Multisensor Image Fusion, Mining, and Reasoning: Rule Sets for Higher-Level AFE in a COTS Environment

Marianne Chiarella, David A. Fay, Richard T. Ivey, Neil A. Bomberger, and Allen M. Waxman
Cognitive Fusion Technology Directorate
Fusion Technology & Systems Division
ALPHATECH, Inc.
Burlington, MA, USA
{mchiarella, fay, rivey, neilb, waxman}@alphatech.com

Abstract – *Cultural features of interest for map making are often defined by qualitative non-local spatial relationships. Coastal shoreline, for example, is the boundary where land, including islands, shares edge with open water. Previously, we reported on our Neural Fusion system for multisensor image fusion and mining for local features. In this paper, we introduce a rule-based post-mining stage that infers higher-level feature descriptions from local mining results. Our methods for fusion and mining are based on concepts derived from neural models of visual processing. They capture local spectral and spatial information for semi-supervised pattern learning & recognition. With the added capability of non-local geospatial reasoning, we demonstrate feature extractions that are not possible using local pattern recognition alone, like river features through cloud occlusions. We have incorporated these methods into a commercial software platform for imagery exploitation. We summarize the approach and user interfaces, and demonstrate the integration of geospatial reasoning with our Neural Fusion and Mining toolset.*

Keywords: Sensor fusion, image fusion, image mining, target recognition, feature extraction, rule-based reasoning.

1 Introduction

In map production, the geospatial analyst applies feature-specific extraction guidelines, map-specific rules for capture conditions, and general rules-of-thumb to create an accurate and readable final vector product. Many extraction rules describe complex features that only exist as combinations of lower-level components. Extraction guidelines also dictate whether features are converted to points, lines, or polygons in the final vector product. For example, width determines if *river* is converted to a double-line or single-line drainage feature. In some cases, the analyst will extend a feature through an obstruction, knowing that the feature exists, but is occluded (e.g., *clouds* over *river*). In other cases, the feature is extracted as a large connected area even though the feature does not exist everywhere, as in small clearings in large crop areas. Spatial reasoning capabilities are required to assist in such extractions that go beyond the local spatial operators used in image mining [1]-[5]. We present a rule-based approach that combines the results of image fusion and image mining (local pattern recognition) with higher-level spatial reasoning.

Section 2 summaries our neural methods for image fusion (adaptive contrast enhancement, opponent-color contrast, multi-scale contour completion, and multi-scale texture enhancement), and image mining (semi-supervised pattern learning & recognition). In Section 3, we report progress on integrating our *Neural Fusion* methods [1]-[2] for multisensor and multispectral image fusion and interactive image mining into a commercial GIS software environment, ERDAS *Imagine*. This integration allows us to leverage all of *Imagine*'s existing geospatial tools and quickly reach an entire user community already familiar with *Imagine*. *Neural Fusion* is packaged as a set of add-on modules for *Image Conditioning*, *Image Fusion*, extraction of spatial *Context Features*, and interactive *Image Mining*. Together, these modules create a work flow enabling a user to create vector products of foundation features (e.g., roads, rivers, and forests) and highlighted target detections from raw multisensor or multispectral imagery.

We introduce a new module to the *Neural Fusion* toolset for *Ground Truthing and Performance Assessment* in Section 4, and present the results of a comprehensive performance assessment over a variety of feature classes. We show that *Neural Fusion* probability of detection is typically above 98% with false alarm probability below 0.5%. We extend *Neural Fusion* capabilities in Section 5 with the adoption of a rule-based post-mining stage with the goal of more closely emulating results produced by the human geospatial analyst.

2 Neural Architectures

Our image fusion methods are motivated by biological examples of Visible/IR fusion in snakes [8], and color processing in the primate visual system [9][10]. We have, in the past, used these methods for both realtime fusion of multiple night vision sensors [6][7], and off-line exploitation of surveillance imagery including EO, IR, MSI, HSI, and SAR modalities [1]-[5].

Image fusion is based on the visual processes of spatial- and spectral-opponent networks, as realized in the retina and primary visual cortex. This serves to contrast enhance and adaptively normalize the wide-dynamic range imagery and also decorrelate the cross-modality interactions that occur in two stages, as shown in Figure 1. These opponent-color interactions can also be thought of as spectral contrast enhancement, as well as estimation of reflectance slope and curvature (i.e., they approximate local spatial derivatives with normalization). The opponent interactions are modeled using center-surround shunting neural networks [11][12] in a local feed-forward

architecture to form a nonlinear image processing operator. This same operator is used repeatedly in the fusion architecture shown in Figure 1.

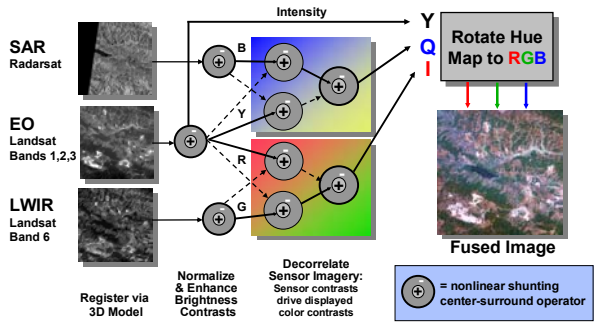

Fig. 1. Double-opponent color architecture for image fusion constructed from center-surround shunting neural networks. This architecture can be modified to fuse imagery from two to five different sensors or spectral bands. (Radarsat and Landsat imagery provided by NGA.)

