AFRL-SR-AR-TR-09-0083

REPORT DOCUMENTATION PAGE

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
28-02-2008	Final Technical Report	15 APR 2006 - 30 NOV 2008

4. TITLE AND SUBTITLE
INFORMATION AND MOTION PATTERN LEARNING AND ANALYSIS USING NEURAL TECHNIQUES

5a. CONTRACT NUMBER
FA9550-06-C-0018

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Brad RHODES, Neil BOMBERGER, Majid ZANDIPOUR, Denis GARAGIC, James DANKERT

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
BAE Systems
Advanced Information Technologies
6 New England Executive Park
Burlington, MA 01803

8. PERFORMING ORGANIZATION REPORT NUMBER
Report No. TR-2392

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research/NL
875 North Randolph Street
Suite 325, Rm 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Distribution A – Approved for public release

13. SUPPLEMENTARY NOTES
Nil

14. ABSTRACT
This work addressed the development and application of neural models for higher-level information fusion at Levels 2+/3 according to the JDL Data Fusion Group Process Model. We explored several new concepts based on insights from neural processing, learning, and representation. Building on some initial prior work under AFOSR sponsorship, we continued investigation of mechanisms to rapidly and incrementally learn models of normal behavior exhibited by moving tracked entities. These models form the basis of anomaly detection (as deviations from normal) and prediction of future behavior. Our approaches require no ground truth or operator input, but the learned models can be refined through operator feedback if such is available. To investigate the broader applicability of the approaches we discovered in this program, we investigated their efficacy with respect to a distinct domain - that of complex object recognition. We successfully demonstrated that the representations produced by our learning techniques could effectively recognize objects on the basis of their constituent parts. This technical report summarizes progress during the period 15 April 2006 - 30 November 2008.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					Brad RHODES
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER *(include area code)*
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL	10	781-262-4206

INFORMATION & MOTION PATTERN LEARNING & ANALYSIS USING NEURAL TECHNIQUES

Final Technical Report: 15 April 2006 – 30 November 2008

Brad Rhodes, Neil Bomberger, Majid Zandipour, Denis Garagic, James Dankert
Adaptive Reasoning Technologies Directorate
Fusion Technology and Systems Division
BAE Systems, Advanced Information Technologies

28 February 2009

ABSTRACT

This work addressed the development and application of neural models for higher-level information fusion at Levels 2+/3 according to the JDL Data Fusion Group Process Model. We explored several new concepts based on insights from neural processing, learning, and representation. Building on some initial prior work under AFOSR sponsorship, we continued investigation of mechanisms to rapidly and incrementally learn models of normal behavior exhibited by moving tracked entities. These models form the basis of anomaly detection (as deviations from normal) and prediction of future behavior. Our approaches require no ground truth or operator input, but the learned models can be refined through operator feedback if such is available. To investigate the broader applicability of the approaches we discovered in this program, we investigated their efficacy with respect to a distinct domain – that of complex object recognition. We successfully demonstrated that the representations produced by our learning techniques could effectively recognize objects on the basis of their constituent parts. This technical report summarizes progress during the period 15 April 2006 – 30 November 2008.

Keywords: Higher-Level Information Fusion, Situation Assessment, Neural Networks, Associative Learning, Knowledge Structure Discovery and Reasoning, Activity Pattern Learning, Anomaly Detection, Behavior Prediction

1 PROGRAM SUMMARY

The objective of this research was to develop cognitively-inspired methods of creating, extending, and exploiting models for situation awareness from data collected across time, location, and multiple sources. This program addressed the development and application of neural models of information fusion at Levels 2+/3 according to the *JDL Data Fusion Group Process Model* (Figure 1) [1, 2]. We investigated multiple domains simultaneously in order to develop new capabilities to address critical military needs, and to be able to work toward generalized methods applicable across multiple problem areas. The challenge was to develop approaches that can fuse diverse data sets and knowledge, succeed in assessing situations, generate expectations of what follows, and support reasoning activities. Our approach emphasizes neural models of information processing, representation, learning, and classification and applies methods of associative learning to address these challenges.

