

Title: **Identification of Adversarial Activities: Profiling Latent Uses of Facilities from Structural Data and Real-time Intelligence**

Suggested Topics:

Topic 2: Networks and Networking

Topic 7: Network-Centric Experimentation and Analysis

Topic 3: Modeling and Simulation

Authors:

Darby E. Grande, Ph.D.

Georgiy M. Levchuk, Ph.D.

E. Webb Stacy, Ph.D.

Aptima Inc.,

12 Gill Street, Suite 1400

Woburn, MA 01801

Phone: 781-935-3966x225, x267

Fax: 781-935-4385

e-mail: dgrande@aptima.com,

georgiy@aptima.com, wstacy@aptima.com

Martin Kruger

Office of Naval Research

975 N Randolph Street, Suite 1425

Arlington, VA 22203-1995

Phone: 703.696.5349

e-mail: Martin_kruger@onr.navy.mil

Correspondence:

Darby E. Grande

Aptima Inc.

12 Gill Street, Suite 1400

Woburn, MA 01801

Phone: 781-935-3966x225

Fax: 781-935-4385

e-mail: dgrande@aptima.com

Abstract

Intelligence Preparation of the Battlefield (IPB) provides critical support for military planning and decision making. In both offensive and defensive operations, the IPB process gives a decision maker information about the enemy, including potential courses of action, as well as information about the environment in which he is working. Physical structures that can support repetitive crimes (such as IED supply chains or illegal drug trafficking) provide important information for C2 planning. Activities of interest must be situated *somewhere*, and the physical structures present in any given location change slowly. Knowledge of those structures and their capabilities therefore provide an effective lens through which to view activities, and therefore an effective means for attacking the problem.

In this paper we discuss the Facility Identification via Networks with Adaptive Links (FINAL) technology that Aptima is developing to find facilities associated with adversarial actions and discover the intent for their use. Based on algorithms that perform probabilistic network pattern identification from partial knowledge about network nodes, links, and their attributes, FINAL profiles the use of facilities by combining networks of data describing actual conditions and more abstract network models of repetitive crimes.

Introduction

A drug cartel offloads its cargo from a merchant ship to an unoccupied warehouse in Boston. A gun running organization makes a large transaction in an old farmhouse in a sparsely populated area in Georgia. An insurgent organization manufactures IEDs in a small house in northeastern Iraq. Each of these operations has its own “invisible” supply chain. Each link in the supply chain must operate in some physical space within some geophysical location, but the supply chains need to stay invisible in order for the operations to achieve success. Identifying, monitoring, and rendering inoperable the physical structures required to mount these operations is a promising way to intervene and eventually destroy these operations.

Unfortunately, the collected data contains large information gaps: much of the critical information is not collected due to limited sensors and the concealment of activities; the amount of deception increases as adversaries learn over time; and the actions of adversarial and neutral groups blend together to prevent clear distinctions. In addition, in most cases there is too much information, it is too fragmented, and the interrelationships are too complex. Attempts to understand incoming information and to relate it to past experience can baffle even seasoned experts. In previous research for DARPA, Aptima tackled an analogous problem—recognizing the roles of and relationships among individual actors and their resources to predict the structure and intent of networked adversaries in order to develop targeted counteractions against them (Levchuk and Chopra, 2005). In empirical studies, our automated network identification technology, called NetSTAR (Network, Structure, Activities, and Roles), outperformed unaided human analysts in accuracy of identification and was able to handle significantly larger levels of uncertainty and data complexity (Levchuk et al., 2006; Levchuk et al., 2007).

Sponsored initially by a Phase I SBIR award from the Office of Naval Research, Aptima is developing the **Facility Identification via Networks with Adaptive Links (FINAL)** tool to support human analysts in finding facilities associated with adversarial actions and discovering the intent for their use. The framework is based on the NetSTAR technology, which performs probabilistic network pattern identification based on partial knowledge about network nodes, links and their attributes. To identify the facilities used for adversarial activities, FINAL performs probabilistic matching of observed facility networks against hypotheses of action patterns performed by adversaries. The facility or **data network** is a labeled graph of facilities, their links, and node/link attributes drawn from available structural and operational information about buildings, roads connecting them, suspected and observed criminal activities in the area, and other concrete intelligence information. The activity or **model network** patterns form the hypothesis set; they are networks created with the help of subject matter experts and/or captured from histories of repetitive adversarial actions. The matching of a hypothesized model network

pattern against the facilities of observed data network corresponds to node-to-node **activity-to-facility mapping**, thus identifying the intended past, current, or future use of the facilities.

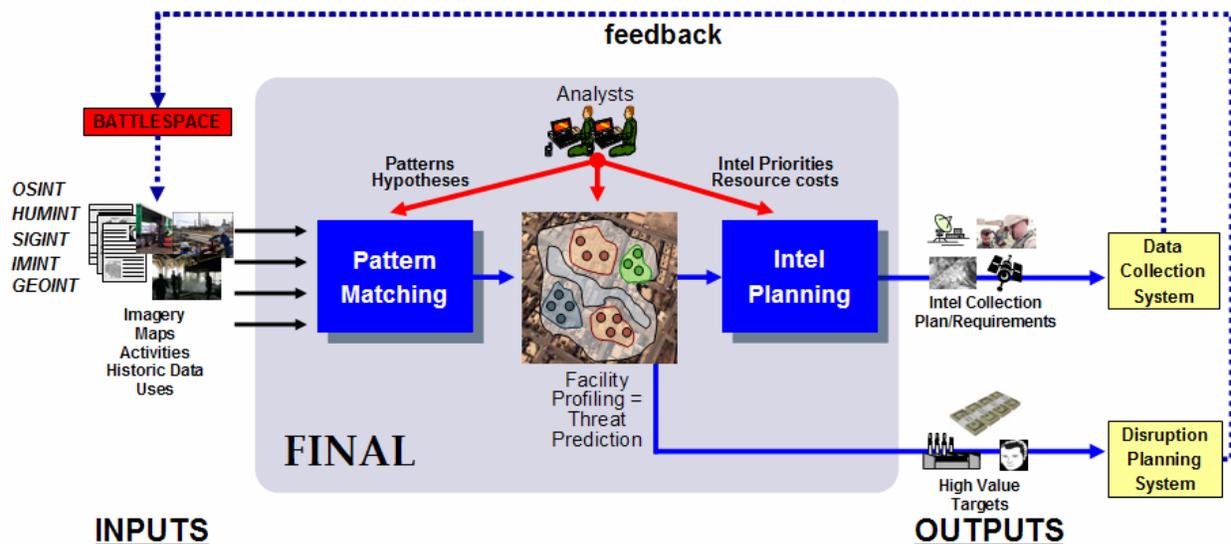


Figure 1: Components and Workflow of FINAL Technology

The matching is performed using probabilistic attributed-structural pattern consistency scoring based on maximum likelihood or maximum a-posteriori estimation, so that the hypothesized model network patterns which best explain (or best match) the observed facilities of the data networks score the highest. These patterns are used to assess the true uses of facilities. Since this approach uses probabilistic pattern matching and is not constrained to perfect matches, our model is *robust to uncertainty* in the data and allows discovery of *facility groupings for common activity involvement* and the *association of facilities with multiple actions*. In addition, the hypothesis testing approach is well-suited to guide additional intelligence collection activities to maximize the information gain (uncertainty reduction).

As shown in Figure 1, the workflow of the FINAL technology begins with the network pattern matching, results in a threat prediction map from which information from confident, or highly-scoring, matches may be used as inputs for disruption planning, and matches with lower confidence may be used to develop intelligence collection plans based on FINAL’s knowledge of the missing discriminatory information. In the first phase of this work, we have focused on developing the facility prediction algorithms and a simulation environment in which to test them. In future phases, we plan to advance the sophistication of our input (data) treatment as well as the interactive visualization capabilities for the results.

In the sections below, we further discuss the model and data network foundations of FINAL, introduce our solution approach, present results from initial use case study, and describe our intended future directions and extensions for this work.

Data and Model Networks

The key components of FINAL are its *data* and *model networks*. As shown in Figure 2, these networks are the inputs to the pattern-matching algorithms. We represent the observations in the form of a network, called the data network, where each node corresponds to a specific facility, and the links correspond to the relationships among the facilities. Nodes and links are quantitatively defined through their sets of attributes or features indicating:

- the **capabilities** of the facilities and their connecting roads; and
- the **uses** of the facilities and their connecting roads.

