Spatial Profiling with Adversarial Process Modeling

Alex Reeve, Dan Fu
Stottler Henke Associates
San Mateo, CA
areeve@stottlerhenke.com, fu@stottlerhenke.com

ABSTRACT

Geographic profiling (GP) techniques for crime analysis have proven useful for identifying the locations where serial killers dwell. In this paper we examine the application of geographic profiling techniques to an organized group of individuals, such as drug dealers and insurgents; in particular, tackling the problem of predicting which facilities in an urban area might support clandestine activities such as drug processing or bomb making. GP techniques assume a single perpetrator whose only observable actions are punctuated killings. In contrast, clandestine organizations involve several distributed individuals who communicate, coordinate, make plans, and execute. Most of their actions, potentially observable such as phone calls, are seemingly innocuous. Through the use of a simulated intelligence stream, we combine GP techniques with plan recognition technology. We advocate a recognition approach which exploits a wide range of knowledge about the group, including the methods of operation, preferences, constraints, and relationships with other like-minded groups. In turn, GP techniques can be augmented with more sophisticated distance metrics using derived geo-spatial attributes, such as cost-of-travel and perceived route risk. We then discuss approaches to fuse all information into a predictive model for each group. This model estimates the risk of future activity based on current observations of group presence. This estimated risk is used to generate actionable products such as security force search paths and prioritization of intelligence collection requests. Finally, we evaluate accuracy of the approach in the presence of noise and incomplete data.

ABOUT THE AUTHORS

Alex Reeve an AI Software Engineer at Stottler Henke Associates, Inc. Mr. Reeve received his Bachelor's degrees in Mathematics and Computer Science from Cornell University in May 2005. He has a background in Artificial Intelligence, including Machine Learning, Artificial Neural Networks, and Expert Systems. His research interests include data mining and geospatial analysis.

Dan Fu is a group manager at Stottler Henke Associates. He joined nine years ago and has worked on several artificial intelligence (AI) systems including AI authoring tools, wargaming toolsets, immersive training systems, and AI for simulations. Dan was principal investigator on an intelligent agents project to create AI middleware for simulations and videogames. The result was SimBionic®, which enables users to graphically author entity behavior for a simulation or videogame. Dan holds a B.S. from Cornell Univ. and a Ph.D. from the Univ. of Chicago, both in computer science.

Spatial Profiling with Adversarial Process Modeling

Alex Reeve, Dan Fu
Stottler Henke Associates
San Mateo, CA
areeve@stottlerhenke.com, fu@stottlerhenke.com

GEOGRAPHIC PROFILING

The development and widespread acceptance of computerized geographical information systems (GIS) have opened the doors to advanced, data-rich spatial analysis algorithms. This analysis is no longer limited by a lack of data. Instead, the advancement of analytical algorithms now relies on the ability to isolate relevant data, detect key patterns and fuse multiple sources of knowledge into a single coherent model. Of particular interest is the spatial analysis of historical data to detect key locations or patterns of a particular activity or process. These patterns may be described by proximity, with events centered around a specific location, or by attributes of the locations themselves.

Ideally the patterns will be useful for predicting the future behavior of an adversarial agent. Patterns must be able to handle three aspects of real-world behavior: activities may only be partially observable, the exact intent unknown, and actors may be purposefully modify their behavior to escape detection. A classic example of such analysis is the search for serial offenders in criminal investigation.

The methods of Geographic Profiling (GP) show promise for predicting key locations and patterns based on sets of incidents. GP is based on the field of Environmental Criminology (EC) [Brantingham and Brantingham 1990], which focuses on criminal spatial behavior, or how crime relates to particular places. It attempts, for instance, to explain the spatial distribution of offenders and offences. *Activity space*, the area of mobility and familiarity for a criminal offender, is a primary concept. GP derives "base of operations" estimates from serial crime data, primarily through the use of distance decay functions, which characterize the offender's hunting pattern in terms of its distance from residence. It is important to note that "the subjective psychological perception of distance is just as critical as the objective physical space involved" [Rossmo 1999]. That is, these measurements of distance are not purely geographic, but capture the offender's behavior with respect to geographic features, which may present particular attractions or barriers to particular offenders. These are special distances, specific to the offender's process.

