Exposing the Black Box: Intuitive Representation of ARTMAP Networks

Joshua New and Jian Huang

Abstract— Interactive learning systems have long been used for directed knowledge discovery, volume segmentation, and transfer function design. Such systems are traditionally treated as black boxes in which a concise summary of the information encoded is not readily available, forcing professionals to leverage visualization of system weights. The inability of a user to understand learned patterns inhibits the ability to make sense of an agent's exhibited behavior. In this paper, we present an online learning system based upon Adaptive Resonance Theory which exhibits several practical strengths over traditional learning systems as well as a new mechanism to translate trained networks into intuitive compound boolean range queries. These queries can subsequently be used in combination with parallel coordinate plots to facilitate sense-making of the learned multivariate patterns. We showcase the efficacy of the system by demonstrating it on multivariate jet combustion data as well as tumor segmentation from MRI.

Index Terms—Learning systems, compound boolean range queries, large multivariate data visualization.

1 Introduction

As the size of modern datasets continues to grow, the problems of knowledge discovery, feature specification and tracking, as well as hypothesis testing becomes increasingly intractable. Fully automated filtering and computational tools are useful to aid in the process, but can rarely be used holistically within a given domain-specific context. Indeed, visualization rests upon the assumption that no matter how good pattern recognition and automation is, the best it can be is semi-automatic within the context of the entire scientific process; there is no magic to jump from fuzzy concepts to fully substantiated and verifiable specifications. In this paper, we seek to leverage visualization and cognitive processing to train the computer which patterns are deemed interesting. Furthermore, we seek to translate a traditionally black box learning system into a series of intuitive representations that can be then be used to facilitate understanding and promote scientific advances.

There are a wealth of learning systems which have been used within the visualization community for everything from autonomous pattern detection to transfer function design. While every system has its own strengths and weaknesses, we have elected to use a fuzzy learning system based upon Adaptive Resonance Theory (ART) due to several technical strengths and few weaknesses outlined in Section 3.2. ART is a mathematical model designed to address the "stability-plasticity" dilemma in which learning systems must be stable enough to avoid castastrophic forgetting but plastic enough to continuously learn. ART-based learning systems accomplish this by growing in size as new experiences are introduced but also retains perfect memory of past experiences. In this paper, we present an interactive segmentation system in which a user can succesively refine segmentation results from a heterogeneous network of learning systems using intuitive brushing of interesting image locations.

While codifying human knowledge into autonomous agents can be very useful, it has often been the case that extracting learned knowledge from an agent has been notoriously difficult. ART-based systems, like most learning systems, are often treated as black boxes in which the learned properties are encoded in a nearly-indecipherable series of edge weights. Several attempts have been made at visualizing neural nets in order to comprehend the reasons for exhibited behavior. In this paper, we present a method for converting SFAM networks into a representation as compound boolean range queries which can be used to intuitively and precisely identify learned categories and also presented in the context of multivariate relationships using parallel coordinate plots.

In the remainder of this paper, we first give a summary of related works in Section 2. In Section 3 we introduce the technical details of our contributions. Finally, our results and discussion are provided in Section 4, and then concluded in Section 5.

• *All authors are with the Department of Electrical Engineering and Computer Science of the University of Tennessee, Knoxville. Email: {new, huangj}@cs.utk.edu.*

Manuscript received 30 July 2007.
Under review.

2 Background

2.1 Segmentation with Learning Systems

Segmentation is defined as the act or process of dividing into segments. In image segmentation, there is a feature vector per pixel location which contains multiple features corresponding to the values for each of the variables available for that pixel location. Volume segmentation can effectively be extruded from image segmentation over multiple slices.

Machine learning is defined as the ability of a machine to improve its performance based on previous results. Machine learning systems allow automatic segmentation through "clustering" feature vectors into categories/classes. Machine learning systems come in three flavors; in increasing order of capability they are unsupervised, reinforced, and supervised. An unsupervised system is provided no hint to the correct classification, a reinforced system is provided a good/bad hint to the correct classification, and a supervised system is given the correct answer.

