

PREDICTIVE ANALYSIS FOR C4ISR

Dr. H. Van Dyke Parunak*
NewVectors LLC
Ann Arbor, MI 48105

Mr. Jacob Crossman
Soar Technology
Ann Arbor, MI 48105

Mr. Sam Hamilton
Lockheed Martin Orincon
San Diego, CA 92121

Dr. William McEneaney
University of California at San Diego
La Jolla, CA 92093

Dr. Wesley A. Milks
Applied Research Associates, Inc.
Orlando, FL 32817

Mr. John Sloan
Alion Science and Technology
Fairfax, VA 20033

Dr. Boris Stilman
Stilman Advanced Strategies
Denver, CO 80202

ABSTRACT

DARPA's Real-time Adversarial Intelligence and Decision-making (RAID) program [1] has demonstrated a promising new capability to predict enemy location and intent in dynamic urban combat environments. This capability may significantly improve the Blue commander's decision processes by increasing his situational awareness and tactical team coordination capabilities. Experimental results obtained over the past two years indicate that a single commander in a simulated urban combat environment assisted by RAID outperforms a 5-person senior staff of military Subject Matter Experts. These results also indicate that RAID predictions and recommendations can improve the mission planning process by providing a previously unavailable level of predictive analysis. This paper will address several team performance factors that are improved by RAID and their impact on the mission planning process, present results from the RAID Experiment 4 (July 2006), and describe a key technology extension that is needed for improving the real-time situational awareness data provided to RAID.

1 WHAT BENEFITS DOES RAID PROVIDE?

RAID integrates intelligence about Red and Blue operations and the state of the environment to provide continuously updated predictions of the current and future state of the battlefield. Multiple predictions are available, and currently include "most dangerous" and "most likely" predicted states. Currently, OPFOR predictions and BLUE COA recommendations are provided for up to two hours into the future for Company sized forces within two minutes on a 3.2 GHz CPU.

RAID can manipulate both descriptive and normative information, both as input and as output. As input, RAID can base its predictions on observations of *past* behavior, declarations of *intended* actions, or combinations of these. For example, in real-time support of a commander, it can integrate the current information state, i.e., a distillation of the ongoing observations of Red activity and a Blue COA into a running prediction of the battlespace. But it can also be employed to monitor the progress of combat between two sides based purely on observations of their behavior.

As output, RAID can produce descriptions of what units *are likely to do*, or suggestions of what units *could do*. Thus it can provide real-time action recommendations to a Blue unmanned combat vehicle, or play out both sides of a scenario based only on given initial conditions.

As currently configured, RAID produces descriptive predictions for Red and normative COA's for Blue. Its near real-time input consists of its own translation of the Blue mission statements or commander's guidance, Blue force tracking, and observed Red behavior.

2 HOW DOES RAID WORK?

RAID consists of three technical components or modules, each focusing on a different aspect of predictive battlespace awareness. In the current configuration (Fig. 1), the Deception Reasoning Module (DRM) processes the sensor feed from the battlefield into likely laydowns of Red units that have, and that have not yet, been detected. The Adversarial Reasoning Module-A (ARM-A) discerns their intentions and extrapolates these into the most likely future behaviors. The Adversarial Reasoning Module-S (ARM-S) computes the most dangerous future Red behaviors, and generates a suggested Blue COA to counter the most dangerous projection of Red.

2.1 Deception Reasoning Module

The Deception Reasoning Module (DRM) estimates the current locations of units based on the interactions of three analytical processes (Fig. 2). Its basic approach is to look forward in time to anticipate likely outcomes and identify the value of these outcomes to both Blue and Red, given their respective missions. Then it reports the laydowns that lead to selected conditions, such as most dangerous or most likely.

DRM's three interacting processes are an Urban Gaming Module, a Deception Analysis Engine, and a Game Theory Module.

The Deception Analysis Engine is the heart of the DRM. It receives clustered observations processed from intelligence assets and uses a Bayesian propagation to identify likely states of the world that would be consistent with these observations. Then it weights different possible outcome states on the basis of risk to formulate its output.

