The Effects of Representation and Juxtaposition on Graphical Perception of Matrix Visualization

Xiaotong Liu
The Ohio State University
liu.1952@osu.edu

Han-Wei Shen
The Ohio State University
shen.94@osu.edu

ABSTRACT

Analyzing multiple networks at once is a common yet dif-
ficult task in many domains. Using adjacency matrices for
this purpose, however, can be effective because of its superior
ability to accommodate dense networks in a small area. We
evaluate various representations and juxtaposition designs for
visualizing adjacency matrices through a series of controlled
experiments. We investigate the effects of using square matri-
ces and triangular matrices on the speed and accuracy of per-
forming graphical-perception tasks. Based on human sym-
metric perception, we propose two alternative juxtaposition
designs to the conventional side-by-side juxtaposition, and
study how users perform visual search and comparison tasks
regarding different juxtaposition types. Our results show that
the matrix representations have similar performance, and the
matrix juxtaposition types perform differently. With the de-
sign guidelines derived from our studies, we present a com-
pact visualization termed TileMatrix for juxtaposing a large
number of matrices, and demonstrate its effectiveness in ana-
lyzing multi-faceted, time-varying networks using real-world
data.

Author Keywords
Networks; Information Visualization; Adjacency Matrices

ACM Classification Keywords
H.5.2 [Information Interfaces and Presentation]: User Inter-
faces – Graphical user interfaces

INTRODUCTION

A network is an abstract data type that represents entities as
nodes and their relationships as edges. Examples include
social networks, computer networks, biological networks,
and organizational networks. Network visualization has be-
come an important research topic aiming to gain an effec-
tive overview of complex relational data [4, 8, 14, 15, 16].
Introduced by Jacques Bertin [6], adjacency matrices offer
an interesting alternative to conventional node-link diagrams,

which can suffer from visual cluttering due to node overlap-
ping and edge crossing. An adjacency matrix shows how
nodes are connected together through the intersection of the
corresponding row and column. When the connections are
undirected, the adjacency matrix is symmetric with respect
to the main diagonal. Consequently, the same information is
shown in both the upper and lower triangular matrices. Ad-
jacency matrices have been shown more readable than node-
link diagrams for many graphical-perception tasks, particu-
larly when networks are dense [2, 10, 19]. Since many real-
world networks are naturally dynamic and associated with
multiple attributes, analyzing multiple networks at once is a
common yet difficult task. Therefore, visualizations that aid
users to compare and contrast multiple networks are of great
importance.

Juxtaposition is an effective visual design that encourages
side-by-side visual comparison of multiple facets of a com-
plex data set, without overplotting or occlusion that may oc-
cur in *superimposition*, which overlays many objects in a sin-
gle visualization. Although conventional side-by-side jux-
taposition, or small multiples [6, 29], has been applied to
adjacency matrices for comparative analysis [4], it predom-
inantly relies on the use of the viewer's memory and attention
shifts to make connections between repeated objects [12].
When too much of the comparative burden is added onto the
viewer's mental effort, he or she can fail to detect changes
even if the information is well represented in the visualiza-
tion [25]. Designing appropriate juxtaposed visualization for
adjacency matrices, which allows visual processing to con-
nect patterns across multiple matrices with less mental effort,
is still an open problem. Furthermore, an effective represen-
tation of multiple adjacency matrices is an instance of fun-
damental visualization research: making more effective use
of display space to increase the amount of data with which
users can effectively work [13, 29]. However, it remains un-
clear how different matrix representations and juxtaposition
designs affect the ability of users to quickly and reliably un-
derstand the underlying information.

In this paper, we evaluate the representation and juxtaposi-
tion designs for visualizing adjacency matrices through a se-
ries of controlled experiments. We investigate the effects of
matrix representation on the speed and accuracy of perform-
ing graphical-perception tasks. Based on *human symmetric
perception*, an automatic visual process that forms an integral
part of perceptual organization [30], we propose two alter-
native juxtaposition designs to the conventional side-by-side
juxtaposition, and study how users perform visual search and

comparison tasks regarding these juxtaposition types. Our results show that triangular matrices are as effective as square matrices, and the juxtaposition types (*side-by-side*, *back-to-back*, and *complementary*) perform differently. With the design guidelines derived from our studies, we present a compact visualization termed *TileMatrix* for juxtaposing a large number of matrices, and demonstrate its effectiveness in analyzing multi-faceted, time-varying networks using real-world data.

RELATED WORK

Previous research related to this work can be classified as follows: matrix visualization, graphical perception, and comparative visualization.

Matrix Visualization

The node-link diagram and adjacency matrix are the most popular for network visualization. Node-link diagrams are generally suitable for sparse networks while adjacency matrices are more effective for dense networks, as shown by Ghoniem et al. [10], Keller et al. [19] and Alper et al. [2]. Many research efforts have recently been made to enhance the usability and readability of matrix-based network visualization. MatrixExplorer [15] couples a node-link diagram and an adjacency matrix representation of the same network by juxtaposed views. MatLink [16] augments matrix representation with links connecting nodes in lines and columns to reconcile the difficulty of path-related tasks in matrix visualization [10]. NodeTrix [14] visualizes community structures (dense sub-networks) as adjacency matrices, which are then reconnected with links that represent the sparse parts of the network. Dinkla et al. [8] exploited the structural characteristics of gene regulatory networks to design compressed adjacency matrices. Behrisch et al. [5] placed matrices of high similarity together in a projection space. However, such automated approach cannot reveal multiple facets of the data, and the order in time is lost after projection, which can make it difficult to identify trends of networks over time. Cubix [4] stacks adjacency matrices over time to form a 3D space-time cube, and supports interactive transitions to projected 2D juxtaposed views. While most previous works employed square matrices for visualizing networks, we explored the effects of triangular matrices and their use in designing compact juxtaposed visualization for showing more readable networks within a limited display space.

Graphical Perception

Graphical perception is defined as the visual decoding of information encoded on graphical displays [7]. Cleveland and McGill [7] established a scientific foundation of graphical methods for data analysis and representation through human graphical perception. Since then, evaluating user performance on graphical perception has been an important research topic in the fields of human computer interaction and information visualization. Lohse [21, 22] developed a cognitive model for perception and understanding of graphs, and performed an empirical study to compare computer-simulated graphical perception based on his model with the actual user performance. Gillan and Lewis [11] also proposed a perceptual model and investigated human interaction with graphs,

by asking users to accomplish visual comparison tasks using various visualizations such as line graphs, scatter plots, and stacked bar graphs.

