

Advances in the Use of Neurophysiologically-based Fusion for Visualization and Pattern Recognition of Medical Imagery

Mario Aguilar

Knowledge Systems Laboratory
MCIS Department
Jacksonville State University
Jacksonville, AL, U.S.A.
marioa@ksl.jsu.edu

Joshua R. New

Knowledge Systems Laboratory
MCIS Department
Jacksonville State University
Jacksonville, AL, U.S.A.
newj@ksl.jsu.edu

Erion Hasanbelliu

Knowledge Systems Laboratory
MCIS Department
Jacksonville State University
Jacksonville, AL, U.S.A.
erioni@ksl.jsu.edu

Abstract - *The ever increasing number of image modalities available to doctors for diagnosis purposes has established an important need to develop techniques that support work-load reduction and information maximization. To this end, we have improved on an image fusion architecture first developed for night vision applications. This technique, presented at Fusion 2002, utilizes 3D operators to combine volumetric image sets while maximizing information content. In our approach, we have combined the use of image fusion and user-defined pattern recognition within a 3D human-computer interface. Here, we present our latest advances towards enhancing information visualization and supporting pattern recognition. We also report on results of applying image fusion across a variety of patient cases. Finally, we have also begun the assessment of pattern recognition based on 2D vs. 3D fused image features. Initial results indicate an advantage to fusing imagery across all three dimensions so as to take advantage of the volumetric information available in medical data sets. A description of the system and a number of examples will serve to illustrate our ongoing results.*

Keywords: Medical imaging, fusion, pattern recognition, data mining, 3D visualization.

1 Introduction

As powerful imaging techniques grow ever more pervasive, medical experts often find themselves overwhelmed by the large number of images being produced. In addition, many of these images are significantly complementary and often lead to an increase in workload. Experts typically have to view and analyze multiple images which force them to follow a tedious scanning procedure. Inadvertently, this work overload leads to a decrease both in the quantity and quality of care that can be provided.

At Fusion 2001 [1] and 2002 [2], we reported on our extensions of a biologically-based fusion architecture [3] implemented to combine and enhance information content from up to four medical modalities (i.e. MRI, CT, PET,

SPECT, etc.). By providing a single fused image, this architecture can increase user efficiency and facilitate decision making. Past results have demonstrated the potential for user performance improvements and workload reduction.

In parallel to this effort, a learning system is being developed to provide an interface to train the computer to perform automated segmentation and preliminary diagnosis. Instead of focusing our efforts on developing a general learning engine, we are pursuing a targeted learning design whereby users help define the context and constraints of a specific recognition task. This task-specific recognition system can be encapsulated as an autonomous agent that can in turn be used to identify and flag or highlight areas of interest that may be present in other areas of the image, other slices, or in other patients as is the case for large image databases. This system, we hope, could further improve quality of care by aiding in the process of pre-screening, efficient diagnosis, and/or by detecting previously undiagnosed issues in patients' historical files.

The third component of the system, an interactive system for the visualization of fusion and pattern recognition results, was developed for exploring the results in both two and three dimensions. This system utilizes the inherent three-dimensional information of the original imagery to create volumes for a more intuitive presentation and interaction with the user. This allows for localization of the task-relevant image features and spatial context.

1.1 Neurophysiologically-motivated fusion

The visual system of primates and humans contains three types of light sensors known as cones, which have overlapping sensitivities (short, medium, and long wavelengths). It is through the combination of these three sensor sets that we obtain our perception of color. Circuitry in the retina is functionally divided into two stages. The first one utilizes non-linear neural activations and lateral inhibition within bands to enhance and normalize the inputs. The second stage utilizes similar neural components in an arrangement of connections that

lead to between-band competition that in turn produces a number of combinations of the three original bands [4]. This last stage of processing enhances the complementary information that exists in each of the bands (e.g. a spectral decorrelation operation).

The fusion architecture is motivated by this basic connectivity. Each of the processing stages is implemented via a non-linear neural network known as the *shunt* (described in section 2.1) [5]. The resulting non-linear combinations lead to information decorrelation not unlike what is usually targeted by principal component analysis techniques. However, in neuro-physiological systems and in our architecture, the non-linear operator has a very narrow spatial window providing a better-tuned decorrelation. In addition, the operator is modulated by more globally defined statistical characteristics of the input that produce normalization, smoothing, and between-band calibration.

2 Methods

For this investigation, we used both morphological and functional imaging modalities [6] obtained from the Whole Brain Atlas (WBA) web site [7]. The imagery obtained was spatially registered across modalities by the original authors of the WBA. Here, the imagery obtained included functional MRI and SPECT modalities. Furthermore, the MRI imagery included three different modalities (PD, T1, and T2) that provide distinct information content (see [1] for details.) In total, there were four different modalities available for each image slice in the data set.

