Embedded Merge & Split: Visual Adjustment of Data Grouping

Ali Sarvghad *, Bahador Saket *, Alex Endert, and Nadir Weibel

Abstract—Data grouping is among the most frequently used operations in data visualization. It is the process through which relevant information is gathered, simplified, and expressed in summary form. Many popular visualization tools support automatic grouping of data (e.g., dividing up a numerical variable into bins). Although grouping plays a pivotal role in supporting data exploration, further adjustment and customization of auto-generated grouping criteria is non-trivial. Such adjustments are currently performed either programmatically or through menus and dialogues which require specific parameter adjustments over several steps. In response, we introduce Embedded Merge & Split (EMS), a new interaction technique for direct adjustment of data grouping criteria. We demonstrate how the EMS technique can be designed to directly manipulate width and position in bar charts and histograms, as a means for adjustment of data grouping criteria. We also offer a set of design guidelines for supporting EMS. Finally, we present the results of two user studies, providing initial evidence that EMS can significantly reduce interaction time compared to WIMP-based technique and was subjectively preferred by participants.

Index Terms—Data Visualization, Direct Manipulation, Embedded Merge & Split, Data Grouping, Embedded Interaction

1 Introduction

Data grouping is the process in which relevant data is gathered, simplified, and expressed in a summary form [10, 24]. For example, if we have data about a set of cars, we can group their prices into a smaller number of price groups. Alternatively, we could group the cars based on their models or manufacturers. The goal for data grouping is to get more information about particular groups based on specific variables such as age, ethnicity, or income. Data grouping is a commonly used operation in data visualization. Grouping data and representing it visually is especially valuable in exploratory analysis [16, 48]. It helps people to effectively understand the underlying distribution of data and investigate patterns and relationships [15].

Many popular visualization and data analysis tools such as Tableau [1] support *automatic* creation and presentation of data groups using the inherent groups for categorical variables or arbitrary groups (e.g., bins) for numerical variables. When exploring data, users may need to *adjust* the default data groupings created by the tool according to their evolving needs and interests [48]. For example, an HIV researcher investigating the relationship between HIV patients' age and risk-factors may discover that patients in 20-30 and 30-40 age groups exhibit similar risk factors (e.g., needle sharing), hence deciding to combine the two groups creating a single 20-40 group (while keeping the rest of the age groups intact).

It is currently non-trivial to adjust data grouping criteria in many of the existing visualization and data analysis tools. In some tools such as Tableau, adjustment of data grouping criteria is supported using the Window/Icon/Menu/Pointer (WIMP) model which requires specific parameter adjustments over several steps. For example, Figure 1 shows the sequence of steps required by Tableau Public [1] for changing the default binning of the variable Age. The WIMP-based model can incur extra execution and cognitive costs especially as the set of available operations increases [32, 37, 42]. In other data analysis tools such

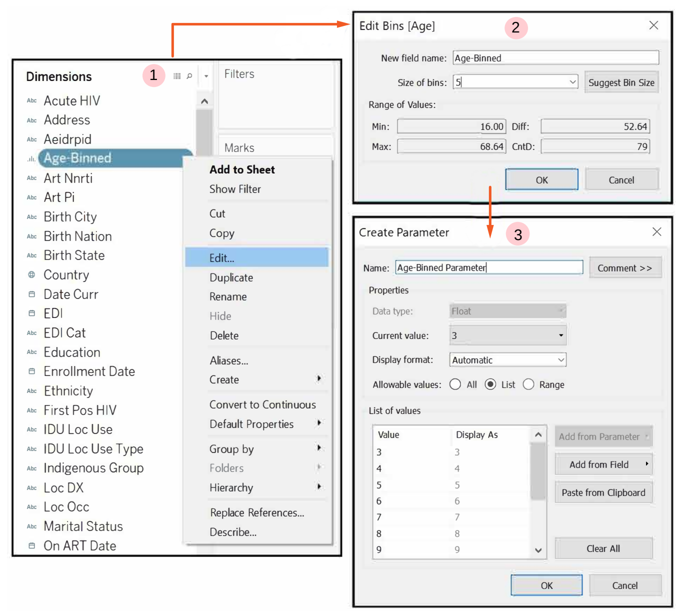

Fig. 1. Screenshots of Tableau Public V10.1 illustrate the sequence of actions required to change binning criteria. **1)** Users first need to select the variable and then select the Edit command from the pop-up menu. **2)** In the Edit Bins dialogue, users can input new size for bins. **3)** Users might also move to the next dialogue for further customization of binning.

as SAS [5], adjustment of data grouping criteria is mainly supported through programming. For example, to bin variable Age in SAS using unequal bin ranges, a user needs to programmatically select the variable Age, define the cut points for each bin range, execute the code, and visualize the results. Thus, the lack of an *intuitive and fast* interaction technique for adjustment of data groupings can increase execution costs and impede or even inhibit exploration. This is in contrast to the nature of exploratory analysis that, as noted by Card et al. [16], thrives on iteration and speed of exploration.

In this work, we introduce **Embedded Merge & Split (EMS)**, a new "embedded interaction" [42] technique for adjusting criteria of data groupings that are represented using linear axis encodings (e.g., width). EMS enables users to directly interact with visual glyphs and directly manipulate the visual encoding channel that is utilized for representing data grouping. For instance, a user can drag and extend the width of a bar in a histogram to increase the range of the values presented by the bar. In response to the user's alteration of visualization, the system reconfigures new grouping of data values and reconstructs the view to new specifications.

*Ali Sarvghad and Bahador Saket contributed equally to this work.

- *Ali Sarvghad is with University of California San Diego. E-mail: asarvgha@ucsd.edu.*
- *Bahador Saket is with Georgia Institute of Technology. E-mail: saket@gatech.edu.*
- *Alex Endert is with Georgia Institute of Technology. E-mail: endert@gatech.edu.*
- *Nadir Weibel is with University of California San Diego. E-mail: weibel@ucsd.edu.*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

We integrated the EMS technique to Avantgarde, a visual data analysis tool currently used by HIV researchers at the University of California San Diego (UC San Diego). We then carried out two studies to measure the effectiveness of the EMS technique. In the first study, we conducted a comparative experiment (Tableau vs. Avantgarde) with 12 participants. Results of the first study indicate that people were significantly faster in merging, splitting, and changing the number of groups using the EMS technique. In the second study, we ran a qualitative analysis with five HIV researchers to collect qualitative feedback and observational data on how experts perform adjustment of data grouping. Results of the second study show that participants found EMS consistent with their mental model, fluid, and easy to use and learn. This study also surfaced some of the potential challenges in implementing EMS. Our main contributions in this work are:

- Introducing Embedded Merge & Split, an embedded interaction technique for adjusting data grouping criteria.
- Presenting results of two studies (quantitative and qualitative), providing initial evidence that the EMS technique was subjectively preferred and quantitatively comparable to traditional WIMP-based technique.
- A set of design guidelines to indicate how future work might adapt the EMS technique to other visualization types.

2 Motivation

The motivation of this work stems from an ongoing project in which we have been collaborating with HIV researchers at the University of California San Diego for 19 months to design and build solutions for supporting exploratory analysis of HIV data. In one particular type of analysis, researchers create multiple univariate distribution graphs (histograms and bar charts) at the same time to investigate patterns and relationships of various demographic (e.g., Age), social (e.g., Marital-status), clinical (e.g., Count of T-Cells) and geographic (e.g., Zip-code) factors that impact HIV transmission. This is a highly iterative process and researchers frequently update and adjust groupings according to their evolving interest in data.

To adjust data grouping criteria, these HIV researchers currently use different set of tools and programming languages including R [4], SAS [5], and ArcGIS [3]. According to them, the current interaction model for adjusting groupings is rather cumbersome, time-consuming, and occasionally error-prone. For instance, one of the HIV researchers mentioned that: "*...[To change binning] I need to open the menu and manually enter the new age ranges. It's kinda **tedious***." In another example, a researcher mentioned that "*I wish to **quickly** combine patients in various early and late treatment groups...I have to write code and execute it. [...] another thing is that if I made a **mistake** like forgot a semicolon, I would only know when it's done [i.e., after execution of code].*" Based on our observations as well as researchers feedback, adjustment of data groupings is rather cumbersome in existing analysis tools. Depending on the data analysis tool used to group data, these researchers either need to write and execute scripts or go through multiple GUI widgets that are provided on a separate control panel.

