Beyond Horizon Graphs : Space Efficient Time Series Visualization with Composite Visual Mapping

Ali Jabbari, Renaud Blanch, Sophie Dupuy-Chessa

Beyond Horizon Graphs : Space Efficient Time Series Visualization with Composite Visual Mapping
Encodage visuel composite pour la visualization compacte des series temporelles

Ali Jabbari
Univ. Grenoble Alpes
LIG
F-38000, Grenoble, France
ali.jabbari@univ-grenoble-alpes.fr

Renaud Blanch
Univ. Grenoble Alpes
LIG
F-38000, Grenoble, France
renaud.blanch@imag.fr

Sophie Dupuy-Chessa
Univ. Grenoble Alpes
LIG
F-38000, Grenoble, France
sophie.dupuy@imag.fr

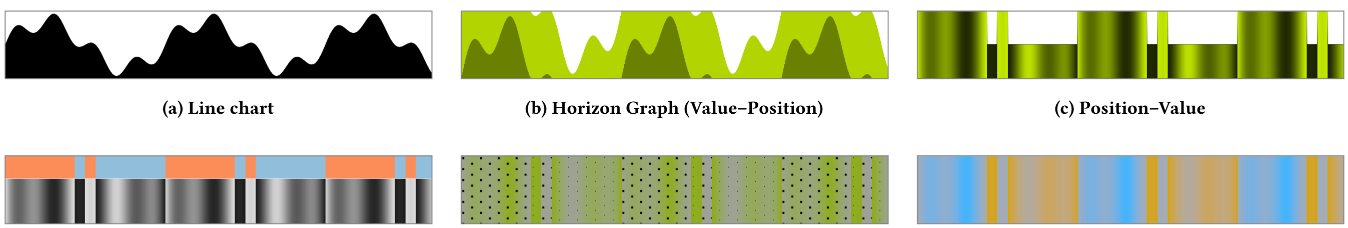

(a) Line chart **(b) Horizon Graph (Value–Position)** **(c) Position–Value**

(d) Hue–Value **(e) Texture–Saturation** **(f) Hue–Saturation**

Figure 1: Complete selection of visual mappings examined in this study displaying a dummy data set. All mappings in the first row use one geometric visual channel. Mappings in the second row exclusively use optical channels. Except line chart which was included as the reference point, all other techniques use composite visual mapping with modulo division as the transformation function.

ABSTRACT

Restricted screen space is a limit to visualization of time series regardless of the medium. To address this challenge, we introduce new space-efficient visual designs for time series based on an approach similar to the well established Horizon Graph, namely "composite" visual mapping. In this approach, each data attribute is decomposed into two components and then each component is mapped onto a separate visual channel. Our visual designs consist in different combinations of geometric and optical visual channels. We compare our propositions with Horizon Graph across different chart heights and we measure accuracy and speed of users in a discrimination and estimation task. Our results show that although Horizon Graph perform best at larger chart heights, our propositions demonstrate same levels of accuracy in small chart heights. Moreover, at least one of our propositions has significant advantage over Horizon Graphs in terms of speed. Based on our findings, we propose design guidelines for using composite visual mapping and combinations of optical visual channels in limited vertical screen resolution settings.

CCS CONCEPTS

• **Human-centered computing → Visualization design and evaluation methods**; *Empirical studies in visualization*;

KEYWORDS

information visualization, visual mapping, time series

RÉSUMÉ

Les dimensions restreintes d'écran posent une limite à la visualisation des séries temporelles, quel que soit le support. Pour relever ce défi, nous introduisons de nouvelles conceptions visuelles peu encombrantes pour des séries temporelles. Avec cette approche de mapping dite "composite", chaque attribut de données est décomposé en deux composants, puis chaque composant est mappé sur un canal visuel distinct. Nos conceptions visuelles consistent en différentes combinaisons de canaux visuels géométriques et optiques. Nous comparons nos propositions avec le fameux Horizon Graph à travers différentes hauteurs de diagrammes et nous mesurons la précision et la vitesse des utilisateurs dans les tâches de discrimination et d'estimation. Nos résultats montrent que même si Horizon Graph entraîne meilleurs performances à des hauteurs de diagramme plus grandes, nos propositions démontrent les mêmes niveaux de précision dans les petites hauteurs de diagramme. De plus, au moins une de nos propositions a un avantage significatif sur Horizon Graphs en termes de vitesse. Basé sur nos résultats, nous proposons des consignes pour l'utilisation du mapping visuel composite et des combinaisons de canaux visuels optiques dans les cas de résolution verticale limitée.

MOTS-CLEFS

visualization d'information, encodage visuel, series temporelles

1 INTRODUCTION

A time series is a set of observations (i.e., events or measured values) indexed in temporal order. Time series are widely used in science, engineering, finance, and industrial contexts. Visualization of the data collected in this form is a vital tool for experts in many domains (e.g., climatology, finance, production monitoring) for making decisions, detecting events and abnormalities, and forecasting a future event or trend.

Often, the visualized data are displayed in a visualization dashboard on a work station's screen. In many scenarios involving visualization of multidimensional data (e.g., simultaneous monitoring of tens or hundreds of production processes), the human analyst is interested in viewing and comparing a multitude of time series at the same time on a limited screen size. Besides stationary dashboards, handheld devices (e.g., touch screen tablets) have emerged in recent years as a new medium for visualization which offers more mobility at the expense of less screen real estate. One common challenge at the core of both situations (stationary dashboards with multiple time series, and mobile devices) is the limit of the available screen space in terms of exploitable pixels. In spite of the previous work on space filling techniques, space efficient visualization of time series still remains a current challenge due to the ever increasing dimensionality and volume of data,.

The visual designs proposed in this work implement an alternative approach to visual mapping, namely *composite visual mapping* [17]. In such visual mappings, data values are decomposed into sub-values which can communicate additional information on different facets of data. Previous studies on Horizon Graphs have shown the advantages of such visual mapping over more familiar line charts [16].

In an earlier systematic study of composite visual mappings [17], we demonstrated that other composite visual mappings can roughly attain the accuracy levels similar to those of Horizon Graphs. However, studies on Horizon Graphs have shown that such *size-related* visual mapping which use geometric visual channels (e.g., size and position) are heavily impacted by the vertical size reduction [16].

In this work, we are interested in using composite visual mapping for visualization of time series data in small physical sizes. Tufte advises for charts with greater length than height because it helps elaborate the workings of the causal variable (time) in more detail [30]. Hence, it is quite common to reduce the height of a chart while preserving the same length (temporal window), in order to increase the space filling properties of the visualization and thus, fit more data (e.g., more data dimensions) in the same space. However, vertical reduction of chart height has not the same effects on all visualization techniques. For instance, Cleveland [8] showed how the aspect ratio of a line chart can impact user's perception of trends in data.