There are several common techniques and problems with segmenting data. Segmentation techniques can be broken into two general categories: manual, semi-automatic, and automatic. In manual segmentation, users must draw the desired borders onto the raw image. While manual segmentation is often highly accurate, this process takes much time, can be highly fatiguing, prone to errors, and obfuscate reproducibility. Automatic segmentation promises to address many of these issues. A short list of automatic segmentation methods includes Active Shape/Contour Models, Adaptive Segmentation, Bayesian Grouping, probabilistic neural networks, and Fuzzy clustering techniques. An analysis of these automatic segmentation methods and related problems is beyond the scope of this work, but the interested reader is referred to [6] showcased in the domain of brain segmentation from MRI. All known automatic segmentation methods suffer in varying degrees from problems such as variable imaging or simulation parameters, signal noise, overlapping intensities for continuous data points mapping to an implicit grid, partial voluming effects, gradients from discontinuities between successive slices or timesteps, and many other domain-specific concerns exist which make segmentation a "confidence"-related task. For this reason, semi-automatic methods are often used in which a human expert can leverage com-

putational tools to negotiate the tradeoffs between conflicting effects to mark up data with a much-reduced workload over a fully manual segmentation.

2.2 Adaptive Resonance Theory

Adaptive Resonance Theory (ART) is a mathematical framework based upon models of the hippocampus and neocortex developed by Carpenter and Grossberg in the 70s [7]. Many connectionist networks at the time suffered from the Stability-Plasticity Dilemma, which states the trade-off between a learning system stable enough to preserve learned patterns and yet plastic enough to learn new ones [8]. Adaptive Resonance Theory was developed to overcome this dilemma and has since served as a host for a plethora of neural network architectures, each demonstrating varying capabilities.

The ART1 class of architectures, developed in 1983, established the first ART-based neural network and performs unsupervised learning for binary input patterns [4]. The ART2 class of architectures, developed in 1987, included the ability to recognize analog vectors in which features are codified to a floating point between 0 and 1 [2]. ART3 [3] and ARTMAP [5], developed ART3 and ARTMAP, developed in 1987 and 1991 respectively, are members of the ART2 class of architectures along with dozens of other modern variants. The ART2 architecture variant known as Simplified Fuzzy ARTMAP is the one that was utilized in this study due to its computational efficiency, interactive performance, and many other properties that will be detailed later.

Simplified Fuzzy ARTMAP (SFAM) [12] is a fast, online/interactive, incremental, supervised learning system for analog signals. Fuzzy means that SFAM utilizes fuzzy learning rules for activation and selection of simulated neurons. SFAM can learn at a custom rate, but the fast learning rule is used because the simple fuzzy learning rules minimize the computation required for learning. SFAM is essentially a two-layer neural network that is specialized for pattern recognition, capable of learning every training pattern with very few iterations. The network starts with no connection weights, grows in size to suit the problem, uses simple learning equations, and has only one user-selectable parameter. In this system, the input vectors correspond to multiple metrics which defines the relationship for each pair of attributes from a set of multivariate data. SFAM is particularly well suited to this problem and we circumvent the dependence on an adjustable "vigilance" parameter and the dependence on the order of the input by using a voting scheme of heterogeneous networks.

3 System Components

As the resolution and number of variables increases in modern multivariate datasets, the ability to precisely identify interesting patterns in the dataset becomes increasingly intractable. This is exacerbated by the fact that one could compute a large number of derivative variables. Human cognitive ability could be overloaded in regard to remembering and using various combinations of metrics for a plethora of relevant scenarios. As is common with initial investigation of new data, experts may not even know what is important until they see it.

To address these situations, we have developed an intuitive user interface for data analysis and present the system diagram in Figure 2. It includes a learning system which takes inputs from the user brushing on a rendered slice of data to define which data points are interesting based upon the scientific question under investigation. The learning system then discovers which combinations of variables are useful in finding those relationships. The user is then presented with an overlay of the learned patterns based on the user selections for further refining the segmentation. The final trained networks can then be saved and reloaded when similar investigation is necessary in potentially new data. The networks can also be converted for use in many traditional visualization schemes such as: compound boolean range queries to quantify learned categories, parallel coordinate plots for qualitative assessment of multivariate trends, and transfer function design.

