Agent-based Simulation of Geo-Political Conflict

Glenn Taylor[1], Richard Frederiksen[1], Russell R. Vane, III[2], Edward Waltz[2]

[1]Soar Technology, Inc., 3600 Green Court Suite 600, Ann Arbor, MI 48103
[2]General Dynamics – Advanced Information Systems, 1400 Key Blvd, Suite 700, Arlington, VA 22209
glenn@soartech.com, rdf@soartech.com, russ.vane@gd-ais.com, ed.waltz@gd-ais.com

Abstract

The intelligence analysis task of anticipating crises and providing decision makers with reasonable (supportable, explainable) possible futures is extremely difficult. To perform this task, a team of analysts must consider the political, economic, and military aspects of national governments and non-governmental organizations. This paper describes an agent-based simulation framework, the Advanced Global Intelligence and Leadership Experiment (AGILE), for building executable models for conducting regional analysis. AGILE is a full-featured simulation framework that enables the specification of simulation parameters, models and agents, lets the user define and run simulations under varying conditions, and enables post-run analysis of the results. In this paper, we describe the system implementation, and some examples of its use in a pseudo-operational setting.

Introduction

Intelligence analysis includes abductive processes, in which analysts propose alternative competing hypotheses to explain current evidence. Long-term analysis of foreign nations requires the analysis of state interactions, especially where those interactions lead to instability, crisis, or conflict. In order to explore potential future scenarios, an analyst must understand the political, economic, military, diplomatic, and social aspects of those countries, with special focus on the major decision-makers and their motivations, capabilities, and tendencies.

This already difficult task is compounded by many circumstances. The information available to analysis is typically sparse, and buried within a vast amount of irrelevant, and often intentionally misleading, information. The analytic process is resource bounded – too few analysts must analyze a large amount of information in a short timeframe. In considering multiple hypotheses, limited resources may lead to considering only the most likely few, which might eliminate low-probability, high-impact hypotheses. While some regions have been analyzed for years, regional instability and regime change may require

analysts to quickly familiarize themselves with new decision-makers in the regions.

Several tools and techniques are currently available to an analyst to help perform the task of state and leadership analysis, including social network analysis, role-playing ("Red Teaming") and wargaming. However, most of these techniques are performed manually, with only basic tools such as databases and search engines used to assist. Recent government R&D programs for the intelligence community have focused on dealing with the overwhelming amount of information available to analysts, such as text summarization, query answering, and information retrieval, as illustrated by programs such as NIMD (ARDA 2002a) and AQUAINT (ARDA 2002b). Less attention has been paid to providing tools to help analysts with the abductive reasoning processes used in the analysis of feasible future scenarios.

This paper describes an implemented system meant to address this shortcoming. The Advanced Global Intelligence Leadership Environment (AGILE) is a simulation environment in which analysts may model the dynamics of a geopolitical crisis, and play out the kind of role-playing and wargaming that is typically conducted manually. AGILE enables a team of analysts to make explicit their tacit assumptions about a state leadership in the form of an executable agent-based model, and use that model to generate and analyze hypotheses about interacting nations. This work focuses not only on the capabilities of the simulation system, but also how the system fits into a larger analytic process. As such, a significant portion of this effort is in making the system usable and intuitive to intelligence analysts. This work is very much along the lines of scenario planning (Schwartz 1991) and fits within work on the practice of intelligence analysis (Heuer 1999).

System Overview

AGILE is a complete simulation environment, where a user can set up, run, and analyze simulations meant to represent a particular region in the world. In order to reflect the analytic process of exploring possible scenarios, we have developed AGILE to provide intelligence analysts with a "simulation workspace" for encoding knowledge,

generating hypotheses, and drawing conclusions to support their reasoning. Specifically, AGILE is designed to meet the following goals:

- Support abductive reasoning about complex nation-state relationships
- Allow analysts to model nation-state interactions and processes at a high level of abstraction
- Support the collaborative nature of intelligence analysis

The national-strategic level of interest to regional analysis has been one of our overriding design considerations. Rather than being focused on making precise predictions about future events, regional analysis is interested in larger trends, toward crisis or conflict. Low-level, tactical behaviors are less important when trying to understand these trends. As with any simulation, the goal is to model at a reasonable level of abstraction, ignoring unnecessary details, without sacrificing meaning or clarity. The design of AGILE has been greatly influenced by these ideas. For example, in our simulations built to date, numerical models represent high-level, aggregate information such as average military tension in a region, or gross domestic product trends. Similarly, agents tend to represent aggregate decision-making entities such as branches of a government, rather than individual decision-makers.