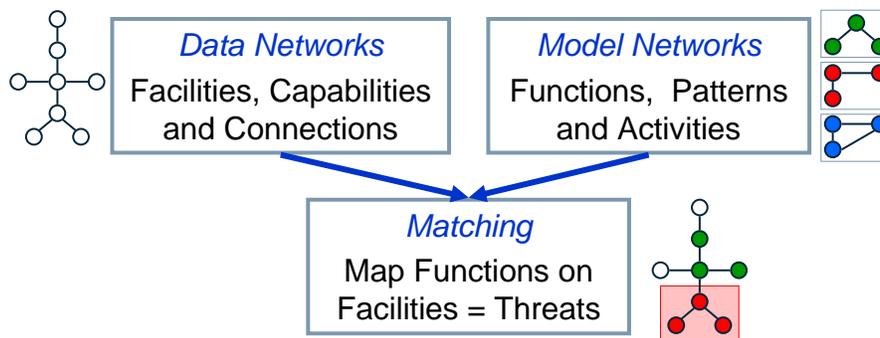


Figure 2: Data and Model Networks

Capabilities influence the potential uses of the facilities and are referred to as *enablers* of criminal acts in the criminology literature (Brantingham & Brantingham, 1995). For example, the size of a building, the number of floors, the number and size of building entrances, and the height of the ceiling may indicate whether the building can be used as a school, an automobile repair shop, a weapons storage facility, or residential quarters. Similarly, a road’s width and pavement quality, the road’s length, and the amount of security on the road can be indicators of whether the road can be used to smuggle drugs, transport weapons, or deliver disaster relief supplies.

Previous *uses* of facilities can forecast potential similar criminal activities; such uses are referred to as *attractors* in criminological science. For example, a facility used to store perishable foods could be used effectively to store weapons components. Also, current facility uses by adversaries represent the signals about true facility employment, and uses by non-adversarial groups may indicate attractive opportunities for criminals to use the facility for their purposes; such data are referred to as criminal activity *generators*. For example, facilities used as schools may be targets for kidnapping and drug trafficking, facilities used by local police may be targets for VBIED attacks, and residential buildings may be used to store small quantities of weapons because this provides good concealment and discourages direct targeting by opposing forces (Brantingham & Brantingham, 1995).

Accordingly, we construct two types of networks – a *capability network* and a *facility use network*, - and overlay them to form a facility data network (shown in Figure 3). There are two sources of data that can be used to construct capability and use networks: static data sources and dynamic data sources. Static sources refer to the data elements available a priori about the area of interest. These include satellite imagery already collected, information about the setup of the security checkpoints and sensor locations, maps of the area and intended facility use data (locations of hospitals, schools, police stations, markets, shops, etc.). Static data sources may also include historic information about previous facility uses. Some of these data may be outdated and need to be updated using more recently collected intelligence. Finally, static data sources allow the extraction of structural information about facility networks – geo-spatial locations, routes, sizes and other capability features of the man-made structures.

Dynamic sources refer to data gathered by active intelligence collection missions. These include patrols, searches, interrogations and interviews with locals, dynamic UAV imagery, etc. This intelligence can capture which facilities are currently in use, who uses them, and for what purposes. This information can be used to connect facilities when the same vehicles or people are spotted near them, when documents discovered during searches indicate the flow of materials among facilities, and when interviewers point to the relationships between people and activities at different facilities. As discussed further below, intelligence collection can be guided to enable discovery of missing features critical to improving predictions of facility employment (e.g., imagery of the building from another vantage point is needed to determine the number of the floors).

We also define the types of features that describe the facilities, their uses, and the connections among them. At present, after initial prototype stage of tool development, we continue to explore available sources for these varied data types. In some cases, we believe that attribute (capability) data are readily

available, but an important part of our future research is continuing to explore the technologies that could automate the facility feature extraction. For example, both static and dynamic data sources include various types of imagery collected from different vantage points and different ranges of the electromagnetic spectrum. There are technologies that allow finding what is in the image (shapes of buildings, materials of structures, patterns of roads, etc.), but coupling the images with symbolic map data can significantly improve the ability to identify the structural (capability) networks. The maps carry context for image exploitation and provide relationship-adjacency of buildings, connectivity of roads and traversability of waterways. Maps, however, are sometimes outdated, and being able to find the differences between existing maps and more recent imagery is necessary for performing map updates and may require active intelligence collection. Adding other data sources, such as original construction plans, architectural specifications, phone books, etc., will improve our ability to capture diverse attributes of a facility network.

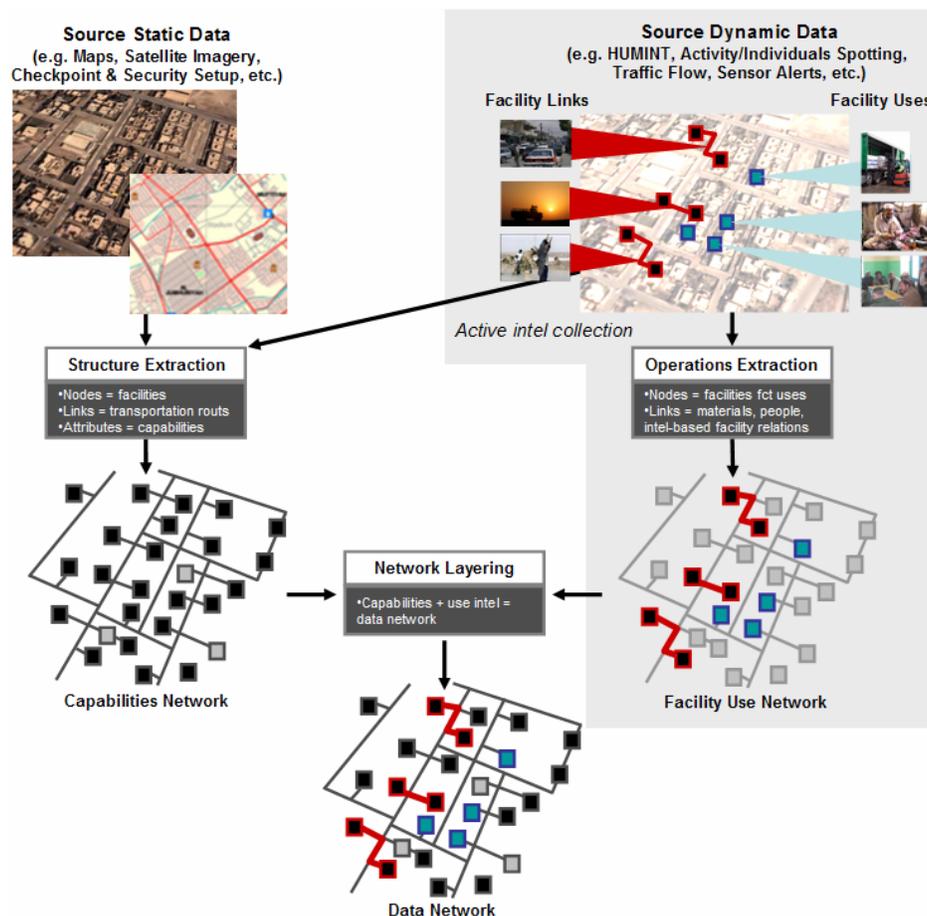


Figure 3: Data Network Construction

The *model networks* consist of nodes, indicating the functions or activities that adversaries might perform, and the necessary links between them. The nodes and links of model networks are associated with attributes that can relate them to the nodes and links of the data network, and generally will represent concepts from the building-knowledge ontology. Thus, the conceptual basis for both data networks and model networks is the same. The nodes of a model network, which are the functions performed by adversarial groups using man-made structures, can be associated with functional requirements and even with specific facilities and geo-spatial information. For example, a generic function of “weapons assembly” can have a definition of capabilities of the facility that can be used to perform this function. Then, the links of a model network can represent the activities that tie the facility

functions together: for example, the function of weapons assembly would require bomb making materials to be transported to the assembly location from their storage location.