In response to these challenges, we aimed to design an interaction technique that enables *intuitive and fast* adjustment of data groupings criteria. While the motivation of this work stems from challenges raised by the HIV researchers, we emphasize that data grouping has a widespread application in several domains (e.g., machine learning [33], databases [18], and genetics [47]).

3 Background

Previous research has mainly focused on designing/optimizing algorithms for grouping data [9] and exploring applications of data grouping in different domains (e.g., [13, 18, 19, 33, 47]). However, relatively less work has been done to investigate and improve interaction techniques for adjusting and customizing data groupings in visualizations. In this section, we first briefly discuss applications of data grouping in different domains including data visualization. We then present the primary interaction technique used in the majority of existing visualization tools for adjusting data grouping. We then discuss applications of embedded interaction paradigm and explain how this paradigm inspires our proposed technique.

3.1 Data Grouping

Data grouping is an operation with widespread application in several domains such as machine learning (e.g., [13, 33]), data visualization (e.g., [8, 48]), database (e.g., [18]), and genetics (e.g., [47]). There are multiple ways to group raw data prior to using it. For example, database languages provide us with commands such as `Group By` to group the data by one or more columns. Libraries such as D3 enable us to group our data using methods such as `nest`, before visualizing the data. We can also use interactive data wrangling tools [27] to group our data before using it. In all these cases, we apply groupings to summarize our data in a specific form based on specific data variables.

3.2 Data Groupings in Existing Visualization Tools

Many existing visualization tools such as Tableau [1] and MS Excel support creation and presentation of grouped data. In response to user specifications, these tools perform data grouping and present the data. For example, in Tableau, we can specify our interest in having a visualization that has patients' ages on the x-axis and their population on the y-axis. In response, the system automatically groups patients' ages into a smaller number of age bins and represent the grouped data using a histogram visualization. The system automatically performs data grouping based on statistical properties and inherent classes of data and presents it.

As previous work [48] indicates, during visual data exploration processes, users constantly need to adjust the predefined data groupings created by visualization tools based on their evolving needs and interests. However, a majority of the existing visualization tools such as Tableau [7] and Spotfire [6] require users to adjust data groupings by going through a broad set of menus provided on control panels (WIMP-based technique). The WIMP-based technique can incur extra execution and cognitive costs especially as the number of available operations increases [32, 37, 42]. Interfaces with a large number of execution steps may deter efficient use by inducing additional cognitive loads [32]. Advanced adjustment of data grouping is only supported by highly specialized data analysis softwares (e.g., SAS [5]) and visualization authoring tools such as Lyra [43]. However, utilizing this class of tools typically requires advanced knowledge of visualization design and/or programming.

3.3 Direct Manipulation Interfaces

Direct manipulation interfaces support performing direct actions on the visual objects of interest [44]. Actions in direct manipulation interfaces are simple and support continuous flow of interaction, and immediate visual feedback is provided in response to physical actions. For instance, dragging a slider to navigate a timeline is a form of direct manipulation if the visualization updates in real time. Beaudouin-Lafon defines interaction instruments as mediators between a user and an object of interest [12]. For example, in interactive data visualizations, sliders can be used as instruments for filtering. A large body of previous work [21, 29, 34, 37] has highlighted the necessity for minimizing the distance between the interaction source and the target object.

Recently, there has been an increasing trend in enabling users to directly manipulate the visual representation rather than mainly relying on additional graphical widgets such as menus (e.g., [17, 22, 29, 31, 38, 39, 41]). For example, DimpVis is a recent system that allows users to directly interact with the length, angle, and position of the visual representations, as a means for temporal navigation [29]. In DimpVis, users can adjust the height of a bar to see its value at different moments in time. Interactive legends are controls that enable users to filter data by directly interacting with visual glyphs used on the legends [39]. Saket et al. [41] enable users to directly manipulate the graphical encodings used in visual representations as a method for providing visual demonstrations. In another study, Chevalier et al. [17] present Histomages, an interaction technique that enables users

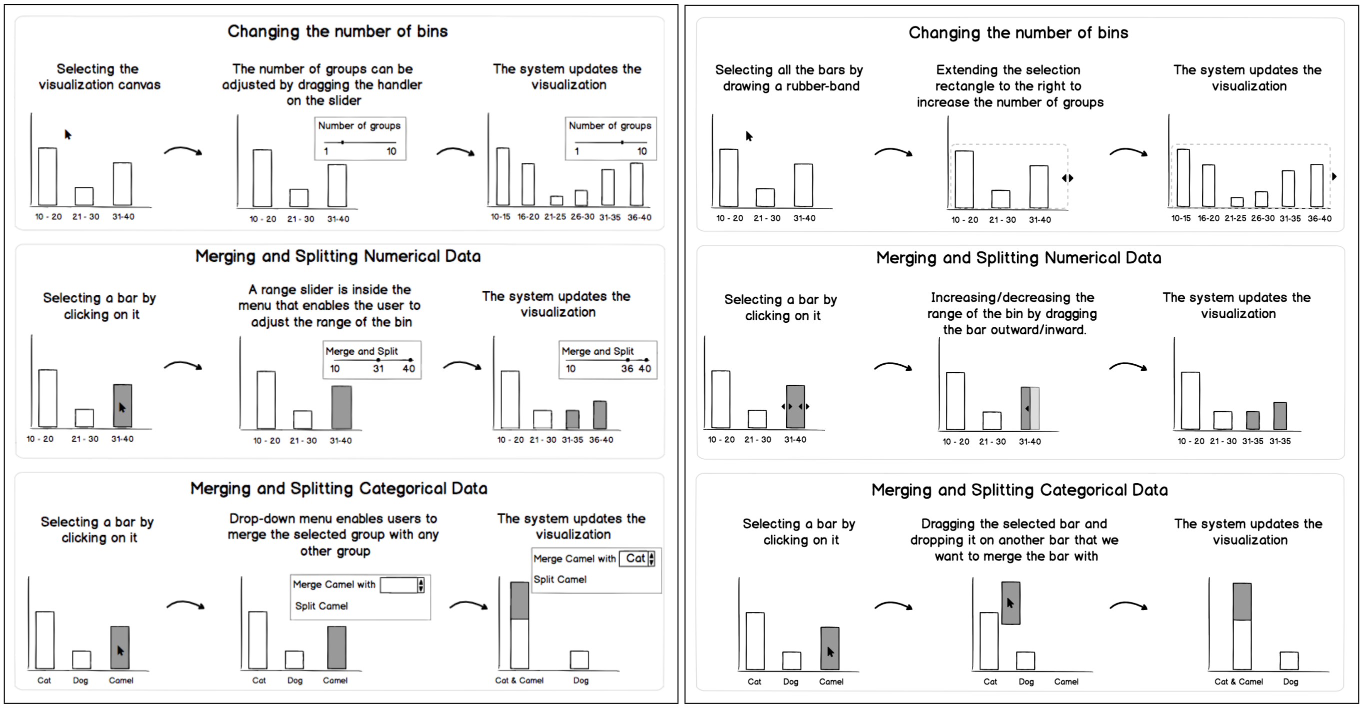

Fig. 2. This figure shows sketches of two of the designs considered during our design process. **Left:** This design uses an in situ pop up to enable users to adjust data grouping criteria. **Right:** This design uses embedded interaction to support adjustment of data grouping criteria.

to change the properties of pixels of an image (e.g., saturation) by directly interacting with the visual glyphs represented in a histogram.

Several projects from the visual analytics community also encouraged direct manipulation of visual representations as means of steering the parameters of the underlying models used in visualization tools (e.g., [14, 22, 28]). For instance, InterAxis enables users to directly interact with the length of a bar in a bar chart to adjust the relative weights of data attributes in the system [28]. Other studies allow users to change the distance between visual glyphs to steer distance and similarity functions [14, 22]. In each of these techniques, adjustment of the interactive graphical encodings implies an intent to change the result of a computation, rather than changing the data value directly.

