



User-Driven Support for Visualization Prototyping in D3

Hannah K. Bako
hbako@cs.umd.edu
University of Maryland
College Park, MD, USA

Alisha Varma
alishav@umd.edu
University of Maryland
College Park, MD, USA

Anuoluwapo Faboro
afaboro@umd.edu
University of Maryland
College Park, MD, USA

Mahreen Haider
mhaider1@umd.edu
University of Maryland
College Park, MD, USA

Favour Nerrise
fnerrise@umd.edu
University of Maryland
College Park, MD, USA

Bissaka Kenah
bkenah@umd.edu
University of Maryland
College Park, MD, USA

John P. Dickerson
john@cs.umd.edu
University of Maryland
College Park, MD, USA

Leilani Battle
leibatt@cs.washington.edu
University of Washington
Seattle, WA, USA

ABSTRACT

Templates have emerged as an effective approach to simplifying the visualization design and programming process. For example, they enable users to quickly generate multiple visualization designs even when using complex toolkits like D3. However, these templates are often treated as rigid artifacts that respond poorly to changes made outside of the template's established parameters, limiting user creativity. Preserving the user's creative flow requires a more dynamic approach to template-based visualization design, where tools can respond gracefully to users' edits when they modify templates in unexpected ways. In this paper, we leverage the *structural similarities* revealed by templates to design resilient support features for prototyping D3 visualizations: **recommendations** to suggest complementary interactions for a users' D3 program; and **code augmentation** to implement recommended interactions with a single click, even when users deviate from pre-defined templates. We demonstrate the utility of these features in Mirny, a design-focused prototyping environment for D3. In a user study with 20 D3 users, we find that these automated features enable participants to prototype their design ideas with significantly fewer programming iterations. We also characterize key modification strategies used by participants to customize D3 templates. Informed by our findings and participants' feedback, we discuss the key implications of the use of templates for interleaving visualization programming and design.

CCS CONCEPTS

• **Human-centered computing** → **Visualization toolkits**; *User interface programming*.



This work is licensed under a Creative Commons Attribution International 4.0 License.

IUI '23, March 27–31, 2023, Sydney, NSW, Australia
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0106-1/23/03.
<https://doi.org/10.1145/3581641.3584041>

KEYWORDS

templates, visualization prototyping, visualization design, programming, visualization tool

ACM Reference Format:

Hannah K. Bako, Alisha Varma, Anuoluwapo Faboro, Mahreen Haider, Favour Nerrise, Bissaka Kenah, John P. Dickerson, and Leilani Battle. 2023. User-Driven Support for Visualization Prototyping in D3. In *28th International Conference on Intelligent User Interfaces (IUI '23)*, March 27–31, 2023, Sydney, NSW, Australia. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3581641.3584041>

1 INTRODUCTION

Creating thoughtful and captivating data visualizations necessitates careful consideration of multiple design alternatives as designers program their visualizations. This process often requires an iterative exercise of refining one or more design options until an optimal visualization design is achieved i.e., *visualization prototyping* [3, 10, 43, 66].

This paper explores the interplay between *visualization programming*, i.e., creating interactive visualizations via languages, and *visualization design*, i.e., creating visualizations via design tools, in supporting the *visualization prototyping* process. With visualization programming, a user can create dozens of interactive visualizations just by permuting a few core variables within a single program. We see these benefits in well-designed languages like D3 that have thousands of visualization programs observed online [7, 8]. However, users must first be knowledgeable of the underlying language, which can be a huge hurdle for individuals with limited or no programming experience [8, 57, 66]. For example, even though interactions are considered essential to defining information visualization [16], interactive components such as brush filters, panning and zooming, etc., are notoriously difficult to implement in D3 even for experienced users [59, 75]. In contrast to visualization programming, visualization design encourages users to explore how to visualize data using visual elements which may be outside the bounds of traditional programming contexts [70]. For example, Data Illustrator [38] and Lyra/Lyra2 [56, 75] enable users to quickly create highly customized and detailed visualizations. That being

said, the process of evaluating multiple alternative designs can still be an arduous task in design tools since the user must edit the designs themselves to generate alternatives.

Visualization programming and visualization design contribute complementary yet distinct benefits to the visualization prototyping process that users struggle to achieve with only one or the other. Particularly, it is difficult to balance out the programming effort needed to explore the space of possible visualization designs and accompanying interactions. The difficulty of maintaining this balance points to a core question that drives our work: How can we make it easier for users to prototype their interactive design ideas through complex visualization languages like D3?

One solution is to abstract visualization examples into general-purpose programming templates [28, 45], and generate code recommendations based on these pre-defined structures (e.g. Falx [67]). In prior work, programming templates are often treated as pre-defined *static* structures, where users can only modify fixed design parameters to create visualizations. However, users may be interested in changing multiple parts of a template such as aesthetics (e.g., color schemes), data encodings, interaction styles (e.g., slider versus brush filters), and mark types. Moreover, these parameters are chosen by the tool developers, not their users. Additionally, prior work finds that designers often dissect existing examples to extract new functionality they want to build on [3, 8]; for example, users could start working with a bar chart example but eventually transform it into a completely different visualization. By limiting how users can modify a template, we hamper these creative design activities.

In this paper, we propose a *dynamic* approach to template-based design for interactive visualizations. Rather than restricting how templates can be modified, our goal is to identify the *design patterns* encapsulated within these templates, and develop automated features to identify and extend these patterns as D3 users program in real-time. These features must be *resilient*, meaning they should work regardless of whether the user is populating template code or writing code from scratch. To achieve this, we leverage a suite of D3 templates derived in our prior work [4] to build two automated support features for D3: (1) a user-driven recommendation model that suggests suitable interaction[s] to add to a user's visualization code and (2) a code augmentation feature that can integrate code snippets directly into live user code. In this way, we can help users *preserve their creative design thinking* as they code, and work towards closing the gap between visualization design and visualization programming. To demonstrate the utility of these features, we implemented them in Mirny, a design-focused environment for rapid prototyping of interactive visualizations in D3 (shown in Figure 1).

We evaluated these features through a study with 20 D3 users prototyping interactive visualizations in Mirny. We find that using the automated features implemented in Mirny, participants were able to prototype design ideas in less time and with fewer programming iterations compared to their typical D3 implementation workflows. We observed participants modifying templates in a variety of interesting ways, which we synthesize into three high-level modification strategies: changing the visual styles (e.g. mark colors or mark types), introducing data transformations (e.g., to aggregate or sort data), and removing or adding new data encodings.

In summary, this paper makes the following contributions:

- We present a user-driven *model for recommending interactions* users can implement within a D3 prototype
- We develop a code augmentation feature to *automatically integrate interaction templates* into live D3 code.
- We implement these features in Mirny and evaluate Mirny through a *user study with 20 D3 users*. Our results show Mirny enabled participants to implement interactive D3 visualizations in less time and with fewer iterations.
- We synthesize *three high-level modification strategies* used by designers to customize visualization templates.

Using our findings, we suggest research opportunities to enhance template-based design approaches within visualization languages. All our data and code are shared on OSF ¹

2 RELATED WORK

Our work draws from research in 1) visualization authoring and recommendations, 2) interactive visualizations, and 3) template and code generation. We review the relevant literature for each topic below.

2.1 Visualization Authoring and Recommendation Tools

There are a plethora of visualization authoring tools available to users such as specification-based languages (e.g [13, 33, 58, 60]), automated tools (e.g [56, 62, 75]), or graphic-driven tools that rely on GUI widgets (e.g [38, 75]). Specification languages are expressive but require considerable programming skills to use [19, 31, 64]. On the other hand, graphic-driven tools eliminate the need for programming. However, by eliminating the need for specifications, they take away the control from the user [30] and still have a learning curve [57].

To help users navigate the space of design options, recommendation tools generate design alternatives for users to choose from, rather than have users create alternatives from scratch [19, 31, 39, 41, 47, 68, 69, 71, 73]. More recently, research has explored the space of generative design tools which allow users to generate and evaluate design alternatives for glyphs [14] and dashboards [40]. Some have explored the data-agnostic design of visualization prototypes, eliminating the data constraints that designers often face when coming up with visualization designs [63, 65].

However, these systems do not account for current practices that visualization designers engage in when iterating over visualization designs. Recent work finds that when designers iterate over visualization designs, they often integrate multiple design ideas [3] or integrate code fragments from multiple examples into a single design [4, 8]. In this regard, we deviate from prior work by shifting our focus from the design of static visualizations from a single template to instead exploring how we can support the integration of multiple code templates during visualization prototyping.

