

CrimeNet Explorer: A Framework for Criminal Network Knowledge Discovery

JENNIFER J. XU and HSINCHUN CHEN
University of Arizona

Knowledge about the structure and organization of criminal networks is important for both crime investigation and the development of effective strategies to prevent crimes. However, except for network visualization, criminal network analysis remains primarily a manual process. Existing tools do not provide advanced structural analysis techniques that allow extraction of network knowledge from large volumes of criminal-justice data. To help law enforcement and intelligence agencies discover criminal network knowledge efficiently and effectively, in this research we proposed a framework for automated network analysis and visualization. The framework included four stages: network creation, network partition, structural analysis, and network visualization. Based upon it, we have developed a system called *CrimeNet Explorer* that incorporates several advanced techniques: a concept space approach, hierarchical clustering, social network analysis methods, and multidimensional scaling. Results from controlled experiments involving student subjects demonstrated that our system could achieve higher clustering recall and precision than did untrained subjects when detecting subgroups from criminal networks. Moreover, subjects identified central members and interaction patterns between groups significantly faster with the help of structural analysis functionality than with only visualization functionality. No significant gain in effectiveness was present, however. Our domain experts also reported that they believed CrimeNet Explorer could be very useful in crime investigation.

Categories and Subject Descriptors: H.4.2 [Information Systems Applications]: Types of Systems—*Decision support (e.g., MIS)*; H.3.4 [Information Storage and Retrieval]: Systems and Software—*Performance evaluation (efficiency and effectiveness)*

General Terms: Algorithms, Design, Experimentation, Performance

Additional Key Words and Phrases: Law enforcement, crime investigation, knowledge discovery, social network analysis, concept space, clustering, complete-link algorithm, multidimensional scaling, visualization, precision and recall

This research has primarily been funded by the National Science Foundation (NSF), Digital Government Program, “COPLINK Center: Information and Knowledge Management for Law Enforcement,” #9983304, July 2000-June 2003, and the NSF Knowledge Discovery and Dissemination (KDD) Initiative.

Authors’ address: Department of Management Information Systems, The University of Arizona, Tucson, AZ 85721; email: {jxu,hchen}@eller.arizona.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2005 ACM 1046-8188/05/0400-0201 \$5.00

1. INTRODUCTION

Since the tragic events of the September 11, 2001, attacks on the United States, academics around the world have been called on for possible contributions to the uncovering of terrorist networks to prevent future attacks. Internationally, both citizens and authorities have realized that knowledge about the structure of terrorist networks and how those networks operate will be a key factor in winning the so-called “netwar.”

Terrorism, together with such crimes as drug trafficking, gang-related incidents, fraud, and armed robberies, is called *organized crime*. Organized crime often requires conspiracy. In drug trafficking, for example, offenders cooperate with each other to carry out a series of illegal activities ranging from supplying, transportation, distribution, and sale of drugs to money laundering. Criminals may even form their own enterprises that operate like businesses, within which different individuals or teams are responsible for different tasks. These interrelated offenders constitute criminal networks.

Law enforcement and intelligence agencies have long realized that knowledge about criminal networks is important to crime investigation and may to a large extent shape police efforts [McAndrew 1999]. A clear understanding of network structures, operations, and individual roles can help law enforcement and intelligence agencies develop effective control strategies to prevent crimes from taking place.

However, criminal network analysis currently is primarily a manual process, usually consuming much time and human effort at each stage of the knowledge discovery process (data processing, transformation, analysis, and visualization). Although some existing tools provide visual representations of criminal networks to assist investigation, they lack structural network analysis functionality that may offer a deeper insight into the structure and organization of criminal enterprises.

To help law enforcement and intelligence agencies discover criminal network knowledge efficiently and effectively, we proposed in this research a framework for automated network analysis and visualization. This framework includes four major stages: network creation, network partition, structural analysis, and network visualization. Based on this framework, we developed a system called *CrimeNet Explorer* that incorporates several advanced techniques (a concept space approach, social network analysis methods, etc.) for automatically detecting subgroups from a network, identifying central members, and extracting interaction patterns between subgroups.

The remainder of the article is organized as follows: Section 2 introduces the background of criminal network analysis; Section 3 reviews existing network analysis tools and social network analysis techniques; Section 4 presents and explains the CrimeNet Explorer framework in detail. System evaluation is discussed in Section 5 and Section 6 concludes the article and suggests future research directions.

2. BACKGROUND

When analyzing criminal networks, crime investigators often focus on characteristics of the network structure to gain insight into the following questions

[McAndrew 1999; Sparrow 1991]:

- Who is central in the network?
- What subgroups exist in the network?
- What are the patterns of interaction between subgroups?
- What is the overall structure of the network?
- Which member's removal would result in disruption of the network?
- How do information or goods flow in the network?

Knowledge of these structural characteristics can help reveal vulnerabilities of criminal networks and may have important implications for crime investigation.

2.1 Implications of Structural Network Analysis

Usually, criminal network members who occupy central positions should be targeted for removal or surveillance [Baker and Faulkner 1993; McAndrew 1999; Sparrow 1991]. A central member may play a key role in a network by acting as a leader who issues commands and provides steering mechanisms or serving as a gatekeeper who ensures that information or goods flow effectively among different parts of the network. Removal of these central members may effectively disrupt the network and put the operation of a criminal enterprise out of action.

In addition to studying roles of individual members, crime investigators also need to pay special attention to subgroups in criminal enterprises. Each subgroup or team may be responsible for specific tasks. Group members have to interact and cooperate to accomplish the tasks. Therefore, detecting subgroups in which members are closely related to one another can increase understanding of a network's organization.

Moreover, groups may interact with each other in such a way that interactions and relationships may reveal certain patterns. For example, one group may have frequent interactions with one other specific group but seldom interact with the rest of the network. When interaction and relationship patterns between groups are found, the overall structure of the network can become more apparent. Indeed, different structures have different points of vulnerability. Intelligence regarding the overall structure of a network can help law enforcement and intelligence agencies develop the most effective strategies to disrupt that network.

2.2 Special Network Structures

Different criminal network topologies such as chain structure, star/wheel structure, and complete/cliقة structure [Evan 1972; Ronfeldt and Arquilla 2001] require specific disruptive strategies. A chain structure consists of members (individuals or groups) that are connected one by one so that information or goods must flow from one member to its neighbor before getting to the next. In a star structure, members are all connected to a central member who acts as a leader or hub. In a complete network, all members are fully connected with one

another so that communication between any two members can be carried out directly. A star structure is a centralized network, whereas chain and complete structures are considered decentralized networks [Baker and Faulkner 1993; Freeman 1979]. To disrupt a centralized network, removal of the central member(s) can cause the network to fall apart. A decentralized network, however, is more difficult to disrupt and more resistant to damage.

Although criminal network knowledge has important implications for crime investigation, little research has been done to develop advanced, automated techniques to assist with such tasks [Klerks 2001; McAndrew 1999; Sparrow 1991]. In the next section we review existing network analysis and visualization tools and introduce several new techniques that could be used for network analysis.

3. LITERATURE REVIEW

Existing network analysis tools used by law enforcement and intelligence agencies mainly focus on network visualization and do not have much structural analysis capability. Such a limitation might be successfully addressed by several methods from social network analysis research.

3.1 Existing Network Analysis Tools

Klerks [2001] categorized existing criminal network analysis tools into three generations.

3.1.1 First Generation: Manual Approach. Representative of the first generation is the Anacapa Chart of Harper and Harris [1975]. In this approach, an investigator first constructs an association matrix by examining data files to identify associations between criminals. Based on this association matrix, a link chart can be drawn for visualization purposes. The criminal having the most links to other people may be placed at the center of the link chart, indicating his/her importance in the network. The investigator then can study the structure of the graphical portrayal of the network to discover patterns of interest.

Krebs [2001], for example, mapped a terrorist network comprised of the 19 hijackers in the September 11 attacks. He first examined publicly released information reported in several major newspapers to gather data about relationships among the terrorists. He then manually constructed an association matrix to integrate these relations and drew a terrorist network depicting possible patterns of interactions based on the matrix (see Figure 1).

Although such a manual approach is helpful for crime investigation, for very large data sets its use becomes extremely ineffective and inefficient.

