Automated Detection of Terrorist Activities through Link Discovery within Massive Datasets

Christopher M. Boner

Metron, Inc.
11911 Freedom Drive, Suite 800
Reston, VA 20190
boner@metsci.com

Abstract

This paper describes link discovery technology that is designed to detect threat activities by extracting and piecing together transactional evidence from massive datasets that are composed mostly of noise and clutter. The approach is an integration of several innovative component technologies, including partial pattern matching, hypothesis evaluation and hypothesis merging.

Background

Terrorist organizations differ from nation states in that their members and assets are less conspicuous, harder to identify and to track. They do not have a uniformed military, well-known bases or huge arsenals of military assets and weapons. The threat posed by terrorist organizations stems from the covert network that enables them to operate "under the radar."

To carry out attacks terrorists must plan, recruit, travel, acquire assets and capabilities, and execute a number of other tasks that are particular to the activity. Executing these tasks requires transactions and associations among the terrorists and entities that carry them out. Each terrorist plan can be characterized by the structures of data, or *signatures*, that result from the execution of its tasks and subtasks.

These signatures constitute more than a simple trail of evidence to be followed. Much of the exploitable information may be embedded in the structure of the network of transactions and associations. Effective counterterrorism requires link discovery technologies that exploit differences between threat signatures and signature-like structures that arise from simple noise together with organized legitimate activities, or *clutter*.

Commercially available link analysis tools such as Analyst Notebook™ and Visual Links™ are primarily visualization aids. To be effective, data must be winnowed down to, at most, thousands of links among hundreds of entities, most of which are relevant to detecting terrorist relationships and activities.

Technology Objectives

This paper describes an approach to automated link discovery and threat activity detection within massive datasets that may have extremely small signal-to-noise ratios. The proposed technologies employ probabilistic methods that manage uncertainty to handle potentially incomplete, unreliable and inconsistent data.

It is assumed that there will be at least two types of given information: (i) databases containing detected transactions and known associations, and (ii) one or more patterns of terrorist activity. Each threat pattern expresses the possible hierarchies of tasks and subtasks by which the plan may be executed and describes, for each task, the possible signatures that may result from its execution.

The technology can also utilize optional input information: (iii) knowledge about terrorist groups and their members; (iv) patterns of clutter activity; and (v) data collection characteristics such as source reliabilities and probabilities that certain types of transactions will be observed given attributes and relationships of the entities involved.

The main objective is to utilize all given information to generate a ranked list of hypotheses, each of which indicates how discovered structures of linked evidence support the occurrence of the hypothesized threat activity.

Innovative Component Technologies

The approach outlined herein is an integration of partial pattern matching, hypothesis evaluation, hypothesis merging and group detection technologies. The first three component technologies were developed, integrated and evaluated by researchers at Metron, Inc. They are the focus of this paper. Figure 1 illustrates the functional flow and evolution of hypotheses during a run of the integrated component technologies against transactional data to discover instances of a particular pattern of terrorist activity.

