Building Social Networks in Persistent Video Surveillance

Daniel T. Schmitt
Stuart H. Kurkowski
Michael J. Mendenhall, Member, IEEE
Graduate School of Engineering and Management
Air Force Institute of Technology
Wright-Patterson AFB, Ohio 45433
Email: [daniel.schmitt, stuart.kurkowski, michael.mendenhall] @afit.edu

***Abstract*—Social networks are a beneficial analysis tool in counterterrorism and counterinsurgent activities. The difficulty lies in the amount of time and resources it takes to construct a social network. By exploiting existing 24-hour overhead persistent video, we can build a social network from vehicle to vehicle and vehicle to building interactions. This paper demonstrates building a social network from vehicle tracks based on their interactions in an urban environment. From this social network we can see relationships among actors and their locations of interest. This information provides additional intelligence about terrorist activities to exploit them.**

I. Introduction

Social networks are a beneficial analysis tool in counterterrorism, surveillance, and counter-surveillance activities. Specific successes applying social networks to counterinsurgency and terrorism include the search and capture of Saddam Hussein [1] and the discovery and arrest of a Canadian terrorist cell in 2006 [2]. The shortfall of applying social networks to counterterrorism is that social networks often take considerable time and resources to construct. An automated system that can process persistent video from a 24-hour overhead source, track vehicles to generate vehicle tracks, and relate the vehicle tracks to each other would construct a social network to free analysts for other critical tasks. In addition, such a system could further allow for the constructing of social networks on individuals before they are identified as suspicious, saving the lead time required to construct them. In this way, if a suspicious individual is identified, an intelligence analyst can proceed immediately with checking who the individual has interacted with and where they have been operating based on the persistent video and the social network constructed by the system.

II. Problem Definition

The problem addressed in this article is the extraction of the inherent information contained in the interactions of Global Positioning System (GPS) tracks and the environment as represented by a Geographical Information System (GIS) (like roads, buildings, etc.). The product is a visualization of the interactions between tracks and the environment where the visualization is a social network.

The approach is to use parameters as defined in Section IV to define interactions between vehicles and buildings. These interactions form the basis for populating links from vehicles to vehicles and vehicles to buildings in a social network. This network will then be displayed and evaluated for its value in understanding the activities that occurred.

III. Workload

The workload used was persistent video data over an urban area. The video is post-processed to detect moving vehicles denoted as tracks (vehicular tracking in and of itself is a separate research effort outside the scope of this paper). Tracks are listed in a common coordinate system to relate to other static information from a GIS. To focus on the social network aspect of the research, GPS track data from the same area was used to represent vehicular track input to our model. The GPS data [3] is a good representative of the persistent video data because GPS data logs position coordinates at nearly the same time interval as persistent video images are captured (about 0.5 sec). Both sets of data are also not without fault, as both GPS data and tracks generated from persistent video can have areas of missed detections caused by system errors or occlusions.

The entirety of the GPS data used totaled 14 different vehicles and 25.5 hours of total driving time captured by vehicle-born waypoint devices. In addition, it is known that of the 14 different tracks generated, some vehicles worked together in completing assigned scenarios, while other vehicles did not interact for the purpose of creating random driving. While the specific scenarios are not known, the workload includes both elements of known interaction and elements of known non-interaction that can be investigated and verified after the social network is constructed.

IV. Parameters and Definitions

In order to build interactions between vehicles and between a vehicle and a building, we must first define the parameters

ISI 2009, June 8-11, 2009, Richardson, TX, USA

that will be used in time and space to describe these interactions. The following sections define the necessary length of time and distance for these links.

A. *Definition of a stop*

Defining a stop is the first necessary step. A stop was defined when a track stayed within a six meter radius for over 30 seconds. The six meter distance allowed some error in tracking, but no more error than the approximate length of a vehicle. Thirty seconds was an arbitrary choice that could be adjusted depending on the intended use of the social network and the vehicle behavior that is desired to be detected (i.e. stop signs, stoplights, traffic jams, or parked cars). These parameters were tested with the urban data and may need to be tweaked to apply to a different environment.

B. *Defining vehicle to building links*

Vehicle to building links were constructed whenever a vehicle stopped in front of a building. There were two approaches to determine which building is being stopped in front of: a Nearest Neighbor approach, and a distance approach.

1) Nearest Neighbor Approach: A Nearest Neighbor approach takes a stopped track, and finds the building that is closest to that point. This approach is the most straight forward approach. It is flexible for different areas (those that have lots of buildings, and those that have few), and reduces each stop to only one building. This is helpful in keeping the social network uncluttered with unnecessary and redundant information.