3.1.2 Second Generation: Graphics-Based Approach. Second-generation tools are more sophisticated because they can produce graphical representations of networks automatically. Most current criminal network analysis tools belong to this generation; among them are Analyst's Notebook, Netmap, and Watson.

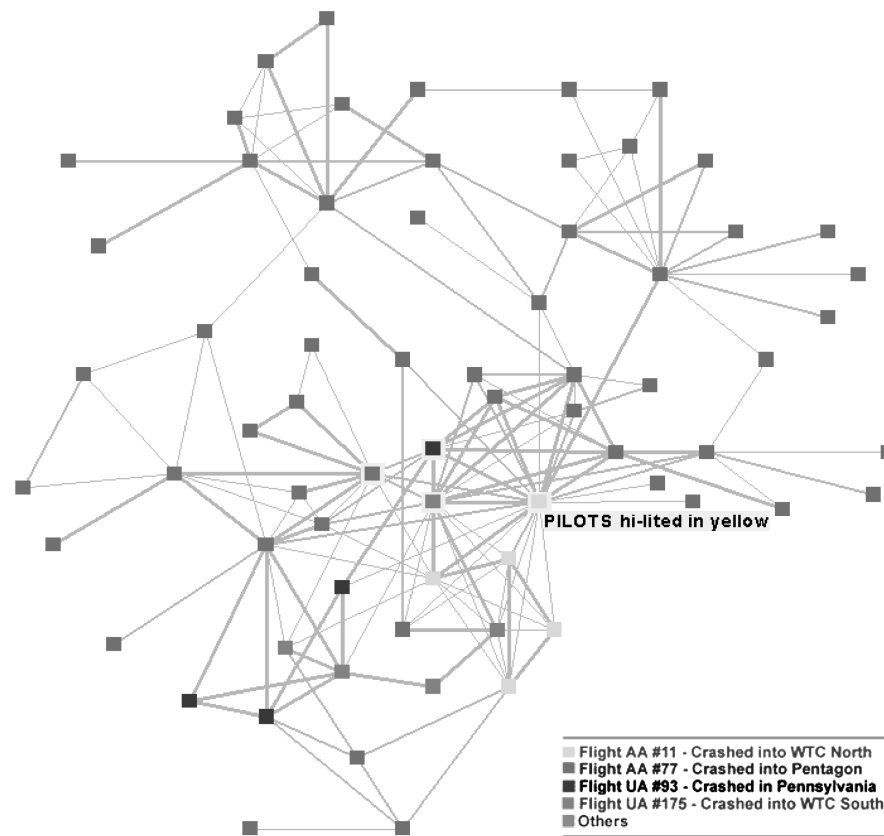


Fig. 1. The terrorist network surrounding the 19 hijackers on September 11, 2001. (Source: <http://www.orgnet.com>).

Analyst's Notebook has been widely employed by law enforcement in the United States and The Netherlands [Klerks 2001]. Like the first-generation approach, Analyst's Notebook relies on a human analyst to detect criminal relationships in data and can automatically generate a link chart based on relational data stored in a spreadsheet or text file. It uses icons to distinguish between different types of entities (e.g., persons, bank accounts, companies, addresses, etc.) and allows a user to drag those icons around to rearrange the network layout. For example, an icon representing a key person can be dragged to the center of the chart, and less important icons can be placed on the periphery.

Similarly, Netmap provides network visualization functionality (see Figure 2(a)). The system lays out entities of various types on the perimeter of a circle and places straight lines between entities to represent links. By examining the links, an analyst may discover useful patterns of interactions and associations hidden behind the network. Netmap has been adopted in the FinCEN system at the U.S. Department of the Treasury to analyze patterns of financial transaction data to detect money laundering [Goldberg and Senator 1998].

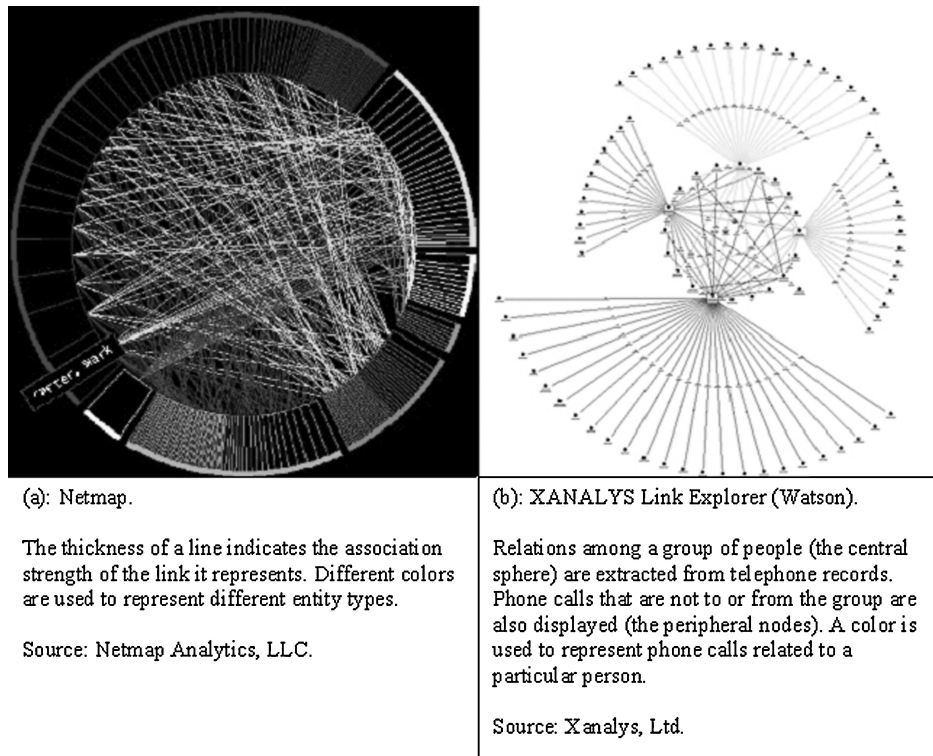


Fig. 2. Second-generation criminal network analysis tools.

Another second-generation tool called *XANALYS Link Explorer*, previously called *Watson* [Anderson et al. 1994] can search and identify possible associations between persons by querying databases (see Figure 2(b)). Given a person's name, *XANALYS Link Explorer* can automatically form a database query to search for related persons. The related persons found are linked to the given person and the result is presented in a link chart.

Although second-generation tools are capable of visualizing criminal networks, their sophistication level remains modest because they offer little structural analysis capability. The analysis burden is still on human crime analysts.

3.1.3 Third Generation: Structural Analysis Approach. No existing tool is sophisticated enough to be categorized as being of the third generation. Tools of this new generation are expected to provide more advanced analytical facilitation that helps discover structural characteristics of criminal networks: central members, subgroups, interaction patterns between groups, and the overall structure.

Prior research has recognized that social network analysis approaches are promising techniques for studying criminal networks [Klerks 2001; Krebs 2001; Ronfeldt and Arquilla 2001; McAndrew 1999; Sparrow 1991].

3.2 Social Network Analysis

Social network analysis (SNA) is used in sociology research to analyze patterns of relationships and interactions between social actors in order to discover an underlying social structure [Berkowitz 1982; Breiger 2004; Scott 1991; Wasserman and Faust 1994; Wellman 1988]. A number of quantitative SNA methods have been employed to study organizational behavior, interorganizational relations, citation patterns, computer-mediated communication, and many other domains [Galaskiewicz and Krohn 1984; Garton et al. 1999; Kleinberg 1999]. SNA has recently been recognized as a promising technology for studying criminal organizations and enterprises [McAndrew 1999; Sparrow 1991]. Studies involving evidence mapping in fraud and conspiracy cases have recently been added to this list [Baker and Faulkner 1993; Saether and Canter 2001]. These studies, however, focused only on central network members and did not identify subgroups and interaction patterns in criminal networks.

In SNA studies, a network is usually represented as a graph which contains a number of nodes (network members) connected by links (relationships). Structural analysis methods in SNA fall into two categories: relational and positional [Burt 1980], both of which are relevant to the study of criminal networks [McAndrew 1999].