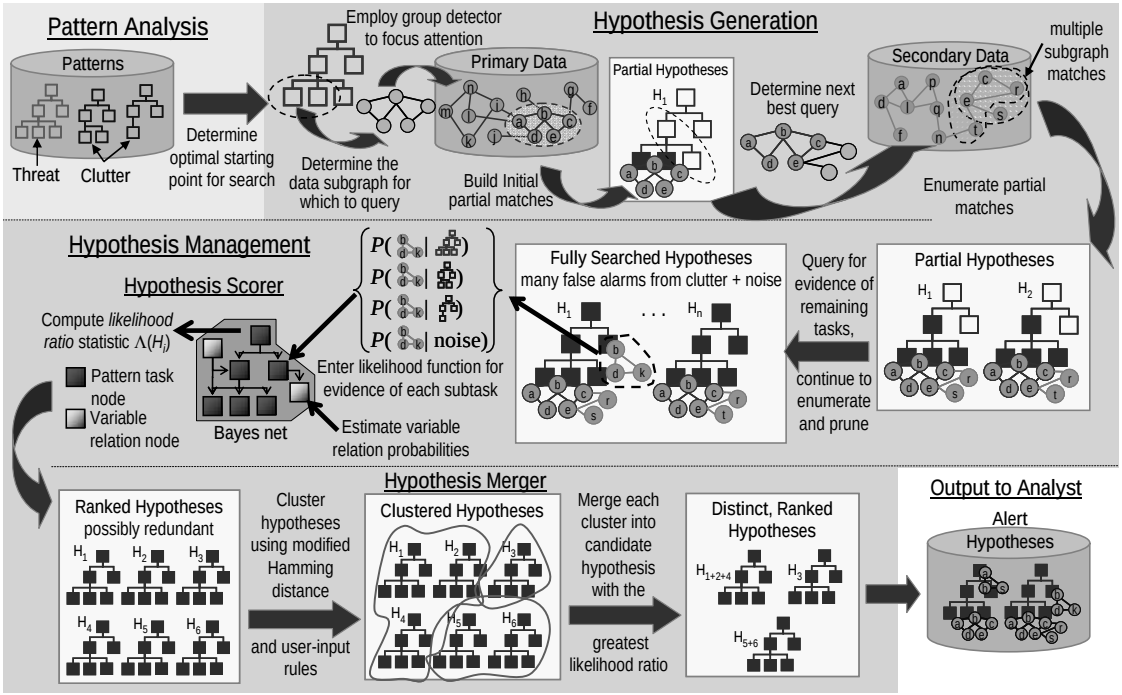

Figure 1. Functional flow and hypothesis evolution of integrated terrorist activity detection technologies

Partial Pattern Matching

The goal of the partial pattern matcher is to link evidence to discover partial matches to patterns of organized threat activity. It generates hypotheses about threat activities that specify the overall pattern of activity (e.g. developing chemical or nuclear warfare capabilities) and identify the tasks and subtasks by which the activity is (believed to be) carried out. Likely role players are identified and supporting evidence is provided for each hypothesized task and for the overall activity.

Any algorithms to discover threat signatures and piece them together to infer organized threat activities may be computationally intense. In many ways, this problem is similar to the subgraph isomorphism problem which is known to be NP-complete (Garey and Johnson 1979). One striking difference is that instead of merely answering the question of whether there is a subgraph of a given type in the data graph, the challenge at hand is to find all instances. Moreover, all of the transactions (analogues of subgraph edges) are not likely to be faithfully recorded in the data, so the present problem is analogous to inferring the existence of all subgraphs of a given type from the observed data graph that may contain only bits and pieces of each subgraph. There are many factors that impact the computational complexity of partial pattern matching algorithms, including but not limited to dataset size, "observability," connectivity, signal-to-noise ratio, reliability, and distinctiveness of patterns of threat activity compared to clutter.

The partial pattern matcher employs a directed search strategy to guard against a combinatorial explosion of the number of partial match hypotheses that are generated. It begins with an analysis of threat patterns, clutter patterns (if available), and the noise model to determine the optimal starting point for each search. A good query for a structure of data is one that is likely to arise from the execution of the threat pattern but unlikely to arise from clutter and noise. Subsequent queries attempt to grow existing partial hypotheses by searching for signatures of remaining tasks that are least likely to arise from clutter and noise. Multiple matches to each query are enumerated and explored in a tree search.

The author and collaborators formulated and published a mathematical theory of detection on structured, linked data (Boner et al. 2004), including a formula for the ratio of the probability that a structure of data will arise from a "threat" random graph model to the probability that the structure will arise from noise. Many of the random graph results and insights are relevant to the implementation of the partial pattern matcher's search/query optimization strategy.

To further reduce the search space, the partial pattern matcher has been integrated with group detection technologies that analyze transactional data to (i) expand and revise a priori hypotheses about terrorist groups and their members, and (ii) to hypothesize new threat groups and members. Such technology is employed by the partial pattern matcher to focus attention on likely threat collaborators. In this mode, only transactions and

associations that involve at least one member of a hypothesized threat group are considered when linking evidence to produce partial matches to threat patterns.