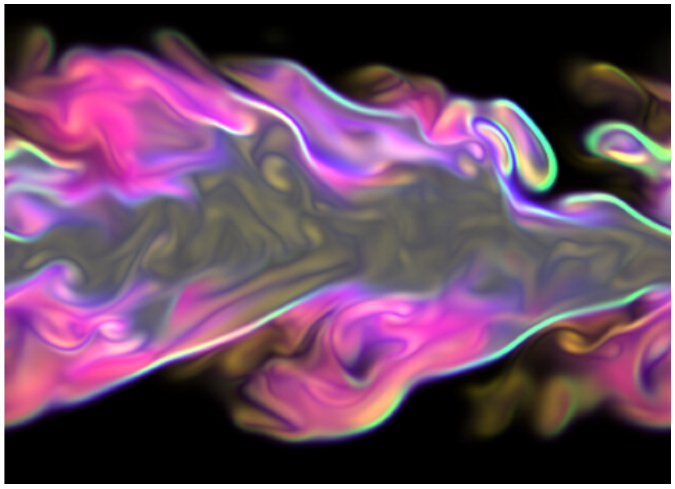

Fig. 1. Shader combination of 5 variables of jet combustion data.

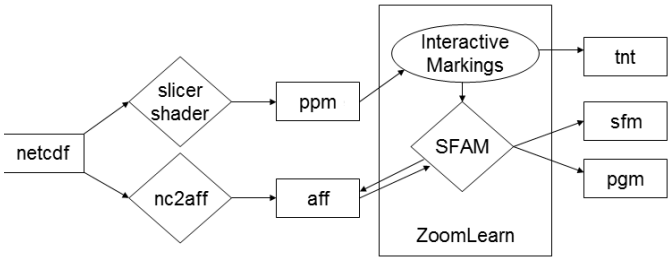

Fig. 2. Learning system determines areas of interest via user-in-the-loop interaction.

3.1 Shader-enhanced Visualization

Users of scientific visualization tools often do not know precisely what is novel, or what specific combination of variables constitute a particular feature of interest until it is seen. One of the most important elements for effective multivariate visualization is appropriate transfer function design. However, this can be difficult as the number of variables in common datasets grows. In this paper, we use a custom shader to blend 5 variables of the jet combustion dataset to create a single RGB image.

3.2 Effective SFAM Utilization

While there are an abundant number of various learning systems, we felt that the SFAM system provided several key advantages that make it particularly suited for this task. First, it is an online/interactive and incremental learning system meaning that it can both learn and classify user selections as a user successively refines the system's performance. This circumvents the laborious training time and scrapping/retraining necessary to incorporate constantly changing training datasets common for approaches such as backpropagation neural networks. Second, it is fast meaning it is able to learn in 5 epochs what takes most learning systems 1000s of epochs to learn. This can be used to provide real-time feedback as a user drills down on a specific pattern of interest. Third, it is a supervised network based upon analog processing which incorporates fuzzy learning rules to model uncertainty. Fourth, in order to speed computation, SFAM compliment-codes the incoming feature vector which doubles the feature vector length by subtracting each incoming feature from unity. This is a strength as it directly encodes the fact that users may be as interested in the absence of data than its presence.

There are also a couple disadvantages of SFAM which need to be addressed. First, there is a "vigilance parameter" which can be set

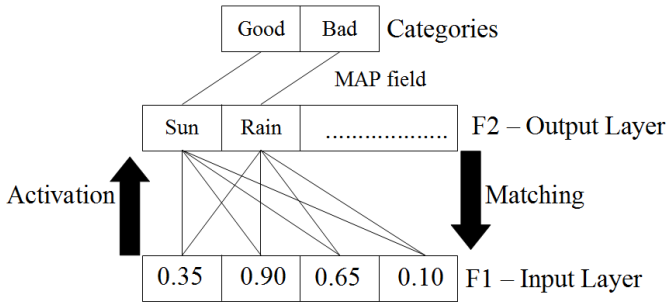

Fig. 3. Structure of an ARTMAP network.

from 0 to 1 and is typically set to 0.7 (but varies widely depending on the application). Vigilance corresponds to the "generality" of the classification, where a value near 0 means very general and a value near 1 means very specific. For example, the same object may be classified as a human at a vigilance of 0.3, a male at 0.5, and John Smith at 0.8. Second, we use fast learning which keeps the learning rate at 1 without sacrificing any recognition ability. However, this does introduce instability in that the learning system is sensitive to the order of the input vectors. To ameliorate these problems, we introduce a voting scheme utilizing 5 heterogeneous networks to establish a level of confidence. For our voting system, we use three SFAMs [12] with vigilances of 0.75, 0.675, and 0.825 that are trained on the same sequence of input data while a fourth and fifth network with vigilances of 0.75 are trained on different input sequences.

