Heterogeneous ARTMAPs for Image Segmentation

Joshua R. New
Knowledge Systems Laboratory
MCIS Department
Jacksonville State University
Jacksonville, AL, U.S.A.
newj@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 a system for use in image segmentation that will allow radiologists to train the computer to recognize the presence of abnormal brain tissue present in Magnetic Resonance Images (MRIs). A voting scheme of ART neural networks has been developed that varies learning system parameters in order to provide robust segmentation quality. In this effort, we will also discuss a 3D shunting operator used to combine volumetric image sets while maximizing information content, thus aiding in image segmentation by enhancing the quality of the features. We also report on results of applying image segmentation across a variety of patient cases. A description of the system and a number of examples will serve to illustrate our ongoing results.*

Keywords: Medical imaging, pattern recognition, data mining.

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.

In order to reduce this overload, a learning system is being developed to provide an interface to train the computer to perform automated segmentation and preliminary diagnosis for MRIs. 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.

In order to provide this learning system with adequate features by which to accurately and robustly segment MRIs, image processing techniques have been applied to maximize the information present in a patient's MRIs. The neurophysiologically-motivated image processing technique provides the machine learning system with a wealth of information on the relationships between images which can then be utilized by any data mining or pattern recognition system for the task of segmentation.

2 Segmentation

Segmentation is defined as the act or process of dividing into segments, specifically in biology. In this study, we are interested in the segmentation of abnormal tissue from MRIs, which is typically performed by radiologists. With pattern recognition systems, it is possible to use automated processing techniques to segment areas of abnormal tissue, thus directing the attention of the radiologist. This will aid in the process of prescreening, efficient diagnosis, and detecting previously undiagnosed issues in patients' historical files.

There are many different approaches available for image segmentation. At the simplest level, there are a number of properties per pixel called "features". When combined together to codify all knowledge for a pixel, the string is called a "feature vector". For example, the pixel (0,0) of an image is typically described by a feature vector consisting of three features: a red value, a green value, and a blue value.

Segmentation methods use pattern recognition to establish which feature vectors are similar to one another, thus "grouping" or "clustering" similar vectors into the same category. In the next section, we will review the general types of machine learning systems available for determining the similarity of pixels in an image.

3 Machine Learning Systems

Machine learning is defined as the ability of a machine to improve its performance based on previous results. In

this study, the machine is the segmentation system and its performance is the quality of its segmentation. Machine learning often occurs through a system that can modify itself so that on a subsequent execution, even with the same input, a better output is produced. There are three very general categories of learning that is common in today's machine learning systems; they are unsupervised, reinforced, and supervised learning.

3.1 Unsupervised Learning

In unsupervised learning, as its name implies, the learning system is not provided a "supervisory signal" which defines the correct answer/classification for the given input/feature vector. Typically, such systems take all input vectors and determine the similarity among them and then assign similar vectors to the same group. The "similarity" parameter is typically provided by the user and can be changed to generate more groupings of very similar pixels or fewer, more general groupings. In essence, unsupervised systems are a mapping from inputs to "categories" which are effectively labels for each category's constituent vectors.

As an example of unsupervised learning for image segmentation, let's say we have an image of different types of fruit on a table. Let each pixel of the image be a feature vector which consists of three features – the pixel's red value, blue value, and green value. When the image is passed to an unsupervised learning system, the pixels which constitute a typical apple will be very different from those that constitute a pear, since apples are red while pears are very green. In an image of fruit, the system should be able to distinguish between pears and apples; that is, to segment them from one another. In the same way that "redness" denotes apple and "greenness" denotes pear, so also can "MR-T1ness" denote fat and "MR-T2ness" denote water.

There are several unsupervised machine learning systems in use today. Some such systems are k-means systems [1], self-organizing maps [2], and ART neural networks. While ART will be discussed in some detail in the next chapter, an analysis of the other systems is beyond the scope of this paper. It is interesting to note that unsupervised systems are arguably not machine learning systems since no input is given to them so that they may improve their performance on subsequent classifications.

