# VizLinter: A Linter and Fixer Framework for Data Visualization



Qing Chen, Fuling Sun, Xinyue Xu, Zui Chen, Jiazhe Wang, and Nan Cao

Fig. 1. Four example cases of VizLinter. In each case, the chart on the left is the original visualization with errors, and the chart on the right is the one modified by VizLinter.

**Abstract**— Despite the rising popularity of automated visualization tools, existing systems tend to provide direct results which do not always fit the input data or meet visualization requirements. Therefore, additional specification adjustments are still required in real-world use cases. However, manual adjustments are difficult since most users do not necessarily possess adequate skills or visualization knowledge. Even experienced users might create imperfect visualizations that involve chart construction errors. We present a framework, VizLinter, to help users detect flaws and rectify already-built but defective visualizations. The framework consists of two components, (1) a visualization linter, which applies well-recognized principles to inspect the legitimacy of rendered visualizations, and (2) a visualization fixer, which automatically corrects the detected violations according to the linter. We implement the framework into an online editor prototype based on Vega-Lite specifications. To further evaluate the system, we conduct an in-lab user study. The results prove its effectiveness and efficiency in identifying and fixing errors for data visualizations.

Index Terms—Visualization Linting, Automated Visualization Design, Visualization Optimization

## **1** INTRODUCTION

Visualization tools, such as business intelligence software and programming toolkit, have gained plenty of users as visualization becomes popular across various domains [27]. Common visualization tools often require users to set specifications of visualizations either through programming syntax or user interfaces. However, these operations demand users to have solid background knowledge in data analysis and visualization development. The rising demand for creating ap-

- Qing Chen, Fuling Sun, Xinyue Xu, Zui Chen, and Nan Cao are with Intelligent Big Data Visualization Lab at Tongji University.
  E-mails: jane.qing.chen@gmail.com, {fulingsun, xuxinyue}@tongji.edu.cn, {zuic10a, nan.cao}@gmail.com. Nan Cao is the corresponding author.
- Jiazhe Wang is with Ant Group. Email: jiazhe.wjz@antgroup.com

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.xxxxxx

propriate visualizations triggers the research interest in developing automated systems to recommend and build expressive and efficient visualizations, especially for users with limited expertise. Most existing automated systems are either rule-based or machine-learning-based. Rule-based systems program visual encoding principles into rule sets to automate the generation of visualizations. They can recommend visualizations based on input data characteristics [62], intended tasks [13], and user behaviors [20]. Machine-learning-based systems learn relationships between data and the corresponding legitimate visualizations directly [24]. Although current recommendation systems are capable of fast-forwarding source data into proper visualizations, they are not always compatible with input data or the visualization requirements due to the limitation of fixed encoding rules and training data. Therefore, manual specifications are inevitable in practice. However, since users of automated systems tend to have inadequate expertise, it is difficult for them to make manual changes while following well-established visual design rationales.

Although manual adjustments are required when automated results cannot fulfill user needs, few existing works are capable of detecting errors and suggesting corrections for manual inputs. Some prior works are able to highlight problems [23, 25] or to lint, that is, flagging errors in the code for visualizations [34], but they do not provide any solutions to the detected problems. Other works concentrate on auto-completing user input specifications and recommending optimal alternatives by resolving predefined constraints in an end-to-end manner [32,37]. However, they do not present any explanations or suggestions on how to fix imperfect visualizations.

Therefore, programmatic solutions are still absent to detect breached visualization principles and resolve the violations with optimum operations. To remedy this absence, we construct a linter, which originally represents static code analysis tools to flag bugs or issues in the programming code, to detect violated rules in the visualization using Answer Set Programming (ASP) [31], a declarative constraint programming language which allows us to model high-level design knowledge into logical facts. An ASP program consists of Prolog-style rules a :- L\_1, L\_2, not L\_3, not L\_4., where a is the head of the rule and L\_i are literals. The rule states that the head a is derived if L\_1 and L\_2 are true while L\_3 and L\_4 cannot be true. Specifically, the head of each rule in VizLinter is the rule id in the rule base and corresponding parameters, and each literal represents an attribute of the visualization. ASP enables easy maintenance and iteration on the rule base [11], currently refined from the Draco knowledge base [37]. After violated rules are detected, we formulate a fixer for the broken rules using linear optimization with a designated reward function.

In this paper, we propose a novel linter and fixer framework for data visualization. It detects infringed rules in visualizations and provides solutions to those broken rules by optimizing through linear programming. We implement our framework for Vega-Lite-based visualizations, as Vega-Lite is a widely used language in the visualization community. Our implementation follows the pipeline as shown in Fig. 2. It first translates input Vega-Lite specifications into ASP facts. The linter then checks against the facts via an ASP solver and detects the violated principles. Next, the fixer performs the action selected by the optimization algorithm and resolves the broken rules. To demonstrate the usefulness of the proposed framework, we develop a prototype and conduct a user study to assess its efficacy. According to the feedback, our framework is confirmed as helpful by providing ad-hoc prompts and solutions during the creation of visualizations. It can benefit not only business intelligence tool users but also more skilled practitioners and researchers. It is considered as accessible for many existing tools and systems to adapt as an extension and to assist direct development in Vega-Lite coding syntax. Besides, the study results indicate the pedagogical value of our framework as users are able to learn visualization specification rules through the interaction. While the current functionality is already valuable as an automatic error detection and correction tool, it is expected to lint and fix aesthetic or semantic-related issues as well in the next iterations. To summarize, our contributions are as follows:

- A linter-and-fixer framework. We present a novel linter-and-fixer framework for data visualization, which can examine for erroneous specifications and resolve them automatically.
- A visualization linter. We refine the design principle constraints in Draco [37] as a concrete rule base and construct a visualization linter to detect violations using ASP.
- A visualization fixer. Using our proposed reward function and the operation transition costs from GraphScape [29], we formulate an optimization algorithm using linear programming to fix the violated rules detected by the linter.
- A user study. We evaluate the effectiveness of our framework through a user study on the prototype system. Through the study, we validate the effectiveness of our framework in practice and collect instructive reflections for future improvement.

