Evaluating the Utility of Multiple Workspaces and Easy Chart Creation for Visual Analytics

Swaroop Panda[1][0000−0003−4786−0713] and Wolfgang Stuerzlinger[2][0000−0002−7110−5024]

[1] Department of Computer and Information Sciences, Northumbria University, Newcastle Upon Tyne, United Kingdom
[2] School of Interactive Arts and Technology, Simon Fraser University, Burnaby, BC, Canada

Abstract. Visual Analytics involves interactive visual exploration of data sets. Interaction methods allow data analysts to explore different aspects and investigate various hypotheses. This work investigates if chart creation and multiple tabs enable analysts to pursue multiple hypotheses or analysis paths better than other approaches. To evaluate this, new features were added to an existing Visual Analytics prototype with a multi-tab design. Two user studies with professional data analysts confirmed the utility of our approach within the Visual Analytics process. Analysts can easily work within multiple workspaces, enabling them to frame and track multiple hypotheses about the data and pursue different exploratory analysis paths, ultimately yielding more comprehensive analysis outcomes.

Keywords: Visual Analytics · Dashboards · User Studies

1 Introduction

Existing visual analytics (VA) systems allow users to interact with datasets through interaction methods like selecting, filtering, or drilling down. These methods enable iterative exploration and analysis of datasets. Many VA tools support a common interaction pattern where analysts focus on a single visualization and build an overview (dashboard) from multiple visualizations [27]. However, current systems lack support for other aspects of the knowledge generation process [31,18].

Current VA tools lack support for concurrent states during analysis, as modifications to visual encoding or filtering globally affect the application. Systems, like Voyager 2 [37], provide an initial overview and a single-state interface for exploring data. Still, they don't allow non-destructive evaluation of how different options affect previous analysis paths and results. Analysts often have to redo part of their analysis or backtrack on previous interactions to pursue multiple alternative analysis tasks, which destroys the history corresponding to other branches of exploration. Support for concurrent workspaces would enable

users to create multiple analysis paths that can be explored in parallel without affecting previous work.

Supporting users in their search for answers, locating findings, and organizing them efficiently is a valuable feature in visual data exploration [34]. DynSpace [10] supports multiple workspaces for visual analysis. The multi-tab interface simplifies different analysis processes and facilitates data and chart copying. The authors [10] investigated the benefits of multiple concurrent workspaces in a visual analysis system with professional data analysts. However, they lacked facilities for rapid chart creation and evaluated how the functionality of pursuing multiple parallel hypotheses affects the analysis process.

In this work, we added new features to DynSpace to facilitate rapid chart creation. We evaluated these features in a first user study and found that they substantially facilitate the pursuit of hypotheses about the data. In a second user study, we compared two versions of DynSpace (*DynSpace SINGLE* and *DynSpace MULTIPLE*) to see if multiple tabs help users pursue multiple alternate hypotheses. Our results confirmed that users find multiple tabs useful for organizing their analysis process when pursuing multiple hypotheses.

The work's main contributions are: 1) evaluating easy chart creation features' utility for the VA process through a user study, and 2) exploring how multiple tabs help users efficiently pursue multiple VA hypotheses and analysis paths.

2 Background

Our work is informed by prior work in VA on dashboards, support for the VA process, and multiple tab interfaces.

2.1 Dashboards

Dashboards, popular visual data displays on a single screen, enable quick data and property viewing and monitoring [12,39,11,3]. In VA, dashboards are often defined as faceted analytical displays with multiple charts and interaction techniques for sufficient analysis [12]. Current VA software like Tableau [27] and SAP Lumira support such dashboards.

Tableau [27] and SAP Lumira support creating dashboards from individual charts through interactions like filtering, drilling, sorting, and dragging. These interactions provide detailed data insights. While these systems allow multiple dashboards, they maintain a consistent global filter state, limiting analysis scope. Our work enables analysts to manage independent workspaces for a dataset, potentially with different filters for each workspace.

VizDeck presented drag-and-drop dashboard construction from recommendations [22]. Data Voyager [36] supported breadth-first exploration with multiple recommended charts in the initial view. Related work focused on chart recommendations, e.g., [5,35,15,33]. Data Voyager 2 [37] added drill-down, filtering, and related actions but lacked multiple dashboards.

2.2 Supporting Nonlinear Visual Analytics Processes

Many existing VA solutions offer users only a single workspace, which biases analysis towards a linear workflow. Despite this, analysts often backtrack, reuse parts of previous workflows, and investigate alternatives and concurrent scenarios, incurring additional interactions like long undo sequences and manual result storage. While commercial VA systems like Tableau facilitate iterative data exploration and analysis, they lack support for the entire knowledge generation process [31], including activities like backtracking and investigating multiple parallel data analyses or hypotheses. After all, analysts risk losing track of their work and forgetting insights in a purely iterative process, making backtracking beneficial.

Several VA systems have been developed to aid users in comprehensively exploring, analyzing, and understanding data [31,7,26,27,16,8,30,19,29]. Some work discussed design trade-offs for developing a VA history tracking mechanism [19]. A few systems presented mechanisms to track users' interactions and store them in a history for review, recall, and retrieval [19,9,38,40,13]. Others focused on supporting and capturing insight provenance [28]. HARVEST tried to automatically capture user intents and provide insight provenance based on actions [17,16]. Yet, little work has explored the idea of multiple concurrent analyses or alternative hypotheses [24], where users can explore different scenarios, backtrack, copy parts of one analysis to another, or go through different analysis phases while still revisiting each phase.