We hypothesize that the degree to which a visualization technique is impacted by reduction of its height is at least partially related to its underlying visual mapping. While we acknowledge the incredible strengths of size and position as visual channels, we hypothesize that size-independent visual mappings can make up for shortcomings of size-dependent mappings in small heights. To verify this point, we propose composite visual mappings which use combinations of optical visual channels. We compare them with the existing geometric mappings, i.e., Horizon Graphs and line charts across different chart heights. Through several user experiments, we show that our visual designs achieve the same levels of accuracy in small sizes and in some cases, can improve the speed of the users.

The structure of this paper is as follows : in Section 2 we explain the theoretical context of our study and we present the related work on space-efficient visualization of time series with a focus on visual mapping. In Section 3, we present our visual designs and some other possible variations and we justify our design choices. In Section 4, we explain our experimental approach for a set of experiments that we describe in Section 5 and 6. Based on the results from our experiments, we propose design guidelines in Section 7, and finally we point to the limitations and perspectives of our work in Section 8.

2 BACKGROUND AND RELATED WORK

In this section, we first discuss the theoretical context of the current study and our approach to visual mapping. Later, we review the literature with a focus on existing visualization techniques that are closely related to this work.

2.1 Theoretical Context

Screen size is a visualization criterion that is directly correlated with the amount of data and thus, the information that can be displayed at once [32]. This makes screen size a physical limit to the amount of information that can be communicated through a visualization and marks the importance of space efficiency in visualization techniques.

In previous works, several approaches have been proposed for improving space efficiency and reducing visual clutter. Ellis et al. [12] classify these approaches into three groups :

- **Appearance :** This group includes techniques that control how much data appear on the screen by sampling (e.g. [11]), filtering (e.g. [1]) changing marks' size or opacity (e.g. [3]), and clustering (e.g., [14]) ;
- **Spatial distortion :** This can be achieved by displacing marks (e.g. [20]), topological distortion (e.g., zooming and Fish-eye lenses), space-filling (e.g., Tree-Maps [29]), and pixel-oriented techniques (e.g., [19]) ;
- **Temporal :** This category refers to animation techniques (e.g., [10]) which display different data within a temporal interval.

Each of these approaches has its benefits and compromises. We acknowledge that evaluation of all aforementioned techniques is out of the scope of this study. Instead, we are interested in space efficient designs for time series visualizations with the most familiar form i.e., horizontal charts and we focus our work on improvements on visual mapping of data by means of combinations of visual channels. In this regard, our approach is more close to space-filling and pixel-oriented techniques.

Visual Channels

In *Semiology of Graphics* [4], Bertin defines the original visual channels as position, size, shape, value (or luminosity), color, orientation, and texture. This list was later extended by Mackinley [22] to include 13 channels and it has been extended ever since. Chen et al. [15] present a categorization of visual channels which include more than 30 channels in four groups :

Geometric Channels allow to distinguish visual structures by their geometric properties. They include :

- Size/length/width/height/depth/thickness/area/volume
- Orientation/angle
- Shape
- Curvature
- Smoothness

Optical Channels allow to distinguish visual structures by means of different optical effects, many of which depend on the observer's perceptual interpretation. This category includes :

- Color/hue/chroma/saturation
- Intensity/brightness/value/luminosity
- Transparency/opacity
- Surface texture/line style
- Lighting effects/halos/shading/shadow
- Optical effects such as focus, blurring and distortion
- Implicit motion/motion blur patterns
- Explicit motion/animation

There are also **Relational Channels** which communicate relationships between visual structures, i.e., spatial and topological relationships and **Semantic Channels** which make use of visual languages to communicate certain concepts.

Visual Mapping

In the visualization reference model proposed by Card et al. [6], visual mapping is the process in which focus data attributes are linked to visual structures. Appropriate visual mapping is thus critical to the effectiveness and expressiveness of any visualization. In its conventional form, visual mapping consists in associating each data attribute to a separate visual channel. Familiar time series visualization techniques, e.g., line charts, rely mostly on geometric channels to encode information. A few earlier works have examined the use of optical channels. For instance, Lam et al. [21] implement a redundant mapping of time series values on color in compact low resolution view, and on both position and color in higher resolution. Such a mapping was intended to attract users' attention to high and low values of data, especially in small vertical resolutions. In another work [24], McLachlan et al. implement a color scale for indicating threshold levels. Filling visualization cells with such color scales allowed users of their system to easily identify the series above the threshold, especially in compact views.

Meanwhile, studies in psychology, specifically the selective attention theory, assert that the human visual system is capable of concentrating on one visual stimulus of an image (e.g., height of a bar) while ignoring all other stimuli (e.g., the color filling the bar) [26]. In cartography, studies on map symbols composed of two separate visual channels have confirmed that users can perceive the information mapped on one visual channel while effectively filtering the information mapped onto the other visual channel [25].

In the scientific and information visualization domain also, many techniques benefit from the combination of visual channels. Chen et al. [7] have done a survey of visual *multiplexing*, or combination of visual channels, in the visualization domain and suggest it as a common phenomenon in visualization that provides an inherent means to access under-utilized screen space. By reporting evidence from psychology and implemented visualization techniques, Chen et al. conclude that human visual system can effortlessly decode many combinations of visual channels.

The visual designs proposed in this study use *composite visual mapping* in which, each data attribute is decomposed into two components and then, each component is mapped onto a separate visual channel. Such an approach to visual mapping allows to separate different aspects of data and hence, increase the extracted information. This can be achieved by different decomposition functions. For example, modulo operation separates the information of the order of magnitude from smaller value fluctuations. Or, frequency separation can extract trend, seasonality, and noise components of the data.

A few existing visualization techniques use this approach to visual mapping. In an earlier work [17], we systematically explored possible combinations of visual channels. By choosing appropriate combinations of visual channels, this approach makes it possible to communicate separable information on multiple facets of data. In this way, it allows the user to focus on a desired aspect of data while ignoring other aspects. In addition, this method reduces overuse of a singe visual channel (e.g., due to noisy data) or underuse of a visual channel (e.g., due to small range of value changes in that visual channel).

2.2 Existing Techniques

Here, we present some of the existing time series visualization techniques that are closely related to our study and composite visual mapping. Based on the strengths and limitations of these techniques, we suggest our visual designs.

