# Chapter Ten, UrbanSim: Using Social Simulation to Train for Stability Operations

Ryan McAlinden, David Pynadath, Randall W. Hill, Jr.
University of Southern California – Institute for Creative Technologies
{mcalinden, pynadath, hill}@ict.usc.edu

## Abstract

As the United States reorients itself towards to a period of reduced military capacity and away from large-footprint military engagements, there is an imperative to keep commanders and decision-makers mentally sharp and prepared for the next 'hot spot.' One potential hot spot, *megacities*, presents a unique set of challenges due to their expansive, often interwoven ethnographic landscapes, and their overall lack of understanding by many western experts. Social simulation using agent-based models is one approach for furthering our understanding of distant societies and their security implications, and for preparing leaders to engage these populations if and when the need arises. Over the past ten years, the field of social simulation has become decidedly cross-discipline, including academics and practitioners from the fields of sociology, anthropology, psychology, artificial intelligence and engineering. This has led to an unparalleled advancement in social simulation theory and practice, and as new threats evolve to operate within dense but expansive urban environments, social simulation has a unique opportunity to shape our perspectives and develop knowledge that may otherwise be difficult to obtain.

This article presents a social simulation-based training application (UrbanSim) developed by the University of Southern California's Institute for Creative Technologies (USC-ICT) in partnership with the US Army's School for Command Preparation (SCP). UrbanSim has been in-use since 2009 to help Army commanders understand and train for missions in complex, uncertain environments. The discussion describes how the social simulation-based training application was designed to develop and hone commanders' skills for conducting missions in environs with multifaceted social, ethnic and political fabrics. We present a few considerations when attempting to recreate dense, rapidly growing population centers, and how the integration of real-world data into social simulation frameworks can add a level of realism and understanding not possible even a few years ago.

### Keywords

Social simulation, mission command, POMDP, counterinsurgency training

## Motivation

Back in 2006, the United States was decisively engaged in major operations in Iraq and Afghanistan. Though traditional offensive-defensive operations remained prevalent, the challenges of fighting against organized yet surreptitious insurgencies and factions drew widespread attention. Today, it is widely accepted that future military leaders will face similarly stressful and demanding situations that are, in many cases, not covered by standard tactics and doctrine (Smith, 2005). These operations, which combine both lethal and non-lethal aspects of warfare, have been referred to as “armed social work,” in which military forces attempt to “redress basic social and political problems while being shot at” (Kilcullen, 2006). The overarching challenge is to develop leaders who possess adaptive expertise and function effectively in complex environments, and to prepare them for novel situations unlike any they may have experienced in the past.

The School for Command Preparation (SCP) at Ft. Leavenworth is the primary Army institution for preparing newly-selected Battalion Commanders for all types of missions, including those centered on protecting and empowering indigenous populations who may be experiencing national security threats of their own. The school has an imperative to ensure their curriculum is updated with topics and material that best positions commanders for success once downrange. To address this challenge back in 2009, USC-ICT, in partnership with SCP, Army Research Laboratory (Human Research and Engineering Directorate and the Simulation and Training Technology Center), and Army Research Institute, developed an instructional software suite for military commanders and their staffs to practice directing and coordinating operations with a “stability-focused” component. The UrbanSim Learning Package (or UrbanSim for short) focuses predominantly, but not exclusively, on military operations in support of the local citizenry and government that take place after primary offensive and defensive efforts have concluded. Applying the principles from Guided Experiential Learning (GEL) (Clark, 2004), UrbanSim was designed, developed, and deployed with a strong pedagogical focus. The resulting learning objectives called for a complex, dynamic, yet highly realistic simulated environment, which brought about the need to employ agent-based research technologies and transition them to software that would eventually be used in a classroom setting.

## Mission Command & UrbanSim

The foundations for commanding in the Army are framed around the precepts of mission command (MC): Understand, Visualize, Describe, Direct, Lead and Assess. A seminal precept of MC requires the commander to blend the art of command and the science of control, focusing on the human dimension of military operations as opposed to technological solutions. A toolkit of commander competencies helps to feed core fundamentals, including MC domain knowledge, communication, decision-making, adaptability, self-awareness and self-assessment. These competencies cannot be learned solely out of a book or as a set of rules. Instead they require practical, tacit skills which typically are developed through experience, time, and with the help of feedback from mentors, superiors, peers and subordinates.