2.3 Multi-Tab Interfaces

While research on multi-tab web browsing has focused on traditional multi-tab browsing, our work focuses on web applications with multiple tabs [20]. A tabbed interface with explicit support for user-declared goals reduces disorientation and aids information gathering, as shown in a study comparing it to traditional multi-tab browsing [23]. This motivates our analysis of the benefits of multiple workspaces for VA processes. A single workspace does not encourage side-by-side comparisons [14], which is crucial for investigating multiple competing hypotheses. Javed and Elmqvist [21] define a composite visualization space as compositions of multiple visualizations in the same space. Their design patterns provide a background for our work.

3 DynSpace MULTIPLE: An Improved Version of DynSpace

To aid readers, we present a summary of DynSpace [10], as we build on it. DynSpace lets users create multiple workspaces to express and reflect on their analytical workflows, like testing hypotheses. However, the authors of that work never evaluated their system with such activities.

3.1 Design Considerations & Interface Design

Separate Analysis Paths Multiple workspaces allow data analysts to explore different analysis paths easily. Each workspace should maintain a separate dataset state (for filters, views, encoding, and analysis). Parallel workspaces can support different analysis phases, like data wrangling (e.g., for outlier detection), data overview, or clustering.

Shared Context Since each workspace maintains a different state, supporting common information shared between workspaces, beyond raw data, is still needed. For instance, shared bookmarks with associated state information can help a data analyst track concurrent analyses.

Easy Navigation between Workspaces Data analysts must quickly switch between workspaces to compare alternate analyses or hypotheses. They may investigate an interesting path, seek additional perspectives, or backtrack to previous results. When stuck, they may restart the analysis while reusing previous workflow parts. Efficient methods for navigating concurrent workspaces and transferring information are essential.

Revisiting and Retrieving Systems should enable interactions to track multiple analysis threads, to support data analysts investigating multiple hypotheses. This allows users to review, recall, reuse, and revisit their steps. Systems should also offer affordances for browsing through previous analysis states and restoring them non-destructively.

Like other VA tools, DynSpace [10] supports both open-ended exploration and focused data analysis through multiple coordinated views for each workspace, chart recommendations, and manual chart creation. The user interface contains three main panels: data dimensions, workspaces, and bookmarks Figure 1. DynSpace is a web application implemented in JavaScript with jQuery and Gridstack.js for the layout. For drawing charts, it uses Vega.js [32] on top of D3.js [4]. Like other VA tools, DynSpace supports open-ended exploration and focused data analysis through multiple coordinated views, chart recommendations, and manual chart creation. The user interface has three main panels: data dimensions, workspaces, and bookmarks. DynSpace is a web application implemented in JavaScript with jQuery and Gridstack.js for layout. It uses Vega.js on top of D3.js for chart drawing.

Data Dimension and Charts DynSpace classifies dataset dimensions into categorical, numerical, and temporal categories. Each dimension is color-coded for visual identification. Clicking a dimension highlights charts containing it. The "Explore" icon creates a new workspace with charts containing that dimension. Dimensions are draggable for creating charts or filters. DynSpace's original version supported only histograms, bar charts, and scatterplots to cater to non-expert users and support user studies. Bar charts can be sorted in ascending order using the "sorting" icon at the top-right. Hovering over a data point shows details, including statistical information like the mean and median.

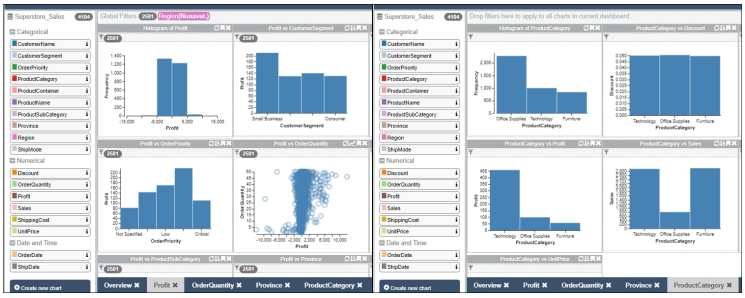

Fig. 1. Two example states of the DynSpace's tab bar in (separate) analyses where multiple workspaces are used to pursue multiple, potentially concurrent analysis paths.

Multiple Workspaces To organize their work workflow as needed, DynSpace allows users to create as many workspaces as they require, and switch between, rename, and delete them. This enables the reuse and retrieval of analysis paths and results. The main panel has multiple workspaces in a multi-tab interface for multiple analysis paths. Each workspace has charts to display specific aspects of analysis. DynSpace supports brushing and linking to explore relationships between charts. Users can create, switch between, rename, and delete workspaces to reuse and retrieve analysis paths and results.

Once loaded into DynSpace, a dataset generates an initial workspace with overview charts of randomly selected dimensions, providing a starting point for analysis.

Workspaces can contain multiple charts, which users can copy to another workspace by dragging and dropping onto the tab title. This action preserves the chart's filters. This feature allows users to explore different hypotheses or questions in different workspaces. Users define chart axes by dropping dimensions into corresponding areas. DynSpace automatically sets or resets the chart type based on the dropped dimensions. For instance, dropping a numerical dimension with a categorical one creates a bar chart, while two categorical dimensions result in a scatter plot. Adding another dimension to a scatter plot converts it to a colored one. Workspaces use a grid-like layout with a default size of 16 (4x4) charts. Users can resize charts for more detail and minimize layout changes during rearrangement or resizing. Charts snap to other charts and grid positions to reduce effort and maintain the grid.