The Deception Analysis Engine draws on the Game Theory Module [2] for two categories of information. First, its Bayesian propagator invokes the Game Theory Module in a prediction loop to explore how a given laydown might evolve in the future. Second, in assessing the risk associated with each laydown, it engages in a state evaluation loop to determine the risk associated with a given set of states.

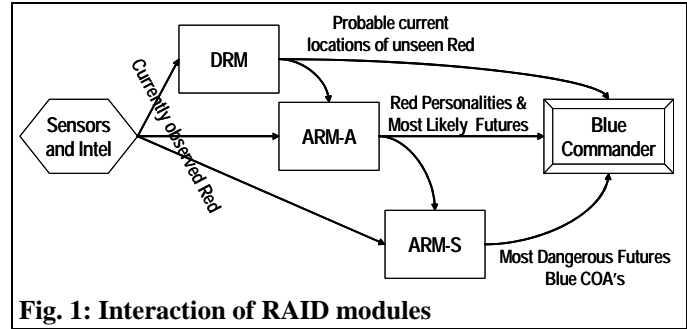


Fig. 1: Interaction of RAID modules

The Game Theory Module in turn draws on the Urban Gaming Model for environmental constraints on warfighters, such as 3-D lines of sight, waypoints, and movement connectivity. The Urban Gaming Model derives this information from a database containing information on terrain and multiple elevation structures.

By applying this analysis to observable enemy activities and intelligence reports, the DRM publishes three laydowns on each cycle (out of many that it explores), including location, strength, and weapon type of each Red unit. It can predict units that have never been seen, as well as likely current locations of units previously sensed but currently invisible, and includes predictions of most probable IED locations that would be most advantageous to Red.

2.2 Adversarial Reasoning Module-A

The Adversarial Reasoning Module-A (ARM-A) learns units' intentions and emotional state from their past behavior, using a statistical analysis of threat regions, a logical analysis of their goals, and a swarming simulation of their interaction with their environment, and produces predictions of their likely behavior. The ARM-A can also ingest historical data that is currently available in the operational environment via typical intelligence reporting.

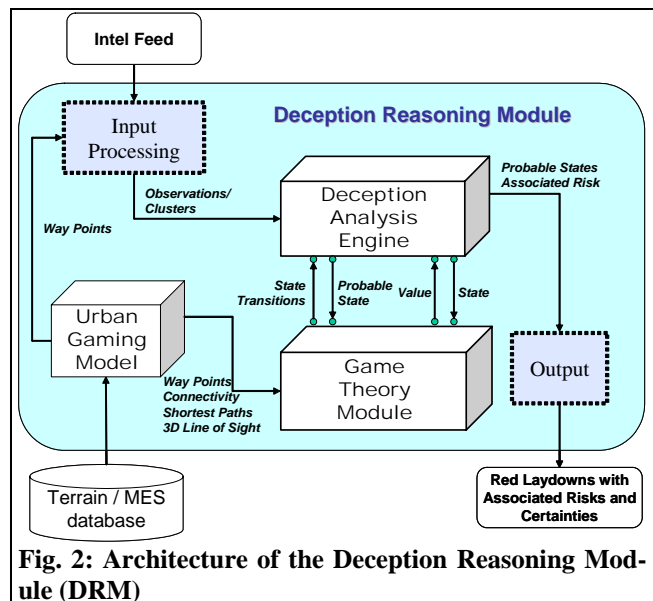
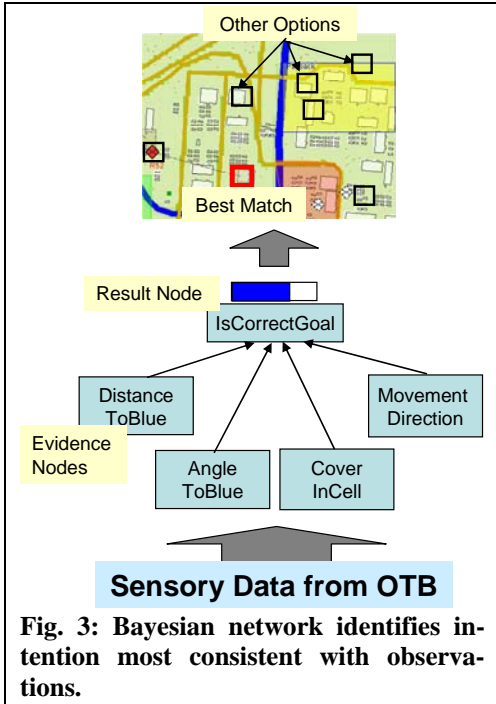


