14th ICCRTS

"C2 and Agility"

Title:

Instantiation of a Sensemaking Agent for use with ELICIT Experimentation

Topic 6: Modeling and Simulation (Agent simulation)

Authors:

Mary Ruddy (Parity Communications), Danielle Wynn and Jimmie McEver (Evidence Based Research, Inc.)

Point of Contact: Mary Ruddy

Name of Organizations:

Parity Communications, Inc.
22 Bartlett Avenue
Arlington, MA 02476

Evidence Based Research, Inc.
1595 Spring Hill Road, Suite 250
Vienna, VA 22182

Telephone: 617-290-8591

E-mail Address: mary@parityinc.net

Title of Paper

Instantiation of a Sensemaking Agent for use with ELICIT Experimentation

Abstract

Development of and experimentation with ELICIT (Experimental Laboratory for Investigating Collaboration, Information-sharing and Trust) is an ongoing activity of the U.S. DoD Command and Control Research Program (CCRP) within the Office of the Assistant Secretary of Defense for Networks and Information Integration (OASD/NII). A recent CCRP-sponsored effort has resulted in the design and instantiation of a configurable Sensemaking Agent to enable agent-based simulation experiments to investigate the effects of agent personality on team behavior and performance. Agent-based ELICIT (abELICIT) models multiple stages of information processing and collaboration behavior undertaken as the agents share, receive, and integrate information elements to determine attributes (who, what, when, where) of a pending attack. The agent instantiation includes modeling of the agent's situational awareness. This paper will discuss specifics of the agent instantiation and provide suggestions for how the agent could be used in future ELICIT experiments.

Unclassified

Keywords: ELICIT, experimentation, edge organization, collaboration, agent

ELICIT Overview

Background

The United States Department of Defense Command and Control Research Program (CCRP) of the Office of the Assistant Secretary of Defense for Networks and Information Integration (OASD/NII) is engaged in developing and testing principles of organization that enable transformation from traditional hierarchy-based command and control practices toward the transference of power and decision rights to the edge of the organization. The need for agility in Information Age militaries is becoming increasingly important. As discussed in Understanding Command and Control (Alberts and Hayes, 2006), in an era of complex coalition civil-military operations, understanding how to organize for agility not just within a specific organization but also across differing organizations and cultures is a key to success.

Historically, there has been a shortage of formal experimentation data on the efficacy of so-called "edge" vs. traditional organizational styles. In order to enable and foster the conduct of formal experiments in the principles of organization, the CCRP created and continues to sponsor and maintain an experimentation environment known as ELICIT: the Experimental Laboratory for Investigating Collaboration, Information-Sharing and Trust. As part of this program, the CCRP (www.dodccrp.org) sponsored the development of a Java-based software platform that can be used to run multi-user experiments focused on information, cognitive, and social domain phenomena.

Experiment

The ELICIT experimentation platform has configurable scenarios that focus on the task of discovering the "who", "what", "where", and "when" of a future incident (e.g., a terrorist attack). Information in the form of "factoids" is provided periodically to each of the participants during a trial. The factoids and their distribution are structured so that no one participant receives all the information necessary to perform the task; thus, information sharing is required in order for any participant to be able to determine a solution to the ELICIT problem. Experimentation treatments can be varied considerably to represent an extremely wide range of organizational/command and control approaches. The environment generates detailed transaction logs; these, together with participant surveys that can be administered either prior to a trial (for calibration), after a trial, or *in situ*, can be used to measure information sharing, collaboration behaviors and situational awareness, as well as a variety of value metrics including the ability of individuals and teams to correctly identify the future adversary attack and the time required to do so.

ELICIT Program Resources for Ongoing Enhancements and Experimentation

The CCRP facilitated the formation of an international Community of Interest (COI) to accelerate the application of ELICIT toward the development of a better understanding of network-centric concepts, approaches, and capabilities as well as power-to-the-edge principles. The COI enables members to build upon each others ELICIT-related work, collaborate on campaigns of experimentation, and develop a shared body of knowledge.

Information about ELICIT, the COI as well as papers related to ELICIT can be found on the CCRP Website at http://www.dodccrp.org/html4/elicit.html.

New Sensemaking Agent

As part of the ongoing efforts, the ELICIT software platform continues to be enhanced based on input from the research community. This paper describes a significant new capability: the instantiation of a Sensemaking Agent for use in ELICIT experimentation. Previously, the ELICIT web-based software platform had been enhanced to allow agents to be used as participants in experiments, and a stub agent had been created to validate this capability. This paper focuses on a new Sensemaking Agent which has been developed. The Sensemaking Agent is a moderately intelligent, highly configurable tool designed to significantly increase the efficiency and effectiveness of the overall ELICIT research initiative.

Motivation for Enhancement, and Place in the Evolution of the ELICIT Software Platform

Since its inception, the evolution of ELICIT has been driven by the CCRP's vision of creating a platform for advancing the state of the art of experimentation in support of new approaches to command and control. Having an intelligent agent capability for use in ELICIT experimentation has been a long term goal.

One of the ongoing trends of the evolution of the ELICIT platform has been to provide an increasing level of configurability so the community of ELICIT researchers can configure ELICIT to meet their own specific research needs, while still leveraging a common platform that allows for the aggregation of a common body of research data for joint analyses. Researcher input to ELICIT enhancements has been very important. Numerous researchers have indicated the desirability of adding the ability to allow software agents to participate in ELICIT experiments.

One of the most expensive and time consuming aspects of conducting formal experiments on the effectiveness of varying organizational structures is recruiting human subjects. Furthermore, the ELICIT parameter space is enormous (there are over 130,000 different communications network configurations possible for a 17-person experiment, not to mention other relevant variables, such as trust, sharing personalities, team tactics, techniques and procedures (TTPs), etc. Even if human subjects were easy to come by, it is simply not feasible to use human experimentation to explore the features of this vast of a space. While human-in-the-loop experiments can not and should not be completely replaced by agent-based experiments, agents can be an important tool in an overall campaign of experimentation, enabling experimenters to conduct exploratory simulation runs to identify interesting regions of the problem space to help make the most of the limited numbers of human trials that can be conducted.

Agent participation can be of two types: Hybrid human and agent experiments and agent-only experiments.

Combination human/agent experiments

Arranging for a suitable number of qualified human participants to be available at one time is often challenging, especially for smaller organizations. Frequently, all of the recruited participants are not actually available at the scheduled time due to unforeseen circumstances. The ability to have an agent take the place of a missing participant greatly increases scheduling flexibility. In situations where ELICIT experiment trials are used pedagogically as part of a classroom exercise, it is also very useful to be able to run a set of experiment trials regardless of the number of persons registered for the class. (e.g., a class with 32 students can run a standard "edge" v "traditional hierarchy" baseline ELICIT treatment with the default 17-person setup if two agents are used to bring the number of participants to the required 34.)