The UrbanSim training package specifically targets the need for practicing these skills and techniques. The application targets trainees’ abilities to maintain situational awareness, anticipate second and third order effects of actions and adapt their strategies in the face of difficult situations. It allows commanders and their staffs to develop skills in executing the “art of mission command” in counter-insurgency (COIN) or stability operations environments. The application includes two components: (1) a self-paced Primer; and (2) a computer game-based practice environment;. The training exercise typically takes one full day to execute. It may either be conducted individually in a classroom setting with a lead instructor (as is done at SCP), or in a Staff Exercise with different players assuming the roles of a Battalion Staff (BN CDR, S2, S3).

UrbanSim adheres to the GEL model by providing learners with a complete cognitive foundation required to conduct complex, dynamic operations ranging from high OPTEMPO, highly kinetic and security-focused, to lower-profile, governance or development-focused missions. The terminal learning objectives (TLOs) of the experience include:

1. Achieve and maintain situational awareness and understanding in a complex environment
2. Balance offense, defense and stability. Understand the role of intelligence and reconnaissance security / raids
3. Anticipate 2nd/3rd order effects of decisions; tactical effects with strategic consequences
4. Reinforce doctrinal principles of “Shape, Clear, Hold, Build”

These TLOs are exercised through various stages in the application, discussed in detail below. The enabling learning objectives (ELOs) are core MC topics that, when exercised at various phases of the game, help satisfy the TLOs. ELOs for UrbanSim include:

1. *Mission Overview* – understanding and interpreting higher-headquarters (HQ) intent; higher-HQ lines of effort (LOEs); and higher-HQ information requirements (CCIR)
2. *Mission Analysis* – understanding the landscape in the area of operations (AO) (e.g. political, economic and military networks, key individuals, organizations and groups)
3. *Mission Plan* – being able to author a tractable, realistic commanders intent, formulating and monitoring LOEs, CCIRs and measures of effectiveness (MOE)
4. *Mission Execution* – directing action of subordinate units in support of a desired end state
5. *Mission Assessment* – being able to self-assess performance along the LOEs and MOEs over time

The UrbanSim application has been used to train Soldiers in a variety of institutional settings to include SCP’s Tactical Commanders Development Program (TCDP); various Captains’ Career Courses; the Engineer Basic Officer Leadership Course; and the AMEDD Senior Leaders Course at Joint Base San Antonio. UrbanSim has been used successfully to stimulate battalion-level, battle staff exercises and to stimulate training for Company Intel Support Teams (CoIST) for Active and National Guard units at Ft. Hood, the Joint Maneuver Training Center (JMTC), and the California Na-

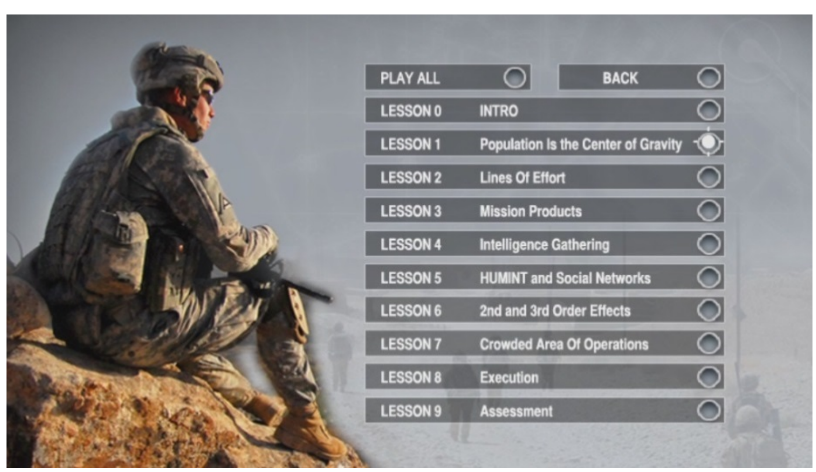

tional Guard. The training package was transitioned from the R&D community to the Army Games for Training (AGFT) Program in November 2011 and is available for distribution Army-wide via the Army's MilGaming web portal. UrbanSim also transitioned to the Army Low-Overhead Training Toolkit (ALOTT) Program in December 2011 and was fielded at Joint Base Lewis/McChord in June 2012 as part of the ALOTT New Equipment Training (NET) program.