The famous primeval example of GP is John Snow's death map, produced during the 1854 London Cholera outbreak [Johnson 2006]. It showed all deaths from cholera on a city map, which turned out to cluster around a single point of origin. At that point stood a public water pump, from which tainted water was being drawn by the city's residents. When that pump was disabled, the outbreak abated. This map was not the entire story, however. Prior to the outbreak, John Snow already had a waterborne model for the transmission of cholera in mind, one which countered the prevailing view of airborne transmission, and which suggested a particular process model based on his knowledge of the city's workings: sewage, draining into river water, distributed to certain tanks by certain water companies, and finally provided to individuals through public pumps. His map led him merely to a city block; from there, the process model led him to the point of intervention, which was a bit of infrastructure crucial to the process: the guilty pump.

Limitations of GP

Current GP models incorporate some infrastructural knowledge, in buffer zones around the offender's residence, within which offenders tend not to act, and catchment areas for crime locations, which are the areas surrounding certain kinds of buildings, such as schools, where the offender may prefer to operate. These produce, in effect, simple spatial preference functions for the serial criminal offender, when leaving home to hunt for a victim.

While GP has had its successes and shows promise for future development, it also possesses significant known limitations. For such profiling to be effective, the data must satisfy four distinct characteristics [Rossmo 2005]:

1. The case involves a series of at least five events
2. The agent has a single stable anchor point

3. The agent is using an understood method
4. The relevant geographical area (the "backcloth") is reasonably uniform

In contrast, clandestine organizations involve several distributed individuals who communicate, coordinate, make plans, and execute. Any bare-bones GP-based approach is insufficient to detect distributed activity as it fails the second constraint of a single anchor-point. While related activities may still have strong geographical centers, the resulting jeopardy surface within a geo-profile will be skewed towards the center of the relevant anchor-points. A more advanced approach must be able to take into account the multiple anchor-points of a spatially separated organization, and to predict which locations are the most spatially relevant.

The constraint of a near-uniform backcloth also fails in a large number of domains, again reducing the effectiveness of traditional GP-based approaches. In urban environments, locations differ dramatically in terms of accessibility, security, population, and other numerous attributes which may be tied intimately into agent activity. Geographical features, too, tend to negate the uniformity constraint with the presence of water bodies and impassable terrain. Ideally, the non-uniformity of the backcloth would be exploited, rather than ignored, to facilitate detection of key patterns of activity in the geo-spatial data.

ADVERSARIAL MODELING

The proposed approach models agents through *preference functions*. These preference functions maintain and analyze characteristics of the adversarial individual or group and the complex process that results in activity. Preference functions also exploit spatial characteristics that support individual parts of that process, and impose constraints in time, space, and resources. The preference knowledge base is built on actor, event, and geographical ontologies for the representation and aggregation of information at multiple levels of specificity. This results in a multidimensional characterization of the adversarial event, as the end result of a specific process by a specific actor or group, or type of actor. Verma and Lodha note that the adversarial event has five dimensions, including space, time, law, offender, and target or victim [2002], and that all of these should be analyzed interdependently.

Preference modeling operates on the assumption that certain attributes of a location positively or negatively impact its selection for use by a given agent. If an adversary's location selection criteria can be learned from past activity, they can then be used to predict future locations by that adversary. Any available geographical data can be exploited to construct these predictive models, drawing upon heterogeneous sources such as

- municipal records – zoning, building construction type, owner, assessed value, construction date, etc.
- geographical data – water bodies, temperature, elevation
- census data – known residents, residents' occupations, genders, etc.
- utility records – electric and water bills, phone numbers, etc.
- intelligence – neighborhood tribal/religious loyalties, etc.