3.3 SFAM Learning

In order to understand this paper's contribution in converting a trained SFAM network to intuitive representations, some technical details of how SFAM learns is necessary. SFAM takes as input a series of K objects each codified by an N-element, compliment-coded floating point vector. SFAM's structure is such that it has an input layer, an output layer, and a mapfield which connects output layers to the specified supervisory signal. To learn, SFAM alternates between three phases for each input: output node activation, top-down pattern matching, and categorical mapping.

First, SFAM calculates output node activation. There are as many nodes in the input layer as there are features in the final, compliment-coded feature vector. The number of nodes in the output layer grows as inputs are classified. The nodes in the input layer and output layer are fully connected by weight-based connections initialized to 1. The manipulation of these weights is what allows feature classification. The activation function for the jth output node is defined by $A_j(i) = |I \wedge w_j|/(\alpha + |w_j|)$ where α is typically 0.0000001 and the $\wedge$ operator is the "fuzzy AND", simply corresponding to the minimum such that $(I \wedge w_j) = min(I, w_j)$. The winning node is defined as the one with the highest activation.

Second, SFAM calculates the top-down pattern matching. Once the winning node has been established, a match function is used to determine if the activated category classifies the feature vector sufficiently (if learning should occur). The match function is defined by $M = |I \wedge W_j|/N$. If $M \geq \rho$, where rho is the vigilance parameter, then the system is said to be in a "state of resonance"; that is, the output node j is good enough to encode input I and node j's weights are updated by adjusting the top-down weight vector to $w_j = \beta * (I \wedge w_j) + (1 - \beta) * w_j$ where β is the learning rate. In this paper, we use the "fast learning" rule in which $\beta = 1$ and thus $w_j = I \wedge w_j$. If $M < \rho$, a "mismatch reset" occurs, vigilance is increased to $\rho = M + \varepsilon$, where $\varepsilon = 0.0001$ and the second-highest activated output node is matched against the new vigilance. If this second output node meets the new vigilance requirement, its weights are adjusted to codify the current input. If it fails, a new output node is created with top-down weights equal to the compliment-coded feature vector. This secondary mismatch reset is what makes the SFAM system grow to recognize genuinely new input features.

Third, SFAM performs categorical mapping from output nodes to meaningful classifications specified by the supervisory signal. The system has taken inputs and learned to classify all of them as output nodes, but there can be several output nodes which constitute a single categorical idea which introduces the necessity of a "MAP field". Let us use the context of a typical training problem in which SFAM must determine from a Cartesian coordinate pair whether it is inside or outside a circle. If SFAM has already learned that one given coordinate is inside the circle, a very close coordinate that lies outside the circle might "match" the first coordinate. However, the supervisory signal provided notifies the system that this vector is outside the circle. This forces a "category mismatch", which triggers a "mismatch reset", forcing the creation of a new output node which is mapped to the "outside the circle" category. As may be gleaned by this example, SFAM output node weights correspond to centroids for clusters whose range is a function of the activation weights that serve to carve out hypercubes in the potential solution space which can be mapped to compound boolean range queries. However, we discuss a more general data-driven method in the section below.

3.4 Data-Driven Query Extraction

While deriving the mathematics for converting an SFAM network to a set of compound boolean range queries, we developed a more general mechanism that could apply to any classification system by adopting a purely data-driven approach. For example, a single SFAM network can perform learning and subsequent classification on every pixel of an image resulting in the classification of every pixel based on a specific output node ID as shown in Figure 6 and Figure 7. These classification results are PGM files colored by SFAM output node but could just as easily be constructed using SFAM "MAP field" categories, another clustering technique or learning system, or even combined results of multiple heterogeneous networks. We then use these clustered images to construct compound boolean range queries in a purely data-driven way that is classification system agnostic.

We have written a pgm2brq converter that takes as input a PGM image of classification values and the original data. The algorithm simply finds the min and max for each of the attributes of each location and outputs a single compound boolean range query for each classification ID. This method is simple, SFAM-agnostic, and runs in $O(N)$ time where N is the number of classified data values.

There are a couple assumptions and properties to this approach that are worth mentioning. First, it assumes that there is a unique classification at each position and therefore is amenable primarily to winner-take-all classification schemes. Second, compound boolean range queries innately carve out hypercubes in the dataspace; therefore, any clustering or learning system which partitions the space in another way will result in compound boolean range queries that will overlap in dataspace for multiple categories. There are many methods to address each of these concerns, but is considered beyond the scope of this paper.