Drag-and-drop Filters DynSpace supports two types of filters: global and local. Global filters apply to all charts in a workspace, while local filters apply only to a single chart. Both types are displayed as small, rounded, color-coded rectangles containing the dimension name and filter extent. Global filters are displayed in the workspace filter bar, while local filters are displayed in the chart filter bar. Multiple local filters can be applied to a chart simultaneously. Filter details can be accessed by hovering.

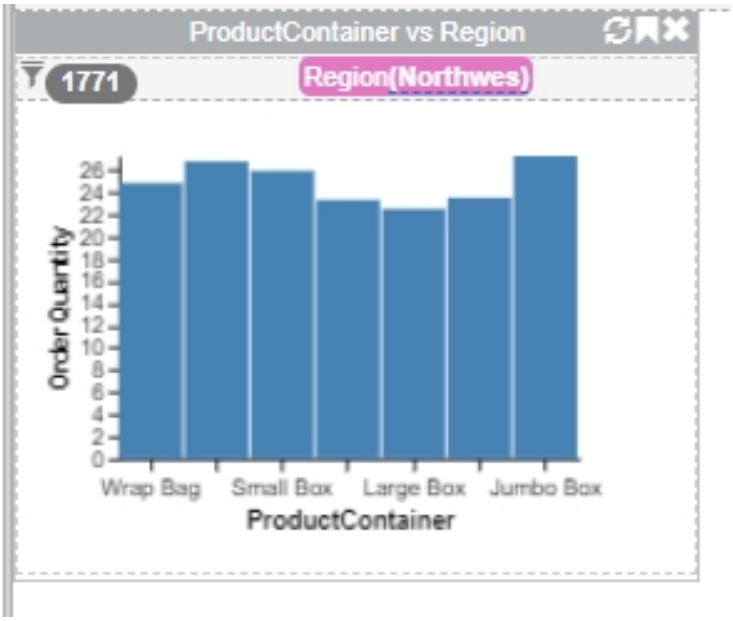

Fig. 2. A chart container showing a bar chart of the ProductContainer and the OrderQuantity Dimensions with a filter on the Regions dimension. The container includes three icons at the top-right corner: transpose axes, bookmark, and close chart (from left to right). The filter icon is at the top left, part of the local filter bar.

Context Bookmarks DynSpace lets users save or bookmark charts using the bookmark icon. This helps users track insights and maintain context. Bookmarks appear in a panel ordered by creation, showing thumbnails, dimensions, filters, annotations, and a context button. Clicking the button displays the workspace state at bookmark creation, aiding the recall of specific insights.

DynSpace MULTIPLE eases chart creation in DynSpace by removing the restriction of dropping dimensions into specific axes. It follows VizInteract [6] and automatically creates an appropriate chart based on the dimension's properties. DynSpace MULTIPLE shows a histogram for a single dropped dimension, making it faster for users to visualize the data.

We added support for transposing (swapping axes) to DynSpace MULTIPLE. Users may want the same dimension to occupy the same axis in all charts in a workspace. This reduces the need to recreate charts compared to the original version.

4 DynSpace MULTIPLE Usage Scenario

This section details how analysts can benefit from DynSpace MULTIPLE's multiple workspaces. The following use case is based on the Superstore Sales dataset [2], with 18 dimensions and 4104 data records.

Fig. 3. Dimension drop areas in the original version of DynSpace and DynSpace MULTIPLE, respectively.

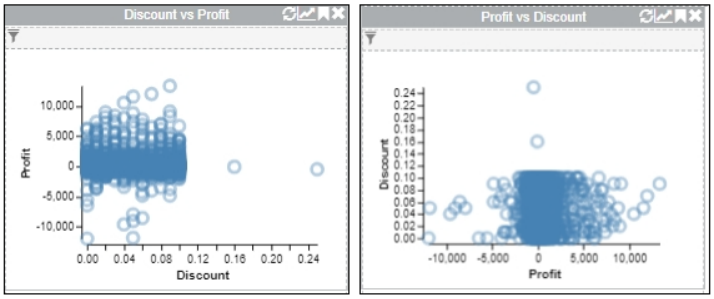

Fig. 4. The axes of a chart are swapped using the transpose option. The transpose icon is left-most in the icon bar at the top-right.

Jane, a data analyst, prepares a report on current store sales. She wants to understand how sales relate to other data factors. She uses the "explore" icon beside the sales dimension to create a new tab with sales-related charts to do this. To easily spot potential relationships, she transposes the charts where sales are vertical to match her preference.

Jane identifies potentially interesting charts in the workspace, such as Sales vs. Profit, Sales vs. Shipping Cost, Sales vs. Region, and Sales vs. Discount, and plans to analyze them. She opens a new workspace and re-creates the charts. Jane identifies the three regions with the highest sales by sorting the Sales vs. Region bar chart. She then visually compares the product category's relationship with profit and sales by dropping the product category dimension onto the Sales vs. Profit chart.

Subsequently, Jane uses the "explore" icon beside the "Shipping Cost" dimension to explore the relationship between Shipping Cost and other variables. She then becomes interested in how attributes like product containers and regions impact both Sales and Shipping Costs. To analyze this, she copies relevant

Fig. 5. Exploring the sales variable in the dataset by charting its relationships with all other dimensions.

Fig. 6. Some charts were transposed to increase visual consistency.

charts into a new workspace, transposes the axes to align the variables, and then examines the relationships.

