

The *Hats* Information Fusion Challenge Problem

Clayton T. Morrison and Paul R. Cohen
Information Sciences Institute
University of Southern California
4676 Admiralty Way, Suite 1001
Marina del Rey, CA, 90292 U.S.A.
{clayton,cohen}@isi.edu

Abstract - *We describe the Hats Simulator as an information fusion challenge problem. Hats is a virtual world in which many agents engage in individual and collective activities. Most agents are benign, some intend harm. Agent activities are planned by a generative planner. Playing against the simulator, the goal of the analyst is to identify and arrest harmful agents before they carry out their plans. The simulator provides both scalar and categorical information. Information fusion tasks in the Hats domain include assessing information value, choosing information collection strategies, tracking individuals and resources, identifying events, hypothesizing group membership, ascribing suspicion, and identifying plans. After each game, the analyst is assessed a set of scores including the cost of acquiring information, the cost of falsely accusing benign agents, and the cost of failing to detect harmful agents. The simulator is implemented and currently manages hundreds of thousands of agents.*

Keywords: Simulation, information cost and quality, tracking and surveillance, sensor management, sensor modeling

1 Introduction

The Hats Simulator [1, 2] was designed originally to meet the needs of academic researchers who want to contribute technology to Homeland Security efforts but lack access to domain expertise and classified problems. Most academic researchers do not have security clearances and cannot work on real data, yet they want to develop tools to help analysts. In any case, real data sets are expensive: They cost a lot to develop from scratch or by “sanitizing” classified data. They are domain-specific, requiring classified expertise to understand. Because data sets are expensive, many that have been made available to researchers are relatively small and the patterns to be detected within them are fixed, few, and known, so working with these data sets is a bit like solving a single “Where’s Waldo” puzzle. Sometimes there also is the problem that real data sets model “signal” (terrorist activities) not “noise” (everything else), while extracting signal from noise is a great challenge. Data sets in general are static, whereas data

become available to analysts over time. It would be helpful to have a data feed, something that generates data as events happen.

To validate information fusion tools, we need a generator of a variety of multi-signal information sources, all of which represent meaningful, goal-directed behavior, but only a small portion of which represents the target activities we care about. The generator should be parameterized for experimental purposes (e.g., varying the distinctiveness of target activities, to make them more or less easily recognizable); and it should come up with novel activities, requiring analysts and their tools to both recognize known patterns and reason about unanticipated suspicious patterns.

The Hats Simulator is home to hundreds of thousands of agents (hats) which travel to meetings. Some hats are covert terrorists and a very few hats are known terrorists. All hats are governed by plans generated by a planner. Terrorist plans end in the destruction of landmarks. The object of a game against the Hats simulator is to find terrorist task forces before they carry out their plans. One pays for information about hats, and also for false arrests and destroyed landmarks. At the end of a game, one is given a score, which is the sum of these costs. The goal is to play Hats rationally, that is, to catch terrorist groups with the least combined cost of information, false arrests, and destroyed landmarks. Thus Hats serves as a testbed not only for analysts’ tools but also for new theories of rational information fusion that take into account information source assessment and cost in the context of analysis goals and decisions. Hats encourages players to ask only for the information they need, and to not accuse hats or issue alerts without justification.

The Hats simulator is lightweight: Agents have few attributes and engage in few elementary behaviors; however, the number of agents is enormous, and plans can involve simultaneously many agents and a great many instances of behaviors. The emphasis in Hats is not domain knowledge but managing enormous numbers of hypotheses based on scant, often inaccurate information. By simplifying agents and their elementary behaviors, we de-emphasize the domain knowledge required to identify terrorist threats and emphasize covertness, complex group behaviors over time, and the frighteningly low signal to noise ratio.

