

The Pattern is in the Details: An Evaluation of Interaction Techniques for Locating, Searching, and Contextualizing Details in Multivariate Matrix Visualizations

Yalong Yang
Virginia Tech
Blacksburg, VA, USA
yalongyang@vt.edu

Carolina Nobre
Harvard University
Boston, MA, USA
cnobre@seas.harvard.edu

Wenyu Xia
Carnegie Mellon University
Pittsburgh, PA, USA
wenyux@andrew.cmu.edu

Robert Krueger
Harvard University
Boston, MA, USA
krueger@seas.harvard.edu

Fritz Lekschas
Harvard University
Boston, MA, USA
lekschas@seas.harvard.edu

Hanspeter Pfister
Harvard University
Boston, MA, USA
pfister@seas.harvard.edu




Figure 1: In a multivariate matrix visualization (MMV) each cell represents a *data point* that is associated to multiple time points, attributes, or other data points (through aggregation). To study a *region of interest*, an analyst typically needs to *locate* the region, *search* for details within the region, and identify *contextual* patterns using details. Multiple interaction approaches can be used to conduct this exploration but which one is most effective for different tasks: *focus+context*, *pan&zoom*, or *overview+details*?

ABSTRACT

Matrix visualizations are widely used to display large-scale network, tabular, set, or sequential data. They typically only encode a single value per cell, e.g., through color. However, this can greatly limit the visualizations' utility when exploring multivariate data, where each cell represents a data point with multiple values (referred to as *details*). Three well-established interaction approaches can be applicable in multivariate matrix visualizations (or MMV): *focus+context*, *pan&zoom*, and *overview+detail*. However, there is

little empirical knowledge of how these approaches compare in exploring MMV. We report on two studies comparing them for locating, searching, and contextualizing details in MMV. We first compared four focus+context techniques and found that the *fish-eye lens* overall outperformed the others. We then compared the fish-eye lens, to pan&zoom and overview+detail. We found that *pan&zoom* was faster in locating and searching details, and as good as overview+detail in contextualizing details.

CCS CONCEPTS

• Human-centered computing → Empirical studies in visualization; Empirical studies in HCI.

KEYWORDS

Multi-level navigation, multivariate, matrix, focus+context, overview+detail, pan&zoom



This work is licensed under a Creative Commons Attribution International 4.0 License.

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9157-3/22/04.
<https://doi.org/10.1145/3491102.3517673>

ACM Reference Format:

Yalong Yang, Wenyu Xia, Fritz Lekschas, Carolina Nobre, Robert Krueger, and Hanspeter Pfister. 2022. The Pattern is in the Details: An Evaluation of Interaction Techniques for Locating, Searching, and Contextualizing Details in Multivariate Matrix Visualizations. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*, April 29–May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, ?? pages. <https://doi.org/10.1145/3491102.3517673>

1 INTRODUCTION

Plotting a series of data points in a regular two-dimensional grid—a *matrix visualization*—is a space-efficient approach for visualizing large-scale and dense network [32, 86], tabular [59, 60], set [72], or sequential data [1, 15, 44]. In a matrix visualization, a cell typically only encodes a single value of a data point, e.g., through color. However, for multivariate data, multiple attributes or values (called *details* hereafter) are associated with each data point. We refer to the matrix visualization of multivariate data as *multivariate matrix visualization* (MMV). MMV are widely used in various applications. For example, analysts frequently use them to explore temporal data [6, 7, 10, 12, 30, 83, 89], such as ecologists studied multi-year international food trade through MMV [43], and biologists studied dynamic Bayesian networks with MMVs to model probabilistic dependencies in gene regulation and brain connectivity across time [81]. Additionally, MMVs can also be used to show multiple attributes of a data point [38, 63, 72, 88] or aggregated values from details [24, 26, 49, 51]. For instance, MMV can help pathologists interpret multiclass classifications by visualizing multiple class probabilities at once [63] in histopathology [85], and MMV can be used to help analyze complex multi-variate geographic data [33].

Exploring MMV requires people to investigate the *details* in each single cell, which is usually challenging because each matrix cell's display space is limited and often cannot show all data points in full detail. To enable analysts to effectively explore MMV, a common strategy is to selectively visualize the details of a subset of data points. To this end, three general interaction approaches can be used for MMV: focus+context (or lens), pan&zoom, and overview+details. In this work, we consider MMV where matrix cells change their representation from a single to a multi-value visualization (i.e., from a single color to a line chart) with these interaction techniques. However, adapting these interactions to MMV is not trivial as the MMV's special characteristics need to be taking into consideration. Focus+context magnifies a selected region (referred to as the focus) within the context to show it in greater detail. To make space for the magnified region, the surrounding area (referred to as context) is compressed in size. Not all focus+context techniques are suitable for MMV. The distortion from many focus+context techniques, like a pixel-based fisheye lens, produce irregular shapes of cells that may prohibit effective exploration in MMV. On the other hand, Responsive matrix cells [38], Mélange [27, 28], LiveRAC [55], and TableLens [68] are representative focus+context techniques that are applicable to MMV. Overview+detail technique provides two spatially-separate views with different levels of detail. One view shows the details, and the other offers the context. For example, Burch et al. [17] used overview+detail to facilitate the exploration of MMV. Pan&zoom presents the visualization at a certain detail level while enabling the user to zoom into the visualization and pan

to other regions. For instance, TimeMatrix [89] provides pan&zoom for the users to navigate a MMV in different levels of details.

Focus+context, pan&zoom, and overview+details have been extensively compared in various applications [9, 18, 21, 71, 77, 87] (more details in Sec. 2). However, mixed results were found about their effectiveness, indicating the applied scenarios might largely influence their performance. Thus, it is not applicable and reliable to compile guidelines for MMV only from prior results. Yet, to the best of our knowledge, there is no user study comparing them in the context of MMV. To close this gap, we conducted two extensive user studies to compare the effectiveness of different interaction techniques for MMV. Our goal is to better understand how people interactively explore multivariate details associated with data points in MMV. Thus, in our evaluation, we did not vary the visual encoding and used a simple visualization within each matrix cell to reduce the complexity and potential confounding factors. To this end, we chose a line chart to visualize a time series in each matrix cell, as exploring temporal data is one of the most frequently reported applications for MMV [6, 7, 10, 12, 30, 83, 89], and line charts are a widely-used technique for visualizing temporal data. We are especially interested in the effectiveness of different interaction techniques for navigating MMV and retrieving details from matrix cells, as it is the unique aspect that distinguishes MMV from univariate matrix visualizations. After analyzing the literature [2, 5, 10, 33, 62], taxonomies [47, 48, 57, 61, 82, 87] and real-world applications [43, 63], we derived and tested three fundamental *interaction* tasks that cover a wide range of MMV use cases: *locating* a single cell and then inspecting the details inside; *searching* a region of interest (ROI) of multiple cells to find the cells that match the target pattern; and *contextualizing* patterns using details, which requires inspecting both the details and the context. These three tested tasks can act as “primitive” interactions to serve more sophisticated visual analytic scenarios.

Given the diversity of focus+context techniques [27, 28, 38, 55, 68], the many ways their distortions could impact the perception and task performance, and their overall good performance in some applications [9, 35, 74], we compared different lenses in **our first study**. To identify representative lenses, we followed Carpendale's taxonomy for distortion [19] and identified four lenses: Cartesian lens [73] applies non-linear orthogonal distortion; TableLens [68] with orthogonal distortion (step and stretch); and a fisheye lens technique that is adapted to matrix visualizations [19, 70]. Overall, the results indicate that fisheye lens performed as well as or better than other techniques in the tested tasks. Participants also rated the fisheye lens as the easiest technique for locating matrix cells.

Our second study compared the fisheye lens — the overall best performing focus+context technique from the first study — against a pan&zoom and an overview+detail technique. We found pan&zoom was faster than focus+context and overview+detail techniques in locating and searching for details and as good as overview+detail in contextualizing details. Pan&zoom was also rated with the highest usability and lowest mental demand in almost all tasks. Our results contribute empirical knowledge on the effectiveness of different interaction techniques for exploring MMV. We also discuss promising improvements over existing techniques and potential novel techniques inspired by our results.

2 RELATED WORK

A foundation of exploring MMV is enabling interactive inspection of multiple levels of detail. Several interaction approaches have emerged for this purpose, such as focus+context, overview+detail, and pan&zoom. Cockburn et al. distill the issues with each approach [21]: focus+context distorts the information space; overview+detail requires extra effort for users to relate information between the overview and the detail view; pan&zoom leads to higher working memory as the users can only see one view at a time. Yet, it is still unclear to what extent these findings apply to MMV.

Focus+context (or lens) techniques. A common group of focus+context techniques is lenses, introduced by Bier et al. [13, 79] as generic see-through interfaces between the application and the cursor. Lenses apply magnification to increase the detail in local areas. Lenses can further reveal hidden information, enhance data of interest [45], or suppress distracting information [25, 40]. While emphasizing details, matrix analysis tasks can necessitate all cells of the matrix to be concurrently visible. To achieve this, Carpendale et al. [19] discuss various distortion possibilities with smooth transitions from focus to context in rectangular uniform grids (matrices). Depending on the data, different spatial mapping techniques can be advantageous. Bifocal Display [3] introduces a continuous one-dimensional distortion for 2D data by stretching a column in focus and pushing all other columns aside. The TableLens technique [68] distorts a 2D grid in two dimensions: stretching the columns and rows of the cell in focus only (non-continuous) and shifting the remaining non-magnified cells outward. LiveRAC [55] adapts the idea of TableLens in showing time-series data. Document Lens [70] offers 3D distortion of 2D fields. M elange [28] is a 3D distortion technique to ease comparison tasks. It folds the intervening space to guarantee the visibility of multiple focus regions. Responsive matrix cells [38] combine focus+context with semantic zooming to allow analysts to go from the overview of the matrix to details in cells. Given the diversity of focus+context techniques, we tested the effectiveness of four representative lenses derived from Carpendale et al.'s taxonomy [19]: Cartesian lens [73], two TableLens variations [68], and an adapted fisheye lens [70].

Evaluating Focus+context techniques. Most previous studies on focus+context concentrate on parameter testing and different types of focus+context techniques have not been compared empirically. McGuffin and Balakrishnan [54] investigated the acquisition of targets that dynamically grow in response to users' focus of attention. In their study with 12 participants, they found that performance is governed by the target's size and can be predicted with Fitts' law [31]. Gutwin et al. [34] found that speed-coupled flattening improved focus-targeting when using fisheye distortion with 10 participants. However, fisheye techniques can also introduce reading difficulties. To alleviate this issue, Zanella et al. [91] showed that grids aid readability in a larger study with 30 participants. Finally, Pietriga's study [64] with 10 participants compared different transitions between focus+context and found that gradually increasing translucence was the best choice. Most previous studies were also with a small number of participants. With 48 participants, our study is less outlier-prone and potentially has a smaller margin of error.

Overview+detail techniques. Prominent examples for 2D navigation are horizontal and vertical scrollbars with thumbnails [20]

and mini-maps [90], as well as more distinct linked views [69] with different perspectives for overview and details. MatLink [36] encodes links as curved edges to give detail at the border of the matrix for improving path-finding. Lekschas et al. [49] propose an overview+detail method to compare regions of interest at different scales through interactive small multiples. In their system, each small multiple provides a detailed view of a small local matrix pattern. They later show that this approach can be extended to support pattern-driven guidance in navigating in MMVs [50]. Co-CoNutTrix [41] visualized network data using NodeTrix [37] on a high-resolution large display. We used a standard overview+detail design in our second user study, where we placed the overview and detail view side-by-side, and the user can interactively select the ROI in the overview to update the detail view.