Figure 1: AGILE high-level architecture

AGILE is composed of three major components: a world model for simulating processes in the environment; agents (and their interactions) for simulating decision-making; and a graphical user interface that allows specification, execution, and analysis of the simulation (see Figure 1). The simulation environment consists of multiple countries (represented as groups of agents) interacting through the world model.

For simplicity, we have designed AGILE to run as a turn-based simulation, with two phases: the world model processes and updates its state, then the agents go through decision-making and post actions to the world model. After all agent decision-making is complete, the world model processes the agent actions, and the cycle repeats. The simulation runs for a user-specified number of turns before completing. Each turn corresponds roughly to a week of activity. Currently, the both the world model and the agents operate deterministically.

The typical process of running a simulation includes specifying the model's initial conditions, and then (via built-in wizards) defining variations on those initial conditions, such that a wide array of simulations can be quickly generated to explore a large space of possible scenarios. This variation also enables the user to account for uncertainties in the starting conditions, and supports a kind of sensitivity analysis on the models. AGILE currently allows automatic parametric variation of at most two initial world model conditions, and arbitrary manual variation of any system parameters, including initial conditions, agents and events.

The following subsections describe the major components of AGILE in detail, and how they serve to meet the objectives outlined earlier.

World Model

The world model (WM) represents the governing rules of geopolitical conflict, and provides the environment for agent interaction. The WM computes change over time, and maintains state from turn to turn. The WM can be composed of sub-models; the instance we describe later is composed of sub-models that correspond to economic, military, diplomatic, and social domains. Most importantly, the WM defines how all of the domains are interconnected. For example, if the level of nationalistic propaganda from Country A is very high, the apprehension level in Country B may rise, which, in turn, may dampen the Country B's economy. These model dependencies allow us to reflect the deep interconnections between nations in a region.

The world model is a well-defined, separable component within the system architecture, and can be replaced by models of different fidelity. We have developed an infrastructure for defining many different instances and types of world models, including numeric, symbolic, agent-based or hybrid models, to act as the underlying engine for the AGILE simulation. However, to date, we have built only numerical world models, composed of variables and their relationships.

During the world model update phase of a turn, the variable values are updated, and changes are propagated to other affected variables based on the WM definition. A subset of the variables in the WM serves to define actions for agent activity (an action is simply a change to a single variable in the world model). In addition to fielding the actions of the agents, the world model allows the injection of exogenous events – timed changes to the system that are not the result of agent modifications. For instance, an earthquake can be modeled using an event to modify regional effects, such as lowered stock markets and heightened economic apprehension. Events also let the user model simple forces external to the particular target countries without having to add several agents to model a country; for example, in the models described below, we model the presence of the US military as events, instead of by creating a suite of US agents.

Agent Models

Agents embody decision-making aspects of nation-state interaction, and are the driving force behind the dynamics of the system. We adopt here a strong notion of agency, as given by (Bratman 1987), who defines agents as autonomous entities that maintain goals, knowledge, beliefs, and obligations, and can interact with their environment and with other agents. The agent system is built using the Soar agent architecture (Rosenbloom et al. 1993), a general architecture for intelligent systems, which provides mechanisms for problem solving, knowledge inferencing, and goal-directed behavior. Soar has been used successfully to model a wide variety of human behavior, including behavior for military simulation (Tambe et al. 1995; Hill et al. 1997; Jones et al. 1999; Taylor et al. 2001). We selected Soar as the basis for the agents in this system due to its track record in representing complex, human-like decision-making.

While Soar, as a general agent system, provides many capabilities, we do not expect the end-users to be experts in developing knowledge-based systems. Therefore, we must allow a user to describe agents in familiar terms, such as goals, beliefs, and relationships to other agents. For our system, it was necessary to abstract away from the low-level details of Soar programming by exposing only the artifacts necessary to specify an agent. In order to do this, we essentially developed an agent architecture on top of the basic Soar architecture. The result is the two-tier architecture shown in Figure 2.
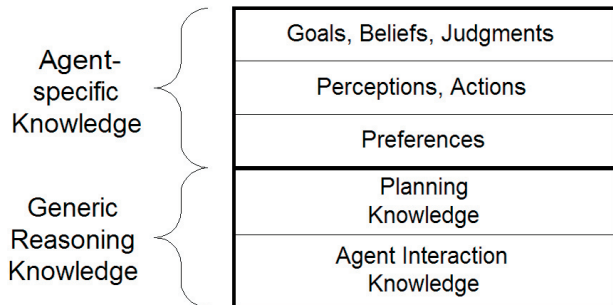

