

Exploring the Design Space of BioFabric Visualization for Multivariate Network Analysis

J. Fuchs¹  F. L. Dennig¹  M.-V. Heinle¹  D. A. Keim¹  and S. Di Bartolomeo¹ 

¹University of Konstanz, Germany



Figure 1: visualization displays the number of interactions in Game of Thrones — a multivariate network, with attributes on nodes (age of the character, number of episodes they appear in) and attributes on edges (number of interactions per season). On the visualization’s left side are five edge encodings, followed by four kinds of on-node encodings. At the bottom and right are juxtaposed visualizations for edges and nodes.

Abstract

The visual analysis of multivariate network data is a common yet difficult task in many domains. The major challenge is to visualize the network’s topology and additional attributes for entities and their connections. Although node-link diagrams and adjacency matrices are widespread, they have inherent limitations. Node-link diagrams struggle to scale effectively, while adjacency matrices can fail to represent network topologies clearly. In this paper, we delve into the design space of BioFabric, which aligns entities along rows and relationships along columns, providing a way to encapsulate multiple attributes for both. We explore how we can leverage the unique opportunities offered by BioFabric’s design space to visualize multivariate network data — focusing on three main categories: juxtaposed visualizations, embedded on-node and on-edge encoding, and transformed node and edge encoding. We complement our exploration with a quantitative assessment comparing BioFabric to adjacency matrices. We postulate that the expansive design possibilities introduced in BioFabric network visualization have the potential for the visualization of multivariate data, and we advocate for further evaluation of the associated design space. Our supplemental material is available on osf.io.

CCS Concepts

- **Human-centered computing** → Empirical studies in visualization; Graph drawings;

1. Introduction

Network visualizations are useful for understanding and exploring relationships between real-world entities. For instance, researchers commonly utilize node-link diagrams and adjacency matrices for displaying social [ACJM03], communication [BFN04], or biological networks [BMGK08], among many others. In addition, many real-world networks contain nodes or edges with one or multiple attributes. Multivariate network visualizations are helpful for representing network typologies showing the general structure, such as densely connected groups and sparse links between entities together with associated attributes. In Biology, for instance, visualizing network topologies together with experimental multivariate data helps to understand cellular processes [PLS*12] by visualizing protein interactions, chemical reactions, and catalyzing enzymes.

A key task for multivariate network visualizations is effectively visualizing network topologies along with their node or edge attributes. In a social network, an example task can be identifying the oldest person and their longest-lasting friendship with another person. However, there is a trade-off between effectively visualizing the network topology and displaying the associated multivariate attributes. Thus, researchers proposed various tailored visualizations and encoding for displaying multivariate networks, such as node-link diagrams and matrices, where additional information is visualized alongside the basic visualizations by juxtaposing the additional attributes in a table or encoding them directly in the cells of matrices.

For a comprehensive overview of multivariate network visualization, one can refer to the work of Nobre et al. [NMSL19], which delineates these visualizations into categories such as node-link, implicit tree, and tabular layouts. Specifically, the paper points out the efficacy of tabular layouts like BioFabric [Lon12] for delineating sparse networks that are laden with a multitude of node and edge attributes, which can be cumbersome to interpret in sparse matrix forms. BioFabric itself is highlighted for its ingenuity, offering a unique method to encode attributes of both nodes and connections.

BioFabric reimagines the visualization of complex networks—imagine a fabric woven with data where each row in a grid corresponds to a network point, and each column signifies a connection. A connection is marked by a colored spot on the grid, crafting a visual fabric that simplifies the identification of patterns and relationships within the network.

Despite the promise shown by BioFabric, it remains underexplored, particularly in the breadth of design possibilities it presents for embodying multivariate attributes. Because of this grid-like structure, additional attributes for nodes and edges can be visualized on the borders of the network visualization, as visible in Figure 1. Our research seeks to bridge this gap with an in-depth exploration of the design potential within BioFabric. Examining how multivariate data is visually encoded, we intend to scrutinize the influence of different design methodologies on the interpretability of the visualization. This endeavor not only aims to affirm BioFabric's utility but also strives to push the envelope in visualizing complex data structures. Therefore, we contribute:

- A design space exploration of BioFabric for multivariate networks.

- A discussion regarding how to effectively encode multivariate attributes into BioFabric visualizations.
- A quantitative comparison of BioFabric to matrix visualizations.

Our supplemental material, which includes an interactive implementation, the user study, the results from the study, the analysis and an appendix with some further observations, is available on osf.io, at https://osf.io/8j42t/?view_only=ab69785be4bc48608d64c2a0550f4b04.

It should be taken into account that the effectiveness of BioFabric is also affected by node and edge ordering [VBP*21, DBPB*22] — an aspect that is not considered in the scope of this paper. Instead, we focus on the encoding of additional data dimensions into the visualization, which can be applied to any ordering of nodes and edges.

2. Related Work

Multivariate Network Visualizations: In their survey, Nobre et al. [NMSL19] surveyed and categorized eleven multivariate network visualization approaches for displaying network topologies and their associated attributes. According to the authors, the most frequent multivariate network representations are node-link layouts, including *on-node/edge encodings*, *attribute-driven faceting*, and *attribute-driven positioning*. The recommended usage of multivariate node-link layouts is for small networks (< 100 nodes) with a limited number of node and edge attributes, for instance, on node-encodings for social networks [ACJM03] or overlaid edge bar-charts [SSSE16]. Implicit tree layouts are useful for hierarchies, including *inner nodes & leaves* and *leaves* visualizations, such as TreeMaps [SDW09] or SunBurst [SZ00] visualizations. Yet, such implicit tree layouts do not support displaying multiple attributes simultaneously.

Tabular layouts display the network topology and the attributes in a table-based layout, encompassing *adjacency matrices*, *quilts* [BW11], *massive sequence views* [CHZ*07], and *BioFabric* [Lon12]. For instance, BioFabric visualizes nodes as horizontal rows and vertical lines (columns) representing edges between nodes. The advantage of such tabular layouts is that the row- and column-based visualizations can be easily used to depict additional attributes. Moreover, Nobre et al. discuss different view operations, describing how multiple coordinated views can be combined through *juxtaposed*, *integrated*, and *overloaded* views. In particular, the authors emphasize that BioFabric has a unique potential for visualizing a large set of attributes using different view operations, such as a juxtaposed view, similar to *Liniage* [NGCL19] and *Juniper* [NSL19], which leverage a Tree+Table approach for multivariate network analysis. However, up to this day, the design space for visualizing multivariate network data with BioFabric is insufficiently explored and empirically studied.

Evaluation of Multivariate Network Visualizations: There is a large body of empirical studies for network visualizations without multivariate attributes [YAD*18]. However, most empirical studies focus on the relationship between network topology and tasks. For instance, Ghoniem et al. [GFC05] compared *topology-based tasks* on node-link diagrams and adjacency matrices for networks with

varying sizes and densities. Only a few empirical studies investigate multivariate networks and *attribute-based tasks*.

First, Alper et al. [ABHR*13] investigated comparison tasks for weighted networks and showed that adjacency matrices outperform node-link diagrams for overloaded views. Next, Abuthawabeh et al. and Wybrow et al. [WEF*13, ABZD13] compared and showed that participants could identify the same graph structures for multi-type edges in adjacency matrices and parallel node-link visualizations. Schöffel et al. [SSSE16] compared node-link diagrams encoding multiple edge attributes as bar charts and partially filled bar charts, showing their usefulness for comparison tasks. In another publication by Nobre et al. [NWLH20], the authors compared node-link diagrams with on-node encodings with adjacency matrices with juxtaposed tables. Their crowd-sourced empirical study showed that both approaches perform similarly on edge attribute tasks. Moreover, the authors showed that node-link diagrams with few attributes are suitable for path- and neighbor-related tasks and that adjacency matrices help identify clusters, providing less distraction for larger sets of attributes.