3.4 Embedded Interaction for Adjusting Data Groupings

Embedded interaction [42] is formulated as an interaction technique in which users directly interact with visual glyphs (bars in a histogram) used in visual representations rather than widgets and menus to perform a task. The goal of embedded interaction is to tighten the gap between a user's intent and the execution of that intent. Embedded interaction is inspired by direct manipulation [44] and instrumental interaction [12] models that aim to make interaction more natural, intuitive, and predictable, resulting in easier to learn and use applications [37].

In this paper, we present a novel technique for advanced customization of data groupings based on embedded interaction. Instead of going through control panels, our proposed technique enables users to customize groupings via direct manipulation of visual glyphs (e.g., changing the width of a bar in a histogram to combine its range with the next bar). The system interprets the changes in visualization, updates the grouping criteria, and presents a new view in real-time.

4 Iterative Design Process

Taking a user-centered approach [36], we started the design process by a grounded investigation of current practices, needs and challenges. We conducted several interviews and group discussions with four infectious diseases specialists at UC San Diego researching HIV transmission in San Diego County. We actively took notes during these interviews and group discussions. We then read through our notes to obtain a general sense of the data and thinking about organization of the challenges these specialists encountered. After reading the notes, we identified the meaningful text segments and assigned a code word or phrase that accurately describes the meaning of the text segment (e.g,. "the large number of steps"). The coding process was an iterative process with three passes by two coders in which the coders developed and refined the codes. As a result of this process, we identified three challenges pertaining to data grouping adjustment:

1. **Menus that adjust data groupings are spatially far from visualizations in the user interface:** A few of the HIV researchers grouped their data and visualized it using a data analysis tool called *ArcGIS* [3]. Similar to many data analysis tools, adjustment of grouping in ArcGIS require users to go through a series of GUI elements (i.e., menus) that are presented on a control panel spatially far from the main visualizations. We noticed that the HIV researchers found this model of interaction challenging because they had to constantly shift their attention from the visual features of interest when adjusting the groupings.

2. **A large number of steps required to carry out the interaction:** To carry out a simple action such as grouping the data based on a specific variable, WIMP-based technique typically requires the following sequence of actions [12]: 1) Selecting the object of interest by clicking it; 2) Selecting a command from a menu or keyboard shortcut; 3) Filling in the fields of a dialog box; and 4) Clicking the OK button to see the result. HIV researchers reported that a large number of operations required by ArcGIS slows down their exploration process.

3. **Lack of fast and incremental feedback:** HIV researchers also reported that neither SAS nor ArcGIS provided incremental immediate feedback as they adjusted data grouping criteria. To verify that their intended changes were made successfully, the researchers need to execute the code (SAS) or press the OK button (ArcGIS) first and then check the outcome.

Based on the challenges raised by the HIV researchers, we iteratively created different designs that would address the challenges and evaluated them through multiple discussions among ourselves (first and second authors) as well as the HIV researchers. Among four design alternatives considered during our design process, two of the designs best met the challenges initially encountered by the researchers. Figure 2 shows two designs that best met researchers requirements.

First design – Our first design uses an in situ pop up to enable users to change the number of groups (first row) and merge and split numerical and categorical data (second and third rows). This design decreases the spatial offset by placing "interaction instruments" [12] in close proximity to visualization. It supports adjustment of grouping by enabling users to merge and split the groups interactively. As users adjust the grouping, the visualization will be updated to provide feedback. However, based on our discussions with the HIV researchers, we decided that this design does not considerably reduce the number of steps required to carry out grouping adjustment in comparison to the WIMP-based model used in current tools. See Figure 2-Left.

Second design – Our second design utilizes embedded interaction [42] to support adjustment of data groupings. In this case, direct interaction with visualization eliminates the need for external menus and dialogues and completely removes the spatial offset between visualization and "interaction instruments" (i.e. visual glyphs). It supports adjustment of grouping by enabling users to merge and split the groups interactively. Finally, this design provides immediate feedback as user manipulate the visual glyphs representing a group. See Figure 2-Right.

The embedded technique (second design) was well received by our target users (HIV researchers) because interaction required no additional use of menus and widgets. For instance, one of the HIV researchers commented that *"this one seems more organic and straightforward."* Therefore, we implemented the embedded technique and conducted an initial usability assessment to gather early empirical user feedback. The following section presents the results of this preliminary evaluation.

4.1 Preliminary User Feedback

We conducted an initial evaluation in a small study with four participants (P1-P4). All the participants (HIV researchers) were familiar with histograms, bar charts and the concept of data grouping. Participants were first given a short tutorial of how the prototype worked, during which they were encouraged to ask questions as needed. After introducing the technique to the participants, they created a histogram showing the distribution of ages of HIV patients. Data was binned using equal 10 years intervals, and we asked each participant to perform six tasks ($2 \times$ merging, $2 \times$ splitting, and $2 \times$ changing the number of groups) to the histogram, while thinking-out-loud. Next, participants created a bar chart showing the distribution of patients' ethnicity information. We asked each participant to perform four tasks ($2 \times$ merging and $2 \times$ splitting) to the bar chart visualization. This was a think-aloud study and participants were asked to verbalize their thoughts.

Overall, all the participants believed that embedded technique was intuitive and easy to understand and use. The participants suggested additional features such as adding some visual aids for selecting very small bars. In addition, they requested a way for selecting and adjusting multiple groups at the same time. We built all the additional features into later implementations of our interactive technique that we are going to present in the next section.

When designing our technique, we initially thought it would be more natural if the system updates the visualization constantly as users are merging or splitting data groups. However, our participants reported confusion resulted from constant updates on the entire view. Thus, we decided to only update the entire view when the user stops dragging. Instead, upon interacting with a bar, we provide a tooltip to show the current range of a group represented using the bar. The tooltip gets updated constantly as the user drags the bar to the left or right to adjust data grouping criteria.

In the next section, we introduce Embedded Merge & Split, discuss a set of general design guidelines for implementing this technique, and provide two example implementations for Bar charts and histograms.

5 Embedded Merge & Split

Embedded Merge & Split (EMS) allows users to directly interact with and manipulate linear axis visual glyphs to adjust data grouping criteria. In response to user's changes of visualization, the system reconfigures new grouping of data values and reconstructs the view to reflect new specifications.

EMS can be primarily used to support three types of adjustments: (1) **Merge groups**, (2) **Split groups**, and (3) **Change the number of groups**. Merge is the act of combining two or more groups completely or partially. Split is the act of dividing a group into smaller sub-groups. Changing the number of groups increases or decreases the number of groups. This adjustment is more common for numerical data.

5.1 Design Guidelines

In creating EMS, we established the following set of design guidelines.

G1: Provide Visual Assistance. Provide signifiers/visual aids [36] that enable users to understand the *where*, *how* and *what* dimensions of their interaction. Where, indicates the location(s) that user can interact with the visualization; how, conveys information about physical action (e.g., direction of dragging); and what, indicates the effect of carrying out the interaction on data grouping (e.g., increase or decrease). Signifiers provide direct perception of possibilities for action without which it may not be immediately obvious to the user how to interact with the visualization to adjust data grouping.

G2: Support Embedded Interaction. Enable users to directly manipulate the visual encoding channel that represents grouping information (e.g., width for a bar chart, angle for a pie chart, etc). Leveraging visual encodings as a method for user interaction has several advantages including not requiring users to shift their attention from the visual features of interest when interacting [42] and simplifying the visualization interface by obviating the need for additional control panels or widgets [34, 37].

G3: Real-time Visual Feedback. Inspired by the design guidelines provided by previous work [36], EMS technique should constantly send back information about what changes have been made and what has been accomplished, allowing users to continue data analysis process.

In the next two subsections, we present the implementation of EMS for histograms and bar charts.

5.2 EMS for Histograms

Histograms are one of the most commonly used visualizations for representing distribution. Histograms are suitable for investigating the shape of the distribution as well as the values [23]. Each bar in the histogram represents a bin where the width of the bar typically shows the range of the values and the height of the bar shows the frequency (count).

Merge – To merge two or more bins using the EMS technique, users first click on a bar to select it. In response to users' action, the selected bar becomes highlighted to visually confirm users' selection. Triangular icons on the outside of the selected bar signify that the bar can be extended to the left or right **(G1)**. Additionally, upon selection of the bar, a label appears on the top of the selected bar to show its current range. A user can drag the triangular icons to left/right to alter (decrease/increase) the number of bins **(G2)**. The system simultaneously recomputes and renders the histogram according to new specifications. The width of the altered bar will change in proportion to the new range **(G3)**. Figure 3-a illustrates the process of merging bins in a histogram using the EMS technique.