Graphical-perception tasks measure how users interpret the information encoded on graphical displays. The effects of visual variables for various graphical-perception tasks have been explored in the literature. Heer et al. [13] conducted two controlled experiments to measure the effects of chart size and layering on user performance while performing visual discrimination and estimation tasks. Javed et al. [18] evaluated user performance for visual comparison, slope, and discrimination tasks for multiple line graph visualizations. Fuchs et al. [9] investigated user performance of different temporal glyph designs in a small multiple setting. Most recently, Albers et al. [1] measured user performance of several visualization design choices for matching visual encodings to various aggregate comparison tasks. Our work attempts to empirically evaluate the effects of visual representation and juxtaposition on graphical perception of adjacency matrices.

### Comparative Visualization

Visual analysis often involves visual comparison of multiple objects. Gleicher et al. [12] described juxtaposition, superimposition, and explicit encoding as three generic design approaches to visual comparison, and suggested that multiple designs can be combined to created hybrid comparative visualization. Meanwhile, Javed and Elmqvist [17] divided the design space of composite visualization into juxtaposition, superimposition, overloading and nesting. Jacques Bertin [6] and Edward Tufte [29] introduced small multiples, the most popular juxtaposition design in both the scientific literature and the mass media. Such side-by-side design has been shown more effective than animation for dynamic networks [3] and trends [24]. Tominski et al. [28] proposed shine-through and folding interactions for visual comparison besides side-by-side comparison. Far more examples are given by Gleicher et al. [12] in a survey of over 110 references on visual comparison for information visualization. Our goal is to obtain a more general understanding about the design of alternatives juxtaposition for adjacency matrix visualization, not just the conventional side-by-side juxtaposition.

STUDY 1: EVALUATING MATRIX REPRESENTATION

As stated in the beginning, a symmetric adjacency matrix can be represented in either a square matrix (Figure 1 (left)) or a triangular matrix (Figure 1 (right)). The goal of the first experiment is to determine the impact of the matrix representation on the user's graphical perception: how does the choice of square or triangular matrices affect user performance?

Tasks

To keep the experiment manageable in time and effort for the participants, we did not include complex tasks related to finding paths and common neighbors since they can be difficult for the given matrix visualization [10]. Rather, such tasks can be accomplished more easily with node-link diagrams [10] or with additional visual cues [16]. Since we were interested in studying the effects of two different representations, tasks that are less relevant to the representation, such as counting
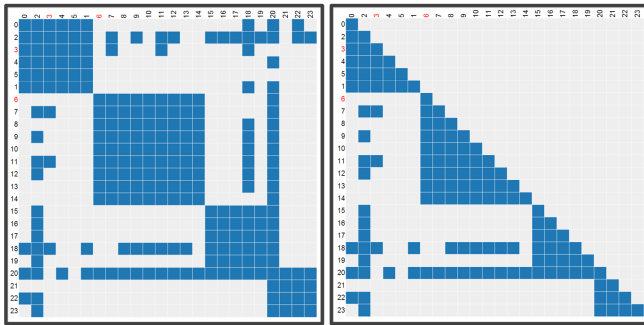
Figure 1. Adjacency matrix representations. Left: a square matrix representation. Right: a triangular matrix representation.

the number of nodes or finding a node given its label [10], were also excluded in our experiment. Finally, three generic tasks were selected in an attempt to capture both overview and detail use cases of matrix visualization:

(T1) *How many communities can you identify in the matrix?*

(T2) *Which node has the maximum number of neighbors?*

(T3) *Are the two specified nodes connected in the matrix?*

T1 is an overview graphical-perception task that represents a common use case of matrix visualization for community detection in networks [14]. T2 is an exploratory visual search task that requires browsing through the visualization to find information. T3 is a confirmatory visual search task that asks the users to make judgments once they locate the specified targets. T1 and T2 are related to the interpretation of blocks (communities) and lines (a node's neighbors), which are categorized as structural features by Mueller et al. [23]; T2 and T3 were also used in previous studies [10, 19].

Hypotheses
(H1) *The square matrix will perform as well as the triangular matrix for overview perception (T1).* For viewing the shapes of communities in a matrix, we believe that squares and triangles should be equally easy for perception.

(H2) *The square matrix will outperform the triangular matrix for exploratory search (T2).* For square matrices, users scan through rows or columns to view the neighbors of one node; for triangular matrices, users need to follow an L-shaped path to view the corresponding neighbors. We expect that following a straight line in rows or columns will have better user performance than following an L-shaped path.

(H3) *The square matrix will outperform the triangular matrix for confirmatory search (T3).* Square matrices encode one connection between two nodes twice in the visualization, while triangular matrices represent the connection exactly once. We predict that the repetition of such information will have a direct effect on user performance.

Experiment Design
The study was conducted as a within-subjects experiment with 2 experimental conditions (matrix representations) and

10 repetitions for each condition. For each repetition, the participant was presented with only one condition. We counterbalanced the selection of condition in the 10 repetitions so that each participant performed the same number of repetitions for both conditions while the choice of condition is random.

In the spirit of classic graphical perception experiments [9, 13, 18], we evaluated the different visual representations alone, disabling selecting, highlighting, zooming, and other interactive operations. Ghoniem et al. [10] have shown that in most cases, matrix visualization was insensitive to size and density variation and no interaction between size and density was found for the matrix representation. Hence, while the size and density of the matrices vary across repetitions for creating diverse experimental datasets, the effects due to size and density were not explicitly studied in our experiment.

We follow the convention of perception studies [13, 18, 20] in using synthetic data to allow control over the characteristics of experimental datasets. We first generated a random number of communities of varying sizes, then added a random number of edges for randomly selected nodes. In order to eliminate any ambiguity with respect to T2, which is to find the most connected node, we added an extra 20% of edges to the most connected node in every matrix. We labeled the nodes numerically according to the order of their creation, and ordered the nodes in such a way that those from the same community are placed together to preserve the community structures in the matrix. In this way, we obtained repetitions of varying number of nodes (from 20 to 50) and edges (density from 0.2 to 0.5), as well as varying number of communities (from 3 to 6). Specifically, the labels of the two specified nodes for T3 are highlighted in a distinct color. Figure 1 shows one example of the datasets we used for this experiment.

