

10th International Command and Control
Research and Technology Symposium

The Future of C2

Title:

**Tools for Real-Time Anticipation of Enemy Actions
in Tactical Ground Operations**

Topic: Decisionmaking and Cognitive Analysis

Alexander Kott^a, Michael Ownby^b

^aDefense Advanced Research Projects Agency,
3701 N. Fairfax Drive, Arlington VA, USA 22203

^bSolers, Inc., 1611 N. Kent St, Arlington VA, USA 22209

Distribution Statement "A" (Approved for Public Release, Distribution Unlimited).

Abstract

DARPA has recently undertaken a research project titled Real-time Adversarial Intelligence and Decision-making (RAID), which provides in-execution predictive analysis of probable enemy actions. A particular focus of the program is tactical urban operations against irregular combatants – an especially challenging and operationally relevant domain. The RAID program leverages novel approximate game-theoretic and deception-sensitive algorithms to provide real-time enemy estimates to a tactical commander. In doing so, the RAID program is addressing two critical technical challenges: (a) adversarial reasoning: the ability to continuously identify and update predictions of likely enemy actions; (b) deception reasoning: the ability to continuously detect likely deceptions in the available battlefield information. Realistic experimentation and evaluation is driving the development process using human-in-the-loop, wargames to compare humans and the RAID system. This paper provides a discussion of the techniques and technologies chosen to perform the adversarial and deception reasoning. It also provides details about the experiments and experimentation environment that are used to demonstrate and prove the research goals.

Introduction

The Information Exploitation Office (IXO) of the Defense Advanced Research Projects Agency (DARPA) has recently undertaken a research project titled Real-time Adversarial Intelligence and Decision-making (RAID)[1], to build tools capable of in-execution predictive analysis of probable enemy actions. A particular focus of the program is tactical urban operations against irregular combatants – an especially challenging and operationally relevant domain.

The RAID program leverages novel approximate game-theoretic and deception-sensitive algorithms to provide real-time enemy estimates to a tactical commander. In doing so, the RAID program is addressing two critical technical challenges: (a) adversarial reasoning: the ability to continuously identify and update predictions of likely enemy actions; (b) deception reasoning: the ability to continuously detect likely deceptions in the available battlefield information. Although many types of military operations can greatly benefit from the capabilities outlined above, the RAID program is focusing on a well-circumscribed, intentionally narrow but still very challenging domain: in-execution, tactical combat of largely dismounted infantry (supported by armor and air platforms) against a guerilla-like enemy force in urbanized terrain. Realistic experimentation and evaluation drives the development process using the human-in-the-loop, Army OTB (OneSAF Testbed) wargame to compare humans and the RAID system. The products of the program have potential for transition to Army military intelligence and battle command systems.

Motivation

In a number of recent publications, US military leaders have called for the development of techniques and tools to address the twin challenges of adversarial and deception reasoning.

The US Air Force community uses the term *predictive battlespace awareness* [2, 3] while a related term, *predictive analysis*, is used in the US Army community [4]. Both refer to future techniques and technologies that would help the commander and staff to characterize and predict likely enemy courses of action, to relate the history of the enemy's performance to its current and future actions, and to associate these predictions with opportunities for friendly actions and effects. Both communities have pointed out the lack of technologies, techniques and tools to support predictive analysis and predictive battlespace awareness.

Recent years have seen the emergence of capable tools for generating friendly courses of action, e.g., the Mixed Initiative Control of Automateams (MICA) [5] and Course of Action Development and Evaluation Tool (CADET) [6] programs. Given the definition of available assets, terrain, and enemy information, missions and rules of engagement, such tools generate a detailed, optimized plan of actions (Figure 1), allocate and task-organize the assets, schedule the actions with respect to applicable time constraints, and estimate the outcome of the operation. They even identify expected reactions of the enemy as well as suitable counteractions of the friendly forces.




Figure 1. Advanced automated tools for battle planning and management are able to produce detailed and optimized resource allocations, task times and movement routes. However, they do not yet explicitly reason about enemy counteractions. [7]

Overall, such tools have demonstrated the ability to perform on par with, or better than, the human staff; they help produce complex planning products dramatically faster (orders of magnitude faster) yet without loss of quality, as compared to the conventional, manual process. The battle plans produced by such automated tools look to human reviewers rather sophisticated and insightful. However, this emerging generation of battle planning and management tools also exhibits serious shortcomings, particularly in adversarial reasoning.

