

# Redundant is Not Redundant: Automating Efficient Categorical Palettes Design Unifying Color & Shape Encodings with *CatPAW*

Chin Tseng

University of North Carolina at Chapel Hill  
Chapel Hill, NC, USA  
chint@cs.unc.edu

Ghulam Jilani Quadri

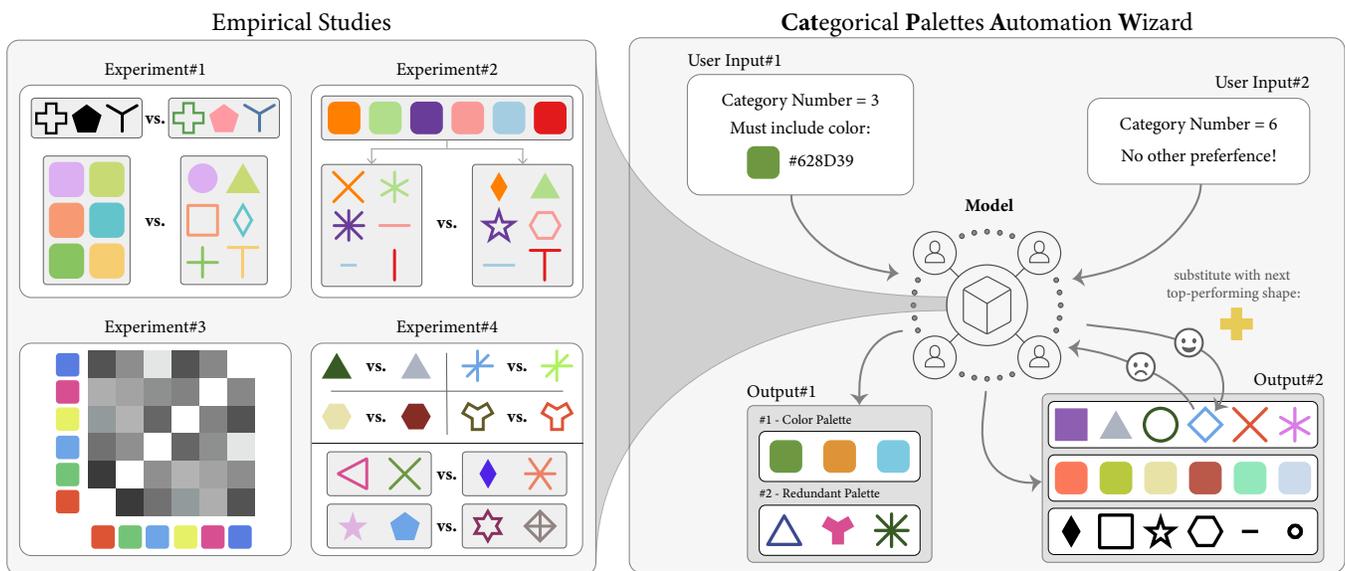
University of Oklahoma  
Oklahoma City, OK, USA  
quadri@ou.edu

Arran Zeyu Wang

University of North Carolina at Chapel Hill  
Chapel Hill, NC, USA  
zeyuwang@cs.unc.edu

Danielle Albers Szafr

University of North Carolina at Chapel Hill  
Chapel Hill, NC, USA  
danielle.szafr@cs.unc.edu



**Figure 1:** We present *CatPAW*, a web-based tool for generating effective categorical palettes using a data-driven model built from four crowd-sourced experiments. Users can specify requirements—such as the number of categories, must-include colors or shapes, and palette type (color-only, shape-only, or redundant). Based on performance estimates drawn from empirical data, *CatPAW* suggests a set of optimal palettes adhering to the provided constraints. Users can interactively swap out undesired elements, and *CatPAW* will suggest high-ranking alternatives by considering the current palette composition. Rounded rectangles represent color-only palettes.

## Abstract

Colors and shapes are commonly used to encode categories in multi-class scatterplots. Designers often combine the two channels to create redundant encodings, aiming to enhance class distinctions. However, evidence for the effectiveness of redundancy remains conflicted, and guidelines for constructing effective combinations are limited. This paper presents four crowdsourced experiments evaluating redundant color–shape encodings and identifying

high-performing configurations across different category numbers. Results show that redundancy significantly improves accuracy in assessing class-level correlations, with the strongest benefits for 5–8 categories. We also find pronounced interaction effects between colors and shapes, underscoring the need for careful pairing in designing redundant encodings. Drawing on these findings, we introduce a categorical palette design tool that enables designers to construct empirically grounded palettes for effective categorical visualization. Our work advances understanding of categorical perception in data visualization by systematically identifying effective redundant color–shape combinations and embedding these insights into a practical palette design tool.



This work is licensed under a Creative Commons Attribution 4.0 International License.  
*CHI '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2278-3/2026/04  
<https://doi.org/10.1145/3772318.3791244>

## CCS Concepts

• **Human-centered computing** → **Visualization systems and tools**; **Empirical studies in visualization**; **Empirical studies in HCI**.

## Keywords

Categorical perception, color, shape, redundant encoding, multi-class scatterplots, visualization effectiveness, perceptual mechanism

### ACM Reference Format:

Chin Tseng, Arran Zeyu Wang, Ghulam Jilani Quadri, and Danielle Albers Szafir. 2026. *Redundant is Not Redundant: Automating Efficient Categorical Palettes Design Unifying Color & Shape Encodings with CatPAW*. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3772318.3791244>

## 1 Introduction

Effective palettes for representing categorical data are crucial for data visualization [74]; however, we have limited guidance as to how to effectively encode categorical data [27, 79]. Categorical encodings principally use color or shape *palettes*, predefined collections of shapes or colors that are then assigned to different categories [86]. Designers frequently differentiate categories using both color and shape simultaneously to increase each category’s perceptual distinctiveness. Such *redundant encoding* is assumed to increase people’s abilities to distinguish between different categories and has been widely applied in real-world scenarios from scientific papers [82] to business reports [14]. However, the benefits of redundant encoding are, in part, contentious, with some studies showing little to no impact on effectiveness [23], others showing notable benefits for given tasks [51], and still others showing that the channels themselves may interact, suggesting more complex effects for redundancy [68]. Part of this tension may arise from the lack of empirically-grounded guidance for designing effective categorical visualization [81]: categorical palettes are largely designed by heuristics, leading to potential inconsistencies in encoding effectiveness. Existing guidance for redundant palette design treats the two channels individually, either helping designers build a color palette [26] or a shape palette [79] and assuming their pairing will lead to improved visualizations. We lack clear guidance for how to effectively pair color and shape to maximize the potential benefits of redundancy.

In this paper, we explore redundant encoding design to understand its utility and derive empirical guidance for effective redundant palette design. Part of the challenge in reconciling color and shape for palette design is that the two spaces are computationally incompatible. Color is sampled from continuous numerical spaces; therefore, it can fit with a diverse range of design and creation processes. For example, individual color palettes can be created based on manual selection [32], mathematical color metrics [26], or even machine learning models [36] sampling over continuous spaces. However, shape lacks a common space or set of agreed-upon features and is less likely to have a universal descriptive language [12, 86]. Even coarse-grained features describing shape fail to explain shape palette effectiveness [79]. Past approaches for

constructing shape palettes rely on preselected shapes or on empirical data used to build pairwise accuracy models [79]. Therefore, even though the use of redundant encoding can improve accuracy for tasks like visual segmentation and grouping [51], we still lack an understanding of how to effectively design redundant palettes and how specific redundant encodings (i.e., color-shape pairings) affect task performance in data visualization [81].

We address this knowledge gap using a data-driven approach, collecting empirical data about expert-crafted palettes to generate a method for crafting effective redundant palettes. We developed this method in a series of four crowdsourced experiments asking people to compare correlations across different categories. We structure these experiments to first evaluate the effectiveness of redundant encodings and then establish effective color-shape pairings in multi-class scatterplots across varying numbers of categories. Our studies first investigate the effectiveness of redundancy for designer palettes (Experiment 1) and then explore the interactions of shape and color in redundant palette design (Experiment 2). Finally, we construct pairwise accuracy matrices for colors in redundant mappings (Experiment 3) and for subsequent interactions between colors and shapes (Experiment 4) to develop a statistical model measuring the performance of color-shape pairings to guide effective redundant palette design.

Our findings indicate that redundant encodings improve accuracy in assessing class-level correlations, in line with both heuristic guidance and past studies [51]. However, redundancy only significantly benefited performance in visualizations with moderate category numbers (five to seven categories). Furthermore, performance varied between different combinations of color and shape palettes, emphasizing the importance of choosing appropriate combinations of redundant encoding design. Notably, effective redundant palette design is not as simple as blending high-performing color and shape palettes: combining top-performing shapes and colors did not consistently result in top-performing redundant encodings.

Using the data from our studies, we developed a statistical model that generates effective redundant palettes based on the number of categories present in a visualization. We incorporate this model into a design tool, **CatPAW—Categorical Palette Automation Wizard**—which assists designers in creating optimized categorical palettes. **CatPAW** supports color-only, shape-only, and color-shape redundant encodings tailored to different category numbers and user preferences for color or shape. By unifying the color and shape channels, the tool facilitates more informed and efficient palette design. Overall, this work contributes to a more integrated understanding of redundant encoding for categorical perception in data visualization and offers an actionable framework for encoding categorical data.

**Contribution:** We empirically investigated the effectiveness of redundant encodings, specifically color and shape pairings, in data visualizations. Our findings help reconcile contradictory evidence on the effectiveness of redundancy, showing that redundancy can benefit categorical analysis, but that this benefit is in part dependent on the number of categories as well as the palette design. Our results built up a deeper understanding of how redundant encodings improve people’s perception and overall task performance in visualizations and suggest that optimal redundant palette design should consider the interplay between color and shape. Driven by

those insights, we created an automated palette generation tool that leverages the data generated in our experiments to aid designers in creating categorical encodings.

## 2 Related Work

Providing empirical evidence for design heuristics is crucial for effective visualization design [19, 43, 74, 81]. In this section, we review prior research on the use of redundant encodings in data visualization, focusing on three key areas: graphical perception, the role of color and shape in categorical perception, and redundant encoding in data visualization.

### 2.1 Graphical Perception of Scatterplots

Graphical perception refers to people’s ability to interpret and understand information presented in graphical form [50, 74]. Graphical perception studies explore how visual encoding and design elements in visualizations, such as position [37], axis [45], angle [84], and color [61], are perceived and processed by people to provide empirical guidance for visualization design. Starting from Cleveland and McGill’s initial work four decades ago [16], prior research on graphical perception identified a hierarchy of perceptual tasks, demonstrating that certain visual encodings are more effective than others for conveying quantitative information [52].

Scatterplots are among the most widely used visualization types in graphical perception research [56]. Designers often use scatterplots to encode multiple categories of data in the same space, enabling direct comparisons between different categories of data [50, 59, 74]. As a result, scatterplots have been used to derive a variety of insights on perceptual efficiency in visualization design, both for improved scatterplots and for general data representation practices [52, 56, 56]. Graphical perception experiments have used scatterplots to explore effective visualization design in a diverse range of tasks, including value judgments [23], correlation [31], comparison [22], trend [17], clustering [66], and uncertainty [60], and visual encodings, such as color [48], shape [12], aggregations [89], and texture [1]. Among different types of scatterplots, multiclass scatterplots, our focus in this work, have played a pivotal role in understanding the visual data communication process [21].

### 2.2 Color and Shape in Categorical Perception

Categorical perception involves the brain’s ability to categorize continuous stimuli into discrete groups [24]. In the context of data visualization, we take advantage of these processes to represent different classes of data by applying encodings that intuitively create distinct visual groups [18, 78, 80]. Color and shape are the channels most commonly used to differentiate between categories [74, 81]. People can efficiently process categorical differences in color [78] and shape [79], but the effectiveness of these encodings can vary depending on several factors [81].

Research in cognition and vision science describes how different color attributes, such as hue [9], saturation [91], and lightness [63], affect our ability to perceive and interpret categorical information. Hue is often considered a primary color feature to consider for categorical data, with best practices emphasizing mapping classes to distinct hues [27, 32, 71]. Predominantly hue-varying categorical color palettes can achieve high levels of accuracy in tasks like

comparing category means [78]. However, hue in isolation is often insufficient for effective visualization design. Lightness and saturation also contribute to people’s abilities to distinguish categorical data [58, 80]. Beyond the fundamental perceptual attributes of color, its psychological and cultural associations can also influence how viewers interpret categorical data visualizations. Color carries semantic or affective meaning that may be more important than common perceptual factors for visualization [5, 44, 61, 92]. For example, warm colors like red and orange might be associated with urgency or excitement [42], while cool colors such as blue and green may convey calmness or stability [65]. Although not the central focus of this study, these factors can influence user interpretation, preference, and overall engagement with the data; however, their utility and design typically must align with the semantic content of the data or its associated context.

