

Dupo: A Mixed-Initiative Authoring Tool for Responsive Visualization

Hyeok Kim , Ryan Rossi , Jessica Hullman , and Jane Hoffswell 

Abstract—Designing responsive visualizations for various screen types can be tedious as authors must manage multiple chart versions across design iterations. Automated approaches for responsive visualization must take into account the user’s need for agency in exploring possible design ideas and applying customizations based on their own goals. We design and implement Dupo, a mixed-initiative approach to creating responsive visualizations that combines the agency afforded by a manual interface with automation provided by a recommender system. Given an initial design, users can browse automated design suggestions for a different screen type and make edits to a chosen design, thereby supporting quick prototyping and customizability. Dupo employs a two-step recommender pipeline that first suggests significant design changes (*Exploration*) followed by more subtle changes (*Alteration*). We evaluated Dupo with six expert responsive visualization authors. While creating responsive versions of a source design in Dupo, participants could reason about different design suggestions without having to manually prototype them, and thus avoid prematurely fixating on a particular design. This process led participants to create designs that they were satisfied with but which they had previously overlooked.

Index Terms—Visualization, responsive visualization, mixed-initiative authoring

1 INTRODUCTION

Creating responsive visualizations—versions of the same visualization adapted for different screen types—is essential to communicate with a broader audience. However, authoring responsive versions can be tedious, often requiring additional design iterations. A common approach is to (semi-) finalize a larger design (e.g., for desktop) and then transform it for smaller screens (e.g., tablets, smartphones), or vice versa. This process often involves making substantial design transformations beyond simple resizing [19, 20]. Authors may try to simplify this process by falling back on design strategies they have used in the past, even if they are not optimal [28]. Even authors who are well versed in a variety of responsive techniques typically face a time-consuming process of manually testing and evaluating designs one by one.

Consequently, some recent work proposes automated approaches to responsive design [22, 42] or visualization retargeting [10, 44]. Nevertheless, in many contexts authors may want to more flexibly switch between agency and automation, such as newsrooms and public data reporting. While the ability to automate certain edits, like repositioning legends and rescaling size encodings, may be widely useful, authors may still need to manually edit the visualization in certain ways that are hard to effectively automate, such as rewriting text annotations to be more concise. A mixed-initiative approach provides flexibility by allowing users to accept automated recommendations from a system and/or make their own manual edits, and has been applied in a number of visualization systems [17, 27, 29, 41] and design tools [31, 36, 39].

We designed and implemented Dupo, a mixed-initiative authoring tool where users can create responsive communicative visualizations both through custom edits and automated assistance with flexible artboard management. The goal of Dupo is to support authors in exploring and reasoning about different responsive versions by making it quicker and easier to prototype design alternatives. **Dupo implements a two-step design recommendation pipeline (i.e., *Exploration* and *Alteration*)** to support design exploration with reduced redundancy. Users can first retrieve and explore more significant responsive design suggestions after or while creating a version of a visualization, such as changes to the layout, encoding choices, and data transformations like (dis)aggregation and filtering (*Exploration*). For each design rec-

ommendation, users can request further suggestions with more subtle changes to axis labels, legend position, etc. (*Alteration*). Users can also manually change the responsive versions at any time while using Dupo. The graphical user interface of Dupo supports flexible edit propagation not only across multiple artboards for manual edits [19] but also between users’ custom edits and the system’s automated suggestions.

We evaluated Dupo with six experienced responsive visualization authors who brought their own prior visualization designs to our study in order to explore new responsive versions using Dupo. Our participants overall believed Dupo could benefit their day-to-day responsive visualization tasks by enabling rapid and high-fidelity prototyping of possible responsive design alternatives. Participants were able to seamlessly make custom edits to the automated designs, demonstrating that Dupo can support progressive mixed-initiative authoring. We observed that Dupo supported participants in exploring different suggestions and reasoning about them as inspiration for future designs. We finally discuss research opportunities to apply our mixed-initiative approach to other responsive visualization methods, such as template-, programming-, and graphic-based creation, and outline next steps for future software for responsive visualization within adaptive systems.

2 RELATED WORK

Our work is motivated by prior work on designing responsive visualizations and mixed-initiative tools for visualization authoring.

2.1 Responsive Visualization Authoring

Responsive visualization design refers to creating multiple versions of a visualization to fit to different screen sizes and device types [4, 19, 20] (e.g., a laptop with a keyboard and track pad; a smartphone with a touch interface). Prior work identifies *programmatically difficulties*, *artboard management*, *design exploration*, and *changes to takeaways* as some of the major challenges in adapting visualizations for different screens.

While creating responsive visualizations demands both cross-device design and development expertise, earlier work mainly focused on programmatic techniques using D3.js [5]. This programmatic support was important particularly when there was limited tooling for responsive visualization. For example, Hinderman [18], Körner [25], and Andrews [4] demonstrate the technical feasibility of creating responsive visualizations using D3.js. The JavaScript library R3S.js [26] provides direct programming interfaces for responsive visualization using D3.js.

Authors who use GUI-based authoring tools (e.g., Adobe Illustrator with ai2html [38], Datawrapper [9], or Flourish [1]) need to manage multiple *artboards* (drawing or design areas) for different device types. Authors often have a set of pre-defined screen breakpoints for responsive versions (commonly including desktops, tablets, and smartphones for the Web environment) to avoid unpredictable errors arising due to dynamic resizing on different devices [4, 19]. An easy approach to

• Hyeok Kim and Jessica Hullman are with Northwestern University. E-mail: {hyeok, jhullman}@northwestern.edu.

• Jane Hoffswell and Ryan Rossi are with Adobe Research. E-mail: {jhoffs, rrossi}@adobe.com.

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.xxx/TVCG.201x.xxxxxx

handling these multiple artboards is to simply finalize an initial design for a certain screen type, and only then create other responsive versions that exhibit a small set of transformations, such as resizing [19, 20]. However, this sequential approach makes it burdensome to run design iterations on non-initial versions of the visualization, even though they are not necessarily less important to the overall success of the design. Instead, Hoffswell et al. [19] suggest transferring edits from one artboard to the other artboards and flexibly toggling such edit transfers in order to support simultaneous design iterations on multiple artboards. Commercial tools like ZingCharts [45] and Datawrapper [9] allow authors to specify conditional settings (e.g., altering axis positions for smartphone screens) for a bundle of templates for different screen types.

Ideally, exploring design alternatives can help an author identify better ideas [23, 24, 28], yet manually drafting and managing more than a few design alternatives is often time-consuming. Thus, computational (semi-) automation is often recommended [23], and prior work has proposed automation approaches to responsive visualization [22, 42, 43]. Business intelligence tools like Power BI [30] and Tableau [37] offer presets to make it easier to create day-to-day responsive visualizations for data analytics. MobileVisFixer [42] uses machine learning (ML) techniques to retarget non-responsively created desktop visualizations for mobile screens. LQ^2 [43] introduces an ML-based (pairwise) ranking model for visualization layout given a chart size (e.g., a larger number of bars or thicker bars for a wider chart size). Kim et al. [22] propose a recommendation pipeline for responsive visualizations using Answer Set Programming (ASP) that expresses a search problem in terms of facts, rules, and constraints [14]. As a means of representing such automated approaches, Kim et al. [21] present Cicero, a declarative grammar for specifying responsive transformations.

When exploring automated design alternatives, authors still need to reason about how effectively the responsive versions maintain high-level design goals such as intended takeaways for viewers [10, 12, 20, 34, 44]. For example, resizing a chart can change the aspect ratio, impacting the implied visual trend or effect size, and changing the visual encodings to adjust graphical density (e.g., changing a complex line chart to a heatmap) may change the relative discriminability of points. Reasoning about changes to a visualization “message” or set of key takeaways can be difficult because the same responsive design transformation strategy may result in different levels of message preservation depending on the underlying data distributions (e.g., a relatively uniform distribution is less impacted by changes to chart size). Kim et al. [22] propose a set of measures that approximate changes to visualization insights between responsive versions in terms of data identifiability, comparison discriminability, and trend recognition. ViSizer [44] and Di Giacomo et al. [10] suggest automated resizing algorithms by preserving important visual features using a significance grid.

While prior work individually addresses various challenges in authoring responsive visualizations, translating those techniques into an authoring system is an important step to actually support users. For example, prior automated approaches produce one final outcome [42] or too many outcomes with redundancy [22], limiting users’ ability to explore designs. Informed by these works, we implement three types of automated recommender pipelines (*Exploration*, *Alteration*, and *Augmentation*) and incorporate them with flexible customization within a unified end-user system for authoring responsive visualizations.

2.2 Mixed-initiative Visualization Authoring

Prior work on visualization tools (e.g., for ML-driven analytics [39], exploratory data analysis [7, 27, 40, 41], business intelligence [37], dashboard design [29, 32], chart design [17], and infographics [8]) has applied a mixed-initiative approach where users can make manual changes or apply system suggestions. One benefit of mixed-initiative approaches for authoring is that it expedites design iterations by reducing the human effort involved in prototyping. However, prior recommenders for responsive visualization design [22, 42] are intended to be one-shot recommendations, lacking considerations for progressive design processes where users make manual edits and request automation iteratively [17, 27, 29, 40]. Our work outlines a recommender pipeline suited for mixed-initiative authoring of responsive visualizations.