### The Primer

The first component of the experience, the UrbanSim Primer, provides the requisite conceptual and task knowledge required for the learner to lead a full-scale stability operation, from analyzing background information via target folders and intelligence briefings, to coordinating the actions that are carried out in support of achieving a desired end state. Taking the form of an interactive tutorial, the UrbanSim Primer is divided into nine lessons, each of which contain a narrative, interview segments from former Commanders, and assorted practice exercises as a means of demonstrating specific tasks to the learner. Taking approximately one to two hours to complete, the self-paced Primer prepares the learner for the second application, the more complex UrbanSim Practice Environment.

### Practice Environment

The UrbanSim Practice Environment is a game-based tool that allows a learner to *plan, prepare, execute*, and *assess* a full stability operation. Similar to a turn-based strategy game (such as Civilization or Age of Empires), the learner directs subordinate units in the game to take action with and against agents (i.e. non-player characters, NPCs) in a virtual environment, and attempts to successfully complete a mission using the products/strategies learned in the UrbanSim Primer and in the classroom. Actions in the game are taken against key individuals, groups, and structures in an area of operation (AOR) with the intent of reaching the desired end state. Each turn-cycle in the game represents one day in simulation time, though actions can take multiple turns (i.e., days), and can be interrupted if conditions in the world do not allow the action to complete (e.g., money runs out to construct a school). Upon completion of a scenario, the learner is brought to a debrief phase where a summary of the mission is presented for the learner to evaluate their progress.

The game is driven by an underlying socio-cultural behavior model, coupled with a novel story engine that injects events and situations based on real-world experiences of former commanders. It also includes an intelligent tutoring system, which provides guidance to trainees during execution, as well as after action review capabilities.

The fundamental scoring mechanism of the game (i.e., how well the player does) is via the LOEs. There are six primary LOEs: Civil Security, Governance, Host Nation Security Forces (HNSF), Essential Services, Information Operations and Economics. Every action in the game has associated effects, both first and second-order, on one or more lines of effort. The value of the LOEs changes over the course of the game and at the end of the scenario are summarized for the player to see where their focus areas were (security, governing, developing), and whether they were aligned their desired end state. These LOEs, as well as all other secondary scoring values in the game (MOEs, CCIRs, etc), are determined by the underlying behavior model, described below.

## UrbanSim's Social Model

The technical challenge of any game or simulation AI is that it must be both expressive but application friendly. Having a multi-tiered, self-organized and self-steering behavior system may in theory be realistic and desirable, but can quickly become intractable either in terms of performance or usability. Therefore one must carefully determine what the requirements are for any behavior system driving an underlying experience. For UrbanSim, this meant balancing expressivity and realism of the model with authorability of scenarios. To accomplish this, UrbanSim uses two separate but coupled AI technologies: PsychSim and the Story Engine. The integration of these technologies is via UrbanSim's system architecture, which follows a data-driven distribution model where AI components work together in a synchronous cycle. Each cycle begins when a learner specifies a set of actions to be executed by subordinate units for the given turn. These actions are then sent to an intelligent tutoring system for evaluation, which may initiate a question-answer tutoring dialogue with the learner. Once this dialogue is complete, the learner commits the actions and the simulation cycle is executed and repeated.

### PsychSim

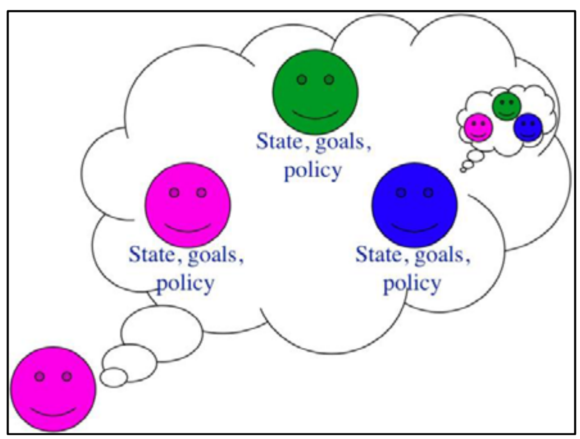


PsychSim is a multi-agent system developed by the University of Southern California (USC) that models beliefs about others to affect behavior of simulation entities. It is a framework for social modeling and simulation that has been used in a range of domains from analysis and planning to basic research on human behavior (Wang et al, 2012). In UrbanSim, entities are modeled using PsychSim and can represent both key individuals and aggregate-level features such as organizations, tribes, geographic regions or structures. The decision to aggregate was due to performance and the objectives of the trainer. UrbanSim is not a mission rehearsal tool, nor was it needed to include an accurate, validated social simulation. Instead, it is intended to teach critical thinking skills in commanders in environs that are representative of where they might eventually deploy. Though in some cases this may require a highly-detailed, realistic model down to the individual level, for the purposes of this trainer that requirement was never specified.