This information can be stored and managed by the GIS as multiple data layers, and new layers can be added as new data sources become available. Each geographical location can thus be thought of as having a vector of associated attribute values

$<a_1, a_2, a_3, \ldots, a_n>$ (With values normalized to the range [0,1])

There are multiple ways to evaluate the significance of a series of locations, represented by attribute vectors. One approach to determining the significance of attributes is to calculate significant variations between the activity mean and the population mean. For each attribute a_i we calculate the standard score,

$$z_i = (x_i - \mu_i) / \sigma_i \quad (1)$$

where x is the activity mean, μ is the population mean, and σ is the population standard deviation. The sign of this score represents whether deviations are positively or negatively correlated with the occurrence of the given activity. The magnitude describes the strength of the correlation.

The score of each attribute is updated after every profiled event. To compute the profile, we calculate the likelihood of each point as the product of the coefficient vector and the deviation vector:

$$L(p) = z \cdot d, \text{ with } d_i = (x_{i,p} - \mu_i) / \sigma_i \quad (2)$$

where $x_{i,\mathbf{p}}$ represents the value of attribute a_i at point **p**. This value goes up when the deviation of the attribute

at point **p** matches the current correlation. **L(p)** essentially represents the likelihood of a future occurrence of the specific activity at the location **p**.

For example, a series of *Festival* events occur at a specified set of locations. Upon analysis, the *population* attribute seems to be statistically higher in the set of event locations when compared to the average population of the area. A stronger statistical deviation of a given attribute implies stronger predictive power. In this case, **L(p)** is higher in areas which contain a significantly higher population than average. Another attribute, *crime rate*, is seen to be generally somewhat lower than the area average. Thus, **L(p)** is also highest at points of low crime rate. The primary areas identified by the profiling algorithm are thus high population, low crime rate locations.

Activity/Process Modeling

Our approach is also based on the hypothesis that we can improve our predictive models by exploiting knowledge about the methodology of a specific activity. The second extension is the use of *process models* to capture specific adversary tactics, plans, or methodologies. A process model is essentially a *hierarchical task network* (HTN), which consists of a collection of events connected by links. Events can be concrete, directly observable actions (for example, a telephone call, posting of propaganda to a website, or a meeting between several individuals) or abstract occurrences (e.g., procuring drug supplies). Links between events denote temporal or spatial constraints between those events (e.g., event B must start after event A and occur within one mile of event C).

Because process models are hierarchical, the nodes within a process model can also refer to other process models. These referenced process models act as "sub plans" within the larger adversary plan. For instance, several drug dealing events might be nested within another process model representing a larger drug campaign by a given adversary. By allowing the analyst to break down complex adversary plans into smaller reusable chunks, HTNs yield a more concise and maintainable model of adversary behavior.

The hierarchical nature of process models has a secondary advantage. Events in the incoming event stream typically describe concrete adversary actions. It would be useful, however, to be able to abstract over those events to a "big-picture" view of an adversary's activity, to enable analysis of hostile activity at the level of plans and goals. With an HTN-based representation, low-level plans can naturally be aggregated into higher-level plans. In addition, all events referenced in the process model library are arranged in a type hierarchy (e.g., both "phone call" and "face-to-face meeting" events are subtypes of the abstract "contact" event). This enables generalization within process models, which can increase their coverage of adversary behavior. For example, a process model that contains an abstract "B contacts E" event will match both "B telephones E" and "B emails E" events. Note that the abstraction-over-activities capability provided by process models strongly complements the abstraction-over-actors capability provided by the actor ontology.

Every event in a process model can be further annotated with attributes that describe, among other things:

- The impact of a process model (e.g. a murder has a high impact, while a phone call has a low impact)
- Whether an event is optional or must occur in the process model
- Which attributes of an event must be present in order to match the process model (which allows the system to weed out "frivolous" matches that contain little or no usable information)
- The importance of each event in the process model (roughly its value as an indicator)
- Temporal constraints that must hold between events in the process model (e.g. event A must happen within 30 days of event B)
- Constraints on the actors, locations, and resources that participate in a process model
- Whether the system should attempt to predict a given event (to avoid predicting banal or insignificant events)
- Type constraints on attributes for events within a process model
- The estimated observability of a given event

The approach can be used to incrementally build evidence for detection of adversary activity. As events flow into the system from the event stream, relevant process models are matched. (Note that some of these process models may have already been partially matched against previous events.) The system then fleshes out these process models with any available attributes of the newly-matched event – for example, the location and time where it took place and the actors involved. (Any of these attributes may be missing.) Each of these fleshed-out process models represents a *hypothesis* by the system about the current activity of a given adversary. The confidence of a given hypothesis is a function of the number of matched nodes and the confidence of each matching event.