The model networks can represent both (1) patterns of potential adversarial activities, and (2) the patterns of specific facilities utilization. In the latter case, the data about ID, location, or a type of facility used to perform a function can also be associated with the node of model network.

Formulation of Facility Use Prediction

In this section we discuss the algorithms developed to identify the uses of facilities. Formally, *function* (or *activity*) *pattern* is represented as a graph $G_M = (V_M, E_M, A_M)$ – a **model network** where V_M is a set of function/action nodes, E_M is a set of links between them, and A_M is a set of attributes on links and nodes ($A_M = \| a_{ij}^M \|$, and a_{ii}^M is attributes vector for node v_i^M and a_{ij}^M is attribute vector for link between nodes v_i^M and v_j^M). Observed data are aggregated to a **data network** – a graph $G_D = (V_D, E_D, A_D)$. Here, V_D is a set of observed facilities, and E_D is a set of observed relationships among them, and A_D is a set of attributes on facilities and their connections.

We need to discover the *mapping from functions to facilities* – that is, from the nodes of the model graph to nodes of a data graph. This is accomplished by finding an assignment matrix $S = \| s_{ij} \|_{i \in V_M, j \in V_D}$, where $s_{ij} = 1$ if model node (activity/function) i is mapped to data node (facility) j . This mapping means that facility j is used to perform function i , and essentially corresponds to how the enemy selects facilities for a specific activities in their operations.

We find an assignment matrix S that maximizes the likelihood function $P(G_D | G_M, S)$ equal to the probability that the observed facilities network (data network) has been generated by the hypothesized function network (model network) given the assignment of facilities to functions (mapping between nodes of data and model graphs). While direct optimization of the likelihood is infeasible, an approximate solution can be obtained. Several approaches have been explored by our team, including:

- (i) matrix norm minimization;
- (ii) quadratic assignment problem;
- (iii) structural consistency measure relaxation;
- (iv) Hidden Markov random fields.

We began with the matrix norm minimization, implementing the algorithm in a prototype software environment to showcase the model capabilities for finding the hostile facilities. From here we moved to solving the quadratic assignment formulation of the problem, as this allows us to better model the errors in the data. In the quadratic assignment formulation, we reduce the problem to

	$\min_S Q(S) = \frac{1}{2} \sum_{km:ij} s_{ki} s_{mj} \cdot c_{km:ij} + \sum_{ki} s_{ki} \cdot c_{k:i}$ $s.t. \sum_i s_{ki} = 1$	[1]
--	--	-----

where the function $c_{km:ij}$ scores the mismatch between the links (k, m) in the model and (i, j) in the data network and $c_{k:i}$ scores the mismatch between nodes k in the model and i in the data.

There are several solutions to the problem [1]. One such solution that achieved a very high accuracy and showed allowable complexity has been reported in (Rangarajan, Yuille, and Mjolsness, 1999). The algorithm, called graduated assignment, is a mixture of the continuous gradient-based optimization, matrix

scaling, and a soft assignment, and this is the solution approach we have employed in the initial prototype implementation of FINAL. Graduated assignment algorithm iteratively finds a continuous

approximation matrix $p_{ki}[n+1] = \frac{\exp\{-\beta \cdot \tilde{p}_{ki}[n]\}}{\sum_r \exp\{-\beta \cdot \tilde{p}_{kr}[n]\}}$, where $\tilde{p}_{ki}[n] = \sum_{m,j} p_{mj}[n] \cdot c_{km:ij} + c_{ki}$. The

mapping is then found using soft-max principle from matrix $p_{ki}[n]$. The β is gradually increased, resulting in $p_{ki}[n]$ approaching 0-1 matrix. The graduated assignment algorithm reduces the complexity and enables iterative update of the solution with incoming intelligence data.

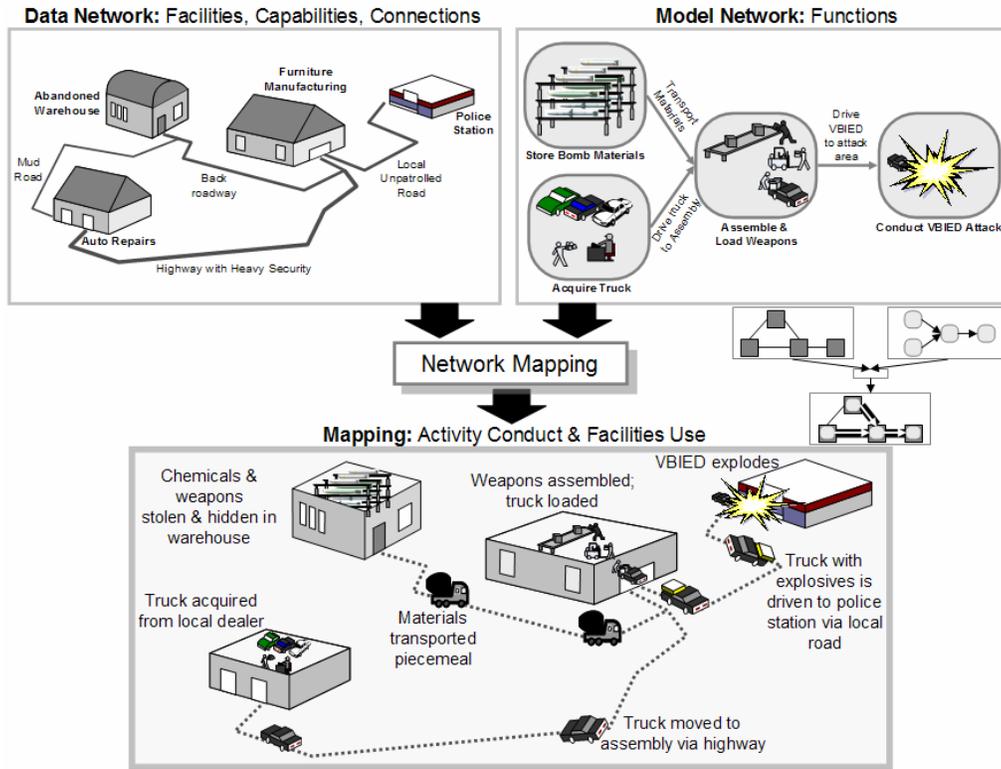


Figure 4: Example of Facility Profiling from Action-Facility Mapping

Upon each application of the mapping algorithms, not only do we obtain the correspondence of tracked facilities to specific activities/functions in each hypothesized behavior pattern, but we can also rank-order these associations for each hypothesis using values of likelihood function $P(G_D | G_M, S)$. This function is often normalized to prevent scalability issues when the number of feasible observations is very large.

The attributed network matching algorithm allows performing the following predictions:

- **Outcome 1: Facility utilization profiling.** Given a model G_M of the activity/function pattern that needs to be explored and the data about facilities G_D , the algorithm comes up with single or multiple mappings S of the functions to facilities that “best explain” obtained observations (shown in Figure 4). The mapping of function to facility indicates the assessment that the function/activity can be performed at the facility as a part of the planned activity/function pattern. Each mapping S is scored with the likelihood function $P(G_D | G_M, S)$. If multiple functions are mapped to a single facility, it indicates that a facility is involved in several potentially illegal activities.

- **Outcome 2: Facility groupings and terrain classification.** When multiple facilities are matched to the activity pattern, they represent a set of resources, or a supply chain, used by adversaries over time. Understanding how these facilities are used and which are the most critical ones will allow for better disruption of the criminal activities. When multiple mappings are developed that indicate several possible alternative function/action executions (that is, multiple facilities, and not a single one, are forecasted to be employed to conduct the activity), this will indicate to the intelligence analysts that potentially either all the facilities need to be shut down, or that additional intelligence gathering need to take place.
- **Outcome 3: The behavior and abnormalcy identification.** The outcome of mapping algorithm can also be used to score each model of potential hostile and benign activities. The algorithm can thus distinguish which behaviors are happening in the environment of interest – what method of attack is ongoing, or whether the normal activity is happening instead of adversarial one. We can use a maximum a-posteriori estimation $P(G_M, S | G_D)$ to score each model, where

$$P(G_M, S_M | G_D) = \frac{P(G_D | G_M, S_M)P(G_M, S_M)}{P(G_D)}. \text{ Computation of the } P(G_{M_j}, S_{M_j}) \text{ can be done}$$

approximately using the frequencies that the model behavior (M_j) has been happening in the past and the decision about facility employment that adversaries were making (S_{M_j}).