Fig. 2: Architecture of the Deception Reasoning Module (DRM)



ARM-A itself has three components: Statistical Anomaly Detection (SAD), Knowledge-Based Intention Projection (KIP), and Behavior Evolution and Extrapolation (BEE).

KIP [3] estimates a unit's internal intentions from its external behavior. A set of spatial-temporal heuristics suggests possible intentions consistent with each unit's observed behavior. Then a Bayesian belief network, which learns its weights from past battles, evaluates each candidate intention against the observed evidence and identifies the most likely one given the current context (Fig. 3). Thus it estimates where units likely will want to go, consistent with their force's doctrine and RAID's understanding of Red's overall mission.

SAD analyzes the topography, known and likely unit locations, and performs an aggregate statistical analysis to determine the threat level that Red or Blue units will perceive at different locations in the battlespace. It thus estimates where units will want *not* to go.

BEE [5] integrates the estimates of KIP and SAD with each unit's distinctive personality, style, and emotional condition [4] (**personal** influences) by evolving a model of the unit's personality against observations (Fig. 4, left side). It incorporates **dynamical** influences (the result of interactions with the environment, including other warriors) by extrapolating the behavior resulting from the learned personality, exploring multiple possible futures concurrently (Fig. 4, right side) to determine the unit's most likely future behavior.

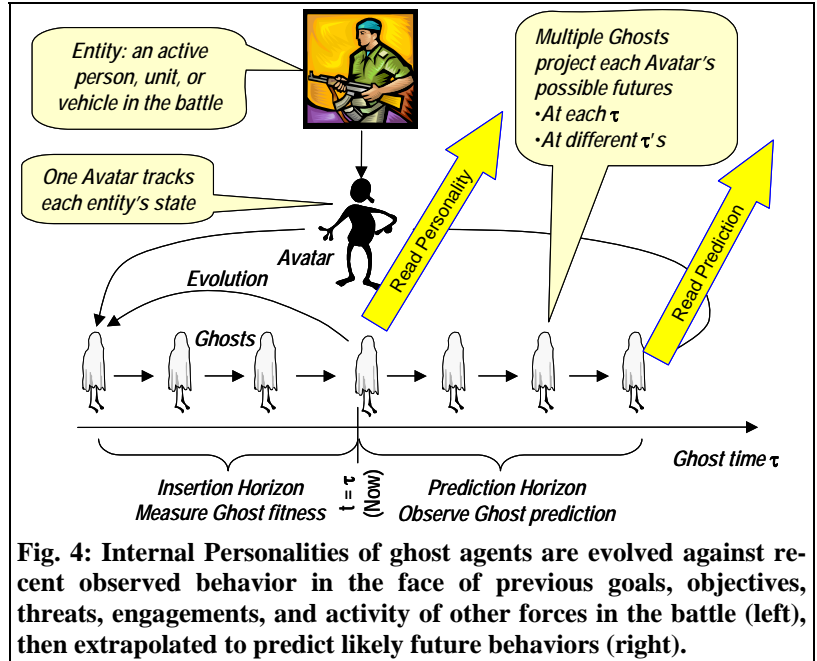
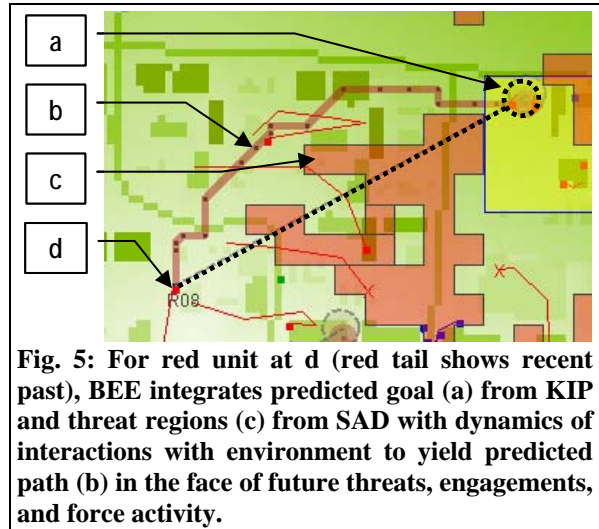