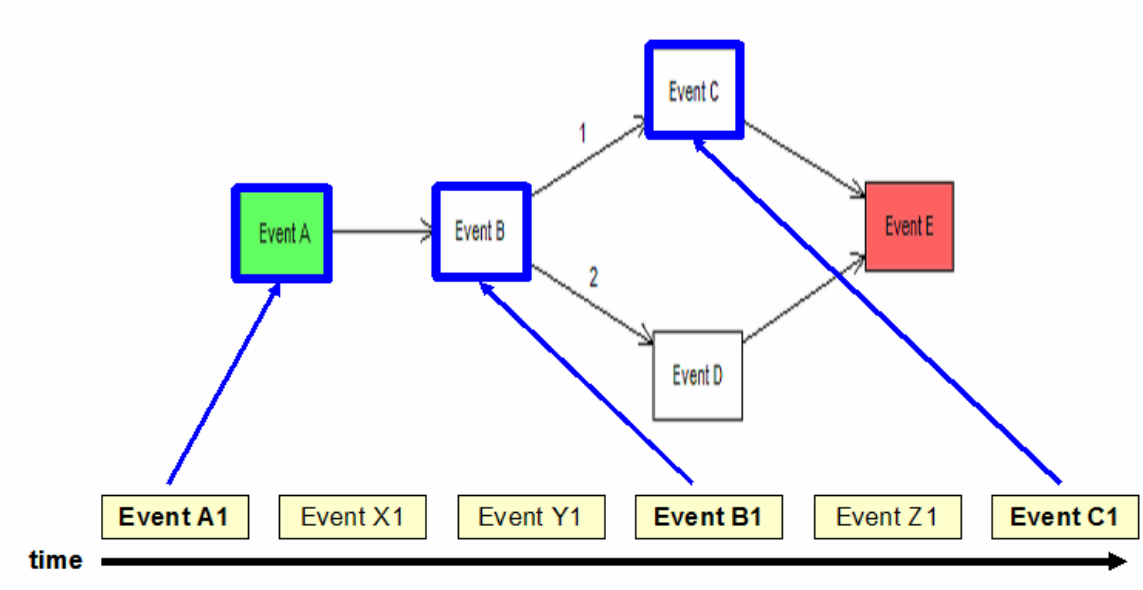

Figure 1. Matching incoming events against a process model

When a hypothesis achieves a certain confidence level, the system generates an abstract event representing that hypothesis and inserts it in the event stream. These abstract events can match against process models just like normal concrete events. This process of abstraction enables the analyst to focus on the high-level activities of adversaries rather than becoming overwhelmed by the tremendous number of low-level events coming in from the data stream.

SIMULATION

For the purposes of the evaluation, we used a simulation which mimics the features of real data streams, including noise, missing data, and corrupt data. It also simulates multiple collection sources with varying confidence and observability levels. The simulator is domain-independent, taking as input a "domain theory" that describes the processes and actors to be simulated. For the purposes of this effort, we developed a domain theory based on domestic drug trafficking, which possesses all of the characteristics necessary to test the detection of organized adversarial activity in the presence of noise and incomplete data. Our domain theory was based on criminology literature and features all elements of the domestic drug trafficking operation, including transport, stockpiling, retail, and inter-gang warfare. It also models the organizational structure of a drug distribution organization, from the "boss" to the street-level dealer. This domain theory, when fed through the simulator, produces a primary data set of 2192 events covering 12 months.

This event set, which represents the *observed* activity of the drug traffickers rather than ground truth, is streamed chronologically into the prototype to simulate a real-time event stream. It is important to note that the prototype does not have access to simulator ground truth (that is, what actually happened in the simulation run) nor to the underlying domain theory. It must base its predictions entirely on the distorted and noisy observed view of the adversary that it gets from the event stream. We also developed a library of process models representing the behavior of the target adversary – in this case the various members of a drug trafficking organization.