Simulation Environment and Use Case

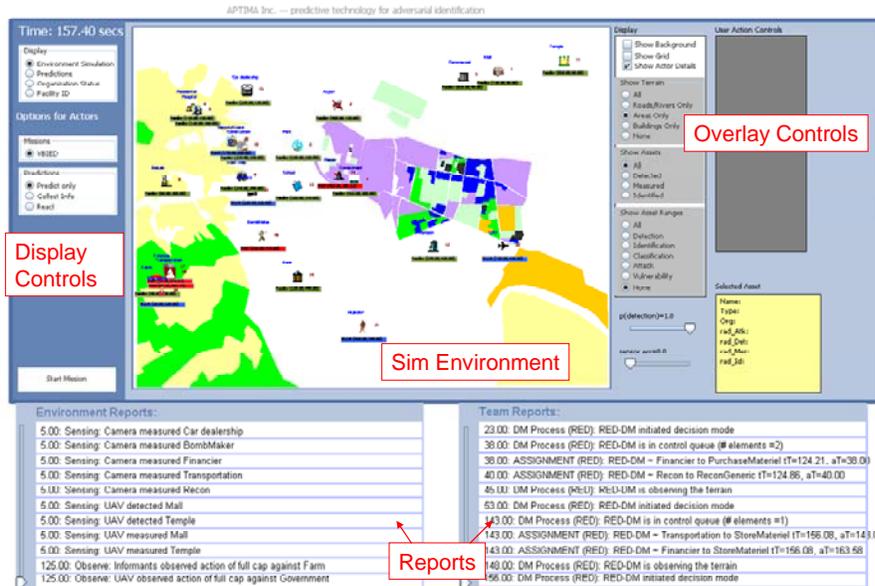
In order to test the facility identification algorithms, we have developed a simulation environment for the purpose. This environment and its support of our illustrative use case are described here.

As an initial prototype implementation for FINAL, we have developed a testbed in which we mimicked the RED mission execution using distributed team decision making for task selection and resource allocation. The virtual simulation component of the testbed (Figure 5a) was given a selected RED mission and actors. Simulation then generated the feasible RED mission execution events – equivalent to the true mission plan and mapping of activities to facilities. This information was hidden from FINAL predictive component that received the noisy data from the environment as observed by BLUE sensors (which were part of the virtual simulation).

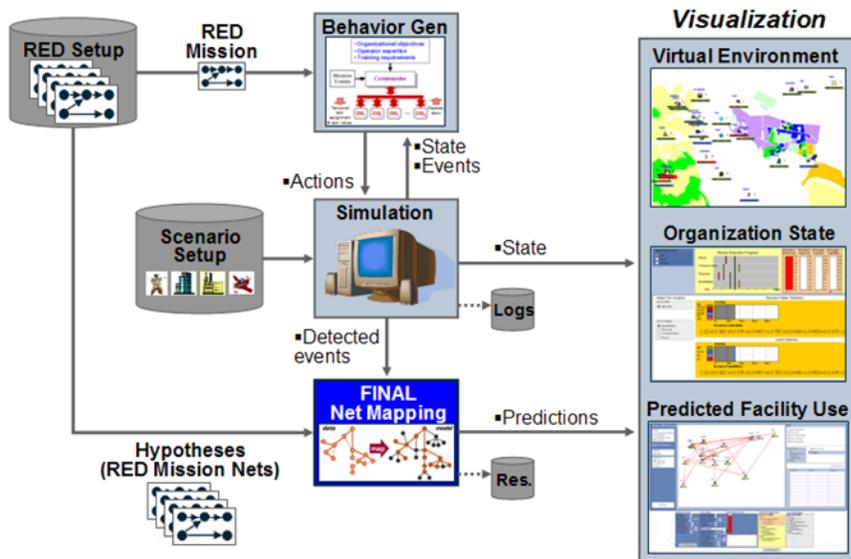
The FINAL prototype (Figure 5b) includes the simulation, prediction, and visualization components. Simulation is instantiated with the database containing the scenario setup (all terrain objects, including facilities, RED and BLUE actors, their attributes, etc.). It is driven by the behavior generation algorithms responsible for selection of tasks and actions for RED and BLUE actors. The behaviors are driven from the definition of the mission for RED team which was stored in a separate database. The multiple possible missions for RED team also served as the hypotheses/model networks for FINAL network mapping predictive algorithm. Developed visualization components show the virtual environment evolution, the state of RED and BLUE organizations (actors status, mission progress, etc.) and the predictions generated by FINAL algorithm.

The FINAL predictive component generates the estimate of the RED’s mission plan and action-facility mapping, and this estimate was compared to the simulator’s output as the ground truth.

To populate the simulation environment for the use-case, we developed a model of adversarial attack consisting of the creation and detonation of a dirty bomb in a fairly urban area. We began by brainstorming the steps that might be taken to accomplish such an attack, as well as the facilities and links (roads) that might be required. This initial story of how the attack might be accomplished is shown in Figure 6.



(a) Virtual Environment Simulation Visualization



(b) Prototype Components

Figure 5: FINAL Prototype

We hypothesized the existence of a RED commander who coordinates hostile activities: conducting surveillance at several potential targets, obtaining required materials from a number of different sites, coordinating transportation for the materials and the surveillance, assembling the weapon, and executing the attack. These steps require a number of different facility types, as well as connections between them for transportation. One could imagine a number of variations on this set of activities, as well as several components of this story working together for benign purposes. In the future work, we will therefore expand the collection of hypothesized function networks to include such variations and “normal” activities.

The first step in creating an environment with which to test the FINAL algorithms has been to construct the facility network. This is the network of capabilities and uses described above. For this initial prototype testing, we chose the map of a town area to serve as our static data layer. From this we manually extracted the network of facilities and assigned attributes.

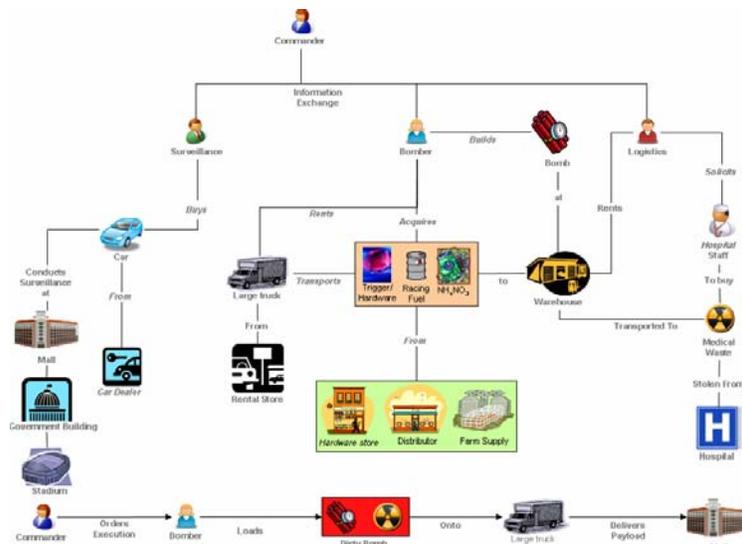


Figure 6: Use Case Storyline

During facility network construction, we also created a road network and assigned attributes to the facilities and roads. These attributes can be obtained from aerial/satellite imagery (e.g., footprint of the building; the roof type; the fact that it is a sport, entertainment, or a parking lot; the surrounding neighborhood; the tree density and landscaping properties; etc.), close-up imagery (e.g., wall size and material construction; the building height and number of the flows and windows; any visual signs on the building about its use; etc.), labeled maps (e.g., original construction information; intended uses; etc.), historic data (e.g., police incident reports in or around facility; public records; cultural and ethnic information about the neighborhood; socio-economic data profile of inhabitants/attendees/neighbors; etc.), and current reports (e.g., spotting of suspect people and activities around the building obtained from HUMINT reports or local informants; etc.). In our use case, we developed a collection of attributes based on the data schema used in a future potential integration site for FINAL.

To formalize this activity, we have created a list of attributes, list of facilities, and set of action patterns for our use case.