3.2.1 Relational Analysis. Relational analysis focuses on relationships and interactions between network members. It is often used to identify central members or to partition a network into subgroups. In such studies, links usually are weighted by relational strength.

Several centrality measures can be used to identify key members who play important roles in a network. Freeman [1979] provided definitions of the three most popular centrality measures: degree, betweenness, and closeness.

Degree measures how active a particular node is. It is defined as the number of direct links a node k has:

$$C_D(k) = \sum_{i=1}^n a(i, k), \quad (1)$$

where n is the total number of nodes in a network, and $a(i, k)$ is a binary variable indicating whether a link exists between nodes i and k . A network member with a high degree could be the leader or “hub” in a network.

Betweenness measures the extent to which a particular node lies between other nodes in a network. The betweenness of a node k is defined as the number of geodesics (shortest paths between two nodes) passing through it:

$$C_B(k) = \sum_k^n \sum_j^n g_{ij}(k), \quad (2)$$

where $g_{ij}(k)$ indicates whether the shortest path between two other nodes i and j passes through the node k . A member with high betweenness may act as a gatekeeper or “broker” in a network for smooth communication or flow of goods (e.g., drugs).

Closeness is the sum of the length of geodesics between a particular node k and all the other nodes in a network. It actually measures how far away one node is from other nodes and is sometimes called *farness* [Baker and Faulkner 1993; Freeman 1979]:

$$C_c(k) = \sum_{i=1}^n l(i, k), \quad (3)$$

where $l(i, k)$ is the length of the shortest path connecting nodes i and k .

Another type of relational analysis is to partition a network based on the strength of relationships between network members. Because criminals often form groups or teams to commit crimes, such an approach can help detect subgroups in a large criminal network.

Two methods have been employed for network partition in SNA studies: matrix permutation and hierarchical clustering [Arabie et al. 1978; Wasserman and Faust 1994]. The purpose of matrix permutation is to rearrange rows and columns of a matrix so that members who occupy adjacent rows (or columns) can be organized into the same group. To use the matrix permutation method, a network must be represented as an $n \times n$ matrix in which both rows and columns represent n network members. The value of a matrix cell, (i, j) , is set to be the relational strength between members i and j . A zero value means that i and j are not directly related. Some SNA studies simply set cell values to be either 1 or 0 to indicate presence or absence of a relationship [Wasserman and Faust 1994]. Since matrix permutation is inherently an NP-hard problem, many SNA studies use hierarchical clustering methods [Arabie et al. 1978]. Hierarchical clustering will be reviewed in Section 3.2.3.

3.2.2 Positional Analysis. Unlike relational analysis, positional analysis examines how similarly two network members connect to other members. The purpose of positional studies is to discover the overall structure of a social network. The key approach is blockmodeling [Breiger et al. 1975], which includes two steps: network partition and interaction pattern identification.

In blockmodeling, a network is first partitioned into positions based on a structural equivalence measure [Lorrain and White 1971] rather than on the relational strength that is used in relational analysis. Two nodes are structurally equivalent if they have identical links to and from other nodes. Since perfectly equivalent members rarely exist in reality, this measure is relaxed to indicate the extent to which two nodes are substitutable in structure [Wasserman and Faust 1994]. A position thus is a collection of network members who are structurally substitutable, or in other words, similar in social activities, status, and connections with other members. Position is different from the concept of subgroup in relational analysis because two network members who are in the same position need not be directly connected [Lorrain and White 1971; Scott 1991]. In Figure 3, for example, nodes 1 and 2 are in the same position because they connect with other nodes (3, 4, and 5) in the same way. However, they are not directly connected to each other. Similarly, nodes 3 and 4 are in the same position.




Fig. 3. Positions in a social network.

Blockmodeling also employs hierarchical clustering to partition a network [Burt 1976]. However, because it is based on structural equivalence measure rather than on relational strength, the resulting clusters are positions rather than subgroups.

To model interaction patterns between positions, blockmodel analysis compares the density of links between two positions with the overall density of a network [Arabie et al. 1978; Breiger et al. 1975; White et al. 1976]. Link density between two positions is the actual number of links between all pairs of nodes drawn from each position divided by the possible number of links between the two positions. In a network with undirected links, for example, the between-position link density can be calculated by $d_{ij} = \frac{m_{ij}}{n_i n_j}$, where d_{ij} is the link density between positions i and j ; m_{ij} is the actual number of links between positions i and j ; and n_i and n_j represent the number of nodes within positions i and j , respectively. The overall link density of a network is defined as the total number of links divided by the possible number of links in the whole network, that is, $d = \frac{m}{n(n-1)/2}$, where m is the total number of links, and n is the total number of nodes in the network. Notice that for an undirected network the possible number of links is always $n(n-1)/2$.

A blockmodel of a network is thus constructed by comparing the density of the links between each pair of positions, d_{ij} , with d : a between-position interaction is present if, $d_{ij} \geq d$, and absent otherwise. Blockmodeling therefore reduces a complex network to a simpler structure by summarizing individual interaction details into relationship patterns between positions [White et al. 1976]. As a result, the overall structure of the network becomes more evident.

3.2.3 Hierarchical Clustering. Although they are based on different measures, both relational and positional analysis in SNA may employ hierarchical clustering to partition a network.

Hierarchical clustering usually groups a set of objects into a series of nested clusters based on similarity/dissimilarity between objects [Aldenderfer and Blashfield 1984; Johnson 1967; Lance and Willams 1967]. When used in relational analysis, hierarchical clustering treats relational strength as a similarity measure. Therefore, the resulting clusters represent subgroups

whose members are closely related. When applied in positional analysis, on the other hand, hierarchical clustering uses structural equivalence to measure similarity and resulting clusters represent positions whose members are similar in the way they connect to other members.

Hierarchical clustering generates a cluster hierarchy represented by a dendrogram, in which clusters are merged at successively less restrictive values of similarity. The advantage of hierarchical clustering is that a network can be partitioned into different numbers of clusters at different similarity levels. At the most restricted level, every object occupies a cluster, whereas at the least restricted level, all objects fall into one big cluster. With this feature, the underlying structure of a network can be analyzed at different levels of detail. The disadvantage of hierarchical clustering, on the other hand, is that each object can be assigned to only one cluster at a specific level of similarity [Wasserman and Faust 1994]. There is no overlap between clusters.

Various methods have been developed for hierarchical clustering, such as single-link [Sibson 1973; Sneath 1957], complete-link [Defays 1977; King 1967], and Ward's algorithm [Ward 1963]. Both single-link and complete-link methods merge the two most similar clusters into one larger cluster one at a time. However, they calculate between-cluster similarity differently. Among the three methods, the complete-link algorithm is the most popular because it gives more homogeneous and stable clusters than the others [Jain and Dubes 1988; Lance and Williams 1967; Jain et al. 1999].

3.2.4 Visualization of Social Networks. Many SNA studies use multidimensional scaling (MDS) to visualize social networks [Wasserman and Faust 1994]. MDS is a data analysis technique that seeks to provide a visual representation of proximities (dissimilarities) among objects so that objects that are more similar to each other are closer on the display and objects that are less similar to each other are farther apart [Kruskal and Wish 1978; Young 1987].

There are two types of MDS methods: metric for quantitative data and nonmetric for qualitative data [Kruskal 1964; Torgerson 1952; Young 1987]. For both approaches, Kruskal's STRESS [Kruskal 1964], which measures the goodness-of-fit when projecting higher-dimensional data onto a lower-dimensional display, is the objective function to be maximized.

SNA studies employ MDS in both relational and positional analysis of social networks [Breiger et al. 1975; Burt 1976; Freeman 2000; Wasserman and Faust 1994]. When applied to a relational analysis, MDS uses relational strength as a measure of proximity and outputs an x - y coordinate for each object on a two-dimensional plane so that closely related members are also close visually. When applied to positional analysis, MDS uses the structural equivalence between members as a proximity measure so that members who are structurally substitutable are close together on the display.

Recently, a network layout algorithm called *spring embedder* has been employed to visualize social networks [Freeman 2000]. This algorithm treats links connecting nodes as springs through which nodes attract and repulse each other. Nodes settle down when the total energy carried by the springs




Fig. 4. A framework for automated criminal network analysis and visualization.

is minimized. A spring embedder usually is intended to achieve pleasing visual effects such as evenly distributed nodes and a minimal number of crossing links [Davidson and Harel 1996]. The network layout generated by a spring embedder might be quite different from that generated by MDS because of their different objective functions and node location handling mechanisms.

