



Mixing Linters with GUIs: A Color Palette Design Probe

Andrew McNutt , Maureen C. Stone, and Jeffrey Heer 

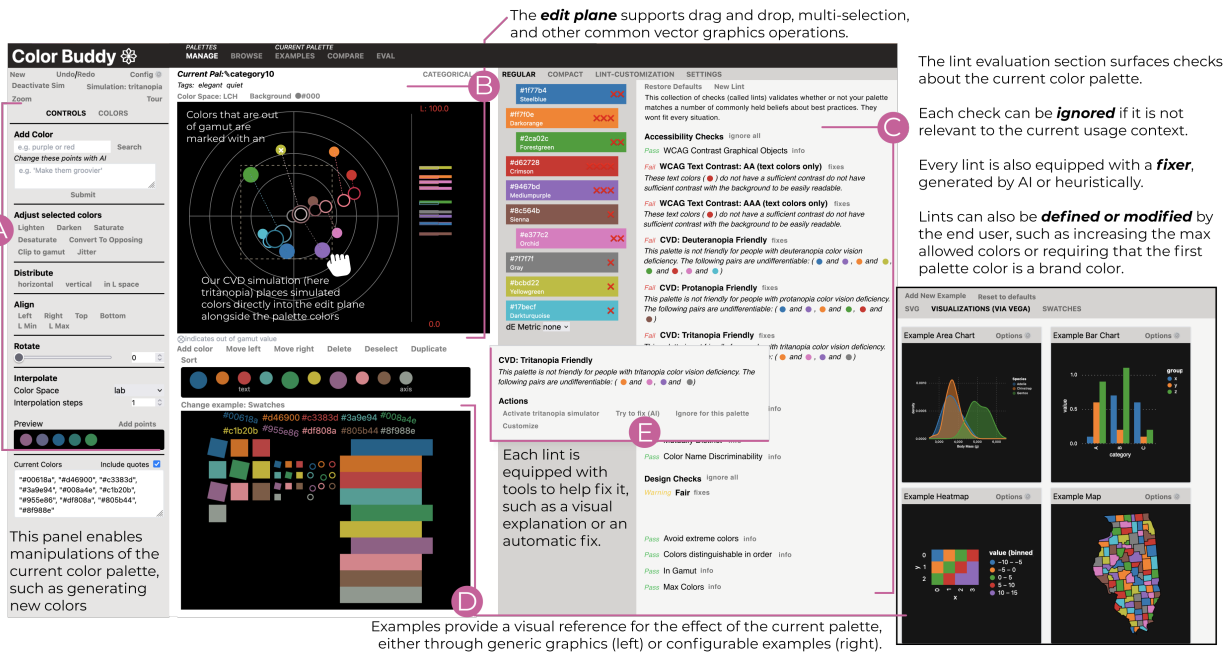


Fig. 1: Our design probe, COLOR BUDDY, explores how a linter can be meaningfully integrated into a color palette design GUI.

Abstract—Visualization linters are end-user facing evaluators that automatically identify potential chart issues. These spell-checker like systems offer a blend of interpretability and customization that is not found in other forms of automated assistance. However, existing linters do not model context and have primarily targeted users who do not need assistance, resulting in obvious—even annoying—advice. We investigate these issues within the domain of color palette design, which serves as a microcosm of visualization design concerns. We contribute a GUI-based color palette linter as a design probe that covers perception, accessibility, context, and other design criteria, and use it to explore visual explanations, integrated fixes, and user defined linting rules. Through a formative interview study and theory-driven analysis, we find that linters can be meaningfully integrated into graphical contexts thereby addressing many of their core issues. We discuss implications for integrating linters into visualization tools, developing improved assertion languages, and supporting end-user tunable advice—all laying the groundwork for more effective visualization linters in any context.

Index Terms—Linters, Color Palette Design, Design Probe, Reflection

1 INTRODUCTION

How do you know what you have done is right? In nuanced domains (like visualization) checking new designs can be challenging. It can require navigation of potentially conflicting design adages, easy-to-miss details, and subtle relationships between disparate components. Reconciling these various guidelines with personal preferences can be difficult, as advice can be conflicting and desires may not be yet fully understood or quantifiable. Various approaches [17] in visualization try to ensure that a given program or chart is correct, or at least not bad. These include checklists, style guides, chart metrics, theories, automated critique, and recommenders. Each approach has trade-offs,

giving agency (and responsibility) to their operator by varying degrees.

Of particular interest are visualization linters [3, 19, 31, 46, 53]. These systems draw an analogy with static evaluation tools from software development of the same name (essentially spell checkers for code). They are interpretable (each *lint rule* checks for a single property), they are customizable (rules can be adjusted or deactivated to taste), automated (simplifying their application), are easily pliable to the grammar-based visualization systems, and can even help evaluate automatically generated code [56]. This unique blend of properties is unavailable in other forms of assistance, such as black-boxed automated recommendations. Despite their potential, linters face several significant issues.

Primarily Text-Based. Linters are a tool for end-users. In visualization, they typically check the static properties of a program written in Vega-Lite [69] or another specification language. While useful, designing linters to target *programs* rather than GUIs excludes all but those skilled in languages like Vega-Lite from this form of assistance. The number of such users is small compared to those that do analysis in tools like Excel or Tableau, leaving many unaided. While there have been explorations of lint-style analysis in GUIs [43, 75], these checks are typically outside of the primary work cycle, excluding interaction between the linter and the rest of the development environment.

- Andrew McNutt is with the Universities of Washington and Utah. E-mail: TODO
- Maureen Stone is with the University of Washington E-mail: mcstone@cs.washington.edu
- Jeffrey Heer is with the University of Washington E-mail: jheer@uw.edu

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

Impolite/annoying. Linters are often seen as annoying. CSS Lint warns its users that it “Will hurt your feelings* (And help you code better)” [91]. It can be a frustrating (or even impolite [84]) experience seeing a huge wall of errors while you are trying to accomplish a task—particularly in formative task stages before such feedback will be useful. While this discomfort can be useful, in that it pushes users to internalize the expressed guidelines, making it unavoidable can lead to rejecting the tool entirely [55]. Yet the basic conversational model that linting supports (in which items are iteratively dealt with, as in a dialog) is clearly beneficial and is reflected in the design of numerous tools (e.g., AI UIs [56]). Strategies for navigating this tension are not well established, likely impeding wider usage.

Un-useful analysis. Finally, lints are valuable for the quality of their analysis, which has so far been limited. Some linters [53, 57] include suggestions to avoid certain chart types (e.g., the debated adage to avoid pie charts [42]), while others [3] merely espouse style preferences, such as avoiding redundant encodings. Similarly, many checked errors are unlikely to appear without the user having made substantially unusual choices, which were likely intentional. For instance, many linters [19, 31, 54] check if the y-axis is pointed in the conventional direction, to prevent issues like the infamously deceptive Florida Gun Deaths chart. Smart defaults typically guide the designer away from this type of design, and so deciding to use it is likely a signal of designerly intent. Finally, with one exception [46], past linters have been agnostic to analytic context (as represented by the noisy signal of chart type); instead focusing broadly on all chart forms.

This work explores these issues via a design probe in the small, yet still intricate, domain of color palette design. Just as how color is seen as a microcosm of visual perception [64], we see palette design as being representative of many of the problems in visualization. Both domains have rich internal structures centered in graphical domains for which there are sprawling collections of guidelines and constraints to navigate. In addition, palette design is a challenging and common task on its own [25], despite a wealth of tool support (e.g., [24, 37, 59, 67, 93]).

To this end, we developed a color palette linter that evaluates discrete color palettes, covering categorical, diverging, and sequential types. It can consider many usage issues, including those related to accessibility (e.g., contrast and color vision deficiency, or CVD), stylistics (e.g., avoiding “ugly” colors), and design (e.g., ensuring that sequential palettes are perceptually sequential). Our linter is tightly woven into its surrounding GUI called COLOR BUDDY (also created for this study), supporting visual explanations, integrated checks, and advice that can be ignored, tuned, or summoned to task and taste. It is powered by a domain-specific language (DSL) that uses quantifiers and predicates augmented with palette-related facilities, such as simulating CVD. These can be created by end-users as *user defined* lints, such as ensuring that a brand color is present. See Sec. 8 for code and a link to a demo. We reflect on our probe in an interview study with design professionals (Sec. 5), as well as a critical reflection [68] (Sec. 6).

We contribute a characterization of the problems, strategies, and opportunities for visualization linters. We present a typology of linter features (Sec. 3) and identify necessary directions for additional research (Secs. 6-7). By focusing on color palette design, we engage with our critiques of prior linters unburdened by a larger breadth of visualization concerns. GUI-based systems can meaningfully weave lints into a core part of the workflow, although their integration needs to be mindful of the different parts of a task being pursued. Having a rich body of theory to draw on supports the evaluation of complex properties—such as effective affective usage and visual accessibility. These findings can be meaningfully returned to visualization—such as by developing theories that support complex assertions. This work lays the groundwork for more effective visualization linters in any context.

2 RELATED WORK

We draw on prior works on supporting work with colors, linting in general, and linting for visualization specifically.

2.1 Color Tools

There have been a wide variety of tools developed to help users pick colors, design palettes, and analyze them. These include a wide range of tasks, covering end-user facing tools [2, 24, 50, 59] to automated palette generation [30, 40, 45, 90]. Nardini et al.’s [62] automated color map improvement is especially closely related to our notion of automated fixes, although we naively handle a generic case whereas they handle rigorously handle several specific cases. Jalal et al. [37] characterize the design space of color pickers, highlighting several axes along which tools tend to operate. We draw on designs from many of these tools, but merely as a setting for our linting explorations. For instance, we employ designs from both mixed-initiative [30, 49, 67] and direct manipulation [2, 63, 74, 76] driven palette construction. In terms of Zhou’s survey of continuous color map studies [93], our work falls into the “rule-based methods” category, à la PRAVDAColor [11].

More closely related to our work, are those studies that consider evaluating palettes. Bujack et al. [16] synthesize the space of guidelines for continuous color maps into concrete mathematical formulas. We note that color palette design involves more nuance than merely being a simpler version of gradients, as the colors often are required to perform a variety of roles (e.g., text or background) in addition to carrying data. Viz Palette [50] checks if colors are discriminable for different sizes and color name discriminability—explicitly drawing on the metrics used in Colorgical [24]. Some tools provide accessibility tools, such as for CVD or contrast—although this support is limited [80]. We draw on these prior surveys to formulate our list of lints, which we check via consultation with professional designers. Closely related, cols4all [79] evaluates palette properties (e.g., does a categorical palette have equally expressive colors), but is focused on palette selection rather than design.