The images and modalities were chosen so that the fused images would contain vital information of both a morphological and functional nature; however, the system can be easily applied to any combination of imaging modalities. Functional modalities have traditionally been used for detecting the presence or progress of brain defects such as tumors, while structural modalities allow for localization of these defects.

An interim file format was developed which stores the relevant information for a given image set such as case number, image sizes, number of slices, case description, etc. By doing so, the user can easily access the desired case and obtain all vital information associated with the patient. To increase graphical update performance, all slices are loaded into memory when the case of interest is opened. We envision utilizing this format as a necessary stage towards integrating the fusion and learning system with medical imagery database systems.

The system is implemented as a pipeline of three main stages: fusion, learning, and exploration. Each stage is implemented as an independent interface mode in a graphical user-interface (GUI) (e.g. a tabbed window) to assist the user in understanding and organizing workflow. The processing sequence is implied by the order in which the modes are enabled. This follows a logical order by

allowing fusion of original image modalities, computer segmentation and identification of the fused results, and then exploration of the learned results only in that order. A thorough discussion of the GUI is outside of the scope of this paper. In the following subsections, we describe progress we have made in the development of image fusion, pattern recognition, and visualization techniques.

2.1 Image Fusion

As discussed previously and in [1], modern medical imaging techniques such as MRI and SPECT produce multi-slice image volumes. For this reason, we have extended the shunting operator described in [5] to three dimensions in order to account for information content between slices. Furthermore, it is desirable to process such image sets and perform image fusion while exploiting the information present in the third dimension. The modified shunting equation is as follows:

$$\dot{p}_{ijk} = -Ap_{ijk} + (B - p_{ijk})C[G_c * I^C]_{ijk} - (D + p_{ijk})[G_s * I^S]_{ijk} \quad (1)$$

Here, the activity of cell p (i.e. the resulting pixel value) is defined in terms of the contrast enhancement and dynamic-range normalization defined by a three-dimensional neighborhood around the cell. In this equation, we introduce two Gaussian terms: the first one modulating the excitatory inputs to cell p_{ijk} , and the second one modulating its inhibitory inputs. This extension required the use of a three-dimensional Gaussian kernel as follows:

$$G(x, y, z) = \frac{1}{4\pi^2 \sigma^3} e^{-\frac{(x^2 + y^2 + z^2)}{2\sigma^2}} \quad (2)$$

The combination of both kernels (G_c and G_s) in equation 1 leads to a three-dimensional operator that we represent as an icon of concentric spherical volumes. In the figures to follow, the differentiation of color of each spheroid is used to denote the positive/negative relationship of the center and surround areas of the kernel (i.e. signal polarity). In addition, the particular hue indicates the source of its input (i.e. which modality).

The basic fusion architecture consists of two distinct processing stages. In the first one, as in the retina [4], we utilize a shunting operation to obtain within-band image enhancement and normalization. This produces contrast enhancement, dynamic range calibration, and normalization of input images. In this version however, the input to the 3D shunting operator is a stack of image slices. Here, one slice is processed at a time with the shunt operating on the current slice as well as those slices immediately above and below. The result is a slice with enhanced contrast characteristics due to between-slice context information.

The second stage adopts the use of the same non-linear neural network operator to produce between-band decorrelation, information enhancement, and fusion. These stages for the two-band fusion of MRI are illustrated in Figure 1.

The shunt combinations of the second stage, as shown in Figure 1, provide two unique sets of information-rich images. These combinations perform paired decorrelations between the input image modalities. For instance, in the case of the top combination (red-center and green-surround in Fig. 1) the result is the decorrelation of the first modality (T1) from the second modality (T2). In other words, it enhances information that is present in band 1 but not in band 2. The resulting images from both decorrelations are then mapped to the red and blue channels of the color display. The enhanced T1 image was chosen as the input to the green channel in order to highlight information from that modality.

Another processing stage may be introduced prior to producing the final color image that remaps the color assignments from those derived by the fusion process. Here, a mapping from RGB to HSV space allows the user to manipulate the appearance of the image (e.g. hue remap) to obtain a more pleasing coloring scheme. The modified HSV values are mapped back to RGB to be used in generating the final color fused image.

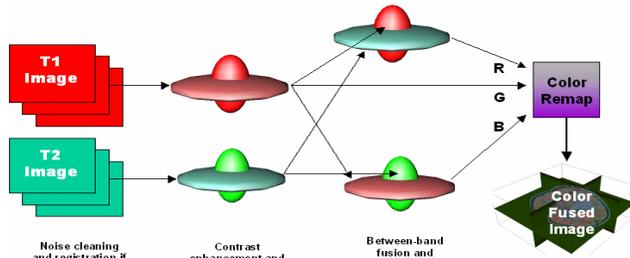


Figure 1. Two-band fusion architecture used for processing functional MRI imagery.

