

C2Views: Knowledge-based Colormap Design for Multiple-View Consistency

Yihan Hou¹, Yilin Ye^{1,2}, Liangwei Wang¹, Huamin Qu², and Wei Zeng^{1,2†}

¹The Hong Kong University of Science and Technology (Guangzhou)

²The Hong Kong University of Science and Technology

Abstract

Multiple-view (MV) visualization provides a comprehensive and integrated perspective on complex data, establishing itself as an effective method for visual communication and exploratory data analysis. While existing studies have predominantly focused on designing explicit visual linkages and coordinated interactions to facilitate the exploration of MV visualizations, these approaches often demand extra graphical and interactive effort, overlooking the potential of color as an effective channel for encoding data and relationships. Addressing this oversight, we introduce C2Views, a new framework for colormap design that implicitly shows the relation across views. We begin by structuring the components and their relationships within MVs into a knowledge-based graph specification, wherein colormaps, data, and views are denoted as entities, and the interactions among them are illustrated as relations. Building on this representation, we formulate the design criteria as an optimization problem and employ a genetic algorithm enhanced by Pareto optimality, generating colormaps that balance single-view effectiveness and multiple-view consistency. Our approach is further complemented with an interactive interface for user-intended refinement. We demonstrate the feasibility of C2Views through various colormap design examples for MVs, underscoring its adaptability to diverse data relationships and view layouts. Comparative user studies indicate that our method outperforms the existing approach in facilitating color distinction and enhancing multiple-view consistency, thereby simplifying data exploration processes.

CCS Concepts

• **Human-centered computing** → **Visualization systems and tools**;

1. Introduction

Multiple-view (MV) visualization is a popular method for effective visual communication and data analysis, often seen in dashboards [BFAR*22] and visual analytics systems [CZL*21]. Representing relations among views is essential to facilitate cross-view tasks and enhance visual analysis. However, effectively depicting these relationships in multi-view settings can be challenging, particularly when dealing with complex data. Current approaches rely on explicit visual linkages [CC07, SNK*22] or interactions [Wea09, CSC*22] to indicate coordination, potentially causing cognitive overload, highlighting the demand for implicit methods.

Implicit methods, such as view layout [LHD18, WZW*23] and data operations [Wu22], help maintain view consistency. Color, a fundamental visual channel that operates independently of form and position, is well-suited for encoding data and indicating MV consistency simultaneously. Effective colormaps enhance the effectiveness of single-view (SV) visualization and improve the consistency among multiple views. However, designing a satisfactory col-

ormap for SV visualization is challenging because it requires considering multiple criteria, such as color distinctness and harmony. Designing colormaps for MV visualizations, which can have various types of data interrelations among views, adds complexity.

Existing authoring tools such as Tableau [tab] support iterative colormap design for individual views, but this manual process is challenging. Novice users are often lack expertise in design principles for multi-view settings, leading them to unwittingly violate consistency rules [QH17]. Even experts find balancing SV clarity and MV consistency arduous. Common color design pitfalls appear even in well-crafted dashboards, such as those featured in Tableau’s “Viz of the Day” community. A frequent issue is ambiguity and confusion from inconsistent color usage (e.g., teal might represent different data in different views in Figure 1(a)). A more complex issue arises when insufficiently distinct parent colors lead to problematic hierarchical palettes. In Figure 1(b), two shades of blue represent the parent categories “Coal” and “Oil”. While distinct individually, the sub-types’ palettes overlap perceptually, hindering comparison across the two hierarchies.

Automated colormap design can mitigate manual pitfalls, but current methods focus on SV visualizations, neglecting MV needs.

† Wei Zeng is the corresponding author




Figure 1: Common pitfalls for colormap design in MV visualization: a) different data with the same colormap, b) similar hierarchy colormaps for adjacent views.

Our semi-automatic approach addresses this by balancing SV effectiveness with MV consistency, tackling two challenges. First, applying design principles for MV colormaps involves reconciling conflicting rules. For instance, the need for discriminative colors for SV effectiveness [LFC*21, GLS17] often clashes with the demand for uniform hues to represent hierarchical relationships across views. Simply choosing one rule over the other leads to sub-optimal design, making the balance of these competing demands a challenging trade-off that necessitates a multi-objective optimization approach to achieve a satisfactory solution. Second, handling MV complexity requires a robust way to model intricate data-view relationships, influenced by factors like data and layout. While existing colormap recommendation approaches for MV visualizations [SCL*22] focus on consistency among views, they overlook detailed inter-view data relationships. To unify these factors into a computable format, a structured graph representation is essential.

To address these challenges, we categorize MV colormap design considerations into: 1) SV effectiveness, including local discriminability, and 2) MV consistency, including global discriminability, hue uniformity, continuity and spatial proximity (Sect. 3.2). We propose a knowledge-based graph representation that unifies the relevant design components, with colormap, data, and view represented as entities and relationships among them represented as relations (Sect. 4.1). Based on this model, we define design criteria and formulate an optimization problem, solved via a genetic algorithm with Pareto optimality to generate colormaps across all views (Sect. 4.2). The automatic approach is complemented with an interactive interface for user-intended refinements (Sect. 4.3). We present case studies that demonstrate the feasibility and expressiveness of our approach in different scenarios (Sect. 5.1). We also conducted a user study, and the results showed that our method outperforms the existing colormap recommendation approach, particularly in accomplishing MV consistency (Sect. 5.3).

In summary, our work makes the following contributions:

- Identification and distillation of MV colormap design considerations, covering SV effectiveness and MV consistency.
- A knowledge-based graph representation unifying design components, including colormap, data, view, and complex design criteria from relations of the components.
- Development of *C2Views*, a semi-automatic approach using a Pareto-enhanced genetic algorithm for colormap generation, complemented by an interactive refinement interface, validated through user studies and interviews. The source code and all sup-

plementary materials are publicly available at <https://github.com/Sunnary2604/C2Views>.

2. Related Work

Multi-View Visualization. MV visualization is widely used in visual analytics for data analysis [Rob07] and in dashboards for visual communication [BFAR*22, YHL*25]. Designing effective MV visualizations requires consideration from various perspectives. Qu and Hullman outlined encoding constraints, including data cardinality, available encoding resources, and chart layout, to achieve consistency across multiple views [QH17]. Additional factors identified in the visualization community include data relationships [SSAZ21, KGB21], view composition [Wu22, CZL*21], interaction coordination [CSC*22], and display modality [LKD19, ZCH*23]. Maintaining view consistency is a primary goal in MV visualizations, which can be achieved through explicit visual linkages (e.g., [CC07, SNK*22, Wea09]) or implicit designs such as layout and color [DFSB98, Wea04, GRW*00]. Explicit methods risk visual clutter and interaction overload, while implicit designs leverage existing visual elements to avoid these issues. For example, Wen et al. [WZW*23] used proper layouts to facilitate immersive analytics involving multiple views. In our work, we adopt an implicit approach, focusing on colormap design to enhance multi-view consistency and using the large design space of colors to align with complex data relationships.

Colormap Design. Color is a key visual channel for information perception and has attracted attention in visual perception [Hol11, MSF22] and graphic design [OAH11, GSX*22]. In visualization, color is frequently used to encode data and facilitate user understanding [SSM11, ZH15]. Colormaps define the mapping from data values to colors, and extensive studies have examined perception and its implications for colormap design. To ensure effectiveness, proposed guidelines include avoiding rainbow colormaps [LH18, RS21], creating data-dependent designs [HB03], optimizing perceptual discriminability [LSS13, FWD*17, LFC*21] and contrast [MSK14, MK15], ensuring harmony [PFC18], and following color semantics [LFK*13, SS15].

Despite these principles, applying them in practice is challenging, especially for novices. As a result, many efforts have been made to simplify the colormap design process. Approaches fall into three categories: *template-based* approaches (e.g., ColorBrewer [HB03], VSUP [CMH18]) offer pre-curated colormaps for users to choose; *rule-based* approaches (e.g., PRAVDAColor [BRT95], Palettator [LFC*21] and CCC-Tool [NCB*21]) identify specific colormap design rules and develop algorithms for automatic creation; and *data-driven* approaches extract color mappings from existing designs using techniques such as color legend detection [PMH18], deep learning [YZF*22], or color-space structure analysis [SWS19]. These approaches are incorporated into libraries and tools like d3.js [BOH11], Tableau [tab], and PowerBI [pow]. Other research has also extended colormap design to dynamic contexts like animations and multi-scale models [WMW*16].

However, most studies focus on single-view visualizations, whilst little research exists on colormap design for MVs. Some early works focus on maintaining color coordination among MV vi-

sualizations [DFSB98, Wea04, GRW*00]. The most relevant is ColorCook [SCL*22], a rule-based colormap optimization for dashboards guided by effectiveness and expressiveness. However, it treats all views uniformly, generating a single parallel palette, and thus overlooks the specific relational structure between views and more complex data relationships like hierarchies. In contrast, our work comprehensively summarizes data and layout relations among views and formulates these relations in a knowledge-based graph to enable automated MV colormap design.