Shape is another fundamental visual channel for categorical data encoding [86]. Shape is considered an effective preattentive attribute [38], meaning that it can be processed by the visual system quickly and without conscious effort, making it particularly useful for conveying categorical information at a glance. The perceptual discriminability of shapes is crucial for effective categorical encoding; however, we lack formalized computational spaces for reasoning about shape. Past efforts have categorized widely-used shapes in data visualization using broad categories, such as open and closed [12]. Q-Tons [85] offer a linear sequence of shapes for ordinal data, but, like most shape palettes, are hand-designed to capture the idea of an "increasing" shape value. Despite the lack of standard spaces for reasoning about shape, shape remains a common choice for visualization, often preferred to color due to its increased accessibility for people with color vision deficiency (CVD) [71]. Guidelines for selecting shape palettes for data visualization are limited, but suggest considering the distinctiveness of shape and ease of visual processing [18]. However, existing shape palette design relies heavily on intuition and the designer’s preferences in real-world practices. More recently, *Shape It Up* explored a data-driven approach to create effective shape palettes based on experimental data collected over a series of categorical data analysis tasks [79].

Past research into both shape and color perception for visualization usually considers each channel individually [81]. While limited past evidence shows the two channels may interact [68], these studies focus on analyzing small differences in color better suited to continuous data. Despite their common pairing [82], we lack an understanding of how these channels may interact in categorical encodings. In this work, we aimed to combine color and shape together to measure their performance, better understand the influence of redundancy on categorical encoding, and build more efficient palette design tools.

### 2.3 Redundant Encoding in Data Visualization

Redundant encoding refers to the practice of using more than one visual variable to represent the same underlying data. A widely-used redundant encoding design combines color and shape. For instance, one category might be represented by red circles, while another is shown as blue squares. People employ redundant encodings largely to make it easier to distinguish between different categories of data

[51]. Redundant encoding with color and shape can also offer improved accessibility, particularly for individuals with CVD [4]. For viewers who have difficulty distinguishing between certain colors, the shape of the data point can serve as an alternative visual cue, allowing them to still effectively differentiate between categories [40]. Conversely, for individuals with typical color vision, the color can reinforce the information the shape conveys, potentially leading to faster and more accurate perception [40]. However, employing redundancy means that there are fewer visual channels available to encode data and can increase the complexity of the visualization.

Redundant encoding, especially for scatterplots, is available in a number of commercial and research toolkits. For example, Tableau [77], R [54], and Matlab [39] all allow people to combine colors and shapes in one encoding; however, these combinations reflect the default shape and color palettes and require manual adjustment by the user to change. The *scatterHatch* package in the *R/Bioconductor* [29] is specifically designed to create plots for massive single-cell data that use redundant coding by combining colors with patterns. This package aims to enhance the accessibility and distinguishability of large-scale single-cell spatial data in bioinformatics, which typically contain large numbers of categories. However, *scatterHatch* primarily focuses on bioinformatics data, and the encoded patterns are usually not exactly shapes but their variations, such as a textured plot with obscured dots, which is not that often in real-world visual data communication [8, 21, 74].

Despite its popularity, redundant encoding remains controversial. Previous research has demonstrated that redundant encodings can improve the speed and accuracy of typical vision science tasks, including data segmentation and grouping [51]. However, alternative studies have found little to no benefit for redundant shape and color encoding in aggregation tasks [23]. Detractors also argue that redundant encoding comes at a significant cost. For instance, encoding too much or irrelevant information can clutter the visualization and hinder rather than help the viewer’s understanding [20]. In the context of scatterplots, which naturally require higher cognitive processing [16, 53], using too many visual dimensions simultaneously, even if intended as redundant encodings, may make the graph appear more complex and more difficult to read. Shape and color may interact perceptually. That is, shape may change perceptions of color and color of shape [68], making it harder to merge these two channels without better models of their interaction. However, we lack informative guidance for crafting redundant encodings; instead, most popular tools typically pair the user’s chosen color and shape palettes without regard to the interactions of the associated colors and shapes.

Our research builds upon the existing tensions surrounding redundant encodings to better characterize effective redundant encoding design. We achieve this by systematically investigating specific color-shape redundant pairings across a range of conditions. We begin by confirming the effectiveness of redundancy and then develop a series of pair-wise matrices describing color-, shape-, and combined color-shape effectiveness in redundant encoding design using empirical data. We use these matrices to design a tool to help people create effective redundant palettes.

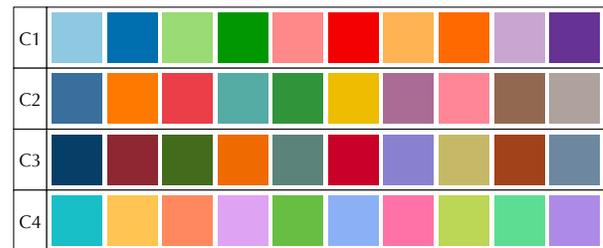
### 3 Experiment One: Do Redundant Encodings Help?

Our first study examined how the choice of non-redundant encodings (color-only or shape-only) versus redundant encodings (color-and-shape) impacts people’s abilities to reason across multiclass scatterplots. We performed a crowdsourced study measuring people’s performance in comparing category correlations over varying category numbers ( $N = 2-10$ ).

We hypothesized that: **H<sub>1</sub>: (a) Redundant encodings would improve overall performance, and (b) differences in accuracy between redundant and non-redundant encodings will be larger as category number grows, but may plateau at a certain category number.** Redundant encodings use an additional channel to encode categorical data. We expect people may benefit from having an extra reference for differentiating categories [51, 86]. People may be relatively strong at distinguishing between categories for smaller numbers of categories [81], such as the three category sets in Gleicher et al.[23]; however, we expect that the benefits of added perceptual cues emerge as the number of categories (and, as a result, difficulty of categorical analysis) increases. As performance tends to be asymptotic with respect to the number of categories in single encodings [81], we anticipate the benefits of redundant encodings will likewise eventually saturate as the visualizations increase in complexity [78, 79].

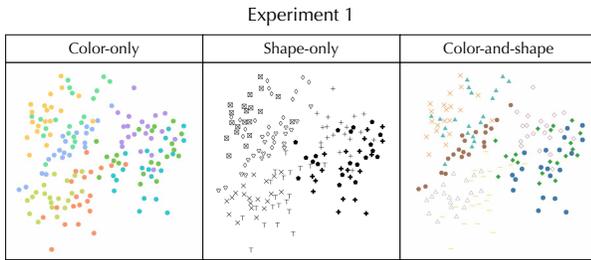
#### 3.1 Experimental Design

**3.1.1 Task.** We employed a correlation judgment task, which is one of the most widely employed tasks in previous studies of multiclass data visualization [31, 41, 55], and is frequently employed in a range of real-world applications, including data reporting [57] and journalism [2]. The task asked participants to identify the category in a multiclass scatterplot with the highest correlation. While this task likely engages similar perceptual mechanisms as averaging [75, 78], it represents another common use of scatterplots where estimation accuracy depends on people’s ability to differentiate between categories [59].



**Figure 2: The four color palettes used in Experiment 1 and 2. Each palette has 10 colors, and drawn from (1) Color-Brewer/Paired [32], (2) Tableau/Tab10 [77], (3) Stata/S2 [70], and (4) Carto/Pastel [13]**

**3.1.2 Stimuli Generation.** Participants estimated category correlations across a series of scatterplots using either colors, shapes, or color-and-shape to differentiate categories. We generated each



**Figure 3: Example stimuli from Experiment 1: scatterplots encoding categorical data using color-only, shape-only, and redundant (color-and-shape) encodings. All three visualizations use the same data distribution with 8 categories. Correlation comparison tasks were used in Experiment 1.**

scatterplot as a 400×400 pixel graph using D3, as shown in Figure 3. Each scatterplot was rendered to a white background and two orthogonal black axes with 13 unlabeled ticks.

We used four color palettes found in existing tools, shown in Figure 2, to encode categorical data in scatterplots: 1. ColorBrewer/Paired [32], 2. Tableau [77], 3. Stata/S2 [70], and 4. Carto/Pastel [13]. We chose these color palettes from previous studies [78] and based on their differences in  $L^*$  variance,  $L^*$  magnitude, and all-pairs perceptual distance [67], which are significant factors that may impact categorical perception. Such selection can also reduce bias from different color palettes.

We encoded shape using 39 shapes collected from a previous study [79]. The 39 shapes were from popular visualization tools: D3 [10], Tableau [77], Excel [49], R [54], and Matlab [39]. For redundant encodings, we preconstructed palettes by randomly selecting both colors and shapes from the above collections for each trial. For instance, in a scatterplot with five categories, five colors were randomly chosen from a specific color palette, while five shapes were randomly selected from the set of 39. These colors and shapes were then randomly paired, with each color assigned to a single shape, and each combination represented a unique category. The randomization helps minimize bias related to the ordering of colors and shapes or their pairings.

Points were rendered in a 6×6-pixel window. Point locations were generated using NumPy’s random multivariate method [30], ensuring that the target category had the highest Pearson correlation coefficient (ranging from 0.8 to 0.95), while the second-highest category maintained a correlation at least 0.2 lower than the target, consistent with thresholds used in past studies [79]. All other categories had a randomly-selected correlation below this threshold. Each category consistently contained 20 points. To minimize confounding effects from spatial overlap, we applied jittering to any overlapping points. This point generation strategy was applied consistently across all experiments.

To evaluate whether redundant encodings enhance performance on multiclass scatterplots across varying numbers of categories, we generated datasets containing  $N = [2 - 10]$  categories. We use an upper bound of 10 categories as this aligns with recent studies of categorical encoding [78, 79], reflects the size of the smallest tested color palette (Tableau 10), allows us to keep the study tractable in terms of the number of tested conditions, and,

in piloting, we found that larger category numbers exhibited little change in over, consistent with past findings on Fetchner’s Law in categorical scatterplots [81]. As an exhaustive comparison of all colors and shapes in our encoding pool is prohibitively large, for each category number, we generated 20 different sets for each encoding: 20 color palettes, each consisting of five sets of  $N$  random colors selected from each of the four color palettes, and 20 shape sets, each generated by randomly selecting  $N$  different shapes from the pool of 39 shapes. We paired these two sets of colors and shapes to create the set of redundant encodings. These resulting 540 stimulus designs (3 encoding types × 20 encoding sets × 9 category numbers), which were divided into 10 task groups. The category numbers and encoding type were equally distributed across task groups, with each containing 54 stimuli.

**3.1.3 Procedure.** Our experiment was structured into three phases: (1) obtaining informed consent, (2) task description and tutorial, and (3) the formal study. Initially, participants provided informed consent in accordance with our IRB protocol and then supplied demographic information. Next, they were introduced to the correlation task description alongside example scatterplots demonstrating varying levels of correlation. The example of study UI is shown in the OSF Supplements. Participants were asked to complete a tutorial with three practice questions, each with a different encoding type (color, shape, and redundant encoding). They were required to answer all the tutorial questions correctly before moving forward to ensure a clear understanding of the task.

During the formal study, participants performed our target task (“Identify which category is the most correlated”) for 57 stimuli presented sequentially. The trials were displayed in a randomized order. At the start of the study, each participant was randomly assigned to one of 10 task groups. Participants had 20 seconds to respond to each stimulus, after which time the answer was marked as incorrect and the study advanced to the next trial. To ensure valid participation, we included three engagement checks. These engagement checks were stimuli with two or three categories that had at least a  $\Delta r = 0.4$  difference in their Pearson correlation coefficient (see the OSF Supplements for examples). We randomly placed the engagement checks throughout the 54 formal trials.

**3.1.4 Participants.** We recruited 105 participants through Amazon Mechanical Turk (MTurk), requiring a minimum 95% approval rating and residency in the US or Canada. Five participants who failed more than one engagement check were excluded from the analysis. The final data included 100 participants (72 male, 28 female), ranging in age from 20 to 65 years, with 10 participants assigned to each task group. All participants reported having normal or corrected-to-normal vision. Our experiment took 10 minutes on average, and each participant was compensated \$1.60 for their time.

Crowdsourcing can introduce inherent variation in stimulus appearance. This variation is ecologically consistent with visualization use and aligns with a wealth of past graphical perception research. These trade-offs limit our abilities to make precise claims about human vision, but do enable us to draw conclusions about graphical perception appropriate to the context of visualization design and use. We follow best practices for crowdsourced study design [7, 33]. To maintain consistent stimulus sizes, we restricted

**Table 1: ANOVA results for category number and encode type (color-only, shape-only, color-and-shape). Significant effects are indicated by bold text.**

Experiment 1	DF <sub>1</sub>	DF <sub>2</sub>	F-value	p-value
Category number	8	792	116.35	<b>&lt;.0001</b>
Encode type	2	198	27.32	<b>&lt;.0001</b>
category number * encode type	16	1584	1.68	<b>0.043</b>

participation to devices no smaller than standard laptops or monitors. We additionally leveraged colors that were robustly different in crowdsourced environments at our tested stimulus size [73].