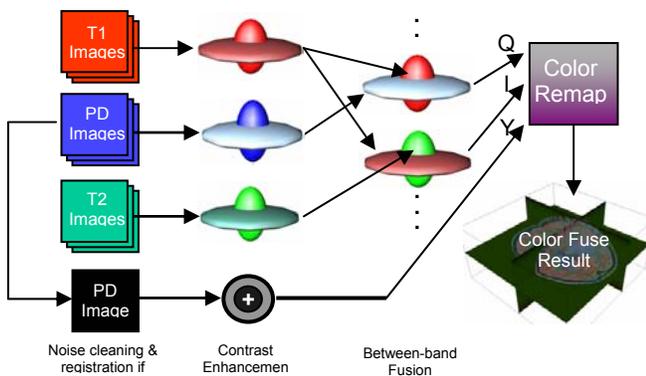


Figure 2. Three-band MRI fusion architecture.

The architectures for three- and four-bands have undergone a significant change. Specifically, we made the observation that while the 2D fusion process did not provide a sufficiently differentiated fusion result, due to

its more limited convolution kernels, it does preserve a significant amount of high-frequency information. This difference becomes quite significant when it comes to enhancement of structural information. On the other hand, we find that the decorrelation and information enhancement of the 3D fusion architecture is crucial to the quality of the final fusion results. Hence, in order to capitalize on the benefits of both of these strategies, we have created a hybrid fusion system in which all of the multi-modal decorrelation and information enhancement is performed via the 3D shunt kernels. In addition, the brightness channel of the final color fused image is defined by the 2D shunt enhanced imagery. We present the specific connectivity in the following two paragraphs.

The 3-band fusion architecture as modified from [1] is illustrated in Figure 2. Here, the first stage of processing is as before, where each of the input bands is separately contrast enhanced and normalized. Then, two between-band shunting operations produce distinct fusion products. The first one decorrelates the information between bands 1 (T1) and 2 (PD). The second does it for bands 3 (T2) and 2. The resulting fused images are then mapped to the I and Q (also known as red-green and blue-yellow) components of the YIQ color space of the image. The Y or achromatic component is derived from the 2D enhanced PD image (notice the 2D kernel icon in Fig. 2) that provides the most faithful structural details.

The architecture for four-band fusion between MRI and SPECT is shown in Figure 3. Here, the second stage of processing produces the decorrelation between T1-weighted and SPECT, as well as between T2-weighted and SPECT. Notice that the decorrelation is done in both directions for each of the pairs. In fact, what is not shown is that at the second stage, 12 different pairing are produced (all combinations between the 4 modalities). While only two of these are used here to produce the color image, the remaining pairs are used in conjunction with the learning system. There, all 12 combinations are exploited as features from which the learning system can extract recognition patterns. Finally as in the 3-band system, to preserve the high resolution of the 2D imagery, the 2D shunted PD image is mapped to the Y, or achromatic channel.

2.2 Image Fusion Visualization

One of the most important issues that arises in applications of medical image processing is the availability of effective visualization tools. In fact, in addition to the obvious difficulty of managing multi-modality visualization, it is also necessary to devise

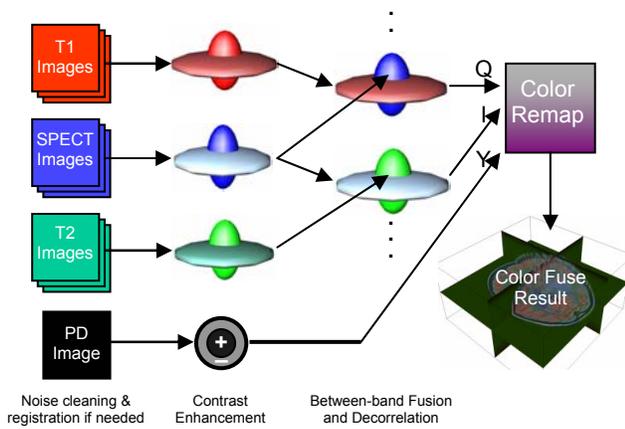


Figure 3. Four-band MRI/SPECT fusion architecture.

methods for visualizing the volumetric information provided by multi-slice image modalities.

Most medical image modalities are inherently 3-dimensional. In particular, these modalities are recorded as multi-plane (slice) images across the organ of interest. While it is natural to associate the individual images as being the main tool used by radiologists in diagnosis, it is also true that medical specialists rely heavily on the information available across adjacent slices. Hence, it is routine practice for radiologists to utilize multi-slice arrays in order to gain volumetric contextual information and assess the extent of features in all three dimensions.

