PeckVis: A Visual Analytics Tool to Analyze Dominance Hierarchies in Small Groups

Darius Coelho, Ivan Chase, and Klaus Mueller, *Senior Member, IEEE*

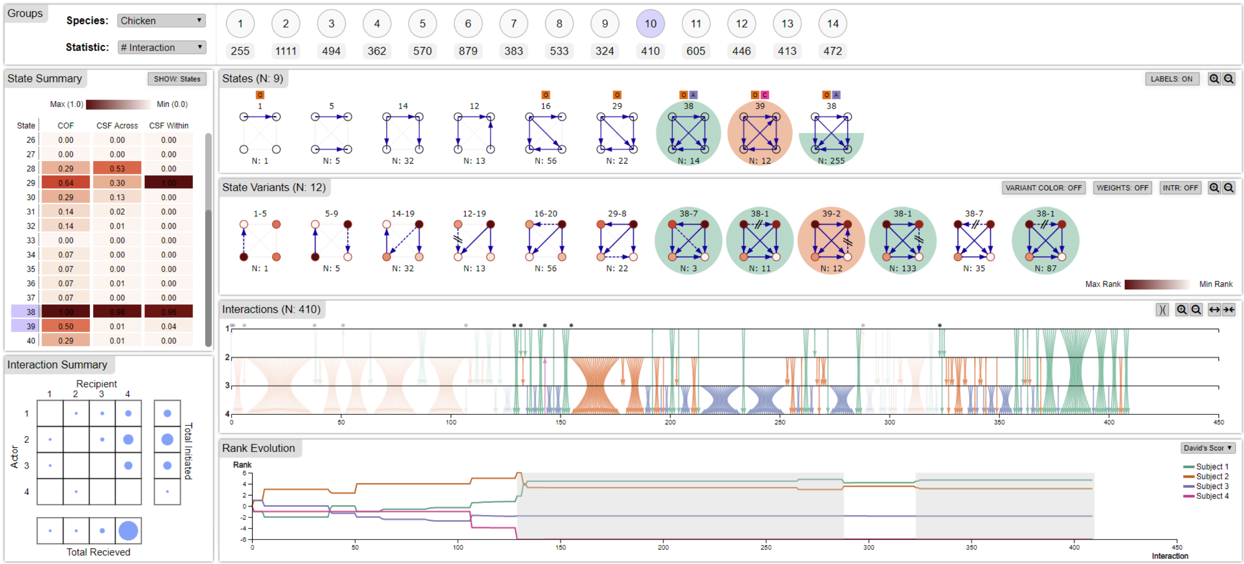

Fig. 1: The *PeckVis* interface for analyzing dominance hierarchies formed by a single group of subjects. At the top is a control panel to select a group for analysis, here chicken group 10 (blue color) is selected. Below this are two panels on the left that include a heatmap table to summarize hierarchies formed and a balloon plot below it to summarize the interactions in the group. To the right of these panels (from top to bottom) is the state sequence and state variant sequence that represent the hierarchies formed and below these is the music notation and rank evolution chart that shows the raw interaction data and how the ranks of subjects change with every interaction. Here, the user selected state 38 and 39 for inspection via the heatmap. The corresponding states, state variants, and interactions were highlighted in the charts to the right of the heatmap. Further, the user deselected a state variant (second to last in the variant sequence) which causes the partial highlight in last state of the state sequence. Section 6.1 details this group's analysis.

Abstract— The formation of social groups is defined by the interactions among the group members. Studying this group formation process can be useful in understanding the status of members, decision-making behaviors, spread of knowledge and diseases, and much more. A defining characteristic of these groups is the pecking order or hierarchy the members form which help groups work towards their goals. One area of social science deals with understanding the formation and maintenance of these hierarchies, and in our work we provide social scientists with a visual analytics tool - PeckVis - to aid this process. While online social groups or social networks have been studied deeply and lead to a variety of analyses and visualization tools, the study of smaller groups in the field of social science lacks the support of suitable tools. Domain experts believe that visualizing their data can save them time as well as reveal findings they may have failed to observe. We worked alongside domain experts to build an interactive visual analytics system to investigate social hierarchies. Our system can discover patterns and relationships between the members of a group as well as compare different groups. The results are presented to the user in the form of an interactive visual analytics dashboard. We demonstrate that domain experts were able to effectively use our tool to analyze animal behavior data.

Index Terms—Visual analytics, interaction sequence, dynamic graphs, time series, dominance hierarchy

1 Introduction

Social behavior in animals and humans is inferred from how individuals interact with each other. Some animals are reclusive while others are highly social and form complex social organizations through a diverse set of interactions. The formation of organizations or groups is important as a group can accomplish objectives that an individual cannot and that help the species as a group survive and reproduce. Understanding a group's objectives in most cases can be relatively easy through direct observation. But understanding the formation of a group is a more difficult task. By studying the social interactions experts can infer the causes for actions and evolution of a species.

Social groups vary in size and can contain as few as two members or be extremely large containing millions of members. Regardless of the size, the group members in most species tend to form hierarchies by competing with each other for rank. The ranking they form is called a *dominance hierarchy*. These hierarchies are important as they

- *Darius Coelho and Klaus Mueller are with the Computer Science Department, Stony Brook University. E-mail: {dcoelho, mueller}@cs.stonybrook.edu.*
- *Ivan Chase is with the Department of Sociology, Stony Brook University. E-mail: ivan.chase@stonybrook.edu.*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

define the roles of the group members which in turn help domain experts understand a group's decision-making process, the spread of knowledge and diseases, and other phenomena. Our work focuses on the formation of these dominance hierarchies as they are informative of the dynamics within a group.

We worked with a social scientist who was interested in analyzing the formation of dominance hierarchies in small animal groups. In this domain, the size of small groups is defined to be three to twelve members as this size range is representative of naturally occurring groups. As groups get larger, animals are unable to recognize other members and are therefore unable to maintain a stable hierarchy [29]. The expert's team collected data by observing groups from three species - chickens, mice, and fish. Each group consisted of four subjects and produced a dataset that contains a sequence of time-stamped interactions between two members. These datasets contain hundreds to tens of thousands of interactions over a two day period and are almost impossible to analyze by just inspecting the raw data.

To simplify the analysis, the interaction sequences are reduced to a series of time-varying networks where members are the nodes and their interactions are represented as directed links between the nodes. These network sequences represent the different hierarchies formed over time. Experts run statistical analyses on these network sequences to extract information about the hierarchy formation, however, these analyses only provide summaries and do not communicate the nuances of the hierarchy formation. Specifically, they do not show how hierarchies change or repeat over time. Using a visual analytics approach can speed up the analysis process and provide insights that may have otherwise been overlooked. Unfortunately, most available network visualization tools are designed to analyze large online social networks which are unsuitable for analyzing the small animal networks or hierarchies. This is primarily due to the tools only visualizing the networks without the context of the interaction sequences which provide details about how the networks were formed. Additionally, social scientists were able to better understand simple node-link representations with minimalistic enhancements over the advance techniques available in existing tools.

We address this issue with *PeckVis*, an interactive visual analysis system for social scientists to analyze the formation of pecking orders in small groups consisting of up to six subjects. The system integrates domain-specific algorithms and metrics with a visual analytics dashboard that allows users to interactively explore and investigate the formation and maintenance of dominance hierarchies. The visual representations employed by *PeckVis* adapts existing network and sequence visualization techniques to the needs of social scientists. *PeckVis* also enables users to compare different groups for commonalities and differences in their hierarchy formation using these visual representations. Finally, we demonstrate our system's capabilities by having domain experts use *PeckVis* to analyze animal interaction data they have collected. Additionally, we demonstrate how our system can be applied to other situations in which hierarchies are formed such as debates.