The outputs of all the opponent operators shown in Figure 1 form registered image layers that contribute to the growing registered data stack shown in Figure 2. In addition, we extract local spatial context features from the enhanced imagery, opponent-color imagery, and (if available) 3D terrain data, in the form of extended contours, periodic textures, grayscale variance, and surface steepness and curvatures. Visual processing models of oriented contrast and contour on multiple scales are implemented as neural networks based on Gabor receptive fields [13]-[15] to extract these local context features. These context features augment the layered set of feature imagery that resembles a stack of spectral imagery (which can also be incorporated into the stack), and can be visualized as shown in Figure 2. A vertical cut through the stack corresponds to a feature vector at each pixel, which will serve as input to the neural pattern classifier described below.

Through the use of a simple point-and-click graphical user interface or GUI (to be illustrated in the next section), feature vectors can be selected to correspond to examples and counter-examples of targets of interest in the fused imagery. This training data is then fed to a pattern learning and recognition network in order to train a target detector (i.e., search agent) that can then be used to search extended areas for more of those targets. We utilize a combination of a modified Fuzzy ARTMAP neural network [16][17][18] and a modified version of an algorithm which exploits the signatures of both target and non-target context [5], to discover a much reduced decorrelated feature subspace that is sufficient to learn the training data selected. This serves to accelerate the search for targets, highlights salient features of the target in context, and has implications for multi-sensor tasking. The process of training data selection, pattern learning and subspace projection occurs very quickly, and is suitable for an interactive environment in which a human operator rapidly teaches the computer to find targets of interest.

A simplified version of the modified Fuzzy ARTMAP neural network is shown in Figure 3. It consists of a lower ART module, which performs unsupervised clustering of feature vectors into categories, and an upper layer in which the learned categories form associations with one or the other class for a target of interest. This approach enables the network to learn a compressed representation of the target class in the context of non-targets in the scene. That is, it learns to discriminate the target from the context, using the multisensor data and spatio-spectral features. A target of interest is typically represented by a few learned categories, as are the non-targets. Also shown are the match-tracking attentional pathways that modulate the matching criterion (vigilance) when predictive errors occur during training [17], leading to further category search or creation of new categories. Simplified Fuzzy ARTMAP [18] combines unsupervised category learning with supervised class association in hierarchy.

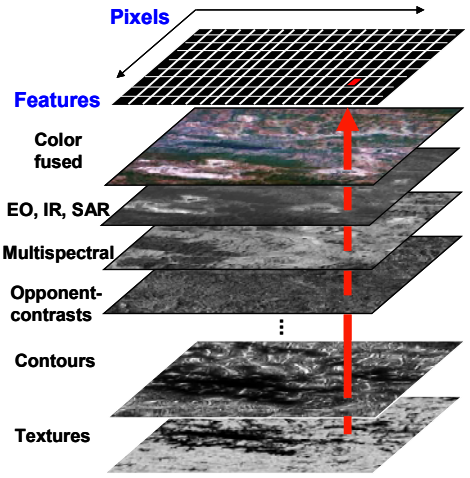

Fig. 2. Layered stack of registered, processed, multisensor imagery and spatial context features. A vertical cut through the stack corresponds to a feature vector at each pixel.

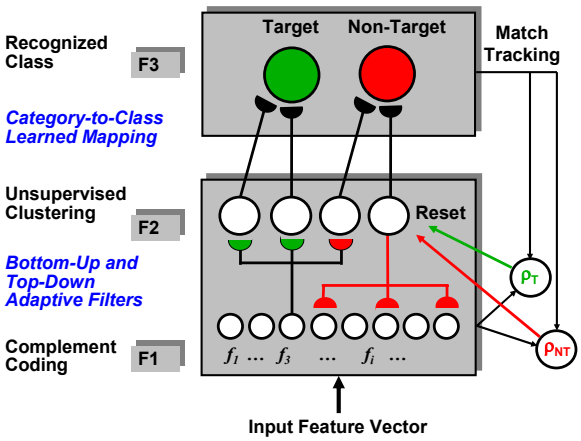

Fig. 3. A modified version of the simplified Fuzzy ARTMAP [18] neural network used for interactive training of search agents that discriminate targets from non-targets in context.

3 COTS Implementation

In order to enable widespread dissemination and use of our neural methods for multisensor/spectral image fusion and mining, we have refined and enhanced our prototype system [4], and reimplemented it within an extensible commercial image exploitation environment, ERDAS *Imagine*. This allows us to take advantage of the significant software infrastructure and capabilities of the *Imagine* suite, while supporting technology transfer to users already familiar with this exploitation environment.

The process workflow for image fusion and mining is shown in Figure 4, where the red and green boxes correspond to our own modules, and the blue boxes are modules developed from existing *Imagine* functionality. These modules are reflected in the software toolbar and pop-up menu, shown in Figure 5.

