree of visual angle. At the luminance levels obtainable with a typical cathode ray tube (2–200 cd/m^2), the spatial contrast sensitivity of the eye peaks at less than 5 cycles/degree and falls off by a factor of 10 at 14 cycles/degree (*34,37*). Unless the contrast of adjacent pixels is very high, a 2 × 2 square of adjacent pixels tends perceptually to merge into a larger "megapixel."

The perceptual merging concept explains the seeming invisibility of the checkerboard and the reduction in brightness and vividness of interleaved images. When one slice is hidden, its mate is displayed interleaved with a black background. Averaging of black and colored pixels reduces the brightness of the larger unit. Similarly, averaging of gray with colored pixels results in a merged color with lower saturation, although it may be either lighter or darker than the original.

Perceptual merging of adjacent pixels is also a potential reason for the apparent camouflage of low-contrast structures. Signal detection theory and experiments have studied the ability of human observers to detect a simple signal superimposed on a uniform background, i.e., the threshold of detectability is described by the Weber ratio of target-to-background luminances. The merged ratio of signal-to-background luminances is reduced when target images are interleaved with uniform backgrounds. If the signal-to-background luminance ratio is near the Weber threshold ratio in a pure (noninterleaved) image, interleaving of a uniform mask image with the target image may cause the signal to fall below the threshold of detectability.

As noted earlier, Table 2 indicates that the camouflage phenomenon reduces the ability of an observer to detect gray objects at the minimum gray level of contrast that can be rendered in interleaved images. The brighter the camouflage color is, the greater the reduction in detection. For several target/mask combinations, the reduction in detection rate was shown to be significant at an alpha level of 0.05 in a one-tailed test of the hypothesis that no change occurred in the detection rate. Also of note in Table 2 is the fact that there is little difference in the detection of signals in noninterleaved images and those interleaved with the minimum mask luminance (only one observer exhibited a significant reduction). This is analogous to the case of hiding the masking volume, as shown with clinical data in the left column of Figure 4. This result is interesting because it complicates the procedure of interpreting merged datasets if it is necessary to examine each dataset as noninterleaved images to detect subtle features. The 2AFC data suggests that the detection of low-contrast features may not be significantly degraded if the masking volume is simply hidden, i.e., without restoring the missing pixels in the target volume.

Because clinical images exhibit both random and structured noise that further obscures low-contrast features, the camouflage effect may grow more pronounced with real objects. However, when two images are not equally "important," i.e., when—as is the case with these PET/MRI merges—the MRI merely serves as a structural template for the PET image, the camouflage effect may not be of great concern. In spite of this potential shortcoming of the image interleaving technique, its many advantages (ease of use, visually pleasant images, high contrast-resolution and familiar color scales) make it the method of choice for the display of merged multimodality images on eight-bit color displays.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Veterans Affairs and by National Institutes of Health grants NS25563, NS25701, and DA07428. Dr. Rehm is supported by research training grant 2T15LM07041 from the National Institutes of Health.

REFERENCES

1. Mazziotta JC, Koslow SH. Assessment of goals and obstacles in data acquisition and analysis from emission tomography: report of a series of international workshops. *J Cereb Blood Flow Metab* 1987;7:S1–S3.
2. Barillot C, Lemoine D, Le Briquer L, Lachmann F, Gibaoud B. Data fusion in medical imaging: merging multimodal and multipatient images, identification of structures and 3D display aspects. *Eur J Radiol* 1993;17:22–27.
3. Alfano B, Brunetti A, Ciarmiello A, Salvatore M. Simultaneous display of multiple MR parameters with "quantitative magnetic color imaging". *J Comput Assist Tomogr* 1992;16(4):634–640.
4. Crowe EJ, Sharp PF, Undrill PE, Ross PG. Effectiveness of colour in displaying radionuclide images. *Med Biol Eng Comput* 1988;26:57–61.
5. Wells MG, Sharp PF, Law ANR. Principles and appraisal of combined images in NMR. *Med Biol Eng Comput* 1989;27:277–280.
6. Hemminger BM. Isoluminance: a color technique for visualizing multivariable medical image data. In: Kim Y, ed., *SPIE medical imaging VII: image*