Line charts. The line chart is the most common time series visualization which uses vertical position as the visual channel for encoding values and horizontal position for encoding time. A line chart is constructed by connecting the points successively sampled over time. Often, the area below the line is filled in order to increase contrast.

Displaying multiple line charts in a shared space can introduce visual clutter and make it difficult to distinguish different time series. Several techniques have been proposed for packing multiple time series in a constrained space. Small Multiples [30] reduces the visual clutter by splitting the space into individual charts, one for each time series. However, this technique reduces chart heights by dividing the vertical space. On the other hand, Stacked Graphs [5] and Braided Graphs [18] use shared space to display multiple time series. The main problem with these techniques is that comparison between time series becomes very difficult as the number of series increases.

Horizon Graphs. Horizon Graphs [28] are one of the best known compact visualization techniques for time series. Horizon Graphs combine the compactness of color mapping with spatial resolution in order to preserve details. They reduce space by dividing the chart into several bands and superposing them onto each other to create a layered form.

Horizon Graphs implement a composite visual mapping where modulo transformation $v = sgn(v).(a \times M + b)$ decomposes each data attribute to a sign of the data attribute's value (negative or positive), quotient a and reminder b. Each resulting component is then respectively mapped onto hue, saturation, and vertical position (Figure 1b).

The increase in data density enables the visualization system to display more data in a limited space. However, the deformations may obscure some patterns in data and the mental unstacking of layered charts may involve cognitive overload in some cases. Heer et al. [16] showed that Horizon Graphs are more accurate than line charts in discrimination and estimation tasks. On the other hand, they demonstrated for both Horizon Graphs and line charts that the estimation error increases as chart height decreases.

Later, Perin et al. [27] introduced *Interactive Horizon Graphs* in which, they implement pan and zoom interactions to achieve an interactive baseline. More recently, Federico et al. [13] introduced *Qualizaon graphs*, a compact time series visualization that is based on Horizon Graphs. While it inherently integrates qualitative abstractions, it does not improve space efficiency further than Horizon Graph.

Given the strengths, limitations, and the general appeal of Horizon Graphs in space-efficient visualization of time series, we set Horizon Graphs as the comparison point and reference of our study.

3 EVALUATED VISUAL DESIGNS

Our initial selection of visual mappings consisted in six visualization techniques (see Figure 1) from which, line charts and Horizon Graphs were adopted from the existing techniques. In the other four techniques in form of [Visual Channel 1]–[Visual Channel 2], we used composite visual mapping where the quotient of the modulo division was mapped onto [Visual Channel 1], and the remainder onto [Visual Channel 2] similar to our reference point, the Horizon Graphs.

(a) **Line chart :** We included the line chart as the most basic and familiar time series visualization technique. As suggested by Tufte [30], a shaded, high contrast line chart might be better that a floating thin line.

(b) Horizon Graph : As the most relevant technique to our approach, Horizon Graphs are included in our list as the reference point. Many coloring schemes are conceivable for differentiating layers in Horizon Graphs. Hue variation for different layers can suffer from the lack of natural ordering of hues. We adopted the more common color value sequencing due to its better contrast than saturation sequencing. In this way, Horizon Graph uses a composite mapping of Value–Position. We selected the final configuration by conducting a pilot test described in Section 5.

(c) Position–Value : This mapping actually is the reverse of the composite mapping in Horizon Graphs. This design was included to examine the effect of the order of the visual channels in a composite mapping.

(d) Double-banded (Hue–Value) : We are aware of interactions between colors and graphical elements and their effects on visual perception [2]. Hence, this design spatially separates the two visual channels into two dedicated bands. Continuous values of the remainder of modulo division naturally need more spatial resolution than discrete values of the order of magnitude. Hence, we have encoded the order of magnitude on the upper band, which has half the width of the lower band for smaller fluctuations. The upper band can be filed with color value sequencing or discrete hues with equal color value resulting in Value–Value and Hue-Value composite mappings, respectively. We chose the final design among other alternatives after running a pilot test detailed in Section 5.

(e) Texture-Saturation : We decided to include this design in our list because of the particularity of the texture channel. Although considered as an optical channel, its impact is directly related to the geometric properties of its motifs. In our opinion, this positions texture on the borderline between optical and geometric visual channels and worth studying for spatial efficiency.

(f) Hue–Saturation : There is no intuitive order of hues that is acceptable to everyone. Therefore, continuous variation of hue ("the rainbow") should be avoided for quantitative data. In discrete quantities (e.g., order of magnitude) the effect is more tolerable but still is context and user dependent. For example, consider two green and red hues used in a unspecified context. While some users may associate green to higher values and red to lower values (as in battery charge indicators), others may attribute red to values exceeding a threshold, and greens to lower values below that limit. Also, users with color deficiency may misinterpret the data mapped on hue. We used a blue-orange color scheme for this design to avoide these effects.

Among the visual designs detailed above, line charts, Horizon Graphs, and Position–Saturation rely on geometric visual channels (i.e., position/size), Hue–Saturation and the Double-banded design rely solely on optical channels and finally, we consider Texture-Saturation on the borderline of the two categories. This varied selection of the visual designs was intended to to examine the effect of the type of visual channels in space-constrained visualizations.

4 APPROACH AND METHODS

We conducted a series of experiments on our proposed visualization designs with several objectives in mind. First, we wanted to reproduce one finding of prior studies on Horizon Graphs [16] which suggests that Horizon Graphs perform better than simple line charts as the most basic form of time series visualization. We had a number of hypotheses :

H1 Horizon Graph performs better than simple line chart in discrimination and estimation task thanks to composite mapping (reproduction of previous studies).

H2 Performance of Horizon Graph and other mappings which use geometric channels is impacted by the increase in chart height (Probably due to spatial effects i.e., physical size, banking angle, etc.).

H3 In small chart heights, our designs which rely on optical channels outperform Horizon Graph in terms of accuracy and speed.

4.1 Experimental Procedure

We intended to compare our visual designs with Horizon Graphs as the reference point for composite mapping, and reproduce some findings in the study by Heer et al. [16]. Therefore, we decided to replicate their experimental protocol. Regardless of the experiment, at each trial the participant viewed a pair of charts. The two charts displayed two distinct time series using identical visual design. At fixed positions, one chart was marked with T and the other with B (see Figure 2). At each trial, participants first performed a discrimination task in which they were asked to report whether T or B points to a larger value. The time elapsed before a participant selected the corresponding radio button was registered as discrimination time. Participants could change their answers and the extra discrimination time was taken into account. Next, participants performed an estimation task where they were asked to estimate the absolute value difference between the two points. They gave their estimations using a continuous slider without tick marks in order to avoid anchoring effects [23]. Once the participant clicked on the "Next" button, the time elapsed since the discrimination task is recorded as the estimation time. At this point, a new trial replaces the old one and participants could no longer go back to review their answers.