Pan&Zoom techniques. The literature distinguishes between geometric and semantic zooming. The former specifies the spatial scale of magnification. Van Wijk and Nuij summarize smooth and efficient geometric zooming and panning techniques and present a model to calculate a solution for optimal view animations [80]. Semantic zooming, by contrast, changes the level of detail by varying the visual encoding, not only its physical size [16]. Lekschas et al. [49] categorize interaction in matrices into content-agnostic and content-aware approaches. Content-agnostic approaches, such as geometric panning&zooming, operate entirely on the view level, while the latter *"incorporate the data to drive visualization."* ZAME [26] and TimeMatrix [89] are content-aware techniques that rely on semantic zoom. It first reorders [75] rows and columns to group related elements and then aggregates neighboring cells dependent on the current zoom level. Horak et al. [38] provide both geometric and semantic zooming in matrices. However, their technique has not been empirically evaluated. In our second user study, following a widely used design (e.g., Google Maps), we tested the pan&zoom condition, which allows the user to scroll the mouse wheel to continuously zoom in and out a certain region of the matrix.

Evaluating focus+context, overview+detail, and pan&zoom. These three interaction techniques have been extensively evaluated in various applications, but not in the context of MMV. Baudisch et al. [9] found that focus+context had reduced error rates and time (up to 36% faster) over pan&zoom and overview+detail for finding connections on a circuit board and the closest hotels on a map. Similarly, Gutwin et al. [35] concluded fisheye views to be advantageous over overview+details and zooming for large steering tasks. Shoemaker and Gutwin [74] also found fisheye lens superior over standalone panning and over zooming for multi-point target acquisition on images. On the other hand, R onne and Hornb æk [71] had opposite findings for locating, comparing, and tracking objects on geographic maps. They found fisheye had the worst performance, while overview+detail performed best. These previous studies had mixed results for different applications, and none of them were conducted in the context of MMV. Most similar to our second study is the study with 12 participants by Pietriga et al. [65] for multiscale search tasks. They found overview+details superior over the fisheye lens, while both techniques outperformed pan&zoom. They tested the conditions in a matrix-like application but with no multivariate details in the cells. They also only tested one searching task, and the tested fisheye lens was the classical one with non-linear radial distortion, which breaks the regularity of the grid.

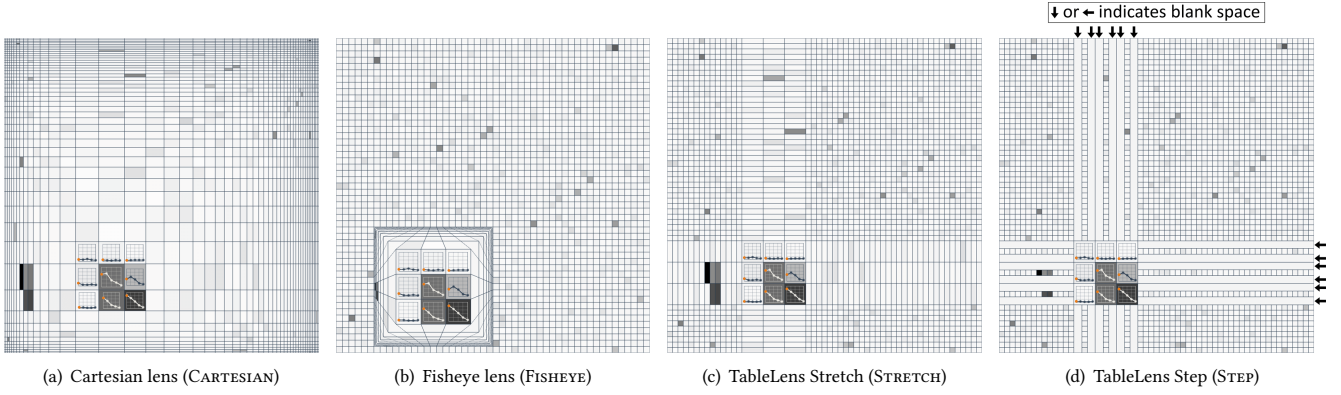


Figure 2: Study 1 visualization conditions with 50×50 matrices: four interactive lenses tested in the user study. An interactive demo is available at <https://mmvdemo.github.io/>, and has been tested with Chrome and Edge browsers.

3 STUDY 1 – DIFFERENT LENSES IN MMV

This first study is intended to address the gaps in the literature described above in terms of how different distortions can impact perception and task performance in using focus+context (or lens) techniques for exploring MMV. As this study was the first to compare different lenses in MMV, there was little empirical knowledge about the user performance with different lenses. Thus, our first study is exploratory rather than confirmatory. We pre-registered the study at <https://osf.io/dxsr5>. Meanwhile, test conditions are demonstrated in the supplementary video, and detailed results of statistical tests are provided in supplementary materials.

3.1 Experimental Conditions

Using lenses to explore MMV selectively enlarges an ROI of the matrix, so that the enlarged cells have enough space to show the detail. These enlarged cells are also referred as focus cells. To make space for the focus cells, lenses introduce two types of distortions: *focal* and *contextual* distortion. Focal distortion applies to cells at the inner border of the lens. Contextual distortion, on the other hand, applies to cells outside the lens.

Unlike lenses in a map or images, lenses in MMV have more constraints. While, there are more flexibility with their elementary units in a map or images, the cells in MMV are all the same size and laid in a regular grid (i.e., rows and columns are orthogonal to each other). Additionally, according to Carpendale et al. [19], gaps are also considered an important distortion characteristic. In summary, we identified three characteristics for the distortions of lenses in MMV: **regularity**—whether the cells are rendered in a regular or orthogonal grid, **uniformity**—whether the cells are sized uniformly, and **continuity**—whether the cells are laid out continuously. Those characteristics can be used to model both the *focal* and *contextual* distortions in the lenses. We chose four lenses for our study, according to Carpendale et al.’s distortion taxonomy [19]:

CARTESIAN distorts the entire matrix continuously such that the cells are proportional sized based on their distances to the cursor (Fig. 2(a)). In **CARTESIAN**, cells in focal and contextual regions are all in a regular grid but sized differently.

FISHEYE magnifies the center part of the focus and shrinks the surrounding area around the lens’s inner boundary. The focal cells need to be rendered in a regular grid and sized uniformly. To

Factor	Cartesian	Fisheye	TableLens (Stretch)	TableLens (Step)
Focal distortion				
Regularity	✓	Partial	✓	✓
Uniformity	X	Partial	✓	✓
Continuity	✓	✓	✓	✓
Contextual distortion				
Regularity	✓	✓	✓	✓
Uniformity	X	✓	Partial	✓
Continuity	✓	✓	✓	Partial

Table 1: Characteristic Comparison: of the four tested lenses.

continuously embed the focal region inside the contextual region, distortion must be applied to the focal area’s *transition area*. As a result, cells inside the transition area are rendered irregularly and are sized differently (see Fig. 2(b)).

STRETCH and **STEP** are two variations of TableLens that enlarge a fixed number of rows and columns around the focal point and uniformly compress the remaining matrix. **STRETCH** stretches the enlarged rows and columns on either of the two axes (Fig. 2(c)); **STEP** preserves the cells’ aspect ratio by adding white space around enlarged rows and columns on either of the two axes, which introduces discontinuities (see the blank space in Fig. 2(d)). **STRETCH** and **STEP** both have regular and uniform cells in the focal region.

We summarize the characteristics of the four tested lenses in Table 1. The four tested lenses cover a variety of characteristics, and the study is to investigate how those characteristics affect perception and interaction performance.

3.2 Data

We used time series as the multivariate data in our study, as it is widely used [6, 10, 12, 30, 83, 89] and has not been empirically tested in the context of MMV before. We generated task datasets consisting of three dimensions (x , y , and t): with x and y as the rows and columns in the matrix and t as the number of time instances in the time series. For a particular value of t , the x and y dimensions are shown as a traditional univariate matrix, which we refer to as the *context*. The t dimension is revealed interactively upon placing the lens over a focus region. Enlarging the cells under a lens’s focal area provides space for displaying this t dimension as a line chart, which we refer to as the *focus*. Each data set contains $x \times y \times t$ values, and each cell contains t values as multivariate details.




Figure 3: Data Generation. To generate a unique dataset for each trial, we following this 3-step pipeline. First, we sampled the location of target patterns. Second, for each target pattern location we sampled cluster of target pattern types. Finally, for each cell we generate a 5-point temporal pattern instance based on the cell’s assigned pattern type.

Pattern type	Data distribution
Upward	Beta with $\alpha=2$ and with $\beta=1$
Downward	Beta with $\alpha=1$ and with $\beta=2$
Tent	Beta with $\alpha=4$ and with $\beta=4$
Trough	Beta with $\alpha=\frac{1}{3}$ and with $\beta=\frac{1}{3}$
Background	Fréchet with $\alpha=5$, $s=1$, $m=0$

Table 2: Pattern Distribution Functions: Pattern instances are generated from the histogram of 100 sampled values from the associated data distribution. The probability density function of each pattern type is shown in Fig. 4.

We included two data sizes: **SMALL** with $50 \times 50 \times 5$ and **LARGE** with $100 \times 100 \times 5$. We decided to test large matrices because small matrices have enough space for each cell to show the multivariate details constantly, and interaction is less necessary for them. We also chose to study the scalability in terms of the matrices’ size and keep the detail’s size (i.e., the number of time instances) unchanged.

The goal of the tasks is to evaluate and compare the temporal patterns that arise along the third dimension (i.e., the *details*). As shown in Fig. 3, our data generation consisted of three steps: first we sampled a matrix of different pattern *types*, then we expanded the target pattern types into clusters, and finally sampled the actual pattern *instances* for each cell. To avoid memory effects and ensure that participants would have to inspect the patterns under the focal area, we sampled pattern instances from five distinct distributions (Fig. 4) inspired by temporal patterns described by Correll and Gleicher [22]: upward, downward, tent, trough, and background.

In the first step, we created a matrix of pattern types. In the beginning, the matrix contained only background pattern types (Fig. 3.1). We then randomly placed non-target pattern types into the matrix (Fig. 3.2). We added these non-target patterns as lightweight distractions and made the final dataset more realistic. We then randomly sampled a position for the target pattern type. Further, to make it easier to locate the target cells, we sampled a cluster of target pattern types using a 2D Gaussian distribution centered on the previously determined target pattern type location (Fig. 3.3).

Finally, for each cell, we generated a pattern instance by randomly sampling 100 values from different distributions (Table 2) and aggregated them in a 5-bin histogram. This approach created patterns instances that differ slightly in shape and magnitude while still being distinct enough to avoid ambiguity (Fig. 4). This approach strikes a balance between predictability and generality.

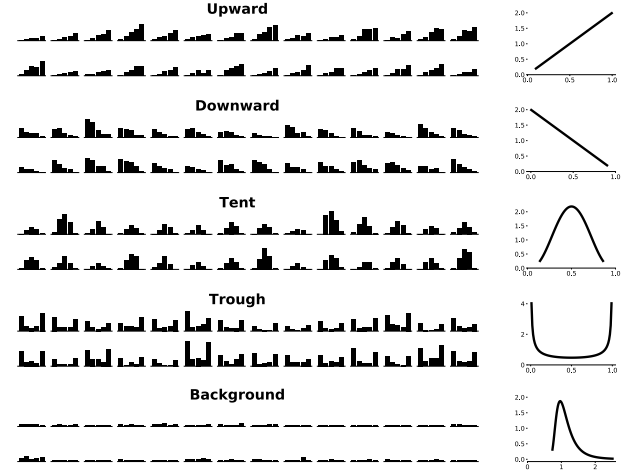


Figure 4: Pattern Types. Example pattern instances for each of the five pattern types: upward, downward, tent, trough, and background. The instances slightly differ in their shape and magnitude to mimic realistic data. On the right we plot the probability density function of each pattern.