¹https://osf.io/97mru/?view_only=c89bc8b4f1a64799ad4453213a255744

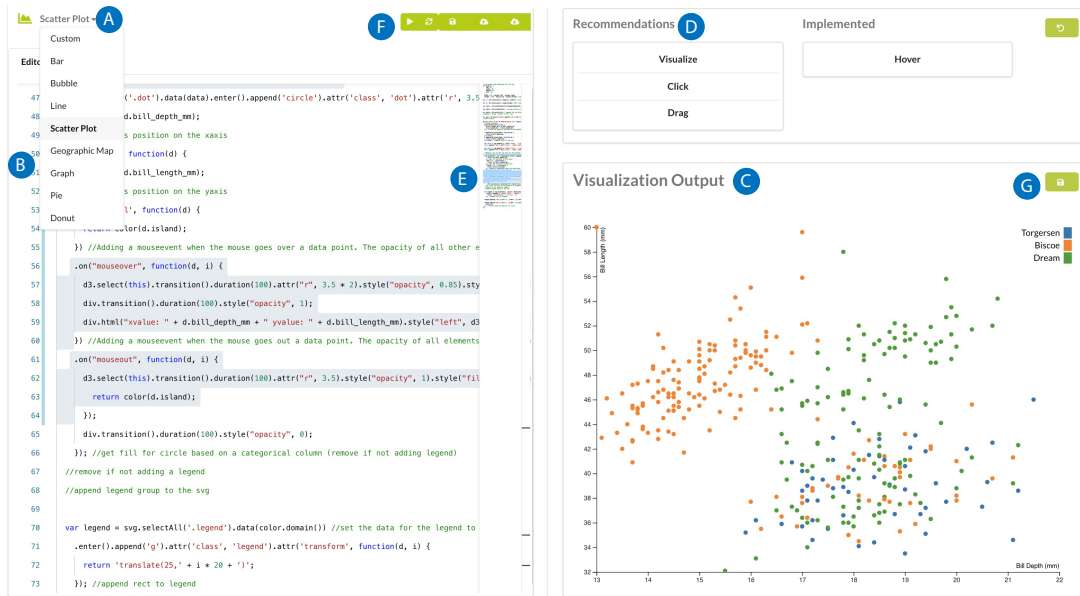


Figure 1: An example of an interactive visualization designed using Mirny. A user can select a visualization such as a *scatterplot* to create from the Template Panel (A). The code responsible for creating the *scatterplot* is displayed within the Editor (B) and the corresponding visualization is rendered in the Visualization Panel (C). A list of recommended interactions is displayed to the user in the Recommendation Panel (D). A user can then choose a recommended interaction such as a *Hover* and the code responsible for implementing the interaction is added to the interface and highlighted for the user to view and modify (E). Once a user is satisfied with their visualization, they can export the code and visualization using the Controls (F) and (G).

2.2 Interactive Visualizations

Interactions are critical for helping users gain a deeper understanding of a dataset through manipulation [20, 55, 72]. However, visualization toolkits like Chart.js² and recommendation systems generally focus on variation in visual encodings and only provide default interactions such as tooltips and clicks [41, 68]. Interaction+ [46] addresses this challenge by extracting information about visual objects from web pages and using this information to provide a variety of interactions for manipulating the existing visualization. However, the interaction implementation process in Interaction+ is still abstracted from users. Lyra 2 allows users to specify interactions through demonstration [75] but hides the underlying specification. Additionally, the space of appropriate interactions for a visualization design is vast [4, 72] and users often find it hard to identify what interactions to include in their designs and how to implement these interactions [8]. Therefore, we observe a need to support rapid exploration and specification of interactions for complex languages like D3.

We aim to help users design and prototype interactive visualizations by giving users access to the actual code. To do this, we need to model the space of possible interaction designs, which is often represented through interaction task taxonomies; we defer to existing surveys for more details, e.g., [15, 24]. In our previous work, we developed a taxonomy of D3 interactions [4] based on a subset of tasks from Brehmer and Munzner’s typology [15] to model the space of interactions a user may include in their visualizations.

²A simple Javascript library for creating charts (<https://www.chartjs.org/>)

2.3 Templates and Code Generation

General code reuse through templates has been advocated for [36, 37] and leveraged to enable semantic-based code search and code completion [53]. Work such as Reiss’s code search approach [53] has informed the development of tools for finding relevant code snippets (e.g., [27, 34, 44]) and recommending significant code templates to complete a user’s program (e.g., [48, 50, 52]). Within the visualization community, certain tools use templates to generate visualization designs [42, 45] and styles [28]. For example, McNutt et al. use parameterized declarative templates as a means of abstracting declarative visualization grammars to promote reuse, exploration, and simplification of the visualization programming process [45]. Templates have also been used in tools like Wrex [21], Falx [67], and Mage [35] to provide readable code in response to users’ interactions with GUI widgets. However, current work focuses on templates as static artifacts that offer limited control to users. We introduce a *dynamic* approach that detects template-like patterns within a user’s visualization program and can integrate design recommendations automatically, even when the user writes code from scratch.

3 MOTIVATING USAGE SCENARIO FOR MIRNY

Sandra is a student who wants to create some interesting visualizations for the final project in her data science course. Although this course was broad in scope, it lacked depth in terms of learning interactive visualization techniques. Sandra has heard that D3 can

support a wide variety of interactive visualizations. Although she has limited experience using D3, Sandra wants to create a D3 visualization, preferably with a simple dataset to start, such as the palmerpenguins dataset [25]. Sandra decides to use Mirny, shown in Figure 1, to program her visualizations. Sandra loads Mirny and imports her dataset, she chooses a Scatterplot from Mirny's Template Panel (Figure 1a). Mirny automatically generates the D3 code for a Scatterplot from the data as shown in (Figure 1b), and the resulting visualization is rendered in the output panel (Figure 1c). Mirny recommends a ranked list of common interactions for Scatterplots (Figure 1d).

Sandra sees nothing wrong with the scatterplot, but she wonders what the distribution of penguins is across the three islands recorded in the dataset. For her second prototype, she selects a bar chart from Mirny's Template Panel. Mirny automatically generates the D3 code for a bar chart a populates the Editor (Figure 1a), saving Sandra the effort of figuring out how to change her design to a different visualization type. Sandra edits the code to use the island attribute for the x-axis and the count of penguins for the Y-axis. The overall visualization looks okay, but Sandra still isn't satisfied. She navigates back her Scatterplot prototype using the Template Panel and changes the color encoding from the species attribute to the island attribute.

Sandra likes the current visualization design but wants to add an interaction. Sandra clicks on the first item from the recommendation panel (*hover*) and Mirny automatically extends her current D3 program with the corresponding code, highlighting the new code block and providing a summary of its behavior in the Editor (Figure 1e). The new code displays additional information about a data point when a hover event is triggered. However, Sandra notices that some of the data points are really close to each other. She also adds a zoom interaction from the recommendation panel, so her viewers can zoom into tightly packed sections of the visualization. Once Sandra completes her visualization prototypes, she exports the code from the editor as a JS file and saves the visualization output in SVG format.

Using Mirny, a D3 user like Sandra can quickly prototype multiple interactive visualizations. Mirny eases the cumbersome process of programming D3 code, enabling users to explore visualization design alternatives using a single interactive interface. Although this example focuses on a new D3 user (Sandra), it can also easily apply to a more advanced user. For example, a professional analyst who has used D3 in several projects may find it tedious to edit an old example from a previous project, and may still prefer to fill in one of Mirny's templates instead. We discuss the experiences of both new and more advanced D3 users in section 9.

4 PRELIMINARY STUDY: LESSONS FROM ANALYZING D3 EXAMPLES ONLINE

The goal of our work is to streamline the process of prototyping visualization designs with complex toolkits like D3. To do this, we first need to understand *what* kinds of visualizations D3 users create and *how* these users implement these visualizations in D3. In our prior work [4], we qualitatively analyzed 2500 D3 examples shared on GitHub [51], BLocks.org [11] and Observable [12]. Through this analysis, we investigated common implementation patterns D3

users adopt when creating interactive visualizations. In this section, we summarize the key results from this analysis and describe how these results influence design considerations for the Mirny system. **Takeaway 1: A Small Subset of Visualizations are Commonly Implemented in D3.** To better understand which visualizations tend to be implemented in D3, we counted the incidence of visualization types across our corpus using Battle et al.'s D3 visualization taxonomy [4, 7]. Our observations reveal that 80% of the D3 visualizations shared online can be accounted for by just five visualization types (see Figure 2a): Bar charts, Geomaps, Line charts, Scatterplots, and Force Directed Graphs. Our findings suggest that *a small subset of visualizations cover the majority of D3 users' needs.*

Takeaway 2: D3 Implementation Patterns are Consistent Within Visualization Types. In a subsequent round of analysis, we divided our corpus by visualization type and analyzed the code for similarities and differences across examples, such as differences in the rendering strategy or differences in which interactions were implemented [4]. We found that in most cases, *the underlying structure of the D3 code examples did not vary.* In other words, the underlying structure and API calls remained the same within each visualization type. As a result, we observe clear patterns in how different D3 users implement the same types of visualizations, with variations made to suit users' design preferences. We also found that D3 users copy code from existing examples, and even attribute the original sources in their visualization programs, corroborating observances of code reuse in the D3 community by prior work [8]. To make our findings actionable for the visualization community, we curated and shared templates for 8 popular D3 visualization type³.

Takeaway 3: Interactions are Commonly Implemented in D3 and Vary by Visualization Type. Interactions are known to be challenging to implement in D3 [8, 59], and thus are a promising focus for design templates. To better understand how D3 users currently implement interactions, we classified observed interactions in our corpus using a subset of Brehmer and Munzer's multi-level typology of visualization tasks [15]. We observed 6 common interaction widget types: *Brush*, *Click*, *Drag*, *Hover*, *Visualize*, and *Zoom*. 52% of the examples we classified contained at least one or more combinations of interaction widgets (e.g *Hover*, *Click*, *Zoom*).

First, our findings show that interactions involving the selection of data points such as *Hover* are used in the vast majority of interactive examples (62%, see Figure 2b). Second, some interactions seem to be coupled with specific visualization types; for example, D3 network graphs often support the drag interactions. Finally, the total interactions implemented varied by visualization type. For instance, network graphs tended to have at least one interaction whereas Line charts tended to have no interactions implemented. This suggests that in D3, *interaction programming is dependent on the visualization design process, requiring support features that are both design-aware and code-aware.*

4.1 Design Considerations

Based on these results and our personal experiences using D3, we provide three design considerations for Mirny.