20 subjects (14 males, 6 females) were recruited for this experiment. The subjects were graduate students, aged 24 to 32 years. Half of the subjects have backgrounds in computer science, and most are using a computer more than 20 hours per week. All subjects have normal or corrected-to-normal vision. 40% of the subjects said they were familiar with matrix visualization.

We deployed the experiment on the web using HTML5 and JavaScript. Each subject used their own machine and browser, and the visualization images were scaled to fit the corresponding screen space before the test. Due to our within-subjects design, each user performed tasks regarding different representation types on the same screen he or she had. Hence, the influence of screen size has been alleviated when studying the effects of representation types. Prior to the experiment, subjects viewed a tutorial that gave a basic explanation of the two matrix representations, and they performed some training tasks to get familiar with the user interface of the experimental system. We asked the subjects to answer as quickly as possible while trying to make answers accurate. After the subjects finished tasks for all repetitions, they were asked to participate in a semi-structured interview.
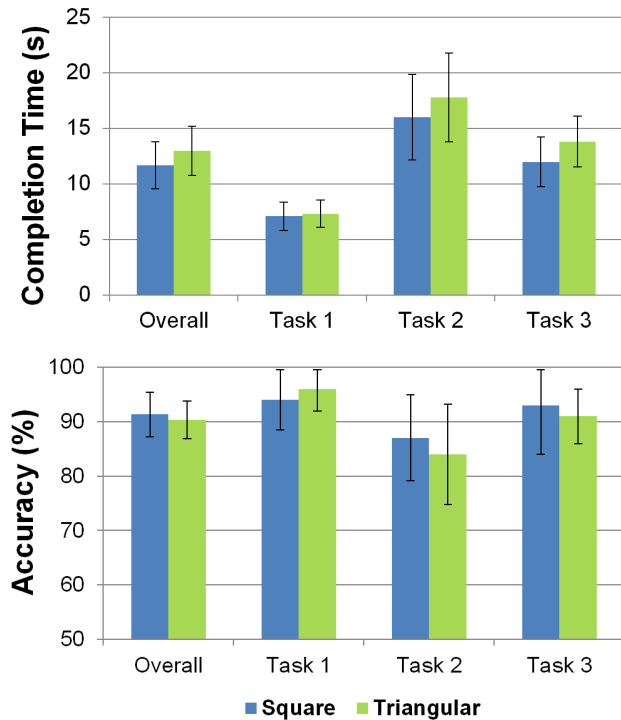
Figure 2. Mean completion time and accuracy for Study 1.

Table 1. RM-ANOVA analysis of results for Study 1.

	Factor	$F_{1,19}$	p
Overall	Time	0.82	0.37
	Accuracy	0.16	0.69
T1	Time	0.08	0.79
	Accuracy	0.40	0.53
T2	Time	0.48	0.49
	Accuracy	0.29	0.60
T3	Time	1.56	0.22
	Accuracy	0.18	0.68

Results and Discussion

We measured the time participants needed to complete each task and the correctness of their reported answers. User performance measures were evaluated using Repeated Measures Analysis of Variance (RM-ANOVA) to test for significant effects. For each participant's performance, we used the average of the repetitions for the following analysis. We report completion time and accuracy in Figure 2 and our RM-ANOVA analysis in Table 1.

Our first hypothesis was that the square matrix would perform as well as the triangular matrix for overview perception. This hypothesis was confirmed. We found that the accuracy and completion time were almost the same across the two representations, which indicates that users were able to interpret communities in both representations.

One surprising outcome of our study is that the matrix representation type has no significant effect on the completion time of exploratory search tasks like finding the most connected node in a matrix. In other words, it seems that participants did not become significantly slower when following the L-shaped path to view one node's neighbors. Participants commented that tracking neighbors for one node was a little bit hard at start, but their speed got faster after the training trials.

We also hypothesized that the square matrix would outperform the triangular matrix for confirmatory search. Contrary to our hypothesis, our quantitative results did not show significant differences between the matrix representations. Rather, these results indicate that the participants performed equally well in the confirmatory search, regardless of the matrix representations.

Overall, results from our first experiment revealed that users were able to interpret triangular matrices correctly, and their task performance with the triangular matrix was comparable to that with the square matrix.

STUDY 2: EVALUATING MATRIX JUXTAPOSITION

The goal of the second experiment is to investigate the effects of juxtaposition types in the presence of multiple adjacency matrices. More specifically, we are interested to see how users perform visual search and comparison tasks under different juxtapositions. Is there a benefit to introducing more types of juxtapositions other than conventional side-by-side juxtaposition?

In our previous experiment we found that square and triangular representations of adjacency matrix had comparable accuracy and completion time, so we removed square matrices from consideration in this experiment and focused on comparing triangular matrices of different juxtaposition types.

Matrix Juxtaposition

We divide the design space of comparative juxtaposition into three general types, based on how the relationships between the related parts of different adjacency matrices are encoded:

Side-by-side juxtaposition (SID), also known as small multiples [29], repeats the same representation multiple times without any modification to the design (Figure 3 (left)), and has been applied to various domains such as system management, quality control, and medical record analysis [29]. This *translational juxtaposition* is the conventional design of comparative juxtaposition.