Figure 2: Agent knowledge architecture

The bottom layer consists of capabilities that govern the basic behavior of an agent in its environment and within a group, including low-level message handling, action execution, goal monitoring and management, belief inferencing, and plan generation and execution. This layer is essentially the same for every agent defined in this system (though there are some differences based on role). The upper layer is the data for the lower layer – it contains the knowledge that specializes an agent's behavior to a certain domain of activity. There are five basic types of knowledge defined in the upper tier:

- The actions available to the agent
- The agent's goals and preferences
- Beliefs to help the agent determine appropriate values for actions (action beliefs)
- Beliefs for the agent to determine the effects actions have on the world model (cause-effect beliefs)
- Beliefs for evaluating states, including the current and future projected states (judgment beliefs)

The basic agent activity revolves around meeting goals: an agent senses information from the world model, evaluates the current state, including what goals are met and unmet, plans for any activities that might help meet its unmet goals, and executes actions to meet its goals. Am agent's particular domain-specific knowledge guides the basic problem-solving process. Figure 3 below shows the execution cycle of an agent, with the knowledge that is applied at each step. Because the base architecture is still Soar, knowledge in both of the upper and lower tier is ultimately in the form of rules; the user specifies knowledge in the upper layer in an architecture-neutral format, and that knowledge is translated to Soar rules automatically.


Figure 3: A single agent's decision-making process

Each country may be composed of multiple agents acting as a group to solve problems. Within a group, each agent has a specialty role, and acts within a larger group decision-making context. Each agent role roughly corresponds to a sub-domain of the world model. There is no limit to the number of countries in a model, and so no limit to the number of agents within the system – however, the system enforces a policy that no two agents can have write access to the same world model variables.

We have so far implemented only a single organization structure, which we have called the "management" style of decision-making – a single moderator (the "leader") and any number of subordinates ("advisors"). In this organization, each advisor maintains a level of influence with the leader, and the leader gives more deference to the suggestions of agents with higher influence. An advisor's most important unsatisfied goal is nominated as an issue to the leader, who collects issues from the entire group (and itself), and prioritizes them by order of agent influence. One at a time, each issue is then suggested to the whole group for problem solving – all the agents in the group can suggest a way to deal with the issue. The leader collects all of the suggested actions, and selects the one that looks most promising, partly based on influence of the suggesting agent. This selected action is then passed through a review process, in which each agent in the group is allowed to vote for or against that action, based on analyzing the action against its own goals and beliefs. The leader tallies all these votes, and if the action passes the review, the agent that suggested it schedules

that action for execution. This process continues until all of the suggested issues have been dealt with for a given turn. Figure 4 illustrates the lead-advisor relationship. This multi-agent action selection and review process is similar in spirit to the FIPA Contract Net Protocol (FIPA 2002), except that all participants in this process are stakeholders, and so have a say in whether an activity happens or not. In fact, all inter-group communication in AGILE is governed by a set of well-defined communicative acts that reflect the role distinctions of leader and advisors.
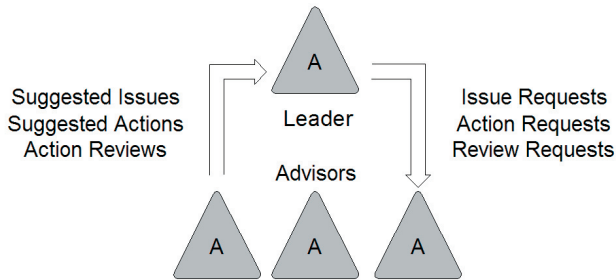

Figure 4: AGILE's muti-agent organization

In addition to intra-country communication, agents in one country can interact explicitly with agents in another by means of diplomatic communications. Two countries can establish diplomatic agreements in which both sides agree to change some goal or activity; for example, limiting the presence of foreign troops near the border or putting a cap on the tariff levels imposed on goods. These agreements begin as requests for one country to change its behavior, which enters into the normal within-country decision-making process as a new issue. If both sides have such requests, and both sides find the terms amenable, a formal agreement is established, and new goals are formed on both sides to accommodate it.

