ColorCook: Augmenting Color Design for Dashboarding with Domain-Associated Palettes

YANG SHI, Intelligent Big Data Visualization Lab, Tongji University, China
SIJI CHEN, Intelligent Big Data Visualization Lab, Tongji University, China
PEI LIU, Intelligent Big Data Visualization Lab, Tongji University, China
JIANG LONG, Huawei Technologies Co. Ltd., China
NAN CAO, Intelligent Big Data Visualization Lab, Tongji University, China

Visualization dashboards serve as an information presentation that uses a tiled layout of key metrics visualized in charts for collaborative decision-making. Existing work has developed tools and techniques for computational color design. Much of these efforts have focused on selecting effective color palettes for independent charts while few attempts have been made to support the expressive color design of multiple coordinated charts in dashboards. In this work, we describe ColorCook, an interactive system that helps design expressive and effective dashboard colorings using domain-associated palettes. ColorCook employs an integrated color workflow for dashboarding, consisting of color selection, assignment, and adjustment. We evaluated ColorCook through a crowdsourcing experiment and a user study. The results of our evaluation indicated that ColorCook is useful for effective and expressive color design.

CCS Concepts: • **Human-centered computing → Visualization systems and tools**; • **Information systems → Crowdsourcing**.

Additional Key Words and Phrases: Color design, Visualization dashboards, Crowdsourcing, Design support tools

ACM Reference Format:
Yang Shi, Siji Chen, Pei Liu, Jiang Long, and Nan Cao. 2022. ColorCook: Augmenting Color Design for Dashboarding with Domain-Associated Palettes. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 433 (November 2022), 25 pages. https://doi.org/10.1145/3555534

1 INTRODUCTION

Visualization dashboards, as an information presentation, use a tiled layout of charts to support collaborative decision-making and condition monitoring [48, 57, 62]. They are ubiquitous in almost every industry, transforming data lakes and warehouses into an accessible overview of an organization – its current state and goals – at a glance [16]. In dashboard design, *color* receives priority among various visual elements; it serves as both an important visual channel to encode data [13] and a pivotal feature to communicate the characteristics of information design [68]. However, creating dashboard colorings is challenging, especially for general users, as a well-designed coloring should be both effective and expressive. It is often achieved by a multi-step color workflow, including how to select a palette that best represents the idea of the content being colorized, how to assign the selected colors to the appropriate segments of charts, and how to arrive at a satisfying coloring.

Authors' addresses: Yang Shi, Intelligent Big Data Visualization Lab, Tongji University, China; Siji Chen, Intelligent Big Data Visualization Lab, Tongji University, China; Pei Liu, Intelligent Big Data Visualization Lab, Tongji University, China; Jiang Long, Huawei Technologies Co. Ltd., China; Nan Cao, Intelligent Big Data Visualization Lab, Tongji University, China.

2573-0142/2022/11-ART433 $15.00
https://doi.org/10.1145/3555534

One potential solution to avoid being overwhelmed by hundreds of design decisions is leveraging computational color design to automate the color workflow. For example, Lu et al. [28] have proposed to integrate palette generation and color assignment into a single procedure for independent charts. LADV [31] has been proposed to generate dashboards by imitating the visual style such as layout and color of example images. Although their approaches facilitate color exploration and presentation, colorizing multiple coordinated charts in a dashboard has been largely overlooked, where the constraints on color consistency such as "*same field, same value-color mapping*" are expected to be maintained [44]. In addition, much of these efforts select palettes based on the effectiveness of color encoding without considering its expressiveness, that is, how to select palettes that can best communicate data content within use-context.

Motivated by the aforementioned challenges, we present *ColorCook*, a design support tool that facilitates expressive and effective color design for dashboarding using both model- and data-driven approaches. To design ColorCook, we formulated a general color workflow of color design, including color selection, assignment, and adjustment. Specifically, we collected a dataset of 115 domain-associated palettes, which serves as a palette library in ColorCook. The library allows users to specify the domain of dashboards and augment its expressiveness through color selection. As a user selects a palette from the library, ColorCook can assign the selected colors to coordinated charts that the user has created for an intended dashboard. During the process of automatic color assignment, ColorCook ensures the effectiveness of color encoding and follows the established rules in visual perception. The system also allows the user to adjust the resulting dashboard coloring by suggesting colors based on a palette rating model [19] trained on a human color preference dataset. We evaluated ColorCook on its ability to support color design via a crowdsourcing experiment and a user study. The results of our evaluation showed that ColorCook has a good performance in terms of usefulness, ease of use, and quality of coloring. By analyzing both the quantitative and qualitative results of the evaluation, we discuss the design implications on developing design support tools geared towards expressive information design.

The main contributions of this work are as follows:

- We designed ColorCook to support the expressive and effective color design of multiple coordinated charts within a dashboard. ColorCook ensures color consistency, harmony, and discriminability, which can also be applied to collaborative dashboard coloring.
- We collected a manually curated dataset of 115 domain-associated palettes for the first time. The dataset was validated for its specificity via a crowdsourcing experiment and is available at https://idvxlab.com/colorcook/.
- We conducted an evaluation consisting of a crowdsourcing experiment and a user study. The results of the evaluation indicated that ColorCook can support color design and increase the likelihood of creating high-quality dashboard colorings.

2 RELATED WORK

To motivate our research, we review the literature on design of multiple views and dashboards, color palettes and color datasets, and color assignment in data visualization.

2.1 Design of Multiple Views and Dashboards

Visualizations are frequently designed and presented as multiple views. Prior research has studied various design aspects such as color, layout, and interaction of multiple views. For example, Qu et al. [44] suggested a set of consistency constraints in color and xy encodings such as "*different fields, non-overlapping palettes*". They also conducted a wizard-of-oz study to investigate how visualization authors perceive and achieve such encoding consistency across multiple views.

Chen et al. [10] explored the composition and configuration of visualization in multiple views by extracting images of visualization tools in visualization publications. Langner et al. [22] proposed a set of design considerations (visual and interaction consistency, use of visual variables, and multi-user requirements) and interaction techniques for multiple views on wall-sized displays.

Dashboards place a set of charts in multiple views to facilitate drawing insights from large datasets [48, 59, 62]. To learn the design space of dashboards, Sarikaya et al. [48] analyzed 83 examples and documentation of dashboards to understand their functional design, purpose, audience, visual features, and data semantics. Vázquez-Ingelmo et al. [59] analyzed 23 papers that address the tailoring capabilities of dashboards and classified them by tailored solutions, including customization, personalization, and adaptation. Prior research has also investigated various methods to support constructing dashboards [12, 65]. For example, Yalçın et al. [65] introduced Keshif that enables drag-and-drop authoring of dashboards from raw data. Its automated model supports the streamlined exploration of tabular data. More recently, LADV [31] uses a deep learning-based model to enable rapid prototyping for dashboards. Specifically, it generates dashboards by imitating the chart types, layout, and color palette of the input image. While the aforementioned work has investigated automatic coloring, few work takes into account coordinated charts in dashboards. Specifically, our work ensures color effectiveness for multiple coordinated charts by following color constraints such as color consistency proposed by the aforementioned work. Another related constraint on color encoding is that charts involved in a dashboard often show different data types which require different data encodings; discrete palettes encode category values while continuous color palettes convey ordinal or quantitative differences [42]. Thus, our color assignment algorithm starts from a palette and expands it for multiple charts with different data types such as quantitative and qualitative data to maintain color consistency.

2.2 Color Palettes and Color Datasets

As a combination of colors, a color theme or a color palette can influence the quality or characteristics of information design. Prior work has proposed various methods to extract [24], generate [34, 63], and suggest et al. [41, 52] color palettes. For example, Lin and Hanrahan [24] proposed to extract color themes from images using a regression model trained on human-extracted theme data. Color Crafter [54] can generate continuous palettes by analyzing patterns in the structure of designer-crafted palettes. Phan et al. [41] used Gaussian process latent variable models to capture color styles embedded in fine art collections and suggest color palettes of different artistic styles such as Van Gogh or Raphael.

While the recent advances of deep learning and crowdsourcing, a number of large-scale color datasets have been constructed to support data-driven design applications. Munroe [35] introduced a color dataset containing 954 RBG colors, while each color is described using a color name based on crowdsourcing user judgment. Kobayashi [20] established a multi-color dataset, which includes 1,170 three-color palettes categorized by their characteristics such as cute and old-fashioned. O'Donovan et al. [37] curated a color dataset of 10,743 five-color palettes, while each palette was rated by MTurk workers. As a subsequent work, Kita and Miyata [19] proposed a palette rating model that considers human aesthetic preferences. Bahng et al. [6] collected a color dataset of 10,183 five color palettes and each palette is associated with its corresponding text description such as neon and skin tones. The aforementioned color datasets characterize palettes by direct color names such as maroon or by imagery that evokes a particular set of colors such as autumn. To the best of our knowledge, no assets or datasets of domain-associated palettes are currently available, especially related to the use-context of dashboards. We curated such a dataset by extracting palettes from context-specific illustrations and validated for their specificity via a crowdsourcing experiment.

2.3 Color Assignment in Data Visualization

Color is one of the most critical visual channels in data visualization. Prior work has investigated design factors that impact color perception such as mark size [55], mark shape [53], spatial frequency [45], and background color [49]. Also, predefined color pickers for visualization design such as PRAVDAColor [8] and ColorBrewer [14] have been developed based on design experience.

Colormaps or color palettes are commonly featured in visualization for mapping data values. Approaches to extracting [67], generating [69], optimizing [13], and evaluating [9, 25, 46] colormaps have been proposed. Color assignment in visualization has focused on exploring the effectiveness of color encoding [42], perceptual discriminability [28, 56, 60], and semantic consistency [15, 50]. For example, Poco et al. [42] extracted either a discrete or continuous color legend from a bitmap visualization image and then automatically recolored this image to ensure the effectiveness of color encoding. Lu et al. [28] applied an integrated approach for the creation and assignment of colors to categorical visualizations that maximize the perceptual discriminability between classes. Setlur et al. [50] used a linguistic approach to generating a semantically resonant color palette for colorable terms in a given dataset. While the aforementioned work has shown to be helpful in designing effective visualization colorings, they have mainly focused on parts of the color workflow for independent charts. The goal of our work is to support the whole color workflow for multiple coordinated charts in a dashboard, including color selection, assignment, and adjustment. Such an integrated workflow can support experimenting with different color palettes for rapid prototyping.