The Hats Simulator consists of the core simulator




Figure 1: Information Broker Interface to the Hats Simulator

and an information broker. The information broker is responsible for handling requests for information about the state of the simulator and thus forms the interface between the simulator and the analyst and her tools (see Figure 1). Some information has a cost, and the quality of information returned is a function of the “algorithmic dollars” spent. Analysts may also take actions: they may raise beacon alerts in an attempt to anticipate a beacon attack, and they may arrest agents believed to be planning an attack. Together, information requests and actions form the basis of scoring analyst performance in identifying terrorist threats. Scoring is assessed automatically and serves as the basis for analytic comparison between different analysts and tools. The simulator is implemented and manages the activities of hundreds of thousands of agents.

The following sections outline the Hats domain, including how we generate populations of hats and how the planner schedules meetings for hats to attend. We describe the information request framework, the actions the analyst may take, and scoring. We then analyze the space of relevant hypotheses and their relations in fusion tasks. We conclude with a discussion of the future of the Hats Simulator.

2 The Hats Domain

The Hats Simulator models a “society in a box” consisting of many very simple agents, hereafter referred to as hats. (Hats get its name from the classic spaghetti western, in which heroes and villains are identifiable by the colors of their hats.) The Hats society also has its heroes and villains, but the challenge is to identify which color hat they should be wearing, based on how they behave. Some hats are known terrorists; others are covert and must be identified and distinguished from the benign hats in the society.

Hats is staged in a two-dimensional grid on which hats move around, go to meetings and trade capabilities. The grid consists of two kinds of locations: those that have no value, and high-valued locations called beacons that terrorists would like to attack. All beacons have a set of attributes, or vulnerabilities, corresponding to the capabilities which hats carry. To destroy a beacon, a task force of terrorist hats must possess capabilities that match the beacon’s vulnerabilities, as a key matches a lock. In general, these capabilities are not unique to terrorists, so one cannot identify terrorist hats only on the basis of the capabilities they carry.

The Hats society is structured by organizations. All

hats belong to at least two organizations and some hats belong to many. Terrorist organizations host only known and covert terrorist hats. Benign organizations, on the other hand, may contain any kind of hat, including known and covert terrorists.

2.1 Population Generation

Hats populations may be built by hand or generated by the Hats Simulator. Because the constitution of a population affects the difficulty of identifying covert terrorists, population generation is parameterized. There are four sets of population parameters. The first set specifies the total number of known terrorists, covert terrorists and benign hats in the population. Another set defines the number of benign and terrorist organizations. Not all organizations have the same number of members, so a third set of parameters assigns the relative numbers of hats that are members of each organization, represented as a ratio among organizations. For example, the ratio 2:1:1 means that the first organization has twice as many members as the other two. Finally, hats may be members of two or more organizations. An overlap parameter determines the percentage of hats in each organization that are members of two or more other organizations. Since hat behaviors are governed by their organization membership, as we will see in the next section, organization overlap affects how difficult it is to identify covert terrorist hats. To generate populations with hundreds of thousands of hats and thousands of organizations, we use a randomization algorithm that estimates organization overlap percentage and membership ratios while matching the total number of organizations and hats in the population. When the population is generated, each hat is assigned a native capability that they will carry throughout the duration of the simulation, and a set of traded capabilities that are temporary, expiring after some number of ticks (e.g., within 40 ticks). Hats are also assigned random locations in the Hats grid world. An efficient, approximate population generation algorithm is described in [3].

2.2 Meeting Generation

Hats act individually and collectively, but always planfully. In fact, the actions of hats are planned by a generative planner. Benign hats congregate at locations including beacons. Terrorist hats meet, acquire capabilities, form task forces, and attack beacons. The purpose of the planner is to construct an elaborate “shell game” in which capabilities are passed among hats in a potentially long sequence of meetings, culminating in a final meeting at a target. By moving capabilities among hats, the planner masks its intentions. Rather than directing half a dozen hats with capabilities required for a task to march purposefully up to a beacon, instead hats with required capabilities pass them on to other hats, and eventually a capable task force appears at the beacon.

