NetVisia: Heat Map & Matrix Visualization of Dynamic Social Network Statistics & Content

Robert Gove§†, Nick Gramsky§, Rose Kirby§, Emre Sefer§, Awalin Sopan§†, Cody Dunne§†, Ben Shneiderman§† and
Meirav Taieb-Maimon†‡*

§Department of Computer Science, †Human-Computer Interaction Lab
University of Maryland, College Park, MD, USA
‡Ben-Gurion University of the Negev, Beer-Sheva, Israel
{rpgove, ngramsky, rkirby, esefer, awalin, cdunne, ben}@cs.umd.edu, meiravta@bgu.ac.il

Abstract—Visualizations of static networks in the form of node-link diagrams have evolved rapidly, though researchers are still grappling with how best to show evolution of nodes over time in these diagrams. This paper introduces NetVisia, a social network visualization system designed to support users in exploring temporal evolution in networks by using heat maps to display node attribute changes over time. NetVisia's novel contributions to network visualizations are to (1) cluster nodes in the heat map by similar metric values instead of by topological similarity, and (2) align nodes in the heat map by events. We compare NetVisia to existing systems and describe a formative user evaluation of a NetVisia prototype with four participants that emphasized the need for tooltips and coordinated views. Despite the presence of some usability issues, in 30-40 minutes the user evaluation participants discovered new insights about the data set which had not been discovered using other systems. We discuss implemented improvements to NetVisia, and analyze a co-occurrence network of 228 business intelligence concepts and entities. This analysis confirms the utility of a clustered heat map to discover outlier nodes and time periods.

Keywords-Dynamic networks; information visualization; network evolution; network visualization; social networks; user interfaces

I. Introduction

In recent years, visualization techniques for the analysis of network evolution have lagged behind the rapid growth of social media and electronic databases. Few tools currently exist for the visualization of networks that change over time, often termed "dynamic" or "temporal" networks. Researchers are interested in questions about metrics that measure different statistical elements about networks. Potential questions include: Which people have been in a social network for the longest time? For any given time period, who was connected to the most other people? Who were the most connected people in their third year after being added to the network? Analysis of temporal network evolution could yield important insights about the world's increasingly ubiquitous networks.

Typical approaches to visualizing temporal evolution involve node-link diagrams that morph or change as the state of the network varies over a period of time. Other researchers have tried a matrix visualization approach, which can show network metrics explicitly. However, these solutions alone are not adequate for answering all types of questions regarding temporal changes to network metrics as they do not show all time slices simultaneously.

This paper introduces NetVisia, a novel system using heat maps and matrices to visualize network node attributes over time. NetVisia's main contributions are heat maps that can be clustered according to node attributes and aligned by temporal events. Section IV describes a formative user evaluation with four participants who performed an insight-based exploration task and two timed tasks. Overall, participants were able to use NetVisia to discover insights about social network data despite some usability problems. Section V further demonstrates NetVisia's utility in a case study analyzing a business community's evolving patterns of interest.

II. Related Work

Node-link diagrams are the most common social network visualizations, used by more than 30 of the first 100 projects on VisualComplexity.com [1]. They present a topological view of a network where nodes appear as points connected by lines representing their relationships. Often the nodes are positioned using force-directed layouts like Fruchterman-Reingold [2]. One basic way of showing dynamic networks is to create a supergraph from all time slices, lay out the supergraph, and use dynamic filters to interactively flip through time slices. An example is TempoVis [3], which also highlights changes between time slices. The generated supergraph can be quite large, resulting in much unused space in each graph and poor layout readability. Alternatively, we can lay out each slice separately and intelligently animate the transitions between them [4]. This approach is computationally inefficient and, more importantly, causes significant user effort to track changes.

Two complementary solutions are to animate and highlight changes between graphs and to minimize the changes as much as possible [5]. One way to do both is to use a continuous, animated force-directed layout as in SocialAction [6], where modifications to the graph structure with a time slider are immediately reflected in the layout. These layouts are unstable though, and slight modifications can lead to different equilibriums. Foresighted Layout [7] considers all time slices simultaneously, creating a supergraph optimized

by the live times of individual nodes and edges. They provide a configurable trade-off between readability and stability, and dynamic filters or animations are used to move between individual time slices. The main difficulty with using node-link diagrams is that users must view visualizations of each time step to track changes. Node-link Difference Maps [8] can show a comparison of two time slices but not more.

Matrix representations of the network like MatLink [9] can show an overview or single time slice of the network. Matrix views can be extended with glyphs inside the matrix cells, like POVNA's time series space-filling curves [10]. TimeMatrix [11] presents a more scalable matrix approach to dynamic network visualization using bar chart and matrix glyphs. It uses several interaction techniques for showing aggregate information, including semantic zooming, hierarchical aggregation, glyph overlays, and dynamic filters. These can provide several high-level overviews of social networks.

Instead of showing several slices of a network in two dimensions, 3D and 2.5D approaches try to show slices stacked on top of each other with width or color showing changes over time [12]. Another way to show changes over time is to connect the slices with colored or sized links [13]. However, these approaches suffer from perspective problems and occlusion. Hadlak et al. [13] provide an interesting alternate visualization, which is a time hierarchy embedded inside regions of a map using a point-based layout. Like the space-filling curves in the matrix cells, these embedded hierarchies may not scale well.