3 PRELIMINARY STUDY

To understand users' current practice of designing dashboard colorings, we conducted a preliminary study with five participants, including one visualization researcher (R1), two data analysts (A1, A2), and two information designers (D1, D2). All the participants have more than five years of professional experience. We first asked each participant to create a dashboard with a focus on expressive color design. The participants were allowed to select their preferred dashboard tools and datasets of interest to accomplish the task. After the dashboarding session, we conducted a series of interviews with the participants. In each interview, we asked them the following questions: (1) what color workflow they adopted for creating a dashboard coloring, (2) what parts of the workflow were easy or difficult to manage with existing tools, and (3) what guidelines or techniques they used to better support the workflow. The dashboarding session lasted for about 0.5 hours while the interview session took approximately 0.5 hours for each participant. Each interview was audio-recorded for subsequent analysis.

3.1 Findings and Design Requirements

To analyze the qualitative data collected from the interviews, we transcribed the audio recordings and then organized the data based on the color workflow, including color selection, assignment, and adjustment.

Color Selection. We found that selecting a palette often serves as the first step in the color workflow. Interestingly, the two information designers also attempted to define the "personality" or "characteristics" of a dashboard using color in this step. D1 suggested, "*color plays an important role in identifying the personality of a dashboard, and an attractive personality makes a dashboard an expressive one.*" D2 noted, "*when designing a dashboard to present education data, I'd use bright colors to make it look creative and playful. For a business dashboard that should look elegant and formal, I'd use dark blue or green... Unfortunately, I didn't find any dashboard tools providing palettes associated with application domains.*" Existing dashboard tools such as Tableau [47] provide predefined palettes categorized into discrete, sequential, and diverging palettes to help encode qualitative or quantitative

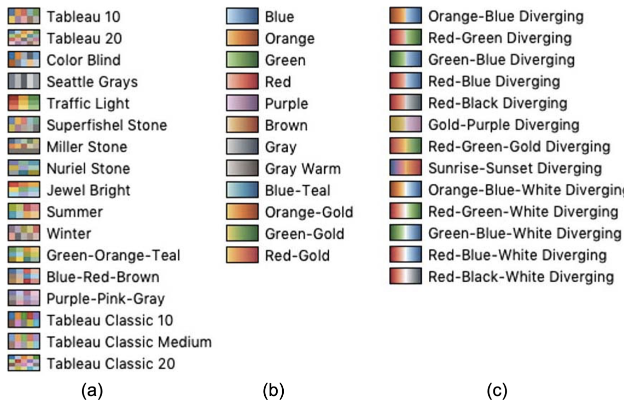

Fig. 1. Color palettes provided by Tableau [47]: (a) *discrete palettes* can be used to encode qualitative values, each representing a distinct category. (b) *sequential palettes* can be used to encode quantitative values while (c) *diverging palettes* can be used to encode quantitative values with a mid-point such as zero. Sequential and diverging palettes constitute continuous palettes.

data values (Fig. 1). While such palette libraries can provide guidance for effective color encodings, the association between a palette and the domains of a dashboard is largely ignored.

Color Assignment. After selecting a palette, the participants proceeded to color assignment, which is often automatically completed by existing dashboard tools such as Tableau [47] and PowerBI [33]. When asked about the guidelines or rules that automated colorization should follow, all the participants emphasized the importance of color harmony or compatibility, "*it affects the overall look and leaves an instant impression.*" Color discriminability or contrast is another factor that affects visual perception. D2 noted, "*although these two different shades of blue look harmonious when placed next to each other, this color combination didn't help differentiate one data attribute from the other.*" We also found that the participants were not satisfied with the color consistency in the resulting dashboard colorings. R1 explained, "*this tool uses a simple colorization strategy: the first series in the bar chart adopts the first color in the palette; the second series, the second color; and so on... In cases where a specific data attribute was colored differently in a succeeding chart in a dashboard, the audience is likely to misidentify this attribute.*" Overall, our study identified a requirement for automating the color assignment process taking into account the data constraints of the coordinated charts as well as coloring constraints of the data types.

Color Adjustment. As automated color assignment can generate unsatisfactory results, the participants usually performed color adjustment as a follow-up step. The participants noted that existing tools lack the capability of suggesting compatible colors based on the selected palettes. A2 said, "*when I selected an eight-color palette and used more than eight data attributes to draw a chart, it'd re-use the colors in the palette instead of suggesting new colors for extra attributes.*" As alternatives, the two designers used online color pickers such as Adobe Color [2] to explore monochromatic or complementary colors, claiming "*it's an effective way to extend the palette from eight colors to ten colors or replace certain colors in the palette with new ones.*" Another strategy observed was experimenting with different palettes until a satisfactory outcome was obtained. R1 stated, "*if I find two or more colors in the result are unsatisfying, I'd change to another palette although it means starting all over again.*" However, both strategies may impair the current color workflow and waste a disproportionate amount of attention.

Based on the findings regarding color selection, assignment, and adjustment derived from the preliminary study, we respectively established three design requirements to inform the development of ColorCook.

DR1 Providing context-specific color palettes that help depict the domains of a dashboard. Our system should present a library of such domain-associated palettes for users to expressively communicate the data content of an intended dashboard through color selection.

DR2 Generating dashboard colorings that ensure the effectiveness of color encoding, that is, color assignment should extend the selected palette to discrete, sequential, or diverging palettes based on data types and ensure color consistency across multiple charts. Also, color assignment should follow established rules in visual perception regarding color harmony and discriminability.

DR3 Allowing users to adjust one or more of the colors in the dashboard colorings by suggesting harmonious colors. Thus, users can extend the current palette to include more colors or replace certain colors in the palette with new ones.

4 DOMAIN-ASSOCIATED PALETTE DATASET

According to the preliminary study, ColorCook first requires a dataset of domain-associated palettes as a palette library for color selection (**DR1**). To create such a dataset, we first identified common domains of dashboards by investigating existing dashboard tools. We then collected graphic designs (illustrations) that depict scenarios of the identified domains and extracted corresponding color palettes from these designs. We also filtered and validated the dataset from three aspects, including distinctiveness, aesthetics, and specificity. The final dataset contains 115 manually curated palettes and is released to support the design of future data-driven color design tools. The dataset is accessible via https://idvxlab.com/colorcook/.

4.1 Domain Identification

To identify common domains of dashboards, we began by surveying existing dashboard tools. We used search keywords such as "dashboard tool/software" and "visualization dashboard" on Google and processed the returned results of good relevance. We also added a selection of dashboard tools recommended by opinionated lists such as "the best dashboard tools of 2020" and by the participants from the preliminary study. In doing so, 52 popular dashboard tools were selected.

Then, we checked the official website of each tool and explored its dashboard templates which are often categorized by "industry" or "application domain". By analyzing these categorizations used by the 52 tools, we identified 65 distinct categories of domains. We then merged related categories and organized them based on their industry groups proposed by the North American Industry Classification System (NAICS) [36]. Finally, 42 categories of domains were identified. We sorted these categories by frequency and used the top 11 categories (*finance, marketing, technology, healthcare, project management, retail, manufacturing, sports, social media, education, business*) that have been used by more than 20% of the 52 tools.

4.2 Illustration Collection

Based on the 11 identified domain categories, we attempted to find color palettes associated with these domains. We initially thought that the dashboard templates provided by the 52 dashboard tools and visualization datasets [67] could be the potential sources of domain-associated palettes. However, by examining the palettes extracted from these visualizations, we found that their quality varies and their difference in color between domains was not apparent. For example, Qlik [1] uses the same color palette containing blue and red for both *marketing* and *finance*. Inspired by previous work [32] that uses typeface as a visual characteristic of a city, we thought of using context-specific graphic designs (illustrations) as alternative sources. Illustrations, as a visual explanation of a specific concept or scenario, possess a variety of color combinations carefully manipulated by professional designers.

To collect illustrations that depict scenarios of the 11 domains, we first used the names of these domains as search keywords on well-known high-quality asset libraries such as Ouch [27] and Isoflat [17] and obtained 5040 illustrations in total. Subsequently, two researchers with design-related backgrounds independently analyzed the illustrations to check their aesthetics and relevance to specific domains. In terms of aesthetics, we investigated the visual elements of the illustrations such as color and scale, following the guidelines of graphic design [30]. To ensure that an illustration is directly relevant to a specific domain, we decided on a set of visual objects that each domain should include after analyzing the illustrations. For example, illustrations of *business* may depict businessmen having conversations or making presentations (Fig.2 (h)) whereas illustrations of *education* may display students reading books or throwing mortarboards in the air (Fig.2 (g)). However, *marketing*, *project management*, and *retail* were found to be relatively abstract concepts and we failed to associate concrete objects to these concepts in their illustrations. Thus, these three domains were eliminated. These visual objects were used as codes to select domain-associated illustrations. During the coding process, we met for three sessions to refine the codes and resolve disagreements. Finally, we independently coded the illustrations using the final coding scheme (Table 1) and reached Cohen's Kappa of 0.86. We then discussed the mismatches and reached a 100% consensus. It is worth noting that one illustration may be categorized into multiple domains based on its visual content. In accordance with the two inclusion criteria regarding aesthetics and relevance, we identified 548 out of 5040 illustrations. We then extracted the corresponding color palettes from these illustrations, one per illustration.

Domain	Visual Objects (Codes)
Finance	coin, dollar sign, banknote, credit card, stock exchange
Technology	robot, keyboard, projector, scanner, VR glasses, outer space
Healthcare	stethoscope, doctor, medical box, heart
Manufacturing	robotic arm, conveyor belt, goods, industrial worker, safety helmet
Sports	athlete, ball games, sports equipment
Social Media	social media apps, the icons of like/comment/bookmark
Education	book, children, characters, stationery
Business	businessman, dashboard, office, meeting room

Table 1. The coding scheme for each domain category accompanied by a list of representative visual objects.

4.3 Palette Filtration and Validation

We filtered and validated the domain-associated palette dataset on its distinctiveness, aesthetics, and specificity. In accordance with the three inclusion criteria, we identified 115 out of 548 palettes to be included in our final dataset.