³The curated templates can be found at https://osf.io/k58bp/?view_only=72fa3798bbba4263b5ad662b26a70cb3

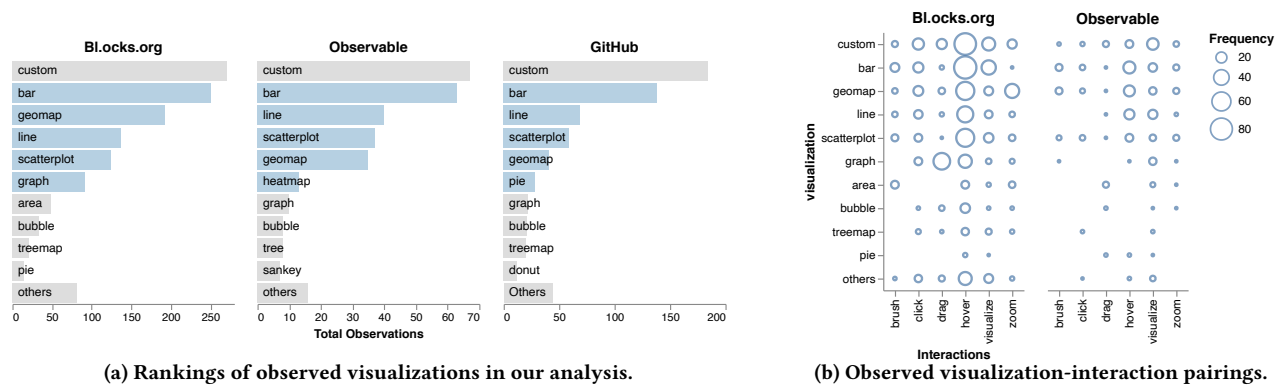


Figure 2: Results from the analysis carried out on the three example corpus [4] (a) presents the frequency of the top visualization types and their corresponding interactions observed in our datasets. Visualizations with 1% and less frequency are captured in the "others" category. (b) shows the number of times different interaction types were implemented for each visualization type in our corpus.

C1: Prioritize familiar interfaces/workflows. Past research highlights a need for tools to be integrated into users' workflows [8, 21] and familiar interfaces [23]. Furthermore, compatibility with other development platforms was one of the original design goals of D3 [13]. For these reasons, we prioritize compatibility with users' existing workflows and preferred interfaces, reducing the burden of introducing yet another new tool to do their design prototyping process. Specifically, Mirny is designed to use either visualization templates *or* existing code as input and makes it easy to export code so it can be used in other tools such as Observable [12]. We summarize the Mirny interface design and explain how it supports visualization prototyping in section 5.

C2: Design adaptive recommendations to match current practice. Our prior work reveals that with the widespread implementation of D3 visualizations comes myriad conventions for when and how to implement interactions for specific D3 visualization types [4]. To design effective interactions, D3 users have to navigate a complex space of existing design examples to identify compatible implementations. By providing automated recommendations for compatible interactions (inspired by existing recommenders, e.g., [26, 61, 64]), we can help users quickly prototype their envisioned designs, rather than forcing them to scour the internet for relevant examples. We present our recommendation engine to support interaction prototyping in section 6.

C3: Support automated code augmentation. Prior work [8, 29] has established code reuse as a prominent implementation strategy employed by D3 users. In our past work, we developed code templates for common data visualizations and interactions that D3 users implement [4]. However, the templates alone are not enough as users may have to modify or add to the templates to realize their target designs. To facilitate template modification and customization, users need automated integration of template[s] into existing user code, which we refer to as *code augmentation*. This way, a user can prototype multiple interactive visualization designs with a simple mouse click. We explain how we support automated code augmentation in section 7.

5 MIRNY SYSTEM OVERVIEW

We implement two automated features in Mirny (recommendations and code augmentation) to facilitate interactive visualization prototyping in D3. These features are situated in a design-focused interface structured for integration with existing user workflows (addressing design considerations C1 and C2). In this section, we describe Mirny's overall design (see Figure 3).

The Mirny web interface provides a user with an interactive graphical user interface through which they can interact with the Mirny system. The web interface consists of four core views: **Templates Panel** (Figure 1A), **Editor** (Figure 1B), **Visualization Panel** (Figure 1C) and the **Recommendation Panel** (Figure 1D).

Templates Panel. The Templates Panel is designed to help D3 users quickly generate an initial visualization prototype for a given dataset. Revisiting our example in section 3: after Sandra uploads her dataset, she selects a Scatterplot from Mirny's list of templates. Mirny uses this template to identify attributes with compatible data types and generate an initial prototype visualization, which Sandra can view in the Editor (code) and Visualization Panel (output). Mirny supports templates for the 8 most popular visualization types identified in our prior work [4]: Bar charts, Line charts, Scatterplots, Geographic Maps, Force Directed Graphs, Pie Charts, Donut charts, and Bubble plots. Mirny also leverages templates for 6 interaction types (Hover, Click, Brush, Drag, Zoom, and Visualize), which are displayed in the Recommendation Panel.

Editor. The Mirny Editor is a live JavaScript editor tailored for prototyping interactive D3 visualizations. Given our focus on integration with existing user workflows, the Editor is designed in a similar style to complementary editors, such as those in Vega-Lite [59] and Observable [12]. For example, when Sandra modifies the code for her Scatterplot prototype in Mirny, it is as if she is making the change with similar editors, smoothing Mirny's learning curve.

Visualization Panel. Mirny displays the visual output for the current prototype in the Visualization Panel. This panel updates live and is fully interactive, so when a user modifies the code, they can quickly test the results in the Visualization Panel. For example,

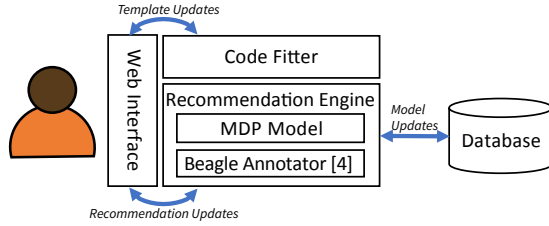


Figure 3: Mirny’s architecture. The Mirny interface captures user requests for interactions and recommendations. Using the Code Fitter, it generates the appropriate D3 specification and renders the results for the user via the Interface. Requests are also sent to the prediction model to receive recommended interactions for the user.

when Sandra adds a Hover interaction to her Scatterplot, she can immediately test the output using this panel.

Recommendation Panel. In parallel to the user’s coding process, Mirny also recommends compatible interaction designs that could be added to the current prototype. These recommendations are displayed in the Recommendation Panel, ranked by their relevance to the target visualization type. Mirny also shows which interactions have already been implemented, such as Sandra’s selected Hover interaction. The recommendations are powered by Mirny’s recommendation engine, described in section 6. Users can click on any recommendation and it will automatically be integrated into their current prototype using Mirny’s code augmentation feature, described in section 7.

6 RECOMMENDING CODE SNIPPETS TO MAKE D3 VISUALIZATIONS INTERACTIVE

Given the prevalent combinations of visualization and interaction types observed in our preliminary analysis (see section 4), D3 users are clearly interested in *interactive* visualizations. However, it may not be obvious to D3 users which interactions complement a given visualization. Revisiting our motivating example, Sandra may not know which interactions to implement to enhance her Scatterplot, such as Hover, Zoom, or Brush.

One approach to solving this problem is to *automatically recommend compatible interactions* anchored on the implementation practices of past D3 users, for example, by training machine learning models to predict complementary interactions for a given visualization (addressing design consideration C2). However, users’ programming behaviors and design preferences evolve as their interests shift over time [26, 64]. Thus, a one-shot approach to model training risks generating stale results as users advance their visualization projects. To keep up with the evolving nature of user preferences, our recommendations must also *adapt* to shifts in user behavior over time. Hence, we leverage Markov Decision Process (MDP) [9] to model users’ interaction preferences and recommend complementary interactions to users as they program D3 visualizations. Here, we provide an overview of the prediction models that drive Mirny’s recommendations.

6.1 Predicting Which Interactions to Implement

Modeling Interaction States. To make recommendations, we first need an understanding of the current state of a user’s visualization code and use that information to predict what interactions to recommend to the user. We define the MDP for our recommendation engine as a set of *states* S , where each state $s \in S$ is a set of the interactions that have already been implemented in the user’s code $\{i_1, i_2 \dots i_n\}$. Given that the user could pick any number of interactions to implement, we create a separate state for every possible combination of interactions that we observed from Bl.ocks.org. Furthermore, this interaction set often starts out empty as D3 users seem to implement functional visualizations before incorporating interactions. For this reason, we make the empty state the initial state for our model.

For any given state $s \in S$, our MDP model produces a ranked list of valid interaction recommendations, ordered by how likely they are to be implemented based on observed frequencies of interaction types among our D3 examples (see section 4).

Adjusting the Model Based on User Feedback. Given our interaction states, we model potential interaction implementation steps through a set of *user actions* A , where each action $a \in A$ represents a user’s reaction to a recommended interaction i_r . The user can react in one of four ways, given the user’s current state s , and the corresponding future state s' :

- **Accept** the recommendation by clicking on the corresponding button in the Mirny interface (see Figure 1D), thereby moving to state s' , given by the equation:

$$P[s, i_r, s'] = \frac{\text{observations}(s')}{\text{observations}(s)}, \quad (1)$$

where *observations* computes the total observations of s (i.e., the combination of interactions) in our corpus.