Agent-only experiments

Agent-only experiments can enable orders-of-magnitude increases in the number of ELICIT experimentation trials that can be run. This can enable broader exploration of ELICIT-related phenomena of interest (for which appropriate agent behaviors have been instantiated) and for the generation of larger samples of ELICIT data from which broad trends and key factors can be identified that may be beyond the reach of small-N studies.

In addition, theories of behavior and performance can be validated by constructing agents with specific capabilities (personalities) and confirming the resulting outcome by running experiments using agents with those personalities. The ability to construct agents with specific behavior styles (information hoarders, members of a hierarchy with a narrow task focus, etc.) enables the design of experiments to investigate which organization types are more resilient under specific staffing scenarios.

Agent-based experiments are intended to augment and complement human trials -- not to replace them. Depending on the level of validation needed, insights gleaned from agent-only trials will need to be confirmed and refined with human-in-the-loop experimentation trials. Because ELICIT supports human-only, agent-only and combination human/agent experiments on a common platform it provides the ready ability to identify an experimentation space with human-in-the-loop experiments, drill down and explore a large number of variables with agent-only trials to identify the most promising scenarios, then run human-in-the-loop experiments for those most promising conditions.

Objectives of Sensemaking Agent

The goal of the design process for the Sensemaking Agent was to describe and instantiate a moderately intelligent, configurable agent. This agent not only needs to be able to take the place of a human participant in ELICIT experiments, but to actually form a mental model of the information in the factoids received and of the members of the group in which it was operating so it develops its own situational awareness. That is, as an agent participates in an experiment it needs to generate situational awareness that can be drawn upon to make decisions about behavior. The agent's behavior must also depend on the scenario in which it is operating, so that it behaves differently, as appropriate, in different

scenarios (i.e., the agent must not be scripted; that is, it must be able to respond to the scenario as it unfolds according to its "personality"). The Sensemaking Agent needs to be able to formulate ELICIT Identify messages based on awareness and understanding of the factoids to which it has been exposed.

A sub-objective of the Sensemaking Agent was that humans interacting with properly configured Sensemaking Agents as part of ELICIT experiment trials should not be able to tell that some of the experiment participants are agents – the agents should pass a "Turing test" [1] in the ELICIT context. Since the interaction of ELICIT agents with other ELICIT participants is limited to those actions that are allowed by the ELICIT platform (sharing factoids, posting factoids to information websites, receiving factoids from other participants, etc.), the agent behavior needed is much more narrowly focused than that of any agent that would actually compete in a more broadly defined (and more traditional) Turing test.

Still, for purposes of the ELICIT campaign of experimentation, it was desirable for the agent be able to perform all the actions (beyond interaction behaviors) of a human participant and since we are studying shared awareness, for the agent to actually construct an awareness of the adversary attack scenario based on factoids to which it had been exposed. This represents an innovative capability, since agent behavior is usually controlled by a (more or less) narrow set of rules.

Furthermore, in addition to having this awareness or "sense making" capability the agents also have additional configurable variables to define their personalities and style of social interaction with the other experiment participants. Using these variables, agents can be configured to operate in human timeframes, rather than just computer timeframes, show human levels of variability and human personality traits such as information hoarding, preferences for different kinds of sharing, trust in others, and how trust affects interactions (among others). We will draw on these Sensemaking agent capabilities to provide some examples in this paper.

Another sub-objective was that the Sensemaking Agent should be configurable so that its behavior can be entirely deterministic, so that if the same set of agents were run with the same task scenario and in the same roles in the same organization style, the results would always be the same. This is important to facilitate initial validation and calibration of the agent logic, and to enable employment of agents with activity profiles that can be held constant from trial to trial. That said, the Sensemaking Agent should also be configurable to behave stochastically where appropriate to the analysis and experiment.

---

[1] The phrase Turing test classically referrers to a test of human intelligence proposed by Alan Turing in his 1950 paper "Computing Machinery and Intelligence". Wikipedia describes the Turning test as "A human judge engages in a natural language conversation with one human and one machine, each of which tries to appear human. All participants are placed in isolated locations. If the judge cannot reliably tell the machine from the human, the machine is said to have passed the test. In order to test the machine's intelligence rather than its ability to render words into audio, the conversation is limited to a text-only channel such as a computer keyboard and screen."

It is important to note that the goal of creating the Sensemaking Agent, was to develop an Agent that could participate in ELICIT experiments to improve the state of the art and the state of the practice of experimentation in support of new approaches to command and control. Enhancing the state of the art of agent simulation in general was not one of the goals.

Instantiation Approach

At a high-level, the agent's configured personality and the factoids it receives determine what actions it takes, the order in which it takes these actions, and how it executes the actions. For example, we can describe the ELICIT Sensemaking Agent's behavior from the perspective of the Sensemaking Agent itself. The agent sees the world through the lens of its inbox. Factoid messages are periodically distributed to it. It also receives factoid messages shared by other participants (human and or agent) and factoid messages that it reads when it looks at information system "websites". When the agent checks its inbox, it may be focused on looking at a particular type of information, such as the target of an adversary attack. It selects a factoid and decides whether it should first think about distributing this new information (social processing) or whether it should first analyze the information (information processing). If the agent decides to analyze the information first, it parses the information and stores it in its memory. It then determines if this new piece of information is sufficient to change its mental model or awareness of the situation (awareness processing). If the new information is significant, it may make a cascade of iterative changes to its mental model. After analyzing and making sense of the new information, the agent then decides if the information should be further distributed. For example, the agent may decide to post the information to a relevant information system website, or the agent may realize that it knows enough to attempt to identify the who, what, where and when of a specific adversary attack.

The logic processing for the Sensemaking agent follows is designed to handle the specific tasks that a human agent would perform. This logic processing is comprised of 6 areas. The first thing that an agent does is check its inbox.

Message Selection – The inbox is the agent's main source of new information, so whenever there is new information in the inbox, the agent (like a human) wants to look at it. Like a human, each agent has its own style. Some agents use a last-in-first out (LIFO) approach and look at the most recently received message next. Others use a first-in-first out (FIFO) approach. Like a human, the agent can have limitations in its ability to handle its inbox. Agents can be configured to forget some of the information that is sent to them. Also like a human, agents can be overwhelmed if their inbox gets too full, and not look at all the messages.