Most of the existing related work empirically investigated network topology tasks, with a few studies investigating multivariate analysis tasks on node-link diagrams and adjacency matrices. BioFabric was originally designed and evaluated in the domain of Bioinformatics through case-studies [Lon12]. Valdivia et al. used BioFabric to visualize hypergraphs and evaluated their design with a small scale usability study [VBP*21]. Until today, none of the possible alternative multivariate network visualizations described by Nobre et al. [NMSL19] have been comparatively evaluated, including BioFabric. Thus, we explore techniques to encode multivariate data with BioFabric, describing edge encodings for a high number of attributes and providing empirical guidance on the visualization design for multivariate analysis tasks.

3. Techniques to Encode Multivariate Data on BioFabric

BioFabric visualization primarily represents *nodes as horizontal lines* and *edges as vertical line segments* connecting these horizontal lines. This method was developed to help better understand and represent large-scale biological networks, as traditional node-link diagrams can become overly cluttered and challenging to interpret [Lon12]. BioFabric combines the table-line structure of adjacency matrices with the visual appearance of node-link diagrams and, thus, reads like a combination of the two. Relations are shown via edges, just like in a node-link diagram. However, horizontal lines represent nodes, allowing for a tabular arrangement like in an adjacency matrix. The combination of these techniques can offer a vast amount of opportunities to encode additional, multivariate data on the visualization.

In their survey about multivariate data visualization, Nobre et al. [NMSL19] explored a few options for this specific task — encoding multivariate data on BioFabric. In this chapter, we delve deeper and expand their proposed techniques and the affordances offered by them, categorizing the design options into three categories: *embedding*, *transformation*, and *juxtaposition*.




Figure 2: Example of a double encoding in BioFabric. Edge attributes are simultaneously encoded as embedded bars on-edge lines and in the juxtaposed table.

3.1. Embedding

The *embedding* of multivariate attributes is defined as *integrating* glyph-like attribute visualizations into a network's topology visualization. An example is to display a histogram-like glyph on the edges of BioFabric, as shown in Figure 2. These attribute visualizations can be subject to *translation*, *rotation*, and *scaling* to place them on top or close to network entities, i.e., nodes and edges.

Embedded on-node encoding: The overall idea is to include a small visualization on top of node lines to encode attributes. A similar concept exists in node-link diagrams where nodes are replaced or substituted by visualization such as bar and line charts [JKS06, NMSL19], although in node-link — because nodes are not represented as lines — the possibilities for encoding are very different.

Adding more complex representations requires more space for a detailed analysis. However, increasing network size (i.e., the number of nodes) shrinks the visualization space since more horizontal lines (i.e., rows) must be displayed. The tabular topology layout introduces such a limitation, which is also true for the adjacency matrix. As a result, complex visualizations or detailed glyphs can hardly be added to networks with many entities.

Embedded on-edge encoding: Like node attributes, edge attributes can be represented as an embedded visualization. Schöffel et al. introduced an encoding for multiple edge attributes in node-link diagrams [SSSE16]. The most supportive variation encodes these as bars of varying height beside each other and on a common baseline. However, this technique is limited for standard node-link diagrams since edge crossings can introduce clutter affecting the attribute visualization, which may overlap. Thus, integrating this encoding into node-link diagrams is only effective on planar graphs.

Due to the parallel arrangement of edges, BioFabric avoids this restriction and allows for a smooth integration of on-edge encoding. The orthogonality provides more clarity to the viewer when differentiating the lines. All bars can be arranged in the same direction. Thus, the attribute values cannot be falsely interpreted as negative, which may be another problem in the node-link diagram. The on-edge

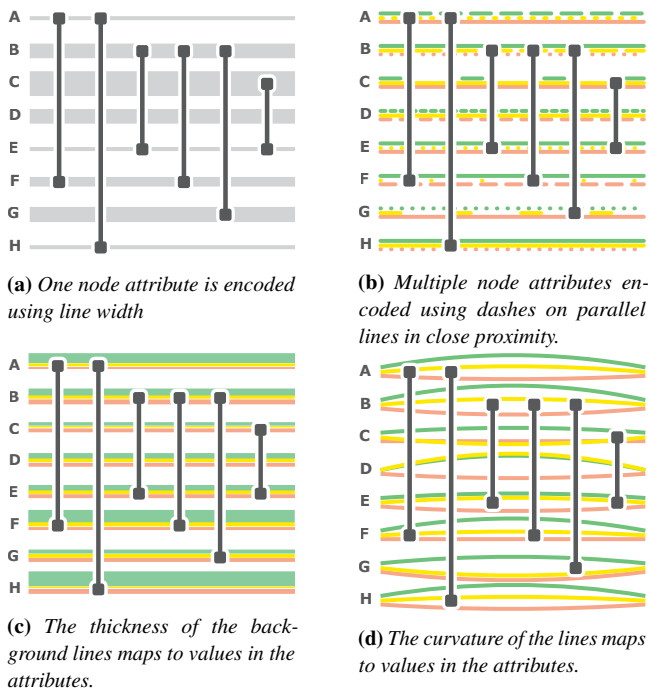


Figure 3: Examples of transformed node encodings in BioFabric.

encoding is even more suitable for the adjacency matrix. Multiple edges can be represented in the same row, sharing the same baseline. This simplifies the value comparison of edge attributes when analyzing a node's neighborhood. However, increasing the number of vertices in the adjacency matrix negatively influences the available space to display an on-edge encoding.

3.2. Transformation

Our definition of *transformation* is the modification of the visual appearance of edges or nodes. Attributes might alter, for example, the *color*, *texture*, *thickness*, or *saturation* of a network entity, i.e., edge or node [BDT21]. We want to highlight that their position and start and end points remain the same.

Transformed node encoding: BioFabric allows for a different set of node transformations since nodes are represented as horizontal lines. Figure 3a provides examples of the on-node encoding of one attribute.

Increasing network size shrinks the visualization space for size-based encodings. Therefore, the network properties must be considered during design to allow a good differentiation of the encoded attributes. The horizontal node lines in BioFabric provide a unique possibility to visualize attribute encodings. One node attribute transforms the line by varying width, color, texture, or other visual attributes (see Figure 4). This approach can be extended to the representation of multiple attributes. In Figure 3b, we show straight parallel lines with varying dashes and differentiate them by color. Independent of its dashed texture, each bundle of lines has a noticeable distance to its neighboring node lines, allowing a clear comparison.




Figure 4: Transformed node and edge encoding refers to using visual variables like color, texture, shape, thickness, saturation, and transparency to represent information on the edges of a visualization [BDT21]. Embedded on-edge encoding describes labels and glyphs added to enhance the representation of multiple attributes.

However, this node attribute encoding loses expressiveness with an increasing number of node attributes or nodes.

The best use of those visual variables depends on the underlying data. For example, the curvature is influenced by the number of possible values and might cause more overplotting if attributes share the same value (see node D in Figure 3d); on the other hand, texture or dashed lines are influenced by the number of attributes; and width is influenced by both the number of attributes and the possible values.