2 Related Work

PeckVis uses graphs to represent and analyze the influence of temporal relationships among group members on their hierarchy. Our system represents interaction sequences as a sequence of networks with nodes representing the actors and edges representing the interactions. Hence, most relevant to our work is the visual analysis of dynamic networks.

Most of the work on visualizing networks has focused on the analysis of large networks. Landesberger et al. [37] provide a good review of these works. Their review shows that coupled with a variety of interaction techniques, the most frequently used representations in prior graph visualization work have been node-link diagrams [31] [24], adjacency matrices [1] [15], or a combination of both [25]. Works such that of Gou et al. [23] have created dashboards that use these representations supported by other charts to explore networks. Both representations have their advantages; the node-link diagrams are intuitive and better suited for path following tasks whereas the adjacency matrix avoids edge crossings and node overlapping which leads to better readability of large graphs. The properties of these representations have been confirmed in studies by Ghoniem et al. [21] and Alper et al. [3].

While the representations discussed are sufficient on their own for network exploration, the task of comparing different networks needs more sophisticated techniques. Metrics can be computed over graphs to describe their differences. These metrics can then be used to analyze and guide the visual comparison of graphs as shown by Kairam et al. [26], and Freire et al. [18]. Alper et al. [3] studied the use of node-link and an adjacency matrix to compare networks describing brain connectivity and reinforced the fact that node-link representations are better for comparing small graphs. Furthermore, similarity metrics can be computed over a larger number of graphs or sub-graphs which then help the user in selecting similar or dissimilar graphs for comparison as in Landesberger et al. [36].

A special case of graph comparison related to our work is that of dynamic graphs, where changes in a graph's structure are compared over time. Beck et al. [8] provide a very recent review of the techniques used to visualize dynamic graphs. Of particular interest is work that juxtaposes different states of a network over time representing them as small multiples [17] [33] [5] [35]. Farrugia et al. [33] use a small multiples display to both show the evolution of an egocentric network and to compare different egocentric networks. In DiffAni [33] and SmallMultiPiles [5] dynamic graphs were aggregated and displayed as small multiples. This approach of aggregation into intermediate graphs over time inspired techniques used in our system. Most similar to our work is that of Velhow et al. [35] who visualized dynamic hierarchies in graph sequences by using an adjacency matrix that integrates hierarchical group structure along with icicle plots. Additionally, they use a flow metaphor and color encoding to visualize changes. Archambault et al. [4] create difference maps of a network at two time slices in a single graph encoding the differences with color. Analogous to this approach we compute differences and use markings and color to highlight differences in consecutive graphs.

Most of the network visualization work discussed has been applied to large social networks, computer networks, or biological structures. Our work focuses on interactions between individuals in small groups and specifically targets animal behavior, an area that has received little attention in the visualization field. However, there are some relatable works which we discuss here. DiMicco et al. [13] used visualizations to review the turn-taking patterns in face-to-face meetings. They can deduce social trends such as dominance, extroversion, and endorsements but have no procedure to inform a concrete reasoning into the formation of these trends. Alallah et al. [2] also visualized face-to-face meetings to review the decision-making process. They used a Gantt chart-like representation to plot and compare a user's behavior. Cao et al. [10] supports the exploration and summarization of user interactions with an interactive visualization. They represent time for a particular user with a horizontal axis with vertical lines representing the user's interactions much like Chase's music notation [11]. Recently, Fu et al. [19] devised a visual analytics system to interactively explore, compare, and track conversation groups in online forums; their work demonstrates the usefulness of glyphs to represent groups.

While social scientists have developed a variety of mathematical models and analyses to investigate hierarchy formations, they have very few suitable visual analytics tools to aid their efforts. We believe that aspects of the works discussed here can be adapted with modifications to help other scientists, as predicted by Beck et al. [8]. Specifically, the use of node-link and adjacency matrix representations arranged in a small multiples display is generally useful to compare network structures. In addition, we were inspired by the dashboard approach of incorporating connected displays of timelines and other metrics to support the graph analysis. In the remainder of this paper, we discuss our approach of combining these techniques into a single system and the evaluation of this system.

3 Mechanisms and Metrics for Dominance Analysis

Dominance is a very important concept in the study of social behavior, and dominant behaviors form dominance hierarchies within a group. Work by social scientists has led to a variety of analysis techniques that quantify these hierarchies. We introduce these methods and algorithms, that we further adapted for *PeckVis*, below.

3.1 Ranking

A variety of ranking algorithms have been proposed that place individuals at different levels in a hierarchy. Unfortunately, these algorithms are not universally accepted as they all have some limitations [20] [7]. For this reason, *PeckVis* provides the user with two widely used ranking algorithms - Davids Score [12] and Elo ranking [16] - each with their own benefits and allows the user to interactively switch between them.

David's score is an interaction matrix based method for computing rank in social science. It has been shown to overcome problems with other interaction matrix based ranking methods in the field [20]. David's score measures the overall success of an individual by weighting each dyadic success measured by the unweighted estimate of the other individuals' overall success, thus taking into account the relative strengths of the other individuals. The Elo-ranking method was initially generated to rank chess players, but it has gained popularity and is now used to compute rankings in a variety of fields. Neumann et. al. have adapted it to compute the ranks in a dominance hierarchy [32]. They show that the Elo-ranking has benefits over matrix based methods. They include the elimination of certain data limitations such as a lower-bound on the number of interactions and the flexibility to extract scores at any point in time, thus making it easier to visualize.

Additionally, we provide users with naive ranking methods, such as the cumulative sum of interactions initiated, the cumulative sum of interactions received and a combination of both over time. These do not accurately represent the social rank of an individual but they do provide useful information, they give users a general idea about an individuals activity over time.

3.2 State Sequence Analysis

The ranking methods inform us of the hierarchy established by the subjects in a group. However, they do not explain what led to the formation of the hierarchy. Addressing this issue, researchers have analyzed the interaction data by modeling them as a dominance network. A review of these methods was presented by Doreian [14]. In our work, we extend the methodology of Lindquist et al. [30] who combined the ranking techniques and network analysis to model hierarchy formation. They trace a group's hierarchy development by inspecting intermediate dominance configurations that subjects in the group form over time. These configurations are aggregations of a set of interactions and long interaction sequences are aggregated into a sequence of configurations. In their work, Lindquist et al. only analyzed the structure of hierarchies, that they call "states", while ignoring the identities of the subjects that formed the hierarchy. In our work, we include the identities as well and call these identity dependent hierarchies "state variants". The remainder of this subsection discusses the procedure followed to aggregate interaction sequences into configurations and how these configurations are annotated to convey important information to the analyst.

3.2.1 States and State Variants

A group's interactions can be aggregated to form multiple dominance configurations called "state variants". The structurally unique configurations that groups form are called "states". More specifically, a "state" refers to the structural form of the configuration while ignoring the identities of the subjects whereas a "state variant" is more specific and considers these identities. Thus two state variants that only differ in vertex labeling map to a single state. It should be noted that the vertex labeling cannot be arbitrary, it must be consistent and comparable across groups. To achieve this we label subjects by their eventual ranks, i.e., their ranks after their group's last recorded interaction. Thus subjects labeled '1' in different groups are the eventually highest ranked subjects in their groups.

