

Balancing Systematic and Flexible Exploration of Social Networks

Adam Perer, Ben Shneiderman

Abstract— Social network analysis (SNA) has emerged as a powerful method for understanding the importance of relationships in networks. However, interactive exploration of networks is currently challenging because: (1) it is difficult to find patterns and comprehend the structure of networks with many nodes and links, and (2) current systems are often a medley of statistical methods and overwhelming visual output which leaves many analysts uncertain about how to explore in an orderly manner. This results in exploration that is largely opportunistic. Our contributions are techniques to help structural analysts understand social networks more effectively. We present *SocialAction*, a system that uses attribute ranking and coordinated views to help users systematically examine numerous SNA measures. Users can (1) flexibly iterate through visualizations of measures to gain an overview, filter nodes, and find outliers, (2) aggregate networks using link structure, find cohesive subgroups, and focus on communities of interest, and (3) untangle networks by viewing different link types separately, or find patterns across different link types using a matrix overview. For each operation, a stable node layout is maintained in the network visualization so users can make comparisons. *SocialAction* offers analysts a strategy beyond opportunism, as it provides systematic, yet flexible, techniques for exploring social networks.

Index Terms—Social networks, interactive graph visualization, attribute ranking, coordinated views, exploratory data analysis.

1 INTRODUCTION

Understanding networks is an inherently difficult process. It is difficult to visualize, navigate, and most problematic, find patterns in networks. Despite all of these challenges, the network perspective is appealing. Network analysts focus on relationships instead of just the individual elements; how the elements are put together is just as important as the elements themselves. Prior to this perspective, social research focused largely on attributes and neglected the social part of behavior (how individuals interact and the influence they have on each other) [12]. Using techniques from the social network community, analysts can find patterns in the structure, witness the flow of resources through a network, and learn how individuals are influenced by their surroundings.

In practice, a network visualization of a domain can be a messy one, particularly when the network is large. Visualizations are useful to leverage the powerful perceptual abilities of humans, but overlapping links and illegible labels of nodes often undermine this approach. This is not to suggest that researchers studying networks are completely lost. There is a rich history of techniques from sociology to graph theory that allow analysts to find interesting features in networks. Analysts might seek a tight-knit community of individuals, or the gatekeepers between them, or the most centrally powerful entities – and there are a variety of sophisticated algorithms for finding these traits.

More mature fields, such as field biology, have developed systematic methods to train novices and ensure consistency among analysts. The methods are complete and repeatable, so if two analysts are presented with the same data, they should reach the same conclusion. However, in the social networks field, different networks need to be analyzed differently. The spread of an epidemic among villages is not necessarily the same as a spread of a financial crisis on world markets [34]. Since there is no systematic way to

interpret networks, users need to be able to flexibly explore features to discover patterns.

We present *SocialAction*, which we believe is more than just a YASNAT (“yet another social network analysis tool”) because it balances systematic and flexible exploration. To help users systematically examine measures, *SocialAction* applies attribute ranking and coordinated views to identify extreme-valued nodes (Section 3). The main idea is that users can select a structural feature (an interesting ranking criterion that measures an aspect of the network they are interested in), and all nodes will be ranked according to that criterion.

Seo and Shneiderman used a more elaborate strategy, called the rank-by-feature framework to create a successful tool for exploring multidimensional data called Hierarchical Clustering Explorer (HCE) [27]. HCE had coordinated views and ranked 1D and 2D attribute distributions. A series of case studies were conducted which showed that HCE helped biologists understand gene activity from microarray data [28]. Data analysts were able to interpret the many dimensions by selecting a feature that interested them, such as correlation, gaps, or outliers, and finding important data points or clusters.

Inspired by the power of ranking and coordinated views, *SocialAction* enables users to rank nodes using ordered lists and visually coded node-link diagrams (Section 4). When these become too crowded, users can aggregate nodes based on link structure to reduce complexity (Section 5), and isolate types of links (Section 6). *SocialAction* organizes these techniques in a way that enables analysts to systematically explore social networks.

2 RELATED WORK

Freeman suggests that social network analysts seek to uncover two types of patterns in networks: (1) those that reveal subsets of nodes that are organized into cohesive social groups, and (2) those that reveal subsets of nodes that occupy equivalent social positions, or roles [11]. There is a large body of work over the past 60 years to uncover such patterns. *Social Network Analysis: Methods and Applications*, by Wasserman and Faust, is perhaps the most widely used reference book for structural analysts [33]. The book presents a review of network analysis methods and an overview of the field.

- Adam Perer is with the Department of Computer Science, University of Maryland, College Park, MD 20742. E-Mail: adamp@cs.umd.edu.
- Ben Shneiderman is with the Department of Computer Science, University of Maryland, College Park, MD 20742. E-Mail: ben@cs.umd.edu.

Manuscript received 31 March 2006; accepted 1 August 2006; posted online 6 November 2006.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

Visualizations of social networks have been used to aid SNA from the beginning [13]. The visualization of networks is important because it is a natural way to communicate connectivity and allows for fast pattern recognition by humans. However, there are great challenges when visualizing networks [9, 18]. There are many layout algorithms that attempt to calculate the position of each node and the curve of each link to minimize link crossings and adhere to aesthetic principles. These algorithms fall short, however, when the number of nodes is larger than several hundred and the large number of overlapping links makes it hard to judge connectivity [31].