2.2 Linting

In software engineering linting typically refers to tools that identify syntactic or stylistic errors in the text of a program. While these checks are often conducted via command line tools (such as part of a test suite), they are just as often integrated into the editing environment itself, for instance highlighting parts of the code that have behaviors marked as either as *errors* or *warnings*. Each assertion about the code, referred to as a *rule*, checks for a particular behavior or property. For instance, *max-len* enforces that all lines of code be shorter than a predefined maximum. Beyond stylistics, lint-like systems surface type errors, check for anti-patterns (such as duplicated code not abstracted into a function [44]), highlight inaccessible patterns, and security leaks [83], among many other errors that might be easily missed.

A central usability feature of lints is that the end-user can deactivate or modify rules. This can be on the scale of an entire project (e.g., lines must have ≤ 120 characters) or deactivating the rule for a particular line (e.g., allowing a long string one time). *Allowing the user to break the rules when appropriate is essential*, as the task model espoused by a set of particular rules will not cover all use cases.

Linters are sometimes perceived as annoying until a user internalizes the code practices espoused by that linter [46, 55]—which can be a useful way to learn a new domain. Li et al. [47] note similar properties in creativity support tools. However, internalizing a rule is closely related to automation bias, particularly when linters are equipped with automatic fixers. Prompting disengagement from critical thinking can be troublesome, but it can be a useful way to teach best practices.

The idea of linters has been used outside of code. write-good [22] guides its documentation authors toward the active voice, while HaTe Detector [86] checks for harmful terminology (e.g., *blacklist*). Data Linter [34] looks for issues in machine learning data. ExcelLint [4] uses entropic program analysis to identify issues in spreadsheets. This highlights how lints can simplify usage of complex metrics: lints are often configured by experts and employed by users. Our work is related to these but translated to a GUI. Textlets [26] explores constraint enforcement mechanisms for legal writing, such as a maximum word count. Farsight [82] provides a lint-like experience for identifying harms that AI prompts might cause—although they envision it through the lens of usable security-style alerts. These works draw from validation in end-user software engineering, which Ko et al. [41] survey.

Linters	Checkable	Customization	Blameable	Fixable
AI Chains [88]	Yes	No	No	Yes
Chart Linter [3]	Yes	Yes	No	Yes
EVM Check [38]	Partial	No	No	No
GeoLinter [46]	Yes	No	No	Yes
LitVis [87]	Yes	Yes	Yes	No
Mirage linter [54]	Yes	No	Partial	No
Perceptual Pat [75]	Partial	Yes	No	No
VisGrader [33]	Yes	No	No	No
vislint_mpl [53]	Yes	Yes	No	No
Visualint [31]	No	No	Yes	No
Visualization Linter [92]	Yes	No	Partial	Yes
VizLinter [19]	Yes	No	Partial	Yes
PALETTE LINT (this work)	Yes	Yes	Yes	Yes

Fig. 2: Past visualization linters provide varying levels of support for different functionalities central to the user experience of linting.

2.3 Linting for Visualization

Visualization is an especially appealing domain for linter-style tools because its users are rarely visualization experts, there are numerous guidelines and best practices to navigate, and it is easy to make subtle mistakes [54]. First described by Meeks [57] in 2017, *visualization linting*, applies the idea of gentle correctable assistance to visualization. A number of linters were subsequently developed, covering matplotlib (via vislint_mpl [53]), chart.js (via Visualization Linter [92]), Vega-Lite (via Mirage linter [54]), and AntV (via Chart Linter [3]). VizLinter [19] expanded visualization linting to include both fixers as well. These assertions generally cover stylistics or data type usage (e.g., bar charts should use nominal data). Qu and Hullman [66] provide an enriched variation of these models by developing constraints for dashboard-like systems. We draw on the design of these systems in our work but in the domain of color palette construction. Visualint [31] explored the visual presentation of lint failures, finding that in-situ chart annotations aided understanding of issues—a design strategy we echo. Closest to our work is GeoLinter [46] for choropleth maps, which provides more specific guidance than prior linters by drawing upon a richer model of usage context. For example, GeoLinter analyses include topics such as color-aware data binning. Similarly, we narrow our domain scope to provide more useful analyses. VisGrader [33] employs a related testing strategy to validate student submissions in a visualization course, although their tests are bound to the specific assignments in that course.

These methods have been extended beyond the specific framing of linting. Wu et al. [88] explore the use of LLM-based prompt chaining to analyze and critique Vega-Lite programs. We also make use of LLMs, but purely as a source of generative recommendation (following Shi et al. [73]), rather than evaluation. Perceptual Pat [75] offers automated commentary on visualizations via computer vision-based methods. They observe that users like the validation, but struggle with some automation biases. LitVis [87] provides lint-like guidance to the visualization design process (rather than the designed artifact) that prompts users to reflect on the values expressed in their design. EVM Check [38] supports end-user checking of data interpretations via expressions of statistical models, although they do not provide assertions to guide evaluation. We build upon these prior explorations to understand the linter design space.

3 A LADDER OF LINTING SUPPORT

Linters include a variety of feature types, each of which offers a different discrete level of complexity of support forming a *linting ladder*. Here, we identify common patterns of support by comparing prior visualization linters (broadly defined) with common software linters such as eslint [77] or SQLFluff [21].

1. **Checkable:** *Assertions can be evaluated*
2. **Customizable:** *Checks can be deactivated, locally or globally, and their parameters can be adjusted*
3. **Blameable:** *The cause of the issue can be identified*
4. **Fixable:** *Automated fixes can be given*

Lint
Affect: Dark reds and browns are not positive [5]
Affect: Negative palettes should not have light colors, particularly greens [5]
Affect: Playful affects can have light blues/beiges/grays [5]
Affect: Saturated not appropriate for (calm serious trustworthy) [5]
Avoid Extreme Colors [60]
Avoid Tetradic Palettes [60]
Avoid Too Much Contrast with Background [60]
Axes Should have Low Contrast with Background [6]
Background Desaturation Sufficient [60]
Blue Should Nameable as Blue [27]
CVD: (Deuter Prot Tritanopia Friendly [78]
Color Name Discriminability [27]
Colors Distinguishable in Order [16]
Diverging Palette In Appropriate Order (Definitional)
Even Distribution in (Hue Lightness) [16]
Fair [79]
In Gamut [16]
Color Distinctness: (Thin Medium Wide) Size Objects [24]
Max Colors [60]
Mutually Distinct [16]
No ugly colors [48]
Prefer yellowish or bluish greens [60]
Require color complements (Color Theory)
Right in black and white [14]
Sequential Palette In Appropriate Order (Definitional)
WCAG Contrast Graphical Objects [78]
WCAG Text Contrast: (AA AAA) [78]

Fig. 3: The lints included in our design probe. Some guidelines come from multiple sources, but we prefer surveys for referential simplicity. Lints with common structures are merged and represented as (a | b).

A check from a software linter like eslint would satisfy each of these rungs. For instance, the check *max-len* would (1) check if the property is violated—i.e., it evaluates if all lines of code are less than a maximum allowed length. Offer the facility to be ignored if necessary (2), such as when using a long configuration string. Describe where the error occurred, such as by highlighting the failing line in an IDE (3). And finally offer an automated fix (4), such as inserting line breaks. These affordances allow linters to traverse the full range of Ceneda et al.’s guidance taxonomy [17]. A lint can help orient (e.g., this is something to watch for), direct (e.g., this issue must be fixed), or prescribe (e.g., this automatic fix is right). This underscores their potential: users can opt-in to however much help they want in a given moment.

In Fig. 2 we categorize how all known visualization linters (or closely related systems) satisfy each of these rungs of support. For instance, early linters (e.g., vislint_mpl and Chart Linter) do not provide detailed feedback. Most linters offer some degree of end-user customization, although it is often limited. Visualint and the Mirage linter explore visual explanation (i.e., blame) of chart issues, although the latter is unable to identify specifically why a particular aspect failed. VizLinter and GeoLinter have fixers that can automatically address identified issues (both of which observed substantial automation bias associated with those fixers). Contrastingly, EVM Check and Perceptual Pat are closely aligned with linters but do not actually make testable assertions—rather they are tools the user is expected to use to make mental comparisons. LitVis has the full tooling support of VSCode, but its reflection-driven lints can not be automatically fixed. Notably, most visualization linters skip the customization and blame rungs, forgoing essential parts of what makes linters effective end-user tools.

Additional rungs could be added to this ladder. For instance, counterfactual (or what-if) analysis might be added after fixable. For instance, Farsight [82] generates *expected harms* for identified issues in AI prompts—although this is the sole example of this type we found. Our ladder is based on trends in observed linters (in visualization and otherwise). We leave later ladder expansion to future work.

4 DESIGN PROBE: PALETTE LINT

The central contribution of this work is an exploration of the structure of visualization linting as seen from the microcosm of color palette design. To support this goal we conducted a design probe in which we

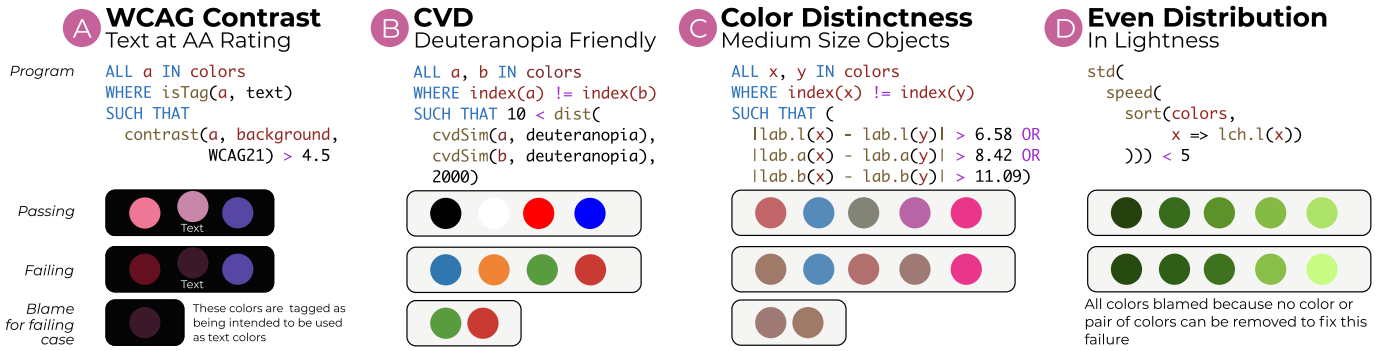


Fig. 4: Example lints in PALETTE LINT’s SQL-like summary text. (A) checks if text objects meet contrast standards for text legibility. (B) compares CVD simulated color pairs using a dE_{2000} heuristic [81] that palettes should be glanceably different. (C) uses values from Gramazio’s d3-jnd [24] to test colors for distinctness as a function of size. (D) checks for even distribution via Bujack et al.’s [16] notion of local speed (i.e., sequential differences).