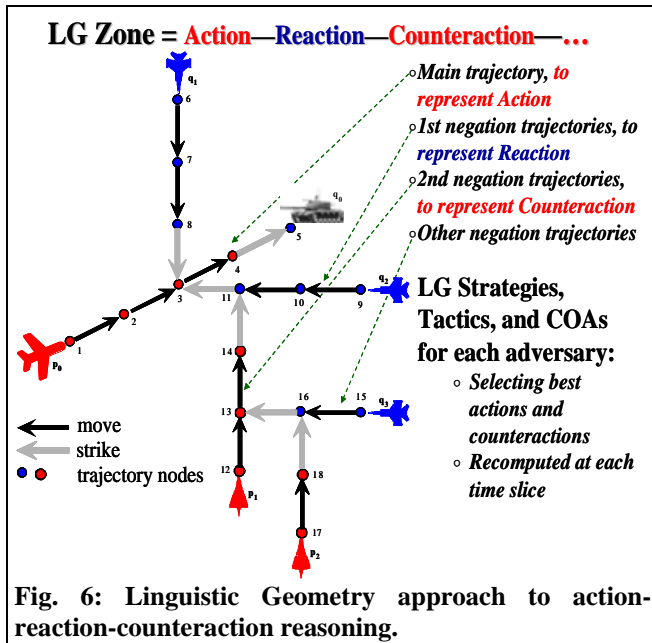
Fig. 5 shows how KIP predicts *where* a unit wants to go and *what* it wants to do, SAD warns about high-threat regions to *avoid*, and BEE integrates these with its own personality induction and emulation to yield high-quality predictions.

2.3 Adversarial Reasoning Module-S

The Adversarial Reasoning Module-S (ARM-S) predicts future behavior of Red and generates courses of action for Blue, using a constructive game theoretic algorithm. The heart of ARM-S is the Game Solving Tool (GST).

The GST provides fast and scalable solutions for mission planning and execution via generation of advantageous strategies, tactics, and COAs to guide the units' actions. For RAID, the GST generates long-term plans,





i.e., predictive Blue/Red COAs, thus serving as the “brain” behind the software oracle that predicts the future for human adversarial teams, Blue and Red, in real time.

Grounded in Linguistic Geometry (LG) [7, 8], the GST supports wargame construction in the form of the LG hypergame, and serves as both an experimentation tool, and a COA generator for “what-if” analysis. Experiments with the GST are conducted by varying mission objectives, abort conditions, rules of engagement, etc. The start state of the game can be selected either by manually placing entities on the abstract board, or by receiving it from an external simulator, e.g., from OTB (OneSAF Testbed Baseline). After the start state is established, for each Blue mission within the wargame, the GST generates the centerpiece LG constructions called LG zones (representing action-reaction-counteraction reasoning, Fig. 6). From the LG zones, the GST extracts the initial COA for each mission, as well as predictive COAs for Red. At the beginning of the engagement, the initial LG-based COAs are utilized to provide advice for the Blue Commander. As the mission progresses, the COAs are updated at regular intervals and/or upon the Commander’s request by taking into account the known changes of entity locations and states, Blue losses/gains, Red actions, and the Blue Commander’s preferences.