Automated Visualization Design. Recent years have seen growing interest in automated visualization design, which falls into two main categories: data-driven and rule-based approaches. Data-driven methods apply machine learning and generative models to recommend designs from large visualization datasets [YHH*24], and have proven effective for dashboards [WXL*21, LLW*24] and colormaps [SWS19, YZZ*21]. However, they heavily rely on large and high-quality datasets, which are challenging to obtain for MV visualizations. Datasets such as VIS30K [CLL*21] and other real-world datasets (e.g., [CZL*21]) contain MV examples but lack colormap labels. In contrast, rule-based approaches formalize design guidelines into quantifiable constraints and solve them via optimization, such as Voyager [WMA*15] and Draco [MWN*19]. Nonetheless, defining appropriate rules for automated visualization design is labor-intensive, especially for MV visualizations. For instance, DMiner [LLW*24] mined over 600 rules from existing MV designs. As the complexity of rules increases, the number of design constraints that must be satisfied also increases significantly, making optimization more challenging.

To manage the complexity of rule-based systems, we employ a knowledge-based graph representation. While this approach has been successfully used to model intricate rules in general visualization design [GSGC08, CE19, LWZ*22], it has not been applied to colormap design specifically. We take advantage of the representation of relations among colormap, data, and view entities to optimize colormaps for MV.

3. Scope and Design Considerations

This section clarifies the scope of the work (Sect. 3.1), followed by design requirements and considerations (Sect. 3.2).

3.1. Method Scope

Effective color usage is essential for conveying information and context in common MV applications like dashboards and visual analytics systems. We conducted systematic reviews on colormap design in visualization [SSM11, ZH15] and identified various design perspectives that need to be considered, including:

- *Color consistency in MV visualization.* Our primary focus is on MV visualizations, with a particular emphasis on color consistency across views based on data relationships. Nevertheless, we also recognize the importance of SV effectiveness, including considerations to ensure data correspondence and maintain graphical integrity in SV visualizations.
- *Data mapping.* Besides mapping data values, color can also be a powerful tool for conveying meanings and semantics [LFK*13,

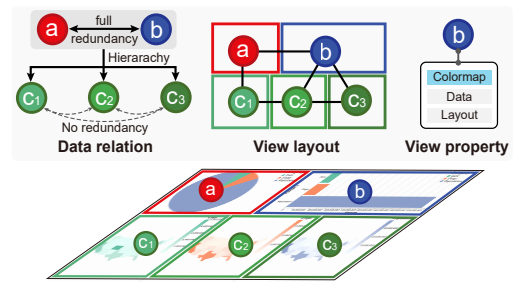


Figure 2: The main considerations for colormap design in MV visualizations in this work include data relation and view layout.

SGS*19, KGB21, SCL*22]. However, it is important to be aware of potential limitations, such as dissimilar background colors can significantly impact color perception and hinder effective data understanding [Mit15]. This work focuses on color mapping for data values rather than on reflecting meaning or semantics.

- *1D colormap.* The focus of this study is on designing 1D colormaps for individual views rather than creating bi- or multi-variate colormaps for multiple views simultaneously. While there are scenarios in which bi-variate data must be presented together (e.g., [ZLL*21]), a variety of 2D colormaps (e.g., VSUP [CMH18]) have been developed to meet this need. However, these 2D colormaps have been criticized for being challenging to interpret, particularly when there are multiple levels in both attributes [BSM*15].

3.2. Design Requirements and Considerations

The manual design process often begins with selecting a colormap based on foundational rules like perceptual distance and effectiveness. However, for multi-view visualizations, this is usually followed by a complex series of manual adjustments, including assigning, checking, and tuning colors to ensure cross-view consistency [SCL*22]. The aim of this work is to automate this complex procedure, considering the following two main requirements:

- R.1 Single-View (SV) Effectiveness:** To comply with the effectiveness principle [Mun14], the choice of colormap should convey unambiguous information. Numerous design guidelines have been proposed for colormap design to ensure effectiveness and assist users in tasks like identifying values. Although this work is intended for MV visualizations, achieving effectiveness in a single view is a main consideration.
- R.2 Multiple-View (MV) Consistency:** MV visualizations provide flexibility for cross-view data analysis, such as comparing values in juxtapositioned views [Gle18]. Enhancing underlying relationships among views can assist users in gaining a comprehensive understanding of the data. This study aims to enhance view consistency using color, an implicit visual channel that requires no extra interaction or visual workload.

To translate our design requirements into a computable model, we first consider the core components of the design space, including colormap, data, and view components, along with their relationships (Figure 2). Based on this structure, we formulate our high-

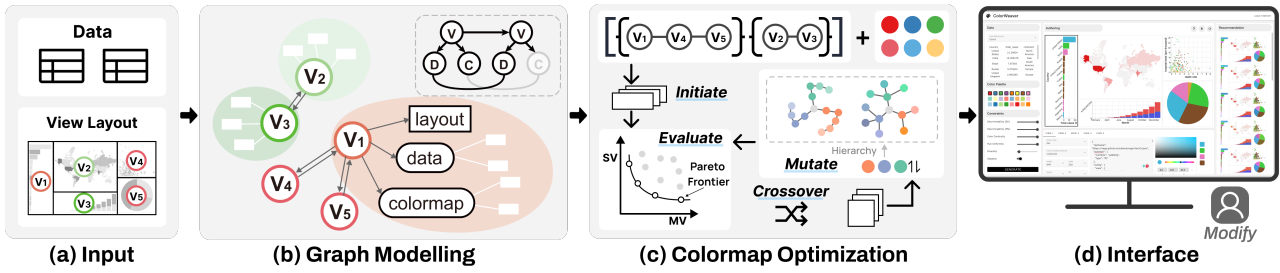


Figure 3: The proposed two-stage framework for colormap design on MV visualizations. (a) Input from the user, (b) knowledge-based graph construction representing the relationship between colormap, data, and view, (c) A knowledge-based graph optimization process for deriving an optimal colormap for multiple views, and (d) Interactive User Interface for customizing and adjusting the colormap used in each view.

Table 1: Design requirements and their corresponding criteria.

Requirement	Criteria	Goals
R1: SV Effectiveness	C1: Local Discriminability	Identify marks [GLS17, LFC*21]
R2: MV Consistency	C2: Global Discriminability	Keeping consistence [QH17, KGB21]
	C3: Hue Uniformity	Enhancing hierarchy relation [Tdj14]
	C4: Continuity	Equal presence of data [CWM09, SCL*22, KYKL14]
	C5: Spatial Proximity	Balance constraints [QH17]

level requirements into the following five design criteria, which are then quantified as mathematical metrics in Sect. 4.2.1.

Many studies have focused on enhancing the effectiveness of color encoding for SV visualization, which mainly focuses on color discriminability (C1).

C1. Local Discriminability: The ability to quickly and accurately identify a graphical mark and perceive differences among marks is a critical design criterion when representing categorical data [GLS17, LFC*21]. Though influenced by other factors like name difference, color discriminability is primarily determined by the perceptual distance between color pairs in a uniform color space.

In contrast, there are fewer guidelines available on designing colormaps for MV consistency, particularly use colormaps to represent cross-view relationships in MV. To address this gap, we reviewed the relevant literature on MV design and identified key considerations related to *data relationship*, including full/partial/no redundancy and hierarchy relations, as well as *view layout* for the spatial distance between views. We then related these considerations to colormap design and distilled a set of desirable criteria as follows:

C2. Global Discriminability: The consistency constraint model advocates for the use of consistent encoding for identical data and distinct encoding for different data [QH17, KGB21]. Specifically, we use identical colors for the full/partial redundancy relations data. Sticking to the consistency model, data attributes that have no redundancy relations should be assigned distinct colors. The criteria not only considers color discriminability within a single view but also requires that the colors be distinguishable from each other in multiple views. To achieve this, the same metric for a perceptual distance of color discriminability can be used.

C3. Hue Uniformity: Hierarchical relationships across views are

prevalent, connecting visual elements in one view to several in another [SSAZ21]. To illustrate these connections, especially for categorical elements, the tree colormap can be utilized [Tdj14]. Effective tree colormaps deploy perceptually uniform hues to clarify hierarchical links among different elements. On the other hand, sequential colormaps with varying lightness levels but the same hue are recommended for ordinal attributes of the child data.

C4. Continuity: The Gestalt laws of continuity suggest that consistent use of colors can improve user preference and create a sense of unity [Wer38, KYKL14, SCL*22]. We prioritize the effective, even representation of attributes, avoiding the distortion caused by overly bright colors. Balancing color continuity and color discrimination can often be a challenge. However, we strive to achieve a balance between the two by adhering to the principles set forth in Colorlogical [GLS17].

C5. Spatial Proximity: Meeting all of the above criteria within a limited color space can be challenging. Studies indicate that the layout of multiple views (MVs) can affect the constraints among them [QH17]. For instance, global distinctness may not be as crucial when two views are not placed side by side. To address this, we calculate the spatial proximity between views and use it to assign weight to the aforementioned criteria.