Interaction sequences are aggregated into a sequence of state variants which are further aggregated into a sequence of states. Each state variant is essentially a directed graph with nodes representing the subjects in a group and links representing interactions between subjects. A new state variant is added to the sequence when an unobserved interaction occurs or an existing interaction is reversed. When a new state variant is created, the relationships from the prior state variant are propagated to it, new relationships are added, and reversed relationships are replaced.

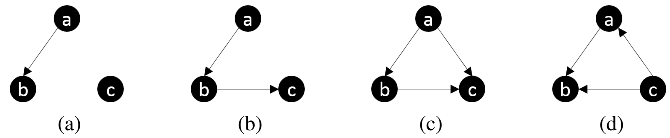

Fig. 2: The progression of a state variant sequence based on four interactions: *a* dominates *b*, *b* dominates *c*, *a* dominates *c* and finally a reversal *c* dominates *b*.

Consider the example shown in figure 2; there are three individuals labeled *a*, *b*, and *c*. In the first interaction *a* dominates *b*, creating the state variant in figure 2a which continues to exist while *a* keeps attacking *b*. When a new interaction *b* dominates *c* occurs, a new state variant is created and the relationship *a* dominates *b* is carried over to it as shown in figure 2b. Similarly new state variants are created when *a* dominates *c* in figure 2c and *c* dominates *b* in figure 2d . It should be noted that reversed relationships are replaced as in the last state variant where *b* dominates *c* was replaced by *c* dominates *b*. Additionally, the last two state variants map to a single state as they are structurally identical - they both have one subject dominating the other two and one of those two dominate the other. A more detailed example of how interaction sequences map to state variant sequences which in turn map to state sequences is illustrated in the supplementary material.

States and state variants allow the experts to split their analysis into two main stages. States allow experts to first identify different types of hierarchies or states such as states where one subject is dominated by all others or states with graph cycles. By analyzing these states, experts can make judgments about the stability or competition within a group. Additionally, these states serve as a high-level feature for comparing multiple groups. State variants allow the expert to dig deeper into the states, they show which subjects form a hierarchy based on their eventual rank. For example, experts find that in some groups the second ranked individual was at top of a one dominate all state but as time progressed the first ranked individual was at the top in that state.

3.2.2 State Annotation

States represent the different types of hierarchies a group can form and the number of possible hierarchies exponentially increases with the number of subjects. To ease the process of analyzing these large numbers of hierarchies, we precompute all possible states an *N* member group can form and build a lookup table for future use. Due to the combinatorial complexity of generating the lookup tables, we limit group size to six subjects which can form 21,479 different states. We then annotate these states based on their structural characteristics, these annotations are described as follows.

Linear States: A linear state is one in which a subject is not attacked by any subject it attacks. In other words, the most dominant individual is never attacked, the second most dominant is attacked only by the first and so on. For example, states in figure 2c and 2d are linear states. Linear states tend to be the most common state and inform domain experts of clear hierarchies where no subject is competing with a more dominant subject.

One Dominate All: States in which one subject in a group is dominating all other active subjects are called *One Dominate All (ODA)* states. The states in figure 2c and 2d are examples of *ODA* states with subjects *a* and *c* being the most dominant respectively. An observation of a large number of specific *ODA* state variants informs experts of groups with very dominant individuals.

All Dominate One: States in which one subject in a group is dominated by all other active subjects are called *All Dominate One (ADO)* states. The states in figure 2c and 2d are also *ADO* states with subjects *c* and *b* being the dominated subjects respectively. An observation of a large number of specific *ADO* state variants informs experts of groups with very submissive individuals.

Intransitive States: In most cases, subjects tend to form transitive or linear hierarchies. But occasionally the subjects or a subgroup of subjects form a cyclic relationship known as *intransitive relationships*.

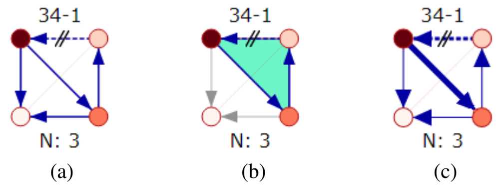

Fig. 3: The state variant representation in its basic form (a) where node positions indicate a subject's eventual rank (decreasing from left to right and top to bottom), node color indicates instantaneous rank, links represent relations where dashed links represent new relations and the // marker indicates a reversal. Users can toggle the highlights of cycles (b) and edge weighting for interaction count (c).

For example, if the relationship *b* attacks *c* in the state shown in figure 2d were to be reversed it would form an intransitive relationship among the subjects with actor *a* dominating *b*, *b* dominating *c*, and then *c* dominating *a*. Intransitive states are also non-linear states and are indicative of competition and instability of rank among subjects.

Occurrence and Stability: In addition to labeling the type of states, we compute two important metrics for each state and state variant - *Class Occurrence Frequency (COF)* and *Class Stability Factor (CSF)* - to perform the comparison. The *COF* measures the occurrence of states across groups in a species. It is the fraction of groups that formed a specific hierarchy during the interactions. For example, if a state occurs in 6 of 10 groups it would have a *COF* of 0.6. A value of 1 indicates that every group forms that hierarchy and 0 indicates that the hierarchy is never formed. The *CSF* informs us of the stability of states that occur versus other states in an *N*-link class. It is the ratio of time spent in a particular *N*-Link hierarchy to the time spent in all *N*-Link hierarchies. For example, if groups had 85% their interactions in 3-Link configurations belonging to a particular hierarchy then the *CSF* for that hierarchy would be 0.85.

4 Visual Representations in *PeckVis*

The data analyses methods yield a rich set of results, but they can be overwhelming for the user to review, especially when dealing with a large number of groups and interactions. To efficiently communicate this information to the users, we extended existing visualization techniques as described below.

4.1 Visualizing Hierarchy Sequences

The state and state variant sequence is the sequence of hierarchies derived from a group's interaction sequence. These hierarchies or states are essentially graphs that can be visualized with a variety of existing techniques. We experimented with multiple techniques such as node-link diagrams, matrix representations, and chord diagrams. The existing approach by Vehlow et al. [35], who used a matrix representation complemented by icicle plots, can be applied to our problem but the domain experts were not comfortable with the representation for multiple reasons. First, the experts preferred a node-link representation as they previously worked with this representation and found it more difficult to follow changes across adjacency matrices. Second, Vehlow et al. used a single representation to represent the topology (state) and the hierarchy structure (state variant), the experts found this representation to be overloaded and preferred it split into two parts - one to analyze states and one to analyze state variants. Finally, Vehlow et al. only show the instantaneous hierarchy in each matrix, the experts we work with wanted both the instantaneous and eventual hierarchies encoded. Additionally, our task involved visualizing small groups and based on previous studies node-link diagrams were suitable. Thus we represent the state and state variant sequence as a sequence of node-link diagrams with additional visual encodings that communicate various characteristics of the hierarchy at a particular point in time.

State Variant Sequences: The domain expert we work with represents state variant sequences with hand-drawn directed node-link diagrams similar to those in figure 2. This representation only encodes the structure of the current hierarchy and identity of subjects, however, we update the node-link representation with additional visual encodings to represent the current and eventual rank of the subjects, the newest link, and a reversed link. Additionally, we label each state with its ID, variant and the total number of interactions it represents. An example of a single state variant is shown in figure 3 and a state variant sequence is shown in figure 1.