2) Distance Approach: A distance approach takes a track, and finds all buildings within a certain distance of where that track stopped. While this is computationally more efficient than the Nearest Neighbor approach, it is very dependent on a consistent distance between stop points and buildings. A distance threshold set too low returns no building to vehicle links. A distance threshold set too high returns several buildings per stop, increasing the clutter within the social network.

In constructing vehicle to building links, we used a Nearest Neighbor approach for stops that were longer than five minutes. This choice of duration is arbitrary and can be increased or decreased depending on the granularity of results required. The choice of using a Nearest Neighbor approach is deliberate to include one building per stop, and not over clutter a network with several buildings per stop.

C. *Defining vehicle to vehicle links*

Vehicle to vehicle links were constructed when two vehicles stopped for longer than five minutes within 100 meters of each other. The choice of time again was arbitrary and can be increased or decreased depending on the granularity of results required. Setting the distance to 100m was used as an estimate of a typical parking lot size in an urban area. When data sets get large, this distance could be decreased to reduce computation time required to evaluate a set of results, at the sacrifice of detecting the desired interactions between vehicles.

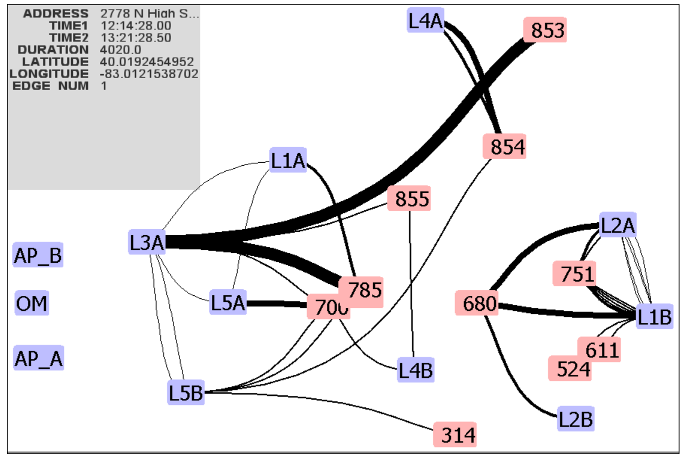

Fig. 1. Social graph with buildings organized in geo-space

V. Results

The results of constructing the social network using the Prefuse visualization package [4] for the GPS tracks is shown in Fig. 1. Each named node represents a vehicle track (blue), each numbered node represents a building (pink), and each link represents an interaction between vehicles and between vehicles and buildings that was at least five minutes long. Those links that are more heavily weighted represent durations that were greater than 15 min, 30 min, and 60 min. Multiple links represent multiple interactions. When a link is selected, the information about that interaction is displayed in the upper left corner. The buildings are organized as they would be in geo-space in the same relative locations to one another as scaled by the size of the display window showing visually the distance proximity of buildings from one another. When a node is selected, an additional window is opened (see Fig. 2). This window is interactive and allows a user to search for an actor or building by name. It also animates and realigns the graph around any node that is selected. This view is beneficial if there is interest in a certain individual to see what connections exist for that particular individual.

VI. Analysis

From the social network output we can deduce information not known prior to the social graph construction, but verified afterwards. For example, the area around building 785 and 700 (center of Fig. 1) was an area of interest for drivers *L3A*, *L1A*, *L5A*, and *L5B* who all interacted with at least one partner, and also all visited this same location. Relating to the map, the area that this refers to is a parking lot on the northeast side of the campus football stadium. This parking lot was used as a coordination meeting place for the four actors prior to their activities.

The other interconnected group represented by vehicles *L2A*, *L1B*, and *L2B* (right side of Fig. 1) coordinated a separate activity. Specifically they coordinated a staged event with *L2A* and *L1B*. What is also picked up in the social network are two primary buildings of interaction, buildings