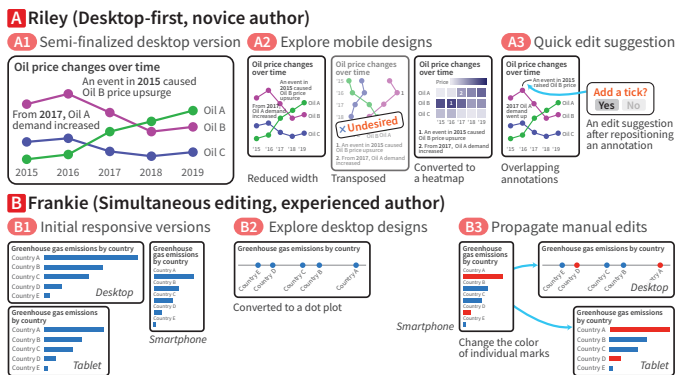


Fig. 1: Two proposed usage scenarios for a mixed-initiative responsive visualization design tool. **A** A novice author creating responsive designs in a desktop-first manner. **B** An experienced author editing responsive artboards simultaneously.

3 USAGE SCENARIOS

To motivate our design considerations for mixed-initiative responsive visualization authoring tools, we propose two usage scenarios in which a visualization author uses a mixed-initiative tool, based on prior formative findings [19, 20] and recommender approaches [22, 42].

Desktop-first, novice author: Riley (Figure 1 A) is a novice visualization author with limited experience in creating responsive designs. Riley uses a desktop device during the development process, and thus decides to start by designing a desktop visualization. After roughly finishing the desktop design of a line chart about oil price changes over time **A1**, Riley explores mobile versions suggested by the mixed-initiative tool **A2**. When examining the recommendations, Riley notices that some of the suggestions violate a design guideline in Riley’s organization that discourages putting a temporal field on a vertical axis in a line chart. To filter that case out, Riley marks it as “undesired,” so that the system does not recommend similar cases next time Riley retrieves design recommendations in this scenario. Riley thinks two versions, one with the reduced width and the other converted to a heatmap, seem useful. Riley selects both of them to discuss with their team. Riley realizes that some annotations are overlapping for these designs and decides to fix the position and width of annotations to avoid overlap and make them more readable **A3**. By doing so, one of the annotations is repositioned further away from the data point that it is indicating, so the tool suggests adding a tick between that distant annotation and the corresponding data mark. Riley accepts the quick edit suggestion.

Simultaneous editing, experienced author: Frankie (Figure 1 B) is an experienced visualization designer. Frankie prefers to create responsive designs simultaneously by managing multiple artboards together, which makes inspecting and comparing takeaways between versions easier. Frankie starts with artboards for desktop, tablet, and mobile, which are mandated screen breakpoints in Frankie’s organization. After creating a horizontal bar chart about greenhouse gas emissions by country **B1**, Frankie realizes that the desktop version looks too wide and does not use the large screen space efficiently. Thus, Frankie retrieves some design recommendations using the mixed-initiative tool for the desktop version to reduce the time it takes to prototype different ideas **B2**. Frankie finds a one-dimensional dot plot that spreads the quantities (previously encoded by the bar length) along the same horizontal axis, which better uses the landscape aspect ratio of desktop devices without distorting the original takeaway about the rank of those countries. Frankie selects the dot plot version to further evaluate alongside the original bar chart. Frankie now wants to highlight two countries to indicate some important points **B3**. Thus, Frankie changes the color of the marks for those countries (in the mobile version) to be contrasting. This edit is propagated to the desktop and tablet versions.

4 DESIGN CONSIDERATIONS

To motivate the design of our mixed-initiative responsive visualization authoring tool, Dupo, we derive the following four design considerations inspired by prior work and our usage scenarios (section 3).

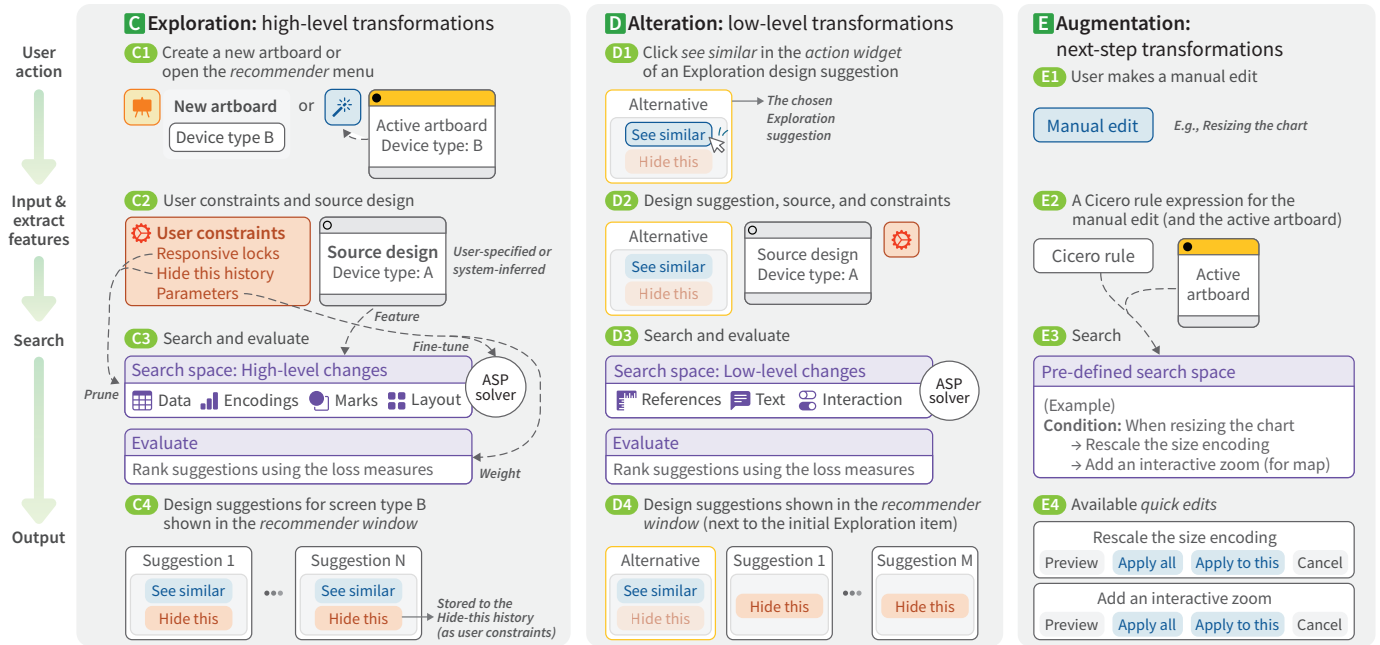


Fig. 2: Dupo’s three recommendation pipelines: *Exploration*, *Alteration*, and *Augmentation*. Initiated by a user action, each pipeline gathers the input data, runs a search, and returns the output suggestions. **C** The user can run the *Exploration* pipeline to see high-level transformations (e.g., to data, marks, encodings, and layout) by pressing the *recommender button* from the *toolbar* or creating a new artboard **C1**. Then, Dupo takes the source design and user constraints as input **C2**. Dupo generates alternatives with high-level changes and ranks them **C3**. These design alternatives are presented in the *recommender window* **C4**. If the user clicks the *hide this* button of an alternative, Dupo stores the alternative design in the *hide this history* as a user constraint for next time the user runs the *Exploration* pipeline. **D** The user can run the *Alteration* pipeline to see low-level transformations (e.g., to references, text, interactions) by selecting *see similar* for an alternative design **D1**. Dupo takes the alternative of interest, the user constraints, and the source design as input **D2**. Dupo then populates alternatives for low-level changes and evaluates them **D3**. The *Alteration* suggestions then appear next to the initial *Exploration* design **D4**. **E** The *Augmentation* pipeline is initiated after a user makes a manual edit **E1**, at which point Dupo takes its Cicero rule as input **E2** and searches a pre-defined search space **E3** to suggest *quick edits* that are commonly applied with the manual edit **E4**.

Support design exploration (DC1): A benefit of automation is quickly enumerating and ranking recommendations from a large design space, so that users can efficiently explore many different options and avoid design fixation. To not overwhelm users with subtle differences, a recommender should minimize redundancy among its suggestions [31,41]. To support reasoning about the critical trade-off between preserving visualization messages and maintaining graphical density [20, 22], a recommender and its interface should employ methods capable of evaluating how well designs preserve important patterns or insights.

Support user control (DC2): Responsive visualization authors need to retain their agency by being able to control and understand the automation process [19]. In particular, users should be able to fine-tune the systems’ recommendation methods (e.g., conveying which suggestions are less useful to guide future recommendation strategies). To support authors in making well-informed decisions about recommended alternatives, a mixed-initiative system should provide easy-to-interpret descriptions of what each suggestion is about and why it is suggested.

Support progressive authoring (DC3): Mixed-initiative systems should support an iterative design process that facilitates seamlessly switching between customization and automation [17,27, 29,40,40]. To enable progressive authoring, users should be able to easily retrieve recommendations on demand, and recommended visualizations should be customizable as needed. In addition, a progressive authoring interface should support multiple ways to revisit and compare prior designs.