Back-to-back juxtaposition (BAC), which reverts the order of the rows or columns in one of two matrices to form a symmetric composition (Figure 3 (middle)). This design is motivated by human symmetric perception, which is an automatic visual process that forms an integral part of perceptual organization [30]. Such *symmetric juxtaposition* has been shown effective in identifying similarity and contrasts in computer-aided diagnosis. For example, radiologists are using the differences between the symmetric juxtaposition of left and right breasts in mammograms (photographs of breasts made by X-rays) to help detect certain malignant breast cancers [26].
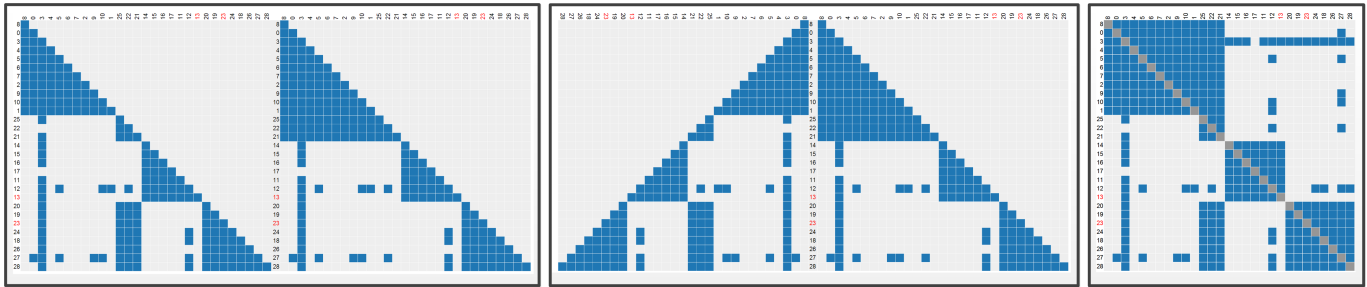
Figure 3. Adjacency matrix juxtapositions. Left: side-by-side juxtaposition (SID). Middle: back-to-back juxtaposition (BAC). Right: complementary juxtaposition (COM).

Complementary juxtaposition (COM), which takes two triangular matrices together to form a compact square matrix. Consequently, the rows and columns in the two matrices are complementary, and the positions of cells in the two matrices are symmetric with respect to the main diagonal (Figure 3 (right)). In other words, two triangular matrices form an asymmetric square matrix. This design is a variant of back-to-back juxtaposition, and only applies to triangular matrices.

We think these three juxtaposition types are *fundamental*: they define the juxtaposed relationship of two matrices (translational juxtaposition for **SID** or symmetric juxtaposition for **BAC** & **COM**), representing the basic cases for comparative visualization. More importantly, they provide the building blocks that can assemble complex juxtaposition of many more matrices. In other words, the three juxtaposition types can be combined to create hybrid juxtaposed visualization.

Tasks

Typical visual comparison tasks often require users to scan and compare related elements simultaneously. Three tasks, modified from the tasks selected in our previous experiment, were included in this experiment:

(T4) *Does the largest community have the same number of nodes in the following matrices?*

(T5) *Does the most connected node have the same number of neighbors in the following matrices?*

(T6) *How many times (0, 1, or 2) are the two specified nodes connected in the following matrices?*

In each task, users were asked to identify multiple targets within two matrices presented in one of the three juxtaposition types, and determine the similarity or difference of the targets. As in our prior experiment, these tasks were selected to represent overview and detail use cases of matrix visualizations, with a balance between task complexity and suitability.

Hypotheses

(H4) *For tasks involving comparison of structures and patterns in matrices (T4, T5), BAC & COM (symmetric juxtaposition) will outperform SID (translational juxtaposition) in task completion time.* The strength of symmetric perception is that it allows easier comparison of shapes across objects. We believe that the mental alignment due to symmetric perception will result in better completion time by connecting patterns and reveal their differences across matrices.

(H5) *When searching and comparing a specific target in matrices (T6), SID, BAC and COM will have comparable performance.* Comparing a specific target requires accurately locating the same item across matrices. We predict that symmetric juxtaposition requires more mental effort to locate a specific target across matrices. However, we think that the easier mental alignment due to symmetric perception compensates this additional effort, resulting in a performance comparable to that obtained using translational juxtaposition.

Experiment Design

The experiment used a 3 (juxtaposition) $\times$ 3 (task) within-subjects design with 9 repetitions. For each repetition, the participant was presented with only one condition. We counter-balanced the repetitions so that each participants performed the same number of conditions for three conditions while the selection of condition in each repetition was random.

We followed the design choices we made in our previous experiment. We first generated matrices that have varying numbers of nodes (from 20 to 50) and edges (density from 0.2 to 0.5), and varying numbers of communities (from 3 to 6). Then for each generated matrix, we randomly assigned several nodes to a different community and added edges for randomly selected nodes to get a new matrix. The two matrices were then grouped using one of the three juxtaposition types. Figure 3 shows one example of the datasets we obtained for this experiment.

We recruited 28 subjects (17 males, 11 females, aged 24 to 39 years), who are graduate students or professionals from the fields of computer science, electric engineering, chemistry, statistics, geographic information science, to name a few. Half of the subjects had previously participated in the first experiment. 36% of the subjects said they had experience with multiple matrix visualization. The experiment follows a similar procedure to the prior experiment. After the subjects finished all tasks, they were asked to rate their satisfaction with **SID**, **BAC** and **COM** in each task on a questionnaire (e.g., *Which view do you think is the best to identify the size difference of the largest community? (A: side-by-side; B: back-to-back; C: complementary; D: no preference)*). Finally, they were asked to participate in a semi-structured interview.
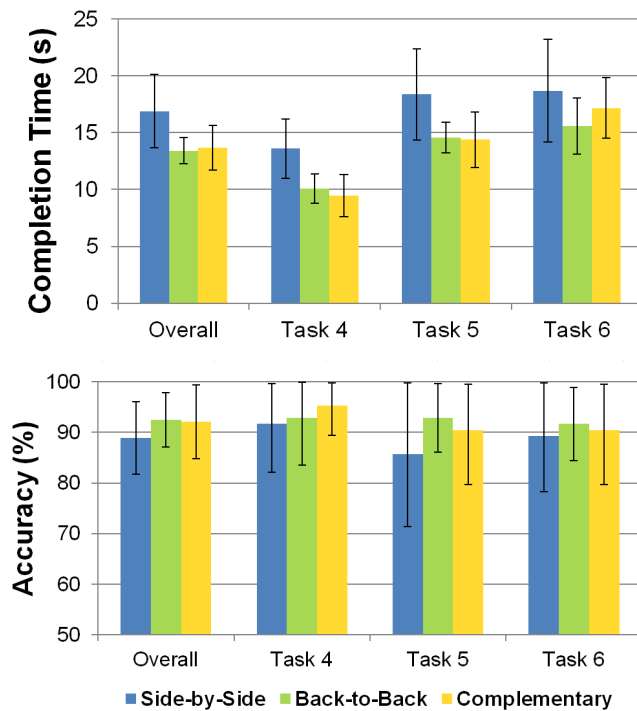
Figure 4. Mean completion time and accuracy for Study 2.

Table 2. RM-ANOVA analysis for Study 2 (s:SID, b:BAC, c:COM).