3.3 Interactions and Tasks

The participants were asked to interact with the MMV, which by default showed a slice of our 3D dataset as a univariate matrix visualization (i.e., a heatmap for one of the five time instances, see Sec. 3.2 for details). We used a continuous color schema from white to black to encode matrix cells. Darker cells indicate higher values. Such a color schema is colorblind-friendly. Upon moving the mouse cursor over the MMV, the lens enlarges an area to show the details of the time series as a line chart. In each chart, the line connects five dots representing the five values. The dot representing the value of the background color is additionally highlighted to represent the currently selected time instance. Participants can switch the time instance by clicking on a respective dot in an embedded line chart. To clearly present interactive line charts while still keeping the context cells legible, we conducted a series of internal tests to find an appropriate combination of the parameters for *the number of cells to be enlarged* and *the magnification factor*. We ensure the size of the enlarged area and each line chart to be consistent across different lenses. For the two data sizes we tested, the enlarged cell is in the same size: four times the side length as the original cell in **SMALL** data and eight times in **LARGE** data. A line chart in such a size can be reasonably interpreted and interacted by users. We kept the ratio of the enlarged cells the same as the matrix, i.e., 1:1, and decided to enlarge 3×3 matrix cells as the focal area to show their line charts. Increasing the number of enlarged cells or magnification factor makes it challenging to interpret the color of the surrounding cells, even for screens with a standard resolution. For example, on a Full-HD (or 1080p) screen, we used 800×800 pixels to visualize a 100×100 matrix for **LARGE** data, and the size of a context cell is 5×5 pixels in **FISHEYE**, **STRETCH**, and **STEP** when the lens is on top of the matrix. Some context cells in **CARTESIAN** are even smaller. According to our internal tests, interpreting colors in context cells smaller than 5×5 pixels is difficult. Fig. 2 demonstrates the tested interaction techniques for enlarging their focal areas.

The most basic interaction for exploring an MMV includes three steps: first finding the cell(s) of interest, then moving the cursor towards them to enlarge them, and finally checking their embedded details. Additionally, in some cases, users must inspect both the focal and context areas. Past research in HCI and visualization proposed taxonomies [47, 48, 57, 61, 82, 87] for these interactions and conducted studies in various applications (see Sec. 2), but not with MMV. We analyzed previous work to break down the fundamental MMV interactions into four components. First, **wayfinding** is the process of searching targets cells. Second, **travel** refers to the act of moving the mouse cursor to the target cells. Third, **interpretation** is the activity of interpreting the targets' visual encoding. And finally, **context-switching** refers to re-interpreting a changed view, for example, when updating visualization through interaction or moving the focus to a different part of the view. We then designed three tasks to cover different aspects of the identified components. Since we used the same visual encoding for all conditions (i.e., the line chart), we did not expect a noticeable performance difference in *interpretation*. Thus, our focus is on evaluating the other three. In the following, we first describe the study tasks with practical examples in the context of multi-year population data for counties in the United States in an MMV. This data is easily accessible and understandable. Each matrix cell represents a county and contains multi-year population data. The cells are typically placed according to their relative geographic locations, which is similar to the tile map representation [56] used by Wood et al. [83]. We then discuss the rationale and motivation of our task choices.

In the first task (LOCATE), we asked participants to click on a specific cell highlighted with an orange outline. Our goal is to test how the distortion influences the participants' perceptual ability to locate a specific cell. Thus, we remove the highlighting as soon as the participants move their cursor into the matrix. For accessibility, the highlighting reappears once the cursor is moved outside the matrix. Additionally, we followed Blanch et al.'s [14] approach and added visual distractors, i.e., non-target patterns (Sec. 3.2) in our case. A frequent operation in analyzing population data is to investigate the temporal trend of a given location, like *"what is the temporal trend of Middlesex County, MA in the last five years?"*

The LOCATE task was designed to inspect the **travel** component in interactions. Locating and selecting an element is the most common task in graphical user interfaces and visualizations and is a primitive visualization interaction [57, 76]. It is also a standard task tested in many user studies (e.g., [42, 71]). Fitts' law provides a way to quantify the performance of basic target selection [76]. However, the standard model does not consider the lenses' distortion effects. This task aims to investigate how different types of distortion influence the performance of locating target.

For the second task (SEARCH), we asked the participants to search for the cell with the highest single value among a cluster of cells, which is a 7×7 region for a 3×3 lens in the study. Since we test the ability to locate a cell in the LOCATE task, we decided to permanently highlight the search area with an orange outline. To enforce the use of the lenses, we pre-selected a value of the time series that does not reveal the target patterns. Only when the user employs the lens the relevant details of the multivariate pattern will be visible. An example of this task in population analysis can

be *"Within New England (a set of counties), which county has the largest number of population of a single year in the last five years?"*

The SEARCH task involves both the **travel** and **wayfinding** components. *Wayfinding* is an essential step for any high-level visual analytics task [57]. In order to find the cell with the highest value, the participants had to inspect and compare the details of multiple cells. Similar tasks have been tested in other contexts, for example, by Pietriga et al. [65] for multiscale searching and Jakobsen & Hornbæk for geographic maps [71]. We expect that the different lens distortions will influence the performance of *wayfinding*, especially in an interactive scenario. It is impractical to only test *wayfinding* performance without physically *traveling* to the targets. Thus, we included both these two components in this task.

In the third task (CONTEXT), we asked the participants to find the largest cluster at the time instance where a given cell reaches its highest value. The participants needed to move the mouse cursor to a cell highlighted with an orange border and click on the dot representing the highest value. The representation of MMV (i.e., the heatmap) will then be updated to the time instance corresponding to the clicked dot. Subsequently, several clusters in the sizes between 5×5 to 7×7 of dark cells appeared in the matrix, and participants were asked to select the largest one. For instance, a practical use case in population analysis can be *"At the year when the population of Orange County, FL reaches its peak value, where is the largest region with high population?"*

The CONTEXT task includes the **travel**, **wayfinding**, and **context-switching** components. Context-switching frequently happens in interactive visualization and multi-scale navigation and has been tested in various scenarios [66, 67, 71, 87], but not with MMV. In MMV, users have to *switch their context* in many scenarios, e.g., when enlarging the cells to show the line charts, changing the time instance, and moving their focus between the focal and contextual areas. We expect different types of distortion will influence the *context-switching* performance. Again, it is unrealistic to test *context-switching* without *travel* and *wayfinding*. Thus, we include all these three components in this task.

3.4 Experimental Design

We included two factors in the user study: LENS and SIZE. The LENS had four different lenses, as described in Sec. 3.1. The SIZE had two data sizes as described in Sec. 3.2. The experiment followed a full-factorial within-subject design. We used a Latin square (4 groups) to balance the visualizations but kept the ordering of tasks consistent: first LOCATE, then SEARCH, and finally CONTEXT. Each participant completed 48 study trials: 4 visualizations × 2 data sizes × 3 tasks × 2 repetitions. The entire study has been tested on common resolution settings (FHD, QHD and UHD).

Participants. We recruited 48 participants on Prolific (<https://www.prolific.co>). All participants were located in the US and spoke English natively. To ensure data quality, we restrict participation to workers who had an acceptance rate above 90%. Our final participant pool consisted of 19 female, 26 male, and three non-binary participants. Out of those participants, twelve had a master's degree, 16 had a bachelor's degree, 14 had a high school degree, and six did not specify their education levels. Finally, 4 participants were between 18-20, 17 participants were between the age of 21-30, 18 participants were between 31-40, five participants were

between 41-50, four participants were above 50. We compensated each participant with 9 USD, for an hourly rate of 12 USD.

Procedures. Participants were first presented with the consent form, which provided information about the study's purpose and the procedure. After signing the consent form electronically, the participants had to watch a short training video (1 minute and 13 seconds) that demonstrated how to read and interact the MMV.

Participants completed the three tasks one by one based on a Latin square design. Prior to working with a new lens, we showed a video demonstrating how to interact with the matrix using the current lens. Each (visualization \times task) block started with two training trials followed by study trials. Before each training trial, we encouraged participants to get familiar with the visualization condition and explicitly told them they were not timed for the training. We also ensured that participants submitted the correct answers in training trials before we allowed them to proceed. Before starting the study trials, we asked the participants to complete the trials "as accurately and as quickly as they can, and accuracy is more important," and informed them that these trials were timed. To start a trial, participants had to click on a "start" button placed in the same location above the MMV. This ensured a consistent cursor starting point and precisely measured the task duration. The visualization only appeared after clicking the start button.

After each task, participants were asked to rate each visualization's perceived difficulty and write their justifications. We collected the demographic information as the final step. The average completion time was around 45 minutes.

Measurements. We collected the following measurements during the user study: **Time.** We measured the time in milliseconds from the moment the user clicked on the start button until they selected an answer. **Accuracy.** We measured the participants accuracy as the ratio of correct over all answers. **Perceived Difficulty Rating.** After the user completed a task with all four lenses we asked the participants to rate "how hard was performing the task with each of the visualizations?" on a 5-point Likert scale ranging from *easy* (1) to *hard* (5). The questionnaire listed visualizations in the same order as presented in the user study with figures. **Qualitative Feedback.** We also asked participants to optionally justify their perceived difficulty ratings in text.

Statistical Analysis. For dependent variables or their transformed values that met the normality assumption (i.e., time), we used *linear mixed modeling* to evaluate the effect of independent variables on the dependent variables [8]. Compared to repeated measure ANOVA, linear mixed modeling does not have the constraint of sphericity [29, Ch. 13]. We modeled all independent variables (four visualization techniques and two data sizes) and their interactions as fixed effects. A within-subject design with random intercepts was used for all models. We evaluated the significance of the inclusion of an independent variable or interaction terms using log-likelihood ratio. We then performed Tukey's HSD post-hoc tests for pair-wise comparisons using the least square means [52]. We used predicted vs. residual and Q-Q plots to graphically evaluate the homoscedasticity and normality of the Pearson residuals respectively. For other dependent variables that cannot meet the normality assumption (i.e., accuracy and perceived difficulty rating), we used the *Friedman test* to evaluate the effect of the independent variable,

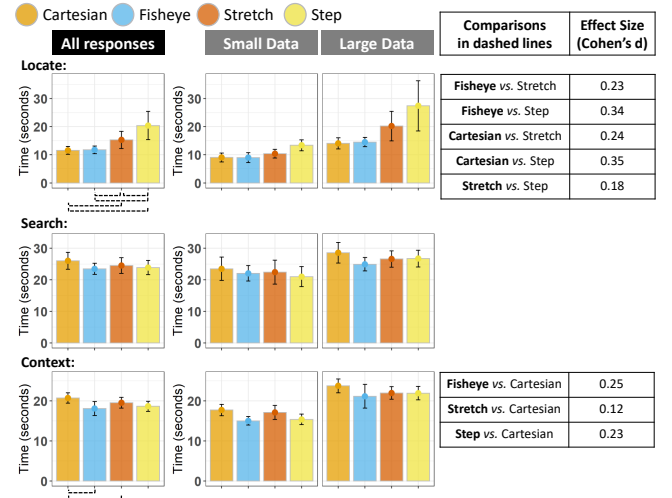


Figure 5: Time by task and in different data sizes. Confidence intervals indicate 95% confidence for mean values. Dashed lines indicate statistical significance for $p < .05$. Tables are showing their effect sizes.

as well as a Wilcoxon-Nemenyi-McDonald-Thompson test for pair-wise comparisons. Significance values are reported for $p < .05$ (*), $p < .01$ (**), and $p < .001$ (***), abbreviated by the number of stars in parenthesis. Numbers in parentheses indicate mean values and 95% confidence intervals (CI). We also calculated the Cohen's d as indicators of effect size for significant comparisons.