Several approaches attempt to more efficiently use available display space by distorting the graph. Fisheye techniques allow users to examine a focus area in great detail, but also tend to obscure the global structure of networks, e.g. [21, 23]. Multiscale graph abstraction is another technique that preserves global structure, however navigation is difficult because clusters are explicitly contracted and expanded, e.g. [2, 26]. Recent work combines these two approaches with topological fisheye views to reduce the number of displayed nodes while preserving the network structure [14]. Van Ham and van Wijk also combine distortion strategies for highly connected, small-world networks [32].

There are a number of software tools designed to help analysts understand social networks, such as [5, 7, 8]. These tools often feature an impressive number of analysis techniques that users can perform on networks. However, they are also often a medley of statistical methods and overwhelming visual output that leaves many analysts uncertain about how to explore in an orderly manner. SNA is an inherently deductive task, and a user’s exploratory process can be distracted by having to navigate between separate analysis and visualization packages.

Recently, there have been several projects focusing on improving interactive exploration with networks. Among them, *GUESS* is a

novel graph exploration system that combines an interpreted language with a graphical front end [1]. *TreePlus* allows users to explore graphs using more comprehensible enhanced tree layouts [22]. *NetLens* allows users to explore an actor-event network using iterative queries and histograms [19]. Ghoneim et al. presented the promise of using matrix-based visualizations instead of node-link diagrams [15]. *JUNG* is a JAVA toolkit that provides users with a framework to build their own social network analysis tools [25].

3 RANKING AND COORDINATED VIEWS FOR SOCIAL NETWORK ANALYSIS

Numerous measures have been proposed by structural analysts to statistically assess social networks [33]. However, there is no systematic way to interpret networks, as measures can have different meaning in different networks. Analysts want to be certain they are not overlooking critical facets of the network. *SocialAction* uses attribute rankings and coordinated views to help analysts go beyond opportunistic discoveries.

The attribute rankings and coordinated views combines principles of information visualization, such as overview, ordered lists, color coding, and dynamic queries, to allow users to systematically examine the breadth of structural measures on a network. *SocialAction* allows users to change measures rapidly, and the rankings are presented in an ordered list, as well as color in the network visualization. When measures are changed, node positions in the network layout remain the same, so analysts will not lose their orientation. These rankings allow users to gain an overview, filter nodes, and find outliers using ordered lists, scatterplots, and visually-coded node-link diagrams (Section 4).

When networks become large, ordered lists become quite long and network layouts become illegible. *SocialAction* alleviates this problem by allowing users to aggregate nodes based on link

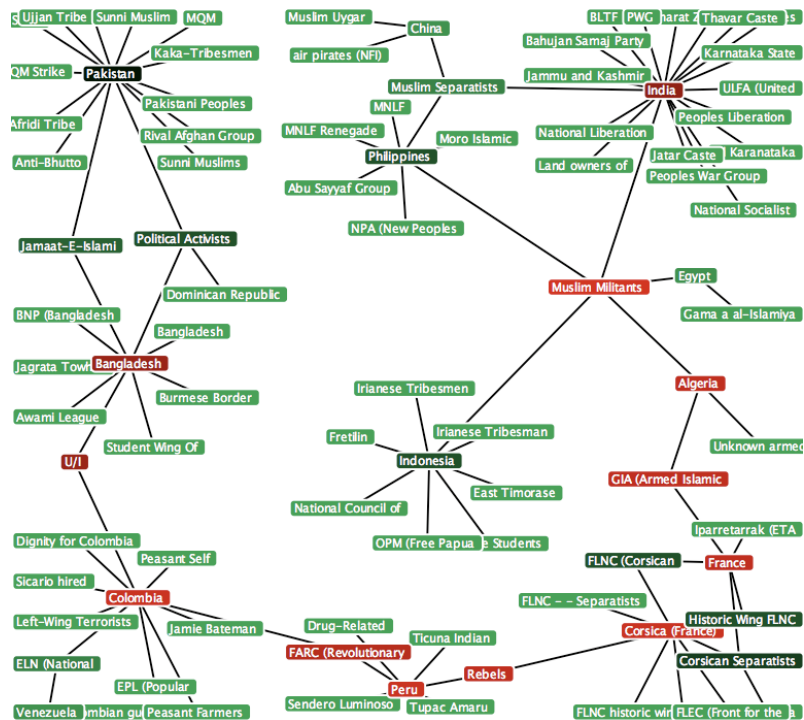
Rankings

Betweenness Centrality

The number of shortest paths between pairs of nodes that pass through a given node.