Each experiment was divided into several blocks. The order of trials in each block was randomized using the Fisher-Yates shuffle algorithm [9] to ensure an unbiased permutation. Participants could take a break between the blocks and resume the experiment when they desired. We instructed the participants to keep the same distance from the screen for all trials.

The experimental interface was created with JavaScript and all experiments were run in Google Chrome web browser version 62.0. We used a 27-inch Apple iMac at its native 5120×2880 pixel resolution with its brightness set on maximum and we conducted all experiments in a room with constant ceiling lighting and no sunlight. We will describe the number of participants and trials for each experiment in its corresponding section.

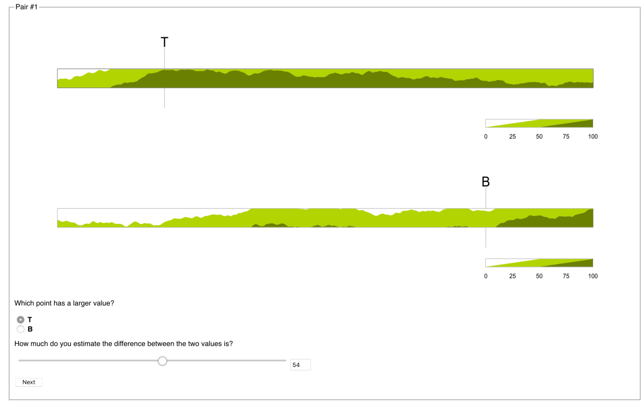

Figure 2: The experimental interface used through this study. Here, a sample trial with Horizon Graphs.

4.2 Chart Generation

In all experiments, each chart represented a distinct time series generated by running a moving average smoothing over a constrained random walk. We used constrained random walks to control the distribution of start/end values of the series, as well as the markers' values. We admit that this type of data encompasses only a fraction of common time series and one may imagine other types of time series, e.g., with bursty or noisy characteristics. However, we believe that smoothed random walks are a relevant type of test data for the decomposition function that we have chosen (modulo division) based on the tasks that we evaluate (discrimination and estimation). However, other types of data (e.g., noisy data) can be used for more relevant tasks (e.g., trend detection) with different decompositions (e.g., frequency separation).

In all time series, values varied between 0 and 100. The charts were generated on HTML5 Canvas elements using JavaScript. In an attempt to ensure perceptual uniformity, we chose CIELUV color space for chart generation.

5 PILOT STUDIES

Prior to our main experiment, we conducted a series of preliminary studies in order to further refine our proposed visualization techniques. We are aware of the statistical weakness of the results from such small samples, yet we consider them as an objective aid to improve our final experimental design.

5.1 Pilot I : Mapping Selection

For statistical reasons, we had to reduce the number of variables in our experimental design. Toward this aim, in a first pilot study we tried to identify the most pertinent composite mappings among the six initial visual mappings in Figure 1. This experiment also allowed us to compare composite mappings with line chart as the reference point. All charts had the same length of 2000 pixels and we tested four chart heights (128, 64, 32, and16 pixels). At 128 pixels height, the chart measured 233×15 mm on the screen. These sizes correspond to the sizes used in the earlier study on Horizon Graphs [16].

We conducted this experiment with 5 participants (4 male and 1 female) with the experimental protocol described in Section 4.1. The experiment consisted in 120 trials per subject (6 mappings × 4 sizes × 5 trial per condition), divided into 5 experimental blocks.

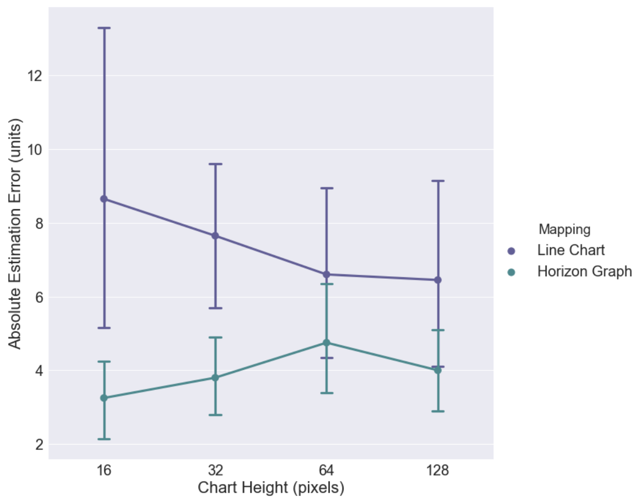

Figure 3: Absolute estimation error for line charts and Horizon Graphs in pilot I. Brackets show confidence intervals at 95 percent.

For all visual designs, discrimination accuracy averaged above 95%, so we focused on estimation accuracy rates. Among our four candidate designs, Hue-Sat had the highest levels of estimation error, while showing no advantage in terms of speed. Based on these results (although not statistically significant) and the discussions with participants, we decided to eliminate this mapping in the main experiment.

Comparing line charts and Horizon Graphs (Figure 3), we observed one of the effects already reported in [16]. As in the preceding study, accuracy decreased at smaller chart heights for both line charts and Horizon Graphs. However, this effect was most prominent for line charts. Given this observation but due to the small sample of results, we cautiously consider our first hypothesis **(H1)** as valid and we exclude line charts from our final experiment and keep Horizon Graph as the reference point.

5.2 Pilot II : Color Scheme Selection

We conducted the second pilot experiment to choose the color schemes for Horizon Graphs and our double-banded design. We chose shades of green (varying color value) as a candidate color scheme in addition to monochromic (no hue effect and maximum color contrast) for Horizon Graphs and the double-banded design. We also included a third color scheme with red and blue hues (with constant and equal color value) for the double-banded design (the upper band) resulting in a Hue-Value composite mapping.

We conducted a pilot test on the 5 color schemes (2 for Horizon Graphs and 3 for the double-banded) with 3 participants that have not had participated in the previous pilot. We tested four chart heights (1, 1/2, 1/4, and 1/8 scale factors) where the tallest chart corresponded to 128 pixels and physical size of 15 millimeters on our screen. In Horizon Graphs, we saw no difference between green gradients and chromatic in terms of accuracy. This can be explained by the fact that with both color schemes, users solely relied on the (equal) difference of color value and presence of hue had a minimal role. Yet, in the discussions with participants they reported their preference for colored Horizon Graphs over the chromatic version. Thus, we selected green color scheme for Horizon Graphs in the main experiment.