Each organization has a generative planner that plans tasks for its members. Hats that are currently

participating in a task are reserved; hats not currently part of a task are free. At each tick, each organization has a chance of beginning a new task. When a new task is started, the Hats meeting planner creates a task force, a subset of hats selected from the free hats of the organization. The size of a task force is controlled by a parameter. The planner next selects a target location in the Hats world. With some probability, that location may be a beacon, otherwise a random location is selected. If a beacon is selected as the target, the goal of the task is to bring to that location the set of required capabilities that match the vulnerabilities of the beacon. If the location is not a beacon, a random set of required capabilities is selected as the set to bring to the location.




Figure 2: Example of a generated meeting tree. Each box represents a meeting and contains a list of participating hats. Arrows indicate planned temporal order of meetings.

Task force members may or may not already possess the required capabilities; usually they don't. The planner creates a set of meetings designed to ensure that the task force acquires all of the required capabilities before going to the target location. This is accomplished by constructing a meeting tree that specifies meetings and their temporal order. Figure 2 shows an example meeting tree, where boxes represent planned meetings among hats and arrows represent the planned temporal partial order of meetings. The tree is “inverted” in the sense that the arrows point from leaves inward toward the root of the tree. Parent meetings, where arrows originate, are executed first. When all of the parent meetings of a child meeting have completed, then the child meeting happens. This ensures that none of the hats that participate in the child meeting are busy in other meetings. Meeting execution means that the hats participating in the meeting begin moving toward the meeting location. The final, root meeting takes place at the task target location and includes all of the task force hats. The locations of the other meetings in the tree are selected randomly.

Initially, the meeting tree is skeletal, containing meetings whose only participants are the task force members themselves. From the organization's remaining free hats, the planner selects a second group of resource hats that carry required capabilities not currently carried by the task force. Resource hats are randomly assigned to existing meetings, and trades of required capabilities from resource to task force hats

are scheduled to take place during the meeting. The planner finishes tree construction by adding “decoy” meetings, trades and other free hats not directly involved in moving required capabilities to the goal. A constraint maintained throughout tree construction is that at least one hat from a parent meeting will go on to meet in a child meeting. These hats will either be task force members, resource hats carrying required capabilities to trade in the next meeting, or will be decoy hats arriving from decoy meetings.

Completed meeting trees are added to a queue of pending tasks. At each tick, the simulator engine searches the task queue for meetings with no currently executing parent meetings. These meetings are then assigned to a queue of currently executing meetings and the participant hats are incrementally moved toward the meeting location. When all of the participants have arrived at the meeting location, the meeting itself lasts for two ticks, after which all hats not participating in more meetings are set “free” and become available to participate in new meetings.

Meeting trees typically have a depth of 2 to 7. The frequency of new tasks depends on both the probability of starting a new task as well as the number of hats in each organization.

3 The Information Broker

As an analyst playing the game, your job is to protect the Hats society from terrorist attacks. You need to identify terrorist task forces before they attack beacons, but you also need to avoid falsely accusing innocent hats. The only way to do this successfully is to gather information about hats, identify meetings, track capability trades and form hypotheses about the intentions of groups of hats. The information broker provides information about the state of the Hats world. The information broker will respond to questions such as “Where is *Hat*₂₇ right now?” It will also provide information by subscription to analysts' tools, which in turn make information broker requests. For example, a tool might process requests like, “Identify everyone *Hat*₂₇ meets in the next 100 steps,” or, “Tell me if *Hat*₂₇ approaches a beacon with capabilities C_1 , C_7 or C_{29} .”

Some information is free, but information about states of the simulator that change over time is costly. The quality of the information obtained is determined by the amount paid. The following two sections describe the two central components of the request framework: the cost of information and noise. Together, these components make the Hats simulator an experimental environment for studying the economics of information value in the context of intelligence analysis.