3.2 Reinforcement Learning

In reinforced learning, the learning system is not provided the correct answer/classification, but is instead given a "reward function". This reward function allows a more general measure of "good" and "bad" performance to be provided to the system. From this function, reinforcement learning can not calculate error between the actual output and the correct output; however, the system can decrease the "cluster size" for "bad" classifications in order to increase performance. This is common in many real-world applications in which the correct answer is not precisely known, but behaviors or actions which lead to failure are realizable.

As an example, consider the same example as before: a picture of a bowl of fruit, except this time we'll make random pixels semi-transparent. This time each pixel's feature vector consists of four features – red, green, blue, and alpha. Alpha is a measure of the opacity of an image, being completely opaque at, say, 255 and completely transparent at 0. If the image were passed to an unsupervised learning system, the system may likely assign several pixels to a category that corresponds to "transparent pixel". When the image is passed to a reinforcement learning system, we can notify the system it is incorrect for labeling a pixel as a "transparent pixel" since transparency doesn't change what the image is of but simply how much of it is visible. In this way, the system will adapt to decrease the size/importance of the "transparent pixel" category, thus establishing that the alpha feature is a confounding/inessential variable.

3.3 Supervised Learning

In supervised learning, the learning system is provided the correct answer/classification. That is, the system is not only given a series of input vectors to classify, but the correct outputs for those input vectors as well (aka. the "supervisory signal"). One may ask, if you have the correct output for the data, why do you need to classify the inputs anyway? Well, there are two phases to supervised learning systems. In the first, the supervised data set is sent to the system for "training". Once the system has been trained, it can then be used to "classify" new data that doesn't have a corresponding supervisory signal.

In the supervised framework we pursue, expert radiologists can "segment" a very small portion of an image with a few swipes of the mouse. The supervised learning system can then be trained on these inputs and used to classify the rest of the image. In this way, the expert can interactively guide the segmentation of the learning system to classify only objects/areas of interest. It is interesting to note that later, the learning system will in turn guide the radiologist's attention by highlighting areas of interest in an image.

Obviously, supervised learning is the most powerful of the three types of machine learning systems due to the foreknowledge of the correct classification, even if only for a small subset of the data. This knowledge allows the learning system to compute an exact output error and correct itself during training. There are several supervised machine learning systems in use today. Some such systems are Kalman filters [3], decision trees [4], and neural networks [5].

It is important to note that an accurate supervisory signal be provided or else the errors will be reflected in the subsequent classification by the supervised learning system. Since radiologist's are trained for image

segmentation (under a different name for sure), and arguably define what is correct (at least for their patients), it is safe to rely upon the supervisory signal provided by a trained radiologist. Furthermore, this method promises to save radiologists time by codifying their knowledge into a machine learning "agent", thus pushing some of their responsibilities onto the computer.

4 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 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. In total, there were four different modalities available for each image slice in the data set. These images provide vital information of both a functional and morphological nature. Functional modalities have traditionally been used for detecting the presence or progress of brain defects such as tumors, while morphological (or structural) modalities allow for localization of these defects.

The four modalities per slice are then passed to an image feature creation system. This system is based on the visual system of primates and is modeled by operations of the shunt operator, which has been extended to take into account the volumetric nature of MRIs. The current system produces 16 grayscale images from the four original modalities. These images are then used as input into a machine learning system.

The machine learning system created is a supervised system that uses ARTMAP neural networks for quick training and classification times. A supervisory system was utilized in order to allow radiologists to train the computer system to recognize abnormal brain tissue present in MRIs. ARTMAP neural networks are supervised, online/interactive learning systems that allow training in seconds (rather than days for back-propagation neural networks) and immediate feedback to the radiologist for correction or refinement.