The current and eventual rank of the subjects is encoded with color and node placement respectively, this allows the experts to view the current hierarchy in the context of the eventual hierarchy. In our design the eventual rank decreases from left to right and top to bottom, thus in the example, the two nodes at the top represent individuals with an eventual rank of 1 (left) and 2 (right) and nodes at the bottom represent rank 3 (left) and 4 (right). The nodes are colored with a sequential red color scale to represent current rank. The subject with the highest rank at the end of the current state is assigned a dark red color and the lowest ranking subject is assigned a light red color. This encoding allows the user to easily track the evolution of the group member's ranks as well as identify states where ranks have changed. For example, in figure1 we see that in the first state shown the subject eventually ranked 3rd (shown with node position) is ranked the highest (shown with node color) in the first state. Over the next six states the subject eventually ranked 2 takes over the highest rank. And in the remaining states, the subject eventually ranked 1 dominates all other subjects with some competition from the 2nd ranked subject. As the state space grows users can observe the ranks and the distance between ranks changing by visually comparing the node colors. When used in conjunction with the ranking chart (discussed below) users can get a finer view of the distances between ranks.

The links between nodes represent relationships between individuals and states are created when a new relationship occurs or an existing relationship is reversed. Initially, we used color to mark new interactions and reversals but users found the representation to be too confusing. Instead, we chose to indicate the newest interaction in a state with a dashed line and reversals were indicated by placing a double slash marker at the middle of a dashed line. We also encode the number of interactions between individuals with line thickness. This is shown in figure 3. We allow the user to toggle this weighting of links as it overloads the representation and this information was rarely required by the experts. This approach follows the juxtaposition combined with explicit encoding as described in Gleicher et al. [22] by sequentially presenting the different configurations over time and explicitly marking the differences between them.

Finally, the experts wanted to easily locate the cycles or intransitive triads in the states as they are indicative of rank contention. We show these cycles by graying out non-contributing edges and filling the space between contributing edges with a green highlight as shown in figure 3b. This was helpful for the user when scrolling through large state spaces as it revealed trends such as the same cycle being formed repetitively which experts found to be interesting behavior. The user is allowed to toggle this functionality as well since it can be distracting when examining at other features of the state variant sequence.

State Sequences: As described previously, states represent the structural form of dominance hierarchies while ignoring the identities of the subjects in those hierarchies. Thus the representation for states is identical to that of state variants with the only difference being that the nodes and their positions do not encode any information and new interactions and reversals are not marked. However, we do keep the node positions consistent for node-link diagrams representing the same state. We also use labels above states to indicate the type of hierarchy (ODA, ADO, cyclic and linear). An example of a state sequence is shown in figure 1.

4.2 Visualizing Interactions

The state variant sequences inform us of the different hierarchies formed and the order in which they were formed. However, they do not inform us of the way in which they were formed, that is, what sequence of interactions led to a given hierarchy. This information can be obtained by inspecting the interaction sequences. We use two visual representations

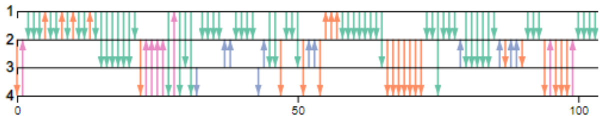

(a)

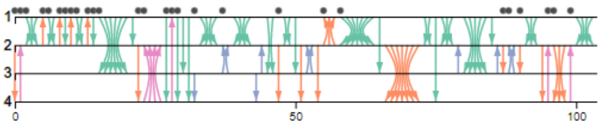

(b)

Fig. 4: The original music notation developed by Chase (a) and the enhanced version (b) that adds edge bundling to emphasize bursts and markers at the top to indicate the formation of new states.

to assist experts in exploring the interactions among subjects. The first is a music notation that is a visual encoding of the raw interaction data and the second is a balloon plot that provides an overview of all the interactions in a group.

Music Notation: The music notation was developed by Chase [11] for social scientists to view raw interaction data. The representation helps users to quickly identify patterns in the data, the most important being bursts. Bursts are sub-sequences in which one subject continuously dominates one or more subjects in a group. Additionally, the chart can indicate the distribution of interactions among subjects. The chart represents each subject in a group with a horizontal line or axis. For every interaction, an arrow is drawn from the actor's horizontal line to the recipient's line. The arrows are colored based on the color assigned to an actor. The user is allowed to choose between positioning the arrows based on clock-time or by their count. Positioning by time allows users to see if the time of day had an effect on the actors. Positioning by interaction count keeps the space between the arrows constant and positions them in the order they occurred. This makes it easier for users to observe bursting behaviors. The horizontal axes are ordered from top to bottom based on the eventual rank of the actor determined by the ranking algorithm selected by the user.

We further build upon the music notation developed by Chase. We emphasize bursts by using edge bundling and we indicate positions at which new state variants were formed with markers above the top axis. An example is shown in figure 4. The edge bundling cause bursts to look like hourglasses, with wider hourglasses indicating larger bursts. In addition to indicating where states are formed, the number of markers and the distance between them is indicative of the stability in the group. Fewer markers with large distances between them indicate that the group and the hierarchies it forms is stable.

Balloon Plot: The interactions in a group can be summarized with an interaction matrix which informs users of the total number of interactions between any two subjects in a group. We chose to use a balloon plot over a stacked bar chart and adjacency matrix to represent the interaction matrix. An example is shown in figure 1 in the bottom left panel. We did not consider the use of a node-link diagram here as users may overlook missing relationships [34]. The domain expert preferred this representation over a traditional adjacency matrix and heatmap representation used for large networks as the expert felt the circle size was easier to compare instead of color intensity (although we believe that this may change as the network size grows). The balloon chart rows represent actors and the columns represent recipients. There is an additional disjoint row and column which represent the totals of each column and row. Each cell contains a circle whose size represents the number of interactions between the corresponding actor and recipient. The representation allows the user to quickly gain insight into the distribution of interactions among group members. Users can easily identify which actor initiates the most interactions and how many individuals it interacts with thus giving them a summary of the interactions.

4.3 Other Representations

In addition to visualizing the state and interaction sequences, we provide experts with visual representations that support their analyses. The representations are used to keep the expert aware of subjects ranks, group similarity, and to communicate other informative statistics.

Rank Evolution Chart: We represent the ranks computed by the selected ranking algorithm after every interaction with a multi-line plot as shown in figure 1 in the bottom right panel. Lines are plotted with a step-function and colored based on the color assigned to the actors. The lines tend to diverge quickly when hierarchies form at a fast pace within a group. However, if actors compete for rank then the lines intersect or follow an alternatively diverging and converging path. Users can configure this chart to present ranks from the desired ranking algorithm which is also reflected with node colors in the state space.

MDS Plot: It is difficult for users to gauge the similarity of groups based on their state space or other metrics computed for each group such as the counts of different states or the time spent in these different states. To address this issue we used multidimensional scaling [28] to represent each state or state variant sequence as a point in 2-dimensional space. The method shows sequences that are similar closer together; the results for our animal datasets are shown in figure 5 in the bottom left panel. To compute the distance between the sequences we used the dynamic time warping (DTW) distance [9] which measures the similarity between time sequences. Here we represent each sequence as a time series by giving each state or configuration a value and plotting this value at every timestamp the group is observed in a particular configuration. This process allows us to account for the duration a group remains in a particular configuration. Some groups form their hierarchies at different speeds and DTW accounts for this and measures similarities even though time series have different speeds and levels of acceleration and decelerations. Sometimes, however, the expert may want to disregard the order in which states occur and measure similarity purely based on the occurrence of states and the number of interactions in each state. In this case, we create a multidimensional dataset where each column represents a state or state variant and for each group, we record the amount of time spent in each state and record it in the appropriate column. We then compute the similarity between groups using the Euclidean distance measure and show it in an MDS plot.