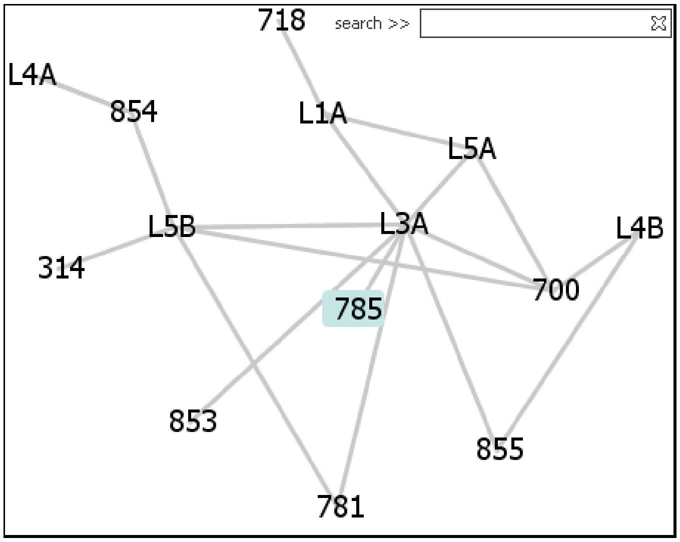

Fig. 2. Radial view of social network graph

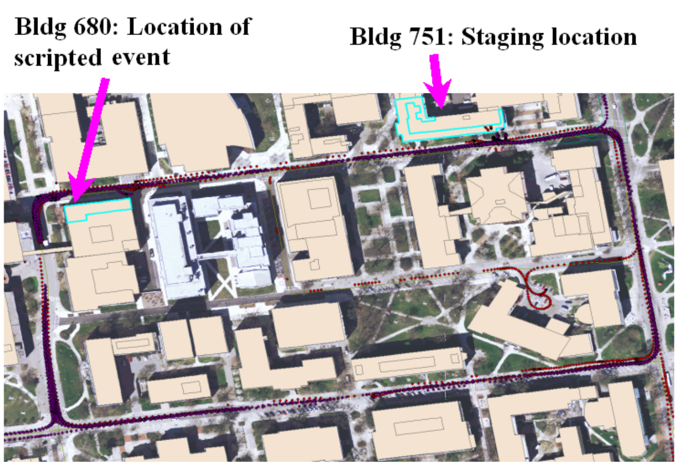

Fig. 3. Area of interest during scripted event

751 and 680. In Fig. 3 it shows the GIS view of this area. After the social network was constructed it was verified that building 751 was a building of interest used as a staging area for an event that occurred in the area of building 680.

Referring back to the social graph in Fig. 1 we are able to see that three tracks were designated as independent random driving tracks (*AP_A*, *AP_B*, and *OM*) that did in fact not interact with anyone. The fact that links were made between actors when there was interaction, and no links were made with actors that did not interact show that the social network is able to relate vehicle to vehicle interactions and vehicle to building interactions effectively providing additional context to an unfolding situation.

Since other data sets are likely to be much larger than the data set used here, additional enhancements were made to the social network viewer. We added the ability to filter out buildings and building interactions in order to focus primarily on vehicle to vehicle interactions. These filters are interactive and can be turned on or off by the user.

VII. Conclusion

In this paper, we presented results on building a social network from vehicle tracks in an urban environment. We showed that the resulting output was useful in effectively relating actors together, and relating them to buildings of interest. By processing the information and displaying it in this manner we were able to determine which actors were acting in concert with one another, and where they were working. Specifically we were able to determine which actors were random actors (no interaction), which actors were working together on scripted events, and the staging location for the event as well as the event location. This model is also open to tailoring for specific areas and interactions as defined by time and distance thresholds. The application of this technique is beneficial in an environment where there is a large amount of collected persistent video. This work can then focus on relating terrorist actors together in a time sensitive manner as is the case in Iraq and Afghanistan enabling analysts to track interactions between buildings, vehicles, and groups of vehicles. The end result is a time savings to the analyst that before was focused on building a social network and now can focus on exploiting an automatically generated social network.

If more data was collected over a series of days, patterns of life will begin to emerge both on an individual (their workplace, home, and other locations) and on an aggregate level (areas that are frequented by many people, and areas that are rarely visited). On an individual level, we can apply social network analysis techniques to learn if groups are organized hierarchical or not by using the duration of interactions or number of interactions as weighted values of the social network links. On an aggregate level, we can use the information to determine what locations are normal and abnormal to be visited by a population. This would allow us to define behaviors that may appear abnormal by their route or location giving us more information to present to intelligence analysts.

Acknowledgment

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government.

References

[1] "Sociological skills used in the capture of Saddam Hussein," 11/18/2008. [Online]. Available: http://www2.asanet.org/footnotes/julyaugust05/fn3.html
[2] "Connecting the dots – social network analysis of 9-11 terror network," 11/18/2008. [Online]. Available: http://www.orgnet.com/prevent.html
[3] S. D. Air Force Research Laboratory, "Osu gps data," 28 Oct 2007.
[4] "prefuse — interactive information visualization toolkit," 12/8/2008. [Online]. Available: http://prefuse.org/