In summary, SNA offers several structural analysis techniques that can be used to extract criminal network knowledge. However, existing network analysis tools are not sophisticated enough to employ these techniques. To analyze a criminal network, an investigator has to extract information about criminal relationships from data, create a network representation, and perform structural analysis manually to identify central members, to detect subgroups, and to discover interaction patterns among groups. It is highly desirable to automate the whole process of criminal network analysis so that knowledge can be extracted more efficiently and effectively.

4. CRIMENET EXPLORER: A FRAMEWORK FOR DISCOVERING CRIMINAL NETWORK KNOWLEDGE

To facilitate criminal network knowledge extraction, we proposed a framework that incorporates several structural analysis and visualization methods. Based on this framework we developed a system called *CrimeNet Explorer* that can be categorized as a third-generation network analysis tool. We also wanted to find out whether criminal network knowledge can be discovered effectively and efficiently by using the analytical functionality provided by the system.

Figure 4 presents the proposed framework, which includes four major stages: network creation, network partition, structural analysis, and network visualization.

4.1 Network Creation

Criminal-justice data collected from crime incident reports, telephone records, surveillance logs, financial transaction records, and other sources usually do not store explicit information about criminal relationships. The task of extracting relational information from raw data and transforming it into a networked format could be quite labor-intensive and time-consuming.

To address this problem, we employed a concept space approach [Chen and Lynch 1992] to create networks automatically [Chen et al. 2003; Hauck et al. 2002]. The concept space approach was originally employed in information retrieval applications for extracting term relations in documents. It uses cooccurrence weight to measure the frequency with which two words or phrases

appear in the same document. The more frequently two words or phrases appear together, the more likely it will be that they are related.

The criminal-justice data used in our research consisted of crime incident summaries provided by the Tucson Police Department (TPD). We treated each incident summary (database records specifying the date, location, persons involved, and other information about a specific crime) as a document and each person's name as a phrase. We then calculated cooccurrence weights based on the frequency with which two individuals appeared together in the same crime incident. We assumed that criminals who committed crimes together might be related and that the more often they appeared together the more likely it would be that they were related. As a result, the value of a cooccurrence weight not only implied a relationship between two criminals but also indicated the strength of the relationship [Hauck et al. 2002].

With the concept space approach, criminal relationships therefore could be extracted from crime incident data and transformed into a networked format automatically. Resulting networks were undirected, weighted graphs in which nodes represented individual criminals and cooccurrence weights of links represented relational strength. It is worth mentioning that the concept space approach has both advantages and disadvantages for extracting relations. On one hand, the weight of a link was normalized to a range between 0 and 1, better than the simple cooccurrence count. More importantly, the distribution of cooccurrences was extremely skewed. More than 90% of the criminal pairs resulted from a one-time cooccurrence and a small portion (around 2.4%) of pairs cooccurred 10 times or more. The concept space approach, which penalized extremely large cooccurrences [Chen and Lynch 1992], helped prevent the link weights from being skewed. On the other hand, the concept space approach is limited since the relational strength can be affected by other factors such as crime type. For example, a cooccurrence relation in a gang-related crime in which a large number of criminals participated might not be as strong as a relation in an auto-theft crime in which only two criminals were involved.

We also observed that the network generated might not necessarily be a single connected graph that contained all criminals in a set of data. This might be due to the fact that some criminal enterprises might not have any connection with other criminal organizations. It could also be caused by the incompleteness of the data [McAndrew 1999].

The networks created were stored in a database table in which each tuple specified a pair of criminals and an associated cooccurrence weight. These cooccurrence weights would be used later in both structural analysis and network visualization.

4.2 Network Partition

With data expressed in a networked format, we employed hierarchical clustering to partition a network into subgroups based on relational strength.

We used a complete-link algorithm since it was less likely to be subject to the chaining effect [Jain et al. 1999]. Existing complete-link algorithms

```

Form a cluster for each node;
while at least one between-cluster distance is less than infinite do
  currentCluster = an arbitrary cluster;
  found = false;
  while not found do
    find the nearest neighbor, C, to the currentCluster;
    if isRNN(C, currentCluster) then
      merge C and currentCluster;
      found = true;
    else
      currentCluster = C;
    end while
  end while
end while

```

Fig. 5. The pseudocode of the modified version of the RNN-based complete-link algorithm.

vary in space and time complexity [Day and Edelsbrunner 1984; Defays 1977; Voorhees 1986]. Although clustering was an offline operation that did not necessarily require high speed in our framework, we took into consideration that online dynamic clustering would be needed under some circumstances in the future. Therefore, time complexity was our primary criterion for algorithm selection. The algorithm we chose was an RNN-based complete-link algorithm that used the reciprocal nearest neighbor (RNN) approach developed by Murtagh [1984]. It took $O(n^2)$ time and $O(n^2)$ space and was significantly faster than other algorithms that typically required $O(n^3)$ time [Roussinov and Chen 1999].

Cooccurrence weights generated in the previous stage were first transformed into distances/dissimilarities. Since we were employing a complete-link algorithm, the distance between two clusters was defined as the distance between the farthest pair of nodes drawn from each cluster.

Initially, the algorithm treated each node as a cluster and then arbitrarily selected a cluster and incrementally built for it a nearest-neighbor chain (NN-chain). In an NN-chain, each cluster was the nearest neighbor of its previous cluster. A chain terminated with two clusters that were the nearest neighbor of each other. The two nearest clusters were then merged into a larger cluster and the cluster hierarchy (dendrogram) was updated. The algorithm kept merging nearest clusters until all the nodes were merged into one big cluster. The resulting hierarchy had multiple levels and each level corresponded to a specific partition of a network.

Since the previous stage created multiple disjoint networks, we modified the algorithm to make it generate a separate cluster hierarchy for each network. The hierarchies generated were stored in a database for later use. Figure 5 presents the pseudocode of the modified algorithm.

4.3 Structural Analysis

In our framework, central member identification and blockmodeling in the structural analysis stage were online operations performed by request.

We used the three centrality measures (degree, betweenness, and closeness) to identify central members in a given subgroup. The degree of a node could be obtained by counting the total number of links a node had to all the other group members. A node's score of betweenness and closeness required computing the shortest paths (geodesics).

In our implementation, Dijkstra's [1959] classical shortest-path algorithm was used to compute the geodesics from a single node to every other node in a subgroup. Given an undirected graph representing a subgroup i that consisted of n_i nodes, applying the algorithm $n_i - 1$ times could generate the shortest paths between all pairs of nodes in the subgroup. Betweenness of a specific node k was thus obtained by counting the number of geodesics between the other nodes passing through node k . Because running Dijkstra's algorithm once took $O(n_i^2)$ time, the overall time complexity for calculating betweenness of nodes in the subgroup i was $O(n_i^3)$.

There are specific algorithms for all-pair shortest-path calculations such as Dantzig's [1960] and Floyd's [1962] algorithms. These algorithms' time complexity is also $O(n^3)$. The advantage of using the Dijkstra's algorithm was that, by the time all the geodesics for a specific node were found, the computation of the closeness of that node was also finished, because the closeness was simply the sum of the length of the geodesics. Thus, closeness was a "byproduct" of betweenness and was obtained with no extra cost.

To extract between-group interaction patterns and the overall structure of a criminal network, we performed blockmodel analysis. Unlike general blockmodel analysis in SNA research that revealed interaction patterns between network positions based on the structural equivalence measure, our blockmodel analysis examined relationships between subgroups based on the relational strength measure. We decided on this approach based on interviews with the crime investigators from TPD and evidence that crime investigators often are more interested in interaction patterns between subgroups rather than between positions.

Blockmodeling in this framework therefore was used to identify interaction patterns between subgroups discovered in the network partition stage. At a given level of a cluster hierarchy, we compared between-group link densities with the network's overall link density to determine the presence or absence of between-group relationships.