Such node transformations are also common in node-link diagrams by coloring nodes or changing their size. Varying the graphical representation of rows or columns in an adjacency matrix will negatively impact the tabular setup of the visualization. Therefore, only interactive approaches allow for a structural distortion of the matrix.

Transformed edge encoding: Edges have a high visual impact in node-link approaches. Many applications using node-link diagrams transform the visual appearance of edge lines to communicate additional information like direction [vdEvW14], uncertainty [GHL15], or dynamic changes [ABHR*13]. These design variations can be adopted since BioFabric also communicates relations via vertical lines. Different design alternatives can be seen in Figure 1 or Figure 4. Since edges may vary in length, a color encoding might be perceived differently due to the strong interplay between color and size.

3.3. Juxtaposition

We define the *juxtaposition* of attribute visualizations as a combination of displaying the network topology and an accompanying attribute visualization with an *unambiguous alignment* of node and edge representations and their respective attributes. Since BioFabric enforces a tabular layout on networks, it offers the opportunity to juxtapose more visualizations right next to the network topology. Therefore, we can separate the network topology from the multivariate attributes with juxtaposed views. The usefulness of this approach has already been shown with adjacency matrices representing addi-

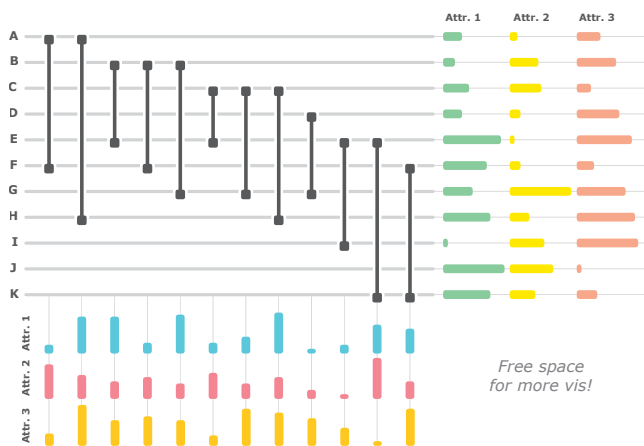


Figure 5: Juxtaposed visualizations show node attributes (right) and edge attributes (bottom) with aligned tabular views. The figure also illustrates the definitions of the measures used throughout Section 3.

tional node attributes [NWHL20]. Node-link diagrams do not offer such an arrangement.

Unlike adjacency matrices, in BioFabric, it is also possible to represent edge attributes in a separate view. Thus, without interaction, the tabular topology layout allows the alignment of node and edge attributes in a juxtaposed view (Figure 5). Since the representation of attributes and the topology do not overlap, multiple edge or node attributes can be displayed, offering a scalable approach regarding the number of attributes. However, with an increasing number of nodes, the space to represent edge attributes shrinks. Likewise, increasing the number of edges limits the space to show node attributes. Another major benefit of this view is the application of node- and edge-orderings based on an attribute value’s value through simple table interactions, e.g., selecting an attribute in a juxtaposed visualization.

Juxtaposition is flexible, so any visualization that is tabular in structure can be used to accompany BioFabric. If both a node and an edge attribute visualization are used to show multivariate attributes, free space is available at the bottom right corner to display additional information, for example, *detail-on-demand* or *summary* visualizations (Figure 5).

3.4. Other Design Considerations

In this section, we discuss other design considerations which do not deal with the representation of multivariate data, but are important to take into account when designing BioFabric visualizations.

Double Encoding: Each attribute visualization technique can be used to embed the representation properties more than once. Figure 2 combines the encoding of attributes on the vertical line and juxtaposed tables. Although this perceptually complementary view is redundant, it provides support for analysis tasks [CBDM17]. Viewers can choose the preferred technique, and the views can be used to counter-check their findings. Other combinations are also possible, such as juxtaposition with transformation or embedding and transformation. This high degree of flexibility is impossible in node-link

diagrams or adjacency matrices. However, double encoding may have distinct drawbacks: increasing the required screen space, especially for juxtaposed views, and increasing clutter and decreased readability.




Figure 6: Shadow links added are colored in light gray. The red polygon highlights a staircase pattern.

Shadow Links: Given the approach used in BioFabric, the edges incident on a node are, by design, distributed along the full length of the horizontal node line. One disadvantage of this approach is that an edge is more closely tied visually with only one of the endpoint nodes and can be conceptually disconnected from the other node. To address the issue of edges of a single node being potentially visually far apart, Longabaugh [Lon12] introduced shadow links — copies of the original edge, such that every edge is drawn twice.

Shadow links allow all edges of a node to be in one location next to each other, enabling a more direct visual tracing and comparison of the degree of nodes by the size of these groups. The main drawback, however, is that the width of the BioFabric plot doubles with respect to the number of edges and reduces the space available for edge encodings.

Background Highlighting: The background of a node or edge line can be highlighted, as it is typical for tables. According to the “Common region” Gestalt Law [Ban94], coloring the background helps to group rows or columns. Therefore, this design variation provides guidance while following a node or an edge line, either for retrieving an associated value or determining the neighborhood. This technique can also be applied to adjacency matrices if the cells do not contain a color-transformed edge encoding.

Interactive Highlighting: Interactive highlighting or filtering is another technique explored in previous literature [LTdSPR17, EHBvW14, vdEHBvW13] in combination with BioFabric visualization. A straightforward approach is to highlight connected nodes when hovering over an edge (Figure 7a) and/or highlight incident edges when hovering over a node (Figure 7b). This method of highlighting makes path tracing easier since lines that need to be followed are more accessible to focus on. Especially, highlighting incident edges of a node is an alternative to shadow links.

Reordering Among the elements that most definitely will influence the perception of network topology in BioFabric, there is the ordering of the nodes and edges. The problem has been mentioned — although without proposed solutions — already in previous research [VBP*21, DBPB*22]. Indeed, the superficial and inattentive combination of the order of the two sets of network elements can produce unnecessarily longer, harder-to-track edges — while, conversely, using clever sorting strategies might create structures in the visualization that can help highlight topological features, such as the staircase pattern shown in Figure 6. While we believe that BioFabric could benefit from many matrix reordering strategies, it is important to consider the unique properties of the visualization at hand, and we believe that this task would require its own in-depth paper. For this reason, we deem this to be future work.

Space Considerations In a comprehensive discussion about design decisions related to BioFabric, it would also be important to consider




Figure 7: Examples of interactions on BioFabric visualization.

the amount of data that can be displayed on such a visualization and how the available screen estate influences the effectiveness of the various encodings. For instance, the number of edges and nodes to be displayed will determine the maximum size allowed for embedded and transformed encodings, which can influence the separability of their colors and/or shapes [SS19].

4. Quantitative Evaluation

We conducted a quantitative user study to reveal the strengths and weaknesses of BioFabric for multivariate network analysis. Our inspiration was drawn from Nobre et al.'s survey, which emphasized the unexplored potential of utilizing BioFabric for displaying edge attributes with juxtaposed views [NMSL19]. For the experiment design, we took inspiration from the most recent experiments about node-link diagrams and adjacency matrices from Nobre et al. [NWHL20]. We partially replicated their evaluation by transferring the visual encodings to an adjacency matrix and BioFabric. Since we see the high potential of BioFabric in communicating multiple edge attributes, we focused solely on this aspect in our experiment. Furthermore, interaction techniques were not supported since the ability to highlight the network topology would considerably reduce the difficulty of completing analysis tasks.