We have adopted the metaphor of beliefs, desires, and intentions to describe the characteristics of agents and how they behave in their environments. However, we have not adopted the full extent of the BDI paradigm as described in (Rao and Georgeff 1991). Instead, we have found that describing agents in terms such as beliefs and desires (goals), similar to what is described in (Dennett 1987), has an intuitive appeal, and finds resonance with our users. While Soar-based models have not traditionally been described in these terms, it is not unreasonable to do so, especially in this case where we have built a knowledge architecture atop the traditional Soar architecture to support these artifacts directly.

Graphical User Interface

The primary purpose of the AGILE graphical user interface is to provide an environment within which an analyst can generate and test hypotheses in simulation. Within this environment, the analyst can create and modify agents,

agent groups, world models, and pre-defined events; quickly set up and run multiple simulations; and analyze the results. The user can also annotate artifacts in the system, such as with comments or references, in order to fit the model into the larger analytic process. As analysis is ultimately a collaborative process, the user can also share these artifacts with other analysts, including model components (agents, world models), results, and annotations. For the discussion here, there are two major relevant GUI capabilities: the specification of simulation parameters, especially agent behavior; and understanding simulation results, especially the agents' decision-making processes.

Because of the two-tier organization of the agent as described earlier, the user only needs to specify the information in the upper tier, since the underlying processes are the same for all agents, and are not exposed to the user. The user must define the actions, goals, and beliefs for an agent, all in terms of the variables defined in the world model. Additionally, the user specifies descriptions for each of these knowledge types, which are used later in the explanation of the results. The GUI assists in this entire process, using various interfaces to simplify what could otherwise be quite complicated.

As an example, suppose the user wishes to define an agent to manage the diplomatic affairs of a country. Suppose also the world model contains, among other things, a variable called Propaganda representing the diplomatic rhetoric level toward another country, with value range (0.0, 1.0). The Propaganda level directly effects the Apprehension level of the economy – high Propaganda means high Apprehension. The user assigns this action to the Diplomat, and creates an accompanying belief that says how to change the Propaganda level under different conditions. The user then assigns a goal to the Diplomat agent to keep the Propaganda level below 0.5, otherwise too much Apprehension dampens the overall economy, and has a diminishing effect on the other country. Next, the user creates agent beliefs that describe the relationship between Propaganda and Apprehension, so the agent will know how to judge the effects of applying the action to the world model. The agents we have developed so far require only a few such goals and accompanying beliefs in order to exhibit interesting behavior. There is no real limit to the number of beliefs and goals an agent can have; however, since goals currently correspond only to variables in the world model, there is typically a correlation between the number of goals an agent will have, and the number of variables in the world model for which it is concerned.

After a simulation has been run, the GUI allows the user to analyze the results in order to relate them back to the analyst's original hypotheses. We have developed a variety of displays for the different types of data generated by the simulation. For example, the results of multi-variation runs can be displayed by a 3-dimensional surface, and the dynamics of the world model are easily described as 2-

dimensional timelines, with each variable shown as a single graph covering the duration of the simulation. Two types of agent information overlays annotate the 2-dimensional timeline: regions of the parameter space covered by agent goals, and changes in world model variables that were the direct result of agent actions and events.

Another level of detail describes the agent activity using textual descriptions in "narratives." Summary narratives include information about the decision-making within a country, turn-by-turn including the issues that were brought up and how they were dealt with. Detailed narratives show information about a decision on an agent-by-agent basis, where agent activity is given side-by-side in columns – which issues were suggested, what actions were suggested to fix that issue, how the agents reviewed other actions, and which action was finally selected and executed (see Figure 5). To give further insight, this narrative also links to details about the precise beliefs an agent applied when making these choices.

	China Diplomat	China Military
Propose Actions	**Assess Status**	**Assess Status**
	Belief: If Taiwan statehood too high, increase propaganda to militant levels (Propaganda) *details*	*Belief:* Increased forced incidents to lower statehood (Military_Incident_Intensity) *details*
	Goal: Want Taiwanese Statehood Low	*Goal:* Want Taiwanese Statehood Low
	Propose Action Set Propaganda to -0.094 (L)	**Propose Action** Set Military_Incident_Intensity to 1.25 (L)

Figure 5: Snapshot of Detailed Narrative Swimlanes

All of the results displays, from the 3d surface plots, to the 2d timelines, down to the individual agent decision calculus, are interconnected to allow the user to drill down for more detail about a particular simulation. This ability to easily navigate to any level of detail enables the user to quickly understand the full dynamics of the simulation.