4.1 Image Feature Creation

The image processing technique has been presented in the context of image fusion at Fusion 2001 [8], 2002 [9], and 2003 [10], where we reported on our extensions of a biologically-based fusion architecture [11] implemented to enhance information content from up to four medical modalities (i.e. MRI, CT, PET, SPECT, etc.). By providing a multitude of other images, this architecture can increase the capabilities of a machine learning system to delineate tissue types.

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. This last stage of processing enhances the complementary information that exists in each of the bands (e.g. a spectral decorrelation operation).

This connectivity is modeled by processing stages via a non-linear neural network known as the *shunt* [12]. 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.

As discussed previously, 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 [12] 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:

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

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 retinal model proposed in [11] consists of two distinct processing stages. In the first one, 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 and information enhancement. These stages for the two-band processing 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 Figure 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 can then operate as information-rich features for a machine learning system.

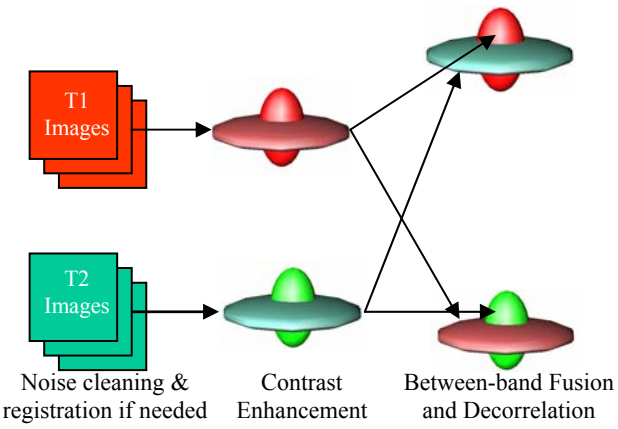

Figure 1. Two-band architecture used for processing functional MRI imagery.

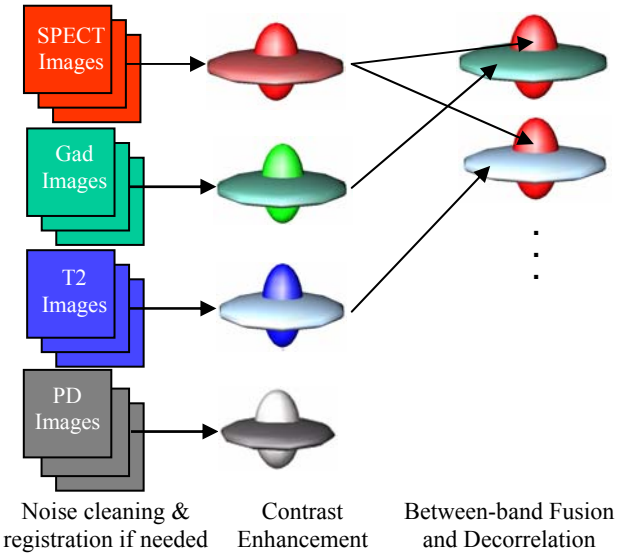

Figure 2. Three-band MRI fusion architecture.

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 [9] 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 image products. The first one decorrelates the information between bands 1 (SPECT) and 2 (Gad). The second does it for bands 1 (SPECT) and 3 (T2). This process is then repeated for single-opponent results between all modalities, producing twelve distinct images (combinations between the 4 modalities) even though only two are shown in Figure 2.

At this point, there are 12 single-opponent results and the 4 shunted originals, thus amounting to 16 grayscale images. These images correspond to the 16 features for each pixel that can be used for recognition by a machine learning system.

4.2 ARTMAP Machine Learning

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 [13] and later implemented in the context of multi-sensor fusion in [14], [15] and [16]. 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 contextual-zoom GUI is implemented to allow the expert to easily 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 [14]-[16], this information is used to train an ART-based supervised neural network [17] known as Fuzzy ARTMAP [18] to recognize or segment the image based on the user's input. This Fuzzy ARTMAP system is implemented using the simplified learning rule in order to allow fast training and classification performance. 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 [15] and [16] 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 it 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 [14]-[16] 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.