$Palette = [Color] + Background\ Color$
 $\left\{ \begin{array}{l} \text{Type} \\ \text{Context Tags} \end{array} \right. ::= \{Categorical \mid Sequential \mid Diverging\}$
 $\left\{ \begin{array}{l} \text{Type} \\ \text{Context Tags} \end{array} \right. ::= \{Calm, Exciting, Scatter\ plot, \dots\}$
 $Color = colorSpace(\dots coordinates)$
 $\left\{ \begin{array}{l} \text{Connections} \\ \text{Semantic Tags} \end{array} \right. ::= \{tint(Color), tone(Color), complement(Color), \dots\}$
 $\left\{ \begin{array}{l} \text{Connections} \\ \text{Semantic Tags} \end{array} \right. ::= \{Brand, Accent, Axis, Blue, \dots\}$
 $Lint = Program\ (Fig.\ 6) + Required\ Context\ Tags$
 $+ Presentational\ Metadata + Allowed\ Types$

Fig. 5: Our model of color palettes and the lints that evaluate them.

developed a linter for color palettes called PALETTE LINT. The core design loop was driven by a months-long dialogue between the first and second author (an expert in color design). It is supported by models (Fig. 5) of the domain and of the assertion space. It is housed in a GUI called COLOR BUDDY (Sec. 4.1). Its design draws on previous works, intentionally mirroring our linting ladder (Sec. 4.2).

PALETTE LINT’s purview is reusable discrete color palettes, as in Tableau 10 or corporate style guides. Our task is not about picking colors, but instead about designing color palettes. The particular values are less important than the relationships between colors and so most of the lints do not check if a specific color is included, but such analyses are possible. This includes evaluation of properties like accessibility (e.g., CVD friendliness), usability (e.g., sequential palettes being perceptually sequential), and design guidelines (e.g., 10 colors max).

We collected 35 lints (Fig. 3) via an informal survey of color palette design tools that provide checks (e.g., cols4all [79], Colorgical [24], and VizPalette [50]), suites of guidelines about color maps [16] and palettes [60], as well as color perception studies [5, 27]. While we try to be comprehensive in this collection, we note that (like any group of lints) it will be incomplete, although a rigorous survey would be useful future work. Beliefs about effective design will change, accessibility measures will change [78], and new use cases will arise. Lints address this variance by modularity: new checks can be introduced as needed.

In essence, PALETTE LINT is a simple interpreter that takes in a program (a lint) and data to operate over (a palette). The result of executing the program for that data is a boolean, indicating pass or fail. Each lint also includes metadata that assists with user understanding of that result, such as a string explaining the failure, the seriousness of the error, and so on. We include automatic fixers that draw on this information (as well as data and program) to generate fixes. Next, we detail our model of palettes and then our model of assertions.

Domain modeling. At the core of our linter is a model (per Fig. 5) that shapes the space of possible expressions and usages that are possible in our probe. This modeling is analogous to the task modeling that drives visualization tools like Vega-Lite [69]. Each palette is made up of a list of colors. Both colors and palettes include contextual properties.

The palette properties are influenced by their intended usage context. For instance, the **Type** partition captures common [25] intended usages. Our initial set of lints has 4 unique *Categorical* lints, as well as a unique lint for *Sequential* and *Diverging*. Outside of these basic usages,

$Expressions$
 $\left\{ \begin{array}{l} \text{Expressions} \\ \text{Quantifiers} \end{array} \right. e ::= b \mid \neg e \mid e \wedge e \mid e \vee e \mid q \mid \oplus_b^{v,v} \ v \ v$
 $\left\{ \begin{array}{l} \text{Expressions} \\ \text{Quantifiers} \end{array} \right. q ::= (\forall \mid \exists) \ var_i \in (arr \ st \ e), \ e$
 $Values$
 $\left\{ \begin{array}{l} \text{Values} \\ \text{Arrays} \\ \text{Variables} \end{array} \right. v ::= b \ (bool) \mid c \ (color) \mid n \ (num) \mid var$
 $\left\{ \begin{array}{l} \text{Values} \\ \text{Arrays} \\ \text{Variables} \end{array} \right. \mid \oplus_n^{n,n} \ n \ n \mid \oplus_v^{arr} \ arr \mid \oplus_n^{c,c} \ c \ c \mid \oplus_v^c \ c$
 $\left\{ \begin{array}{l} \text{Arrays} \\ \text{Variables} \end{array} \right. arr ::= [v] \mid \oplus_{arr}^{e, arr} \ e \ arr$
 $\left\{ \begin{array}{l} \text{Arrays} \\ \text{Variables} \end{array} \right. var ::= arr \mid v$
 $Operations \ (\oplus_{input\ types}^{output\ types})$
 $\left\{ \begin{array}{l} \text{Val Ops} \\ \text{Array Ops.} \\ \text{Color Ops.} \end{array} \right. \oplus_b^{v,v} ::= = \mid | \mid < \mid > \mid \sim \mid \dots$
 $\left\{ \begin{array}{l} \text{Val Ops} \\ \text{Array Ops.} \\ \text{Color Ops.} \end{array} \right. \oplus_n^{n,n} ::= * \mid + \mid - \mid \% \mid absDiff \mid \dots$
 $\left\{ \begin{array}{l} \text{Val Ops} \\ \text{Array Ops.} \\ \text{Color Ops.} \end{array} \right. \oplus_v^{arr} ::= count \mid sum \mid min \mid max \mid std$
 $\left\{ \begin{array}{l} \text{Val Ops} \\ \text{Array Ops.} \\ \text{Color Ops.} \end{array} \right. \mid mean \mid first \mid last \mid extent \mid \dots$
 $\left\{ \begin{array}{l} \text{Val Ops} \\ \text{Array Ops.} \\ \text{Color Ops.} \end{array} \right. \oplus_{arr}^{e, arr} ::= map_\theta \mid sort_\theta \mid filter_\theta \mid speed \mid \dots$
 $\left\{ \begin{array}{l} \text{Val Ops} \\ \text{Array Ops.} \\ \text{Color Ops.} \end{array} \right. \oplus_n^{c,c} ::= dist_\theta \mid dE_\theta \mid contrast_\theta$
 $\left\{ \begin{array}{l} \text{Val Ops} \\ \text{Array Ops.} \\ \text{Color Ops.} \end{array} \right. \oplus_v^c ::= toSpace_\theta \mid cvdSim_\theta \mid name_\theta \mid isTag_\theta$

Fig. 6: PALETTE LINT’s abstract syntax. Each lint consists of a (potentially) nested single boolean expression, e . Palette values (e.g., background and colors) are accessed during evaluation as variables. θ denotes function parameters or arguments, such as contrast algorithm in Fig. 4.

palettes include a huge variety of different contexts, such as being of a *calm* affect or being used in a *scatter plot*. Rather than model all such usages, we instead include a notion of **Context Tags**. These sets of properties act as filters, curating what lints should be applied. To wit, only including the rule *Affect: Saturated Colors Not Appropriate for Calm* palettes [5]) to palettes marked as calm. We demonstrate this idea through 6 tags each with a unique lint. The background color holds a special place in the palette because it defines the visual context in which the colors are used. Other roles might be elevated to this level of importance (e.g., text color), but as these will not appear in every palette, we push such specification into our color model.

The colors that make up the palettes are lightweight extensions to familiar models of color (consisting of color space and coordinates) and some metadata. Colors can serve a variety of different purposes in a palette depending on context and user. Just as with palettes we instead capture this through a polymorphic notion of **Semantic Tags**. These capture properties like “fixed brand color”, being a type of blue, or being an axis or text color. Among these, if there is a straightforward representation for that tag, then we include it in the edit plane scatter plot—such as “text” changing the text color and “axis” changing the axes in Fig. 1. These properties are accessible from within the lints, supporting assertions like *Colors Marked as Blue Should be Namable as Blue* or *Brand Color Should Be First*. Palettes often have internal structure, such as from color harmonies like complements or tints. We model this as **Connections** for completion but forgo implementation as it bears little impact on the analysis of our probe. Related properties can be measured via rules like *Require Color Complements*.

Assertion Modeling. To support our linting we developed a formalized representation of the process of making assertions about color palettes through a lightweight domain-specific language (DSL) synonymously called PALETTE LINT. It consists of quantifiers and predicates that operate over palettes and related functions, such as *simulating CVD* ($cvdSim$), or identifying the English name for a color (*name*). Each lint includes an assertion as a program (Fig. 4), metadata (e.g., description, error/warning), and what parts of Fig. 5 it applies to. Fig. 6 shows the abstract syntax while the docs show the concrete JSON syntax as well as the implementation of each check (Sec. 8).

This collection of color operations was selected to support the 35 lints we identified, providing some external validity of this function domain and the expressivity of this language—see language docs, (Sec. 8) for implementation. Using this collection of operations, PALETTE LINT can model all checks that designers noted in our interview study (Sec. 5). More operations could be added, but this set is sufficient to demonstrate the types of problems this assertion model can address. PALETTE LINT also includes functions for manipulating palettes modeled as arrays, such as aggregate (\oplus_{arr}^{arr}) and unary array operations ($\oplus_{arr}^{e, arr}$). In addition to familiar functions like filter or sort, $\oplus_{arr}^{e, arr}$ includes an analog to Bujack et al.’s [16] local speed operator. This supports checks like *Even Distribution Among Lightness* via assertions like Fig. 4D.

PALETTE LINT can not express every possible check. For instance, Bujack et al.’s [16] notion of triangle side difference for intuitive order measure can not be expressed because user defined local variables are not supported. This can not be correctly modeled [15] by naively checking for channel monotonicity as we do in our naive proof-of-concept *Sequential Palette In Appropriate Order*. This is a weakness of our design, but offers an approximation that can be ignored by the user when incorrect. Similarly, it is unclear if PALETTE LINT can test for rainbow or cyclic color maps. However, as these maps are primarily problematic for gradients, this case is not critical to our context.

4.1 GUI: COLOR BUDDY

Linting is not a tool that is typically used alone. In software engineering, lints are used as part of a suite of different tools that include systems like interactive development environments, unit tests, documentation, and autocomplete. Similarly, color palette design does not occur in isolation, and so we weave our linter closely into its containing GUI, COLOR BUDDY. Here we detail our probe and highlight several ways in which we have integrated our linter into our GUI.