Rank	Node	Type
2,516.00	Muslim Militants	Terrorist Group
2,436.50	Corsica (France)	Country
2,413.00	Colombia	Country
2,368.00	Peru	Country
2,280.50	France	Country
2,239.00	Algeria	Country
2,226.00	Rebels	Terrorist Group
2,214.00	GIA (Armed Islamic Group)	Terrorist Group
2,124.00	FARC (Revolutionary Armed For...	Terrorist Group
1,718.00	Bangladesh	Country
1,656.00	U/I	Terrorist Group
1,598.00	India	Country
1,063.00	Pakistan	Country
798.00	Corsican Separatists	Terrorist Group
704.00	FLNC (Corsican National Libera...	Terrorist Group
704.00	Historic Wing FLNC	Terrorist Group
637.00	Indonesia	Country
614.00	Political Activists	Terrorist Group
596.00	Philippines	Country
520.00	Jamaat-E-Islami	Terrorist Group
330.00	Muslim Separatists	Terrorist Group
276.00	ELN (National Liberation Army)	Terrorist Group
187.00	Venezuela	Country
187.00	China	Country
94.00	Egypt	Country
0.00	Dignity for Colombia	Terrorist Group
0.00	Jamie Bateman Canon Front	Terrorist Group
0.00	Sendero Luminoso	Terrorist Group
0.00	Jamiat-ul-Mujahideen	Terrorist Group
0.00	Timorese Students	Terrorist Group

(a) Ordered list of 97 nodes in the largest connected component of the terrorism network in 1996. The nodes are ranked according to their betweenness centrality.



(b) Network visualization of the same 97 nodes, colored according to their ranking. The nodes with highest betweenness rankings, sometimes referred to as “gatekeepers”, are painted red.

Figure 1.

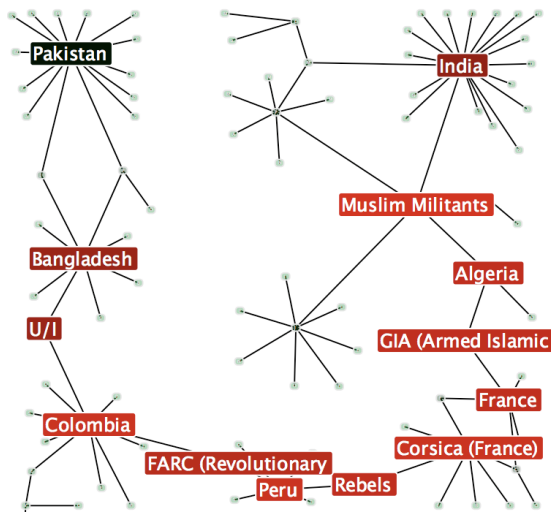


Figure 2. Users can adjust the double range slider to filter nodes that are not of interest. In this graphic, the nodes of Figure 2 that do not have a “betweenness centrality” ranking score of at least 1000 become faded and their labels are removed. The labels of nodes that meet the ranking criteria can be increased by the user. This allows users to focus on the type of nodes they are interested without ignoring the overall structure.

structure. This allows analysts to compress a network or examine communities that are of interest (Section 5).

Some social networks contain multiple types of links. *SocialAction* allows users to systematically iterate through them while maintaining node layout stability. A matrix overview is also provided to help discover patterns across different link types, such as temporal evolution (Section 6).

SocialAction helps structural analysts understand social networks that are difficult to study using current software. The authors have partnered with sociologists early in the design phase to ensure we are meeting the needs of those we wish to aid. One partner is working with a dataset consisting of terrorist groups committing over 70,000 terrorist attacks across the world spanning 27 years. This network is being assembled by the Center of Excellence for Behavioral and Social Research on Terrorism and Counter-Terrorism, with the goal of developing strategies for disrupting the formation of terror networks and minimizing the impact of future attacks. We illustrate our techniques using this network, to suggest how our approach applies to real data.

4 RANKING SOCIAL NETWORKS

The Visual Information Seeking Mantra [29] suggests that designers present users with an overview first. The entire social network is visualized as a node-link diagram on the left, and a corresponding list of nodes is presented on the right. The nodes are positioned using a force-directed layout approach, a generic layout algorithm common in many network visualization packages [9]. Users can select nodes in either the ordered list or node-link diagram, and they will be highlighted in both views. In some cases, such as very small networks, this display alone may allow the users to make discoveries. Many networks tend to be larger, such as our example, and this is where the attribute rankings comes into play.

According to our interviews, the methodology of social network analysts suggests to ignore the individual attributes of nodes and instead look at their structural attributes for meaning. Nodes can have meaning derived from their position in the network, as nodes can be isolated or connected to many other nodes. *SocialAction* allows users to rank nodes by their structural position by choosing a ranking of interest from a drop-down menu. Sample choices are:

- *bary center*: the total shortest path of a node to all other connected nodes [25]
- *betweenness centrality*: how often a node appears on the shortest path between all other nodes [6]
- *closeness centrality*: how close a node is compared to all other nodes [10]
- *cut-points*: the subgraph becomes disconnected if this node is removed
- *degree*: the number of nodes a node is connected to. (For directed graphs, rankings based on in- and out-degree also exist)
- *HITs*: a “hubs and authorities” importance measure [20]
- *power-centrality*: how linked a node is to the rest of the network [4]

When users select a ranking, all of the nodes are ranked according to this criterion in the ordered list. Each ranking is assigned a corresponding color, ranging from green to black to red, based on its value. This helps illustrate each node’s position among all ranked entities. The network visualization also paints each node with this color.

Figure 1 illustrates *SocialAction*’s technique on a subgraph from the global terrorism network. This network is two-mode, which mean it has two different types of nodes: terrorist groups and countries. In this network, a link exists if a terrorist group committed an attack in that country. In this example, betweenness centrality was selected as the ranking criterion. This feature is often used to detect “gatekeepers” between highly connected subgraphs. The nodes are ordered by ranking in Figure 1a. In this example, the “Muslim Militants” group has the highest score. The network visualization, colored according to these rankings, is shown in Figure 1b. Groups with high betweenness appear red in this network.