In an effort to present the fusion results in an intuitive interface, a cranial volume was computed from the PD modality of the patient. This volume can be created from any modality, but PD is most useful since it displays high-resolution structural information. By embedding the fusion imagery into the generated volume, contextual navigation of the fusion stage output is made possible. The user can manipulate the viewable slices of the fused imagery via a slider. To increase slice-change performance, all slices are loaded into memory when the window opens and only the slice of interest is rendered within the volume. We present a detailed description and sample screenshots in section 3.1.

2.3 Contextual Zoom Visualization

The use of a zoom tool for detailed analysis of imagery is now common place in imaging applications. There are two important issues of concern when introducing a zoom window to the interface. The first is what is known as the screen real estate problem [8]. This refers to the fact that in order to introduce the zoom window, a portion of the computer screen has to be allocated to this window at the expense of other important areas of the interface. The second issue arises when the solution to screen area allocation is resolved by superimposing the zoom window on top of the main image. Figure 5a and 5b present this approach. In Fig. 5a, an AOI has been outlined, and in Fig. 5b a zoom window of that area has been super-

imposed on the original image. The consequence of this technique is that important areas of the image become occluded. Finally, both of these approaches reduce efficiency as they force the user to refocus to different areas of the screen and maintain a mental record in order to get the “full picture”.

These issues are of great relevance to the information fusion community because their solution can address the fact that while computers can store and manipulate vast amounts of information they have a comparatively small screen on which this data can be viewed. Sheelagh Cpendale [8] proposed a solution to these issues with what she terms Elastic Presentation Space. In this technique (Fig. 5c), the image is warped so as to obtain an effective zoom of the AOI while also producing a smooth transition to the non-zoomed area. In our system, we utilized a software product developed by IDELIX Inc. [9] which readily generates a variety of useful warping methods and GUI elements. In section 3.3, we present examples illustrating our modifications to the interface that accommodate user interaction with the pattern recognition tool in order to train the system.

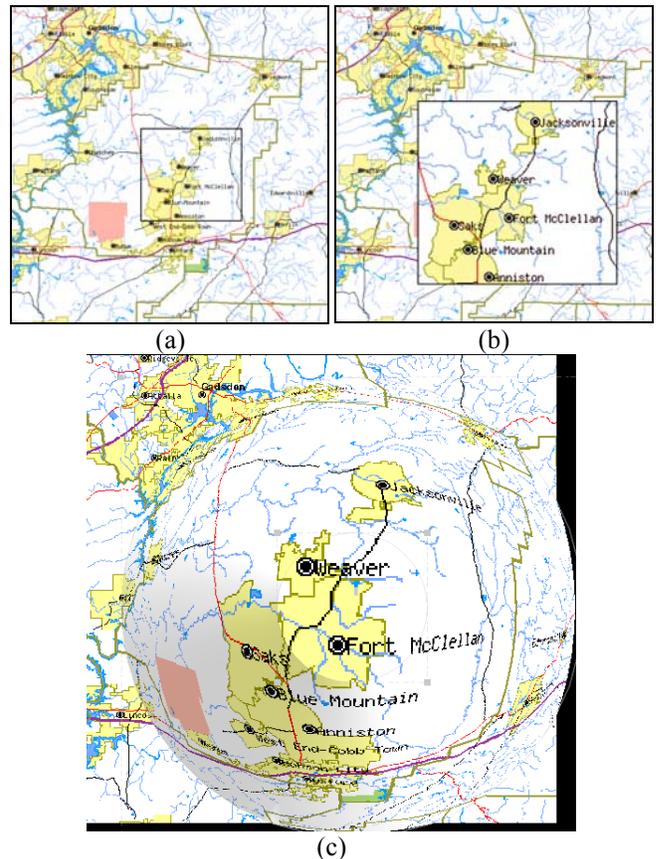


Figure 4. Comparison of zooming-in-place techniques. (a) original image with outline around AOI. (b) Typical zoom-in-place approach leading to occlusion. (c) Contextual zoom technique preserves all information in the original image.

2.4 Fusion-based Pattern Recognition

The purpose of this stage of the system is to use principles of image mining and supervised learning to enable the computer to identify objects in the imagery. The system is based on the proposed approach in [10] and later implemented in the context of multi-sensor fusion in [11], [12] and [13]. This is accomplished by allowing the user to define features of interest by selecting representative pixels on the imagery. This information is then utilized by a machine-learning system to establish an appropriate mapping between inputs and desired landmark identity. In other words, the system establishes input-based recognition patterns that uniquely identify the characteristics of interest to the user. A GUI is implemented to allow the expert to mark pixels as targets or non-targets for recognition/segmentation. By doing this, the learning system can leverage off the expert to solve the specific task of identifying the sampled structure across all available modalities. As in [11]-[13], this information is used to train an ART-based supervised neural network [14] known as Fuzzy ARTMAP [15] to recognize or segment the image based on the user's input. In addition, the resulting train network can be encapsulated and saved as an agent which can later be loaded to pre-screen images by highlighting areas of potential interest for the user.