First, we developed a set of attributes using three classes (Figure 7):

- *attributes describing structural capabilities of facilities*: these data can be collected from analyzing imagery by automated or manual means, using information such as radar scans and intelligence data about availability of resources at facilities, etc.
- *attributes describing current events and facility uses*: this information is dynamic and can be obtained from human collection teams, UAV data feeds, etc.
- *attributes describing previous uses of facilities*: this is historic information that could have been obtained from the past events in the area of interest.

These attributes formed the vectors that capture information about facilities and their links (observed attributes) and actions and their relationships (ground truth) as shown conceptually in Figure 8. For simplicity of analysis in the current prototype stage, we focused only on two types of attributes – capabilities and current events/uses.

To lay the ground work for algorithm testing, we have modeled 4 types of noise in the data that made the observed attribute vectors different from ground truth:

(1) Event miss: Events about the uses of facilities are captured by BLUE sensors (SIGINT, HUMINT, IMINT), and not all such events might be detectable at some facilities. For example: *Facility was used to*

hold a meeting between terrorists, but there was no UAV/patrol at the time in the area. As an outcome, all attributes from the missed event are excluded from analysis.

(2) **Attributes miss:** Sensors (humans, data processing algorithms, etc.) might miss an attribute present in the incoming data/event. For example: *LIDAR data were incorrectly analyzed by the image classification algorithm.* As an outcome, correct attribute was missed and excluded from the analysis.

(3) **Attributes/event deception:** Sensors (humans, data processing algorithms, etc.) might falsely perceive that an attribute was present in the incoming data/event or might falsely add an event due to deceptive information that has never occurred. For example: *Analyst, based on studied imagery, reported the presence of a hide-out at the construction site.* As an outcome, incorrect attribute is added as input and is used for analysis.

(4) **Attributes errors:** Sensors (humans, data processing algorithms, etc.) might incorrectly assess the value of an attribute in the incoming data/event. For example: *Analyst, based on studied imagery, reported that the building had large footprint, while building had medium-to-small footprint.* As an outcome, incorrect attribute value is used as input for analysis.

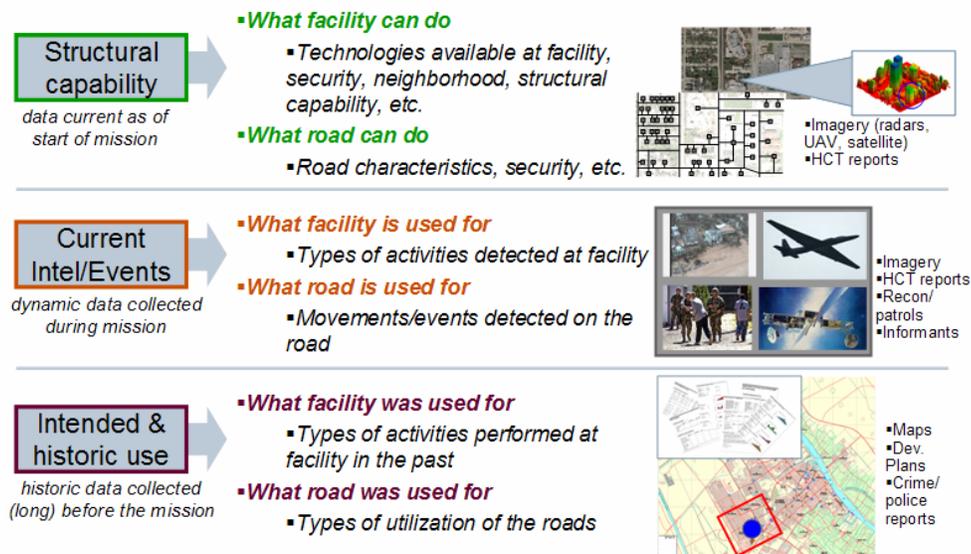


Figure 7: Attribute Classes

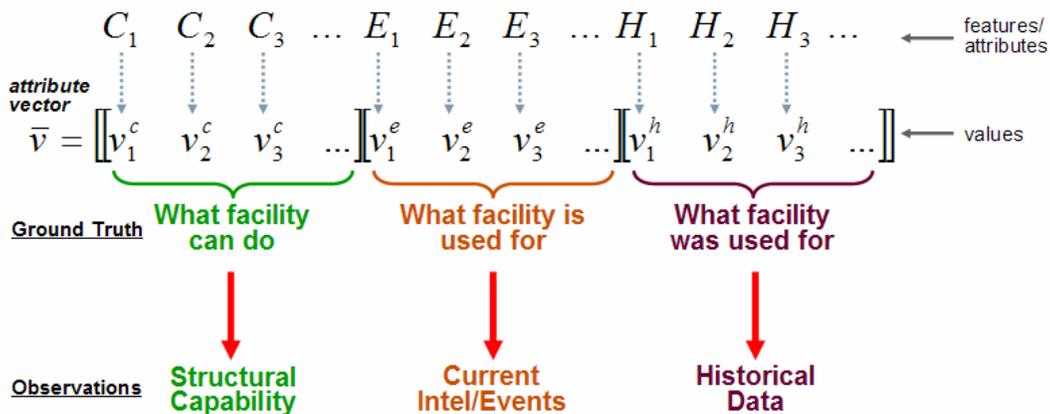


Figure 8: Information Capture in the form of Attribute Vectors

In our use case, we defined 16 facilities of interest in the terrain. Their attributes are shown in Figure 9.

13th ICCRTS-2008 “C2 for Complex Endeavors”

For example, BioLab had size=1, security=2, had materials (capability=2) that could be purchased by RED (vulnerability of money=2), could be attacked (vulnerability of attack = 1), reconnoitered (vulnerability of recon = 1), could be used to attempt to store materials but no capability of doing so successfully.

The RED team in the environment was represented by 4 actors that could be either single people or small cells (see Figure 10). Actor capabilities were defined to match them with tasks that were part of RED’s mission as described in next section. BLUE actors were represented by BLUE human collection teams and aerial sensors.

The second component of the Use Case is the set of action patterns to be matched against the facilities network. These are the hypothesized activities supported by the facilities that RED could execute as part of their plan. For the initial use case introduced above, we developed an adversarial action pattern to analyze which describes 5 tasks/activities in preparing for and carrying out a dirty bomb attack, as shown in Figure 11. The set of action patterns we are continuing to develop will contain variations upon this use case baseline action graph, alternative adversarial attack possibilities, as well as benign actions that may explain the data and observations we collect and analyze.

	Capability (what it CAN do)													Vulnerability (what can be done TO it)												
	SZ	SEC	STR	MAT	TEC	KNW	MON	ATK	VAL	REC	TRS	SZ	SEC	STR	MAT	TEC	KNW	MON	ATK	VAL	REC	TRS				
BioLab	1	2	0	2	1	0	0	1	0	0	0	0	0	0	1	0	1	2	1	0	1	1				
Mall	2	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	1	1	0	1	1	1				
Airport	3	0	0	0	0	0	0	0	2	0	0	0	0	0	0	0	1	3	0	1	1	1				
Park	3	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1	1	1				
Residential	1	3	1	1	0	0	0	0	0	0	0	0	0	0	1	0	0	0	1	0	1	0				
Commercial	1	0	1	0	1	0	0	1	0	0	0	0	0	0	0	1	2	1	0	1	0	0				
Farm	3	3	3	1	0	0	0	0	0	0	0	0	0	2	0	0	1	0	0	1	1	1				
Government	1	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	2	0	1	0	0				
PublTrnsp	1	1	0	0	0	0	0	0	1	0	0	0	0	0	0	1	1	0	1	0	1	0				
Hospital	2	1	0	3	0	0	0	0	0	0	0	0	0	1	0	1	0	1	0	1	0	0				
Construction	2	2	1	0	2	0	0	1	0	0	0	0	0	2	0	1	0	1	0	1	1	1				
Temple	2	4	2	0	0	0	0	0	0	0	0	0	0	2	0	0	0	1	0	1	1	1				
Mansion	1	4	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	1	1				
School	1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	1	0	1	0	1	0				
Bank	0	0	0	0	0	0	2	0	1	0	0	0	0	0	0	0	1	1	0	1	0	0				
Car dealership	1	2	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0				