ManyNets [14] utilizes tabular views to compare time slices of networks by their network statistics, with each slice a separate row. Instead of disjoint time slices, Falkowski et al. [15] use overlapping slices and cluster nodes in each slice by the topology. They show line charts for the temporal development of a selected cluster using metrics for stability, density, cohesion, Euclidean distance, correlation coefficient, and group activity. They then find clusters that have overlapping membership, which are then clustered to find related communities. Their examples show communities split by major temporal events, but do not show long-term evolution well. Rosvall and Bergstrom found more compelling clusters by combining bootstrap resampling with significance clustering, then visualizing the changes with alluvial diagrams [16]. The ribbons in the alluvial diagrams represent clusters over time and show movement between them like mergers and divergences.

Clustered heat maps have also been used to give overviews of attribute distributions [17], with clusters computed by similar attributes instead of by topology. Sopan et al. conducted a usability study which suggested that heat maps are beneficial for showing an overview of the data and identifying outliers and clusters of similar data.

III. NetVisia: Visualization of Network Evolution

In their taxonomies of network evolution analysis tasks, both Yi et al. [11] and Ahn et al. [18] list three entities of interest: (1) analysis of change at global level, (2) analysis of temporal changes within subgroups, and (3) analysis of temporal associations between nodes and edges. These tasks were considered as NetVisia was designed and implemented. NetVisia provides an analysis of the network at the global level by utilizing the Workspace Window, which can cluster and aggregate nodes to better analyze related nodes and cliques. Finally, the node-to-node adjacency matrices let users analyze associations between nodes and edges in a given time period (or periods if time binning is used).

NetVisia is a system built in Java that uses matrices and heat maps to visualize networks over discrete time periods. NetVisia has four components to view and control the network: (1) a heat map for analyzing specific social network metrics over all time periods, (2) previews of the workspace heat maps, (3) a control panel for controlling the views of the data, and (4) a pair of time-specific views that provide node details for selected time periods. These are shown in Fig. 1.

A social network analyst at the University of Maryland evaluated a paper prototype of NetVisia. The analyst's comments affirmed the interface design and provided feedback on system terminology and system capabilities.

Workspace Window. The Workspace Window is the main working area. The network across all time points is presented in a heat map: Columns on the horizontal axis are time points, and rows on the vertical axis are nodes. Data is visualized through the heat map by varying the color intensity to indicate the node's metric value during each time period. Black indicates the node is not present in the network during the time period represented by that column. For example, for the degree metric, black indicates that the node does not exist in the network during that time period. Gradients vary from white to red indicating the variation in values from low to high; for the degree metric, white shows low degree while red demonstrates high degree. The sharp contrast between black and the color scheme allows users to see where nodes were active at any given time and the change in gradient color across time periods visually indicates how nodes' metric values changed over time.

The two Time Details Views at the bottom can show matrices which capture the edges and their respective metrics. Using tools in the Control Panel to the right, users can arrange, filter or bin according to the time the nodes appear in the Workspace Window. The filters are visual filters so that metric values do not change after users apply filtering.

Preview Window. Because users can only view one metric at a time in the Workspace Window, it might not be obvious which metric could be "interesting." The Preview Window allows users to see small previews of each network metric heat map, which can guide users to metrics with interesting patterns and evolution. Clicking on a particular heat map populates the Workspace Window with that metric or attribute.