3 HOW HAS RAID PERFORMED?

RAID has been evaluated in two series of controlled experiments. The first series took place in July 2005 in Orlando, FL. The most recent round took place in July-August 2006 with help from the Battle Command Battle Lab and the US Army Command and General Staff College (CGSC) at Ft. Leavenworth, KS. The Ft. Leavenworth experiment consisted of 18 pairs, or 36 individual, simulation wargames. Data collected from these experi-

ments supported a statistically significant conclusion that the RAID software was, on average, as accurate as or better than a competent human staff at pinpointing the true locations of concealed insurgent teams and estimating their future positions / strength. The results also supported a conclusion that the RAID runs of the experimental pairs were more effective in terms of accomplishing the Blue mission. These results also demonstrate a strong statistical significance.

3.1 Experimental Setup

These experiments pitted a Blue dismounted company, supported by five Stryker vehicles and the ability to request up to 4 rotary wing aircraft from higher echelons, in an urban fight against a well prepared force of hostile militias and insurgent cadres. Unfolding in a capital of a developing country, the battles were reminiscent of events in places like Mogadishu, Najaf, and Fallujah.

Human Red and Blue commanders engage one another in an urban environment (Fig. 7), mediated through a computer simulation OTB. Their commands are conveyed to OTB’s simulations of their units through human role players acting as the platoon level commanders dedicated to each force. OTB plays out effects of these commands, including

- the movements of units
- the outcome of engagements
- the visibility of units to one another, based on the urban terrain

OTB reports the state of the battle every 20 seconds. To simulate the fog of war, the data feed is filtered. Instead of reporting the locations of all units, the feed re-

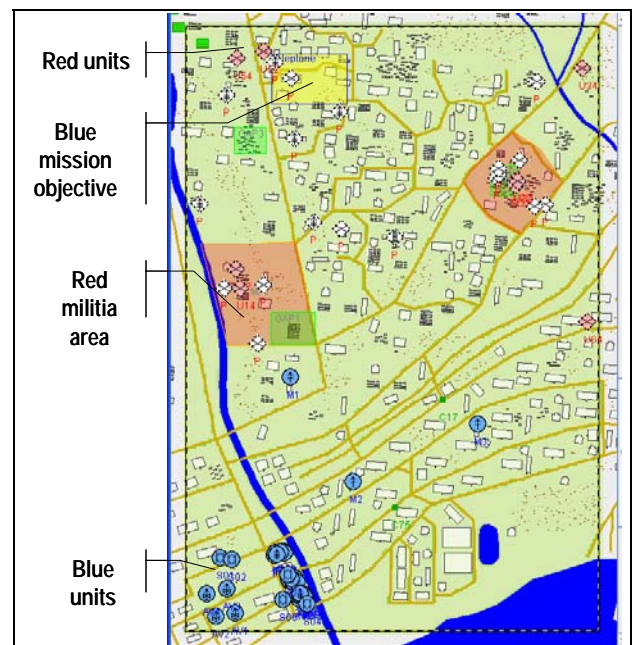


Fig. 7: A portion of RAID’s urban landscape.

ports only those units that are visible to Blue and is delayed by 2 minutes of wargame time.

A specially trained team played Red as an intelligent, innovative, hard-to-template adversary. Red's creative use of urban terrain, rapid movement in the familiar city, concealment, deceptions, ambushes, IEDs, RPGs, MANPADS, heavy machine guns, infiltration and civilian spies challenged the simulated Blue forces and their human commanders.

As each wargame progresses, RAID continually generates a series of running estimates--enemy goals, strengths, positions, positions of key weapon systems, movements--looking forward as far as 10 to 120 minutes. Independently, another set of predictive running estimates is produced by a competent staff, most of them Majors with operational experience in the Mideast.

3.2 Metrics

RAID is evaluated against a human staff, in two main areas: battle outcome and prediction accuracy.

Battle Outcome.—The outcome of a battle (typically one to two hours in length) is evaluated using a scenario-specific weighted sum of 7 scoring factors;

- Red attrition
- Blue Force Protection
- Remaining distance to the Blue objective
- Time to complete the mission
- Extent of Red strength occupying key terrain
- Facility protection
- Collateral Damage

These factors are tracked in real time, weighted and summed to provide an overall run score. The official run score is assigned as the run score at the time the Exercise Director declares the run is complete. Run completion occurs when the Blue commander completes his mission or the experiment run time limit is reached. The factor weights depend on the type of Blue mission: Attack, Defend, or Move. For example, in an Attack mission, Red attrition is weighted highly. In a Defend mission, Asset protection is more important, while in a Move mission, Blue strength is dominant.