Support flexible artboard management (DC4): Users may have different preferences in creating responsive visualizations (e.g., desktop-first, simultaneous editing). Hoffswell et al. [19] provide four design guidelines for flexible artboard management. Users should be able to quickly preview and edit multiple artboards simultaneously as well as customize a specific artboard and propagate edits between designs. To support progressive authoring (DC3), such flexibility should also enable propagation between custom edits and automated design suggestions.

5 RECOMMENDATION PIPELINE

To facilitate **design exploration (DC1)**, Dupo provides design suggestions using three pipelines: *Exploration*, *Alteration*, and *Augmentation*, as shown in Figure 2. For the main recommendation workflow, Dupo provides design suggestions in two steps: *Exploration* and *Alteration*. The *Exploration* pipeline (Figure 2 **C**) suggests design alternatives with high-level changes mainly to mark types, layout, data transformations, and encodings. The *Alteration* pipeline **D** suggests design alternatives with subtle, low-level changes to text elements (e.g., annotations, titles, and labels), tooltips, references, etc. The differentiation into these two steps aims to reduce redundancy during **design exploration (DC1)**, and enable fine-grained refinement of preferred recommendations (**DC2**).

Dupo also provides a simplified *Augmentation* pipeline **E** that recommends a single responsive transformation as a possible next step in response to manual user edits. These recommendations (also known as *quick edits*) encode commonly co-occurring design patterns from the prior responsive visualization literature [19,20].

5.1 Exploration: High-level Transformations

1. Input: The *Exploration* pipeline takes a source design and set of user constraints as input **C2** in order to provide **user control (DC2)** over the outcome of the recommendation pipeline.

1a. Source design. Dupo uses the source design to determine inapplicable transformations while maintaining the underlying data characteristics (e.g., data type), aesthetic choices (e.g., encoding scales), and text contents. For example, if the source design is a pie chart, then transformation strategies for a bar chart (e.g., transposing) are omitted from the search space unless there is a change to the mark type. During the evaluation stage, ranking measures compare each suggested outcome with the source design in terms of how well they preserve the design information. The source design also helps maintain consistency, which reflects common responsive visualization authoring practices where authors apply transformations to an initial version to produce other versions, as described in the usage scenario (Figure 1 **A**).

1b. User constraints. The *Exploration* pipeline leverages three types of user constraints as input: (1) *responsive locks* that provide explicit preferences about what elements cannot be updated; (2) *hidden recommendations* that provide implicit feedback about what transformations are undesirable; and (3) *recommender parameters* that provide abstract feedback on the outputs and weights applied by the recommender. First, the user can apply two types of *responsive locks*: *element-locks* and *position-locks*. An *element-lock* allows the user to specify elements they want to prevent from being removed and/or restyled. For example, if the user *element-locks* a legend, then Dupo ensures that all alternatives will include a legend, but still allows for repositioning the legend within the design. To prevent repositioning of the element, the user can specify a *position-lock*. We distinguish between an *element-lock* and a *position-lock* to enhance the degree of freedom for controlling the recommender given that authors do not always modify and reposition elements together [20]. Using these *responsive locks*, Dupo prunes suggestions with the locked elements from the search space.

Second, when viewing recommendations in Dupo, the user can choose to *hide this suggestion* (i.e., mark it as undesired) to implicitly update the behavior of the recommender. Dupo then records the unique transformations of that hidden suggestion (i.e., the transformations that do not appear in any other recommendations at the time) and removes them from the search space for subsequent rounds of recommendation.

Third, the user can provide recommender parameters including the maximum number of suggestions (at each time of request), the weights for the ranking measures, and how “drastic” the transformations are. Here, non-drastic transformations are those that only apply to resizing and transposing; drastic transformations include those that update encodings and layout beyond transposing. For example, if the maximum number of suggestions is five and the drastic parameter is 0.6, then three suggestions will include drastic changes.

2. Search space generation. When the *Exploration* pipeline is initiated, Dupo generates a search space C3 that combines predefined responsive transformation techniques with the user constraints outlined above. Technically, the search space consists of facts describing the user input, rules encoding the responsive transformations, and hard user-specified constraints, expressed using the ASP grammar [14]. Dupo uses Clingo (an ASP solver) [13] to compute the search space for a set of suggestions. These suggestions are translated into Cicero specifications [21] so that the Cicero compiler can convert the design accordingly.

These pre-defined techniques include high-level transformation rules for mark type, layout (rows and columns; small multiples), data transformations (e.g., filtering or aggregation), and encoding channels C3. When required for well-formedness, these techniques include changes to text elements (e.g., repositioning annotations when transposing the axes). Transformation strategies for the *Exploration* search space are motivated by several goals: minimize changes between responsive designs, avoid overplotting, fit to the new aspect ratio, and maintain graphical density (i.e., the white space and overlapping area). For each rationale, we encode transformation strategies observed in prior work [19–21]. For example, Dupo suggests aggregation when the outcome device size is smaller than the source to avoid overplotting. We document these rationales and strategies in the Supplementary Material.

Transformation strategies encoded in our search space are applied based on the device and responsive authoring direction (e.g., mobile-first, desktop-first). Chart resizing and transposing axes, for example, are the default strategies across those conditions. While Dupo suggests further *Exploration* suggestions, it only suggests those default transformations for tablets in a desktop-first direction unless requested by the user. Dupo suggests moving annotations out of the chart for desktop-first whereas it suggests moving titles into the chart for mobile-first.

3. Ranking measures. After generating design alternatives from the search space and before returning them to the user, Dupo evaluates them using several ranking measures in comparison with their source design in terms of “message” and density C3. The ranking measures are motivated by a trade-off from prior work between preserving takeaways and adjusting density in responsive visualizations [20]. We use the task-oriented loss measures for *identification*, *comparison*, and *trend*

proposed by Kim et al. [22] to approximate changes to a visualization’s support for those tasks as a proxy for “message.” We added a *text loss* that estimates changes to the text elements (whether removed or changed), given their importance in communicative visualizations. For density, we included an *overplotting ratio* (as a proportion of the overplotted area when elements are overlapping) and an *occupation ratio* (as a proportion of the non-white space out of the entire chart area), inspired by the clutter reduction method proposed by Ellis et al. [11].

Provided as a way to express users’ different priorities, these six measures (four insights and two density) are combined into a single scalar value as a weighted sum. By testing them on example use cases, we heuristically determined the initial weights in a way that prioritizes trend insight. Users can also **manually fine-tune the weights** in the interface (DC2). While we do not claim these measures as our contribution, we provide additional details in the Supplementary Material.

5.2 Alteration: Low-level Transformations

The user can request *Alteration* suggestions for each *Exploration* suggestion (Figure 2 D1). The *Alteration* pipeline suggests low-level changes to text elements (e.g., title, annotations), references, tooltips, etc. that make relatively subtle modifications to the design. Given that some *Exploration* suggestions may already include low-level changes in order to ensure well-formedness, the *Alteration* pipeline takes the *Exploration* suggestion as input to prune duplicate transformations from the search space, in addition to leveraging the source design and user constraints D2. The search space encodes a variety of the low-level changes documented by prior work [19, 20] C3. For example, *Alteration* from phone to desktop suggests adding labels for a quantitative axis, internalizing a short title, and adding a tooltip. Users can request *Alteration* multiple times for the same *Exploration* suggestion.

5.3 Augmentation: Next-step Transformations

Prior work indicates that some responsive strategies are commonly applied together [20], so Dupo suggests *quick edits* that recommend next-step transformations for a manual user edit. As shown in Figure 2 E, whenever a user makes a manual edit E1, Dupo passes the Cicero rule representing the edit E2 to a pre-defined search space E3, and then suggests applicable *quick edits* in the user interface E4. To identify relevant suggestions, the *Augmentation* search space considers the direction of the edit and the intended device type. For instance, Dupo suggests a *quick edit* for fixing the tooltip position at the bottom when a tooltip is *added* (direction) for a *phone* version (device type). Dupo supports a variety of *quick edits* including adding or removing tick lines after repositioning annotations and moving the title text into the chart area when it is added for desktop or tablet versions, for example. We outline the entire set of *quick edit* rules in the Supplementary Material.

6 DUPO USER INTERFACE

Based on our design considerations (section 4), we implement Dupo, a mixed-initiative authoring tool for responsive visualization. Dupo supports designing communicative visualizations with common visual encodings, annotations, and interactivity (including zoom+pan, tooltips, a brush-based context view, and an interactive legend for selecting marks on hover). The following sections describe the key features of Dupo, as illustrated by the walkthrough in Figure 3; we provide a user manual with a video demo detailing all of the functionality of Dupo and a walkthrough video in the Supplementary Material.

6.1 Overview

WALKTHROUGH: *Kris starts using Dupo by manually drafting an initial bar chart for a smartphone screen in a mobile-first manner F1.*

As shown in Figure 4, Dupo has two key components: an *editing interface* G for manually editing the visualization through pre-defined marks (with layers) and drag-and-drop interactions, and a *recommender window* K to present the responsive recommendations for the current version. The *editing interface* of Dupo consists of a *toolbar* G1, *editing panel* G2, *navigation bar* G4, and *artboard area* G5.

