International Journal of Cartography

ISSN: 2372-9333 (Print) 2372-9341 (Online) Journal homepage: https://www.tandfonline.com/loi/tica20

# The influence of domain expertise on visual overviews of spatiotemporal data

Susanne Bleisch, Matt Duckham & Chris Pettit

**To cite this article:** Susanne Bleisch, Matt Duckham & Chris Pettit (2017) The influence of domain expertise on visual overviews of spatiotemporal data, International Journal of Cartography, 3:2, 166-186, DOI: 10.1080/23729333.2017.1294820

**To link to this article:** https://doi.org/10.1080/23729333.2017.1294820

Published online: 16 May 2017.

Submit your article to this journal

Article views: 88

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=tica20

INTERNATIONAL JOURNAL OF CARTOGRAPHY, 2017
VOL. 3, NO. 2, 166–186
https://doi.org/10.1080/23729333.2017.1294820

Check for updates

# The influence of domain expertise on visual overviews of spatiotemporal data

Susanne Bleisch [a], Matt Duckham[b] and Chris Pettit[c]

[a]Institute of Geomatics Engineering, FHNW University of Applied Sciences and Arts Northwestern Switzerland, Muttenz, Switzerland; [b]School of Science, RMIT University, Melbourne, Australia; [c]Faculty of the Built Environment, University of New South Wales, Sydney, Australia

**ABSTRACT**

Overviews of spatiotemporal data are acknowledged to play an important role in visualization in initiating and supporting geovisualization and exploratory data analysis (EDA). However, relatively little research has focused on the visual overviews themselves, and their potential impacts on EDA outcomes. In a user study, we evaluated the influence of different levels of domain knowledge on the usefulness of four distinct types of static visual overview of spatiotemporal data. Beyond simply orienting users, our results indicate that visual overviews can be important in gaining insights into a data set, for example, in learning about metadata. Although subjects without domain knowledge struggled to judge the quality of their findings, they were as successful at identifying interesting patterns in the data as those with domain expertise. Our results suggest that detailed background knowledge of a data set can actively hinder EDA. Being already familiar with their own data sets, our results highlight the tendency of data experts to disregard findings that do not match their pre-existing domain knowledge. Based on these findings, our conclusions identify a range of potential avenues for future work, including the use of visual overviews that deliberately do not, from first view, reveal the context of the data they show. This later approach could help in cases where domain experts need to see their data with 'fresh eyes', and detect interesting patterns in spatiotemporal data before relating the findings to specific knowledge about the data sets and the domain.

**RÉSUMÉ**

Les vues d'ensemble sur les données spatiotemporelles sont reconnues pour jouer un rôle important en visualisation pour débuter et faciliter la géovisualisation et l'analyse de données exploratoire (ADE). Pourtant assez peu de recherches se sont focalisées sur les vues d'ensemble et leurs impacts sur les résultats des ADE. Au cours d'une étude utilisateurs, nous avons évalué l'influence du niveau de connaissance du domaine sur l'efficacité de quatre différentes vues d'ensemble sur des données spatiotemporelles. Au-delà d'orienter simplement les utilisateurs, nos résultats montrent que les vues d'ensemble visuelles peuvent être importantes pour mieux comprendre les données, par exemple pour l'apprentissage des métadonnées. Bien que les

**ARTICLE HISTORY**
Received 29 May 2015
Accepted 10 February 2017

**KEYWORDS**
Visual overviews; exploratory data analysis; evaluation; usefulness; domain knowledge; spatiotemporal data; movement analysis

**CONTACT** Susanne Bleisch susanne.bleisch@fhnw.ch
© 2017 International Cartographic Association

personnes n'ayant pas de connaissance du domaine aient eu des difficultés pour juger de la qualité de ce qu'ils ont trouvé, ils ont trouvé avec autant de succès les patterns intéressants dans les données que les experts du domaine. Nos résultats suggèrent que des connaissances de base détaillées sur les données peuvent entraver l'ADE. Etant déjà familiarisés avec leurs propres données, nos résultats mettent en évidence la tendance des experts sur les données de ne pas considérer les résultats qu'ils trouvent lorsque ces résultats ne correspondent pas à leur connaissance antérieure du domaine. A partir de ces résultats nous identifions un ensemble de directions possibles pour les travaux à venir, en incluant l'utilisation de vues d'ensemble qui, de façon délibérée, ne révèleraient pas dans un premier temps le contexte des données qu'elles représentent.

## 1. Introduction

Exploring a large spatiotemporal data set to find salient patterns and identify new information can be a daunting task. This may be especially true in cases where users have limited domain knowledge about the data set or set out with ill-defined research questions. Visual overviews are acknowledged to play an important role in supporting geovisualization and exploratory data analysis (EDA), both in initiating exploration and in orienting users throughout their investigations (Few, 2009; Hornbæk & Hertzum, 2011; Keim, 2001). The visual information seeking mantra (Shneiderman, 1996), for example, starts with an overview of the whole data set before offering tools for filtering or selecting subsets of the data for answering specific questions.

While the information seeking mantra (Shneiderman, 1996) is often used as a guiding principle for visualization design (Craft & Cairns, 2005), there is a paucity of the literature regarding the importance and perceived usefulness of using visual overviews for generating ideas and hypotheses. Haroz and Whitney (2012) have highlighted the difficulties humans face in making sense of large amounts of data displayed at once, but do not address their effectiveness in triggering further questions for exploration.

The qualitative study presented in this paper focuses on the influence of varying levels of domain knowledge on the perceived usefulness of four types of static visual overviews of spatiotemporal data, as a prelude to EDA. The static visual overviews present in its entirety a large ecological data set about the movement of hundreds of fish over a period of years. Our focus is on the generation of hypotheses and ideas as a starting point for further EDA, which needs a certain purpose (Andrienko & Andrienko, 2006) and thus initial hypotheses. The confirmation of hypotheses (cf. so-called overview-plus-detail interfaces, Baudisch, Good, Bellotti, & Schraedley, 2002) would be a natural next step of the data exploration process, but is not part of this study.

Therefore, our goal in this paper is to address the following questions:

- What are the differences in the usage of visual overviews by user groups with varying domain knowledge?
- What can we learn from the differences in perceived usefulness in order to benefit most from visual overviews?

### 1.1. Application domain and data set

Our analysis was conducted with a large spatiotemporal data set about the movements of fish in the Murray River, Australia (Bird, Lyon, Nicol, McCarthy, & Barker, 2014). The Murray River is a river of considerable economic and ecological significance, heavily used for irrigation purposes. The data set concerned the movement of approximately 1100 radio-tagged fish moving between 24 river zones between 2006 and 2011. The monitored area included several tributaries and a dammed lake with a total length of more than 300 km. Figure 1 shows a schematic of the different river zones in the river system, labeled a–x, with the locations of the 18 radio-telemetry logging towers that partition the river into zones indicated with triangles.