	Factor	$F_{2,54}$	$p_{s,b,c}$	$p_{s,b}$	$p_{s,c}$	$p_{b,c}$
Overall	Time	4.81	**0.01**	**0.01**	**0.03**	0.77
	Accuracy	0.59	0.56	0.30	0.42	0.91
T4	Time	8.47	**0.001**	**0.003**	**0.001**	0.49
	Accuracy	0.32	0.73	0.82	0.41	0.57
T5	Time	4.30	**0.02**	**0.02**	**0.03**	0.86
	Accuracy	0.73	0.49	0.25	0.49	0.63
T6	Time	1.46	0.24	0.12	0.46	0.26
	Accuracy	0.10	0.91	0.64	0.84	0.81

Table 3. Subjective user preference for tasks in Study 2.

	SID	BAC	COM	None
T4	0%	46%	50%	4%
T5	10%	36%	40%	14%
T6	21%	21%	18%	40%

Results and Discussion

We report completion time and accuracy in Figure 4. Statistical significance was tested for all groups and each pair respectively. Table 2 summarizes all results and highlights the significant ones. The qualitative feedback is shown afterwards in Table 3.

In H4 we hypothesized that **BAC** & **COM** would be faster than **SID** due to their symmetric alignment for detecting changes of structures and patterns in matrices. This hypothesis was confirmed. We found that **SID** took significantly longer time than **BAC** & **COM**. This result validates our design goals: facilitating identification of the repeated patterns and differences of connectivity between juxtaposed matrices. Subject ratings (Table 3) confirmed that most participants preferred either **BAC** or **COM** in accomplishing T4 and T5. Participants commented that **BAC** and **COM** are helpful in connecting the left and right matrices, while **SID** requires more mental effort in certain cases.

Our results confirmed H5: **BAC** & **COM** would not have a negative impact on performance when searching and comparing a specific target in matrices. We found no significant difference between the juxtaposition types in task completion time or accuracy. Participants do not seem to have a strong preference in choosing a particular juxtaposition for T6. Accuracy seems to stabilize at a high rate for each task in any juxtaposition. It indicates that the participants were equally careful in visual comparison, regardless of juxtaposition type.

DESIGN IMPLICATIONS

Based on our experimental results, we offer the following implications for designing adjacency matrix visualizations:

Triangular representation does not hamper graphical perception of adjacency matrices. One unexpected result was that the triangular matrix neither slowed down the task completion time nor hurt the accuracy. The triangular representation cuts the size of the matrix in half without any observed downside, as long as the viewer learns how to interpret the triangular matrix.

Symmetric juxtaposition rather than translational juxtaposition should be preferred for detecting changes of structures and patterns. The mental alignment due to symmetric perception is able to connect patterns and structures across matrices, which facilitates visual comparison with less mental effort.

Complementary juxtaposition is beneficial for optimizing utilization of display space. Since complementary juxtaposition doubles the data density (data marks per display area) compared to other juxtaposition types, we advocate its use when space constraints warrant.

TILEMATRIX: CREATING COMPACT VISUALIZATION BY TILING THE MATRICES

With the design implications, we devise a compact visualization — *TileMatrix*, coupling the side-by-side juxtaposition, back-to-back juxtaposition and complementary juxtaposition to display a large number of adjacency matrices. The TileMatrix representation is inspired by the physical act of laying tiles to cover floors, walls, ceilings and roofs. A *tile* is generally designed in an interlocking pattern so that multiple tiles fit together to have an aesthetic appearance. In TileMatrix, one tile is a triangular matrix. Every two adjacent matrices are placed in either back-to-back juxtaposition or complementary juxtaposition, resulting in a hybrid juxtaposition that also includes side-by-side juxtaposition . In the following sections,

we first discuss typical usage scenarios for displaying a large number of matrices, and then describe our TileMatrix system. We demonstrate its effectiveness in a case study using real-world data, and report initial user feedback.

Usage Scenarios

A weighted network associates a weight with every edge in the network, representing the strength of the connection between the entities [2]. In many network applications, the weights of edges can be quite complex and dynamic. On the one hand, the weight of the connection between the entities can change as multiple facets are associated with the connection. For instance, consider a collection of National Basketball Association (NBA) players as the entities, who are connected either strongly or weakly based on the similarity of their performance. Each player is associated with multiple performance facets: points, rebounds, assists, steals, blocks, and many others. Therefore, players can have different weights of connections on different facets. On the other hand, the weight of a connection between the entities can change over time. This is also seen in the NBA players example since their performance often vary in different years.

Visual analysis of such multi-faceted, time-varying weighted networks is challenging as the number of networks grows in both data and temporal domains. Although we can always reduce the number of facets using dimension reduction techniques or aggregate networks over time, we learn little about the heterogeneity nor the dynamics of the network. The reduced dimensions do not always have a semantic interpretation, and the temporal differences are hidden in temporal aggregation. Consequently, we understand neither how to interpret the relationships from multiple perspectives nor how certain temporal trends are formed.

TileMatrix offers one alternative solution to visualizing multifaceted, time-varying weighted networks without losing information — the matrices of networks are tiled simultaneously in two directions of the display space: matrices in different facets are tiled in columns while matrices over time are tiled in rows. In this way, the viewer can examine and compare networks of multiple facets at a particular time step (horizontally), as well as networks of a particular facets over time (vertically). The repeated patterns and differences in the relationship of entities can be identified, in both data and temporal domains.