Message Screening – Once an agent has selected a message, it needs to make some initial decisions about what to do with the message. The first thing an agent asks itself is "have I seen this message before?'' That is, is it a duplicate? Since each agent is tasked with working on one or more of the task areas (identifying the who, what, where and when of the adversary attack) the agent also asks itself if the factoid relates to its current area(s) of interest. If a factoid isn't a duplicate and it is relevant to an area of interest,

then the agent may analyze the message further. This part of the agent's logic also keeps track of what steps the agent has performed on a new message.

If there are no new messages that are unscreened, the agent may look for additional information on one of the available information websites.

Do I Share before I process? – When given a new interesting piece of information, some human's first reaction is to tell others. Other humans prefer to undertand the information and what it means before sharing it. Agents are like this too.

Information Processing – Like a human, the agent reads (parses) each factoid it receives to determine what data the factoid contains and whether the factoid provides any new information. The factoidsets which contain the messages that drive the ELICIT scenarios can include agent readable "factoid translation tables". These translation tables effectively allow the agent to parse the provide information in the factoids. The information processing step enables the agent to "know" what information it has received, by actually storing the data contained in each factoid "in its head." The agent builds and updates a set of state tables containing information that it knows about the situation. For example, consider the following illustrative factoid:

> **The Azure, Brown, Coral, Violet, or Chartreuse groups may be planning an attack.**

Information Processing of this factoid may result in the agent's adding five new group names to the table of potential groups that may be involved in planning the attack. Subsequently processed factoids may add to this list or eliminate some of these groups. As with a human, the interpretation of the data depends on the information that the agent already knows. If the agent is already aware of these five groups, it is possible that this factoid provides no new information to the agent. The tables used in Information Processing were modeled from those that humans drew on scratch paper during human-in-the-loop experiments.

The agent also keeps track of who has sent the agent information, and considers how important it thinks that information is. After an agent has learned enough new information, it assesses what it knows in a step called Awareness Processing.

Awareness Processing – Periodically after a human has learned a number of new things, it reassess what it knows about its situation. As with a human, the Sensemaking agent actually becomes aware of or "makes sense of" its environment. This sense making capability is a new kind of behavior for an agent. Like a human, the Sensemaking agent doesn't just act at random; rather, it models how it views its world and uses this model to make decisions and determine its actions. The agent asks itself if it knows enough to identity any of the who, what, where or when components of the potential adversary attack. As with a human, awareness processing is an iterative step. Sometimes combining information from several different factoids allows the agent to come to a new awareness of the situation. If an agent is able to eliminate all but one option in a particular situation

area, (i.e. the who, or the group planning the attack), it may move on to focus on learning more about another less well known aspect of the attack, such as the target (what).

There are also several variables to govern whether the agent thinks it knows enough to identify the adversary attack. These include whether the agent is a guesser and how much information the agent needs to see before it thinks it understands the situation.

Social Processing – Humans are inherently social, and so are agents. The agent decides what it should do with its new factoid in terms of passing it along to the rest of the group. It may decide to share a factoid with other member(s) of its group, or it may decide to distribute the factoid via a website or it may decide to do nothing.

For example, depending on an agent's configured sharing personality, an agent may share a new factoid with everyone the agent has access to, or the agent may only share the information with its supervisor, or the agent may hoard the information if it seems important. There are a number of configurable variables to control different aspects of its sharing personality. These include the Agent's propensity to share and its sharing modality (which governs the tradeoff between sharing information with an individual and posting information to a website so that many can see it). An agent can also be configured so that it always shares information with other specific participants.

In addition, the agent also has a simple model of its relationships with the other members of its group. It maintains a trust matrix, governed by past interactions with the other participants in the experiment, that can be used to determine which participants can be trusted in the future.

If after processing all the information available to it, the agent still isn't confident that it can identify the who, what, where and when of the adversary attack, the agent can determine that it needs to seek additional information. The tendency of the agent to seek new information, rather than to wait passively for information to come to it, is also modeled. Social processing initiates Pull attempts by generating one or more messages to websites, instructing them to send any new factoids to the agents. In the future, this logic may be extended to enable agents to send messages requesting information of certain types to other participants, as well.

[Note that the agent also uses the Social Processing routine to "send" any identification attempt made to the server. This process was instantiated in this way to facilitate a potential future capability for agents to send identification attempts to each other, and to allow agents to incorporate others' identification attempts into their awareness.]

Agent state data – The core of the agent's "brain" is an area where it stores state data. This includes

- Information about all the messages it has received
- Opinions it has formed about the who, what, where and when of the adversary attack.

- Opinions it has formed about the other members in its group, for making decisions on engaging in reciprocity, etc.

- What task(s) it is currently focused on, etc.

This information is stored in a series of state tables that are updated as the agent interacts with its environment, receives and processes information and uses this information to develop situational awareness.

Architecture

Since the goals of the Sensemaking agent were not merely to create an agent that could participate in an ELICIT experiment, but to create a sense making agent, it was necessary to model the various steps that a human takes when participating in an ELICIT experiment. Figure 1 shows the relationships between the various types of decision making processing that the agent performs.