The data itself consisted of tuples containing a letter indicating the river zone (or NA before a fish was radio-tagged or after its radio-tag had ceased to respond) for each day between 1 January 2006 and 31 December 2011 (2190 days) and for each radio-tagged fish (1131 fish), a total of more than 2.5M records. The data set was originally collected and analyzed in order to evaluate the effect on native fish populations of conservation activities in certain river sections. However, it was expected that this data set could be a rich source of further information and insights about fish behavior in the Murray River. Thus, exploring these data holds the potential to reveal new knowledge about fish behavior and to generate hypotheses for further analysis beyond the specific research questions that the study set out to address.

## 2. Materials

In general, a visual overview 'is constructed from, and represents, a collection of objects of interest' (Greene, Marchionini, Plaisant, & Shneiderman, 2000). Visual overviews generally make no or minimal assumptions about the specific exploration tasks. Craft and Cairns (2005) state that visual overviews paint 'a "picture" of the whole data entity' and allow seeing patterns 'from a vantage point that comprises the whole view' (Craft & Cairns, 2005, p. 111). Examples of different types of visual overviews surveyed in Hornbæk and Hertzum (2011) include tag clouds, tree maps, graph overviews and fisheyes. This selection shows that visual overviews may include some form of aggregation. However, aggregation

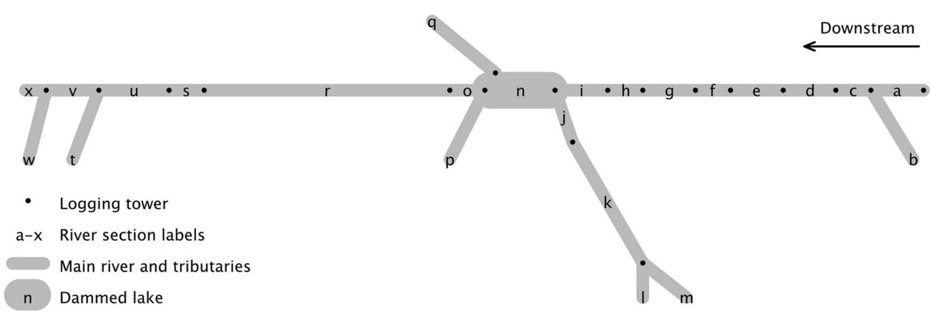


**Figure 1.** Schematic map of the monitored part of the Murray River and its tributaries. The river is partitioned into 24 different river zones (a–x) by 18 logging towers (dots). Zone n is the dammed Lake Mulwala.

or visibility of 'major components' or 'significant features' (Craft & Cairns, 2005, p. 111) can also be achieved using selected visual variables and data-dense displays. In this section, we introduce the design and characteristics of the data graphics used in our study to provide visual overviews.

### 2.1. Defining visual overviews

For this study, the visual overviews are designed to show the complete data set in a single two-dimensional display, through the use of different visual variables, such as color, size and layout. We do not hypothesize about the data and thus only use aggregation or suppress aspects of the data set when otherwise it is impossible to show all the data in the same view. We try to achieve visual overview and aggregation by employing data-dense representations. We also restricted our visual overviews to static and two-dimensional views, on the grounds that any interactive, 3D or animated capabilities immediately move us away from overviews of the entire data set into affording users capabilities to explore detail. Such further exploration of details is a natural next step when the overviews have allowed generating hypotheses and ideas about the data and when the iterative process of EDA (Andrienko & Andrienko, 2006) has started.

Thus, in achieving static two-dimensional visualizations of the space, time and attributes in our data set, it was required to suppress certain aspects or characteristics of the data in favor of others. No assumptions about the tasks to be fulfilled or the users' intentions when exploring the data were made, except that we focused on the primary dimension of general interest: movement (i.e. moving fish is more interesting than stationary fish). Thus in a functional view of the data set (Andrienko & Andrienko, 2006), fish movement is the characteristic, while space and/or time serves as references (cf. Table 1).

Besides this fundamental setting, we adopted the following guidelines in designing our visual overviews:

- reuse existing conventions, including color coding of existing data graphics already in use by the data experts;
- employ data aggregation only when necessary to construct static 2D representations of our movement data;
- keep the layouts simple and reusable with minimal formatting of the data, to avoid biases that might be introduced through specific optimization of the visual overviews and to maximize the applicability of results to different data sets;
- minimize the use of labeling and legends, as the visual overviews will be explained and used in an interactive setting where questions are encouraged; and
- aim for a clear and legible display on A3 paper (taking into account the natures of overview visualization where the single data point may not need to be clearly visible but the users should gain a comprehensive view of the data set at once).

The last of these guidelines, format for A3 paper, was for both practical reasons (to make it easy for subjects to discuss the visual overviews without needing any specialized equipment) and to avoid any unnecessary confusion or unwarranted expectation of interactivity by the users (e.g. no one expects a paper to be able to zoom or pan).

**Table 1.** Visual encodings of the data set dimensions (time, space and fish ID) in the four visual overviews (V1–V4). The data characteristic – fish movement – is encoded differently in all visualizations (see the last column).

| | Data | | | Dimension of interest/ data characteristic |
|---|---|---|---|---|
| | Time | Space | Fish (ID) | Fish movement |
| V1 | *x*-axis | Color-coded | *y*-axis | Color changes along horizontal lines |
| V2 | *x*-axis | Linear arrangement of zones on *y*-axis | Line through space and time | Vertical lines |
| V3 | Suppressed | Schematic, not to scale map; *x*- and *y*-axis for relative locations | Grouped and aggregated, bar length | Star plot arms |
| V4 | Suppressed | Linear arrangement of zones on *y*-axis (origin) and *x*-axis (destination) | Grouped and aggregated, cell shading/value | Cell shading/values |

Table 1 gives an overview of the visual encodings of the data characteristics (fish ID or movement) and the referential components (space and/or time). All visualizations show the dimension of interest – fish movement – differently (see the last column of Table 1).

### 2.1.1. Visual overview V1

Figure 2 (minimized to fit page width) shows the location of each fish during the analyzed time span. The *y*-axis shows the randomly ordered fish IDs. The *x*-axis shows time. The colors indicate the different river zones, and reuse the color coding employed by the data collectors in their internal documents. Thus, following a single horizontal line shows the whereabouts of a single fish over time through color-coded space.

### 2.1.2. Visual overview V2

Figure 3 (minimized to fit page width) emphasizes fish movement between river zones over time. The *y*-axis shows simplified space as a linear arrangement of the river zones labeled a–x (Figure 1). The *x*-axis shows time as in visualization V1. Thus, the two

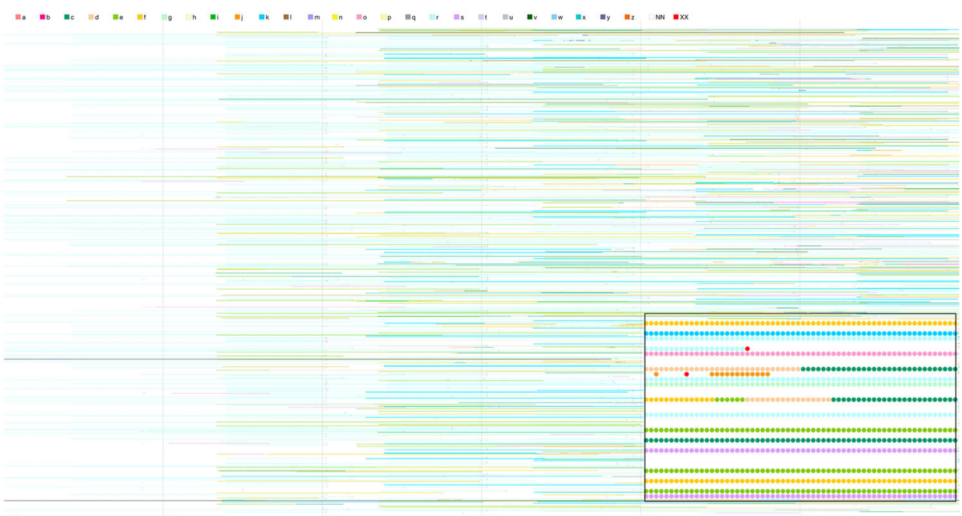


**Figure 2.** Visualization V1 (minimized to fit page width, zoomed-in inset) showing location (color-coded river zones) of all fish (randomly ordered fish IDs on *y*-axis) through time (*x*-axis).