Distinctiveness. We mandated that each of the extracted palettes should present six or more distinct colors to constitute a discrete palette. Sequential and diverging palettes were excluded as they can be automatically generated from the discrete palettes using our algorithms (Section 5.3) to ensure color harmony. To evaluate the distinctiveness of the palettes, we conducted gradation analysis [19] by examining if a palette displays linearity in the order of colors. Specifically, we computed the gradation of each palette in CIELAB color space to determine if the order of colors decreases or increases monotonically. The palettes with average gradation below 0.05 were considered to have a higher probability of presenting continuous colors and were thus removed from our palette dataset.

Aesthetics. We also evaluated the aesthetics of the palettes using a palette rating model [19]. The model learns human aesthetic preferences from a large-scale dataset of human-rated palettes [38] and is able to rate a palette with any number of colors. The model rates the palettes from one to five and we removed palettes whose ratings were below two.

Specificity. To ensure that each palette in our dataset is perceived as relevant to a certain domain, we conducted a crowdsourcing experiment by recruiting 151 participants aged between 18 and 67 (M = 26.44, SD = 63.58) from Prolific [43]. Each participant was paid $5.00 for rating 20 palettes and providing explanations. The experiment consisted of three sessions. Session 1 introduced our research intent and the structure of the experiment. In Session 2, the participants were randomly presented with a group of 20 palettes and instructed to rate the association between a palette and each of the eight domains using a 5-point Likert scale. Here, 1 denoted "not at all relevant" while 5 denoted "extremely relevant". After rating each palette, the participants were also asked to explain their options. Ratings with no reason provided or reasons that failed to explain options were regarded as unqualified and thus excluded. We iterated the process to ensure that each palette had received 10 qualified answers. Session 3 required the participants to fill out a demographic survey. The three sessions lasted about 0.5-1 hour for each participant.

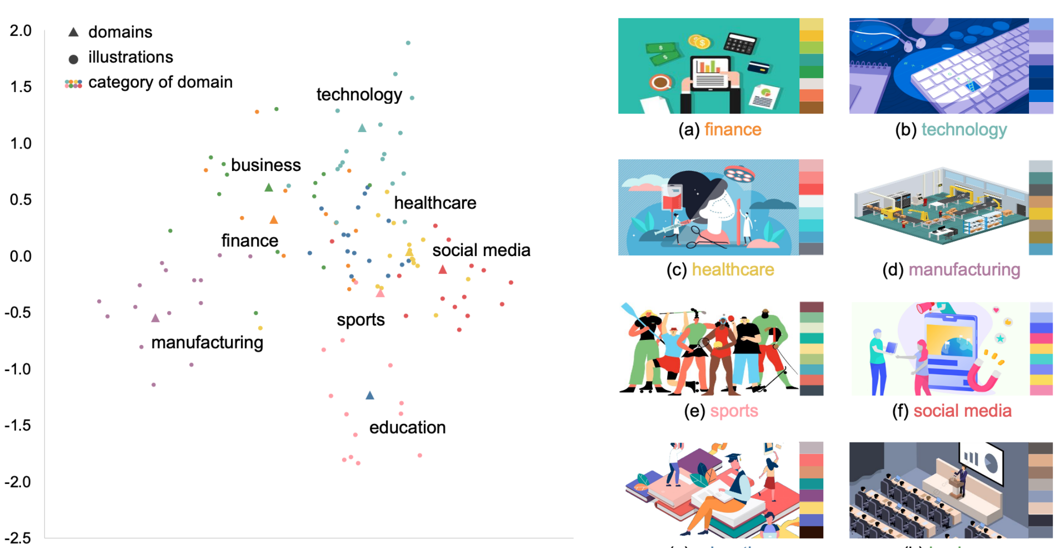

Fig. 2. The dataset is visualized using Correspondence Analysis (CA). Each palette is represented as a dot while each domain is illustrated as a triangle. The proximity between a dot and a triangle shows their association. The top-rated palette from each domain and its corresponding illustration are shown on the right, which is represented as the nearest dot to its corresponding triangle.

To analyze the results of the experiment, we first used Correspondence Analysis (CA) to visualize the joint relationships between the palettes and the domains by transforming user ratings into categorical data [21]. Specifically, we associated a domain to a palette if the domain had earned a score above 4 and received qualified explanations, such as e.g., "*this palette looks very professional and calm, it evokes a feeling of strictness, so I rated business a higher score*" and "*it's a very bright palette, it reminds me of the sky and earth and open-air activities, which my mind associates with sports*". The results of CA show that the total inertia equals 0.49 (chi-square = 947.36, df = 798, p < .000), suggesting a statistically significant association between the palettes and the domains. Fig. 2 presents the results in a two-dimensional diagram, along with the top-rated palette from each domain showing on the right side. In the diagram, the eight domains are represented as triangles while the 115 palettes are represented as dots. We used the color of each triangle and dot to encode the category of its domain. The proximity between a dot and a triangle shows their association while the proximity between two dots shows their similarity. Specifically, a closer distance between an illustration and a domain suggests that the illustration is more likely to display color features related to the domain. For example, we can observe that the participants sometimes confuse the palettes of *sports* with those of *education* as both domains prefer colorful palettes. We also found that the palettes associated with the domains of *business* and *finance* are located close to each other, suggesting that the palettes preferred by the two domains share similar characteristics. We then used Intraclass Correlation Coefficients (ICC) to estimate the consistency among ratings [64]. Table 2 shows that the participants from the study formed relatively consistent judgments on domain-specific palettes (ICC(2,k) = .67).

Items	Intraclass Correlation[b]	95% Confidence Interval		F Test with True Value 0			
		Lower Bound	Upper Bound	Value	df1	df2	Sig
Finance	0.616[c]	0.502	0.713	2.604	114	1026	0.000
Technology	0.731[c]	0.652	0.799	3.721	114	1254	0.000
Healthcare	0.743[c]	0.667	0.808	3.896	114	1026	0.000
Manufacturing	0.752[c]	0.678	0.814	4.024	114	1026	0.000
Sports	0.663[c]	0.563	0.748	2.965	114	1026	0.000
Social media	0.662[c]	0.561	0.747	2.957	113	1017	0.000
Education	0.646[c]	0.541	0.735	2.827	114	1026	0.000
Business	0.569[c]	0.441	0.678	2.320	114	1026	0.000

Table 2. The results of Intraclass Correlation Coefficients (ICC), indicating moderate reliability.

To further understand how the color design of palettes affects the perception of domains, we calculated the Spearman's correlation (ρ) between the mean of each HSV color channel and the mean of the domain scores [68] rated in the crowdsourcing study (Fig. 3). In each diagram, a dot encodes the mean color of a palette and a linear least-square fitting line is added to highlight the correlation trend. For example, we observe a decreasing trend of the fitting lines for mean Value in *manufacturing* and *technology*, suggesting a preference for dark colors. To analyze the favored colors of a domain, we also quantized each HSV color channel into 20 bins and ran the linear RankSVM model on each of them [68]. The weights are displayed as a color stripe above each diagram in Fig. 3, where shades of red (or shades of blue) indicate that the corresponding numeric intervals on the x-axis have a positive (or negative) effect on the perception of a specific domain. The saturation of the shades encodes the extent of the effects. For example, we can see that warm colors (e.g., red to orange to yellow) play a dominant role in designing for *manufacturing* while *technology* shows a strong preference using darker and colder colors (e.g., blue and purple).

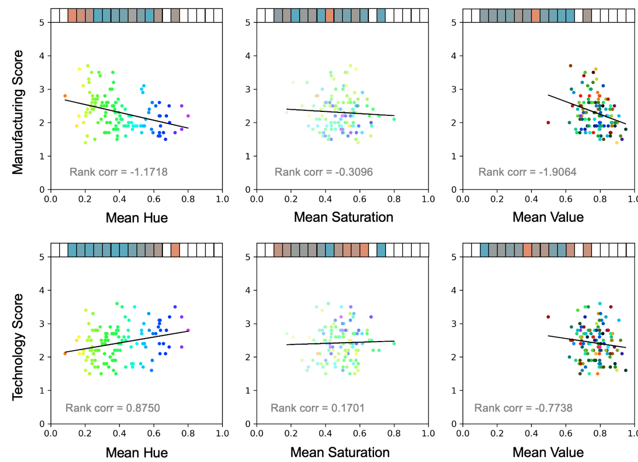

Fig. 3. Correlation between HSV colors and two domains, including manufacturing and technology.

5 COLORCOOK SYSTEM DESCRIPTION

In this section, we introduce ColorCook, a design support tool that facilitates the color design of dashboards in an expressive and effective manner. We present how each of ColorCook's components addresses the design requirements (**DR1**-**DR3**) described in Section 3, accompanied by details on the user interface, interaction, and algorithm.

5.1 User Interface

The user interface of ColorCook consists of four main components: (1) a visualization component for creating charts based on the input dataset (Fig. 4: A, B), (2) a palette library for selecting color palettes associated with the domains of dashboards (Fig. 4: C), (3) a canvas for exploring and adjusting the resulting dashboard colorings (Fig. 4: D, E), and (4) a bookmarking capability for capturing design variations displayed on the canvas (Fig. 4: F). Specifically, the Data Panel allows a user to upload a dataset, whose data fields are automatically extracted and displayed. The Charts Panel provides nine chart types (bar chart, pie chart, line chart, area chart, scatterplot, radar chart, heatmap, treemap, and map), each shows a set of available channels for visual encoding. The Color Palettes Panel displays icon buttons of the eight identified domains in alphabetical order (business, education, finance, healthcare, manufacturing, social media, sports, and technology), each shows the associated palettes from our dataset. The toolbar at the top-left corner contains icon buttons of update, bookmark, save, and trash, supporting exploring, bookmarking, saving, and deleting dashboard colorings, respectively, while the Color Picker facilitates adjusting colors. The Bookmarks Panel is a collapsible panel that lists snapshots of the dashboard colorings manually captured by the user for subsequent reference and comparison.

5.2 ColorCook in Action

We describe how ColorCook works in action through a usage scenario. The scenario highlights the key interactions and functions of ColorCook by illustrating how a user creates dashboard colorings.