**3.1.5 Analysis.** We used accuracy as our primary dependent measure. To compare the performance between non-redundant encodings (color-only or shape-only) and redundant encodings (color-and-shape), we analyzed the resulting data using an ANOVA with category number and encoding type as independent variables. Table 1 summarizes our results. The anonymized data and results for our study can be found on OSF Supplements<sup>1</sup>.

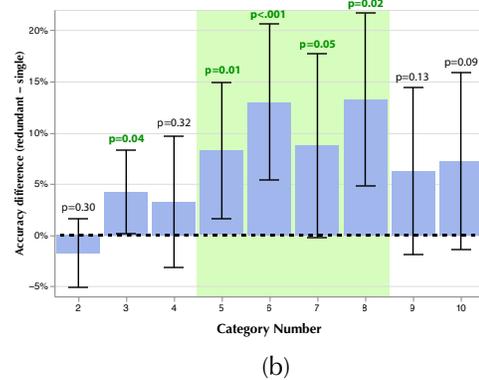
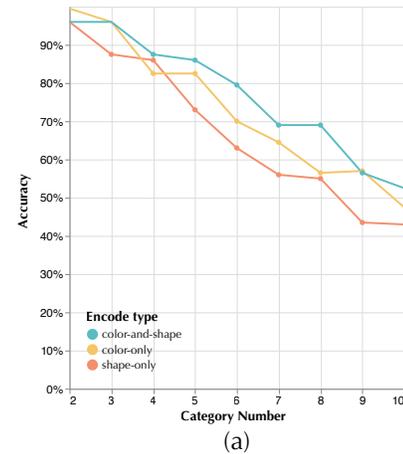
## 3.2 Results

Our analysis revealed that both category numbers ( $F(8, 792) = 116.35, p < .0001$ ) and encoding types ( $F(2, 198) = 27.32, p < .0001$ ) significantly affected correlation comparisons (Table 1). Figure 4 (a) displays the average accuracy separated by encoding type. Notably, redundant color-and-shape achieved the highest average accuracy at 76.9% (95% CI=[74.9, 78.8]), representing an improvement of 4% over color-only encoding (72.9%, 95% CI=[70.8%, 74.9%]) and 9.9% over shape-only encoding (67%, 95% CI=[64.8%, 69.2%]).

To better understand how redundancy benefits vary across category numbers, we combined color-only and shape-only into a single non-redundant baseline in Figure 4 (b). On average, redundant encodings achieved an accuracy of 76.9% compared to 69.4% for non-redundant encodings, with the largest difference observed at 8 categories (13.2%) and the smallest at 2 (-1.75%). Redundant encodings improved performance most consistently when the number of categories ranged from 5 to 8.

To further quantify these differences, we conducted Welch’s  $t$ -tests at each category number to compare the redundant encodings against non-redundant encodings (Figure 4 (b)). The results show that redundant encodings significantly outperformed non-redundant encodings for 3 categories ( $t = 2.06, p = 0.04$ ), 5 categories ( $t = 2.44, p = 0.01$ ), 6 categories ( $t = 3.36, p < .001$ ), 7 categories ( $t = 1.93, p = 0.05$ ), and 8 categories ( $t = 3.11, p = 0.02$ ). No significant differences were observed at 2 ( $t = -1.03, p = 0.30$ ), 4 ( $t = 0.99, p = 0.32$ ), 9 ( $t = 1.51, p = 0.13$ ), or 10 ( $t = 1.65, p = 0.09$ ) categories. These findings align with the trend in Fig. 4(a) that redundancy is most beneficial when visual complexity is moderate (5–8 categories), but less helpful when category numbers are too low or too high. Overall, these results support  $H_1$ : applying redundant encodings enhances performance on correlation-estimation tasks. However, the benefit diminishes when the number of categories is below 5 or exceeds 8.

<sup>1</sup>[https://osf.io/4up7j/?view\\_only=43e5681b71fc4fdb92617b8cef27c60](https://osf.io/4up7j/?view_only=43e5681b71fc4fdb92617b8cef27c60)



**Figure 4: Results from Experiment 1. (a) Line chart showing average accuracy across category sizes (2 to 10), separated by scatterplots using color-only, shape-only, and color-and-shape (redundant) encodings. (b) Bar chart with 95% CI showing the accuracy difference between redundant and non-redundant encodings (redundant minus single-channel: color-only or shape-only) across category numbers. The green background highlights the category range where redundant encoding provides the largest performance gains. For each category number, we conducted a Welch’s  $t$ -test; statistically significant p-values are shown in bold green text, indicating conditions where redundant encoding significantly outperforms single-channel encodings. Overall, redundant encoding yields higher accuracy than either color-only or shape-only encodings, with the most substantial improvements observed in the 5–8 category range.**

## 4 Experiment Two: Interactions between Color Palettes and Shape Palettes

While Experiment 1 demonstrated that redundant encodings can enhance people’s performance in interpreting scatterplots, it remains unclear whether the specific method of creating redundant encodings also impacts task accuracy. In other words, the way designers pair colors and shapes may influence people’s ability to distinguish between classes. Smart et al. [68] demonstrated perceptual interactions between colors and shapes, highlighting the

importance of their pairing. However, their study focused solely on single shapes and small changes in color rather than categorical designs in multi-class scatterplots, which require reasoning over larger numbers of marks and leveraging greater color differences. In Experiment 2, we investigate how redundant coding palettes influence performance. We use six shape palettes generated from *ShapeltUp* [79] and the four color palettes from Experiment 1 to measure the influence of color-shape pairings.

We hypothesized that: **H<sub>2</sub>: mappings between color palettes and shape palettes impact participants' accuracy in correlation estimation.** Because shapes differ considerably in their colored areas (e.g., filled circles versus unfilled circles), different color attributes may either enhance or hinder performance. For example, more similar colors within a palette may be robust if paired with filled shapes that have more overall color area, but less so with thin, line-based shapes with less color area [68]. If the mapping between color and shape palettes does not matter, we expect that the relative performance of redundant palettes changes as a function of each channel's individual performance (e.g., color palette X always outperforms color palette Y if X and Y are mapped to the same shape palette). However, if mapping matters, we expect that the choice of both color and shape palette will change their relative effectiveness.

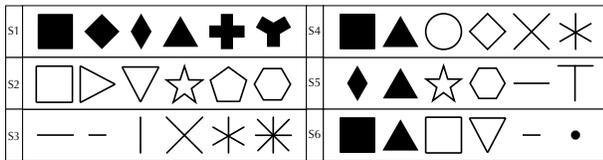


Figure 5: The six shape palettes used in Experiment 2. Each palette consists of six shapes and was generated using *ShapeltUp* [79].

## 4.1 Experiment Design

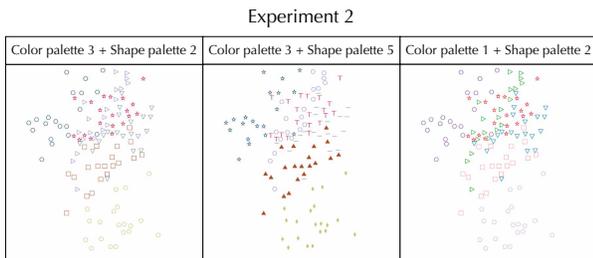


Figure 6: Example stimuli from Experiment 2: three scatterplots using different combinations of color and shape palettes, shown with the same data distribution of 6 categories. Correlation comparison tasks were used in Experiment 2.

**4.1.1 Task & Stimuli Generation.** We used the same task, point distribution, and number of points as in Experiment 1. As shown in Figure 2 and Figure 5, we selected 4 color palettes and 6 shape palettes. We elected to understand mapping effects at the scale

of palettes to reflect design practices used in current visualization authoring tools where people choose palettes which the tool then pairs rather than selecting individual color-shape pairings, which we investigate in more detail in Experiments 3 and 4. The color palettes were chosen for their distinct color attributes (e.g.,  $L^*$  magnitude), consistent with Experiment 1. The shape palettes were generated using *ShapeltUp* [79] and included one filled shape palette, one unfilled shape palette, one open shape palette, and the three highest-recommended palettes selected from a pool of 39 shapes.

We fixed the number of categories at six, since this value is the midpoint of our tested range in Experiment 1 (i.e., between 2 and 10) and exhibited some of the most prominent improvements from redundant encodings. By combining the 6 shape palettes (Figure 5) with the 4 color palettes (Figure 2), we generated 24 combinations of color-and-shape redundant palettes. Within color-and-shape palettes, we randomly paired colors and shapes, with one color assigned to one shape, representing a unique category. Finally, we generated 10 stimuli for each of the 24 color-and-shape palettes, resulting in a total of 240 stimulus designs. These were then evenly divided into five task groups, each containing all 24 palettes with two stimuli for each color-shape palette combination.

**4.1.2 Procedure & Participants.** We followed the same general procedure as in Experiment 1. Each participant completed 51 trials (48 formal trials and 3 engagement checks), presented sequentially in a random order. We recruited 61 participants under the same criteria as in Experiment 1. Six participants who failed the engagement checks were excluded from the analysis, resulting in 55 participants (24 female, 31 male; age range: 28–65), with 11 participants per task group. On average, each participant spent 10 minutes on the experiment and was compensated \$1.60 for their time.

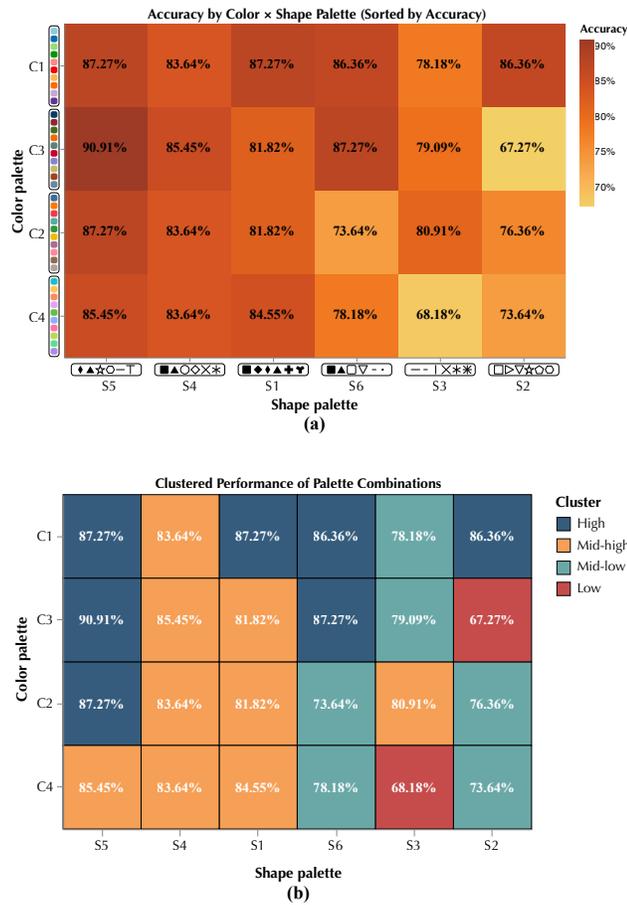
Table 2: ANOVA results for color palette and shape palette. Significant effects are indicated by bold text.

Experiment 2	DF <sub>1</sub>	DF <sub>2</sub>	F-value	p-value
Shape palette	5	270	6.79	<b>&lt;.0001</b>
Color palette	3	162	2.99	<b>0.029</b>
shape palette * color palette	15	810	1.83	<b>0.026</b>

## 4.2 Results

We analyzed how the choices of color palettes, shape palettes, and their combinations influence response accuracy using a two-factor ANOVA and Tukey's HSD for posthoc analysis. Our analysis revealed that color palettes, shape palettes, and their pairings all significantly affected correlation comparisons (see Table 2). Figure 7 (a) displays a heatmap where the x-axis represents six different shape palettes and the y-axis represents four different color palettes, with cell values indicating the average task accuracy for each color-shape redundant encoding. The highest average accuracy (90.91%) was achieved with color palette No. 3 paired with shape palette No. 5, whereas the lowest accuracy (67.27%) was also observed with color palette No. 3 when paired with shape palette No. 2.