4 Results and Discussions

4.1 Datasets

The first dataset we utilize is a simulation of turbulent combustion from a jet engine created by Sandia National Lab and made available through the SciDAC Institute for Ultra-Scale Visualization. This dataset consists of a 480x720x120 volume with 122 timesteps of 5 variables: OH (hydroxy radical), chi is scalar dissipation rate, hr, mixture fraction of air to fuel, and vorticity of . A custom transfer function and shader, as described in section 3.1, was used to combine all 5 variables into the color image used for segmentation. The database used for machine learning is simply the 5 normalized variables from the original dataset. One of the domain-specific goals of this dataset is to determine the location of flame boundaries along which extinction and reignition of the jet flame occurs based on underlying physics of chemical reactions.

The second dataset includes medical imagery, mostly consisting of MRIs of the brain, obtained from the Whole Brain Atlas web site [11].

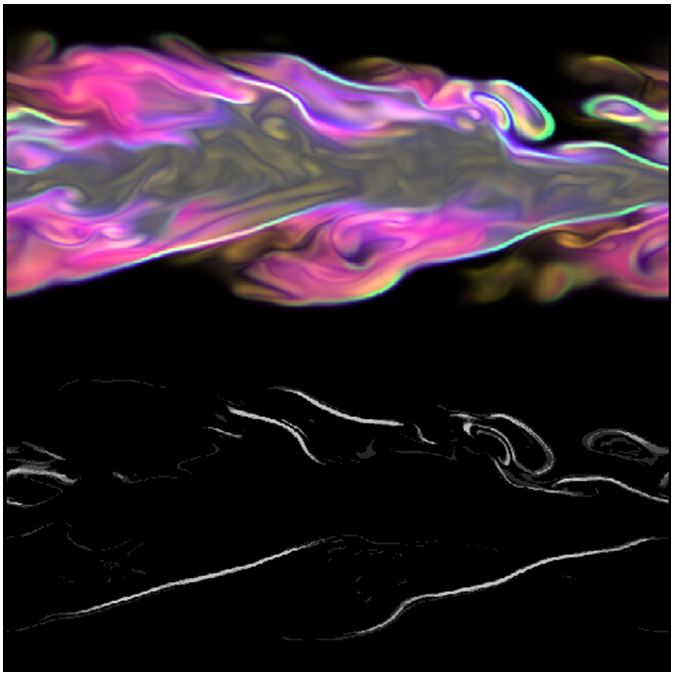

Fig. 4. Segmentation of flame boundaries in the jet combustion dataset.

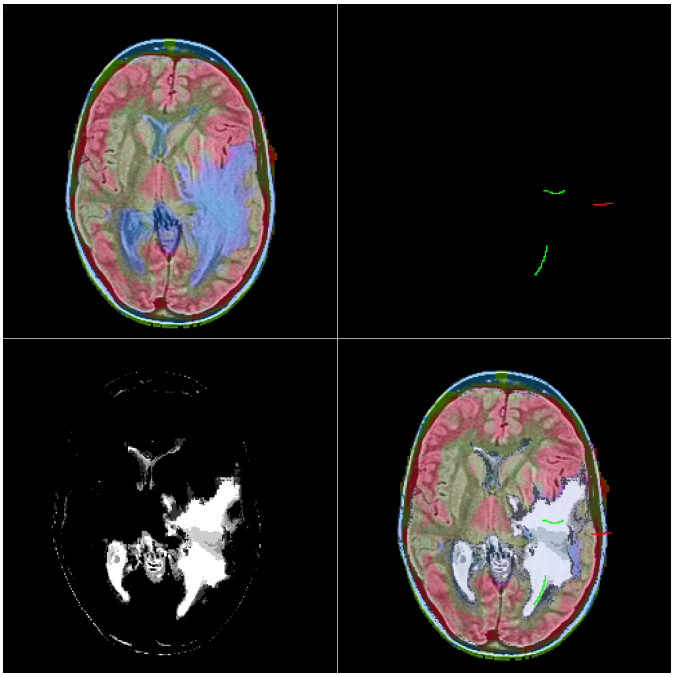

Fig. 5. Segmentation of tumor in MRI dataset.

This dataset consists of 256x256 image slices from the brain of a patient suffering from metastatic bronchogenic carcinoma. To generate the database of information for the SFAM-based segmentation, we took the spatially registered PD, T1, T2, and SPECT imaging modalities and used image processing with the 3D shunt operator [1] to create 16 opponency values for each pixel location. A sample of 3 of these opponencies was used as YIQ channels and then mapped to RGB chromatic space for display of multiple modalities simultaneously. One of the domain-specific goals is to segment the unhealthy brain tissue for planning of decompressive surgery.