Each mission is conducted in two configurations. In the non-RAID configuration, the Blue commander has a human staff of three intelligence officers and 1 operations officer to assist in situation assessment and tactical decisions. In the RAID configuration, RAID replaces the staff, and the commander is limited in how much he can deviate from the COA that RAID generates. The overall mission score achieved by the human staff is compared with that achieved by RAID within each pair of experimental runs.

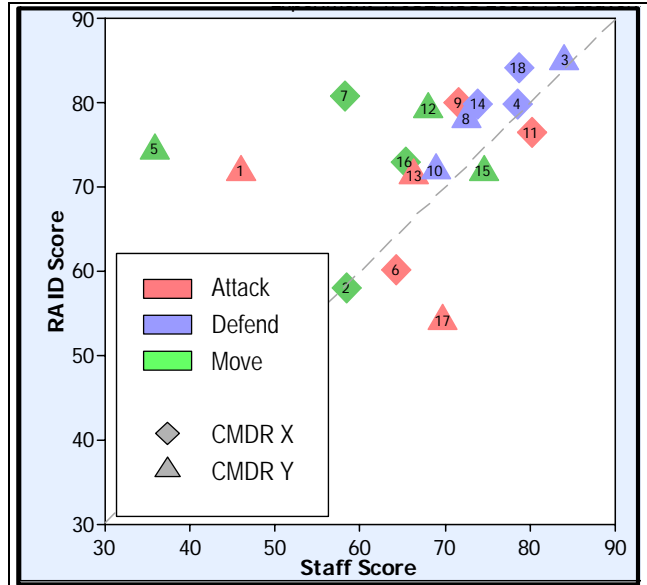


Fig. 8: RAID vs. human staff on paired runs. Points above the line favor RAID.

Predictive Accuracy.—The second score is a comparative analysis of the accuracy of the RAID and human staff predictions at two time horizons (0 and 15 minutes). During a run, the human staff is asked to predict the locations of Red units approximately every 15 minutes, and the accuracy of RAID's predictions are compared with those of the humans. The accuracy is measured in two ways. A mathematical metric is used to measure the extent of error in each prediction and results in a value between 0 (low) and 1 (high) being assigned as the prediction accuracy for that prediction. The second measure is a stylized version of Circular Error Probable (CEP) that has more tactical meaning than the pure metric measurement. As interpreted for RAID, the CEP measure of prediction accuracy is defined as the minimum error distance from the predicted locations such that there is a 50% overlap with the ground truth strength.

3.3 Results

Fig. 8 shows RAID's performance on battle outcome in comparison with the human staff. Eighteen pairs of runs were constructed with similar objectives and constraints, and one run of each pair was run with RAID, the other with the human staff. Each point plots the staff score of a given pair against the RAID score of the same pair. Thus points above the diagonal line favor RAID. In only four runs did staff achieve a higher mission score than RAID. Two cases were comparable, while in twelve cases, RAID outperformed the human staff. A traditional analysis with a T-test shows that RAID outperforms staff with a p-value of 98%. Because the T-test assumes normality of the underlying distribution, we also apply the non-parametric Wilcoxon rank-sum test, and find significance at the level of 97.3%.

In 18 RAID runs, RAID generated 1405 predictions at each of two time horizons (0 and 15 minutes), while in 18 non-RAID runs, staff generated 102 predictions. Fig. 9 shows a box-and-whisker plot of the CEP measures, in meters, of these predictions. The box covers the interquartile range with a line at the median, whiskers extend to the most distant data points within 1.5 of the interquartile range from the edge of the box, squares show outliers within 3 interquartile ranges, and stars show more distant outliers. Because the error measure is the CEP radius, a lower score reflects superior prediction. RAID’s median score even at 15 minutes is lower than either Staff median. The Wilcoxon test shows that the difference between the H15 scores is significant at the 99.76% level, while that between the H0 scores is significant at more than 99.999%.