The PsychSim architecture is rooted in the Theory of Mind (ToM) principle. ToM is the ability to attribute [mental] states to oneself and others, and to understand that others have beliefs, desires and intentions that are different from one's own (Premack, 1978). In PsychSim this refers to agents that have subjective perspectives on others, and are able to potentially predicts others' actions/reactions, but also be able choose actions for themselves that will change the beliefs of others. Agents also have the ability to communicate, distort and hide information to influence others. In addition to ToM, decision-theoretic reasoning plays a pivotal part in PsychSim. This reasoning states that agents are free to pursue their own goals based on their values and beliefs. Often agents are presented with conflicting goals, or choices that must be made under uncertainty. In both cases, the agents will weigh the tradeoffs and make the best decision given the situation.

To do so, each agent generates its beliefs and behavior by solving a partially observable Markov decision process (POMDP). A POMDP consists of *state*, *actions*, *transition*, *observations*, and *reward*. The *state* of a POMDP in UrbanSim represents various features of the different entities (e.g., a structure's capacity, a group's military power, a person's political support), each a real-valued number from -1 to 1 (e.g., 1 means that the structure is functioning at 100% of capacity). The *actions* are the choices available to each agent (e.g., repair a structure vs. patrol a neighborhood), and the *transition* represents the effects of these different choices on the state. The *observations* capture the probability that certain states and actions are hidden from certain agents. The *reward* function represents what each agent seeks to achieve in the world (e.g., maximizing its own security, minimizing an enemy's military power).

Given a set of such POMDP models for the entities in the scenario, each corresponding PsychSim agent can use standard algorithms to compute its best course of action (Kaelbling, 1998). These algorithms operate by projecting the effects of candidate actions into the future, aggregating the reward resulting from those effects (as well as the effects of the anticipated responses by the other entities), and selecting the action with the highest expected reward.

In UrbanSim, the POMDP models of the underlying 'society' were created by non-technical subject-matter experts (SMEs) specializing in Iraqi and Afghan cultures. The baseline UrbanSim scenario, al 'Hamra, contains 92 non-player agents. This decomposes to over 1400 features and another 4700 possible actions. This quickly grows exponentially to almost 450,000 possible effects of actions. In any given turn, there are approximately 1000 different actions that the player can choose from, and 1100 possible responses for the agents.

One of the core features of PsychSim, and key to several of the learning objectives of the game, is the ability to generate causality chains of actors to capture both intended and unintended effects of agent actions. Many AIs in games and simulations cover *who, what, where* and *when* quite well. However the *why* (which elicits causality) has proven allusive due to the complexity of modeling human cognitive function in the minds of non-player characters. For entertainment game AI this rarely poses a problem. NPCs are often tactical in their behavior and do not require elaborate decision-making capability to execute actions like selecting their weapon, moving to contact, and even basic formation control. However for social simulation-based training applications, agents represent individuals and groups in a society with myriad beliefs, desires and intentions that must work together to produce coordinated, plausible action. Moreover, these actions must have meaningful effects that can be described to the user in a way that adds training value. In UrbanSim, this is accomplished through causality chains. Causality chains help establish concrete linkages between an agent's actions and their effects, and if the linkages extend multiple turns and are part of a fully connected societal graph, it allows us to address issues related to first-order, second-order and third-order effects.

## Story Engine

Even with a tight control over authoring in terms of the available actions to PsychSim agents and their goal structure, the scale and scope quickly becomes difficult to manage. Though agents were taking actions that were plausible and contributed towards the pedagogical experience of the application, there were specific instances that project SMEs (former commanders) wanted to highlight to students as they played through the experience. This was difficult to force with a complex multi-agent system, and eventually led to the development of the Story Engine. The Story Engine was specifically designed for instructors and SMEs to incorporate real-world events and situations in the game. These events could be strung together to form stories that would play out over multiple turns.

The Story Engine uses as input variable states from PsychSim agents. The figure below presents an example where PsychSim agents took action to kill an Iraqi police officer. The event checks for when this condition occurs and then launches a storyline that involves conducting an investigation and working alongside the police chief to determine what happened. These events were authored to always occur, regardless of what action(s) the users or agents in the game take. The Story Engine is intended to convey key teaching points related to the learning objectives of the game.