4.1. Experiment Factors

We considered the factors *visualization technique*, *task*, *size* and *density* of a network. Based on the recent experiment conducted by Nobre et al. [NWHL20], we chose the adjacency matrix (AM) and BioFabric (BF) due to their similar tabular structure avoiding overplotting. The authors used a bar chart to visualize node attributes in the node-link diagram. This representation can be replicated as an edge encoding for AM. We embedded a bar chart within the intersection cell to visualize multiple edge attributes. Nobre et al. further aligned a table to the rows of AM to represent node attributes. We transferred this approach to BF and visualized each edge attribute in a separate column of a juxtaposed table. Figure 8 illustrates an example of both techniques from our study.

We reordered the nodes using the Reverse Cuthill McKee-Algorithm [GL81]. This method uses a heuristic for reasonable bandwidth reduction at a very low computational cost [dOC15]. As a result, distances between adjacent nodes in AM and lengths of




Figure 8: Visualization design for our user study. The examples visualize a network of friendships with 10 nodes and a density of 10%. AM (left) encodes four edge attributes in a bar chart within the cells. BF (right) encodes multiple edge attributes in a juxtaposed table below the topology visualization.

the edges in BF are minimized [VBP*21]. BF further provides the opportunity to vary the ordering of edges. We tried to provide an optimal ordering for each task to benefit most from this advantage for BF. For topology-based tasks, we sorted the edges based on the appearance of adjacent nodes. For tasks focusing on only one edge attribute, we sorted the edges of the specified attribute in descending order. The edges in tasks including two attributes are represented in descending order of the first specified attribute. Edges that agree in their values on the first attribute are sorted in descending order of the second attribute. We further applied the ordering based on adjacent nodes when edges agreed on all specified attributes. Feedback from our pilot studies implied that this approach allows a better estimation of the attribute value ranges in large and dense networks. Figure 9 presents an ordering example based on two attributes.

For each task, the visualization was drawn at a resolution of 2050×1300 pixels. The juxtaposed table for visualizing edge attributes in BF occupied one-third of the provided space. We refrained from visualizing attributes when the analysis did not involve the inspection of attributes. In both techniques, we designed the representations of the network topology to remain monochrome. This allows the use of color for the visualization of edge attributes. We employed a qualitative color map from *ColorBrewer* [BH] to differentiate the attributes within an edge.

4.1.1. Tasks

In our experiment, we aimed to analyze various attributes within a network. As such, we narrowed down our task selection to those focusing on attribute-based analysis [LPP*06, PPS14]. While earlier user studies mainly focused on analyzing node attributes or a limited set of edge attributes, we focused on multiple edge attributes.

T1 - Neighborhood Inspection: Given the visualized network, participants are asked to select all adjacent nodes of a target node. This task involves identifying the target node, scanning all outgoing edges, and retrieving the source nodes.




Figure 9: Interface of our user study. The panel on the left provides the task description, legend, and answer field. The visualization of the network is displayed to the right. In this example, participants on the BF condition had to solve task 3 on a network with 50 nodes and 10% density. The correct answer is selected on the left.

T2 - Extreme Value Detection: Given the visualized network, participants have to select the adjacent node of a target node with the second highest data value for a specified attribute. After identifying the target node, this task requires comparing the target attribute value across all adjacent edges of the target node. The values are judged based on the length of the corresponding vertical bar. Then, the source node has to be selected. As identifying the highest value is trivial on BF when the edges are ordered accordingly, we refrained from using that task.

T3 - Within-Edge Comparison: Given the visualized network, we ask participants to select all adjacent nodes of a target node whose edge has a higher value in a first specified attribute than in a second specified attribute. Similar to task 1, all outgoing edges of the target node must be scanned, but the node selection includes comparing the attribute values based on the lengths of the bars.

4.1.2. Network Size & Density

We created synthetic data of undirected networks without self-loops using Python's *networkX* package [HSS08]. To test the scalability of each visualization technique, we used network data of varying sizes and densities. We selected network properties according to both a comprehensive survey on network visualization [YAD*18] and our experimental limitations. Since interaction or aggregation techniques were not allowed, we came up with 20 nodes for the small network, 50 for a medium-sized network, and 80 for the largest network.

In our experiment, density is defined as $d = \frac{e}{n(n-1)}$ where n denotes the number of nodes, and e is the number of edges [YAD*18]. Given previous studies on attribute-based tasks, we chose a minimal density of 2.5%, a medium density of 6.25%, and a maximum density of 10% [NWHL20, CBDM17].

To provide more context for the participants, we decided to rep-

resent networks of friendships. A person in this network is illustrated as a node. An edge represents the friendship between two persons. Each friendship is described by four edge attributes (i.e., years of friendship, distance, interactions per week, and common hobbies) whose values were assigned randomly. We chose a number equal to the numerical node attributes evaluated in Nobre et al. to allow the hypothetical transfer of results [NWHL20]. Like in related experiments, we varied the attribute values between 6 levels [ABHR*13, CBDM17].

4.2. Experiment Design

We used a split-plot design with the variable *visualization technique* as the between-subject to directly compare the designs (i.e., AM and BF) on the same tasks. The within-group design was applied to the independent variables *task*, *size* and *density*. The dependent variables were the measured *accuracy* and *completion time*.

4.2.1. Hypotheses

Based on experience, previous pilot studies, and literature [Lon12, NMSL19, NWHL20], we derived the following hypotheses:

H1 - Scalability Hypothesis: BF is more scalable concerning the size and density of the network than AM. This hypothesis is based on Longabaugh's work, which introduced BF to address scalability issues of networks [Lon12]. Adding an edge to a visualized network does not decrease the quality of the existing topological structures. Thus, the authors describe BF as a precise technique for displaying networks apart from their scale. In direct contrast, they mention that visualization space in AM increases quadratically with the number of nodes. As most networks are relatively sparse, this space usage is considered less efficient.

H2 - Neighborhood Hypothesis: AM outperforms BF in accuracy and completion time on the neighbor search when no edge attributes

are involved. Although both techniques allow for a systematic approach to scan adjacent edges, we expect *BF* to be more error-prone and take more time to retrieve a neighbor as multiple lines with changing directions must be followed. Nobre et al. share the same assumption in their extensive survey about multivariate network visualization [NMSL19].

H3 - Edge Sorting Hypothesis: *The judgment on attribute values on BF results in a better accuracy and completion time than in AM.* The topological structure of *BF* benefits from sorting the edges by attribute values. Also, the separate view provides more space to visualize attributes. Compared to embedded bar charts, the sorted table representing node attributes is strongly favored when searching for an extreme value [NWHL20].

H4 - Within-Edge Comparison Hypothesis: *For tasks involving the comparison of attributes within individual edges, AM is more accurate and requires less completion time than BF.* The bar chart in the cells of *AM* supports a direct comparison of multiple attributes as the bars are arranged beside each other and use the same baseline. In *BF*, we expect the effect of sorting the edges on multiple attributes to be outweighed by the juxtaposed layout of the edge attribute visualizations. This hypothesis is based on previous research on embedded bar charts that resulted in better performance when comparing attributes of neighbors [NWHL20].