The overall accuracy on the correlation comparison task was 81.59% (95% CI=[80.11%, 83.07%]). The individual average accuracies



**Figure 7: (a)** Heatmap showing the average accuracy for each combination of color and shape palettes. The x-axis lists the six shape palette IDs and the y-axis lists the four color palette IDs (corresponding to Figure 5 and Figure 2, respectively). Both axes are sorted by average accuracy to facilitate comparison across palette combinations. Each cell displays the mean task accuracy for the associated palette pairing. As shown, the same color palette can produce different accuracies depending on the shape palette with which it is paired, and vice versa, illustrating strong palette-mapping effects. **(b)** Clustered visualization of the same color-shape combinations using a four-cluster Ward hierarchical solution, clustering color-shape palette combinations based on performance. Each cell is assigned to one of four performance clusters—High, Mid-high, Mid-low, or Low—revealing distinct families of palette pairings. The cluster structure highlights that neither color nor shape palettes are consistently effective on their own: several palettes appear in both high- and low-performing clusters depending on the pairing.

for the six shape palettes, in the order shown in Figure 5, were 83.86%, 75.9%, 76.59%, 84.09%, 87.72%, and 81.36% with  $\sigma = 4.21\%$ . For the four color palettes, the individual average accuracies were 84.84%, 80.6%, 81.96%, and 78.93% with  $\sigma = 2.16\%$  in Figure 2 order. The variability in accuracy across shape palettes was slightly greater than color palettes.

We calculated the variance of accuracies across all color palettes using the same shape palettes and across shape palettes using the same color palettes. The results show that shape palette No. 2 had the largest variance in accuracies across the four color palettes, and color palette No. 3 (Stata/S2) had the largest variance across the six shape palettes. These findings suggest that certain shape or color palettes may be more sensitive to pairing choices and need closer attention during design. Furthermore, our ANOVA results (see Table 3) indicate a significant interaction effect between color and shape palettes, supporting  $H_2$ : choices of how to pair color and shape palettes influence task performance. As shown in Figure 6 and Figure 7 (a), color palette No. 3 (Stata/S2), when paired with shape palette No. 5, achieved the overall highest accuracy. In contrast, pairing the same color palette with shape palette No. 2 (unfilled shapes) resulted in the overall lowest accuracy. On the other hand, pairing shape palette No. 2 with color palette No. 1 (ColorBrewer/Paired) improved accuracy by nearly 20% over its pairing with the Stata/S2 color palette.

We additionally used Ward’s method to cluster the palette combinations based on performance, allowing us to examine broader performance patterns across the 24 color-shape combinations (Figure 7 (b)). The method resulted in four performance clusters. The resulting clusters show that neither color nor shape palettes perform consistently well across all pairings; several palettes appear in both high- and low-performing clusters. Only one shape palette and no color palettes appeared entirely within any single cluster. This reinforces the interaction effect between color and shape observed in the ANOVA and further supports  $H_2$ : performance depends on how color and shape palettes are paired, not just the performance of the palettes in isolation.

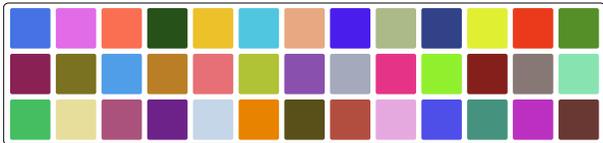
We further examined how color and shape palette attributes influence task performance in redundant pairings. Our findings show that the unfilled (shape palette No. 2) and open (shape palette No. 3) shapes resulted in lower accuracies when paired with color palettes characterized by either high lightness ( $L^*$ , palette No. 4–Carto/Pastel) or low lightness (palette No. 3–Stata/S2). In contrast, these shape palettes achieved higher accuracies when paired with color palettes exhibiting greater perceptual distances (palettes No. 1–ColorBrewer/Paired and 2–Tableau). Overall, mixed-shape-type (filled+unfilled+open) palettes (palettes No. 4, 5, and 6) achieved higher average accuracies, particularly in improving accuracies with lighter or darker color palettes (palettes No. 3–Stata/S2 and 4–Carto/Pastel).

## 5 Experiment Three: Pairwise Distance Matrix for Colors

Experiments 1 and 2 found that redundant encodings not only enhance task performance but that the specific color-shape palette pairings also influence accuracy. To make these findings actionable,

we aimed to develop an empirically-grounded model that generates effective redundant color-shape palettes based on the number of categories visualized. To achieve this goal, we conducted two separate experiments: one to build a pairwise accuracy matrix for colors (Experiment 3) to complement existing shape palette design measures and another to explore efficient color-shape pairings (Experiment 4). We separated these experiments because constructing a pairwise matrix for 39 colors and 39 shapes would involve over two million comparisons. Instead, this approach allows us to first identify effective color sets, which can be sampled for the next experiment. While this approach does potentially exclude useful shape-color combinations, it allows our sampling process to focus on empirically-validated color combinations, increasing the likelihood that the modeled space captures high-performing palette designs.

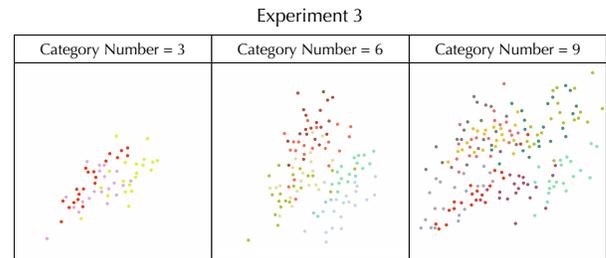
While past work has compared the effectiveness of different categorical palettes [26, 46, 78, 80], these studies focus on comparing entire palettes rather than specific pairwise color relations. Our objective is to collect data to model how specific color pairings influence categorical analysis. We hypothesized that: **H<sub>3</sub>: The relationships between colors will influence people’s abilities to compare correlations.** Given that various studies [26, 78, 80] have emphasized the importance of selecting appropriate colors for encoding categorical data in scatterplots, we expect that different color combinations will impact performance. Furthermore, different color combinations may be differently robust in varying numbers of categories [78].



**Figure 8: The 39 colors used in Experiment 3. These colors were generated using k-means clustering from the CIELAB color space. Each pair exceeds the just-noticeable difference (JND) value for the mark sizes used in our study.**

## 5.1 Experiment Design

**5.1.1 Task & Stimuli Generation.** The task, point distribution, and point number were generated using the same methodology as in Experiments 1 and 2. Following the approach in Tseng et al. [79], we focused on performance differences between shape pairs by constructing a distance matrix. Unlike shapes in a discrete space, colors exist in a continuous space. Past work has attempted to bin colors into a smaller and more manageable set of readily distinguishable, representative values. However, the suggested colors largely vary—from sets of eight [88] to thirteen [6] to twenty-two [28]. Further, these characteristic colors may fail to meet users’ needs and preferences for color aesthetics [63, 64], semantics [62], or other design factors [5]. Therefore, following prior work [26, 36], we derive our core colors by mathematically sampling across colorspace to ensure the diversity and representativeness of colors. To ensure sufficient and evenly distributed colors for our model, we employed k-means clustering in the CIELAB color space, as implemented by



**Figure 9: Example stimuli from Experiment 3, showing scatterplots with 3, 6, and 9 categories encoded using color only. This experiment was designed to construct pairwise accuracy matrices for 39 colors, capturing how effectively different color pairs support correlation estimation.**

iWantHue [47]. CIELAB is designed to be perceptually uniform, encompasses all perceivable colors, and has been widely used in past visualization research [15, 26, 69, 76, 83, 87].

We first sampled colors within the ranges:  $L^* = [25, 100]$ ,  $a^* = [-128, 127]$ , and  $b^* = [-128, 127]$ , using 5 steps for  $L^*$  and 2 steps for both  $a^*$  and  $b^*$ . Colors darker than  $L^* = 25$  were excluded, as the CIELAB space tends to contain more colors at lower lightness values, which can lead to an unbalanced sample, and design heuristics discourage using colors in this range for visualization [71]. This process yielded 38,734 valid color samples. We then applied k-means clustering to extract 200 representative colors. To ensure perceptual distinguishability between color pairs, we used d3-jnd [25, 72] to validate just-noticeable differences (JND) between each pair for our target mark size. For each of the 200 colors, we computed the number of valid JND pairings and selected the maximum subset in which all pairs were noticeable differently [26]. This step eliminated 163 colors, resulting in a refined set of 37 colors.

To further balance the set, we analyzed color name distribution using C3 [34]. While the overall distribution was balanced, we observed a lack of orange hues. To address this, we manually added two orange colors from C3 [34], bringing the final color set to 39. We also used d3-jnd [25, 72] to ensure each color is sufficiently perceptible against our tested white background. The resulting palette is shown in Figure 8, and example stimuli are shown in Figure 9(a).

We generated the tested color combinations based on the approach proposed by Tseng et al. [79] for shapes, which ensures a balanced distribution of pairwise samples. Using this method, we created 810 color combination sets (90 color combinations  $\times$  9 category counts). These sets were then organized into 15 task groups, with each group containing an equal distribution across category counts (6 color combination sets  $\times$  9 category counts = 54 stimuli per group).

**5.1.2 Procedure & Participants.** We recruited 167 participants through MTurk, excluding 17 who failed the engagement checks. The final analysis included data from 150 participants (107 male, 43 female; ages 21 to 65), with 10 participants assigned to each task group. The recruitment criteria and general procedure were consistent with those used in Experiments 1 and 2. The study

took approximately 10 minutes to complete and each participant received a compensation of \$1.60. This study followed the same procedure as Experiment 1 (see Section 3).

**5.1.3 Analysis.** Our primary dependent measure was pairwise accuracy, a metric employed in previous work [26, 79]. This metric was computed by evaluating the accuracy of each scatterplot containing a given pair of colors. For example, if a scatterplot having blue, orange, and green was answered correctly, the pairs blue–orange, blue–green, and orange–green each received one correct count, whereas an incorrect response contributed zero. Pairwise accuracy for each color pair was then calculated as the number of correct responses divided by the total number of trials in which that pair appeared. Because the correlation-estimation task requires participants to visually identify all categories and compare their correlations simultaneously, every category in the stimulus contributes to the overall judgment. Misattributing marks to the wrong category or failing to find a mark of a given category can change the perceived correlation [3]. For example, failing to find a cluster of outlying blues because it becomes masked by a cluster of green points will increase the perceived correlation of blue and may, as a result, change the relative perceived correlation of all three categories. This scoring approach therefore captures how well each color-shape pair supports category differentiation within a multi-class context, while avoiding the need for an impractically large number of isolated pairwise trials. We conducted a three-factor ANOVA to examine the effects of category number, color pair, and color attributes on accuracy. A summary of the results is presented in Table 3.

**Table 3: ANOVA results for category size and color pair. Significant effects are indicated by bold text.**

Experiment 3	DF <sub>1</sub>	DF <sub>2</sub>	F-value	p-value
Category size	2	148,497	2165.50	<b>&lt;.0001</b>
Color pair	740	147,759	1.22	<b>&lt;.0001</b>
category size * color pair	1480	147,019	0.76	1

**Table 4: ANOVA results for color attributes. Significant effects are indicated by bold text.**

Color Attribute	DF <sub>1</sub>	DF <sub>2</sub>	F-value	p-value
L* magnitude	1	148,498	119.79	<b>&lt;.0001</b>
C* magnitude	1	148,498	14.25	<b>0.0002</b>
h* magnitude	1	148,498	1.15	0.28

## 5.2 Results

As shown in Table 3, our analysis revealed significant effects of both category size ( $F(2, 148497) = 2165.5, p < .0001$ ) and specific color pairs ( $F(740, 147759) = 1.22, p < .0001$ ) on performance in the correlation comparison task. The overall accuracy was 78.14%. Accuracy varied across category sizes, with small ( $N = 2, 3, 4$ ), medium ( $N = 5, 6, 7$ ), and large ( $N = 8, 9, 10$ ) categories having accuracies of 91.59%, 78.33%, and 64.48%, respectively, reflecting a decrease in

accuracy as the number of categories increased. We summarized these results using heatmaps presented on an interactive website<sup>2</sup>, allowing detailed exploration of accuracy differences across color pairs. The results show *Perceptual Distance* does not significantly impact performance; however, this is very likely due to the fact that we already optimized it in the stimuli generation by ensuring color differences exceeded necessary JNDs. Above a certain threshold, the differences between colors are large enough that making them larger does not meaningfully improve discriminability [34].