Jane quickly identified how variables impact each other and top-performing regions using DynSpace MULTIPLE. By preserving her workflow through charts in different workspaces, she can compile the information into a report summarizing insights for her boss.

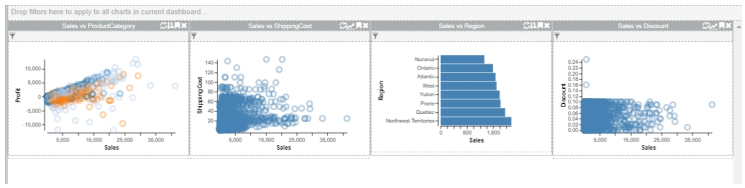

Fig. 7. Jane explores some attributes related to the Sales dimension.

5 DynSpace SINGLE: A restricted version of DynSpace MULTIPLE

5.1 Design Considerations & Interface Design

To explore how users leverage multiple tabs in a VA system, we modified DynSpace to restrict tab usage. In DynSpace SINGLE, users can only access the overview tab and cannot create or switch tabs. A new "clear" button clears all charts in the current tab. For instance, users can clear existing charts before creating new ones or adding new charts to existing ones. All other features, like quick chart creation, remain the same as in DynSpace MULTIPLE.

6 DynSpace SINGLE Usage Scenario

Here is how an analyst can use DynSpace SINGLE for the same analysis process as DynSpace MULTIPLE. We focus on highlighting the differences between the two versions.

During her analysis, Jane notices important relationships like Sales vs. Profit, Sales vs. Shipping Cost, Sales vs. Region, and Sales vs. Discount. She clears the overview tab charts and quickly re-creates them by dragging and dropping.

She wants to explore Shipping Cost vs. Product Container and Shipping Cost vs. Region in more detail. Before proceeding, she clears the Sales charts and recreates the Shipping Cost charts from memory.

Fig. 8. A snapshot of DynSpace SINGLE. There is no option for the user to create new charts and there is an option to clear all charts in the tab.

7 Study A - Easy Chart Creation

The primary objective of our initial user study was to explore how analysts benefit from easy chart creation in DynSpace MULTIPLE during data exploration and hypothesis generation. This question was previously unexplored for DynSpace [10]. Our current study addresses two key questions: *Does DynSpace MULTIPLE's improved interaction methods facilitate chart creation for data analysis?* and *What patterns emerge when analysts use DynSpace MULTIPLE for various data analysis tasks?*

We used a qualitative approach to understand chart creation and usage patterns with DynSpace MULTIPLE, and a case study approach to learn about users' analysis strategies when exploring data.

7.1 Participants

We recruited 20 participants (18 male, 2 female). Most (60%) were professional data analysts (business analysts/consultants, data scientists), while the remaining 40% were research assistants and graduate students. Participants were 23 to 29 years old.

7.2 Procedure and Apparatus

The remote user study, conducted via Zoom, involved participants using laptops and Google Chrome to interact with DynSpace MULTIPLE. Some participants allowed screen capture of their interactions. The main investigator guided participants through tasks, demonstrated features, and provided support. Observations were collected. The study averaged 45 minutes.

The study began with collecting participant background information using VA tools. Each participant then received a short training session on DynSpace MULTIPLE's capabilities. The main tasks involved creating charts, exploring data, and analyzing it. An interview was conducted at the end to clarify observations, collect feedback, and elicit ideas for future improvements.

7.3 Dataset and Tasks

The user study used the car dataset [1], containing information for 428 cars in 14 dimensions. New to all participants, the dataset allowed for detailed exploration and analysis while being small enough for meaningful conclusions within the allocated time. Participants were presented with the task scenario, dataset explanation, and dimensions, and then asked to use DynSpace MULTIPLE to create charts for specified tasks.

Table 1 details each task's purpose and design, aligned with our main research question. These tasks involve exploring the dataset by creating and manipulating various charts.

Scenario/Task	Purpose
Your friend Susan wants to buy a new car. She wants to explore options. Through appropriate charts, help her explore various aspects of buying a new car.	Introduction
Susan wants to know how the cars vary in terms of price, horsepower, or weight.	To check how users create a chart (specifically frequency-histograms).
Susan is keen to know the relationship between vehicle type and horsepower.	To check whether the user can create a bar chart.
Susan would like to know the relationship between engine size and the length of a car. Once she sees this, she also wants to know how this relationship changes across vehicle types.	To test the creation of a scatter plot and then a colored scatter plot.
Susan has an eye for detail. She wants to understand how the dealer cost relates to other properties by always showing these costs on the vertical axis.	To evaluate the transpose capability and how participants use workspaces.

Table 1. The tasks participants performed in each session, along with their purpose.

7.4 Data Collection and Analysis

Investigators observed participants, recorded some of them, and conducted open-ended interviews. An open-coding approach categorized patterns, followed by thematic analysis. The analysis focused on how participants used DynSpace MULTIPLE to complete tasks. The interview data also revealed participants' opinions.

7.5 Study Findings

The analysis focused on participants' data analysis process and workspace usage patterns.

Participants easily explored data with simple and quick chart creation. They tried unsuccessfully to drop a dimension onto a chart area to change its display, but dropping a third dimension onto a scatter plot created a colored scatter plot. However, they quickly realized that creating a new chart helped them identify answers by creating one beside an existing related one.

The study's frequently performed actions included creating new charts and workspaces, switching between workspaces, and using bookmarks. Observing interactions and clustering them, we identified four specific patterns used with multiple workspaces.