4. C2Views

This section outlines the two-stage framework for achieving color consistency in an MV visualization, including *Graph Modeling* (Sect. 4.1) that formulates the design considerations and related components in a graph, and *Colormap Optimization* (Sect. 4.2) that optimize colormaps to balance SV effectiveness and MV consistency. Moreover, we offer an *Interactive Interface* (Sect. 4.3) that enables user-driven refinement based on their preference.

4.1. Graph Modeling

To model the complex components and relationships for automated multi-view colormap design, we propose a holistic knowledge-based graph. This graph encodes instance-specific facts, such as data relationships and view layouts, inferred from user inputs to build a model for optimization. The “knowledge” here refers to a priori design principles from visualization literature, which are applied to derive color relations and complete the graph structure,

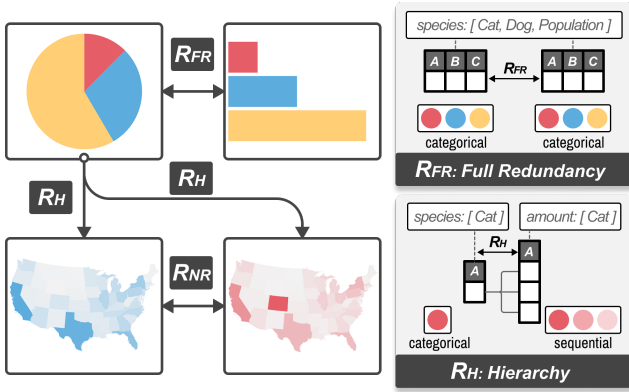


Figure 4: Example of data-data relations for color palette generation, including full redundancy and hierarchy.

rather than being directly encoded. This process yields a computable model tailored to each specific visualization instance.

Definition of Entity. In the knowledge-based graph, three types of entities are defined as follows:

- **View.** An MV visualization consists of a set of views, denoted as $\{V_i | i \in \{1, \dots, n\}\}$, where V_i represents an individual view and n represents the view count. Each view V_i includes properties of a bounding box and chart type predefined by users. The bounding boxes determine the layout relationship of the views.
- **Data.** An MV visualization can represent various datasets or multiple perspectives of the same dataset [Rob07]. For simplicity, we focus on the view level and refer to the data encoded by a single view as $D(V_i)$. The data $D(V_i)$ may consist of multiple fields, but at least one field must be encoded using colors.
- **Colormap.** As previously mentioned, this work uses colors for encoding data without taking into account color semantic mapping. Thus, each view utilizes a colormap to represent a single field in the corresponding data, denoted as $C(V_i) \rightarrow D(V_i)$. A colormap $C(V_i)$ can either be a *discrete* or *continuous* type, where *discrete* colormaps consist of a limited number of colors, while *continuous* colormaps include a gradual progression of colors without abrupt changes. To maintain consistency, we model both discrete and continuous colormaps as an ordered list of m colors, denoted as $C(V_i) := \{c_j\}_{j=1}^m$, for which we evenly sample m colors from continuous colormaps.

Definition of Relations. Within a single view, the relationship between the view V_i , its corresponding data $D(V_i)$, and the colormap $C(V_i)$ is straightforward. Our focus is on the cross-view relationships between views and data. To model these relationships, we first form a topology graph by connecting all adjacent views. This topology graph allows us to quantify the view-view relations $R(V_i, V_j)$ between two views V_i and V_j based on their distance. In contrast, the data-data relations $R(D(V_i), D(V_j))$ between $D(V_i)$ and $D(V_j)$ can be more complex. We conducted a literature review and identified several common types of data-data relations, including:

DR.1 Full Redundancy: Two data sets $D(V_i)$ and $D(V_j)$ are considered to have a *full redundancy* relation if and only if their color-encoding data fields are exactly the same, *i.e.*,

$\forall d \in D(V_i) \rightarrow d \in D(V_j)$ and vice versa. For example, in Figure 4, the pie chart and bar chart share the same data fields, which are ‘species’ and its variables [‘cat’, ‘dog’, ‘population’]. Therefore, these data sets require identical colormaps to maintain consistency.

DR.2 Partial Redundancy: Two data sets $D(V_i)$ and $D(V_j)$ are regarded as *partial redundancy* relation if they have some overlapping but are not identical, *i.e.*, $D(V_i) \cap D(V_j) \neq \emptyset$ and $D(V_i) \neq D(V_j)$. To meet the consistency requirement, the color encoding for a variable existing in both data sets should be identical in both views. On the other hand, to meet the *global discriminability* requirement (C2), the colors should be distinct.

DR.3 Non-Redundancy: Two data sets $D(V_i)$ and $D(V_j)$ are regarded as *no redundancy* relation if they have no overlaps, *i.e.*, $D_c(V_i) \cap D_c(V_j) = \emptyset$. For example, in the two maps shown in Figure 4, colors are used to encode ‘cat’ and ‘dog’ separately, and there is no overlap in the color mapping between the two maps. In such cases, the colormaps shall be distinct from each other to meet the *global discriminability* requirement (C2).

DR.4 Hierarchy: The relation occurs when one data set, $D(V_i)$, is subordinate to or derived from another data set, $D(V_j)$. For example, in the bottom-left map shown in Figure 4, colors are used to encode the number of ‘cat’ variables in each state, which is only an element of the data in the pie chart and the bar chart. The *hierarchy* relation is directional and should be considered when determining the ordering of colormap assignment (Sect. 4.2). The colormaps should also satisfy the *hue uniformity* requirement (C5) with perceptually uniform hues to indicate the hierarchy relation. For instance, map’s red sequential colormap is consistent with the red color used in the pie chart and the bar chart.

Graph Construction. With the definitions of entities and relations, the graph G is automatically constructed for each input multi-view instance. The process begins by iterating through all pairs of views (V_i, V_j) to infer their relationships based on the user’s input. View-view relations are determined by calculating the spatial distance between their bounding boxes, adding an edge $E_V = (V_i, V_j)$ for adjacent views. Concurrently, the system infers data-data relations by analyzing the underlying data fields for patterns such as Full Redundancy (shared fields) or Hierarchy (aggregations). Once the relationship type R_{ij} is identified, a corresponding typed edge $E_D = (D(V_i), R_{ij}, D(V_j))$ is added to the graph. This iterative process results in a comprehensive graph that programmatically captures all necessary constraints for optimization.

4.2. Colormap Optimization

Optimization begins by translating our design considerations into a series of measurable metrics.

4.2.1. Metric Formulation

We first leverage the following color metrics:

- **Color Difference.** To measure the color difference, which is crucial for local discriminability (C1) in a single view and

global discriminability (C2) across multiple views, we use the CIEDE2000 color difference formula [SWD05]. Here we measure the local discriminability of categorical colors $C(V_m)$ in a single view V_m as:

$$f_{l_{dis}}(V_m) = \sum_{i=1}^{|C(V_m)|} \sum_{j=1, j \neq i}^{|C(V_m)|} DE_{00}(c_i, c_j). \quad (1)$$

A large value of $f_{l_{dis}}(V_m)$ indicates that more distinct colors are used in V_m . Similarly, we compute the global discriminability of colors in two views V_m and V_n as:

$$f_{g_{dis}}(V_m, V_n) = \sum_{i=1}^{|C(V_m)|} \sum_{j=1}^{|C(V_n)|} DE_{00}(c_i, c_j). \quad (2)$$

A high value of $f_{g_{dis}}(V_m, V_n)$ indicates that more distinct colors are used in V_m and V_n .

- **Hue Uniformity.** To maximize the distance between different child views in the hierarchy (C3), we need to generate the parent view so that the color in the HSL space has the largest hue difference. The hue uniformity f_{hu} is defined as below:

$$f_{hu}(V_m, V_n) = \min(|H(c_i) - H(c_j)|), c_i \in V_m, c_j \in V_n, \quad (3)$$

where $H(\cdot)$ stands for the hue value of a color. A larger value of f_{hu} indicates that the two most similar primary colors of the subview are as far away from each other as possible, which is preferred.

- **Color Continuity.** To meet the requirement of color continuity across the whole MV (C4), we measure the lightness difference in the CIELAB color space between colors in a colormap.

$$f_{con}(V_m, V_n) = \sum_{i=1}^{|C(V_m)|} \sum_{j=1}^{|C(V_n)|} (|L(c_i) - L(c_j)|), \quad (4)$$

where $L(\cdot)$ stands for the lightness values of a color. A smaller value of f_{con} indicates that more continuous colors are used.

Spatial Proximity. For criteria involving two views, namely $f_{g_{dis}}$, f_{hu} , f_{con} we weight the constraints based on view layout. Specifically, we take into account the spatial proximity (C5) between a pair of views V_m and V_n by incorporating the view layout into the weight function ω_{sp} .

$$\omega_{sp}(V_m, V_n) = \frac{1}{D(V_m, V_n)}, \quad (5)$$

where $D(\cdot)$ is the distance between two views. We organize all the views into a graph structure according to their adjacency relations and measure the spatial proximity between two views in terms of the shortest distance between two entities in the graph. Smaller ω_{sp} indicates the spatial proximity between two views is weaker, and the corresponding constraints could be loosened.