Figure 9: FINAL Use Case --- Observed Facilities and their Attributes

	Capability (what it CAN do)												
	SZ	SEC	STR	MAT	TEC	KNW	MON	ATK	VAL	REC	TRS		
BombMaker	0	0	0	0	0	1	0	0	0	0	0		
Financier	0	0	0	0	0	0	1	0	0	0	0		
Transportation	0	0	0	0	0	0	0	0	0	0	1		
Recon	0	0	0	0	0	0	0	0	0	1	0		

Figure 10: FINAL Use Case --- RED Actors and their Capabilities

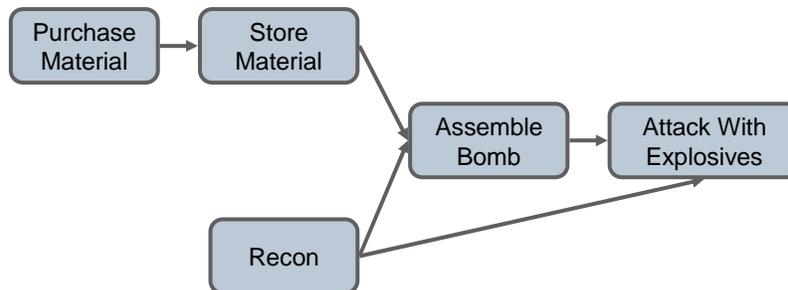


Figure 11: FINAL Use Case --- Dirty Bomb Attack Action Graph

The simulated mission and activity execution works as follows. The resource requirements attributes of tasks indicate the resources that RED actors must bring to facility. For example, in order to assemble the bomb RED actors must possess the knowledge and materials for bomb making. As the result, to execute

the activity, RED team selects the actors that have corresponding capabilities. The facility requirements indicate the resources that facilities should have for successful task completion. For example, to assemble the weapon the size of facility should be at least 1 and technology must be available for assembly. Thus, to execute the activity, the RED team selects the facility that matches task facility requirements with its resource capabilities.

Similarly to facilities and activities, the roads and task relationships were defined and their attributes specified. The details are available upon request.

Results

We conducted several analysis of the sensitivity of the FINAL predictions of RED mission plan and allocation of activities to facilities to the level of noise in the data and types of parameters. We have used the simulator to generate five (5) different random mission execution allocations for the RED mission (different mappings). We then conducted 10 random runs for each action-facility allocation generated by the simulator, which resulted in different levels of the noise. The total of 50 thus obtained samples was averaged for each noise parameter set to generate “average accuracy” measure equal to the number of action-facility allocations corrected identified by the FINAL algorithm.

Analysis 1: Effects of Missing Intelligence Data

We ran the FINAL algorithm before the mission and during the mission executed by RED. This gave us six (6) sample points against the progress of the mission. This progress is equivalent to the level of missing dynamic data, and we thus were able to assess the sensitivity of FINAL’s accuracy in predicting the facility-to-action allocation to the amounts of missing event data. In this particular example, all data that were detected contained no noise.

From the results shown in Figure 12, we see that FINAL is able to achieve >66% recognition accuracy under 80% missing data. We also noticed that the ability to develop multiple solutions (solution choice (b)) provides significant (20%-40%) improvement in accuracy as compared to having a single solution at the higher missing data levels, but the effect is reduced as the amount of obtained intelligence increases. We also note that the algorithm is not perfect even when all the data are available. The reason for this is that the mission execution generated by the simulation component of our testbed developed solutions that were not optimal in terms of utility of mission execution. Since FINAL algorithm is using optimization-based matching, in some instances it tends to value the intelligence about the capabilities of facilities higher than the intelligence of their observed use.

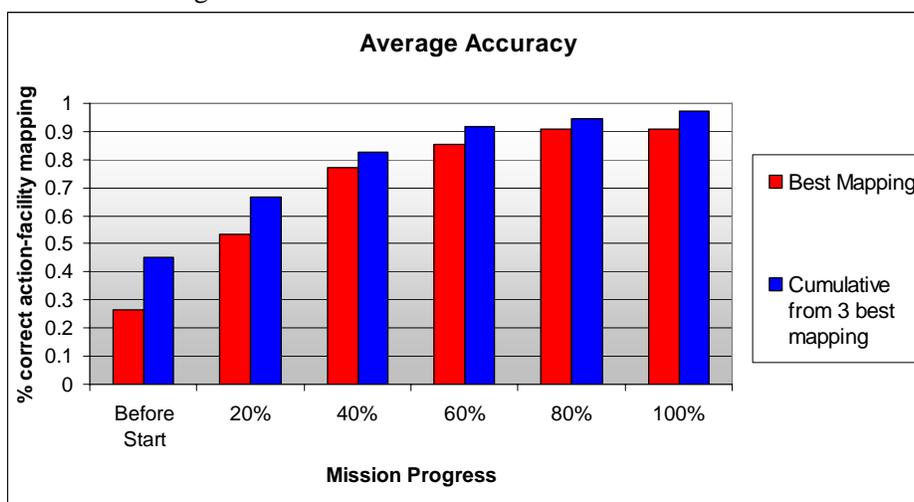


Figure 12: Sensitivity to Missing Event Data

We conclude that FINAL provides the ability to predict the potential uses of facilities when only scarce information is available about what RED are doing, and that the effects of different attributes needs to be studied in more detail.

Analysis 2: Effect of Missing Attributes

After studying the impact of missing event data, we decided to analyze an impact of detecting the event but missing some of its attributes/features. Since most of the attributes in our defined use-case were essential, missing the attributes would have a significant impact on our ability to map the facility to an activity.

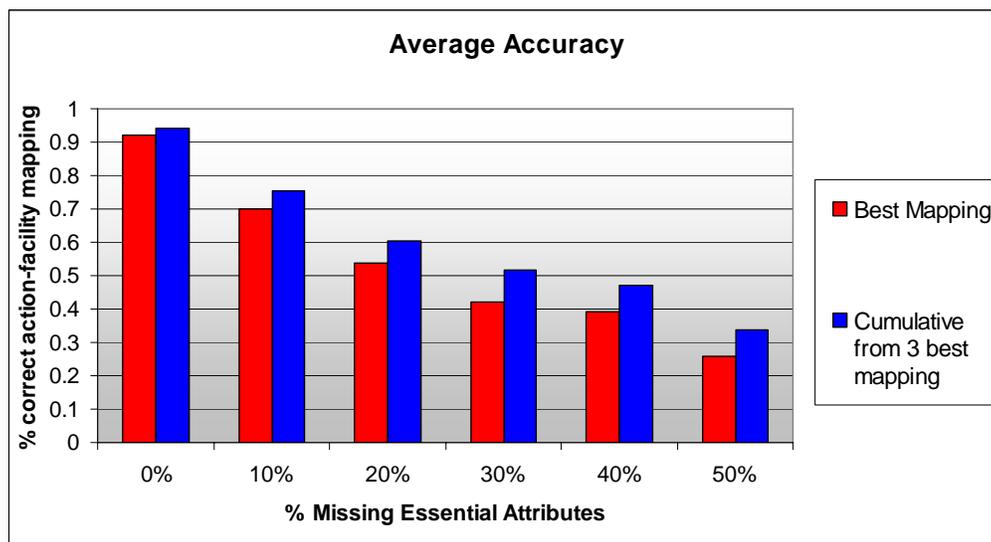


Figure 13: Sensitivity to Missing Attributes

We notice (Figure 13) that the impact of missing attributes is more pronounced than the impact of missing events. The reason for this is due to the penalties for missing key attributes and the type of information that is affected: while missing events affected only facility use attributes, current analysis also changes the capability attribute vector by missing the attributes of this set. As the result, the data get more skewed and cannot be matched well to the activities.

Analysis 3: Importance of Attributes

From the second study, we learned that the presence of some attributes might affect differently the accuracy of the algorithm. We thus ran a similar analysis looking at the accuracy in cases we restrict the use (and corresponding penalties) of the attributes. We compared (Figure 14) three cases: (i) using only structural attributes; (ii) using only current events/intel about uses of facilities; and (iii) accounting for all attributes.