4.4 Network Visualization

To map a criminal network onto a two-dimensional display, we employed MDS to assign a location to each node in a network of n nodes, given the corresponding $n \times n$ distance matrix. Since distances transformed from cooccurrence weights were quantitative data, we selected Torgerson's [1952] classical metric MDS algorithm. This algorithm first transformed the distance matrix into a scalar product matrix \mathbf{B} by double-centering. It then solved the singular value decomposition (SVD) problem for \mathbf{B} to generate an $n \times n$ matrix \mathbf{X} , the first two columns of which stored the coordinates of the n nodes.

The key step in this algorithm was SVD, which could be solved efficiently using the library routine provided by Press et al. [1992].

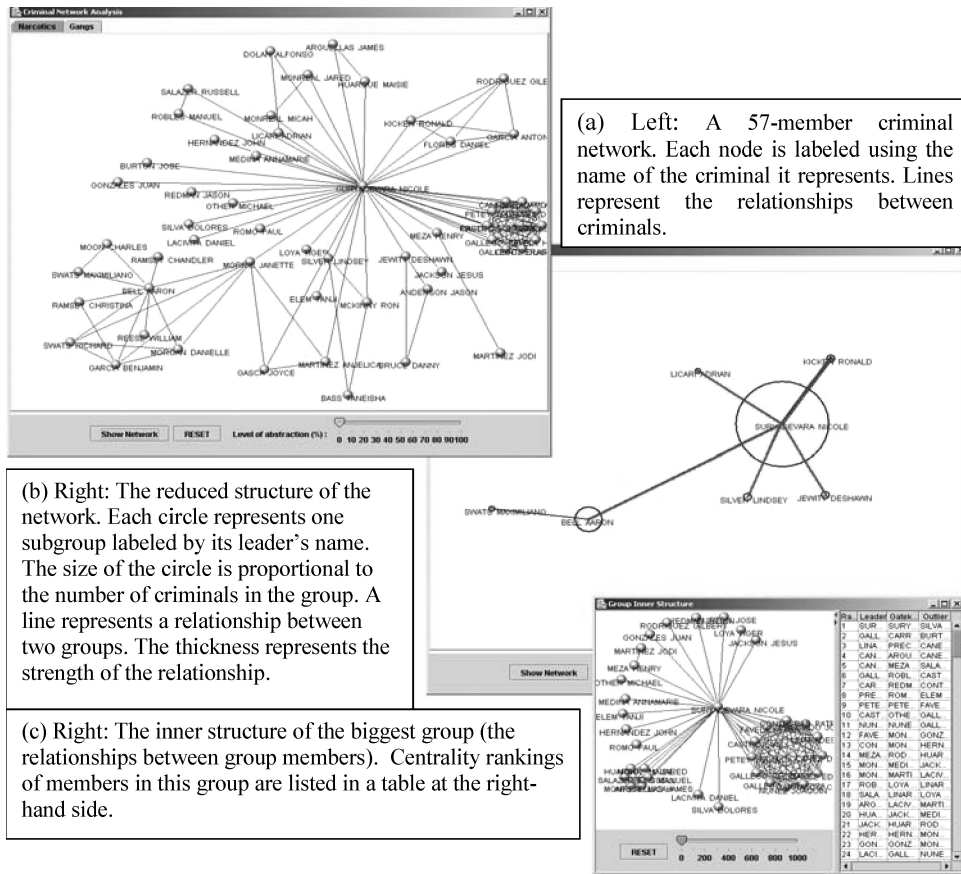


Fig. 6. CrimeNet Explorer. (In this example, the network appeared to be a star structure after performing blockmodel analysis. The vulnerability of this network, therefore, lay in the central members.)

4.5 CrimeNet Explorer

Our CrimeNet Explorer was developed based on the proposed framework. In the system a graphical user interface was provided for easy interaction between a user and the system. Figure 6 shows screen shots of the system interface. Each node was labeled with the name of the criminal it represented. Criminal names were scrubbed for data confidentiality. A straight line connecting two nodes indicated that the two corresponding criminals committed crimes together and thus were related.

To find subgroups and interaction patterns between groups, a user could adjust the “level of abstraction” slider at the bottom of the panel. A high level of abstraction corresponded with a high distance level in the cluster hierarchy. At any level of abstraction, a circle represented a subgroup. The size of the circle was proportional to the number of criminals in the subgroup. To view how group members were connected within a subgroup, a user could click on the corresponding circle to bring up a small window depicting the group’s inner

structure. At the same time, rankings in terms of the three centrality measures of the group members were listed at the right-hand side of the small window.

Straight lines connecting circles represented between-group relationships. The thickness of a line was proportional to the density of the links between the two corresponding groups. Such a design was different from general blockmodel analysis, which treats a low link density as an indicator of the absence of a between-group relationship. We thought that the absence of a line between two subgroups might possibly cause a user to infer mistakenly that there was no actual link connecting members from the two groups. We therefore kept a line between two groups as long as there was a link between members from the two groups. Our design decision could be more informative than the treatment in general blockmodel analysis for crime investigations.

5. SYSTEM EVALUATION

As discussed previously, the purpose of proposing a framework in this research was to employ advanced structural analysis and visualization techniques to help discover valuable criminal network knowledge. The major advantage of CrimeNet Explorer over existing network analysis tools was its structural analysis capabilities.

We conducted system evaluation to answer the following research questions:

- Will the system detect subgroups from criminal networks correctly?
- Will the structural analysis functionality help extract structural properties of criminal networks more effectively and efficiently?

Prior to the system evaluation, we carefully examined our TPD datasets and found that networks generated from them varied in size and structure.

5.1 Narcotics and Gang Networks

We extracted two datasets from the TPD databases: (a) incident summaries of narcotics crimes from January 2000 to May 2002, and (b) incident summaries of gang-related crimes from January 1995 to May 2002. Both narcotics and gang-related crimes were organized crimes likely to have been committed by networked offenders. We chose a longer time period for gang data because in each year there were substantially fewer gang-related crimes than narcotics crimes.

We analyzed the sizes of the networks generated from the two datasets. The narcotics dataset consisted of 12,842 criminals who were from 2628 networks. The gang dataset consisted of 4376 criminals from 289 networks. Both datasets contained a single large network (e.g., the 502-member network in the narcotics dataset) and a large number of small networks with fewer than 20 members. The biggest gang network was much larger than the biggest narcotics network although the gang dataset contained fewer criminals. Table I provides network-size statistics of the two datasets. Further examination of the incident summaries revealed that members in the large networks (those having more than 20 members) were mostly serial offenders and possibly came from various criminal organizations. In contrast, small networks (those having fewer than

Table I. Sizes of Networks Generated from the Two Datasets

	2–20 Members	21–100 Members	> 100 Members
Narcotic networks	2618	9	1 (a 502-member network)
Gang networks	284	4	1 (a 2595-member network)

20 members) consisted primarily of “one-time” offenders and would probably be less interesting for a study of criminal organizations and enterprises.

In addition to network size, we examined network structures using the blockmodeling function of CrimeNet Explorer. Because it was quite difficult to display the biggest networks in the two datasets on a screen, each having several hundred members, we analyzed only the structures of networks with 21–100 members. We found that the two types of networks had distinguishing structural patterns:

- Two out of the four gang networks studied had a star structure similar to the example in Figure 6. The third network had a chain of stars. The fourth network had a star structure with each branch being a smaller star or a clique and its overall structure looked like a snowflake.
- All nine narcotics networks had a chain structure. Three of these networks were chains of stars. One network had a circle in the middle of the chain.

Analysis of network size and structure revealed that gang networks tended to be bigger and more centralized, whereas narcotics networks were smaller and more decentralized. This finding implied that different strategies could be used to disrupt the two types of networks.

We selected a 60-member narcotics network and a 24-member gang network and used them in a subject study to evaluate CrimeNet Explorer.

5.2 Experimental Design

To address our research questions, we conducted a controlled laboratory experiment that allowed us to evaluate system performance. Thirty students from the Department of Management Information Systems at the University of Arizona participated in the experiment. We used students rather than crime investigators as research subjects based on two considerations. First, it was difficult to recruit a sufficient number of crime investigators because of their busy work schedules. Second, although our system was designed for criminal network analysis, finding structural patterns from networks of nodes was not a domain-specific task. Student subjects should be able to perform the tasks assigned to them even without domain knowledge in crime investigation.