3.1 The Cost of Information

Some information from the broker is free. This includes information about the population (who the known terrorists are), the simulator world (world-map dimensions), and some event bookkeeping (locations of at-

tacks, a list of currently arrested hats). Other types of information require payment and the amount paid sets a base probability that is used to determine the accuracy of the information. In the current implementation, increasing accuracy requires exponentially more “algorithmic dollars”; Equation 1 maps payment to probability. The same function is applied to every payment-based request.

$$probability = 1 - \left[\frac{1}{\log(payment)} \right] \quad (1)$$

3.2 Making Requested Information Noisy

Modeling noise is a topic suitable for an entire research program. Among the issues to consider are: Can the analyst request the same information multiple times, and if so, how does information quality change? Can the analyst get accurate information about events that occurred in the past? Can the analyst tell which information sources are reliable by asking the same question several times? And so on. We have started simply, imposing the constraint that the analyst may request a particular piece of information only once. This means that the analyst must select the level of payment for the information at the time of the request and there is no going back once the request is made. Information that updates from one tick to the next, such as the current location of a hat, may be requested again at the next tick. However, information that is fixed in time, such as when or where a meeting took place, can be requested only once. This model allows us to avoid, for now, the thorny issue of how to “noise up” multiple requests for the same information.

Using the payment function described in the previous section, a payment amount is mapped to a “base” probability p . With probability p , the information requested is returned in its entirety or subject to noise.

Eight basic information requests can be made: the hats at the location (if any), participants in a meeting, capabilities carried by a hat, capability trades, meeting times, hat death time, meeting locations and hat locations. The first five requests return lists of things, such as hats, capabilities, times, etc. The latter three are scalar values. How we add noise to responses to these requests depends on the type of thing requested as well as whether they involve single elements or lists of elements.

For scalar values, noise is modeled by sampling from a normal distribution, where the mean is the location or time of the requested item, and the variance is a function of the size of the Hats world (for locations) or the amount of time since beginning the simulation (for times). Adding noise to reports about a hat or capability requires sampling from the original set of hat and capability IDs defined for the scenario.

Information about lists of elements is made noisy in two stages. First, the list itself is modified by discarding or adding elements. Then, with probability p , each element of the resulting list is subject to a noisy process such as selecting an element at random instead of

the true list element (e.g., selecting a random hat ID).

Information that is requested about events or entities that do not exist are subject to noise in much the same way as above. If the request is determined to be non-noisy, then a query accurately returns the response that the requested information does not exist. If, however, the requested information is to be made noisy, then random information of the same type requested is returned.

3.3 Exporting Data

The Hats Simulator and Information Broker are designed to provide an online data feed and allow for interaction between the analyst and simulation. However, we also have implemented facilities to export batch data from the Information Broker. Hats data stored in the Information Broker can be exported as perfect information (ground truth) or noisy data sets. Application of noise for batch data works differently because there is no analog of online requests with user-varied payment levels. Noise is applied to exported data in three ways: exclusion of perfect information, inclusion of false information and corruption of perfect information. The level and type of noise is parameterized. Exported Hats data has been used in several projects, including DARPA EELD and AFRL EAGLE Program mini-TIEs, and controlled experiments with Iterative Node Classification [4], Group Finder [5], Proximity [6], Aleph [7], and a variety of social network analysis tools.

4 Actions

In addition to requesting information, the analyst playing the Hats game can also change a beacon’s alert level and arrest hats. Both actions affect an analyst’s performance score (discussed in Section 5).

4.1 Raising Alerts

We may not be able to stop an attack, but if we know it is coming, we can prepare and minimize loss. This is the inspiration behind modeling alerts. Each beacon can be in one of three alert levels: off (default), low, or high. These correspond to the conditions of no threat, a chance of an attack, and attack likely. The analyst decides which level a beacon alert is set to, but the Hats Simulator keeps track of alert states over time and whether an actual attack occurs while the state is elevated. The simulator keeps statistics including counts of hits (occurrences of attacks during elevated alerts) and false positives (elevated alerts that begin and end with no beacon attack occurring). The goal of the analyst is to minimize the time beacon alerts are elevated. High alerts are more costly than low ones. On the other hand, if an attack does occur on a beacon, a high alert is better than a low alert, and a low alert is better than none.