Color Selection (DR1). Imagine that the Sales Director, Alice, from an e-commerce company, is preparing to present a weekly sales report to a group of stakeholders. Her goal is to create an informative and expressive dashboard using ColorCook for the purpose of connecting the

Fig. 4. The user interface of ColorCook consists of four main components: (1) users upload their datasets by dragging CSV files into the *Data* Panel (A) and create charts accordingly using the *Charts* Panel (B), (2) users select domain-associated color palettes for their dashboards in the *Color Palette* Panel (C), (3) users view the resulting dashboard colorings on the *Canvas* (D) and adjust them using the *Color Picker* (E), and (4) the *Bookmark* Panel (F) presents snapshots of the dashboard colorings manually captured by users in their design process.

stakeholders to market demand. To start with, Alice uploads a sales dataset by dragging a CSV file into the Data Panel (A), where a set of data fields extracted from the dataset is also displayed. She then uses a bar chart in the Charts Panel (B) to depict weekly sales, where the x-axis, y-axis, and color, respectively, encode date, sales, and product. In the Color Palettes Panel (C), Alice selects the category of 'business' and browses the corresponding domain-associated palettes. She also compares the palettes from the light mode and that from the dark mode, and then decides that an eight-color palette from the dark mode can better communicate the personality of her design. After selecting the palette, she then clicks the 'Generate' button and views the result on the Canvas (D).

Color Assignment (DR2). Alice observes that in the bar chart, each pair of neighboring segments is encoded by discriminable colors, which helps differentiate between the information of different products. She then continues to add more charts to the canvas. She can also move, resize, and delete charts, or clear the canvas if necessary. During the process, Alice visualizes sales by province using a map, where a sequential palette generated from the selected eight-color palette is automatically applied. On the map, she can easily identify the provinces with the highest or lowest sales as the contrast between different color scales is readily apparent and thus improves readability. Following this, Alice creates a heatmap to present the data concerning customer ratings. She observes that ColorCook uses orange and sky blue from the current palette as two endpoints to generate a diverging palette and to encode negative and positive ratings, respectively. Alice also explores different coloring results with the current palette by clicking the 'Update' icon button on the top of the Canvas. Finally, Alice edits the title of each chart and refines the layout of the dashboard.

Color Adjustment (DR3). After determining a satisfactory coloring, Alice would like to add a feeling of elegance and modernity to the overall look of the dashboard by replacing specific colors in the palette. To do this, Alice directly selects a segment colored in purple on the canvas, where all the segments encoding the same data value are highlighted and the Color Picker (E) pops up. The Color Picker allows Alice to set the parameters of R, G, B channels and suggests a set of harmonious colors based on the current palette. The legend of each chart is also interactive; Alice can click a specific legend item marker to adjust the color of its associated visualization segments. To perform color adjustment, Alice can also edit the current palette in the Color Palette Panel (C) by deleting, adding, or replacing specific colors. Throughout the design process, Alice is able to bookmark the evolution of a dashboard coloring as it is refined. She can retrieve a bookmark from the Bookmarks Panel (F) to continue fine-tuning that coloring.

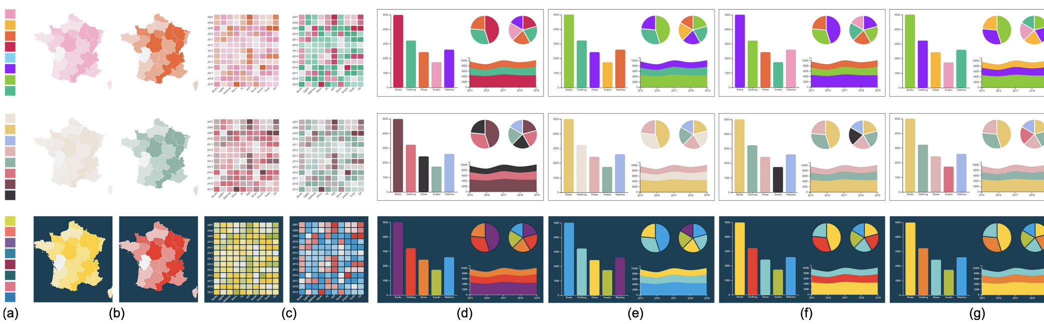

Fig. 5. Given a color palette (a), (b) and (c) show the resulting colorings using sequential and diverging palettes generated from (a), respectively. (d) and (e) show the effect of color harmony term. (f) and (g) show the effects of adding color discriminability and continuity terms, respectively.

5.3 Algorithm

The algorithm of ColorCook consists of three major components, including the components of palette extension as well as color assignment (**DR2**) using a model-driven approach and the component of color suggestion (**DR3**) using a data-driven approach. Given an M-color discrete palette $C = \{C_1, ..., C_M\}$, the objective of the algorithm is to assign the selected colors to an N-field dataset $D = \{D_1, ..., D_N\}$, satisfying both color encoding principles and color perception theories.

Palette Extension. First, the size of the M-color discrete palette C is automatically extended if its size, M, is smaller than the number of data fields, N. In other words, the M colors in C are not enough to encode the N qualitative values if each color is required to represent a distinct category. The color extension method applies the component of color suggestion which will be explained later in this subsection. Then, sequential and diverging palettes are generated from the discrete palette C to support encoding quantitative data types. To select a color C_k from the palette C as the endpoint of a sequential color scale, we applied the heavy-light model HL [40], which is defined as follows:

$$HL = -2.1 + 0.05(100 - L_{C_k}) \tag{1}$$

where L_{C_k} denotes the lightness of C_k in CIELAB color space. Once the heaviest color C_k in C is identified, we can generate its sequential color scale that ranges from light gray to C_k by varying the parameters of L, A, and B in the color space. Fig. 5 (b) shows the results of selecting a color from an eight-color palette (Fig. 5 (a)) for generating a sequential palette which is used to colorize a map. For each palette, the map coloring on the left is ranked lower than the one on the right. We can see that using a heavier color as the endpoint can make each color step more visually perceptible when the intensity of the color decreases.

In contrast to a sequential color scale, a diverging color scale requires two colors with different hues as its endpoints and uses a neutral color in the middle separating the two. In our algorithm, we used two complementary colors as the endpoints of a diverging color scale as they can better visually distinguish positive from negative values. To select a pair of complementary colors C_i and C_j from C, their difference in hue should be constrained to a certain range:

$$| H_{Ci} - H_{Cj} | \in [180^\circ - 60^\circ, 180^\circ + 60^\circ] \tag{2}$$

where H_{Ci} and H_{Cj} respectively denote the hue of C_i and C_j in HSV color space. Similar to a sequential color scale, we generate the transition of a diverging color scale from one endpoint C_i, through a neutral middle (light gray), and to the other endpoint C_j. If such pairs are not observed in C, we instead select one endpoint C_i by finding a color that can act as the endpoint of a sequential color scale and then generate its complementary color C_j. Fig. 5 (c) shows the results of selecting two colors from the palette (Fig. 5 (a)) for generating a diverging palette which is used to colorize a heatmap. For each palette, the heatmap coloring on the left is ranked lower than the one on the right. We can see that using a pair of complementary colors as the endpoints allows users to more easily tell apart one side of quantitative values from the other.

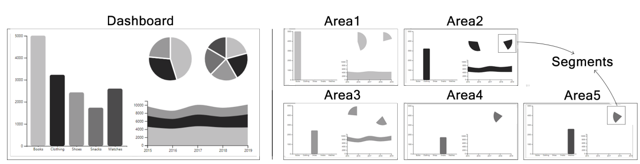

Fig. 6. A dashboard consists of five disjoint areas. Each area is a collection of one or more segments encoding the same data attribute.

Color Assignment. Colorizing a dashboard is about assigning the M-color palette to the N-field dataset, based on the constraints on color consistency. Accordingly, we define a *segment* as a component of a chart in a dashboard that can be color encoded (e.g., a sector in a pie chart, a bar in a bar chart) while an *area* as a set of segments encoding the same data field, as shown in Fig. 6. Here, we relate dashboard areas to data fields as their numbers are equal. An area consisting of one or more segments is assigned to the same color while a dashboard is a collection of such areas. By extending the method of pattern colorization in graphic design [18], we formulated the color assignment problem as follows: given an M-color palette $C = \{C_1, ..., C_M\}$ and an N-area dashboard $A = \{A_1, ..., A_N\}$, the goal is to find a mapping from A to C that maximizes established rules in color perception including color harmony and discriminability. Formally, our formulation is defined as finding a mapping function σ s.t. $\sigma(A_i) = C_j, i \in [1...n], j \in [1...m]$, which maximizes the objective function as follows:

$$E = \sum_{A_i, A_j \in A} E_h(A_i, A_j, \sigma) + E_d(A_i, A_j, \sigma) + \lambda \cdot E_c(A_i, A_j, \sigma) \tag{3}$$

Here, the first term E_h denotes color harmony, which is produced when two or more colors seen in neighboring areas causing a pleasing effect [39]. E_h is measured by computing the chromatic, lightness, and hue for each pair of areas A_i and A_j in CIELab color space. Fig. 5 (d) and (e) show the effects of the color harmony term, where (d) is ranked lower than (e). The second term E_d denotes color discriminability, which describes the difference in lightness between neighboring areas. E_d is measured by computing the L_2 distance between each pair of areas A_i and A_j on the L channel of CIELab color space. We also added name difference [28] to E_d by computing the cosine distance between the probability distributions of all color name pairs. Fig. 5 (f) shows the effect of adding the color discriminability term based on (e). We can see that (f) better enhances the legibility of the dashboard content when compared to (e). We also consider lightness contrast between foreground and background. Note that when preprocessing the palettes, color discriminability was used to help decide if a palette is suitable for light mode, dark mode, or both. In the algorithm, both E_h and E_d encourage stronger effects between adjacent areas by involving the length of the border between the areas into consideration.

The last term E_c in Eq. 3 controls the additional constraint of color continuity. E_c is a variant of the Gestalt laws of continuity and is the opposite of E_h. E_c encourages adjacent areas to have similar lightness values. Fig. 5 (g) shows the effect of adding the color continuity term based on (f). When comparing the (g) and (f), we can observe that (g) presents a smoother change in color between neighboring areas. In our implementation, we used a stochastic optimization process to find the globally optimal solution that maximizes the energy among all possible colorings.