Participants: We recruited 28 volunteers from our local student population to participate in our study. 17 identified themselves as male, 11 as female. All participants had normal or corrected-to-normal vision and did not report color weakness. Their ages ranged from 20 to 31 years (M: 22.96, SD: 2.56). Most participants finished high school. Only four had obtained their Bachelor's, and one participant had a Master's degree. 13 participants were currently engaged in the field of computer science. The others represented diverse academic backgrounds (biology, economics, law, mathematics, sports, and psychology). Based on a self-assessment, almost all participants were generally familiar with computers (M: 4.04, SD: 0.96, on a range where 1 means not familiar at all, and 5 means very familiar). However, fewer participants had experience with data visualization (M: 2.57, SD: 1.26). While our recruitment process favored university students, we aimed to balance their academic backgrounds and prioritize individuals with strong computer knowledge but limited experience in data visualization, representing a broader population.

Procedure: The study was conducted in a quiet setting at our university. The participant sat in front of a 27 inch display of 2560×1440 pixels. At approximately 60 cm distance, participants interacted with a web-based application (Figure 9) using a mouse. The experimenter guided the participants through an introduction, introducing them to the dataset. This structured approach allowed for intermediate examples of identifying nodes, tracing edges, and assessing edge attributes. In the condition of *BF*, the experimenter additionally drew the participants' attention to the ordering of edges, which we took advantage of in the study.

After this introduction, the experimenter explained the structure of the study. We divided the experiment into three sessions, each presenting a different task type. We began with the topology-based task, continued with the task on one attribute, and finished with comparing two attributes. Each session started with an introduction to the task in the web-based application, giving caution to possi-

ble misunderstanding. Nevertheless, the experimenter was present throughout the participation and encouraged participants to ask questions if a task was unclear. The participants got familiar with the task type by solving three training tasks on networks of 20 nodes and a density of 2.5%. These tasks were designed to train them on the technique and become comfortable with this study's interface. To ensure their understanding, participants had to answer at least two of the training tasks correctly to proceed with the study. After the training, nine tasks of the same type were presented in increasing order of difficulty. We first increased the density on a fixed-sized network, then continued with the next larger size. Overall, this resulted in 36 trials. Finally, a form collected demographic information and feedback on the study anonymously. We compensated each volunteer with 10€ for one hour. We estimated the completion based on pilot experiments.

5. Results

This section reports significant results ($p < .05$) from our quantitative analysis. We recorded the performance in terms of accuracy and completion time. The accuracy was calculated by summing the number of correctly selected nodes, subtracting the number of falsely selected nodes, and dividing by the total number of correct nodes. We report mean values (M) for both measures with a 95% bootstrapped confidence interval (CI). We use the format $[LL, UL]$ where LL is the lower limit, and UL is the upper limit of the interval.

As assessed by Levene's [Lev60] and Shapiro-Wilk [SW65] tests (with $p < .05$), the data on accuracy and completion time is non-parametric. Thus, we used Spearman's rank correlation coefficient [Dan87] to explore the relation of the performance to *size* and *density*. Since the tasks varied in purpose and accomplishment, we analyzed the differences between the *visualization techniques* within each *task* independently. We based our analysis on Mann-Whitney U tests [MW47] and reported U statistics, p -values, and the Common Language Effect Size [MW92]. Detailed results are reported and discussed in Table 1.

5.1. Scalability

Figure 10 illustrates the effect of network size and density on each technique's performance. Independent of the task type, the performance of *BF* seems less scalable with respect to network *size* than *AM*. The accuracy of participants working with *BF* decreased drastically when the network *size* was increased from 50 to 80 nodes. The computed Spearman's ρ revealed that *size* and accuracy had a moderate correlation ($\rho = -0.41, p < .0001$). Further, Figure 10 indicates a linear relationship between the network *size* and completion time on *BF*. These factors correlated strongly ($\rho = 0.59, p < .0001$). For *AM*, the accuracy and the network *size* correlated only weakly ($\rho = -0.15, p = .003$). However, the increase in network *size* led to increasing completion time. We identified an even stronger correlation compared to *BF* ($\rho = 0.74, p < .0001$).

BF was also more affected by network *density* than *AM*. Participants working with *BF* were less accurate and required more time when the *density* was increased. Tests revealed that the number of edges correlated weakly with accuracy ($\rho = -0.27, p < .0001$) and time ($\rho = 0.25, p < .0001$). *AM* was less affected in the completion time ($\rho = 0.14, p = .0067$).




Figure 10: Effect of network size and density on the performance independent of the task. The error bars are confidence intervals.

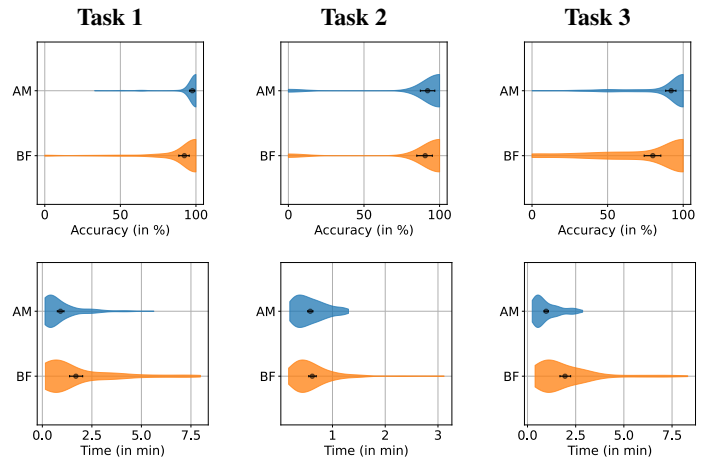


Figure 11: Experiment results regarding accuracy and completion time for all three tasks. The mean is encoded as a point, and the error bars indicate confidence intervals.