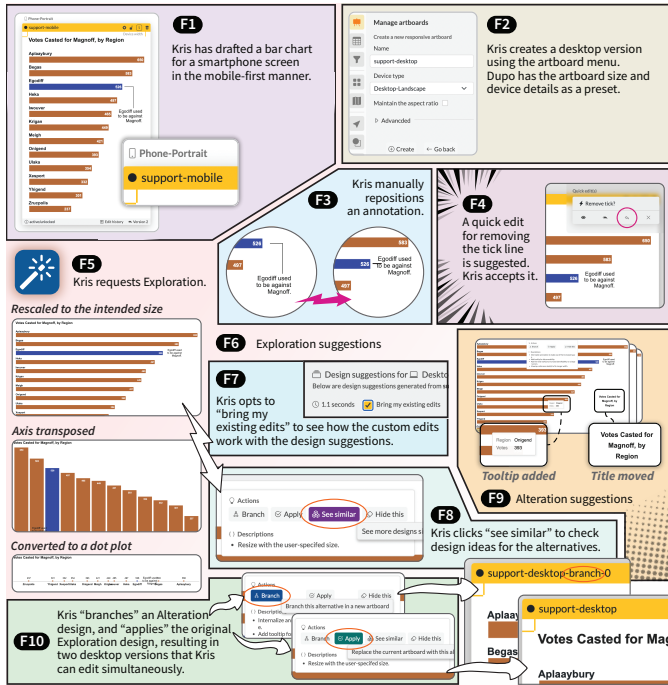


Fig. 3: A walkthrough example for creating a responsive visualization using Dupo in a mobile-first manner. Kris starts by drafting an initial bar chart for a mobile phone, and then uses Dupo's three recommendation pipelines to refine the design. A video walkthrough of this example is included in the Supplementary Material.

As the main control menu of Dupo, the *toolbar* offers functionalities for manual edits (e.g., updating properties of the *mark* or *layout*) and system features (e.g., *preview* and *export*). A user can access the menus by using the *toolbar*, pressing a keyboard shortcut, or double-clicking a chart element, which opens an *editing panel* specific to the chosen menu. As shown in Figure 4 G, for example, when the *mark* menu (🔴) is selected, the *editing panel* offers options for customizing the mark type and details (e.g., use a point on line marks) and mark property encoding channels (e.g., fill color or mark size). Following the drag-and-drop variable assignment in popular tools like Tableau Software, the user can drag a field from the floating *data widget* (G3) and drop it on the field form (or shelf). Dupo also enables limited direct manipulation for repositioning annotations directly from the *artboard area*.

The *navigation bar* offers undo and redo options and a search form for quickly accessing relevant menus. At the bottom right corner, the user can access each artboard from the artboard list (*quick artboard access*) (G10), which scrolls to the selected artboard in the *artboard area*.

6.2 Artboard Management

WALKTHROUGH: Kris creates a desktop version using the artboard menu by selecting the “Desktop-Landscape” preset, which has a default artboard size and relevant device details already defined (F2).

The *artboard area* (Figure 4 G5) displays responsive artboards, and allows a user to **see and edit multiple artboards simultaneously** (DC4), inspired by Hoffswell et al. [19]. An *artboard interface* (H) consists of a *header* (H1), *view area* (H2), and *status bar* (H4). The *header* includes the name of the artboard and indicates its status (H8), as well as providing buttons for artboard settings (e.g., artboard size and targeted device size), lock, solo-lock, and deletion (H7). The *view area* (H2) of an artboard displays the responsive design. On the top and left offsets of the *view area*, *device size indicators* (H3) (i.e., the yellow ticks), analogous to rulers in graphical software, help the user to check whether the content overflows the screen of the intended device.

The *status bar* (H4) of an artboard provides access to the *edit history* and *version list* to allow authors to revisit previous designs (DC3). If a user clicks the *edit history* (H5) button, then they can see the list of manual and automated edits. To help users distinguish the source of

edits for better overall **control** (DC2), Dupo uses different icons: a pencil icon (✎) for manual edits, a cursor icon (👉) for direct manipulation, a star icon (★) for edits from design suggestions, and a lightning icon (⚡) for augmented quick edits. A user can undo each edit by clicking the cancel button (✖). In the *version list* (H6), a triangle icon (▶) indicates the current version, and the star icon (★) denotes a version suggested by Dupo. When a user wants to compare the current version with a previous one, they can preview a version with the eye icon (👁). If the user wants to revert to a previous version, then they can “check out” a version with the check-mark icon (✓). The user can save the current artboard at any time as a new version in the *version list* to provide more **user control** over the state of the responsive versions (DC2).

Given that Dupo supports **edit propagation across multiple artboards** (DC4; c.f. [19]), it is important to enable the user to **control and see related status information** (DC2). Each artboard maintains an *activeness* status to indicate whether it is the source of changes for edit propagation (‘from’) and a *lock* status to denote whether it is receiving propagated edits (‘to’). An *active* artboard (which applies to only one artboard at a time) refers to the artboard that a user is making changes to. Each custom edit made to the active artboard is propagated only to the other *unlocked* artboards. The user can activate an unlocked artboard by clicking it, indicated by the yellow *header* color and in the *status bar*. The lock button in the artboard *header* (H7) toggles and shows the lock status (locked: 🔒, unlocked: 🔓). To edit a single artboard without edit propagation, the user can click *solo-lock* to lock all the other artboards (S). The *status bar* (H4), opacity, and header color (H8) of an artboard cues the *lock* and *activeness* status at a glance.

6.3 Edit Augmentation (Quick Edits)

WALKTHROUGH: Kris manually moves an annotation closer to the corresponding data mark for the desktop version (F3). Dupo then suggests a quick edit to remove the tick line connecting them as their relationship looks clearer (F4), which Kris accepts.

After a user makes a manual edit, Dupo suggests *quick edits* that are commonly combined with the user’s edit for other responsive visualization use cases [20] (*Augmentation*), as shown in Figure 4 J. Applicable *quick edits* appear on the top right corner of the *artboard area* (G7), with options to preview a *quick edit*, apply it to the current artboard or all the unlocked artboards, and cancel it.

6.4 Responsive Design Recommendations

As described in section 5, Dupo distinguishes the primary recommendation workflow into a two-step approach: high-level changes to mark types, layout, data transformations, and encodings (*Exploration*) and low-level transformations regarding references, annotations, text elements, and more (*Alteration*). Responsive recommendations for the current *active* visualization are shown in the *recommender window* (K).

6.4.1 Design Exploration

WALKTHROUGH: To continue refining the design, Kris clicks the recommender button to request new Exploration suggestions for the desktop version (F5–F6). Before exploring them in the recommender window, Kris first applies the previous edits (F7) to all the suggestions.

After (semi-) finalizing an artboard for a certain screen type (i.e., after indicating a dataset, a layout, and a mark layer), the user can explore design alternatives by creating another artboard for a different screen type or by clicking the *recommender button* (🔍 in Figure 4 G8) in the *toolbar*. Dupo then opens the *recommender window* (K) consisting of the *control panel* (K3) and design suggestions. The *recommender window* initially shows design suggestions for the current active artboard that consist of combinations of changes made to high-level visualization elements like visual encodings, layout, and data (i.e., *Exploration* suggestions). The order of the suggestions implicitly encodes Dupo’s estimate of the effectiveness of the designs, as described in subsection 5.1. To support **progressive authoring** (DC3), Dupo lets the user apply the edits from the current active artboard (if applicable) to the design suggestions using the *control panel*. To do so, Dupo applies the user’s manual edits (as Cicero rules) to the automated transformation rules so that those manual edits are displayed.