4.1 Filtering by Rankings

In line with the next step of the Visual Information Seeking Mantra [29], *SocialAction* allows users to zoom and filter, since users’ perceptual clarity improves when the number of visualized elements is limited [18]. Users can freely zoom into sections of the graph to improve clarity by dragging the right-mouse button. *SocialAction* also allows users to filter the nodes in both the ordered list and the network view based on their rankings.

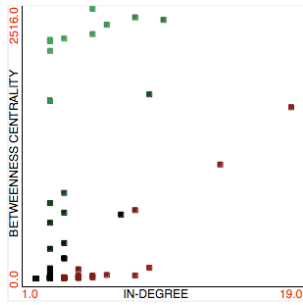
Users can dismiss portions of the network that do not meet their criteria using a double range slider. Users are also able to use the filter to fade the nodes to keep the networks full structure intact (Figure 2). In this example, a user selected the betweenness centrality measure and the left bar of the range slider was dragged to the right until it reached the value of 1000. All nodes that do not have a betweenness centrality measure of at least 1000 are faded and are no longer labeled. The nodes that meet the criteria are now prominently displayed with larger labels and allow users to focus on them. If the users believe the remaining nodes are still a distraction, the user can have them removed entirely.

By allowing users to filter based on rankings that are important to them, the network becomes more manageable in terms of legibility, as the number of nodes and link crossings will be reduced. It also allows users to spot the phenomena of interest across an entire network.

4.2 Comparing Rankings with Scatterplots

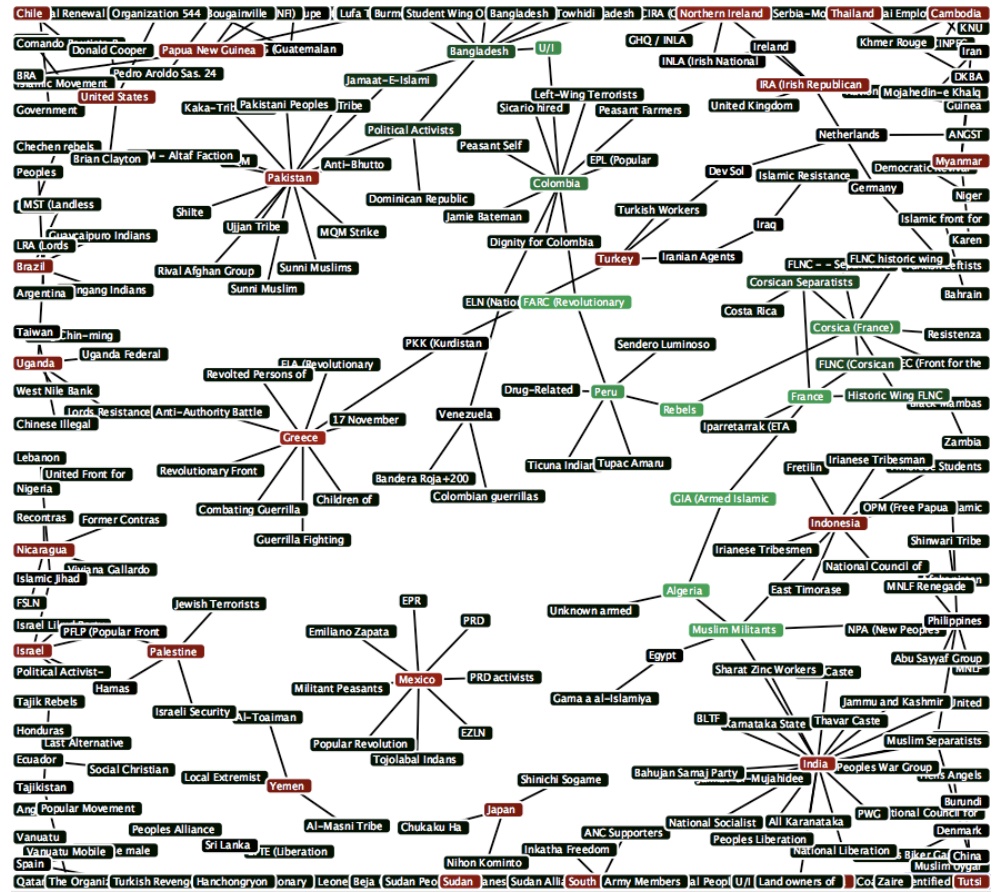
Structural analysts may be interested in the nodes that meet criteria across two rankings. *SocialAction* presents this two-dimensional projection as a scatterplot. A scatterplot reveals the form, direction and strength of a relationship between two features, in addition to identifying outliers easily. Users can select two features that form the axes for a scatterplot.

A scatterplot of 276 nodes is shown alongside the network it represents (Figure 3). Users can select any of the ranking features to be the horizontal and vertical axes. In this example, the horizontal axis is in-degree and the vertical axis is betweenness centrality.



(a) Scatterplot plot of 276 nodes

Figure 3. *SocialAction* allows users to rank nodes by two different features in a scatterplot. The colors of nodes in the network visualization are determined by the scatterplot position. This allows users to find nodes exhibiting characteristics they seek, as well as outliers. For instance, nodes with low-degree but high betweenness centrality are colored bright green. These nodes can be quickly spotted even in the otherwise unkempt network visualization.