The simulation also consisted of a micro-GIS to provide geographic query capability to the prototype. Our GIS models 40,000 map locations in a single urban area. It has a grid-based road model that includes roads varying in size from alleys to highways to provide a simple model of accessibility. It also stores arbitrary attributes for each facility on the map. For the purposes of the prototype, we implemented 10 attributes:

1. property value
2. building size
3. zoning (retail, office, residential, industrial)
4. owned/rented
5. number of residents
6. average resident age

7. average resident income
8. accessibility to major roads
9. crime rate
10. police presence

The prototype can query the GIS for the attributes of any facility as well as for the distance (using various metrics) between two facilities. The prototype does not have access, of course, to the underlying preference functions that simulated actors use to select facilities. This preference model is learned automatically from the event stream.

Two separate and methodologically distinct activities were simulated within our model drug trafficking domain. Gang killings were a low-frequency activity most likely to occur in areas of overlapping gang territory, with minimal regard to spatial attributes. Drug deals represented a relatively high-frequency event with a preference for locations which combined both low police presence and low income (providing a strong weighting for geospatial preferences), within the bounds of gang territory (representing a softer distance constraint). Because the income and security attributes were significantly non-uniform, this activity tested the quality of the proposed approach in the absence of constraint (4). In addition, drug dealing agents were simulated individually, each with their own anchor point. Thus, the drug deal activity evaluated the approach in the presence of multiple anchor points.

To evaluate the detection of differences between adversaries, two separate organizations were simulated. Two gangs, named "Avon" and "Marlo", each possessed variations in preferences for the drug deal activities. The Avon gang weighted highly accessible areas, while the Marlo gang preferred lower accessibility. In addition to these differences in preference function, the gangs were given spatially separated "territories", which (as mentioned above) affected both activities.

RESULTS

We present here the definition of some terms that are commonly used in the discussion of the various predictive models.

Jeopardy Surface

The height of the profile's jeopardy surface represents the probability that the next event will occur at that point. We define P(**x, y**) as the height at position **(x, y)** and P(**event**) as the height at the location of **event**.

Miss Area

We define the search area for a particular event,

A_{miss}(**event**) = { **(x, y)** such that P(**x, y**) >= P(**event**) }

In other words, the miss area represents the set of points that were identified as being more likely than the location at which the event occurred. Note that this area relies only on the relative ordering of the values P(**x, y**), and not their magnitude.

Hit Rate

The hit rate for a particular event is defined as the ratio of the size of the miss area to the total size of the profile,

HitRate(**event**) = #A_{miss}(**event**) / N

In other words, the hit rate is the percentage of points with a greater or equal surface height than that of the event's location. The best hit rate is nearly 0 (every other point had a lower probability), while the worst is 1 (every other point had a higher probability). Note again that this relies only on relative ordering and not on magnitude.

The following sections summarize the predictive results of the prototype system. Note that hit rates are all calculated *before* the models interpret the data, so the following values accurately represent prediction accuracy.

Generated Profiles

In the following profiles, the red dots represent the locations of known activity events, while the green areas represent the points most likely for the next event to occur.

Figure 2 demonstrates the profiles generated for the *Sell Drug* activity over two separate actor groups. The group used on the first is *Anyone*, and covers all *Sell Drug* events. The profile on the right considers only events associated with the *Avon* gang. Note that the events in the latter are a subset of the events of the former. This activity is an *organized* activity.