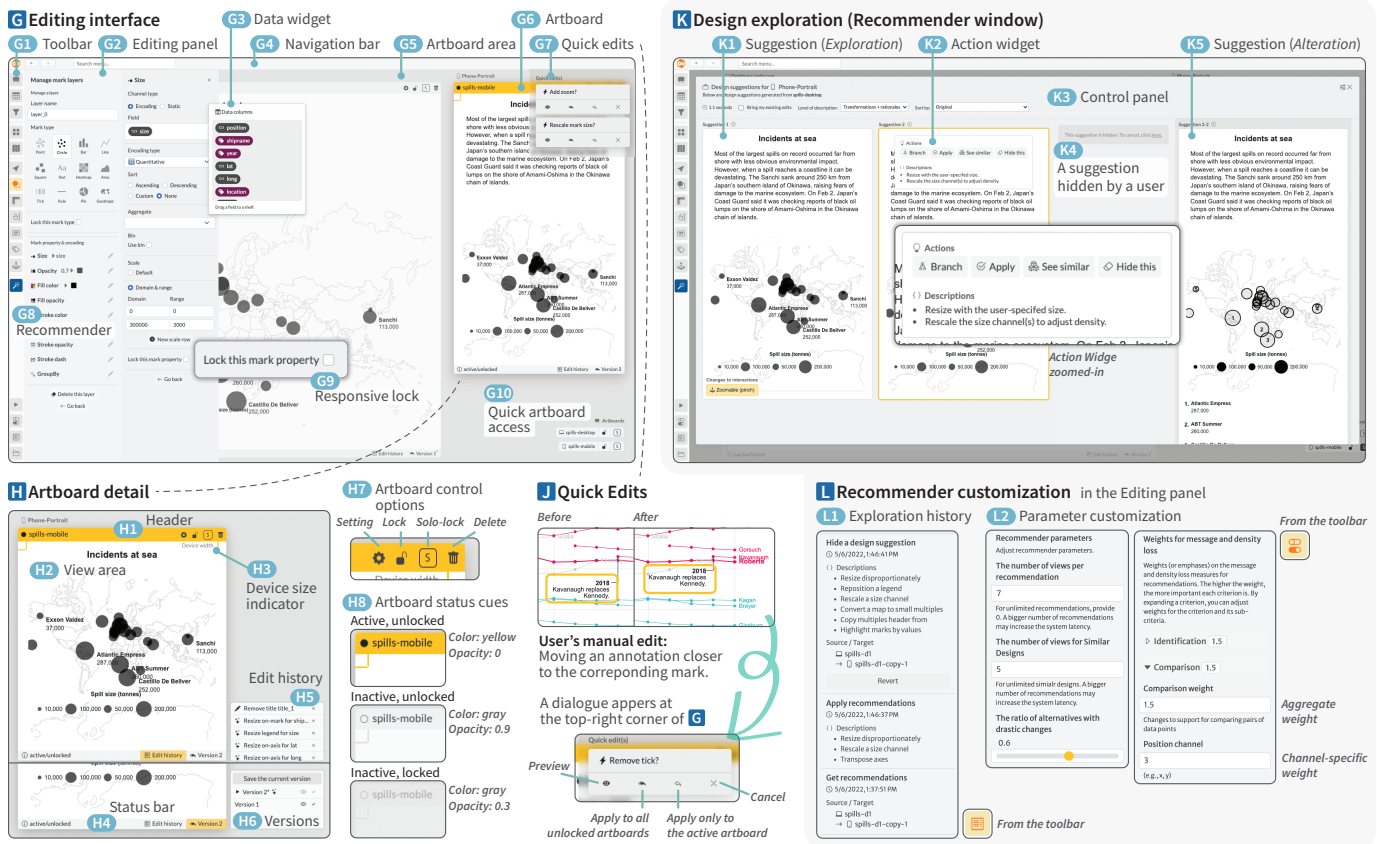


Fig. 4: An overview of Dupo's interface. **G** By selecting the *mark* menu (📌) from the *toolbar* **G1**, the user can edit marks (e.g., mark type and encodings) in the *editing panel* **G2**. To set an encoding channel, the user drags a data field from the *data widget* **G3** and drops it on the desired field. In the *navigation bar* **G4**, the user can undo or redo by clicking the arrow icons and search the system menus. The *artboard area* **G5** displays responsive *artboards* **G6**. The user can load suggestions using the *recommender button* **G8**, and set *responsive locks* on elements that the recommender must keep **G9**. At the bottom right corner, the user can *quickly access particular artboards* **G10**. **H** Each artboard consists of *header* **H1**, *view area* **H2**, and *status bar* **H4**. The *header* provides options to manage its *status* and *settings* **H7** and indicates the *activeness* and *lock status* **H8**. In the *view area*, the *device size indicator* **H3** helps the user to check if the content overflows the intended space. From the *status bar*, the user can access the *edit history* **H5** and *versions* **H6**. **J** Some manual edits by the user prompt Dupo to recommend associated *quick edits*, appearing on the top right corner of the *editing interface* **G7**. **K** The user has requested *design suggestions* **K1** generated for a smartphone version. From the *action widget* **K2**, the user can *branch* each suggestion as a new artboard, *apply* it to the current artboard, request further *Alteration* suggestions, *hide* a suggestion **K4**, and read the rationales for each suggestion. The *control panel* **K3** allows users to *apply their custom edits* to the suggestions and toggle the depth of descriptions in the *action widget*. **L** The user can review the *history* (📄) of their use of the recommender **L1** and customize the recommender parameters **L2** (⚙️) from the *editing panel* **G2**.

6.4.2 Interaction with Design Suggestions

WALKTHROUGH: *Kris finds an appealing recommendation for a resized bar chart, but is not completely satisfied with it. An action widget shows up as Kris hovers over the suggestion. Kris reads the description of the design changes, and clicks see similar to explore other Alteration examples (F8). For example, one suggestion involves moving the title into the chart and adding a tooltip (F9). After exploring design suggestions, Kris decides to branch an Alteration design as a separate version and apply the initial Exploration suggestion to the active artboard (F10).*

When the user hovers on a suggestion, an *action widget* (Figure 4 **K2**) allows them to *load similar* design suggestions, *apply* the suggestion to the current artboard, *branch* the suggestion as a new artboard, and *hide the suggestion*. To help users understand what transformations were applied to the design and why, the *action widget* shows the list of the responsive transformations (e.g., resizing, encoding changes) applied to the suggestion. If the user clicks the *see similar* button in the *action widget*, Dupo shows new alternatives **K5** that include low-level changes to the text placement or style (i.e., *Alteration* suggestions).

The *branch* and *apply* options allow the user to further refine **designs they explored (DC1)** and provide different ways to revisit previous designs to support **progressive authoring (DC3)**. If the user chooses to *branch* a suggestion from the *action widget* **K2**, a new artboard is created with the chosen design, thereby preserving the current artboard and allowing the user to compare and iterate on multiple versions of

the design. If the user chooses to *apply* a suggestion, Dupo saves the existing design in the version history **H6** and updates the artboard with the chosen design. The user can then make manual edits to **progressively** refine the design (**DC3**). When the user clicks *hide this*, Dupo immediately removes the suggestion **K4** and records the recommendation as one not to suggest again. The user can cancel the *hide this* behavior immediately or revert it from the *exploration history* **L1**, in order to better **control** the recommender behavior (**DC2**).

6.5 User Controls

Dupo enables the user to **control the recommender (DC2)** in various ways, including applying *responsive locks*, reviewing the *exploration history*, adjusting *recommender parameters*, and *quick sorting*. The user can indicate visualization elements (e.g., marks or layout) that they want to prevent the recommender from changing (Figure 4 **G9**) within the *editing interface*. The user can review their interactions with the recommender in the *exploration history* menu (📄); this menu shows the user what suggestions they have applied, branched, or hidden, and allows the user to revert *hide this* decisions **L1**. The *preference* menu (⚙️) allows the user to adjust recommender parameters **L2**. As a simplified fine-tuning method, the user can choose a sorting criterion as either identification, comparison, trend recognition, text content, or graphical density from the *control panel* **K3**. In the *control panel*, the user can also adjust the level of detail in the descriptions (i.e., transformations only or transformations with rationales).

6.6 General System Features

Dupo supports common GUI-based user interaction for graphical software, such as color pickers and sliders. To preview responsive versions together, the user can use the *device preview* menu (▶) that rescales visualizations using the pixel per inch (PPI) value of each artboard as a way to quickly verify the output of their responsive views. Dupo lets the user *export* their artboards as an HTML file with media queries (⌘) that ensure that each artboard appears only for the specified browser size and/or aspect ratio (if provided). In addition, Dupo offers several general system *preference* options, such as switching device size indicators on and off, and changing the default artboard presets.

6.7 System Design Refinement

We refined the interface based on feedback from a pilot study with four visualization designers (see the Supplementary Material for details). For example, Dupo originally had two ways to add a responsive artboard: using the *artboard* menu to create a blank artboard and duplicating an existing artboard with a new size. Users could specify the source design for the duplicated artboard as the original one, which was intended to enable users to specify source designs flexibly. However, our pilot participants were confused about having multiple options for the seemingly same task. To consolidate these options, by default we duplicate the source design (the earliest created artboard) when the user creates a new artboard in the *artboard* menu, rather than requiring a separate pipeline to produce this behavior. We also provide an option to change the source design for the recommendation to retain the same flexibility. There were also originally separate buttons for high-level and low-level design suggestions; in the revised system, we moved the button for low-level suggestions to the *action widget* (Figure 4 K2).

6.8 Implementation Details

Dupo renders visualizations using Vega-Lite [35] with our own custom extensions for common Web-based communicative visualization techniques (e.g., text wrapping and positioning annotations relative to corresponding data marks). To support **progressive authoring (DC3)**, it is important to streamline human edits and automated suggestions, so that users can revisit both manual and automated edits. To do so, Dupo uses the Cicero [21] grammar to express both human and automated responsive design transformations. Each Cicero rule is composed of a *specifier* (what to change; e.g., axis, layout), *action* (type of change; e.g., add, modify), and *option* (how to change; e.g., font size, mark type). Whenever there are new Cicero-expressed edits, Dupo compiles it with other existing rules to update the visualization in each artboard. Dupo uses Answer Set Programming (ASP) [14] with the Clingo solver [13] to search the design suggestions, and translates the resulting ASP models (set of transformations) to Cicero rules. Dupo's interface and recommender are backed by Svelte [3] and FastAPI [33].

7 USER STUDY

Motivated by our design considerations (section 4), we evaluated Dupo with six expert responsive visualization authors to answer the following research questions: *How well can authors explore alternative design ideas in Dupo (RQ1)? How do authors combine manual edits and automated design suggestions (RQ2)? How do authors use Dupo's recommender control options while interacting with the action widget (RQ3)?* We also consider general feasibility and usefulness questions: *Are Dupo's design suggestions and final outcomes deemed reasonable by authors (RQ4)? Do authors see Dupo as a feasible tool for their day-to-day responsive visualization tasks (RQ5)? What do authors see as the remaining challenges to address to improve Dupo's usefulness (RQ6)?* For greater ecological validity, our study asked participants to create responsive versions for prior visualizations they had created. Evaluating Dupo with participants' own use cases was also useful for helping us identify small improvements to Dupo to better support certain design situations while keeping the overall two-stage recommendation approach and other key features consistent across participants.