3.5 Results

The accuracy was similarly high across all conditions: on average, 95.3% for LOCATE, 92.3% for SEARCH, and 78.1% for CONTEXT. We did not find any significant differences between LENS and SIZE on accuracy. Therefore, we focus our analysis on the time (Fig. 5), perceived difficulty (Fig. 6), and qualitative feedback.

We found LENS had a significant effect on time in both LOCATE (***) and CONTEXT (***) tasks, but no significant effect in the SEARCH task ($p = 0.163$). We also found SIZE had a significant effect on time in all tasks (all ***). No significant effect has been found in the interaction between LENS and SIZE for all tasks. For the perceived difficulty ratings, LENS had a significant effect in the LOCATE task (**), but not in SEARCH and CONTEXT tasks. All statistical results are included in the supplementary materials.

Quantitative Key Findings.

FISHEYE was the best performing technique. FISHEYE (11.8s, CI=1.4s) and CARTESIAN (11.7s, CI=1.4s) had a similar performance in the LOCATE task, and they both outperformed STRETCH (15.3s, CI=3s) and STEP (20.4s, CI=5s) (all ***). The perceived difficulty ratings mostly aligned with the performance results. I.e., participants rated FISHEYE (2.19, CI=0.33) easier than STRETCH (3, CI=0.33, *) and STEP (3.77, CI=0.34, ***). CARTESIAN (2.71, CI=0.39) was also rated easier than STEP (**). FISHEYE (18s, CI=1.8s) also outperformed CARTESIAN (20.7s, CI=1.3s) in the CONTEXT task (**). Overall, FISHEYE was the best choice for the tested tasks.

CARTESIAN was not ideal for the CONTEXT task. CARTESIAN (20.7s, CI=1.3s) was slower than FISHEYE (18s, CI=1.8s, ***), STEP (18.6s, CI=1.3s, **), and STRETCH (19.5s, CI=1.3s, *). Participants




Figure 6: Perceived difficulty ratings by task. Dashed lines indicate $p < .05$.

also tended to consider CARTESIAN (2.69, CI=0.38) more difficult than FISHEYE (2.27, CI=0.34) and STRETCH (2.46, CI=0.32), but not statistically significant.

STRETCH had advantage over STEP in the LOCATE task. The only performance difference between these two condition is that STRETCH (15.3s, CI=3s) was faster than STEP (20.4s, CI=5s) in the LOCATE task (***). Again, the perceived difficulty ratings aligned with the performance, where participants found STRETCH (3, CI=0.33) easier than STEP (3.77, CI=0.34) in the LOCATE task (*).

All performed similar in the SEARCH task. We did not find an effect of visualization on time or on the perceived difficulty rating. **Qualitative Feedback.**

We also asked participants to justify their perceived difficulty rating after each task. We analyze the collected feedback to get an overview of the pros and cons of each lens.

CARTESIAN was commented to be “natural” by six participants. They found it to be intuitive to have cells that are closer to the cursor to be larger. However, 18 participants complained its distortion. More specifically, five participants found it difficult to know the cursor’s current location within the matrix. Six found the distortion results in unexpected “jump” and made it challenging “to get to the right cell.” One participant also felt “sea-sick”. In the SEARCH task, one participant found the distortion made it hard to “see the boundary of the highlighted region.” In the CONTEXT task, two participants found it tough to “see far away clusters.”

FISHEYE was commented to be “easy to use” by 18 participants. More specifically, nine found it “easy to follow,” four found it “not jump so much” and “more in line with the cursor,” two found it “pinpoint fast,” two found it “easy to locate,” and one found it “easy to know the current location.” Four participants also found “(cells are) the same size outside the fishseye” and easy to “compare the clusters at once” in the CONTEXT task. However, nine participants found it “hard to see the surroundings” due to irregular shapes in the transition area and they sometimes found it hard to precisely identify the highlighted box in the SEARCH task.

STRETCH was reported by four participants who found its regularity to be beneficial: “lined up with the boxes” and “(easy) to keep context in my head.” Three participants explicitly commented it to be “better than the step.” However, 14 participants found it disorienting, like “hard to get my bearing” and “alignment off.”

STEP was found positive with its regularity by two participants. However, ten people found it “disorienting.” Eight also found the empty space in the enlarged row and column confusing, with one specifically pointed out that “the gap breaks out the clusters” in the CONTEXT task. Four reported it challenging for “precise moves.”

3.6 Discussion

Most lenses performed similarly in the first user study, with a few notable differences. We discuss the potential reasons for these differences, and provide guidelines for future lens design in MMV.

Correspondence facilitates precise locating. LOCATE is a fundamental part of many high-level tasks in exploring MMV. In this task, after moving the mouse cursor into the matrix, the context gets distorted. Thus, it is important to find a good entry point to facilitate this task. A common strategy is to enter at the same row or column as the target cell. However, due to distortion, the cursor may not land on the target row or column. We define the difference between the expected and actually hovered row or column after entering the matrix as correspondence. Higher correspondence means less offset and gives the user more predictable interactions. To find out the lenses’ correspondence, we simulated the cursor moving into the matrix from the top and scanning the entire boundary with an incremental one pixel each time. We found FISHEYE and CARTESIAN have a perfect correspondence. However, for STRETCH and STEP, the offsets vary from 0 to 3 cells in 50×50 matrices and from 0 to 6 cells in 100×100 matrices (see supplementary material for details). In summary, the ranking of correspondence is: FISHEYE = CARTESIAN > STRETCH = STEP. The performance and perceived difficulty results align with correspondence, where FISHEYE and CARTESIAN were faster and generally perceived as easier than STRETCH and STEP in the LOCATE task. Appert et al. [4] discussed this in pixel-based lenses and proposed interaction techniques to improve correspondence. However, it is unclear how to adapt their techniques to MMV. On the other hand, their evaluation results partially aligned with our results and confirmed our hypothesis: techniques with higher correspondence have better locating performance.

Discontinuity affects the performance for precise locating. STEP was slower than STRETCH in the LOCATE task. The perceived difficulty also aligned with the time performance, where STEP was considered more difficult than STRETCH in the LOCATE task. The only difference in these two lenses is the way they visualize the enlarged rows and columns: STRETCH stretches them, while STEP aligns the cells in the center and leaves blank space with discontinuity. We conjectured that this discontinuity hinders the ability of precise movement in the MMV, thus degrading the performance of STEP. This is also reflected in participants’ comments, where eight specifically found the “gaps” confusing.

Uniformity facilitates contextualizing patterns. CARTESIAN was the slowest in the CONTEXT task. This task had two components, where the first component is similar to the LOCATE task, and the second component required identifying the cluster with most number of cells in the context. The second component started right after the first one, which means the visualization was still distorted by the lenses. For uniform distortion, the participants only need to compare the areas of the clusters. However, when the context was distorted non-uniformly, comparing the areas of clusters may lead to a wrong answer. As a result, participants had to count the number of cells, which is expected to take longer. Participants might also first remove the distortion by moving the cursor outside, which would prolong the task. From Table 1 and Sec. 3.1, we can see that the ranking of contextual uniformity is: FISHEYE = STEP > STRETCH > CARTESIAN. This ranking aligns with the time performance. In summary, our results suggest that the performance was proportional to the level of contextual uniformity.

Small regions with irregular distortion might not affect performance. Within the lenses, all conditions had perfect regularity, except for the FISHEYE, where the cells in the transition area were not in a regular grid. Despite this irregularity, FISHEYE had the best overall performance. This does not necessarily mean that regularity is not important for lenses in MMV, since FISHEYE only has a limited region that is irregular. Further studies are required to confirm the effect of regularity in other regions (i.e., focal and contextual regions) and at different sizes.

Different distortions do not affect coarse locating. All lenses had similar performance in the SEARCH task. We believe this is due to having a large target region (7×7) in this task. With a large target, participants only needed to coarsely locate a region instead of precisely locating a single cell like in the LOCATE task. As a result, the correspondence, discontinuity and other distortion characteristics do not lead to significant performance difference in coarse locating.

4 STUDY 2 – FOCUS+CONTEXT, OVERVIEW+DETAIL AND PAN&ZOOM IN MMV

This study is intended to address the literature gaps in terms of identifying which is the best interaction techniques among FOCUS+CONTEXT, OVERVIEW+DETAIL and PAN&ZOOM for MMV. We chose FISHEYE as the representative technique for FOCUS+CONTEXT as it was the best performing technique from the first study. We designed our first study to be generalizable for testing interaction techniques in MMV, i.e., the same experimental setups can be used to test interaction techniques other than just lenses. Therefore, we reused many materials from the first study in our second study. Same as the first user study, we designed the second user study as an exploratory study rather than hypothesis testing. This is because the literature had mixed results for the comparisons of the three generic interaction techniques and with little empirical knowledge about the comparison of them in the context of MMV. As a result, there was not enough guidance to generate reliable hypotheses. We also pre-registered this study at <https://osf.io/q4zp9>.

4.1 Experimental Conditions

Same as the first user study, the tested conditions use different ways to selectively enlarge an ROI of the matrix so that the enlarged cells have sufficient display space to show the multivariate details. Unlike the first study, where all conditions superimpose the enlarged ROI (or focused view) within the matrix (or the contextual view), different conditions use different strategies to manage the focused and contextual views in the second study.

FOCUS+CONTEXT: we used the same design of FISHEYE from the first user study (Fig. 2(b)). FOCUS+CONTEXT displays the focused view inside the contextual view.

OVERVIEW+DETAIL: we placed a separate view as the *detail view* on the right of the overview (the matrix). Some designs place one view at a fixed location (e.g., top right corner) inside the other view (e.g., in [71]). However, such a design is not suitable in our case, as it will occlude part of the matrix. Thus, we decided to place the two views side-by-side. In the detail view, the multivariate details (i.e., lines chart in this study) are rendered for a selected ROI. A red box is used to indicate the ROI in the matrix. The user can drag the red box within the matrix, and the detail view will update in

real-time. This tightly coupled design between overview and detail view is suggested by Hornbæk et al. [39]. We set the size of the detail view and the number of line charts to the same as the FISHEYE, i.e., the detail view always renders 3×3 line charts in the same size as they are in FISHEYE. A demonstration of OVERVIEW+DETAIL is presented in Fig. 7(a). OVERVIEW+DETAIL uses a spatial separation between the focused and contextual views.

PAN&ZOOM: the participant can scroll the mouse wheel to zoom in or out an ROI of the matrix continuously. The mouse cursor is used as the center of zooming, and the transitions are animated. When the user zooms into a certain level, where the cells' size is equal to or larger than a threshold, the line charts will be rendered inside the cells. We set the threshold as the size of enlarged cells with line charts in FISHEYE. The user can also pan to inspect different parts of the matrix. The design of PAN&ZOOM follows the widespread map interfaces (e.g., Google Maps), and it is a standard design in many user studies (e.g., in [65, 71, 84]). A demonstration of PAN&ZOOM is presented in Fig. 7(b). PAN&ZOOM uses a temporal separation between the focused and contextual views, i.e., only one zoom level can be viewed at a time.

4.2 Experimental Setups

Experimental Design. Similar to the first study, we have two factors: TECHNIQUE (see Fig. 7) and SIZE. Each participant completed 36 study trials: 3 visualizations×2 data sizes×3 tasks×2 repetitions.