BAE SYSTEMS Advanced Information Technologies
6 New England Executive Park, Burlington, MA 01803, USA
Telephone (781) 273-3388

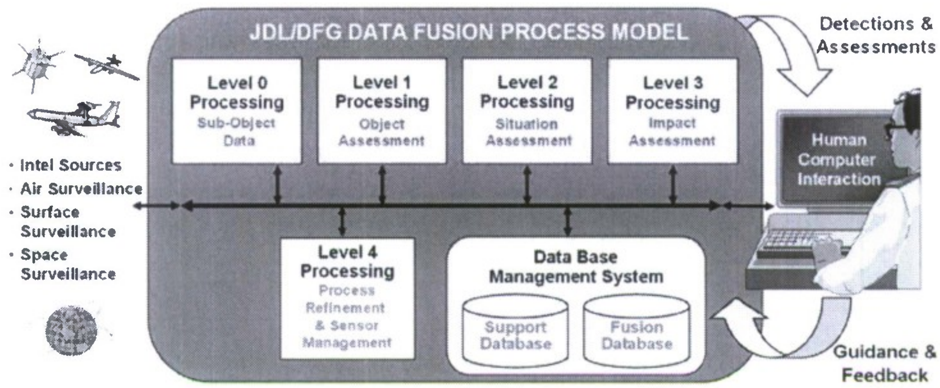

Figure 1. The data fusion process model, as developed by the Joint Directors of Laboratories / Data Fusion Group, distinguishes between multiple levels of data and information fusion (together with sensor management), benefiting from human-in-the-loop guidance as necessary.

Scientifically, our neural-based approach is unique and exciting because it does not rely on statistical assumptions about the observations, the domain, or the representation. This allows us to learn from sparse data, with or without analyst guidance, where the underlying representation must evolve to accommodate a non-stationary real-world environment, yet still maintain stability of the learned model as well as a continuously improving recognition capability. Knowledge models have node-link structures that map well to a neural modeling approach.

This program addressed two major thrusts, leveraged the results from these thrusts for multiple applications, and aimed to transfer the developed technology to relevant transition partners. The main application was the learning of high-level representations of motion pattern behavior to enhance situation monitoring and awareness for purposes of anomaly detection and activity prediction. The first thrust focused on learning of models of normal motion behavior, deviations from which could be detected as anomalies. The second thrust learned motion-based transition models to represent temporal links between motion states, enabling prediction future location from current state. The algorithms developed or extended in these thrusts were then applied to knowledge modeling that supports reasoning applications for higher-level fusion.

Throughout our research efforts, we have transferred technology to a number of agencies and programs. This was achieved through participation in workshops, direct visits and presentations, and transfer of algorithms developed in support of this research to DARPA, DHS, and more recently AFRL programs. The further development of the algorithms under the DARPA programs has supported successful negotiation of Phase I and II Go/NoGo gate criteria, thus providing testament to the technical feasibility and utility of both the general neural inspiration approach to higher-level information fusion and the specific algorithms and representations discovered.

2 SUMMARY OF PROGRAM RESULTS

Our approach builds upon a substantial body of neural modeling and technology in image fusion and mining that we have developed under other programs. The following sections outline the results from each of the major thrusts of the overall program. Each section provides references to the actual publications that present the work in more detail.

2.1 Motion Pattern Learning

Motion pattern learning was one of the research areas that we pursued under our previous AFOSR-funded research program (Neural Methods for Imagery, GMTI, and Information Fusion). Under the IMPLANT (Information & Motion Pattern Learning & Analysis using Neural Techniques) program being reported on here, we built on this research to further enhance these motion pattern learning capabilities. The enhanced capabilities from this AFOSR-funded research were successfully applied to multiple domains (e.g., maritime domain awareness, terrestrial dismount and vehicle monitoring) in order to improve operator situational awareness. These transitions were performed under several different programs, including the DHS-funded USCG Automated Scene Understanding program (SeeCoast), the DARPA-funded Predictive Analysis for Naval Deployment Activities (PANDA) program, and the DARPA/JIEDDO-funded VADER Exploitation Ground Station (VEGS) program.

Normal patterns of motion can be learned for either individual or classes of moving objects. Under our approach, motion patterns can be categorized as either simple motion events (snapshots) or more complex behaviors (e.g., temporal sequences). We have addressed learning at both the simple and complex motion event levels.