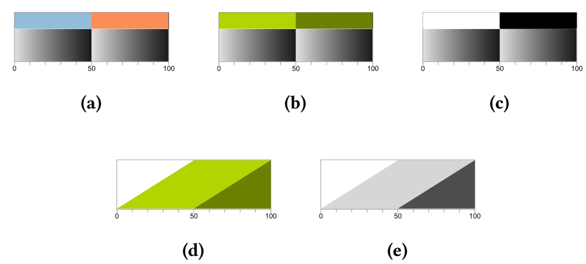

Figure 4: Different color schemes tested in pilot II. We finally selected (a) for the double-banded design and (d) for Horizon Graphs in the main experiment.

Regarding the double-banded design, although the accuracy was similar for all color schemes, participants reported that they found it extremely fast and easy to detect red zones (and therefore, higher values) without confusing the order of magnitude of the two colors. Interestingly, some participants found chromatic double-banded design confusing, as they considered darker regions corresponding to smaller values (i.e., "off" pixels) and lighter regions to larger values (i.e., "on" pixels). Therefore, we selected the double-banded design with blue/red as the final design for the main experiment.

5.3 Pilot III : Texture Selection

In order to finalize our Texture–Saturation design, we conducted a pilot within three candidate textures (see Figure 5). For space efficiency reasons, we aimed to minimize the occlusion of the superimposed motifs. Therefore, we avoided complex motifs or inclined lines which take more space than straight lines or dots. We tested three chart heights (1, 1/4, and 1/8 scale factors) where the tallest chart corresponded to 64 pixels and physical size of 8 millimeters on our screen.

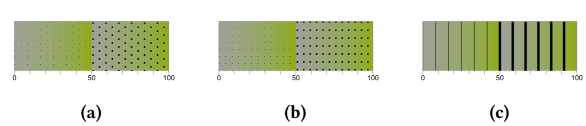

Figure 5: We tested three textures for the Texture–Saturation design. We finally selected (a) for the main evaluation.

Six persons performed the discrimination and estimation task on 45 trials (15 trial per condition). Accuracy was similar for all three cases. Nevertheless, a few observations by the participants convinced us to choose the texture 5a. Some participants reported having discomfort similar to moiré effects [30] with 5b and 5c due to the repetitive pattern. Impact of spatial frequency (e.g., "spreading" effect) on visual and color perception is well documented in the literature [31]. These effects can heavily influence the perception of colors between the bars in 5c. Some users also reported that in some instances, a marker was positioned exactly on a bar which naturally covered the underlying color and so, they based their estimation on the color in adjacent zones.

6 MAIN EXPERIMENT

Thanks to the findings from the preliminary pilot tests, we refined our final experimental design. We consequently removed line charts (Figure 1a) and Hue–Saturation (Figure 1f) for the reasons discussed earlier. The four remaining visual mappings, i.e., Horizon Graphs, Position–Value, Hue–Value, and Texture–Sat were used as they appear in Figure 1. In order to maintain a reasonable number of conditions, we only tested three scale factors (1, 1/4, 1/8) where factor 1 corresponded to a height of 64 pixels and 8 millimeters on our screen. All charts measured 2000 pixels wide corresponding to 233 millimeters in physical size on our screen. We recruited twenty unpaid participants (17 male, 3 female) aged between 22 and 39 (with median age of 27.5) among people working in our laboratory. All participants were degree holders in computer science and had normal or corrected vision and no color deficiency. For all trials, participants were told to keep a fixed viewing distance of about 50 centimeters from the screen.

The participants first engaged in a tutorial consisted in total of 4 (chart) $\times$ 3 (size) $\times$ 3 (trials per condition) = 36 trials per participant. During this stage, participants could familiarize with all conditions and develop their strategies for the demanded tasks. The main evaluation that followed the tutorial consisted in a total of 4 (chart) $\times$ 3 (size) $\times$ 7 (trials per condition) = 84 trials per participant. Trials were divided into four experimental blocks of 21 trials and the order of the trials in each block was randomized for each participant. In order to make sure that participants make use of both visual channels, the comparison points were located in different orders of magnitude in each trial. We also balanced the trials for value difference between the comparison points. Participants performed the experiment with the procedure detailed in Section 4.1.

6.1 Results

For all conditions, discrimination scores averaged 98% or higher. Therefore, we focus on interpretation of other measurements. Linear regression lines in Figure 12 report overall observed trends for each visual mapping regarding different measures. For more in depth analysis of these trends, we used multi-factor analysis of variance (ANOVA) to identify statistically significant differences in performance. When significant differences were found, we used t-tests and Tukey honest significant difference (HSD) to identify pair-wise significance.

Discrimination Time : ANOVAs showed a significant effect of visual mapping on discrimination speed with $F(3, 57) = 15.894, p < 0.001$. Tukey HSD comparisons in Figure 6 showed that Hue–Value is significantly faster than Horizon Graphs regarding this criterion ($t(57) < 0.001$). Furthermore, pair-wise t-test showed that at the smallest chart height (4 pixels), participants performed significantly faster with Hue–Value and Texture–Saturation than Horizon Graph ($t < 0.01$).

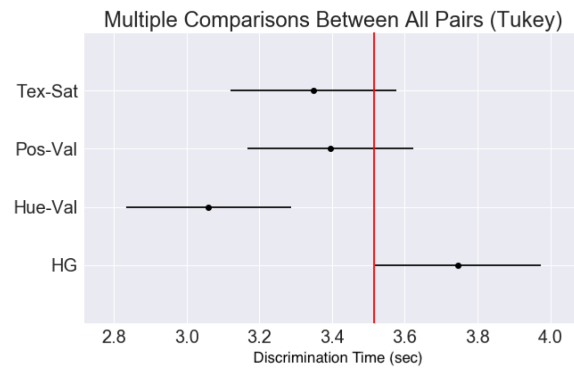

Figure 6: Discrimination time : Tukey honest significant difference (HSD) comparisons show significant difference between Hue–Value and Horizon Graphs (pair-wise $p < 0.001$). The red line is drawn to show the overlapping and lack of significance between Horizon Graphs and other two mappings. Measures are in seconds.