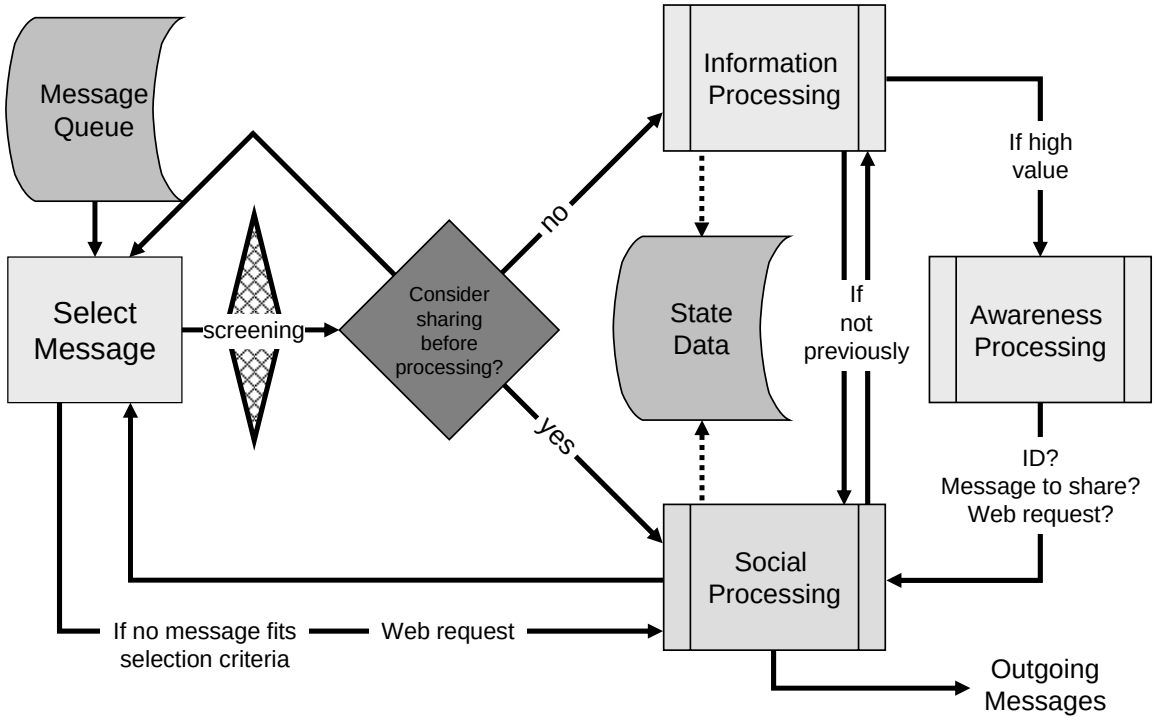

Figure 1. High-level View of Sensemaking Agent Logic Flow

Messages that have been made available to the agent are queued in a message queue, shown in the upper left corner of Figure 1.

The agent is written in Java and structured so we can continue to evolve the software to meet the needs of the research community. Appendix A contains a sample agent configuration including a brief description of the configuration variables that are used to determine an agent's personality. Together these variables model the agent's cognition and how it behaves.

Calibration of Agent Behavior

One of the benefits of the ongoing campaign of experimentation is that the work products can be leveraged in multiple ways. Previous ELICIT based experiments provide a rich set of data on which to model agent behavior. This data includes transaction logs that record and timestamp every action taken by a participant, scratch paper used by experiment participants, and information gleaned from participants in post experiment discussion sessions. Based on this information and experience, we had a strong basis for modeling the wide range of human behaviors expressed in the experiment data and an understanding of the mental models used by humans when participating in the experiments. In particular the software tables used to model an agent's "mental model" of the situation are an abstraction of tables that experiment participants drew on their scratch paper.

Configurable Variables that Govern Agent Behavior

To date, about four dozen configurable parameters have been designed into the Sensemaking Agent. These variables and a sample configuration are available in Appendix A.

The variables fall into several general categories:

- Deterministic variables about an agent's personality/motivation and task focus. For example, a variable that instructs an agent to only focus on identification of the "who" part of the scenario and ignore the what, where and when can be used to model a participant who has been assigned to a specific function and who interprets this assignment narrowly.

- Stochastic variables – variables that can be used to model error rates and behavior variability in agents. Agents configured to have the stochastic variables active can be used to conduct Monte Carlo experiment simulations.

- Time delay variables – variables that govern how long it takes for an agent to perform a behavior. These are generally used to slow down an agent's behavior so that it operates at a more human pace. There are about a dozen different time delay variables that model each step in the agent's decision making processes. For example, there is a variable to control the amount of time an agent takes to decide to share a piece of information and a variable to control the amount of time taken to perform the Share action.

It is anticipated that additional variables will be added in the future to model additional aspects of human behavior. Further details about use of the SenseMaking agent can be found in Appendix B

Next Steps in Agent Calibration and Evolution

The Sensemaking agent has been developed with an architecture that allows us to model a fairly wide range of behaviors we have seen in human-in-the-loop ELICIT experiments,

and leaves us with extension points where additional intelligence could be added. When given all the factoids necessary to correctly identify the who, what, where and when of an adversary attack, the Sensemaking agent can correctly make the identifications. When only provided some of the information, it accurately forms a situational awareness that is proportional to the information that it has processed.

As the Sensemaking agent evolves, validation of the agent is an ongoing effort. As part of developing the agents, comparisons have been made to the behaviors seen in the human-in-the-loop experiments. Agent calibration efforts include:

- Running a series of agent-only experiment trials that cover a full matrix of the key deterministic variables driving the agent's behavior. The resulting data are used to validate the range of behaviors the agent can exhibit and ensure that there are no unintended interactions.

- Developing archetype designs for the hierarchy and edge organizations. The results of these ideal, well behaved networks should be predictable and therefore this approach will test the logic and functionality of the agent's behaviors when comparing the data to hypothesized results.

- Configuring a set of agents that approximate the behavior of specific individual types seen in human-in-the-loop experiments to validate the flexibility of the agent.

Note that the intent of the Sensemaking agent is not only to provide an agent that can substitute for an unavailable human or a human with a specific personality in a live human test, but that also can be used in agent only investigations of general network centric concepts. Network centric concepts and the network centric value chain should apply to both human and non-human groups.

Note also that there is nothing in the Sensemaking agent that is specific to humans. So the Sensemaking agent could also be used to investigate the broad effects of the structures of networked systems on information sharing, collaboration and trust in machine and machine-learning contexts.

To date, the initial Sensemaking Agent development efforts have all been performed on a single ELICIT scenario factoidset. (Factoidset 1-17). Other anticipated efforts include creating Agent readable factoidset translation tables for the other three baseline ELICIT factoid sets.

New Avenues for Experimentation Supported By Enhancement

The new Sensemaking Agent capability greatly expands the number of experiments that can be run so that more variables can be explored and a sufficient number of runs be conducted for each experimental condition so that results will be statistically significant.

In addition, the agent capability opens new avenues for exploring theories of behavior and performance. Results gleaned by human participant runs, could be validated by configuring agents with specific traits and confirming that resulting agent experiment outcomes matched the human results.

Agents could also be used to explore the effect of certain "personality traits" on group performance. For example, how many persons with the "information hoarding" trait does it take to negate a group's performance? Such experiments could help to quantify and validate rich targets for team training.

The agent capability complements and supplements ELICIT experiments done with human participants. It is usually difficult to schedule large numbers of human subjects for experiments and have them all show up when requested. The agent can be used to mitigate these challenges. The agent capability is not intended to replace experiments with human subjects, and promising results from experiments with agents will still need to be validated with humans.

Some of the "softer" aspects of ELICIT experimentation, particularly sharing of free-form text information via post-cards or an external chat system are still outside the scope of the current Sensemaking Agent. Therefore these avenues of research will remain the purview of human subjects for some time. The lower priority on creating agents that can handle sharing of free-form ELICIT data is due to the difficulty of the task. It is not for lack of interest; as the effects of sharing freeform information (that could represent requests for specific information, sharing of theories about the task situation, or suggestions for modes of self organization) are a very fertile area for future ELICIT research. The situational awareness that the new Sensemaking Agent formulates based on awareness and understanding of the factoids to which they have been exposed could be used as a basis for extending the agent's capabilities in this area. As more subtle aspects of sharing, trust and awareness are studied, there will be an on going tension between the models instantiated in the agent implementations and the designs of the experiments that use them.