Control Panel. The Control Panel allows users to control the views by clustering, aggregating, binning, filtering, or zooming in on the heat map. The "Cluster nodes" checkbox clusters the nodes using the Bar-Joseph et al. algorithm [19] (see Fig. 2). Clustering is based on attribute similarity across all time periods in the network, so that nodes with similar values
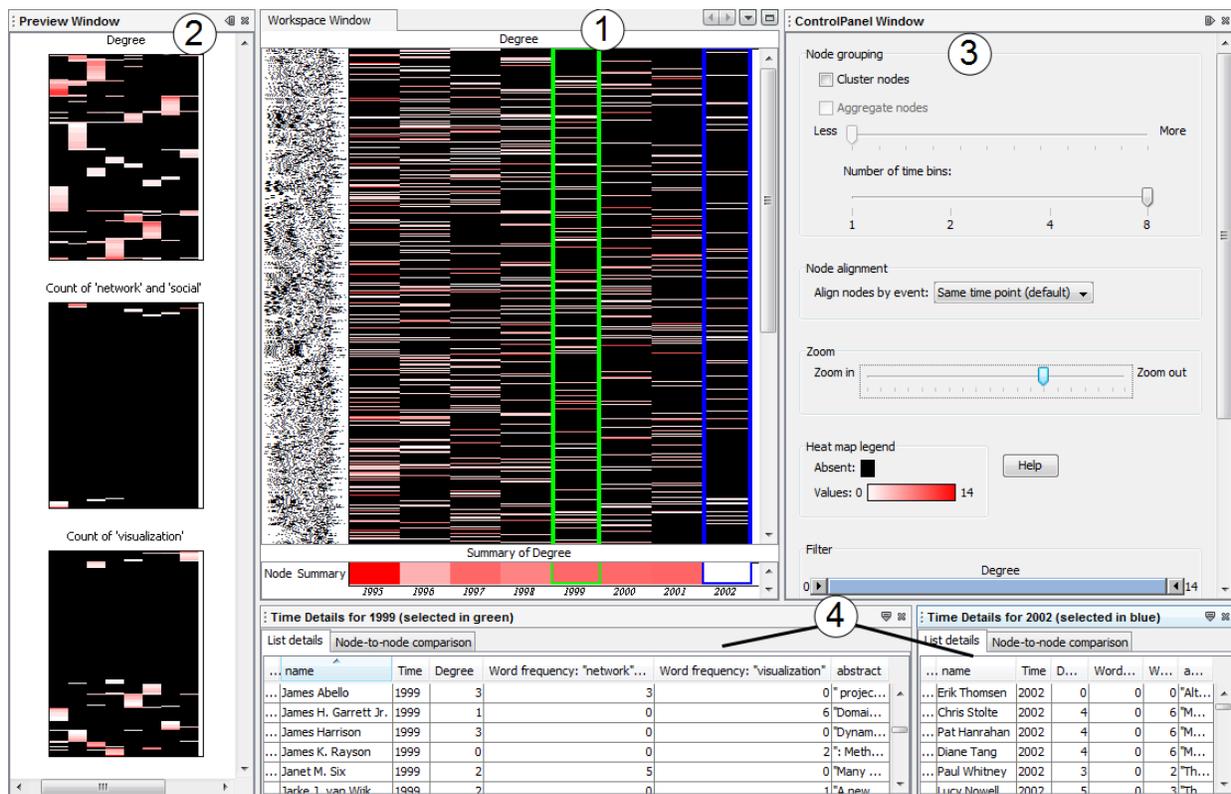
Fig. 1: NetVisia: (1) Workspace Window, (2) Preview Window, (3) Control Panel, (4) Time Details views. This figure shows co-authorship network from the InfoVis 2004 contest data set.

across time will be rearranged to appear next to each other. The default "unclustered" view sorts the nodes alphabetically. The zoom slider allows users to zoom in to have a closer view of particular nodes or zoom out to observe the entire network.

In order to improve readability, NetVisia can aggregate similar nodes and collapse them into a single row after the nodes have been clustered. The height of the new row is logarithmically proportional to the number of nodes used to form the new row, which indicates how many nodes are aggregated in the row while reducing the amount of screen space used by the heat map. Users adjust the aggregation threshold by changing the node aggregation slider, and nodes collapse accordingly.

By default, each column represents one time period. NetVisia can merge adjacent columns by binning time periods into one column. Merging makes the Workspace heat map more compact if users want to see a coarse-grained view.

The control panel contains range sliders to filter nodes by attribute values, which removes them from the heat map, list details, and matrix views. For example, users can filter the views to show only nodes that have at least degree five in at least one time period and have used a certain keyword between three and nine times in at least one time period.

NetVisia also allows users to align nodes by temporal events. This capability shifts rows horizontally according to a temporal event so that the leftmost column now represents the occurrence of an event instead of a specific time bin. For example, aligning by "Joining Network" shifts rows so that the leftmost column represents the first time nodes appeared in the network, the second column is the time immediately after joining, etc (see Fig. 2).

Time Details Views. The bottom two views help users see specific time periods in the network after working in the Workspace Window. The Workspace Window has two column selection indicators (green and blue) that indicate the currently selected time periods. The left mouse button selects a time period in green which is displayed in the left Time Details view, and the right mouse button selects a time period in blue which is displayed in the right Time Details view. The list in the Time Details views shows the attributes for each node in that time period presented in a tabular form that can be sorted by node name or attribute values. In the matrix view, users can see the node-to-node relationships during that time period (see Fig. 3). Example relationships are weighted adjacency matrices, Pearson correlation coefficients, and cosine similarity. NetVisia can visualize other edge attributes in the adjacency matrix view, such as color for word frequency between nodes.

