have the capability of redefining their own local neighborhood at each iteration. Finally, the causal belief map of the agents is updated from *pbest* and *lbest* by adjusting the structure of the nodes and the strength of the causal beliefs between nodes. Experiments (Figs. 2 and 3) are shown comparing simulations with and without adaptation. The degree of coalition (Fig. 4) is measured as the inverse total variance of the population in the 2D space obtained with an iterative *k*-means clustering procedure.

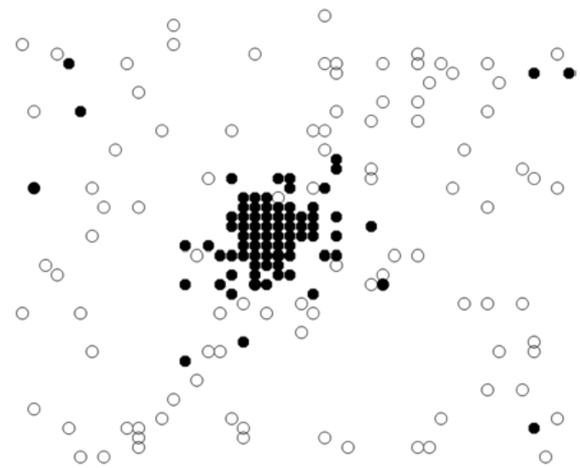

FIGURE 2
Group formation: the filled circles indicate the final state of the simulation after adaptation.

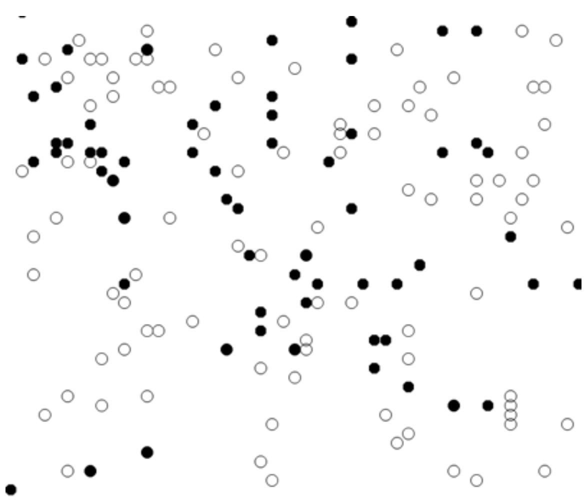

FIGURE 3
Dispersion: the filled circles indicate the final state of the simulation without adaptation.

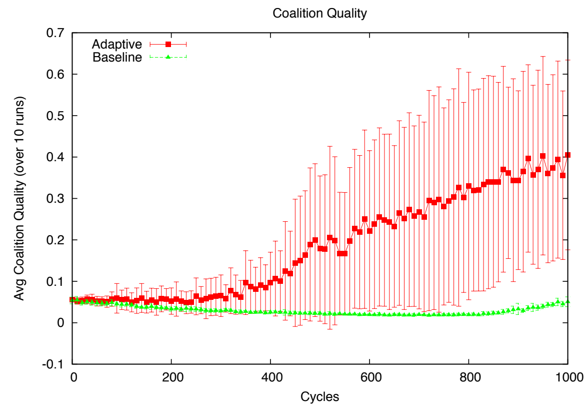

FIGURE 4
Coalition quality comparison.

Summary: This approach has shown how to incorporate belief changes in group formation and provides insight on the mechanism of group formation at a more fundamental level than the stated position and influence of key actors. This approach can be used to model the interactions of agents from different cultures represented by a set of prescriptive rules (e.g., proverbs or narratives). For example, the impact of a foreign presence in a multiethnic society can be modeled, quantified, and evaluated over several time cycles based on the interaction of cognitive agents. Comparisons between initial and final causal belief map variants can provide structural content insights in addition to predictive trends. This agent-based model can be coupled with a discrete-event simulator to synchronize with the timeline of IW war gaming scenarios.

[Sponsored by the NRL Base Program (CNR funded)]

References

[1] "Counterinsurgency," Field Manual No. 3-24, MCWP 3-33.5 (Dept. of the Army, Washington, DC), December 2006.

[2] B. Kosko, *Neural Networks and Fuzzy Systems* (Prentice Hall, Englewood Cliffs, NJ, 1992).

[3] J. Kennedy and R.C. Eberhart, *Swarm Intelligence* (Morgan Kaufmann, San Francisco, CA, 2001).

[4] M. Abramson, "Coalition Formation of Cognitive Agents," in *ICCCD 2008 Proceedings of the Second International Conference on Computational Cultural Dynamics* (AAAI Press, Menlo Park, CA) pp. 2–8.