(b) Network visualization's nodes colored by scatterplot position

SocialAction fits the scatterplot to a linear function. All nodes that appear above this linear cross-section are shaded from black to green, and those below are shaded from black to red. The nodes in the network visualization are painted using the same palette. The views are coordinated, so when a user selects a node in the scatterplot, the corresponding node also becomes highlighted in the network visualization.

Using this scatterplot, users can quickly spot nodes of interest. For instance, suppose an analyst was seeking nodes with low degree (committed attacks in few countries) but high betweenness centrality. The nodes would appear in the upper-left of the scatterplot (Figure 3a). These nodes are also easily detected in the network visualization according to their bright green color (Figure 3b).

5 AGGREGATE RANKINGS FOR COHESIVE SUBGROUPS

When networks become large, presenting rankings for every node and link is not ideal. Ordered lists get quite long and scatterplots become crowded. Filtering by ranking is one solution to this problem but it is not perfect because it ignores nodes that do not meet a certain criteria. *SocialAction* offers subgroup detection to isolate a group of nodes based on their structural properties. In fact, one of the main goals of sociologists studying social networks is to find cohesive subgroups of nodes [11].

There are a variety of techniques to detect subgroups. For disconnected graphs, a subgroup could be defined as each connected component. However, in practice, components are often large and

need to be partitioned themselves into local “communities” of tightly-knit nodes. *SocialAction* includes choices, including a feature to automatically determine communities based on link structure. Newman’s community identification algorithm [24] was chosen because it is fast enough to support interactive real-time adjustments. The implementation of this algorithm in [17] was used. This algorithm was also demonstrated effective in isolating subgroups of personal online social networks in *Vizster* when combined with a slider [16]. Since the algorithm may identify communities at an undesirable granularity, users can move the slider to adjust the state of clustering. Like *Vizster*, *SocialAction* visually presents the communities by surrounding all members with a translucent convex hull. In this example, the color of the convex hull represents the maximum ranking of any of its entities, so communities containing nodes with high betweenness are red. Users can optionally color the communities by minimum and average ranking values, or opt for each community to be assigned an arbitrary, unique color. By default, communities are labeled with a unique integer but users can rename the labels to have more semantic meaning.

In Figure 4a, community detection is enabled on the 97-node network presented earlier in Section 4. The algorithm finds nine communities. In addition to detecting subgroups, *SocialAction* allows users to use the subgroup information to improve clarity in two ways.

The first technique is presented in Figure 4b. If the user is interested in an overview of the structure, users can collapse a subgroup into a single meta-node (linked with meta-edges).

This meta-node, representing the entire subgroup, will be positioned in the center of where the subgroup previously existed. The node's size will be in proportion to the number of nodes it contains. Similarly, the size of the meta-link between nodes will be proportional to the number of links between the groups. The ranking panels (ordered list, scatterplots) treat each subgroup as one entity, and users can search for patterns using the compressed subgroups.

If users are interested in local structure, subgroups can be analyzed in isolation. The system can treat the subgroup as if it is the entire graph, and all ranking panels will be updated accordingly. Further aggregation can be performed on this subgroup, as well. Figure 4c is the result of a user choosing the third community in the upper center of Figure 4a, and then further dividing the community into two subgroups.

After users finish exploring subgroups using either of these techniques, the users can return to the original graph, and all nodes that reappear will keep the position they held when they initially disappeared.

6 MULTIPLEX RANKINGS

When social networks have multiple link types, they are often referred to as multiplex networks. For instance, in the terrorism network, nodes can be connected based on if they committed a terrorist attack in the same area, or used the same weapons, or if they come from the same region. Links can also have temporal characteristics; a link could represent an attack in a certain year. The types of links used depend on what types of questions the analyst is trying to answer.

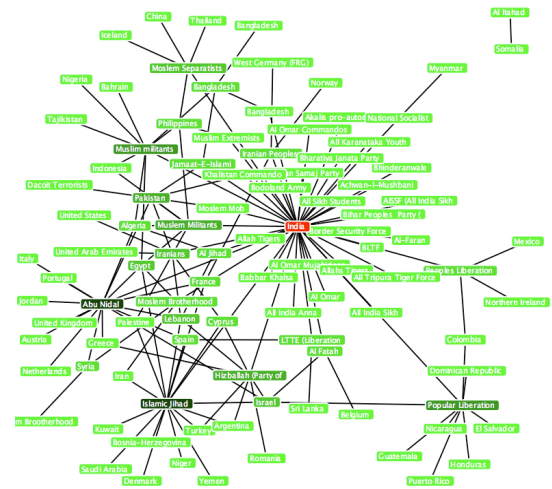
Often, a network will look drastically different based on which types of links are shown (Figure 5). The top image shows all links. The middle shows links between terrorist groups and countries based on if they attacked in the year 1988. The bottom shows links if they attacked in 1989. *SocialAction* allows users to quickly iterate between networks of different link types while keeping the layout stable. In this example, a force-directed layout was used based on the network structure with all links present. Since the layout was not optimized for the individual years, users have the choice to leave them in this position, or have the layout update with smooth animation to reduce the number of edge crossings.