IV. EVALUATION

In order to evaluate NetVisia, we conducted a user evaluation to evaluate: (1) the tool's usability, (2) whether users can find insights using it, and (3) potential improvements.
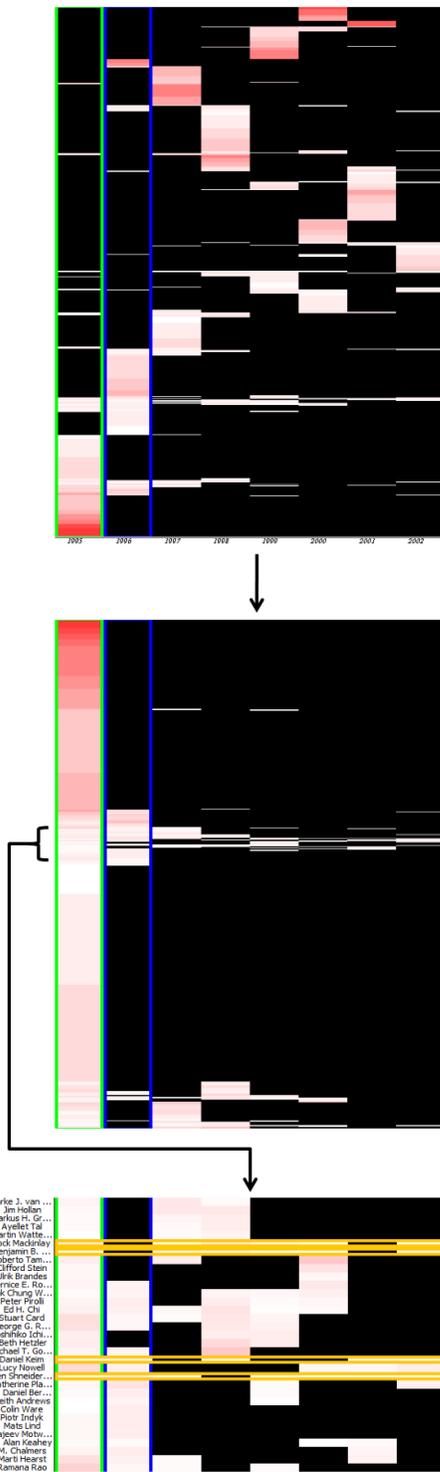
Fig. 2: *Top:* Nodes can be clustered by similar attributes over all time points (compare to the unclustered view in Fig. 1). *Middle:* Nodes can be both clustered and aligned horizontally by when nodes joined the network. In the clustered and aligned view of the InfoVis 2004 co-authorship network, the leftmost column represents the first year an author appeared in the network. We see most authors only wrote papers in one year, and we can easily see outlier authors who wrote papers over multiple years. *Bottom:* Users can zoom in to see more details of individual nodes.
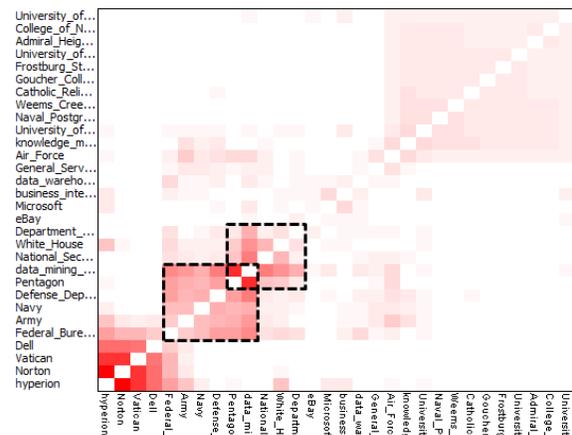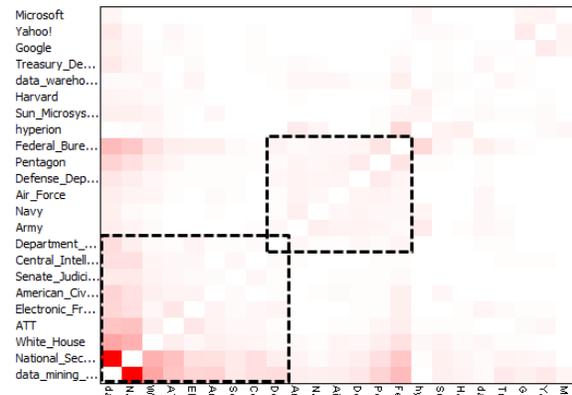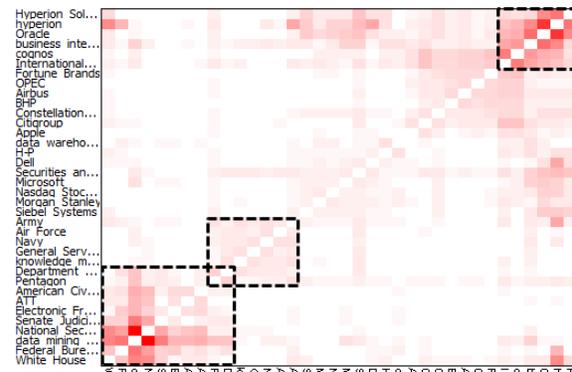

Fig. 3: The Time Details degree matrices for 2005 (a), 2006 (b), and 2007 (c) in the STICK data. Dark red indicates two terms co-occur frequently. Clusters on the diagonal show military and warrantless wiretapping controversy terms in all three years, but 2007 also has a cluster of business intelligence terms related to Oracle's purchase of Hyperion.

A. Participants

We recruited four current computer science graduate students who have already taken or are currently taking a graduate-level class in information visualization. These participants were familiar with network analysis and network metrics, which we also expect from potential NetVisia users. The age range of the participants was 25 to 33 years old. Three participants were male, and one was female. Participants' names are coded as P1, P2, P3, and P4. Our goal was to have enough users to identify most of the obvious usability problems, and evidence suggests that four participants are sufficient for these kinds of studies [20].

B. Experimental Design

Evaluations were performed on an Intel Core i3 2.26 Ghz laptop with 4 GB of RAM. NetVisia was displayed on an external 24-inch monitor at 1920x1200 resolution. The investigators loaded a subset of the InfoVis 2004 contest data set [21] into NetVisia. The subset has been transformed from a citation network to a collaboration network, so that nodes are authors and an edge between them indicates that those two co-authored papers together. This subset contains all authors and references from the InfoVis Conference from 1995 to 2002. In total there are 761 authors, stored as 1034 nodes for time slices with 1569 edges over this eight-year timespan. The smallest granularity of time for this subset is one year. During each year, two nodes have an edge connecting them for each paper they wrote together in that year. This transformation of the data set represents an evolving community of researchers over eight years. Insights from this data set identify patterns of participation, co-authorship, and evolving research interests.