Overall, the pairwise accuracy ranged from 56.84% to 80.5%, with a standard deviation of 3.56%. Our results support **H<sub>3</sub>**: color pairs impact the task performance. Examples of best-performing color palettes are shown in the OSF Supplements. We further examined the relationship between various color attributes and pairwise accuracy using CIELCh [90], which is commonly used for color palette design [26, 71] and comprises a polar representation of the CIELAB colorspace. For all 741 color pairs, we computed the  $L^*$ ,  $C^*$ , and  $h^*$  (Lightness, Chroma, Hue) magnitude/difference, as well as the name difference [34], composite difference (perceptual distance [67]) and decomposed ( $L^*$ ,  $a^*$ ,  $b^*$ ) from CIELAB. Table 4 presents the exploratory ANOVA results, which indicate that the magnitudes in  $L^*$  and  $C^*$  significantly affect task performance, whereas perceptual distance;  $L^*$ ,  $a^*$ ,  $b^*$ , and  $h^*$  differences; and name differences appear to be less influential (shown in OSF Supplements), likely due to all selected colors being sufficiently perceptually distinct by design. These findings are consistent with previous work [78], suggesting that lightness and chroma magnitudes are critical for categorical perception.

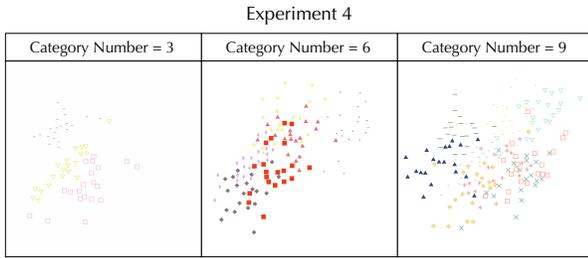
## 6 Experiment Four: Redundant Encoding Matrices

Experiment 2 revealed that the pairing of color and shape palettes significantly affects task performance. Therefore, rather than simply combining the best-performing color and shape palettes, effective redundant encodings must be designed with redundancy in mind. We designed Experiment 4 to explicitly measure the effectiveness of different color–shape pairings. We hypothesized the following: **H<sub>4</sub>: (a) How colors and shapes are paired will influence correlation comparison accuracy, and (b) combining the best-performing color pairs and shape pairs does not necessarily produce optimal encodings.** As Experiment 2 highlighted the importance of pairing at the palette level, we extend this assumption to the individual marker level (e.g., blue triangle vs. blue square). Consequently, although we selected high-performing palettes, we do not expect that random combinations of color and shape elements from those palettes will automatically result in optimal redundant encodings.

### 6.1 Experiment Design

**6.1.1 Task & Stimuli Generation.** We adopted the same task design, point distribution, and point count as in the previous three experiments. Building on the pairwise accuracy matrices obtained in Experiment 3, we developed a palette selection model that takes category number ( $N$ ) as input and outputs high-performing color sets from the pool of 39 colors from Experiment 3. The model operates

<sup>2</sup><https://catpaw-categorical-palette.web.app/pair.html?group=all>



**Figure 10: Example stimuli from Experiment 4, showing scatterplots with 3, 6, and 9 categories encoded using redundant color-shape combinations, where each category is represented by a unique pairing of a color and a geometric shape. This experiment was used to construct pairwise accuracy matrices for all 39 colors  $\times$  39 shapes to assess how combined encodings influence perceptual discrimination.**

by randomly selecting half of the color pool, generating all possible combinations for the given  $N$ , computing the sum of pairwise accuracies for each combination, and identifying the combinations with the highest average pairwise accuracy. This process is repeated ten times to ensure stability. For each category number ( $N = 2-10$ ), we selected the top three color palettes using this approach, and paired them with the top three shape palettes generated using ShapeltUp [79], resulting in 9 color-shape palette combinations per  $N$ .

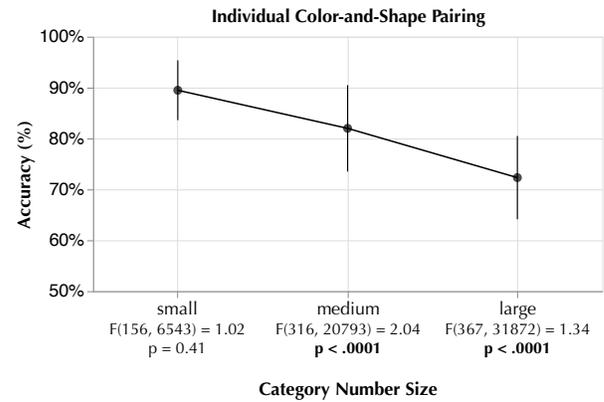
To expand the range of tested color-shape pairings, we designed a pairing strategy rather than relying on random assignment. Exhaustively testing all permutations of color-shape matches is computationally infeasible for higher category counts (e.g.,  $N = 5$  yields  $5! = 120$  possible permutations). For  $N = 2$  and  $N = 3$ , we included all 2 and 6 permutations, respectively. For  $N = 4-10$ , we selected 13 representative permutations for each color-shape palette pair using a greedy algorithm which selects one palette with the lowest cosine similarity compared to selected ones at a time, ensuring a diverse sampling of redundant encodings. This results in 99 ( $2+6+7 \times 13$ ) redundant palettes  $\times$  9 (3 color palettes  $\times$  3 shape palettes) palette combinations, resulting in 891 stimulus designs. We divided these 891 stimuli into 15 task groups, each containing 60 stimuli with an approximately uniform distribution of category number and source palette. Example stimuli are shown in Figure 10(b).

**6.1.2 Procedure & Participants.** A total of 170 participants were recruited via MTurk, with 20 excluded for not passing the engagement checks. This left 150 participants (115 male, 35 female; aged between 28 and 65) whose data were included in the final analysis, with 10 participants assigned to each task group. The recruitment criteria and overall procedure matched those used in Experiments 1, 2, and 3. On average, the study took 10 minutes to complete, and participants were compensated \$2.90 for their time. This study followed the same procedure as Experiment 1 (see Section 3).

**6.1.3 Analysis.** We used task accuracy as the primary dependent variable and conducted two main analyses: (1) individual color-shape pairing accuracy (e.g., yellow triangle, blue cross) and (2) pairwise marker accuracy (e.g., red circle + green square). For the individual color-shape pairing analysis, we calculated accuracy by

evaluating tasks in which a specific color-shape marker appeared. For instance, if a scatterplot containing red circles, green squares, and orange diamonds was answered correctly, each of those three markers received one correct count. Accuracy for each color-shape pairing was then computed as the ratio of correct responses to the total number of appearances.

For pairwise accuracy, we adopted a similar approach but focused on specific pairs of color-shape markers. Accuracy was calculated based on the frequency of correct responses when a particular pair appeared together in a task. Both analyses were stratified by category size, grouped into three bins based on the number of categories—small ( $N = 2-4$ ), medium ( $N = 5-7$ ), and large ( $N = 8-10$ )—to enable sufficient sample sizes across conditions. We conducted an ANOVA to examine the effect of individual color-shape pairings on correlation comparison task performance. Additionally, to investigate whether simply combining top-performing color pairs and top-performing shape pairs leads to effective redundant encodings, we performed a secondary ANOVA to compare pairwise accuracies of both colors and shapes against their joint combinations.



**Figure 11: Results from Experiment 4. Average accuracy for individual color-shape pairings across three category size groups (small, medium, and large). Error bars represent standard deviation. ANOVA results indicate significant effects of individual color-shape pairings on task accuracy for medium and large category sizes. These indicate how color-shape paired will impact the task accuracy, particularly with category numbers larger than four.**

## 6.2 Results

As shown in Figure 11, average accuracy decreases as category size grows, with an overall accuracy of 79.9%. Our analysis revealed that the way in which color and shape are paired to form a single marker has a significant effect on task performance for medium ( $F(316, 20793) = 2.04$ ,  $p < .0001$ ) and large ( $F(367, 31872) = 1.34$ ,  $p < .0001$ ) category sizes. However, no significant effect was found for small category sizes ( $N = 2-4$ ), with  $F(156, 6543) = 1.02$ ,  $p = 0.41$ . Figure 11 also presents the standard deviations of individual color-shape pairings across category sizes, illustrating greater variance in accuracy for medium and large categories compared to

smaller ones. These findings partially support  $H_4$  (a): how colors and shapes are paired influences the accuracy, though this effect appears less prominent for small category sizes. Examples of best-performing color-shape palettes are shown in the OSF Supplements.

To further investigate the importance of mapping choices, we examined the relationship between pairwise marker accuracy in this experiment and the color/shape pairwise accuracies derived from previous models. For each marker pair (e.g., red circle and green square), we computed its task accuracy and retrieved corresponding color pair (e.g., red + green) and shape pair (e.g., circle + square) accuracies from previous matrices. ANOVA results indicate that combining top-performing color pairs and shape pairs simultaneously ( $F(1, 4221) = 36, p < .0001$ ) produce effective redundant encodings only within medium category sizes ( $N = 5 - 7$ ), where redundant encodings are most effective. Across all category sizes, pairwise shape accuracy only showed a moderate effect for large category sizes ( $F(1, 10053) = 4.18, p = 0.04$ ), and we did not find any significant effects of pairwise color accuracy. These findings partially support  $H_4$  (b): combining best-performing color and shape pairs does not always lead to optimal encodings, though such combinations appear to be more effective in the 5–7 category range.

## 7 CatPAW: Categorical Palette Automation Wizard

### 7.1 Comprehensive Categorical Encoding Model

To translate our experimental findings into an actionable framework, we developed a statistical model to evaluate the effectiveness of categorical encodings using shape and color. This model is grounded in a rich set of empirical data, including four pairwise accuracy matrices for 39 colors derived from 14,850 trials in Experiment 3, 839 individual color–shape pairings, and 14,935 pairwise color–shape accuracy samples from Experiment 4. Additionally, we incorporated pairwise accuracy matrices for 39 shapes from Tseng et al. [79] and key findings from Experiment 1 (highlighting the effectiveness of redundant encoding within a moderate category range) and Experiment 2 (showing that mixed shape types generally perform better).

Given a specified category number and encoding type (color-only, shape-only, or color–and–shape), the model operates in the following way: For color-only or shape-only palettes, the model first samples a manageable subset of candidate palettes from the corresponding set of 39 shapes or colors and computes their overall pairwise accuracy based on the selected category number. It then returns the top-performing palette(s). For redundant (color–and–shape) encodings, the model randomly samples combinations of color and shape palettes from each of the 39 color/shape pools and generates a tractable number of color–shape pairing combinations. Each pairing set is evaluated using individual color–shape pairs and pairwise-level accuracy results from Experiment 4. The top-ranked combinations are further assessed based on their color-only and shape-only pairwise accuracy scores, variance in lightness, and variance in shape types. The final output is a ranked list of color–shape paired palettes, each with corresponding performance scores.

### 7.2 Design Tool for Categorical Palettes

We implemented our model in a web-based categorical palette recommendation tool (see CatPAW<sup>3</sup>, as shown in Figure 12). While most visualization tools provide default color or shape palettes for encoding categorical data, they often lack guidance on when to use color-only, shape-only, or redundant (color–and–shape) encodings. Moreover, users may wish to incorporate personal design preferences, such as specific colors or shapes, without sacrificing perceptual effectiveness. Our findings emphasize that choosing appropriate encodings depends on the number of categories and the interaction between visual channels. To support designers in reasoning across these constraints, CatPAW offers a comprehensive palette suggestion framework that helps designers generate effective, customizable palettes backed by empirical data.