System Details. COLOR BUDDY is a web-based GUI application for color palette design. Its design draws on vector graphics tools (e.g., Illustrator) to utilize designer’s familiarity with that tool genre—analogue to how many visual analytics tools lean on data analysts’ familiarity with Tableau-like systems. Palettes are created or modified via tools on the left side of the UI. New colors can be added by clicking on the edit plane (Fig. 1B). Color values can be adjusted by dragging them in the edit plane or by using a WIMP-style drag-box. Following vector manipulation tools, we include tools like align and distribute (Fig. 1A) as well as palette-specific ones like saturate or lighten. The UI supports opportunistic reuse [13], analogue to how a software developer might find a piece of code on StackOverflow and then adjust it to their purposes. To wit, a user might transform a palette from our built-in library (●●●●●●●●) to taste by rotating it to be (●●●●●●●●). Next, palettes are evaluated on the UI’s right side. This includes several approaches to evaluating palettes. There is a means to view the palette on relevant images or visualizations (Fig. 1D), to compare palettes (see appendix), and to evaluate them (our linter). Each assessment approach comes at different stages of the task cycle. For instance, receiving granular lint-level assessment might not be relevant while the user is still exploring the space of possible options for their task—where example images might be more useful. This allows users to select their desired form of feedback at a given moment. The supplement (Sec. 8) documents our probe in greater detail.

Visual Explanations. Next, we observe that providing a tight visual connection between the lints and the rest of the application offers opportunities to provide wider-ranging explanations. For instance, the

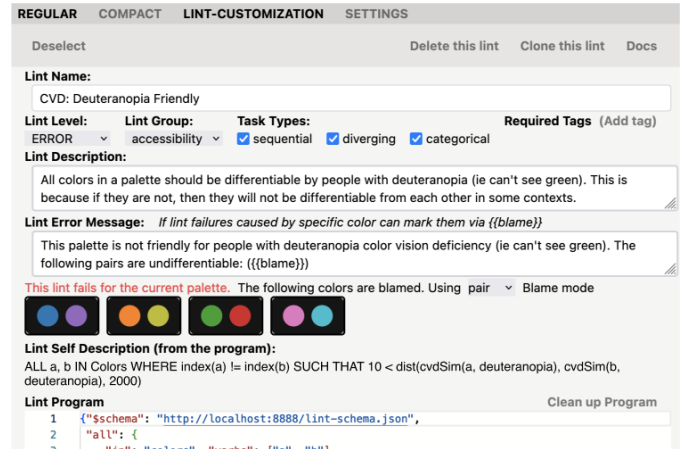


Fig. 7: PALETTE LINT lints can be configured or created by filling out a form within the COLOR BUDDY UI.

blames include visual renderings of the colors suggested as being at fault for a lint failure (as in Fig. 8). Clicking on these circles selects them in the rest of the interface, setting them up for manipulation and correction. We augment this by providing visual explanations of failing lints wherever possible via heuristics. For instance, if a lint description notes a type of CVD then we provide a button to activate the appropriate CVD simulator (as in Fig. 1B). Alternatively if a lint description characterizes something as being an error in a specific color space (e.g., *Even Distribution in Lightness* mentions LCH) then we include a button to switch the current palette to that space. While these are limited, they provide straightforward ways to find support within the tool. Future enhancement might involve automated counterfactuals or lint program visualization.

User Defined Lints. A list of pre-defined lints will rarely capture all of the constraints a user might require. For example, a designer may wish to enforce the presence of a specific brand color but allow for some leeway for perceptually similar shades. This flexibility is central to the philosophy of linting: rather than trying to completely model user intent, we can instead provide consistent, simple, and interpretable signals to shape one’s work. We approach this issue by developing *user defined lints*—assertions defined by the end-user to achieve particular goals. We support user defined lints via end-user manipulable linting programs. Users can manually create a new lint via a form-like UI (Fig. 7), covering the configuration of appropriate task types, affect types, logical PALETTE LINT constraints, and so on. These programs support a range of activities and intents, and allow for far more specificity than our initial list of lints (which are themselves implemented using the same mechanisms as any user defined lint). Lints can also be customized by adjusting parameter values, such as the color contrast algorithm or discriminability thresholds. Our probe does not support parameter adjustment outside of full program editing, but in the future, something like Ivy [51] might simplify lint configuration.

Visual Evaluation. Finally, we observe that other forms of GUI-centered evaluation are possible and provide valuable support in this domain. For instance, if a palette will only be used to create a particular chart type, then the user will be better served by evaluating how it operates on that chart type rather than all charts generically. We support this style of analysis by enabling the user to upload example graphics either specified in SVG or Vega/Vega-Lite that are dynamically recolored—similar to Adobe Color [2] or Höggräfer et al.’s [29] modified version of ColorBrewer respectively. The compare tab (see appendix) further supports visual evaluation via an un-editable second scatter plot that allows comparison between palettes (such as with a reference palette). While reasonably standard [37] these approaches dovetail with the concrete representation of correctness espoused by lints, filling in those gaps when the user is the only appropriate oracle. In future work, ideas like linted reflections [87] or other human-as-oracle techniques like lineups [85] could be usefully incorporated.

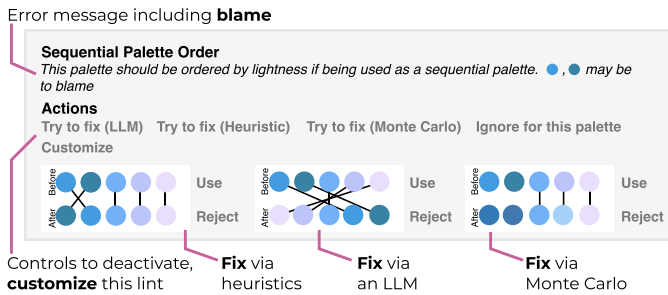


Fig. 8: Each lint comes with several forms of assistance, including checks, blames, fixes (Try to fix), and customizations (Ignore/Customize).

4.2 Climbing the Linting Ladder

Finally, we give a tour of PALETTE LINT via the lint ladder.

Checkable. The first rung of the ladder describes that a property can be evaluated as being true or false, which for us is captured by execution of PALETTE LINT programs. A program is parsed into an abstract syntax tree following Fig. 6. The program is executed via a traversal of the tree. On the downwards pass variables are inserted into the evaluation environment such that operators deeper in the tree can use them. During a loop—as in Fig. 4B’s ALL quantifier—the child nodes (WHERE and SUCH THAT) see versions of the execution environment containing the values of the variables (i.e., a and b), as well as their positional indexes (index(a) and index(b)). The results of these operations are bubbled upwards as return values, with each node eventually collapsing into a boolean expression, yielding a single result for the program. Visually the results of these lints are presented to the end-user as a list, as in Fig. 1C. Each lint is shown as passing, failing, or warning.

Customizable. By reifying our lints into the user interface, we make them straightforward to customize and adjust. Customization can happen either by modifying it (Fig. 7) or by ignoring it (Fig. 8). Modifications include adjusting the description of the lint, the content of the PALETTE LINT program, or the conditions in which it is run (e.g., setting it to be only run on sequential palettes). For instance, the user could raise the number of allowed colors on *Max Colors*. A critical decision to make in designing a linter is what level of granularity will be used which is connected to the depth of customization available. For instance, VisGrader [33] only evaluates specific charts in specific ways. This is well aligned with its purpose of automatically grading d3 charts, but makes those assertions hard to re-use. For us, the unit of input is the whole palette and the relationships contained therein, rather than individual colors, as the latter are frequently adjusted, reordered, or manipulated. Lints can be deactivated per palette or for the entire application. This sets our focus only on analyzing palettes rather than their usage context. This means that we do not evaluate specific charts (as in Lee et al.’s [45] palette optimizations for choropleths) or ensure that colors are semantically resonant with a dataset (à la Setlur et al. [71]).

Blameable. Without a description of why a lint has failed, it can be difficult to diagnose and fix them. Our language automatically generates a description of what colors caused a lint failure, analogous to how code lints would point to a relevant line of code. These descriptions, or *blames*, are surfaced to the end-user in the UI as in Fig. 8. We do this via a simple brute force algorithm that focuses on blaming either single colors or pairs. Respectively these support checks like *all colors should stand out from the background* and *CVD Deuteranopia: Friendly*. Our algorithm iterates across the color palette and checks if a new subset palette (consisting of just a single color or color pair) causes the lint to fail. If this initial “constructive” approach fails, then a “reductive” approach is applied in which each item or pair is removed from the current palette. Multiple colors can be blamed for a failure and so this algorithm must run to check all colors in the palette. This algorithm could be improved by instead doing a tracing analysis in such a way that the errors can be identified analytically [1]. While this would no doubt be faster and more robust, it is out of the scope for this work’s focus on mixing linters with GUIs.

Fixable. Finally, there are two automated means of automatically fixing failing lints. As in Fig. 8, when presented with an error the user can select which of these three methods (labeled “LLM”, “Heuristic”, and “Monte Carlo”) they would like to use. The LLM fixer uses external LLMs (Anthropic Haiku and GPT 4o in spring 2024) to generate fixes. Each lint generates a failure message, which is fed into an LLM along with the current palette. The LLM proposes an update based on the prompt, which the user can accept or reject. The heuristic fixer makes use of fixes defined outside of PALETTE LINT for some built-in (7) lints—e.g., clipping to the hex color gamut for *In Gamut* failures. The Monte Carlo-inspired approach iteratively perturbs colors blamed for a failure until a pass is identified or a time-out is reached. All three approaches do not always fix the error, but we find that Monte Carlo generates passing repairs more frequently than other methods, but update quality varies. See appendix for a performance comparison and more details. Each fixer fixes a single lint error, which can result in cycles in which fixing an error causes another error, such that auto-fixing the second error causes the first one to reappear. Tools like VizLinter [19] try to fix this through solvers. We forgo a “fix all” strategy to limit the already substantial automation bias latent in presenting a fixer, although this approach may still impede creativity.

5 INTERVIEW STUDY: ECOLOGICAL VALIDITY CHECK

We conducted semi-structured interviews with five design professionals (denoted P_x and “quoted”) who have previously made color palettes for visualization. In particular, we sought to answer the questions: *is our usage model appropriate?* and *is our collection of lints aligned with perceived real-world usage?* These interviews covered their experience making palettes and elicited their initial reactions to our probe. Participants were identified from our social networks and were invited to participate via email. They were qualified to participate if they could be identified as having developed data color palettes in a professional capacity. Interviews were conducted over Zoom and lasted about 60 minutes. Participants were not compensated.

Response to our probe was generally positive, with P_1 noting that “we’d love to have had it” for a recent project. These responses may be biased by our connection with participants, so we do not take their responses as evaluation. Instead, we note that their expressed usage model and expected checks are broadly in agreement with our probes.

Usage Model. All participants observed that designing high-quality palettes is a complex, iterative process involving numerous steps. It requires balancing many tradeoffs to achieve a result that is functional, aesthetic, accessible, and satisfies the stakeholders. P_2 summarized this view noting that “the color design process is very messy.”