From Figure 14 we can see that the algorithm accuracy can be improved if we rely solely on the event attributes as the number of essential attributes missing is increased. This effect was mainly due to large penalties that were introduced for missing facility capability attributes. We note however that we should expect that the noise in the capability data sources would be smaller than in online event/intel collection sources. We then hypothesize that reliance on the data elements to predict facility employment must depend on quality of collected intelligence. If such quality can be estimated, we could develop solutions customizable to the situation. This hypothesis is confirmed by running the predictions against the data of varied quality as shown in Figure 15. We see that different attributes uses had varied success at different data noise levels and none was the best for all situations.

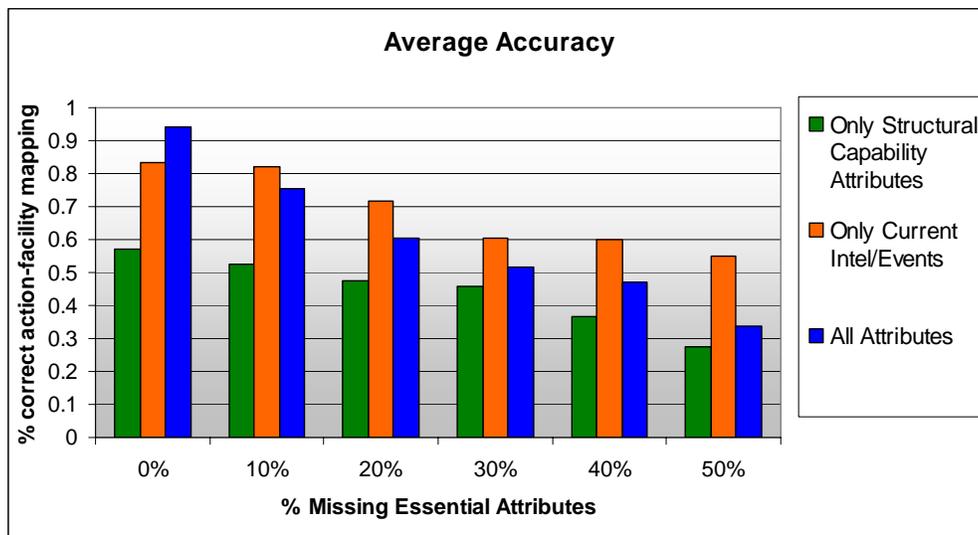


Figure 14: Importance of Attributes

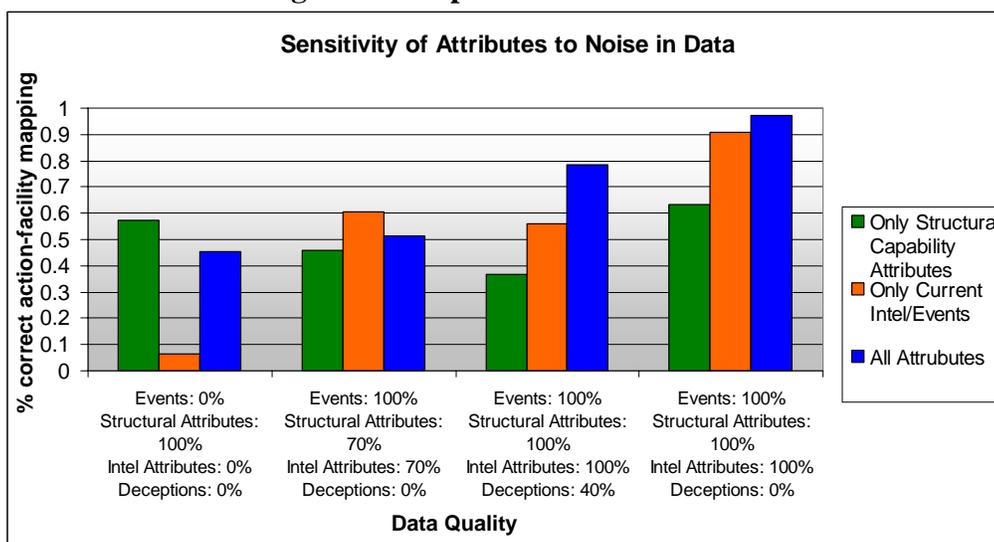


Figure 15: Trade-off in Relying on Attributes for different Data Quality Levels

The reliance on attribute types can be controlled via penalties/weights in the objective function. These weights should thus be selected based on data quality for best performance. If however a technology is developed that can fill the gaps under high missing data conditions, the improvements to the prediction accuracy could be significant.

Analysis 4: Effect of Deceptions

In this study, we increased the probability of deceptive (false) attributes occurring in the data from 0 to 50% with the value of such attributes in the range [-2,2]. As shown in Figure 16, we achieved >70% accuracy under 50% probability of deception. We see that there is a slower decrease in accuracy with decrease in data quality from deceptions due to anchoring the solutions in both static (facilities capabilities) and dynamic (facility use events) information. With introduction of deceptive attributes the pattern in the data was preserved. We conclude that when amount of irrelevant information increases, we can maintain high accuracy by accounting for static and dynamic data elements.

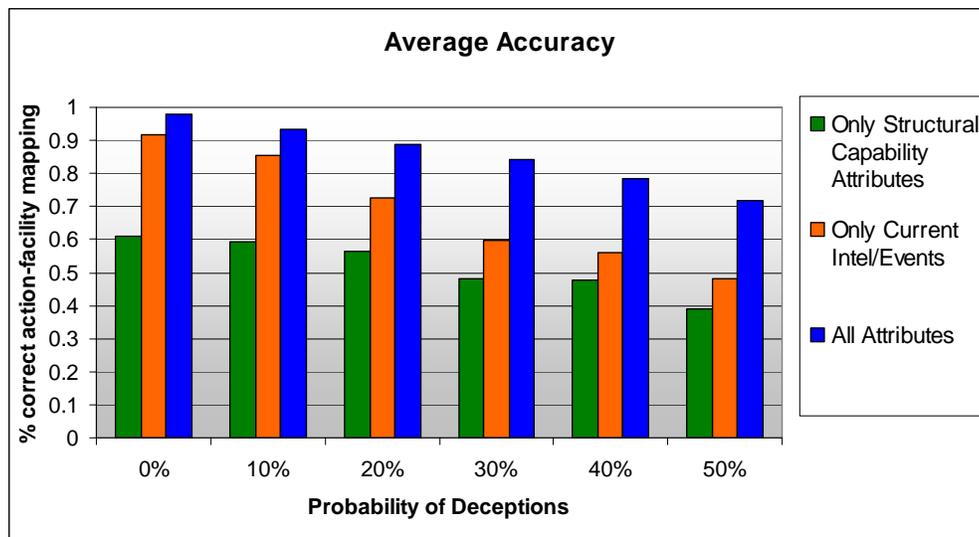


Figure 16: Effects of Deceptions and False Alarms

Future Extension to Guided Intelligence Collection

Due to limited sensors and intelligence collection capabilities, much of the desired data are not initially available for the analysis. A challenge, therefore, is to **identify data elements most critical to facility use prediction**. This can be achieved by assessing which attributes of the facilities and their relationships could help improve the accuracy of identification. In other words, we need to know which data elements can most significantly improve the current information and bring about the largest dichotomy between hypotheses. Choosing a single hypothesis out of the given set can then be done with higher confidence.

The automated intelligence collection planning will provide increased accuracy and confidence of facility use predictions. This approach is unique in that it integrates the prediction with data collection: the data collection is targeting the most critical missing data, which allows the largest increase in the accuracy and confidence of the prediction.

The goal of intelligence collection is to *find information* that would improve understanding of facility employment – or, in other words, **information that would improve accuracy of action-to-facility matching**. The outcome of the intelligence collection planning is the design of the set of actions that need to be performed over time.

In intelligence collection planning, we need first to recognize the differences between the hypotheses of activity patterns in terms of the activity attributes. We would like to make three important points about capturing these attributes and their effects on identification.

The missing data: Current observations provide the data about some attributes of the facility employment, while information about other attributes is missing. Intelligence collection planning is aimed at discovering these missing attributes. Without loss of generality we can assume for now that the current observations have already been accounted for and the set of action networks that need to be explored is reduced according to the a-posteriori threshold. The missing data can contain the *structural information* about the facility, as well as the dynamic *operational intelligence* reports about the current and previous facility uses. The objective is to find such information that is most critical to improve current predictions, determine the reconnaissance actions and their cost, and construct a plan from them. While we do not know the outcomes of the data collection actions, we certainly can know and account for potential outcomes since the actions will be specific at collecting such intelligence.