For the simple event level, we have applied online neural learning methods, derived from the Adaptive Resonance Theory (ART)-based [3] classifier, to construct representations of normal vessel activity in order to detect anomalous activity (i.e., activity that is sufficiently deviant from the patterns of normal behavior). In the final report for our previous AFOSR-funded program, we described our first research efforts to apply these methods for learning at the simple event level [4] based on vessel track data from the maritime domain via Automatic Identification System (AIS) records). In this work, the models were learned in a two-dimensional feature space consisting of 1) a spatial dimension representing proximity to pre-defined spatial regions and a kinematic dimension (e.g., speed).

Since then, we have made several advances in our approach due to research performed under the IMPLANT program. Using maritime AIS data, we have extended our ART-based learning approach to higher dimensions representing full vessel kinematics, including combinations of 2-D spatial position and 2-D velocity (speed and course). This provides a richer representation of normal vessel behavior, and doesn't require the scene to be seeded with pre-specified spatial regions or markers. In addition, we extended our approach to learn context-sensitive models. What is normal may differ between contexts—such as class of vessel, weather conditions, or tidal status—so we extended our learning approach such that it is capable of discovering normalcy for different contexts. Utilizing the speed and performance of ART-based clustering, we developed interactive displays that allow an operator/analyst to help teach the model via simple point and click actions. These advances have been described in [5], and have been applied to develop the ART-Based Clustering (ABC) engine, which was applied to domestic port monitoring for the DHS-funded USCG Automated Scene Understanding (SeeCoast) program, as reported in [6, 7, 8].

In [9] we described further advances to our approach made possible by IMPLANT research. (This paper won the 2007 Wohl Award for the best technical paper by a BAE Systems AIT employee.) In order to improve the deviation alerting accuracy of our normalcy models, we extended the neural network classifier described above to form a hybrid neural/statistical learning method. For small amounts of training data, ART-based clustering is used to learn model categories. When more data is available, learning transitions to a statistical-based approach which incrementally constructs a multidimensional Gaussian (hyper-ellipsoid) model of each category that is relatively insensitive to outliers and learns the normal pattern of behavior independent of the feature dimensions comprising the learning hyper-space. When a new data point falls into a particular category, the network updates its parameters adaptively to the incoming data, and provides an accurate measure of normal/anomalous behavior. When a new data point is outside of all learned categories, then a new category is formed. During classification, the network reports the distance between the data point to its closest category. If this distance is not within the predefined settable threshold, then that point is reported as a deviation from normalcy. The maximum size of each category hyper-ellipsoid is also a predefined location- and dimension-dependent variable, which controls the fineness and size of each category. This allows the learning sensitivity to be defined as a function of spatial location so that, for example, high spatial resolution models are learned in areas (such as ports) with tight patterns of activity, while low spatial resolution models are learned in areas (such as oceans) with looser patterns of activity. This technology was successfully transitioned to the DARPA-funded PANDA maritime domain awareness program, which monitors behavior patterns thousands of commercial shipping vessels on a global scale. Due to advances in our learning approach made possible by IMPLANT, PANDA successfully achieved its Phase I and Phase II Go/No-go gate criteria for behavioral anomaly detection performance.

Under IMPLANT we have also conducted more recent explorations into multi-scale learning [10] and an incremental statistical-based learning approach called Gaussian Adaptive Mixtures [11]. The multi-scale learning described in [10] extends the work described in [9] by not only learning the structure of the model of normal motion, but also learning the appropriate model scale at a given point in time. This is achieved by applying the idea of *scale space* to learn models at a range of scales, and tracking an evidence function at each of these scales. The goal is to use the finest-resolution model supported by the available training data. As more training data is received, the value of the evidence function increases at each scale. When the evidence passes a threshold at a given scale, the active model scale passes from a coarser scale model to a finer scale model. This allows a coarse model to be used when there is sparse training data, and then gradually refine to higher resolution models as more data is received. Initial research into this approach has yielded initially promising results, but more work is required to improve the evidence function used, and to allow scale switching to occur piece-wise across a model, as needed.