Estimation Error : Figure 7 shows distributions of estimation errors by each visual mapping across all chart heights. Even though they look similar, ANOVAs showed a significant effect of visual mapping ($F(3, 57) = 15.021, p < 0.001$). Post-hoc Tukey HSD comparisons identified Texture–Saturation as the only significantly different visual mapping in terms of estimation error averaged across chart heights (see Figure 9). We found a significant difference in estimation accuracy across different chart heights ($F(2, 57) = 15.033, p < 0.001$). Tukey HSD comparisons confirm significant differences between the smallest chart height and the two taller heights (see Figure 8). In order to identify the exact effects of size, we calculated t-tests which confirm that even if the differences in estimation accuracy is insignificant between 1/16 and 1/4 scale factors, Horizon Graph has significant worse accuracy in those chart heights than the tallest chart height (with pair-wise significances $p < 0.05$).

Estimation Time : We found a significant effect of chart height in estimation error ($F(2, 38) = 4.42, p = 0.012$). Tukey HSD comparisons reveal that estimation time decreases with chart height and the difference is significant between the smallest and tallest chart height (see Figure 10). We found no significant effect of visual mapping on estimation time ($F(3, 57) = 1.94, p = 0.14$). However, we found a weak significance of mapping effect between Horizon Graphs and Texture–Saturation at the tallest chart height by looking at t-test results (pair-wise $p = 0.04$).

6.2 Questionnaire Results

Following the evaluation, participants filled in a questionnaire in which, they rated and commented the four mappings. Participants

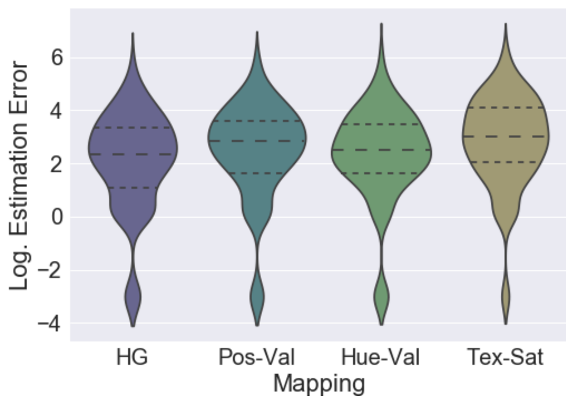

Figure 7: Distribution of logarithmic estimation error by visual mapping across all chart heights

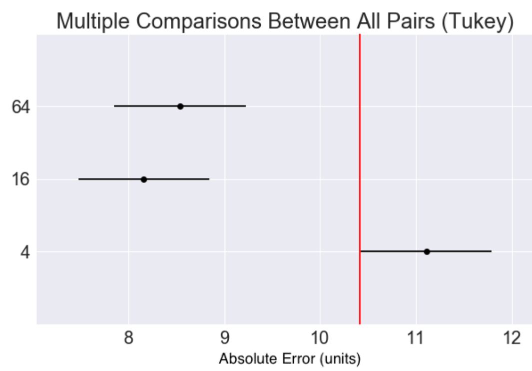

Figure 8: Estimation accuracy by chart height : Tukey honest significant difference (HSD) comparisons show significant higher errors in the smallest chart height. Errors are reported in units.

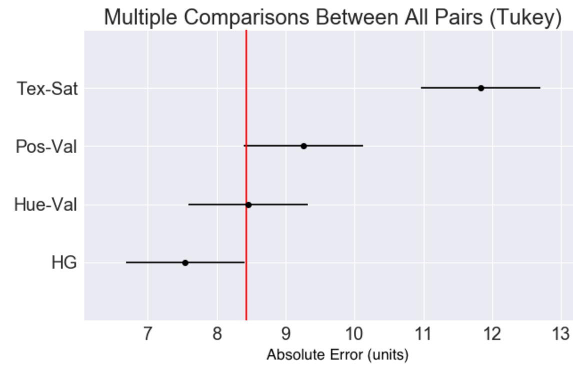

Figure 9: Estimation accuracy : Tukey honest significant difference (HSD) comparisons show no significant difference between Horizon Graphs and other mappings except Texture–Saturation (pair-wise $p < 0.001$). Errors are reported in units.

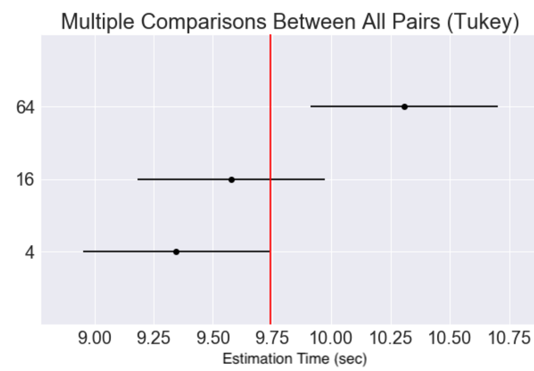

Figure 10: Estimation time : Tukey honest significant difference (HSD) comparisons show faster estimation at smaller chart heights.

were asked to give a note out of 10 to each mapping for discrimination and estimation tasks in small and large chart heights respectively. They also rated their overall satisfaction and their familiarity with each technique.

Horizon Graph was the most familiar mapping for the most of the participants, seconded by Position–Value. Other techniques were much less known to the participants, even though some of the participants had already participated in our pilot experiments. Participants reported Hue–Value as their favorite technique, followed by Horizon Graph (see Figure 11). For the discrimination task, participants in average preferred Hue–Value regardless of chart's height. Horizon Graphs was most favored for estimation task in large chart sizes and Hue–Value in small chart heights.

6.3 Discussion

We accepted our first hypothesis **(H1)** as valid after that we partially reproduced the findings from comparison of Horizon Graphs and line charts in our first pilot experiment. Our results confirm that at the tallest chart heights (1 scale factor), Horizon Graphs have the best estimate accuracy compared to other designs ($p < 0.05$ for all pair-wise comparisons). The accuracy is impacted by reduction of charts' height, yet the effect is not identical for all designs. Accuracy differences between the smallest and tallest chart heights were most significant for Horizon Graphs (pair-wise $p = 0.011$) and Texture–Saturation (pair-wise $p < 0.001$) and not significant for Hue–Value and Position–Value (see also Figure 12a). At the smallest chart heights, Horizon Graph lost its absolute superiority and we showed that accuracy of Hue–Value, Position–Value, and Horizon Graphs converge (as seen in Figure 12a). This validates our second hypothesis **(H2)** that in small chart heights, our visual designs are on a level with Horizon Graphs.