The evaluation used a preliminary version of NetVisia that did not support aligning, aggregating, filtering nodes, or binning time periods.

C. Procedure

Each participant took part in a separate evaluation session after signing the consent form. All participants volunteered and did not receive compensation. Each session began with a training phase, followed by an experimental phase, and concluded with a debriefing. During training, participants watched three training videos and performed training tasks after each. In the experimental phase, participants began with free exploration of the data, followed by performing two tasks. Afterwards, participants filled out a questionnaire and were debriefed. Each session lasted 30–40 minutes.

Training consisted of three videos, each about two minutes in length, and participants performed one training task after each video. The order of videos and training tasks was fixed for all participants.

During experimental phase, participants were asked to think aloud. Following North's recommendations for insight-based evaluations [22], the participants began the experimental phase by performing an undirected search for 10 minutes and describing all of the interesting things they found in the data.

TABLE I: The time participants took to finish each task. The ↓ and ↑ arrows show min and max values respectively.

	P1	P2	P3	P4
Task 1	52s	52s	↓26s	↑2m 26s
Task 2	↑6m 41s	3m 7s	2m 6s	↓1m 33s

D. Results

During the free exploration phase, participant P1 did not generate any insights.

P2 noticed that each year had fewer authors than the previous year. P2 found this insight by comparing two arbitrary years in the Node-to-node comparison view and saw that the matrix in the later year had much fewer nodes. P2 then continued by examining the size of the matrix in each year.

P2 compared two arbitrary years in the Node-to-node comparison view and saw that the matrix in the later year was smaller and had much fewer nodes. P2 then continued by examining the size of the matrix in each year, and noticed that each year had fewer authors than the previous.

P2 noticed three large clusters in the Node-to-node degree comparison for 1995, indicating groups of people who probably worked together a lot in that year. P2 looked at the color intensity in the heat map in the main workspace and commented that "1995 seems to be a hot year."

P2's final insight was that there are many authors who only publish paper in one year. P2 commented that the authors who have papers in six out of the eight years are probably more "impressive" authors whom P2 would like to investigate further.

P3's main insight was that there was an "interesting phenomenon" about clusters in the 1995 Time Details matrix for Pearson coefficients. However, P3 found it difficult to find more details about the groups because the labels were difficult to read and there was no ability to select nodes in the Node-to-node comparison view.

By viewing the heat map in the main workspace, P4 found two authors in 2000 that used the words "network" and "social" a lot that year, but not before or after. P4 had trouble finding these two authors in the degree matrix for 2000, but eventually found them and concluded that they did indeed work together.

1) Timed Tasks: To reduce ordering effects, two participants performed Task 1 before Task 2, and the other two participants performed Task 2 before Task 1. All participants responded correctly to all of the timed tasks. Table I shows the amount of time each participant took to complete each task.

Task 1 ("In 1996, who co-authored with the largest number of people?") required identifying the two authors who both had higher degrees than all the other authors in 1996. P1-P4 selected 1996 in the main heat map, viewed the tabular list details, sorted the table by degree, and then noted the two authors with the highest degree. P4 was also able to identify one of the authors by visually scanning the 1996 column in the heat map.

For Task 2 ("Identify which author used the word 'visualization' in the longest span of consecutive years"), participants

TABLE II: Participants' questionnaire Likert-scale responses. The ↓ and ↑ arrows indicate min and max values respectively.

	P1	P2	P3	P4
Interface	↑8	7	↓1	3
Navigation	↑8	3	↓1	7
Responsiveness	8	↑9	↓6	8
Learnability	↑5	↑5	↓2	4
Terminology	3	↓1	2	↑6

needed to find the author whose row had four consecutive non-black cells. Participants found it difficult to vertically scan the 761 rows in the main workspace heat map for individual rows which spanned many columns. Even after participants found the correct block of cells, they had trouble tracing it back to the label on the left because there were no horizontal grid lines to guide users' eyes. Despite this, all participants successfully identified the correct author.

2) Questionnaire: During debriefing, participants filled out a questionnaire with nine questions: seven closed-ended and two open-ended. The first five closed-ended questions used a 9-point semantically anchored Likert scale—low values are positive responses and high values are negative responses. The investigators designed the Likert-scale questions to include one question each about overall interface design, navigation, responsiveness, learnability, and terminology. The two other closed-ended questions were designed to identify specific features in NetVisia the participants found useful.

Table II shows the Likert-scale question responses. Participants were split on the first question ("The interface is: simple to complex"). For question two ("Navigating the interface is: easy to difficult"), participants were also split, but P2 and P4 had opposite responses from question one. In response to question three ("The interface responds: quickly to slowly"), the participants agreed that the responsiveness is slow. For question four ("Learning to use the interface is: easy to difficult"), participants responded moderately favorably. For question five ("The system's terminology is: informative to uninformative"), participants were mostly favorable, but P4 felt the terminology was mildly uninformative.

Overall, participants appear to have differing opinions of NetVisia. P3 responded favorably to all but the third question, whereas P1 responded negatively to the first three questions, moderately to the fourth, and positively to the fifth.