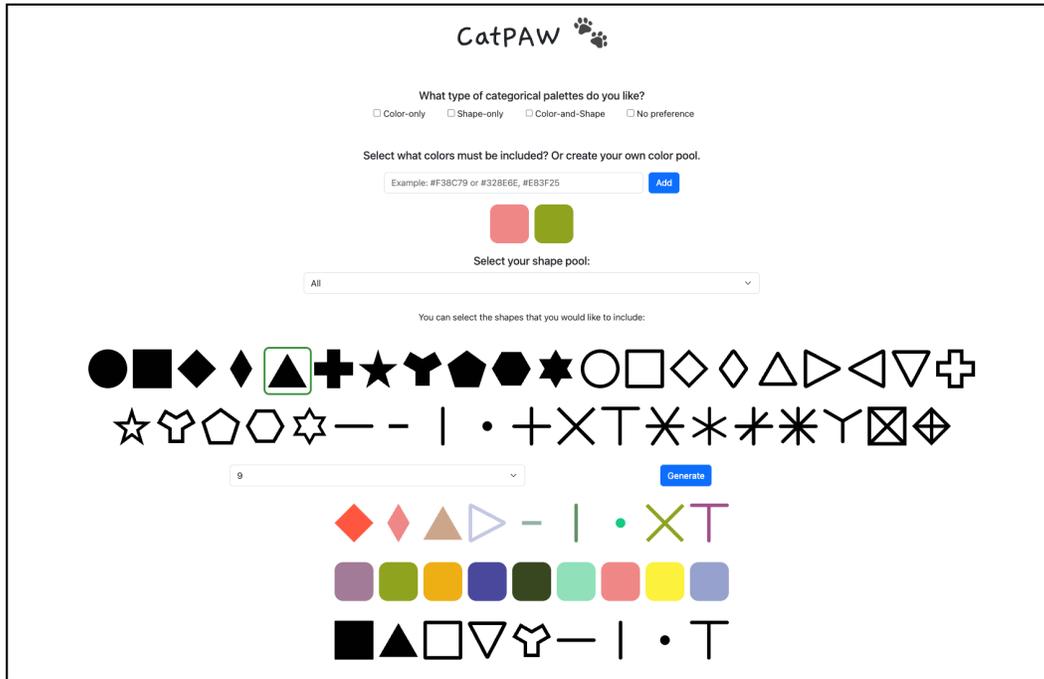
CatPAW allows users to specify what type of categorical palettes they prefer (color, shape, or redundant), and any desired colors, shapes, or color–shape pairs. Based on the selected category number and user constraints, the tool recommends optimal palette configurations. We leverage the fact that color can be expressed in a continuous space to extend the generated palettes to a wider range of colors. We expect that colors within a given distance of a tested color will have comparable performance. Therefore, rather than directly selecting subsets from the original 39 colors, the tool computes palettes using these colors and then randomly samples nearby colors that are within a certain perceptual distance ( $\Delta \leq 30$ ) to the original color to increase design flexibility. We note that this approach provides an approximation of likely performance across the colorspace; however, future studies should derive a more granular model through additional data collection studies. Similarly, when users input custom colors, the tool maps them to the closest representative color and proceeds with the palette generation accordingly. CatPAW provides three core features:

1. *Support for three categorical palette types.* Users can select from color-only, shape-only, or color–and–shape categorical palettes or allow the tool to recommend the most effective encoding strategy based on the provided constraints.
2. *Data-driven palette suggestions based on category number.* The underlying model draws on empirical results from 455 participants across four experiments. It generates optimized palettes based on category number and adapts to the user’s design preferences.
3. *Customizable palettes with user-defined constraints.* Users can specify pools of colors or shapes or define required elements that must appear in the output. After palette generation, users can also remove unwanted elements, prompting the tool to substitute them with the next best-performing alternatives.

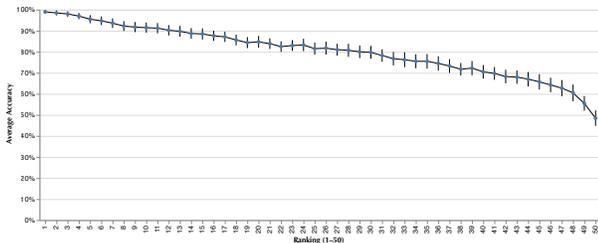
### 7.3 Cross-Measure Validation using the Categorical Encoding Model

Our model supports the unified design of color, shape, and redundant palettes by estimating their performance using three pairwise accuracy matrices. Given a set of candidate palettes, the model assigns a score to each palette and generates a ranking based on the predicted effectiveness. To evaluate the model’s ability to predict performance rankings for redundant palettes, we follow the

<sup>3</sup><https://catpaw-categorical-palette.web.app/>



**Figure 12:** The *CatPAW* interface. Users can specify their preferred palette type (color-only, shape-only, color-and-shape, or no preference), select required colors and/or shapes, and enter the desired number of categories. *CatPAW* then generates candidate palettes, which users can iteratively refine by swapping out unwanted elements with the next ranked option.



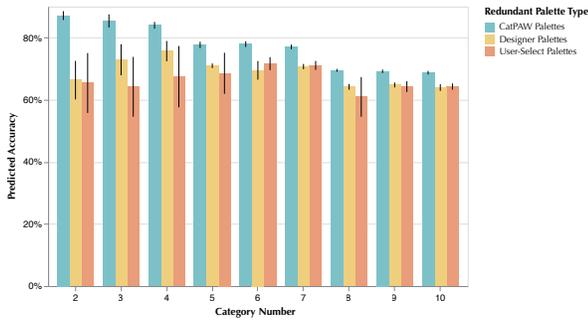
**Figure 13:** Cross-measure validation comparing the model-predicted ranking of redundant categorical palettes with human accuracy from Experiment 4. Our model predicts a ranking (1–50) for redundant palettes within each category number, where a higher rank indicates better expected performance. For validation, we randomly sampled 50 redundant palettes for each category number and repeated this process three times. The plot shows, for each ranking position, the mean human accuracy from our study along with 95% confidence intervals, aggregated across all palettes assigned the same predicted rank. The monotonic decrease in performance provides preliminary evidence that *CatPAW*’s ranking model aligns with empirical performance.

approach used in Tseng et al. [79] and perform a preliminary cross-measure validation comparing model predictions with ground-truth accuracy obtained in Experiment 4.

In Experiment 4, approximately 120 unique redundant palettes were tested for each category number, with each palette receiving 10 responses. For validation, we randomly selected 50 palettes per category number, repeated this sampling three times, and used our model to generate predicted rankings for each set. We then computed the actual mean human accuracy for palettes at each predicted rank. As shown in Figure 13, there is a clear monotonic downward trend: palettes predicted to have lower within-category rankings consistently exhibited lower human accuracy across all category numbers. The correlation between predicted rank and mean human performance was strong ( $r = 0.97, p < .0001$ ), indicating that the model closely aligns with empirical behavior. These findings offer preliminary evidence that our model can faithfully predict relative performance across redundant palettes. However, this validation focuses on ranked comparisons within our dataset; more extensive evaluation across broader datasets and task contexts remains an important direction for future work.

#### 7.4 Evaluation of *CatPAW* against Designer and User-Select Categorical Palettes

To the best of our knowledge, prior work offers no established guidance or design principles for constructing redundant color–shape palettes, despite their increasing use in visualization practice. As a result, there is no existing benchmark or standard which can be used to assess *CatPAW*’s recommendations. To provide a preliminary assessment of *CatPAW*’s palettes, we instead use the categorical encoding model to compare *CatPAW*-generated palettes



**Figure 14: We compared the model-predicted performance of three groups of redundant categorical palettes across category numbers (2–10): *CatPAW*-generated palettes, designer palettes, and user-selected palettes. Designer palettes were derived from widely used visualization systems (D3, Excel, MATLAB, and Tableau); for each system, we paired its color palette with its matching default shape palette and randomly combined these elements to create five redundant palettes per category number. For comparison, we sampled ten redundant palettes generated by *CatPAW* and ten palettes selected directly by users for each category size. Error bars show 95% confidence intervals of the predicted accuracy.**

against two baselines: widely-used palettes from designer tools and palettes that users create on their own. Although *CatPAW* is trained on this model, the goal of this analysis is not to claim a definitive performance advantage, but rather to characterize how *CatPAW*'s data-driven recommendations differ from conventional design choices and to provide initial evidence of its potential utility.

For the designer baseline, we selected one categorical color palette and one shape palette from each of D3, Excel, MATLAB, and Tableau. Because none of these systems provides guidance for how colors and shapes should be combined into redundant palettes, we randomly paired colors and shapes from each tool to construct five redundant palettes for each category number. To construct the user-selected palette baseline, we recruited ten participants from a local college with self-reported experience creating visualizations for professional publication, who reflect representative target users for *CatPAW*. Each participant was asked to design one redundant palette for each category number (2–10) that they believed would perform best for distinguishing categories. Participants had access to the pool of 39 shapes used in *CatPAW* and were free to choose any colors they wished, either manually or by drawing directly from designer palettes. Examples of the palettes from three groups are shown in Figure 15. We then applied our categorical encoding model to generate predicted accuracy scores for every palette based on empirical data and computed the average accuracy for each palette group. As shown in Figure 14, *CatPAW*-generated palettes (77.5% mean predicted accuracy) outperformed both designer palettes (68.4%) and user-selected palettes (66.2%) across category sizes, demonstrating the potential value of data-driven palette construction. The performance gap is most pronounced for smaller category numbers (2–6), where *CatPAW*'s predictions are consistently higher than both baselines. Designer palettes also

outperformed user-selected palettes in these smaller category sizes, though the difference between them narrows as the number of categories increases. Overall, this analysis suggests that *CatPAW* produces redundant palettes that are predicted to support correlation judgments more effectively than those created through traditional design tools or intuitive user choices. These findings provide initial evidence that empirical, model-driven palette generation may offer meaningful advantages.

However, this evaluation should be viewed as an early step toward understanding the comparative performance of *CatPAW*. A more definitive assessment of *CatPAW*'s utility will require a large-scale empirical study that directly measures task performance using *CatPAW*-generated palettes, designer palettes, and user-created palettes with human participants. We see this as an important direction for future work.

## 8 Discussion

Despite being widely used in real-world applications, the role of redundancy in visualization remains contentious, with limited experimental data to demonstrate its utility and no empirical data guiding effective redundant encoding design. Our results confirm prior work demonstrating the benefits of redundant encoding for improving accuracy in visual tasks [51]. Contrary to previous insights on smaller class numbers (2-3) that redundant encoding does not increase performance in multiclass scatterplots [23], we observed that combining color and shape information enhances the accuracy of correlation judgments. However, these benefits depend on the number of categories encoded. We found that redundancy offered the greatest perceptual benefit for between 5 and 8 categories. For larger numbers of categories, the benefits of redundancy may begin to diminish as the task becomes either simply too difficult or the display too complex. This sensitivity to the number of categories may also explain the lack of benefit found in prior studies [23]: scatterplots with two to three categories did not demonstrate any significant benefit from redundancy. For smaller category numbers, people may be better able to leverage the limited capacity of working memory to reason over categories [11, 81], allowing them to attend to each serially and readily process relevant information about each.

The efficacy of redundant designs depends not only on choosing effective color or shape palettes, but also in considering the mapping between color and shape. We found significant interactions between color and shape choices, which resonate with prior work on the limited separability of these visual channels [68, 73]. For instance, unfilled and open shapes paired with high or low lightness color palettes tended to yield lower accuracy. The interplay between color and shape for redundant encoding highlights the importance of a more holistic orientation to visualization design, considering not only visual channels in isolation but also how these channels (and encodings more generally) might be designed to complement one another. Our experiments unify past work that individually focused on color encodings [80] or shape encodings [79] to offer a more holistic approach to design. Our research provides further evidence for these preliminary insights, offers specific guidance on how redundant encodings can mitigate this effect, and results in a universal tool to combine the benefits of both.

	CatPAW-Generated Palettes	Designer Palettes	User-Select Palettes
Category Num = 3			
Category Num = 6			
Category Num = 9			

Figure 15: Example redundant categorical palettes used in the model-based evaluation of palette quality. The figure shows three groups of palettes for category sizes 3, 6, and 9: CatPAW-generated palettes, designer palettes, and user-selected palettes.

### 8.1 Redundancy and Perceptual Processes

To help bridge our findings with work in vision science and to understand the potential generalizability of our findings, we hypothesize several potential underlying perceptual mechanisms that may be at play in our study. By examining our results through the lens of these mechanisms, we offer a set of potential explanations that may explain both why redundant color-shape encodings improve accuracy in categorical perception—particularly within the 5–8 category range—and why specific shape-color pairings matter.

The cognitive load associated with visualization viewing (i.e., the demand the task places on visual working memory) likely plays an important role in performance. Past work has demonstrated that when visualizing fewer than six categories (termed as the *Subitizing* limit), viewers can efficiently process and compare categories using a single channel, such as color or shape, without overloading cognitive resources [81]. However, as the number of categories increases, the discriminability of categories within a single visual encoding channel diminishes, leading to significantly decreased performance. Our results demonstrated that redundant encoding mitigates the reduced performance at the subitizing limit by providing a second perceptual cue, effectively reducing the cognitive effort required to distinguish between categories. Furthermore, the gains associated with redundancy are highest near the *Subitizing* limit [81] (in this case, for 5-8 categories). This indicates that the gains from redundant encoding are highest when less efficient perceptual mechanisms that place more demand on visual working memory are first employed. As a result, visual working memory may more efficiently process redundant encodings despite their increased complexity. Beyond approximately 8 categories, however, the benefits of redundancy may plateau or decline due to visual clutter or interference between channels, consistent with known theories in visual working memory [11].

Second, we observed there are potential interaction effects between color and shape, indicating that not all pairings are equally effective. For instance, unfilled or open shapes paired with very light or very dark color palettes led to lower accuracy. This may

be due to reduced perceptual salience: unfilled shapes have less colored area, which may weaken the effectiveness of color encoding [68]. Conversely, mixed-shape palettes (combining filled, unfilled, and open shapes) performed better across a wider range of color palettes, suggesting that variability in shape type enhances overall discriminability. This supports the idea that effective redundant encoding relies on complementary channel strengths rather than simply doubling the same type of perceptual information [51].