Anyone who is a member of the ELICIT COI will have access to versions of the Sensemaking Agent as they are released. As with all other ELICIT resources and materials, one of the requirements of the Sensemaking Agent's use is that all raw data and results are made available to the community. Additional information about the ELICIT COI can be found at http://www.dodccrp.org/html4/elicit.html, or contact the authors directly.

References

Alberts, D.S., Garstka, J.J., Hayes, R.E., and Signori, D.T., *Understanding Information Age Warfare* Washington, DC: CCRP (2001).

Alberts, D.S. and Hayes, R.E., *Campaigns of Experimentation* Washington, DC: CCRP (2005).

Alberts, D.S. and Hayes, R.E., *Power to the Edge* Washington, DC: CCRP (2003).

Alberts, D.S. and Hayes, R.E. *Understanding Command and Control,* Washington, DC: CCRP (2006).

CCRP, Command & Control Research Program ELICIT Workshop, Vienna, VA (2007).

Leweling, T.A. and Nissen, M.E., "Hypothesis Testing of Edge Organizations: Laboratory Experimentation using the ELICIT Multiplayer Intelligence Game," *Proceedings* International Command & Control Research & Technology Symposium, Newport, RI (June 2007).

Mancik, Eric and McCaskill, Josh, *ELICIT: Tracking Networks,* Assembly, May/June 2007 West Point, pp. 30-31.

Nissen, M.E., "Exploring Edge C2 Approaches: ELICIT and POW-ER Tools," Plenary Presentation, International Command and Control Research and Technology Symposium, Newport, RI (June 2007).

Nissen, M.E., "Hypothesis Testing of Edge Organizations: Specifying Computational C2 Models for Experimentation," *Proceedings International Command & Control Research Symposium.* McLean, VA (2005).

Nissen, M.E. and Buettner, R.R., "Computational Experimentation with the Virtual Design Team: Bridging the Chasm between Laboratory and Field Research in C2," *Proceedings* Command and Control Research and Technology Symposium, San Diego, CA (June 2004).

Nissen, M.E. and Sengupta, K., "Incorporating Software Agents into Supply Chains: Experimental Investigation with a Procurement Task," *MIS Quarterly* 30:1 (2006), pp. 145-166.

Ruddy, M.E., *ELICIT 2.0 Web-based Software Guide*, Parity Communications Inc. under contract to the Naval Post Graduate School Center for Edge Power, Arlington, MA, (2007)

Ruddy, M.E., *ELICIT 2.2.1 Web-based Software Guide For Human Experiment Directors* Parity Communications Inc. under contract to Evidence Based Research, Arlington, MA, (2008)

Ruddy, M.E., *Phase 2 - Final Detailed Design for Agent-based Processing v1.1*, Parity Communications Inc. under contract to Evidence Based Research, Chestnut Hill, MA, (2007)

Ruddy, M.E., and Nissen, M.E., "New Software Platform Capabilities and Experimentation Campaign for ELICIT " Parity Communications, Inc. Under contract to the Naval Post Graduate School Center for Edge Power, Arlington, MA, (2008)

Turning, A., “Computing Machinery and Intelligence”, *Mind*, (1950)

Appendix A - ELICIT Sensemaking Agent Variables and Sample Configuration

The purpose of this appendix is to provide an example configuration of the Sensemaking agent so the reader can get a general understanding its capabilities. It is not a complete documentation of all available options for each variable.

Below is a sample Sensemaking Agent configuration file (SenseMaking_Agent_1h10_10.txt) with annotations in **[bold].** Note that time intervals are in milliseconds.

SenseMaking_Agent_1
<begin agent configuration parameters>
SenseMaking_Agent_1.jar **[Name of the java jar file used by the Sensemaking agent]**
net.parityinc.ccrp.web.agent.impl.SenseMaking_Agent_1 **[Name of the java code file used by the Sensemaking agent]**

readyIntervalDelay|Time interval to click Ready button|10000 **[The interval of time between receipt of the instructions URL and when the agent clicks Ready in milliseconds. Default is 10000 (10 seconds).]**

messageQueueCapacity|Capacity of queue (-1 means unlimited)|-1 **[If message queue capacity is less than the size of the queue, the message is not added to the queue. For example: if it is set to 100 and there are 100 items in the queue already, the message will not be added. A value of -1 means the queue is unlimited.]**

messageQueueTimeRemainInQueue|Time a factoid can remain in queue (-1 means unlimited)|-1
[If the value is set to 10 minutes, and a message has been in the queue for 11 minutes, it will be skipped (not selected.) Value is in milliseconds. A value of -1 means the message queue time is unlimited.]

messageQueueNewerBeforeOlder|If true then newer messages are selected before older|false **[If false, then older messages are selected first. If true, than most recent message is selected first.]**

selectMessageFromQueueDelay|Select message from queue delay|1000
[Amount of time in milliseconds the agent waits before selecting the next message from the message queue.]

shareBeforeProcessing|If true then share message before Processing|true
[If true, perform Share Social Processing before Information Processing the message. If False, then do Information Processing before Share Processing. Default is False.]

postedTypes|PostedTypes|who,what,where,when

[Possible websites to Post to. This is a required condition. If a website is not on this list, the agent can never post to it. (Note that there are additional variables that control posting. (e.g. IsCompetitiveHoarder and PostFactor, as well as organization configuration. In order to be able to Post to the Who website it must be allow by this configuration variable and the relevant organization configuration file.)]

postFactor|PostFactor|1
[This is a stochastic variable that should have a value between 0 and 1. If 1, then if message type is in posted types and all other conditions are true, Post if Post factor is greater than or equal to a random number. Default is 1, which is to always Post.]

postOutOfArea|PostOutOfArea|false **[If message type (from factoid key) is for a type of website that is not available to the agent, and PostOutofArea is True, then Post to first available website. Else, don't Post.]**

shareWithFactor|ShareWithFactor|1 **[This is a stochastic variable that should have a value between 0 and 1. If 1, then if all other Share conditions are true, Share if Share factor is greater or equal to a random value. Default is 1, which is to always Share.]**

sharedTypes|SharedTypes|who,what,where,when
[If message is of a type (as defined in the factoid key) that is listed here, then proceed and check the other conditions for Sharing. Else do not Share.]

shareRelevantAccordingToSiteAccess|ShareRelevantAccordingToSiteAccess|false **[If condition is true and if message type is available for this site and factoid is not of type Expertise or Key (as defined in the factoid key), then do not Share message except with "must share with" players and websites specified by the shareWith and shareWithWebSite variables.]**

shareAccordingToSiteAccess|ShareAccordingToSiteAccess|false
[**If condition is true and if message type is available for this site, then do not Share message (presumably because the agent will have chosen to Post this message rather than Sharing it.)]**

isCompetitiveHoarder|IsCompetitiveHoarder|false
[If set to true, than agent is a competitive hoarder and does not Share Expertise or Key factoids (as defined in the factoid key.)]