- **Export** the visualization code files (see Figure 1E) and exiting Mirny. This is a cue that the user may be satisfied with her visualization and wants to save it for future use, calculated as:

$$P[s, i_r, \text{export}] = \frac{\text{observations}(s')}{\sum_{s \in S} \text{observations}(\text{export}|s)} \quad (2)$$

- **Undo** an added interaction, suggesting the user was dissatisfied with the interaction and wants to remove it:


$$P[s, i_r, \text{undo}] = \frac{\text{observations}(s')}{\sum_{s \in S} \text{observations}(\text{undo}|s)} \quad (3)$$

- **Ignore** all of Mirny’s recommendations I_r by not clicking on any of them in the interface, remaining in the same state s :


$$P[s, i_r, s] = 1 - P[s, i_r, \text{export}] - P[s, i_r, s'] - P[s, i_r, \text{undo}] \quad (4)$$

Each time the user transitions to a different state, the probability distributions for outgoing transitions from s will be recalculated to account for this additional information. We normalize all corresponding probabilities to ensure that they sum to one. Note that since the total interactions that a user will implement are relatively small, these distributions can be computed easily.

Predicting Transitions Between Interaction States. We model how likely the user is to transition from one state s to another s' as a transition probability between states. This is represented



(a) The Mirny recommendation engine models user interaction state using a Markov decision process (or MDP) model. Here is a simplified example of how the Mirny MDP model is structured with two interactions (zoom and hover), explained in subsection 6.2.



(b) Results of the LOO cross-validation of the recommendation model across per interaction type.

Figure 4: Figure highlighting (a) an illustration of the behavior of the MDP recommendation engine and (b) the results of the evaluation of the recommendation model.

as a *transition function* T , which computes the probability that an interaction i recommended in state $s_i \in S$ leads to state $s_j \in S$, or $T(s_i | i_r, s_j)$. Given the user's current state $s \in S$, a recommended interaction i_r , and a corresponding future state $s' \in S$, we define our transition function probability as follows:

$$T_r(s, i_r, s') = P[s, i_r, s'] \quad (5)$$

Note that an interaction will only be recommended if it has been observed before. However, we derive our states and transition probabilities using our observations from section 4, ensuring that our MDP model covers D3 interactions that users often implement.

When a user does or does not choose one of Mirny's recommendations, the MDP is rewarded or penalized accordingly with a positive or negative *reward* R . We assume that the reward R for a user's reaction to recommendation i_r to be: a small positive number if the user **accepts** i_r , a large positive value if the user **exports**, a large negative value if the user **undoes**, and 0 if the user **ignores** i_r . In this way, the MDP can adapt to match users' behavior.

6.2 MDP Example.

We demonstrate the behavior of the recommendation engine using a simplified example based on our motivating scenario with Sandra from section 3. A simplified MDP model is depicted in Figure 4a. In this example, we assume that there are only two possible interactions that Sandra will want to implement: zoom and hover. With two interactions, there will be four possible states for our model, shown in Figure 4a: no interactions implemented (the initial state), zoom has been implemented, hover has been implemented, and all interactions have been implemented (the terminal state).

Assuming Sandra starts in the initial empty state with her Scatterplot, we can calculate initial transition probabilities, where four (fictitious) transition probabilities are provided in Figure 4a. The intuition behind these probabilities is as follows: Sandra may iterate on her Scatterplot before incorporating any interactions (represented with probability 0.1), and Sandra is much more likely to implement hover than zoom interactions (probabilities 0.6 and 0.1, respectively). There is also a possibility that Sandra could implement a hover interaction but then change her mind and click on the undo button (represented with probability 0.2). For the empty

state, our model provides a ranked list with two items, where the first item is hover, and the second item is zoom.

In our motivating example, Sandra clicks on the first recommendation of hover, representing the **accept** reaction. In this case, the MDP model receives a small positive reward for providing a useful recommendation, represented by the blue line in Figure 4a. However, suppose that Sandra **undoes** her last added interaction, this results in a negative reward to the MDP model, represented by the red line in Figure 4a. If Sandra **exports** her visualization after adding the interaction, the MDP model receives a large positive reward for providing helpful recommendation[s] to Sandra.

6.3 Preliminary Model Accuracy Results

To ensure that our recommendations align with our observations from section 4, we performed leave-one-out cross-validation [22] of the MDP model using our coded dataset. For each visualization example from our dataset, assuming the initial empty state, we classify a recommendation as correct only if they include all interactions implemented within that example. Overall, our model has an average accuracy of 76% (see Figure 4b). For certain visualization types, this accuracy can be improved, e.g., for zooming interactions (68% accuracy). We find that our model performs poorly for brush interactions due to the relatively small number of data points for this interaction type in our training data (see Figure 2). Such cases can be improved by adding more input observations for the model to train on. However, this evaluation does not take into account how users' reactions to recommendations may improve the model's predictions over time. We discuss our evaluation with actual D3 users in section 8.

6.4 Predicting Visualization types

We described how the recommendation engine supports interaction predictions when the user selects a visualization template first. However, our implementation in the Mirny system can also support recommendations even when the user does not select a template. To do this, Mirny utilizes the Beagle Annotator developed by Battle et al. [7] to enable real-time prediction of the visualization a user is

currently implementing. With this annotator, Mirny can predict if a user is implementing a Scatterplot, or some other visualization type by extracting the SVG elements of the visualization and running it through the Beagle Annotator’s decision tree classifiers.

7 AUGMENTING USER CODE

Recommending relevant interactions could help D3 users find and prototype interesting interactions for their visualization designs. However, these recommendations alone would not be very useful to visualization designers. Users would still have to search for example implementations and figure out how to integrate the code into their prototypes. A major goal of this work is *to reduce the challenge of integrating code from different sources when prototyping visualization designs*. We envision a user clicking just a single button to automatically add an interaction to her design specification (addressing design consideration C3).

The D3 API supports six interactions (i.e. *Hover, Click, Brush, Drag, Visualize and Zoom*) which have been identified and analyzed in section 4. The underlying code base for these interactions has been manually extracted into pre-generated templates that Mirny uses to automatically generate code for selected interactions. Once a user selects a recommended interaction, we apply a series of code rewriting steps to merge the template into the user’s current code file. Specifically, we translate the user’s code file into its corresponding abstract syntax trees (or ASTs), analyze the trees for potential insertion points, and merge the interaction code template with the user’s code file at a selected insertion point. This process happens in four steps which we describe below.

Step 1: Identify Input Variables for the Interaction Template.

Interactions often result in changes to the state of the visualization. This may include displaying additional information about the selected data point through a tooltip or making modifications to the state of the selected mark (or non-selected marks) such as changing its position on a drag event. To create meaningful code from interaction templates, specific inputs from the user’s code need to be extracted. However, given the abstract nature of code, it is initially unclear which variables in the user’s code file map to the required inputs within a given interaction code template. Considering our running example from section 3, if the hover interaction needs to update certain properties of selected circle marks in Sandra’s Scatterplot, then the code that sets properties for the circle marks, and the data attributes passed as inputs to the mark properties, need to be incorporated into the hover template.

Hence, the first step of the code augmentation process is to identify analogous variables between the user’s code and the interaction template. Since the properties needed for the interaction templates are known beforehand, we can identify the appropriate data attributes in the user’s code by parsing through its AST.

Using our motivating example, when Sandra selects the *Hover* interaction from the recommendation panel, our templates assume that at the end of the *Hover* event, properties like the color and radius of the circle mark would need to be restored to its original value. The Code Fitter parses through Sandra’s code to identify the input variables that are needed to make a viable template, searching within API calls, variables, and function declarations for the needed

components. For instance, suppose Sandra’s code creates and sets the fill of the circle marks in a scatter plot using the following code:

```
svg.selectAll('.bubble')
  .data(data)
  .enter().append('circle')
  ...
  .attr('r', (d) => radius(d.culmenLength));
  .style('fill', (d) => color(d.species));
```

The code fitter parses through the AST of her code to find where the circle marks’ fill attribute is assigned and extract the assigned value (or in this case function) (see Figure 5a).

Step 2: Populate the Interaction Template. Once all the required input variables have been identified, the interaction template is modified to include these input variables. The generic interaction template contains prop variables that are replaced with the appropriate input variables from the user’s code. For instance, our default Hover template changes all marks to a template-specified color. At the end of the hover interaction, this color will need to be reset. Thus all API calls are updated with the right variable names and parameters as needed. In our running example, the mouseout event of the hover interaction will be populated with the original function to determine the circle’s fill as seen in Figure 5b. If an attribute is not found within the code (see description in step 1), a default value is assigned. However, designers have the discretion to change the default values our templates provide after the code augmentation process is completed.

Generic Template	Populated Template
<pre>.on('mouseout', (d) => { d3.select(this) .style("fill", "prop1"); });</pre>	<pre>.on('mouseout', (d) => { d3.select(this) .style("fill", (d) => color(d.species)); });</pre>

Step 3: Identify Anchor Points in the Users’ code for Insertion.

In D3, the specification of interactions is *imperative* [59], hence, the interaction code cannot be placed just anywhere within an existing D3 specification. Before modifying the user’s code, we must first identify appropriate locations within the code to include new interactions. We refer to these locations as *anchor points*.

Anchor points signify nodes within the AST that need to be changed to support a new interaction. For example, to add the Hover interaction to Sandra’s code, a new node that maps to the interaction code has to be created and added to her code at the point where the visual marks are created. Because we know she is creating a Scatterplot, this would translate to the code block that contains the `.append(circle)` API call. ASTs are traversed using a depth-first approach, hence, identifying the anchor point simply involves iterating over the AST, with the objective of finding the node which has the appropriate data binding API call in its child branch. For our hover example, the anchor point would translate to the first node we identify that has the `.append(circle)` expression as a member of its child nodes as seen in Figure 5c.