Third, compared to hue, the significant impact of lightness ( $L^*$ ) and chroma ( $C^*$ ) magnitude in driving pairwise color accuracy suggests that they may be key factors in facilitating rapid visual segmentation between categories. This is consistent with theories and evidence in early visual processing mechanisms that prioritize luminance and contrast over hue [86]. For example, evidence from magnetoencephalography (MEG) data shows that hue reaches a slower peak decoding, about 20 ms after luminance polarity, in visual systems [35]. Therefore, when combined with shape, these attributes may either reinforce or undermine the perceptual integrity of each mark together with the shape’s visual properties (e.g., filled vs. outline), for example, lightness and chroma may drive initial segmentation, and then shape and hue help refine that segmentation.

We note that these hypotheses should be investigated in more controlled settings and using methods tailored to vision science-based inquiry. We offer them as a preliminary means for more broadly and systematically understanding redundancy and categorical encoding and for prompting deeper interdisciplinary inquiry.

### 8.2 Design Implications

We embed the results of our experiments into a system to support designers in creating effective categorical encodings. However, our results also raise several general guidelines for using redundancy effectively:

- **Prioritize Redundant Encodings for Moderate Categories** Utilizing redundant encodings unifying color and

shape generally improves the accuracy of correlation estimation in multi-class scatterplots. This approach is particularly beneficial when dealing with a moderate number of categories (5-8).

- **Consider Color-Shape Pairing Strategy Before Using It** Designers should carefully consider how they pair colors and shapes, especially when creating scatterplots with more than 4 categories. Certain combinations can significantly impact performance. For example, our findings suggest that mixed-shape type palettes (i.e., those employing a combination of filled, unfilled, and open shapes) tend to work well with both light and dark color palettes.
- **Be Aware of the Influence of Color Attributes** Color attributes such as lightness and perceptual distance influence the effectiveness of color-shape pairings. Avoid pairing unfilled or open shapes with color palettes that have high or low lightness.
- **Balance Palette Complexity** While redundant encoding can improve performance, it is important to avoid over-complicating a visualization by introducing unnecessary visual features or over-consuming limited visual channels. Our results suggest that applying redundant encoding in scatterplots with more than 4 categories can improve overall task accuracy. However, the marginal benefits of redundancy appear to plateau as the number of categories increases beyond 8.
- **Prioritize Top-performing Color and Shape Pairs for Moderate Categories** While simply combining top-performing colors and shapes does not guarantee optimal results, designers should consider matching color and shape pairs with better pairwise performance—particularly in scatterplots with a moderate number of categories (around 5-8).

### 8.3 Limitations and Future Work

While our study provides insight into how we can effectively design categorical visualization with redundant encodings, there are several directions for future work that will improve our understanding of redundant encoding and categorical visualization broadly. First, we assess only a subset of colors and shapes and only on a white background, which is among the most common background choices for modern visualizations. We employ a number of sampling strategies to maintain a broad yet tractable sampling of redundant palettes. Future work could explore a broader range of colors, shapes, category numbers, and background colors, as well as scatterplots with heavily overlapping configurations, to assess the generalizability of our findings. While correlation is a common task for multivariate scatterplots [59], understanding performance across other tasks, such as outlier detection or cluster identification, and visualization idioms like bar charts and icon arrays, is also important future work to assess the generalizability of our findings. Such work should also investigate associations between specific shape and color attributes, such as hue variation or shape features, and performance to provide a more nuanced understanding of color-shape interactions.

Our study focuses purely on abstract palette design. Other factors beyond simply color, shape, and category number may influence palette effectiveness. For example, our study did not explicitly account for participant preferences in colors and shapes, which may also impact their responses [64]. We also focused on palettes for abstract, neutral values; however, both shapes and colors can carry semantic meaning that may influence their effectiveness in isolation as well as in redundant pairs [61]. Future research could investigate how these factors influence the effectiveness of redundant encodings and categorical palette design. Further, redundancy can also be applied with ordinal and numeric data. Our findings offer hypotheses about general redundant encodings, but future work should confirm whether our observations generalize to quantitative redundant encodings.

*CatPAW* represents a step towards automated palette generation by considering redundant encodings. While *CatPAW* focuses on leveraging the third dimension for categorical data, it does not explore using colors and shapes to encode different data dimensions simultaneously nor can it currently adapt to new shapes. Future work could focus on refining the underlying model and incorporating more sophisticated optimization techniques and account for accessibility to create even more effective and user-friendly design tools. Further, future work should evaluate *CatPAW*'s performance across a broader range of tasks, factors such as aesthetic preference, and in more formal human subjects experiments. As *CatPAW* is an extensible open-source tool, future results, such as new shapes or colors, can be readily integrated into the underlying data model to generate a wider range of improved predictions.

## 9 Conclusion

We investigated the effectiveness of redundant encoding by systematically evaluating how color and shape pairings influence categorical perception in data visualization. Our findings demonstrated that while redundant encodings can enhance task performance in multiclass scatterplots, their benefits are most evident within a certain range of category numbers (5-8). Furthermore, we highlight the importance of combining colors and shapes, and provide a set of guidelines for designing effective redundant categorical palettes. To translate these insights into an actionable framework, we developed a data-driven model informed by experimental results and implemented it in a web-based tool that enables designers to create categorical palettes with flexible user preferences. We hope this work will encourage future research toward developing comprehensive strategies for categorical data visualization.

## Acknowledgments

We thank the reviewers for their insightful comments. This work was supported by the National Science Foundation under grant No.2127309 to the Computing Research Association for the CIFEL-lows project, NSF IIS-2046725, and NSF IIS-1764089.

## References

- [1] Sven Bachthaler and Daniel Weiskopf. 2008. Continuous scatterplots. *IEEE transactions on visualization and computer graphics* 14, 6 (2008), 1428–1435.
- [2] Emily Badger. 2018. What's the right number of taxis (or Uber or Lyft cars) in a city. *New York Times* (2018).
- [3] Caroline Barras and Dirk Kerzel. 2017. Target-nontarget similarity decreases search efficiency and increases stimulus-driven control in visual search. *Attention, Perception, & Psychophysics* 79, 7 (2017), 2037–2043.
- [4] Beth A Barstow, Jason Vice, Sean Bowman, Tapan Mehta, Seanna Kringen, Peter Axelson, and Sangeetha Padalabalanarayanan. 2019. Examining perceptions of existing and newly created accessibility symbols. *Disability and Health Journal* 12, 2 (2019), 180–186.
- [5] 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*. 1364–1374.
- [6] Brent Berlin and Paul Kay. 1991. *Basic color terms: Their universality and evolution*. Univ of California Press.
- [7] Rita Borgo, Luana Micaleff, Benjamin Bach, F McGee, and B Lee. 2018. Information Visualization Evaluation Using Crowdsourcing. *Comput. Graph. Forum* 37, 3 (2018), 573–595. doi:10.1111/cgf.13444
- [8] Michelle A Borkin, Azalea A Vo, Zoya Bylinskii, Phillip Isola, Shashank Sunkavalli, Aude Oliva, and Hanspeter Pfister. 2013. What makes a visualization memorable? *IEEE Trans. Vis. Comput. Graph.* 19, 12 (2013), 2306–2315.
- [9] David Borland and Russell M Taylor II. 2007. Rainbow color map (still) considered harmful. *IEEE Comput. Graph. Appl.* 27, 2 (2007), 14–17.
- [10] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. 2011. D<sup>3</sup> Data-Driven Documents. *IEEE Trans. Vis. Comput. Graph.* 17, 12 (2011), 2301–2309.
- [11] Timothy F Brady and Joshua B Tenenbaum. 2013. A probabilistic model of visual working memory: Incorporating higher order regularities into working memory capacity estimates. *Psychological review* 120, 1 (2013), 85.
- [12] David Burlinson, Kalpathi Subramanian, and Paula Goolkasian. 2017. Open vs. closed shapes: New perceptual categories? *IEEE Trans. Vis. Comput. Graph.* 24, 1 (2017), 574–583. doi:10.1109/TVCG.2017.2745086
- [13] Carto. 2022. Location Intelligence & GIS for cloud natives. <https://carto.com/>
- [14] Sarun Charumilind, Matt Craven, Jessica Lamb, Adam Sabow, Shubham Singhal, and Matt Wilson. 2022. When will the COVID-19 pandemic end? *McKinsey & Company*: <https://www.mckinsey.com/industries/healthcare/our-insights/when-will-the-covid-19-pandemic-end> (2022).
- [15] Haidong Chen, Wei Chen, Honghui Mei, Zhiqi Liu, Kun Zhou, Weifeng Chen, Wentao Gu, and Kwan-Liu Ma. 2014. Visual abstraction and exploration of multi-class scatterplots. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1683–1692.
- [16] William Cleveland and Robert McGill. 1984. Graphical perception: Theory, experimentation, and application to the development of graphical methods. *J. Amer. Statist. Assoc.* 79, 387 (1984), 531–554.
- [17] Michael Correll and Jeffrey Heer. 2017. Regression by eye: Estimating trends in bivariate visualizations. In *Proc. ACM Hum. Factors Comput. Syst. (CHI)*. 1387–1396. doi:10.1145/3025453.3025922
- [18] Çağatay Demiralp, Michael S Bernstein, and Jeffrey Heer. 2014. Learning perceptual kernels for visualization design. *IEEE Trans. Vis. Comput. Graph.* 20, 12 (2014), 1933–1942. doi:10.1109/TVCG.2014.2346978
- [19] Madison A Elliott, Christine Nothelfer, Cindy Xiong, and Danielle Albers Szafir. 2020. A Design Space of Vision Science Methods for Visualization Research. *IEEE Trans. Vis. Comput. Graph.* (2020). doi:10.1109/TVCG.2020.3029413
- [20] Geoffrey Ellis and Alan Dix. 2007. A taxonomy of clutter reduction for information visualisation. *IEEE Trans. Vis. Comput. Graph.* 13, 6 (2007), 1216–1223.
- [21] Steven L Franconeri, Lace M Padilla, Priti Shah, Jeffrey M Zacks, and Jessica Hullman. 2021. The science of visual data communication: What works. *Psychological Science in the Public Interest* 22, 3 (2021), 110–161.
- [22] Michael Gleicher, Danielle Albers, Rick Walker, Ilir Jusufi, Charles D Hansen, and Jonathan C Roberts. 2011. Visual comparison for information visualization. *Information Visualization* 10, 4 (2011), 289–309.
- [23] Michael Gleicher, Michael Correll, Christine Nothelfer, and Steven Franconeri. 2013. Perception of average value in multiclass scatterplots. *IEEE Trans. Vis. Comput. Graph.* 19 (2013). doi:10.1109/TVCG.2013.183
- [24] Robert L Goldstone and Andrew T Hendrickson. 2010. Categorical perception. *Wiley Interdiscip. Rev.: Cogn. Sci.* 1, 1 (2010), 69–78. doi:10.1002/wcs.26
- [25] Connor Gramazio. 2019. d3-jnd: Perceptual color difference tool for D3.js. <https://github.com/connorgr/d3-jnd>. GitHub repository. Accessed: 2025-03-30.
- [26] Connor C Gramazio, David H Laidlaw, and Karen B Schloss. 2016. Colorogical: Creating discriminable and preferable color palettes for information visualization. *IEEE Trans. Vis. Comput. Graph.* 23, 1 (2016), 521–530. doi:10.1109/TVCG.2016.2598918
- [27] Maxene Graze and Jonathan Schwabish. 2024. Building color palettes in your data visualization style guides. *Journal of the American Medical Informatics Association* 31, 2 (2024), 488–498.
- [28] Paul Green-Armytage. 2010. A colour alphabet and the limits of colour coding. *JAIC—Journal of the International Colour Association* 5 (2010).
- [29] Tejas Guha, Elana J Fertig, and Atul Deshpande. 2022. Generating colorblind-friendly scatter plots for single-cell data. *Elife* 11 (2022), e82128.
- [30] Charles R. Harris, K. Jarrod Millman, Stéfan J van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. 2020. Array programming with NumPy. *Nature* 585 (2020), 357–362. doi:10.1038/s41586-020-2649-2
- [31] Lane Harrison, Fumeng Yang, Steven Franconeri, and Remco Chang. 2014. Ranking visualizations of correlation using weber's law. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 1943–1952. doi:10.1109/TVCG.2014.2346979
- [32] Mark Harrower and Cynthia A Brewer. 2003. ColorBrewer.org: an online tool for selecting colour schemes for maps. *Cartogr. J.* 40, 1 (2003), 27–37. doi:10.1179/000870403235002042
- [33] Jeffrey Heer and Michael Bostock. 2010. Crowdsourcing graphical perception: using mechanical turk to assess visualization design. In *Proc. ACM Hum. Factors Comput. Syst. (CHI)*. doi:10.1145/1753326.1753357
- [34] Jeffrey Heer and Maureen Stone. 2012. Color naming models for color selection, image editing and palette design. In *Proc. ACM Hum. Factors Comput. Syst. (CHI)*. 1007–1016. doi:10.1145/2207676.2208547
- [35] Katherine L Hermann, Shridhar R Singh, Isabelle A Rosenthal, Dimitrios Pantazis, and Bevil R Conway. 2022. Temporal dynamics of the neural representation of hue and luminance polarity. *Nature communications* 13, 1 (2022), 661.
- [36] Matt-Heun Hong, Zachary Nolan Sunberg, and Danielle Albers Szafir. 2024. Cieran: Designing Sequential Colormaps via In-Situ Active Preference Learning. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [37] Matt-Heun Hong, Jessica K Witt, and Danielle Albers Szafir. 2021. The Weighted Average Illusion: Biases in Perceived Mean Position in Scatterplots. *IEEE Trans. Vis. Comput. Graph.* 28, 1 (2021), 987–997. doi:10.1109/tvcg.2021.3114783
- [38] Liqiang Huang. 2020. Space of preattentive shape features. *Journal of vision* 20, 4 (2020), 10–10.
- [39] The MathWorks Inc. 2022. *MATLAB version: 9.13.0 (R2022b)*. Natick, Massachusetts, United States. <https://www.mathworks.com>
- [40] International Organization for Standardization. 2024. *ISO 9186-2:2008 - Graphical symbols — Test methods*. <https://www.iso.org/standard/43484.html>
- [41] Matthew Kay and Jeffrey Heer. 2015. Beyond weber's law: A second look at ranking visualizations of correlation. *IEEE Trans. Vis. Comput. Graph.* 22, 1 (2015), 469–478. doi:10.1109/TVCG.2015.2467671
- [42] Max Kinadeer, William H Warren, and Karen B Schloss. 2019. What color are emergency exit signs? Egress behavior differs from verbal report. *Appl. Ergon.* 75 (2019), 155–160.
- [43] Robert Kosara. 2016. An empire built on sand: Reexamining what we think we know about visualization. In *BELIV*.
- [44] Sharon Lin, Julie Fortuna, Chinmay Kulkarni, Maureen Stone, and Jeffrey Heer. 2013. Selecting Semantically-Resonant Colors for Data Visualization. *Comput. Graph. Forum* 32, 3pt4 (2013), 401–410. doi:10.1111/cgf.12127
- [45] Sheng Long and Matthew Kay. 2024. To cut or not to cut? a systematic exploration of y-axis truncation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [46] Kecheng Lu, Mi Feng, Xin Chen, Michael Sedlmair, Oliver Deussen, Dani Lischinski, Zhanglin Cheng, and Yunhai Wang. 2020. Palettaior: Discriminable Colorization for Categorical Data. arXiv:2009.02969 [cs.GR] <https://arxiv.org/abs/2009.02969>
- [47] Geoffroy Martin. 2015. iWantHue: Colors for Data Scientists. <https://medialab.github.io/iwanthue/>. Accessed: 2025-03-26.
- [48] Adrian Mayorga and Michael Gleicher. 2013. Splatterplots: Overcoming overdraw in scatter plots. *IEEE Trans. Vis. Comput. Graph.* 19, 9 (2013), 1526–1538.
- [49] Microsoft Corporation. 2024. *Microsoft Excel*. <https://office.microsoft.com/excel>
- [50] Tamara Munzner. 2014. *Visualization analysis and design*. CRC press.
- [51] Christine Nothelfer, Michael Gleicher, and Steven Franconeri. 2017. Redundant encoding strengthens segmentation and grouping in visual displays of data. *Journal of Experimental Psychology: Human Perception and Performance* 43, 9 (2017), 1667.
- [52] Ghulam Jilani Quadri and Paul Rosen. 2021. A survey of perception-based visualization studies by task. *IEEE Trans. Vis. Comput. Graph.* 28, 12 (2021). doi:10.1109/tvcg.2021.3098240
- [53] Ghulam Jilani Quadri, Arran Zeyu Wang, Zhehao Wang, Jennifer Adorno, Paul Rosen, and Danielle Albers Szafir. 2024. Do You See What I See? A Qualitative Study Eliciting High-Level Visualization Comprehension. In *ACM CHI*. 1–26. doi:10.1145/3613904.3642813
- [54] R Core Team. 2021. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project>