Data and Tasks. We reused the data from the first study. To avoid learning effects, we used a screening tool from Prolific to limit participants to people who have not seen our first study. We slightly modified the LOCATE task from the first study to adapt to the new interaction conditions. The LOCATE task in the first study asked participants to click on a highlighted cell as it is important for understanding how different distortion from lenses affects precise selection. However, in the second study, OVERVIEW+DETAIL and PAN&ZOOM do not have any distortion. Thus, the previous task can lead to undesired bias. Instead, in the second study, we asked participants to interpret the temporal pattern and select an answer from five options (see Fig. 4). With the adapted LOCATE task, we can compare the effectiveness of interpreting a given cell's detail in MMV, which involves locating the target cell and navigating to details. For the SEARCH and CONTEXT tasks, we believe OVERVIEW+DETAIL and PAN&ZOOM do not introduce performance bias. Therefore, we used the same SEARCH and CONTEXT tasks from the first user study.

Participants. We recruited 45 participants on Prolific. As mentioned, to avoid learning effect, we filtered out participants from the first study at screening stage. All participants were located in the US and spoke English natively. To ensure data quality, we again restrict participation to workers who had an acceptance rate above 90%. Our final participant pool consisted of 16 female, and 29 male. Out of those participants, one had a PhD degree, one had a master degree, 15 had a bachelor degree, 21 had a high school degree, and seven did not specify their education levels. Finally, 7 participants were between the age of 18-20, 22 participants were between the age of 21-30, 11 participants were between the age of 31-40, one participant was between the age of 41-50, and four participants were above 50. We compensated each participant with 7 USD.




Figure 7: Study 2 visualization additional conditions (with 50×50 matrices): (a) OVERVIEW+DETAIL, the *detail view* on the right shows the details of the *red box* in left matrix. The user can drag the red box to update the detail view in real-time. **(b) PAN&ZOOM**, the user can scroll the mouse wheel to zoom in or out a certain region of the matrix. In addition to OVERVIEW+DETAIL and PAN&ZOOM conditions, we used the same design of FISHEYE from the first user study, as demonstrated in Fig. 2(b). An interactive demo is available at <https://mmvdemo.github.io/>, and has been tested with Chrome and Edge browsers.

Procedures. We used similar procedures as in the first study, except after each task, instead of only rating the perceived difficulty, we asked participants to rate the overall *usability*, *mental demand*, and *physical demand* for each visualization. This change is intended to obtain a more nuanced understanding of the perceived effectiveness. The average completion time was around 35 minutes.

Measurements and Statistical Analysis. We collected similar measures as in the first study, including *time*, *accuracy*, and *qualitative feedback*. As described in the *procedures*, we also collected the *subjective ratings* of *usability*, *mental demand*, and *physical demand* for each visualization at each task. We expected that the additional ratings could help us towards a more nuanced understanding of the perceived performance of different techniques. We used the same method as in the first study to analyze the collected data.

4.3 Results

Same as the first user study, the accuracy was high across all conditions: on average, 98.7% for LOCATE, 93.0% for SEARCH, and 74.8% for CONTEXT. We did not find any significant differences on accuracy. Therefore, we focus our analysis on the time (Fig. 8), subjective ratings (Fig. 9), and qualitative feedback.

We found **TECHNIQUE** had a significant effect on time in all tasks: LOCATE (***), SEARCH (***), and CONTEXT (*). We also found **SIZE** had a significant effect on time in SEARCH (*), and a marginal effect in CONTEXT ($p = 0.092$), but not in LOCATE ($p = 0.172$). No significant effect has been found in the interaction between **LENS** and **SIZE** on time for all tasks. In terms of subjective ratings, we found **TECHNIQUE** had a significant effect on usability and mental demand in all tasks (all ***). For physical demand, we found significance in the SEARCH (***), and CONTEXT (**) tasks. All statistical results are included in the supplementary materials.

Quantitative Key Findings

PAN&ZOOM was the best performing technique. In the LOCATE task, PAN&ZOOM (12.1s, CI=1.5s) was faster than FOCUS+CONTEXT (13.8s, CI=1.5s, $p = 0.081$) and OVERVIEW+DETAIL (17.4s, CI=3.3s, ***). In the SEARCH task, PAN&ZOOM (14.5s, CI=1.2s) was faster than FOCUS+CONTEXT (23.0s, CI=1.9s, ***), and OVERVIEW+DETAIL (26.3s, CI=1.7s, ***). In the CONTEXT task, PAN&ZOOM (16.5s, CI=1.1s) had a similar performance as OVERVIEW+DETAIL (15.8s, CI=1.6s), and tended to be faster than FOCUS+CONTEXT (18.4s,




Figure 8: Time by task and in different data sizes. Confidence intervals indicate 95% confidence for mean values. Dashed lines indicate statistical significance for $p < .05$ (black) and $0.05 < p < 0.1$ (gray).

CI=2.0s), but not significant. The subjective ratings mostly aligned with the performance: participants rated PAN&ZOOM with a higher usability and lower mental demand than FOCUS+CONTEXT in all tasks, all ***. Participants also found PAN&ZOOM with a higher usability, lower mental and physical demand than OVERVIEW+DETAIL in the SEARCH task, all ***. Overall, PAN&ZOOM was the best choice for the tested tasks.

OVERVIEW+DETAIL performed well in the CONTEXT task. In the CONTEXT task, OVERVIEW+DETAIL (15.8s, CI=1.6s) was faster than FOCUS+CONTEXT (18.4s, CI=2.0s, *). It also tended to be slightly faster than PAN&ZOOM (16.5s, CI=1.1s), but that was not statistically significant. Again, subjective ratings mostly aligned with the performance results. OVERVIEW+DETAIL was rated to have a higher usability (***), lower mental (**) and physical demand (**) than FOCUS+CONTEXT for the CONTEXT task. OVERVIEW+DETAIL was

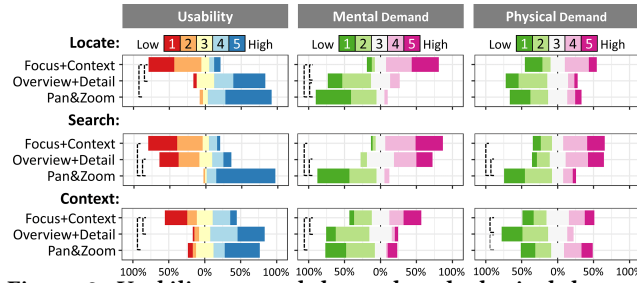


Figure 9: Usability, mental demand, and physical demand ratings by task. Dashed lines indicate $p < .05$ (black) and $0.05 < p < 0.1$ (gray).

also rated to be marginally less physical demand than OVERVIEW+DETAIL for the CONTEXT task ($p = 0.090$).

OVERVIEW+DETAIL was the slowest technique in the LOCATE and SEARCH tasks. Despite its good performance in the CONTEXT task, OVERVIEW+DETAIL was slower than PAN&ZOOM, all ***. OVERVIEW+DETAIL was also slower than FOCUS+CONTEXT (23.0s, CI=1.9s) in the SEARCH task (***), and was marginally slower than FOCUS+CONTEXT (13.8s, CI=1.5s) in the LOCATE task ($p = 0.055$).

FOCUS+CONTEXT received the worst subjective ratings. FOCUS+CONTEXT had the second best performance in the LOCATE task. However, it was rated as with the lowest usability and highest mental demand (all ***). For the SEARCH task, it was rated with a lower usability, higher mental and physical demand than PAN&ZOOM (all ***). For the CONTEXT task, it was again rated with lowest usability and highest mental demand (all > **), and with a higher physical demand than OVERVIEW+DETAIL (**).

Qualitative Feedback

Same as the first study, in addition to quantitative data, we also asked participants to justify their subjective ratings after each task. We analyze the collected feedback to get an overview of the pros and cons of each interaction technique.

FOCUS+CONTEXT was complained to be “difficult for precise selection” and “hard to get where I wanted to be” by 21 participants. 17 participants also considered it to be “disorienting” as “hard to tell where I was.” 11 participants did not feel confident with it, and “had to double check.” Five participants found it “difficult to anticipate the mapping.” In the SEARCH task, eight participants also found using it to “scan a large region is difficult,” which “requires high working memory,” as they need to keep the context in mind. In the CONTEXT task, four participants also commented that the distortion makes it difficult to identify and inspect close clusters.

OVERVIEW+DETAIL was reported to be beneficial for “clearly knowing where you are” by four participants. However, 12 participants also pointed out it “requires a bit of working memory to translate the position in the matrix to the detail view.” Five reported that they “had to double check.” Two participants found it “becomes more difficult for large matrices.” In the SEARCH task, 13 participants found using it to “scan a large region is difficult.” Potentially also because participants need to keep switching between the overview and the detail view all the time. In the CONTEXT task, 11 participants found “having two view at once” helps complete the task.

PAN&ZOOM was found to be “intuitive” and “familiar” by 15 participants. One participant also took advantage of the large number of cells it can enlarge: “there is no need to precisely zoom in.” In the

SEARCH task, 18 participants reported it can show a large number of enlarged cells, and “having all at once” makes this task easy. In the CONTEXT task, 12 participants complained about “the extra physical movements required to zoom in and out.”

5 DISCUSSION

The overarching goal of our studies is to answer the question “Which is the best interaction technique for exploring MMV?”. Our results show that PAN&ZOOM, was as fast as or faster than the overall best performing FOCUS+CONTEXT (i.e., the FISHEYE) and OVERVIEW+DETAIL. Participants also rated PAN&ZOOM the overall best option in terms of usability, and mental and physical demand.

5.1 What leads to different performance for focus+context, overview+detail, and pan&zoom?

Spatial separation of views requires extra time. OVERVIEW+DETAIL was the slowest in the LOCATE and SEARCH tasks. We believe a potential reason is the spatial separation of two views in OVERVIEW+DETAIL. In OVERVIEW+DETAIL, the participants had to interact with the overview, and then inspect the “far away” detail view. In FOCUS+CONTEXT and PAN&ZOOM, the participants have all the information in just one display space. As a result, OVERVIEW+DETAIL was likely to require more eye movements, and potentially introduce extra context-switching cost. Our findings partially align with previous studies [21], where they also found OVERVIEW+DETAIL required more time in some applications.

Spatial separation of views is beneficial for contextualizing details. OVERVIEW+DETAIL was faster than FOCUS+CONTEXT in the CONTEXT task, and tended to be slightly faster than PAN&ZOOM, but not significantly. As mentioned earlier, the CONTEXT task has two components, with the first one similar to the LOCATE task, and the second one in identifying the largest cluster. OVERVIEW+DETAIL was the slowest in the LOCATE task, which means its good performance in the CONTEXT tasks was mainly from identifying the clusters. With OVERVIEW+DETAIL, there is no further interaction needed to finish the second component after the first one. While, for PAN&ZOOM, participants had to zoom out to complete the second component. This is also confirmed by the reported physical demand ratings, where 33 out of 45 participants found PAN&ZOOM required equal (nine participants) or more physical (24 participants) movements than OVERVIEW+DETAIL in the CONTEXT task. On the other hand, compared to OVERVIEW+DETAIL, the distortion in FOCUS+CONTEXT was likely to affect the contextualizing performance, as participants might require extra effort to interpret the distortion. This is also confirmed by the usability, mental demand, and physical demand ratings. In summary, in contextualizing details, the gain of having spatial separation of views outweighs its loss.