Here we illustrate a few of the GUIs for these image fusion and mining tools. Figure 6 illustrates the Image Fusion module user interface, with an area in Monterey, CA, imaged in four spectral bands (red, green, blue, near-IR) and a panchromatic band, being fused into a color presentation. This module, based on opponent-color processing, supports mixed-resolution imagery as collected by the IKONOS satellite.

After image fusion and spatial context features are generated using the various modules, and layered onto the geodatabase we call a *deepfile*, the user begins the process of interactive *Image Mining* by launching the GUI shown in Figure 7. A search agent is easily trained to find targets pointed out (using the mouse) by example and counter-example. Also, a subset of sufficient feature layers is discovered that is able to support 99-100% correct categorization of the selected training examples and counter-examples, and the user is then informed of these important feature layers, as shown in Figure 8. As the user points out mistakes or missed targets in subsequent search areas, the mining system learns to improve its performance and refines the trained search agent. This training process is very rapid. The trained agent can then be used to search for targets over the entire scene.

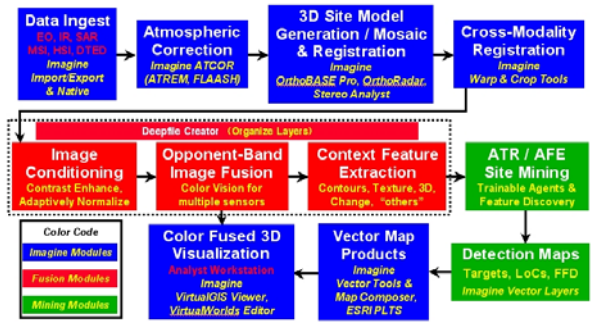

Fig. 4. Workflow for multisensor and multispectral image fusion (red boxes) and mining (green boxes), integrated into the ERDAS *Imagine* environment (blue boxes).

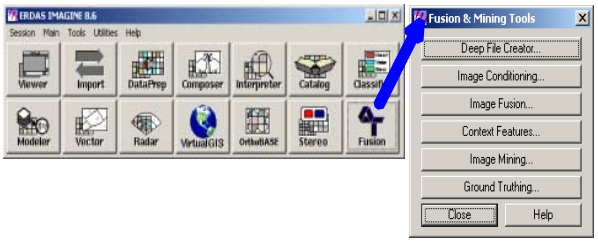

Fig. 5. Toolbar for ERDAS *Imagine*, augmented with a *Fusion* button that launches our menu of *Fusion & Mining Tools*. Each menu button launches one of the corresponding modules.

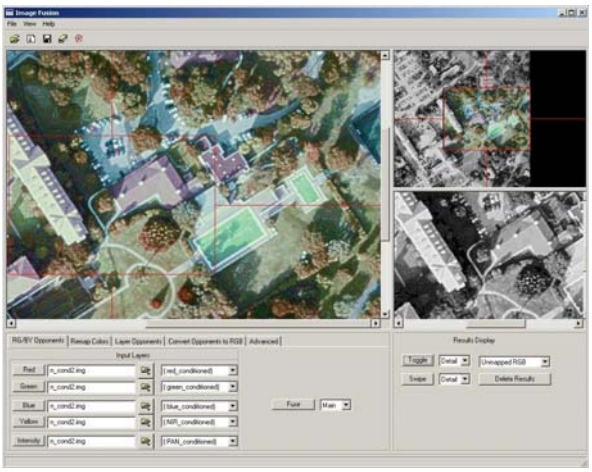

Fig. 6. User interface for image fusion of up to 5 sensors or spectral bands. Opponent-sensor pairs are mapped to human opponent-colors *red vs. green* and *blue vs. yellow*. An additional channel supports a high-resolution intensity band, enabling pan-sharpening. The viewers show a site overview (upper-right), main window (left), and detail area (lower-right window) of Monterey, CA. (ADAR imagery provided by NGA.)

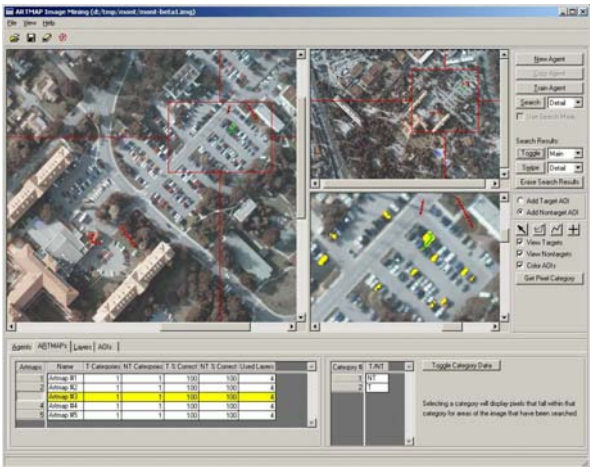

Fig. 7. GUI for interactive *Image Mining* using modified ARTMAP pattern learning & recognition. The user selects training pixels, examples (in green) and counter-examples (in red). Each pixel has an associated feature vector. In this case the agent has learned to find *red cars* (detections

shown in yellow in the lower-right detail window) in Monterey, CA. (**ADAR** imagery provided by NGA.)