- org/
- [55] Ronald Rensink and Gideon Baldrige. 2010. The perception of correlation in scatterplots. *Comput. Graph. Forum* 29, 3 (2010), 1203–1210. doi:10.1111/j.1467-8659.2009.01694.x
  - [56] Ronald A Rensink. 2013. On the prospects for a science of visualization. In *Handbook of human centric visualization*. Springer, 147–175.
  - [57] Hannah Ritchie and Max Roser. 2019. Age structure. *Our World in Data* (2019).
  - [58] Francesca Samsel, Lyn Bartram, and Annie Bares. 2018. Art, affect and color: Creating engaging expressive scientific visualization. In *2018 IEEE VIS Arts Program (VISAP)*. IEEE, 1–9.
  - [59] Alper Sarikaya and Michael Gleicher. 2017. Scatterplots: Tasks, data, and designs. *IEEE Trans. Vis. Comput. Graph.* 24, 1 (2017), 402–412. doi:10.1109/TVCG.2017.2744184
  - [60] Abhraneel Sarma, Shunan Guo, Jane Hoffswell, Ryan Rossi, Fan Du, Eunyeek Koh, and Matthew Kay. 2022. Evaluating the Use of Uncertainty Visualisations for Imputations of Data Missing At Random in Scatterplots. *IEEE Trans. Vis. Comput. Graph.* 29, 1 (2022), 602–612.
  - [61] 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 Trans. Vis. Comput. Graph.* 25, 1 (2018), 810–819. doi:10.1109/TVCG.2018.2865147
  - [62] Karen B Schloss, Laurent Lessard, Charlotte S Walmsley, and Kathleen Foley. 2018. Color inference in visual communication: the meaning of colors in recycling. *Cogn. Res.: Princ. Implic.* 3, 1 (2018), 1–17.
  - [63] Karen B Schloss and Stephen E Palmer. 2010. Aesthetics of color combinations. In *Human Vision and Electronic Imaging XV*, Vol. 7527. SPIE, 365–376.
  - [64] Karen B Schloss, Eli D Strauss, and Stephen E Palmer. 2013. Object color preferences. *Color Research & Application* 38, 6 (2013), 393–411.
  - [65] Karen B Schloss, Christoph Witzel, and Leslie Y Lai. 2020. Blue hues don't bring the blues: questioning conventional notions of color-emotion associations. *Journal of the Optical Society of America A* 37, 5 (2020), 813–824.
  - [66] Michael Sedlmair, Andrada Tatu, Tamara Munzner, and Melanie Tory. 2012. A taxonomy of visual cluster separation factors. *Comput. Graph. Forum* 31, 3pt4 (2012), 1335–1344.
  - [67] Gaurav Sharma, Wencheng Wu, and Edul N Dalal. 2005. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Res. Appl.* 30, 1 (2005), 21–30.
  - [68] Stephen Smart and Danielle Albers Szafrir. 2019. Measuring the Separability of Shape, Size, and Color in Scatterplots. In *Proc. ACM Hum. Factors Comput. Syst. (CHI)*. 669. doi:10.1145/3290605.3300899
  - [69] Stephen Smart, Keke Wu, and Danielle Albers Szafrir. 2019. Color crafting: Automating the construction of designer quality color ramps. *IEEE Trans. Vis. Comput. Graph.* 26, 1 (2019), 1215–1225.
  - [70] Inc. Statgraphics Technologies. 2022. Statgraphics19. <https://www.statgraphics.com/>
  - [71] Maureen Stone. 2006. Choosing colors for data visualization. *Business Intelligence Network* 2 (2006).
  - [72] Maureen Stone, Danielle Albers Szafrir, and Vidya Setlur. 2014. An engineering model for color difference as a function of size. In *Color and Imaging Conference*, Vol. 2014. 253–258.
  - [73] Danielle Albers Szafrir. 2018. Modeling color difference for visualization design. *IEEE Trans. Vis. Comput. Graph.* 24, 1 (2018), 392–401. doi:10.1109/TVCG.2017.2744359
  - [74] Danielle Albers Szafrir, Rita Borgo, Min Chen, Darren J Edwards, Brian Fisher, and Lace Padilla. 2023. *Visualization Psychology*. Springer Nature.
  - [75] Danielle Albers Szafrir, Steve Haroz, Michael Gleicher, and Steven Franconeri. 2016. Four types of ensemble coding in data visualizations. *Journal of Vision* 16, 5 (2016), 11–11. doi:10.1167/16.5.11
  - [76] Danielle Albers Szafrir, Maureen Stone, and Michael Gleicher. 2014. Adapting color difference for design. In *Color and Imaging Conference*, Vol. 22. Society for Imaging Science and Technology, 228–233.
  - [77] Tableau. 2022. Tableau. <https://www.tableau.com/>
  - [78] Chin Tseng, Ghulam Jilani Quadri, Zeyu Wang, and Danielle Albers Szafrir. 2023. Measuring Categorical Perception in Color-Coded Scatterplots. In *Proc. ACM Hum. Factors Comput. Syst. (CHI)*. doi:10.1145/3544548.3581416
  - [79] Chin Tseng, Arran Zeyu Wang, Ghulam Jilani Quadri, and Danielle Albers Szafrir. 2025. Shape It Up: An Empirically Grounded Approach for Designing Shape Palettes. *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS 2024)* (2025).
  - [80] Chin Tseng, Arran Zeyu Wang, Ghulam Jilani Quadri, and Danielle Albers Szafrir. 2024. Revisiting Categorical Color Perception in Scatterplots: Sequential, Diverging, and Categorical Palettes. In *Proceedings of the 26th EG/VGTC Conference on Visualization (EuroVis)*. doi:10.2312/evs.20241073
  - [81] Arran Zeyu Wang, Ghulam Jilani Quadri, Mengyuan Zhu, Chin Tseng, and Danielle Albers Szafrir. 2026. Characterizing Visualization Perception with Psychological Phenomena: Uncovering the Role of Subitizing in Data Visualization. *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS 2025)* (2026).
  - [82] Hongru Wang, Melinda A Yang, Shargan Wangdue, Hongliang Lu, Honghai Chen, Linhui Li, Guanghui Dong, Tinley Tsring, Haibing Yuan, Wei He, et al. 2023. Human genetic history on the Tibetan Plateau in the past 5100 years. *Science advances* 9, 11 (2023), eadd5582.
  - [83] Lujin Wang, Joachim Giesen, Kevin T McDonnell, Peter Zolliker, and Klaus Mueller. 2008. Color design for illustrative visualization. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1739–1754.
  - [84] Yunhai Wang, Zeyu Wang, Lifeng Zhu, Jian Zhang, Chi-Wing Fu, Zhanglin Cheng, Changhe Tu, and Baoquan Chen. 2017. Is there a robust technique for selecting aspect ratios in line charts? *IEEE Trans. Vis. Comput. Graph.* 24, 12 (2017), 3096–3110.
  - [85] Colin Ware. 2009. Quantitative texton sequences for legible bivariate maps. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1523–1530. doi:10.1109/TVCG.2009.175
  - [86] Colin Ware. 2012. *Information visualization: perception for design*. Elsevier.
  - [87] Colin Ware, Terece L Turton, Roxana Bujack, Francesca Samsel, Piyush Shrivastava, and David H Rogers. 2018. Measuring and modeling the feature detection threshold functions of colormaps. *IEEE transactions on visualization and computer graphics* 25, 9 (2018), 2777–2790.
  - [88] Bang Wong. 2010. Points of view: Color coding. *Nature Methods* 7, 8 (2010), 573.
  - [89] Cindy Xiong, Joel Shapiro, Jessica Hullman, and Steven Franconeri. 2019. Illusion of causality in visualized data. *IEEE Trans. Vis. Comput. Graph.* 26, 1 (2019), 853–862.
  - [90] Achim Zeileis, Kurt Hornik, and Paul Murrell. 2009. Escaping RGBland: Selecting colors for statistical graphics. *Comput. Stat. Data Anal.* 53, 9 (2009), 3259–3270.
  - [91] 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.
  - [92] Clementine Zimnicki, Chin Tseng, Danielle Albers Szafrir, and Karen B Schloss. 2023. Effects of data distribution and granularity on color semantics for colormap data visualizations. In *2023 IEEE Visualization and Visual Analytics (VIS)*. IEEE, 11–15.