4.2 Segmentation

The segmentation GUI is intuitive to use and capable of robustly delineating features in the data with only a few swipes of the mouse. In this interface, we use green to denote examples and red to denote counterexamples. Based upon these markings, the SFAM networks analyze the database information and report a gray-scale image recording the confidence of the heterogeneous collection of learning systems of other points similar to those selected by the user.

As can be seen in Figure 4, successful segmentation of flame boundaries is displayed. For this segmentation task, an unusually large number of disjoint points, consisting of 17 examples and 32 counterexamples, was utilized in order to highlight the strength of both the SFAM networks as well as the subsequent extraction of quantitative queries. Although 49 points were used, the SFAM networks created an average of only 7 output nodes per network to encode the diffent material types in the dataset. This results in data reduction and a filtered clustering of the dataset to aid in comprehension and attention direction.

Figure 5 showcases successful segmentation of metastatic bronchogenic carcinoma. This segmentation task involved only 3 swipes of the mouse and transparent network training despite 68 training points encoded using an average of 40 output nodes over 32 complement-coded features. The heterogeous set of trained agents is saved and can subsequently be utilized on a database of patients to scan for images which exhibit similar risk of this disease or for use in prescreening to direct radiologist's attention to the most likely locations of various disease types.

4.3 Transfer Function Design

Segmentation and other classification schemes can be used to create effective transfer function designs of structure within the data. In Figure 4 and Figure 5, we simply show a segmentation confidence overlay. This segmentation overlay could instead be used to modulate opacity for different structures depending on the task at hand. Indeed, assuming that the learning system used provides robust segmentation across slices, this method could be applied to multiple slices of a volumetric dataset for identification of isosurfaces or interval volumes that could be made transparent or highlighted for feature tracking.

In order to learn more about the dataset, we color the datasets by output node of an SFAM networks trained to recognize flame boundaries and carcinoma in Figure 6 and Figure 7. This classification clearly shows very coherent structures in the data in a method very similar to the non-photorealistic technique of toonification. This same technique could be applied to SFAM output nodes for material types, SFAM Map field categories to reduce this knowledge down to the segmentation results, multiple heterogeneous networks trained for a common task as in Figure 4 and Figure 5, or even multiple clustering or classification systems trained for different tasks (such as toonified results for various types of diseases visible to multiple image modalities).

4.4 Query Representation

While autonomous learning systems are useful in many circumstances, domain scientists are often trying to determine precisely which interplay of variables is giving rise to a specific, visual effect. In order to open the black box and allow the user to understand what the system has learned, we have developed a data-driven mechanism for mapping a set of classifications to a set of compound boolean range queries. For the result in Figure 4, we show the data-space centroid corresponding to each grayscale level in Table 1.

In these 10 complement-coded data centroids from the output nodes of an SFAM network, we have a strong delineation of specific types of chemical concentrations and their corresponding location within the dataset. There are 2 example classes which codify somewhat similar chemical properties of the flame boundaries. There are 3 counterexample nodes in which the first cluster consists only of data points inside the flame boundaries near the center of the simulation, the second

Table 1. Cluster centroids for different chemical concentrations in the jet combustion data

Boundary:

Y_OH	chi	hr	mixfrac	vort	!Y_OH	!chi	!hr	!mixfrac	!vort
0.269	0.620	0.632	0.432	0.126	0.731	0.380	0.368	0.568	0.874
0.254	0.125	0.476	0.268	0.073	0.371	0.576	0.311	0.362	0.673

Other:

Y_OH	chi	hr	mixfrac	vort	!Y_OH	!chi	!hr	!mixfrac	!vort
0.308	0.001	0.262	0.411	0.050	0.425	0.980	0.324	0.326	0.703
0.033	0.001	0.155	0.520	0.045	0.067	0.993	0.561	0.127	0.877
0.000	0.000	0.000	0.000	0.000	0.904	0.994	0.945	0.0198	0.767

Fig. 6. SFAM network output node clustering of jet combustion data.