One issue inherent in the mapping derived by Fuzzy ARTMAP is its dependence on the ordering of the training inputs. That is, the order in which inputs are presented greatly affects the recognition results of the system. This produces inconsistencies when attempting to replicate results through a graphical interface as described here. On the other hand, Fuzzy ARTMAP's output is limited to an all-or-none form which limits the amount of information that the user is provided regarding the confidence of the results. To ameliorate this last problem, the authors in [12] and [13] proposed a combined system where five separate Fuzzy ARTMAP systems would be trained simultaneously on separate orderings of the inputs. The combined system is used in a voting scheme to establish the level of confidence from 1 to 5 and used to define the final output of the system. As a consequence of this design, the former problem of order dependence is also addressed as the broader learning sample helps to capture the variability across input orderings.

In the system, the learning module trains to discriminate pixels as belonging to either a target or non-target class. For this reason, we utilized 16 inputs representing the features of a given pixel as derived from the 4-band fusion architecture. Of these, the first four are the 3D shunt-enhanced single modality images while the remaining twelve correspond to the possible combinations of the four original bands.

An additional issue of concern that affects all learning systems is their dependence on well selected set of operating parameters. For Fuzzy ARTMAP, this problem is lessened by having only one main parameter which

affects its performance. This parameter, known as vigilance, defines a threshold level for assigning inputs to a given category. While it is only one parameter, in the context of its target user base, it does remain difficult to define what the most appropriate value should be. We have attempted to address this issue by modifying the proposal in [11]-[13] so as to break the symmetry between the voters in the combined system. In particular, we train the first three voters with different ordering and the same vigilance parameter. The remaining voters are trained with the same ordering as chosen for the third voter but we modified the vigilance parameter of each to be 10% lower and 10% higher than the first three respectively. Hence, the original voting scheme which had five homogenous voters is replaced by a heterogeneous arrangement that increases the breadth of the learning process.

3 Results

3.1 3D Fusion Results

At Fusion 2002, we presented preliminary results that compared the 2D vs. 3D shunting kernels. In those results, we reported on the qualitative advantages of the 3D fusion approach. In particular, we presented comparisons of fusing 2-, 3-, and 4-band imagery with significant improvements on the results when the fusion process included decorrelation operations across image slices.

Due to the success of those results, we proceeded to assess their generality across multiple patient cases. The most important finding was that regardless of the condition of the original imagery or the type of case involved, the fusion results remain consistent in both quality and color mapping. This is relevant as it suggests the ability of relying on consistency to understand the information provided by the fusion process.

An illustration of 4-band fusion is presented in Figure . Here, Figure 6a presents the shunt-enhanced PD, T1, T2 and SPECT images (in clockwise order). Figure 6b and 6c provide a comparison of the 2D vs. 3D fusion results. In 6c, we can once again see the advantages of 3D in supporting improved color fusion results. On the other hand, we found that for some cases like this one, the 2D shunting results preserved a significant amount of high frequency information not present in the 3D results (Fig. 6b). For this reason, a hybrid approach was adopted where the brightness channel of the final image was replaced with the 2D shunt-enhanced PD image. Figure 6d presents the result of fusing the same case with the hybrid architecture. It is clear that the end result captures the information provided by the separate approaches.