Because the partial pattern matcher allows *partial* matches to patterns of activity, many of the hypotheses that it generates may be false alarms whose supporting evidence consists primarily of clutter and noise transactions. Some may be quite sparse, while others may include a lot of evidence in support of tasks that are, by themselves, less indicative of threat activity, such as organized team-communications or team visitations to sites that are vulnerable to attack (both of which are likely to be carried out during many legitimate activities). Thus, it would be advisable to have a method for discriminating between threat and non-threat hypotheses that is more sophisticated and reliable than, say, simply comparing the amount of supporting evidence that each hypothesis contains.

Hypothesis Evaluation

For classical detection problems, the likelihood ratio $\Lambda(\boldsymbol{x}) = \mathbf{Pr}(\boldsymbol{x} \mid \boldsymbol{H_1}) \;/\; \mathbf{Pr}(\boldsymbol{x} \mid \boldsymbol{H_0})$ is an optimal statistic for deciding between two hypotheses $\boldsymbol{H_0}$ and $\boldsymbol{H_1}$, such as target present versus target not present, given information $\boldsymbol{x}$ (DeGroot 1970, Neyman-Pearson lemma).

The hypothesis scorer component computes a likelihood ratio statistic for each threat activity hypothesis and produces a ranked list of the most compelling hypotheses. A Bayesian network representation of the threat activity pattern is utilized to estimate the probability that the structure of supporting evidence is observed given that terrorists are executing the plan divided by the probability of observing the supporting evidence given that there is no terrorist plan, only clutter activities and noise transactions.

The Bayes net, which is constructed automatically from input threat and clutter patterns, encodes all probabilistic and relational information among the tasks, subtasks, role players and other pattern variables. Thus, variable relation probabilities may be factored into the likelihood ratio computation.

For example, the hypothesis scorer component has been integrated with group detection technology developed by researchers at Carnegie Mellon University to estimate the probabilities of the transactional data given that a hypothesized set of team members belong to (a) the same threat group, (b) the same non-threat group, or (c) no common group. The hypothesis scorer has also been integrated with a Bayesian classifier developed by the same CMU team to estimate from training data the probabilities that a set of assets can (a) be used to exploit a vulnerability of the hypothesized target, (b) be applied productively to the target, or (c) neither of the above. Each of these likelihood functions can be easily factored into the probability distributions maintained by the Bayes net. It was demonstrated during multiple experiments that the above mentioned integrations significantly improved the discriminative power of the hypothesis scorer.

Hypothesis Merging

Once hypotheses are scored and ranked, there may still be multiple subsets of hypotheses that pertain to the same ground truth activities. The hypothesis merger employs a modified Hamming distance metric together with user-input rules to cluster hypotheses that are likely to pertain to the same ground truth activity. Each cluster is merged into the single most compelling candidate hypothesis. This reduces the burden on analysts to assimilate like hypotheses and enables them to investigate a greater number of distinct terrorist activity hypotheses.

Experimental Results Summary

During a 2004 technology evaluation of a government counter-terrorism program, the integrated technologies described herein were applied to detect threat activities within simulated datasets that included up to a million transactions among 10,000 entities. Each dataset contained evidence from dozens of threat activities, hundreds of clutter activities, and up to hundreds of thousands of noise transactions.

Our integrated approach achieved greater than 75% recall and 81% precision in an average of five minutes per dataset. Our results were significantly better along every metric than the results of any other effort.

Acknowledgements

I thank my colleagues at Metron, Inc. who collaborated with me to design, develop, integrate and evaluate the component technologies described in this paper, especially Tom Mifflin, Didier Vergamini and Paul Goger. I thank Jeff Schneider, Andrew Moore and Dan Pelleg of Carnegie Mellon University and Hans Chalupsky and Andre Valente of the University of Southern California Information Sciences Institute for sharing their group detection code and expertise.

References

Boner, C.; Godfrey, G.; Mifflin, T.; and Skokan, J. 2004. A Random Graph Model for Terrorist Transactions. In *IEEE Aerospace Conference Proceedings*. Big Sky, MN: IEEE.

DeGroot, M. 1970. *Optimal Statistical Decisions*. New York, NY: McGraw-Hill.

Garey, M. and Johnson, D. 1979. *Computers and Intractability: A Guide to the Theory of NP-completeness*. New York, NY: W. H. Freeman and Company.