Task 1 - Neighborhood Inspection						
Accuracy	<p>Although both techniques provide high accuracy, participants working with AM are significantly ($U = 6701.5, p = .0002, CLES = 0.58$) more accurate ($M = 97.64, CI = [95.71, 99.22]$) compared to BF ($M = 92.33, CI = [88.67, 95.56]$). The results reveal that the difference in performance is affected by size. While it did not differ in small networks, AM outperformed BF on networks with 50 and 80 nodes. Also, AM showed fewer errors than BF on networks with 2.5% and 10% density. Results did not significantly differ on networks with a density of 6.25%.</p>	Size	20	$U = 882.5$	$p = .5135$	$CLES = 0.5$
		50	$U = 697.5$	$p = .0025$	$CLES = 0.6$	
80		$U = 659.0$	$p = .0056$	$CLES = 0.63$		
Density		2.5%	$U = 776.0$	$p = .0245$	$CLES = 0.56$	
6.25%		$U = 819.5$	$p = .1236$	$CLES = 0.54$		
10%		$U = 649.5$	$p = .0036$	$CLES = 0.63$		
Time	<p>Participants were significantly ($U = 10166.0, p = .0001, CLES = 0.64$) slower when working with BF ($M = 1.69, CI = [1.38, 2.03]$) compared to AM ($M = 0.92, CI = [0.75, 1.1]$). Similar to the accuracy, the difference in completion time was affected by size and density. AM required significantly less time than BF.</p>	Size	20	$U = 1347.0$	$p < .0001$	$CLES = 0.76$
		50	$U = 1204.0$	$p = .002$	$CLES = 0.68$	
80		$U = 1267.0$	$p = .0003$	$CLES = 0.72$		
Density		2.5%	$U = 1103.0$	$p = .0243$	$CLES = 0.63$	
6.25%		$U = 1350.0$	$p < .0001$	$CLES = 0.77$		
10%		$U = 1271.0$	$p = .0003$	$CLES = 0.72$		
Task 2 - Extreme Value Detection						
Accuracy	<p>There were no significant differences between the visualization techniques. The overall similar accuracy of AM is $M = 92.06, CI = [87.3, 96.07]$, while BF has $M = 90.48, CI = [84.92, 95.24]$. When considering network size, AM was more accurate than BF for larger networks with 50 and 80 nodes.</p>	Size	20	$U = 882.0$	$p = 1.0$	$CLES = 0.5$
		50	$U = 1008.0$	$p = .0272$	$CLES = 0.57$	
80		$U = 714.0$	$p = .0203$	$CLES = 0.4$		
Density		2.5%	$U = 882.0$	$p = 1.0$	$CLES = 0.5$	
6.25%		$U = 945.0$	$p = .1725$	$CLES = 0.54$		
10%		$U = 777.0$	$p = .1588$	$CLES = 0.44$		
Time	<p>Again, no significant differences existed between visualization techniques. Participants needed slightly more time with BF ($M = 0.62, CI = [0.55, 0.69]$) than with AM ($M = 0.58, CI = [0.53, 0.63]$). The longer completion times are more apparent in smaller networks (20 nodes) or really dense networks (10%), resulting in significant effects.</p>	Size	20	$U = 1253.0$	$p = .0005$	$CLES = 0.71$
		50	$U = 653.0$	$p = .98$	$CLES = 0.37$	
80		$U = 847.0$	$p = .6246$	$CLES = 0.48$		
Density		2.5%	$U = 936.0$	$p = .3161$	$CLES = 0.53$	
6.25%		$U = 653.0$	$p = .98$	$CLES = 0.37$		
10%		$U = 1082.0$	$p = .0372$	$CLES = 0.61$		
Task 3 - Within-Edge Comparison						
Accuracy	<p>Participants were significantly more accurate ($U = 6461.5, p = .0005, CLES = 0.59$) with AM ($M = 92.0, CI = [88.44, 95.14]$) compared to BF ($M = 79.9, CI = [74.27, 85.26]$). The accuracy is further affected by size and density. The results reveal that analysis results in networks with 50 and 80 nodes, or highest density (10%) were more accurate with AM than with BF.</p>	Size	20	$U = 948.0$	$p = .8663$	$CLES = 0.46$
		50	$U = 683.0$	$p = .0066$	$CLES = 0.61$	
80		$U = 490.5$	$p = .0001$	$CLES = 0.72$		
Density		2.5%	$U = 917.5$	$p = .6903$	$CLES = 0.48$	
6.25%		$U = 750.5$	$p = .0759$	$CLES = 0.57$		
10%		$U = 486.5$	$p < .0001$	$CLES = 0.72$		
Time	<p>There is a significant effect ($U = 11843.0, p < 0.0001, CLES = 0.75$) of completion time on visualization technique with participants working with AM ($M = 0.97, CI = [0.85, 1.08]$) being faster than with BF ($M = 1.95, CI = [1.68, 2.23]$). As the results indicate, this is also true for all different sizes and densities.</p>	Size	20	$U = 1544.0$	$p < .0001$	$CLES = 0.88$
		50	$U = 1632.0$	$p < .0001$	$CLES = 0.93$	
80		$U = 1428.0$	$p < .0001$	$CLES = 0.81$		
Density		2.5%	$U = 1264.0$	$p = .0003$	$CLES = 0.72$	
6.25%		$U = 1390.0$	$p < .0001$	$CLES = 0.79$		
10%		$U = 1397.0$	$p < .0001$	$CLES = 0.79$		

Table 1: Accuracy and completion time for each task split by network size and density. Significant results are highlighted in bold font.

6. Discussion

H1 - Scalability Hypothesis: We conjectured that *BF* is more scalable with respect to network size and density than *AM*. However, the analysis revealed that working with *BF* led to more error-prone results on complex networks than *AM*. Representing nodes as horizontal lines instead of points was new to all participants and added additional difficulty to accomplishing the analysis tasks. Also, participants had problems following individual edges. Increasing network size and density shrank the space between the node and edge lines, significantly impacting the visualization. Most participants commented on losing the focus of the target node. They proposed interactive highlighting or alternating colors for the edge lines as a solution. However, this approach requires careful consideration for the color encoding of additional attributes, especially when interactive highlighting is available. Although *BF* uses the visualization space more efficiently compared to *AM* in terms of the number of nodes, it does not preserve the legibility of complex graphs well. This result contrasts the objective of Longabaugh to address the scalability issue [Lon12]. Thus, we cannot confirm **H1**.

H2 - Neighborhood Hypothesis: Our results of task 1 support the recommendation of Nobre et al. on the neighborhood inspection [NMSL19]. When no edge attributes are involved, *AM* outperforms *BF*. We believe retrieving a neighbor in *BF* is impeded because the alignment of multiple lines to trace might change. As some participants mentioned, this is not true for the edges of the upper nodes. In this case, the representation of adjacent edges resembled a waterfall pattern, so a longer edge connects the next neighbor to the right. To reveal these kinds of patterns for all nodes, Longabaugh introduces shadow links [Lon12], which duplicate the number of edges, jeopardizing scalability. In *AM*, the aligned headers help keep track of inspected neighbors. Some participants working with *AM* tried to minimize the error by memorizing the number of cells between multiple adjacent edges and retrieving corresponding neighbors in the header simultaneously. This approach is not possible with *BF*. For small networks, *BF* provides an overview as good as *AM*. Overall, *AM* has a better performance. Therefore, we confirm our hypothesis **H2**.

H3 - Edge Sorting Hypothesis: Across all networks, we could not identify differences in performance between *AM* and *BF* for task 2. This is surprising since participants using *BF* perceived the ordering of edges as helpful for detecting an extreme value. Also, Nobre et al. identified the ordering of attributes as highly beneficial for nodes in *AM* [NWHL20]. Thus, we explain the discrepancy in our findings due to the erroneous differentiation of lines during neighbor retrieval. As the juxtaposed table covered one-third of the visualization space for *BF*, differentiating between node lines was even more challenging than in task 1. Further, we believe that the spatial separation of the edges and their attributes hampered their correct association (see the gap between the topology view and the corresponding attributes Figure 9). Again, an interactive highlighting might help to connect edges with their corresponding attributes better. The integrated approach in *AM* provided more accurate results on large networks and significantly decreased the time for identifying attributes in dense networks. Based on this, we cannot confirm **H3**.

H4 - Within-Edge Comparison Hypothesis: The embedded bar chart in the matrix cells is more amenable when multiple attributes

need to be inspected. *AM* showed better performance in task 3 than *BF*. These results are in line with recent research [NWHL20]. For *BF*, the juxtaposed view counteracted the benefit of sorting the edges on multiple attributes. On large and dense networks, participants were uncertain when following edges at the boundaries of the attribute value ranges. Also, comparing attributes over adjacent rows took longer and resulted in less accurate answers. In *AM*, the encoding as a bar chart made it possible to spot the difference between the target values easily due to the identical baseline. Thus, we can confirm our hypothesis **H4**.

7. Design Considerations

We propose the following design considerations based on our design space exploration, user experiment results, and discussion.

For assessing a node's neighborhood, analysts should consider *AM*. Although *BF* and *AM* share an overall high accuracy, the results of task 1 indicate that *AM* scales better for network size and density. **When analyzing few edge attributes, *AM* should be preferred.** Participants performed better when working with *AM*. This is true for all network sizes and densities.

For small network sizes and densities, both visualization techniques can safely be used. Participants performed equally well with both visualization techniques in all three tasks.