5 Results

5.1 Image Feature Creation Results

At Fusion 2002 [9], we presented preliminary results that compared the 2D vs. 3D shunting kernels. In those results, we reported on the qualitative advantages of the approach. In particular, we presented comparisons of 2-, 3-, and 4-band imagery with significant improvements on the results when the 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 processing results remain consistent in regards to quality. This is relevant as it suggests the ability of relying on consistency to understand the information utilizing a machine learning system.

As an illustration of the improvements of the 3D shunting kernel over the 2D shunting operator, an illustration of 4-band process is presented in Figure 3. Here, Figure 3a presents the shunt-enhanced PD, T1, T2 and SPECT images (in clockwise order). Figure 3b and 3c provide a comparison of the 2D vs. 3D process results when mapping three of the 16 images to the visible spectrum. In Figure 3c, we can once again see the advantages of 3D in supporting improved information and contrast.

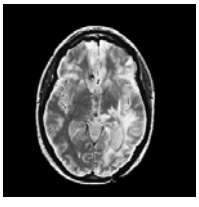

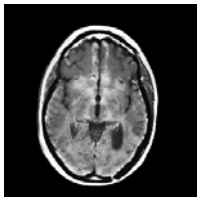

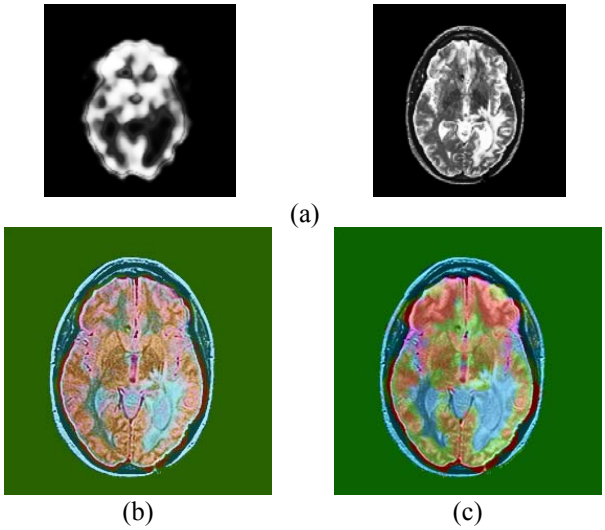

(a)

(b) (c)

Figure 3. Processing results and comparison.
(a) original imagery. (b) 2D fused. (c) 3D fused.

5.2 ARTMAP Machine Learning Results

As explained in a previous section, a supervised machine learning system is the most powerful type when an appropriate "supervisory signal" can be provided. In this system, radiologists define the supervisory signal by making targets and non-targets in an image. This is achieved by left-clicking to mark targets in green and right-clicking to mark non-targets in red. This can be seen in Figure 4. This information is then sent to the network of five ARTMAP voters which then goes through a training process to learn the patterns which constitute the desired association. The network then classifies the entire image and an image overlay is displayed which allows the segmentation to be visually inspected. Due to the voter network, areas that are suspected to be targets by all five voters appears brighter than areas only suspected by one voter. Due to the speed of training and classification, the machine learning system's results can be quickly refined by the radiologist. This is typically done by the radiologist right-clicking on false positives and retraining the system.