Participants describe palette design as consisting of a series of cycles. Constraints are set before the work, such as requiring that a client’s brand colors be used. The palette is then drafted around these constraints using a mixture of color pickers or palette generators (e.g., Colors [12]). Next, it is evaluated on sample graphs, relevant to the current domain. These might be found via client elicitation—e.g., by asking “are they never making maps?” (P_5)—or by using “a pretty large test suite” (P_1) for an in-house tool. Then another (typically external) set of tools is used to check for accessibility—e.g., contrast and CVD. Finally, it is given to stakeholders for feedback. Each of these steps are iterated many times to incorporate comments and balance conflicting requirements. P_1 observed that a good palette can take months, and P_3 noted that “it takes a lot more time than it should.”

This model is closely related to the iterative cyclic checks and evaluation on examples housed in our probe. P_1 , P_3 , P_2 , and P_5 valued consolidating the many components of palette design in a single tool, if for no other reason than that it minimizes tool switching. P_3 noted that “the more tools I have to jump to the more frustrated I get”. This consolidation offers synergistic integration between elements, for instance allowing visual explanations of lint failures (e.g., our CVD simulation).

However, our approach did not fit all usage models. P_4 noted that she “can’t usually play around with the colors” as they are usually predefined by the client. This meant that assembling a categorical palette was similar to evaluating various combinations of colors from

predefined sequential palettes. Based on our interviews and other tools (e.g., Colors [12], Viz Palette [50]) this model seems to be less common, however our probe could be modified to allow this more discrete form of design. Participants noted that palettes are often components of color systems, and so can not be evaluated by themselves as our system assumes. To wit, in visualization this includes axes and gridlines (P_3), text labels (P_5), other palettes (P_1), and so on. We address some of these via Semantic Tags, which allow us to evaluate components independently (as in *Axes Should have Low Contrast with Background*), but managing large collections of lints is non-trivial.

Expected Evaluations. Participants were generally enthusiastic about having lints specifically for color palettes. To wit, P_1 , P_4 , and P_5 had developed or modified ad hoc tools to simplify parts of the checking process—highlighting that our lints may be able to fulfill an unmet need. P_1 and P_5 stressed that our lints serve as an educational resource, with P_5 noting that a lay user “could learn all the key things they need to know, but they don’t know”. Our lint’s reporting style is not unfamiliar, with all participants comparing it to prior tools, such as Viz Palette’s [50] Color Report and contrast tools [35, 61]. P_5 suggested that our probe surpasses those works by offering automated fixing.

Participants expressed several different properties as being important to check. All participants expressed the need to pay attention to and include brand colors. Similarly, all participants highlighted the need to consider CVD in their designs. P_2 and P_4 emphasized the importance of contrast compared to CVD, as the latter is rarer but more studied due to its prevalence in white men. P_1 , P_3 , and P_4 had considered ensuring that colors looked good on different devices but did not have a formal means of evaluating, other than trying out those devices. We capture a related quantity to this via *Out of Gamut*. P_4 and P_3 sometimes check grayscale to circumvent this issue, which we check via *Right in Black and White*. P_1 noted that he tried to ensure that colors had straightforward name discriminability, which we assess via *Color Name Discriminability*. When asked to ideate on potential checks, participants augmented this list with checks for other visual disabilities such as “astigmatism” (P_3), client preferences about particular colors (P_5), color harmonies (P_3), and interactions with specific fonts (P_5). P_3 went on to offer several guidelines, such as “try not to do white backgrounds on mobile, because it’s just glaring” and “stay away from pure black”. PALETTE LINT supports many of these suggestions, such as via tags like “mobile” and creating purpose-specific lints. Others are out of gamut (such as astigmatism or font interaction) because there is not a sufficient model to draw on to test those properties.

Our reification [10] of lints as end-user configurable entities seems to be well aligned with participant expectations. For instance, P_3 highlighted the need to be able to adjust the values in a rule like *No Ugly Colors* to address situational changes, such as audience or background color. Even straightforward lints like contrast checkers might necessitate adjustments: while P_2 stressed the value of WCAG (referring to the contrast model used by the Web Content Accessibility Guidelines [78]), P_4 preferred APCA [61]. Rather than taking an opinionated stand about which algorithm is most effective, our probe supports all common ones. Similarly, the natural modularity of linting seemed to match expectations about palette design, with P_4 observing that “it’s definitely an evolution, and I’m sure it’s going to change a lot. We’ll have a different practice in like five years, even hopefully, it changes.”

Of course, there are limits to lints in this context. For instance, the guidelines they implement are only effective if those guidelines are expressible. To wit, P_3 noted that she had “a very hard time articulating this as the general rule” for what makes a palette effective. Using something like the espoused-preference model explored in CIERAN [30] may be valuable, although this may forgo explainability. P_5 observed that color usage is situational and impedes general-purpose analysis, musing “I’m gonna put text inside the bubbles in my scatter plot, do I need to use white text or black text?...that doesn’t seem like a simple check.” Our model of evaluating palettes, rather than the in-situ effect of palettes, precludes this type of contextual analysis. Future work might explore a richer model of color interaction, but our simple model seems appropriate to probe visualization linters.

6 ANALYSIS: TECHNICAL DIMENSIONS OF VIS LINTERS

Next, we use our probe to reflect on visualization linters via a critical reflection [68]. We structure this discussion via the Technical Dimensions of Programming Systems [36] (TDoPS), which is a descriptive framework that considers how programs are used and created across clusters of dimensions (*Denoted* and *characterized*). We use TDoPS because its focus on the UI around programs is aligned with the artifact-as-program perspective elicited via the linting metaphor.

Interaction & Feedback Loops. Which loops in the system are overlapping and how far apart are the corresponding gulfs of evaluation? An essential missing piece from many prior visualization linters is a model of the relationship between the linter and the rest of the workflow: when is it used, how does it connect to its surroundings, and so on. Our probe explored an integration that emphasizes connection to the surrounding context in a way that supports palette design and, per our interview, is closely aligned with the natural palette design cycle our participants use. For example, connecting lints to both explainable notions of blame and subsequent fixes can make it easier to address a lint violation. Considering the context of lint usage is thus valuable and can guide the types of checks to include. In contrast, prior visualization linters have had task models like “editing a Vega-Lite program”. While this is a common task, it is rarely an end unto itself. It may be more effective to consider lints that arise in larger tasks, such as data exploration (e.g., data ambiguity errors) or presentation (e.g., stylistic judgments).

Within such contexts, the feedback loop tightness is generally prescribed by the medium. Code linters are understood as offering nearly instantaneous feedback, although the timing of that feedback has been observed as being an important component of the perceived annoyingness of the design [55]: too fast and they will be perceived as abrasive and over-eager, too slow and they will be seen as out of touch. Our probe centers on nearly real-time feedback, as whenever a change is made the lints quickly update. This allows for tight cycles of viewing and adjustment, which can be useful for internalizing the rules [55]. In contrast, it may be useful to explore slower feedback for some topics, à la Perceptual Pat [75]. Slow checks occur naturally as a part of any design process (for color or visualization), such as by talking to stakeholders (as noted by all our participants). Characterizing the design space of such checks is valuable future work for expanding linter utility.

Errors. What are they and how are they handled? Surfacing errors in the domain of analysis is the core of the linting interaction loop.

The quality of our lints is supported by the quite substantial body of knowledge from color perception and color design practice. For instance, being able to check for an *Even Distribution in Lightness* (Fig. 4) draws from the well-foundedness of the luminance and Bujack et al.’s [16] formulation of local speed. Like many before us [18], we highlight that there is a need for greater theoretical development to make comparable assertions in visualization. Developing an evaluative theory operable by humans (for interpretability) and computers (for automation) is important for making linters that can analyze more complex properties in more contexts (e.g., overplotting or readability). Such theories are not necessary for assembling sets of effective lints, but may be useful for identifying new ones and managing those collections.

Like prior work, we naively surface lints to the user as a long list of errors as in Fig. 1C. This can be potentially overwhelming, which P_1 summarized by noting “there’s a lot going on here”. In contrast, code linters do not present a wall of annoying errors within an IDE, in part because the edit plane of the IDE is so large. A color palette (or a visualization) is a relatively small visual object compared to a code base, and so the errors can crowd the presentation. Prior in-situ solutions [31] likely do not scale beyond a couple of errors. For instance, stacking multiple CVD simulations (Fig. 1) could quickly lead to clutter. Developing means of ranking and hierarchically presenting errors is an important area for future work.

Beyond linting’s natural automation biases (particularly those with fixers [46]), these checklist-like presentations can give rise to perceived false negatives: if no issues appear then the user may falsely equate this with a conclusion that there are no issues with the palette. We contend that it is critical to calibrate user’s trust in the lints so that they feel

empowered to disregard lint advice. This is similar to how scaffolds in creativity support tools help users bootstrap knowledge until they are ready to creatively break the rules [47]. Our probe attempts this with text labels that highlight that lints are imperfect (Fig. 1C top), but the efficacy of these strategies remains unassessed.

Notation. *What notations are present and how do they interrelate?* In this and other linters, there are two primary notations: that of the thing being analyzed (visualization, code, color palettes) and that of the mechanism of analysis and assertion (typically heuristic checks, AST traversals, and our PALETTE LINT respectively). Between these two families, assertions are typically not given as much attention. Our exploration of the form of that assertion structure is a component of our contribution—in particular showing that a bespoke domain-centered model of assertions is possible and sufficiently expressive to capture many rules. To that end, while we found using the language of quantifiers and predicates to be effective, metaphors like search (à la SQL) or if-thens (à la RulePad [58]) might have been successful. The most closely related prior approach is the use of Answer Set Programming (ASP) in VizLinter and the recommender Draco. Prior work [89] has noted that ASP can be difficult to develop and understand, sometimes necessitating secondary tools for interpretability [70]. PALETTE LINT’s assertion style operates differently than ASP—in that modeling and assertion are fully distinct, rather than being interwoven through ASP modeling—but it is not yet clear if this proves more ergonomic than ASP. We highlight that this reification [9, 10] of an assertion language opens interesting new doors to wrapping user interfaces around these programs in easy-to-understand ways. We highlight usability evaluation of these languages as important future work.

Conceptual Structure. *What is the shape of the notations at play and how do they relate to user goals?* The structure of a linter is a simple one: something is checked as being right or wrong. We found that this model works well with color palettes specifically and graphical interfaces in general, as it can be expressed through a visual presentation that is straightforward to interpret. Moreover, we found that, for designers, this interaction was familiar as all our participants had previously used tools that contained checks of some kind—although not to the extent that PALETTE LINT offers. While powerful, this evaluation-based perspective can be limiting. For instance, it does not include support for non-evaluative categorization, as all observed behaviors are either marked as passing or failing. In contrast, cols4all [79] evaluates if a palette is “vivid.” While this could be modeled as an ignorable lint (e.g., *Palette Should Be Vivid*), requiring opt-out imposes different cognitive burdens than categorizing labels. Potentially useful for visualization linting, are past modifications to lint structure like *neutral* or *info* tests [23] or embedding tutorials into linters [65].