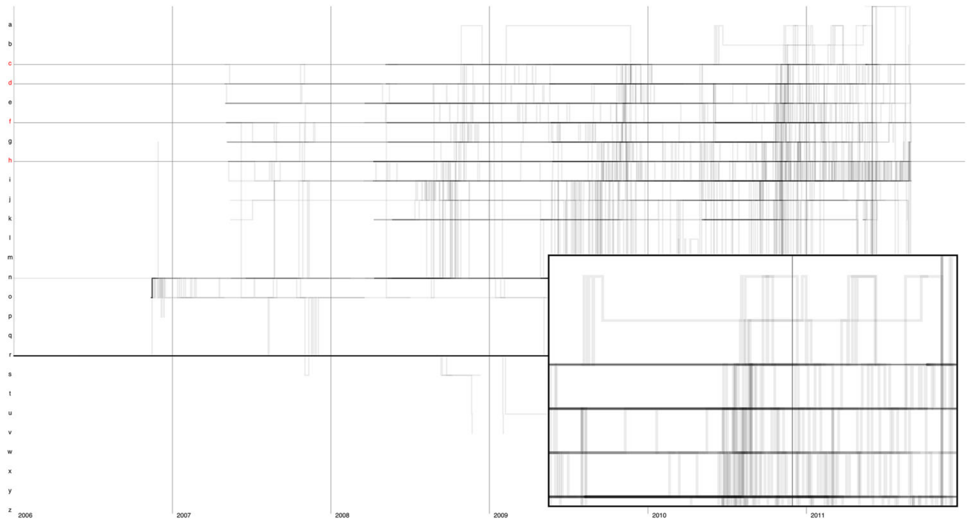


**Figure 3.** Visualization V2 (minimized to fit page width, zoomed-in inset) showing fish movement (transparent gray lines) between river zones (*y*-axis) over time (*x*-axis).

visualizations V1 and V2 can be vertically aligned and analyzed in comparison. In V2, each fish is plotted as a transparent gray line as it moves between zones (vertical lines) or stays in specific river zones (horizontal lines) over time. All fish paths over time are plotted on top of each other. Thus, darker lines in a horizontal direction mean more fish being stationary, while darker lines in the vertical direction mean more fish engaged in changing between river zones. The visualization was inspired by plots of single fish paths that were already being used by the domain experts to learn about the behavior of selected fish.

### *2.1.3. Visual overview V3*

Figure 4 (minimized to fit page width) uses a schematic map (i.e. not to scale) of the river as the background (cf. Figure 1). Each river zone is centrally marked by a colored dot (using the same colors as in visualization V1). The movement of fish between river zones is aggregated over the whole six-year time period. For each river zone, a star plot indicates how many fish have left that specific river zone and moved to any of the other river zones

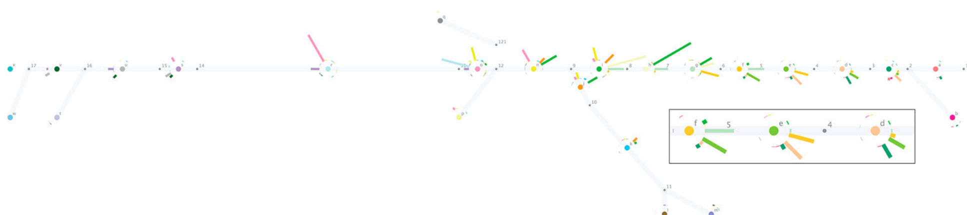

**Figure 4.** Visualization V3 (minimized to fit page width, zoomed-in inset) shows the number of fish (aggregated over time) moving from a specific river zone to any of the other river zones through star plots. The arm lengths of the star plots indicate numbers of fish, while the colors relate to the river zone the fish are moving to.

(one star plot arm per ‘moved-to’ river zone). The length of the star arms relates to the absolute numbers of fish moving out, while the color of the arms shows to which river zone they moved to.

### *2.1.4. Visual overview V4*

Figure 5 (minimized to fit page width) is an origin–destination (OD) matrix, typically used by transport planners, for example, to analyze travel behaviors of people, moving through a city. In the visual overview in Figure 5, the OD matrix shows the number of fish moving from one river zone to another (numbers aggregated over the whole six-year time period). The *y*-axis shows the linear arrangement of the river zone labels where fish movement originated and the *x*-axis show the destination river zones. The total numbers of fish moving are labeled in the cells of the OD matrix and, additionally, the cells are shaded in gray