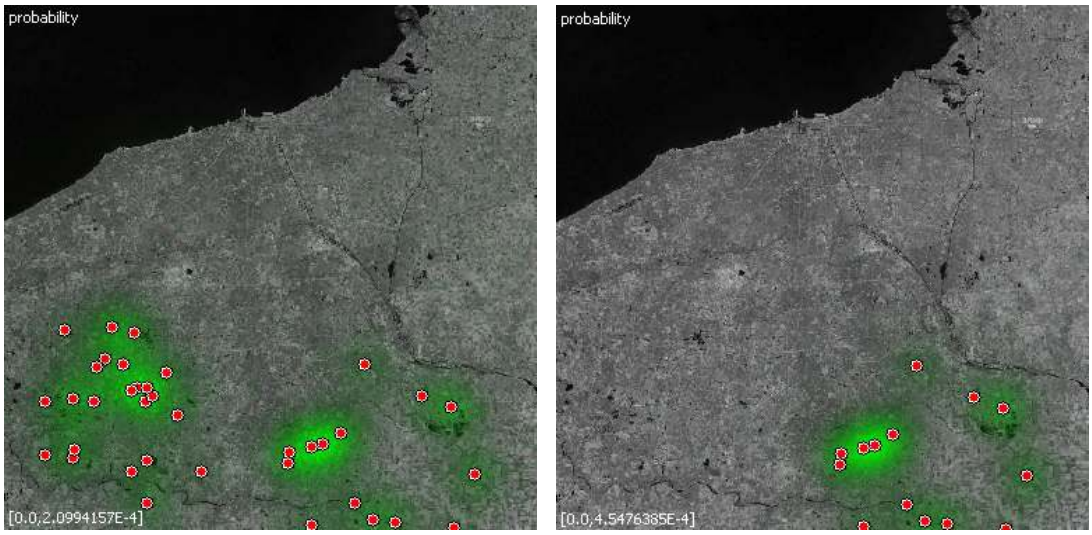

Figure 2: Profile for the Sell Drug activity, (left) over all actors, and (right) restricted to the Avon gang

Prediction Error

Figure 3 displays the hit rate for the profile ("Anyone", "Murder").

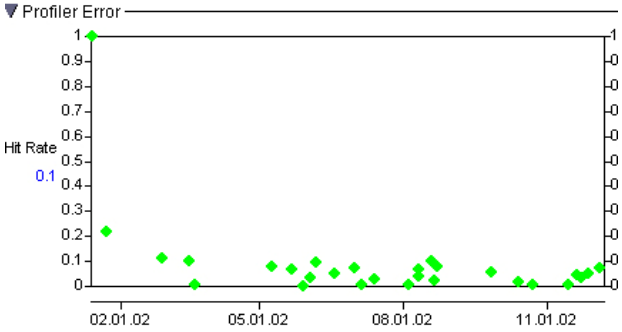

Figure 3: Prediction error for the Murder activity

Convergence to the best hit rate occurs rapidly. The steady hit rate likely indicates that the underlying actor preference model is not changing over time, or is changing only slowly. This convergence indicates that the given data and models are adequate to profile *Murder* activity.

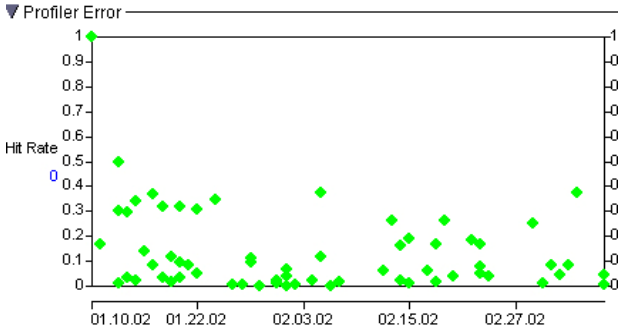

Figure 4: Prediction error for the Sell Drug activity over all actors

Figure 4 displays the hit rate for the profile ("Anyone", "Sell Drug"). In other words, the graph shows the error in answers to the question, "Where will the next *Sell Drug* event occur?" This profile never converges to a good hit rate, but its average hit rate hovers between .15 and .2. The variance in error is high, indicating that the given models are inadequate to predict this particular activity.

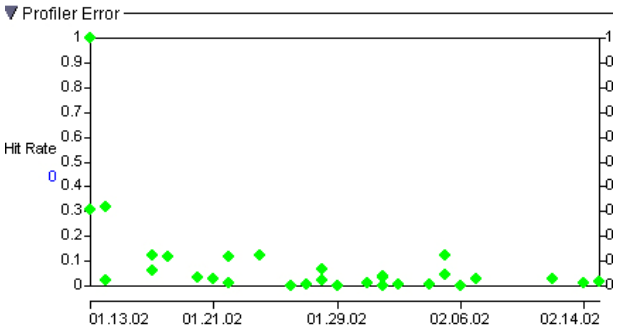

Figure 5: Prediction error for the Sell Drug activity within the Avon gang

Figure 5 displays the hit rate for the profile ("Avon", "Sell Drug"). This profile takes another look at the *Sell Drug* activity, but only profiles the set of events associated with a specific organization – the *Avon* gang. Thus, it displays the error in answers to the question, "Where will the next *Sell Drug* event associated with the *Avon* gang occur?"