Summary Statistics: Finally, we compute certain summary statistics for each group which we communicate to the user with bar charts. These statistics are computed over the state sequences and are useful for comparing groups. Most important to the expert were the *COF* and *CSF* values of states and state variants for each species and each group. We communicate these values to the expert with a heatmap table as shown in figure 1 in the middle panel on the left. As the number of subjects increases, the number of states exponentially increases and the heatmap allows users to quickly learn about the kind of states and the time spent in them, within and across groups. Additionally, statistics such as the number and percentage of states categorized by their annotations, for example, the counts of *ADO*, *ODA*, and intransitive states are communicated with bar charts.

5 The *PeckVis* Interface

The visual representations discussed in the previous section are individually very effective at communicating hierarchies and interactions. However, combining and linking these representations into a single interactive interface allows the user to gain finer insights into a group's hierarchy formation and maintenance. The experts we worked with wanted to perform two main tasks - (1) analyze how individual groups form and maintain hierarchies, and (2) compare multiple groups to learn about similarities in hierarchy formation. To achieve this, we provide experts with two dashboards to analyze a single group and to compare multiple groups. The dashboards and supported interactions are discussed as follows.

5.1 Analyzing a Single Group

Analyzing a single group involves the tasks of identifying the different hierarchy structures, how subjects are positioned or ranks change in these structures and the characteristics of interactions within these structures. We enable such analyses by laying out the state and state variant sequence representation, the music notation, rank evolution, and balloon plot into a dashboard as shown in figure 1.

At the top is a panel containing circles that represent each group for a particular species and below each circle is a label communicating a group statistic such as the total number of states or interactions in a group. The species and the statistic can be changed using the drop-down menus at the left end of the panel and clicking circles selects the group to be inspected. The space below this panel is split horizontally into two sections. The left section contains two panels, the first is a heatmap table that communicates a summary of the states such as their occurrences and stability (top panel) and the second is a balloon plot communicating the summary of interactions. In the right section, we stack the selected group's state sequence, state variant sequence, music notation and rank evolution chart from top to bottom. This stacking order allows the user to analyze the data with an overview to detail approach by moving from the state sequence which shows the types of hierarchies to the more detailed state variant sequence which shows how subjects are ranked in those hierarchies and finally to the music notation that communicates the sequence of interactions in the hierarchies.

Every representation in the dashboard is linked to each other. If a user selects a state, all state variants that were aggregated into it are highlighted. The corresponding interactions are also highlighted in the music notation and the corresponding range is highlighted in the rank evolution chart. This is reflected in figure 1 where the states 38 and 39 were selected from the state summary panel on the left. This led to the corresponding states, state variants and interactions being highlighted. It should be noted that states may be partially highlighted if some of the state variants they encompass are deselected. For example, the last state in figure 1 is partially highlighted as a state variant (the second last in the state variant sequence) it encompasses is deselected. A selection can be made from any chart with ranges corresponding to that selection highlighted in other charts. Left clicking a state or state variant selects that particular item in the respective sequences, but right clicking selects all identical state or state variants in the sequence. For example right clicking a state labeled 38 will highlight all states labeled 38 in the sequence. To differentiate states, we use different colors to highlight them based on the state label. For example in figure 1 state 38 and 39 are highlighted with green and orange, respectively. A large number of states can occur in a sequence and we may not have enough colors to represent them. In this case we reuse colors. Theoretically, this can mislead the user however we ensure that two consecutive states do not use the same color. Thus when the same color is seen in consecutive state variants, it indicates that it maps to the same state. Additionally, each state is labeled with its state and variant above it.

We also provide controls to toggle options for each visual representation in the dashboard. For example, users may want to zoom into the music notation or deactivate its edge bundling. The state variant sequence encodes a lot of information and may be overloaded. Thus we allow the user to toggle the encodings such as the edge weighting, intransitive triads, and state type labels. We also allow the user to switch between ranking algorithms.

5.2 Comparing Multiple Groups

In addition to analyzing the behavior of a single group, experts also need to compare groups. They compare groups to find commonalities and differences in the types of hierarchies formed, the amount of rank contention and the number of interactions and how they occur. We support such analyses through a second dashboard that uses a small multiples display of the various visual representations previously discussed. Additionally, the dashboard includes a set of tools and an MDS plot in the panel on the left which help in selecting similar or dissimilar groups. An example the dashboard with a small multiples display of the state sequences is shown in figure 5.

To compare interactions and rankings between groups, the user can use a small multiple display of the balloon plots, rank evolution charts, and the music notations. Through the balloon plots users can identify groups that are either common in nature such as those where the top ranked subject performs the most interactions followed by the second ranked subject and so on. They can also find unusual groups such as those where the second-ranked individual performs the most interactions but the highest ranked subject earned its rank by dominating this second-ranked dominant subject with fewer interactions. Comparing multiple rank evolution charts in a single display allows the user to differentiate between groups based on rank contention. For example, in some groups rank is determined early indicated by a clear separation among subjects at the start, while other groups display continuous rank contention among a subgroup of subjects. Finally, comparing music notations allows the user to first compare the number of interactions that occur in each group. Users can either compare the density of the chart if local scaling is selected or the length of the chart if global scaling is selected. The music notations also allow users to visually compare bursting behaviors such as burst frequency and duration.

To compare the different hierarchies groups form and the sequence in which they were formed, users employ a small multiples layout of the state and state variant sequence. We provide users with tools to highlight a state in one group and check if it is highlighted or occurs in other groups. We also compute common sequences of user-specified lengths and highlight them across groups. By using the same coloring strategy used in the single group analysis, users were able to quickly recognize repeating patterns within and across groups. Repeated observations of two distinct colors appearing consecutively across states is indicative of frequent state or hierarchy transitions. Users can then further inspect these transitions with the state variant sequence to see if they are being caused by similarly ranked individuals or if it is different across groups. For example, users found that some state transitions that were common across groups usually occur due to subjects ranked 2 and 3 competing with each other. Additionally, we provide supporting tools to explore these state and state variant sequences. As in the single group analysis, we provide a heatmap of the COF and CSF for all groups and we also provide a list of common state transitions of a user specified length. These tools can be used to highlight states across multiple groups that can be further inspected by the user.

6 Case Study: Analyzing Dominance In Animals

To evaluate *PeckVis* we had two experts in the social science domain analyze their data with our tool. Both experts were social science professors and one of them was the expert we worked with to design the tool. Their goal was to analyze the formation of dominance hierarchies in three animals: mice, chicken, and fish. They had 14 groups of mice, 14 groups of chickens, and 17 groups of fish, each consisting of 4 subjects. The experts had three main tasks - exploring the formation of a hierarchy within a group, comparing groups of the same species and comparing the behaviors between two species.