Split – This action enables the users to divide a bin into smaller sub-bins. The implementation is very similar to that of the merge, with the exception that inner triangular icons are used to signify and alter the width of the bar. Figure 3-b provides an example of split action with details.

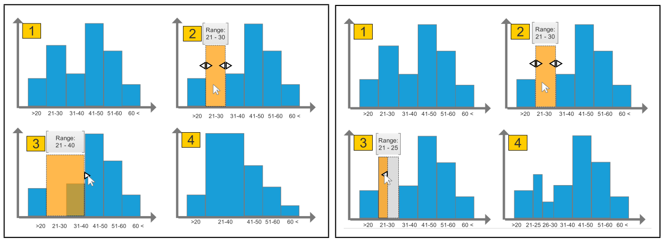

(a) (b)

Fig. 3. Implementation of EMS for histograms. **(a) Merge:** 1) Initial state of bins representing age distribution of a group of people; 2) Selecting a bar by clicking on it; 3) Dragging the bar to the right to increase the range of the bin; 4) Updated histogram. The width of the altered bar corresponds to the new bin range. **(b) Split:** 1) Initial state; 2) Selecting a bar by clicking on it; 3) Reducing the range of the bin by dragging the bar inwards; 4) Updated histogram.

Changing the number of bins – To change the number of bins using the EMS technique, a user draws a rubber-band rectangle to select the histogram. After the selection, the rectangle covers the entire histogram with two triangular icons signifying where and how to interact with the histogram **(G1)**. Additionally, a label appears on the top of the selected histogram to show its current range. The user can drag the triangular icons to left/right to alter (decrease/increase) the number of bins **(G2)**. The system continuously provides feedback about the current number of bins as the user drags the rubber-band to the left or right using a label placed on top **(G3)**. When dragging is stopped, the system computes the new bin ranges according to the number of bins and updates the histogram accordingly. See Figure 4 for more details.

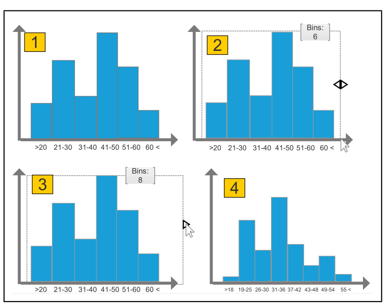

Fig. 4. Changing the number of bins: 1) Initial state. 2) Selecting all the bars by drawing a rubber-band. 3) Extending the selection rectangle to right to increase the number of bins. 4) Updated histogram.

5.3 EMS for Bar Charts

Bar charts are used to compare individual values (e.g., frequency) between several groups [23]. Bar charts use size to encode scalar values in height and position to encode grouping information. Typically, bar charts are used for categorical variables and categories inherent in data are used for grouping values.

Merge – To merge two or more groups using the EMS technique, a user starts the operation by selecting a bar to be merged . Upon doing this, the selected bar will be highlighted to confirm the user's selection. Triangular icons appear in the middle of the bar to guide the user's interaction (**G1**). Selected bar (i.e., group) can be dragged using either left or right signifier towards a target bar to be merged with (**G2**). While dragging the bar, the label at the top is updated continuously to indicate that the desired target is reached (**G3**). When the user drops the selected bar at the target, system recomputes new grouping of values and then reconstructs and presents the bar chart. The combined bars will be shown as a stacked bar. Users can also select multiple bars and combine them with a target. See Figure 5-a for more details.

Split – To split two groups that are merged into a bar, a user first selects the bar that needs to be split. The selected bar will be highlighted to confirm user selection, and signifiers appear in the middle of the bar to guide user interaction **(G1)**. Users can drag the bar upward or downward to indicate their interest in splitting the bar **(G2)**. When the user stops dragging, the system will reconstruct the bar chart, and a bar will be added to present the new grouping **(G3)**. Figure 5-b illustrates the process of splitting bins in a bar chart using the EMS technique.

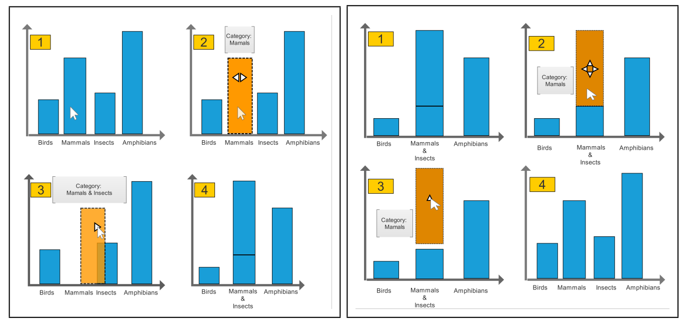

(a) (b)

Fig. 5. Implementation of EMS for bar charts. **(a) Merge:** 1) Initial state. 2) Selecting a bar by clicking on it. 3) Dragging the bar towards a target bar to be merged with. 4) Updated bar chart. The combined bars are presented as a stacked bar. **(b) Split:** 1) Initial state. 2) Selecting a bar by clicking on it. 3) Splitting the group by dragging the bar upward. 4) Updated bar chart. The width of the two new bars, resulting from the splitting of the original bar, are proportional to their ranges.

Changing the number of groups – In histograms, a numerical variable is grouped into a smaller number of bins. With numerical variables, grouping is subjective. For example, while one might decide to group the Age variable into 5 bins where the range of each bin is 10, another might group the variable into 10 bins where the range of each bin is 5. Thus, in histograms, we can change the number of bins by altering the range of each bin (the higher the number of bins, the smaller the range of each bin). On the other hand, grouping of categorical variables is usually performed based one the inherent categories that exist in data (e.g., Animal Types [Mammal, Bird, Reptiles]). Therefore, changing the number of groups for categorical variables is not supported by the current version of EMS.

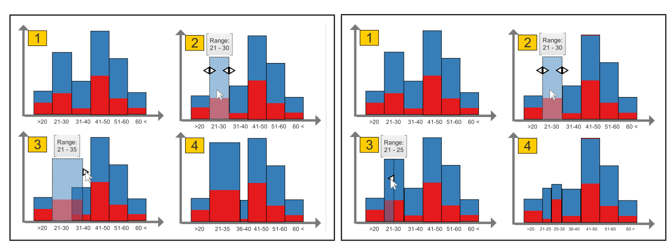

(a) (b)

Fig. 6. Handling color mappings while merging and splitting in histograms. **(a) Merge:** 1) Initial state. 2) Selecting a bar by clicking on it. 3) Dragging the bar to the right to increase the range of the bin. 4) The system colors the merged bins accordingly. **(b) Split:** 1) Initial state. 2) Selecting a bar by clicking on it. 3) Reducing the range of the bin by dragging the bar inwards. 4) The system colors the split bins accordingly.

5.4 Handling Color Mappings

The EMS technique can also support adjustment of data grouping in bar charts and histograms, in which color is mapped to a second variable. For both bar charts and histograms, upon adjusting data grouping (merging, splitting, or changing the number of groups), the system immediately computes the new coloring for the specified groups and updates the visualization. For example, Figure 6 illustrates how the EMS technique handles coloring during adjustment of data grouping.

6 Evaluation

To evaluate our technique, we implemented EMS in Avantgarde (Figure 7), an exploratory data analysis tool currently used by HIV researchers at UC San Diego. Next, we conducted two studies to evaluate effectiveness and usability of our technique. In the first study, we conducted a controlled lab experiment with 12 participants (undergraduate and graduate students) and compared adjustment of data groups between Avantgarde and Tableau Public (see Figure 8). In the second study, we performed a qualitative assessment with five HIV researchers to collect subjective feedback and observational data.

We separated the two studies due to a lack of access to a large number of expert participants, reducing the statistical power of our analysis. In the first study, we used students who provided us with a large enough pool of participants to collect enough quantitative data for statistical analysis. The second study provided qualitative feedback by domain experts who frequently perform data groupings using histogram and bar chart visualizations.