The Gaussian Adaptive Mixtures approach described in [11] improves upon our earlier work by more integrating the best features of the neural- and statistical-based learning approaches, and improving the statistical-based learning approach by introducing multidimensional probability density components to represent class density using an adaptive mixture of such components. The number of components in the adaptive mixture algorithm, as well as the values of the parameters of the density components, is estimated from the data. The network utilizes a recursive version of the Expectation Maximization (EM) algorithm to minimize the Kullback-Leibler information metric by means of stochastic approximation combined with a rule for creation of new components. Learning occurs incrementally in order to allow the system to take advantage of increasing amounts of data without having to take the system offline periodically to update models. This approach retains the advantages of our prior approach (incremental online learning, accuracy of anomaly detection) while more effectively rejecting outlier data from model categories during learning. More work is required to improve learning for this approach in the case of sparse training data.

For learning of complex behaviors, we use an approach that applies research in neural associative learning of temporal sequences. The basic research that underlies this method is described in [12]. Representing complex behaviors in this way affords a predictive capability. In the final report for our previous AFOSR research program, we described some preliminary research in which we applied neural associative learning to learn to predict future vessel location from current vessel behavior [5]. In that work, online learning produced weighted links between position/velocity states that are used to estimate vessel transitions over pre-specified intervals (e.g., 15 minutes). The limitation of this work is that the prediction interval must be pre-specified, and can only take on a single value. In addition, predicted spatial location in this work was only represented at a single spatial scale, whether the vessel is in a port, in a coastal area, or in the ocean.

In more recent work funded by IMPLANT, we have demonstrated the capability to learn transition sequence models across different spatial scales, as appropriate for the local spatial region [13, 14]. We have demonstrated that this capability affords a more accurate predictive capability when the predictive spatial scale is tuned to the local region. For the PANDA program, we have successfully transitioned this capability for learning multi-scale transition sequence models to the problem of predicting vessel location in the global maritime domain. In addition, we have extended the capability to enable prediction over a specified range of time intervals, rather than a single time interval. The PANDA program passed the Phase I and Phase II Go/No-go gate criteria for prediction, demonstrating improved prediction accuracy over time and out-performing predictions calculated through dead reckoning.

The research on motion pattern learning described above is currently being applied to the DARPA-funded Video and Image Retrieval and Analysis Tool (VIRAT) program. The objective of VIRAT is to develop a system that enable operators to select a video clip in which a motion behavior (e.g., a person walking, or a car executing a U-turn) is being exhibited. The system then searches a video database and attempts to locate clips in which similar motion behaviors are present. This is a new application of our motion learning approaches and requires extensions of the transition sequence learning to higher-order temporal sequences.

In addition, these learning methods developed under IMPLANT form the basis of our approach for a new AFRL-funded SSA (Space Situation Awareness) program. Our effort is called Multi-INT Fusion For Space Situation Awareness (MIFFSSA) and aims to apply these methods to learn models of normal behavior in multiple INT domains as well as correlations between them. These normalcy models will then used to detect anomalous behavior that may predict an attack on U.S. space assets.

2.2 Knowledge Modeling and Reasoning

Knowledge modeling to support reasoning was another main area of focus under IMPLANT. Our approach for this research area is based on extending the hyper-elliptical normalcy learning [9] and neural associative learning methods [12] that we have applied to the motion pattern learning problem domain, as described above. Methods were developed that applied to higher level fusion and decision making, including knowledge discovery and spatial context reasoning to enhance situational awareness. This work was successfully applied to the DARPA-funded Urban Reasoning and Geospatial ExploitatioN Technology (URGENT) program.

Our first approach was an algorithm based on probabilistic associative learning and hyper-elliptical category algorithms [15]. This algorithm learned associations between scene and complex objects and their primitive components with and/or without *a priori* information. Furthermore, the spatial relationships between the simple components and their probabilities were learned incrementally.

The associative learning algorithm used a biologically-inspired approach based on a Hebbian learning rule [5, 12, 13, 14]. During training the system learned connections between simple objects, complex objects comprised of simple objects, and scenes. By using a data-dependent learning rate the connection strengths were able to track the conditional probabilities between the different objects. Hence objects that had strong connections occurred together with high probability and objects that had weak connections occurred together with low probability. This learning approach has a number of attractive properties. First, more frequent combinations of simple and scene objects are rapidly learned, as indicated by stronger connections. Second, random/infrequent combinations will cause learning when they occur, but will also be unlearned through weight decay when they do no occur. This provides the system with some ability for noise tolerance. Third, the system is also able to maintain multiple sets of models for different contexts.