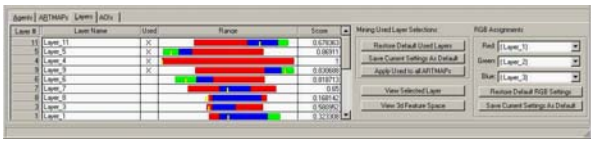

Fig. 8. The Layers tab in the *Image Mining* GUI shows the system discovered that only 4 of the 20 spatio-spectral layers were needed to find the *red car* targets in context. The GUI supports the user in understanding the feature space, the categories generated, and their relation to the training data. Interactive 3D viewing of data clusters and learned categories is supported.

Fig. 9. IKONOS imagery has been fused and combined with spatial context features for mining. A first-time user rapidly trained the system to search for iron ore. This *detection mapping* capability is part of the *Image Mining* module and supports morphological filtering of the detections and conversion to vectors for map-making. Iron ore detections have been outlined by vectors in yellow, and train cars are found carrying ore to the plant (as revealed in the lower-right detail window). (IKONOS imagery provided by ERDAS/Leica Geosystems.)

Following the search process, local detections can be spatially filtered using morphological processing and then converted to vectors for editing, attributing, and map-making. Multiple search agents are trained and used in order to construct maps of multiple classes. To add new target classes to a map, one need only train a new search agent for the new class, and multiple agents could be constructed simultaneously by multiple operators exploiting the same data set. Figure 9 illustrates the GUI used to convert raster detections into vectors for map-making. This example utilizes IKONOS imagery (4-band MSI at 4-meter resolution in conjunction with a panchromatic image at 1-meter resolution) of a steel plant outside Charleston, South Carolina. In this example, a first-time user trained an agent to find piles of iron ore (outlined with yellow vectors), and discovered trains carrying ore to the plant (as indicated in the lower-right detail window). Because *Neural Fusion* is tightly integrated into the ERDAS *Imagine* environment, we are able to take advantage of *Imagine's* support for processing image subsets so as to enable processing very large format imagery using a tiling strategy. Throughout the processing chain, only those portions of an image or deepfile that are needed are kept in memory at any one time.

4 Performance Assessment

A new *Ground Truth and Performance Assessment* module was recently developed for the *Neural Fusion* tools. The user manually ground truths the scene features using drawing tools provided on the GUI as shown in Figure 10. Different colors are assigned to different scene features, such as *orange* for ore piles, *yellow* buildings, and *pink* industrial ground. The ground truth is used in conjunction with *Image Mining* detection results to assess the performance of the mining. The Evaluation tab of this module allows the user to choose different evaluation methods, depending on the type of feature and accuracy in ground truthing. Mining performance is calculated for each class in the image in terms of detection and false alarm probabilities. Results are displayed for each class separately (in the class table), as well for the entire scene.

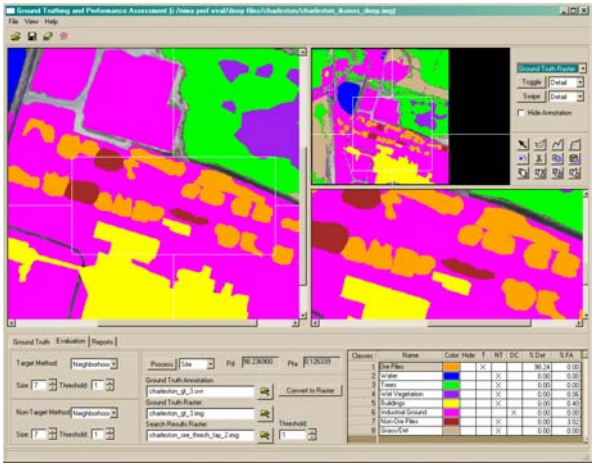

Fig. 10. The *Ground Truth and Performance Assessment* module is used here to evaluate detection performance for *iron ore* extraction in this Charleston, South Carolina IKONOS scene. The ground truth classes are shown in pseudo-color, and the results of the performance assessment are displayed in text fields on the tab window.

Neural Fusion performance was evaluated over a wide range of foundation features including *crops*, *trees*, *rivers*, *piers*, *urban areas*, *vegetation* (trees and grassland, or groups of crops), *submerged vegetation*, *boats*, *cars*, and *ore piles*.. The majority of the imagery analyzed was IKONOS, but other sensors have been utilized as well including fused Landsat & Radarsat, Hyperion and HyMap hyperspectral sensor systems, and simulated MTI reflective bands from AVIRIS hyperspectral. For this study, 15 data sets were manually ground truthed for comparison with the aided feature extraction results. The results are summarized in Figure 11. The system characteristically performs at a probability of detection

above 98% and false alarm rate below 0.5% for land cover and drainage features. Aside from roads, which are sometimes confused with rooftops made of similar materials, the biggest loss in performance was due to difficulty in manual ground truthing. The poorest performance on vegetation features, for example, were submerged vegetation and sparse trees in suburban scenes, both very difficult to ground truth.