6.1 Quantitative Evaluation

In the first study, we performed a within-subjects controlled lab experiment and compared EMS (using Avantgarde) against WIMP-based (using Tableau Public version 10.5.1) for adjustment of data grouping criteria. Based on feedback collected from our preliminary studies and our observations, we anticipated that EMS would reduce the interaction time for merge, split, and change the number of bins tasks. As such, we hypothesized that users would be significantly faster **(H1)** and more accurate **(H2)** in performing these tasks using Avantgarde.

We used Tableau Public as the baseline condition for two main reasons: 1) it is publicly available and offers interactive support for merge, split, and changing the number of bins for both numerical and categorical variables using a WIMP-based model, and 2) to avoid potential biases in designing an in-house baseline that could affect the internal validity of our experiment.

6.1.1 Participants and Setting

We recruited 12 participants (7 male, 5 female) aged 21-29 years with normal vision (not color blind). Participants were undergraduate and graduate students with science and engineering backgrounds. All the participants were familiar with bar charts, histograms and the concept of grouping, and none had prior experience with interactive adjustment of grouping in a data analysis tool. We used a single workstation and a 24-inch monitor (1920×1200 pixel resolution) for this study.

6.1.2 Tasks

To measure participants' performance, we asked them to perform three types of tasks. Below, we describe each type of task.

- **Merge groups:** Participants were asked to combine two or more groups completely or partially. *For example, can you combine age groups 20-30 and 31-40?*
- **Split groups:** Participants were asked to divide a group into smaller groups. *For example, can you split the age group 40-50 into two 40-45 and 46-50 groups?*
- **Change the number of groups:** Participants were asked to increase/decrease the number of groups. *For example, can you change the number of groups to 5?*

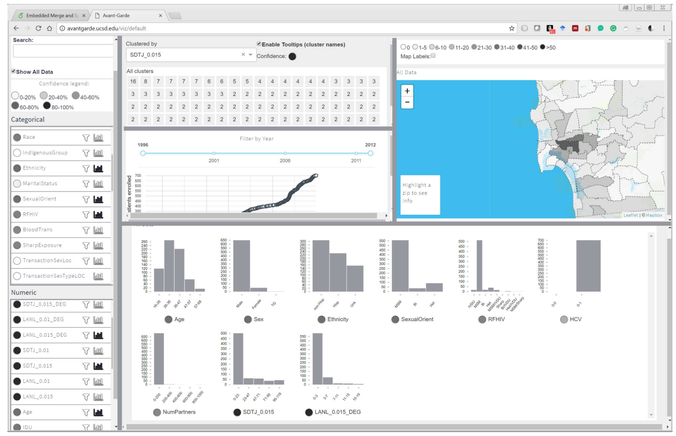

Fig. 7. Avantgarde tool Interface. Avantgarde is an exploratory data analysis tool currently used by HIV researchers at the University of California San Diego.

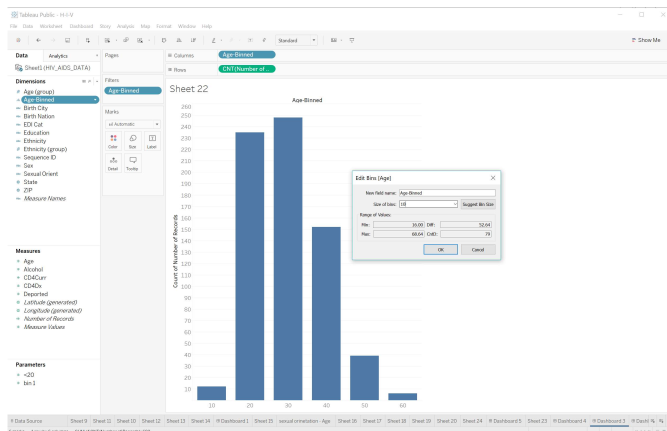

Fig. 8. Screen shot of Tableau Public. Tableau is one of the current leaders in the visualization tool market. Figure shows a user adjusting grouping criteria for the variable Age-Binned.

To create initial visualizations and design tasks, we used a numerical (Age) and a categorical (Ethnicity) variable from the HIV dataset that was made available by our collaborators. The histogram and bar chart visualizations created using the Age and Ethnicity variables contained 10 (each bar representing a different age range) and 7 bars (each bar representing a different ethnicity) respectively.

6.1.3 Procedure

Each evaluation session was started by randomly assigning a participant to either Tableau or Avantgarde (counter-balanced). Next, depending on the tool, participants would receive a 15 minutes introduction into tool, tasks, and data. Training for each tool included the set of features that were necessary for adjustment of grouping criteria. After the introduction, participants were asked to complete six training tasks (2 merges + 2 splits + 2 changing the number of bins) as quickly and accurately as possible. The participants were encouraged to ask questions during this stage (we did not record the time and accuracy during the training session). All tasks were printed on a sheet of paper. Each time the interviewer selected a task randomly and asked the participants to perform the task. The participants were not allowed to move to the next training question unless they answered the question correctly. After the training session, participants immediately proceeded to the main experiment, where they performed 15 tasks (6 merges + 6 splits + 3 changing the number of bins). The tasks were presented in a random order. Before performing each task, participants were given a visualization as a starting point. This way we made sure that all the participants performed each task starting from the same visualization. We recorded interaction time and accuracy for each task. After the completion of the tasks and a short break, the same procedure was performed for the alternative tool.

6.1.4 Data Analysis and Results

To address our hypotheses (**H1** and **H2**), we needed to test how the different tasks were performed using each interaction technique in terms of time and accuracy. To analyze the differences among the various types of tasks, we first calculated separate mean performance time and accuracy for all trials. That is, for each participant, we averaged outcome values of trials for the type of task. We initially planned to take into account both performance time and accuracy in our analysis. However, the participants performed all the tasks correctly using both techniques, so we excluded accuracy from our analysis. Initially, the assumption of normality was not satisfied for performance time. However, the normality was satisfied for log transformation of time values. So, we treated log-transformed values as our time measurements. We then conducted a paired sample t-test to analyze the collected data. The main effect of interaction technique indicates which technique produces the best performance, regardless of the task.

A paired sample t-test indicated that EMS ($M = 0.49, SD = 0.08$) was significantly faster than the WIMP-based technique ($M = 1.03, SD = 0.09$) across all tasks ($t_{(11)} = 17.6, p < .001, \eta_p^2 = 0.89$). Overall, EMS was 0.54 seconds faster than the WIMP-based technique across all tasks. Our results also show that the participants are significantly faster in merging, splitting, and changing the number of groups using the EMS technique. Thus, our results confirm our hypothesis (**H1**). The performance time results are summarized in Table 1.

Table 1. Means of completion times (M) in seconds for the Embedded Merge & Split (EMS) and WIMP-based (WIMP) techniques. The standard deviation (STD) is indicated as well. Significant differences in completion time are indicated by $\star$. Significantly faster results are highlighted in bold.

TASK TYPE	EMS	WIMP	*p-value*	*t-test*	η_p^2
Merge groups	**M=0.45** **STD=0.08**	M=1.02 STD=0.08	$< .001 \star$	$t_{(11)} = 21.1$	0.96
Split groups	**M=0.34** **STD=0.12**	M=1.19 STD=0.09	$< .001 \star$	$t_{(11)} = 19.9$	0.97
Change the number of groups	**M=0.68** **STD=0.05**	M=0.89 STD=0.12	$< .05 \star$	$t_{(11)} = 5.51$	0.75

6.2 Qualitative Evaluation

Our qualitative study had two main goals: (1) collect qualitative feedback on EMS features and design, and (2) collect observational data on how experts perform adjustment of data grouping on bar charts and histograms using the EMS technique.

6.2.1 Participants and Setting

We recruited five expert HIV researchers (4 male, 1 female, all with graduate degrees). They had not participated in our preliminary evaluation of EMS and were also not involved in the design of EMS. They were all familiar with the concept of data grouping, bar charts, and histograms, and had previous experience with grouping with least one data analysis tool (e.g., SAS, ArcGIS, R). All the participants were also familiar and had used the Avantgarde tool before incorporating the EMS technique. During the entire study participants used a computer with 24-inch screen.