The next two questions attempted to quantify which of NetVisia's features and network metrics participants found useful (the results are shown in Fig. 4). In Fig. 4(a) we see that not all participants found the Preview Window as useful as the cluster nodes and node list details views. As seen in Fig. 4(b), based on our limited sample size, users do not appear to think that there is much difference in usefulness between the network metrics.

The last two open-ended questions asked participants to discuss additional features, comments, and suggestions to improve NetVisia. Their comments had several common themes: Improving responsiveness, adding tooltips for values and labels in the heat map and matrix, highlighting selected nodes in one view across other views, and shrinking the size of the
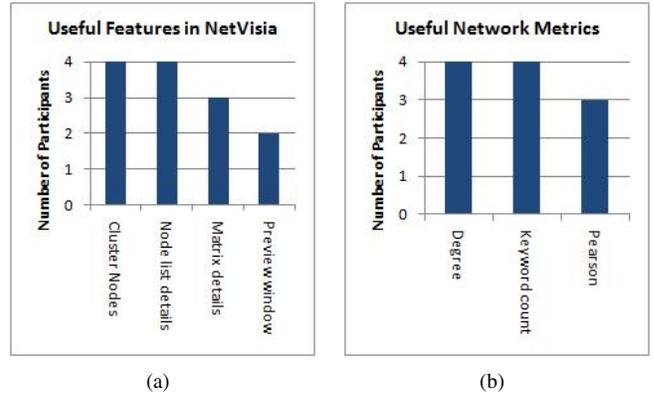

(a)   (b)

Fig. 4: Questionnaire results for who found which features useful (a), and who found which metrics useful (b).

preview heat maps so that they are all visible at once (in the prototype version of NetVisia used in the evaluation, the heat maps were much larger than shown in the figures here).

E. Insight Comparison: User Study and InfoVis 2004 Contest

After the user study, we examined the InfoVis 2004 contest papers to compare the data insights from our user study participants to the insights from the contest participants. Our user study participants discovered many of the insights noted in the InfoVis 2004 contest papers; however, because our version of the data set is a co-authorship network, the user study participants could not gain insights regarding citations, or relationships between specific papers and authors.

The contest papers gained three insights that our user study participants did not gain: the contest papers discovered (1) the top cited authors with the largest number of co-authors, (2) the number of authors who tended to have coauthors, or who started co-authoring less over time, and (3) groups of authors who co-author frequently. P2 gained an insight we could not find in the InfoVis 2004 contest papers: P2 discovered that the number of authors declined from year to year. Our user study participants were also able to identify authors who used certain keywords over a continuous span of time.

Several features in NetVisia were not implemented for the user study: aligning by event, aggregating, filtering, and time binning. Despite this, participants still successfully used NetVisia to answer questions about the data set and discovered new insights in a short amount of time.

F. Summary and Outcomes

During free exploration, participants usually began by exploring the degree metric, and did not usually have time to explore the word frequency metrics before time expired.

Performing the clustering operation on the co-authorship network took around eight seconds of computational time, which slowed down participants in the tasks. Consequently, task-completion times might be shorter if the clustered views were pre-computed and cached in the system. Data sets with fewer nodes yield much faster performance.

Participants' comments suggest that NetVisia needs tighter integration between the different views of the data. Adding tooltips to show node details, and allowing users to select a node in any view and show it selected in other views could improve the analysis capabilities of NetVisia. Some participants asked if NetVisia can perform textual searches (e.g., for node names), which is not currently supported.

Overall, users were able to discover trends and outliers that might be difficult to identify in node-link diagrams. For example, participants were able to identify an author who used the word "visualization" in four consecutive years, and P2 was able to see that the number of authors publishing per year was shrinking over time. This demonstrates NetVisia's potential utility for analyzing the evolution of other networks such as Nation of Neighbors[1] (NoN), which is a social computing network that enables neighborhood citizens to share local information on crime and suspicious activity in their neighborhoods. (Analyzing the NoN social network evolution could reveal interesting insights regarding community formation and crime reporting.)

Based on these results we made the following modifications:

- Added aligning, filtering, aggregating, and binning.
- Added tooltips to show node labels and node values in the heat map.
- Added selection highlighting to the heat map.
- Added coordinated selections between the Workspace heat map and the node list details views.

V. Case Study

After making improvements to NetVisia based on the results of the user study, we analyzed a dataset provided by the STICK team at UMD[2]. This is a data set of business intelligence concepts and organizations (entities) from the ProQuest news database from 2000–2009, representing the evolving patterns of interest for participants in a business community. The concepts were manually defined by domain experts, and the organizations were automatically extracted using named entity recognition. The data set is 1298 nodes (228 concepts and entities over 10 years) with 17,508 edges (each edge is one co-occurrence of two concepts or entities in the same document). This data set is a co-occurrence network, so the minimum degree is necessarily 1 and the maximum degree is 491.

Fig. 5 shows the clustered degree heat map filtered to show only nodes with degrees of at least 45 in at least one year. The vertical axis shows all the nodes which have not been filtered, and the horizontal axis shows the degree of each node in each year. The column summary at the bottom shows the average degree for each year, revealing that the average number of co-occurrences have maxima in 2001 and 2007.