When there is sufficient space, designers should introduce shadow links for *BF*. Qualitative feedback has shown that participants prefer a staircase pattern for neighborhood inspection. The grouping of related edges facilitates reading neighborhoods.

With an increasing number of attributes, juxtaposed views might be preferable. Overplotting becomes more likely when many attributes are visualized using embedded or transformed views, irrespective of the chosen visualization technique (*AM* or *BF*).

Independent of the visualization technique, interactive highlighting should be implemented. In the qualitative assessment, participants commented on the usefulness of interactive highlighting to avoid losing track of following individual edges.

8. Conclusion & Limitations

In this paper, we explored the design space of BioFabric for multi-variate network analysis. We conducted an experiment with 28 participants to assess the performance of BioFabric for edge attribute-based analysis tasks compared to the well-established adjacency matrix. The results show that the adjacency matrix is more scalable and better suited for adjacency tasks with and without additional edge attributes. However, both visualizations achieved an overall high accuracy. From our qualitative feedback, the ordering of edges in BioFabric was helpful for the inspection of attributes. However, the benefits of juxtaposed tables proved in other studies did not prevail in our experiment. Therefore, our results did not support the potential ascribed to BioFabric [NMSL19].

Our design exploration offers various visual encoding and arrangement possibilities for BioFabric. Although previous experiments inspired our designs, alternative design considerations could produce different results. This is especially true for double encodings, which profit from redundant mappings [Kop79], or sorting the nodes and edges differently. In our experiment, the chosen tasks focused only

on edge attributes. In multivariate network analysis, node attributes and the combination of the two might also be interesting. Especially for the latter use case, BioFabric might profit from the possibility of sorting two juxtaposed tables (i.e., nodes and edge attributes). The underlying data in our user study contained up to four edge attributes. However, when increasing the number of attributes, we expect the juxtaposed tables used in BioFabric to profit from their aspect ratio. At the same time, integrated representations like the bar charts in adjacency matrices would drastically decline in visual quality. In other visualizations, interaction techniques are helpful for different tasks. Thus, we suggest the integration of interaction possibilities like visual highlighting are provided.

In the future, we want to investigate the performance of BioFabric in other topology- and attribute-based tasks when increasing the number of attributes and testing alternative design variations. The application of edge- and node-sorting algorithms, as already discussed for Massive Sequence Views [EHBvW14, vdEHBvW13] also seems promising. This study complements the multivariate network analysis research field with the first quantitative evaluation of BioFabric and a set of design considerations for practitioners.

Acknowledgement

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2117 – 422037984 and TRR 161 (Project A03) – Project-ID 251654672. Open Access funding enabled and organized by Projekt DEAL.

References

- [ABHR*13] ALPER B., BACH B., HENRY RICHE N., ISENBERG T., FEKETE J.-D.: Weighted graph comparison techniques for brain connectivity analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2013), CHI '13, Association for Computing Machinery, p. 483–492. doi:10.1145/2470654.2470724. 3, 4, 7
- [ABZD13] ABUTHAWABEH A., BECK F., ZECKER D., DIEHL S.: Finding structures in multi-type code couplings with node-link and matrix visualizations. In *2013 First IEEE Working Conference on Software Visualization (VISOFT)* (2013), pp. 1–10. doi:10.1109/VISOFT.2013.6650530. 3
- [ACJM03] AUBER D., CHIRICOTA Y., JOURDAN F., MELANCON G.: Multiscale visualization of small world networks. In *IEEE Symposium on Information Visualization 2003 (IEEE Cat. No.03TH8714)* (2003), pp. 75–81. doi:10.1109/INFVIS.2003.1249011. 2
- [Ban94] BANERJEE J. C.: *Gestalt Theory of Perception*. M.D. Publications Pvt. Ltd., 1994, pp. 107–109. 5
- [BDT21] BLUDAU M., DÖRK M., TOMINSKI C.: Unfolding edges for exploring multivariate edge attributes in graphs. In *21st Eurographics Conference on Visualization, EuroVis 2019 - Short Papers, Porto, Portugal, June 14-18, 2021* (2021), Byska J., Jänicke S., Schmidt J., (Eds.), Eurographics Association, pp. 17–19. doi:10.2312/evp.20211070. 4
- [BFN04] BALL R., FINK G. A., NORTH C.: Home-centric visualization of network traffic for security administration. In *Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security* (2004), pp. 55–64. 2
- [BH] BREWER C. A., HARROWER M.: ColorBrewer: Color Advice for Maps. <https://colorbrewer2.org/>. Accessed online: August, 2022. 6
- [BMGK08] BARSKY A., MUNZNER T., GARDY J., KINCAID R.: Cerebral: Visualizing multiple experimental conditions on a graph with biological context. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1253–1260. 2
- [BW11] BAE J., WATSON B.: Developing and evaluating quilts for the depiction of large layered graphs. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2268–2275. doi:10.1109/TVCG.2011.187. 2
- [CBDM17] CHANG C., BACH B., DWYER T., MARRIOTT K.: Evaluating Perceptually Complementary Views for Network Exploration Tasks. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2017), Association for Computing Machinery, p. 1397–1407. doi:10.1145/3025453.3026024. 5, 7
- [CHZ*07] CORNELISSEN B., HOLTEN D., ZAIDMAN A., MOONEN L., VAN WIJK J. J., VAN DEURSEN A.: Understanding execution traces using massive sequence and circular bundle views. In *15th International Conference on Program Comprehension* (2007), IEEE Computer Society, pp. 49–58. doi:10.1109/ICPC.2007.39. 2
- [Dan87] DANIEL W. W.: The Spearman Rank Correlation Coefficient. *Biostatistics: A Foundation for Analysis in the Health Sciences* (1987). 8
- [DBPB*22] DI BARTOLOMEO S., PISTER A., BUONO P., PLAISANT C., DUNNE C., FEKETE J.-D.: Six methods for transforming layered hypergraphs to apply layered graph layout algorithms. *Computer Graphics Forum* 41, 3 (2022), 259–270. doi:https://doi.org/10.1111/cgf.14538. 2, 5
- [dOC15] DE OLIVEIRA S. G., CHAGAS G. O.: A Systematic Review of Heuristics for Symmetric-Matrix Bandwidth Reduction: Methods not based on Metaheuristics. In *The XLVII Brazilian Symposium of Operational Research (SBPO)* (2015). 6
- [EHBvW14] ELZEN S., HOLTEN D., BLAAS J., VAN WIJK J. J.: Dynamic network visualization with extended massive sequence views. *IEEE Transactions on Visualization and Computer Graphics* 20, 08 (aug 2014), 1087–1099. doi:10.1109/TVCG.2013.263. 5, 11
- [GFC05] GHONIEM M., FEKETE J.-D., CASTAGLIOLA P.: On the readability of graphs using node-link and matrix-based representations: A controlled experiment and statistical analysis. *Information Visualization* 4, 2 (2005), 114–135. doi:10.1057/palgrave.ivs.9500092. 2
- [GHL15] GUO H., HUANG J., LAIDLAW D. H.: Representing uncertainty in graph edges: An evaluation of paired visual variables. *IEEE Transactions on Visualization and Computer Graphics* 21, 10 (2015), 1173–1186. doi:10.1109/TVCG.2015.2424872. 4
- [GL81] GEORGE A., LIU J. W.: *Computer Solution of Large Sparse Positive Definite Systems*. Prentice Hall, 1981. 6
- [HSS08] HAGBERG A. A., SCHULT D. A., SWART P. J.: Exploring Network Structure, Dynamics, and Function using NetworkX. In *Proceedings of the 7th Python in Science Conference* (Pasadena, CA USA, 2008), Varoquaux G., Vaught T., Millman J., (Eds.), pp. 11 – 15. 7
- [JKS06] JUNKER B. H., KLUKAS C., SCHREIBER F.: Vanted: A system for advanced data analysis and visualization in the context of biological networks. *BMC Bioinformatics* 7, 1 (2006), 109. doi:10.1186/1471-2105-7-109. 3
- [Kop79] KOPALA C. J.: The use of color-coded symbols in a highly dense situation display. *Proceedings of the Human Factors Society Annual Meeting* 23, 1 (1979), 397–401. doi:10.1177/1071181379023001100. 10
- [Lev60] LEVENE H.: Robust Tests for Equality of Variances. *Contributions to Probability and Statistics. Essays in Honor of Harold Hotelling* (1960), 278–292. 8
- [Lon12] LONGGABAUGH W. J.: Combing the hairball with biofabric: a new approach for visualization of large networks. *BMC Bioinformatics* 13, 1 (2012), 275. doi:10.1186/1471-2105-13-275. 2, 3, 5, 7, 10