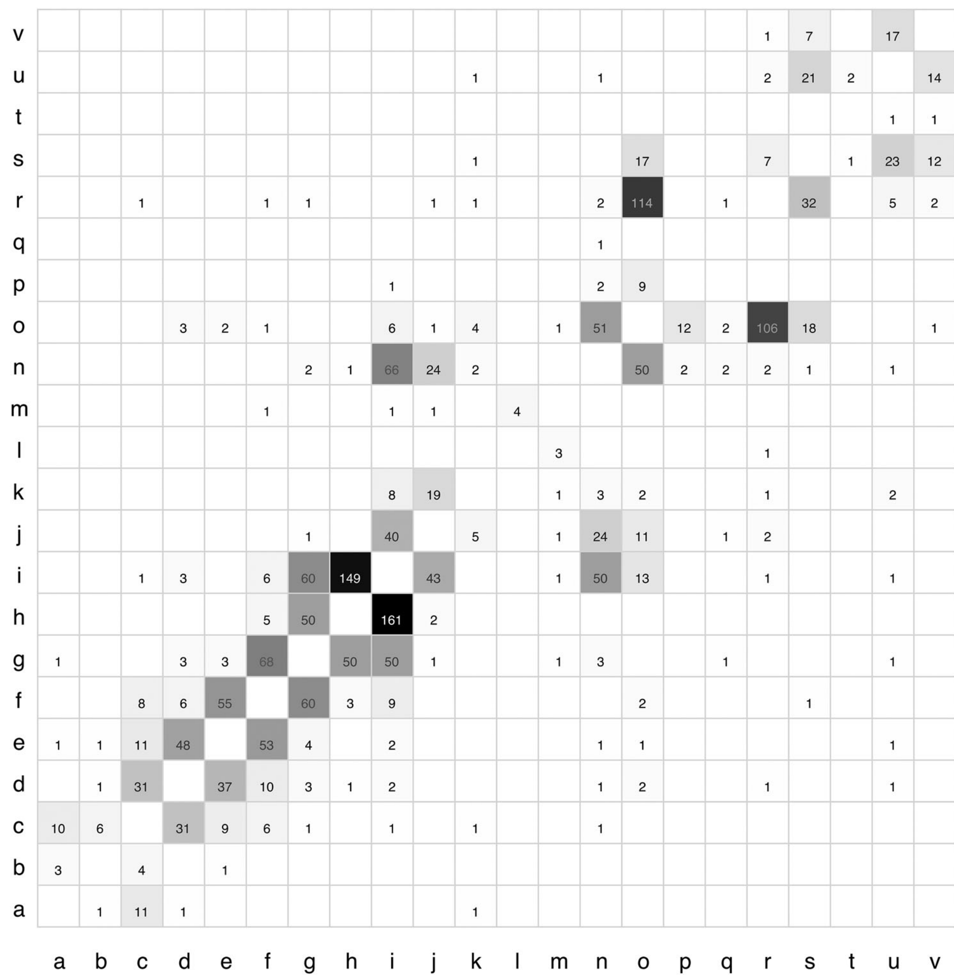


**Figure 5.** Visualization V4 (minimized to fit page width) shows the number of fish (aggregated over time) moving from a specific river zone to any of the other river zones in an OD matrix. The *y*-axis labels the river zones where fish movement originates from; the *x*-axis shows the destination river zones.

values relating to those numbers. The values represented in V4 are the same as the values of the star plot arms in visualization V3 (cf. Figure 4). The linear arrangement of the Murray River allows employing an OD matrix. For more complex river systems, an OD map (Wood, Dykes, & Slingsby, 2010) could be used.

## 3. Methods

In his seminal book on EDA, Tukey (1977) acknowledges the value of domain knowledge. He compares EDA to detective work that is facilitated by knowing the setting in which a crime takes place. For example, if you have no idea where a subject is likely to have put his or her fingers, then you may not find a single fingerprint (Tukey, 1977). Similarly, it is difficult to formulate findings using EDA without domain knowledge. However, Tukey (1977) also states that pictures are of most value when they force us to see what we did not expect to see. Thus, it seems likely that people may gain an overview of a data set and begin to generate hypotheses using EDA with only limited domain knowledge. While then the findings cannot be interpreted in the specific subject-matter context, an open mind and an interest in the data set or the application domain may still help learning and building knowledge from the visualizations. Hartwig and Dearing (1979) argue that basic adherence to the two principles 'openness' and 'skepticism' will maximize learning from data through EDA. Thus, in this study, by comparing the exploratory visual data analysis pursued by users with varying levels of domain knowledge, we set out to gain an understanding of the influence of the context knowledge on exploratory visual data analysis.

### *3.1. Participants*

In the user study reported here, a total of 20 participants with four different levels of domain knowledge participated in five focus group sessions. Such workshop sessions with experts have been shown to be an effective way in eliciting feedback on data visualization products, for example, in exploring landscape and climate change and land-use change scenarios (Pettit, Raymond, Bryan, & Lewis, 2011; Wissen, Schroth, Lange, & Schmid, 2008). The four different groups consisted of the following participants:

- DE1 (domain experts with data knowledge): One group of participants consisted of the domain experts (three participants, aquatic ecologists) who collected the data set used in this study. These participants had already analyzed the data for its original specific purpose, but expressed a general interest in learning more about what the data could tell them more broadly about common or unusual fish behavior in the Murray River.
- DE2 (domain experts without specific data knowledge): A second group consisted of domain experts (three participants, aquatic ecologists) who were not involved in the data collection process or any previous analysis of the data. Their interest was in learning generally about fish behavior in rivers.
- ND1 (PhD students): The third group of participants consisted of PhD students who did not have specific domain knowledge in the field of river ecology, but did have an interest in data analysis and/or visualization in general. They participated in two separate

focus group sessions with four participants in each session, but the statements of the two groups were combined for analysis.

- ND2 (senior researchers): Finally, a fourth group of participants consisted of senior researchers (six participants) with expertise in data analysis and/or visualizations in general, but without specific domain knowledge.

Two participants (one PhD student and one domain expert without data knowledge) have impaired color vision. The visualizations were not specifically optimized to account for this.

### 3.2. Focus groups

Focus groups are small groups moderated to ensure specific topics being discussed by the group (Grumbein & Lowe, 2010). The focus group methodology was used to collect qualitative data about how the study participants worked with our four visual overviews. The collection of data through focus groups has the advantage of being a less formal setting than a one-to-one interview situation and participants have the opportunity to build upon one another's comments.

All participants from a single focus group session belonged to the same group of participants in terms of the level of their domain knowledge. All focus group meetings took about 60 minutes and were structured as follows. First, the participants were introduced to the data set, including data collection and original purpose of the data (cf. Section 1.2). The domain expert group was also introduced similarly, even though they presumably already had this knowledge. Then the four visualizations were shown in order V1–V4 and explained (cf. Section 2.2). This introduction took about 10–12 minutes, the difference mainly being caused by the number of questions asked in some groups during the introduction. For the rest of the time, the participants were encouraged to work with the visualizations, for example, to discuss the visualizations and the data they showed, talk about any interesting patterns, report any hypotheses about the fish behavior they would like to explore further and any suggestions for different visualizations or tools they would like to use to examine the data set.

During the group discussions, any question asked by the participants was answered by the focus group moderator if the answer was known. This led to some groups learning more about the data set than others. Further input included asking the group for preferences or occasionally pointing them to specific parts of the visualizations to encourage group discussions. Usually, asking for preferences was enough to keep the discussions going.

With the participants' consent, video recording (from a static camera on the table) was used to capture the participants' precise statements and to allow for the reconstruction of the knowledge built up through discussion and the use of the visualizations, that is, including participants sketching on the paper printouts of the visualizations.

### 3.3. Qualitative data analysis

The video recordings of the five focus group sessions were transcribed and then qualitatively analyzed looking for general comments and insights (i.e. a 'nontrivial discovery', Plaisant, Fekete, & Grinstein, 2008). The concept of gaining insight (North, 2006) can be used to measure usefulness. Reported insights can be compared in relation to quantity, foci, depth, complexity and relevance of the insights gained, as suggested by North (2006). However, the collected data revealed that in-depth insights were rarely reported. This is to be expected given that only static overview data visualizations were used, rather than interactive exploratory visualization tools, which would allow a more in-depth analysis of data. Thus, a more diversified emerging coding scheme (Stemler, 2001), inductively developed from the focus group transcripts, was adopted. To identify the categories detailed below, we segmented the transcripts in statements and grouped them according to their characteristics. In an iterative review process, the categories were refined and condensed by combining similarities and separating unique characteristics until we arrived at the following seven main categories:

- *Finding*: The statements closest to insight were 'findings' where people reported something they saw in the visual overview. Although most findings did not contain enough depth to be called insights, the 'findings' category includes all findings from superficial ones to complex and unexpected insights.
- *Questions*: This category comprises two subcategories. One contains specific questions about the visual overview, the data and the data collection. Questions were generally answered if the answer was known. The second subcategory contains all statements that were phrased as questions but may also belong to other categories such as findings or improvements (e.g. 'Do fish skip zones?' is a finding phrased as question).
- *Analysis*: The analysis category includes all ideas and suggestions for further analysis of the same data, for example, together with additional data, with other visualizations, or other methods and techniques.
- *Improvement*: Suggestions for improvement of the visualization were often but not always mentioned together with a problem (see below).
- *Problem*: Reports of something awry in the visual overview, the data or with potential interpretations as a problem.
- *Context*: Statements about context information and knowledge. This category included information provided about the participants' background, their ability or the data set (especially by the experts).
- *Preference*: Statements formulated as either a positive or a negative preference, or a comparison for one or more of the visual overviews.