Global Cost Formulation with Pareto Optimality. We normalize all the costs above to $[0 - 1]$ and define the global cost function for colormap for the whole MV. The goal is to find a colormap m that effectively balances the trade-offs between single view effectiveness (C_{SV}) and multiple view consistency (C_{MV}), each defined as the weighted sum of the individual costs as follows:

$$C_{SV}(m) = \sum_{i=1}^{|V|} \omega_{dis} f_{dis}, \quad (6)$$

$$C_{MV}(m) = \sum_{i=1}^{|V|} \sum_{j=1, j \neq i}^{|V|} \omega_{sp} (-\omega_{g_{dis}} f_{g_{dis}} - \omega_{hu} f_{hu} + \omega_{con} f_{con}), \quad (7)$$

where ω_{sp} applied to criteria involving two views. The weights ω_d , ω_{con} , ω_{hu} , and $\omega_{g_{dis}}$ are all positive and reflect the relative importance of each factor; they are set to 1 by default. We then employ a Pareto optimality approach to define the global cost function G . A colormap m is said to be Pareto optimal if there is no other colormap m' such that it has a lower cost in both C_{SV} and C_{MV} . The Pareto front P is the set of all such Pareto optimal colormaps and is defined as:

$$P = \{m \in M \mid \nexists m' \in M : (C_{SV}(m') < C_{SV}(m)) \wedge (C_{MV}(m') < C_{MV}(m))\}, \quad (8)$$

where M is the set of all possible colormaps. From this Pareto front, we select colormaps that best satisfy the design requirements and constraints, ensuring that the final colormap is both effective for single views and consistent across multiple views.

4.2.2. Colormap Optimization

We determine the coloring sequence according to the cross-view relationships in the graph, prioritizing views without dependencies such as hierarchical relations. We employ a genetic algorithm to efficiently explore the complex, multi-objective search space that results from generating colors in the continuous color space rather than from a discrete palette. A genetic algorithm is particularly suitable for this task due to its population-based search strategy and flexible encoding. The algorithm takes the graph structure as input and iteratively optimizes the colormaps for all views in the multi-view visualization. The following steps outline the procedure involved in this iterative process:

Population Initialization. The genetic algorithm commences by initializing a population of colormaps, each corresponding to the unique data fields for color encoding within the MV. ColorBrewer and Tableau palettes, while not specifically crafted for MVs, offer a distinctive and broad coverage of the color space, ranging from 3 to 12 colors, making them good enough for initial seeding. Predefined palettes are utilized directly when the required color count is within their limit. When the necessary number exceeds what is available from predefined sets, we procedurally generate additional colors to supplement and enlarge the initial colormap selection.

Fitness Score Evaluation. All individuals in the population undergo a fitness score evaluation determined by a global cost function and Pareto optimization. Colormaps that fail to meet hard constraints, such as insufficient color differentiation, are immediately rejected. This process yields a Pareto front comprising a set of colormaps that optimally balance single-view effectiveness with multiple-view consistency.

Selection and Crossover. To preserve high-quality colormaps, we employ an elitist strategy where a set of top-performing individuals are selected from the Pareto front to form the elite pool for the next generation. We then perform crossover by randomly pairing colormaps from the Pareto front and swapping dominant colors with a 50% chance. This process passes on advantageous features while infusing the new generation with variation.

Mutation. The genetic algorithm plays a pivotal role in the optimization process by ensuring variety within the population while preserving the beneficial traits from previous iterations. This is achieved through two key operations: perturb and inheritance.

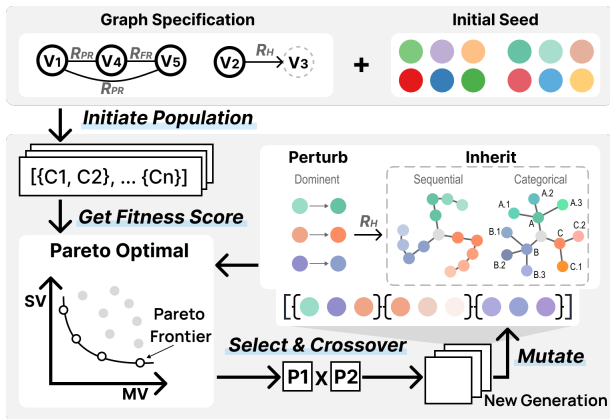


Figure 5: Optimizing colormaps via a Pareto-enhanced genetic algorithm. Inputs include multi-view specifications and an initial color palette. The iterative process yields colors that balance SV effectiveness with MV consistency through Pareto optimization.

- **Perturb:** This phase is fundamental for adding diversity to the population. By adjusting the dominant colors of the existing colormaps within a predetermined range of steps, we maintain a delicate balance. This controlled tweaking of the colormap’s colors is essential for fostering diversity and retaining the successful attributes identified in earlier rounds.
- **Inherit:** For hierarchical relations, the operation essentially involves generating a set of related colors inherited from a dominant color. Specifically, we adopt different approaches based on whether data in the child node are sequential or categorical or sequential. For sequential data, we modify the luminance and chroma in the HCL color space while keeping the hue consistent with the single color input to generate the corresponding sequential colormap. For categorical data, we draw inspiration from Tree Colors [TdJ14] for hierarchy-structured data, where the initial color is centered in the HCL color space, and the hue is adjusted within a certain range to generate colors with a similar hue. Figure 5(a) illustrates a general case of this process, where the three marked colors (A-green, B-blue, and C-orange) represent the initial colormap. We expand horizontally on the color ring to generate three categorical colormaps based on the hue of the parent color, i.e., A.1-3 inherited from A, B.1-3 inherited from B, and C.1-2 inherited from C. We reorder the colors in each generated colormap to ensure the distance between the two adjacent visual marks is uniformly perceived.

After executing multiple iterations of the genetic algorithm, a set of Pareto optimal colormaps emerges. These colormaps are not only fine-tuned for individual view efficiency but are also consistent across the entire MV visualization, providing a superior user experience and enhanced data comprehension.

4.3. Interactive User Interface

While our automated approach generates optimized colormaps, relying solely on automation may not be sufficient. Manual refinement is essential to address subjective needs like aesthetics or branding, which are difficult to capture with perceptual metrics. To

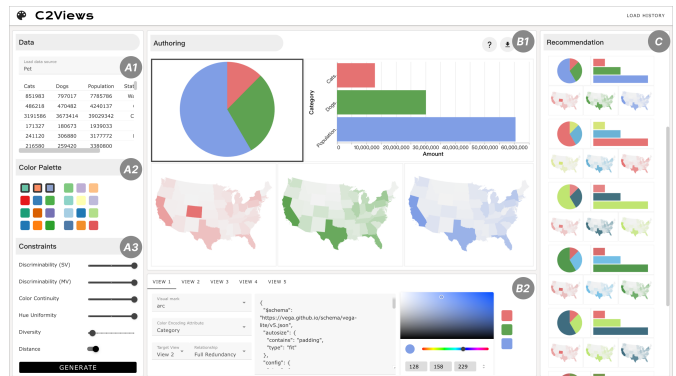


Figure 6: System interface consists of a Control Panel (A), an Authoring Panel with editing detail (B), and a Gallery (C).

support this, our interactive interface supplements automation with manual control, allowing users to fine-tune the resulting colormaps. As shown in Figure 6, the interface has three components:

- **Control Panel.** Users can load their dataset, view, select from predefined color palettes, and alter constraint settings. This panel presents a table view featuring the basic properties of the chosen dataset (Figure 6(A1)). Predefined color palettes are available in the color panel (Figure 6(A2)). Users can tweak the weight of each constraint in the constraint panel (Figure 6(A3)). Once settings are satisfactory, users can generate colormaps for the target MV visualization by clicking on the “GENERATE” button.
- **Authoring Panel.** This panel showing the MV visualization with the chosen colormap applied (Figure 6(B1)). The editing panel (Figure 6(B2)) allows users to fine-tune the color settings for each view. Using a tabbed interface, users can select a view to inspect its properties, which display the visual mark being encoded, the colormap type, and the automatically calculated data relationship. For direct modification, the panel provides a color picker to adjust specific colors, as well as an editable code view of the Vega-Lite specification.
- **Colormap Gallery.** Upon generation, the gallery displays a set of Pareto-optimal colormap solutions (Figure 6 (C)). Each represents a different trade-off and serves as a high-quality starting point for refinement. A user can select any of these options to apply it to the main Authoring Panel for detailed adjustments.

To preserve consistency during manual refinements, the system employs an automatic propagation mechanism. When a user modifies a color in the Authoring Panel, the system references the knowledge-based graph to instantly recalculate and update colors in all linked views based on their data relationship. This lightweight approach ensures consistency is preserved across the entire multi-view display without the need for a full, computationally expensive re-optimization.

5. Evaluation

This section presents a comparative feasibility case study of our approach against ColorCook [SCL*22] (Sect.5.1). Following this, we compare the results generated through our approach with those by ColorCook in a user study (Sect.5.3).