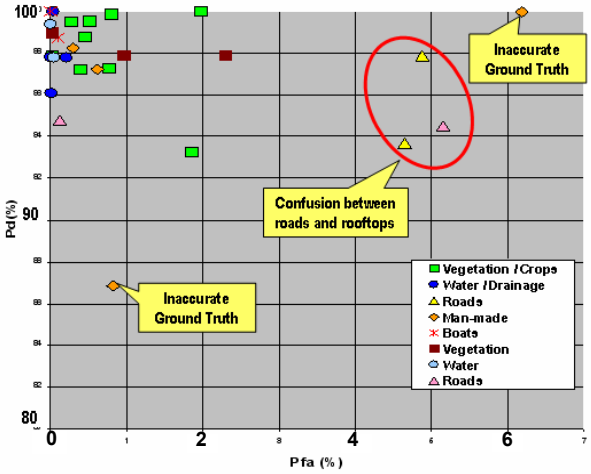

Fig. 11. Feature Extraction Performance across 15 data sets is typically above 98% with false alarms below 0.5%.

5 Rule Sets for Higher-Level AFE

The human geospatial analyst draws from multiple sources of information to extract features for map-making. In addition to multi-source imagery, they include feature-specific extraction and attribution guidelines, map-dependent capture conditions, and general analyst rules-of-thumb. In order for assisted feature extraction tools to gain acceptance for map production, the assisted results must closely emulate those of the human analyst. While *Neural Fusion* provides reliable search results for constituent map features, a post-mining stage is necessary to assist in composing the final result for more complex or abstract features.

Many of the most frequently extracted features for map making are actually best described as combinations of related features. *Coastal shoreline*, for example, is defined as the boundary line where *land*, including *islands*, shares an edge with *open water*. Moreover, each constituent feature has extraction guidelines that govern capture conditions to ensure that proper data densities are attained for the associated map scale. An *Irrigation Ditch* in *Cropland*, for example, would not be extracted if it did not meet a minimum length specification. Integrating rules for capture conditions into *Neural Fusion* assists the human analyst by providing a more standardized result that will require less vector editing in map production. In addition to assisting with complex feature extractions and reducing vector editing time, rule-based reasoning is applied here to overcome anomalous AFE situations such as occluded features, e.g., clouds or trees over *river* feature.

A rule-based post-mining stage is a natural extension to *Neural Fusion*. It provides the framework to incorporate feature-specific capture conditions, complex geospatial feature descriptions, and analyst rules-of-thumb without compromising demonstrated *Neural Fusion* performance (summarized in Section 4 above). The rule base (or knowledge base) is created by translating extraction specifications into a set of IF-THEN conditional statements. An example is shown in Figure 12 for the case of extracting submerged vegetation near the coastal shoreline. The IF part of the rule contains the condition that must be satisfied in order for the THEN part of the rule to take effect. As with any rule-based system, once the THEN part of the rule is satisfied it can be used in the IF part of another rule, and so on [19][20]. The rules are triggered, for the most part, by search results supplied by the *Neural Fusion Image Mining* module. Take the rule set in Figure 12, for example, where the goal is to classify submerged vegetation close to the coastal shoreline. Three *Neural Fusion* search agents are used. They are *land*, *open water*, and *submerged vegetation*. These results are locally filtered and additional information is extracted for *land boundary* and *coastal buffer zone* using *Neural Fusion detection filtering* in conjunction with *Imagine's* image processing and geospatial measurement tools. The final hypothesis is *Submerged Vegetation near Coastal Shore*.

```
IF  Submerged Vegetation
    AND   Coastal Shoreline
    THEN  Submerged Veg near Coastal Shore

IF  Coastal Buffer Zone
    AND   Open Water
    THEN Coastal Shoreline

IF  proximity to Land boundary < 100 m
    THEN  Coastal Buffer Zone

IF  Land agent detects with high confidence
    AND area > 400,000 sq m
    THEN  Land

IF  Open Water agent detects with high
          confidence
    THEN  Open Water

IF  Submerged Vegetation agent detects with
          high confidence
    THEN  Submerged Vegetation
```

Fig. 12. Rule Set for Submerged Vegetation near Coastal Shoreline.

Figure 13 is the result of applying this rule set to a small sample of IKONOS data of Beach Haven, NJ. Capture conditions for proximity to shoreline and land mass filter size are incorporated into the rule set. Notice that as a result of the feature-specific capture condition for land mass size, the small area of land near the tip of the island is not included in the coastal shoreline feature. The result highlights only the submerged vegetation feature that satifies the rule for *Coastal Buffer Zone* (within 100m of land).

In order to formulate a rule set, working knowledge is required about the *Neural Fusion* system as well as domain knowledge about geospatial analysis. *Neural Fusion* expertise is required because most of the variables

used in the rule set are products of *Neural Fusion* processing, in particular search results from the *Image Mining* module. Search results are typically filtered using morphological processing in the *detection filtering* capability in the *Image Mining* module. *Detection Filtering* provides a front-end to rule-based processing in much the same way as the *Neural Fusion* front-end prepares data for *Image Mining*. *Detection filtering* creates feature centerlines (thinning), outlines (boundary), and filtered results (erosion/dilation, open/close, despeckle/inverse despeckle, group). Custom filter schemes can be saved for reuse on similar extractions. Other variables used in constructing feature-specific rule sets are created using *Imagine's* modules for GIS analysis (proximity, contiguity, thematic attributes) and image manipulation (subtraction, union, intersection). One additional source of data that is important to mention is the *Neural Fusion* front-end *Context Features* module. *Context Features* provides local neighborhood spatial information including boundary contours, textures, and terrain elevation, slopes, and curvatures.