Color Suggestion. We used the color palette rating model proposed by Kita et al. [19] to suggest harmonious colors for a given color palette. The model learns human aesthetic preferences by extracting 121 features from more than 10K human-rated color palettes and learning the weights for each feature vector. To suggest colors, HSV color space is first explored where candidate colors are constrained to those that are essentially different from the colors in the given palette. The candidate colors are then selected by rejection sampling and added to the given palette. Last, each of the palettes is rated by the model, and the 12 top-rated colors are recommended to users.

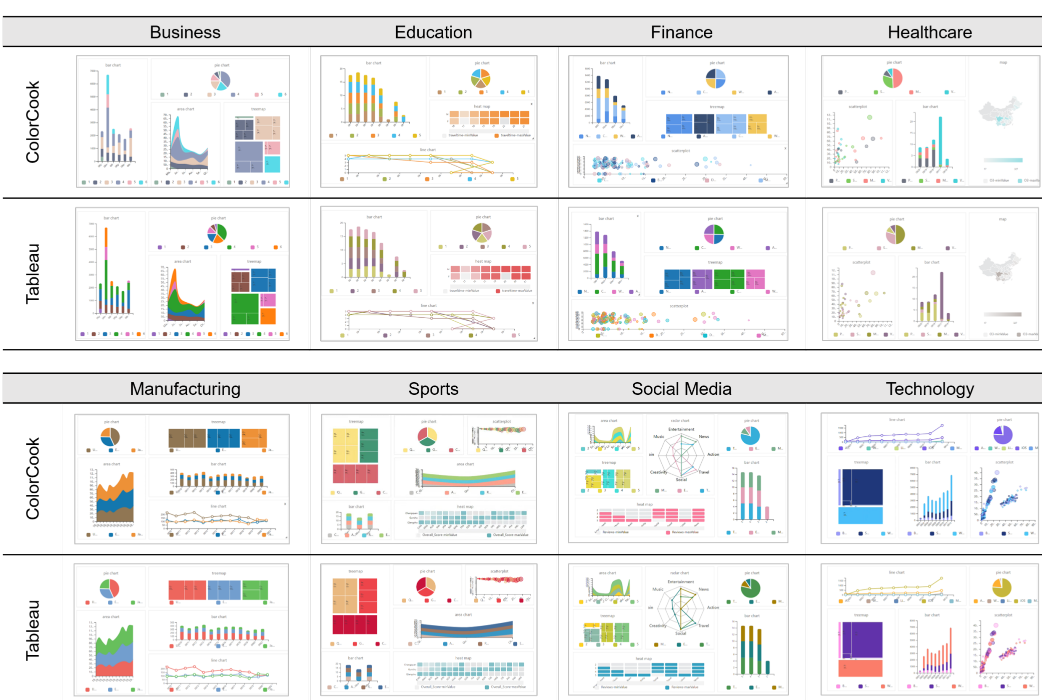

Fig. 7. Dashboard colorings generated using our domain-associated palettes and Tableau palettes.

6 EVALUATION

The evaluation of ColorCook includes a crowdsourcing experiment and a user study. We first conducted a crowdsourcing experiment with 80 participants to estimate the effectiveness of our domain-associated palettes when they are used to colorize dashboards. In the experiment, the participants were instructed to rate the dashboard colorings generated using our palettes and baseline palettes. To further understand how ColorCook is used by professionals, we then carried out a user study with 10 participants. In the user study, the participants were tasked with using the system to create dashboard colorings for specific domains.

6.1 Crowdsourcing Experiment

We compared our palettes with the baseline, Tableau palettes, which are considered as both effective and professional by the experts from our preliminary study. Specifically, we randomly selected one palette from each of the eight domains in our dataset, resulting in eight palettes in total. Then, we used each of the eight palettes to colorize a dashboard which was created using a dataset related to its domain. For example, a business palette was applied to a dashboard visualizing a business dataset. We also randomly selected eight palettes from Tableau palettes (Fig. 1 (a)) and used these palettes to colorize the same set of domain-related dashboards. The colorization process was automatically performed using our algorithm and generated 16 dashboard colorings, with eight using our palettes and eight using Tableau palettes, as shown in Fig. 7.

6.1.1 Methodology. We recruited 80 participants aged between 18 and 35 (M = 29.00, SD = 1.00) for our second crowdsourcing experiment. The experiment consisted of three sessions, of which Session 1 and 3 were similar to the corresponding sessions of our first crowdsourcing experiment. In Session 2, the participants were presented with the 16 dashboard colorings, whose order was randomized. The participants were instructed to rate the quality of the dashboard colorings using a 5-point Likert scale and the measurements include *color harmony*, *color discriminability*, the *association* between a dashboard coloring and the domain of the dashboard, the *expressiveness* of dashboard colorings, and *storytelling capability*. After rating each coloring, the participants were also encouraged to explain their options. Each participant was paid \$5.00 for completing the task. The three sessions lasted about 15-25 minutes for each participant.

6.1.2 Analysis and Results. We conducted a paired t-test to examine if there is a significant difference among each measurement. Fig. 8 shows the quantitative results.

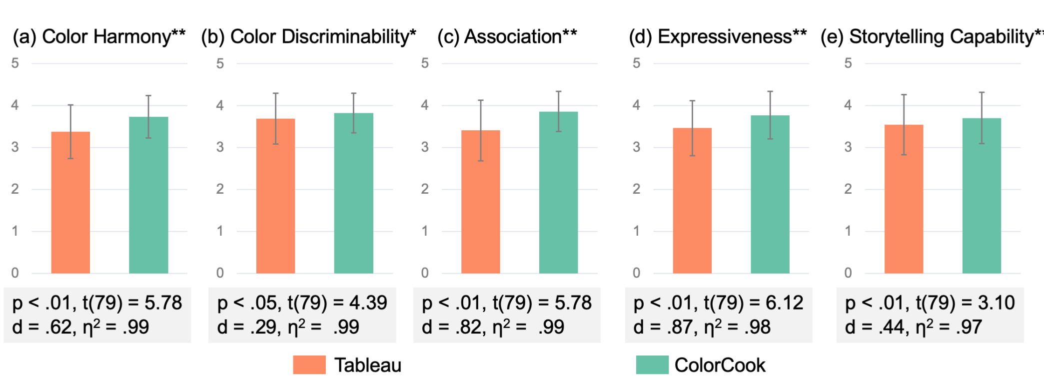

Fig. 8. Means and standard errors of each measurement in Tableau palettes and our palettes conditions ($*$: $p < .05$, $**$: $p < .01$).

Compared to Tableau palettes, our palettes were rated significantly higher in terms of *color harmony* (Fig. 8 (a)) (ColorCook: M = 3.73, SD = .51, Tableau: M = 3.37, SD = .64) and *color discriminability* (Fig. 8 (b)) (ColorCook: M = 3.82, SD = .47, Tableau: M = 3.69, SD = .60). Fig. 8 (c) also shows that significant difference in *association* is found between our palettes and Tableau palettes (ColorCook: M = 3.86, SD = .48, Tableau: M = 3.40, SD = .72). That reason is most of the participants thought that our palettes can remind them of imagery related to a specific domain, "*red, green, and yellow in this sports dashboard remind me of lanes, lawns, and sunshine.*" Dashboard colorings using our palettes can augment their *expressiveness* (Fig. 8 (d)) (ColorCook: M = 3.77, SD = .56, Tableau: M = 3.46, SD = .65), which were perceived as more "*reader-friendly*". The *storytelling capability* (Fig. 8 (e)) (ColorCook: M = 3.70, SD = .61, Tableau: M = 3.54, SD = .72) of ColorCook palettes were significantly better than Tableau. The reason is that semantic information of colors is able to help the participants understand the insights and stories behind the dashboards.

6.2 User Study

To evaluate the effectiveness of ColorCook, we conducted a within-subject user study.

6.2.1 Methodology. In our user study, we recruited participants via the social media platforms of our lab. Our recruitment material demonstrated that we were looking for visualization researchers or practitioners who have dashboarding experience. As a result, 10 participants were recruited and their demographic information is shown in Table 3.

Alias	Age	Gender	Occupation	Affiliation	Experience
P1	24	F	Graduate student	Visualization lab in a university	1~2 years
P2	26	F	Interaction designer	UX department in a bank	1~2 years
P3	24	F	Interaction designer	UX department in a integrated technology enterprise	1~2 years
P4	29	F	Data analyst	Data department of a commercial company	1~2 years
P5	27	M	Visualization researcher	Research institute of a commercial company	2~5 years
P6	25	M	Front-end developer	Data department in a commercial company	2~5 years
P7	29	M	Visualization researcher	Research institute of a commercial company	2~5 years
P8	30	F	Visualization & Interaction designer	UX department in a commercial company	2~5 years
P9	33	M	Visualization expert	Research institute of a commercial company	5+ years
P10	33	M	Visualization & Front-end developer	Data department of a commercial company	5+ years

Table 3. The demographic information and dashboarding experience of our study participants.

Data. We provided the participants with eight public datasets [26], each related to one of the eight domains. We ensured that the eight datasets are of similar complexity regarding their column count and row size. According to previous studies [3, 5], we also provided the participants with a set of extracted data facts (e.g., distribution, rank, trend) as well as available visualization types [61] for each dataset rather than raw data to keep the focus of the user study on color design.

Tasks and Procedure. Our study task instructed the participants to create a dashboard with a focus on color design. The participants were encouraged to decide on one of the eight datasets and create a dashboard based on the corresponding data facts and visualization types. They were also instructed to choose a palette based on the domain of the selected dataset. While exploring the system, the participants could use the think-aloud protocol to add any comments and suggestions they might have.

The study began with a 10-minute introduction explaining the goal of our study and the details of the datasets. In the task, we first demonstrated the use of our system and asked the participants to carry out a 15-minute warm-up dashboarding exercise using a mock dataset. At the end of the study, we collected dashboards from the participants and conducted a 20-minute semi-structured interview. After that, each participant filled out a brief demographic survey and completed a questionnaire using a 7-point Likert scale. The study lasted about 0.5-1 hour for each participant. The interaction process, completion time, and semi-structured interview were recorded for subsequent analysis.