## 2 RELATED WORK

In this section, we draw on prior work on recommendation systems, linting, and optimization models for visualization.

#### 2.1 Recommendation Systems for Visualization

Visualization recommendation systems accommodate users by performing the heavy-lifting jobs, such as data analysis and insight extraction. They can automatically generate and suggest visualizations for a given dataset. Earlier studies tend to recommend visualizations based on predefined rules. They derive rules from statistical features [46, 54, 61], visual effectiveness measures [9, 14], and user behaviors [20] or tasks [43]. Recent works benefit from machine learning and other algorithms [64]. VizML automatically infers five design choices for input datasets using a machine learning algorithm derived from data characteristics [24]. ClustMe, developed by Abbas *et al.*, ranks cluster patterns of monochrome scatterplots using a perceptual visual quality measure to match human judgment [2]. Shi *et al.* present a sequencing model that recommends optimal sequences for different tasks [47]. Kim *et al.* introduce a directed-graph model for visualization sequence recommendation based on chart similarity [29].

Our framework presents a solution to amend visualization with flaws compared to these studies on fast-forwarding visualization recommendation. It allows users to decide if they want to adopt the recommended adjustment and has more pedagogical value. Although our framework is also rule-based, we model our rulesets using Answer Set Programming. Therefore, unlike traditional rule-based recommendation systems whose rule bases are immutable, our framework allows timely updates and modifications on the rules.

## 2.2 Linting for Visualization

The idea of linting is borrowed from computer science, which means static code analysis tools to flag bugs or errors in the programming code. Recently, linting has been applied beyond traditional programming languages. Before being applied in visualization, the concept of linting has been used in data wrangling as a validation approach to eliminate errors in provided datasets. CheckCell [5], ExceLint [4], and Uni-Detect [55] undertake data validation errors in tabular data. Muslu et al. [40] and Brachmann et al. [10] furthermore consider potential problems with data curation and large data sets. Since data wrangling is an important procedure in creating visualizations, works on data linting are also valuable for later research on visualization linting. Meanwhile, in the visualization field, the verification and evaluation of visualizations are also highly relevant to visualization linting. Kirby and Silva emphasize the significance of verifying the accuracy of visualization techniques [30]. Prior works undertake the evaluation of visual quality, algorithm performance, and user experience of visualization with different approaches [19,26]. The idea of "visual linting" was first introduced by McNutt and Kindlmann to apply linting in visualization [34]. They define visualization lint as a framework, that takes in a visualization or its production code, evaluates whether it will pass, and explain why it fails. Their linter targets only readability-related problems in charts, those causing charts difficult for humans to observe [34]. Our work, instead, focuses on errors that make charts illegal or prevent them from rendering properly, rather than aesthetic problems such as readability issues of a valid visualization. After linting being introduced into visualization, some research takes on different perspectives to study the application of linting in visualization. In Metamorphic VegaLite Analyzer, metamorphic testing is applied to validate if a visualization can properly represent the input data [35]. However, they do not pay special attention to critical errors. Besides, due to reliance on bootstrapping and other statistical techniques, Metamorphic VegaLite Analyzer tends to show limited performance in terms of speed. Since our framework uses Answer Set Programming for rule inspection, which involves less computation, it is able to detect errors without latency.

More recently, Hopkins *et al.* design and evaluate VisuaLint, which displays five chart construction errors in visualization using sketchy annotations instead of conventional cues in text messages [23]. While the works mentioned above share some similar design considerations with our work, they focus more on the design- and readability-related issues, without any suggestions or functions to help users correct existing errors in the visualization.

#### 2.3 Optimization for Visualization

Visualizations are usually not perfect upon generation due to various factors, such as data volumes or display sizes. Therefore, optimization can help visualizations accommodate different situations and remain



Fig. 2. The pipeline of VizLinter includes two key modules: the linter and the fixer. Given the visualization specification with errors, it goes through the linter and the fixer to get the revised specification.

good quality. Since data quality has a significant impact on the quality of visualization, some research focuses on the data preparation for visualization. For example, Wen and Zhou provide a model to dynamically derive a set of optimal data transformations for the target visualization [59]. Many other works concentrate on the optimization of charts directly. Some target problems for specific types of charts. Ellis and Dix focus on clutter reduction in parallel coordinates using sampling techniques [16]. Trutschl et al. also tackle point occlusion issues in parallel coordinates but using a self-organized map instead [53]. Wang et al. improve multi-class scatterplot by optimizing color assignment for class separation [56]. Heer and Agrawala develop multi-scale banking, enhancing the visual perception of segment orientations in line charts [21]. Byron and Wattenberg present a wiggle-minimizing method to optimize the overall shapes of stacked graphs [12]. Tang et al. develop an authoring tool for storyline visualization, PlotThread, which consists of an AI agent to optimize the layout [50]. Existing research also covers the optimization for pixel-based charts [28], radial charts [51], node-link diagrams [57], matrices [6], treemaps [48], geographical maps [8], and text-based charts [45]. Other research resolves general issues regardless of chart types. Draco is a constraint-based system to score visualizations based on the violations against design rules and recommend the top-scored candidate [37]. Lin et al. contribute Dziban extending on Draco, which additionally optimizes the synthesized charts to accord with the anchor charts [32]. Wu et al. address the responsive display issues of SVG-based visualization on mobile screens using reinforcement learning [63].