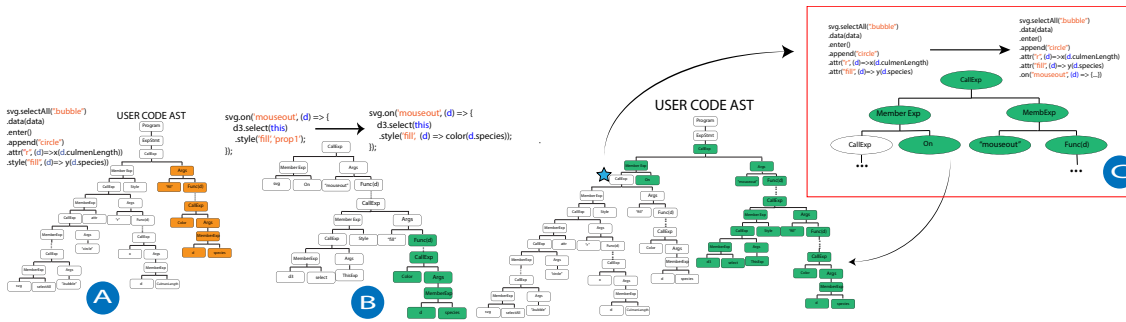


Figure 5: Given a recommended code snippet, e.g., a template for “Hover” interactions in D3, the Mirny system supports *automatic code augmentation*, where the user’s D3 code is mapped to an Abstract Syntax Tree (AST) that Mirny uses to populate and insert the code snippet. In (A), Mirny detects the “fill” variable in the user’s current code (highlighted in orange). In (B), Mirny uses the “fill” variable to automatically populate the “Hover” template (highlighted in green). Finally, in (C), Mirny inserts the populated “Hover” template into the user’s code at a detected anchor point in the AST (marked with a star).

Anchor points for most visualizations generally can be found where the data binding to visual marks occurs. However, for complex interactions such as *Zoom* or *Brush* interactions helper functions and new visual elements need to be added to the user’s code, as such multiple anchor points are needed. Helper functions need to be accessed from anywhere in the D3 script (i.e. have a global scope), hence their anchor point is by default the body of the entire script. Elements, such as rects, to enable zooming or brushing over sections of the visualization need to be added over the entire SVG, as such they use the node that binds the <SVG> element to the web page, as their anchor point.

Step 4: Insert the Populated Template into the User’s Code. Once a template has been populated and the anchor point has been found, the template code is then added as a new node in the user code at the selected anchor point. As a result, the populated template code becomes a new sub-tree within the AST of the user’s code. The template’s root node can either be *appended* or *pre-pended* to and in some cases *replace* the anchor point.

The nature of the template node and the semantic correctness of the code are taken into consideration when determining which of the insertion actions to perform. Code that introduces a new code block after the anchor point typically will use an *append* operation, whereas new code that goes before the anchor point will use a *pre-pend* operation. For instance, adding new code that defines a tooltip for the Hover interaction will be included before we call any hover event. As such, the code for the tooltip will be added before the anchor point using a *pre-pend* operation. The *replace* operation is usually used when adding new code as part of an already existing code sequence in the user’s code. This is due to the fact that ASTs are traditionally traversed depth-first, so the new code will have to become a parent node to the anchor point. In our motivating example for the *Hover* interaction, the code that handles the mouseover and mouseout events will need to be added to the entire anchor point’s branch. This will require a *replace* operation to make the template code’s AST a parent node of the anchor point as seen in Figure 5c. Finally, Mirny converts the new AST back into code and updates the Editor to display the updated code file.

```
svg.selectAll('.bubble')
  .data(data)
```

```
.enter().append('circle')
...
.attr('r', (d) => radius(d.culmenLength));
.style('fill', (d) => color(d.species));
.on('mouseover', (d) => { ... });
.on('mouseout', (d) => { ... });
```

Limitations with Babel. The code augmentation feature helps users add interactions to their live D3 code in a single click. However, the code fitter uses the Babel compiler [49] to transform D3 code into its corresponding AST, as a result, there is a limitation to its operation. The Babel compiler requires that the code to be transformed is both semantically and syntactically correct. Hence, the code augmentation feature will not work if a user selects an interaction on incorrect D3 code. In these cases, Mirny will alert the user about the issues and not complete the code augmentation process.

8 EVALUATION: USER STUDY DESIGN

In previous sections, we introduced two automated features with the aim of supporting template-based visualization prototyping. We implemented these features in a design-focused environment called Mirny. In this section, we focus on evaluating if these automated features lead to measurable improvements in how D3 users prototype interactive visualizations [54, 57]. Our study was preregistered on aspredicted.org⁴.

Participants. We recruited 20 participants from mailing lists and social networking sites (4 female, 15 male, and 1 non-binary) between the ages of 18 to 44. To ensure that all participants had a sufficient understanding of D3 to complete the study tasks under time constraints, we required participants to have at least 3 months of experience working with D3.js. 11 participants had 3 months to 1 year of experience, 3 had 2 - 3 years of experience, 4 had 4 - 5 years of experience and 2 had 5+ years of experience. Participants self-report being proficient in creating visualizations. Each session took approximately 1 hour and 30 minutes and participants were compensated with a \$20 Amazon gift card.

⁴Pre-Registration available at <https://aspredicted.org/t2hb3.pdf>

Tasks. To observe how participants use these automated features in various visualization design contexts, participants were asked to perform two tasks for the study [54, 68]. Participants were allowed to create whatever visualization or interaction type they chose as long as they met the requirements for each task.

- *Task One: Targeted visualization design:* For this task, participants were asked to use the Cherry Blossoms dataset to create a visualization meeting specific design requirements. The output visualization[s] had to meet two requirements: target users of the output visualization should be able to (a) identify correlations between two specific variables, and (b) interact with the data points in the visualization.
- *Task Two: Exploratory visualization design:* Participants were asked to choose between the popular Cars⁵ and Iris [1] datasets, and then brainstorm what type of visualization they would like to create. Then, participants were asked to program this visualization and include at least one interaction in their visualization.

Experiment Conditions. Participants experienced two different Mirny setups across our tasks, resulting in four total experiment conditions ($2 \text{ tasks} \times 2 \text{ interface setups}$). In one setup (baseline), participants only had access to Mirny’s visualization templates, but not Mirny’s interaction recommendations. In the second setup, participants had access to all of Mirny’s features. Condition order was counterbalanced across the visualization tasks. We tuned our experiment conditions through a pilot study, described below.

Pilot. We conducted a pilot study to fine-tune our initial protocol. One important parameter we sought to tune was task difficulty, i.e., should participants be asked to create D3 visualizations from scratch as a task condition (the common case), or should participants be provided with Mirny’s templates? In the pilot study, all participants (2 female, 1 male) were familiarized with Mirny and were asked to perform the tasks described above. For task one, we provided access to Mirny’s visualization templates for two of the three pilot participants, while one participant created visualizations from scratch. Participants were able to use both the predefined templates and recommendations for task 2.

We observed that the participants who were provided with templates only used Mirny’s templates while the third participant copied code from a D3 example they found online. The participant who had to write their D3 code from scratch spent 45 minutes programming a single visualization and was unable to implement any interactions, thus unable to complete the task. The participants who were provided with templates took less than 15 minutes to create a single interactive visualization for the same task. These results show that it would have been difficult to conduct a one-hour experiment if participants had to create visualizations from scratch, and the experiment time could be drastically reduced if templates were always provided. Hence, we decided to provide templates for all tasks and experiment conditions. However, users could always opt to write their own code from scratch if desired.

Protocol. Participants were required to complete a demographic survey to determine eligibility for the study. Each session started with a Pre-study session where participants were familiarized with the study and signed the consent form. They then proceeded to

a training session where the experimenter walked them through the interface functionality via a prerecorded demo of the interface. Participants were then given 5 minutes to explore the full Mirny interface themselves and ask any questions they had about the interface. Once participants were comfortable with the interface, they proceeded to complete our two study tasks. The experimenter limited guidance to only situations where the user reached a complete standstill and was unable to figure out the next step. Participants were allowed to use any external resource they wanted such as Stack Overflow or Blocks.org. No time limits were placed for any of the tasks. Once all the tasks were completed, the participants completed a short exit survey and interview.

Data Collection For each session, an experimenter observed the participants while they performed their tasks and took notes. Logs of participants’ interactions with Mirny were also recorded through the Mirny back-end, which captured all click and keyboard events. We collected data from the post-task survey on participants’ ratings of Mirny’s features to capture participant feedback. Finally, video and audio were recorded to capture participants’ interactions with the interface and responses to the exit interview questions.

9 EVALUATION: ANALYSIS & RESULTS

Here we present three analyses of our study data: measuring session duration and design prototypes[s] (subsection 9.1), classifying template usage and modification (subsection 9.2), and summarizing participant feedback (subsection 9.3).

9.1 Measuring Session Duration and Design Prototypes

We formed two hypotheses to test how Mirny’s code recommendation and augmentation support the visualization prototyping process:

- **H1:** Users produce a higher number of interactive visualizations provided with visualization templates and interaction recommendations, versus only visualization templates.
- **H2:** Users spend less time implementing interactive visualizations when provided with visualization templates and interaction recommendations, compared to when users only have access to visualization templates.

To test our hypotheses, we collected data on the number of visualizations created, the number of interactions participants implemented for each visualization (**H1**), and the time taken to implement a fully interactive visualization, i.e., time taken to render a single visualization + time taken to implement a single interaction (**H2**). We fit linear mixed effects models [6] with the condition and condition order as fixed effects and participant and dataset as random effects.