4.2 Arresting Hats

Analysts can also issue arrest warrants for hats in order to prevent beacon attacks. Arrests are successful only when the targeted hat is currently a member of a terrorist task force. Attempted arrests under any other conditions, including hats that are terrorists but not currently part of a terrorist task force, result in a false arrest (a false positive). Under this model, a hat can be a terrorist but not be guilty of any crime. Unless terrorist hats are engaged in *ongoing* terrorist activities, their arrest incurs penalties. While this is a simple model, it places realistic constraints on the analyst’s choice of actions.

Successful arrests do not guarantee saving beacons. A beacon is only attacked when some subset of members from a terrorist task force successfully carry the capabilities matching the target beacon’s vulnerabilities to a final meeting at on that beacon. It is possible to successfully arrest a terrorist task force member but the other terrorist task force members still have the capabilities required to attack the beacon. However, if the analyst successfully arrests a terrorist task force member carrying required capabilities that no other task force member has, then the final meeting of the task force will take place but not result in an attack. This is counted as a beacon save.

5 Scoring Analyst Performance

The Hats Simulator and Information Broker together provide an environment for testing analyst tools. The object of the game is to identify terrorist task forces before they attack beacons. Three kinds of costs are accrued:

1. The cost of acquiring and processing information about a hat. This is the “government in the bedroom” or intrusiveness cost.
2. The cost of falsely arresting benign hats.
3. The cost of harm done by terrorists.

The skill of analysts and the value of analysis tools can be measured in terms of these costs, and they are assessed automatically by the Hats Simulator as analysts play the Hats game. At the end of a simulation run, a final report is generated that includes the following four categories of scores:

1. Costs: the total amount of “algorithmic dollars” spent on information from the Information Broker.
2. Beacon Attacks: including the total number of terrorist attacks that succeeded and the total number of attacks that were stopped by successful arrests.
3. Arrests: the number of successful arrests and the number of false arrests (false positives).
4. Beacon Alerts: the number of low and high alert hits and false positives.

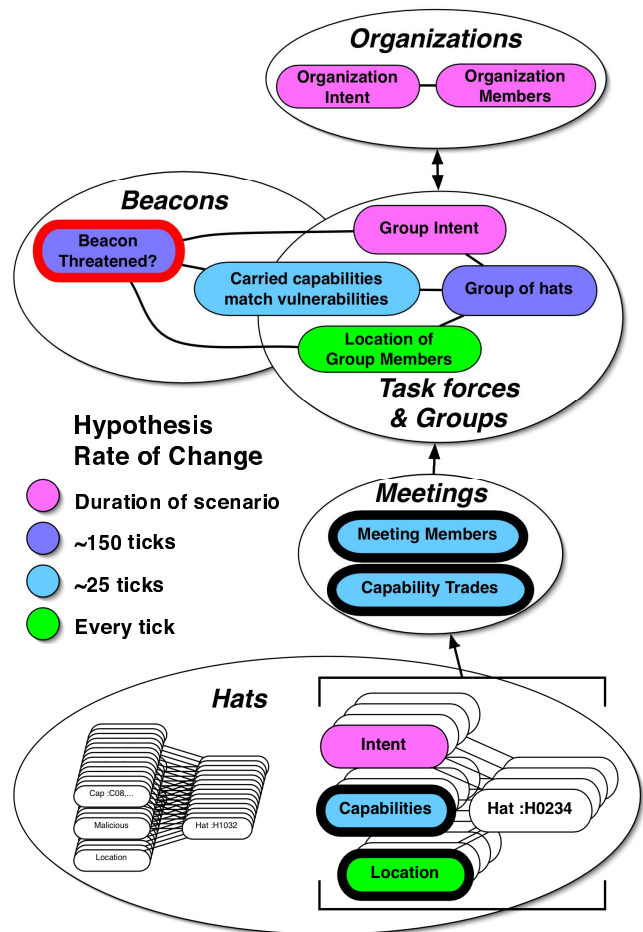


Figure 3: Schematic representation of Hats domain hypothesis types, their relations, and the approximate duration that instances of hypothesis types remain valid.