These categories do overlap in some instances, but all statements were only assigned one category. For example, participants occasionally stated the same aspect as a finding or as a question. However, even in these cases, it was interesting to consider why one group of participants framed the same idea as a question, while another stated it as a finding. Suggestions for improvement or further analysis could in some cases be an indication of a finding. For example, it might represent a finding when a participant sees a pattern in the visualizations and would like to explore that further by using additional data or an improved visualization.

Additionally, the statements in the seven coding categories above were grouped as either being general statements without specific reference to a visual overview; referring specifically to one of the four visual overviews or comparing at least two visual overviews in one statement. The statements were also grouped topic-wise per visualization and across visualizations. This facilitated the evaluation of which groups commented on what topic and in what form (i.e. some groups reported findings in a topic area or suggested improvements implicitly based on findings, while other groups phrased their findings in the same topic area as questions). The categorized statements were mainly analyzed qualitatively for their contents and meanings. Additionally, the qualitative statements were also analyzed quantitatively to gain an understanding of the approximate density of the discussions in the different participant groups.

### 3.4. Measuring and validating usefulness

Usability evaluation methods generally require defined tasks and goals based on which the performance of a system can be measured (Tobon, 2005). Measuring usefulness when the task is EDA – to gain an understanding and generate ideas, with no predefined questions with known right/wrong answers – is more difficult. In this case, standard measures such as task completion time and task accuracy (Chen & Yu, 2000) cannot be used. In relation to the seven categories of statements identified in Section 2.5, we define the following analysis structure for usefulness.

- Most importantly, we want to gain an indication of the variety of findings that were reported and whether participants stated ideas for further analysis. These aspects can also be roughly grouped into deeper or more superficial thoughts.
- Additionally, we looked at the second question subcategory (findings phrased as questions) and improvement suggestions. We presume that findings phrased as questions suggest a lack of trust in one's own findings. Potentially these indicate missing context knowledge that would allow evaluation of the finding. The suggestion of improvements to the visual overview may stem from a finding in that participants would like to explore further. However, suggestions for improvements can also be made purely based on preferences or features of visualizations seen elsewhere (e.g. change colors or add interaction possibilities) without a specific desire to follow up on a finding. Differentiation between these two types of improvement suggestions can be difficult.
- Finally, problem, context or preference statements can also indirectly include findings and new ideas, although this is more rare.

In analyzing and finally reporting the results in the following sections, a 'thick' description approach (Geertz, 1973) has been employed. The statements were aggregated per participant group and topic, but detail about the content of the statements and the level of context knowledge of each participant group was maintained. The framing of statements as belonging to one of the above-defined categories gives information about how the context knowledge influenced the participants' work with the visualizations, and thus the reported results.

Most of the analysis was completed on the encoded participant statements. The approach of using thick descriptions allowed as much context information as possible to be kept and finally summarized in the results section. However, the findings and interpretations were then also cross-validated with the help of the original participants' statements from the video transcripts. This ensures that the subtleties of the original discussions and statements were not lost through the coding and summarizing process.

## 4. Data analysis and results

The following subsections summarize the results of the qualitative data analysis. We have endeavored to retain as much content as possible. However, in the interests of brevity, the reported content was selected with a view to minimizing repetition and in the context of discussion depth and relevance to the participant groups.

### *4.1. General statements*

The domain expert group with data knowledge (DE1) discussed the visualizations most intensely (with on average most relevant statements per participant), followed by the domain experts without specific data knowledge (DE2). The PhD students' (ND1) and the senior researchers' (ND2) discussions of the visualizations resulted in a similar but smaller numbers of statements. Figure 6 compares the average numbers of statements per participant between the different participant groups.

**DE1**: The participants of group DE1 were most interested in the data analysis. This interest was manifest in that they looked for out-of-the-ordinary movement in environmental contexts, such as high water flow. Their stated aim was to learn more from an already well-known data set. They tended to restate and reconsider their knowledge about the data, the data collection process and their previous analyses. However, they often dismissed their findings and interesting aspects of the visualizations they reported by, for example, belittling them or citing context information that makes them seem unlikely.

**DE2**: The general statements of group DE2 were mainly questions about the characteristics of the data. They commented on their difficulty in effectively working with the visual overviews as they felt they did not have enough background knowledge about the specific data set. However, they were generally interested and did see patterns in the visual overview. They stated they would like to explore the patterns further, for example, by linking them to ecological or biological facts.

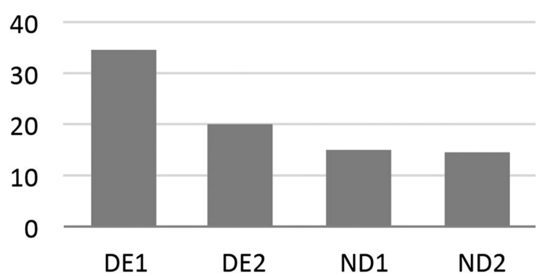


**Figure 6.** Indication of the average numbers of relevant statements per participant in each participant group (DE1–ND2) to gauge discussion intensity. The data experts (DE1) discussed the visualizations most intensely, followed by the domain experts (DE2). The PhD students (ND1) and senior researchers (ND2) discussed least intensely in terms of average number of statements.

**ND1**: The participants of group ND1 had very few questions or comments about the data and data collection in general. As described in the sections below, they queried and discussed the data and their collection in relation to specific visual overviews.

**ND2**: The general statements of group ND2 were mainly questions targeted at learning about the data, the data collection, environmental contexts and the original purpose of the data as well as specific research questions. Group ND2 wanted to add contextual information, such as additional environmental data sets, to all the visual overviews without giving specific information about how and why to do so. ND2 was the only group stating that visual data mining is 'hard work', and especially so as no specific questions needed to be answered. They remarked that it is difficult to show all the data in one view and that there is always suppression of some aspects.

### 4.2. Visual overview V1

**DE1**: DE1 participants did not see anything new in V1. They had worked with a similar data layout already and could relate what they saw to their contextual knowledge about the fish and the data capture. DE1 participants remarked on the arbitrariness of the colors (which they had originally chosen). However, they did point out some interesting aspects in the visualizations, but again discarded them as not important or as visual artifacts; for example, the thicker lines that may be visible when two fish are at the same time in the same river zone.

**DE2**: DE2 participants used V1 to learn about the data set and reported their learning mostly as findings. They identified some problems, such as the line ordering, the large number of zones and the resulting surfeit of different colors. DE2 participants argued that V1 was a dense display that is hard to follow and extract movement or patterns. They misinterpreted that the number of fish increases over time.