Based on the results of our pilots, templates were provided for all participants. As a result, we did not statistically test the number of visualizations explored by each participant since all participants used the templates. On average participants created 2 visualizations for each design task. We find a significant effect for interface layout on the number of interactions [**H1**: ($\chi^2(3, N = 20) = 13.406, p = 0.011$)] and time to implement an interaction [**H2**: ($\chi^2(1, N = 20) = 3.85, p = 0.0499$)]. Recommendations + code augmentation led to *1.9 more interactions implemented per visualization* compared to

⁵<https://www.kaggle.com/abineshkumark/carsdata>

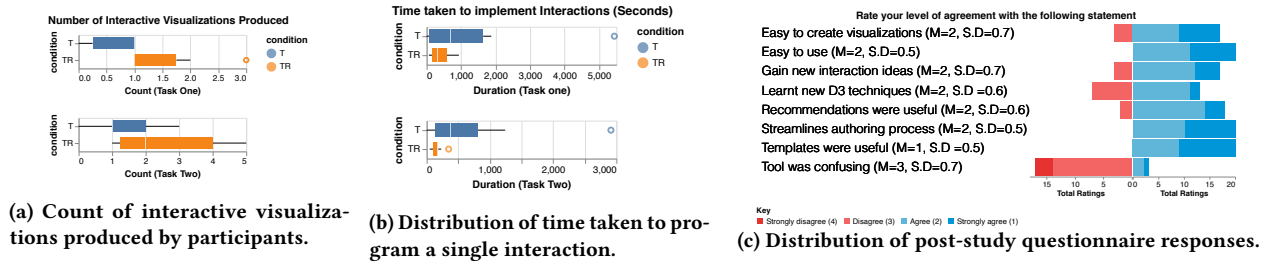


Figure 6: Results from analysis of data collected during evaluation of Mirny. T = Templates only, TR = Templates + Recommendations.

only 0.7 interactions when using only templates. Participants also *implemented interactions 3.2x faster* when recommendations and code augmentation were provided (5 minutes versus 16 minutes for templates only). Our survey results show that participants found the recommendations to be useful (Median(M)=2/Agree, $\sigma = 0.6$) and instrumental in helping them gain new ideas of interactions to implement (Median(M)=2/Agree, $\sigma = 0.6$).

9.2 Classifying Template Usage and Modification

To assess the efficacy of our automated features in helping users program visualizations, we formed two hypotheses:

- **H3:** Users produce a higher number of executable visualizations when provided with visualization templates compared to when users program visualizations from scratch.
- **H4:** Users find D3 more interesting and more useful when provided with both visualization templates and interaction recommendations, versus only visualization templates.

However, based on the results of our pilot study we needed to provide templates for all participants, affecting our analysis methods. While we do not quantitatively test **H3**, we have informal support for **H4** based on our observations of user behavior and feedback from participants.

Participants default to templates to complete tasks. We expected that some participants would want to write their own code and provided an option to support this. Contrary to our expectations, *every participant in the study* defaulted to using Mirny’s templates, regardless of D3 expertise. Most participants settled on Scatterplots as their final design. We suspect that some participants may have been influenced by the word ‘correlation’ in task one which primed participants to choose scatterplots as their final design. Unfortunately, this behavior was unintentional and the fixation on scatterplots could have been mitigated by a more careful examination of the task description. Nevertheless, participants often explored other chart types such as line charts, bar charts, and pie charts before deciding what visualization to use.

Template Modification Strategies. Our findings amplify prior observations that templates are a great starting point for users but require modifications to fit users’ needs. We observed a total of 33 distinct code modifications to Mirny’s visualization templates (n=22) and interaction templates (n=11). The majority of these modifications involved changes to accommodate *stylistic preferences* of participants (n=22), such as changes to the fills for marks and

their associated color schemes, changes to the displayed layout of visual elements such as text, or changes to add new behavior to interactions. For example, P16 modified the code for a hover interaction to remove the default tooltip and instead change the opacity and size of other circles outside of the species class of the selected data point(s). P16 also modified the code to increase the size of all circles in the selected species class to make them more prominent.

We observed 6 occurrences of participants modifying Mirny’s visualization templates to *transform data* such as aggregating, sorting, or grouping data attributes. For instance, P15 modified the bar chart template to aggregate the number of cars produced each year and group them by country of origin. We also observed 5 modifications to *introduce new data encodings* or visual marks to the templates, which resulted in particularly interesting visualization redesigns. For example, P19 introduced an encoding mapping the LastFrost attribute to circle size in their Scatterplot implementation; more examples can be seen in Figure 7.

Unexpected uses of templates. We observed an unexpected use case for templates by some participants while completing the design tasks. A small subset of our participants (n=4) used an initial template to *perform exploratory data analysis*, then chose a different template later on to present the results. For instance, when completing task two, P7 used the bar chart template to analyze the distributions of various data attributes. Then, they used the resulting insights to select attributes to present in their final Scatterplot.

9.3 Participant Feedback

Participants had varied reactions to the automated features provided in the Mirny system. In particular, participants often explained that the templates helped them integrate their design thinking into their programming process because Mirny provided starter code which helped save time finding suitable code. Participants mentioned that on their own, they would have to spend hours finding example code, which is known to be a major design challenge for D3 users [8] and even visualization designers in general [3]. For instance, P11 comments “those [templates] are pretty valuable because there’s always that boilerplate of code you would have to write if you don’t have those templates.” Similarly, P10 said “I think it’s good that it’s smart enough to do this. And normally, doing these interactions is not easy. For example, I wouldn’t know how to implement brushing. I know, it’s a very common thing. But I believe every time I want to do it, I have to search for it online.” Participants mentioned that having access to the code for both the visualizations




Figure 7: Examples of visualization designs created by participants during the experiment. (A), (B), (C), and (E) are visualizations of the Cherry Blossoms dataset depicting different attributes in the dataset such as yearly records of January Snow and February Snow (A), day of peak bloom vs January Temperature (B), etc. (D) and (F) are visualizations of attributes from the cars dataset.

and interactions helped them to improve their understanding of the corresponding D3 designs.

The ability to quickly iterate over different interaction prototypes also allowed them to weigh their design options to understand and identify which designs were suitable for the task. This feature also helped participants to avoid the sunk cost fallacy that would have otherwise made them stick to a specific design only because of the time they had invested in creating their design. P20 states “...instead of having to look for separate things myself, all these recommendations are in one place so it’s very easy to try out which interaction is suitable for which kind of visualization. So I just click and then I can undo and then be like, okay, I probably don’t want this. Whereas if I were to implement them myself, I would want to use something that I made, because I already spent time on it. So I want to use it, even though I probably shouldn’t.”

More experienced participants explained that access to code may not be useful for them due to their existing knowledge of D3, but also agreed that an environment similar to what Mirny provides would be useful for novices who are just starting to learn D3. For example, P9 instructs a visualization class and mentions that teaching D3 is a difficult task as students typically need to learn HTML, CSS, JavaScript, and D3 all at the same time, hindering students’ ability to prototype different visualization designs to understand their perceptual pros and cons. “I also teach people d3. And I think this would be great...Having something like this, that creates the chart and the code and comments the code, so you can read the code. They [students] would like cry tears of joy to have this.”

Participants welcomed the automated features demonstrated in Mirny, however, our approach is not perfect, and participants made suggestions for further refinement. Participants expressed the need for insight on why certain recommendations are made, e.g., P10 said “instead of having the recommendation shift, just give more insight as to why this is being recommended, basically”. Participants also suggested including more widgets to support further customization of D3 code, e.g., P15 said “...maybe adding something like the components, to support selecting an axis or legend. If those are also given as part of templates to completely customize code, that might be more helpful”. Participants also found the live updates feature provided in Mirny to be a hindrance when debugging their code. Participants suggested adding a run button to allow them to update the visualization output when desired.

10 DISCUSSION AND FUTURE WORK

In this work, we present an approach to visualization prototyping that relies on template-based recommendation and augmentation. Our results demonstrate that our approach allows users to generate multiple visualization and interaction designs with significantly less effort and time. Our evaluation highlights strategies that participants engage in to modify templates to fit their design goals. Participants’ feedback also shows that our approach saves users valuable time that they would have otherwise spent seeking design implementations; allowing them to invest their time exploring and evaluating design alternatives. We reflect on the implications of our findings and future directions for research on understanding visualization design, supporting user creativity and autonomy as well as advancing template-driven design.

Understanding visualization design reasoning. Visualization designers often start building their designs from an example or template [3], however, these templates are just a starting point for further design exploration. We find that when provided with visualization templates, participants would often make changes to implement their desired design ideas. We find that the key reasons why participants *augment our templates* are to: (1) accommodate their stylistic preferences, (2) perform data transformations, and (3) introduce new encoding[s] to the visualization. Modifying templates or examples is a common phenomenon amongst designers, yet the visualization community has hardly explored how designers build on and source ideas for their visual designs. In the future, we need to deeply investigate and understand the designer’s reasoning during the design process to understand how the context of use and design tasks influence the changes a user makes to design a template. Furthermore, more research is needed to understand when and how code augmentation fits into the visualization prototyping process. Finally, research also needs to explore the role datasets and stylistic preferences play in the visualization design process.

Balancing automation and user agency in visualization generation. The recent rise in program synthesis tools like GitHub’s Copilot has made it possible to automatically generate code [17]. However, as with all automated systems, users lose control over the code specifications that are generated as these tools often generate buggy or outdated code that a user neither understands nor knows how to fix [2, 18]; a limitation acknowledged in Copilot’s documentation [32]. While our approach relies on the same principles of automated code generation, its code generation process is built on manually curated code templates that incorporate domain-relevant expertise, safeguarding users from some limitations of automatic code generators. However, we acknowledge that the manual nature of Mirny’s template curation limits its scalability to more examples. Future work could explore ensemble strategies that incorporate the strengths of both approaches. For instance, one could integrate Mirny’s domain-specific data on D3 visualizations and its recommendation engine with Copilot to learn from a wider array of visualization and general programming examples to generate recommendations.