6 The Space of Hats Hypotheses and Fusion Tasks

As mentioned at the beginning of Section 3, successful play in the Hats game requires tracking and fusing a variety of information types. Analysts rely on reports returned by the information broker to construct a model of observed events that can explain and predict hat behaviors. The goal of this analysis is to identify threats to beacons and, if possible, arrest terrorist task force hats before they carry out an attack.

Figure 3 provides a schematic representation of the variety of classes of hypotheses, along with how hypotheses are related to one another and the approximate rate at which they change.

Each colored oval in the figure represents a hypothesis type, and these are grouped into five larger categories that include hypotheses about hats, meetings, groups (task forces), beacon threat and organizations. Links between ovals indicate dependency relations between hypothesis types. In general, information flows up from the bottom the figure. The four ovals with bold black outlines indicate information that may be requested directly from the information broker. For in-

formation about the location and capabilities carried by a hat, the analyst need only specify the hat ID. For information about meeting members and capability trades, the analyst must specify the time and location of the meeting. This makes requests for information about meetings depend on first identifying where and when a meeting took place. Further up, a hypothesis that there exists a task force of hats depends on an observed pattern of meetings between hats; hats that are members of a task force are, in turn, likely members of the same organization.

The primary target hypothesis concerns whether a beacon is currently threatened. (This hypothesis is highlighted with bold red outline.) This hypothesis depends on information that includes belief in the existence of a group, whether the capabilities shared by group members overlap with the beacon’s vulnerabilities, the location of the group members relative to the beacon, and finally the intent of the group.

Figure 3 also shows that information supporting hypothesis changes at different time scales. The figure represents for different relative time scales using colors; the legend on the left indicate the approximate time at which a specific hypothesis would likely need to be updated, in simulator ticks. For example, the location of a hat changes almost every tick (green), the set of capabilities a hat is carrying updates roughly every 25 ticks (blue), but the intent of a hat (whether it is terrorist or benign) remains constant throughout the game (pink).

7 Conclusion

Despite the simplicity of the domain ontology, Hats provides a number of nontrivial information fusion challenges. Taken as a whole, the Hats fusion challenge includes the following dimensions:

- Information access is not uniform. The analyst must determine both what queries are worth making as well as navigate the tradeoff between information cost and quality.
- Hats represents events unfolding over time. Rather than selecting queries over a static database of information, the analyst must consider how information changes and update her representation of the current state of the simulation. Furthermore, the rate at which information changes over time varies depending on the type of information.
- Hats generates a variety of types of information that include both scalar and categorical values: locations (of hats, beacons and meetings), resources (capabilities), behaviors over time (hat movements, capability trades, plans to transport capabilities to a target), and hidden group (task force) and organization membership.
- All behaviors are planned and therefore meaningful in the sense that they are goal-directed. These

goals cannot be inferred from just one type of information; identifying plans requires incorporating information about the individuals involved in the plan (group membership), member locations (vicinity to a beacon or heading to another meeting), capabilities carried and traded, and hypotheses about intention (suspicion based on association with other known terrorists).