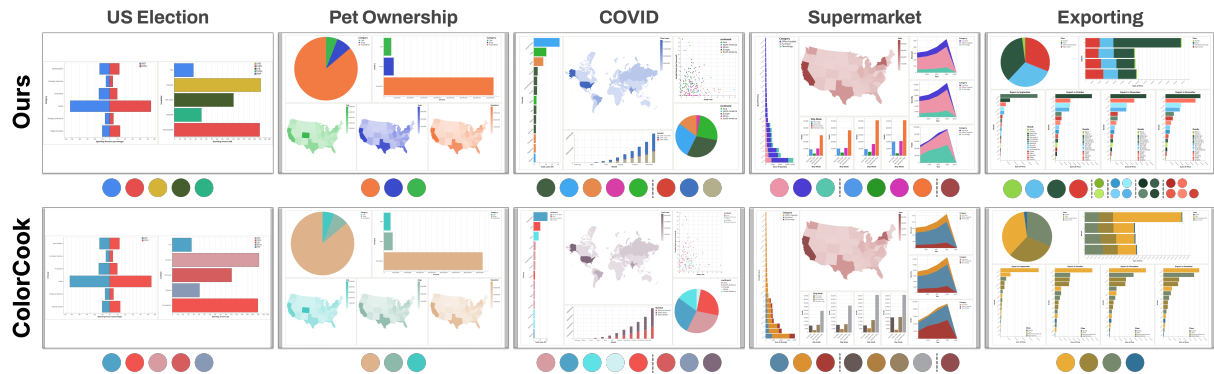


Figure 7: Examples of the result generated by ColorCook and C2Views, with varying view counts, cross-view relationships, and data domains. The used color palettes are grouped and shown below. More details are in the supplementary material.

5.1. Case Study

Our method’s feasibility and effectiveness are demonstrated across five cases with varying view counts, cross-view relationships, and data domains, as shown in Figure 7, benchmarking against the established ColorCook [SCL*22] method. For a fair comparison, all visualizations were generated automatically: we applied ColorCook’s assignment algorithm using its most distinctive domain-appropriate palette, while our results were generated without manual refinement.

Partial Redundancy (Case 1). This two-view visualization of US election data exhibits a partial redundancy relationship (Figure 7(A)). ColorCook’s generated less distinctive colormap for the right view, likely due to its reliance on specific dominant colors for domain semantics. Furthermore, it is designed for a fixed number of categories (8 or 10), which limits flexibility for datasets with different category counts. In contrast, our approach uses a broader hue spectrum to clearly differentiate the five parties across both views.

Hierarchical (Case 2). The second case introduces an MV visualization consisting of five views that provide an overview and detailed representation of pet ownership data, as shown in Figure 7(B). As with the first case, ColorCook’s color choices for the two sectors are less distinguishable with closing hue and low saturation, complicating sector comparison and the correlation of corresponding sectors between the top and bottom views. Our method, however, uses a broader color spectrum with strong hue uniformity, improving the readability of the hierarchical relationship.

Use Multiple Group (Case 3 & Case 4). These cases demonstrate scalability of our method with more views and complex relationships. Case 3 presents an MV visualization comprising six views, divided into two color encoding groups. The left bar chart, scatter plot, and pie chart use a 5-color palette to represent continent information with partial redundancy. The bottom bar chart and map are color-coded to illustrate COVID symptoms, forming a hierarchical relation. Increasing the number of views escalates the challenge of maintaining color distinctiveness, given the limited color space. Unlike ColorCook, which treats all colorable sectors uniformly, our method considers both single-view and cross-view distinctiveness, facilitating a more effective color design. Case 4 features a nine-view visualization with three color groups. The left bar charts

and the area charts employ colors to encode different categories of goods. Four small multiples use color to represent shipping modes, while the map uses distinct colors for total sales. ColorCook introduces confusion by using a similar red for unrelated data in the map and an area chart. Our method avoids this by ensuring hue uniformity and color distinctiveness across all views, minimizing color overlap for clearer distinction.

Vertical hierarchical (Case 5) The final case presents a six-view MV visualization of export data, with a vertical hierarchical relationship: the top two views display total exports by main class, while the bottom four provide detailed breakdowns for each goods type. Lacking explicit support for hierarchies, ColorCook incorrectly applies a full redundancy pattern, using the same colors for main classes at both levels. In contrast, our method is designed to accentuate the hierarchical relationship between views, utilizing distinguishable colors for the top views and applying a tree colormap for the detailed bottom views.

One difference between the two methods is that ColorCook favors lower saturation for aesthetic reasons and employs similar hues for domain-specific palettes. In contrast, our approach emphasizes colormap effectiveness, resulting in higher saturation to enhance color distinctiveness.

5.2. Quantitative Analysis

To objectively compare our method with the baseline, we evaluated the case studies using three metrics: Worst-Case Discriminability (WCD) for single-view clarity, the Parallel Relationship Score (PRS) for non-hierarchical consistency, and the Hierarchical Quality Score (HQS) for hierarchical encodings. For all metrics, a higher score indicates better performance, mathematically defined can be found in the supplementary material.

The results, summarized in Table 2, demonstrate a significant and consistent advantage for our approach. Our method achieved a WCD score of 32.84, substantially outperforming ColorCook’s 11.61. A similar improvement was observed for parallel relationships, where our PRS score (27.64) was nearly three times higher than the baseline’s (9.85). The most critical difference was in evaluating hierarchical structures. Our method achieved a positive HQS score of 1.50, whereas the baseline scored zero. This is because the

baseline method does not account for hierarchical data, assigning the same color to all sub-categories and failing to visually represent the data's structure. In contrast, our positive score quantitatively demonstrates a successful encoding of these complex relationships.

Table 2: Quantitative comparison of our method against the baseline across three proposed metrics.

Method	WCD (\uparrow)	PRS (\uparrow)	HQS (\uparrow)
ColorCook	11.61	9.85	0.00
Ours	32.84	27.64	1.50

5.3. User Study

We further conducted a user study to compare the results generated through our approach and those by the baseline methods of ColorCook [SCL*22].

5.3.1. Study Design

Data. We crafted a testing dataset comprising four diverse cases, each varying in complexity based on the number of views, the relation between views, and the underlying data domain. All visualizations were generated automatically without manual refinement. For the ColorCook component, we selected the most distinctive color palette appropriate for the data domain from their available dataset, aligning with their emphasis on color distinctiveness and aesthetic appeal to ensure a fair comparison. Using these palettes, we then applied ColorCook's color-assignment algorithm to the MV visualizations. To ensure participants concentrated on the color design, we minimized textual distractions by eliminating non-essential labels from the visualizations. Essential text, including legends, was standardized across both methods in terms of font size to maintain readability. In total, we have 8 distinct color designs for the 2 methods, with 4 cases each.

Participants. Our study enlisted 16 participants, balanced by gender, with ages ranging from 21 to 28 years ($M=24.9$, $SD=1.87$). They came from diverse fields, including statistics, computer science, media art, and design, contributing a wide range of perspectives. All participants reported experience in using or designing visualizations, with 9 participants who claimed to know little or had general knowledge, while 7 described themselves as familiar or very familiar with the subject. In terms of color design, 10 had little or general knowledge, and 6 were familiar with the domain. This diversity aimed to provide insights across different levels of expertise in visualization and color design.

Procedure. The study was conducted in a quiet, controlled environment on a desktop computer, following four stages: introduction, experiment, questionnaire, and semi-structured interview. Participants first received a 5-minute overview of the research goals, key MV color design principles, and the role of color in illustrating cross-view relationships. The evaluation metrics in Figure 8 were also introduced. After confirming their understanding, participants reviewed MV visualizations colored by both our method and ColorCook in a counterbalanced order, completing two single-view tasks (*i.e.*, identifying maxima) and two cross-view tasks (*i.e.*, multi-step




Figure 8: Quantitative results of the subjective ratings of different approaches on the cases.

inference using color cues). To mitigate the learning effect, we slightly modified the tasks to avoid repetition, such as changing the retrieval range of values. Afterward, participants filled out a questionnaire on a scale of 1-5 to evaluate the color designs based on the established metrics, with the designs anonymized as Design 1 and Design 2 to avoid bias. The questionnaire consisted of 6 questions, with Questions 1-2 focusing on the ease of completing single-view and cross-view tasks, respectively. Questions 3-5 assessed the colormap regarding color distinctiveness, relationship clarity, information conveyance, and preference. Participants were encouraged to think aloud during the questionnaire phase. After completing four cases, a semi-structured interview was conducted to gather qualitative feedback on the overall experience and practicality of the color designs. The verbal responses were recorded and transcribed for further analysis. Each session lasted about 40 minutes, and participants received \$10 compensation.

5.3.2. Result and Analysis

From the user study we obtained 640 responses (16 participants \times 2 designs \times 4 cases \times 5 questions) from the user study. We apply the Wilcoxon test to measure the significance level of the difference between the two methods, detailed in Figure 8. Significant values are reported for $p < .05$ (*), $p < .01$ (**), $p < .001$ (***). The results are summarized as follows:

Task Completion Ease (Q1 & Q2). Our approach demonstrates a significant advantage over ColorCook in facilitating task completion ($d.f. = 63$, $p < .001$). Specifically, our method ($M=4.22$, $SD=0.65$) outperforms ColorCook ($M=3.25$, $SD=0.98$) for single-view tasks, also demonstrating greater stability. Cross-view tasks, which are inherently more complex due to the need for understanding relationships between different views, also see our method outperforming ColorCook ($M=3.78$, $SD=0.98$ vs. $M=2.98$, $SD=1.07$), underlining its enhanced efficacy in MV visualization exploration.