Linting by itself does not support generative interactions, such as to overcome the blank page problem [68]. We circumvent this issue by designing for opportunistic re-use [13] (e.g., start from a well-formed palette)—which is aligned with palette design practice: “for me it’s a lot of borrowing and cheating” (P₅). Involving similar design patterns seems to be a productive means to capture needs not met directly by linters (such as those addressed by generative solutions). An unexplored use of lints would be to use them as selectable components of a generative system—highlighting the reciprocal nature between evaluative and generative tools (e.g., recommenders [30, 67]). The essential differences between these tool types are not well understood in terms of results or user experience (e.g., is one more annoying?). This work does not interrogate these questions, but in future work we intend to use COLOR BUDDY as a platform to explore them.

Customization. *How can programs be modified?* A key experiment in our system is the reification [10] of lints as end-user manipulable programs; that is, we make them deeply customizable. Beaudouin-Lafon et al. [9] argue that such a formalization of a concept is useful as a means to empower the user to understand and create with a tool. Subjectively, we find that this integration makes it straightforward to experiment with new types of assertions, as the palettes can be adjusted and results checked. More concretely, our probe shows that this degree of customization is possible in a GUI-based linter and thus suggests

that similar strategies might apply to visualization linters more generally. This seems to be especially appropriate for our domain. For instance P₃ observed that a particular color ●, “works really well on dark, but you cannot use it on white because it really does look like baby puke”, highlighting the need to be able to adjust rules like *No Ugly Colors* to context. Lints do not need to be modeled as programs, but this form does seem to be a natural way to offer extensive customization. Software linters can be tuned by adjusting properties like max line length. Similarly, our lints might be combined to express multiple functions through configuration, for instance a *Sufficient Contrast with Background* rule capturing each possible algorithm and threshold. However, this freedom may make those rules difficult to configure or discover. Instead, relevant configurations might be materialized as separate lints (e.g., *APCA Contrast* vs. *WCAG Contrast*). Balancing lint customization and convenience is an area for future work.

Complexity. *How is complexity dealt with through design and automation?* Complexity in linters comes from their use and management. Our linter hides substantial complexity behind its seemingly simple pass/fail checks in the GUI. Substantial domain knowledge may be surfaced as direct assertions (e.g., *CVD: Deuteranopia Friendly*). However, as P₅ noted about Viz Palette [50], fixing the identified issues can be difficult without relevant experience. Our probe approaches this via visual explanations (such as our CVD simulation, Fig. 1B) and automated fixers. Moreover, integration with a GUI makes it possible to not only report but also help fix the problem identified by lint. Our linter, as with any, is driven by a growing body of knowledge. Managing large collections of knowledge is difficult, and arises in systems like Draco [70, 89] as well. This modularity can be positive. Per P₂ and P₄ what makes an effective palette will change. Linters accommodate these changes by adding new lints or replacing old ones as needed. However, growing a vast body of lints may be unsustainable, as rules may be conflicting, have differing assumptions, or may require protracted maintenance to keep up-to-date. This is analogous to the problems that gave way to the second AI winter when the limits of expert systems became manifest [39]. We again suggest that substantial theoretical innovation may address some of these issues. Being able to derive (or at least approximate) knowledge bases could make their maintenance and development simpler.

Adaptability. *What socio-technical (e.g., learnability) dimensions are considered?* Linters are social objects. For instance, the documentation for the SQL linter SQLFluff [21] includes guidance on navigating the integration of this type of validation into teams. Enforcing a set of styles ossifies a set of beliefs about how things should work, and should be done carefully to align with tasks and teams. In our linter we allow the user to add or remove lints as they wish, but team-driven collections of lints (as in eslint [77]) could support different situations or standards. Similarly, no technical object operates in complete isolation; rather, it is supported by the surrounding ecosystem of tools, such as auto-complete and documentation. However, these elements do not come for free. For instance, we designed PALETTE LINT as a JSON DSL because extant tooling supports in-situ documentation, auto-complete, and type hints at the cost of a more cumbersome syntax. Similarly, we surface the implementation of all the built-in lints both for documentation and for opportunistic reuse [13]. While these features are likely helpful for novices [55], they do not necessarily make our DSL easy to learn which likely has a substantial learning curve associated with it as it captures diverse concepts like predicate logic and color theory. In the future, we intend to explore how we might flatten this curve through modification to the language (e.g., via macros or simplifying concepts), structure editing [52], or automatic lint generation (such as via programming-by-demonstration or LLMs)—in addition to traditional learning resources.

7 DISCUSSION

We presented a design probe that explores the complexities of visualization linters by taking color palette design as a microcosm of visualization. We found that linters can be straightforwardly integrated into GUI-based contexts and tools, but adjustment to the affordances of those ecosystems is critical for appropriate use. For instance, having linters “always on” (or visually present) is acceptable in the vast visual

space of a code base, but can be overwhelming in the small plane of a color picker. We identified a series of basic features that linters tend to have, which we used to analyze and structure our probe.

While our probe seems to be useful (but improvable), the primary contribution of this work is not this system. Instead, we seek to answer the trio of criticisms of prior linters that opened this paper. We first observed that linters are **primarily text-based**, which may leave out many users. Following conventional color palette design tools, our probe was GUI-based and allowed us to demonstrate that a linter could be meaningfully integrated into a GUI tool. While there is still ample design territory to explore, our probe suggests that this integration can be fruitful, and, more importantly, that linters need not be bound exclusively to text. Moreover, all our participants were familiar with lint-like checks in end-user tools (e.g., sufficient contrast in a tool like WebAIM [35]) suggesting that users can be receptive to this type of integration and guidance. While many issues are addressable via guard rails or smart defaults, complex domains like color or visualization have ample space for the subtle errors that linters excel at catching.

Next, we noted that most linters are **impolite/annoying** by presenting huge walls of errors. Our probe explored incorporating task context (such as the palette type) and allowing users to select when and how they want linting advice—such as by focusing on examples rather than lints. We found these to be helpful ways to reduce visual cognitive load. These issues can thus be meaningfully reduced via design rather than being intrinsically bound to the medium. We found that opt-in control was sufficient to demonstrate a reduction in visual load, but models like probabilistic interrupts [32] might reduce it further still.

Finally, we observe that linters often feature **un-useful analysis**. By working in a domain with richer theoretical foundations and a cleaner set of results, we found that it was possible to straightforwardly present a meaningful, complex analysis. For instance, checking for properties like *Even Distribution in Lightness* (Fig. 4D) as well as easy-to-miss accessibility checks like *WCAG Contrast* (Fig. 4B). Crucially, being able to evaluate closer to usage (than to specification) allowed us to shift our analysis focus from the interface to the interaction [8]. To put it more simply: linters wedded to a particular domain—such as colors or choropleths or charts generally—offer a unique opportunity to validate semantics rather than merely the means they are formed in—going from *is the way in which I have produced something correct?* to *is what I have done correct?* DSL tooling (e.g., Vega-Lite’s JSON Schema) typically captures most typos and surface-level errors. Instead, linters have the opportunity to surface domain knowledge by acting as proxies for the more difficult problem of semantic evaluation.

Limitations. Like any work, ours suffers a number of limitations. We did not experimentally explore the effect of linting on color palette design. The purpose of our probe was to traverse the linting design space, rather than prove that it improves palette design. In future work we intend to compare the quality of palettes produced with related approaches. While our interview study provides evidence that our model of color palette usage aligns with that of professional palette designers, this is not the same as an empirical evaluation using our tool. Further, that study may be limited by selection bias as we drew from our professional networks to identify participants. Participants spoke about their practices being common, but our sample may bias this perspective. We attempted to circumvent these issues by using an external theory as a means of reflection (TDoPS [36]), but our results may be biased by this research-as-instrument approach. Similarly, while our lints covered all of the topics (excluding those not supported by current theory) that participants expressed as elements that they considered in their work our survey of these lints was not conducted formally. Conducting such a survey is valuable future work. Our probe had non-trivial conceptual and visual complexity, which may be difficult to approach for some users. Reducing this complexity may cause some of our results to differ, but our probing of the GUI-based linter design space did not require a perfected UX. However, we intend to continue to improve PALETTE LINT’s technical features and make it more approachable to non-expert users, as well as evaluate how such users (as well as designers) interact with it. We will also explore integration of PALETTE LINT with systems like Figma.

The next N linters. This work does not try to be simply yet another linter, but instead seeks to chart a trajectory for what future works and systems in this area might look like. Here we sketch a conceptual recipe we followed in this work, also observed in other works.

We first identified a **(1) problem domain** from which we collected appropriate guidelines, heuristics, and best practices. It can be important during this phase to engage with a domain expert. For instance, Lei et al. [46] engaged with a geographer (one of their paper authors). Similarly, one of the authors of this work is a noted expert in the field of color. We suggest that engaging with a domain expert is a valuable way to provide a layer of ecological and domain validity. Such collaboration can help answer questions like: *what checks are critical?* and *when is a set of checks sufficient?* Of course, no one expert can characterize appropriate usage for an entire domain of knowledge, so validating those results (as we did in Sec. 5) is important.

We then selected a **(2) problem space representation** consisting of standardized color representations, such as CIELAB and RGB hex colors, enhanced with a small amount of structure (Fig. 5). Similarly, tools like GeoLinter [46] or VizLinter [19] use Vega-Lite [69] as their problem space model. Much of a tool’s utility is embedded into this representation, so appropriate selection is important. The design of these representations is critical and may benefit from drawing on theoretical models (as Vega-Lite does with the Grammar of Graphics) or expanding them (such as via a model of visualization semantics) as well as human-centered design methods (such as PLIERS [20]).

Next, we developed an **(3) assertion representation** for the problem space, namely PALETTE LINT—creating an instrumental interaction-style relationship [7] between lint and color palette. In prior systems this representation is implicit, typically large collections of heuristics implemented in a general-purpose language. VizLinter [19] is an exception, drawing on answer-set programming. The effectiveness of the problem and assertion representations guides the system’s capability to climb the lint ladder (Sec. 3) in a user-understandable way. In contrast, `vislint_mpl` [53] analyzed matplotlib programs, which have an inconsistent API, limiting the scope of possible checks. Future work might further explore co-designing assertion and domain representations.

We developed a **(4) usage context** for the linter. We found that it was useful to integrate closely with the rest of the expected usage context—which in this case was the exploratory graphical setting of COLOR BUDDY. Software linters are dynamically integrated throughout the code authoring experience; they are not a standalone tool used without reference or interaction with other parts of the process. Unlike previous visualization linters, we intentionally connect ours with other forms of automation and support. This integration offered us fertile ground on which to build better domain support while giving a broader fixing space (including both visual hints and auto-fixers).