In Section 3.3 we noted experimental studies that have used Hats data to test a variety of data mining and information fusion tools. We are also using Hats as a challenge problem domain in the COLAB system [8], a laboratory environment designed for analyst training and the study of distributed analyst collaboration. Within COLAB, multiple human analysts in different physical locations can collaborate to build hypotheses of unfolding Hats scenarios using a variety of tools. COLAB consists of two components: an instrumented analysis working environment built on a blackboard system, and a web-based user interface that integrates the TRELIS [9] hypothesis authoring and management tool with a query language. On the blackboard, analysts collect data and perform analyses using relational database queries and a set of analysis tools that include group identification and suspicion scoring algorithms. Analysts may also define automated agents, instructing them to look for conditions that may arise in data, or carry out routine tasks such as automated queries for tracking specific hats. Evidence extracted from data analysis on the blackboard may then be incorporated in hierarchical structured arguments in the TRELIS tool, combining evidence sources with user-supplied free text. TRELIS arguments may be shared between analysts and collaboratively authored. COLAB has been integrated with the Hats Simulator and serves as a prototype mixed-initiative information fusion system.

We are told by intelligence analysts that Hats has many attributes of “the real thing.” Some say in the same breath that Hats ought to have other attributes, for instance, telephone communications, rapid transportation of hats around the board, different kinds of capabilities, and so on. We resist these efforts to make Hats more “realistic” because for us, the purpose of Hats is to provide an enormously difficult detection problem without the overhead of building rich (and probably classified) models of real domains. No doubt Hats will change over time, but we will strive to keep it simple. The other goal that guides our development of Hats is what we might call the “missing science” of intelligence analysis. To the best of our knowledge, in the current climate, analysts penalize misses more than false positives. This sort of utility function has consequences – raised national alert levels, lines at airports, and so on. Hats is intended to be a simulated world in which analysts can experiment with different utility functions. It is a laboratory in which scientific models of intelligence gathering, filtering, and use – models based on utility and information theory – can be tested and compared.

8 Acknowledgements

The Hats Simulator was conceived of by Paul Cohen and Niall Adams at Imperial College in the summer of 2002. Dr. Cohen implemented the first version of Hats. Work on this project was funded by the Air Force Research Laboratory, account number 53-4540-0588.

References

- [1] P. R. Cohen and C. T. Morrison. The hats simulator. In *Proceedings of the 2004 Winter Simulation Conference*, 2004.
- [2] C. T. Morrison, P. R. Cohen, G. W. King, J. J. Moody, and A. Hannon. Simulating terrorist threat with the hats simulator. In *Proceedings of the First International Conference on Intelligence Analysis*, 2005.
- [3] A. C. Hannon, G. W. King, C. T. Morrison, A. Galstyan, and P. R. Cohen. Population generation in large-scale simulation. In *Proceedings of AeroSense 2005*, 2005.
- [4] A. Galstyan and P. R. Cohen. Identifying covert sub-networks through iterative node classification. In *Proceedings of the First International Conference on Intelligence Analysis*, 2005.
- [5] J. Adibi, P. R. Cohen, and C. T. Morrison. Measuring confidence intervals in link discovery: a bootstrap approach. In *Proceedings of the ACM Special Interest Group on Knowledge Discovery and Data Mining (ACM-SIGKDD-04)*, 2004.
- [6] J. Neville and D. Jensen. Supporting relational knowledge discovery: Lessons in architecture and algorithm design. In *Papers of the ICML 2002 Workshop on Data Mining Lessons Learned*, 2002.
- [7] A. Srinivasan. The aleph manual. Technical report, Oxford University, <http://www.comlab.ox.ac.uk/oucl/research/areas/machlearn/Aleph/>, 1999.
- [8] C. T. Morrison and P. R. Cohen. Colab: A laboratory environment for studying analyst sensemaking and collaboration. In *Proceedings of the Tenth International Command and Control Research and Technology Symposium (10th ICCRTS)*, 2005.
- [9] T. Chklovski, V. Ratnakar, and Y. Gil. User interfaces with semi-formal representations: a study of designing argumentation structures. In *Proceedings of the Conference on Intelligent User Interfaces (IUI05)*, 2005.