Color Distinctiveness (Q3). In the realm of color distinctiveness, our method ($M=4.39$, $SD=0.76$) significantly outperforms ColorCook ($M=2.84$, $SD=1.08$) ($d.f. = 63$, $p < .001$). This is due to ColorCook's reliance on domain-specific and aesthetically pleasing colormaps, which may limit the range of available colors. In contrast, our method emphasizes colormap effectiveness for discerning data attributes, resulting in clearer, more distinct color schemes.

Relationship Clarity (Q4). Our approach demonstrates enhanced clarity in visualizing view relationships ($M=4.05$, $SD=0.98$) compared to ColorCook ($M=3.23$, $SD=1.10$), with a significant differ-

ence (d.f. = 63, $p < .001$). Such difference is mainly due to ColorCook using similar colors for non-related data attributes, which may lead to confusion.

Preference (Q5). There is no significant difference in user preference between our method ($M=3.66$, $SD=0.96$) and ColorCook ($M=3.19$, $SD=1.22$). However, our method still maintains a slight advantage in user preference, indicating that our method is generally well-received by participants.

5.3.3. Feedback

In general, participants reported that using color to represent cross-view relationships implicitly was helpful, as it aided their understanding of the relationships between the data presented in each view and facilitated the completion of tasks. Here are some highlights of the user feedback.

C2Views’s Effectiveness. Most participants acknowledged the effectiveness of our approach in supporting task completion ($N=14$) since the generated colormaps have “higher construal” and “support identity different element”, making it effectively convey information ($N=13$). As P5 noted, “The bright color allows me to associate information with the corresponding legend quickly.” For multiple views, colormaps generated by C2Views support quick linking between views, and “support finding the corresponding information at a glance” (P4, P15). Despite the capability of accomplishing tasks using ColorCook colormap, it was reported to use colors of low saturation and similarity, which hinders efficient information seeking. “It is hard for me to gain information in a single view, and I will be more likely to make mistakes in the cross-view. Using distinct colors will make me less nervous, and it will be easier to gain the information whether I am identifying differences or discovering relations”, P4 reported.

Balancing Aesthetics with Information Efficiency. The majority of the participants prioritize the effectiveness and efficiency of information seeking and would like to use colormap generated by C2Views in real practice ($N=9$). They perceive it as a means to “minimize data interpretation efforts” and “swiftly access desired information”, deeming it “sufficient for data exploration”. Despite being less effective, ColorCook is appreciated for its aesthetics and domain relevance ($N=4$), especially by designers and artists. They note that they would modify the palette for improved distinction prior to real-world usage, as the existing colormap could impede task fulfillment. An insightful observation from P8 highlighted that designers are particularly sensitive to subtle color differences, and ColorCook is distinct enough for them. In such cases, they tend to use the harmonious one. Similarly, P6 advocated for considering aesthetics when the discriminability is sufficient. P2 added that the trade-off between aesthetics and effectiveness hinges on context: aesthetics take precedence in posters or presentations, whereas effectiveness is crucial for analysis.

Navigating View Relationships with Color Guidance. Participants mentioned the generated color by C2Views can help them get the connections between views quickly ($N=6$). They can easily understand the full/partial redundancy, as well as the hierarchical color for sequential data. However, they reported that the tree colormap should be carefully used. Although half of the participants

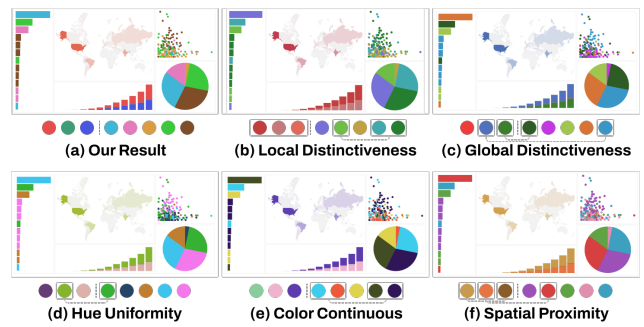


Figure 9: Example results of colormap generation with different metrics omitted from the cost function. (a) colormap considers all metrics. (b) - (f) are colormap results when omitting local distinctiveness, global distinctiveness, hue uniformity, continuity, and spatial proximity, respectively.

said they could understand the tree colormap, it still confused them at first glance and took time to understand ($N=9$). They consider the color is “different” beyond the “hierarchical relationship” and prefer using consistent color across the hierarchical category data. P14 noted that the effectiveness of the tree colormap hinges on the audience’s expertise. It may be complex for the general public but offers simultaneous overview and detail to informed users, proving particularly beneficial for administrators.

6. Discussion

In this section, we reflect on our approach, demonstrate its possible application scenarios, and discuss the limitations and future work.

Contribution of Metrics. Our methodology incorporates multiple criteria, and assessing the impact of these metrics on the outcome is essential. We investigate the effects of excluding specific metrics from the cost function:

- **Local distinctiveness.** Exclusion of this metric results in color repetition within the same view, leading to a dominant color per view due to the requirement of global distinctiveness for color variation across views (Figure 9(b)).
- **Global distinctiveness.** Without this metric, colormaps may overlap between views (Figure 9(c)), causing confusion despite well-distributed colors in each view. For example, views may inadvertently share a “green, blue” color scheme.
- **Hue Uniformity.** Omitting hue uniformity in a sequential colormap causes monopolization of hues, leaving other views to utilize colors with the same hue but altered lightness and saturation. This can make it challenging to associate the colormap with its respective view (Figure 9(d)).
- **Continuity.** The absence of this metric leads to an imbalanced color distribution (Figure 9(e)), making some colors disproportionately conspicuous and potentially distorting the intended focus, in agreement with Shi et al. [SCL*22].
- **Spatial Proximity.** The aforementioned metrics can be conflicting. Discarding spatial proximity can lead to an inability to satisfy the established criteria and risks contravening both local and global distinctiveness (Figure 9(f)).

The results confirm that all factors successfully constrain suboptimal outcomes, highlighting their roles in the optimization process.

Graph Representation and Optimization. The current optimization approach generates colormaps for all views simultaneously, which may neglect the complex relationship between views. In contrast, our data relation-based optimization algorithm considers the relationship between views by utilizing a knowledge-based graph's (entity, relation, entity) representation. This approach determines the color generation strategy based on the data relation and optimizes the global cost to generate colormaps that reflect the relationships between views. We employ a genetic algorithm with Pareto optimality to optimize and balance the trade-offs between different metrics. This is more effective than using a single weighted sum cost function, which may lead to suboptimal solutions.

Implicit vs. Explicit Encoding. The participants agreed that color is an important and useful channel to implicitly encode data relationships. Particularly, they think the consistency of colormap serves as a visual cue for them to grab the key information without reading the text. However, they also point out that many people are accustomed to combining text and color to reassure themselves of the reliability of perceived information. "While I can infer the relationship between the views from the colors, I still need words to verify my hypothesis, and explicit representation can make me more confident", P7 commented. This reflects that implicit representations are intuitive, but the combination of explicit and implicit representations is more desirable in practice. Therefore, for the MV design, it is essential to balance implicit and explicit elements, conveying the information precisely and concisely.

Generalization and Application. The feedback from our user study highlights the importance of balancing explicit text and implicit color encoding. Our implicit color encoding approach has potential benefits for various scenarios. In common cases, our approach enables users to notice data relationships between views while allowing text to confirm their assumptions. Moreover, in extreme cases where text is unavailable, it could facilitate them to grab the general idea without reading the text. For example, elderly people and people with visual impairment may find it difficult to read the text in the legend clearly. Language barriers may also prevent users from understanding information from multiple views in cross-language environments.

Limitation and Future Work. First, we address limitations related to our method's scope and implementation. Our current focus on categorical and sequential colormaps overlooks other variants like diverging colormaps and excludes considerations of aesthetics, semantics, and mark type. Similarly, we do not model subjective user preferences, such as user age and cultural background [MTK*24]. Our optimization-based approach, while effective at balancing complex trade-offs, also lacks the inherent explainability of rule-based systems like KG4Vis [LWZ*22], it can generate an optimal result but cannot provide human-readable rules explaining why. On the implementation side, we identify data relationships through specific encodings, which may not scale well or uncover inter-column relationships. Our use of Vega-Lite also covers basic mark types but omits complex marks and glyphs.

Second, our evaluation has its own limitations. The user study is

limited by the confounding of color saturation, as our emphasis on discriminability inherently produces more saturated colors than the baseline. Consequently, we cannot fully disentangle a preference for effectiveness from a potential preference for saturation alone. Lastly, although we developed an interactive interface, its usability has not been formally evaluated.