Using this recipe and our ladder of support we suggest that linters in other domains might flourish both within visualization (covering more domains and chart types) as well as outside of it (such as for notebooks [56] or \LaTeX documents). For instance, being able to evaluate dashboards would be useful, such as by developing an expanded version of Qu et al.’s [66] work. Similarly, an LLM prompt linter (as Henley argues for [28]) and as a sibling to Farsight [82]) could be powered by Shankar et al.’s Spade [72]. McNutt et al. [56] argue for the potential of data science and notebook linters.

Across these and other domains, we suggest that the explainable and modular advice linters provide is a valuable way to evaluate the semantics of those contexts—particularly when mixed with GUIs.

ACKNOWLEDGMENTS

We thank our participants for their insights and perspectives. We are appreciative of reviewer’s commentary, as well as the help given by Ravi Chugh, Junran Yang, Leilani Battle, and the rest of the UW IDL. This work was supported by the Moore Foundation.

8 SUPPLEMENTARY MATERIAL

Our materials are available at osf.io/geauf, and include our code, a video figure, and an appendix. Docs for our language can be found at color-buddy-docs.netlify.app. A live demo is at color-buddy.netlify.app

REFERENCES

- [1] U. A. Acar, A. Ahmed, J. Cheney, and R. Perera. A core calculus for provenance. *Journal of Computer Security*, 21(6):919–969, 2013. doi: 10.3233/JCS-130487 6
- [2] Adobe. Adobe Color. color.adobe.com, 2024. 2, 5
- [3] AntV. Advisor.lint. ava.antv.antgroup.com/en/api/advice/advisor-lint. 1, 2, 3
- [4] D. W. Barowy, E. D. Berger, and B. Zorn. ExcelLint: automatically finding spreadsheet formula errors. *Proceedings of the ACM on Programming Languages*, 2(OOPSLA):1–26, 2018. doi: 10.1145/3276518 2
- [5] L. Bartram, A. Patra, and M. Stone. Affective color in visualization. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1364–1374, 2017. doi: 10.1145/3025453.3026041 3, 4
- [6] L. Bartram and M. C. Stone. Whisper, don’t scream: Grids and transparency. *Trans. Visual Comput. Graphics*, 17(10):1444–1458, 2010. doi: 10.1109/TVCG.2010.237 3
- [7] M. Beaudouin-Lafon. Instrumental interaction: an interaction model for designing post-wimp user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 446–453, 2000. doi: 10.1145/332040.332473 9
- [8] M. Beaudouin-Lafon. Designing interaction, not interfaces. In *Advanced Visual Interfaces*, pp. 15–22, 2004. doi: 10.1145/989863.989865 9
- [9] M. Beaudouin-Lafon, S. Bødker, and W. E. Mackay. Generative theories of interaction. *ACM Transactions on Computer-Human Interaction*, 28(6):1–54, 2021. doi: 10.1145/3468505 8
- [10] M. Beaudouin-Lafon and W. E. Mackay. Reification, polymorphism and reuse: three principles for designing visual interfaces. In *Advanced Visual Interfaces*, pp. 102–109, 2000. doi: 10.1145/345513.345267 7, 8
- [11] L. D. Bergman, B. E. Rogowitz, and L. A. Treinish. A rule-based tool for assisting colormap selection. In *Visualization Conference*, pp. 118–125. IEEE, 1995. doi: 10.1109/VISUAL.1995.480803 2
- [12] F. Bianchi. Coolers. coolers.co. 6, 7
- [13] J. Brandt, P. J. Guo, J. Lewenstein, M. Dontcheva, and S. R. Klemmer. Two studies of opportunistic programming: interleaving web foraging, learning, and writing code. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1589–1598, 2009. doi: 10.1145/1518701.1518944 5, 8
- [14] M. Brown. Tableau Research leader and resident color expert Maureen Stone talks R&D at tableau. Tableau Blog. tableau.com/blog/how-maureen-stone-makes-data-more-accessible-tableau-research. 3
- [15] R. Bujack, T. L. Turton, D. H. Rogers, and J. P. Ahrens. Ordering perceptions about perceptual order. In *Scientific Visualization Conference*, pp. 32–36. IEEE, 2018. doi: 10.1109/SciVis.2018.8823772 5, 12
- [16] R. Bujack, T. L. Turton, F. Samsel, C. Ware, D. H. Rogers, and J. Ahrens. The Good, the Bad, and the Ugly: A Theoretical Framework for the Assessment of continuous colormaps. *Trans. Visual Comput. Graphics*, 24(1):923–933, 2018. doi: 10.1109/TVCG.2017.2743978 2, 3, 4, 5, 7, 12
- [17] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski. Characterizing guidance in visual analytics. *Trans. Visual Comput. Graphics*, 23(1):111–120, 2016. doi: 10.1109/TVCG.2016.2598468 1, 3
- [18] M. Chen, G. Grinstein, C. R. Johnson, J. Kennedy, and M. Tory. Pathways for theoretical advances in visualization. *IEEE computer graphics and applications*, 37(4):103–112, 2017. doi: 10.1109/MCG.2017.3271463 7
- [19] Q. Chen, F. Sun, X. Xu, Z. Chen, J. Wang, and N. Cao. Vizlinter: A linter and fixer framework for data visualization. *Trans. Visual Comput. Graphics*, 28(1):206–216, 2021. doi: 10.1109/TVCG.2021.3114804 1, 2, 3, 6, 9
- [20] M. Coblentz, G. Kambhatla, P. Koronkevich, J. L. Wise, C. Barnaby, J. Sunshine, J. Aldrich, and B. A. Myers. PLIERS: a process that integrates user-centered methods into programming language design. *ACM Transactions on Computer-Human Interaction*, 28(4):1–53, 2021. doi: 10.1145/3452379 9
- [21] A. Cruickshank. Rolling out SQLFluff with a new team. docs.sqlfluff.com/en/stable/teamrollout.html, 2023. 3, 8
- [22] B. Ford. Write Good: Naive linter for English prose. github.com/btford/write-good, 2019. 2
- [23] Github. Building CI checks with a GitHub App. docs.github.com/en/apps/creating-github-apps/writing-code-for-a-github-app/building-ci-checks-with-a-github-app, 2024. 8
- [24] C. C. Gramazio, D. H. Laidlaw, and K. B. Schloss. Colorgorical: Creating discriminable and preferable color palettes for information visualization. *Trans. Visual Comput. Graphics*, 23(1):521–530, 2016. doi: 10.1109/TVCG.2016.2598918 2, 3, 4
- [25] M. Graze and J. Schwabish. Building color palettes in your data visualization style guides. *Journal of the American Medical Informatics Association*, 31(2):488–498, 06 2023. doi: 10.1093/jamia/ocad084 2, 4
- [26] H. L. Han, M. A. Renom, W. E. Mackay, and M. Beaudouin-Lafon. Textlets: Supporting constraints and consistency in text documents. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020. doi: 10.1145/3313831.3376804 2
- [27] J. Heer and M. Stone. Color naming models for color selection, image editing and palette design. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1007–1016, 2012. doi: 10.1145/2207676.2208547 3, 4, 12
- [28] A. Henley. The pain points of building a copilot. austinhenley.com/blog/copilotpainpoints.html, 2024. 9
- [29] M. Hogräfer and H.-J. Schulz. ReVize: A Library for Visualization Toolchaining with Vega-Lite. In M. Agus, M. Corsini, and R. Pintos, eds., *Smart Tools and Apps for Graphics - Eurographics Italian Chapter Conference*. Eurographics, 2019. doi: 10.2312/stag.20191375 5
- [30] M.-H. Hong, Z. N. Sunberg, and D. Albers Szafir. Cieran: Designing Sequential Colormaps via In-Situ Active Preference Learning. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–14, 2024. doi: 10.48550/arXiv.2402.15997 2, 7, 8
- [31] A. K. Hopkins, M. Correll, and A. Satyanarayan. Visualint: Sketchy in situ annotations of chart construction errors. In *Computer Graphics Forum*, pp. 219–228. Wiley, 2020. doi: 10.1111/cgf.13975 1, 2, 3, 7
- [32] E. Horvitz. Principles of mixed-initiative user interfaces. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 159–166, 1999. doi: 10.1145/302979.303030 9
- [33] M. Hull, V. Pednekar, H. Murray, N. Roy, E. Tung, S. Routray, C. Guerin, J. Chen, Z. J. Wang, S. Lee, M. Roozbahani, and D. H. Chau. VIS-GRADER: Automatic Grading of D3 Visualizations. *Trans. Visual Comput. Graphics*, 2023. doi: 10.1109/TVCG.2023.3327181 3, 6
- [34] N. Hynes, D. Sculley, and M. Terry. Data linter: Lightweight, automated sanity checking for ml data sets. In *NIPS ML Sys Workshop*, 2017. 2
- [35] Institute for Disability Research, Policy, and Practice. WebAIM: Contrast Checker. webaim.org/resources/contrastchecker, 2024. 7, 9
- [36] J. Jakubovic, J. Edwards, and T. Petricek. Technical Dimensions of Programming Systems. *Art, Science, and Engineering of Programming*, 7(3), 2023. doi: 10.22152/programming-journal.org/2023/7/13 7, 9
- [37] G. Jalal, N. Maudet, and W. E. Mackay. Color portraits: From color picking to interacting with color. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 4207–4216, 2015. doi: 10.1145/2702123.2702173 2, 5
- [38] A. Kale, Z. Guo, X. L. Qiao, J. Heer, and J. Hullman. EVM: Incorporating model checking into exploratory visual analysis. *Trans. Visual Comput. Graphics*, 2023. doi: 10.1109/TVCG.2023.3326516 3
- [39] H. Kautz. The third AI summer: AAAI Robert S. Engelmore Memorial Lecture. *AI Magazine*, pp. 105–125, 2022. doi: 10.1609/aimag.v43i1.19122 8
- [40] N. Kita and K. Miyata. Aesthetic rating and color suggestion for color palettes. In *Computer Graphics Forum*, pp. 127–136. Wiley, 2016. doi: 10.1111/cgf.13010 2
- [41] A. J. Ko, R. Abraham, L. Beckwith, A. Blackwell, M. Burnett, M. Erwig, C. Scaffidi, J. Lawrance, H. Lieberman, B. Myers, M. B. Rosson, G. Rothermel, M. Shaw, and S. Wiedenbeck. The state of the art in end-user software engineering. *ACM Computing Surveys (CSUR)*, 43(3):1–44, 2011. doi: 10.1145/1922649.1922658 2
- [42] R. Kosara. Understanding Pie Charts. eagereyes.org/pie-charts. 2
- [43] Y. S. Kristiansen, L. Garrison, and S. Bruckner. Semantic snapping for guided multi-view visualization design. *Trans. Visual Comput. Graphics*, 28(1):43–53, 2021. doi: 10.1109/TVCG.2021.3114860 1
- [44] A. Kucherenko. jscpd: Copy/paste detector for programming source code. github.com/kucherenko/jscpd, 2023. 2
- [45] S. Lee, M. Sips, and H.-P. Seidel. Perceptually Driven Visibility Optimization for Categorical Data Visualization. *Trans. Visual Comput. Graphics*, 2013. doi: 10.1109/TVCG.2012.315 2, 6
- [46] F. Lei, A. Fan, A. M. MacEachren, and R. Maciejewski. GeoLinter: A Linting Framework for Choropleth Maps. *Trans. Visual Comput. Graphics*, 2023. doi: 10.1109/TVCG.2023.3322372 1, 2, 3, 7, 9
- [47] J. Li, E. Rawn, J. Ritchie, J. Tran O’Leary, and S. Follmer. Beyond the Artifact: Power as a Lens for Creativity Support Tools. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–15, 2023. doi: 10.1145/3586183.3606831 2, 8