We did not observe significant impact of chart height on discrimination time and Hue–Value was significantly faster than Horizon Graphs across all chart heights (see Figure 12c for trends). Similar to findings in previous works, we found that estimation time decreases

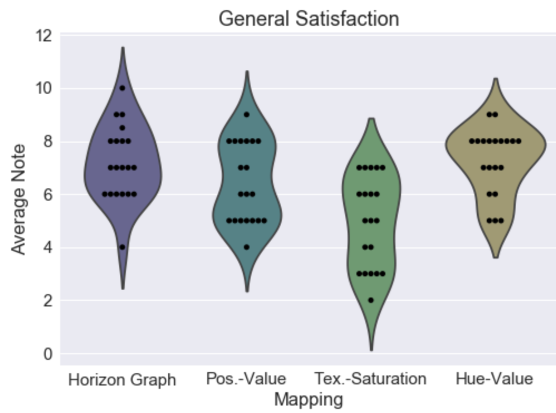

Figure 11: General Satisfaction for each visual design. Each black point represents one participant's answer.

as chart height decreases and that it converges for all mappings in the smallest chart hight(see Figure 12b). This partially confirms our third hypothesis **(H3)** that in small chart heights mappings such as Hue–Value offer the same levels of accuracy as Horizon Graphs and yet, offer advantages in terms of speed.

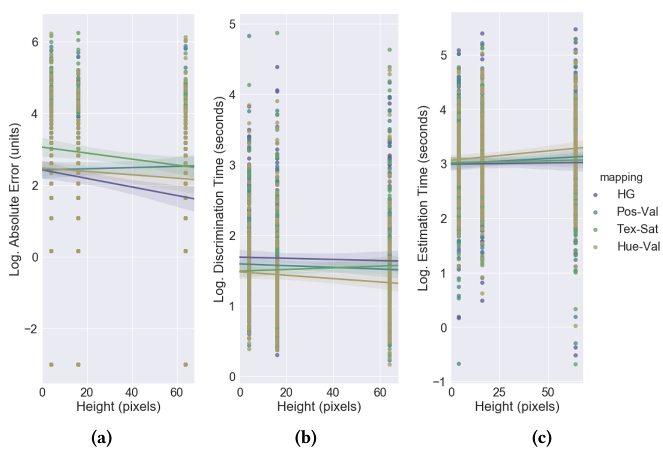

Figure 12: Linear regression lines with 95% confidence intervals for different mappings across the tested chart heights showing trends for (a) estimation errors, (c) discrimination time, and (b) estimation time.

7 DESIGN IMPLICATIONS

Based on the presented experimental results and discussions with participants of our study, we offer the following design recommendations for effective use of the introduced visual designs and composite visual mapping in general.

Optical visual channels gain advantage as chart size decreases

Users of visualization systems have naturally more confidence in comparing physical attributes like size and position than perceptual attributes like color. At least for the task we tested, we showed that accuracy difference is reasonable in large sizes and insignificant in small sizes. We believe that optical visual channels can be used more widely in cases where perception of size-related attributes are impacted, i.e., when chart size decreases.

Using texture in composite mapping is tricky

Blending texture and variation of other optical channels should be considered with a grain of salt. Many motifs for texture can unnecessarily occlude information, introduce visual clutter, or create uncomfortable vibratory effects. Some other's may look too minimalistic and difficult to distinguish in small sizes. Also, texture can seriously interfere with perception of the underlying color.

Hue coding is fast, but context should be considered

Mapping information on hue channel can increase discrimination speed even in small chart sizes. However, interpretation of hues is context dependent and some hues (e.g., bright reds and yellows) are easier to spot. For example, participants in our experiment found it extremely easy to spot a red region as it raises *a red flag*. As the participants were asked to report the higher value, the selected order of hues (reds > blues) was in coherence with their task.

Geometric channels can still be used in small chart heights

Horizon Graph is still a powerful space-efficient technique for time series visualization and further improvements in screen resolutions can ease the use of geometric visual channels in small screen sizes. However, size-related visual channels may suit better for discrete information (e.g., order of magnitude) or data with less fluctuation (e.g., trend component of a time series) in small chart sizes.

8 LIMITATIONS AND FUTURE WORK

Our design recommendations are valid with some limitations. Due to experimental and statistical constraints, we only investigated a common but narrow task of comparison and estimation and only examined one type of data. Visual analysis of time series involves many other tasks that have not been evaluated in this study. For future studies, other decompositions (e.g., frequency separation instead of modulo division) can be evaluated for corresponding tasks (e.g., trend detection) with their commonly associated data (e.g., very detailed or noisy data). Also, we only evaluated a specific type of data with only positive values. The negataive values in those cases may require an extra visual channel.

We learned from our evaluations and discussions that the new visual designs are easy to learn by novice users and mostly positively received by participants. We should remember that users are more familiar with geometric representations of data and that further training may improve their performance with our proposed designs.

We are also aware that further studies are clearly needed to evaluate scalability issues, i.e., larger number of layers and visualization of larger and higher dimensional datasets. While our study report important trends on the effects of chart height and visual mapping, more participants with more diverse backgrounds can improve the statistical robustness of our analysis.

9 CONCLUSIONS

In this work, we presented new visual designs for space-constrained time series visualization. Similar to Horizon Graphs, these visual designs combine two visual channels to visualize the same data attribute, in order to separate different facets of data and increase the information bandwidth. We pit our designs against Horizon Graph in discrimination and estimation tasks and observed that while mappings with geometric visual channels (i.e., Horizon Graph) have superior performance in larger chart heights, they are significantly impacted by chart height reduction. In smaller chart heights, our proposed visual designs that mainly use optical visual channels have proven accuracies similar to Horizon Graph while being faster in some cases (i.e., Hue–Value).

We note that despite the general mistrust of optical channels for quantitative visualization due to their perceptual nature, they perform as good as geometric visualizations in space-constrained visualizations in the tasks we evaluated. We offered design recommendations based on our findings, yet more work is clearly needed to evaluate effectiveness of these designs with different types of data and other scenarios.

RÉFÉRENCES

[1] Christopher Ahlberg, Christopher Williamson, and Ben Shneiderman. 1992. Dynamic queries for information exploration : An implementation and evaluation. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI '92)*. ACM, ACM, 619–626.

[2] Josef Albers. 2006. *Interaction of color.* Yale University Press.

[3] Benjamin B Bederson, Ben Shneiderman, and Martin Wattenberg. 2002. Ordered and quantum treemaps : Making effective use of 2D space to display hierarchies. *ACM Transactions on Graphics (TOG)* 21, 4 (2002), 833–854.

[4] Jacques Bertin. 1983. *Semiology of graphics : diagrams, networks, maps.* University of Wisconsin press.