**ND1**: ND1 participants learned about the data set and the data collection. Their findings were mostly framed as questions. They discussed extensively the problem of visual impression, for example, the effect of different colors or the dense arrangement of lines. The participants found that the encoding of space through colors was not optimal, and found information about fish movement difficult to extract.

**ND2**: ND2 participants used V1 to learn about the data and reported their process either as findings or as questions. ND2 was the only group reporting, based on V1, that fish were mostly stationary, that is, lines do not change color often. But ND2 participants also reported that the colors were not ideal and that it was difficult to follow the lines. They stated that the ordering of the lines could potentially cause problems through resulting visual artifacts that could lead to misinterpretation. The participants suggested several improvements, such as a more detailed time axis, ordering, zooming or highlighting, but cautioned that interaction may not be useful because data are unstructured.

In summary, the participants across the groups reported the strengths of visualization V1 as including enabling learning about the data set (especially about the fish tagging events), when batteries expire, the number of fish over time and fish location; that it contained the most dense data and that it was easily understood. The main reported weaknesses included that the river was subdivided into too many zones; the color encoding for those river zones was not optimized and information about movement was difficult to extract (to do so, single horizontal lines needed to be followed and each color-change noted).

### 4.3. Visual overview V2

**DE1**: DE1 participants noted many findings about seasonal patterns and how these relate to the knowledge they already had about fish and the data set. They questioned patterns that to some degree contradicted their previous knowledge. DE1 participants found new interesting patterns, but cautioned that they had not investigated those previously, and that those patterns might also have been found using other methods. DE1 participants asked for improvements in relation to previously used visual overviews (e.g. showing single fish or groups of fish) and proposed ideas for exploring the data further (e.g. overlaying environmental context data). This indicates it could be advantageous to combine geographical visualization techniques with the visual overviews developed in this research.

**DE2**: The participants of DE2 worked most intensively with V2 and reported many findings in relation to seasons or specific events, such as floods. The participants stated they would like to explore those findings further, for example, by overlaying additional data or being able to highlight data or mark seasons.

**ND1**: ND1 participants identified patterns in V2 and tried to find some explanations for these patterns based on their limited domain knowledge. The participants in ND1 suggested further explanations for the patterns, for example, calculating correlations for seasonal patterns over the years. They expressed the view that it was not obvious what might be a visual artifact and what might constitute an interesting fish movement pattern. The group's suggestions for improvement included the overlay of a single or a group of fish and the encoding of the density of movements.

**ND2**: The participants of ND2 found some seasonal patterns in V2, but remarked that fish are mostly stationary. They also tried to find some explanations based on their limited context knowledge. The group did not suggest improvements, but did question the usefulness of the visualization for showing interesting information.

In summary, the subjects reported as strengths of visualization V2 that it focused on movement; and it allowed the analysis of movement patterns over time, for example, seasonally. By contrast, the subjects reported as the primary weaknesses of V2 that the visualization was missing information about the direction of movement; that the linear arrangement of river sections did not exactly correspond to the tree-like structure of the river, potentially leading to visual artifacts; that the density of moving fish was difficult to judge and that additional layers of contextual data could be useful in enabling EDA.

### 4.4. Visual overview V3

**DE1**: DE1 participants, influenced by previous research questions, examined the bars to learn about net gain/loss of fish in the river zones. They stated they would like to see net gain/loss directly and also to see proportions of fish represented. The group identified that movement to adjacent zones was prevalent, as would be expected, and thus did not examine this information further. The participants said that the visualization looked easy to understand, but matching the colors was tricky (they suggested that a flow chart might help). They also observed that the different colors together with bar length evoked different impressions of importance.

**DE2**: DE2 participants mainly identified problems or had questions, especially in connection with the color usage, the suppression of time and the direction of the bars. The participants suggested adding time and changing the visual encoding of the data.

**ND1**: ND1 participants reported that fish movement was mainly between adjacent zones and interpreted the bar directions as indicating movement direction. ND1 participants also suggested adding time to see seasonal or yearly patterns and to generally avoid color. Using the visual overviews, the participants made the finding that some fish were skipping zones.

**ND2**: ND2 participants asked for the visual overview to include net gain/loss or proportions of fish. The participants in ND2 thought that the color choice was problematic as some colors were too prominent or could be misinterpreted. They commented that time was suppressed, bar angles did not indicate movement direction and the river layout was not to scale. ND2 participants also found that fish skip zones and remarked that they did not have enough context knowledge to analyze the data represented.

In summary, the groups identified the following strengths in visualization V3: that it offered information about spatial location even though it was schematic and not to scale; that it was easy to understand and made fish skipping zones visible and that it suppressed time. The main weaknesses identified by the groups included that it seemed too simple; that it did not contain enough information for an in-depth analysis and that it suppressed time.

### 4.5. Visual overview V4

**DE1**: DE1 participants took some time to understand this visualization, but then worked intensively with it. In particular, the group identified the symmetry of up- and downstream movements and the patterns of fish movement beyond an adjacent zone as interesting. The participants stated that they would like to filter the fish data, zoom in, to add time and to add additional environmental data in order to learn what may be the cause of long-distance movement. The participants also found the visualization somewhat difficult to read, as the river system is not completely linear.

**DE2**: The participants in DE2 also remarked on the symmetry of up- and downstream movement. They remarked that there seemed to be progressively darker shading (more fish) toward the center of the visualization, something that they argued should be explored further. The participants found that the visualization easily showed the whole data set and allowed the exploration of fish movement between zones. They also identified that time was suppressed and that the river system was not completely linear.

**ND1**: ND1 participants found a balance between up- and downstream movement. They suggested adding proportions or calculating the correlation of the symmetry.

**ND2**: Group ND2 looked at this visualization but did not discuss it. They spent most time discussing the other visualizations. ND2 participants only remarked upon the suppressed time in V4.

In summary, the groups' comments on visualization V4 highlighted the following strengths: that the visualization focused on movement from one river zone to another; and it allowed judgments about net gains/losses through the a/symmetry of the display. On the other hand, the groups identified the following weaknesses: that the tree-like structure of the river system was mapped to a linear structure; and that time was suppressed.

### 4.6. Misinterpretations

A prominent issue discovered in the groups' statements was the lasting impression of misinterpretations of the visual overviews. For example, in visual overviews V1 two groups, domain experts without data knowledge (DE2) and PhD students (ND1), initially reported that it looked as if there were more fish tagged toward the end of the six-year monitoring period. In fact, as the participants were informed in response to this observation, the number of tagged fish was lower in the first year, but was similarly high for all the following years with some seasonal variation. However, the idea persisted and came up again in the discussion of visualization V2 as well as occasionally for V3 and V4. Additionally, the two color-vision impaired participants became interested in the idea of the bar directions in visualization V3 being meaningful, even though bar directions did not indicate the direction of movement. Since those participants have difficulties interpreting color, their reaction of assuming meaning in another encoding is sensible. However, it is problematic if false interpretations persist with the user. It remains an open question as to how to follow up on findings to overcome such apparently valid but incorrect impressions.