Visualization design is still a very subjective process that requires human input and creativity, evidenced by our participants’ desire to modify the generated visualization designs. Current visualization recommenders (e.g., Show Me [41], Voyager [68, 69],

etc.), often treat system-generated designs as final output with little leeway (if any) for modifying the recommended designs, limiting user creativity. In the future, research needs to explore how automated visualization authoring tools can simultaneously preserve user agency and individual design preferences in the designs they generate or recommend [5]. A potential approach is to explore how users can be incorporated into the design generation process providing an opportunity for the user and recommenders to co-design data visualizations.

Modular visualization construction. Inspiration for visualization designs often comes from multiple sources and designers often piece together components from these sources into their designs [3, 74]. Participants like P15 in our study also expressed interest in templates that break visualization elements into modular components to support further customization. This request echoes the sentiment of other D3 users shared online [8]. Future research on the use of design templates should allow for dialog between users, templates, and authoring tools to improve the visualization programming experience. For instance, modularized templates for each part of a visualization (e.g axes, marks) could allow users to mix and match designs.

Supporting complex template transformations. We acknowledge that our choice to focus on only the common D3 visualization types and code integration for interactions is modest. However, this work is only a starting point in supporting dynamic template transformations for visualization design. We observed during our user study that users may be interested in changing a template from one visualization form to another. For instance changing a line chart to a connected scatter plot (see Figure 7b). An important direction for future work is to explore extending our automated features to support more complex template transformations such as changing the structure and arrangement of visual marks in response to data transformations. For instance, how can we dynamically transform a template that implements a bar chart to a stacked bar chart when a user performs data aggregation?

11 CONCLUSION

In this paper, we present a new approach to supporting visualization prototyping when using complex toolkits like D3. We contribute two template-driven automated features to help users program interactive D3 visualizations with less time and effort: (1) a *recommendation engine* for suggesting complementary interactions to add to a D3 visualization, and (2) *automatic code augmentation* to incorporate suggested code snippets to live user code, even when users are *not* following a pre-defined template. We demonstrate these features in Mirny and in a user study with 20 D3 users, we find that they enable users to create interactive visualizations faster and in fewer iterations compared to the typical D3 development workflow (i.e., when automated features are not available). We also observe three key strategies used by D3 users to customize visualization templates and unexpected uses of visualization templates during design. Our work highlights important takeaways on the use of dynamic templates and D3 users' perspectives that we hope will influence future work on visualization prototyping.

ACKNOWLEDGMENTS

We would like to thank members of the UMD Human-Computer Interaction Lab, the Human Data Interaction Group, the UW Interactive Data Lab, and our reviewers for their invaluable feedback. This work was supported in part by NSF award # IIS-1850115 and a VMware Early Career Faculty Grant.

REFERENCES

- [1] E Andersen. 1935. The Irises of the Gasp e Peninsula. *Bulletin of the American Iris Society* 59 (1935), 2–5.
- [2] Owura Asare, Meiyappan Nagappan, and N. Asokan. 2022. Is GitHub's Copilot as Bad As Humans at Introducing Vulnerabilities in Code? <https://doi.org/10.48550/ARXIV.2204.04741>
- [3] Hannah K. Bako, Xinyi Liu, Leilani Battle, and Zhicheng Liu. 2023. Understanding how Designers Find and Use Data Visualization Examples. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2023), 1048–1058. <https://doi.org/10.1109/TVCG.2022.3209490>
- [4] Hannah K. Bako, Alisha Varma, Anuoluwapo Faboro, Mahreen Haider, Favour Nerrise, Bissaka Kenah, and Leilani Battle. 2022. Streamlining Visualization Authoring in D3 Through User-Driven Templates. In *2022 IEEE Visualization and Visual Analytics (VIS)*. 16–20. <https://doi.org/10.1109/VIS54862.2022.00012>
- [5] Calvin S. Bao, Siyao Li, Sarah G Flores, Michael Correll, and Leilani Battle. 2022. Recommendations for Visualization Recommendations: Exploring Preferences and Priorities in Public Health (*Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*). ACM, Article 411, 17 pages. <https://doi.org/10.1145/3491102.3501891>
- [6] Dale J Barr, Roger Levy, Christoph Scheepers, and Harry J Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language* 68, 3 (2013), 255–278.
- [7] Leilani Battle, Peitong Duan, Zachery Miranda, Dana Mukusheva, Remco Chang, and Michael Stonebraker. 2018. Beagle: Automated extraction and interpretation of visualizations from the web. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [8] Leilani Battle, Danni Feng, and Kelli Webber. 2022. Exploring D3 Implementation Challenges on Stack Overflow. In *2022 IEEE Visualization and Visual Analytics (VIS)*. 1–5. <https://doi.org/10.1109/VIS54862.2022.00009>
- [9] Richard Bellman. 1957. A Markovian Decision Process. *Indiana Univ. Math. J.* 6 (1957), 679–684. Issue 4.
- [10] Alex Bigelow, Steven Drucker, Danyel Fisher, and Miriah Meyer. 2014. Reflections on How Designers Design with Data. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces (Como, Italy) (AVI '14)*. Association for Computing Machinery, New York, NY, USA, 17–24. <https://doi.org/10.1145/2598153.2598175>
- [11] Michael Bostock. 2016 (accessed November, 2019). *Popular Blocks*. <http://blocks.org>
- [12] Michael Bostock and Melody Meckfessel. 2016 (accessed July, 2021). *Observable*. <https://observablehq.com/>
- [13] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2301–2309.
- [14] Matthew Brehmer, Robert Kosara, and Carmen Hull. 2022. Generative Design Inspiration for Glyphs with Diatoms. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2022), 389–399. <https://doi.org/10.1109/TVCG.2021.3114792>
- [15] Matthew Brehmer and Tamara Munzner. 2013. A Multi-Level Typology of Abstract Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2376–2385. <https://doi.org/10.1109/TVCG.2013.124>
- [16] Stuart K Card, Jock D Mackinlay, and Ben Shneiderman. 1999. Using vision to think. *Readings in information visualization: using vision to think* 579 (1999).
- [17] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. <https://doi.org/10.48550/ARXIV.2107.03374>
- [18] Arghavan Moradi Dakhel, Vahid Majdinasab, Amin Nikanjam, Foutse Khomh, Michel C. Desmarais, Zhen Ming, and Jiang. 2022. GitHub Copilot AI pair