First, these tools have no means to take into account the emotional and cognitive aspects of the battle. Real human warriors, at all levels of responsibility, have beliefs, emotions, desires, biases, preferences, etc., that contribute much into their plans and actions. These emotional and cognitive aspects are complex, and they change dynamically as the battle unfolds. Today's tools do not reason on such factors. They also do not take into account the inevitable errors and cognitive limitations of the humans in real-world warfare.

Second, today's tools do not explicitly look ahead, wargame or game-solve their plans with a view that the enemy may also have insight into friendly courses of action and may counteract it. Generally, they plan backward from the key events pre-defined in the human-generated high-level course of action. Unlike humans, such tools do not attempt to invent (even in a limited sense) the strategy of the battle. They merely fill in the details (albeit important and complicated) in the outline of the adversarial encounter envisioned by the human. While acceptable in some applications, this shortcoming may not be in many others.

Third, much of warfare is based on deception and concealment. Human commanders explicitly and continuously pay attention to the possibility that the enemy would employ deceptive actions or conceal actual actions and assets in order to manipulate the friendly understanding of the current and future events. Human commanders and staff planners also develop and employ deception and concealment in their own plans. Today's tools (Figure 2) do not explicitly reason about such issues.

Finally, the current generation of C2 tools does not take into consideration a very important factor: the impact of decision-making processes and organizations on the enemy (and friendly) actions. There are complex and influential dynamics in command decision-making, in communications, in propagation of uncertainty, errors, confusion, trust and fears through the formal and informal networks of the leadership of multiple units and echelons. All of this is outside the scope of today's tools.

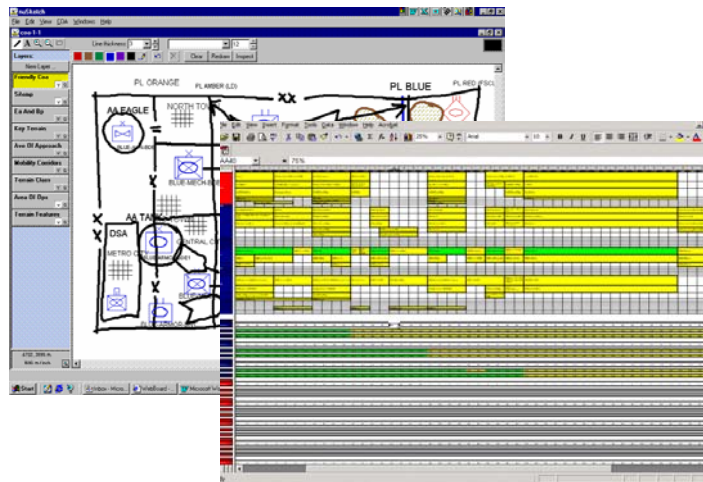


Figure 2: Battle planning tools like CADET [6] generate a detailed battle plan while performing the action-reaction-counteraction analysis. However, they rely on a pre-specified overall scheme of maneuver, and do not reason about deception.

It is these shortcomings in the state of art that the RAID program aims to mitigate.

Adversarial Reasoning

RAID focuses on approaches and challenges in what may be encompassed by the term adversarial reasoning: computational solutions to problems of determining the state, intent and actions of one's adversary, in an environment where one strives to effectively counter the adversary's actions.

The subtopics within this subject include: belief and intent recognition, opponent's strategy prediction, plan recognition, deception discovery, deception planning, and strategy generation. From the engineering perspective, the applications of adversarial reasoning cover a broad range of practical problems: military planning and command, military and foreign intelligence, anti-terrorism and domestic security, law enforcement, information security, recreational strategy games, simulation and training systems, applied robotics.

Naturally, adversarial reasoning is particularly important to the domain of military operations. In military command and control, the challenge of automating the reasoning about intents, plans and actions of the adversary would involve the development of computational means to reason about the future enemy actions in a way that combines:

- ◆ the enemy's intelligent plans to achieve his objectives by effective use of his strengths and opportunities;
- ◆ the enemy's perception of friendly strengths, weaknesses and intents;
- ◆ the enemy's tactics, doctrine, training, moral, cultural and other biases and preferences;
- ◆ the impact of terrain, environment (including noncombatant population), weather, time and space available;
- ◆ the influence of personnel attrition, ammunition and other consumable supplies, logistics, communications, sensors and other elements of a military operation; and
- ◆ the complex interplay and mutual dependency of friendly and enemy actions, reactions and counteractions that unfold during the execution of the operation.