However, while previous work could suggest optimal visualizations, they are not able to optimize erroneous visualizations with violated design rules, especially when violations prevent charts from compiling. The target of our work is to provide common solutions to various types of visualizations. Specifically, the ideal framework is aimed to improve a given visualization by revising the incorrect specifications. Moreover, to provide the most efficient modification, we adapt the transition cost collected in GraphScape [29]. A detailed description of the proposed algorithm is presented in Section 4.2.

## **3** DESIGN REQUIREMENTS

In this section, we collect the common issues that practitioners have encountered and summarize the design goals accordingly.

## 3.1 Common Issues in Visualizations

To better understand the common issues that practitioners encounter and what kind of assistance might be helpful, we conducted semi-structured interviews with two experts. One is an experienced visualization researcher who has ten-year programming experience and visualization knowledge; the other is the team leader of a data intelligence department in a well-established IT company. Both experts have practical experience in developing data visualization systems for various domains such as business intelligence and social network analysis. In the interviews, the experts noted that creating a visualization requires specifying marks and channels, the two primary components of visual encoding [7]. However, simply declaring marks and channels will not guarantee a successful encoding. The complex conditions behind it, such as data characteristics, can lead to failed visualizations if not compatible with stated specifications. We frame the issues within high-level visualization grammars, such as Tableau/Polaris [49] and Vega-Lite [44]. These high-level grammars are favored in exploratory visualization due to the preference of conciseness over expressiveness [22]. In high-level visualization grammars, declarations of *marks* and *encoding channels* are required to create visualizations for a given data set. Declaration of *marks* is to specify the mark type (e.g., point or line) for the visualization. Declaration of *encoding channels* includes the specification of the expressed data field (column) and how it is expressed: whether using size or color channel and if any data transformation (e.g., binning, aggregation, logging, scale) or visual transformation (e.g., stacking) are applied. Consequently, four common issues are categorized from the expert interviews and previous literature related to visualization principles [14, 33, 39, 58] and visualization grammar [60].

**I1. Incompatibility issues within each encoding channel.** Withinencoding issues illustrate illegal mapping between the selected data field (column) and the encoding channel [14, 39]. They can lead to failed visualizations when, for instance, a wrong data type is declared, a chosen encoding channel is incompatible with the selected column, or a chosen aggregation is incompatible with the selected column.

**12.** Incompatibility issues across multiple encoding channels. Even if each individual encoding is correct, the visualization can still be problematic due to conflicts across encoding channels [14, 60]. One possible cause is that an encoding specification is not compatible with another. It can be attributed to situations such as duplicated usage of the same channel or identical columns used in both axes.

**I3.** Incompatibility issues between encoding channels and marks. Potential failures can occur when the chosen mark is not compatible with the selected encoding channel [14, 39], for example, when the size channel is chosen with marks other than points or text.

**I4. Typo issues.** It is natural for humans to make illegal declarations or typo errors during manual input. However, any minor typo mistakes can prevent a visualization from valid rendering.

According to the above issues, we refine a collection of 41 rules fundamental to construct legitimate charts. In this paper, we regard a chart as invalid or illegal once it violates one or more rules of the four common issue categories, and possibly fails to render. Illegal chart specifications differentiate from "readability-related" problems mentioned by McNutt *et al.* [34] in that they may prevent charts from rendering or presenting correct data. However, "readability-related" problems focus on charts rendered successfully with correct data presentation, while other issues, such as a lack of title or annotation or a misuse of colors, occur. Such issues are more related to the aesthetics aspect of visualizations, which are excluded in our current work. While our current coverage intersects with the one defined by Hopkins *et al.* [23], we do not take the issues regarding dual axis, legends, and ineffective colors into account since the primary focus of this paper is the legibility of visualization.

## 3.2 Design Goals

Apart from the common issues in visualization, we also interviewed the experts on how they would utilize external support to help find and modify existing issues. Three design goals are thus derived as follows. **G1. Facilitate visualization developers with error indicators.** Our primary target users are visualization developers with inadequate experience and skilled developers who want to improve their efficiency and accuracy. For those users, concurrent error disclosure could speed up the development process, especially when they are not able to detect problems. Therefore, we need to offer developers the option to see what and where errors are in the specification. In our framework design, this consideration resolves to a linter that inspects errors in visualizations and displays details of the errors.

**G2.** Automate the fixing operations on detected errors. Error indicators alone cannot fully speed up development or ensure correctness because it still takes time and effort to seek solutions for the errors. When making manual adjustments, developers may choose operations in a roundabout way. Consequently, we design a fixer to tackle the problem in the most efficient way. We regard the best solution as the one that involves fewer actions while fixing more errors. In Section 4.2, we present the reward function and the corresponding optimization algorithm based on this logic.

**G3.** Support easy integration and iteration. Many existing business intelligence tools or visualization systems are built on standard visualization grammars. Consequently, we design the liner and the fixer with grammar-based visualization specifications as input. Since Vega-Lite is a standard grammar in visualization research, we use it as the basic data input into our framework. Besides, we choose to form the linter using Answer Set Programming (ASP) so that the ruleset can be easily updated according to different scenarios or domains [11]. Moreover, the fabrication of our linter-and-fixer framework can serve as design guidance for existing tools and systems with similar functionalities.