Each subject participated in four sessions: demographic survey, training, testing, and posttest questionnaire. The demographic survey focused on subjects’ background information such as gender, age, and computer experience. The training session was designed to help subjects understand the major concepts (e.g., subgroups, central members, etc.) and gain hands-on experience with the system. During the testing sessions, subjects performed nine tasks on each of two test networks. They then completed a posttest questionnaire on

which they reported their attitudes toward the system’s ease-of-use and their satisfaction with the system’s functionality.

The 18 tasks used in the experiment were divided into three types: (1) detecting subgroups in a network, (2) identifying interaction patterns between subgroups, and (3) identifying central members within a given subgroup.

5.2.1 Task I: Subgroup Detection (Clustering). We wanted to learn through task I whether our system could achieve performance comparable to that of untrained users when partitioning a network into clusters (subgroups). We asked a domain expert (a detective who had served in law enforcement for more than 20 years) to provide partitions of the two test networks based on his knowledge of narcotics and gang-related crimes. His partitions were used as “gold standards” to evaluate clustering results generated by our system and subjects who represented untrained users.

There has not been a generally accepted metric for evaluating clustering results [Jain and Dubes 1988]. We selected for our experiment the clustering precision and cluster recall metrics developed by Roussinov and Chen [1999]. These two measures examined whether or not a pair of documents was put in the same cluster by human subjects and by the system [Sahami et al. 1998]. Based on the same rationale, we defined our cluster precision and recall as follows:

$$\begin{aligned} & recall_{system} \\ &= \frac{\text{number of node pairs in both system partition and expert partition}}{\text{number of node pairs in expert partition}}, \end{aligned} \quad (4)$$

$$\begin{aligned} & recall_{human} \\ &= \frac{\text{number of node pairs in both human partition and expert partition}}{\text{number of node pairs in expert partition}}, \end{aligned} \quad (5)$$

$$\begin{aligned} & precision_{system} \\ &= \frac{\text{number of node pairs in both system partition and expert partition}}{\text{number of node pairs in system partition}}, \end{aligned} \quad (6)$$

$$\begin{aligned} & precision_{human} \\ &= \frac{\text{number of node pairs in both human partition and expert partition}}{\text{number of node pairs in human partition}}. \end{aligned} \quad (7)$$

We developed two hypotheses to compare the clustering results from our system and the human subjects:

- H1: The system and subjects will achieve different clustering *recall*.
- H2: The system and subjects will achieve different clustering *precision*.

Since hierarchical clustering generated nested partitions for a network, we selected the partition containing the same number of clusters as in the expert’s partition to be the system’s clustering result. During the experiment, subjects were asked to partition a given network into the same number of clusters as in the expert partition. Although both the system and subjects generated the same number of clusters, they could assign different node pairs in a cluster, resulting in different recall and precision.

5.2.2 *Tasks II and III: Interaction Pattern and Central Members Identification.* Because the major advantage of CrimeNet Explorer was its structural analysis capability in addition to its network visualization functionality, we were interested in comparing subjects' performances under two experimental conditions: (1) structural analysis plus visualization, and (2) visualization only.

We considered two general information systems performance metrics [Jordan 1998]:

Effectiveness

= total number of correct answers a subject generated for a given type of tasks, (8)

Efficiency

= the average time a subject spent to complete a given type of tasks. (9)

Since the system could automatically identify interaction patterns between subgroups and central members within a subgroup, it was expected that a subject could achieve higher efficiency and effectiveness with the help of structural analysis functionality than with only visualization functionality. Specifically, we developed four hypotheses to compare the performance under two experimental conditions:

- H3: A subject will achieve higher *effectiveness* for interaction pattern identification tasks using the system having both structural analysis and visualization functionality than with that having visualization functionality only.
- H4: A subject will achieve higher *effectiveness* for central member identification tasks using the system having both structural analysis and visualization functionality than with that having visualization functionality only.
- H5: A subject will achieve higher *efficiency* for interaction pattern identification tasks using the system having both structural analysis and visualization functionality than with that having visualization functionality only.
- H6: A subject will achieve higher *efficiency* for central member identification tasks using the system having both structural analysis and visualization functionality than with that having visualization functionality only.

Our domain expert validated answers to all the questions for tasks II and III. To eliminate a learning effect, the orders of experimental conditions and tasks were randomized for each test network.

For task II, subjects were asked to answer two questions regarding the interaction patterns between subgroups:

- Given two subgroups, determine whether they were related.
- Given three subgroups (e.g., A, B, and C), determine whether group A had more interactions with group B than with group C.

For task III, subjects were asked to identify central members with the highest degree. We did not assign tasks of identifying central members with the highest betweenness and closeness because these two measures required

Table II. Clustering Recall and Precision

	Human	System
Recall	0.86 (0.07)	0.93 (0.00)
Precision	0.77 (0.03)	0.91 (0.00)

computation of shortest paths, which were difficult for subjects to find under the visualization-only condition. We therefore included only degree for fair comparison between the two experimental conditions.

For tasks II and III, subjects were encouraged to complete the tasks as quickly as possible. Each subject's task completion time was recorded. On average, it took a subject 30–45 min to complete all 18 tasks.

5.3 Results and Discussion

5.3.1 Quantitative Analysis: Clustering Recall and Precision. H1 and H2 were supported. Paired t -tests showed that the system's clustering recall and precision were significantly higher than subjects' (recall: $t = 4.39$, $p < 0.001$; precision: $t = 5.33$, $p < 0.001$). Table II gives the recall and precision rates of the system and the subjects. Numbers in parentheses are standard deviations.

We believe that the difference in clustering recall and precision resulted from visual clues that subjects relied on when performing clustering tasks.

- In our experiment, the domain expert based his partitioning of the test networks on his knowledge of network members and grouped criminals who frequently hang together in the same clusters. His judgment of clusters was not affected by visual clues from the network layouts.
- The system neither had domain knowledge nor was affected by visual clues from the network layouts. Thus, partitioning of the networks depended entirely on link weights (relational strengths). Since relational strength was determined by the frequency with which two criminals committed crimes together, it could relatively accurately reflect reality. Therefore, partitions generated by the system closely resembled the expert's partitions.
- Untrained subjects had to rely entirely on relative locations of nodes in the visual display of networks to determine relational strength between criminals. Visual clues thus could affect subjects' judgment heavily. When a network display was distorted (caused by dimensionality problem in the MDS algorithm), a subject actually could group weakly related criminals into one cluster if they appeared to be close visually. The test networks used in our experiment suffered from the distortion problem, which may have caused the clustering recall and precision by subjects to be worse than the clustering recall and precision of our system.

In regard to *effectiveness*, H3 and H4 were not supported. We performed paired t -tests for both tasks II and III to compare the effectiveness under the two experimental conditions (task II: $t = 1.41$, $p > 0.05$; task III: $t = 1.80$, $p > 0.05$). Such results implied that the analysis functionality did not help to achieve a significantly higher effectiveness. Table III shows the results.

Table III. Effectiveness

	Visualization Plus Analysis	Visualization Only
Task type 2	3.90 (0.31)	3.73 (0.59)
Task type 3	3.30 (1.02)	3.20 (1.13)

Table IV. Efficiency

	Visualization Plus Analysis	Visualization Only
Task type 2	7.13 (2.19)	12.10 (4.81)
Task type 3	6.24 (3.85)	26.93 (12.45)

Such a result could be for two reasons:

- For both tasks II and III, a subject could obtain a correct answer by counting lines on the network display under the visualization-only condition. For example, to compare the frequency of interactions between one group (A) and the other two groups (B and C), a subject could count the number of lines between A and B, and the number of lines between A and C. A simple comparison of these two numbers would suggest which two groups had more frequent interactions. As long as the subject was careful, he/she could find the correct answer.
- The two testing networks used in this experiment were not very large, making these two types of tasks relatively simple.

In regard to *efficiency*, the paired *t*-tests for efficiency comparison supported both H5 and H6 (task II: $t = 6.92$, $p < 0.001$; task III: $t = 10.66$, $p < 0.001$). This means that subjects could achieve significantly higher efficiency under the visualization-plus-analysis condition than under the visualization-only condition. Table IV shows the efficiency statistics.