To learn context-sensitive models of spatial relationships among simple objects observed in a scene, we leveraged the neural network classifier that incrementally constructs a multi-dimensional Gaussian (hyper-ellipsoid) model of each object's spatial map of spatial relationships to all other objects in a scene [9]. The Hyper-Elliptical Learning and Matching algorithm (HELM) learns spatial relationships between objects in an object-centric frame of reference that is relatively insensitive to outliers.

The two learning/representation elements employed for knowledge modeling can be trained incrementally and simultaneously using ground truth data. During the detection process the associative learning algorithm is used to select an initial set of possible scenes based on the simple objects that are observed. HELM is then used to evaluate the spatial relationships between the objects. Hence, both the co-occurrence and spatial relationships between objects are used to extract higher-level knowledge about the scene.

Building on this basic approach we developed a hybrid reasoning approach that leverages the results from our state-of-the-art approaches to scene segmentation and object recognition, by combining extracted detections and features at multiple spatial and specificity levels through complementary use of associative learning and Bayesian network paradigms [16]. This neurally-inspired learning technique fuses low-level features, identified objects, high-level context, and spatial constraints to more accurately determine the nature of specific scene level targets. The inputs to the system consist of the outputs from a set of individual object detectors. The hybrid reasoning engine uses the outputs from these individual detectors to produce a belief map of the scene and feedback that the individual object detectors can use for tuning.

The first component of this contextual reasoning algorithm is the associative learning algorithm that was previously described. The output from this algorithm is taken as a belief map of possible scenes given the set of simple objects that are present. The belief map unrolls into the second component of our hybrid contextual reasoning algorithm, a Bayesian network, which is utilized to initialize the learning of spatial relationships between constituent objects over the space of possible relationship networks using standard structural learning algorithms. Since a Bayesian network constitutes a complete probabilistic model of the variables in a domain, the network contains the information needed to answer all probabilistic queries about these variables.

One of the benefits of this top-down reasoning strategy, which utilizes scene and spatial context information, is to provide a regularization mechanism to reduce the amount of ambiguity and resulting false alarms (misclassifications) associated with classifying objects/scenes individually. Furthermore, this knowledge discovery is done without any *a priori* information. We were able to demonstrate detection accuracy improvements when our approach reasons over outputs from individual object detectors compared to the accuracy obtained when only the outputs from the individual detectors are used.

3 ACCOMPLISHMENTS, PERSONNEL, & PUBLICATIONS

3.1 Accomplishments

- Modified neural associative incremental learning mechanism to learn links – as conditional probabilities – between co-occurring data elements for form hierarchical knowledge structures.
- Enhanced on-the-fly learning of normalcy models of tracked entity behavior to capture cluster statistics.
- Demonstrated detection of anomalous behavior using maritime (AIS) data in various port regions (New York, NY, Miami, FL, Portsmouth, VA).