## 4 THE VIZLINTER FRAMEWORK

Guided by the design goals summarized in Section 3.2, we design and illustrate the pipeline of the framework in Fig. 2. First, the user inputs a piece of visualization specifications (in our implementation, the inputs are in Vega-Lite JSON format). The framework then transforms the specifications into a list of facts that ASP can handle (G3). Next, the linter built upon ASP examines the facts and detects violations. The corresponding errors are highlighted in the original specifications (G1). For the fixer part, by applying linear optimization programming, it automatically selects an optimal set of operations that resolve the detected violations (G2). Finally, the solutions are translated back to a refined piece of Vega-Lite JSON specifications. In this section, we first present the construction of the linter, including the constraints and the answer sets we developed. Then, we describe the algorithm, the reward, and the cost functions used to build the fixer.

#### 4.1 Visualization Linter

The linter is made up of two components, a set of predefined rules visualization should obey and an Answer Set Programming (ASP) solver detecting contravened rules.

## 4.1.1 Linting Rules

A visualization linter, borrowing the concept of "linting" from computer science, is a tool built on a collection of rules that should be obeyed. It functions similarly to traditional linting tools in programming, such as ESLint [1]. Like ESLint, VizLinter also examines the violation of rules; however, visualization linting rules include not only syntax errors but also visualization principles. Consequently, the coverage of the ruleset is a determinant of the performance of the linter. The essential principles, which any legal visualization should adhere to, fabricate the basic ruleset for our linter's structure. We refine a collection of 41 rules according to the common issues summarized in Section 3.1 and embed them in the linter.

In order to minimize human efforts and potential errors caused by user input, we directly parse data range and cardinality from original data sets. Consequently, some rules related to user input about data properties in Draco are not considered in our collection. Some rules defined in Draco have identical coverage or one's coverage subsets another. In this case, we discard rules with duplicated coverage and keep rules that cover more cases. Besides, we simplify some complex constraints by breaking them down into more straightforward rules in the fixer. For instance, in Draco, a constraint requiring *stack* to come along with a discrete *color* channel is divided into three distinct rules: one takes care of the absence of color channel when no encoding channel is specified; one deals with the absence of color channel when size is specified as an encoding channel; the third handles when a non-discrete color channel is used. Such decomposition aids in the development of a coherent ruleset and the assignment of resolving operations. Finally, we elaborate a collection of 41 rules as the linter's initial rule base, with a complete list available online<sup>1</sup>.

## 4.1.2 ASP for Linting

We follow the same approach as in Draco to formulate visualization principles using ASP since it enables us to construct a stable model for our rule base. When there is a need to introduce new rules or update the complete ruleset for specific use cases, one can simply add or rewrite the constraints for the intended purposes and not have to re-construct the whole framework again.

Atoms, literals, and rules are the three building blocks for ASP programming [11]. Atoms are the elementary propositions that may be true or false; literals are atoms A or not A; rules are formed by atoms as a :- L\_1, L\_2, not L\_3, not L\_4., where a is the head of these rules and L\_i are literals. If all literals are true, then the head a is derived; otherwise, a is not established. Particularly, if a rule has only the head but no literal, such as a :- ., it is called a fact, indicating that the head is unconditionally true, and such a rule can be abbreviated as a. On the contrary, a headless rule, such as :- L\_1, L\_2. without head a in the rule, is regarded as a constraint, which means L\_1 and L\_2 cannot be true at the same time.

In the linter, we present the refined visualization rules in the *rule* format of ASP. For example, the rule hard(bin\_and\_aggregate, C) : $bin(E,_)$ , aggregate(E,\_), channel(E,C). points to encoding E with channel C using both bin and aggregate. In this rule, a predicate hard with two parameters, the rule id bin\_and\_aggregate and the channel C, is used to express the violation. Channel C indicates in which encoding channel that the rule is violated. Differently, hard constraints in Draco are modeled as a set of headless rules, that is, constraints of ASP, preventing the automatic generation of illegal visualizations that violate these constraints, rather than detecting any violations from the given visualization.

The first step in our linter's process is to translate visualization specifications (Vega-Lite JSON in our case) into ASP programs of facts (*rules* without *literals*), characterizing how the given visualization is formed. We develop the translator based on Draco to extract the attributes such as mark type, encodings, and their properties. We also incorporate additional functionalities, such as automatically detecting *fieldtype* from data. Based on the translated visualization facts and the defined ruleset, we run an ASP solver to find out rules whose literals are all satisfied by the facts, in other words, the errors in the visualization.

## 4.2 Visualization Fixer

After acknowledging problems in the visualization, it is still difficult and time-consuming for inexperienced developers to understand the meaning of each linted problem and modify the visualization specification. To better facilitate visualization developers, we propose an optimization algorithm to automatically provide optimal solutions to the detected issues in a given visualization.