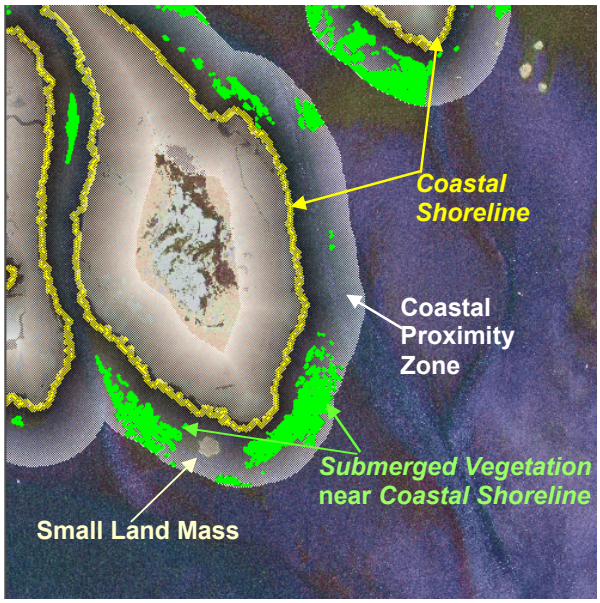

Fig. 13. *Submerged Vegetation* extracted near *Coastal Shoreline*. This composite image shows the color-fused pan-sharpened IKONOS image with overlays of images processed by the rule set of Fig.12. (IKONOS imagery provided by Naval Research Laboratory.)

The example in Figure 14 uses the Boundary Contour System (BCS) result from *Context Features* to fill small gaps in a river extraction. BCS in *Neural Fusion* is a fast, bottom-up, multi-scale boundary enhancement that approximates Boundary Contour System theory [14]. The gaps most probably resulted from insufficient training (i.e., missed variability of data) during the *Image Mining* stage. To correct this result, the user could supply more training data and search again. This may work for this example, but not at all for gaps caused by obscurations (e.g., clouds or overhanging trees). Many iterations may be required to train the learning system for anomalies in the feature of interest. The imperfect raster extraction could be converted to vector (the red outline in Figure 14a) and edited using vector editing tools. This approach is unattractive because it is labor-intensive and therefore very time consuming. Using the rule-based approach, we construct a rule for *River with gaps*. The rule states that IF a pixel is in close proximity to the centerline of the high confidence *river* search result AND there is a strong BCS response, THEN the pixel is *river*. Applying this rule in a post-mining stage enables the missed river pixels to be recovered quickly with no additional training.

Once defined, the rule set is added to a knowledge base using *Imagine's Knowledge Engineer*, a COTS visual programming tool for designing knowledge bases for use with *Imagine's Expert Classifier* (www.gis.leica-geosystems.com). The *Expert Classifier* works like a typical backward-chaining system (i.e., a decision tree) and results in the most likely hypothesis chosen in the final classification. The decision tree shown in Figure 14e is the rule set representation in the *Knowledge Engineer*. The blue boxes are the conditions, the yellow boxes are rules, and the green box is the terminal hypothesis.

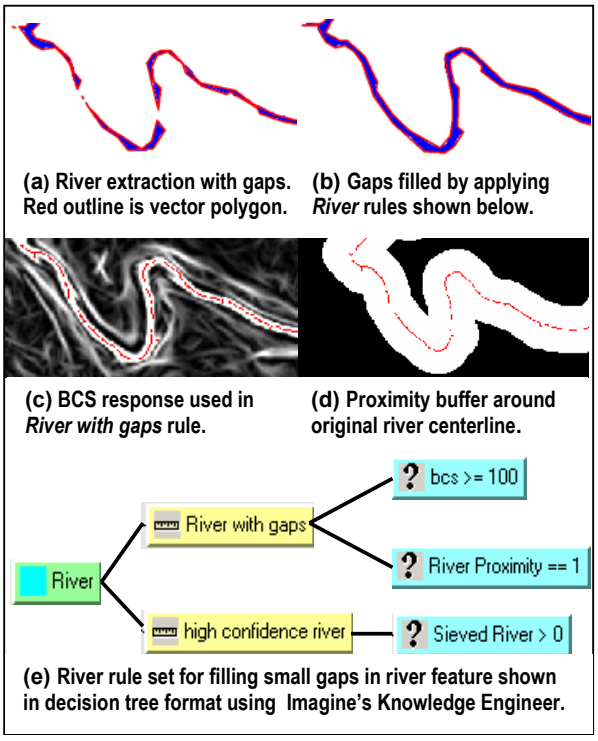

Fig. 14. Filling Small Gaps in a *River* Extraction. A Boundary Contour image from the Neural *Fusion Context Features* module is tested in the *River with Gaps* rule.