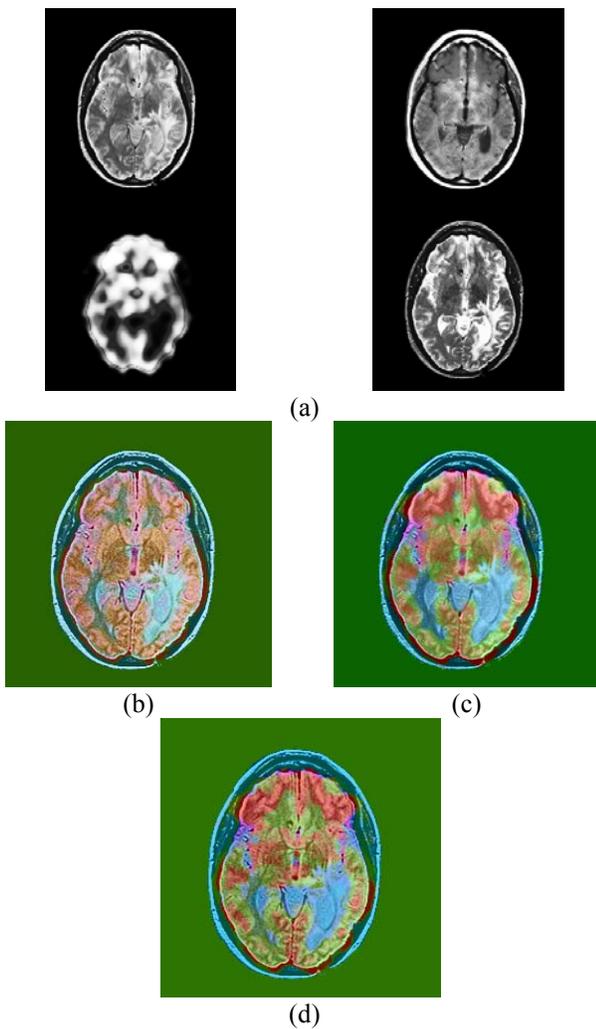


Figure 5. Fusion results and comparison across approaches. (a) original imagery. (b) 2D fused. (c) 3D fused. (d) hybrid fusion results.

3.2 Image Fusion Visualization

The visualization stage, as can be seen in Figure 6, is used to display a multitude of information in a powerful and intuitive interface. Group A contains the generated skull with the slices of interest embedded within the skull for contextual navigation. These embedded slices consist of the selected fusion image combined with a transparent SPECT overlay which emphasizes the crucial metabolic information of that modality. Also, a transparent overlay of highlighted regions is used to help the user identify the results that the learning agent has produced. This group can intuitively be rotated, zoomed, and panned with a 3-button mouse. Group B contains the usual slice slider which is used to update groups A, E, and F. Group C consists of the SPECT transparency slider. This slider allows customized view of the SPECT overlay in group A so that the desired amount of metabolic information may be displayed. Group D consists of our viewing options for the exploration stage. Here the user may select the desired fusion result for display in groups A and F, remove the

skull if it is not needed or hinders the view, remove the raw images in group E, or remove the SPECT overlay in group A. Group E consists of the original images for the slice of interest. These images are selectively displayed to permit referring back to the original information used in obtaining the fused result. Group F consists of the fusion slice of interest in the middle, as well as, the slices above and below which help the user identify the extent and characteristics of AOIs.

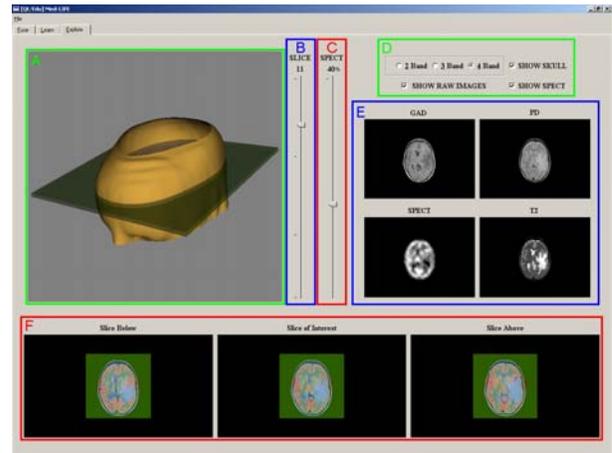


Figure 6. Visualization interface.

3.3 Contextual Zoom

As explained in a previous section, the contextual zoom functionality was introduced to address the important issues of screen real estate and unhindered task performance. While the zooming tool is useful for analysis and image understanding, the most important role of this tool in the system is in supporting user interaction with the learning system. Specifically, the user must be able to select targets and non-targets in order to train the learning system. Because the features are defined at the pixel level, it is crucial that the user be able to unambiguously “mark” the relevant areas. Figure 7 presents the resulting contextual-zoom window and an example of marking targets (in green) and non-targets (in red) on one of the image slices.

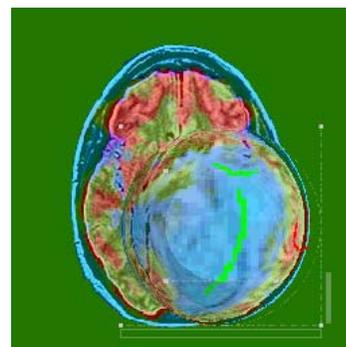
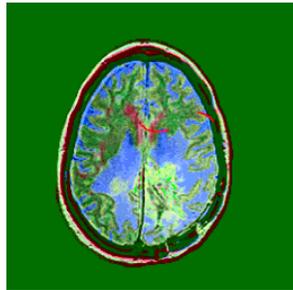


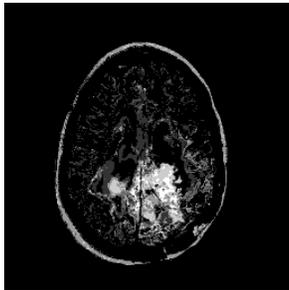
Figure 7. Contextual zoom in the learning interface.