6.2.2 Procedure

At the beginning of each evaluation session, the participants received an introduction to EMS in Avantgarde (5 minutes) followed by a practice round (5 minutes). The participants were then asked to use Avantgarde for 30 minutes to perform a short data analysis session using the familiar HIV dataset. We asked participants to think-aloud during the analysis session and audio recorded their thoughts. We also instructed the participants to use bar charts and histograms to investigate relationships and patterns between different factors that impact HIV transmission. In particular, we asked them to adjust data grouping in both bar chart and histogram at least multiple times during their data analysis process. During this phase, we tried to avoid interrupting the participants as much as possible. However, we sometimes needed to remind the participants that this is a think-aloud study and encouraged them to verbalize their thoughts.

This phase of our study concluded with a follow-up interview, in which we asked participants about what they liked and disliked about the EMS technique. This was to allow the participants to convey their feedback and ideas and in order to solicit potentially unexpected insights. We audio-recorded our interview with the participants. We also took notes during the session recording our observations of user actions.

6.2.3 Data Analysis

To analyze the audio material collected during the qualitative study, we fully transcribed data from the interviews and observations. The coders (first and second authors) then read the transcribed data from the interviews. After reading the data, the authors identified the meaningful text segments and assigned a code word or phrase that best explains the meaning of the text segment. During the coding phase, we mainly focused on processes of collecting anecdotal evidence describing participants positive and negative opinion on *ease of learn*, *ease of use*, and *usefulness* of the EMS technique. In the rest of this paper, we use R1 to R5 to respectively denote the researchers one to five who participated in the evaluation.

6.2.4 Results

Overall, all five HIV researchers who participated in this study found the EMS technique easy to use and effective in performing tasks related to adjustment of data grouping criteria. However, a few of them experienced difficulties in selecting very small and overlapping visual glyphs. In the rest of this section, we categorize and discuss the findings of our qualitative study in more details.

Consistency with user mental model

Participants found the design of merge and split interactions consistent with their mental model and expectations. R3 mentioned: *"it feels natural and intuitive to drag the border to make the group larger and smaller. [...] this is what I would expect to happen."* R1 noted that *"breaking up the groups pulling them up is just like breaking the Lego pieces by pulling them apart."* As previous work also suggests [36, 37], consistency and natural mapping between user's intent and the actions required for performing the intent is important in designing new interactions. Consistency also improves learnability and ease of use of an interaction [37].

The design of signifiers were also consistent with user perception of how to use them and what associated outcome would be. For example, R5 mentioned that *"I like the tiny triangles that pop. They are very useful. Telling me to move it to left or right."*

Seamless and fluid interaction

Fluid interaction promotes users to stay in the flow of their analytic process [21]. Based on our observations during the study and participants feedback, the EMS technique supports less deviation from the analysis by reducing time and the number of steps required to adjust data grouping criteria. For instance, R1 mentioned that *"it [EMS] gives me control over this [adjusting grouping]. I can focus more on my analysis."*. R2 noted that *"here you can do it interactively and see the bar chart move over and immediately see you have successfully done it, but in SAS I have to create a new variable, specify what I want to combine, and run it."* R3 expressed that *"[...] this is so much faster than what I do in R."*. Fluid and seamless interaction also increase user's sense of control and engagement [37]. R4 also mentioned that *"it's cool to see that bar chart changes immediately. I don't have to wait to see if I got it right"*. R5 mentioned that: *"most of the churning that I have done has been in SAS, it's nowhere near this user friendly."*

Difficulty in selecting very small targets
R1 and R3 sometimes created bar charts and histograms with more than 20 groups (each bar represented a group). In bar charts/histograms with a large number of bars, the bars were relatively narrow and it was difficult for the participants to select a specific bar with a short height. For example, as one of the participants (R3) noted *"I want to combine A1, A2 and A3, but it's a bit tricky to get A1 because it's tiny."* Although this did not result in an interaction dead-end and users were able to eventually select the target bar, difficulty in interacting with small visual glyphs used in visualizations may cause user frustration and anxiety. Going forward, we envision designing techniques for selecting and adjusting small visual glyphs (we will elaborate this point in more details later in the discussion section).

Difficulty in selecting overlapping targets
When using EMS, R3 mentioned: *"I'm wondering how can I do this [selecting a target] if they [visual glyphs] are on top of each other [overlap]?"* While in our design of EMS for bar charts, two groups will never overlap (they will always be stacked on the top of another), this raises a higher level challenge in adapting EMS where we have visualizations with high information density (i.e. the amount of information encoded versus the amount of unused space [35]). Very high concentration of visual glyphs may hinder the selection and manipulation of objects of interest due to occlusion. This might affect users' speed while adjusting data grouping criteria using the EMS technique.

7 Discussion

Based on our findings we further reflect on EMS, focusing on current limitations, possible solutions, and other possible generalizations.

7.1 Narrowing the Gulfs of Execution and Evaluation

The results of our studies showed that the EMS technique can significantly reduce the time required to adjust data grouping criteria and improve the overall user experience while adjusting grouping criteria. We believe that EMS results in a better usability and user experience by narrowing the gulfs of *"execution"* and *"evaluation"* first introduced by Hutchins et al. [25].

The gulf of execution is defined as difference between the intentions of users and what the system allows them to do or how well the system support those actions. The EMS technique lowers the gap between the intention of the user and how it requires the user to perform the task. EMS superimposes the interaction on the visual glyphs representing groups. Thus, making it easier for users to mentally map their intentions with actions required for performing those intentions. For example, to merge two groups in a bar chart, users can drag a bar representing one group and dropping it on another bar representing a different group. This is very similar to how we group items in a real world.

The gulf of evaluation is the degree to which the design provides representations that can be directly perceived and interpreted in terms of the expectations and intentions of the user. The EMS technique lowers the gulf of evaluation by constantly providing feedback about user changes to groupings. For example, as users drag a bar representing a bin in a histogram to increase the range of the bin, the system updates the label in real-time.

7.2 Selecting Small and Overlapping Targets

As we discussed earlier, selecting small and overlapping targets, and adjusting them is one of the challenges in the EMS technique. One possible design solution to mitigate this issue can be the use of "surrogate objects" [26] (i.e. objects that users can interact with instead of the real domain objects) to select small or overlapping visual glyphs. Upon selecting such glyphs, we can envision multiple ways to enable users to adjust the selected glyphs. For example, we could enlarge the selected glyph temporary to increase users accuracy and facilitate its adjustment of the glyphs. However, the effectiveness of using surrogate objects and temporary enlargement as methods for accommodating selection and adjustment of small and overlapping targets remains to be formally studied.

Difficulties in selecting and adjusting small visual glyphs also exist in other interaction techniques (e.g., [29]) and tools (e.g., [41]) that enable user direct interaction with visual representations. As more studies and visualization tools enable direct manipulation of visual glyphs, this motivates us to encourage the community to investigate different ways to help users in selecting and directly manipulating small and overlapping visual glyphs.

7.3 Enabling EMS for Trellis Visualizations

When using EMS, R1 noted: *"does the rest of the views [other histograms] change when I am merging bins here?"* and R5 asked *"if I combine these [two age bins] here, would they also be combined on the map?"* These comments motivated us to think about an interesting research direction that is expanding the EMS technique to trellis visualizations [23]. Trellis visualizations (also known as small multiples [46]) represent a dataset using multiple graphs of the same type side by side to avoid clutter and ease the process of data exploration. In many cases, in a trellis visualization, the graphs differ according a single variable. Thus, visualizations share at least one data variable. In such cases, it is not clear if the system should apply the adjustments made to grouping criteria locally or it should apply the updates globally to all graphs accordingly. For example, imagine we have a trellis visualization where it shows two bar charts side by side. Both bar charts have the same categorical attribute on the x-axis, but different attributes on the y-axis. Should merging two bars in one bar chart results in changes in that specific bar chart or both bar charts? One possible solution to resolve such ambiguity is to recommend applying changes locally and globally, and let the user decide. An example of this could be visualization by demonstration in which the system recommends set of potential options based on a given demonstration [41].

7.4 Defining Customized Groups

While using EMS to adjust grouping for categorical data, we noticed that few of the HIV researchers wanted to define new groups that did not exist initially. For example, R4 mentioned that *"I wonder if I could create a new group for A-1.0 patients [i.e. one of the categories expressed by the EDICat variable] with only patients that use drugs."* In such cases, these HIV researchers wanted to define these groups based on their domain expertise. The current version of the EMS technique does not provide functionalities for defining and labeling new groups. However, in future one interesting line of work is to investigate intuitive and fast interaction techniques to enable users to define and label new groupings.