Adversarial reasoning is the process of making inferences over the totality of the above factors.

Although many of the problems inherent in adversarial reasoning have been traditionally seen as belonging to the field of game theory, we argue that practical adversarial reasoning calls for a broader range of disciplines: artificial intelligence planning, cognitive modeling, control theory, and machine learning in addition to game theory. An effective approach to problems of adversarial reasoning must combine contributions from disciplines that unfortunately rarely come together. One of RAID's research objectives is to explore important close relations between ideas coming from such diverse areas.

Adversarial reasoning is broader than the military domain, and the benefits of RAID's research results extends beyond C2 applications. The applied communities (practitioners, engineers, developers) interested in adversarial reasoning certainly include military planners and analysts as well as the intelligence community. Also, those who develop applications and processes related to anti-terrorism and domestic security and law enforcement would share similar interests. Other, less obvious communities of practitioners include those concerned with financial fraud detection and information security. They would also benefit from a better understanding of what and how the opponent thinks while preparing and executing financial fraud or intrusions into an information system.

Developmental and Experimental Approach

The RAID system (Figure 3) is composed of two major components: the Adversarial Reasoning Module and the Deception Reasoning Module. The purpose of the Adversarial Reasoning Module is to generate, either on-demand or in response to battle situation changes, predictions of Red (enemy) actions and assumption about Blue (friendly) actions. Continually observing the evolution of the battlefield and the evolution of the predictions made by the Adversarial Reasoning Module, the Deception Reasoning Module infers possible concealed enemy force elements or movements of elements, incorrectly identified enemy assets,




Figure 3: Key components of the RAID program.

decoys, and actions designed to mislead friendly forces.

The development of the RAID system is driven by a rigorous schedule of increasingly difficult and realistic experiments using the OTB wargaming simulation system. The purpose of the experiments is to explore the ability of RAID to make effective estimates of enemy actions and assumptions about friendly counteractions (move-countermove reasoning), as compared to a human staff.

To focus the experimental development process, the RAID program concentrates on a particularly demanding and operationally-challenging domain: tactical operations in an urban environment against dismounted irregular combatants. The complex urban terrain offers a high density, as well as fragmentation, of threats and opportunities for forces [8]. Further, the terrain itself is dynamic because it is continually modified by human actions (barricades in the streets, holes in the walls, etc.). The presence of non-combatants on the battlefield must be explicitly considered and collateral damage minimized. Fire and maneuver of forces are not the only actions that must be carefully considered. Intelligence gathering, communications, and logistics (including casualty evacuation) are tightly coupled with fire and maneuver. The scale of the computational problem is immense and yet solutions must be generated in near real-time.

The RAID experiments aim to approximate the complexity of the target environment to the best extent possible. The core of the experimentation testbed, the Combat Simulation System, is based on the proven Army simulation and training system, OTB. Certain modifications to the existing system's interfaces and entity behaviors are being implemented to meet the needs of RAID experimentation.

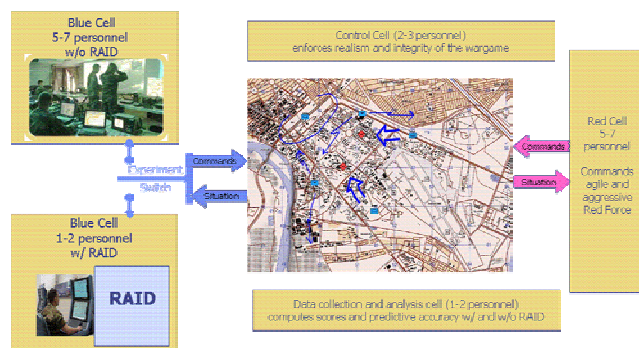


Figure 4. RAID experiment setup.

In the experiments, the RAID system performs the following functions: reads Red/Blue situation from the Combat Simulation System; accepts guidance from the Blue commander (priorities, key objectives, etc.); on demand, estimates the likely actions of Red and assumed actions of Blue for the next X minutes of wargame time; completes every new estimate rapidly; presents the estimate to Blue commander as overlay graphics.

A typical series of experiments (Figure 4) consists of a number of benchmark games (without RAID) and a similar number of test games (with RAID). Control of Red entities is performed by a Red cell of 5-7 experienced human wargamers. Control of Blue entities in the benchmarking games is performed by a Blue cell of 5-7 human wargamers. Control of Blue entities in the test games is by a Blue cell of only 1-2 human wargamers supported by RAID.