[5] Lee Byron and Martin Wattenberg. 2008. Stacked graphs - Geometry & aesthetics. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1245–1252. DOI : http://dx.doi.org/10.1109/TVCG.2008.166

[6] Sk Card and Jock Mackinlay. 1997. The structure of the information visualization design space. *Information Visualization, 1997. ...* (1997), 92–99. DOI : http://dx.doi.org/10.1109/INFVIS.1997.636792

[7] M Chen, S Walton, K Berger, J Thiyagalingam, B Duffy, H Fang, C Holloway, and A E Trefethen. 2014. Visual Multiplexing. *Comput. Graph. Forum* 33, 3 (2014), 241–250.

[8] William S Cleveland. 1993. *Visualizing data.* Hobart Press.

[9] Wikipedia contributors. 2017. Fisher–Yates shuffle — Wikipedia, The Free Encyclopedia. (2017). https://en.wikipedia.org/w/index.php?title=Fisher%E2%80%93Yates_shuffle&oldid=811076565 [Online ; accessed 12-December-2017].

[10] Oscar de Bruijn and Robert Spence. 2000. Rapid Serial Visual Presentation : A Space-time Trade-off in Information Presentation. In *Proceedings of the Working Conference on Advanced Visual Interfaces.* ACM, 189–192.

[11] Alan Dix and Geoff Ellis. 2002. By Chance - Enhancing Interaction with Large Data Sets Through Statistical Sampling, In Proceedings of the Working Conference on Advanced Visual Interfaces. *Proceedings of the Working Conference on Advanced Visual Interfaces* (2002), 167–176. DOI : http://dx.doi.org/10.1145/1556262.1556289

[12] Geoffrey Ellis and Alan Dix. 2007. A Taxonomy of Clutter Reduction for Information Visualisation. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1216–1223.

[13] Paolo Federico, Stephan Hoffmann, Alexander Rind, Wolfgang Aigner, and Silvia Miksch. 2014. Qualizon Graphs : Space-efficient Time-series Visualization with Qualitative Abstractions. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces (AVI '14).* ACM, ACM, 273–280. DOI : http://dx.doi.org/10.1145/2598153.2598172

[14] Ying-Huey Fua, Matthew O Ward, and Elke A Rundensteiner. 1999. Hierarchical parallel coordinates for exploration of large datasets. In *Proceedings of the conference on Visualization'99 : celebrating ten years.* IEEE Computer Society Press, IEEE Computer Society Press, 43–50.

[15] Charles D Hansen, Min Chen, Christopher R Johnson, Arie E Kaufman, and Hans Hagen. 2014. *Scientific visualization : uncertainty, multifield, biomedical, and scalable visualization.* Springer. DOI : http://dx.doi.org/10.1017/CBO9780511975721.020

[16] Jeffrey Heer, Nicholas Kong, and Maneesh Agrawala. 2009. Sizing the horizon : the effects of chart size and layering on the graphical perception of time series visualizations, In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. *CHI '09Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, April 4–9, 2009* (2009), 1303–1312. DOI : http://dx.doi.org/10.1145/1518701.1518897

[17] Ali Jabbari, Renaud Blanch, and Sophie Dupuy-Chessa. 2018. Composite Visual Mapping for Time Series Visualization. In *2018 IEEE Pacific Visualization Symposium (PacificVis).* IEEE, 116–124.

[18] Waqas Javed, Bryan McDonnel, and Niklas Elmqvist. 2010. Graphical perception of multiple time series. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 927–934. DOI : http://dx.doi.org/10.1109/TVCG.2010.162

[19] D A Keim. 2002. Designing pixel-oriented visualization techniques : Theory and applications. *IEEE Transactions on Visualization and Computer Graphics* 6, 1 (2002), 59–78.

[20] Daniel A Keim and Annemarie Herrmann. 1998. The gridfit algorithm : An efficient and effective approach to visualizing large amounts of spatial data. In *Proceedings of the Conference on Visualization '98 (VIS '98).* IEEE, IEEE Computer Society Press, Los Alamitos, CA, USA, 181–188.

[21] Heidi Lam, Tamara Munzner, and Robert Kincaid. 2007. Overview Use in Multiple Visual Information Resolution Interfaces. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1278–1285.

[22] Jock Mackinlay. 1986. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics* 5, 2 (1986), 110–141. DOI : http://dx.doi.org/10.1145/22949.22950

[23] Justin Matejka, Michael Glueck, Tovi Grossman, and George Fitzmaurice. 2016. The Effect of Visual Appearance on the Performance of Continuous Sliders and Visual Analogue Scales. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2016), 5421–5432. DOI : http://dx.doi.org/10.1145/2858036.2858063

[24] Peter Mclachlan, Tamara Munzner, and Florham Park. 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).* ACM, 1483–1492.

[25] Elisabeth S. Nelson. 2000. Designing Effective Bivariate Symbols : The Influence of Perceptual Grouping Processes. *Cartography and Geographic Information Science* 27, 4 (2000), 261–278. DOI : http://dx.doi.org/10.1017/CBO9781107415324.004 arXiv:arXiv:1011.1669v3

[26] Elisabeth S. Nelson. 2002. The Impact of Bivariate Symbol Design on Task Performance in a Map Setting. *Cartographia* 37, 4 (2002), 61–78. DOI : http://dx.doi.org/10.1017/CBO9781107415324.004 arXiv:arXiv:1011.1669v3

[27] Charles Perin. 2013. Interactive Horizon Graphs : Improving the Compact Visualization of Multiple Time Series. *CHI '13 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), 3217–3226. DOI : http://dx.doi.org/10.1145/2470654.2466441

[28] Takafumi Saito, Hiroko Nakamura Miyamura, Mitsuyoshi Yamamoto, Hiroki Saito, Yuka Hoshiya, and Takumi Kaseda. 2005. Two-tone pseudo coloring : Compact visualization for one-dimensional data. *Proceedings - IEEE Symposium on Information Visualization, INFO VIS* (2005), 173–180. DOI : http://dx.doi.org/10.1109/INFVIS.2005.1532144

[29] Ben Shneiderman. 1992. Tree Visualization with Tree-Maps : 2-D Space-Filling Approach. *ACM Transactions on graphics (TOG)* 11, 1 (1992), 92–99.

[30] Edward R Tufte and PR Graves-Morris. 1983. *The visual display of quantitative information* (2 ed.). Graphics press Cheshire, CT.

[31] Brian A Wandell. 1995. *Foundations of vision.* Sinauer Associates.

[32] Steffen Wenz. 2009. Beyond-the-Desktop Interactive Visualizations. *Trends in Information Visualization* (2009), 54.