More cells showing details lead to better search performance. In the SEARCH task, PAN&ZOOM can show more enlarged cells with line charts. PAN&ZOOM can treat the entire space of MMV as the focal area, as a result, more enlarged cells showing details can fit in the space. In our user study, a maximum of roughly 10×10 cells can be presented with details. On the other hand, FOCUS+CONTEXT and OVERVIEW+DETAIL only represent 3×3 cells with details in the study, and more travels (or scanning) was required to complete this

task. This is also confirmed by the subjective rating of usability, mental demand, and physical demand, where PAN&ZOOM clearly received the best ratings. Increasing the number of cells in detail for FOCUS+CONTEXT and OVERVIEW+DETAIL can potentially improve their performance. For FOCUS+CONTEXT, however, a larger focal area also introduces more distortion. Moreover, it is not possible to have a focal area as large as the entire matrix like in PAN&ZOOM. On the other hand, it can be straightforward to increase the size of detail view in OVERVIEW+DETAIL. However, more screen space will be required, which may not be an option for scenarios with limited screen estate. Another potential reason might be that the majority of users are more familiar with PAN&ZOOM compared to other techniques, as it is a standard interaction technique in many web applications (e.g., Google Maps, photo viewers). Our results partially aligned with Yang et al [87], which found that PAN&ZOOM had better performance than OVERVIEW+DETAIL. Our results differ from the study by Pietriga et al [65], where they found FOCUS+CONTEXT and OVERVIEW+DETAIL outperformed PAN&ZOOM. One possible reason is that they only consider one navigation task, and did not consider multivariate details. Interpreting multivariate details and switching between the focal and contextual areas can introduce additional effort for different conditions.

Distortion results in bad user experience. FOCUS+CONTEXT was rated lowest on usability and highest in mental demand for almost all tasks, despite its generally good performance in the LOCATE and SEARCH tasks. An interesting fact is that the FOCUS+CONTEXT technique used in the second study (i.e., FISHEYE) was rated as the easiest technique in the first study. However, when compared to the regularity and uniformity in OVERVIEW+DETAIL and PAN&ZOOM, participants clearly disliked the FISHEYE. These results were not found in previous studies [9, 35, 74]. A possible explanation is that prior studies employed applications, where the regularity and uniformity are not important. However, keeping rows and columns in a regular grid is critical for matrix visualization and should be considered for the design of MMV.

5.2 Generalization, Limitations and Future Work

Interaction Techniques. In our first study, we followed Carpendale et al's taxonomy [19], and tested four representative lenses. In the second study, we compared the best performing lens (focus+context) to overview+detail, pan&zoom. Those are the most widely-used techniques, and are likely to be among the first choices when designing interactions for MMV. Thus, we believe our selected techniques cover a wide range of interaction techniques for MMV, and provide practical guidance on selecting the most effective and applicable technique. There are other interaction techniques that can be adapted to MMV, like insets [50], editing values, aggregating values across cells and adapting visualizations to the aggregated data [38] and re-ordering the matrix [11]. Our study is meant as a first assessment of the fundamental interactions for MMV. Including these techniques in a future study can obtain a more comprehensive understanding of the effectiveness of interaction techniques in MMV but beyond the scope of this paper.

Our results and discussion can inspire improvements to existing approaches and generate potential new techniques. In the first study, we found correspondence necessary for precise locating.

TableLenses particularly suffered from its low correspondence. One possible way to increase the correspondence in TableLenses is to dynamically move the entire matrix based on the mouse cursor to compensate for its row or column offsets. However, such a design requires extra screen space and may confuse the users. Future design can also consider adapting 3D distortion and techniques, like the Perspective Wall [53] and Mélange [28] to MMV. In the second user study, we found PAN&ZOOM was overall a good option, but OVERVIEW+DETAIL had a similar performance as PAN&ZOOM and was rated lowest on physical demand in the CONTEXT task. To have the benefits of both, the two techniques could be combined. However, adding an overview to a zooming interface leads to mixed results in the literature [39]: some found it useful for navigation, and some found it unnecessary [58, 87].

The performance of FOCUS+CONTEXT was not ideal in our studies. One potential way to improve lenses is to allow the users to select multiple focal areas, which has been explored in some applications [28, 38, 50]. However, as the first controlled study for MMV, we decided to have only one focal area to focus on the basic MMV interactions. Further tests are required to understand the effectiveness of these techniques. Meanwhile, we believe our results can be partially generalized to multi-focal interactions. The tested LOCATE and SEARCH tasks investigated the *wayfinding* and *travel* components. We conjectured that adding the multi-focal feature to our tested conditions would not significantly change our results of these two tasks, as they are basic interactions and do not require participants to investigate multiple areas of interest. Having multiple focus areas is likely to severely change the distortion for the CONTEXT task, and additional investigation is required for checking the *context-switching* component. Moreover, one key motivation of having multiple focus areas is to reduce *the number of travels* [28, 50, 87]. Our tested tasks did not explicitly test *the number of travels* component, and it should be systematically explored in future studies.

Technique Configurations. We discuss the rationale and limitation of the chosen parameters for the tested techniques:

The size of the focal area. We chose the parameters for lenses to ensure that the users can interact with the enlarged cells while the contextual cells are still legible on screens with standard resolutions (Sec. 3.3). With higher resolution, other settings could be tested, and we expect larger focal areas to be beneficial for some tasks (e.g., the SEARCH task). The focal area for OVERVIEW+DETAIL is not as constrained as FOCUS+CONTEXT, but a larger focal area will require more screen space, and we intentionally kept their sizes consistent to reduce confounding factors. One future direction is to investigate the effect of focal area size on different interaction techniques.

Lens on-demand. Providing the ability to switch on and off the lenses is likely to improve their correspondence. However, one key motivation for using lenses is interactive exploration [78, 79], where the location of the targets is not known upfront. Our studies were designed to investigate the performance of interactive exploration and simulate different interaction components of the exploration scenario in the tasks. Additionally, allowing the participants to turn on and off the lenses might bring extra complexities to the interactions, which could potentially affect their performance. Providing extra training might reduce this side effect but significantly increases the user study time. However, experts can get familiar with

enabling/disabling lenses with less time constraint in real-world applications, and its effectiveness should be evaluated.

Dragging interaction in OVERVIEW+DETAIL. There are two ways to select the focal area in OVERVIEW+DETAIL: point-and-click and drag-and-drop [87]. Point-and-click requires fewer steps, while drag-and-drop provides a better estimation of the interaction [46]. It is unclear which is a better choice for MMV. We chose drag-and-drop in our study because the point-and-click method conflicts with our target selection interaction. A future study is desired to compare the effectiveness of these two methods for MMV.

Embedded Visualization and Tasks. In our studies, we tested time series data, one type of widely used multivariate data. Our tested tasks focus on the *interactions* to locate, search, and contextualize multivariate details. These tasks were chosen to investigate the tested conditions' *wayfinding*, *travel*, and *context-switching* performance. We intentionally lower the difficulty of *interpreting* the embedded visualization, so that the participants do not need deep knowledge about a particular type of visualization and can focus on the interactions. Changing the embedded visualization is likely to affect the interpretation performance but will likely bring minimum influence on the *wayfinding*, *travel*, and *context-switching* performance. Thus, we expect our findings on the effectiveness of different interaction techniques to be partially generalized to MMV with other embedded visualizations. Future studies are required to confirm our hypothesis. Horak et al [38] demonstrated embedding different types of visualizations for different cells. Such an adaptive design can facilitate complex data analysis process and should be tested in the future. We also plan to study more specific MMV applications with more sophisticated and high-level tasks.

Scalability. We identified three potential effects related to scalability that were not fully investigated in our studies:

The number of data points and pattern types in the line chart. Inspired by Correll & Gleicher [23], we used five primitive temporal patterns in our study. We did not include more complicated patterns, as we want to focus on studying the performance of different interactions. We also chose to have five data points in each cell, as this is enough for representing all selected temporal patterns while still allowing interaction within the line chart. Interpreting more complicated patterns or increasing the number of points in the line chart is likely to increase the difficulty constantly for all testing conditions. Thus, our findings can still provide helpful guidelines for selecting the appropriate interaction technique. However, further investigations are required to confirm this conjecture.

Size of the matrix. We tested two different sizes of matrices. We believe the tested sizes are representative as they can be reasonably rendered and interacted on a standard screen. We found that the performance of almost all conditions decreased in the larger data set. However, we did not find significant evidence that one specific condition resists the increasing data size better than others. Future studies are required to investigate MMV in larger data sets.

Size of target regions. In the SEARCH task, we used 7×7 as the size of the target regions, which was larger than the size of the lenses, so that participants had to move the lenses to fully explore it. In the CONTEXT task, we controlled the range of cluster sizes (from 5×5 to 7×7) to make the task less obvious and more challenging for participants. We cannot find any literature indicating a significant effect of cluster size, and it should be tested in the future.

6 CONCLUSION

We have presented two studies comparing interaction techniques for exploring MMV. The findings extend our understanding of the different interaction techniques' effectiveness for exploring MMV. Our results suggest that pan&zoom was the overall best performing technique, while for contextualizing details, overview+detail can also be a good choice. We also believe there is potential to improve the design of lenses in MMV, for example, reducing the influence of distortion through lensing on demand. To provide structured guidelines for future research and design, we discussed the effect of correspondence, uniformity, irregularity, and continuity of lenses. Our results indicate that high correspondence, uniformity, and continuity led to better performance for lenses. Future lens design should take these metrics into account. Another potential future direction is to investigate hybrid techniques, such as adding an overview to a zooming interface or providing interactive zooming inside the lenses. In summary, we believe there is much unexplored space in MMV, and our study results and discussion can potentially lead to improved and novel interaction designs in MMV.

ACKNOWLEDGMENTS

This work was partially supported by NSF grants III-2107328 and IIS-1901030, NIH grant 5U54CA225088-03, the Harvard Data Science Initiative, and a Harvard Physical Sciences and Engineering Accelerator Award.

REFERENCES

- [1] Simon Anders. 2009. Visualization of genomic data with the Hilbert curve. *Bioinformatics* 25, 10 (2009), 1231–1235. <https://doi.org/10.1093/bioinformatics/btp152>
- [2] Gennady Andrienko, Natalia Andrienko, Peter Bak, Daniel Keim, Slava Kisilevich, and Stefan Wrobel. 2011. A conceptual framework and taxonomy of techniques for analyzing movement. *Journal of Visual Languages & Computing* 22, 3 (June 2011), 213–232. <https://doi.org/10.1016/j.jvlc.2011.02.003>
- [3] Mark D Apperley, I Tzavaras, and Robert Spence. 1982. A bifocal display technique for data presentation. (1982). <https://doi.org/10.2312/eg.19821002>
- [4] Caroline Appert, Olivier Chapuis, and Emmanuel Pietriga. 2010. High-precision magnification lenses. In *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*. ACM Press, Atlanta, Georgia, USA, 273. <https://doi.org/10.1145/1753326.1753366>
- [5] Benjamin Bach, Pierre Dragicevic, Daniel Archambault, Christophe Hurter, and Sheelagh Carpendale. 2017. A Descriptive Framework for Temporal Data Visualizations Based on Generalized Space-Time Cubes: Generalized Space-Time Cube. *Computer Graphics Forum* 36, 6 (Sept. 2017), 36–61. <https://doi.org/10.1111/cgf.12804>
- [6] Benjamin Bach, Nathalie Henry-Riche, Tim Dwyer, Tara Madhyastha, J-D Fekete, and Thomas Grabowski. 2015. Small MultiPiles: Piling Time to Explore Temporal Patterns in Dynamic Networks. *Computer Graphics Forum* 34, 3 (2015), 31–40. <https://doi.org/10.1111/cgf.12615>
- [7] Benjamin Bach, Emmanuel Pietriga, and Jean-Daniel Fekete. 2014. Visualizing dynamic networks with matrix cubes. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Toronto Ontario Canada, 877–886. <https://doi.org/10.1145/2556288.2557010>
- [8] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting Linear Mixed-Effects Models Using **lme4**. *Journal of Statistical Software* 67, 1 (2015), 47 pages. <https://doi.org/10.18637/jss.v067.i01>
- [9] Patrick Baudisch, Nathaniel Good, Victoria Bellotti, and Pamela Schraedley. 2002. Keeping things in context: a comparative evaluation of focus plus context screens, overviews, and zooming. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. Association for Computing Machinery, New York, NY, USA, 259–266. <https://doi.org/10.1145/503376.503423>
- [10] Fabian Beck, Michael Burch, Stephan Diehl, and Daniel Weiskopf. 2014. The State of the Art in Visualizing Dynamic Graphs. *EuroVis - STARs* (2014), 21 pages. <https://doi.org/10.2312/EUROVISSTAR.20141174>
- [11] Michael Behrisch, Benjamin Bach, Nathalie Henry Riche, Tobias Schreck, and Jean-Daniel Fekete. 2016. Matrix Reordering Methods for Table and Network