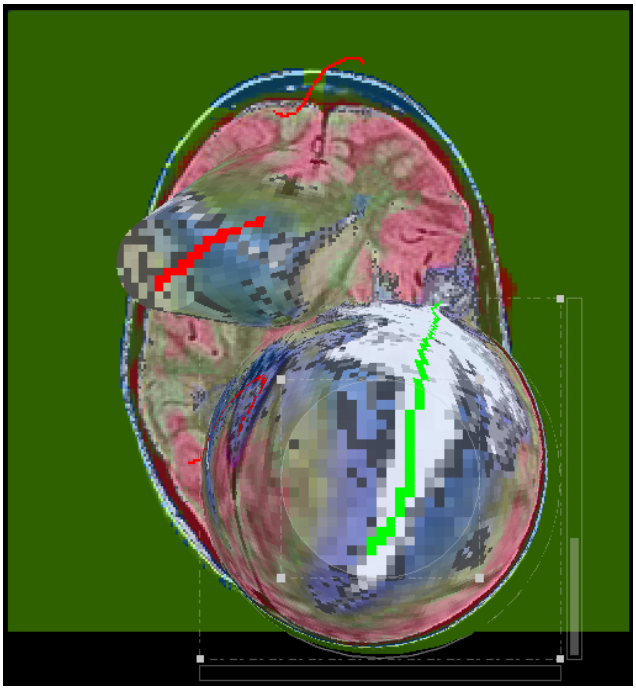

Figure 4. Contextual Zoom of Supervisory Signal

Figure 5 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 Figure 5a). The selection was made based on expert assessment presented in [7].

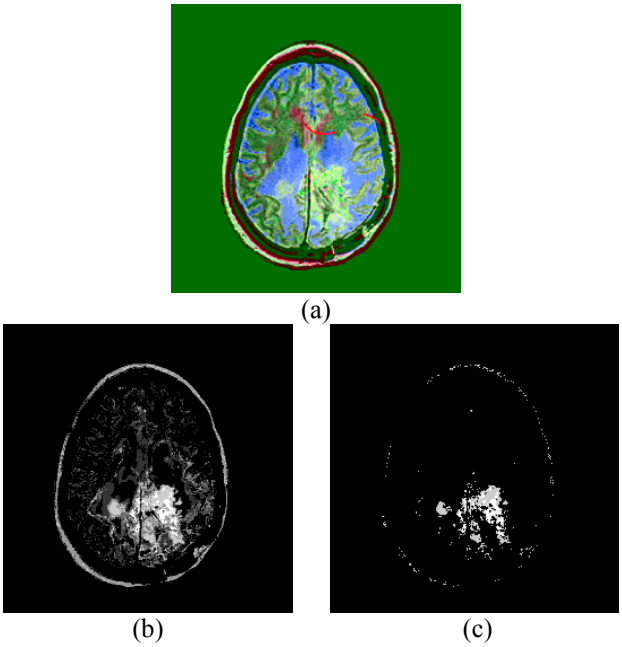

(a)

(b) (c)

Figure 5. Segmentation of parietal lesion caused by an anaplastic astrocytoma. See text for details.

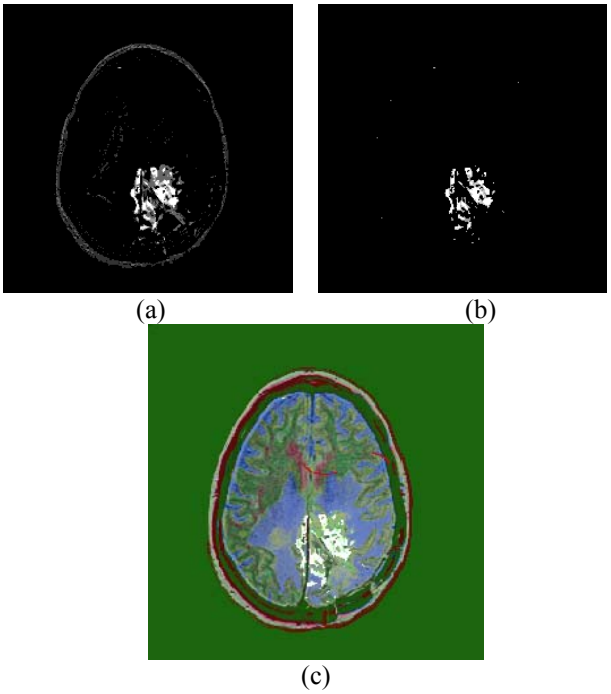

Figure 6. Segmentation as in Figure 5 using heterogeneous voting system. See text for details.