The results implied that, with the help of structural analysis functionality, subjects could identify interaction patterns among subgroups and the central members in a given subgroup significantly faster. Under the visualization-plus-analysis condition, a subject did not have to count lines manually to identify interaction patterns between groups because a straight line between two groups implied the presence of a between-group interaction. At the same time, the thickness of the line indicated the frequency of the interaction. In addition, the degrees of all group members were computed by the system so that a subject could find the one with the highest degree directly from the centrality table on the interface.

In summary, the structural analysis functionality provided by the system could significantly improve the efficiency of network analysis tasks although the gain in effectiveness was not significant. Moreover, the system could identify subgroups of a network significantly better than untrained subjects.

5.3.2 Qualitative Feedback. Most subjects reported that features provided by the system were easy to learn and easy to use. For example, it was easy to adjust the slider to view different partitions at different abstract levels; it was convenient to visualize the inner structure of a subgroup in a small window.

The table used to list degree rankings of group members was similar to an Excel spreadsheet and easy to understand.

Subjects' negative comments about the system were primarily concerned with network layout and network partitions:

- Network layout.* Many subjects felt that the network was too cluttered in some areas where nodes were so close to each other that labels were overlapped and hard to read.
- Network partition.* Most reported the difficulty of deciding where to put nodes that had many connections to nodes from different groups. They said they wished overlapped groups could be allowed so that some very popular nodes could belong to more than one group. However, hierarchical clustering algorithms always generated mutually exclusive clusters that did not overlap.

Our domain expert also provided positive feedback. He said he had enjoyed using our system and believed that our CrimeNet Explorer could be very useful for crime investigation in the following ways:

- Increasing work productivity.* With the structural analysis functionality of CrimeNet Explorer, a large amount of investigation time could be saved.
- Assisting training for new crime investigators.* New investigators who did not have sufficient knowledge about local criminal organizations could use the system to grasp the essence of the networks and crime history quickly. They would not have to spend a significant amount of time studying hundreds of incident reports.
- Suggesting investigative leads that might otherwise be overlooked.*
- Assisting prosecution.* Known relationships between individual criminals and criminal groups would be helpful to the prosecution when seeking to prove guilt in court.

Overall, the results of the quantitative and qualitative analysis showed that our system developed based upon the framework could be efficient and useful for extracting criminal network knowledge from large volumes of data.

6. CONCLUSIONS AND FUTURE WORK

Network analysis is important for understanding the structure and organization of criminal enterprises. Advanced, automated techniques and tools are needed to extract knowledge about criminal networks efficiently and effectively. Such knowledge could help intelligence and law enforcement agencies enhance public safety and national security by developing comprehensive disruptive strategies to prevent and respond to organized crimes such as terrorist attacks and narcotics trafficking. We proposed in this research a framework for automated criminal network analysis and visualization that includes the major stages of a network analysis process: network creation, network partition, structural analysis, and network visualization.

Our main contribution is the proposal of a framework to guide knowledge discovery in the criminal network analysis domain. We identified and incorporated in it various techniques to automatically extract valuable criminal network

knowledge from large volumes of data. Most of these techniques originated in other disciplines and initially were not intended for knowledge discovery. For example, the concept space approach, originally designed to generate automated thesauri from textual documents, was used in our framework to identify criminal relationships from crime incident summary data. The blockmodeling approach in SNA research was designed for validating theories of social structures and focused on interactions between “positions” of network members who were similar in social status and roles. We used the blockmodeling approach to extract interaction patterns among criminal groups in which members were closely related.

Based on this framework, we developed a system called *CrimeNet Explorer*, which has structural analysis functionality to detect subgroups, to identify between-group interaction patterns, and to identify central members of subgroups. Quantitative evaluation of the system demonstrated that subjects could achieve significantly higher efficiency with the help of structural analysis functionality than with only network visualization. No significant gain in effectiveness was present, however. Feedback from the subjects and the domain expert showed that CrimeNet Explorer was very promising and could be useful for crime investigation. We plan to incorporate CrimeNet Explorer into the COPLINK research that is seeking to build an integrated environment for information and knowledge management in the law enforcement and intelligence domains [Chen et al. 2003].

In addition, our framework and its associated techniques could be applied to other domains. In Web mining research, for example, our framework could be used to locate high-quality Web pages that are “hubs” or “authorities” [Kleinberg 1999] or to identify Web communities that consist of Web pages created by people with similar interest [Gibson et al. 1998].

On the other hand, CrimeNet Explorer has limitations. First, using the concept space approach to extract criminal relationships is very simplistic. In future implementations we will incorporate other domain knowledge and heuristics such as crime-type-dependent relational strengths. Other data sources such as incident report narratives and phone records can also help generate and validate criminal relationships. These improvements will depend on sophisticated knowledge engineering techniques. Second, our CrimeNet Explorer only considered criminal-criminal relationships. We plan to analyze “criminal activity networks” which consist of not only people but also entities such as locations, weapons, vehicles, and organizations. Third, MDS did not always generate a pleasing network layout. High STRESS values caused by insufficient dimensionality could distort a network and cause nodes to be unevenly distributed on the display. To address this problem, we plan to use a spring embedder algorithm, which may produce nicer renderings of networks.

Other future work could be done to improve the system by adding a time dimension to the system to permit the prediction of how a network is likely to evolve over time. We would also like to apply our framework to other domains such as Web structure mining and organizational research.

An important issue is the legal and ethical dimensions of this type of knowledge discovery application. The potential negative effects of intelligence

gathering and analysis on the privacy and civil liberties of the public have been well publicized [Cook and Cook 2003]. There exist many laws, regulations, and agreements governing data collection, confidentiality, and reporting which could directly impact the development and application of criminal network knowledge discovery. During implementation we carefully excluded all victims from our data sets. As a result, only individuals who were identified as suspects or arrestees in previous crimes were included in the criminal networks. We also suggest that a hypothesis-guided, evidence-based approach be employed when any intelligence gathering system is used. Proper investigative and legal procedures need to be strictly followed. It is neither ethical nor legal to “fish” for potential criminals from diverse and mixed crime, intelligence, and citizen-related data sources.

ACKNOWLEDGMENTS

Special thanks go to Dr. Ronald Breiger from the Department of Sociology at the University of Arizona for his kind help with the initial design of the research framework. We would like also to thank the following people for their support and assistance during the entire project development and evaluation process: Dr. Daniel Zeng, Michael Chau, and other members at the University of Arizona Artificial Intelligence Lab. We also appreciate the critical and important comments and suggestions from personnel from the Tucson Police Department: Lieutenant Jennifer Schroeder, Sergeant Mark Nizbet of the Gang Unit, Detective Tim Petersen, and others.

REFERENCES

- ALDENDERFER, M. S. AND BLASHFIELD R. K. 1984. *Cluster Analysis*. Sage Publications, Beverly Hills, CA.
- ANDERSON, T., ARBETTER, L., BENAWIDES, A., AND LONGMORE-ETHERIDGE, A. 1994. Security works. *Sec. Manage.* 38, 17–20.
- ARABIE, P., BOORMAN, S. A., AND LEVITT, P. R. 1978. Constructing blockmodels: How and why. *J. Math. Psych.* 17, 21–63.
- BAKER, W. E. AND FAULKNER R. R. 1993. The social organization of conspiracy: Illegal networks in the heavy electrical equipment industry. *Amer. Soc. Rev.* 58, 837–860.
- BERKOWITZ, S. D. 1982. *An Introduction to Structural Analysis: The Network Approach to Social Research*. Butterworth, Toronto, Ont., Canada.
- BREIGER, R. L. 2004. The analysis of social networks. In *Handbook of Data Analysis*, M. A. Hardy and A. Bryman, Eds. Sage Publications, London, U.K. 505–526.
- BREIGER, R. L., BOORMAN, S. A., AND ARABIE, P. 1975. An algorithm for clustering relational data, with applications to social network analysis and comparison with multidimensional scaling. *J. Math. Psych.* 12, 328–383.
- BURT, R. S. 1976. Positions in networks. *Soc. Forces* 55, 93–122.
- BURT, R. S. 1980. Models of network structure. *Ann. Rev. Soc.* 6, 79–141.
- CHEN, H. AND LYNCH, K. J. 1992. Automatic construction of networks of concepts characterizing document databases. *IEEE Trans. Syst. Man Cybernet.* 22, 885–902.
- CHEN, H., ZENG, D., ATABAKHSH, H., WYZGA, W., AND SCHROEDER, J. 2003. Coplink: Managing law enforcement data and knowledge. *Commun. ACM* 46, 28–34.
- COOK, J. S. AND COOK, L. L. 2003. Social, ethical and legal issues of data mining. In *Data Mining: Opportunities and Challenges*, J. Wang, Ed. Idea Group Publishing, Hershey, PA, 395–420.
- DANTZIG, G. 1960. On the shortest route through a network. *Manage. Sci.* 6, 187–190.