Overview Tab as a focal point Many participants used the overview tab as a focal point for completing the analytics tasks. The overview tab was regularly revisited for identifying the kind of charts available in DynSpace MULTIPLE

(P2, P3, P5, P11), visually scanning its content to get more information or guidance for future exploration (P1, P8, P9), or even testing chart container features, like transposing axes or bookmarks (P12, P14). A single participant, P10, started working on the analysis in the overview tab but later switched to a new workspace to complete it. In the subsequent interview, P10 suggested that *"I thought I could do all the analysis in the [overview] itself and would not require any more tabs. But with so many charts in the tab, I remembered that I could move to a new [workspace]."* Overall, participants used the overview tab in one form or another as a focus area.

One Workspace for a Single Task This pattern, using a single workspace for each different task, was observed in the majority of participants, 11 of 20 (P1, P2, P6, P7, P8, P9, P11, P13, P16, P18, P19). Participants created a new tab whenever they encountered the instruction for a new task. In the subsequent interview, one of the participants stated that *"I found it intuitive to open a new tab for every new [task] question. That way I can recheck them or correct them if needed."* Another participant suggested that *"It is easy and clean to create new charts on a new tab."*

All-in-One Workspace This pattern, performing all the tasks in a single workspace, was observed in a subset of participants, 6 of 20 (P3, P4, P5, P12, P14, P15). Here, participants opened a new tab and then completed a whole series of tasks within that workspace (typically tasks T1-T3). Participants P12 and P15 started creating charts manually for tasks T4 and T5, but at some point during their analysis recalled the option to use the *explore* button – which opens a new tab/workspace — and then used it to complete the tasks. P12 later mentioned that *"I thought I needed to explore the relationships by drawing the charts manually. I had forgotten the explore option."* P15 stated that *"I thought the explore button [was] only to provide information about the type of car. I later recollected that it is useful in exploring relationships."*

Working Across Multiple Workspaces Three participants, P10, P17, and P20 used two or three workspaces to complete the analysis tasks. These participants did not relate each workspace directly to a task. P17 completed tasks T1 and T2 in a single new tab while they switched to a new tab for T3. P20 completed T1 in an initially empty workspace and then completed T2 and T3 in a single new tab. P20, in the interview, stated that *"I thought there was no need to waste space on a new tab when I could do the analysis in a single tab."*

7.6 Participant Feedback

Almost 90% of the participants asked for clarification on the dataset and its dimensions for better performance of the tasks. One of the participants suggested *"If I could have had a better understanding of the dataset and the variables I*

would have definitely completed the analysis in a better way.". Another suggested that *"More description of the data is required to understand and perform the tasks better."*

There was also further feedback on the design of DynSpace MULTIPLE. P12 suggested that *"There are too many [unfamiliar] features in the [DynSpace MULTIPLE system]. A little bit of annotation could help. I had forgotten about the explore button and kept on creating multiple charts."* Another participant suggested that *"Labelling will help understand [] DynSpace MULTIPLE better."*

7.7 Other Observations

Participants initially followed the experimenter's demonstration, but their analysis diverged when they investigated challenges. Most used a different workspace for each task, while the rest completed multiple tasks in the same tab.

8 Study B - Single vs. Multiple Workspaces

In this study, we compared analysts working with a single or multiple workspaces. We used the two versions of the system, DynSpace SINGLE and MULTIPLE, as introduced above. DynSpace SINGLE supports a single tab, while DynSpace MULTIPLE allows users to create multiple tabs for tasks (as in study A).

8.1 Participants

We recruited 20 participants (18 male, 2 female). 7 had participated in the earlier user study. Most (80%) were professional data analysts (business analysts and data scientists), while the remaining 20% were graduate students. Participants aged 25 to 30.

8.2 Procedure and Apparatus

The remote user study, conducted via Zoom, involved participants using laptops to interact with DynSpace SINGLE and DynSpace MULTIPLE with Google Chrome. Some participants allowed screen capture of their interactions with DynSpace versions. As in the previous study, the investigator guided participants, demonstrated features, offered help, and collected observations. The study averaged 50 minutes.

8.3 Dataset and Tasks

Two datasets were used: the cars dataset from the earlier study and the world dataset created from Gapminder and Wikipedia data. The world dataset includes continents, countries, regions, the population, mortality rates, per capita GDP, carbon emissions, and employment rates for 197 countries across 12 dimensions.

Each participant performed two analysis tasks in a factorial design, where they encountered both system versions (DynSpace MULTIPLE and DynSpace SINGLE) and both datasets in a counterbalanced manner to address potential learning effects.

Although slightly different, the tasks for both datasets were chosen to require the same functionality, so that users followed the same work path irrespective of the dataset. This is reflected in the *purpose* of the tasks in the two datasets.