Measurements. Our questionnaire consisted of 17 items. Item 1-6 were an attempt to capture the aspects of usefulness and ease of use regarding the system. We selected items for our questionnaire from the criteria commonly used for user interface usability and user experience evaluations [29, 58], including *usefulness, effectiveness, satisfaction, and desirability, ease of use*, and *ease of learning*. Item 7-12 measure the quality of dashboard colorings considering the characteristics of the tasks: the *effectiveness* of color encoding, *color harmony*, *color discriminability*, the *association* between a dashboard coloring and the domain of the dashboard, *expressiveness*, and *storytelling capability*. We also included five additional criteria, items 13-17, related to creativity support [11]: *enjoyment, exploration*, *expressivity*, *immersion*, and *results worth effort.*

6.2.2 Analysis and Results. Based on the analysis, we now report both the quantitative and qualitative results of the user study. Fig. 9 shows the quantitative results of the questionnaire rated by the participants. By analyzing qualitative feedback, we attempted to understand the participants' comments in more depth and highlight their thoughts behind the quantitative results.

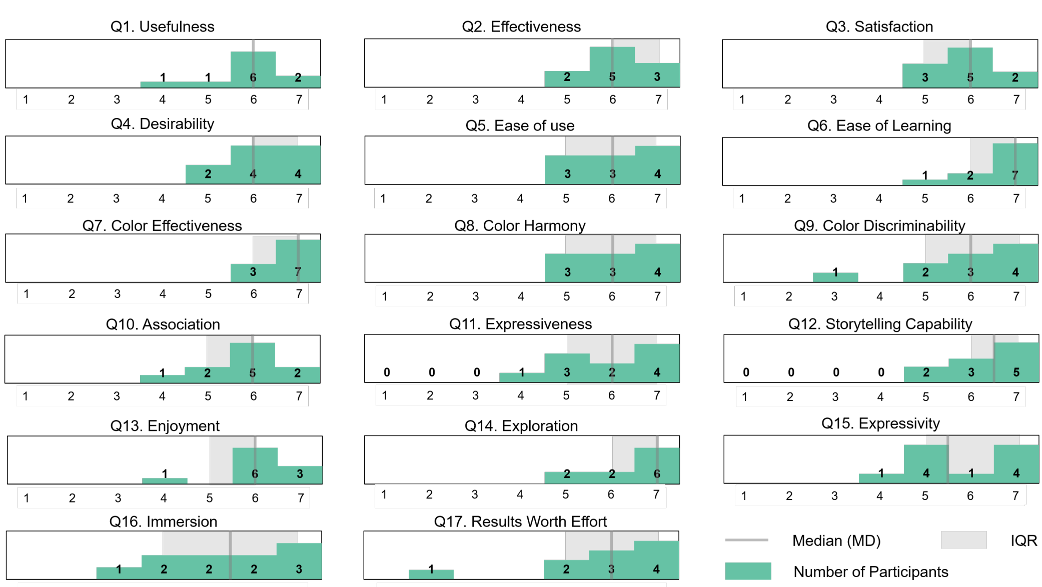

Fig. 9. The participants' responses to overall experience, coloring quality, and creativity support of ColorCook. Each participant answered on a 7-point likert scale with the questionnaire.

Color Manipulation. When analyzing the log data of the participants' color workflow, we counted the average frequency of different operations (#*palette switching*: 0.35, #*coloring updating*: 0.04, #*local color changing*: 0.59, #*bookmarking*: 0.32). We can see that the frequency of *coloring updating* is low, as most of the participants were satisfied with the results of our color assignment algorithm. In terms of *local color changing*, we observed that the participants with a design background were more likely to adjust the colors themselves for customization, while those without a design background were more likely to use the colors recommended by ColorCook for efficiency. One participant suggested, "*intelligence and flexibility are well-balanced in ColorCook which would attract border audience.*"

ColorCook is Useful and Easy to Use. Overall, ColorCook received good ratings in terms of *usefulness* (*MD* = 6.00, *IQR* = 0.00; Q1), *effectiveness* (*MD* = 6.00, *IQR* = 1.00; Q2), *satisfaction* (*MD* = 6.00, *IQR* = 1.00; Q3) and *desirability* (*MD* = 6.00, *IQR* = 1.00; Q4), all having a median rating of 6. 40% of the participants applauded the desirability of ColorCoook, giving the highest rating. The participants rated *ease of use* (*MD* = 6.00, *IQR* = 2.00; Q5) and *ease of learning* (*MD* = 7.00, *IQR* = 1.00; Q6) as good. 70% of the participants rated 7 for *ease of learning* and one of them said, "*uploading datasets, generating dashboard colorings are one-click operation. The process was finished in less than a second and largely reduced my learning costs.*"

ColorCook Can Generate High-Quality Results. Overall, ColorCook performed well in the aspect of generating high-quality results. Almost all of the participants rated ColorCook from 5 to 7 in the *effectiveness* of color encoding (*MD* = 7.00, *IQR* = 1.00; Q7), *color harmony* (*MD* = 6.00, *IQR* = 2.00; Q8) and *color discriminability* (*MD* = 6.00, *IQR* = 2.00; Q9). However, one participant gave a low rating as 4 in color discriminability. He explained that color in ColorCook had lower saturation, which might lead to lower discriminability. Regarding the *association* (*MD* = 6.00, *IQR* = 1.00; Q10) between a dashboard coloring and the domain of the dashboard, two participants gave the highest rating as 7 as ColorCook palettes allow them to relate to various domains such as metal and grass. The participants also acknowledged ColorCook in terms of *expressiveness* (*MD* = 6.00, *IQR* = 2.00; Q11) and *storytelling capability* (*MD* = 6.50, *IQR* = 1.00; Q12), suggesting it "*effectively communicates data being colorized*" (P3).

ColorCook Provides Creativity Support. ColorCook received positive ratings in terms of *enjoyment* (*MD* = 6.00, *IQR* = 1.00; Q13), *exploration* (*MD* = 7.00, *IQR* = 1.00; Q14), *expressivity* (*MD* = 5.50, *IQR* = 2.00; Q15), *immersion* (*MD* = 5.50, *IQR* = 3.00; Q16), and *results worth effort* (*MD* = 6.00, *IQR* = 2.00; Q17). Specifically, *exploration* received the highest median score and the smallest IQR, which suggests that ColorCook can inspire new ideas, "*ColorCook helps fast prototyping of visualization dashboards"* (P9).

ColorCook Applies an Integrated Color Workflow. Most participants gave positive feedback on the effectiveness of the color workflow in ColorCook which is "*smooth*" and "*practicable*". P3 said, "*this system seems useful. If Tableau integrated it into its chart authoring interface, people would probably use it and like it.*" P7 noted, "*categorizing palettes by domains is a brilliant idea, which can help communicate the characteristics of the dashboard. It also greatly reduces the time I spent on selecting colors.*" P6 said, "*ColorCook supports 'global' color encoding; the palette I selected was used to colorize all the charts in the dashboard. In this case, I don't have to manually set it for individual charts. Such a process is usually time-consuming and draining, especially as the number of data attributes increases.*" P10 commented on the function of color adjustment, "*I can see diversity in the colors suggested by ColorCook, and I can always find a satisfying one for my dashboard.*" P3 had a different experience for Tableau and Colorcook, "*it always takes me a lot of time to adjust the colors for each area individually with Tableau, which is frustrating. The color workflow in ColorCook is quite simple and intuitive that I can get the job done quickly.*"

ColorCook is Effective on Expressive Color Design. During the interview with each participant, we presented the two dashboards he or she respectively created using ColorCook and Baseline as references, as shown in Fig. 10. 7 out of 10 participants noted that ColorCook is capable of delivering better results when compared to Baseline. P6 and P8 provided positive feedback on the effectiveness of color encoding in ColorCook, as it "*automatically applies different palettes based on data types and these palettes constitute a harmonious color design*" and "*maintains color consistency across multiple charts*" (Fig. 10 (a)). Also, 5 out of 10 participants addressed the usefulness of the domain-associated palettes provided by ColorCook. "*Manufacturing has serious tones like gray scales. I want the design* (Fig. 10 (b)) *to create serious vibes* (P6)." "*I selected this palette in technology as it feels calm and cold... This coloring leaves an impression of cool aspects of technology such as AI and robot*" (P7).

ColorCook Offers Satisfying User Experience. All the participants were satisfied with the user experience provided by ColorCook. In terms of interaction, the instant feedback provided by ColorCook improve efficiency, with P5 saying "*it's more user-friendly compared to existing dashboard tools, I can achieve a satisfying result with a few tweaks and save time for experimenting with different color schemes.*" In addition, nearly half of the participants mentioned that ColorCook offers them great freedom of configuration based on intelligence, P8 noted, "*automated tools often lack of flexibility, whereas ColorCook has great DIY capabilities and even gives me some recommendations when I change local colors.*" In terms of functionality, P5 said, "*what I really like about the system is that it provides the function of bookmarking; I can easily compare and retrieve previous steps of my current design process.*" P2 reflected, "*I noted that the system provides palettes for both dark mode and light mode. Such categorization is helpful when you have to design dashboards for different screens and devices. Enabling dark mode on large screens, where dashboards are meant to display, can improve visual ergonomics by reducing eye strain.*"

ColorCook Serves as a Design Assistant. All the participants agreed that ColorCook can assist and augment human creative capabilities. P3 thought it helps the process of ideation and exploration, "*it works as an intelligent color design assistant and you get what you want within a second; the result looks great from both perspectives of visualization design as well as graphic design.* " P7 suggested that ColorCook provides design support as it "*lowers the barrier to information design*" and "*augments expressivity by finishing the task with one-click and allowing users to spend more time on refinement*". P6 noted, "*personally, I think that creating a satisfying dashboard coloring in a very short time can build a sense of fulfillment. It serves as a good starting point and then I can add some personalities*

Fig. 10. Dashboard colorings created with ColorCook by (a) participant 3 and (b) participant 9, respectively, from the user study.

to my design". However, P8 and P10 expressed their concern that automated colorization might conflict with data exploration, "*the process of color encoding provides an opportunity for users to understand data facts in more depth but now it's somehow been automated*". The participants also suggested promising research directions that ColorCook can follow. For example, P1 noted, "*it'd be interesting to see how the system works for colorizing multidimensional data.*" P6 implied, "*color discriminability should be improved based on chart types. I'm looking forward to seeing its effect on heat map or even graphs.*"

ColorCook Helps Data Communication in Collaboration. Color plays an important role in dashboards, which are main vehicles for data communication within organizations. On one hand, domain-associated dashboard colorings can help the audience understand the data insights, resulting in immersion in data-driven storytelling. P6 has a deep understanding of that, "*I once made a dashboard for a conference. I found that purple or blue appeared elegant and professional, and the audience can emerge themselves in the vast universe while watching my dashboard. It could be easier to create that atmosphere with ColorCook.*" On the other hand, the ability of ColorCook to represent the same data attribute with the same color across different charts makes the data more contextually coherent and the communication in collaboration more understandable. P10 said, "*"ColorCook tied different charts together and made my presentation clearer. When I wanted to highlight a field I just hovered over the corresponding data category and all the areas associated with it were highlighted. Such a function really helps when presenting data to my clients during the meeting.*"

7 DISCUSSION

In this section, we discuss design implications that have arisen through the course of developing and evaluating ColorCook, as well as limitations in our current work.