Evaluation

We have used AGILE to model a few scenarios for testing and for demonstrating its capabilities to potential users. First, with the help of analysts, we modeled a particular historical scenario, the 1995-1996 Taiwan Straits Crisis, based on open source data, to see if AGILE was able to model the important events. Later, we staged a hypothetical crisis between China and Taiwan in 2005, and worked with analysts in a simulated Red Team to use AGILE as the workspace for their hypothesis generation and testing.

Taiwan Straits Crisis of '95-'96

This scenario detailed the largely diplomatic conflict that ensued over the Chinese perception that Taiwan was making a bid for independence. The crisis involved rising diplomatic and military tensions, including Chinese missile test firings across the Taiwan Straits as a way to subtly threaten Taiwan into backing down on its perceived move toward independence.

In order to model these events, we developed a simple numerical world model that captures the relationships and interconnections between economic, social, military, and diplomatic processes. We created a mapping from the historical record into the terms of the simulation in order to create a reference model of events for comparison. The reference model captures the historical data at the same level of abstraction as the simulation, so we can compare historical cases against simulation outputs. The important variables we modeled included the China's level of perceived Taiwanese independence, the movement of troops into and out of the theater, the level of military incidents in the Straits, the diplomatic rhetoric level (propaganda), and economic indicators such as GDP and apprehension. Additionally, we included the presence of US troops as part of the numeric model, since we were not interested in US decision-making in this exercise. We then developed a set of agents, five for each of China and Taiwan – one for each of the four world model domains, and one leader – and assigned those agents with the goals and beliefs that were suggested by our data sources. A major Chinese goal was keeping Taiwanese independence in check without going to war; a major Taiwanese goal was to see how high it could raise its perceived independence without invoking a drastic response from China.

We initialized the world model with historically informed parameters, and generated a parametric variation space around the initial conditions to account for our uncertainties in those parameters. We ran each simulation for 65 turns (equivalent to weeks) to match the rough span of time taken by the actual historical events. Using this multi-simulation method of 100 simulations, we were able to generate a wide range of system behaviors, a handful of which matched quite closely to the actual historical events, and whose underlying rationale fit well against the historical data. Figure 6 represents a comparison of the reference model we created (left) against the results of a single simulation run (right). A visual comparison of the reference versus evaluation models demonstrates the close similarities exhibited by the agent behavior and the model dynamics. We have begun to develop tools that help identify these individual runs within the multiple simulations. One example is the Interesting Run Finder, which allows the user to find simulation runs that match certain criteria, such as when a variable's value falls within a defined range. Qualitative descriptors in the model definition identify meaningful ranges of values – for
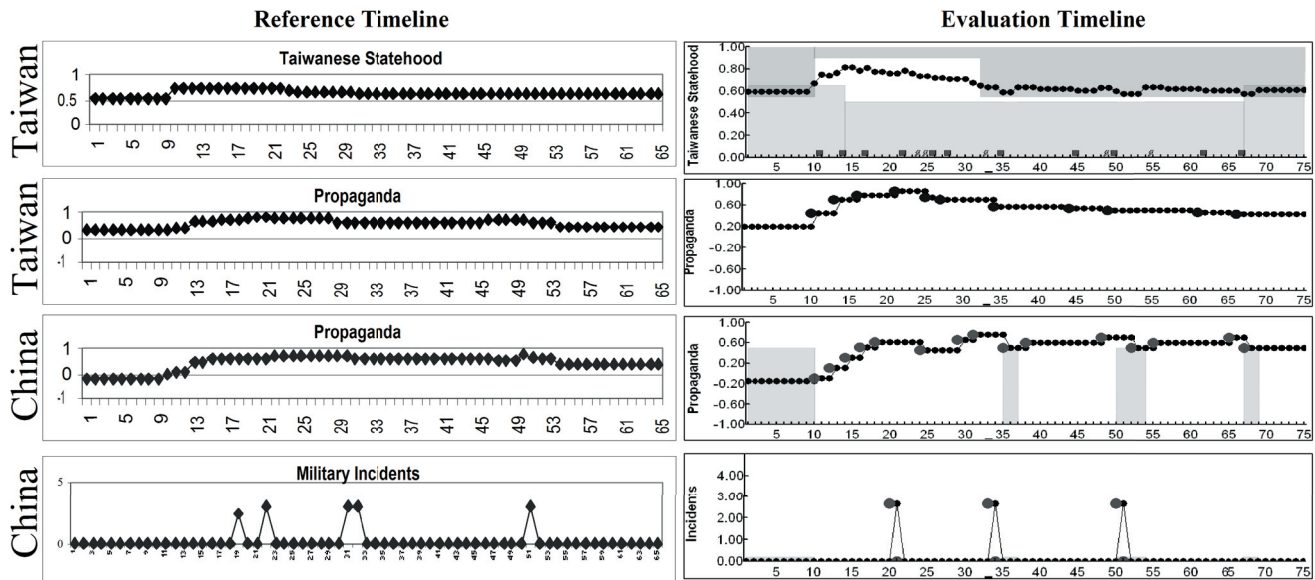
Figure 6: Comparison of Reference Model to '95-'96 Taiwan Straits Crisis Model