In all series of experiments, overall complexity of the problem is varied by adjusting number and granularity of Red/Blue units, and restrictions on the set of available actions and weapon types. It is estimated that the complexity measure (the size of the search space) of the

Phase I problem exceeds $10^{8,000}$, rapidly growing even further in the following phases. For comparison, a chess game's complexity is on the order of 10^{32} .

For the purposes of these experiments, success of the Blue force is measured by the rate of progress toward the mission accomplishment (e.g., advancing to or clearing the specified objective); Red personnel casualties; avoidance of friendly losses and collateral casualties. Success of the Red force is measured by delaying the Blue force and causing Blue casualties. Success of the Blue force is considered an effectiveness metric, and scores are compared between the benchmark and test games.

Comparison between the benchmark series of experiments and the test series also provide other rigorous quantitative measures of RAID's capabilities compared to those of human experts. Among the experiment metrics, accuracy is particularly important: the number of wrong predictions made by RAID, expressed as a fraction of total predictions and compared statistically to the same measure of human expert performance. Typical predictions refer to tangible estimates used in the practice of military intelligence, such as location, strengths and actions of an enemy unit at a particular time interval in the future. Wrong predictions also include false positives – Red actions that are predicted but do not occur and false negatives – Red actions that occur but are not predicted.

Key Technology Themes

Three themes are particularly salient in adversarial reasoning. Faced with an intelligent adversary, a decision maker, whether human or computational, often must begin by using the available information in order to identify the intent, the nature and the probable plans of the adversary. Hence the first key theme of adversarial reasoning – opponent's intent and plan recognition. Further, a capable adversary is likely to conceal his plans and to introduce crucial deceptions. Therefore, the second theme – deception discovery – focuses on detection of concealments and deceptions. Finally, having made progress in identifying an adversary's intent and guarding himself against possible deceptions, the decision maker has to formulate his own plan of actions that takes into account potential counteractions of the adversary – and this is the third theme, strategy formulation.

Recent years have seen a dramatic rise in the capabilities of techniques relevant to adversarial reasoning, making potential solutions relevant, for the first time, to problems of a practical scale and complexity. The 1950's and 1960's saw critical developments in the understanding of Game Theory, a key element of adversarial reasoning. The game problems are tremendously more complex than those for systems without antagonistic inputs. Until recently, game formulations of practical problems, with the attendant level of detail and scale, resulted in a degree of complexity that could not be satisfactorily handled. Today, however, there are claims of computational techniques that offer the promise of robustness and scalability appropriate for practical applications. Furthermore, there has been a dramatic rise in the maturity of technical approaches that address the cognitive aspects of adversarial reasoning, particularly the means to model how an adversary perceives a situation, reflects on what the opponent might perceive and do, and decides on a course of action.

Technology Theme: Recognition of the Opponent's Intent

In formulating the predictions, RAID takes into account such factors as high-level objectives, the intents and preferences of the friendly and enemy commanders, physical capabilities and needs of the assets available to both sides, mutual influence of actions of Blue and Red forces, terrain, weather, non-combatants, cultural and doctrinal aspects, psychological factors affecting troops and commanders, prior evolution of the operation, etc. It is critical for RAID to consider cognitive and emotional factors. Real human warriors, at all levels of responsibility, have beliefs, emotions, desires, biases, and preferences, that contribute much into their plans and actions. They are complex, and they change dynamically as the battle unfolds.

One of the technical approaches RAID is exploring for this purpose is an integration of cognitive modeling tools (Figure 5) with pheromone-like modeling of the combatant's perception of the battle's threats and opportunities. This approach uses a cognitive framework [10] to model a fighter's cognitive and emotional state (beliefs, desires, and intents) and includes cultural and doctrinal preferences. This cognitive framework is then augmented with pheromone-based algorithms [9], using "ghosts" to traverse multiple trajectories in the solution space to converge on a solution.

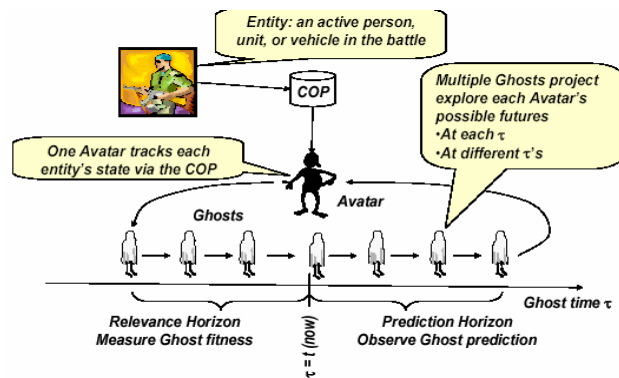


Figure 5. One approach - a cognitive framework with pheromone-based algorithms [9] that extrapolate an agent's past behaviors and mental states into the future.