6.1 Analyzing a Single Group

A user typically starts off the analysis of a single group by first inspecting the heatmap table showing the COF and CSF values of the hierarchies or states. There he or she can quickly receive an overview of the occurrence of different states or state variants and their stability in the group. Now if the user is interested in exploring highly stable states he or she may select states or variants with high CSF values. On the other hand, selecting states with low CSF values highlights unstable states that are indicative of rank contention. For example, in figure 1 the expert selected chicken group 10 for examination as every subject in the group changed rank during the course of the interactions. Next, he selected state 38 and 39 as he was interested in complete hierarchies which are six link states where every subject interacts with every other subject and he wanted to inspect the maintenance of these states.

Digging deeper, he inspected the state and state variant sequence. Through the state sequence, he observed that state 38 appears to encompass most of the interactions although the group forms state 39 for a brief duration. Upon further investigation, he learned through the state variant sequence that in fact two variants of state 38 - 38-1 and 38-7 - were formed with 38-1 being more stable. By inspecting the node colors he learned that subject one was gaining rank over subject two. But subject two fights back by attacking one and causes the creation of states 38-7. There was also some retaliation from subject four against two in state 39-2. Upon toggling the edge weighting, the expert found that the retaliatory actions did not seem to contribute too much to the state variants. To learn more about these interactions, the

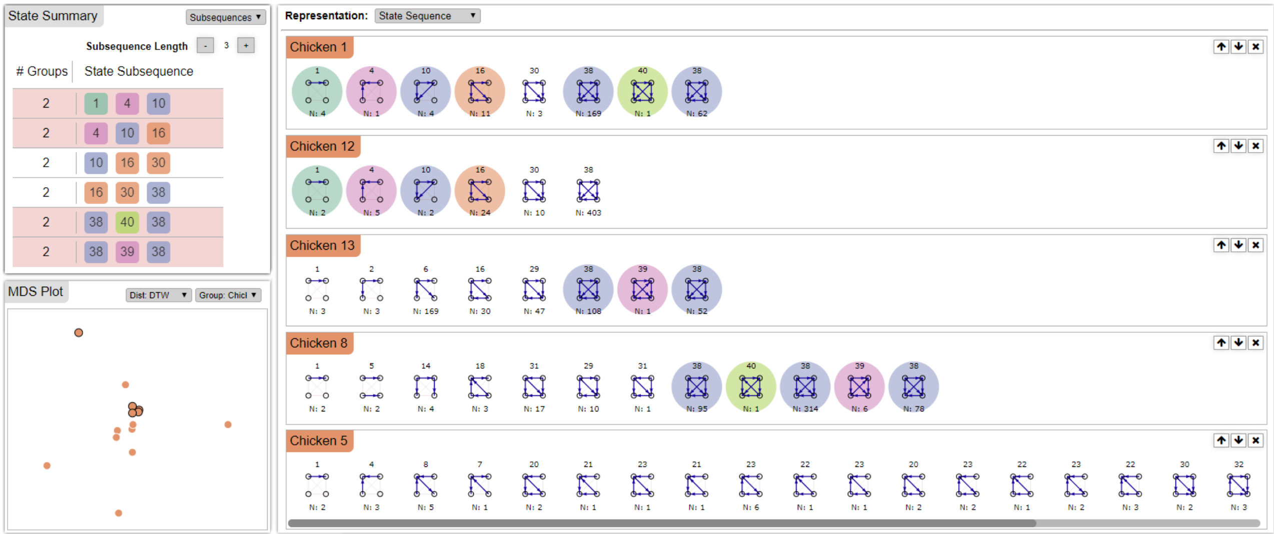

Fig. 5: The *PeckVis* interface for comparing dominance hierarchies between multiple groups of a species. On the left, we have two panels - a summary panel and an MDS plot below it. Here, with the MDS plot the user has selected four similar groups (1, 12, 13, and 8) highlighted at the center and an outlier group (group 5) highlighted at the top left of the plot. The summary panel above the MDS plot displays a list of the common state sub-sequences across all selected groups. The user can change it to show the COF and CSF heatmap. The panel on the right contains a small multiples display of the user choice of visual representation, in this case the user has selected the state sequence. To investigate the similarities between the sequences, the user selects the sub-sequences in the list which in turn highlight the corresponding states in the small multiples display. Now comparing the highlighted states, the user can investigate the similarities as discussed in section 6.2

expert selected the states with retaliations. He then inspected the music notation for the corresponding interactions. Here he saw that in each state the retaliation occurs only once, this led him to believe that if these interactions did not occur, the group would have formed an extremely stable six-link hierarchy in state 38-1. Additionally, he saw that the retaliatory states highlighted portions of the rank evolution chart in which lines intersect, that is rank contention occurs.

The experts analyzed groups from several species following a similar procedure and made multiple findings. They found that in mice and fish groups the top-ranked subject committed a bulk of the interactions but in the case of chicken, some groups had the second-ranked subject committing the most interactions but the top-ranked subject in these groups was actually dominating this second-ranked subject. They also found that chicken committed very few retaliatory interactions and formed stable hierarchies without much competition, the mice were a little more competitive but formed stable hierarchies as time progressed. On the other hand, fish were very active and competitive. The experts made the findings by examining the state variant sequence in conjunction with the music notation and rank evolution chart. They also used the state sequence to help navigate the state variant sequence. The state variant sequences for fish were extremely long and formed hundreds of state variants, however, the experts were engaged in the analysis and readily spent long periods of time exploring entire sequences in detail. They found that except for the most dominant fish, the other subjects kept competing with each other. This was evident through the rank evolution chart with the lines for these subjects constantly crossing each other. But on deeper inspection of the state variants, they found that these competing fish actually formed intermediate stable hierarchies for certain periods of time with each fish taking turns dominating the others for extended periods of time.

6.2 Comparing Multiple Groups

After investigating the individual groups and developing a hypothesis about the formation of hierarchies, the experts used the inter-group analysis tools to compare groups and validate their hypotheses. They started by selecting groups from the MDS plot for display in the small multiples layout. The experts first selected groups that were clustered together in the MDS plot and later moved to the outliers to investigate what was different in their hierarchy evolution. The experts first investigated the groups' balloon plots to check for similarities between groups based on the number of interaction by subject rank. They then moved to the rank evolution charts to check for similarities in rank contention, for example, in the case of chicken most groups formed clear hierarchies early on but some outliers had subjects that were interchanging ranks.

Next, the experts compared the groups based on the states they formed. Using the small multiples display of state sequences the experts were quickly able to identify common states among groups and what states made a group an outlier. For example figure 5 shows four similar chicken groups (1, 8, 12. and 13) and one outlier (5), selected through the MDS plot based on DTW distance between the state sequences. Among the similar groups, all of them have a majority of their interactions in state 38. They also have common state transitions such as in groups 1 and 12 which have identical sequences except for the last two states in group 1. Additionally, groups 1 and 8 have the state transition 38-40-38 while groups 13 and 8 have transition 38-39-38 in common. On the other hand, group 5 was an outlier as it had a longer state sequence which contained many states of which some were never formed by any other group, this was also confirmed by investigating the state's COF values and not having it sub-sequences appear in the panel on the left. Experts explored the state variant sequence in a similar manner. They also used the music notation but only superficially to explore bursting behaviors and compare the frequency of state formations by concentrating on the markers at the top of the music notation.