#### 4.2.1 Actions for Rules

To validate the legitimacy, we invited two visualization experts to define feasible actions to solve each rule. In Table 1, all embedded actions are listed with a detailed explanation of what they can do with the visualization. Because each rule can correspond with several declarations in the specification, breaking the conditions for its establishment is the most effective way to resolve it. For example, if an encoding misuses the log scale, one can remove the log scale in the encoding or even change the encoded data field to another quantitative data field. Following such methodology, each rule has at least one feasible action and a maximum

<sup>1</sup>https://github.com/VizLinter/VizLinter-rules

Туре	Action	Meaning		
Mark	CHANGE_MARK	Change the mark type for the visualization.		
	ADD_CHANNEL	Add one encoding channel in the visualization.		
	CHANGE_CHANNEL	Change the encoding channel in the visualization.		
	REMOVE_CHANNEL	Remove the encoding channel in the visualization.		
Encoding	ADD_FIELD	Add data field in the encoding channel.		
	CHANGE_FIELD	Change the data field used in the encoding channel.		
	REMOVE_FIELD	Remove the data field declared in the encoding channel.		
	CHANGE_TYPE	Change the type of the encoding channel.		
	BIN	Discretize data values of the encoding channel into a set of bins.		
	REMOVE_BIN	Discard data binning in the encoding channel.		
	AGGREGATE	Perform aggregation on data values in the encoding channel.		
Transformation	CHANGE_AGGREGATE	Change aggregation function on the data values in the encoding channel.		
	REMOVE_AGGREGATE	Remove aggregation on data values in the encoding channel.		
	STACK	Stack data values in the encoding channel.		
	REMOVE_STACK	Remove stacking on data values in the encoding channel.		
	LOG	Apply logarithmic transformation on data values in the encoding channel.		
Scale	REMOVE_LOG	Remove logarithmic transformation on data values in the encoding channel.		
Scale	ZERO	Require a zero-baseline in the scaled domain of the encoding channel.		
	REMOVE_ZERO	Waive the zero-baseline requirement in the scaled domain of the encoding channel.		
	CORRECT_MARK	Change the illegal mark type to the closest correct one.		
	CORRECT_CHANNEL	Change the illegal channel to the closest correct one.		
<b>Typo Correction</b>	CORRECT_TYPE	Change the illegal data type to the correct one according to data values.		
	CORRECT_AGGREGATE	Change the illegal aggregation function to the closest correct one.		
	CORRECT_BIN	Change the illegal bin number to the correct one.		

Table 1. Action Space of VizLinter

of five actions, where a rule is established corresponding to several different situations.

#### 4.2.2 Score of Action

In single-problem cases, it is straightforward to find a solution. However, in situations where there are multiple problems, potential solutions are rarely unique. To find the optimal action sets to resolve the violated rules, we construct an optimization algorithm that considers the reward and the cost of actions.

We define each action's reward as its contribution to the visualization, that is, how many problems in the visualization can be solved after completing it. The reward is composed of two parts. First, we evaluate the proportion of problems solved in the visualization:

$$reward^{+}(a) = \frac{|R_i - R_{i+1}|}{|R_i|}$$
 (1)

where  $R_i$  and  $R_{i+1}$  represent the violated rulesets of visualization before and after the specific action *a*. Here we use the relative complement of  $R_i$  with respect to  $R_{i+1}$  to represent the solved problems. However, there are occasions where such actions can bring unintended consequences, one of which is violating new rules. For example, a visualization violating the rule that x-axis and y-axis cannot perform count aggregation simultaneously has two possible actions, removing count aggregation from the x-axis or y-axis. However, suppose the channel where the aggregation is removed happens to have no declared data field. A newly violated rule will then derive that an encoding channel should declare the data field or use count aggregation.

$$reward^{-}(a) = \frac{|R_{i+1} - R_i|}{|R_{i+1}|}$$
(2)

In this case, we punish the action *a* by the proportion discrepancy between  $R_{i+1}$  and  $R_i$  to show the new violated rules caused by it. Hence, the overall reward of an action *a* is computed by subtracting its input from its side effect. To concentrate on the positive effect of an action, we set the default weight w = 0.05 to tolerate the side effects brought by an action.

$$reward(a) = reward^{+}(a) - w \times reward^{-}(a)$$
 (3)

When making corrections, actions modify visualizations in different ways. The transition cost for the action, adapted from GraphScape [29], is defined to model the changes between visualization  $V_i$  and  $V_{i+1}$  resulting from the action. If multiple actions can solve an error in the visualization, then the action with the lowest transition cost is the best solution for the problem.

$$cost(a) = transition\_cost(V_i, V_{i+1})$$
 (4)

The overall score of an action is measured by its reward and cost. After trials during experiments, we set the default weight  $\alpha = 0.8$ ,  $\beta = 0.2$  to encourage the following algorithm to pay more attention to the benefits of actions:

$$score(a) = \alpha \times reward(a) - \beta \times cost(a)$$
 (5)

## 4.2.3 Optimization Algorithm

After calculating the scores of all candidate actions for rectifying a given visualization, the next problem to consider is which sequence and combination of actions can fix the errors with optimal scores.

Assume that a visualization has *n* broken rules  $r_1, r_2, ..., r_n$ . Each rule  $r_i$  has  $m_i$  candidate actions  $a_{i,1}, a_{i,2}, ..., a_{i,m_i}$ , implying that this rule has  $m_i$  possible solutions. Meanwhile, each action  $a_{i,j}$  has its corresponding score calculated by Formula 5. In the algorithm, we set the candidate action  $a_{i,j}$  as a binary variable to indicate whether to select it eventually.