7.1 Methods

Participants and Use Cases: We recruited six graphics reporters (E1–E6) who regularly published responsive visualizations on media outlets, snowballing via social media (Twitter, Mastodon) and Slack (News Nerdy, Data Visualization Society). We asked participants to provide candidate use cases from their own work. We chose one of these examples that fit the scope of Dupo (in terms of supported visual encodings and interactivity) and exhibited variations in visual encoding combinations and layout. As shown in Figure 5, participants contributed a map (E1), small multiples of line charts (E2), a layered and annotated line chart (E3), a stacked bar chart (E4), a 2D dot plot with many annotations (E5), and a grouped bar chart (E6). For interactivity that Dupo does not support (E3: parallax; E4: view toggling), we took a single representative view (E3: conclusion chart; E4: default view).

Study Procedure. Before each session, we asked two background questions about the design process behind each participant's contributed visualization and their preferred authoring tools. During each remote-control Zoom session, the participant first went through a step-by-step guided training (creating a bar chart for desktop and exploring design suggestions for a phone) to get familiar with Dupo. Then, participants spent 30 minutes completing the main study task: to create responsive versions for their desktop design (which we replicated in Dupo) while thinking aloud. We asked them to create the phone version first and then work on tablet, thumbnail, and print versions if time allowed. While Dupo supports different directions of responsive design (e.g., mobile-first, simultaneous editing), we decided to fix the desktop-first direction for two main reasons: (1) all participants reported desktop-first as their typical design direction when asked prior to the study, and (2) our goal was to observe how they use Dupo's design suggestions along with manual editing, rather than how Dupo affects the direction of creation.

After the main task, participants self-evaluated their work and then completed a 15-minute interview. For self-evaluation, participants rated their satisfaction on a scale from 1 (Very unsatisfied) to 5 (Very satisfied): *Given the limited design time you had, how satisfied are you with this draft?* We followed up the self-evaluation by asking about any further edits they had wanted to make if given more time. In the structured interview, we asked participants' about their overall reactions to seeing the design suggestions, comparison between Dupo and their day-to-day tasks, and Dupo's features that they would want for their usual tools. We then asked about the usefulness of the *Exploration* suggestions for the phone version by showing them the suggestions again. Lastly, we asked about suggestions for general system improvement. Each participant was compensated with USD 60 (gift card). Detailed study protocols are included in the Supplementary Material.

7.2 Results

We analyzed the think-aloud and interview transcripts using top-down thematic analysis [6] given the research questions outlined above. We analyzed the think-aloud transcripts, system logs, and screen recordings together to identify design exploration patterns. We find that Dupo supports exploring a range of responsive transformations and creating satisfactory designs, yet our participants suggested improvement for direct manipulations and different ways to show *Alteration* suggestions.

Support for Design Exploration (RQ1): During the study sessions, participants made two to four responsive versions across different screen sizes that they were pleased with over a relatively short amount of time (about 30min), thereby demonstrating that Dupo supports creating multiple outcomes. We observed participants reasoning about different design suggestions in various ways. We tagged the task recordings with observations of how participants appeared to be reasoning about the designs as they interacted with the *recommender window*. First, we observed that participants reviewed *most of the suggestions*, rather than just one or two that initially appealed to them, suggesting that the recommendation sets were compelling.

Second, participants frequently compared two or more suggestions, explicitly reflecting on which one was better for their purposes. For example, comparing suggestions for a stacked bar chart with small multiples for a tablet version, E4 noted that “[The multiple bar chart]’s

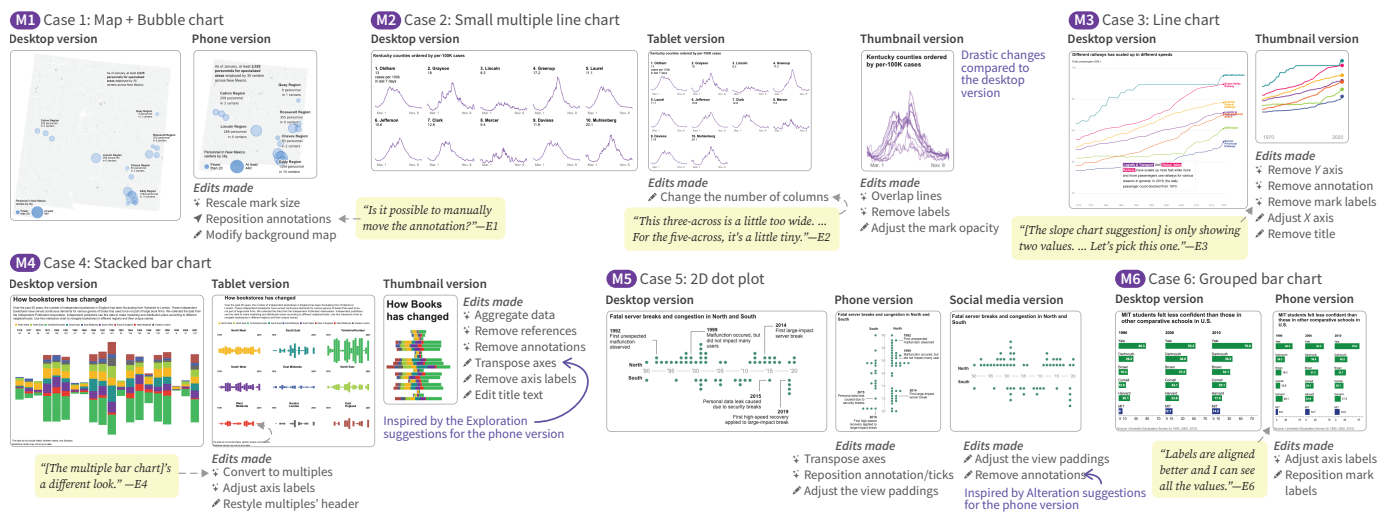


Fig. 5: Selective design outcomes from the user study with six expert responsive visualization authors. Participants provided their own visualizations for use in the study. For anonymization purposes, we replaced participants' data sets while maintaining the cardinality and data type, edited the text content but preserved annotation length, and altered the color schemes. A star icon (★) indicates an edit from the recommendations, a pencil icon (✎) denotes a manual edit made by a participant, and a cursor icon (☞) represents a manual edit using direct manipulation.

a different look. It's an emphasis on the [regions] as opposed to the overall [volume] of the [independent book stores] like you get here [with the stacked bar chart]." E4 then chose the small multiples (Figure 5 M4, Tablet). Participants also confirmed the need for authoring agency over small changes by considering possible manual edits to the suggestions, for example, by asking "Is it possible to manually move the annotation?" (E1; M1, Phone) or by saying "I'm gonna take out all the annotations for... social [media]" (E5; M5, Social).

Overall, participants seemed relatively confident in the design choices they made, which they articulated by summarizing their decisions. For instance, E3 first summarized that a slope chart suggestion for his thumbnail version "is only showing two values" and then said, "let's pick this one [with more detail]" because he had already noted that he wanted to keep the details (M3, Thumbnail).

In terms of fixation on well-known designs, participants in general made a phone version that seemed similar to their original phone versions. This fixation is not necessarily surprising given that they had already gone through rigorous design iterations in a newsroom-like setting. E2 mentioned, "It's hard to disconnect from my original decisions I made." However, participants considered designs with more drastic changes for the versions that they had previously paid less attention to (e.g., tablet, thumbnail). For example, given a small multiple line chart, E2 overlaid those lines for a thumbnail (Figure 5 M2). With a grouped bar chart for desktop, E6 chose a heatmap for a thumbnail because it seemed to "make people curious" about the content, though she still wanted further aesthetic changes. We observed design suggestions for a previous version sometimes influencing participants' designs for later versions. For example, E4 manually transposed the axes for the thumbnail M4, inspired by similar Exploration suggestions for the phone version. Likewise, E5 removed annotations for the social version M5, pointing out Alteration (see similar) suggestions for the phone version.

Mixed-initiative Usage Patterns (RQ2): Given that participants immediately saw design suggestions upon creating a new responsive artboard, they tended to choose a design suggestion first before making manual edits. Participants generally spent less time using the recommender ($m=3.0\text{min}$, $s=1.9$) than making manual edits ($m=5.3$, $s=5.8$) per version. Yet, they took relatively less time to make minimal or already brainstormed manual edits, such as E2's column adjustment for small multiples (Figure 5 M2, Tablet) and E5's social version inspired by the suggestions for the phone (M5, Social).

Dupo was used to automate changes to overall sizes (scaling the chart size, text size, etc.), layout (transposing), and encoding, while participants tended to manually edit text elements (annotations and titles) and reference objects, which are essential in narrative visualization.

For example, E4 changed the title text for text alignment and justification. After making an edit, participants tended to closely examine every element to ensure that all important information was preserved, which helped them identify what they might need to edit next. After moving annotations in the phone version, for example, E5 realized that he needed to further adjust the spacing. E6 requested recommendations after making some manual edits, and then branched an Alteration suggestion. Then, E6 made more edits to the two candidate artboards simultaneously via edit propagation, before finalizing one of them.