- Visualization. *Computer Graphics Forum* 35, 3 (2016), 693–716. <https://doi.org/10.1111/cgf.12935>
- [12] Michael Behrisch, James Davey, Fabian Fischer, Olivier Thonnard, Tobias Schreck, Daniel Keim, and Jörn Kohlhammer. 2014. Visual Analysis of Sets of Heterogeneous Matrices Using Projection-Based Distance Functions and Semantic Zoom. *Computer Graphics Forum* 33, 3 (2014), 411–420.
 - [13] Eric A. Bier, Maureen C. Stone, Ken Pier, William Buxton, and Tony D. DeRose. 1993. Toolglass and magic lenses: the see-through interface. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*. 73–80.
 - [14] Renaud Blanch and Michael Ortega. 2011. Benchmarking pointing techniques with distractors: adding a density factor to Fitts' pointing paradigm. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Vancouver BC Canada, 1629–1638. <https://doi.org/10.1145/1978942.1979180>
 - [15] Carles A. Boix, Benjamin T. James, Yongjin P. Park, Wouter Meuleman, and Manolis Kellis. 2021. Regulatory genomic circuitry of human disease loci by integrative epigenomics. *Nature* 590, 7845 (Feb. 2021), 300–307. <https://doi.org/10.1038/s41586-020-03145-z>
 - [16] Maged N. Kamel Boulos. 2003. The use of interactive graphical maps for browsing medical/health Internet information resources. *International Journal of Health Geographics* 2, 1 (Jan. 2003), 1. <https://doi.org/10.1186/1476-072X-2-1>
 - [17] Michael Burch, Benjamin Schmidt, and Daniel Weiskopf. 2013. A Matrix-Based Visualization for Exploring Dynamic Compound Digraphs. In *2013 17th International Conference on Information Visualisation*. IEEE, London, United Kingdom, 66–73. <https://doi.org/10.1109/IV.2013.8>
 - [18] Stefano Burigat, Luca Chittaro, and Edoardo Parlato. 2008. Map, diagram, and web page navigation on mobile devices: the effectiveness of zoomable user interfaces with overviews. In *Proceedings of the 10th international conference on Human computer interaction with mobile devices and services - MobileHCI '08*. ACM Press, 147. <https://doi.org/10.1145/1409240.1409257>
 - [19] Sheelagh Carpendale, David J Cowperthwaite, and F David Fracchia. 1997. Extending distortion viewing from 2D to 3D. *IEEE Computer Graphics and Applications* 17, 4 (Aug. 1997), 42–51. <https://doi.org/10.1109/38.595268>
 - [20] Richard Chimera. 1998. Value Bars: an information visualization and navigation tool for multi-attribute listings and tables. (Oct. 1998). <https://drum.lib.umd.edu/handle/1903/376>
 - [21] Andy Cockburn, Amy Karlson, and Benjamin B. Bederson. 2009. A review of overview+detail, zooming, and focus+context interfaces. *Comput. Surveys* 41, 1 (Jan. 2009), 2:1–2:31. <https://doi.org/10.1145/1456650.1456652>
 - [22] Michael Correll and Michael Gleicher. 2016. The semantics of sketch: Flexibility in visual query systems for time series data. In *2016 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 131–140.
 - [23] Michael Correll and Michael Gleicher. 2016. The semantics of sketch: Flexibility in visual query systems for time series data. In *2016 IEEE Conference on Visual Analytics Science and Technology (VAST)*. 131–140. <https://doi.org/10.1109/VAST.2016.7883519>
 - [24] Tuan Nhon Dang, Hong Cui, and Angus G Forbes. 2016. MultiLayerMatrix: visualizing large taxonomic datasets. In *EuroVis Workshop on Visual Analytics (EuroVA)*. The Eurographics Association. 6 pages.
 - [25] Geoffrey Ellis, Enrico Bertini, and Alan Dix. 2005. The sampling lens: making sense of saturated visualisations. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems (CHI EA '05)*. Association for Computing Machinery, New York, NY, USA, 1351–1354. <https://doi.org/10.1145/1056808.1056914>
 - [26] Niklas Elmqvist, Thanh-Nghi Do, Howard Goodell, Nathalie Henry, and Jean-Daniel Fekete. 2008. ZAME: Interactive Large-Scale Graph Visualization. In *2008 IEEE Pacific Visualization Symposium*. IEEE, Kyoto, 215–222. <https://doi.org/10.1109/PACIFICVIS.2008.4475479>
 - [27] Niklas Elmqvist, Nathalie Henry, Yann Riche, and Jean-Daniel Fekete. 2008. Melange: space folding for multi-focus interaction. In *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*. ACM Press, Florence, Italy, 1333. <https://doi.org/10.1145/1357054.1357263>
 - [28] Niklas Elmqvist, Yann Riche, Nathalie Henry-Riche, and Jean-Daniel Fekete. 2010. Mélange: Space Folding for Visual Exploration. *IEEE Transactions on Visualization and Computer Graphics* 16, 3 (May 2010), 468–483. <https://doi.org/10.1109/TVCG.2009.86>
 - [29] Andy Field, Jeremy Miles, and Zoë Field. 2012. *Discovering statistics using R*. Sage publications.
 - [30] Maximilian T Fischer, Devanshu Arya, Dirk Streeb, Daniel Seebacher, Daniel A Keim, and Marcel Worring. 2021. Visual Analytics for Temporal Hypergraph Model Exploration. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb. 2021), 550–560. <https://doi.org/10.1109/TVCG.2020.3030408>
 - [31] Paul M. Fitts. 1954. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology* 47, 6 (1954), 381–391. <https://doi.org/10.1037/h0055392>
 - [32] Mohammad Ghoniem, Jean-Daniel Fekete, and Philippe Castagliola. 2005. 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.
 - [33] Sarah Goodwin, Jason Dykes, Aidan Slingsby, and Cagatay Turkay. 2016. Visualizing Multiple Variables Across Scale and Geography. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan. 2016), 599–608. <https://doi.org/10.1109/TVCG.2015.2467199>
 - [34] Carl Gutwin. 2002. Improving focus targeting in interactive fisheye views. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. Association for Computing Machinery, New York, NY, USA, 267–274. <https://doi.org/10.1145/503376.503424>
 - [35] Carl Gutwin and Amy Skopik. 2003. Fisheyes are good for large steering tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. Association for Computing Machinery, New York, NY, USA, 201–208. <https://doi.org/10.1145/642611.642648>
 - [36] Nathalie Henry and Jean-Daniel Fekete. 2007. MatLink: Enhanced Matrix Visualization for Analyzing Social Networks. In *Human-Computer Interaction – INTERACT 2007 (Lecture Notes in Computer Science)*, Cécilia Baranauskas, Philippe Palanque, Julio Abascal, and Simone Diniz Junqueira Barbosa (Eds.). Springer, Berlin, Heidelberg, 288–302. https://doi.org/10.1007/978-3-540-74800-7_24
 - [37] Nathalie Henry, Jean-Daniel Fekete, and Michael J. McGuffin. 2007. NodeTrix: a Hybrid Visualization of Social Networks. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (Nov. 2007). <https://doi.org/10.1109/TVCG.2007.70582>
 - [38] Tom Horak, Philip Berger, Heidrun Schumann, Raimund Dachsel, and Christian Tominski. 2021. Responsive Matrix Cells: A Focus+Context Approach for Exploring and Editing Multivariate Graphs. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb. 2021), 1644–1654. <https://doi.org/10.1109/TVCG.2020.3030371>
 - [39] Kasper Hornbæk, Benjamin B. Bederson, and Catherine Plaisant. 2002. Navigation patterns and usability of zoomable user interfaces with and without an overview. *ACM Transactions on Computer-Human Interaction* 9, 4 (Dec. 2002), 362–389. <https://doi.org/10.1145/586081.586086>
 - [40] Kasper Hornbæk and Erik Frøkjær. 2001. Reading of electronic documents: the usability of linear, fisheye, and overview+detail interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*. Association for Computing Machinery, New York, NY, USA, 293–300. <https://doi.org/10.1145/365024.365118>
 - [41] Petra Isenberg, Sheelagh Carpendale, Anastasia Bezerianos, Nathalie Henry, and Jean-Daniel Fekete. 2009. CoCoNutTrix: Collaborative Retrofitting for Information Visualization. *IEEE Computer Graphics and Applications* 29, 5 (Sept. 2009). <https://doi.org/10.1109/MCG.2009.78>
 - [42] Waqas Javed, Sohaib Ghani, and Niklas Elmqvist. 2012. Polyzoom: multiscale and multifocus exploration in 2d visual spaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 287–296. <https://doi.org/10.1145/2207676.2207716>
 - [43] Thomas Kastner, Karl-Heinz Erb, and Helmut Haberl. 2014. Rapid growth in agricultural trade: effects on global area efficiency and the role of management. *Environmental Research Letters* 9, 3 (March 2014), 034015. <https://doi.org/10.1088/1748-9326/9/3/034015>
 - [44] Peter Kerpedjiev, Nezar Abdennur, Fritz Lekschas, Chuck McCallum, Kasper Dinkla, Hendrik Strobel, Jacob M. Luber, Scott B. Ouellette, Alaleh Azhir, Nikhil Kumar, Jeewon Hwang, Soohyun Lee, Burak H. Alver, Hanspeter Pfister, Leonid A. Mirny, Peter J. Park, and Nils Gehlenborg. 2018. HiGlass: web-based visual exploration and analysis of genome interaction maps. *Genome Biology* 19, 1 (Aug. 2018), 125. <https://doi.org/10.1186/s13059-018-1486-1>
 - [45] Robert Krüger, Dennis Thom, Michael Wörner, Harald Bosch, and Thomas Ertl. 2013. TrajectoryLenses—A Set-based Filtering and Exploration Technique for Long-term Trajectory Data. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 451–460.
 - [46] Harsha P. Kumar, Catherine Plaisant, and Ben Shneiderman. 1997. Browsing hierarchical data with multi-level dynamic queries and pruning. *International Journal of Human-Computer Studies* 46, 1 (Jan. 1997), 103–124. <https://doi.org/10.1006/ijhc.1996.0085>
 - [47] Heidi Lam. 2008. A Framework of Interaction Costs in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (Nov. 2008), 1149–1156. <https://doi.org/10.1109/TVCG.2008.109>
 - [48] Joseph J LaViola Jr, Ernst Kruijff, Ryan P McMahan, Doug Bowman, and Ivan P Poupyrev. 2017. *3D user interfaces: theory and practice*. Addison-Wesley Professional.
 - [49] Fritz Lekschas, Benjamin Bach, Peter Kerpedjiev, Nils Gehlenborg, and Hanspeter Pfister. 2018. HiPiler: Visual Exploration of Large Genome Interaction Matrices with Interactive Small Multiples. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 522–531. <https://doi.org/10.1109/TVCG.2017.2745978>
 - [50] Fritz Lekschas, Michael Behrisch, Benjamin Bach, Peter Kerpedjiev, Nils Gehlenborg, and Hanspeter Pfister. 2020. Pattern-Driven Navigation in 2D Multiscale Visualizations with Scalable Insets. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan. 2020), 611–621. <https://doi.org/10.1109/TVCG.2019.2934555>
 - [51] Fritz Lekschas, Xinyi Zhou, Wei Chen, Nils Gehlenborg, Benjamin Bach, and Hanspeter Pfister. 2021. A Generic Framework and Library for Exploration of Small Multiples through Interactive Piling. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb. 2021), 358–368. <https://doi.org/10.1109/TVCG.2020.3030371>