capture, formatting, and display. Bellingham, WA: SPIE;1993;1897:325–335.
7. Weiss KL, Stiving SO, Herderick EE, Cornhill JF, Chakeres DW. Hybrid color MR imaging display. *AJR Am J Roentgenol* 1987;149:825–829.
8. Ney D, Fishman EK, Dickens L. Interactive multidimensional display of magnetic resonance imaging data. *J Digit Imaging* 1990;3:254–260.
9. Hawkes DJ, Hill DLG, Lehmann ED, et al. 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 SM, eds., *3D imaging in medicine*. Berlin: Springer-Verlag; 1990:241–251.
10. Evans AC, Beil C, Marrett S, Thompson CJ, Hakim A. Anatomical-functional correlation using an adjustable MRI-based region of interest atlas with positron emission tomography. *J Cereb Blood Flow Metab* 1988;8:513–530.
11. Barillot C, Lemoine D, Gibaud B, Toulemont PJ, Scarabin JM. A PC software package to confront multimodality images and a stereotactic atlas in neurosurgery. *SPIE Medical Imaging IV: Image Capture and Display* 1990;1232:188–199.
12. Bidaut L. Composite PET and MRI for accurate localization and metabolic modeling: a very useful tool for research and clinic. In: Loew M, ed., *SPIE Medical Imaging V: Image Processing* 1991;1445:66–77.
13. Valentino DJ, Mazziotta JC, Huang HK. Mapping brain function to brain anatomy. In: Sneider R, Dwyer S, Jost R, eds., *SPIE Medical Imaging II* 1988;914:445–451.
14. Zhang J, Levesque ML, Wilson CL, et al. Multimodality imaging of brains structures for stereotactic surgery. *Radiology* 1990;175:435–441.
15. Day RH. *Human perception*. Sydney: John Wiley & Sons; 1969:76–79.
16. Porter T, Duff T. Compositing digital images. *Comput Graphics* 1984;18: 253–259.
17. Valentino DJ, Cutler PD, Mazziotta JC, et al. Volumetric display of brain function and brain anatomy. In: Sneider R, Dwyer S, Jost R, eds., *SPIE Medical Imaging III: Image Capture and Display* 1989;1091:212–220.
18. Ehricke H-H, Laub G. 3D-visualization of intracranial vessels and brain anatomy in magnetic resonance imaging. In: Loew M, ed., *SPIE Medical Imaging IV: Image Processing* 1990;1233:60–66.
19. Evans AC, Marrett S, Collins L, Peters TM. Anatomical-functional correlative analysis of the human brain using three dimensional imaging systems. In: Sneider R, Dwyer S, Jost R, eds., *SPIE Medical Imaging III: Image Processing* 1989;1092:264–274.
20. Evans AC, Marrett S, Torrescorzo J, Ku S, Collins L. MRI-PET correlation in three dimension using a volume-of-interest (VOI) atlas. *J Cereb Blood Flow Metab* 1991;11:A69–A78.
21. Levin DN, Hu X, Tan KK, et al. The brain: integrated three-dimensional display of MR and PET images. *Radiology* 1989;172:783–789.
22. Pelizzari CA, Chen GTY, Spelbring DR, Weichselbaum RR, Chen CT. Accurate three-dimensional registration of CT, PET and MR images of the brain. *J Comput Assist Tomogr* 1989;13:20–26.
23. Fleming JS, Britten AJ, Perring S, Keen AC, Howlett PJ. A general software package for the handling of medical images. *J Med Eng Technol* 1991;15:162–169.
24. Wahl RL, Quint LE, Cieslak RD, et al. "Anametabolic" tumor imaging: fusion of FDG PET with CT or MRI to localize foci of increased activity. *J Nucl Med* 1993;34:1190–1197.
25. Foley JD, Van Dam A, Feiner SK, Hughes JF. *Computer graphics: principles and practice*, 2nd ed., Reading, PA: Addison-Wesley; 1991:592–595, 754–755.
26. Burns B, Shepp BE. Dimensional interactions and the structure of psychological space: the representation of hue, saturation, and brightness. *Percept Psychophys* 1988;43:494–507.
27. Wyszecki G, Stiles WS. *Color science: concepts and methods, quantitative data and formulae*, 2nd ed., New York: John Wiley & Sons; 1982.
28. Levin DN, Pelizzari CA, Chen GTY, Chen C-T, Cooper MD. Retrospective geometric correlation of MR, CT, and PET images. *Radiology* 1988;169: 817–823.
29. Le Grand Y. *Light, colour and vision*, 2nd ed., London: Chapman and Hall; 1968:220–225.
30. Kosslyn SM. The psychology of visual displays. *Invest Radiol* 1989;24:417–419.
31. Anderson JR, Strother SC, Xu X-L, Bonar DC, Rottenberg DA. Error bounds for five image registration techniques based on high resolution MRI. In: Uemura K, Lassen NA, Jones T, Kanno I, eds., *Quantification of brain function: tracer kinetics and image analysis in brain PET*. Amsterdam: Elsevier Science Publishers; 1993:429–435.
32. Pratt WK. *Digital image processing*. New York: John Wiley & Sons; 1978.
33. Green DM, Swets JA. *Signal detection theory and psychophysics*. New York: Wiley; 1966.
34. Haber RN, Hershenson M. *The psychology of visual perception*. New York: Holt, Rinehart and Winston; 1973:143–144.
35. Hemminger BM, Johnson RE, Rolland JP, Muller KE. Perceptual linearization of video display monitors for medical image presentation. In: Kim Y, ed., *SPIE Medical Imaging 1994: Image Capture, Formatting and Display* 1994;2164:222–240.
36. Mendenhall W, Beaver RJ. *Introduction to probability and statistics*. Boston: PWS-Kent Publishing; 1991.
37. Hill CR. Perception and interpretation of images. In: Webb S, ed., *The physics of medical imaging*. Bristol: Hilger; 1988:567–583.