7.5 Potential Interpretation Challenges Resulting from Unequal Data Groups

Histograms primarily use height and width encodings to represent the data. Height of the bars is often assigned to a data attribute (e.g., population), so it varies across different bars. However, width of the bars is often the same across the bars (presenting data groups with equal ranges). In the EMS technique, we enable users to directly manipulate the width of bars, as means for adjustment of range of the bins in histograms. For example, we enable users to increase/decrease the range of a specific bin represented using a bar by increasing and decreasing its width. This results in bins with unequal widths (see Figure 3). Representing bars with unequal widths might have advantages and disadvantages. One one hand, it can help users to immediately identify which bins have ranges larger than others by simply looking at the width of bars. On the other hand, this might increases difficulties interpreting and reading the visualization since we are double encoding a single visual glyph (e.g., encoding height of a bar to a data attribute and width of a bar to the range of the bins). While none of the participants in our study reported difficulties in reading or interpreting such bars with unequal widths, more investigations is needed to understand trade-offs in creating unequal groups and its affect on interpretation and readability.

7.6 Generalizing Embedded Merge & Split (EMS)

We implemented and evaluated EMS for bar charts and histograms, providing initial evidence that our technique can facilitate adjustment of data grouping. Bar charts and histogram were selected because of two reasons. First, our target users (HIV researchers) constantly group their data using bar charts and histograms. Second, bar charts and histograms are among the most common visualization types used for representing grouped data [23]. This work is the first step towards exploring the relatively large design space of data grouping in visualization tools. Multiple avenues for future work lie in enriching the EMS technique. We envision expanding the EMS technique to other visualization (e.g., pie chart, treemap) and data types (e.g., hierarchical), and work towards expanding EMS to other devices.

Generalizing EMS requires support for other visualization types (e.g., pie chart, treemap, parallel sets [30]). New questions will then arise, such as how can we enable users to merge or split groups represented using a pie chart or a parallel set? For example, we could enable users to merge and split groups shown as slices in a pie chart by dragging the edge of each slice inward or outward. Additionally, we could apply design guidelines provided in this study to enable adjustment of data groupings in a parallel set visualization. Figures 9 and 10 demonstrate potential ways of implementing EMS to pie charts and parallel sets. While EMS can be implemented in different ways in each of these visualization types, the underlying design guidelines remain the same.

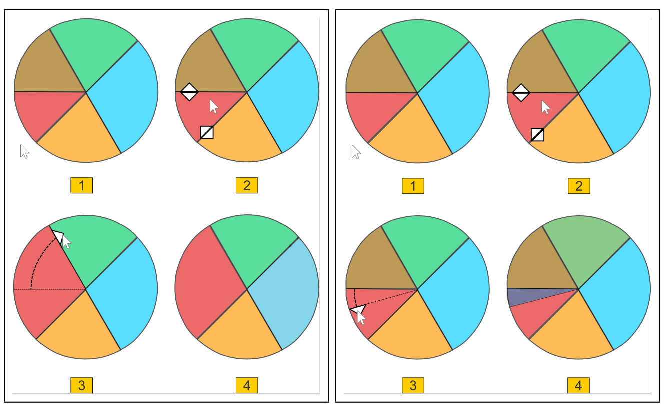

(a) (b)

Fig. 9. A possible implementation of EMS for pie charts: (a) Merge and (b) Split. Similar to our implementation of EMS for bar charts and histograms, signifiers indicate where, and how to interact with the visualization.

Further, we need to investigate ways to apply EMS to visualizations that use data types other than tabular data (e.g., enabling adjustment of grouping of hierarchical data represented using a treemap). Supporting EMS for hierarchical data might require a new set of implementation strategies and design decisions. For example, should we enable users to merge or split two groups that belong to different branches? If so, what does the outcome look like? How are the changes to the data structure communicated to the user? If not, what should our design strategies be to prevent users from merging and splitting groups of different branches?

7.7 Enabling EMS on Touch-Based Devices

Today, touch-based devices offer computing capabilities that are competitive with traditional desktop PCs. A large body of previous work designed visualizations for touch-based interfaces (e.g., [2, 11, 20, 34, 40]). Through directly touching and manipulating a visual representation, touch-based interfaces decrease the separation between the user from the data and visual representation [29]. However, touch-based interfaces introduce a new set of design challenges such as occlusion and selection precision [45]. For example, selecting a visual glyph using a finger on touch-based interfaces might not be as precise as using a mouse on desktop PCs. Going forward, we are interested in understanding what are principles for expanding the EMS technique for touch-based interfaces. We posit that our suggested set of design guidelines for supporting EMS are extendable to touch-base interfaces, but the implementation might differ based on specific affordances of the interface. However, this remains to be formally studied.

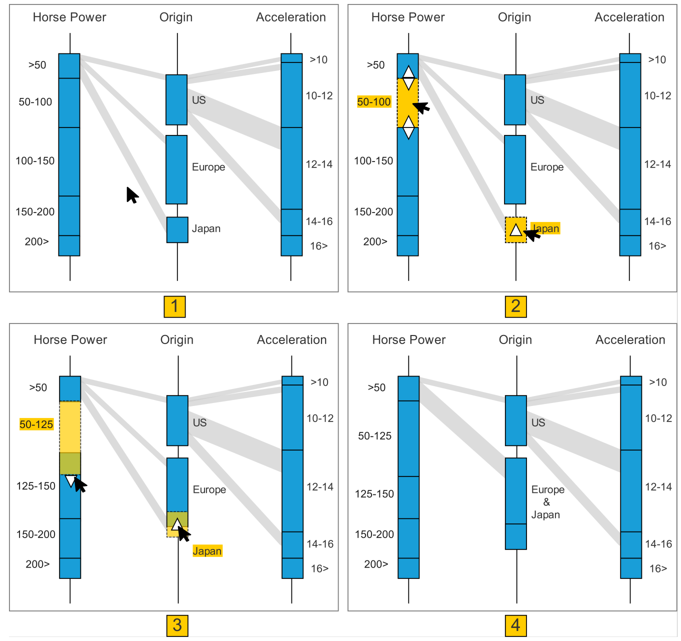

Fig. 10. A parallel set visualization (relations are only partially shown). User can drag and drop to categorical groups to merge them (e.g., Origin). Bins can be combined by dragging the top or bottom border of the rectangular shape representing a bin (e.g., Horse power). In both cases, signifiers guide user where and how to interact with the visualization.

8 Conclusion

In this paper, we introduced Embedded Merge & Split (EMS), a novel embedded interaction technique that enables users to adjust data grouping criteria by directly manipulating encodings used for presenting groups. We implemented EMS for bar charts and histograms, two commonly used visualization types. Results of our quantitative user study comparing EMS with Tableau shows that the EMS technique can significantly reduce the time required to merge, split, and change the number of groups in bar charts and histograms. Results of our qualitative study with five expert HIV researchers indicate that the use of EMS results in a more fluid and natural experience.

All in all, we view this work as the first step towards exploring EMS, and plan to further investigate our technique for other visualizations and data types. We also plan to conduct case studies in collaboration with HIV researchers to investigate if and how providing EMS can effect the process and outcomes of a more holistic exploratory data analysis.

Acknowledgments

We are grateful to HIV researchers at the Center for AIDS Research (CFAR) at UC San Diego, reviewers, study participants, the other members of the Georgia Tech Visualization Lab and UCSD Ubicomp Lab for their helpful insights and comments. We also thank Narges Mahyar, Melanie Tory, and Steve Haroz for the constructive criticism and suggestions for improvements. This project is supported by NIH 1P01AI131385-01 and NSF IIS-1750474.

REFERENCES

[1] Tableau Datasets, https://public.tableau.com/s/resources, 2015.

[2] Vizable by Tableau, https://vizable.tableau.com/., 2015.

[3] Arcgis Software, https://www.arcgis.com/features/index.html, 2018.

[4] R Language, https://www.r-project.org/about.html, 2018.

[5] SAS Software, https://www.sas.com, 2018.