In these examples, the nodes are colored based on their degree ranking. Nodes that do not have any links of the selected type are faded. Users can increase the legibility of nodes with links by increasing their font size with a slider. For instance, 'LTTE' is an active terrorist group in 1988 (attacking India and Spain and painted red), but fades in 1989 due to a lack of activity.

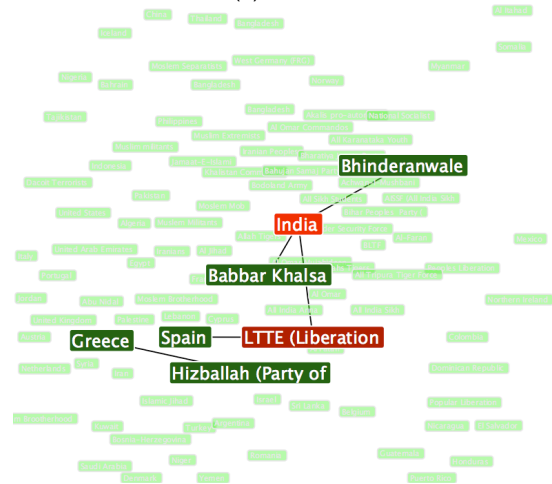
SocialAction also offers a matrix summary so users can spot patterns across many different link types at once (Figure 6). Each node occupies a row, and each column represents a different link type. Each cell is colored based on a node's ranking when only that link type is present. In Figure 6, degree was the selected ranking criterion and the rows are sorted in descending order by degree when all link types are present. For this subgraph, India has the highest degree (most terrorist attacks) when all links are present as well as most years from 1980-1997, as those cells are colored bright red. Countries such as Lebanon, Egypt and Pakistan are dominant in years that India is not. *SocialAction* allows users to flexibly explore multiplex networks. Users can iterate through different link types separately and apply the ranking and aggregation techniques as well (Sections 4,5). Users can also spot patterns across link types using the matrix overview.

7 SUPPORTING SYSTEMATIC EXPLORATORY DATA ANALYSIS

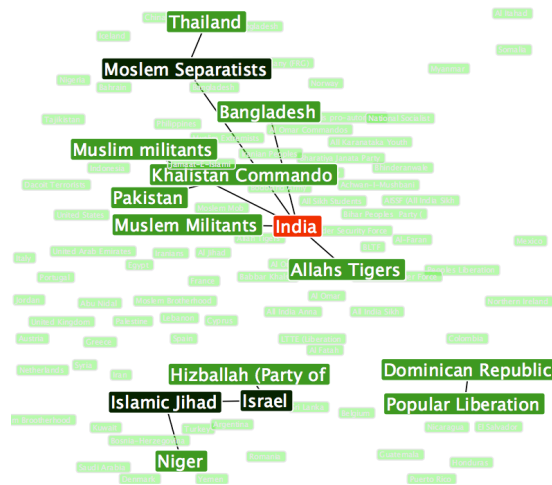
So far, we have described techniques of importance to the SNA community: ranking nodes, finding cohesive subgroups and exploring multiple link types. In addition to providing these features, we also designed the interface to support orderly exploratory data analysis. Users can systematically iterate through the network measures available to examine the range of structural



(a) All links



(b) Links active in 1988



(c) Links active in 1989

Figure 5. Users can iterate through each type of links to find patterns in individual years. The nodes are colored according to their degree ranking, and correspond to the matrix in Figure 7. Nodes that do not have any active ties in a particular time period are faded to improve legibility of the active nodes.

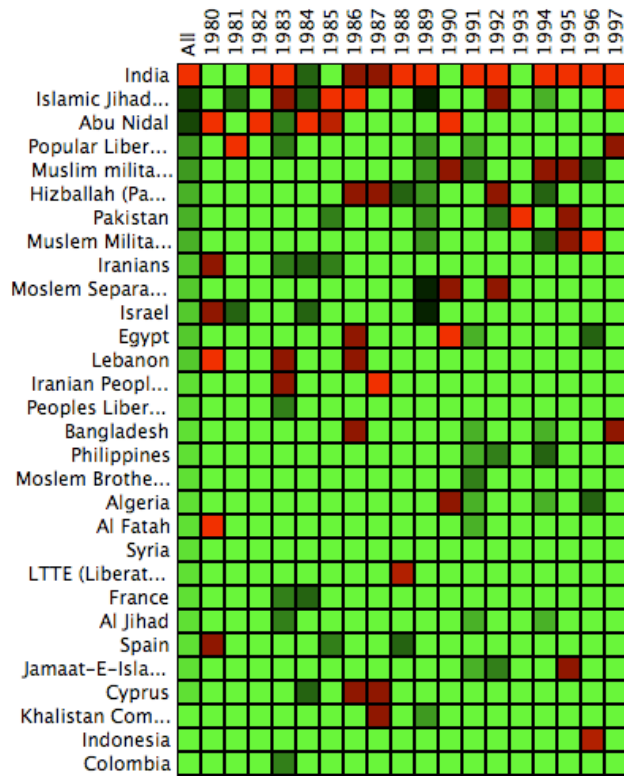


Figure 6. The overview matrix's rows are the top 30 nodes with the highest degree (when all ties are present), sorted in descending order. Each column represents a different type of link present. For each cell, a greater intensity of red implies relative higher degree, whereas greater intensity of green implies lower degree.

properties. The spatial layout of the node-link diagram remains unchanged during this process to preserve the user's mental model of the network. If the network is too large to effectively deal with the entire network, users can systemically iterate through each subgroup and apply the network measures to these smaller groups in order. Finally, if a network is multiplex, a user can systematically iterate through each link type while being provided a matrix overview. Users have the freedom to apply specific techniques to support their hypotheses. However, if they are interested in exploratory data analysis and want to examine the full range of measures, our interface provides an orderly, systematic method for doing so.