Participants tended to make manual edits by testing out a variety of different values for the visual encoding (e.g., by repeatedly changing the x or y value to position an annotation appropriately). When we filtered out manual edits (written as Cicero rules) dealing with the same element (i.e., the same specifier and option properties) from the log of participants' manual edits, the average number of manual edits reduced from 85.3 ($s=82.1$) to 34.8 ($s=20.5$). In contrast, participants applied an average of 28.3 ($s=12.7$) edits by the recommender. We note that each recommender rule changed multiple properties, and participants explored multiple design suggestions with at most fifteen rules. Overall, Dupo enabled participants to focus on edits that needed careful visual inspection, while expediting high-level responsive transformations.

Users' Control and Interpretability (RQ3): Participants interacted with features in the recommender window, and often read the descriptions provided in the action widget. For example, E6 actively used the hide this features, and she valued hide this because "it removes all the noise that you don't need, and you can focus on the limited options that you will get." Using the bring my existing edits feature, E6 said, "[it] is a learning experience... in terms of how I can do better [because I could] compare what I did earlier [manually] and some new suggestions." Meanwhile, E4 wished for more control in selecting a source design by asking, "Can I make a thumbnail based on the mobile version?" instead of the desktop version. E4 also wanted to use the hide this feature independently for each version.

Design Satisfaction (RQ4): Across all sessions, participants created and rated their satisfaction for 17 designs. In general, participants were satisfied (seven "very satisfied" and seven "satisfied") with the versions they created. Some responded "just okay" to their last versions because they did not have enough time to complete an initial version (e.g., E3, Print) or they wished to fine-tune the designs (e.g., E2, E4, Thumbnail).

Participants in general found our design suggestions to be reasonable and realistic potential improvements to their design process. E4 said, "[Dupo] literally made 90% of the chart I would have made by hand on my own for a mobile [phone]." E2 mentioned that his team "actually did mock up a version that was like this [heatmap]." Even though

participants sometimes thought some design suggestions were less visually appealing due to drastic changes from the desktop versions or less preferred label arrangement, all participants seemed able to understand why those examples were suggested. For example, E3 did not like transposing the line chart, yet he acknowledged that “*in some cases, it can also be helpful.*” E6 said some axis-transpose options were not aligned with her intention with the original bar grouping.

When we asked participants to rank the *Exploration* suggestions for the phone version during the follow-up interview, they exhibited several criteria, such as similarity with the source design and well-formedness. For example, E2 ranked the two suggestions with encoding changes (to a heatmap) in the third and fifth place, saying that they were “*missing the point [of] . . . what the desktop one was doing.*” Considering the well-formedness of the design, E6 said, “*Labels are aligned better and I can see all the values*” as she ranked the design suggestion she chose for the phone version (M6) in first place. Participants found some *Exploration* suggestions to be subtle in a way that we did not expect previously. While comparing aggregated and non-aggregated *Exploration* alternatives, for example, E1 was not able to figure out the difference at first. Once we clarified the aggregation, E1 ranked the aggregated version in the first place, saying, “*that’s actually really smart.*”

Feasibility and Remaining Challenges (RQ5 & RQ6): Comparing to their day-to-day tools, participants believed Dupo could help reduce the time required to prototype design ideas by providing high-fidelity designs. For example, E1 noted that “*I [had to] create three [responsive] versions. It takes up a lot of time. Even though you can copy and paste, there’s still a lot of manual work going on.*” E4 said that, “*It wasn’t just a little pencil sketch or a crappy chart that needed work. Basically, it did all that work.*” Pointing out that “*mobile is like a whole different world to think about,*” E2 said Dupo could help to “*get into that mindset much easier.*” E6 mentioned that being able to have different designs for responsive versions, which her tool (Flourish) did not support, would allow her team to consider further encodings like treemaps for desktop versions with simplified mobile versions.

However, those familiar with Adobe Illustrator or similar tools pointed out the needs for extended direct manipulation. For example, E5’s visualization was annotation-heavy, so positioning the annotation and adjusting the spacing accordingly were important tasks. E5 found it inconvenient to make these changes in Dupo due to the limited direct manipulation. Participants also needed additional support for inspecting *Alteration* suggestions as they often included more subtle changes. E3 proposed suggesting “*edits*” rather than “*designs,*” and letting users test out different edits to help them make sense of what is being changed.

8 DISCUSSION

We implemented Dupo, a mixed-initiative tool for creating responsive visualizations, to support design exploration, users’ agency over recommendation procedures, progressive authoring, and flexible artboard management. Dupo provides two-step design suggestions (*Exploration* and *Alteration*), edit *Augmentation* (*quick edits*), various user control options (e.g., the *action widget*), and edit propagation across artboards as well as between manual and automated designs. We evaluated Dupo with six expert data journalists using their own visualization cases, observing that Dupo supports exploring different design ideas, helping users to come up with satisfactory outcomes. Below, we discuss implications to other responsive visualization design approaches and future challenges for mixed-initiative authoring with adaptive systems.

8.1 Extension to Other Responsive Visualization Methods

Dupo employs a graphic-based interface, similar to Adobe Illustrator and Canva, that allows for visually manipulating chart elements; this approach has several benefits with respect to implementing the recommender. First, our approach can regularize the representation of user input for the recommender and hence easily specify the roles of chart elements, such as using our extended Vega-Lite [35] for the visualizations and Cicero [21] for responsive transformations. In contrast, users can have more freedom in representing their designs with programming-based approaches (e.g., writing code for D3.js [5]), which poses more complexity in interpreting different visualization roles. Second, Dupo

can visually demonstrate design suggestions within its graphical interface, whereas programming-based approaches may require additional rendering processes. Third, Dupo supports flexibility in customizing responsive artboards without necessarily maintaining the same mark type, encoding channels, or text elements, compared to the template-based tools (e.g., Datawrapper [9] or Flourish [1]), as pointed out by E6.

Our user study participants generally leverage a diverse set of visualization tools including template-based, programming-based, and graphic-based tools. To better support different authoring experiences, it is important to extend our mixed-initiative approach to different types of tools. As noted above, however, there are several challenges to do so for template-based and programming-based tools, illuminating future research opportunities. Using a template-based tool, for example, users choose a template of fixed responsive designs, and they can only use those designs. Instead, such tools could allow for combining each responsive version from different templates, so that a recommender can search designs more flexibly and users can have more freedom with their designs, as noted by E2 and E6. Doing so would require asking research questions like how to search over a template design space in a way that supports preserving design information across responsive versions. For programming-based approaches, recommenders could employ an extraction model [16] or AI code suggestion [15] for responsive visualization. Such approaches will benefit from real-time bundlers (e.g., Rollup [2]) that renders the outcome as a user makes changes to the code, potentially enabling immediate visual demonstration of design recommendations expressed programmatically.

8.2 Next Step: Adaptive Recommendation

Compared to analytic domains like exploratory data analysis, responsive design involves a high degree of creativity in order to craft and effectively communicate a narrative across different end-user devices. Many organizations that require responsive design may develop their own individual preferences or design patterns that are necessary to incorporate into the design process. In addition, we will discover new design patterns for responsive visualization as the area is evolving, demanding an extended search space. To make this process easier, responsive visualization tools can employ an adaptive recommender that updates its search space by learning design patterns from users. While Dupo offers implicit user constraints using *hide-this*, adaptive mixed-initiative systems can learn and share new responsive techniques across different users. This approach opens up several challenges for future work in addition to developing an adaptive mechanism. For example, such a system may need to determine whether to learn a design pattern of a user and whether one user’s technique would be useful for other users. It is also necessary to ask how to support users to easily control the outcome of the recommenders with a growing search space.

8.3 Limitations and Future Work

Dupo currently considers common mark types, encodings, and interactions supported by Vega-Lite [35]. While we outlined high-level rationales for generating *Exploration* design suggestions, we do not claim that our search space is exhaustive. Thus, future work could extend Dupo with regard to encodings and design suggestions. Next, our user study only considered the desktop-first approach. A longer term user study could observe how mixed-initiative approaches impact responsive visualization authoring in different settings like collaborative creation or simultaneous editing and improve the system with adaptive learning. Lastly, we qualitatively evaluated Dupo to elicit perspectives of expert users. Future work could compare future responsive visualization tools with Dupo as a benchmark.

9 CONCLUSION

We presented Dupo, a mixed-initiative responsive visualization authoring tool, that augments users’ design exploration with recommendations to support authoring multiple artboards while retaining user control. Our qualitative evaluation with six expert users on their pre-existing, real-world responsive use cases shows that Dupo’s two-step recommendation pipeline can support exploring responsive designs by reducing redundancy and tedious work required for prototyping alternatives.

ACKNOWLEDGMENTS

Special thanks to Eunye Koh and Priyanka Nanayakkara for their valuable feedback and support.

SUPPLEMENTARY MATERIAL

In the Supplementary Material (<https://see-mike-out.github.io/dupo-supplementary>), we provide (1) details for our pilot study, (2) the Dupo system documentation (user manual), (3) a video demo of Dupo, (4) a walkthrough video of using Dupo, (5) high-level rationales for each *Exploration* suggestion, (6) details for the ranking measures, (7) the main study protocols, and (8) the study outcomes (visualization designs and suggestions). A public demo of Dupo is available at <https://see-mike-out.github.io/dupo-public>.