3.4 Pattern Recognition

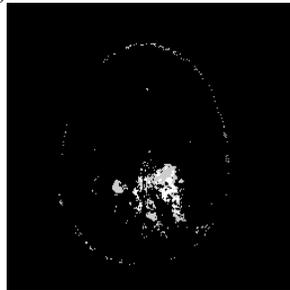
Figure 8 presents results of training the system to identify cortical areas affected by an anaplastic astrocytoma. The areas selected as targets and non-targets are indicated in green and red respectively (short colored lines in Fig. 8a). The selection was made based on expert assessment presented in [7].



(a)

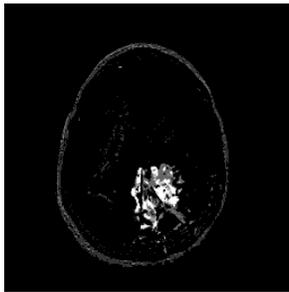


(b)

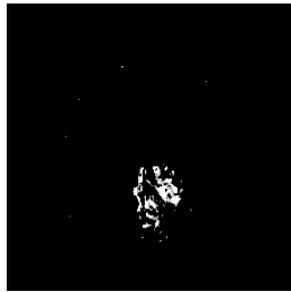


(c)

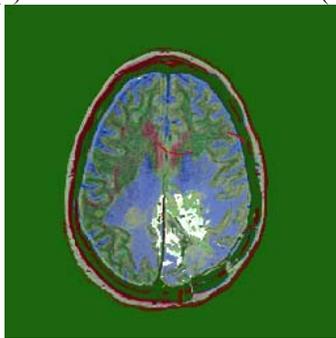
Figure 8. Segmentation of parietal lesion caused by an anaplastic astrocytoma. See text for details.



(a)

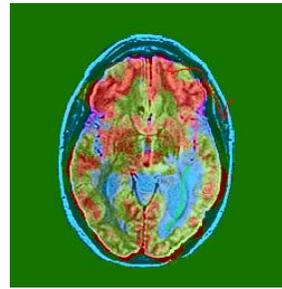


(b)

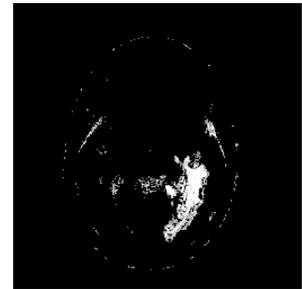


(c)

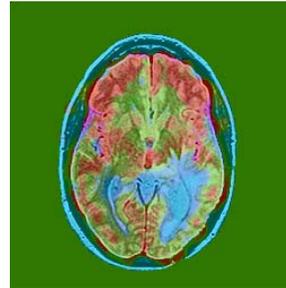
Figure 9. Segmentation as in Fig. 8 using heterogeneous voting system. See text for details.



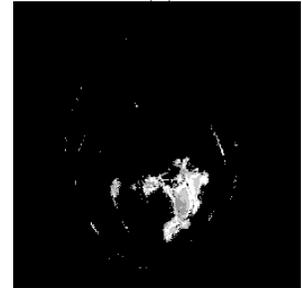
(a)



(b)



(c)



(d)

Figure 10. Cross-slice generalization results. (a) Training inputs (marking on slice 26). (b) Segmentation results (slice 26). (c) Generalization test input (slice 27). (d) Segmentation results (slice 27)

The results of applying the learn patterns across the entire image are shown in Figure 8b. Here, the grayscale value represents the number of voters that identified the given pixel as belonging to the “target” class. In Figure 8c, a threshold has been applied to require high confidence before highlighting a given area. In this case, the threshold is such that 4 out of the 5 voters are required in order to flag a given pixel as a candidate.

As described in section 2.4, we implemented a heterogeneous voting approach where two out of the five voters were trained with modified vigilance parameters. Figure 9 presents the results obtained in segmenting the case shown in Figure 8. Here, Fig. 9a presents the raw output of the system before thresholding for confidence. Fig. 9b presents the final segmentation result. Here, the results are greatly improved eliminating all false alarms generated by the homogenous voting system (Fig. 8c). Figure 9c provides an example of the layering of fusion and pattern recognition results to enhance visualization.

An important benchmark for a learning system is its ability to generalize from the training data. While the previous results attest to this capability given that the training set was a very small portion of the entire slice, it is also important to assess generalization across slices. Figure 10 presents some of the results obtained when the system was trained and tested on a single slice and then used to recognize similar areas on a new slice. Figure 10a illustrates the areas selected for training. Here, target areas are those associated with a cerebral edema in the left hemisphere. The segmentation results on that slice are shown in Figure 10b. To test the generalization abilities of the system, a new slice was provided to the system.