Therefore, our problem can be converted into a Binary Integer Programming(BIP) problem as:

maximize 
$$\sum_{i}^{n} \sum_{j}^{m_{i}} score(a_{i,j}) \times a_{i,j}$$
(6a)

ubject to 
$$\sum_{i=1}^{m_i} a_{i,j} = 1, \quad i \in \{1, ..., n\}$$
 (6b)

$$a_{i,j} = a_{x,y}$$
, if  $a_{i,j}$  and  $a_{x,y}$  are equivalent actions (6c)

$$a_{i,j} \in \{0,1\}, \quad \begin{array}{l} i \in \{1,\dots,n\}\\ j \in \{1,\dots,m_i\} \end{array}$$
(6d)

s

Algorithm 1:	Update	Candidate	Actions	of Each Rule
--------------	--------	-----------	---------	--------------

	<b>Input</b> :action $\alpha$ , original violated ruleset $\mathcal{R}_i$ , violated rule $\mathcal{R}_{i+1}$ after performing action $\alpha$ , candidate actions.	set A				
	of each rule					
<b>Output :</b> updated $\mathcal{A}$						
	/* Traverse each rule $r$ that action $lpha$ has solved	*/				
1	for $r$ in $\mathcal{R}_i - \mathcal{R}_{i+1}$ do					
	<pre>/* Retrieve candidate action list of rule r</pre>	*/				
2	$A \leftarrow \mathcal{A}(r);$					
3	if $\alpha \notin A$ then					
4	update $\mathcal{A}$ by adding $\alpha$ in A;					
5	end					
6	end					
7 return $\mathcal{A}$ ;						

The above programming objective is to maximize the accumulative score of all candidate actions taken to fix the flawed visualization. To prevent redundant operations, Constraints 6b guarantee that one rule is fixed by only one of its candidate actions; therefore, the sum of the assigned value of actions from one rule can only be one. Since each rule's candidate actions are defined respectively as mentioned in Sec. 4.2.1, there may be equivalent actions in different rules, that is, actions with identical names (e.g., REMOVE\_AGGREGATE) performing in the same encoding (e.g., x). Constraints 6c ensure that any two equivalent candidate actions selected by different rules should be assigned with the same value, i.e., picking up identical actions simultaneously or discarding both. It is used to eliminate situations where a rule  $r_2$  uses a different solution when the problem could have been solved by using the adopted action of another rule  $r_1$ . Constraints 6d restrict values of all actions  $a_{i,i}$  to be either 0 or 1, indicating the status of not being selected or being selected respectively. Hence, the objective function 6a is determined solely by the scores of the adopted actions.

By providing a list of candidate actions, we address the above BIP problem by finding the optimal action set with the highest accumulative scores. BIP is an NP-Complete problem, whose solution can be found using the Branch-and-Cut algorithm [41] or various well-established solvers such as SCIP [3] and CBC [17]. We utilize the Python package PuLP [36] to solve the BIP problem in our implementation. PuLP is a widely-used linear programming API for defining problems and invoking multiple external solvers. We conduct all optimization using Python 3.7 with PuLP 2.4 with the default CBC solver on a computer equipped with 2.4 GHz Intel Core i5 and 16GB RAM. All the testing cases for our problem can be solved within seconds.

The above optimization has a relatively good performance, especially when all the violated rules in the visualization and their corresponding candidate actions are independent. However, in some testing cases, candidate action  $a_{\alpha}$  of rule  $r_{\alpha}$  helped to fix the issue brought by rule  $r_{\beta}$ , which was also solved by action  $a_{\beta}$  adopted by the optimization algorithm. When defining each rule's feasible actions as its possible solutions, the experts only consider the rule itself, regardless of the co-occurrence of other rules. Therefore, cases occur when a rule is solved by a non-predetermined action. To address such an issue, we update constraints 6b by expanding candidate actions of Rule  $r_i$  with the actions of other rules that also contribute to solving  $r_i$  as described in Algorithm 1. We then solve the BIP model with the updated constraints. In the optimization result, no more redundant actions are recommended, providing users with simple and straightforward fixing suggestions.

Fig. 1 shows four examples of how VizLinter detects issues in the visualization (shown as Before) and then corrects them (shown as After). The data used in the examples are from the sample dataset of Vega-Lite<sup>2</sup>. Fig. 1(a) and Fig. 1(b) are built based on the *car* dataset, including records of cars and their basic properties. The original visualization of Fig. 1(a-Before) depicts the relationship of Horsepower, Miles\_per\_Gallon of cars and their Origin, where the size channel is not compatible with the nominal data field Origin (I1). VizLinter rectifies it by changing the size channel to the color channel as shown

in Fig. 1(a-After). One rule of **I3** type is violated in Fig. 1(b-Before), where mark type point is not suitable to depict data with stacking. Changing the chart type by modifying mark type to bar corrects the visualization, recommended by VizLinter. Fig. 1(c) and Fig. 1(d) visualize data about daily weather reports from the *seattle-weather* dataset. Fig. 1(c) before modification misuses the size channel with the data field temp\_min containing negative values (**I1**) and performs log transformation incorrectly in the y-axis encoded the data field temp\_max holding negative values (**I1**). VizLinter removes the log transformation in the y-axis and substitutes the size channel with the color channel. For the visualization of Fig. 1(d), both x-axis and y-axis execute count aggregation. As a result, only a single bar is visualized (**I2**). VizLinter discards the aggregation on the x-axis to represent the number of appearances of different weather types.

## 5 IMPLEMENTATION

Following the framework described above, we develop a Python package, vega-lite-linter<sup>3</sup>, which embeds a function lint() to detect errors in Vega-Lite syntax and a function fix() to provide the optimal solutions with the fewest steps. By automatically detecting and resolving errors, vega-lite-linter enables visualization developers to build accurate charts quickly in Python. Meanwhile, fix() provides alternative actions that can resolve each rule and ranks them in the order of scores calculated by Formula 5. Users can refer to this action list and choose the solutions to fulfill their own needs.