8 IMPLEMENTATION

Our system is implemented in Java and integrates several open-source toolkits to take advantage of other contributions. *JUNG* [25] provides our underlying node-link data structures, as well as an implementation of some of the SNA ranking algorithms. *Prefuse* [17] is used for the network visualizations and community algorithm. Finally, *Piccolo* [3] is used for the scatterplot and matrix visualizations.

All of the ranking features perform in real-time for the networks used in this paper. For multiplex networks, the matrix overview can cause a small delay if there are many multiple link types. The authors have tested *SocialAction* with networks of approximately 1000 nodes and there are minimal performance disruptions on a PowerBook G4.

The terrorist network used as a running example was an undirected, multiplex, two-mode network. These attributes are not mandatory, as *SocialAction* can also handle single-mode, non-multiplex, and directed graphs. Networks can be easily imported

into *SocialAction* from a variety of popular formats, such as Pajek and GraphML.

9 DISCUSSION AND FUTURE WORK

Throughout the paper, *SocialAction* has been demonstrated on a network of terrorist groups. This system was not just designed for this domain, but instead to assist analysts struggling to understand networks of interest to them. Network analysis is a serious research effort and cannot be replicated easily in the form of small user studies. The authors believe a series of longitudinal case studies is the most effective way in evaluating our system [30].

So far, two case studies with sociologists have been initiated to begin the evaluation of *SocialAction*. In initial meetings, there was great deal of encouraging feedback when *SocialAction* was introduced. The sociologists have been using tools common in among analysts in their field [5, 8]. When seeing our tool in action, they began to think of ways to analyze their data that they had not previously thought of before. *SocialAction's* highly interactive interface and the systematic guidance encouraged them to explore their data sets in novel ways.

Social networks can range from only a few nodes to a much larger number of nodes. Although our techniques have been tested on networks with thousands of nodes, presenting rankings for each node and link may not be appropriate when the network size drastically increases. The sociologists acting as our design partners have been studying networks of modest size (300-3000 nodes). We will initiate future case studies with analysts studying larger networks to see if our techniques are applicable to their problems, as well. As network size grows different measures may become important, algorithm speed becomes an issue, and screen management is more complex. However, extraction of subgraphs, selection of only key nodes, and aggregation strategies are promising methods.

10 CONCLUSION

Understanding networks is a difficult task. Sociologists have been developing techniques since the 1930s in order to tackle this problem. Despite this lengthy history, there are few tools to support exploratory data analysis. *SocialAction* is a new system that strongly links together the statistical and visual components of a network. However, this system is more than just a bag of features. *SocialAction* is a goal-directed system designed to help users make discoveries. The attribute ranking interface allows users to explore different SNA measures in a systematic yet flexible way to gain overviews, filter nodes, find outliers and visually code the network visualization. Users can interactively aggregate nodes to reduce complexity, find cohesive subgroups, and focus on communities of interest. Users can untangle a multiplex network and use an overview to spot trends across different link types, such as temporal evolution. The network layout is kept stable after each action so users can perceive patterns.

These tasks were carefully chosen to form the foundation of our vision. The goal of *SocialAction* is to mold the currently opportunistic techniques of SNA into a systematic and flexible interface. Fields need systematic techniques to ensure consistency among analysts. *SocialAction* aims to provide the social network community just that.

ACKNOWLEDGEMENTS

We thank Dr. Alan Neustadtl, a professor of sociology, who has been extremely helpful throughout the design of our system. Sociology graduate student Brian Reed also helped clarify needs of social network analysts. We would also like to thank Ben Bederson, Mustafa Bilgic, Lise Getoor, François Guimbretière and Catherine Plaisant for their helpful comments on early drafts. We would also like to thank the Center of Excellence for Behavioral and Social Research on Terrorism and Counter-Terrorism at the University of

Maryland for use of their dataset for examples in this paper. Adam Perer was supported by the Joint Institute for Knowledge Discovery at the Univ. of Maryland.