7 Case Study: Analyzing Debate Data

While our system is primarily designed to investigate hierarchies in animal groups, we can apply it to analyze dominance hierarchies in other situations just as well. One such situation is debates. To demonstrate this, we used our system to analyze aggressive behaviors in the 2016 U.S. presidential debates. For the purpose of this demonstration, we only categorized interrupts during debates as aggressive behavior. We used the transcribed debate data and interruption extraction technique from Stephanie Kirmer's post on Kaggle (https://www.kaggle.com/skirmer/interruptions-at-the-first-presidential-debate). With a more sophisticated analysis, one could extract aggressive statements by the candidates as well. For this analysis, we had two groups, the presidential debate with Donald Trump, Hilary Clinton, the moderator, and the audience, and the vice-presidential debate with Mike Pence, Tim Kaine, the moderator, and the audience.

We approached the analysis in a manner similar to that followed

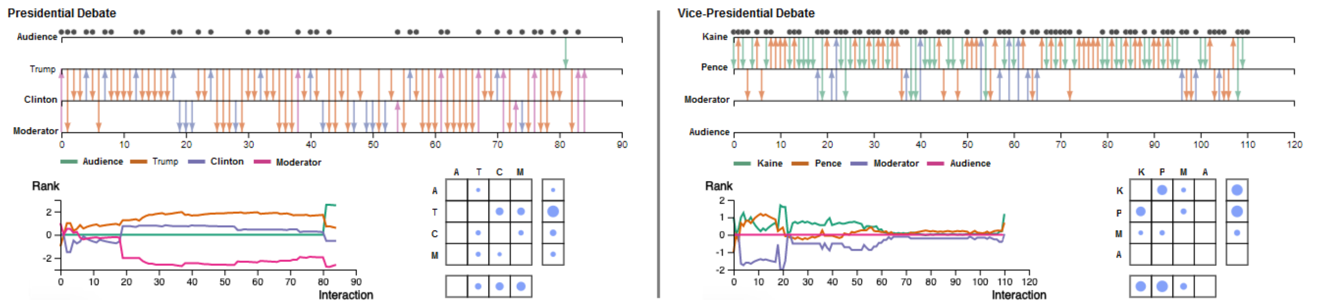

Fig. 6: The music notation representation, Elo ranking chart, and the balloon plot for the presidential (left) and vice-presidential (right) debates. In the presidential debates we see multiple bursts of interrupts from Trump but eventually the two debaters' ranks converged, it also shows he committed the most interrupts. But in the vice-presidential debate the candidates each had one burst of interrupts (between interactions 10-20 and 70-80). However, they both committed a similar number of interrupts and kept interrupting each other thus never creating a clear hierarchy.

by the experts as described in section 6. We first inspected the state sequence and noticed that in the presidential debate a complete six-link hierarchy (state) was formed but in the vice-presidential debate only a three-link hierarchy was formed. On further inspection, we learned that the reason for this was that the audience commits an interruption (of Trump) in the presidential debate but they do not interact in the vice-presidential debate. The interruption is the 81st interaction in the music notation shown in figure 6 (left). Next, we inspected the music notations, rank evolution and balloon plot of the two groups as in figure 6. Here we observed that in the debate with the presidential candidates Trump committed a burst of interrupts at multiple points in the debate. However, in the vice-presidential debates each candidate committed a burst of interrupts at the start of the debate after which the candidates and the moderator never committed more than two sequential interrupts. The ranking chart and the balloon plot informed us that in the presidential debate Trump committed more than half of the total number of interruptions and interrupted Clinton and the moderator equally. This also caused Trump to be ranked very highly during the debate. But towards the end, the ranks of the group members started to converge as the moderator interrupted Trump more frequently. Finally, the audience which stayed neural throughout the debate committed an interrupt against Trump. This event led the ranking algorithm to bump their ranking to the top spot. In the vice presidential debate the candidates interrupted each other equally and thus no single member was dominant for a long period of time. The state variant sequence for both debates was very long with very few stable states informing us that the groups never achieved a distinct hierarchy. This is reflective of highly competitive debates which these were. Finally, as we had just two groups, comparing them with the inter-group representations did not reveal any new information.

8 Discussion

In this work, we presented *PeckVis*, a visual analytics system to inspect the formation and maintenance of dominance hierarchies in small groups. Using our system, experts successfully investigated interaction sequences within groups and the hierarchies these interactions among group members form. We demonstrated the use of a sequence of node-link graphs in conjunction with a music notation and ranking chart in an interactive dashboard to investigate the contribution of interactions to a dominance hierarchy. Additionally, a small multiples layout of these representations assisted the experts in interactively learning about the similarities in hierarchy formation within and across multiple species.

The experts who worked with our system stated that it was intuitive and easy to use. They said that having multiple visualizations of their data interactively linked in a single interface allowed them to better understand the data and speed up their analyses. The experts informed us that by explicitly encoding the instantaneous rank, reversals, intransitive triads and state types in the state variant sequence, they were able to detect and investigate patterns that were hard to find. The enhancements to the music notation also helped the experts connect the state and state variant sequence to the interactions while better highlighting bursts. Additionally, they stated that using the comparative interface along with the user interactions made it easy to inspect the similarity among states. The experts often switched between the comparative interface and single group interface. They would first use the MDS plot to select similar and dissimilar groups and extract their differences in the comparative interface; they then went back to the single group interface to study the differences. For example, if a user found a group with a sequence of uncommon states in the comparative interface, he would investigate this sequence with the single group interface where he could examine the formation and stability of uncommon states. The experts stated that the DTW distance performed well to show groups that were similar based on their state or state variant sequences. However, the similarity measures were not able to help them identify groups that were similar based on their rank evolution or interaction matrix.

While our system was primarily designed to investigate hierarchies in animal groups, we showed that it can be adapted to analyze hierarchies in other situations just as well. The prime candidates would be situations involving competitive behavior by multiple entities such as individuals or teams competing for positions in a sports tournament and online games or to debates as demonstrated in section 7.

9 Limitations and Future Work

PeckVis was designed for social scientists working with small groups. Hence in its current form it does not scale well to the analyses of large groups. As group size increases the number of axes for each subject in the music notation would increase as well making it infeasible to use on a regular computer display. The state space makes use of node-link diagrams to represent the interaction networks but this would be inappropriate for large networks as demonstrated by previous work [21] [27]. A possible solution would be to use an alternate representation such as an adjacency matrix [6]. But these representations also have problems with scaling when comparing extremely large networks. Addressing the issue usually involves the use of glyphs or only showing differences between adjacent graphs both of which can be applied to our system by replacing the state sequence with a glyph or diff sequence.

Currently, *PeckVis* requires the manual selection of a ranking algorithm based on the user's domain expertise. In the future, we aim to devise a mechanism to at least suggest if not auto-select the most suitable ranking algorithm. We also believe that group comparison techniques can be improved. First, we aim to design a better measure of similarity to help experts find groups that have similar rank evolutions and interaction matrices. Second, a small multiples layout will be inefficient for experts to find similarities for a large number of groups. To address this issue we propose merging the state or state variant sequences of similar groups and representing them as a composite sequence highlighting areas of small difference much like in DiffAni [33]. We can then use the small multiples display to compare these composite sequences. Additionally, our evaluation only involved two users. We plan to conduct a more detailed evaluation that involves more users to further test our system design and incorporate refinements if necessary.

Acknowledgments

This work was supported in part by NSF grant IIS 1527200.

REFERENCES

[1] J. Abello and F. Van Ham. Matrix zoom: A visual interface to semi-external graphs. In *IEEE Symposium on Information Visualization (INFOVIS), 2004.*, pp. 183–190, 2004.