Scenario/Task	Purpose
Your friend Susan wants to buy a new car. She wants to explore options. Creating appropriate charts, helps her explore the different aspects of buying a new car.	Introduction
Susan wants to identify the drive types that have the most Horsepower. She then wants to see how the different cars of this type compare with each other with respect to EngineSize and DealerCost.	Single hypothesis (could be true or false)
1) Susan wants to understand how the Vehicle Types SUV, Sports Cars, and Pickups compare in terms of EngineSize and RetailPrice. 2) Among them, she also wants to know how cars with the maximum Engine Size or Retail Price compare with each other to Cylinders and Weight.	Single hypothesis (exploratory)
Susan realizes that Pickups typically have the largest Engine Size. She supposes that 1) Pickups would have the maximum cylinders, 2) maximum weight, and 3) a median retail price (lesser than SUVs but more than Wagons). Susan wants to verify if these are true	Multiple hypotheses (Alternate - any of them can be true or not)
Susan is keen to identify the (rough) average Engine Size of the top 3 cars with maximum CityMPG. She thinks that these 3 cars have either very high HwyMPG and a median Retail Price	Multiple hypotheses requiring the user to revisit an earlier result (Competing - only one hypothesis could be true).
After identifying the Vehicle Type of cars with the largest Engine Size, Susan wants to know the Weight of the 4 cars that have the maximum Retail Price. Within the same range of Retail Price and Weight, Susan looks for the number of available Sports Cars.	Multiple hypotheses (more than two hypotheses with several outcome possibilities)

Table 2. The tasks participants were required to perform during each session in study 2 with the car dataset, along with the task purpose.

8.4 Data Collection and Analysis

Data collection included investigator observations, screen recordings (for some participants), and an open-ended interview. Similar to the first user study, open coding categorized usage patterns, followed by thematic analysis. This study focused on whether tabs affected task completion and how participants used

Scenario/Task	Purpose
Your friend Susan is curious about the current state of the world. As a starting point, she needs your help to explore a dataset about the world and infer some conclusions.	Introduction
Susan wants to know if the continent that has the maximum democracy index also has the maximum life expectancy. For the continents having the highest democracy index and the maximum life expectancy, she also wants to see how countries relate to Military Expenditure and Population.	Single hypothesis (could be true or false)
Susan wants to know how the Northern Africa, Middle Africa, and Southern Africa regions 1) compare in terms of the Child Mortality rate and 2) in terms of Population and per capita GDP.	Single hypothesis (exploratory)
Susan remembers from her education the Asian continent has the largest population. She supposes that Asia would thus have the highest CO2 emissions and potentially also median per capita GDP.	Multiple hypotheses (Alternate - any of them can be true or not)
Within all countries that have CO2 emissions between 15 and 25, Susan wants to identify their approximate range of military expenditure, population, and employment rate	Multiple hypotheses requiring the user to revisit an earlier result (Competing - only one hypothesis could be true).
For the region that has the second largest population, Susan is keen to know the approximate CO2 emissions of the top 3 countries that receive the largest aid. She also wants to know the employment rate and democracy index of these 3 countries	Multiple hypotheses (more than two hypotheses with several outcome possibilities)

Table 3. The tasks the participants were required to perform during each session in study 2 with the world dataset, along with the task purpose.

them. The open-ended interview data also analyzed participants' opinions on tasks and DynSpace's versions.

8.5 Study Findings

The most frequent actions during the study included re-reading, clarifying the problem statement, creating new charts, clearing tabs, and creating new tabs. Through thematic analysis, we identified three specific patterns used by participants in multiple workspaces.

Single tabs vs. Multiple Tabs All users successfully used both system versions to analyze data. However, DynSpace SINGLE participants were frustrated by the implicit need to clear charts after each task, which P1, P6, and P17 explicitly mentioned. This issue worsened with task complexity. DynSpace MULTIPLE

didn't have this issue as participants naturally used multiple tabs to organize their analysis, similar to our first study.

Offering the "clear charts" button was intended to see if clearing the workspace would compensate for the need for multiple tabs. However, several users (P1, P6, P17) preferred multiple tabs for analysis over clearing existing charts and continuing on the same tab.

Participants quickly understood how to generate new tabs, unlike their reluctance to clear charts. None of the multiple-tab users mentioned or displayed any reluctance to use new tabs for analysis. The slight hesitation to clear charts may stem from the concern of losing important information or previous analyses, which is relevant for complex tasks where more information is needed for analysis.

An analysis of the effort involved in revisiting charts supports this. There are three options: recreating a single two-dimensional chart with two drag-and-drop actions, copying an existing chart across workspaces with two workspace switches and a copy action, or switching to a saved workspace with a single action, which also makes all other charts accessible. This makes a multi-tab interface more efficient, especially for analyses involving multiple charts.

Overall, participants found it thus easier to perform their analysis through multiple tabs.

Creating charts Multiple users (P1, P10, P17) expressed a desire for easier chart creation, especially with a single tab for tasks (DynSpace SINGLE). We believe this is due to the complexity of the tasks. As tasks become more complex, users want higher ease of operations to focus on the task, not the interface.

Curiosity about the Data Several users (P3, P5, P7, P11) explored the world dataset in more depth, suggesting that they were curious about the data, even beyond the task questions. Here, it seemed that users intuitively considered several parallel hypotheses and – once they saw the data – used the dashboard to explore and verify which hypothesis held. One of the users commented that *"I thought that Western Europe was the second highest emitter of CO2."* Another user wanted to see how different regions compared in terms of population with each other, e.g., Western Europe and Southeast Asia. The user was familiar with continents and their population, but upon seeing the more specific region variable was encouraged to explore deeper.

Users were more focused on the data than the interface, indicating that both interfaces facilitated exploration. Participants working on DynSpace MULTIPLE (P3, P7, P11) used new tabs to create charts, showcasing their effectiveness in exploring complex data.

Overall Findings Taking all the outcomes of the multiple facets of the study outcomes into account, it is safe to conclude that DynSpace MULTIPLE helped users explore different aspects of the dataset better than the version that offered only a single workspace.