Technology Theme: Discovery of the Opponent's Deceptions

While continually observing the evolution of the battlefield and the evolution of the predictions regarding enemy's future actions, RAID infers possible concealed enemy force elements or movements of elements, incorrectly identified enemy assets, decoys, and actions designed to mislead friendly forces.

If available, RAID uses user-provided estimates regarding overall strength of concealed enemy assets, and types of most likely deceptions. In formulating its estimates of enemy concealment and deception, RAID considers the state of Red knowledge about the Blue, the Red beliefs about Blue sensor capability, the known Red tactics of concealment and deception, the costs and efforts of actions and measures involved in execution of concealment and deception, and the ability of the Red to use non-combatants for the purposes of concealment and deception.

RAID generates several alternative estimates of Red concealments and deception, if multiple alternatives are indeed likely. Each estimate is accompanied by its likelihood, and

assumptions on which it is based. Elements of the estimates are linked to the underlying evidence.

To produce such deception estimates, RAID employs several techniques. One of them is based on a risk-sensitive estimation [11, 12]. Preliminary results indicate this approach is superior to approaches such as Bayes/Kalman filter estimators. The technique is supplemented by utilizing a non-symmetric evaluation function, which models the goals and values of the Blue and Red team, as modified by input from the commander. Another technique is a deception robustness estimator that also takes advantage of the non-symmetric evaluation function. It deals explicitly with the presence of potentially antagonistic action on the part of the opponent, by searching to uncover deception activities that may be part of a long term enemy plan (Figure 6). It recognizes that these adversarial actions may be affecting both the dynamics of the unfolding operation and the observations that are obtained.




Figure 6. A small scale prototype [12] uses deception robustness to intercept a stealthy Red

Technology Theme: Identification of Strategies

The key function of the RAID system is to generate, either on-demand or in response to battle situation changes, predictions of Red actions and assumption about Blue actions. As information regarding battlefield situation (locations, strengths, postures, actions, etc.) of enemy and friendly troops becomes available or changes, either in the deliberate IPB and wargaming mode or during the execution of the operation, RAID generates a new or modified set of predictions, including most dangerous and most likely prediction, each characterized by its likelihood. These detailed prediction look forward anywhere from 30 to 300 minutes (as specified by the user) from the current moment, including sequence of actions, situated in time and space, performed by the enemy force.

Because the actions of Red and Blue forces are closely connected and influence each other, RAID also necessarily generates its estimates of the friendly actions similar to the predictions of enemy actions. These can be seen as assumption or recommendations regarding the friendly course of action. Although the primary function of RAID is to anticipate Red actions, the capability to suggest Blue actions is a natural, valuable byproduct that can be effectively utilized by an integrated C2 /Intel system.

The set of predictions should be broad enough to provide the commander with a sense of possible alternative futures, and yet small enough that it can be rapidly reviewed in the tempo of tactical combat. In particular, RAID should provide a suitable abstraction of each alternative prediction so that it could become a basis for displaying graphically as a rapidly comprehensible, simplified sketch.

To produce such detailed predictions of Red and Blue (coupled) actions, RAID is exploring several technologies. One of them is Linguistic Geometry [13] where an algorithm takes a Blue goal (e.g., capture an objective, destroy a target) and a Blue asset, and constructs a group of multiple sequences of actions that can lead to that goal; then it constructs a group of countermoves that Red can use to counteract the Blue action (the reaction), and then constructs the Blue counteractions (Figure 7) and so on. It is shown that the construction of multiple sequences is rather computationally inexpensive.




Figure 7. A small example of building action, reaction, counteraction clusters in the Linguistic Geometry [13] approach.

All such sequences of actions are then organized into clusters. The algorithm then repeats this process with the remaining goals and assets and generates more clusters. There is a relatively small space of such clusters and the search for the best moves within this space of clusters is dramatically less expensive than the search within the enormous space of individual moves and board positions used in conventional gaming techniques. Such approaches may produce effective solutions for adversarial reasoning problems of a practical scale and complexity as targeted by the RAID program

Future Applications and Deployment

In the future, as a military operation is being executed (Figure 8), the information from the battlespace, such as location, strengths and postures of enemy and friendly troops, is rapidly delivered to a military intelligence system. With today's proliferation of Blue force tracking devices and unmanned air and ground sensors, one envisions that the latency of such information would be measured in minutes. This fused sensor data would identify locations of some of the enemy units and some attributes, such as type, size, and posture. It is understood, however, that the fog of war would remain thick – in spite of the proliferation of sensors and improved fusion techniques. The battlespace information, especially the enemy information, would remain incomplete and potentially deceptive.