### 4.7. Summary and preferences

Comparing the relative numbers of relevant statements (as per coding scheme, cf. Section 2.5) per visual overview type and participant group (see Figure 7), we found that the four groups used the different visual overviews with varying intensities. DE1 participants made only minimal use of V1, a visual overview similar to the one they were already using. But the two groups without domain knowledge (ND1 and ND2) more heavily used V1. These participants reported that it helped them in getting to know and understand the data set better. V2 seemed most popular with the data and domain experts (DE1 and DE2). These participants saw patterns in the fish behavior and either related them to their knowledge (DE1) or suggested further analysis (DE2). V3 was similarly popular with all groups, with the spatial layout of the visualization being commented upon especially positively. DE1 made most use of V4, while ND2 almost completely ignored it. It should be noted that all groups were introduced to the visual overviews in the same order (V1 first, then V2, V3 and V4 last) and often, but not necessarily, discussed the visual overviews in this order as well.

The differences in usage were to some degree also reflected in the participant groups' preference statements, but were also contradicted in some cases. DE1 participants

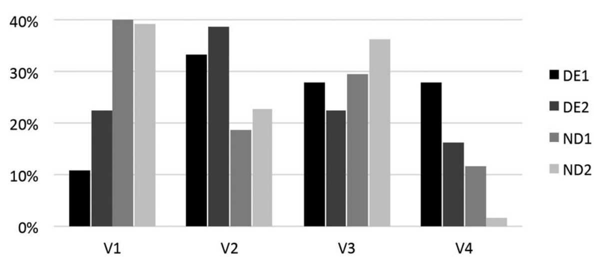


**Figure 7.** Comparison of relative numbers of statements per visualization type (V1–V4) and participant group (DE1–ND2). The differences between the number of statements of the different participant groups are significant at the 99% confidence level ($\chi^2 = 75.73$, df = 12, $p$-value < .001).

commented that V2 was easier to read than V1 and that movements from/to river zones were better grasped from V3 (easy, quick overview based on spatial layout) than from V4. However, they stated a strong positive preference especially for V2 and V4.

Group DE2 not only found V3 and V4 easier to read and follow than the others, but also stated that the usefulness depends on the questions being asked, as those visualizations suppress time. The results indicated that V1 was only moderately useful for general information. The participants identified that V1 was the only visualization that did not focus on the most interesting information (the movement of fish). This movement information was regarded as most useful for answering biological questions. DE2 participants stated they would use V1 for getting an overview, but then work with V2 and V4. But V3 was positively commented upon for its potential use for communication purposes to a wider audience.

Participants of ND1 had similarly strong preferences for V1 (quick overview, most data, but somewhat difficult to read and potentially not so useful), V2 (focus on movement) and V3 (spatial layout, easy to read). V4 got mixed preference statements from ND1 participants, interestingly the visualization was easy to understand, but did not help to identify patterns.

Group ND2 had mixed preferences for V1. These participants argued that V2 was useful as a starting point to learn about the data and its density, but it was overwhelming. ND2 participants had a strong negative preference for V2 (seemed complex and not useful). They also tended to judge V3 negatively as easy to use, but not containing much information. The participants thought the impression of movement evoked by V3 was not reflected in V2, where most fish seemed stationary. ND2 participants stated they would like to use other means to explore and explain the data further (without giving more details about specifically what means).

As described in Section 2.1, the visual overviews were not specifically optimized. The results suggest that the importance of the different data encodings used in the four visualizations varies between the groups. While, for example, the PhD students (ND1) discussed problems with the colors in V1 and V3 extensively, the data experts (DE1) only mentioned that they were familiar with those colors, although other colors could be used. Similarly, for other visual encodings, including the bar length and direction in V3 and partial linear ordering of river zones in V2 and V4, some participants commented on them; others did not appear perturbed.

## 5. Discussion and conclusions

All the visual overviews received strongly positive preference statements from at least one group: each type is apparently useful for answering some questions, learning about and interpreting different aspects of the fish data set. However, we interpreted from our results that visual overviews are different in their usefulness according to the degree of expertise found. More specifically:

- Visual overviews were especially useful for the domain experts without specific data knowledge (DE2). This group reported hypotheses, engaged strongly with the visual overviews and proposed ideas for further analysis of the data.
- Visual overviews were also useful for the PhD students (ND1) albeit for slightly different reasons. Some of their ‘findings’ related to the specific characteristics of the data set or the data collection that they did not previously know about. However, they also

reported findings about fish behavior (the same or very similar as the ones reported by the domain expert groups).
- The visual overviews were to a lesser degree useful for the data experts (DE1). Although this group was highly motivated to learn more about the data, engaged with and discussed the visual overviews most intensively (cf. Figure 6), they always instantly compared everything they identified with their existing context knowledge. DE1 participants seemed to have difficulties in accepting or exploring aspects of the data that did not directly relate to their current context knowledge even if that exploration might potentially extend their knowledge – an explicit goal of EDA.
- The visual overviews were least useful for the senior researchers in this study (ND2) in terms of gaining insight into fish behavior. However, ND2 participants engaged with the visual overviews to learn more about the specific data set, the data collection and potential questions that should be answered.

From these results, we concluded that visual overviews have two distinct areas of usefulness: (1). insight and hypothesis formulation, and (2). learning about the data itself and data collection.

### 5.1. Insight, hypothesis formulation and expertise

Thus, contrary to the expectations as raised by Tukey's (1977) proposition of the importance of domain knowledge in EDA, the results of the data experts (DE1) indicated that detailed knowledge about a specific data set can limit a participant's openness to visual exploration. Even though the data experts were explicitly interested in learning more about the data set, they were either hindered by what they already know about the data or were not convinced in using visualizations to support EDA.

By contrast, the domain experts without data knowledge (DE2), as well as to a lesser extent the PhD students (ND1), used the visual overviews to gain insight into fish behavior and to develop hypotheses for further analysis of the data set. The levels of context knowledge for these two groups were markedly different. The domain experts were certainly more confident in their findings, since they could relate them to domain knowledge. However, the PhD students reported many similar hypotheses and ideas for further analysis. Both groups frequently commented upon patterns and distinctive features in the visual overviews. The PhD students used these comments to ask questions, while the domain experts brought in their context knowledge.

This raises further research questions into what visualization tools, if any, best support data experts in EDA. It might be helpful to data experts to design visual overviews that are richer in content, possibly interactive and multilayered. Another approach would be to *not* make it immediately clear that the experts are viewing their own data. In this way, the experts would need to look for patterns or distinctive features in the visualizations first in a 'context-free' fashion, only later would visualizations then promote experts bringing in their context knowledge. The observation that the data experts discussed most findings and suggestions with V4 potentially supports this later idea. The data experts took some time to understand the visual overview, as it was completely new to them. But after familiarization, the data experts started looking at patterns and discussing different explanations.

In short, we believe there may be value in visual representations of data that enable experts to approach their own data with naïve, fresh eyes, and visualization techniques which allow a richer exploration of the data should be further investigated. As reported by Cartwright, Williams, and Pettit (2007), different visualizations may suite different users, so it is perhaps not surprising that visual overviews are not equally preferred by all participating groups.