- [LPP*06] LEE B., PLAISANT C., PARR C. S., FEKETE J.-D., HENRY N.: Task taxonomy for graph visualization. In *Proceedings of the 2006 AVI Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization* (New York, NY, USA, 2006), BELIV '06, Association for Computing Machinery, p. 1–5. doi:10.1145/1168149.1168168. 6
- [LTdSPR17] LINHARES C. D. G., TRAVENÇOLO B. A. N., DE SOUZA PAIVA J. G., ROCHA L. E. C.: Dynetvis: a system for visualization of dynamic networks. In *Proceedings of the Symposium on Applied Computing* (2017), Seffah A., Penzenstadler B., Alves C., Peng X., (Eds.), ACM, pp. 187–194. doi:10.1145/3019612.3019686. 5
- [MW47] MANN H. B., WHITNEY D. R.: On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics* 18, 1 (1947), 50–60. doi:10.1214/aoms/1177730491. 8
- [MW92] MCGRAW K. O., WONG S. P.: A Common Language Effect Size Statistic. *Psychological bulletin* 111, 2 (1992), 361–365. doi:10.1037/0033-2909.111.2.361. 8
- [NGCL19] NOBRE C., GEHLENBORG N., COON H., LEX A.: Lineage: Visualizing multivariate clinical data in genealogy graphs. *IEEE Trans. Vis. Comput. Graph.* 25, 3 (2019), 1543–1558. doi:10.1109/TVCG.2018.2811488. 2
- [NMSL19] NOBRE C., MEYER M., STREIT M., LEX A.: The state of the art in visualizing multivariate networks. *Computer Graphics Forum* 38, 3 (2019), 807–832. doi:https://doi.org/10.1111/cgf.13728. 2, 3, 6, 7, 8, 10
- [NSL19] NOBRE C., STREIT M., LEX A.: Juniper: A tree+table approach to multivariate graph visualization. *IEEE Trans. Vis. Comput. Graph.* 25, 1 (2019), 544–554. doi:10.1109/TVCG.2018.2865149. 2
- [NWHL20] NOBRE C., WOOTTON D., HARRISON L., LEX A.: Evaluating multivariate network visualization techniques using a validated design and crowdsourcing approach. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2020), CHI '20, Association for Computing Machinery, p. 1–12. doi:10.1145/3313831.3376381. 3, 5, 6, 7, 8, 10
- [PLS*12] PARTL C., LEX A., STREIT M., KALKOFEN D., KASHOFER K., SCHMALSTIEG D.: enrout: Dynamic path extraction from biological pathway maps for in-depth experimental data analysis. In *2012 IEEE Symposium on Biological Data Visualization (BioVis)* (2012), pp. 107–114. doi:10.1109/BioVis.2012.6378600. 2
- [PPS14] PRETORIUS J., PURCHASE H. C., STASKO J. T.: *Tasks for Multivariate Network Analysis*. Springer International Publishing, Cham, 2014, pp. 77–95. doi:10.1007/978-3-319-06793-3_5. 6
- [SDW09] SLINGSBY A., DYKES J., WOOD J.: Configuring hierarchical layouts to address research questions. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 977–984. doi:10.1109/TVCG.2009.128. 2
- [SS19] SMART S., SZAFIR D. A.: Measuring the separability of shape, size, and color in scatterplots. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2019), CHI '19, Association for Computing Machinery, p. 1–14. URL: https://doi.org/10.1145/3290605.3300899, doi:10.1145/3290605.3300899. 6
- [SSSE16] SCHÖFFEL S., SCHWANK J., STÄRZ J., EBERT A.: Multivariate networks: A novel edge visualization approach for graph-based visual analysis tasks. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2016), CHI EA '16, Association for Computing Machinery, p. 2292–2298. doi:10.1145/2851581.2892451. 2, 3
- [SW65] SHAPIRO S. S., WILK M. B.: An analysis of variance test for normality (complete samples). *Biometrika* 52, 3/4 (1965), 591–611. 8
- [SZ00] STASKO J., ZHANG E.: Focus+context display and navigation techniques for enhancing radial, space-filling hierarchy visualizations. In *IEEE Symposium on Information Visualization 2000. INFOVIS 2000. Proceedings* (2000), pp. 57–65. doi:10.1109/INFVIS.2000.885091. 2
- [VBP*21] VALDIVIA P. R., BUONO P., PLAISANT C., DUFOURNAUD N., FEKETE J.-D.: Analyzing Dynamic Hypergraphs with Parallel Aggregated Ordered Hypergraph Visualization. *IEEE Transactions on Visualization and Computer Graphics* 27, 1 (Jan. 2021), 1–13. doi:10.1109/TVCG.2019.2933196. 2, 3, 5, 6
- [vdEBvW13] VAN DEN ELZEN S., HOLTEN D., BLAAS J., VAN WIJK J. J.: Reordering massive sequence views: Enabling temporal and structural analysis of dynamic networks. In *IEEE Pacific Visualization Symposium* (2013), Carpendale S., Chen W., Hong S., (Eds.), IEEE Computer Society, pp. 33–40. doi:10.1109/PACIFICVIS.2013.6596125. 5, 11
- [vdEvW14] VAN DEN ELZEN S., VAN WIJK J. J.: Multivariate network exploration and presentation: From detail to overview via selections and aggregations. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2310–2319. doi:10.1109/TVCG.2014.2346441. 4
- [WEF*13] WYBROW M., ELMQVIST N., FEKETE J., VON LANDEBERGER T., VAN WIJK J. J., ZIMMER B.: Interaction in the visualization of multivariate networks. In *Multivariate Network Visualization - Dagstuhl Seminar #13201, Dagstuhl Castle, Germany, May 12-17, 2013, Revised Discussions* (2013), Kerren A., Purchase H. C., Ward M. O., (Eds.), vol. 8380 of *Lecture Notes in Computer Science*, Springer, pp. 97–125. doi:10.1007/978-3-319-06793-3_6. 3
- [YAD*18] YOGHOORDJIAN V., ARCHAMBAULT D., DIEHL S., DWYER T., KLEIN K., PURCHASE H. C., WU H.-Y.: Exploring the limits of complexity: A survey of empirical studies on graph visualisation. *Visual Informatics* 2, 4 (2018), 264–282. doi:https://doi.org/10.1016/j.visinf.2018.12.006. 2, 7