7.1 Expressive Information Design and Domain-Associated Palettes

Based on the analysis of the crowdsourcing study, we found that the association between a palette and a domain frequently emerges from embodiment, that is, people intend to relate colors to the imagery that evokes tangible or situated feelings. The results of our user study suggested that such embodied palettes can help increase the expressiveness of information design. For example, in a user study task where the participants were instructed to design a dashboard visualizing manufacturing data, one reflected, "*I used a toned down, colorful, artsy palette. It makes me think of raw materials and it surely conveys strength and power for a manufacturing dashboard.*" In our future work, we plan to explore other color factors except for domains that should be considered in expressive information design. For example, affective colors can promote an effective understanding of information in visualization [7]. When designing an advertising poster concerning technology, one can select a palette containing vivid and bright colors. Such a coloring looks exciting and cheerful and thus boosts sharing.

The current palette library for ColorCook was derived from a collection of illustrations. As our dataset is not exhaustive nor representative of the field, the construction of the palette library is considered as an initial step toward understanding the domain-associated colorization of dashboards. Our future work plans to extend the current dataset by collecting more domain-associated palettes from a variety of online sources. For example, company branding is a potential source to explore. The reason is that a company in a certain domain usually builds its branding using a specific color scheme. Such color choices are also perceived as domain-related. We hope that such domain-associated palettes can further contribute to the HCI community and benefit future work by allowing professionals and practitioners to gain more insights into color usage.

7.2 Balance between Constraints and Freedom to Design

According to the quantitative and qualitative feedback received from our user study, the participants were satisfied with the design support provided by ColorCook. The participants suggested that laborious and tedious work in color design such as experimenting with different palettes can be done automatically. Thus, more efforts can be focused on creative activities such as refining and improving the results.

In terms of design constraints, one participant from the expert interview raised an interesting question, "*I can't help wondering what if someone wants to use hundreds of different colors? What's the limit?*" We acknowledge that extending a discrete palette to support the growing demand for a dashboard can become less helpful if badly used. For example, Excel allows users to create charts with dozens of color variations and thus opens a floodgate of badly designed charts. Thus, our future work should involve data cardinality and available encoding resources into consideration to decide if suggesting more colors or cycling current colors is a better strategy. Some of the participants addressed their concerns about the degree of freedom. For example, one participant suggested that a fully automated color workflow may limit users' design decisions and lower their perceived controllability. P10 noted, "*I don't have to worry about aesthetic issues as ColorCook will tell me what're good choices, but somehow I feel it's little space for me to add my own thoughts and ideas.*" Thus, constraints on the option to adjust or add colors should also be carefully designed for future color design tools to leave room for creativity support.

7.3 Generalization to More Visualization Genres

We observed that the participants are mostly positive toward the color workflow supported by ColorCook. They suggested that such automated workflow can provide assistance for color design which usually involves much trial-and-error tweaking. "*It works like your design assistant who prepares and presents various design drafts and you're the boss to decide which one to use*" (P1). Integrating such workflow into ColorCook also allows for rapid prototyping, from which dashboard authors can receive quick feedback on design decisions and experiment with multiple ideas. The color workflow we identified for designing dashboard colorings can also be generalized to other narrative visualization genres such as infographics [21] and data videos [51, 66]. These visualization genres can be accessible to a broader audience and used on social media such as TikTok and Twitter. For example, such a workflow can be modified by designers to maintain color consistency cross charts for infographics or across panels for data comics, as color has the most impact on identifying elements [4]. Similarly, data video designers can apply such a workflow to help highlight attributes in the data and facilitate the understanding of evolving data relationships over time [3].

7.4 Limitations and Future Work

Our current palette extension component for ColorCook was designed to support generating single-hue sequential palettes and diverging palettes composed of two hues. Our future work includes improving the algorithm by enabling generating continuous palettes with multiple hues. Second, some participants noted that the color assignment component overlooks selecting semantically-resonant colors for visualization [23]. For example, when visualizing gender attributes in the data, red and blue can encode female and male categories, respectively. Future work should consider integrating color semantics into the algorithm. Third, the current algorithm for color suggestion does not ensure the recommended colors are still specific to the domain associated with the original palette. Our future work plans to expand our dataset and leverage a learning-based approach to suggesting domain-associated colors. Another interesting research direction to explore is to learn personal color preferences and encourage human-AI co-creation. For example, if a user fine-tunes the colors of his or her design, the system can record the user's adjustments and use such information to guide the subsequent generation. Or, if the user prefers using highly saturated colors for a specific dashboard, the system can predict potential domains of the dashboard and then suggest relevant palettes.

8 CONCLUSION

In this work, we introduced ColorCook, a design support tool that facilitates expressive and effective color design for dashboarding. ColorCook assists in the color workflow consisting of color selection, assignment, and adjustment. Quantitative and qualitative feedback from the evaluation showed that ColorCook can effectively help design high-quality dashboard colorings. Future work includes applying our approach to collaborative dashboard coloring. Such application will help identify the effects of our approach on color consistency when different users apply different color encodings as well as color harmony when they prefer diverse color choices. Also, we hope this work can shed light on the exploration and development of model- and data-driven approaches to expressive information design.

9 ACKNOWLEDGMENTS

Nan Cao is the corresponding author. This work was supported by NSFC 62072338 and NSF Shanghai 20ZR1461500.

REFERENCES

[1] QlikTech International AB. 1993. Qlik. https://www.qlik.com. Accessed: 2022.
[2] Adobe. 2020. Adobe Color. https://color.adobe.com/create. Accessed: 2022.
[3] Fereshteh Amini, Nathalie Henry Riche, Bongshin Lee, Christophe Hurter, and Pourang Irani. 2015. Understanding Data Videos: Looking at Narrative Visualization Through the Cinematography Lens. In *Proceedings of the 2015 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 1459–1468.
[4] Benjamin Bach, Natalie Kerracher, Kyle Wm. Hall, Sheelagh Carpendale, Jessie Kennedy, and Nathalie Henry Riche. 2016. Telling Stories about Dynamic Networks with Graph Comics. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 3670–3682.
[5] Benjamin Bach, Zezhong Wang, Matteo Farinella, Dave Murray-Rust, and Nathalie Henry Riche. 2018. Design Patterns for Data Comics. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 38:1–38:12.
[6] Hyojin Bahng, Seungjoo Yoo, Wonwoong Cho, David Keetae Park, Ziming Wu, Xiaojuan Ma, and Jaegul Choo. 2018. Coloring with words: Guiding image colorization through text-based palette generation. In *Proceedings of the European Conference on Computer Vision*. 431–447.
[7] Lyn Bartram, Abhisekh Patra, and Maureen Stone. 2017. Affective color in visualization. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1364–1374.
[8] Lawrence D Bergman, Bernice E Rogowitz, and Lloyd A Treinish. 1995. A Rule-Based Tool for Assisting Colormap Selection. In *Proceedings Visualization'95*. IEEE, 118–125.
[9] Roxana Bujack, Terece L Turton, Francesca Samsel, Colin Ware, David H Rogers, and James Ahrens. 2017. The Good, the Bad, and the Ugly: A Theoretical Framework for the Assessment of Continuous Colormaps. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 923–933.
[10] Xi Chen, Wei Zeng, Yanna Lin, Hayder Mahdi Al-maneea, Jonathan Roberts, and Remco Chang. 2021. Composition and Configuration Patterns in Multiple-View Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 1514–1524.
[11] Erin Cherry and Celine Latulipe. 2014. Quantifying the Creativity Support of Digital Tools through the Creativity Support Index. *ACM Transactions on Computer-Human Interaction* 21, 4 (2014), 21:1–21:25.
[12] Micheline Elias and Anastasia Bezerianos. 2012. Annotating BI Visualization Dashboards: Needs & Challenges. In *Proceedings of the 2012 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 1641–1650.
[13] Hui Fang, Simon Walton, Emily Delahaye, James Harris, DA Storchak, and Min Chen. 2016. Categorical Colormap Optimization with Visualization Case Studies. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2016), 871–880.