Convergence in this case occurs within just four iterations. The dramatic increase in predictive quality over the previous profile illustrates the importance of organization to this activity. This could have several explanations. First, we can expect a minor increase in prediction error as we increase the scale of our profile. Secondly, it is apparent from the map in Figure 2 that there is a significant territorial bias in this activity. Finally, it is possible that Sell Drug attribute preferences within the Avon gang differ from those of other groups.

Effects of Positioning Error

In general, prediction error increases gradually with the scale of the organization being profiled. Naturally, error increases significantly if the organization contains actors with conflicting preferences. In this way, analysis of the effect of organizational-scale on prediction error can reveal which groups contain actors with similar preferences.

We also evaluate the effect of positioning error on two activities in our dataset. To relate scale, we define our profile area to be 10km x 10km. The leftmost graph shows the hit rates with no error. The middle graphs shows the effect of a 750m standard error and the rightmost graphs shows the effect of a 1.5km error.

Figure 6 shows the effect of positioning error on the *Murder* activity. Recall from the previous section that the *Murder* profiler draws most heavily from the attribute-preference model. As positioning error increases, higher-frequency attributes are "blurred." This reduces the overall effectiveness of the attribute model, and reduces predictive quality. While some level of accuracy still remains with a 750m error, the model is essentially unusable with a 1.5km error.

As shown in Figure 7, the ("Avon", "Sell Drug") profile is significantly more tolerant to positioning error. Though the average prediction error increases with positioning error, the increase is gradual. This is due to the fact that the main indicator is the proximity model, which tends to be a much lower frequency signal.

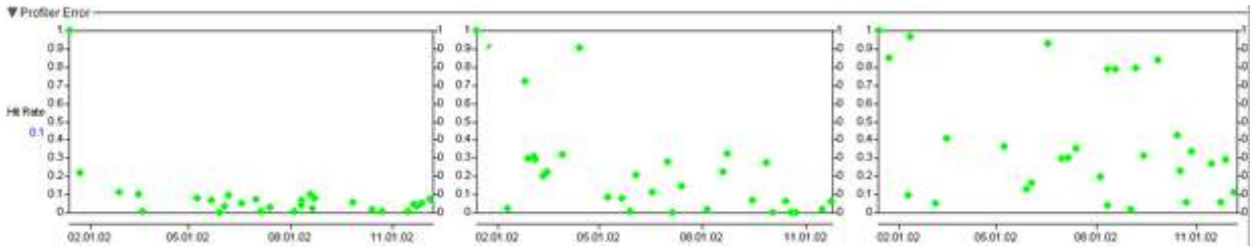

Figure 6 The effect of positioning error on prediction of the Murder activity. (left) no error, (middle) 750m error, (right) 1.5km error

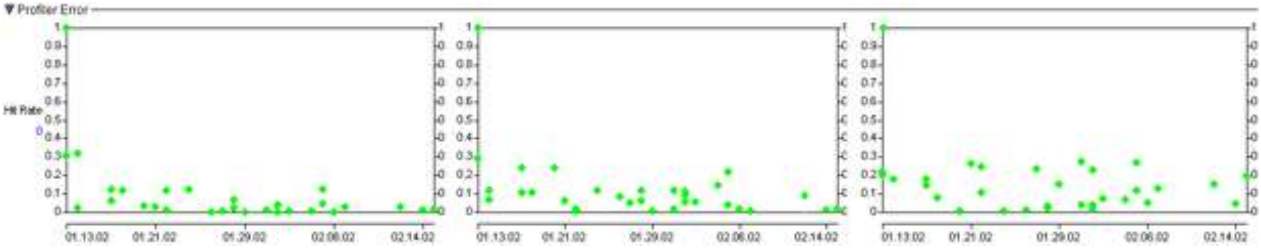

Figure 7: The effect of positioning error on prediction of the Avon-Sell Drug activity. (left) no error, (middle) 750m error, (right) 1.5km error