- Developed associative learning techniques to enable prediction of future vessel location based on current behavior.
- Demonstrated utility of learning approaches in a new domain – complex object recognition
- Developed Java viewer for visualization of associatively learned knowledge structures
- Developed visualization capabilities for learned normalcy models – used Matlab and Google Earth
- Presented paper at SPIE conference in Orlando, FL (April 2006) on motion pattern learning of vessels in port areas (Brad Rhodes).
- Presented paper at the 2006 MSS National Symposium on Sensor and Data Fusion, McLean, VA, USA (June, 2004) on learning of normal vessel behavior and anomaly detection for port scene understanding (Neil Bomberger).
- Presented paper at Fusion 2006 in Florence, Italy (July 2006) on prediction of future vessel location in port and coastal regions (Neil Bomberger).
- Presented paper at AFOSR Workshop on Information Fusion at Fusion 2006 in Florence, Italy (July 2006) on IMPLANT program plans and future directions (Brad Rhodes).
- Presented paper at Situation Management (SIMA) Workshop (at MILCOM2006) on motion pattern learning of normalcy models (Brad Rhodes).
- Two papers published in *Information Fusion* for a special issue on *Concurrent Learning and Information Fusion*:
    - *A new approach to higher-level information fusion using associative learning in semantic networks of spiking neurons* (Neil Bomberger, Allen Waxman, Brad Rhodes, and Nathan Sheldon)
    - *Taxonomic knowledge structure discovery from imagery-based data using the neural associative incremental learning (NAIL) algorithm* (Brad Rhodes)
- Presented paper at SPIE conference in Orlando, FL (April 2007) on automation of port scene understanding using motion pattern learning and assessment (Brad Rhodes).
- Presented paper at Fusion 2007 in Quebec City, Quebec, Canada (July 2007) on probabilistic associative learning (Brad Rhodes).
- Participated in Issues in Higher-Level Information Fusion panel at Fusion 2007 in Quebec City, Quebec, Canada (July 2007) on biologically-inspired approaches to higher level fusion (Brad Rhodes).
- Presented paper at Situation Management (SIMA) Workshop (at MILCOM2007) on motion pattern analysis (Brad Rhodes).
- Published paper at SPIE conference in Orlando, FL (April 2008) on synthesis of Fusion 2007 IF issues panel discussion.
- Presented paper at Fusion 2008 in Cologne, Germany (July 2008) on probabilistic associative learning at multiple scales (Brad Rhodes).
- Presented paper at Fusion 2008 in Cologne, Germany (July 2008) on knowledge modeling of simple object relationships to reason on the existence of complex objects (Brad Rhodes).
- Presented paper at Situation Management (SIMA) Workshop (at MILCOM2008) on motion normalcy leaning at multiple spatial scales (Brad Rhodes).
- Published chapter in Sensor and Data Fusion book (2009) on biologically inspired approaches to higher level fusion.
- Paper accepted to IEEE ICC2009 conference on adaptive mixtures normalcy modeling.
- Paper accepted to *Intelligent Design Technologies* journal on motion pattern learning and prediction.
- Submitted paper to Fusion 2009 in Seattle, WA on hybrid learning and reasoning.
- Transitioned motion normalcy learning, anomaly detection and prediction technology to the DARPA PANDA program – ultimately deployed to ONI.

- Transitioned enhanced motion pattern normalcy model learning and anomaly detection to HSARPA/USCG SeeCoast automated scene understanding prototype system currently deployed at JHOC, Portsmouth, VA.
- Transitioned knowledge modeling and reasoning technology to DARPA URGENT program
- Activity pattern and associative learning technology formed a key discriminator in winning proposal for DARPA's VIRAT program.
- Motion pattern learning, detection, and prediction technology used as major platform in winning proposal for AFRL's ISSA program.

3.2 People involved in research effort

- Brad Rhodes
- Neil Bomberger
- Allen Waxman
- Michael Seibert
- Majid Zandipour
- Alan Gove
- Denis Garagic
- James Dankert
- Michael Hlasyszyn
- Kyle Whitson
- Jonathan Barron

3.3 Publications

Garagic, D., Zandipour, M., Stolle, F.R., & Rhodes, B.J. (Submitted). Hybrid neuron-Bayesian spatial contextual reasoning for scene context understanding. Submitted to *The 11th International Conference on Information Fusion*, Seattle, WA, USA, July 6–9, 2009.

Rhodes, B.J., Bomberger, N.A., Zandipour, M., Garagic, D., Stolzar, L.H., Dankert, J.R., Waxman, A.M., & Seibert, M. (in press). Automated activity pattern learning and monitoring provide decision support to supervisors of busy environments. To appear in *Intelligent Decision Technologies.*

Garagic, D., Rhodes, B.J., Bomberger, N.A., & Zandipour, M. (in press). Adaptive mixture-based neural network approach for higher-level fusion and automated behavior monitoring. Accepted to *IEEE International Conference on Communications,* Dresden, Germany, June 14–18, 2009.

Rhodes, B.J., Bomberger, N.A., Zandipour, M., Stolzar, L.H., Garagic, D., Dankert, J.R., & Seibert, M. (2009). Anomaly detection & behavior prediction: Higher-level fusion based on computational neuroscientific principles. In *Sensor and Data Fusion.*

Bomberger, N.A., Zandipour, M., & Rhodes, B.J. (2008). Adaptive spatial scale for cognitively-inspired motion pattern learning & analysis algorithms for higher-level fusion and automated scene understanding. In *Proceedings of IEEE MILCOM 2008 Military Communications Conference*, San Diego, CA, USA, November 17-19, 2008.

Zandipour, M., Rhodes, B.J., & Bomberger, N.A. (2008). Probabililstic prediction of vessel motion at multiple spatial scales for maritime situation awareness. In *Proceedings of the 10th International Conference on Information Fusion*, Cologne, Germany, June 30 – July 3, 2008.