Based on the above package, we then create a web prototype<sup>4</sup> with the Python backend, as shown in Fig. 3, so that it is easily accessible to any Vega-Lite developer to build and validate visualizations. We customize the interface to better serve the framework, referring to the original Vega-Lite Online Editor [44]. In addition to the basic code editing panel and chart render panel, we integrate a prompt panel for the linting and fixing functions of VizLinter as shown in Fig. 3(d). When clicking the "Inspect Specs" button, any violated rules of the current visualization specification will be shown. Users can correct the specification on their own according to the description of violated rules or check our suggested actions by clicking the "Suggest Revision" button. Users can click the "Preview" button to preview the automatic modification result. The corresponding revised specification will show up in the code panel. "Accept" and "Reject" buttons are provided for users to approve then apply or reject the revisions. Meanwhile, the chart panel will render the visualization with the revised specification.

## 6 EVALUATION

## 6.1 User Study

The user study examined the effectiveness of the proposed framework. Specifically, we invited the participants to use VizLinter and asked whether they agreed with the issues identified by VizLinter and the corresponding solutions. After the study, we analyzed and discussed both the quantitative and qualitative results.

#### 6.1.1 Participants

We invited 20 participants with visualization experience to take part in our study. Twelve of them are from a technology and financial services company (four females and eight males), whose daily work involves developing visualizations with business intelligence tools. The other 8 participants are visualization researchers (six females and two males). The participants' ages range from 23 to 30 (M = 25.6, SD = 2.0).

## 6.1.2 Preparation

We prepared 15 flawed visualizations in Vega-Lite JSON format using the three sample datasets from Vega-lite, including *cars*, *airports*, and *seattle-weather*. Among these visualizations, five included one error, five two errors, and the rest contained three errors. There were 25 distinct rules involved in these cases, covering 61% of those in our framework's rule base.

<sup>&</sup>lt;sup>3</sup>http://vegalite-linter.idvxlab.com/

<sup>&</sup>lt;sup>4</sup>http://vizlinter.idvxlab.com/

© 2021 IEEE. This is the author's version of the article that has been published in IEEE Transactions on Visualization and Computer Graphics. The final version of this record is available at: xx.xxx/TVCG.201x.xxxxxx/



Fig. 3. The interface of VizLinter's web prototype consists of four views: (a) the code editor with highlighted differences from the original Vega-Lite JSON, (b) the optimized chart after revision, (c) the original chart before revision, and (d) the toggles (and detailed messages) of the linter, the fixer and its modification preview.

We developed a user study system adapted from the VizLinter prototype, whose interface was split into two panels to enable direct comparisons. The left panel presented the original Vega-Lite specification with flawed visualization in a read-only mode. The editable code panel on the right was provided for participants to revise and edit on their own, with the corresponding visualization rendered simultaneously. After participants finished revising each question, they were encouraged to click "Inspect Specs" and "Suggest Revision" to view the linter detected violations and the fixer's suggestion by VizLinter. If participants agreed on the suggestions, they could click "Accept Suggestion" to apply the advised actions on the original visualization automatically. After accepting the fixer's suggestion, participants were still allowed to edit the revised visualization further.

#### 6.1.3 Procedure

The study started with a twenty-minute training session about Vega-Lite grammars. We also introduced the four categories of common issues  $(I1 \sim I4)$  within visualizations and a demo case of the upcoming tasks. To verify they understood the instructions, we asked them to try the test questions and answered their questions regarding the user study. Each participant was then given 15 flawed visualizations with the original Vega-Lite specifications. The participants were asked to find out the errors and fix them according to their visualization knowledge and experience. Their revisions in the Vega-Lite specifications were highlighted in green on the editable code panel, and the mutated lines in the original specification were highlighted in red on the original code panel. When the participants completed the task or felt stuck, they were asked to use the linter and the fixer functions of our system. Two questions were then asked: (1) whether they acknowledged the detected errors by the linter and the proposed solutions by the fixer, and (2) whether they would like to make further edits on the improved chart. After each participant finished all the tasks, we conducted a brief interview to collect feedback on general user experience, effectiveness, and suggestions for our VizLinter framework.

## 6.2 Analysis & Results

We obtained participants' revisions and their acceptance decisions for each suggestion provided by VizLinter, as well as their completion time of each question and the correction rate. We first analyzed the effectiveness and efficiency of our framework by computing the completion time, the correction rate, and the acceptance rate. Then, we summarized their feedback from interviews to further evaluate the prototype system.

## 6.2.1 Quantitative Analysis

**Completion Time.** It took the participants 97.7 seconds on average to find and fix errors in each visualization (SD = 28.1). The corresponding mean completion time for correction of the different number of errors (one, two, and three) in the visualization was 101.3 seconds (SD = 23.0), 93.8 seconds (SD = 27.9), and 98.1 seconds (SD = 32.1), respectively. Contrary to our expectations, the error fixing time did not increase with the number of errors alone. Instead, it was related to the difficulty of figuring out solutions and the complexity of manipulating the specification as well. For example, adding a new encoding channel resulted in more steps than simply modifying the aggregation function within an encoding.

*Correction Rate.* We inspected the edited specifications and recorded the number of remaining issues after user revision to calculate the correction rate. The average correction for all the questions was 77% (SD = 19.7%). We also calculated the average correction rate of four types of common issues. Rules related to issues within each encoding channel (**I1**) had a 75.2% (SD = 22.7%) correction rate, while the average correction rate of those related to issues across multiple encoding channels (**I2**) was 70.8% (SD = 20.8%). Rules related to issues between encoding channels and marks (**I3**) gained an average correction rate of 82.5% (SD = 14.4%), and typo issues (**I4**) reached 90% (SD = 5%) in average correction rate. The latter two types of errors in VizLinter seemed easier to resolve than the first two, potentially because some details inside encoding channels were more difficult to be noticed or fixed.