Figure 10c presents the color fused results for the slice immediately above the one used for training. The results of segmenting on this new slice are shown in Figure 10d. As can be seen, the system performs very well in segmenting all of the non-target areas while producing few false alarms (areas around the skull). Ongoing efforts are aimed at reducing these false alarms by utilizing grouping techniques and eliminating isolated false alarms.

4 Summary

We have extended our fusion architecture to address the advantages of both 2D image enhancement and 3D fusion. Our results compare favorably against the results of the original 2D and 3D fusion architectures. Better definition of image details and volumetric information make this approach a significant advance in medical image fusion. We presented a modified voting scheme for combined Fuzzy ARTMAP systems that improves segmentation results. We have initial evidence for robust generalization across slices. In future work, it will be important to validate these findings and investigate generalization across patients.

The pipelined system promises to be a powerful tool in leveraging professional expertise for use with computational techniques. The ease of use and capabilities of this system may provide an essential step towards providing a solution to ease the workload of medical personnel, thereby increasing the quantity and quality of medical care.

5 Acknowledgements

The authors would like to thank IDELIX for their guidance and assistance in the deployment of the PDT software tools. This work was supported by a Faculty Research Grant awarded to the first author by the faculty research committee and Jacksonville State University. Opinions, interpretations, and conclusions are those of the authors and not necessarily endorsed by the committee or Jacksonville State University.

References

[1] M. Aguilar and J.R. New, "Fusion of Multi-Modality Volumetric Medical Imagery", *Proc. of the Fifth International Conference on Information Fusion*, Baltimore, MD, 2002.

[2] M. Aguilar and A.L. Garrett, "Neuro-physiologically-motivated sensor fusion for visualization and characterization of medical imagery", *Proc. of the Fourth International Conference on Information Fusion*, Montreal, Canada, 2001.

[3] M. Aguilar, D.A. Fay, W.D. Ross, A.M. Waxman, D.B. Ireland, and J.P. Racamato, "Real-time fusion of low-light CCD and uncooled IR imagery for color night

vision", *Proc. of SPIE Conf. on Enhanced and Synthetic Vision*, **3364**, 1998.

[4] P. Schiller and N.K. Logothetis, "The color-opponent and broad-band channels of the primate visual system", *Trends in Neuroscience*, **13**, pp.392-398, 1990.

[5] S.A. Ellias and S. Grossberg, "Pattern formation, contrast control, and oscillations in the short memory of shunting on-center off-surround networks", *Biological Cybernetics*, **20**, pp. 69-98, 1975.

[6] M. Aguilar and A.L. Garrett, "Biologically-based sensor fusion for medical imaging." In *Proceedings of SPIE Conference on Sensor Fusion Architectures, Algorithms, and Applications V*, **4385**, 2001.

[7] K.A. Johnson and J.A. Becker, Whole Brain Atlas, <http://www.med.harvard.edu/AANLIB/home.html>, 1999.

[8] M.S.T. Carpendale, "A framework for elastic presentation space", School of Computing Science, Burnaby, BC, Simon Fraser University: 276, 1999.

[9] IDELIX, "PDT3D – Pliable Display Technology – SDK", Vancouver, BC, IDELIX Software Inc., 2003.

[10] T.L. Delanoy, "Toolkit for image mining: User-trainable search tools.", *Lincoln Laboratory Journal*, **8**(2), pp. 145-160, 1995.

[11] W.D. Ross, A.M. Waxman, W.W. Streilein, M. Aguilar, J. Verly, F. Liu, M.I. Braun, P. Harmon, and S. Rak, "Multi-sensor 3D image fusion and interactive search", In *Proc. of 3rd International Conf. on Information Fusion*, Paris, France, 2000.

[12] W. Streilein, A. Waxman, W. Ross, F. Liu, M. Braun, D. Fay, P. Harmon, and C.H. Read, "Fused multi-sensor image mining for feature foundation data", In *Proc. of 3rd International Conf. on Information Fusion*, Paris, France, 2000.

[13] W. Streilein, A. Waxman, W. Ross, M. Braun, F. Liu, and J. Verly, "An interactive mining tool utilizing an extension of Fuzzy ARTMAP for efficient exploitation and enhanced visualization of multi-sensor imagery", In *Proc. of the 4th International Conf. on Cognitive and Neural Systems*, Boston, MA, 2000.

[14] S. Grossberg and G. Carpenter, "The ART of Adaptive Pattern Recognition by a Self-Organizing Neural Network" In *Computer*, March 1988.

[15] G.A. Carpenter, S. Grossberg, J.H. Markuson, J.H. Reynolds, and D.B. Rosen, "Fuzzy ARTMAP: A neural architecture for incremental supervised learning of analog multidimensional maps. *IEEE Transactions on Neural Networks*, **3**(5), 698-713, 1992.