Zandipour, M., Rhodes, B.J., & Bomberger, N.A. (2008). COALESCE: A probabilistic ontology-based scene understanding approach. In *Proceedings of the 11th International Conference on Information Fusion*, Cologne, Germany, June 30 – July 3, 2008.

Kadar, I., Bosse, E., Salerno, J., Lambert, D.A., Das, S., Ruspini, E.H., Rhodes, B.J., & Biermann, J. (2008). Results from levels 2/3 fusion implementations: Issues, challenges, retrospectives and perspectives for the future – An annotated perspective. In *Proceedings of SPIE Vol. 6968: Multisensor fusion methodologies and applications V*, Orlando, FL, USA, March 17-20, 2008.

Rhodes, B.J., Bomberger, N.A., Zandipour, M., Waxman, A.M. & Seibert, M. (2007). Cognitively-inspired motion pattern learning & analysis algorithms for higher-level fusion and automated scene understanding. In *Proceedings of IEEE MILCOM 2007 Military Communications Conference*, Orlando, FL, USA, October 29-31, 2007.

Rhodes, B.J. (2007 - Invited paper). Biologically-inspired approaches to higher-level information fusion. In *Proceedings of the 10th International Conference on Information Fusion*, Quebec City, Canada, July 9-12, 2007.

Rhodes, B.J., Bomberger, N.A., & Zandipour, M. (2007). Probabilistic associative learning of vessel motion patterns at multiple spatial scales for maritime situation awareness. In *Proceedings of the 10th International Conference on Information Fusion*, Quebec City, Canada, July 9-12, 2007.

Rhodes, B.J. (2007). Taxonomic knowledge structure discovery from imagery-based data using the Neural Associative Incremental Learning (NAIL) algorithm. *Information Fusion, 8,* 295-315.

Bomberger, N.A., Waxman, A.M., Rhodes, B.J., & Sheldon, N.A. (2007). A new approach to higher-level information fusion using associative learning in semantic networks of spiking neurons. *Information Fusion, 8,* 227-251.

Bomberger, N.A., Rhodes, B.J., Seibert, M., & Waxman, A.M. (2006). Associative learning of vessel motion patterns for maritime situation awareness. In *Proceedings of the 9th International Conference on Information Fusion*, Florence, Italy, July 10-13, 2006.

4 CONCLUSION, FUTURE DIRECTIONS

Our successful execution of this research program has resulted in:

- Novel, neurally-based techniques for higher-level information fusion (that have been successfully applied via other programs);
- Numerous publications and presentations in scientific and application-specific forums;
- Transition of several outcomes to Air Force and other partners; and
- Identification of additional basic research questions that will continue innovative development stemming from the outcomes of this program.

As indicated in the accomplishments above, we have discovered and enhanced a number of neurally-inspired techniques to address important higher-level information fusion problems and have shown them to be applicable across domains. Moreover, we have been extremely successful in applying those techniques in applied research domains and have passed a number of (always) challenging DARPA Go/NoGo phase gates. Thus, our basic research work is very efficiently directed and produces strong results that are both disseminated broadly to the information fusion research community and transferred to the DoD community via our strong performance when applying these techniques to DARPA-hard information fusion problems. The success of our work has led to a set of capabilities that form a strong foundation for continued basic research. We have identified a set of research questions arising from this program that we believe are important next steps from basic and applied perspectives. We have prepared and submitted to AFOSR a comprehensive proposal for support to conduct a research program focused on these questions. We will continue to seek opportunities to further enhance and apply the results of our research for the benefit of the DoD community. We believe that continued support from AFOSR for our research efforts would be of great benefit to this community.