example, in this model, values of Military Intensity above 0.8 would indicate those runs in which full war broke out.

Encouraging was the fact that with fairly simple agent models (roughly 5-10 goals, 40-50 beliefs, and 5 basic actions per agent), and a fairly simple world model (about 15 interconnected variables), we were able to achieve quite rich, believable system behavior. This initial combined model took roughly one month's time of a software developer to build, test, simulate, and analyze the results. We have not yet run usability tests with end-users, so we do not know how this would translate to the field.

China-Taiwan 2005

In this exercise, the analysts imagined a future scenario where the current Taiwanese president had won a second term in office, and Taiwan's Democratic People's Party had introduced to the legislature an enabling bill for a future referendum on national independence. The goal was to exercise the model's ability to explore what-if scenarios and suggest believable possible futures for analysis. Using the same world model as in the previous example, the analysts suggested initial conditions and brainstormed on the mindsets of the major decision-making entities, which we translated into viable agent beliefs and goals. Over the course of three days, the analysts would suggest changes to the base model as what-if scenarios, which we encoded into the model, and then ran as simulations whose results were presented the following morning.

Figure 7 illustrates Chinese military activity as multi-run surfaces taken from the economics excursions of the exercise. Each surface represents a set of parametrically varied initial conditions for the world model, and each grid location on surface a represents the maximum value of the Chinese Military Incidents variable over a single simulation run – essentially showing China's use of military force to quell

Taiwanese moves toward independence. The four surfaces reflect different agent economic beliefs and goals. As the series progresses, both China and Taiwan increase their level of concern about the economic aspects of their relationship. *(a)* shows the base run for the exercise, indicating generally high military incidents; *(b)* displays the results of increasing the influence of the Taiwanese Economic advisor, with an effect of decreased military incidents; *(c)* shows the effects of adding an additional belief to the Taiwanese Economic advisor to dislike Taiwanese Propaganda increases; and *(d)* displays the results of making every Chinese actor consider (at a medium importance level) the economic effects of their actions. These results indicate that Chinese military coercion decreases as the importance of the economic interests increase. This corresponded to the analysts' hypothesis
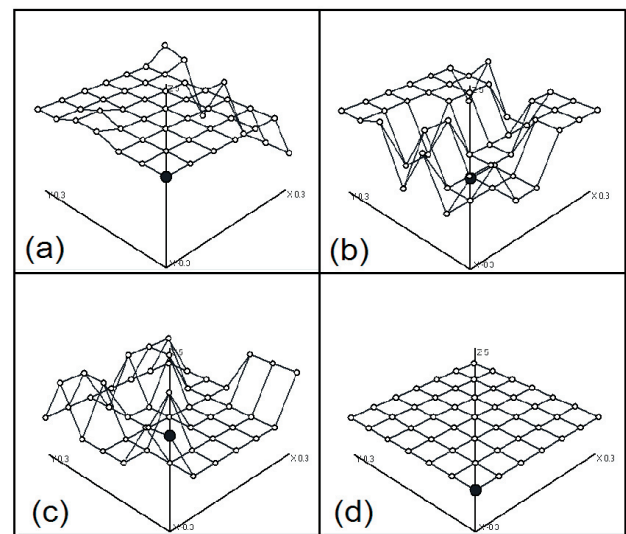

Figure 7: Multi-simulation surfaces showing the effect of growing economic importance on China's military actions

that China has resisted military invasion of Taiwan because of possible negative economic consequences. As with the previous example, we used the Interesting Run Finder to identify runs of interest to the analysts.