[14] Mark Harrower and Cynthia A Brewer. 2003. ColorBrewer.org: An Online Tool for Selecting Colour Schemes for Maps. *The Cartographic Journal* 40, 1 (2003), 27–37.
[15] Jeffrey Heer and Maureen Stone. 2012. Color Naming Models for Color Selection, Image Editing and Palette Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1007–1016.
[16] Michael Hockenhull and Marisa Leavitt Cohn. 2021. Speculative Data Work & Dashboards: Designing Alternative Data Visions. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW3 (2021), 1–31.
[17] IsoFlat. 2019. IsoFlat. https://isoflat.com. Accessed: 2022.
[18] Hye-Rin Kim, Min-Joon Yoo, Henry Kang, and In-Kwon Lee. 2014. Perceptually-based Color Assignment. *Computer Graphics Forum* 33, 7 (2014), 309–318.
[19] Naoki Kita and Kazunori Miyata. 2016. Aesthetic Rating and Color Suggestion for Color Palettes. *Computer Graphics Forum* 35, 7 (2016), 127–136.
[20] Shigenobu Kobayashi. 1981. The aim and method of the color image scale. *Color research & application* 6, 2 (1981), 93–107.
[21] Xingyu Lan, Yang Shi, Yueyao Zhang, and Nan Cao. 2021. Smile or Scowl? Looking at Infographic Design Through the Affective Lens. *IEEE Transactions on Visualization and Computer Graphics* 27, 6 (2021), 2796–2807.
[22] Ricardo Langner, Ulrike Kister, and Raimund Dachselt. 2018. Multiple Coordinated Views at Large Displays for Multiple Users: Empirical Findings on User Behavior, Movements, and Distances. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 608–618.
[23] Sharon Lin, Julie Fortuna, Chinmay Kulkarni, Maureen Stone, and Jeffrey Heer. 2013. Selecting Semantically-Resonant Colors for Data Visualization. In *Proceedings of the 2013 Eurographics Conference on Visualization*. The Eurographs Association & John Wiley & Sons, Ltd., 401–410.
[24] Sharon Lin and Pat Hanrahan. 2013. Modeling How People Extract Color Themes from Images. In *Proceedings of the 2013 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 3101–3110.
[25] Yang Liu and Jeffrey Heer. 2018. Somewhere Over the Rainbow: An Empirical Assessment of Quantitative Colormaps. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 598:1–598:12.
[26] Google LLC. 2010. Kaggle. https://www.kaggle.com/datasets. Accessed: 2022.
[27] Icons8 LLC. 2020. Ouch. https://icons8.com/illustrations. Accessed: 2022.
[28] Kecheng Lu, Mi Feng, Xin Chen, Michael Sedlmair, Oliver Deussen, Dani Lischinski, Zhanglin Cheng, and Yunhai Wang. 2020. Palettailor: Discriminable colorization for categorical data. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 475–484.
[29] Arnold M Lund. 2001. Measuring usability with the use questionnaire12. *Usability interface* 8, 2 (2001), 3–6.
[30] Ellen Lupton and Jennifer C. Phillips. 2008. *Graphic Design: The New Basics* (1st ed ed.). Princeton Architectural Press, Maryland Institute College of Art.
[31] Ruixian Ma, Honghui Mei, Huihua Guan, Wei Huang, Fan Zhang, Changye Xin, Wenzhuo Dai, Xiao Wen, and Wei Chen. 2021. LADV: Deep Learning Assisted Authoring of Dashboard Visualizations from Images and Sketches. *IEEE Transactions on Visualization and Computer Graphics* 27, 9 (2021), 3717–3732.
[32] Ruixian Ma, Wei Wang, Fan Zhang, Kyuha Shim, and Carlo Ratti. 2019. Typeface Reveals Spatial Economical Patterns. *Scientific Reports* 9, 1 (2019), 15946:1–15946:9.
[33] Microsoft. 2016. Power BI. https://powerbi.microsoft.com. Accessed: 2022.
[34] Kenneth Moreland. 2009. Diverging Color Maps for Scientific Visualization. In *International Symposium on Visual Computing*. Springer, Springer Berlin Heidelberg, Berlin, Heidelberg, 92–103.
[35] Randall Munroe. 2010. Color survey results. https://blog.xkcd.com/2010/05/03/color-survey-results/. Accessed: 2022.
[36] LLC NAICS Association. 1993. North American Industry Classification System. https://www.naics.com. Accessed: 2022.
[37] Peter O'Donovan, Aseem Agarwala, and Aaron Hertzmann. 2011. Color compatibility from large datasets. In *ACM SIGGRAPH 2011 papers*. ACM, New York, NY, USA, 1–12.
[38] Peter O'Donovan, Aseem Agarwala, and Aaron Hertzmann. 2011. Color Compatibility From Large Datasets. *ACM Transactions on Graphics* 30, 4 (2011), 63:1–63:12.
[39] Li-Chen Ou and M Ronnier Luo. 2006. A colour harmony model for two-colour combinations. *Color Research & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de la Couleur* 31, 3 (2006), 191–204.
[40] Li-Chen Ou, M Ronnier Luo, Andrée Woodcock, and Angela Wright. 2004. A Study of Colour Emotion and Colour Preference. Part I: Colour Emotions for Single Colours. *Color Research & Application* 29, 3 (2004), 232–240.
[41] Q. Huy Phan, Hongbo Fu, and B. Antoni Chan. 2018. Color Orchestra: Ordering Color Palettes for Interpolation and Prediction. *IEEE Transactions on Visualization and Computer Graphics* 24, 6 (2018), 1942–1955.

[42] Jorge Poco, Angela Mayhua, and Jeffrey Heer. 2017. Extracting and Retargeting Color Mappings from Bitmap Images of Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 637–646.
[43] Prolific. 2014. Prolific. https://www.prolific.co. Accessed: 2022.
[44] Zening Qu and Jessica Hullman. 2017. Keeping Multiple Views Consistent: Constraints, Validations, and Exceptions in Visualization Authoring. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 468–477.
[45] Khairi Reda, Pratik Nalawade, and Kate Ansah-Koi. 2018. Graphical Perception of Continuous Quantitative Maps: The Effects of Spatial Frequency and Colormap Design. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 272:1–272:12.
[46] Khairi Reda and Danielle Albers Szafir. 2020. Rainbows Revisited: Modeling Effective Colormap Design for Graphical Inference. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1032–1042.
[47] Salesforce. 2003. Tableau. https://www.tableau.com. Accessed: 2022.
[48] Alper Sarikaya, Michael Correll, Lyn Bartram, Melanie Tory, and Danyel Fisher. 2018. What Do We Talk about When We Talk about Dashboards? *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 682–692.
[49] Karen B Schloss, Connor C Gramazio, Allison T Silverman, Madeline L Parker, and Audrey S Wang. 2018. Mapping Color to Meaning in Colormap Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 810–819.
[50] Vidya Setlur and Maureen C Stone. 2015. A Linguistic Approach to Categorical Color Assignment for Data Visualization. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2015), 698–707.
[51] Yang Shi, Xingyu Lan, Jingwen Li, Zhaorui Li, and Nan Cao. 2021. Communicating with motion: A design space for animated visual narratives in data videos. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–13.
[52] Maria Shugrina, Wenjia Zhang, Fanny Chevalier, Sanja Fidler, and Karan Singh. 2019. Color Builder: A Direct Manipulation Interface for Versatile Color Theme Authoring. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 456:1–456:12.
[53] Stephen Smart and Danielle Albers Szafir. 2019. Measuring the Separability of Shape, Size, and Color in Scatterplots. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 669:1–669:14.
[54] Stephen Smart, Keke Wu, and Danielle Albers Szafir. 2019. Color Crafting: Automating the Construction of Designer Quality Color Ramps. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 1215–1225.
[55] Danielle Albers Szafir. 2017. Modeling Color Difference for Visualization Design. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 392–401.
[56] Martijn Tennekes and Edwin de Jonge. 2014. Tree Colors: Color Schemes for Tree-Structured Data. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2072–2081.
[57] Melanie Tory, Lyn Bartram, Brittany Fiore-Gartland, and Anamaria Crisan. 2021. Finding Their Data Voice: Practices and Challenges of Dashboard Users. *IEEE Computer Graphics and Applications* (2021), 1–1.
[58] Thomas Tullis and William Albert. 2013. *Measuring the User Experience: Collecting, Analyzing, and Presenting Usability Metrics*. Morgan Kaufmann, Amsterdam.
[59] Andrea Vázquez-Ingelmo, Francísco J Garcia-Peñalvo, and Roberto Therón. 2019. Information Dashboards and Tailoring Capabilities-A Systematic Literature Review. *IEEE Access* 7 (2019), 109673–109688.
[60] Yunhai Wang, Xin Chen, Tong Ge, Chen Bao, Michael Sedlmair, Chi-Wing Fu, Oliver Deussen, and Baoquan Chen. 2018. Optimizing Color Assignment for Perception of Class Separability in Multiclass Scatterplots. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 820–829.
[61] Yun Wang, Zhida Sun, Haidong Zhang, Weiwei Cui, Ke Xu, Xiaojuan Ma, and Dongmei Zhang. 2019. DataShot: Automatic Generation of Fact Sheets from Tabular Data. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 895–905.
[62] Steve Wexler, Jeffrey Shaffer, and Andy Cotgreave. 2017. *The Big Book of Dashboards: Visualizing Your Data Using Real-World Business Scenarios*. John Wiley & Sons, Inc.
[63] Martijn Wijffelaars, Roel Vliegen, Jarke J Van Wijk, and Erik-Jan Van Der Linden. 2008. Generating Color Palettes Using Intuitive Parameters. *Computer Graphics Forum* 27, 3 (2008), 743–750.
[64] Ziming Wu, Taewook Kim, Quan Li, and Xiaojuan Ma. 2019. Understanding and Modeling User-Perceived Brand Personality from Mobile Application UIs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–12.
[65] Mehmet Adil Yalçın, Niklas Elmqvist, and Benjamin B Bederson. 2017. Keshif: Rapid and Expressive Tabular Data Exploration for Novices. *IEEE Transactions on Visualization and Computer Graphics* 24, 8 (2017), 2339–2352.
[66] Leni Yang, Xian Xu, XingYu Lan, Ziyan Liu, Shunan Guo, Yang Shi, Huamin Qu, and Nan Cao. 2021. A Design Space for Applying the Freytag's Pyramid Structure to Data Stories. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 922–932.

[67] Linping Yuan, Wei Zeng, Siwei Fu, Zhiliang Zeng, Haotian Li, Chi-Wing Fu, and Huamin Qu. 2021. Deep Colormap Extraction from Visualizations. *IEEE Transactions on Visualization and Computer Graphics* (2021), 1–13.

[68] Nanxuan Zhao, Ying Cao, and Rynson W. H. Lau. 2018. What Characterizes Personalities of Graphic Designs? *ACM Transactions on Graphics* 37, 4 (2018), 116:1–116:15.

[69] Liang Zhou and Charles D Hansen. 2015. A Survey of Colormaps in Visualization. *IEEE Transactions on Visualization and Computer Graphics* 22, 8 (2015), 2051–2069.

Received January 2022; revised April 2022; accepted May 2022