While continually monitoring the changes in the battlespace state as the information unfolds, the RAID system periodically or on request generates predictions of enemy actions and presents them in a user-friendly,




Figure 8. RAID is intended to provide tactical leaders with real-time predictive estimates of enemy actions – continuously and in-execution.

rapidly comprehensible format. These predictions are delivered to the user, such as a company commander, via computer displays in combat vehicles or on personal digital assistant (PDA)-like devices for dismounted personnel.

RAID output products include several possible alternative enemy courses of action, worked out in requisite detail, ranked in the order of likelihood, and presented as graphic overlays with brief textual notes and with an explanation of assumptions about the friendly course of action. The scope and details of the products are tailored for each individual user, his current situation and area of responsibility. RAID products are designed to be unobtrusive to the user. The user may elect not to see them at all, or to see them occasionally on request; he may use them extensively or ignore them entirely. If time and situation permits, the user, at his discretion, may input to RAID additional information, such as his updated intent and friendly scheme of maneuver for the upcoming phase of the operation or his own estimates of enemy intent. RAID uses this additional information, when available, to fine-tune its predictions. In the absence of such input, RAID makes do with its own assumptions and estimates. In no case does RAID become an additional burden on the user's time and attention.

Although military intelligence is one area where RAID capabilities are particularly relevant, other applications of such technologies are also attractive. Thus, tools for friendly COA preparation and real-time battle management can benefit from RAID's adversarial perspective. Further, developers of military simulation and training systems, as well as developers of commercial entertainment games, are always striving for a more realistic and intelligent opposing force within their respective systems. To a significant extent, they can benefit from advances in adversarial reasoning offered by the RAID program. Finally, a less obvious, but very relevant, area of practical applications is military robotics. In order to survive and be effective in a hostile environment, a robot (e.g., a highly autonomous unmanned aerial vehicle) should reason about the likely actions of its adversaries.

References

- [1] RAID program website, <http://dtsn.darpa.mil/ixo/programdetail.asp?progid=57>
- [2] Interview with General G. S. Martin, *Military Aerospace*, v.2, issue 6, 2003 pp.17-20
- [3] Air Force Handbook, <http://www.af.mil/library/posture/2002handbook.pdf>
- [4] US Army TRADOC, Force Operating Capabilities, Pamphlet 525-66, TRADOC, Fort Monroe, VA, June 2003, <http://tradoc.monroe.army.mil/tpubs/pdf/pams/p525-66.pdf>
- [5] MICA program website, <http://dtsn.darpa.mil/ixo/programdetail.asp?progid=7>
- [6] Kott, A., Ground, L., Budd, R., Rebbapragada, L., and Langston, J., Toward Practical Knowledge-Based Tools for Battle Planning and Scheduling, *Proceedings of the IAAI 2002 Conference*
- [7] M. Ownby, Mixed Initiative Control of Automa-teams (MICA) - a Progress Report, *Proceedings of the Third AIAA Unmanned Unlimited Technical Conference, July 2004*
- [8] US Army, Field Manual 3-06.11, Combined Arms Operations in Urban Terrain, <http://www.globalsecurity.org/military/library/policy/army/fm/3-06-11/index.html>

- [9] H. V. D. Parunak, S. Brueckner, and R. Savit. Universality in Multi-Agent Systems. In *Proceedings of Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, pages 930-937, IEEE, 2004.
- [10] J. Gratch and S. Marsella. A Domain-independent Framework for Modeling Emotion. *Journal of Cognitive Systems Research*, 5(4):269-306, 2004
- [11] W.M. McEneaney and B.G. Fitzpatrick, Control for UAV Operations under Imperfect Information, *Proceedings First AIAA UAV Symposium, 5/22/02*
- [12] W.M. McEneaney, B.G. Fitzpatrick and I.G. Lauko, Stochastic Game Approach to Air Operations, *IEEE Trans. on Aerospace and Electronic Sys*
- [13] B. Stilman, *Linguistic Geometry: From Search to Construction* (Kluwer Academic, 2000)