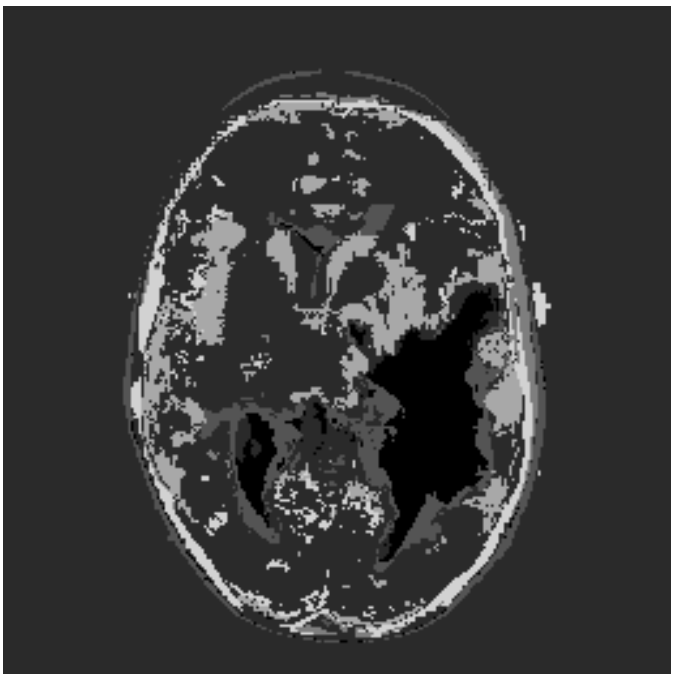

Fig. 7. SFAM network output node clustering of MRI data showing accurate classification of the tumor (black) as well as the surrounding edema.

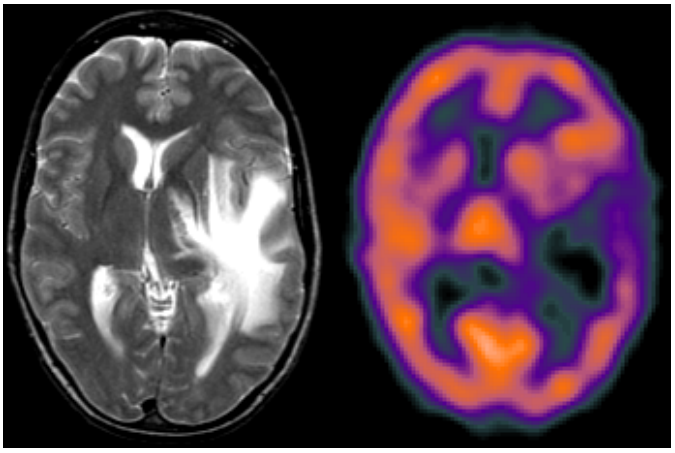

Fig. 8. High proton density and low amounts of blood flow is the single most important database factor in delineating a tumor caused from metastatic bronchogenic carcinoma.

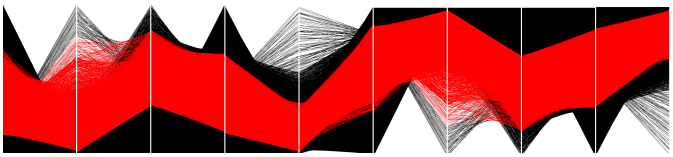

Fig. 9. Parallel coordinate plot of 10 complement-coded features for the jet combustion dataset showing in red all datapoints corresponding to flame boundaries based upon a set of 4 extracted compound boolean range queries.

cluster is outside the flame boundaries, and the third cluster is the black area near the edges of the simulation grid.

When applied to the material types codified by the output nodes of an SFAM network, we also have direct quantitative specifications of those types. For example, the segmented tumor in Figure 5 corresponds to only one output node and thus only one compound boolean range query as shown in Table 2. The tightest range, and therefore the single variable which most concisely represents the area segmented as tumor is the 13th variable with a normalized range of 0.21 corresponds to a 3D shunt operator using functional MR-PD (proton density) in the "on center" channel and metabolic SPECT modality in the "off surround" channel. As can be seen in Figure 8, the tumor has been traced to an area of high proton density but inhibited bloodflow. Upon further inspection, this pattern was confirmed by radiologists from the case details. Use of other variables are necessary to improve segmentation by removing competing regions such as that of the skull.

4.5 Multivariate Representation

In an effort to not only convey qualitative segmentation results or precise quantitative ranges, we also use parallel coordinate plots for relaying multivariate trends in the data. This is important because while the quantitative queries can relay specific features that are of primary importantance for the current segmentation task, it is often difficult for a user to understand the inter-variable dependencies present for a segmentation task. This property becomes readily apparent in the case of the jet combustion dataset in which there is a rich microphysics interplay among all variables to determine areas of the flame boundary.