5 REFERENCES

1. White, F. (1987). Data Fusion Lexicon. *Joint Directors of Laboratories, Technical Panel for C3, Data Fusion Sub-Panel*, Naval Ocean Systems Center, San Diego.
2. Steinberg, A., Bowman, C., & White, F. (1999). Revisions to the JDL Data Fusion Model. In *Proceedings of SPIE Sensor Fusion: Architectures, Algorithms, and Applications III*, 430–441.
3. Carpenter, G.A., Grossberg, S., Markuzon, N., Reynolds, J.H., & Rosen, D.B. (1992). Fuzzy ARTMAP: A neural network architecture for incremental supervised learning of analog multidimensional maps. *IEEE Transactions on Neural Networks*, *3*, 698–713.
4. Rhodes, B.J., Bomberger, N.A., Seibert, M., & Waxman, A.M. (2005). Maritime situation monitoring and awareness using learning mechanisms. In *Proceedings of IEEE MILCOM 2005 Conference*, Atlantic City, NJ, USA, 17–20 October, 2005.
5. Bomberger, N.A., Rhodes, B.J., Seibert, M., & Waxman, A.M. (2006). Associative learning of vessel motion patterns for maritime situation awareness. In *Proceedings of the 9th International Conference on Information Fusion*, Florence, Italy, July 10-13, 2006.
6. Seibert, M., Rhodes, B.J., Bomberger, N.A., Beane, P.O., Sroka, J.J., Kogel, W., Kreamer, W., Stauffer, C., Kirschner, L., Chalom, E., Bosse, M., & Tillson, R. (2006). SeeCoast port surveillance. Invited paper for *Proceedings of SPIE Vol. 6204: Photonics for Port and Harbor Security II*, Orlando, FL, USA, April 18-19, 2006.
7. Rhodes, B.J., Bomberger, N.A., Santos, B.C., & Seibert, M. (2006). SeeCoast maritime domain awareness system. In *Proceedings of 2006 MSS National Symposium on Sensor and Data Fusion*, McLean, VA, USA, June 6-8, 2006.
8. Rhodes, B.J., Bomberger, N.A., Seibert, M., & Waxman, A.M. (2006). SeeCoast: Automated port scene understanding facilitated by normalcy learning. In *Proceedings of IEEE MILCOM 2006 Military Communications Conference*, Washington, DC, USA, October 23-25, 2006.
9. Rhodes, B.J., Bomberger, N.A., Zandipour, M., Waxman, A.M. & Seibert, M. (2007). Cognitively-inspired motion pattern learning & analysis algorithms for higher-level fusion and automated scene understanding. In *Proceedings of IEEE MILCOM 2007 Military Communications Conference*, Orlando, FL, USA, October 29-31, 2007.
10. Bomberger, N.A., Zandipour, M., & Rhodes, B.J. (2008). Adaptive spatial scale for cognitively-inspired motion pattern learning & analysis algorithms for higher-level fusion and automated scene understanding. In *Proceedings of IEEE MILCOM 2008 Military Communications Conference*, San Diego, CA, USA, November 17-19, 2008.
11. Garagic, D., Rhodes, B.J., Bomberger, N.A., & Zandipour, M. (in press). Adaptive mixture-based neural network approach for higher-level fusion and automated behavior monitoring. Accepted to *IEEE International Conference on Communications*, Dresden, Germany, June 14–18, 2009.
12. Rhodes, B.J. (2007). Taxonomic knowledge structure discovery from imagery-based data using the Neural Associative Incremental Learning (NAIL) algorithm. *Information Fusion, 8*, 295-315.
13. Rhodes, B.J., Bomberger, N.A., & Zandipour, M. (2007). Probabilistic associative learning of vessel motion patterns at multiple spatial scales for maritime situation awareness. In *Proceedings of the 10th International Conference on Information Fusion*, Quebec City, Canada, July 9-12, 2007.
14. Zandipour, M., Rhodes, B.J., & Bomberger, N.A. (2008). Probabilistic prediction of vessel motion at multiple spatial scales for maritime situation awareness. In *Proceedings of the 10th International Conference on Information Fusion*, Cologne, Germany, June 30 – July 3, 2008.
15. Zandipour, M., Rhodes, B.J., & Bomberger, N.A. (2008). COALESCE: A probabilistic ontology-based scene understanding approach. In *Proceedings of the 11th International Conference on Information Fusion*, Cologne, Germany, June 30 – July 3, 2008.
16. Garagic, D., Zandipour, M., Stolle, F.R., & Rhodes, B.J. (Submitted). Hybrid neuron-Bayesian spatial contextual reasoning for scene context understanding. Submitted to *The 11th International Conference on Information Fusion*, Seattle, WA, USA, July 6–9, 2009.

BAE SYSTEMS Advanced Information Technologies
6 New England Executive Park, Burlington, MA 01803, USA
Telephone (781) 273-3388