Thus on both measures evaluated, RAID performs as well as or better than human staff with high confidence.

4 WHY IS RAID BENEFICIAL?

As a computer-based tool, RAID offers three benefits to the commander’s decision making: detail, diversity, and objectivity.

RAID can analyze information at a much finer level of *detail* than can a human. It can track each individual unit, or potentially individual combatants, every square meter of a city, and every floor of every building, and update this information based on information arriving several times a minute, several times an hour, or several times a day. This level of detail would quickly overwhelm a human, who must resort to approximations and heuristics.

RAID can consider a more *diverse* range of alternative futures than can a human, and project them farther into the future. In this regard, its advantage is similar to that of the IBM chess computer “Deep Blue” in its historic defeat of Gary Kasparov. The best human chess player can reason only about ten steps into the future; the computer was able to reason twelve steps ahead, analyzing more possible outcomes and using this information to guide its selection of the next step.

RAID can be more *objective* than a human. Humans tend to converge quickly on a working hypothesis, through which new information is then filtered and frequently inappropri-

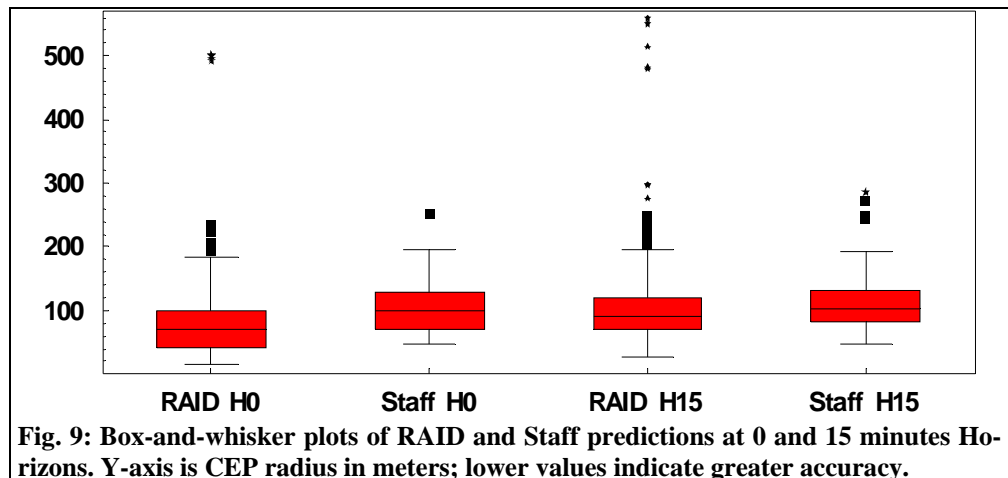
ately discarded. As a result, information that might suggest an alternative hypothesis is frequently rationalized away and is subject to a pattern recognition or decision bias until it becomes overwhelming, leading to a discontinuous “conversion” in the human’s thinking. Because RAID can maintain multiple hypotheses, it can more objectively evaluate each new item of information, and can help the human avoid premature convergence.

5 DEPLOYING RAID

In its developmental environment, RAID receives its data feed from a computerized simulation of the battle, specifically OTB. A major issue in deploying RAID in real combat is obtaining a comparable computer-readable feed of current engagements and emerging combat actions. An important potential solution to this requirement is the FightSight system recently tested by DARPA [6]. FightSight fuses thermal and acoustic sensors to locate and map urban firing events, including where shooters are located, the kinds of weapons they are using, and movements of firepower as the battle evolves. Fig. 10 illustrates the data flow involved in integrating these two technologies, and also suggests how the combination could drive automated platforms, an application that will be discussed in a forthcoming companion paper.

6 CONCLUSION

RAID has demonstrated (in a simulated environment) the potential to provide a MOUT commander with a new and novel insight into the near-term course of an ongoing battle. Current results indicate that it may be able to predict enemy movements more accurately than human experts and generate recommendations that aide the commander’s decision-making process. The digital feed of situation awareness that it requires can be provided by other systems under development, such as FightSight. Together, these technologies may increase the effectiveness of our troops and unmanned systems operating in complex urban settings against highly adaptive adversaries.