Despite these limitations, our adaptable framework is designed for extension. For instance, it could be updated to support users with color vision deficiencies by incorporating CVD-safe metrics, or expanded to include other criteria like name differences in addition to the established CIEDE2000 metric for color discriminability. We also plan to develop more sophisticated techniques for automatic relationship detection and incorporate libraries like d3.js to support an expanded range of visualizations. These efforts will enhance our colormap design approach to be more comprehensive.

7. Conclusion

This work enhances consistency in multi-view visualizations using color encoding, which simplifies designs by removing the need for extra visual connectors and reducing viewer cognitive load. We address two main challenges: maintaining effectiveness within individual views while ensuring overall consistency, and translating these needs into measurable design criteria. To solve these, we propose a framework using knowledge-based graphs for MV visualizations to unify components and establish design parameters. We provide an algorithm for generating optimal colormaps and an interactive interface for user adjustments. Our approach, validated through case studies and a user study, proves more effective and efficient than current techniques. Our work represents significant progress in enhancing MV visualization interpretability and efficiency, while substantially reducing the colormap design burden for both professional designers and non-expert users. Future work will enhance our colormap design approach to be more comprehensive.

Acknowledgment. The authors wish to thank anonymous reviewers for their constructive comments. The work was supported in part by National Natural Science Foundation of China (62172398).

References

- [BFAR*22] BACH B., FREEMAN E., ABDUL-RAHMAN A., TURKAY C., KHAN S., FAN Y., CHEN M.: Dashboard design patterns. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 342–352. doi:10.1109/TVCG.2022.3209448. 1, 2
- [BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D3: Data-driven documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2301–2309. doi:10.1109/TVCG.2011.185. 2
- [BRT95] BERGMAN L., ROGOWITZ B., TREINISH L.: A rule-based tool for assisting colormap selection. In *Proceedings Visualization '95* (1995), pp. 118–125. doi:10.1109/VISUAL.1995.480803. 2
- [BSM*15] BERNARD J., STEIGER M., MITTELSTÄDT S., THUM S., KEIM D., KOHLHAMMER J.: A survey and task-based quality assessment of static 2D colormaps. In *Visualization and Data Analysis* (2015), vol. 9397, International Society for Optics and Photonics, SPIE, p. 93970M. doi:10.1117/12.2079841. 3
- [CC07] COLLINS C., CARPENDALE S.: VisLink: Revealing relationships amongst visualizations. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1192–1199. doi:10.1109/TVCG.2007.70521. 1, 2

- [CE19] CHEN M., EBERT D. S.: An ontological framework for supporting the design and evaluation of visual analytics systems. *Computer Graphics Forum* 38, 3 (2019), 131–144. doi:10.1111/cgf.13677. 3
- [CLL*21] CHEN J., LING M., LI R., ISENBERG P., ISENBERG T., SEDLMIR M., MÖLLER T., LARAMEE R. S., SHEN H.-W., WÜNSCHE K., WANG Q.: Vis30k: A collection of figures and tables from iee visualization conference publications. *IEEE Transactions on Visualization and Computer Graphics* 27, 9 (2021), 3826–3833. doi:10.1109/TVCG.2021.3054916. 3
- [CMH18] CORRELL M., MORITZ D., HEER J.: Value-suppressing uncertainty palettes. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems* (2018), p. 1–11. doi:10.1145/3173574.3174216. 2, 3
- [CSC*22] CHEN R., SHU X., CHEN J., WENG D., TANG J., FU S., WU Y.: Nebula: A coordinating grammar of graphics. *IEEE Transactions on Visualization and Computer Graphics* 28, 12 (2022), 4127–4140. doi:10.1109/TVCG.2021.3076222. 1, 2
- [CWM09] CHUANG J., WEISKOPF D., MOLLER T.: Hue-preserving color blending. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1275–1282. doi:10.1109/TVCG.2009.150. 4
- [CZL*21] CHEN X., ZENG W., LIN Y., AI-MANEEA H. M., ROBERTS J., CHANG R.: Composition and configuration patterns in multiple-view visualizations. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 1514–1524. doi:10.1109/TVCG.2020.3030338. 1, 2, 3
- [DFSB98] DEBORAH F. SWAYNE D. C., BUJA A.: Xgobi: Interactive dynamic data visualization in the x window system. *Journal of Computational and Graphical Statistics* 7, 1 (1998), 113–130. doi:10.1080/10618600.1998.10474764. 2, 3
- [FWD*17] FANG H., WALTON S., DELAHAYE E., HARRIS J., STORCHAK D., CHEN M.: Categorical colormap optimization with visualization case studies. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 871–880. doi:10.1109/TVCG.2016.2599214. 2
- [Gle18] GLEICHER M.: Considerations for visualizing comparison. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 413–423. doi:10.1109/TVCG.2017.2744199. 3
- [GLS17] GRAMAZIO C. C., LAIDLAW D. H., SCHLOSS K. B.: Col-ogorical: Creating discriminable and preferable color palettes for information visualization. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 521–530. doi:10.1109/TVCG.2016.2598918. 2, 4
- [GRW*00] GRESH D. L., ROGOWITZ B. E., WINSLOW R. L., SCOLLAN D. F., YUNG C. K.: Weave: A system for visually linking 3-d and statistical visualizations applied to cardiac simulation and measurement data. In *Proceedings of the 11th IEEE Visualization 2000 Conference* (2000), pp. 489–492. doi:10.1109/VISUAL.2000.885739. 2, 3
- [GSGC08] GILSON O., SILVA N., GRANT P. W., CHEN M.: From Web Data to Visualization via Ontology Mapping. *Computer Graphics Forum* 27, 3 (2008), 959–966. 3
- [GSX*22] GUO S., SHI Y., XIAO P., FU Y., LIN J., ZENG W., LEE T.-Y.: Creative and Progressive Interior Color Design with Eye-tracked User Preference. *ACM Transactions on Computer-Human Interaction* 30, 1 (2022), 1–31. doi:10.1145/3542922. 2
- [HB03] HARROWER M., BREWER C. A.: ColorBrewer.org: An online tool for selecting colour schemes for maps. *The Cartographic Journal* 40, 1 (2003), 27–37. doi:10.1179/000870403235002042. 2
- [Hol11] HOLTZSCHUE L.: *Understanding Color: An Introduction for Designers*, 4th ed. Wiley, 2011. 2
- [KGB21] KRISTIANSEN Y. S., GARRISON L., BRUCKNER S.: Semantic Snapping for Guided Multi-View Visualization Design. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 43–53. doi:10.1109/TVCG.2021.3114860. 2, 3, 4
- [KYKL14] KIM H.-R., YOO M.-J., KANG H., LEE I.-K.: Perceptually-based color assignment. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 309–318. doi:10.1111/cgf.12499. 4
- [LFC*21] LU K., FENG M., CHEN X., SEDLMIR M., DEUSSEN O., LISCHINSKI D., CHENG Z., WANG Y.: Palettaylor: Discriminable colorization for categorical data. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 475–484. doi:10.1109/TVCG.2020.3030406. 2, 4
- [LFK*13] LIN S., FORTUNA J., KULKARNI C., STONE M., HEER J.: Selecting semantically-resonant colors for data visualization. *Computer Graphics Forum* 32 (2013), 401–410. doi:10.1111/cgf.12127. 2, 3
- [LH18] LIU Y., HEER J.: Somewhere over the rainbow: An empirical assessment of quantitative colormaps. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems* (2018), p. 1–12. doi:10.1145/3173574.3174172. 2
- [LHD18] LANGNER R., HORAK T., DACHSELT R.: Vistiles: Coordinating and combining co-located mobile devices for visual data exploration. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 626–636. doi:10.1109/TVCG.2017.2744019. 1
- [LKD19] LANGNER R., KISTER U., DACHSELT R.: Multiple coordinated views at large displays for multiple users: Empirical findings on user behavior, movements, and distances. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 608–618. doi:10.1109/TVCG.2018.2865235. 2
- [LLW*24] LIN Y., LI H., WU A., WANG Y., QU H.: Dminer: Dashboard design mining and recommendation. *IEEE Transactions on Visualization and Computer Graphics* 30, 7 (2024), 4108–4121. doi:10.1109/tvcg.2023.3251344. 3
- [LSS13] LEE S., SIPS M., SEIDEL H.: Perceptually driven visibility optimization for categorical data visualization. *IEEE Transactions on Visualization and Computer Graphics* 19, 10 (2013), 1746–1757. doi:10.1109/TVCG.2012.315. 2
- [LWZ*22] LI H., WANG Y., ZHANG S., SONG Y., QU H.: Kg4vis: A knowledge graph-based approach for visualization recommendation. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2022), 195–205. doi:10.1109/TVCG.2021.3114863. 3, 11
- [Mit15] MITTELSTÄDT S.: *Methods for Effective Color Encoding and the Compensation of Contrast Effects*. Thesis, 2015. 3
- [MK15] MITTELSTÄDT S., KEIM D. A.: Efficient Contrast Effect Compensation with Personalized Perception Models. *Computer Graphics Forum* 34, 3 (2015), 211–220. doi:10.1111/cgf.12633. 2
- [MSF22] MAULE J., SKELTON A. E., FRANKLIN A.: The development of color perception and cognition. *Annual Review of Psychology* 74 (2022), 87–111. doi:10.1146/annurev-psych-032720-040512. 2
- [MSK14] MITTELSTÄDT S., STOFFEL A., KEIM D. A.: Methods for compensating contrast effects in information visualization. *Computer Graphics Forum* 33, 3 (2014), 231–240. doi:10.1111/cgf.12379. 2
- [MTK*24] MORITA I., TAKAHASHI S., KOMO Y., NISHIMURA S., ARIKAWA M., MISUE K.: Color recommendations based on individual differences. In *Proceedings of the International Symposium on Visual Information Communication and Interaction* (New York, NY, USA, 2024), VINCI '24, Association for Computing Machinery. URL: <https://doi.org/10.1145/3678698.3678701>, doi:10.1145/3678698.3678701. 11
- [Mun14] MUNZNER T.: *Visualization Analysis and Design*. CRC Press, 2014. doi:10.1201/b17511. 3
- [MWN*19] MORITZ D., WANG C., NELSON G. L., LIN H., SMITH A. M., HOWE B., HEER J.: Formalizing visualization design knowledge as constraints: Actionable and extensible models in draco. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 438–448. doi:10.1109/TVCG.2018.2865240. 3
- [NCB*21] NARDINI P., CHEN M., BÖTTINGER M., SCHEUERMANN G., BUJACK R.: Automatic improvement of continuous colormaps in euclidean colorspaces. *Computer Graphics Forum* 40, 3 (2021), 361–373. doi:10.1111/cgf.14313. 2
- [OAH11] O'DONOVAN P., AGARWALA A., HERTZMANN A.: Color compatibility from large datasets. *ACM Transactions on Graphics* 30, 4 (2011), 1–12. 2