REFERENCES

- [1] Flourish, n.d. <http://flourish.studio/>. 1, 9
- [2] Rollup, n.d. <https://rollupjs.org/>. 9
- [3] Svelte, n.d. <https://svelte.dev/>. 7
- [4] K. Andrews. Responsive visualisation, 2018. https://mobilevis.github.io/assets/mobilevis2018_paper_4.pdf. 1
- [5] M. Bostock, V. Ogievetsky, and J. Heer. D³ data-driven documents. *IEEE Trans. Visual Comput. Graphics*, 17(12):2301–2309, 2011. doi: 10.1109/TVCG.2011.185 1, 9
- [6] V. Braun and V. Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006. doi: 10.1191/1478088706qp063oa 7
- [7] K. Cook, N. Cramer, D. Israel, M. Wolverton, J. Bruce, R. Burtner, and A. Endert. Mixed-initiative visual analytics using task-driven recommendations. In *Proc. IEEE VAST*, pp. 9–16. IEEE, New York, 2015. doi: 10.1109/VAST.2015.7347625 2
- [8] W. Cui, J. Wang, H. Huang, Y. Wang, C.-Y. Lin, H. Zhang, and D. Zhang. A mixed-initiative approach to reusing infographic charts. *IEEE Trans. Visual Comput. Graphics*, 28(1):173–183, 2022. doi: 10.1109/TVCG.2021.3114856 2
- [9] Datawrapper GmbH. Datawrapper, 2013. <https://www.datawrapper.de/>. 1, 2, 9
- [10] E. Di Giacomo, W. Didimo, G. Liotta, and F. Montecchiani. Network visualization retargeting. In *Proc. IISA*, pp. 1–6. IEEE, New York, 2015. doi: 10.1109/IISA.2015.7388095 1, 2
- [11] G. Ellis and A. Dix. Enabling automatic clutter reduction in parallel coordinate plots. *IEEE Trans. Visual Comput. Graphics*, 12(5):717–724, 2006. doi: 10.1109/TVCG.2006.138 4
- [12] N. Elmqvist and J. Fekete. Hierarchical aggregation for information visualization: Overview, techniques, and design guidelines. *IEEE Trans. Visual Comput. Graphics*, 16(3):439–454, 2010. doi: 10.1109/TVCG.2009.84 2
- [13] M. Gebser, R. Kaminski, B. Kaufmann, and T. Schaub. Clingo = ASP + control: Preliminary report. 2014. <https://arxiv.org/abs/1405.3694>. 4, 7
- [14] M. Gebser, B. Kaufmann, R. Kaminski, M. Ostrowski, T. Schaub, and M. Schneider. Potassco: The potsdam answer set solving collection. *AI Commun.*, 24(2):107–124, 2011. doi: 10.3233/AIC-2011-0491 2, 4, 7
- [15] GitHub. Copilot, n.d. <https://github.com/features/copilot>. 9
- [16] J. Harper and M. Agrawala. Deconstructing and restyling d3 visualizations. In *Proc. UIST*, p. 253–262. ACM, New York, 2014. doi: 10.1145/2642918.2647411 9
- [17] C. Healey, S. Kocherlakota, V. Rao, R. Mehta, and R. St. Amant. Visual perception and mixed-initiative interaction for assisted visualization design. *IEEE Trans. Visual Comput. Graphics*, 14(2):396–411, 2008. doi: 10.1109/TVCG.2007.70436 1, 2, 3
- [18] B. Hinderman. *Building Responsive Data Visualization for the Web*. John Wiley & Sons, 2015. 1
- [19] J. Hoffswell, W. Li, and Z. Liu. Techniques for flexible responsive visualization design. In *Proc. CHI*, pp. 1–13. ACM, New York, 2020. doi: 10.1145/3313831.3376777 1, 2, 3, 4, 5
- [20] H. Kim, D. Moritz, and J. Hullman. Design patterns and trade-offs in authoring communication-oriented responsive visualization. *Comput. Graphics Forum (Proc. EuroVis)*, 40(3):459–470, 2021. doi: 10.1111/cgf.14321 1, 2, 3, 4, 5
- [21] H. Kim, R. Rossi, F. Du, E. Koh, S. Guo, J. Hullman, and J. Hoffswell. Cicero: A declarative grammar for responsive visualization. In *Proc. CHI*. ACM, New York, 2022. doi: 10.1145/3491102.3517455 2, 4, 7, 9
- [22] H. Kim, R. Rossi, A. Sarma, D. Moritz, and J. Hullman. An automated approach to reasoning about task-oriented insights in responsive visualization. *IEEE Trans. Visual Comput. Graphics*, 28(1):129–139, 2022. doi: 10.1109/TVCG.2021.3114782 1, 2, 3, 4
- [23] J. L. Kolodner and L. M. Wills. Case-based creative design. Technical Report SS-93-01, AAAI, 1993. 2
- [24] J. L. Kolodner and L. M. Wills. Powers of observation in creative design. *Design Studies*, 17(4):385–416, 1996. doi: 10.1016/S0142-694X(96)00021-X 2
- [25] C. Körner. *Learning Responsive Data Visualization*. Packt Publishing, 2016. 1
- [26] J. Leclaire and A. Tabard. R3s.js—towards responsive visualizations. In *Workshop on Data Exploration for Interactive Surfaces DEXIS 2015*, pp. 16–19, 2015. 1
- [27] H. Lin, D. Moritz, and J. Heer. Dziban: Balancing agency & automation in visualization design via anchored recommendations. In *Proc. CHI*, p. 1–12. ACM, New York, 2020. doi: 10.1145/3313831.3376880 1, 2, 3
- [28] J. S. Linsey, I. Tseng, K. Fu, J. Cagan, K. L. Wood, and C. Schunn. A study of design fixation, its mitigation and perception in engineering design faculty. *J. Mech. Des.*, 132(4):041003, 2010. doi: 10.1115/1.4001110 1, 2
- [29] R. Ma, H. Mei, H. Guan, W. Huang, F. Zhang, C. Xin, W. Dai, X. Wen, and W. Chen. Ladv: Deep learning assisted authoring of dashboard visualizations from images and sketches. *IEEE Trans. Visual Comput. Graphics*, 27(9):3717–3732, 2021. doi: 10.1109/TVCG.2020.2980227 1, 2, 3
- [30] Microsoft. Power bi, 2011. <https://powerplatform.microsoft.com/en-us/>. 2
- [31] P. O’Donovan, A. Agarwala, and A. Hertzmann. DesignScope: Design with interactive layout suggestions. In *Proc. CHI*, pp. 1221–1224. ACM, New York, 2015. doi: 10.1145/2702123.2702149 1, 3
- [32] Z. Qu and J. Hullman. Keeping multiple views consistent: Constraints, validations, and exceptions in visualization authoring. *IEEE Trans. Visual Comput. Graphics*, 24(1):468–477, 2017. doi: 10.1109/TVCG.2017.2744198 2
- [33] S. Ramírez. FastAPI, n.d. <https://fastapi.tiangolo.com/>. 7
- [34] R. Rosenbaum, J. Zhi, and B. Hamann. Progressive parallel coordinates. In *Proc. PacificVis*, pp. 25–32. IEEE, New York, 2012. doi: 10.1109/PacificVis.2012.6183570 2
- [35] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-lite: A grammar of interactive graphics. *IEEE Trans. Visual Comput. Graphics*, 23(1):341–350, 2017. doi: 10.1109/TVCG.2016.2599030 7, 9
- [36] A. Swearngin, C. Wang, A. Oleson, J. Fogarty, and A. J. Ko. Scout: Rapid exploration of interface layout alternatives through high-level design constraints. In *Proc. CHI*, pp. 1–13. ACM, New York, 2020. doi: 10.1145/3313831.3376593 1
- [37] Tableau Software. Tableau, 2003. <https://www.tableau.com/>. 2
- [38] A. Tse. ai2html, 2011. <http://ai2html.org/>. 1
- [39] E. Wall, S. Das, R. Chawla, B. Kalidindi, E. T. Brown, and A. Endert. Podium: Ranking data using mixed-initiative visual analytics. *IEEE Trans. Visual Comput. Graphics*, 24(1):288–297, 2018. doi: 10.1109/TVCG.2017.2745078 1, 2
- [40] K. Wongsuphasawat, D. Moritz, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE Trans. Visual Comput. Graphics*, 22(1):649–658, 2016. doi: 10.1109/TVCG.2015.2467191 2, 3
- [41] K. Wongsuphasawat, Z. Qu, D. Moritz, R. Chang, F. Ouk, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager 2: Augmenting visual analysis with partial view specifications. In *Proc. CHI*, p. 2648–2659. ACM, New York, 2017. doi: 10.1145/3025453.3025768 1, 2, 3
- [42] A. Wu, W. Tong, T. Dwyer, B. Lee, P. Isenberg, and H. Qu. MobileVisFixer: Tailoring web visualizations for mobile phones leveraging an explainable reinforcement learning framework. *IEEE Trans. Visual Comput. Graphics*, 27(2):464–474, 2021. doi: 10.1109/TVCG.2020.3030423 1, 2
- [43] A. Wu, L. Xie, B. Lee, Y. Wang, W. Cui, and H. Qu. Learning to automate chart layout configurations using crowdsourced paired comparison. In *Proc. CHI*. ACM, New York, 2021. doi: 10.1145/3411764.3445179 2
- [44] Y. Wu, X. Liu, S. Liu, and K. Ma. ViSizer: A visualization resizing framework. *IEEE Trans. Visual Comput. Graphics*, 19(2):278–290, 2013. doi: 10.1109/TVCG.2012.114 1, 2
- [45] ZingSoft. Zingchart, 2009. <https://www.zingchart.com/>. 2