pullFactor|PullFactor|1
[**This is a stochastic variable that should have a value between 0 and 1. Pull if Pull factor is greater than or equal to a random number. Default is 1, which is to always Pull, when otherwise considering a Pull.]**

timeBeforeFirstIdentify|Time before the agent does its first identity (in minutes)|8
[Minimum number of minutes that need to pass before the agent does its first Identify. Default is 10 minutes]

minSolutionAreas|The minimum number of ID tables with some data|1
[The ELICIT scenarios have 4 information areas (who, what, where and when.) If minSolutionAreas is set to 3, then agent must know something about 3 of these areas before it performs an Identify action.]

hasSeenEnoughToIdentify|HasSeenEnoughToIdentify|1
[Only identify if number of messages selected from queue is greater or equal to this number.]

isGuesser|IsGuesser|false **[If agent has several potential answers that it is aware of for an area and if value is true, then agent selects one of these potential answers (it guesses from among the possible answers.) else it doesn't select an answer.]**

isFrequentGuesser|IsFrequentGuesser|false **[If this variable is set to true, the agent ignores timeBeforeFirstIdentify, minSolutionsAreas and hasSeenEnoughToIdentify variables when deciding if it should identify.]**

idConfidencelevel|IdConfidencelevel|1
[Values range from 0-1. Agent only provides an answer in an Identify Action if the value in the relevant state table is greater or equal to this value.]

partialIdentify|Identify if don't have answers for all areas|true
[If set to true, then Agent may Identify if is only has information on less then four of the identify areas (who, what, where, when.)]

propensityToShare|PropensityToShare possible values (low, moderate, high, very high)|high
[Controls agent's willingness to Share information. For example, if propensity to share is very high, agent Posts a message to all websites for which agent has permissions, sends a message to all entities where trust in recipient is not distrust, and sends a message to all entities with must_share flag set to on. If propensity to share is low, then agent sends a message only to entities with must_share flag set to on. Default is low, and its effects depend on the value of the shareModalChoice variable.]

shareModalChoice|ShareModalChoice possible values (both, post dominant, post only, peer to peer dominant, peer to peer only)|both
[This variable governs the interaction and tradeoffs between Sharing a message with an individual and making it available by Posting it to a website. Its effects depend on the value of the propensityToShare variable. For example if shareModelChoice is set to both, and propensity to share is high, then Post message to area (of message website) and send message to all entities where trust in recipient

is not distrust or no opinion, and send message to entities (participants and websites) with must-share flag set to on. Default is both.]

screeningSelectedMessageDelay|Screening selected message (message processing) delay|1000 **[Time interval in milliseconds that the agent delays after selecting a message from the message queue and before proceeding with social/information processing. Default is 0.]**

informationProcessingDelay|Information Processing delay|3000
[Time interval in milliseconds that the agent delays while processing and remembering the information in a message. Default is 0.]

socialProcessingDelay|Social Processing delay|4000
[Time interval in milliseconds that the agent delays while thinking about whether and how to communicate the message to others. Default is 0.]

sharingPostingMessageDelay|Sharing/Posting each Message delay|8000
[**Time interval in milliseconds that the agent delays while actually Sharing or Posting a message. Default is 0.]**

awarenessProcessingDelay|Awareness Processing delay|3000
[Time interval in milliseconds that the agent delays while thinking about the implications of the information that it has received. During this time, an agent may become aware that it has identified the who, what, where or when of an adversary attack. Default is 0.]

determiningKnowledgeNeedsDelay|Determining Knowledge Needs delay|3000
[Time interval in milliseconds that the agent delays while thinking about whether it needs additional information to identify the who, what, where or when of an adversary attack. Default is 0.]

idAttemptDelay|ID Attempt delay|20000
[Time interval in milliseconds that the agent delays when performing an Identify action. Default is 0.]

webRequestDelay|Web Request (Pull)|9000
[**Time interval in milliseconds that the agent delays when deciding to check a website for information (Pull). Default is 0.]**

shareWith|List of players with whom agent may share (-1 means share with all from organization configuration file)|3
[List of players, specified by player role number, with which the agent must Share information for certain shareModalChoice propensityToShare combinations. A value of -1 means Share with all players with whom you have Share capability (as specified in the organization configuration file.) Default is -1.]

shareWithWebSites|List of websites with whom agent must share|when
[List of websites to which the agent must Post information for certain shareModalChoice propensityToShare combinations. Note that website must also be available to the agent as specified in the organization configuration file.]

propensityToSeek|PropensityToSeek possible values (low, moderate, high, very high)|moderate
[Propensity to seek information via Pull action. Default is high. If propensityToSeek is very high, then agent pulls from all available sites, else it only pulls from the agent's primary area. PropensityToSeek also controls the minimum time between Pulls. For example if propensityToSeek is moderate, the minimum time between pulls is 3 minutes.]

minTimeBetweenPulls|If the time since the last pull is not >= minTimeBetweenPulls, do not Pull (in milliseconds), -1 means ignoring this parameter)|20000 [**Minimum time interval in milliseconds that the agent delays between subsequent Pulls (requests of data from a website.) Default depends on value of propensityToSeek. For example if propensityToSeek is moderate, the default is 3 minutes. If you would like the value of propensityToSeek to control the spacing between pulls, then do no include this variable in the configuration.]**

minTimeBetweenShares|If the time since the last Share is not >= minTimeBetweenShares, the agent should wait before it Shares (in milliseconds, -1 means ignoring this parameter)|5000 **[Minimum time interval in milliseconds between subsequent message Share actions.]**

trustInIndividuals|TrustInIndividuals possible values (high, medium, distrust, no opinion)|1=no opinion,2=no opinion,3=no opinion,4=no opinion,5=no opinion,6=no opinion,7=no opinon,8=no opinion,9=no opinion,10=no opinon,11=no opinion,12=no opinion,13=no opinon,14=no opinion,15=no opinion,16=no opinion,17=no opinion **[Initialization of the trust value for each of the other participants in the agent's group. If value is blank or invalid, default to no opinion.]**

trustInWebSites|List of initial values of Trust for web sites. Possible values (high, medium, distrust, no opinion)|who=medium,where=medium,what=medium,when=medium
[Initialization of the trust value representing the agent's trust in each of the websites.]