The most challenging aspect of the rule-based post-mining classification approach is creating the initial set of image variables (or conditions) to feed to the rule set. Depending on the feature extraction task, the difficulty can range from "simple", with the images being readily available outputs from *Image Mining*, to "difficult", requiring advanced image processing techniques using unconventional imagery like the internal byproducts of

the *Neural Fusion* learning system. We saw in the submerged vegetation example (Fig. 13), that image products were easily realized. The feature outline is available from *detection filtering*, and the proximity buffer is a supported *Imagine* operator. Producing images to fill small river gaps was also fairly easy, with the caveat that gaps could be only as large as the expected width of the river. This restriction is a consequence of using the proximity buffer to the centerline of the river search result. A more advanced method for extending the centerline of a river extraction into a gap is required when the gap is larger, in the case of obscurations, for example. Using orientation information from the BCS processing, a directional proximity buffer can be created that provides a better fit to the river feature over greater gap distances. As the gap gets larger, an omni-directional buffer extends in all directions, causing the buffer to widen from the centerline. This results in bulges in the final classification that require further filtering. This is exemplified in the rule set given in Figure 15 for the case of extracting *river* through occlusions. The result for *river* occluded by clouds is shown in Figure 16. This result shows a slight bulge where the river is filled in through the cloud due to an oversized proximity buffer.

```
IF  river agent detects with high confidence
    THEN  pixel is river

IF  pixel is river occluded by cloud
    THEN pixel is river

IF  pixel is river occluded by tree
    THEN pixel is river

IF  pixel is within river buffer zone
    AND   high confidence cloud agent
    THEN  pixel is river occluded by cloud

IF  pixel is within river buffer zone
    AND   high confidence tree agent

IF  river agent detects with low confidence
    AND   pixel is within river buffer zone
    AND   pixel is on a strong Boundary
          Contour System (BCS) response
    THEN  pixel is river

IF  proximity to river centerline is TRUE
    THEN  pixel is within river buffer zone
```

Fig 15. Rule set for extracting *river* through occlusions by clouds and overhanging trees.

With the tools for rule-based reasoning available in COTS software, and domain expertise from the expert geospatial analyst, our focus is on designing rule sets for frequently extracted features and extending image processing capabilities to satisfy those rule sets. Take the vector product for crop patches shown in Figure 17, for example. *Neural Fusion Image Mining* was conducted on Hyperion hyperspectral imagery for healthy crop patches. *Detection filtering* rules were applied to produce a vector result that most closely matches what an image analyst would produce with manual extraction. The analyst provided simple guidelines for crop patch extractions, such as combining adjacent areas of the same crop type into one polygon and filling areas that are small in comparison to the patch size. We created a custom detection filtering (morphology) rule set for crop patches to emulate the analyst's guidelines. The resulting vector product matches the analyst's manual extraction very closely.

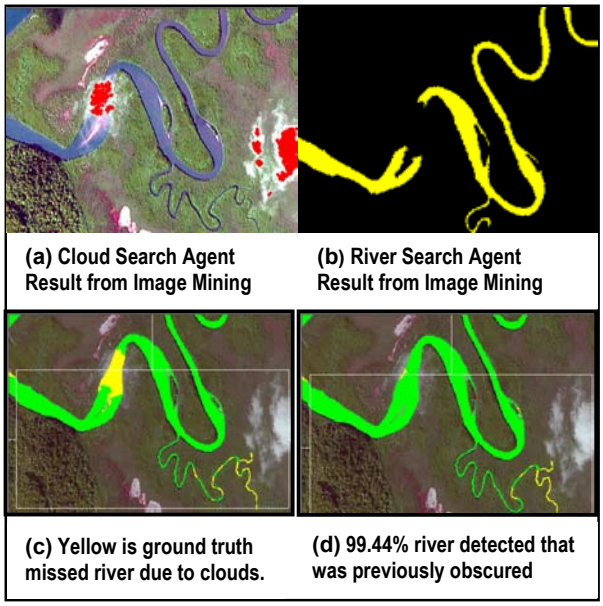

Fig. 16. River extraction through cloud obscurants. (IKONOS imagery provided by NGA.)

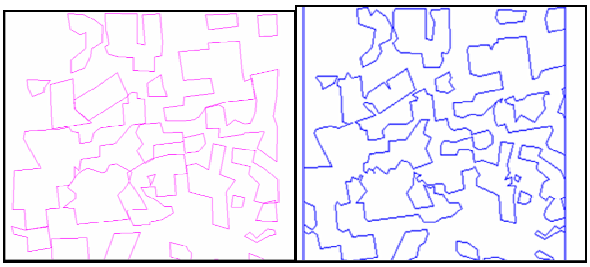

Fig. 17. Comparison between manual (*Left*) and *Neural Fusion* AFE (*Right)* crop patch extractions showing qualitatively very similar results.

6 Summary

We have presented a rule-based post-mining extension to *Neural Fusion*, using tools available within the same *Imagine* COTS environment where *Neural Fusion* is currently integrated. The extension provides a higher-level of reasoning about *Image Mining* results, taking into consideration the relationships between different local features, guidelines for capture conditions, and analyst know-how. We have demonstrated this rule-based reasoning approach by extracting compound features, incorporating capture conditions into the rule set, and creating vector products that more closely resemble those of the human analyst. We applied *Imagine's* Knowledge Engineer tool to translate feature-specific rule sets into decision trees for post-mining classification, and relied mostly on search results and existing capabilities of the *Neural Fusion* system to steer the decisions. We have demonstrated that context-sensitive feature extraction can be achieved through a combination of image fusion, local

image mining, and rule-based geospatial reasoning, as embodied in ALPHATECH's *Neural Fusion*.