Case Study

To demonstrate the effectiveness of TileMatrix, we used our system to explore the NBA statistics dataset [27], which consists of 16 performance facets for NBA players: GP (Games Played), MIN (Minutes Played), PT (Points Scored), AST (Assists), REB (Rebounds), STL (Steals), BLK (Blocked Shots), TO (Turnovers), ORE (Offensive Rebounds), DRE (Defensive Rebounds), FGA (Field Goals Attempted), FGM (Field Goals Made), FTA (Free Throws Attempted), FTM (Free Throws Made), TPA (Three Point Attempted), and TPM (Three Point Made). To view a relatively large number of matrices with TileMatrix, in this study we extracted players that have relatively longer NBA careers (at least 15 years
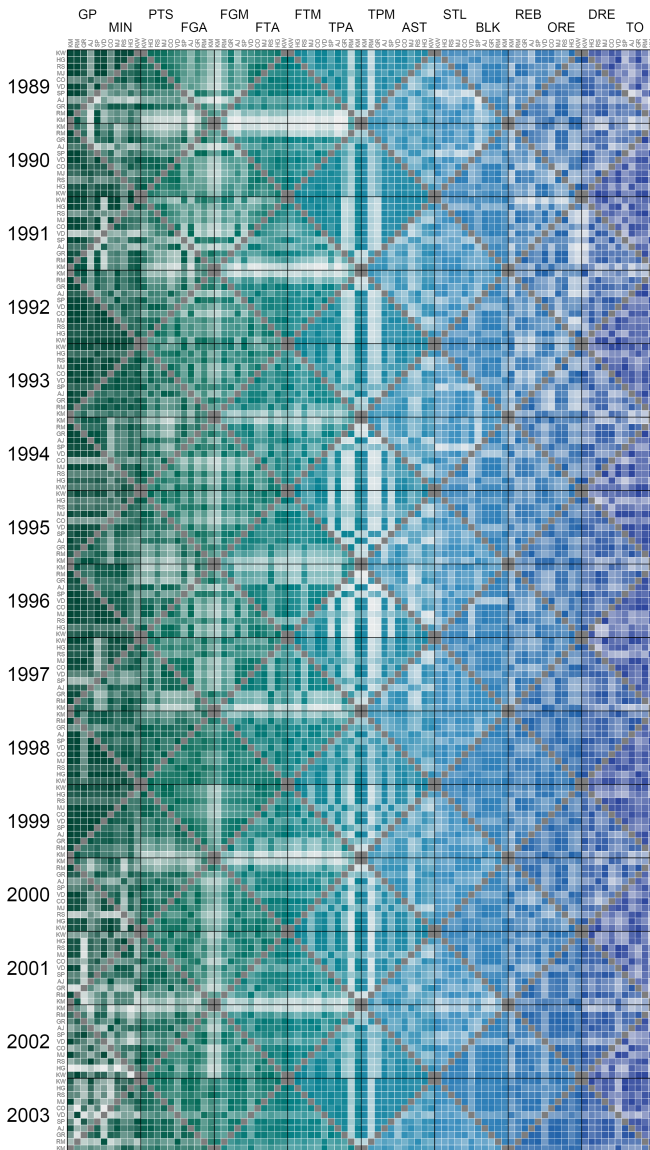
compared to an average career length of 5 years) from 1989 to 2003. This gave us 11 NBA players: Horace Grant, Reggie Miller, Vlade Divac, Avery Johnson, Glen Rice, Karl Malone, Scottie Pippen, Charles Oakley, Kevin Willis, Rod Strickland, and Mark Jackson. The similarity networks for all pairs of players were computed regarding each facet for each year. The corresponding TileMatrix contains $16*15*11*11/2 = 14520$ cells. The color of self-connections is set to grey. Matrices of adjacent facets are assigned similar colors while distant facets are colored differently for overall appearance.

By analyzing the semantic meaning of the performance facets, we roughly categorized them into three groups based on the criteria of the performance measurement: *activity* = {GP, MIN}, *score* = {PT, FGA, FGM, FTA, FTM, TPA, TPM}, *aid* = {AST, STL, BLK, REB, ORE, DRE, TO}. Therefore, we reordered the facets in TileMatrix to place related facets in the same category closer. To highlight the contrasts between highly and weakly connected nodes, we reordered the nodes by the weighted average degree.

Figure 5 shows the resulting TileMatrix visualization after reordering. Facets are labeled on the top of TileMatrix while years are on the left. The names of the players are also labeled accordingly. From the TileMatrix visualization, we had a few key findings on the multi-faceted, time-varying relationships of the NBA players that had relatively longer careers:

Temporal trends regarding a particular facet. To understand the time-varying similarity/dissimilarity of players' performance with respect to one facet, we can simply view matrices form a column of TileMatrix. For example, the left-most column of TileMatrix in Figure 5 shows the time-varying similarity networks based on the number of games the players had played, from which we can observe the following trend: (a) in 1989, 1991, and 1997: one or two players had played quite a different number of games compared to the rest of players, as the their connections to the others were relatively weak (e.g., CO (Charles Oakley) and AJ (Avery Johnson) in 1989, VD (Vlade Divac) in 1991, and SP (Scottie Pippen) in 1997); (b) in 1992, 1993, 1996 and 1999: the number of games played were almost the same among the players, since the matrix cells show very strong connections; (c) during 2000-2003: players started to have distinct degrees of activity, and the differences seemed to reach a peak at the end of this time window (e.g., in 2003, most players were strongly connected to only a few other players while weekly connected to the majority). This could be partly because some players had played for many years, and their performance became diverse compared with their early career.
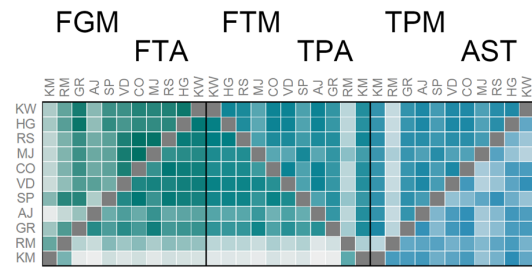
Repeated patterns across multiple facets at a specific time. As implied from our experiments, the symmetric juxtaposition has significant effects on connecting patterns and structures across matrices. To view the multi-faceted similarity/dissimilarity of players' performance at a specific time, we can focus on one row of matrices in TileMatrix. For instance, the first row of TileMatrix in Figure 5 shows the multi-faceted similarity networks of players in 1989. Approximately repeated visual patterns can be identified between the following groups of facets respectively: FGA & FGM, FTA & FTM,
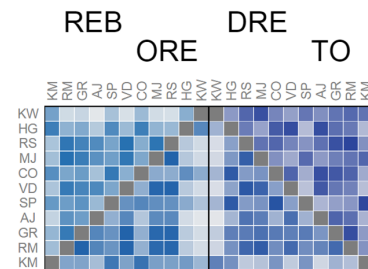
Figure 6. Two zoom-in views of the TileMatrix visualization in Figure 5. (a) Matrices of 6 selected facets (FGM, FTA, FTM, TPA, TPM, AST) in 1989. (b) Matrices of 4 selected facets (REB, ORE, DRE, TO) in 1991.
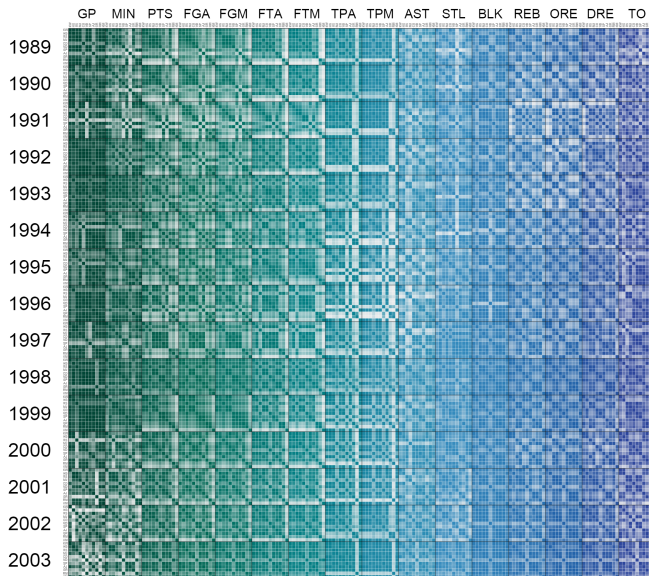
1991 as an example (shown in Figure 6 (b)). We observed that the degree of dissimilarity to other players are different between offensive rebounds and defensive rebounds for HG (Horace Grant), while it roughly stayed the same for KW (Kevin Willis). A comparison between REB and TO for the connections of KW (Kevin Willis) also highlights the contrasts between these two facets: his performance was more different from others on rebounds than on turnovers.