Maritime Threat Detection

D.W. Aha,[1] K.M. Gupta,[2] and B. Auslander[2]
[1]*Navy Center for Applied Research in Artificial Intelligence*
[2]*Knexus Research Corporation*

Introduction: Detecting and preventing small vessel attacks is crucial for protecting Navy personnel and assets in busy maritime locations, as exemplified by the USS *Cole* bombing and related incidents. While some force protection can be performed unaided, watchstanders increasingly depend on decision aids to combat surveillance data overload. Deployed systems for local maritime surveillance perform perimeter defense, raising alerts when a vessel penetrates an "electronic fence" about a shoreline or asset. Although often effective, this approach cannot detect threats based on analysis of vessel behavior outside of this perimeter, nor reason about intent or coordinated threats.

We claim that machine learning methods, and in particular probabilistic relational networks (PRNs), can acquire models of small vessel behaviors to more accurately predict maritime threats. PRNs compactly encode a distribution in a multidimensional space, model variable independencies, and have well-understood mathematical foundations. While we studied methods for classifying the behavior of small maritime vessels[1] and identifying anomalous tracks,[2] no prior research has applied PRNs to small vessel threat assessment. Therefore, for this task we tested three types of PRNs that vary in the relations they can represent and their methods for learning and inferencing.

Algorithms: We chose hidden Markov models (HMMs), conditional random fields (CRFs), and Markov logic networks (MLNs) because they performed well on many tasks. HMMs are graphs whose nodes denote hidden states and whose links denote state transition probabilities. They model the *joint* distribution $p(y_t, x_t)$, for observation x_t and state y_t at time t, and assume that state transitions depend only on the preceding state and observations depend only on the current state. HMM learning and inferencing is performed using the forward-backward and the Viterbi algorithms, respectively. The basic HMM model cannot easily represent local features and spatial relations. In contrast, CRFs model local features using the *conditional* distribution $p(y \mid x)$ and reason with interdependent features, which exist in this task. We used a gradient descent algorithm to learn CRF parameter settings, and the Viterbi algorithm for inference. CRFs cannot model spatial relations (e.g., the distance between two vessels), whereas MLNs can, by combining first-order logic (FOL) with a probabilistic interpretation to represent states. Unlike FOL, MLNs can model domains where constraint violations have low probability but are not impossible. An MLN is a set of pairs (F_i, w_i) where F_i is a FOL formula and $w_i \in [0,1]$. Together with a set of constants C it defines a Markov network $M_{L,C}$ containing one node per grounding of each predicate in the set of possible groundings L and one feature per grounding of each formula $F_i \in L$. This network can vary widely in shape and size depending on its constraints. We used maximum pseudo log-likelihood estimation for generative training and MC-SAT for inference.

Data: Our data were provided by Spatial Integrated Systems,who conducted exercises during Trident Warrior 2010 in San Diego Harbor, in which two autonomous unmanned sea surface vehicles (USSVs) were tasked with blocking two human-controlled boats from "attacking" a high-value unit (HVU). Data streams from USSV-mounted sensors were collected and fused to create the corpus' tracks, which we cleaned and synchronized. We manually annotated each track instance as *Attacking*, *Cruising*, or *Escaping*, and computed four features per instance (e.g., the *Prior Activity* that a vessel performed in the prior time step, and *In Front of HVU*, which denotes whether the boat is bearing on the HVU). This produced two sets of tracks, each of 53 minutes duration, where an instance corresponds to 10 seconds.

Evaluation and Results: We trained and applied the PRNs to predict, at each instance, whether a boat is attacking the HVU, and used precision, recall, the F_1 measure, and run-time to assess performance. We used a twofold cross-validation protocol and included two baseline algorithms: *Default* predicts that every instance is an attack, while *Perimeter Rule* mimics the perimeter defense strategy. We optimized the window sizes used for the PRNs and the triggering distance for *Perimeter Rule*. Table 1 displays some of the results, which provide initial support for our claim. As expected, *Default* had low precision. *Perimeter Rule* performed well for the first set but not the second set, which contains few *Attacking* instances. MLNs attained the highest F_1 scores for both sets, which may reflect their ability to represent spatial relations and learn weight settings from few training instances. However, they were expensive to train and test, whereas the other PRNs were highly efficient.

Discussion: MLNs were accurate, but required substantial trial and error to create a good rule set.

TABLE 1 — Comparative Results for Predicting Attack Instances

Algorithm	Trained on Set #1			Trained on Set #2				
	Precision	Recall	F_1	Precision	Recall	F_1	Time	
							Training	Test
Default	0.04	1.00	0.07	0.02	1.00	0.03		
Perimeter Rule	0.38	1.00	0.55	0.04	0.37	0.08	N/A	0.3
HMM	0.40	0.46	0.43	0.06	0.40	0.10	1.3	0.4
CRF	0.63	0.11	0.19	0.11	0.57	0.18	3.6	0.3
MLN	0.42	1.00	**0.59**	0.11	1.00	**0.19**	82.0	47.0

Methods for automatically learning structure, and more efficient learning and inferencing algorithms, are warranted for applying MLNs to this task. In the future, we will address the topic of coordinated attacks, study intent and plan recognition techniques, and test our algorithms on board USSVs under real-time conditions.