[2] F. S. Alallah, M. Nezhadasl, P. Irani, and D. Jin. Visualizing the decision-making process in a face-to-face meeting. In *IEEE 11th International Conference Information Visualization, 2007 (IV'07).*, pp. 168–176, 2007.

[3] B. Alper, B. Bach, N. Henry Riche, T. Isenberg, and J.-D. Fekete. Weighted graph comparison techniques for brain connectivity analysis. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 483–492. ACM, 2013.

[4] D. Archambault, H. C. Purchase, and B. Pinaud. Difference map readability for dynamic graphs. In *International Symposium on Graph Drawing*, pp. 50–61. Springer, 2010.

[5] B. Bach, N. Henry-Riche, T. Dwyer, T. Madhyastha, J.-D. Fekete, and T. Grabowski. Small multipiles: Piling time to explore temporal patterns in dynamic networks. In *Computer Graphics Forum*, vol. 34, pp. 31–40. Wiley Online Library, 2015.

[6] B. Bach, E. Pietriga, and J.-D. Fekete. Visualizing dynamic networks with matrix cubes. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pp. 877–886. ACM, 2014.

[7] A. Bang, S. Deshpande, A. Sumana, and R. Gadagkar. Choosing an appropriate index to construct dominance hierarchies in animal societies: a comparison of three indices. *Animal Behaviour*, 79(3):631–636, 2010.

[8] F. Beck, M. Burch, S. Diehl, and D. Weiskopf. A taxonomy and survey of dynamic graph visualization. In *Computer Graphics Forum*, vol. 36, pp. 133–159. Wiley Online Library, 2017.

[9] D. J. Berndt. Finding patterns in time series: a dynamic programming approach. *Advances in knowledge discovery and data mining*, pp. 229–248, 1996.

[10] N. Cao, Y.-R. Lin, F. Du, and D. Wang. Episogram: Visual summarization of egocentric social interactions. *IEEE computer graphics and applications*, 36(5):72–81, 2016.

[11] I. D. Chase. Music notation: a new method for visualizing social interaction in animals and humans. *Frontiers in zoology*, 3(1):18, 2006.

[12] H. A. David. Ranking from unbalanced paired-comparison data. *Biometrika*, 74(2):432–436, 1987.

[13] J. M. DiMicco, K. J. Hollenbach, and W. Bender. Using visualizations to review a group's interaction dynamics. In *CHI'06 extended abstracts on Human factors in computing systems*, pp. 706–711. ACM, 2006.

[14] P. Doreian. Actor network utilities and network evolution. *Social networks*, 28(2):137–164, 2006.

[15] N. Elmqvist, T.-N. Do, H. Goodell, N. Henry, and J.-D. Fekete. Zame: Interactive large-scale graph visualization. In *IEEE Pacific Visualization Symposium, 2008 (PacificVIS'08).*, pp. 215–222. IEEE, 2008.

[16] A. E. Elo. *The rating of chessplayers, past and present*. Arco Pub., 1978.

[17] M. Farrugia, N. Hurley, and A. Quigley. A.: Exploring temporal ego networks using small multiples and tree-ring layouts. In *In Proceedings of the 4th International Conference on Advances in Computer-Human Interactions (2011), ACHI*.

[18] M. Freire, C. Plaisant, B. Shneiderman, and J. Golbeck. Manynets: an interface for multiple network analysis and visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 213–222. ACM, 2010.

[19] S. Fu, Y. Wang, Y. Yang, Q. Bi, F. Guo, and H. Qu. Visforum: A visual analysis system for exploring user groups in online forums. *ACM Trans. Interact. Intell. Syst.*, 8(1):3:1–3:21, Feb. 2018. doi: 10.1145/3162075

[20] M. P. Gammell, H. de Vries, D. J. Jennings, C. M. Carlin, and T. J. Hayden. David's score: a more appropriate dominance ranking method than clutton-brock et al.'s index, 2003.

[21] M. Ghoniem, J.-D. Fekete, and P. Castagliola. On the readability of graphs using node-link and matrix-based representations: a controlled experiment and statistical analysis. *Information Visualization*, 4(2):114–135, 2005.

[22] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, 2011.

[23] L. Gou, X. Zhang, A. Luo, and P. F. Anderson. Socialnetsense: supporting sensemaking of social and structural features in networks with interactive visualization. In *IEEE Conference on Visual Analytics Science and Technology (VAST), 2012.*, pp. 133–142, 2012.

[24] J. Heer and D. Boyd. Vizster: Visualizing online social networks. In *IEEE Symposium on Information Visualization (INFOVIS), 2005.*, pp. 32–39, 2005.

[25] N. Henry and J.-D. Fekete. Matrixexplorer: a dual-representation system to explore social networks. *IEEE transactions on visualization and computer graphics*, 12(5), 2006.

[26] S. Kairam, D. MacLean, M. Savva, and J. Heer. Graphprism: compact visualization of network structure. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, pp. 498–505. ACM, 2012.

[27] R. Keller, C. M. Eckert, and P. J. Clarkson. Matrices or node-link diagrams: which visual representation is better for visualising connectivity models? *Information Visualization*, 5(1):62–76, 2006.

[28] J. B. Kruskal and M. Wish. *Multidimensional scaling*, vol. 11. Sage, 1978.

[29] J. Lehmann, A. Korstjens, and R. Dunbar. Group size, grooming and social cohesion in primates. *Animal Behaviour*, 74(6):1617 – 1629, 2007. doi: 10.1016/j.anbehav.2006.10.025

[30] W. B. Lindquist and I. D. Chase. Data-based analysis of winner-loser models of hierarchy formation in animals. *Bulletin of mathematical biology*, 71(3):556–584, 2009.

[31] P. Mutton. Inferring and visualizing social networks on internet relay chat. In *Proceedings. Eighth International Conference on Information Visualisation, 2004. IV 2004.*, pp. 35–43. IEEE, 2004.

[32] C. Neumann, J. Duboscq, C. Dubuc, A. Ginting, A. M. Irwan, M. Agil, A. Widdig, and A. Engelhardt. Assessing dominance hierarchies: validation and advantages of progressive evaluation with elo-rating. *Animal Behaviour*, 82(4):911–921, 2011.

[33] S. Rufiange and M. J. McGuffin. Diffani: Visualizing dynamic graphs with a hybrid of difference maps and animation. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2556–2565, 2013.

[34] F. Van Ham, H.-J. Schulz, and J. M. Dimicco. Honeycomb: Visual analysis of large scale social networks. In *IFIP Conference on Human-Computer Interaction*, pp. 429–442. Springer, 2009.

[35] C. Vehlow, F. Beck, and D. Weiskopf. Visualizing dynamic hierarchies in graph sequences. *IEEE transactions on visualization and computer graphics*, 22(10):2343–2357, 2016.

[36] T. von Landesberger, M. Gorner, and T. Schreck. Visual analysis of graphs with multiple connected components. In *2009 IEEE Symposium on Visual Analytics Science and Technology*, pp. 155–162, Oct 2009. doi: 10.1109/VAST.2009.5333893

[37] T. Von Landesberger, A. Kuijper, T. Schreck, J. Kohlhammer, J. J. van Wijk, J.-D. Fekete, and D. W. Fellner. Visual analysis of large graphs: state-of-the-art and future research challenges. In *Computer graphics forum*, vol. 30, pp. 1719–1749. Wiley Online Library, 2011.