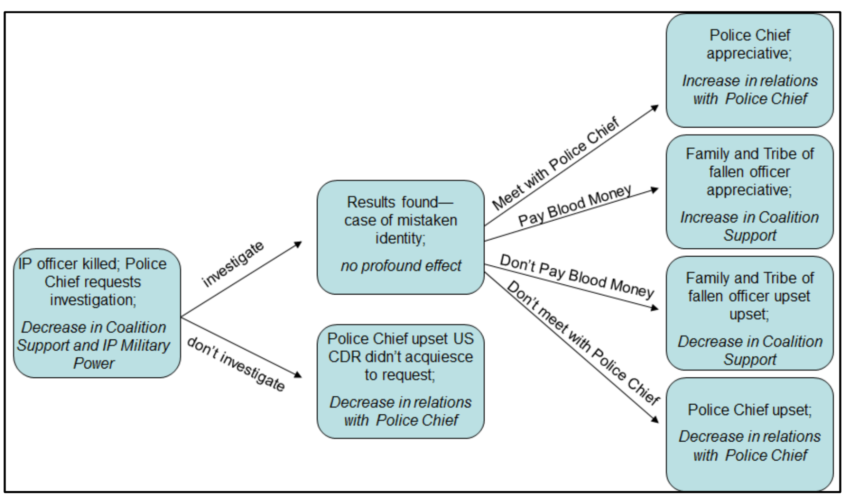


Figure 22: PsychSim Example

The use of dual AI technologies to drive UrbanSim has allowed scenario authors and instructors to tailor the experience for certain audience types. In cases where UrbanSim was used to train operational commanders during Staff Exercises, the heavy reliance on story events derived from similar real-world situations they might encounter was important in helping them and their staff prepare for conditions downrange. In cases where classroom instructors were simply covering the basics of MC with no specific operation or region in-mind, the diversity and richness of the multi-agent system was sufficient for students to gain an understanding of the complexities of the environments in which they may find themselves in one day.

## Way Ahead

The strategic trajectory of the United States in terms of military engagements remains uncertain. It will obviously be influenced by geopolitical currents that we may or may not have influence over. TRADOC contends that the strategic environment will be "characterized by multiple actors, adaptive threats, chaotic conditions, and advanced-technology-enabled actors seeking to dominate the information environment. The Army must be operationally adaptive to defeat these complex challenges and adversaries operating within this environment" (Operational Environments to 2028, 2012). The National Intelligence Council outlines four potential worlds in 2030, influenced by 'megatrends:' individual empowerment; the diffusion of power; demographic patterns dividing the world into zones of population growth and others with stable or even declining populations; and a food/water/energy nexus that will lead to increasing competition for these commodities in places (Global Trends 2030, 2013).

1. "Stalled Engines" (a worst case scenario in which the United States draws inward, globalization stalls, and the risks of interstate conflict increase);
2. "Fusion" (the most plausible best-case scenario in which the United States and China collaborate on a number of issues leading to broader global cooperation);
3. "Genie-Out-of-the-Bottle" (inequalities within and between nations explode and the United States no longer manages world order); and,
4. "Nonstate World" (driven by new technologies, nonstate actors surpass states in confronting global challenges).

As Metz points out in the Strategic Landpower Task Force Report (2013), the most likely opponents of the US military are hybrid compositions of militaries and nonmilitary entities, or 'evolved irregular threats' (Flynn, 2011). They will be highly complex, adapt rapidly, rely on asymmetric methods, and often operate in congested urban areas.

As training and technology continue to evolve alongside emerging threats from this futurescape, one important capability is being able to accurately model the social environments in which we may find ourselves. Though social simulation has morphed significantly since the days of the Von Neumann machine and Conway's Game of Life, investment must continue from a cross-section of disciplines (sociology, psychology, anthropology, computer science) to make social simulations a mainstay in future training solutions. Additionally, with the influx of big data from all corners of the globe via social media, there is a unique opportunity to incorporate it into our modeling approaches. For example, combining data mining with social media analysis techniques, we could adjust non-player agents to make choices based on specific locations: Dhaka and Cairo might have very different responses probabilities to the same situation. Not only has social media been shown to instrument change in the real world (Casilli & Tubaro, 2012), it provides the social simulation community with a valid and useful tool for developing and tuning their models. Research in this space remains scant, though as this data becomes more available and researchers understand its utility (and limitations), we can expect it to be a core foundation of social simulations in the future.