REFERENCES

- [1] E. Adar, "GUESS: A Language and Interface for Graph Exploration", ACM Conference on Human Factors in Computing Systems, 2006.
- [2] D. Auber, Y. Chiricota, F. Jourdan and G. Melancon, "Multiscale Visualization of Small World Networks", IEEE Symposium on Information Visualization, pp. 75-81, 2003.
- [3] B. B. Bederson, J. Grosjean and J. Meyer, "Toolkit Design for Interactive Structured Graphics", IEEE Transactions on Software Engineering, 30, pp. 535-546, 2004.
- [4] P. Bonacich, "Power and centrality: A family of measures." American Journal of Sociology, 92, pp. 1170-1182, 1987.
- [5] S. Borgatti, M. G. Everett and L. C. Freeman, *UCINET 6*, Analytic Technologies, 2006.
- [6] U. Brandes, "A Faster Algorithm for Betweenness Centrality", Journal of Mathematical Sociology, 25, pp. 163-177, 2001.
- [7] U. Brandes and D. Wagner, "visone - Analysis and Visualization of Social Networks", Graph Drawing Software, in M. Junger and P. Mutzel, eds., Springer-Verlag, 2003.
- [8] W. de Nooy, A. Mrvar and V. Batageli, *Exploratory Social Network Analysis with Pajek*, Cambridge University Press, Cambridge, 2005.
- [9] G. Di Battista, P. Eades, R. Tamassia and I. G. Tollis, *Graph Drawing: Algorithms for the Visualization of Graphs*, Prentice Hall, New Jersey, 1999.
- [10] L. C. Freeman, "Centrality in social networks: Conceptual clarification", *Social Networks*, 1, pp. 215-239, 1979.
- [11] L. C. Freeman, "Graphic Techniques for Exploring Social Network Data", *Models and Methods in Social Network Analysis*, in P. J. Carrington, J. Scott and S. Wasserman, eds., Cambridge University Press, Cambridge, 2004.
- [12] L. C. Freeman, *The Development of Social Network Analysis: A Study in the Sociology of Science*, Empirical Press, 2004.
- [13] L. C. Freeman, "Visualizing Social Networks", *Journal of Social Structure*, 2000.
- [14] E. R. Gansner, Y. Koren and S. North, "Topological Fisheye Views for Visualizing Large Graphs", IEEE Transactions on Visualization and Computer Graphics, 11, pp. 457-468, 2005.
- [15] M. Ghoniem, J.-D. Fekete and P. Castagliola, "A Comparison of the Readability of Graphs Using Node-Link and Matrix-Based Representations", IEEE Symposium on Information Visualization, 2004.
- [16] J. Heer and d. boyd, "Vizster: Visualizing Online Social Networks", IEEE Symposium on Information Visualization, 2005.
- [17] J. Heer, S. K. Card and J. A. Landay, "prefuse: A Toolkit for Interactive Information Visualization", ACM Conference on Human Factors in Computing Systems, 2005.
- [18] I. Herman, G. Melancon and M. S. Marshall, "Graph visualization and navigation in information visualization: A survey", IEEE Transactions on Visualization and Computer Graphics, 6, pp. 23-43, 2000.
- [19] H. Kang, C. Plaisant, B. Lee and B. B. Bederson, "NetLens: Iterative Exploration of Content-Actor Network Data", IEEE Symposium on Visual Analytics Science and Technology, 2006.
- [20] J. Kleinberg, "Authoritative sources in a hyperlinked environment." 9th ACM-SIAM Symposium on Discrete Algorithms, 1998.
- [21] J. Lamping and R. Rao, "The hyperbolic browser: A Focus+Context Technique for Visualizing Large Hierarchies", *Journal of Visual Languages and Computing*, 6, 1995.
- [22] B. Lee, C. S. Parr, C. Plaisant, B. B. Bederson, V. D. Veksler, W. D. Gray and C. Kotfila, "TreePlus: Interactive Exploration of Networks with Enhanced Tree Layouts", IEEE Transactions on Visualization and Computer Graphics, 2006.
- [23] T. Munzner, "H3: Laying Out Large Directed Graphs in 3D Hyperbolic Space", IEEE Symposium on Information Visualization, pp. 2-10, 1997.
- [24] M. E. J. Newman, "Fast algorithm for detecting community structure in networks." *Physical Review E*, 69, 2004.
- [25] J. O'Madadhain, D. Fisher, P. Smyth, S. White and Y.-B. Boey, "Analysis and Visualization of Network Data using JUNG", *Journal of Statistical Software*, VV, 2005.
- [26] G. Parker, G. Franck and C. Ware, "Visualization of Large Nested Graphs in 3D: Navigation and Interaction", *Journal of Visual Languages and Computing*, pp. 299-317, 1998.
- [27] J. Seo and B. Shneiderman, "A Rank-by-Feature Framework for Interactive Exploration of Multidimensional Data", *Information Visualization*, 4, pp. 99-113, 2005.
- [28] J. Seo and B. Shneiderman, "Knowledge Discovery in High Dimensional Data: Case Studies and a User Survey for the Rank-by-Feature Framework", IEEE Transactions on Visualization and Computer Graphics, 12, pp. 311-322, 2006.
- [29] B. Shneiderman, "The Eyes Have It: A Task by Data Type Taxonomy for Information Visualization", *Visual Languages*, pp. 336-343, 1996.
- [30] B. Shneiderman and C. Plaisant, "Strategies for Evaluating Information Visualization Tools: Multi-dimensional In-depth Long-term Case Studies", *Beyond time and errors: novel evaluation methods for Information Visualization*, Workshop of the Advanced Visual Interfaces Conference, 2006.
- [31] F. van Ham, *Interactive Visualization of Large Graphs*, Technische Universiteit Eindhoven, 2005.
- [32] F. van Ham and J. J. van Wijk, "Interactive Visualization of Small World Graphs", IEEE Symposium on Information Visualization, 2004.
- [33] S. Wasserman and K. Faust, *Social Network Analysis: Methods and Applications*, Cambridge University Press, 1994.
- [34] D. J. Watts, *Six Degrees: The Science of a Connected Age*, W.W. Norton & Company, New York, 2003.