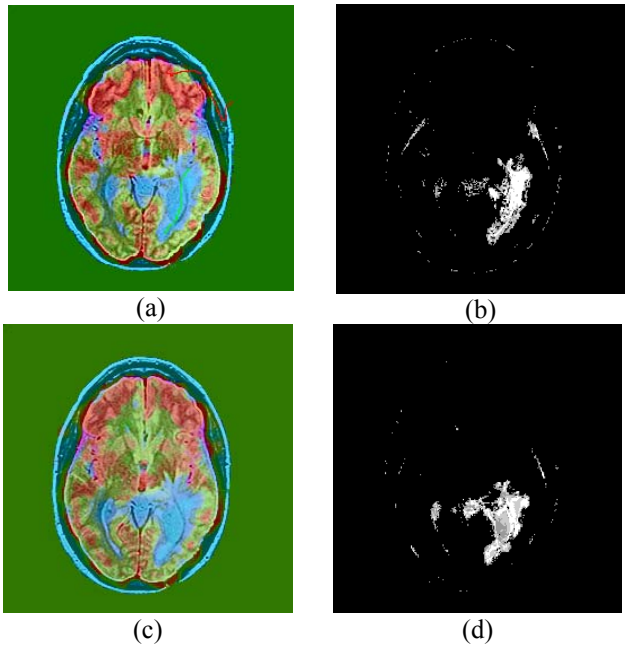

Figure 7. Cross-slice generalization results. (a) Training inputs (marking on slice 26). (b) Segmentation results (slice 26). (c) Generalization test input (slice 27). (d) Segmentation results (slice 27)

The results of applying the learn patterns across the entire image are shown in Figure 5b. Here, the grayscale value represents the number of voters that identified the given pixel as belonging to the "target" class. In Figure 5c, 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.

We implemented a heterogeneous voting approach where two out of the five voters were trained with modified vigilance parameters. Figure 6 presents the results obtained in segmenting the case shown in Figure 5. Here, Figure 6a presents the raw output of the system before thresholding for confidence. Figure 6b presents the final segmentation result. Here, the results are greatly improved eliminating all false alarms generated by the homogenous voting system (Figure 5c). Figure 6c 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 7 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 7a 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 7b. To test the generalization abilities of the system, a new slice was provided to the system. Figure 7c 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 7d. 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.

6 Summary

We have extended our image processing architecture to address the advantages of the 3D nature of MRIs. Our results compare favorably against the results of the original 2D and 3D processing architectures. Better definition of image details and volumetric information make this approach a significant advance in medical image feature creation. 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 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.

Acknowledgements

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] P. Devijver and J. Kittler, Pattern Recognition: A Statistical Approach, Prentice-Hall International, Englewood Cliffs, NJ, 1982.

[2] T. Kohonen, Self-Organizing Maps, Springer, Berlin, Heidelberg, 1995.

[3] G. Welch & G. Bishop, "An Introduction to the Kalman Filter", University of North Carolina at Chapel Hill, SIGGRAPH 2001, Course 8, <http://www.cs.unc.edu/~welch/media/pdf/kalman_intro.pdf>.

[4] V. Berikov & A. Litvinenko, "Methods for Statistical Data Analysis with Decision Trees", 2003, <http://www.math.nsc.ru/AP/datamine/eng/decisiontree.htm>.

[5] M. Egmont-Petersen, D. deRidder, & H. Handels, "Image Processing with Neural Networks – A Review", *Pattern Recognition*, Vol. 35, No. 10, pp. 2279-2301, <http://www.cs.uu.nl/people/michael/Journal-papers/Egmont-PR-Review2002.pdf>.

[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. 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.

[9] 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.

[10] M. Aguilar, J.R. New, and E. Hasanbelliu, "Advances in the use of neurophysiologically-based fusion for visualization and pattern recognition of medical imagery", *Proc. Of the Sixth International Conference on Information Fusion*, Cairns, Australia, 2003.

[11] 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.

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

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

[14] 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.

[15] 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.

[16] 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.

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

[18] 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.