- 2020.3028948
- [52] Russell V. Lenth. 2016. Least-Squares Means: The R Package **lsmeans**. *Journal of Statistical Software* 69, 1 (2016), 33 pages. <https://doi.org/10.18637/jss.v069.i01>
 - [53] Jock D. Mackinlay, George G. Robertson, and Stuart K. Card. 1991. The perspective wall: detail and context smoothly integrated Mackinlay color plates. In *Proceedings of the SIGCHI conference on Human factors in computing systems Reaching through technology - CHI '91*. ACM Press, New Orleans, Louisiana, United States, 173–176. <https://doi.org/10.1145/108844.108870>
 - [54] Michael McGuffin and Ravin Balakrishnan. 2002. Acquisition of expanding targets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. Association for Computing Machinery, New York, NY, USA, 57–64. <https://doi.org/10.1145/503376.503388>
 - [55] Peter McLachlan, Tamara Munzner, Eleftherios Koutsofios, and Stephen North. 2008. LiveRAC: interactive visual exploration of system management time-series data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. Association for Computing Machinery, New York, NY, USA, 1483–1492. <https://doi.org/10.1145/1357054.1357286>
 - [56] Graham McNeill and Scott A. Hale. 2017. Generating Tile Maps. *Computer Graphics Forum* 36, 3 (June 2017), 435–445. <https://doi.org/10.1111/cgf.13200>
 - [57] Tamara Munzner. 2014. *Visualization analysis and design*. CRC press.
 - [58] Dmitry Nekrasovski, Adam Bodnar, Joanna McGrenere, François Guimbretière, and Tamara Munzner. 2006. An evaluation of pan & zoom and rubber sheet navigation with and without an overview. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. Association for Computing Machinery, New York, NY, USA, 11–20. <https://doi.org/10.1145/1124772.1124775>
 - [59] Mario Popolin Neto and Fernando V. Paulovich. 2021. Explainable Matrix - Visualization for Global and Local Interpretability of Random Forest Classification Ensembles. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb. 2021), 1427–1437. <https://doi.org/10.1109/TVCG.2020.3030354>
 - [60] Christina Niederer, Holger Stitz, Reem Hourieh, Florian Grassinger, Wolfgang Aigner, and Marc Streit. 2017. TACO: visualizing changes in tables over time. *IEEE transactions on visualization and computer graphics* 24, 1 (2017), 677–686.
 - [61] Niels Christian Nilsson, Stefania Serafin, Frank Steinicke, and Rolf Nordahl. 2018. Natural Walking in Virtual Reality: A Review. *Computers in Entertainment* 16, 2 (April 2018), 1–22. <https://doi.org/10.1145/3180658>
 - [62] Carolina Nobre, Miriah Meyer, Marc Streit, and Alexander Lex. 2019. The State of the Art in Visualizing Multivariate Networks. *Computer Graphics Forum* 38, 3 (2019), 807–832. <https://doi.org/10.1111/cgf.13728>
 - [63] Adam Pearce. 2020. Communicating Model Uncertainty Over Space. <https://pair-code.github.io/interpretability/uncertainty-over-space/>. <https://pair-code.github.io/interpretability/uncertainty-over-space/>
 - [64] Emmanuel Pietriga and Caroline Appert. 2008. Sigma lenses: focus-context transitions combining space, time and translucence. In *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*. ACM Press, Florence, Italy, 1343. <https://doi.org/10.1145/1357054.1357264>
 - [65] Emmanuel Pietriga, Caroline Appert, and Michel Beaudouin-Lafon. 2007. Pointing and beyond: an operationalization and preliminary evaluation of multi-scale searching. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. Association for Computing Machinery, New York, NY, USA, 1215–1224. <https://doi.org/10.1145/1240624.1240808>
 - [66] Matthew Plumlee and Colin Ware. 2002. Zooming, multiple windows, and visual working memory. In *Proceedings of the Working Conference on Advanced Visual Interfaces - AVI '02*. ACM Press, Trento, Italy, 59. <https://doi.org/10.1145/1556262.1556270>
 - [67] Matthew D. Plumlee and Colin Ware. 2006. Zooming versus multiple window interfaces: Cognitive costs of visual comparisons. *ACM Transactions on Computer-Human Interaction (TOCHI)* 13, 2 (June 2006), 179–209. <https://doi.org/10.1145/1165734.1165736>
 - [68] Ramana Rao and Stuart K. Card. 1994. The table lens: merging graphical and symbolic representations in an interactive focus + context visualization for tabular information. In *Proceedings of the SIGCHI conference on Human factors in computing systems celebrating interdependence - CHI '94*. ACM Press, Boston, Massachusetts, United States, 318–322. <https://doi.org/10.1145/191666.191776>
 - [69] Jonathan C Roberts. 2007. State of the Art: Coordinated Multiple Views in Exploratory Visualization. In *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV 2007)*. 61–71. <https://doi.org/10.1109/CMV.2007.20>
 - [70] George G. Robertson and Jock D. Mackinlay. 1993. The document lens. In *Proceedings of the 6th annual ACM symposium on User interface software and technology - UIST '93*. ACM Press, Atlanta, Georgia, United States, 101–108. <https://doi.org/10.1145/168642.168652>
 - [71] Mikkel Rønne Jakobsen and Kasper Hornbæk. 2011. Sizing up visualizations: effects of display size in focus+context, overview+detail, and zooming interfaces. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*. ACM Press, Vancouver, BC, Canada, 1451. <https://doi.org/10.1145/1978942.1979156>
 - [72] Ramik Sadana, Timothy Major, Alistair Dove, and John Stasko. 2014. Onset: A visualization technique for large-scale binary set data. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 1993–2002.
 - [73] Manojit Sarkar and Marc H. Brown. 1992. Graphical fisheye views of graphs. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '92*. ACM Press, Monterey, California, United States, 83–91. <https://doi.org/10.1145/142750.142763>
 - [74] Garth Shoemaker and Carl Gutwin. 2007. Supporting multi-point interaction in visual workspaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. Association for Computing Machinery, New York, NY, USA, 999–1008. <https://doi.org/10.1145/1240624.1240777>
 - [75] Harri Siirtola. 1999. *Interaction with the Reorderable Matrix*.
 - [76] R William Soukoreff and I Scott MacKenzie. 2004. Towards a standard for pointing device evaluation, perspectives on 27 years of Fitts' law research in HCI. *International journal of human-computer studies* 61, 6 (2004), 751–789.
 - [77] Stefano Burigat and Luca Chittaro. 2013. On the effectiveness of Overview+Detail visualization on mobile devices. *Personal and Ubiquitous Computing* 17, 2 (2013), 371–385. <https://doi.org/10.1007/s00779-011-0500-3>
 - [78] Christian Tominski, Stefan Gladisch, Ulrike Kister, Raimund Dachselt, and Hei-drun Schumann. 2014. A Survey on Interactive Lenses in Visualization. *EuroVis - STARs* (2014), 20 pages. <https://doi.org/10.2312/EUROVISSTAR.20141172>
 - [79] Christian Tominski, Stefan Gladisch, Ulrike Kister, Raimund Dachselt, and Hei-drun Schumann. 2017. Interactive Lenses for Visualization: An Extended Survey. *Computer Graphics Forum* 36, 6 (2017), 173–200. <https://doi.org/10.1111/cgf.12871>
 - [80] Jarke J Van Wijk and Wim AA Nuij. 2003. Smooth and efficient zooming and panning. In *IEEE Symposium on Information Visualization 2003 (IEEE Cat. No.03TH8714)*. 15–23. <https://doi.org/10.1109/INFVIS.2003.1249004>
 - [81] Athanasios Vogogias, Daniel Archambault, Benjamin Bach, and Jessie Kennedy. 2020. Visual Encodings for Networks with Multiple Edge Types. In *International Conference on Advanced Visual Interfaces 2020*. 9.
 - [82] Michelle Q. Wang Baldonado, Allison Woodruff, and Allan Kuchinsky. 2000. Guidelines for using multiple views in information visualization. In *Proceedings of the working conference on Advanced visual interfaces - AVI '00*. ACM Press, Palermo, Italy, 110–119. <https://doi.org/10.1145/345513.345271>
 - [83] Jo Wood, Aidan Slingsby, and Jason Dykes. 2011. Visualizing the Dynamics of London's Bicycle-Hire Scheme. *Cartographica: The International Journal for Geographic Information and Geovisualization* 46, 4 (Nov. 2011), 239–251. <https://doi.org/10.3138/carto.46.4.239>
 - [84] Linda Woodburn, Yalong Yang, and Kim Marriott. 2019. Interactive visualisation of hierarchical quantitative data: an evaluation. In *2019 IEEE Visualization Conference (VIS)*. IEEE, 96–100. <https://doi.org/10.1109/VISUAL.2019.8933545>
 - [85] Yan Xu, Zhipeng Jia, Liang-Bo Wang, Yuqing Ai, Fang Zhang, Maode Lai, I Eric, and Chao Chang. 2017. Large scale tissue histopathology image classification, segmentation, and visualization via deep convolutional activation features. *BMC bioinformatics* 18, 1 (2017), 1–17.
 - [86] Yalong Yang, Tim Dwyer, Sarah Goodwin, and Kim Marriott. 2017. Many-to-Many Geographically-Embedded Flow Visualisation: An Evaluation. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 411–420. <https://doi.org/10.1109/tvcg.2016.2598885>
 - [87] Yalong Yang, Maxime Cordeil, Johanna Beyer, Tim Dwyer, Kim Marriott, and Hanspeter Pfister. 2021. Embodied Navigation in Immersive Abstract Data Visualization: Is Overview+Detail or Zooming Better for 3D Scatterplots? *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb. 2021), 1214–1224. <https://doi.org/10.1109/TVCG.2020.3030427>
 - [88] Andrew Yates, Amy Webb, Michael Sharpnack, Helen Chamberlin, Kun Huang, and Raghu Machiraju. 2014. Visualizing Multidimensional Data with Glyph SPLOMs: Visualizing Multidimensional Data with Glyph SPLOMs. *Computer Graphics Forum* 33, 3 (June 2014), 301–310. <https://doi.org/10.1111/cgf.12386>
 - [89] Ji Soo Yi, Niklas Elmqvist, and Seungyoon Lee. 2010. TimeMatrix: Analyzing Temporal Social Networks Using Interactive Matrix-Based Visualizations. *International Journal of Human-Computer Interaction* 26, 11–12 (Nov. 2010), 1031–1051. <https://doi.org/10.1080/10447318.2010.516722>
 - [90] Veronica Zammito. 2008. Visualization Techniques In Video Games. In *Electronic Visualisation and the Arts*. 267–276. <https://doi.org/10.14236/ewic/EVA2008.30>
 - [91] Ana Zanella, Sheelagh Carpendale, and Michael Rounding. 2002. On the effects of viewing cues in comprehending distortions. In *Proceedings of the second Nordic conference on Human-computer interaction (NordCHI '02)*. Association for Computing Machinery, New York, NY, USA, 119–128. <https://doi.org/10.1145/572020.572035>