In Section 4.4, we were able to use the extracted queries to quantitatively define the range for the most important factor in determining which region(s) constitute the tumor. By representing this data in parallel coordinate space, we are able to see that the range could be cut in half by dropping only 5 datapoints. Therefore, this mechanism can be used in a linked-viewport format in which radiologists can interactively continue to refine the segmentation results via brushing.

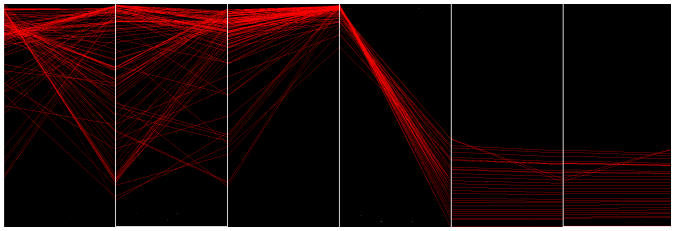

Fig. 10. Parallel coordinate plot of a subset of the variables in the MRI dataset showing in red all datapoints corresponding to tumor.

5 Conclusion

In conclusion, we have provided a heterogenous learning system capable of interactive performance on large data capable of determining which metrics are of interest based upon trends identified by the user. The classification of these networks has been demonstrated for transfer function design of large, real-world datasets. A mechanism has been developed for translating SFAM-based learning systems to an intuitive representation of the patterns learned. Parallel coordinate plots are used to convey these patterns to the user during the interactive process for enhanced hypothesis testing. The results demonstrated clearly show the recognition and summary capabilities of the system for multivariate data.

Acknowledgments

The authors wish to thank Kwan-Liu Ma for the combustion dataset. The work is funded in part by NSF CNS-0437508, and through DOE SciDAC Institute of Ultra-Scale Visualization under DOE DE-FC02-06ER25778.

References

[1] M. Aguilar and J. New. Fusion of multi-modality volumetric medical imagery. In *Proceedings of the 5th Inernational Conference on Information Fusion*, 2002.

[2] G. Carpenter. Neural network models for pattern recognition and associative memory. *Neural Networks*, 2(4):243–257, 1989.

[3] G. Carpenter. Art3: Hierarchical search using chemical transmitters in self-organizing pattern recognition architectures. *Neural Networks*, 3:129–152, 1990.

[4] G. Carpenter and S. Grossberg. A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, 37(1):54–115, 1987.

[5] G. Carpenter, S. Grossberg, and J. Reynolds. Artmap: Supervised real-time learning and classification of stationary data by a self-organizing neural network. *Neural Networks*, 4:565–588, 1991.

[6] V. Caviness and D. Kennedy. Mri brain segmentation: Automatic segmentation. Technical report, 2004.

[7] S. Grossberg. Adaptive pattern classification and universal recording ii: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, 23(4):187–202, 1976.

[8] S. Grossberg. How does the brain build a cognitive code? *Psychological Review*, 87(1):1–51, 1980.

[9] Alfred Inselberg. The plane with parallel coordinates. *The Visual Computer*, 1(2):69–91, 1985.

[10] H. Janicke, M. Bottinger, and G. Scheurermann. Brushing of attribute clouds for the visualization of multivariate data. *IEEE Transactions on Visaulization and Computer Graphics*, 14(6):1459–1466, 2008.

[11] K. Johnson and J. Becker. Whole brain atlas. http://www.med.harvard.edu/AANLIB/home.html, 1999.

[12] T. Kasuba. Simplified fuzzy artmap. *AI Expert*, 8:18–25, 1993.

Table 2. Extracted 32-feature query for tumor in MRI data corresponding to the black region of Figure 8.

[0.245,0.990]	[0.100,1.000]	[0.405,0.998]	[0.000,0.326]	[0.114,0.991]	[0.145,0.916]	[0.560,1.000]	[0.161,0.880]
[0.154,0.505]	[0.208,1.000]	[0.103,0.998]	[0.137,0.992]	[0.789,1.000]	[0.000,0.405]	[0.000,0.376]	[0.000,0.358]
[0.010,0.755]	[0.000,0.900]	[0.002,0.595]	[0.674,1.000]	[0.009,0.886]	[0.084,0.855]	[0.000,0.440]	[0.120,0.839]
[0.495,0.846]	[0.000,0.792]	[0.002,0.897]	[0.008,0.863]	[0.000,0.210]	[0.595,1.000]	[0.624,1.000]	[0.642,1.000]