- programmer: Asset or Liability? <https://doi.org/10.48550/ARXIV.2206.15331>
- [19] Victor Dibia and Çağatay Demiralp. 2019. Data2vis: Automatic generation of data visualizations using sequence-to-sequence recurrent neural networks. *IEEE CG&A* 39, 5 (2019), 33–46.
- [20] E. Dimara and C. Perin. 2020. What is Interaction for Data Visualization? *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 119–129.
- [21] Ian Drosos, Titus Barik, Philip J. Guo, Robert DeLine, and Sumit Gulwani. 2020. *Wrex: A Unified Programming-by-Example Interaction for Synthesizing Readable Code for Data Scientists*. ACM, 1–12.
- [22] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. 2001. *The elements of statistical learning*. Vol. 1. Springer series in statistics New York.
- [23] James Gale, Max Seiden, and Çağatay Demiralp. 2020. Sigma Worksheet: Interactive Construction of OLAP Queries. *arXiv preprint arXiv:2012.00697* (2020).
- [24] Sneha Gathani, Shayan Monadjemi, Alivitta Ottley, and Leilani Battle. 2022. A Grammar-Based Approach for Applying Visualization Taxonomies to Interaction Logs. *Computer Graphics Forum* 41, 3 (2022), 489–500. <https://doi.org/10.1111/cgf.14557> <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14557>
- [25] Kristen B Gorman, Tony D Williams, and William R Fraser. 2014. Ecological sexual dimorphism and environmental variability within a community of Antarctic penguins (genus *Pygoscelis*). *PLoS one* 9, 3 (2014), e90081.
- [26] David Gotz and Zhen Wen. 2009. Behavior-driven visualization recommendation. In *IUI'09*. 315–324.
- [27] Xiaodong Gu, Hongyu Zhang, and Sunghun Kim. 2018. Deep Code Search. In *IEEE/ACM ICSE 2018*. 933–944. <https://doi.org/10.1145/3180155.3180167>
- [28] J. Harper and M. Agrawala. 2018. Converting Basic D3 Charts into Reusable Style Templates. *IEEE Transactions on Visualization and Computer Graphics* 24, 3 (2018), 1274–1286. <https://doi.org/10.1109/TVCG.2017.2659744>
- [29] E. Hoque and M. Agrawala. 2020. Searching the Visual Style and Structure of D3 Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 1236–1245. <https://doi.org/10.1109/TVCG.2019.2934431>
- [30] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the 1999 CHI Conference on Human Factors in Computing Systems*. 159–166.
- [31] Kevin Hu, Michiel A Bakker, Stephen Li, Tim Kraska, and César Hidalgo. 2019. Vizml: A machine learning approach to visualization recommendation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [32] GitHub Inc. [n. d.]. About github copilot. <https://docs.github.com/en/copilot/overview-of-github-copilot/about-github-copilot>
- [33] Plotly Technologies Inc. 2015. *Collaborative data science*. Montreal, QC. <https://plot.ly>
- [34] Yalin Ke, Kathryn T. Stolee, Claire Le Goues, and Yuriy Brun. 2015. Repairing Programs with Semantic Code Search (T). In *IEEE/ACM ASE 2015*. 295–306. <https://doi.org/10.1109/ASE.2015.60>
- [35] Mary Beth Kery, Donghao Ren, Fred Hohman, Dominik Moritz, Kanit Wongsuphasawat, and Kayur Patel. 2020. *image: Fluid Moves Between Code and Graphical Work in Computational Notebooks*. *UIST'20* (Oct 2020). <https://doi.org/10.1145/3379337.3415842>
- [36] Miryung Kim, Lawrence Bergman, Tessa Lau, and David Notkin. 2004. An ethnographic study of copy and paste programming practices in OOPL. In *ISESE'04*. IEEE, 83–92.
- [37] Yun Lin, Guozhu Meng, Yinxing Xue, Zhenchang Xing, Jun Sun, Xin Peng, Yang Liu, Wenyun Zhao, and Jinsong Dong. 2017. Mining implicit design templates for actionable code reuse. In *IEEE/ACM ASE 2017*. IEEE, 394–404.
- [38] Zhicheng Liu, John Thompson, Alan Wilson, Mira Dontcheva, James Delorey, Sam Grigg, Bernard Kerr, and John Stasko. 2018. Data Illustrator: Augmenting Vector Design Tools with Lazy Data Binding for Expressive Visualization Authoring (*Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*). ACM, 1–13. <https://doi.org/10.1145/3173574.3173697>
- [39] Y. Luo, X. Qin, N. Tang, and G. Li. 2018. DeepEye: Towards Automatic Data Visualization. In *2018 IEEE ICDE*. 101–112.
- [40] Ruixian Ma, Honghui Mei, Huihua Guan, Wei Huang, Fan Zhang, Chengye Xin, Wenzhuo Dai, Xiao Wen, and Wei Chen. 2020. Ladv: Deep learning assisted authoring of dashboard visualizations from images and sketches. *IEEE Transactions on Visualization and Computer Graphics* 27, 9 (2020), 3717–3732.
- [41] J. Mackinlay, P. Hanrahan, and C. Stolte. 2007. Show Me: Automatic Presentation for Visual Analysis. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1137–1144.
- [42] Michele Mauri, Tommaso Elli, Giorgio Caviglia, Giorgio Uboldi, and Matteo Azzi. 2017. RAWGraphs: A Visualisation Platform to Create Open Outputs (*CHIItaly '17*). ACM, Article 28, 5 pages. <https://doi.org/10.1145/3125571.3125585>
- [43] Sean McKenna, Dominika Mazur, James Agutter, and Miriah Meyer. 2014. Design activity framework for visualization design. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2191–2200.
- [44] Collin McMillan, Mark Grechanik, Denys Poshyvanyk, Qing Xie, and Chen Fu. 2011. Portfolio: Finding Relevant Functions and Their Usage (*ICSE '11*). ACM, 111–120. <https://doi.org/10.1145/1985793.1985809>
- [45] Andrew M McNutt and Ravi Chugh. 2021. *Integrated Visualization Editing via Parameterized Declarative Templates*. ACM.
- [46] Min Lu, Jie Liang, Yu Zhang, Guozheng Li, Siming Chen, Zongru Li, and X. Yuan. 2017. Interaction+: Interaction enhancement for web-based visualizations. In *IEEE PacificVis 2017*. 61–70.
- [47] Dominik Moritz, Chenglong Wang, Greg L Nelson, Halden Lin, Adam M Smith, Bill Howe, and Jeffrey Heer. 2018. Formalizing visualization design knowledge as constraints: Actionable and extensible models in draco. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 438–448.
- [48] Anh Tuan Nguyen, Tung Thanh Nguyen, Hoan Anh Nguyen, Ahmed Tamrawi, Hung Viet Nguyen, Jafar Al-Kofahi, and Tien N. Nguyen. 2012. Graph-based pattern-oriented, context-sensitive source code completion. In *ICSE'12*. 69–79. <https://doi.org/10.1109/ICSE.2012.6227205>
- [49] Open Source. [n. d.]. *Babel*. <https://babel.js.io/docs/en/>
- [50] Luca Ponzanelli, Gabriele Bavota, Massimiliano Di Penta, Rocco Oliveto, and Michele Lanza. 2014. Mining StackOverflow to Turn the IDE into a Self-Confident Programming Prompter (*MSR 2014*). ACM, 102–111. <https://doi.org/10.1145/2597073.2597077>
- [51] Tom Preston-Werner, Chris Wanstrath, P.J. Hyett, and Scott Chacon. 2008 (accessed July, 2019). *GitHub*. <https://github.com/about>
- [52] Veselin Raychev, Martin Vechev, and Eran Yahav. 2014. Code Completion with Statistical Language Models. *SIGPLAN Not.* 49, 6 (jun 2014), 419–428. <https://doi.org/10.1145/2666356.2594321>
- [53] Steven P. Reiss. 2009. Semantics-based code search. In *ICSE'09*. 243–253. <https://doi.org/10.1109/ICSE.2009.5070525>
- [54] Donghao Ren, Bongshin Lee, Matthew Brehmer, and Nathalie Henry Riche. 2018. Reflecting on the evaluation of visualization authoring systems: Position paper. In *IEEE BELIV 2018*. IEEE, 86–92.
- [55] Bahador Saket, Hannah Kim, Eli T. Brown, and Alex Endert. 2017. Visualization by Demonstration: An Interaction Paradigm for Visual Data Exploration. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 331–340. <https://doi.org/10.1109/TVCG.2016.2598839>
- [56] Arvind Satyanarayan and Jeffrey Heer. 2014. Lyra: An Interactive Visualization Design Environment. *CGF* 33 (06 2014). <https://doi.org/10.1111/cgf.12391>
- [57] Arvind Satyanarayan, Bongshin Lee, Donghao Ren, Jeffrey Heer, John Stasko, John Thompson, Matthew Brehmer, and Zhicheng Liu. 2019. Critical reflections on visualization authoring systems. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 461–471.
- [58] Arvind Satyanarayan, Dominik Moritz, Kanit Wongsuphasawat, and Jeffrey Heer. 2016. Vega-lite: A grammar of interactive graphics. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2016), 341–350.
- [59] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. 2017. Vega-Lite: A Grammar of Interactive Graphics. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 341–350.
- [60] A. Satyanarayan, R. Russell, J. Hoffswell, and J. Heer. 2016. Reactive Vega: A Streaming Dataflow Architecture for Declarative Interactive Visualization. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 659–668.
- [61] Guy Shani, Ronen I. Brafman, and David Heckerman. 2002. An MDP-Based Recommender System. In *UAI'02*. Morgan Kaufmann Publishers Inc., 453–460.
- [62] Chris Stolte, Diane Tang, and Pat Hanrahan. 2002. Polaris: A system for query, analysis, and visualization of multidimensional relational databases. *IEEE Transactions on Visualization and Computer Graphics* 8, 1 (2002), 52–65.
- [63] Theophanis Tsandilas. 2020. StructGraphics: Flexible visualization design through data-agnostic and reusable graphical structures. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 315–325.
- [64] Manasi Vartak, Silu Huang, Tarique Siddiqui, Samuel Madden, and Aditya Parameswaran. 2017. Towards visualization recommendation systems. *ACM SIGMOD Record* 45, 4 (2017), 34–39.
- [65] Romain Vuillemin and Jeremy Boy. 2017. Structuring visualization mock-ups at the graphical level by dividing the display space. *IEEE transactions on visualization and computer graphics* 24, 1 (2017), 424–434.
- [66] Jagoda Walny, Christian Frisson, Mieka West, Doris Kosminsky, Søren Knudsen, Sheelagh Carpendale, and Wesley Willett. 2019. Data Changes Everything: Challenges and Opportunities in Data Visualization Design Handoff. *IEEE transactions on visualization and computer graphics* (2019).
- [67] Chenglong Wang, Yu Feng, Rastislav Bodik, Isil Dillig, Alvin Cheung, and Amy J Ko. 2021. Falx: Synthesis-Powered Visualization Authoring. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [68] K. Wongsuphasawat, D. Moritz, A. Anand, J. Mackinlay, B. Howe, and J. Heer. 2016. Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 649–658.
- [69] Kanit Wongsuphasawat, Zening Qu, Dominik Moritz, Riley Chang, Felix Ouk, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2017. Voyager 2: Augmenting visual analysis with partial view specifications. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 2648–2659.
- [70] Tiffany Wun, Lora Oehlberg, Miriam Sturdee, and Sheelagh Carpendale. 2019. You say Potato, I say Po-Data: Physical Template Tools for Authoring Visualizations. In *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction*. 297–306.

- [71] Cong Yan and Yeye He. 2020. Auto-Suggest: Learning-to-Recommend Data Preparation Steps Using Data Science Notebooks (*SIGMOD '20*). ACM, 1539–1554. <https://doi.org/10.1145/3318464.3389738>
- [72] J. S. Yi, Y. a. Kang, J. Stasko, and J. A. Jacko. 2007. Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1224–1231.
- [73] Zehua Zeng, Phoebe Moh, Fan Du, Jane Hoffswell, Tak Yeon Lee, Sana Malik, Eunye Koh, and Leilani Battle. 2022. An Evaluation-Focused Framework for Visualization Recommendation Algorithms. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2022), 346–356. <https://doi.org/10.1109/TVCG.2021.3114814>
- [74] Jiayi Eris Zhang, Nicole Sultanum, Anastasia Bezerianos, and Fanny Chevalier. 2020. DataQuilt: Extracting Visual Elements from Images to Craft Pictorial Visualizations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376172>
- [75] Jonathan Zong, Dhiraj Barnwal, Rupayan Neogy, and Arvind Satyanarayan. 2020. Lyra 2: Designing interactive visualizations by demonstration. *IEEE Transactions on Visualization and Computer Graphics* (2020).