### 5.2. Learning about the data itself

Additionally, the PhD students (ND1) and especially the senior researchers (ND2) used the visual overviews intensively to learn about the data itself and the data collection process. The visual overviews seemed to provide them with an engaging means to learn about metadata.

However, the visual overview's usefulness for effective metadata communication in different application settings needs further research. The apparent usefulness in this study could have been facilitated by the focus group methodology where the group moderator answered questions regarding the data and data collection.

### 5.3. Future work

Our results indicate that analysis by non-experts has the potential to result in new insights into the data. It is a matter for further research to ascertain whether in fact these insights would only educate naïve users in trends that the data experts already know about, or would lead to truly new knowledge.

Additionally, we conclude that data experts may require either more sophisticated visualizations that include interactivity and contextual information such as geography and environment layers, or visualization designs that initially hide the data context to allow them to see their data with fresh eyes.

Another option would be to use a larger number of visual overviews with randomized visual encodings based on the structure of the data. This may lead to increased visual artifacts in visual overviews, and thus require strong adherence to the 'skepticism' principle of EDA (Hartwig & Dearing, 1979). However, over the course of using several visualizations and rechecking findings, such an approach may help data experts gain insights from their data. Importantly, it may also provide a means to overcome the problem of lasting impressions of misinterpretations (cf. Section 4.6).

The visualizations tested in this study were all static. Additionally, after hypotheses generation, the testing of those with further visualizations and interfaces would be a natural next step. Some participants noted the need for further exploration of the data, especially also by using interactive and connected visualizations. Interactive displays including overview and detail visualizations offer a variety of opportunities for further research regarding the usefulness of them.

## Acknowledgements

We would like to thank all the study participants, without whose time and commitment this research would not have been possible.

## Disclosure statement

No potential conflict of interest was reported by the authors.

## Funding

This work was supported by the Australian Research Council [DP120100072].

## Notes on contributors

***Susanne Bleisch*** is Professor of Geovisualization and Visual Analytics at the FHNW University of Applied Sciences and Arts Northwestern Switzerland. From 2012 to 2014 she was a postdoctoral researcher in the Ambient Spatial Intelligence Research Group at the University of Melbourne. She studied Geomatics Engineering and obtained a PhD in Geographical Information Science from City University London in 2011.

***Matt Duckham*** is Professor and Associated Dean Geospatial Sciences in the School of Science at RMIT University, Australia. Before moving to RMIT University in 2015, he was Professor of Geographic Information Science at the University of Melbourne, Australia. From 2010 to 2014 he was an ARC Future Fellow working in the area of distributed spatial computing and geosensor networks. Prior to moving to Australia in 2004 he was a postdoctoral researcher at the NCGIA (National Center for Geographic Information and Analysis) at the University of Maine, USA.

***Professor Chris Pettit*** is the inaugural Chair of Urban Science at the University of New South Wales. He is the Director of the Master of City Analytics Program and is Associate Director of the City Futures Research Centre.

## ORCID

*Susanne Bleisch* http://orcid.org/0000-0002-4563-060X

## References

Andrienko, N., & Andrienko, G. (2006). *Exploratory analysis of spatial and temporal data: A systematic approach*. Berlin: Springer.

Baudisch, P., Good, N., Bellotti, V., & Schraedley, P. (2002). *Keeping things in context: A comparative evaluation of focus plus context screens, overviews, and zooming*. Proc. SIGCHI Conference on Human Factors in Computing Systems, Minneapolis, MN (pp. 259–266). ACM.

Bird, T., Lyon, J., Nicol, S., McCarthy, M., & Barker, R. (2014). Estimating population size in the presence of temporary migration using a joint analysis of telemetry and capture–recapture data. *Methods in Ecology and Evolution*, *5*(7), 615–625.

Cartwright, W. E., Williams, B., & Pettit, C. (2007). Realizing the literate traveller. *Transactions in GIS*, *11*, 9–27.

Chen, C., & Yu, Y. (2000). Empirical studies of information visualization: A meta-analysis. *International Journal of Human-Computer Studies*, *53*, 851–866.

Craft, B., & Cairns, P. (2005). *Beyond guidelines: What can we learn from the visual information seeking Mantra?* Ninth International Conference on Information Visualisation (IV'05), London (pp. 110–118). IEEE.

Few, S. (2009). *Now you see it: Simple visualization techniques for quantitative analysis*. Burlingame, CA: Analytics Press.

Geertz, C. (1973). *The interpretation of cultures: Selected essays*. New York: Basic Books.

Greene, S., Marchionini, G., Plaisant, C., & Shneiderman, B. (2000). Previews and overviews in digital libraries: Designing surrogates to support visual information seeking. *Journal of the American Society for Information Science*, *51*(4), 380–393.

Grumbein, M. J., & Lowe, P. A. (2010). Focus group. In N. J. Salkind (Ed.), *Encyclopedia of research design* (pp. 501–502). Los Angeles, CA: SAGE.

Haroz, S., & Whitney, D. (2012). How capacity limits of attention influence information visualization effectiveness. *IEEE Transactions on Visualization and Computer Graphics*, *18*(12), 2402–2410.

Hartwig, F., & Dearing, B. E. (1979). *Exploratory data analysis* (Sage Unive.). Newbury Park, CA: Sage.

Hornbæk, K., & Hertzum, M. (2011). The notion of overview in information visualization. *International Journal of Human-Computer Studies*, *69*(7–8), 509–525.

Keim, D. (2001). Visual exploration of large data sets. *Communications of the ACM*, *44*(8), 38–44.

North, C. (2006). Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, *26*(3), 6–9.

Pettit, C. J., Raymond, C., Bryan, B., & Lewis, H. 2011. Identifying strengths and weaknesses of landscape visualisation for effective communication of future alternatives. *Landscape and Urban Planning*, *100*, 231–241.

Plaisant, C., Fekete, J.-D., & Grinstein, G. (2008). Promoting insight-based evaluation of visualizations: From contest to benchmark repository. *IEEE Transactions on Visualization and Computer Graphics*, *14*(1), 120–134.

Shneiderman, B. (1996). *The eyes have it: A task by data type taxonomy*. In Proc. IEEE Visual Languages '96, Boulder, Colorado (pp. 336–343).

Stemler, S. (2001). An overview of content analysis. *Practical Assessment, Research & Evaluation*, *7*(17). Retrieved from http://PAREonline.net/getvn.asp?v=7&n=17

Tobon, C. (2005). Evaluating geographic visualization tools and methods: An approach and experiment based upon user tasks. In J. Dykes, A. M. MacEachren, & M.-J. Kraak (Eds.), *Exploring geovisualization* (pp. 645–666). Oxford: Elsevier.

Tukey, J. W. (1977). *Exploratory data analysis*. Boston, MA: Addison-Wesley.

Wissen, U., Schroth, O., Lange, E., & Schmid, W. A. (2008). Approaches to integrating indicators into 3D landscape visualisations and their benefits for participative planning situations. *Journal of Environmental Management*, *89*, 184–196.

Wood, J., Dykes, J., & Slingsby, A. (2010). Visualisation of origins, destinations and flows with OD maps. *The Cartographic Journal*, *47*(2), 117–129.