Participant Feedback One specific feedback was to remove the create chart button and show a chart container when a dimension is dropped onto the workspace. Multiple users mentioned the innovative uses of the scatterplot, such as selecting regions to zoom and filter. They appreciated the utility of local and global filters and recommended sorting them alphabetically. One user commented that the scatterplot has too many functionalities compared to the bar chart.

9 Overall Discussion

Participants used the overview tab/workspace as an anchor point for exploration and analytic tasks in the first study. In this user study, while many used multiple workspaces, some still relied on a single one. All found it easy to quickly build charts in DynSpace MULTIPLE, validating our design intent.

In the second study, we increased the complexity of analysis tasks. Users preferred using multiple tabs over a single tab with "clear charts" functionality. Especially with a single tab, users wanted an even easier way to create charts. This suggests that chart creation is a crucial task for VA systems.

Overall, the user studies support our hypothesis that analysts can benefit from easy chart creation in a multi-workspace system. Multiple workspaces enable analysts to pursue multiple, independent analytical tasks, i.e., multiverse analysis [25].

10 Conclusion & Future Work

We investigated the benefits of supporting multiple concurrent workspaces in Visual Analytics tools. We presented DynSpace MULTIPLE, an enhanced VA system with a multi-tab interface for easy chart creation with multiple workspaces. User studies showed that multiple workspaces benefit complex analysis tasks. Users preferred opening new tabs (workspaces) to clearing existing tabs (single workspaces) for further analysis. We also identified usage patterns and potential avenues for further investigation, such as refining the interface design, increasing display size, and revisiting implementation choices to better support multiple workspaces in VA tools. While our work focused on data analysts interacting with multiple workspaces for analysis within an hour, the same mechanisms likely support knowledge generation in even more complex VA scenarios, including collaborative analysis.

We plan to add advanced chart types like parallel coordinate plots to visualize multidimensional information and facilitate more complex tasks. Another approach is using a large physical display where data analysts can organize their processes. We will enhance DynSpace MULTIPLE accordingly and conduct further user studies while recording these sessions, to later analyze execution times and completion rates. Organizing charts on a large canvas, whether zoomable or not, can be challenging, especially when rearranging or moving charts to make

space for new analysis paths or enlarged charts. These activities require additional user work unrelated to the main data analysis.

DynSpace MULTIPLE's non-overlapping grid layout offers a compromise between the restrictive single-workspace model and the freedom of a large canvas. Users can enlarge individual charts while avoiding overlap. Contrasting multi-workspace analysis with multiple analyses in slightly different ways and comparing results simplifies tasks, especially when revisiting earlier steps due to new outliers. DynSpace MULTIPLE's ease of tracking and copying threads of analysis and filters ensures overall efficiency for visual analytics tasks.

References

1. Car Dataset (2004), https://perso.telecom-paristech.fr/eagan/class/igr204/datasets
2. Superstore Sales Dataset (2017), https://community.tableau.com/docs/DOC-1236
3. Bach, B., Freeman, E., Abdul-Rahman, A., Turkay, C., Khan, S., Fan, Y., Chen, M.: Dashboard design patterns. IEEE Transactions on Visualization and Computer Graphics **29**(1), 342–352 (2022)
4. Bostock, M., Ogievetsky, V., Heer, J.: D3: Data-Driven Documents. IEEE Transactions on Visualization and Computer Graphics **17**(12), 2301–2309 (2011)
5. Bouali, F., Guettala, A., Venturini, G.: Vizassist: an interactive user assistant for visual data mining. The Visual Computer **32**(11), 1447–1463 (2016)
6. Chakraborty, S., Stuerzlinger, W.: Vizinteract: Rapid data exploration through multi-touch interaction with multi-dimensional visualizations. In: Human-Computer Interaction. INTERACT '21, vol. LNCS 12934, p. 610–632 (Aug 2021). https://doi.org/https://doi.org/10.1007/978-3-030-85613-7_39
7. Cui, W.: Visual analytics: A comprehensive overview. IEEE access **7**, 81555–81573 (2019)
8. Dörk, M., Riche, N.H., Ramos, G., Dumais, S.: Pivotpaths: Strolling through faceted information spaces. IEEE transactions on visualization and computer graphics **18**(12), 2709–2718 (2012)
9. Dunne, C., Henry Riche, N., Lee, B., Metoyer, R., Robertson, G.: Graphtrail: Analyzing large multivariate, heterogeneous networks while supporting exploration history. In: Proceedings of the SIGCHI conference on human factors in computing systems. pp. 1663–1672 (2012)
10. El Meseery, M., Wu, Y., Stuerzlinger, W.: Multiple workspaces in visual analytics. In: 2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA). pp. 1–12. IEEE (2018)
11. Elias, M., Bezerianos, A.: Exploration views: understanding dashboard creation and customization for visualization novices. In: IFIP Conference on Human-Computer Interaction. pp. 274–291. Springer (2011)
12. Few, S.: Information dashboard design: The effective visual communication of data, vol. 2. O'reilly Sebastopol, CA (2006)
13. Gadhave, K., Cutler, Z., Lex, A.: Reusing interactive analysis workflows. Computer Graphics Forum **41**(3), 133–144 (2022)
14. Gleicher, M., Albers, D., Walker, R., Jusufi, I., Hansen, C.D., Roberts, J.C.: Visual comparison for information visualization. Information Visualization **10**(4), 289–309 (2011)