CONCLUSION

Though the above analysis deals with the detection of drug-trafficking activity, the approach has significant relevance to detection of any partially observable agent behavior. For example, a relevant process theory could be developed for the domain of bank fraud or identity theft. Preference functions could be automatically developed for each individual based on purchase history. Spatial constraints could be formed based on the locations of purchases. Deviations in the known patterns would result in warnings of possible adversarial activity.

This initial investigation has turned up a number of promising directions for future research that would build upon and refine the initial approach. There are a number of ways in which the basic process model framework in the prototype could be extended. The initial analysis focused on temporal constraints when detecting related events. Process models may also contain spatial constraints, however, which can specify either a bound on or a preference for minimizing factors such as distance, transit time, safety, or number of checkpoints between two events. If we were to consider the entire process model while generating predictions for a given activity, we could refine our prediction by taking into account spatial constraints between that activity and other activities in the process. This might significantly improve prediction accuracy.

Another limitation of the process modeling approach is that it requires a library of process models to work from. Currently, the system expects that these models are provided by a subject matter expert, probably created manually using some type of software tool. For large numbers of agents with rapidly-evolving behavior, the analyst would become a serious bottleneck and maintaining an up-to-date and comprehensive process model library would become difficult. It would thus be useful to explore approaches to automatically learning process models through analysis of the event stream.

If an existing "seed" library of process models is available, a case-based reasoning (CBR) approach to learning new process models might be feasible. In this approach, when the system fails to produce good matches against existing process models, it tries to generate a new process model that is a better fit. It would first retrieve the existing process model that best fits the events. It would then adapt that process model to improve its fit to the data, either by adding or removing events or by changing constraints between events. The new process model then is added to the library.

Another direction for future work is to investigate methods for learning process models "from scratch," using only historical event data and no pre-existing models. One avenue here is statistical approaches that consider co-occurrence of events and actors. Another avenue is model-based approaches that use domain-specific knowledge to construct process models in causal fashion based on agent goals. Obviously the latter might require a significant knowledge engineering effort that might outstrip the difficulty of simply hand-constructing process models, but there may be approaches that minimize the domain knowledge required.

REFERENCES

Bloom, L. (2003) "Modeling Adaptive, Asymmetric Behaviors." Conference on Behavior Representation in Modeling and Simulation (BRIMS).

Brantingham, P. and Brantingham, P. (1990), Environmental Criminology, Waveland Press.

Giles, Chris and Fedora, Richard (2005), "CrimePointWeb and Environmental Criminology."

White paper, http://www.forensiclogic.com, April 2005.

Jenkins, Brian Michael (2002) Countering al Qaeda: An Appreciation of the Situation and Suggestions for Strategy. RAND Report MR-1620-RC.

Johnson, Steven (2006), The Ghost Map. Riverhead, 2006.

Rossmo, D. Kim (1999), Geographic Profiling. CRC.

Rossmo, D. Kim (2005), An Evaluation of NIJ's Evaluation Methodology for Geographic Profiling Software

O'Sullivan, David (2004) "Too Much of the Wrong Kind of Data: Implications for the Practice of Micro-Scale Spatial Modeling." In Spatially Integrated Social Science, Goodchild, M.F. and Janelle, D.G., eds., Oxford University Press, 2004.

Skov-Petersen, Hans (2001). "Estimation of distance-decay parameters - GIS-based indicators of recreational accessibility." ScanGIS 2001, The 8th Scandinavian Research Conference on Geographical Information Science, 2001.

Snook, B., Zito, M., Bennell, C., and Taylor, P.J. (2005) "On the Complexity and Accuracy of Geographic Profiling Strategies," Journal of Quantitative Criminology, Volume 21, Number 1, March 2005.

Verma, Arvind and S. K. Lodha (2002), "A Typological Representation of the Criminal Event." Western Criminology Review 3(2).