Temporal trends regarding multiple facets. Because of the two dimensional tiling layout in TileMatrix, we are able to examine the temporal trends of adjacent facets, by viewing matrices in adjacent columns as a whole. For example, the columns of TPA and TPM of TileMatrix in Figure 5 shows a trend for the dissimilarity of players' performance regarding these two facets: (a) starting from 1989, only one player RM (Reggie Miller) performed differently from the others; (b) the difference among players became more and more obvious over the years until it reached top in 1996, when three players performed very differently from the others regarding both three point shots and scores; (c) after that, the degree of dissimilarity among the players dropped; (d) finally in 2003, RM (Reggie Miller) once again became the only one that performed differently compared with the other players, just like how the trend started in 1989. Comparisons between other related facets also revealed interesting temporal trends.

Informal User Feedback

We observed six users (5 males, 1 female) viewing the TileMatrix in Figure 5, and a conventional side-by-side juxtaposed view of square matrices (denoted as SidMatrix) in Figure 7. SidMatrix displays the same number of matrices as TileMatrix with the same color encoding. We explained the designs of TileMatrix and SidMatrix to the users, and how the

Figure 5. The TileMatrix visualization of 16-faceted similarity networks of National Basketball Association (NBA) players from 1989 to 2003. Matrices of different facets are tiled in columns, while matrices of different years are tiled in rows. The color opacity encodes the strength of the connection.

and TPA & TPM. From a zoom-in view of the facets in Figure 6 (a), we learned that: if one player (e.g., RM (Reggie Miller)) behaved differently from the others in one facet (e.g., TPA), he was very likely to perform differently from the others in another related facet (e.g., TPM); if the performance of a group of players (e.g., KW (Kevin Willis), HG (Horace Grant), RS (Rod Strickland), MJ (Mark Jackson), CO (Charles Oakley), VD (Vlade Divac)) were similar to each other in one facet (e.g., FTM), such similarity relationship could also be found in another related facet (e.g., FTA).

Contrasts between multiple facets at a specific time. Another implication from our experiments is that the symmetric juxtaposition also helps reveal contrasts between adjacent facets. Take the similarity networks on ORE and DRE in

Figure 7. The side-by-side juxtaposed matrix visualization (SidMatrix) for the same data as TileMatrix in Figure 5.

information of the multi-faceted, time-varying similarity networks of NBA players is presented. Our task was very informal, simply asking them to comment on their understanding and visual comfort of the two visualizations, and how easily they could find interesting trends in each visualization. The two visualizations were uniformly scaled to the same width and put side by side for direct comparison.

As initial feedback, all six participants were able to understand TileMatrix as well as SidMatrix. One general feedback is that a cell in SidMatrix is much smaller and less legible than that in TileMatrix. This is expected since the two visualizations display the same number of matrices, square matrices in SidMatrix take double space compared with triangular matrices in TileMatrix, and thus each cell in SidMatirx is displayed in half size of that in TileMatrix. Participants commented that TileMatrix was easier to see details. They were more comfortable in seeing larger cells with more effective information in TileMatrix. Participants also explained that with TileMatrix, it was easier to compare adjacent matrices when they are in back-to-back or complementary juxtaposition, since the symmetric layout helped them to see the subtle differences, which were not as obvious in SidMatrix. However, some participants reported that TileMatrix required more training effort to follow the L-shaped path when viewing the connections of a particular player, while they were used to following straight lines in SidMatrix. This is consistent with the findings in our first controlled experiment.

Our observations partly explained the tradeoffs of introducing hybrid juxtapositions of triangular matrices in visualizing a large number of networks. While such hybrid design makes better use of the display space, and may be useful for identifying repeated patterns and differences across multiple matrices more comfortably and clearly, it requires more training effort than the conventional side-by-side juxtaposed visualization.

DISCUSSION AND FUTURE WORK

As stated in the beginning, the triangular matrix works when the connections in a network are undirected. As for directed networks, it is not possible to increase the data density by juxtaposing triangular matrices. Our experiments considered a small set of generic tasks for matrix visualization. We believe these findings generalize to a wider range of situations, but have not confirmed this empirically. Also we did not include interactive operations in our experiments but only focused on evaluating the effects of different representations and juxtapositions alone. Still, future work is needed to determine the additional effects when visualization is augmented by interaction.

Our TileMatrix design is effective in viewing the similarity and differences in matrices across facets and time, in particular for comparing matrices placed at nearby locations. On the other hand, using TileMatrix requires training to follow L-shaped paths when viewing one node's neighbors, and thus perhaps requiring more effort when tracing one node's neighbors across multiple matrices. Providing additional interactions such as highlighting one's neighbors across matrices when selecting a node could be helpful. Although the effect of colors used in TileMarix was not explicitly studied in our experiments, the design of TileMatrix can be easily extended, and the task performance with additional factors such as colors will be investigated in the future.