**FASTEST GROWING MEGACITIES IN THE WORLD**
**(Urban Areas with more than 10 million residents)**

| Rank | Geography | Urban Area | Population Estimate | GROWTH (Decade) |
|---|---|---|---|---|
| 1 | Pakistan | Karachi | 20,877,000 | 80.5% |
| 2 | China | Shenzhen | 12,506,000 | 56.1% |
| 3 | Nigeria | Lagos | 12,090,000 | 48.2% |
| 4 | China | Beijing, BJ | 18,241,000 | 47.6% |
| 5 | Thailand | Bangkok | 14,544,000 | 45.2% |
| 6 | Bangladesh | Dhaka | 14,399,000 | 45.2% |
| 7 | China | Guangzhou-Foshan | 17,681,000 | 43.0% |
| 8 | China | Shanghai | 21,766,000 | 40.1% |
| 9 | India | Delhi | 22,826,000 | 39.2% |
| 10 | Indonesia | Jakarta | 26,746,000 | 34.6% |

Figure 23: Top Ten Fastest Growing Megacities

As highlighted in (Kotkin & Cox, 2013), of the top 10 fastest-growing megacities in the world, all are either in Asia or Africa. 10-year growth in these areas is between 35 and 81%, yet our understanding of these regions remains limited. Obvious cultural and social differences make them unique to study, and combine this with complex and opaque political and economic structures, there is a need to find alternative approaches to developing our understanding of these locales. One approach involves using data mining and scrubbing techniques to help a cross-disciplinary team of anthropologists, demographers, social scientists and engineers develop models of populations residing in these areas.

For UrbanSim, work is underway to develop a new set of models and scenarios based on TRADOC's Common Training Scenarios (CTS) framework. The CTS is an expansive set of use cases that cover a variety of operation types from major combat to stability to disaster relief. At its core, CTS attempts to be both broad and deep in its coverage of potentialities. USC-ICT is working alongside social scientists, data miners and military SMEs to develop a stability-focused scenario with a strong non-state and coalition focus. One difference in the design approach from previous scenarios will be the reliance on social media feeds from the area of interest – in this case Georgia, Armenia, Azerbaijan, Turkey (GAAT) – to seed the modeling of the underlying behavior model. The scenario is scheduled for release in mid-2014.

## Acknowledgments

This work was supported by the United States Army Research, Development, and Engineering Command. Statements and opinions expressed in this material are those of the authors and do not necessarily reflect the views of the Government. No endorsement should be inferred.

## References

Casilli, A. & Paola T. (2012). Social Media Censorship in Times of Political Unrest: A social simulation experiment on the UK riots, Bulletin of Sociological Methodology, 115 (1).

Clark, R. (2004) Design Document for a Guided Experiential Learning Course. Final report on contract DAAD 19-99-D-0046-0004.

Flynn, K. (2011). The Evolved Irregular Threat Project, Paper for the Office of the Secretary of Defense. Washington, DC: Noetic Corporation, 2011.

Kaelbling, L.P. Littman, M. L., and Cassandra, M. L. (1998). Planning and acting in partially observable stochastic domains. Artificial Intelligence, 101:99–134.

Kotkin, J. & Cox, W. (2013). The World's Fastest Growing Megacities. New Geography, April 2013. http://www.newgeography.com/content/003630-the-worlds-fastest-growing-megacities.

Kilcullen, D. (2006) Twenty-Eight Articles: Fundamentals of Company-level Counterinsurgency, Military Review, May-June 2006 (pp 103-108).

Metz, S. (2013). Strategic Landpower Task Force Research Report. Accessed 6 October 2013. http://www.strategicstudiesinstitute.army.mil/index.cfm/articles/STRATEGIC-LANDPOWER-TASK-FORCE/2013/10/3.

Operational Environments to 2028: The Strategic Environment for Unified Land Operations. Fort Eustis, VA: Training and Doctrine Command, 2012.

Premack, D. G.; Woodruff, G. (1978). Does the chimpanzee have a theory of mind? Behavioral and Brain Sciences 1 (4) (pp 515–526).

Smith, R. (2005) The Utility of Force: The Art of War in the Modern World. London, UK: Allen Lane, 2005.

Wang, N., Pynadath, D. V., & Marsella, S. C. (2012). Toward automatic verification of multiagent systems for training simulations. In Intelligent tutoring systems (pp. 151-161). Springer Berlin Heidelberg.