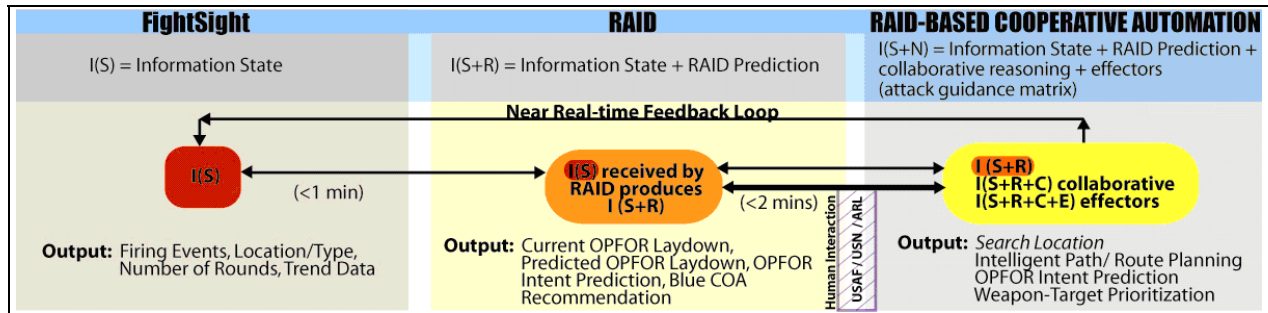


Fig. 10: RAID's potential relationship to other novel technologies

7 ACKNOWLEDGMENTS

This material is based in part upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract Nos. F3062-99-C-0202 and NBCHC040153. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the DARPA or the Department of Interior-National Business Center (DOI-NBC). Distribution Statement "A" (Approved for Public Release, Distribution Unlimited).

8 REFERENCES

- [1] A. Kott and M. Ownby, Tools for Real-Time Anticipation of Enemy Actions in Tactical Ground Operations, *Proceedings of the 10th International Command and Control Research and Technology Symposium – The Future of C2*, June 2005
- [2] W. M. McEneaney and R. Singh. Robustness against Deception. In A. Kott and W. McEneaney, Editors, *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind*, pages 167-208. Chapman and Hall/CRC Press, Boca Raton, FL, 2006.
- [3] P. Nielsen, R. Jones, and J. Crossman. Human Factors in Opponent Intent. In A. Kott and W. McEneaney, Editors, *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind*, pages 23-48. CRC Press, Boca Raton, FL, 2006.
- [4] H. V. D. Parunak, R. Bisson, S. Brueckner, R. Matthews, and J. Sauter. Representing Dispositions and Emotions in Simulated Combat. In *Proceedings of Workshop on Defence Applications of Multi-Agent Systems (DAMAS05, at AAMAS05)*, Utrecht, Netherlands, pages 51-65, Springer, 2005. <http://www.newvectors.net/staff/parunakv/DAMAS05DETT.pdf>.
- [5] H. V. D. Parunak and S. A. Brueckner. Extrapolation of the Opponent's Past Behaviors. In A. Kott and W. McEneaney, Editors, *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind*, pages 49-76. Chapman and Hall/CRC Press, Boca Raton, FL, 2006.
- [6] S. Snarski., K. Scheibner, S. Shaw, R. Roberts, A. LaRow, E. Breiffeller, J. Lupo, D. Neilson, B. Judge, J. Forren. Autonomous UAV-based Mapping Of Large-Scale Urban Firefights. In *Proceedings of SPIE -- Volume 6209*, May 2006
- [7] B. Stilman. *Linguistic Geometry: From Search to Construction*. Boston, Kluwer, 2000.
- [8] B. Stilman, V. Yakhnis, and O. Umanskiy. Strategies in Large-Scale Problems. In A. Kott and W. McEneaney, Editors, *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind*, pages 251-286. Chapman and Hall/CRC Press, Boca Raton, FL, 2006.