reciprocity|Reciprocity possible values (high, low, medium, na, none)|1=none,2=none,3=none,4=none,5=none,6=none,7=none,8=none, 9=none, 10=none,11=none,12=none,13=none,14=none,15=none,16=none
[Initialization of the reciprocity value representing the sense of reciprocity towards the other participants in the agent's group.]

primary|Primary areas of interest. Possible values: who, what, where, when)|what
[Area(s) of interest on which the agent initially focuses.]

secondary|Secondary areas of interest. Possible values: who, what, where, when)|who,where,when
[Other area(s) of interest on which the agent may subsequently focus.]

propensityToShareExternal|If message is not in area of interest, then agent shares it according to sharing preferences with probability = propensityToShareExternal|1
[Stochastic variable. If message area (as defined in the message's factoid key) is not one of the agent's current areas of focus then only Share according to other Sharing preferences if value of random variable is >= this value.]

awarenessProcessingThreshold|If cumulative value of the perceived message value is more or equal to this variable, then start awareness processing.|4
[Cumulative value of additional newly processed information that can accumulate before a new cycle of awareness processing is initiated.]

pullBetweenSitesDelay|Pull between sites delay|1000
[When Pulling (requesting) data from multiple websites, time interval in milliseconds that the agent delays between subsequent Pulls.]

postBetweenSitesDelay|Post between sites delay|500
[When Posting to multiple websites, time interval in milliseconds that the agent delays between subsequent Posts.]

provideRelevance|Provide relevance for posted and shared messages|false
[If the "allow participant ratings" option is turned on in the Configuration page, and value is set to true, then agent provides a relevance rating value.]

provideTrust|Provide trust for posted and shared messages|false
[If the "allow participant ratings" option is turned on in the Configuration page, and value is set to true, then agent provides a trust rating value.]

inactivityPeriod|Inactivity period after which agent should raise an inactivity flag (in milliseconds, if the agent has done nothing but pulls for >= Inactivity period)|0

We are continuing to add to and refine the Sensemaking Agent's configurable variables as we evolve the Sensemaking Agent. For example, the following variable is in development:

futilityThreshold|Time working in an area during which no new messages in that area are processed before moving on to another area|-1

Appendix B - ELICIT Sensemaking Agent Details

For those who are already familiar with ELICIT and in particular the web-based 2.2 version of ELICIT, this section provides additional background on how the Sensemaking agent can be used in conjunction with ELICIT. These appendices are intended to provide an overview of the new Sensemaking Agent capability, and are not meat to substitute for complete instructions on how to configure and use the agent-based version of ELICIT. These instructions will be made available to members of the ELICIT COI with the released version of the agent-based version of ELICIT (abELICIT).

The ELICIT web-based software was enhanced so that it could recognize and run agent-based processes either as part of agent-only experiment runs or as part of hybrid human and agent-based participant experiments. The software has the ability to run multiple specific software agents each of which is configurable. Specific configurations of the Sensemaking agent can be created by downloading an existing Sensemaking Agent configuration file, renaming it, modifying the settings and uploading it using the Configuration Files page.

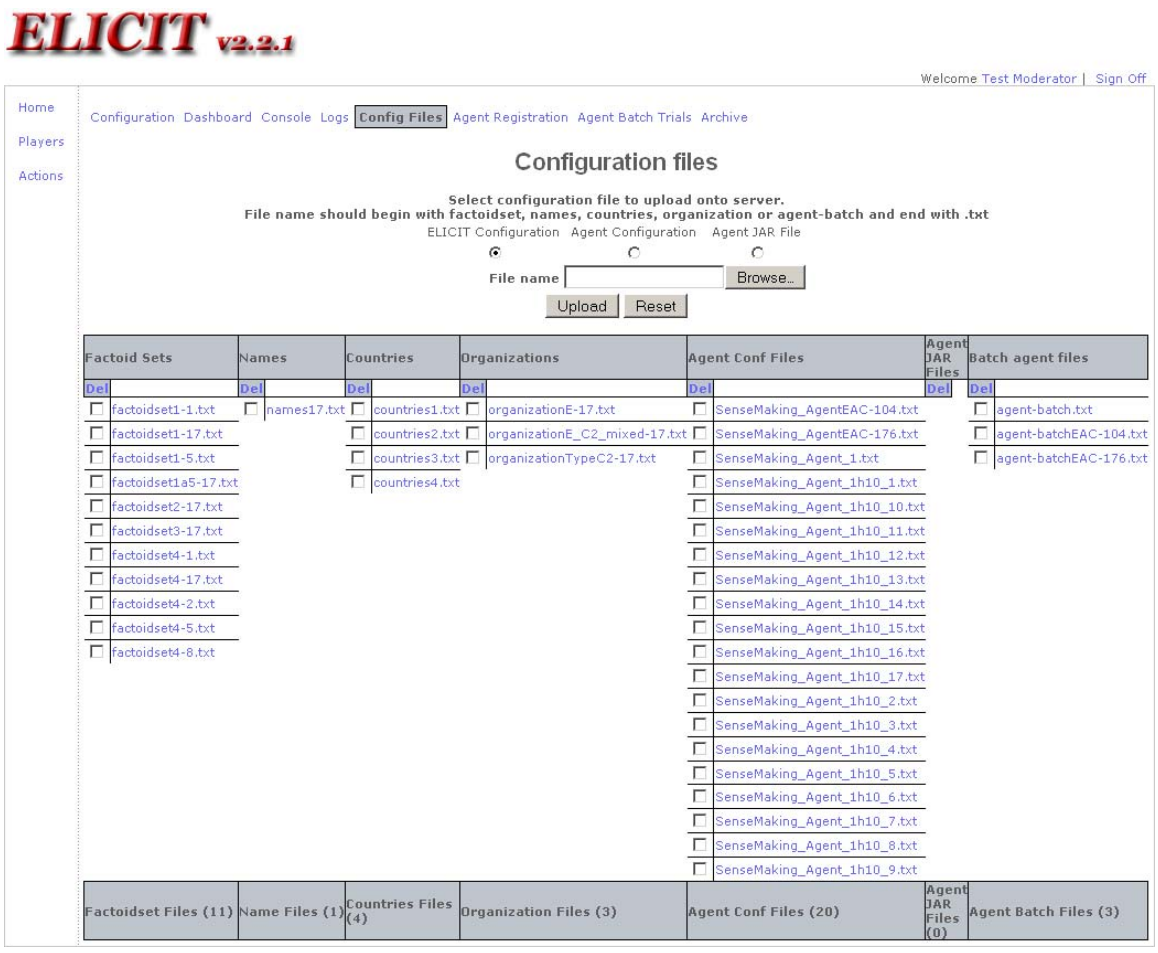

Figure 2. abELICIT Configuration Screen

After the Sensemaking agent .txt configuration files have been uploaded, it is possible to add/register a new agent in the ELICIT server application so that it can participate in an online experiment. It is also possible to run all agent experiment trials in batch mode. To register an agent for an online experiment:

1) Click on the "Agent Registration" option on the Server Moderator Actions page:

2) Select the desired agent configuration file on the "Agent Registration" page:

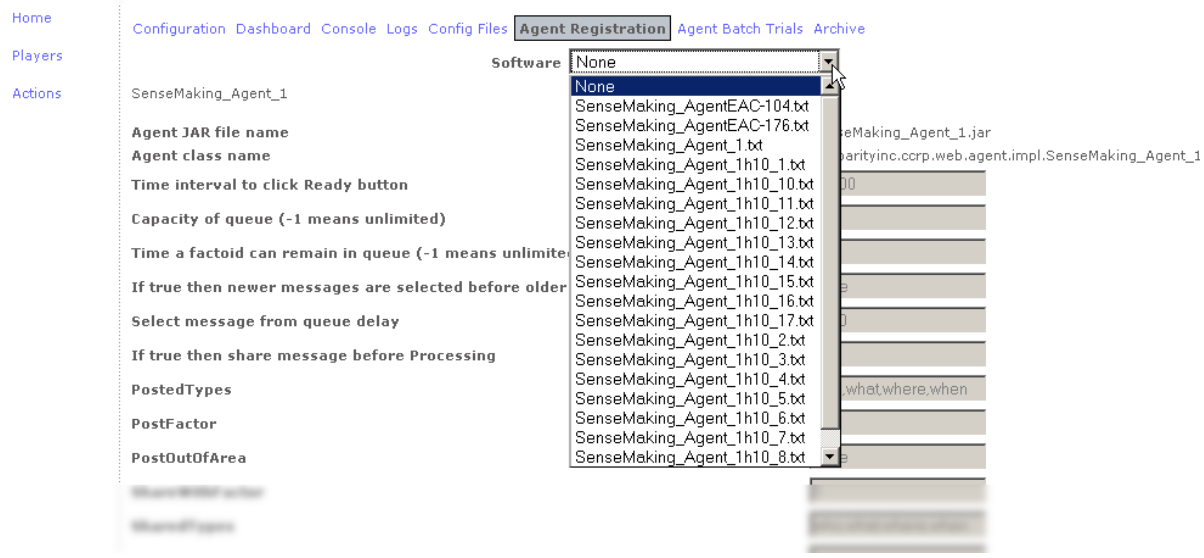

Figure 3. ELICIT Agent Registration Page

The dropdown is populated by the values of the names of agent configuration files that were uploaded to the server via the Configuration File page. A single agent .txt file can be registered multiple times (to represent multiple similar participants) in an experiment trial.

Once an agent has been registered to the server, the agent will appear on the Server dashboard screen. Note that the content of the dashboard fields varies slightly between human and agent participants.

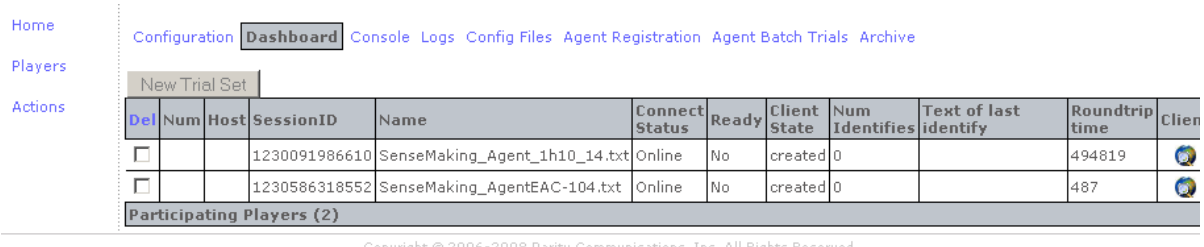

Figure 4. ELICIT Dashboard Page with two agents registered

The Sensemaking Agent can perform all the actions that a human ELICIT participant can perform. In particular, agents can

- Click **Ready** button after configured interval after a trial set is started. Action is triggered by receipt of a New Trial Set instruction URL.
- **Post** a factoid to a website after configurable delay factor. Action is triggered by receipt of a new factoid in the inbox and configured decision processing.
- **Pull** data from a website. Action is triggered as a result of the decision making processing and configurable delay factors.
- **Share a** factoid. Action is triggered as a result of decision making processing and configurable delay factors.
- **Identify** action. Action is triggered as a result of decision making processing and configurable delay factors. Content of Identify is based on the mental model the agent has constructed from the factoids it has seen.

Once received by the ELICIT Server via the Agent API, the server treats these actions as if they were generated by a human participant.

Agents can also be run in batch mode. Multiple trial sets may be defined in an Agent batch. As with the agent configuration, new Agent-batch files can be created by downloading an existing agent-batch*.txt configuration file, renaming the file, modifying the settings and uploading it using the Configuration Files page.

To initiate an Agent batch run,

1) Click on the “Agent Batch Trials” option on the Server Moderator Actions page:

2) Select the desired agent-batch configuration file on the “Agent Batch Trials” page:

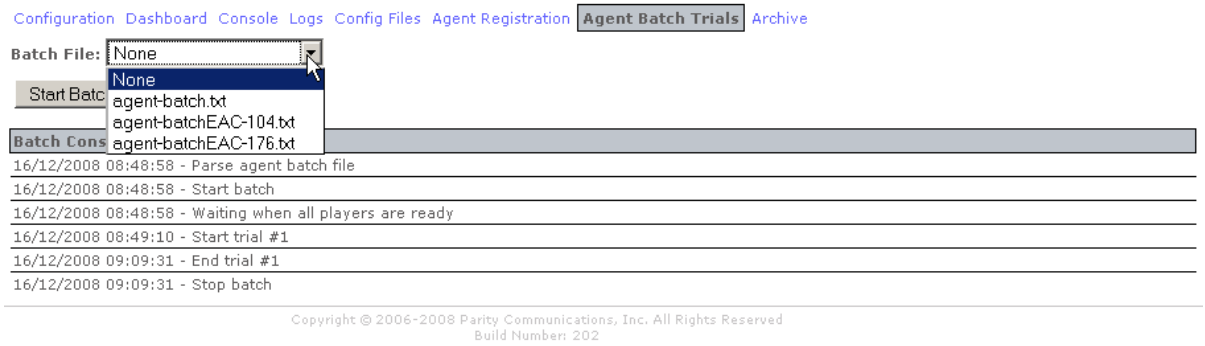

Figure 5. Initiating an ELICIT Batch Agent Trial