Acknowledgments: Many thanks to Spatial Integrated Systems for providing the data we used in this investigation and to the Office of Naval Research (ONR 33) for sponsoring this research.

[Sponsored by ONR]

References

[1] K.G. Gupta, D.W. Aha, and P. Moore, "Case-Based Collective Inference for Maritime Object Classification," in *Proceedings of the Eighth International Conference on Case-Based Reasoning* (Springer, Seattle, WA, 2009) pp. 443–449.

[2] B. Auslander, K.M. Gupta, and D.W. Aha, "A Comparative Evaluation of Anomaly Detection Algorithms for Maritime Video Surveillance," in *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense X,* ed. E.M. Carapezza, *Proc. SPIE* **8019** (SPIE, Bellingham, WA, 2011) 801907.

NRL Flight-Tests Autonomous Multi-Target, Multi-User Tracking Capability

B.J. Daniel,[1] M. Wilson,[1] J. Edelberg,[1] S. Frawley,[2] C. Meadows,[3] S. Anderson,[3] T. Johnson,[3] and M. Duncan[1,4]

[1]*Optical Sciences Division*

[2]*Smart Logic, LLC*

[3]*Space Dynamics Laboratory*

[4]*Office of Naval Research*

Introduction: The Naval Research Laboratory has demonstrated an autonomous multisensor motion-tracking and interrogation system that reduces the workload for analysts by automatically finding moving objects, then presenting high-resolution images of those objects with little to no human input. Intelligence, surveillance, and reconnaissance (ISR) assets in the field generate vast amounts of data that can overwhelm human operators and can severely limit an analyst's ability to generate intelligence reports in operationally relevant time frames. This multi-user tracking capability enables the system to manage the collection of imagery without continuous monitoring by a ground or airborne operator, thus requiring fewer personnel and freeing up operational assets. The Office of Naval Research (ONR)-sponsored multisensor motion-tracking and interrogation demonstration leveraged three systems developed under prior ONR programs: N-WAPSS, a wide-area survey sensor; the ground stations, providing sensor control and motion tracking; and EyePod, an interrogation sensor.

Wide-Area Survey Sensor: The midwave infrared (MWIR) Nighttime Wide-Area Persistent Surveillance Sensor (N-WAPSS) was developed with ONR support for the Angel Fire program and then transitioned to the Air Force Blue Devil program. N-WAPSS is a 16-megapixel, large-format camera that captures single frames at 4 Hz and has a step-stare capability with a 1 Hz refresh rate.[1] The 16-megapixel imagery is compressed on board the sensor using an NRL-developed hardware implementation of JPEG2000 compression. The compressed images are sent to the Tracking Ground Station

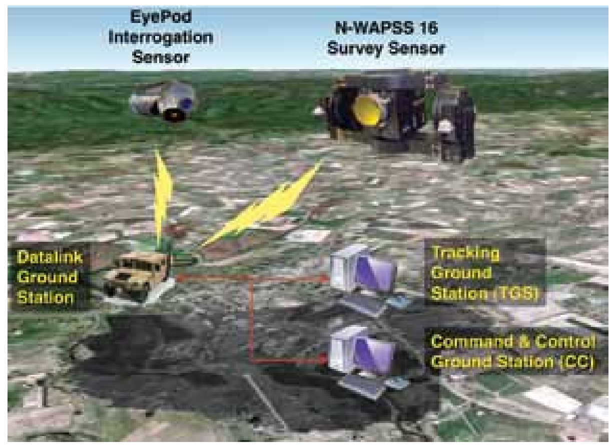

FIGURE 5
Schematic of the sensor network. The N-WAPSS survey sensor downloads compressed images through the data link to the ground stations for tracking and command and control. The control commands are uploaded through the data link to EyePod, the interrogation sensor.

(TGS) via the high-speed Tactical Reachback Extended Communications (TREC) data link provided by the NRL Information Technology Division. See Fig. 5.

Ground Stations: The ground stations are composed of two computers, each with a different task: tracking and command and control. The TGS uses precision geoprojection of the N-WAPSS imagery to detect and track all moving vehicle-sized objects in the field of view in real time[2] (see Fig. 6). The tracks are converted to georeferenced cues and sent via a ground-based network to the Command and Control Ground Station (CCGS). The CCGS manages these cues autonomously and tasks the interrogation sensor to image all selected tracks for target classification and identification.[3] The low-bandwidth CCGS commands are sent to the interrogation sensor via the TREC data link.

Interrogation Sensor: The georeferenced cues are sent to EyePod, a precision jitter-stabilized inter-