Distinguishing activity mappings: Observations about facilities reinforce some activity-to-facility associations, while discarding others. For a case of single model network, we can look at different mappings that score the highest and determine the features that would have the most impact on distinguishing the mapping with new data.

Distinguishing activity patterns: Observations also provide the ability to link/relate the model network activity patterns to one another. This is obtained via mapping the actions in the hypotheses to the facilities. The mapping of actions from different activity patterns to the same facility anchors these actions and allows exploring for their differences. If actions are of the same class, then conducting reconnaissance at the mapped facility would not provide any benefit. If however the actions are of different types, then finding more information at the facility may indicate a stronger matching for one action than another, and in this case lower the feasibility of one activity pattern and increase the feasibility of another.

As an example, the information about the entrances to a certain building might not be known. This information will not provide an improvement to the current identification in case the building is mapped to the activities which are similar, or for which the entrance information is irrelevant. For instance, if building is mapped to be a coffee shop, or a book store, or a diner, then obtaining information about the building entrances would not reinforce either of these mappings. However, if the building is also mapped to storing of bomb materials, then the knowledge that the building has a wide concealed entrance will reinforce the association with this activity and reduce the associations likelihood that it is a coffee shop, a diner, or a book store.

Clearly, in this example we are looking for the nodes and links in the action/function pattern that are different and the information that distinguishes them that have not been collected before. The nodes and links are explored for such difference only if they are mapped to the same facility or connection (a road) in the facility network.

The Intelligence Collection Planning is formulated using the following four steps (see Figure 17):

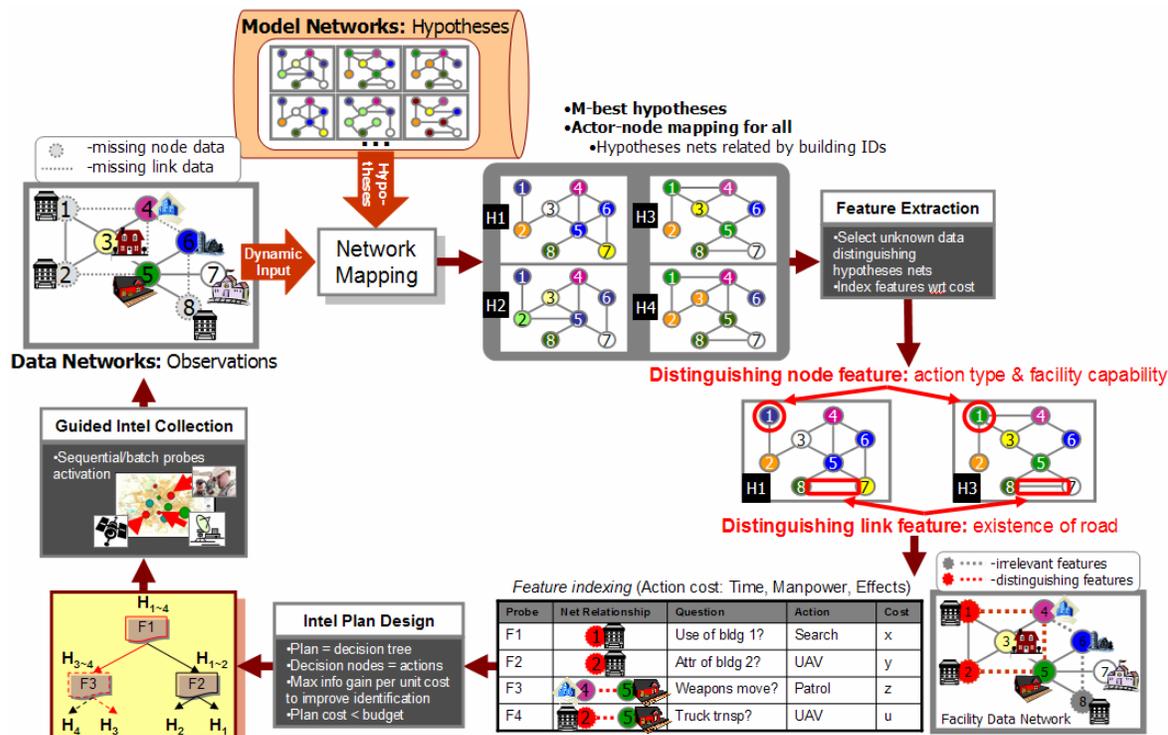


Figure 17: Intelligence Collection Planning Process

Step 1 (Network Mapping Output): First, the data network derived from observations is matched against all hypothesis networks from the libraries. The most probable hypotheses, which have produced similar rankings, can then be further explored. From the network mapping, each hypothesis node is mapped to a data network node, or in other words a specific facility ID. This provides an ability to relate hypotheses to each other.

Step 2 (Feature Extraction): Second, we compare the hypothesis networks with the current data network to see which of the potential data elements are missing. The resulting data elements, which could include missing attributes of nodes and links, are then classified as relevant (ones that distinguish the hypotheses) and irrelevant. The latter subset is discarded, and the former is categorized according to the action that needs to be performed to collect the data element, the cost of this action (e.g., in terms of time, money, effects of the action such as making the enemy suspicious of our presence, etc.), the question it is supposed to answer, and the corresponding network relationship that it will resolve. Each action-data element is then called a “probe” to indicate that it is trying to explore the existence of a specific attribute and its value through active search in the environment.

Step 3 (Intel Plan Design): In FINAL, an intelligence collection plan is represented in the form of a decision tree, where each node corresponds to the probe to be conducted, and each leaf corresponds to the next probe and is selected based on the outcomes of the previous probe. Organization of the probes into a data collection tree is based on maximizing the anticipated information gain which contributes to the ability to distinguish among hypotheses (Levchuk et al., 2007). The features can then be merged together for integrated intelligence collection actions (e.g., a patrol can spot the activities at buildings and identify the specific structural characteristics of the buildings on its path).

Step 4 (Guided Intel Collection): Guided intelligence collection is performed by following the decision tree nodes. Each node defines the action that needs to be performed, and based on the outcome of the action the process moves to the next node in the decision tree.

Conclusion

Our simulation studies showed that the FINAL technology can achieve high recognition accuracy in the presence of significant noise due to missing events, attributes, errors, and deceptive observations. FINAL has been tested on a small-scale example to show the feasibility, and approaches to improve the accuracy of algorithms and the ability to handle large data sets have been identified. In future research and development activities, we plan to focus on addressing the challenges of data set size and quality and develop a fully functional prototype solution with the objective to integrate into the real system of systems currently supporting adversarial analyses.

References

- Brantingham, P., and P. Brantingham. "Criminality of Place: Crime Generators and Crime Attractors". *European Journal on Criminal Policy and Research*, 3(3):1-26, 1995.
- Levchuk, G.M., and K. Chopra, “NetSTAR: Identification of Network Structure, Tasks, Activities, and Roles from Communications”, *Proceedings of the 10th International Command and Control Research and Technology Symposium*, McLean, VA, June, 2005.
- Levchuk, G., Levchuk, Y., and Pattipati, K. “Identifying Command, Control and Communication Networks from Interactions and Activities Observations”, *Proceedings of the Command and Control Research and Technology Symposium*, San Diego, CA, 2006.
- Levchuk., G., F. Yu, H. Tu, K. Pattipati, Y. Levchuk, and E. Entin, “Identifying the Enemy – Part I: Automated Network Identification Model”, *Proceedings of the Command and Control Research and Technology Symposium*, New Port, RI, 2007.
- Levchuk, G., Yu, F., Meirina, C., Singh, S., Levchuk, Y., Pattipati, K., Willett, P., and Kelton, K. “Learning from the Enemy: Approaches to Identifying and Modeling the Hidden Enemy

13th ICCRTS-2008 “C2 for Complex Endeavors”

Organization”, in A. Kott (ed), *Information Warfare and Organizational Decision-making*, Artech House, Norwood, MA, 2007.

Rangarajan, A., A. Yuille, and E. Mjolsness “Convergence Properties of the Softassign Quadratic Assignment Algorithm”, *Neural Computation*, 11(6):1455-1474, 1999.

Acknowledgement

This work was supported in part by the Office of Naval Research under contract N00014-07-M-0200.