- DAVIDSON, R. AND HAREL, D. 1996. Drawing graphs nicely using simulated annealing. *ACM Trans. Graph.* 15, 301–331.
- DAY, W. H. E. AND EDELSBRUNNER, H. 1984. Efficient algorithms for agglomerative hierarchical clustering methods. *J. Class.* 1, 7–24.
- DEFAYS, D. 1977. An efficient algorithm for a complete link method. *Comput. J.* 20, 364–366.
- DIJKSTRA, E. 1959. A note on two problems in connection with graphs. *Numer. Math.* 1, 269–271.
- EVAN, W. M. 1972. An organization-set model of interorganizational relations. In *Interorganizational Decision-Making*, M. Tuite, R. Chisholm, and M. Radnor, Eds. Aldine Publishers, Chicago, IL, 181–200.
- FLOYD, R. W. 1962. Algorithm 97: Shortest path. *Commun. ACM* 5, 345–370.
- FREEMAN, L. 1979. Centrality in social networks: Conceptual clarification. *Soc. Netw.* 1, 215–239.
- FREEMAN, L. 2000. Visualizing social networks. *J. Soc. Struct.* 1. (Electronic journal; go to <http://www.heinz.cmu.edu/project/INSNA/joss/index1.html>.)
- GALASKIEWICZ, J. AND KROHN, K. 1984. Positions, roles, and dependencies in a community interorganization system. *Sociolog. Quart.* 25, 527–550.
- GARTON, L., HAYTHORNTHWAITE, C., AND WELLMAN, B. 1999. Studying online social networks. In *Doing Internet Research*, S. Jones, Ed. Sage Publications, London, UK, 75–105.
- GIBSON, D., KLEINBERG, J. M., AND RAGHA-VAN, P. 1998. Inferring Web communities from link topology. In *Proceedings of the 9th ACM Conference on Hypertext and Hypermedia* (Pittsburgh, PA, June), R. Akscyn, Ed. ACM Press, New York, NY, 225–234.
- GOLDBERG, H. G. AND SENATOR, T. E. 1998. Restructuring databases for knowledge discovery by consolidation and link formation. In *Proceedings of 1998 AAAI Fall Symposium on Artificial Intelligence and Link Analysis* (Orlando, FL, Oct.). AAAI Press, Menlo Park, CA.
- HARPER, W. R. AND HARRIS, D. H. 1975. The application of link analysis to police intelligence. *Hum. Fact.* 17, 157–164.
- HAUCK, R. V., ATABAKHSH, H., ONGVASITH, P., GUPTA, H., AND CHEN, H. 2002. Using Coplink to analyze criminal-justice data. *IEEE Comput.* 35, 30–37.
- JAIN, A. K. AND DUBES, R. C. 1998. *Algorithms for Clustering Data*. Prentice-Hall, Upper Saddle River, NJ.
- JAIN, A. K., MURTY, M. N., AND FLYNN, P. J. 1999. Data clustering: A review. *ACM Comput. Surv.* 31, 264–323.
- JOHNSON, S. C. 1967. Hierarchical clustering schemes. *Psychometrika* 32, 241–254.
- JORDAN, P. W. 1998. *An Introduction to Usability*, Taylor and Francis, Bristol, PA.
- KING, B. 1967. Step-wise clustering procedures. *J. Amer. Statist. Assoc.* 62, 86–101.
- KLEINBERG, J. M. 1999. Authoritative sources in a hyperlinked environment. *J. Assoc. Comput. Mach.* 46, 604–632.
- KLERKS, P. 2001. The network paradigm applied to criminal organizations: Theoretical nitpicking or a relevant doctrine for investigators? Recent developments in the Netherlands. *Connections* 24, 53–65.
- KREBS, V. E. 2001. Mapping networks of terrorist cells. *Connections* 24, 43–52.
- KRUSKAL, J. B. 1964. Nonmetric multidimensional scaling: A numerical method. *Psychometrika* 29, 115–128.
- KRUSKAL, J. B. AND WISH, M. 1978. *Multidimensional Scaling*. Sage, Beverly Hills, CA.
- LANCE, G. N. AND WILLIAMS, W. T. 1967. A general theory of classificatory sorting strategies: II. Clustering systems. *Comput. J.* 10, 271–277.
- LORRAIN, F. P. AND WHITE, H. C. 1971. Structural equivalence of individuals in social networks. *J. Math. Soc.* 1, 49–80.
- MCANDREW, D. 1999. The structural analysis of criminal networks. In *The Social Psychology of Crime: Groups, Teams, and Networks*. D. Canter and L. Alison, Eds. Dartmouth Publishing, Aldershot, UK, 53–94.
- MURTAGH, F. 1984. A survey of recent advances in hierarchical clustering algorithms which use cluster centers. *Comput. J.* 26, 354–359.
- PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T., AND FLANNERY, B. P. 1992. *Numerical Recipes in C*, 2nd ed. Cambridge University Press, Cambridge, UK.

- RONFELDT, D. AND ARQUILLA, J. 2001. What next for networks and netwars? In *Networks and Netwars: The Future of Terror, Crime, and Militancy*, J. Arquilla and D. Ronfeldt, Eds. Rand Press, Santa Monica, CA, 311–361.
- ROUSSINOV, D. G. AND CHEN, H. 1999. Document clustering for electronic meetings: An experimental comparison of two techniques. *Decis. Supp. Syst.* 27, 67–79.
- SAETHER, M. AND CANTER, D. V. 2001. A structural analysis of fraud and armed robbery networks in Norway. In *Proceedings of the 6th International Investigative Psychology Conference* (Liverpool, UK, Jan.).
- SAHAMI, M., YUSUFALI, S., AND BALDONADO, Q. W. 1998. SONIA: A service for organizing networked information autonomously. In *Proceedings of the 3rd ACM International Conference on Digital Libraries* (Pittsburgh, PA, June).
- SCOTT, J. 1991. *Social Network Analysis*. Sage Publications, London, UK.
- SIBSON, R. 1973. Slink: An optimally efficient algorithm for the single-line cluster method. *Comput. J.* 16, 30–45.
- SNEATH, P. H. A. 1957. The application of computers to taxonomy. *J. Gen. Microbiol.* 17, 201–226.
- SPARROW, M. K. 1991. The application of network analysis to criminal intelligence: An assessment of the prospects. *Soc. Netw.* 13, 251–274.
- TORGERSON, W. S. 1952. Multidimensional scaling: Theory and method. *Psychometrika* 17, 401–419.
- VOORHEES, E. M. 1986. Implementing agglomerative hierarchical clustering algorithms for use in document retrieval. *Inform. Process. Manage.* 22, 465–476.
- WARD JR., J. H. 1963. Hierarchical grouping to optimize an objective function. *J. Amer. Statist. Assoc.* 58, 236–244.
- WASSERMAN, S. AND FAUST, K. 1994. *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, UK.
- WELLMAN, B. 1988. Structural analysis: From method and metaphor to theory and substance. In *Social structures: A network approach*, B. Wellman and S. D. Berkowitz, Eds. Cambridge University Press, Cambridge, UK, 19–61.
- WHITE, H. C., BOORMAN, S. A., AND BREIGER, R. L. 1976. Social structure from multiple networks: I. Blockmodels of roles and positions. *Amer. J. Soc.* 81, 730–780.
- YOUNG, F. W. 1987. *Multidimensional Scaling: History, Theory, and Applications*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Received February 2003; revised October 2004; accepted February 2005