15. Golfarelli, M., Rizzi, S.: A model-driven approach to automate data visualization in big data analytics. Information Visualization **19**(1), 24–47 (2020)
16. Gotz, D., Wen, Z., Lu, J., Kissa, P., Zhou, M.X., Cao, N., Qian, W.H., Liu, S.X.: Harvest-visualization and analysis for the masses. poster (2010)
17. Gotz, D., Zhou, M.X.: Characterizing users' visual analytic activity for insight provenance. Information Visualization **8**(1), 42–55 (2009)
18. Heer, J., Agrawala, M.: Design considerations for collaborative visual analytics. Information visualization **7**(1), 49–62 (2008)
19. Heer, J., Mackinlay, J., Stolte, C., Agrawala, M.: Graphical histories for visualization: Supporting analysis, communication, and evaluation. IEEE Transactions on Visualization and Computer Graphics **14**(6), 1189–1196 (2008)
20. Huang, J., White, R.W.: Parallel browsing behavior on the web. In: Proceedings of the 21st ACM conference on Hypertext and Hypermedia. pp. 13–18 (2010)
21. Javed, W., Elmqvist, N.: Exploring the design space of composite visualization. In: 2012 ieee pacific visualization symposium. pp. 1–8. IEEE (2012)
22. Key, A., Howe, B., Perry, D., Aragon, C.: Vizdeck: self-organizing dashboards for visual analytics. In: Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data. pp. 681–684 (2012)
23. Li, L.Y., Chen, G.D.: A web browser interface to manage the searching and organizing of information on the web by learners. Journal of Educational Technology & Society **13**(4), 86–97 (2010)
24. Liu, J., Boukhelifa, N., Eagan, J.R.: Understanding the role of alternatives in data analysis practices. IEEE transactions on visualization and computer graphics **26**(1), 66–76 (2019)
25. Liu, Y., Kale, A., Althoff, T., Heer, J.: Boba: Authoring and visualizing multiverse analyses. IEEE Transactions on Visualization and Computer Graphics **27**(2), 1753–1763 (2020)
26. Ltifi, H., Benmohamed, E., Kolski, C., Ben Ayed, M.: Adapted visual analytics process for intelligent decision-making: application in a medical context. International journal of information technology & decision making **19**(01), 241–282 (2020)
27. Mackinlay, J., Hanrahan, P., Stolte, C.: Show me: Automatic presentation for visual analysis. IEEE transactions on visualization and computer graphics **13**(6), 1137–1144 (2007)
28. North, C., Chang, R., Endert, A., Dou, W., May, R., Pike, B., Fink, G.: Analytic provenance: process+ interaction+ insight. In: Extended Abstracts on Human Factors in Computing Systems. pp. 33–36 (2011)
29. North, C., Shneiderman, B.: Snap-together visualization: can users construct and operate coordinated visualizations? International Journal of Human-Computer Studies **53**(5), 715–739 (2000)
30. Roberts, J.C.: State of the art: Coordinated & multiple views in exploratory visualization. In: Fifth international conference on coordinated and multiple views in exploratory visualization (CMV 2007). pp. 61–71. IEEE (2007)
31. Sacha, D., Stoffel, A., Stoffel, F., Kwon, B.C., Ellis, G., Keim, D.A.: Knowledge generation model for visual analytics. IEEE Transactions on Visualization and Computer Graphics **20**(12), 1604–1613 (2014)
32. Satyanarayan, A., Wongsuphasawat, K., Heer, J.: Declarative interaction design for data visualization. Proceedings of the 27th annual ACM symposium on User interface software and technology pp. 669–678 (2014)
33. Shen, L., Shen, E., Tai, Z., Xu, Y., Dong, J., Wang, J.: Visual data analysis with task-based recommendations. Data Science and Engineering **7**(4), 354–369 (2022)

34. Shneiderman, B.: The eyes have it: a task by data type taxonomy for information visualizations. In: Proceedings of IEEE Symposium on Visual Languages. pp. 336–343. No. UMCP-CSD CS-TR-3665 (1996)
35. Vartak, M., Rahman, S., Madden, S., Parameswaran, A., Polyzotis, N.: Seedb: Efficient data-driven visualization recommendations to support visual analytics. Proceedings of the VLDB Endowment International Conference on Very Large Data Bases **8**(13), 2182 (2015)
36. Wongsuphasawat, K., Moritz, D., Anand, A., Mackinlay, J., Howe, B., Heer, J.: Voyager: Exploratory analysis via faceted browsing of visualization recommendations. IEEE transactions on visualization and computer graphics **22**(1), 649–658 (2016)
37. Wongsuphasawat, K., Qu, Z., Moritz, D., Chang, R., Ouk, F., Anand, A., Mackinlay, J., Howe, B., Heer, J.: Voyager 2: Augmenting visual analysis with partial view specifications. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems. pp. 2648–2659 (2017)
38. Xu, K., Ottley, A., Walchshofer, C., Streit, M., Chang, R., Wenskovitch, J.: Survey on the analysis of user interactions and visualization provenance. Computer Graphics Forum **39**(3), 757–783 (2020)
39. Zhang, L., Stoffel, A., Behrisch, M., Mittelstadt, S., Schreck, T., Pompl, R., Weber, S., Last, H., Keim, D.: Visual analytics for the big data era—a comparative review of state-of-the-art commercial systems. In: 2012 IEEE Conference on Visual Analytics Science and Technology (VAST). pp. 173–182. IEEE (2012)
40. Zhou, Z., Wen, X., Wang, Y., Gotz, D.: Modeling and leveraging analytic focus during exploratory visual analysis. In: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. pp. 1–15 (2021)