- [48] C. Lovers. The ugliest colors. https://www.colourlovers.com/palette/1416250/The_Ugliest_Colors. 3
- [49] K. Lu, M. Feng, X. Chen, M. Sedlmair, O. Deussen, D. Lischinski, Z. Cheng, and Y. Wang. Palettaiolor: Discriminable colorization for categorical data. *Trans. Visual Comput. Graphics*, 27(2):475–484, 2020. doi: 10.1109/TVCG.2020.3030406 2
- [50] S. Lu and E. Meeks. Viz Palette. projects.susielu.com/viz-palette, 2018. 2, 4, 7, 8
- [51] A. McNutt and R. Chugh. Integrated visualization editing via parameterized declarative templates. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–14, 2021. doi: 10.1145/3411764.3445356 5
- [52] A. McNutt and R. Chugh. Projectional editors for json-based dsls. In *Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, pp. 60–70. IEEE, 2023. doi: 10.1109/VL-HCC57772.2023.00015 8
- [53] A. McNutt and G. Kindlmann. Linting for visualization: Towards a practical automated visualization guidance system. In *VisGuides (IEEEVIS Workshop)*, pp. 1–14, 2018. 1, 2, 3, 9
- [54] A. McNutt, G. Kindlmann, and M. Correll. Surfacing visualization mirages. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–16, 2020. doi: 10.1145/3313831.3376420 2, 3
- [55] A. McNutt, A. Outkine, and R. Chugh. A Study of Editor Features in a Creative Coding Classroom. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–15, 2023. doi: 10.1145/3544548.3580683 2, 7, 8
- [56] A. McNutt, C. Wang, R. A. Deline, and S. M. Drucker. On the design of ai-powered code assistants for notebooks. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–16, 2023. doi: 10.1145/3544548.3580940 1, 2, 9
- [57] E. Meeks. Linting Rules for Complex Data Visualization. PlotCon, 2017. youtube.com/watch?v=KEI-Spdaz0. 2, 3
- [58] S. Mehrpour, T. D. LaToza, and H. Sarvari. Rulepad: interactive authoring of checkable design rules. In *European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pp. 386–397, 2020. doi: 10.1145/3368089.3409751 8
- [59] B. J. Meier, A. M. Spalter, and D. B. Karelitz. Interactive color palette tools. *IEEE Computer Graphics and Applications*, 24(3):64–72, 2004. doi: 10.1109/MCG.2004.1297012 2
- [60] L. C. Muth. A detailed guide to colors in data vis style guides. blog.datawrapper.de/colors-for-data-vis-style-guides, 2022. 3, 4, 12
- [61] Myndex Research. APCA Contrast Calculator. myndex.com/APCA. 7
- [62] P. Nardini, M. Chen, M. Böttinger, G. Scheuermann, and R. Bujack. Automatic improvement of continuous colormaps in euclidean colorspace. In *Computer Graphics Forum*, number 3, pp. 361–373. Wiley Online Library, 2021. doi: 10.1111/cgf.14313 2, 12
- [63] Paletton. Paletton, 2018. paletton.com. 2
- [64] S. E. Palmer. *Vision science: Photons to phenomenology*. MIT press, 1999. doi: 10.1016/S1364-6613(99)01430-8 2
- [65] M. Pocock. ts-error-translator. github.com/mattpocock/ts-error-translator, 2024. 8
- [66] Z. Qu and J. Hullman. Keeping multiple views consistent: Constraints, validations, and exceptions in visualization authoring. *Trans. Visual Comput. Graphics*, 24, 2017. doi: 10.1109/TVCG.2017.2744198 3, 9
- [67] A. Salvi, K. Lu, M. E. Papka, Y. Wang, and K. Reda. Color Maker: a Mixed-Initiative Approach to Creating Accessible Color Maps. In *SIGCHI Conference on Human Factors in Computing Systems*, 2024. doi: 10.48550/arXiv.2401.15032 2, 8
- [68] A. Satyanarayan, B. Lee, D. Ren, J. Heer, J. Stasko, J. Thompson, M. Brehmer, and Z. Liu. Critical reflections on visualization authoring systems. *Trans. Visual Comput. Graphics*, 26(1):461–471, 2019. doi: 10.1109/TVCG.2019.2934281 2, 7, 8
- [69] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vegalite: A grammar of interactive graphics. *Trans. Visual Comput. Graphics*, 23(1):341–350, 2016. doi: 10.1109/TVCG.2016.2599030 1, 4, 9
- [70] J. Schmidt, B. Pointner, and S. Miksch. Visual Analytics for Understanding Draco’s Knowledge Base. *Trans. Visual Comput. Graphics*, 2023. doi: 10.1109/TVCG.2023.3326912 8
- [71] V. Setlur and M. Stone. A linguistic approach to categorical color assignment for data visualization. *Trans. Visual Comput. Graphics*, pp. 698–707, 2015. doi: 10.1109/TVCG.2015.2467471 6
- [72] S. Shankar, H. Li, P. Asawa, M. Hulsebos, Y. Lin, J. Zamfirescu-Pereira, H. Chase, W. Fu-Hinthorn, A. G. Parameswaran, and E. Wu. SPADE: Synthesizing Assertions for Large Language Model Pipelines. *arXiv*, 2024. doi: 10.48550/arXiv.2401.03038 9
- [73] C. Shi, W. Cui, C. Liu, C. Zheng, H. Zhang, Q. Luo, and X. Ma. NL2Color: Refining Color Palettes for Charts with Natural Language. *Trans. Visual Comput. Graphics*, 2023. doi: 10.1109/TVCG.2023.3326522 3
- [74] X. Shi, Z. Zhou, J. W. Zhang, A. Neshati, A. K. Tyagi, R. Rossi, S. Guo, F. Du, and J. Zhao. De-Stijl: Facilitating Graphics Design with Interactive 2D Color Palette recommendation. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–19, 2023. doi: 10.1145/3544548.3581070 2
- [75] S. Shin, S. Hong, and N. Elmqvist. Perceptual Pat: A Virtual Human Visual System for Iterative Visualization design. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–17, 2023. doi: 10.1145/3544548.3580974 1, 3, 7
- [76] M. Shugrina, W. Zhang, F. Chevalier, S. Fidler, and K. Singh. Color builder: A direct manipulation interface for versatile color theme authoring. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2019. doi: 10.1145/3290605.3300686 2
- [77] D. Specker. awesome-eslint. github.com/dustinspecker/awesome-eslint, 2023. 3, 8
- [78] J. Spellman, R. B. Montgomery, M. Cooper, S. Lauriat, C. Adams, and A. Campbell. W3C Accessibility Guidelines (WCAG) 3.0, July 2023. 3, 4, 7
- [79] M. Tennekes and M. J. Puts. cols4all: a color palette analysis tool. In *EuroVis (Short Papers)*. Wiley, 2023. doi: 10.2312/evs.20231040 2, 3, 4, 8
- [80] G. W. Tigwell, D. R. Flatla, and N. D. Archibald. ACE: a colour palette design tool for balancing aesthetics and accessibility. *ACM Transactions on Accessible Computing*, 9(2):1–32, 2017. doi: 10.1145/3014588 2
- [81] ViewSonic. What is Delta E? And Why Is It Important for Color Accuracy? viewsonic.com/library/creative-work/what-is-delta-e-and-why-is-it-important-for-color-accuracy, 2021. 4
- [82] Z. J. Wang, C. Kulkarni, L. Wilcox, M. Terry, and M. Madaio. Farsight: Fostering Responsible AI Awareness During AI Application Prototyping. In *SIGCHI Conference on Human Factors in Computing Systems*, 2024. doi: 10.48550/arXiv.2402.15350 2, 3, 9
- [83] webhint. Use webhint to improve your website. webhint.io/docs/user-guide/hints, 2023. 2
- [84] B. Whitworth. Polite computing. *Behaviour & Information Technology*, 24(5):353–363, 2005. doi: 10.1080/01449290512331333700 2
- [85] H. Wickham, D. Cook, H. Hofmann, and A. Buja. Graphical inference for infovis. *Trans. Visual Comput. Graphics*, 16(6):973–979, 2010. doi: 10.1109/TVCG.2010.161 5
- [86] H. Winchester, E. Al Haque, A. Boyd, and B. Johnson. HaTe Detector: A Tool for Detecting and Correcting Harmful Terminology in computing artifacts. In *Symposium on Visual Languages and Human-Centric Computing*, pp. 245–248. IEEE, 2023. doi: 10.1109/VL-HCC57772.2023.00043 2
- [87] J. Wood, A. Kachkaev, and J. Dykes. Design exposition with literate visualization. *Trans. Visual Comput. Graphics*, 25(1):759–768, 2018. doi: 10.1109/TVCG.2018.2864836 3, 5
- [88] T. Wu, M. Terry, and C. J. Cai. AI Chains: Transparent and controllable human-ai interaction by chaining large language model prompts. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 1–22, 2022. doi: 10.1145/3491102.3517582 3
- [89] J. Yang, P. F. Gyarmati, Z. Zeng, and D. Moritz. Draco 2: An extensible platform to model visualization design. In *IEEE Visualization and Visual Analytics (VIS)*, pp. 166–170. IEEE, 2023. doi: 10.1109/VIS54172.2023.00042 8
- [90] L.-P. Yuan, Z. Zhou, J. Zhao, Y. Guo, F. Du, and H. Qu. Infocolorizer: Interactive recommendation of color palettes for infographics. *Trans. Visual Comput. Graphics*, 28(12):4252–4266, 2021. doi: 10.1109/TVCG.2021.3085327 2
- [91] N. C. Zakas and N. Sullivan. CSS Lint. csslint.net, 2011. 2
- [92] W. Zheng and Y. Sherif. Visualization Linter, 2019. 3
- [93] L. Zhou and C. D. Hansen. A survey of colormaps in visualization. *Trans. Visual Comput. Graphics*, 22(8):2051–2069, 2015. doi: 10.1109/TVCG.2015.2489649 2