At the top of the heat map, we see two nodes that have similarly high degree over all years: Hyperion and data mining, both of which are related since Hyperion develops data mining software. Just below those two nodes we see a group of
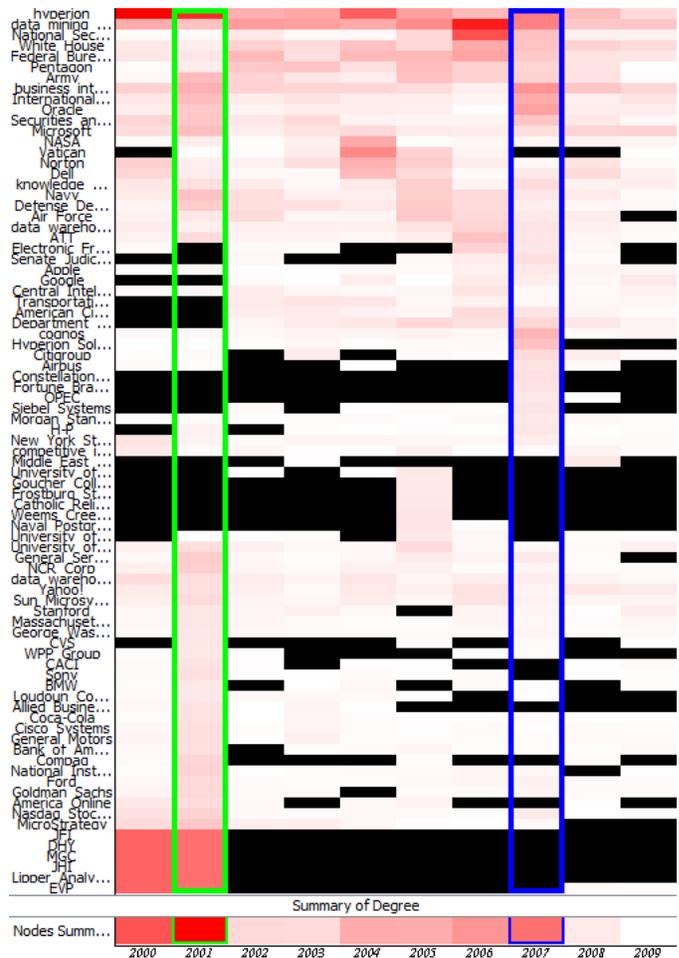
Fig. 5: The clustered heat map of the degree values for the STICK data, filtered to show only nodes with degrees between 45 and 491 in at least one year. The maxima for average degree occurred in 2001 and 2007 (47 and 32 respectively).

government entities that seem to have a high degree after 2001 (White House, FBI, Pentagon, and Army), which could be related to the September 11, 2001 terrorist attacks and the invasion of Afghanistan in the same year. At the bottom of the heat map there is a group that had a very high degree in 2000 and 2001 but suddenly were not mentioned again, except for EVP which did not reappear until 2008. After examining the source material, we discovered that those particular node values in 2000 and 2001 come from stock tables that were last published in 2001, revealing a data cleaning problem as those relationships are not relevant. Just above the bottom cluster is a cluster of nodes which had a low degree (4–31), but then spiked in 2001 (degrees in the range 44–86). In the middle of the heat map, there is a group of nodes where most of them did not co-occur with other nodes except for 2005 when they suddenly had a high degree (43–51).

After selecting the years 2005–2007, the adjacency matrix views provide further details about specific co-occurrences (see Fig. 3(c)). This matrix is filtered to show only nodes with at least degree 40. In the bottom left we see a cluster

of concepts and entities related to the White House and NSA warrantless wiretapping controversy (in 2007 the US Ninth Circuit Court of Appeals heard arguments and issued a ruling regarding a challenge to the warrantless wiretapping). We also see a cluster in the top right of business intelligence concepts and entities, several of which are related to Oracle's purchase of Hyperion Solutions. Near the middle, there is also a faint cluster of military concepts and entities that usually appear.

Discussion. The clustered heat map lead us to identify outlier nodes (such as Hyperion and data mining, or the stock table group) and the column summary reveals outlier years (2001 and 2007). These insights drew our attention to the year 2007 and the nodes Hyperion and data mining, where we discovered two clusters of concepts and entities that occurred together over a sequence of years, and we discovered a third cluster of business intelligence concepts and entities.

VI. CONCLUSION

This paper presents NetVisia, a new approach for visualizing the evolution of numerical node attributes in social networks by using a 3-dimensional clustered view. We demonstrate NetVisia's utility by analyzing a social network and a concept–entity network. We provide a user study demonstrating that novice users can learn NetVisia and generate insights after only a few minutes. Nevertheless, the user study revealed several potential improvements to NetVisia, such as smaller visualizations in the Preview Window, and brushing and linking for selections between the different views. We have since implemented modifications to address these problems. Users also were frustrated with the interface responsiveness for large data sets; however, interface responsiveness was not a problem in our experience with the smaller STICK data set. System performance of social network analytics tools is an area of ongoing research.

A node-link view could also help users examine topological network changes. One avenue of future work is to investigate whether an additional such view is beneficial, e.g. a node-link diagram or arc overlays such as those in MatLink [9].