[6] Spotfire, https://spotfire.tibco.com/, 2018.

[7] Tableau Software, http://www.tableau.com/, 2018.

[8] N. Adrienko and G. Adrienko. Spatial generalization and aggregation of massive movement data. *IEEE Transactions on Visualization and Computer Graphics*, 17(2):205–219, Feb 2011. doi: 10.1109/TVCG.2010.44

[9] W. Amos, J. Hoffman, A. Frodsham, L. Zhang, S. Best, and A. Hill. Automated binning of microsatellite alleles: problems and solutions. *Molecular Ecology Resources*, 7(1):10–14, 2007.

[10] D. Anderson. *Essentials of statistics for business and economics*. Business & Economics, 1989.

[11] D. Baur, B. Lee, and S. Carpendale. Touchwave: Kinetic multi-touch manipulation for hierarchical stacked graphs. In *Proceedings of the 2012 ACM International Conference on Interactive Tabletops and Surfaces*, ITS '12, pp. 255–264. ACM, New York, NY, USA, 2012. doi: 10.1145/2396636.2396675

[12] M. Beaudouin-Lafon. Instrumental interaction: an interaction model for designing post-wimp user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pp. 446–453. ACM, 2000.

[13] C. Beyan, V.-M. Katsageorgiou, and R. B. Fisher. Extracting statistically significant behaviour from fish tracking data with and without large dataset cleaning. *IET Computer Vision*, 2017.

[14] E. T. Brown, J. Liu, C. E. Brodley, and R. Chang. Dis-function: Learning distance functions interactively. In *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 83–92, Oct 2012. doi: 10.1109/VAST.2012.6400486

[15] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in information visualization: using vision to think*. Morgan Kaufmann, 1999.

[16] S. K. Card, G. G. Robertson, and J. D. Mackinlay. The information visualizer, an information workspace. In *Proceedings of the SIGCHI Conference on Human factors in computing systems*, pp. 181–186. ACM, 1991.

[17] F. Chevalier, P. Dragicevic, and C. Hurter. Histomages: fully synchronized views for image editing. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*, pp. 281–286. ACM, 2012.

[18] V. Dhar. Data science and prediction. *Commun. ACM*, 56(12):64–73, Dec. 2013. doi: 10.1145/2500499

[19] S. Diehl and T. S. Statler. Adaptive binning of x-ray data with weighted voronoi tessellations. *Monthly Notices of the Royal Astronomical Society*, 368(2):497–510, 2006.

[20] S. M. Drucker, D. Fisher, R. Sadana, J. Herron, and m. schraefel. Touchviz: A case study comparing two interfaces for data analytics on tablets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pp. 2301–2310. ACM, New York, NY, USA, 2013. doi: 10.1145/2470654.2481318

[21] N. Elmqvist, A. V. Moere, H.-C. Jetter, D. Cernea, H. Reiterer, and T. Jankun-Kelly. Fluid interaction for information visualization. *Information Visualization*, 10(4):327–340, 2011.

[22] A. Endert, P. Fiaux, and C. North. Semantic interaction for visual text analytics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 473–482. ACM, 2012.

[23] S. Few. *Now you see it: simple visualization techniques for quantitative analysis*. Analytics Press, 2009.

[24] H. P. Friedman and J.Rubin. On some invariant criteria for grouping data. *Journal of the American Statistical Association*, 62(320):1159–1178, 1967. doi: 10.1080/01621459.1967.10500923

[25] E. L. Hutchins, J. D. Hollan, and D. A. Norman. Direct manipulation interfaces. *Human-computer interaction*, 1(4):311–338, 1985.

[26] W. Javed, N. Elmqvist, J. S. Yi, et al. Direct manipulation through surrogate objects. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 627–636. ACM, 2011.

[27] S. Kandel, A. Paepcke, J. Hellerstein, and J. Heer. Wrangler: Interactive visual specification of data transformation scripts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pp. 3363–3372. ACM, New York, NY, USA, 2011. doi: 10.1145/1978942.1979444

[28] H. Kim, J. Choo, H. Park, and A. Endert. Interaxis: Steering scatterplot axes via observation-level interaction. *IEEE transactions on visualization and computer graphics*, 22(1):131–140, 2016.

[29] B. Kondo and C. M. Collins. Dimpvis: Exploring time-varying information visualizations by direct manipulation. *Visualization and Computer Graphics, IEEE Transactions on*, 20(12):2003–2012, 2014.

[30] R. Kosara, F. Bendix, and H. Hauser. Parallel sets: interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):558–568, July 2006. doi: 10.1109/TVCG.2006.76

[31] B. c. Kwon, W. Javed, N. Elmqvist, and J. S. Yi. Direct manipulation through surrogate objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pp. 627–636. ACM, New York, NY, USA, 2011. doi: 10.1145/1978942.1979033

[32] H. Lam. A framework of interaction costs in information visualization. *IEEE transactions on visualization and computer graphics*, 14(6), 2008.

[33] M. Lavielle and K. Bleakley. Automatic data binning for improved visual diagnosis of pharmacometric models. *Journal of pharmacokinetics and pharmacodynamics*, 38(6):861–871, 2011.

[34] B. Lee, P. Isenberg, N. H. Riche, and S. Carpendale. Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *Visualization and Computer Graphics, IEEE Transactions on*, 18(12):2689–2698, 2012.

[35] T. Munzner. *Visualization analysis and design*. CRC press, 2014.

[36] D. A. Norman. *Things that make us smart: Defending human attributes in the age of the machine*. Basic Books, 1993.

[37] C. Perin. *Direct Manipulation for Information Visualization*. PhD thesis, Paris 11, 2014.

[38] C. Perin, R. Vuillemot, and J.-D. Fekete. A table!: Improving temporal navigation in soccer ranking tables. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems*, CHI '14, pp. 887–896. ACM, New York, NY, USA, 2014. doi: 10.1145/2556288.2557379

[39] N. H. Riche, B. Lee, and C. Plaisant. Understanding interactive legends: a comparative evaluation with standard widgets. In *Computer graphics forum*, vol. 29, pp. 1193–1202. Wiley Online Library, 2010.

[40] J. M. Rzeszotarski and A. Kittur. Kinetica: Naturalistic multi-touch data visualization. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems*, CHI '14, pp. 897–906. ACM, New York, NY, USA, 2014. doi: 10.1145/2556288.2557231

[41] B. Saket, H. Kim, E. T. Brown, and A. Endert. Visualization by demonstration: An interaction paradigm for visual data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):331–340, Jan 2017. doi: 10.1109/TVCG.2016.2598839

[42] B. Saket, A. Srinivasan, E. D. Ragan, and A. Endert. Evaluating interactive graphical encodings for data visualization. *IEEE Transactions on Visualization and Computer Graphics*, 24(3):1316–1330, March 2018. doi: 10.1109/TVCG.2017.2680452

[43] A. Satyanarayan and J. Heer. Lyra: An interactive visualization design environment. In *Computer Graphics Forum*, vol. 33, pp. 351–360. Wiley Online Library, 2014.

[44] B. Shneiderman. 1.1 direct manipulation: a step beyond programming languages. *Sparks of innovation in human-computer interaction*, 17:1993, 1993.

[45] K. A. Siek, Y. Rogers, and K. H. Connelly. Fat finger worries: How older and younger users physically interact with pdas. In M. F. Costabile and F. Paternò, eds., *Human-Computer Interaction - INTERACT 2005*, pp. 267–280. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.

[46] E. R. Tufte. *Envisioning information*. Graphics press, 1990.

[47] Y.-W. Wu, Y.-H. Tang, S. G. Tringe, B. A. Simmons, and S. W. Singer. Maxbin: an automated binning method to recover individual genomes from metagenomes using an expectation-maximization algorithm. *Microbiome*, 2(1):26, 2014.

[48] J. S. Yi, Y.-a. Kang, J. T. Stasko, and J. A. Jacko. Understanding and characterizing insights: How do people gain insights using information visualization? In *Proceedings of the 2008 Workshop on BEyond Time and Errors: Novel evaLuation Methods for Information Visualization*, BELIV '08, pp. 4:1–4:6. ACM, New York, NY, USA, 2008. doi: 10.1145/1377966.1377971