- [PFC18] PHAN H. Q., FU H., CHAN A. B.: Color orchestra: Ordering color palettes for interpolation and prediction. *IEEE Transactions on Visualization and Computer Graphics* 24, 6 (2018), 1942–1955. doi:10.1109/TVCG.2017.2697948. 2
- [PMH18] POCO J., MAYHUA A., HEER J.: Extracting and Retargeting Color Mappings from Bitmap Images of Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2018), 637–646. 2
- [pow] Power BI. <https://powerbi.microsoft.com>, last accessed on 22/03/2023. 2
- [QH17] QU Z., HULLMAN J.: Keeping multiple views consistent: Constraints, validations, and exceptions in visualization authoring. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (2017), 468–477. doi:10.1109/TVCG.2017.2744198. 1, 2, 4
- [Rob07] ROBERTS J. C.: State of the art: Coordinated & multiple views in exploratory visualization. In *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization* (2007), pp. 61–71. doi:10.1109/CMV.2007.20. 2, 5
- [RS21] REDA K., SZAFIR D. A.: Rainbows Revisited: Modeling Effective Colormap Design for Graphical Inference. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 1032–1042. 2
- [SCL*22] SHI Y., CHEN S., LIU P., LONG J., CAO N.: Colorcook: Augmenting color design for dashboarding with domain-associated palettes. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (2022). doi:10.1145/3555534. 2, 3, 4, 7, 8, 9, 10
- [SGS*19] SCHLOSS K. B., GRAMAZIO C. C., SILVERMAN A. T., PARKER M. L., WANG A. S.: Mapping Color to Meaning in Colormap Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 810–819. 3
- [SNK*22] SUN M., NAMBU RI A., KOOP D., ZHAO J., LI T., CHUNG H.: Towards Systematic Design Considerations for Visualizing Cross-View Data Relationships. *IEEE Transactions on Visualization and Computer Graphics* 28, 12 (2022), 4741–4756. doi:10.1109/TVCG.2021.3102966. 1, 2
- [SS15] SETLUR V., STONE M. C.: A linguistic approach to categorical color assignment for data visualization. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2015), 698–707. doi:10.1109/TVCG.2015.2467471. 2
- [SSAZ21] SUN M., SHAIKH A. R., ALHOORI H., ZHAO J.: SightBi: Exploring Cross-View Data Relationships with Biclusters. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 54–64. 2, 4
- [SSM11] SILVA S., SANTOS B. S., MADEIRA J.: Using color in visualization: A survey. *Computers & Graphics* 35, 2 (2011), 320–333. doi:10.1016/j.cag.2010.11.015. 2, 3
- [SWD05] SHARMA G., WU W., DALAL E. N.: The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application* 30, 1 (2005), 21–30. doi:10.1002/col.20070. 6
- [SWS19] SMART S., WU K., SZAFIR D. A.: Color crafting: Automating the construction of designer quality color ramps. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 1215–1225. doi:10.1109/TVCG.2019.2934284. 2, 3
- [tab] Tableau. <https://www.tableau.com/>, last accessed on 22/03/2023. 1, 2
- [TdJ14] TENNEKES M., DE JONGE E.: Tree colors: Color schemes for tree-structured data. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2072–2081. doi:10.1109/TVCG.2014.2346277. 4, 7
- [Wea04] WEAVER C.: Building highly-coordinated visualizations in improvise. In *IEEE Symposium on Information Visualization* (2004), pp. 159–166. doi:10.1109/INFVIS.2004.12. 2, 3
- [Wea09] WEAVER C.: Cross-filtered views for multidimensional visual analysis. *IEEE Transactions on Visualization and Computer Graphics* 16, 2 (2009), 192–204. doi:10.1109/TVCG.2009.94. 1, 2
- [Wer38] WERTHEIMER M.: Laws of organization in perceptual forms. doi:10.1037/11496-005. 4
- [WMA*15] WONGSUPHASAWAT K., MORITZ D., ANAND A., MACKINLAY J., HOWE B., HEER J.: Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2015), 649–658. doi:10.1109/TVCG.2015.2467191. 3
- [WMW*16] WALDIN N., MUZIC M. L., WALDNER M., GRÖLLER E., GOODSELL D., LUDOVIC A., VIOLA I.: Chameleon - dynamic color mapping for multi-scale structural biology models. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2016), Bruckner S., Preim B., Vilanova A., Hauser H., Hennemuth A., Lundervold A., (Eds.), The Eurographics Association. doi:10.2312/vcbm.20161266. 2
- [Wu22] WU E.: View Composition Algebra for Ad Hoc Comparison. *IEEE Transactions on Visualization and Computer Graphics* 28, 6 (2022), 2470–2485. doi:10.1109/TVCG.2022.3152515. 1, 2
- [WXL*21] WU A., XIE L., LEE B., WANG Y., CUI W., QU H.: Learning to automate chart layout configurations using crowdsourced paired comparison. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems* (2021), pp. 1–13. doi:10.1145/3411764.3445179. 3
- [WZW*23] WEN Z., ZENG W., WENG L., LIU Y., XU M., CHEN W.: Effects of View Layout on Situated Analytics for Multiple Representations in Immersive Visualization. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2023), 440–450. doi:10.1109/TVCG.2022.3209475. 1, 2
- [YHH*24] YE Y., HAO J., HOU Y., WANG Z., XIAO S., LUO Y., ZENG W.: Generative ai for visualization: State of the art and future directions. *Visual Informatics* 8, 2 (2024), 43–66. doi:10.1016/j.visinf.2024.04.003. 3
- [YHL*25] YANG M., HOU Y., LI L., CHANG R., ZENG W.: Dashboard vision: Using eye-tracking to understand and predict dashboard viewing behaviors. *IEEE Transactions on Visualization and Computer Graphics* (2025), 1–16. doi:10.1109/TVCG.2025.3532497. 2
- [YZF*22] YUAN L., ZENG W., FU S., ZENG Z., LI H., FU C. W., QU H.: Deep Colormap Extraction from Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 28, 12 (2022), 4048–4060. doi:10.1109/TVCG.2021.3070876. 2
- [YZZ*21] YUAN L.-P., ZHOU Z., ZHAO J., GUO Y., DU F., QU H.: Infocolorizer: Interactive recommendation of color palettes for infographics. *IEEE Transactions on Visualization and Computer Graphics* 28, 12 (2021), 4252–4266. doi:10.1109/TVCG.2021.3085327. 3
- [ZCH*23] ZENG W., CHEN X., HOU Y., SHAO L., CHU Z., CHANG R.: Semi-Automatic Layout Adaptation for Responsive Multiple-View Visualization Design. *IEEE Transactions on Visualization and Computer Graphics* (2023), 1–15. doi:10.1109/TVCG.2023.3240356. 2
- [ZH15] ZHOU L., HANSEN C. D.: A survey of colormaps in visualization. *IEEE Transactions on Visualization and Computer Graphics* 22, 8 (2015), 2051–2069. doi:10.1109/TVCG.2015.2489649. 2, 3
- [ZLL*21] ZENG W., LIN C., LIN J., JIANG J., XIA J., TURKAY C., CHEN W.: Revisiting the Modifiable Areal Unit Problem in Deep Traffic Prediction with Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021), 839–848. doi:10.1109/TVCG.2020.3030410. 3