The scalability of TileMatrix depends on several factors: the number of attributes (M), the number of time steps (T), and the number of entities per attribute per time step (N). The total number of cells in TileMatrix is $M * T * N * N / 2$. Given a display screen, a tradeoff between N, M and T must be made when using TileMatrix, for example: (1) the network is small so many attributes and/or time steps are shown (as in our case study); (2) the network is medium sized so a medium number of attributes and time steps are shown; and (3) the network is large, thus a few selected attributes and time steps are shown. We note that if the display screen cannot hold all cells, panning and zooming could be provided to view sub-regions on demand. Although we demonstrate the effectiveness of our TileMatrix design with a case study and informal user feedback, the trade-off between design complexity and task performance needs to be better understood in the future.

CONCLUSION

In this paper, we conducted two controlled experiments to assess the performance of adjacency matrices in two representations — square matrices and triangular matrices, and three juxtaposition designs — side-by-side juxtaposition, back-to-back juxtaposition, and complementary juxtaposition. We quantitatively measured speed and accuracy based on generic tasks in each experiment. The results showed that triangular matrices were as effective as square matrices, and the three juxtaposition types performed differently. We showed that back-to-back juxtaposition and complementary juxtaposition are generally a good choice for detecting changes of structures and patterns across matrices due to the mental alignment

of symmetric perception. Based on the design guidelines derived from our studies, we propose a compact visualization termed TileMatrix for juxtaposing a large number of matrices, and show its benefits in analyzing multi-faceted, time-varying networks using real-world data.

ACKNOWLEDGMENTS

REFERENCES

1. Albers, D., Correll, M., and Gleicher, M. Task-driven evaluation of aggregation in time series visualization. In *CHI* (2014).

2. Alper, B., Bach, B., Henry Riche, N., Isenberg, T., and Fekete, J.-D. Weighted graph comparison techniques for brain connectivity analysis. In *CHI* (2013).

3. Archambault, D., Purchase, H. C., and Pinaud, B. Animation, small multiples, and the effect of mental map preservation in dynamic graphs. *TVCG* (2011).

4. Bach, B., Pietriga, E., and Fekete, J.-D. Visualizing dynamic networks with matrix cubes. In *CHI* (2014).

5. Behrisch, M., Davey, J., Fischer, F., Thonnard, O., Schreck, T., Keim, D., and Kohlhammer, J. Visual analysis of sets of heterogeneous matrices using projection-based distance functions and semantic zoom. *Computer Graphics Forum 33*, 3 (2014), 411–420.

6. Bertin, J. Semiology of graphics: diagrams, networks, maps, 1983.

7. Cleveland, W. S., and McGill, R. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American statistical association* (1984).

8. Dinkla, K., Westenberg, M. A., and van Wijk, J. J. Compressed adjacency matrices: untangling gene regulatory networks. *TVCG 18*, 12 (2012), 2457–2466.

9. Fuchs, J., Fischer, F., Mansmann, F., Bertini, E., and Isenberg, P. Evaluation of alternative glyph designs for time series data in a small multiple setting. In *CHI*, ACM (2013), 3237–3246.

10. Ghoniem, M., Fekete, J.-D., and Castagliola, P. On the readability of graphs using node-link and matrix-based representations: A controlled experiment and statistical analysis. *Information Visualization* (July 2005).

11. Gillan, D. J., and Lewis, R. A componential model of human interaction with graphs: 1. linear regression modeling. *Human Factors: The Journal of the Human Factors and Ergonomics Society 36*, 3 (1994), 419–440.

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

13. Heer, J., Kong, N., and Agrawala, M. Sizing the horizon: The effects of chart size and layering on the graphical perception of time series visualizations. In *CHI* (2009).

14. Henry, N., Fekete, J., and McGuffin, M. Nodetrix: a hybrid visualization of social networks. *TVCG 13*, 6 (Nov 2007), 1302–1309.

15. Henry, N., and Fekete, J.-D. Matrixexplorer: a dual-representation system to explore social networks. *TVCG 12*, 5 (2006), 677–684.

16. Henry, N., and Fekete, J.-D. Matlink: Enhanced matrix visualization for analyzing social networks. In *Human-computer Interaction* (2007).

17. Javed, W., and Elmqvist, N. Exploring the design space of composite visualization. In *Pacific Visualization Symposium*, IEEE (2012), 1–8.

18. Javed, W., McDonnel, B., and Elmqvist, N. Graphical perception of multiple time series. *TVCG 16*, 6 (2010).

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

20. Lam, H., Munzner, T., and Kincaid, R. Overview use in multiple visual information resolution interfaces. *TVCG 13*, 6 (2007), 1278–1285.

21. Lohse, G. L. A cognitive model for understanding graphical perception. *Human-Computer Interaction 8*, 4 (1993), 353–388.

22. Lohse, J. A cognitive model for the perception and understanding of graphs. In *CHI* (1991).

23. Mueller, C., Martin, B., and Lumsdaine, A. Interpreting large visual similarity matrices. In *International Asia-Pacific Symposium on Visualization* (2007), 149–152.

24. Robertson, G., Fernandez, R., Fisher, D., Lee, B., and Stasko, J. Effectiveness of animation in trend visualization. *TVCG 14*, 6 (2008), 1325–1332.

25. Simons, D. J., and Levin, D. T. Change blindness. *Trends in cognitive sciences 1*, 7 (1997), 261–267.

26. Tang, J., Rangayyan, R., Xu, J., El Naqa, I., and Yang, Y. Computer-aided detection and diagnosis of breast cancer with mammography: Recent advances. *IEEE Transactions on Information Technology in Biomedicine 13*, 2 (March 2009), 236–251.

27. TODAY, U. Online nba statistics. http://www.usatoday.com/sports/nba/statistics/.

28. Tominski, C., Forsell, C., and Johansson, J. Interaction support for visual comparison inspired by natural behavior. *TVCG 18*, 12 (2012), 2719–2728.

29. Tufte, E. *Envisioning Information*. Graphics Press, 1990.

30. Tyler, C. W. *Human symmetry perception and its computational analysis*. Psychology Press, 2003.