We have not yet conducted a thorough quantitative evaluation of these models. The next step forward will be traditional verification and validation—does the built model match the analyst's theory, and does that theory describe the real world? So far, we have only used face validation by intelligence analysts as a technique for validating the example models, to see that they "make sense." However, face validation is fraught with problems, due to the complexity of the system and human biases. Quantitative evaluation, on the other hand, can be done at multiple levels, from looking at just the numeric outputs from the world model (to know that the system dynamics are valid), to validation of the agent's knowledge and decision-making processes. Human behavior validation is known to be a difficult problem, and while face validation is the most widely used method, other methods have been suggested (DMSO 2001). In the absence of data at the same abstract level as the model (as is the case here), raw historical data can be re-cast into abstract simulation terms as a reference model. From this, we can perform regression testing against the system output. It must be noted, however, that the abstraction from historical data into a reference model is itself subject to interpretation and bias, so one must be careful in comparing to abstracted reference data or any other selective data set. In the case of the Taiwan Straits crisis, there is little data available in a form that could be compared against the output of the abstract model. Likewise, we have not found other models of the Taiwan Straits scenario against which we might evaluate the results of our models. Validation in general remains a challenging problem for these types of systems. Because we intend AGILE to be deployed to analysts, we expect to include extensions to help users evaluate the goodness of the models they build.

Discussion

The process of modeling real phenomena in simulation is always challenging, especially as the level of abstraction goes up – the ability to map from reality to an abstract model and back is often considered a black art. This seems especially true when modeling human decision-making. As with any simulation, there are tradeoffs between usability, fidelity, abstraction, explainability, and predictability. The constraint that analysts need to specify agent behavior illustrates this tradeoff. This constraint impacted the design of our agent system, which effectively limits the types of agents a user can specify in AGILE. That is, whatever capabilities are present in the lower tier of agent definition are the only capabilities the user can exploit when specifying the behavior of an agent. We expect that some new capabilities in the agent would require new types of

knowledge for agent specification, potentially complicating the job for the user. We have begun to understand some of these tradeoffs, but finding a multi-dimensional "sweet spot" remains a challenge.

In building a system like this, it is very easy to fall into a rut of designing something that makes sense to its designers but not its intended users. Not everyone thinks in the same way, so a display that makes sense to an engineer may not be intuitive to an intelligence analyst. Therefore, it is critically necessary to stay in active engagement with these users, to ensure we answer their needs directly, rather than our interpretation of their needs. This is certainly not a new idea, but a project such as this, where it is difficult to get time with our potential end-users, only reinforces its importance.

While working directly with analysts, we have encountered a range of responses to AGILE. Some users were not familiar with computer simulation as a support tool, so did not see a clear benefit to this approach. Others were reluctant to adopt a technology they perceived as possibly supplanting them. Others still were put off by the overhead of setting up and maintaining the simulations over time. Those who saw the benefit of the tool wanted to obtain it immediately, even in an unfinished state, as they saw its benefit not only as a simulation tool, but also as a way to encode and relate their own knowledge to an analysis problem. This only serves to illustrate that the challenges in building new tools are seldom solely technical or scientific.

Relationship to other work

Certainly the idea of using agent-based simulation for what-if analysis is not new – there is a rich literature on decision-making models for military simulation (Pew and Mavor 1998). However, little work has been done in agent-based modeling of national-strategic decision-making, where the military is but one component. Furthermore, agents in simulation have traditionally been difficult to employ, and specifying agent behavior for these domains is still regarded as something of an art. Much of our effort here has been in finding a good balance between the native capabilities of the agent and the ease in user specification of agent behavior

The agent system we describe here has taken inspiration from multiple sources, especially from the Multi-Agent Systems community (Casti 1996; Epstein and Axtell 1996) and organizational modeling (Prietula et al. 1998; Carley et al. 2001). Our approach can perhaps be described as a middle ground between the very simple agents of Complex Adaptive Systems (CAS), and much more sophisticated agent behavior, such as those traditionally built using cognitive architectures like Soar. In fact, CAS has been used to model some nation-level phenomena such as policy evolution (Axelrod 1997). However, our focus has been on the decision-making calculus of organizations and their members, rather than on the evolution of system behavior. Therefore, AGILE agents are more sophisticated than CAS agents. Our placement in that spectrum is driven by the need

for the agents to be, on the one hand, simple to specify and employ, and on the other hand, still capable of rich, goal-directed behavior that is transparent and explainable to the user.