Our main contributions are: (1) a clustered heat map view of the evolution of numerical node attributes over time, (2) aligning nodes by temporal events, and (3) a usability evaluation of NetVisia and a case study using NetVisia to analyze a social network and a concept-entity network.

REFERENCES

[1] B. Shneiderman and A. Aris, "Network visualization by semantic substrates," *TVCG: IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 733–740, 2006.

[2] T. M. J. Fruchterman and E. M. Reingold, "Graph drawing by force-directed placement," *Software: Practice and Experience*, vol. 21, no. 11, pp. 1129–1164, 1991.

[3] J.-W. Ahn, M. Taieb-Maimon, A. Sopan, C. Plaisant, and B. Shneiderman, "Temporal visualization of social network dynamics: prototypes for Nation of Neighbors," in *SBP '11: Proc. 4th international conference on Social computing, behavioral-cultural modeling and prediction*, J. Salerno, S. Yang, D. Nau, and S.-K. Chai, Eds., vol. 6589, 2011, pp. 309–316.

[4] C. Friedrich and P. Eades, "The Marey graph animation tool demo," in *GD '00: Proc. 8th International Symposium on Graph Drawing*, J. Marks, Ed., vol. 1984, 2001, pp. 396–406.

[5] J. Branke, "Dynamic graph drawing," in *Drawing Graphs*, M. Kaufmann and D. Wagner, Eds. Springer Berlin Heidelberg, 2001, vol. 2025, pp. 228–246.

[6] A. Perer and B. Shneiderman, "Balancing systematic and flexible exploration of social networks," *TVCG: IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 693–700, Oct. 2006.

[7] S. Diehl and C. Görg, "Graphs, they are changing: dynamic graph drawing for a sequence of graphs," in *GD '02: Revised papers from the 10th International Symposium on Graph Drawing*, M. T. Goodrich and S. G. Kobourov, Eds., vol. 2528, 2002, pp. 23–31.

[8] D. Archambault, "Structural differences between two graphs through hierarchies," in *GI '09: Proc. Graphics Interface 2009*, 2009, p. 8794.

[9] N. Henry and J.-D. Fekete, "MatLink: enhanced matrix visualization for analyzing social networks," in *INTERACT '07: Proc. 11th IFIP TC 13 International Conference on Human-computer interaction*, 2007, pp. 288–302.

[10] K. Stein, R. Wegener, and C. Shchlieder, "Pixel-oriented visualization of change in social networks," in *ASONAM '10: Proc. 2010 International Conference on Advances in Social Networks Analysis and Mining*, Aug. 2010, pp. 233–240.

[11] J. S. Yi, N. Elmqvist, and S. Lee, "TimeMatrix: analyzing temporal social networks using interactive matrix-based visualizations," *IJHCI: International Journal of Human-Computer Interaction*, vol. 26, no. 11–12, pp. 1031–1051, 2010.

[12] F. Steinbrückner and C. Lewerentz, "Representing development history in software cities," in *SOFTVIS '10: Proc. 5th international symposium on Software visualization*, 2010, pp. 193–202.

[13] S. Hadlak, C. Tominski, H.-J. Schulz, and H. Schumann, "Visualization of attributed hierarchical structures in a spatiotemporal context," *IJGIS: International Journal of Geographical Information Science*, vol. 24, no. 10, pp. 1497–1513, Oct. 2010.

[14] M. Freire, C. Plaisant, B. Shneiderman, and J. Golbeck, "ManyNets: an interface for multiple network analysis and visualization," in *SIGCHI '10: Proc. 28th international conference on Human factors in computing systems*, New York, NY, USA, 2010, pp. 213–222.

[15] T. Falkowski, J. Bartelheimer, and M. Spiliopoulou, "Mining and visualizing the evolution of subgroups in social networks," in *WI '06: Proc. 2006 IEEE/WIC/ACM International Conference on Web Intelligence*, Dec. 2006, pp. 52–58.

[16] M. Rosvall and C. T. Bergstrom, "Mapping change in large networks," *PLoS ONE*, vol. 5, no. 1, p. e8694, 01 2010.

[17] A. Sopan, M. Freire, M. Taieb-Maimon, J. Golbeck, and B. Shneiderman, "Exploring distributions: design and evaluation," University of Maryland, Human-Computer Interaction Lab Tech Report HCIL-2010-01, Apr. 2010.

[18] J.-W. Ahn, C. Plaisant, and B. Shneiderman, "A task taxonomy of network evolution analysis," University of Maryland, Human-Computer Interaction Lab Tech Report HCIL-2011-09, Apr. 2011.

[19] Z. Bar-Joseph, D. K. Gifford, and T. S. Jaakkola, "Fast optimal leaf ordering for hierarchical clustering," *Bioinformatics*, vol. 17, no. suppl 1, pp. S22–S29, 2001.

[20] J. Nielsen and T. K. Landauer, "A mathematical model of the finding of usability problems," in *CHI '93: Proc. INTERACT '93 and CHI '93 conference on Human factors in computing systems*, ser. CHI '93, 1993, pp. 206–213.

[21] J.-D. Fekete, G. Grinstein, and C. Plaisant. (2004) IEEE InfoVis 2004 Contest, The History of InfoVis.

[22] C. North, "Toward measuring visualization insight," *IEEE Computer Graphics and Applications*, vol. 26, no. 3, pp. 6–9, 2006.