Acknowledgements

This work was supported in part by the U.S. National Imagery and Mapping Agency under contract NMA501-03-C-0004. Note: the National Imagery and Mapping Agency has recently been renamed as the National Geospatial-Intelligence Agency (NGA).

References

[1] D.A. Fay, , R.T. Ivey, A.M. Waxman, *Image Fusion and Mining Tools for a COTS Environment*, in Proceedings of the 6th International Conference on Information Fusion, Australia, 2003.

[2] M. Chiarella, D.A. Fay, A.M. Waxman, R.T. Ivey, N.A. Bomberger, *Multisensor Image Fusion & Mining: From Neural Systems to COTS Software with Applications to Remote Sensing AFE*, in IEEE Workshop on Advances in Techniques for Analysis of Remotely Sensed Data, 2003.

[3] A.M. Waxman, D.A. Fay, B.J. Rhodes, T.S. McKenna, R.T. Ivey, N.A. Bomberger, and V.K. Bykoski, *Information fusion for image analysis: Geospatial foundations for higher-level fusion,* in 5th International Conference on Information Fusion, Annapolis, 2002.

[4] A.M. Waxman, J.G. Verly, D.A. Fay, F. Liu, M.I. Braun, B. Pugliese, W. Ross, and W. Streilein, *A prototype system for 3D color fusion and mining of multisensor/spectral imagery,* in 4th International Conference on Information Fusion, Montreal, 2001.

[5] W. Streilein, A. Waxman, W. Ross, F. Liu, M. Braun, J. Verly, and C.H. Read, *Fused multisensor image mining for feature foundation data,* in 3rd International Conference on Information Fusion, Paris, 2000.

[6] D.A. Fay, A.M. Waxman, M. Aguilar, D.B. Ireland, J.P. Racamato, W.D. Ross, W.W. Streilein, and M.I. Braun, *Fusion of multisensor imagery for night vision: Color visualization, target learning and search,* in 3rd International Conference on Information Fusion, Paris, 2000. Also see (same authors) *Fusion of 2-/3-/4-Sensor Imagery for Visualization, Target Learning and Search,* in Proceeds of the SPIE Conference on Enhanced and Synthetic Vision 2000, **SPIE-4023**, pp.106-115, 2000.

[7] D.A. Fay, A.M. Waxman, J.G. Verly, M.I. Braun, J.P. Racamato, and C. Frost, *Fusion of visible, infrared and 3D LADAR imagery,* in 4th International Conference on Information Fusion, Montreal, 2001. Also see (same authors) *Fusion of Multisensor Passive and Active 3D Imagery,* in Proceeds of the SPIE Conference on Enhanced and Synthetic Vision 2001, **SPIE-4363**, pp. 219-230, 2001.

[8] E.A. Newman and P.H. Hartline, *The infrared vision of snakes,* Scientific American, **246** (March), 116-127, 1982.

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

[10] P. Gouras, *Color vision,* Chapter 31 in **Principles of Neural Science** (E.R. Kandel, J.H. Schwartz and T.M. Jessell, editors), pp. 467-480, New York: Elsevier Science Publishers, 1991.

[11] S. Grossberg, Nonlinear neural networks: Principles, mechanisms and architectures, Neural Networks, **1**, 17-61, 1988.

[12] D.S. Levine, **Introduction to Neural and Cognitive Modeling**, 2nd edition, Mahwah, NJ: Lawrence Erlbaum Associates, 2000.

[13] J.G. Daugman, Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters, J. Optical Society of America A, **2**, pp. 1160-1169, 1985.

[14] S. Grossberg and E. Mingolla, Neural dynamics of perceptual grouping: Textures, boundaries and emergent segmentations, Perception & Psychophysics, **38**, 141-171, 1985.

[15] A.M. Waxman, M.C. Seibert, A. Gove, D.A. Fay, A.M. Bernardon, C. Lazott, W.R. Steele, and R.K. Cunningham, *Neural processing of targets in visible, multispectral IR and SAR imagery,* Neural Networks, **8**, 1029-1051, 1995.

[16] G.A. Carpenter, S. Grossberg, and D.B. Rosen, Fuzzy ART: Fast stable learning and categorization of analog patterns by an adaptive resonance system, Neural Networks, **4**, 759-771, 1991.

[17] G.A. Carpenter, S. Grossberg, N. Markuzon, J.H. Reynolds, and D.B. Rosen, *Fuzzy ARTMAP: A neural network architecture for incremental supervised learning of analog multidimensional maps,* IEEE Transactions on Neural Networks, **3**, pp. 698-713, 1992.

[18] T. Kasuba, *Simplified Fuzzy ARTMAP*, AI Expert, pp. 18-25, Nov 1993.

[19] Lauriere, *Expert Systems*, Chapter 7 in **Problem-Solving and Artificial Intelligence**, pp.313-388, Englewood Cliffs, NJ: Prentice Hall, 1990.

[20] D. Patterson, **Introduction to Artificial Intelligence & Expert Systems**, Englewood Cliffs, NJ: Prentice Hall, 1990.