Summary and Future Work

We have developed an agent-based simulation framework called AGILE intended for use by intelligence analysts engaged in understanding the complex relationships in geopolitical conflict. AGILE allows intelligence analysts to hypothesize about the relationships between countries, and to explore events and decisions leading up to conflict. Analysts can encode in these models their knowledge about the workings of the target countries, including the major decision-makers, and in this process make explicit their tacit assumptions about the target. This enables them to more readily share their conclusions and results with other analysts. Where, previous to this work, analysts used static formats such as databases and text reports, the idea of building executable models for estimation represents a new paradigm for intelligence analysis.

While we have not yet deployed AGILE to analysts for everyday use, we have engaged them at various points in the development process, to help us build and validate these models, to demonstrate the system, and to elicit their feedback. Based on those interactions, we are currently improving the system capabilities at all levels to more effectively implement the process of hypothesis generation and testing using an agent-based simulation environment. Overall, this effort has demonstrated the richness of the intelligence analysis domain for future AI research.

Acknowledgements

References

ARDA (2002a). Novel Intelligence From Massive Data (NIMD), Advanced Research and Development Activity.

ARDA (2002b). Advanced Question and Answering for Intelligence, Advanced Research and Development Activity.

Axelrod, R. (1997). The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration. Princeton, NJ, Princeton University Press.

Bratman, M. (1987). Intention, Plans, and Practical Reason. Cambridge, MA, Harvard University Press.

Carley, K. M., J.-S. Lee and D. Krackhardt (2001). "Destabilizing Networks." Connections **24**(3): 31-34.

Casti, J. (1996). Would-Be Worlds: How Simulation Is Changing the Frontiers of Science, John Wiley & Sons.

Dennett, D. C. (1987). The Intentional Stance. Cambridge, MA, The MIT Press.

DMSO (2001). VV&A Recommended Practices Guide, Department of Defense Modeling and Simulation Office.

Epstein, J. M. and R. Axtell (1996). Growing Artificial Societies: Social Science from the Bottom Up. Cambridge, MA, MIT Press.

FIPA (2002). FIPA Contract Net Interaction Protocol Specification. http://www.fipa.org/specs/fipa00029/. [Web document], Foundation for Intelligent Physical Agents.

Heuer, R. J., Jr. (1999). Psychology of Intelligence Analysis. Washington, DC, Center for the Study of Intelligence.

Hill, R. W., J. Chen, J. Gratch, P. S. Rosenbloom and M. Tambe (1997). Intelligent Agents for the Synthetic Battlefield: A Company of Rotary Wing Aircraft. Ninth Conference on Innovative Applications of Artificial Intelligence, AAAI Press.

Jones, R. M., J. E. Laird, P. E. Nielsen, K. J. Coulter, P. G. Kenny and F. V. Koss (1999). "Automated Intelligent Pilots for Combat Flight Simulation." AI Magazine **20**(1): 27-42.

Pew, R. and A. Mavor, Eds. (1998). Modeling Human and Organizational Behavior: Application to Military Simulations. Washington, DC, National Academies Press.

Prietula, M., K. M. Carley and L. Gasser, Eds. (1998). Simulating Organizations: Computational Models of Institutions and Groups. Menlo Park, CA, AAAI Press.

Rao, A. S. and M. P. Georgeff (1991). Modeling Rational Agents within a BDI-Architecture. Second International Conference on Principles of Knowledge Representation and Reasoning, San Mateo, CA, Morgan Kaufmann Publishers.

Rosenbloom, P. S., J. E. Laird and A. Newell, Eds. (1993). The Soar Papers: Research on Integrated Intelligence. Cambridge, MA, MIT Press.

Schwartz, P. (1991). The Art of the Long View: Planning for the Future in an Uncertain World. New York, NY, Doubleday.

Tambe, M., W. L. Johnson, R. M. Jones, F. M. Koss, J. E. Laird, P. S. Rosenbloom and K. B. Schwamb (1995). "Intelligent Agents for Interactive Simulation Environments." AI Magazine **16**(1): 15-39.

Taylor, G., F. Koss and P. E. Nielsen (2001). Special Operations Forces IFORS. Tenth Conference on Computer Generated Forces and Behavioral Representation, Norfolk, VA.