We summarized some cases that the participants failed to correct and analyzed their causes. A typical case is when participants did not notice the error and could not figure out feasible solutions based on their knowledge. For example, Question 2 was corrected by only four participants, where a rule "Use at most 20 categorical colors in the visualization" was violated. Some participants did not recognize this problem due to unawareness of the corresponding design principle or habitual neglect of color overuse in their daily practice. While some participants noticed this error, they still failed to resolve it. This is due to the difficulty of manipulating encoding channels, as optimizing color encoding is also an important research topic for effective visualization



Fig. 4. One case from the user study: (a) the original visualization with errors; (b) the suggested modification result made by VizLinter; (c),(d),(e),(f) four modificated results by different participants.

design [56]. Another rule, "*Channel size is not suitable for data with negative values*," involved in Question 8 and Question 14, also led to low correction rates because only some participants realized that it was illegal to encode size channel for negative data.

Among the rules embedded in VizLinter, some violations do not intervene with visualization rendering due to built-in exception handling of Vega and Vega-Lite. One example is shown in Fig. 3, where two rules "use log scale with non-discrete data" and "use valid aggregation, including count, mean, min, etc." were violated, but the visualization rendering was not affected by the invalid log scale. Such errors could only be found in the codes and sometimes were ignored by the participants. We believe that specifications with unnecessary declarations would be error-prone in creating visualizations. Hence, VizLinter can facilitate users to identify errors that are not obviously reflected in the rendered visualizations.

Acceptance Rate. The average acceptance rate was 90% (SD = 14.2%), indicating that most of the suggestions were adopted. For those recommended solutions with a low acceptance rate, there were four types of scenarios. First, some violated rules were questioned by a few participants. For example, some participants disagreed with the rule requiring the y-axis to start from zero in bar charts. Consequently, no action was taken for this violation. Second, after comparing the revision by VizLinter and by themselves, some participants favored their own revisions. For example, one case shown in Fig. 4(a) involved two violated rules: "Mark bar, tick, line, and area require some continuous variables on x- or y-axis" and "Use both binning and aggregation on the data at the same time is illegal." VizLinter's one-step fixer suggestion was to remove the bin operation on the y-axis shown in Fig. 4(b). While some participants made the same choice as VizLinter, others attempted different solutions and rejected the suggestion provided by VizLinter. A common alternative was to remove the bin operation of both x- and y-axis and to change the mark type from bar to line (Fig. 4(c)) or point (Fig. 4(d)), then to discard the aggregation on the y-axis to depict the distribution (Fig. 4(e)). Another interesting result from one developer was to ignore the former rule and only remove the aggregation on the x-axis, as shown in Fig. 4(f). "I did see such kind of chart at work," explained the participant. For the same visualization with flaws, VizLinter and the participants sometimes gave different solutions for revising the specification to correct the chart. Since the participants' revisions are influenced by their preferences in visualization, it is difficult to evaluate these subjective choices. However, in future studies, it might be possible to assess the revisions based on design considerations such as aesthetics. Third, some recommended actions were regarded as neglecting user intents, therefore rejected by some

participants. For example, in Question 12, the encoding specification broke three rules by binning and aggregating simultaneously on nominal data: "Only use binning on quantitative or ordinal data," "Nominal data cannot be aggregated," and "Use both binning and aggregation on the data at the same time is illegal." The fixer recommended RE-MOVE\_AGGREGATE and CHANGE\_FIELD to correct the violations. On the contrary, the participants took the REMOVE AGGREGATE and REMOVE\_BIN actions to fix the errors to avoid changing the original data mapping. However, this combination broke the framework's constraint that only one candidate action for each rule could be selected since REMOVE\_AGGREGATE and REMOVE\_BIN were both candidates for the rule "Use both binning and aggregation on the data at the same time is illegal." This case indicates one limitation of VizLinter as it takes only one candidate action for a single rule other than combinational actions. The last cause for rejection was that some participants believed there were other solutions outside of the action scope to fix the problem. For example, the suggested revision for Question 2 mentioned above was to remove the whole color encoding since it was not suitable to describe the nominal data over 20 categories. Though it was the best option within the current scope, some participants argued the chart after applying this action was still not good enough, even if it did meet the basic design principles, "I expect another chart type or more aggregation on the color channel instead of deleting it."

After reviewing the edited specifications, we discovered new errors after human revisions. While some revisions partially or fully correct the original problems, new errors were introduced into the visualizations. Twenty-three of all the 300 chart modifications (15 questions × 20 participants) violated rules that did not appear in the original problems. This result illustrates our VizLinter framework's advantage over manual adjustments in that it prevents new problems after modifications to the original specifications.

## 6.2.2 Qualitative Analysis

To further evaluate the usability of VizLinter, we interviewed the participants after the completion of all the tasks.

Overall Performance. All the participants agreed that VizLinter was helpful as it could automatically find errors in visualizations and provide suggestions to fix them. One participant commented, "for those who are not familiar with the visualization guidelines, it is challenging to make a legitimate visualization on their own. Tools as VizLinter would really help a lot." In terms of efficiency, all the participants responded that VizLinter saved a lot of time and made it possible to revise a chart with a single click. They appreciated its efficiency and commented, "this really helps me save time, especially when I attempt to create charts with unfamiliar datasets."

Instructive Value. Some participants mentioned that VizLinter could help quickly learn visual design guidelines during the linting and fixing process. One participant said, "when I employ this tool, I can pick up a lot of visualization principles which I used to neglect." Another participant noted, "this tool expands my ideas, helps me think out of the box and offers some options for modification."

*Suggestions.* In spite of the above positive feedback, we also summarize and discuss the suggestions from the participants.

- Consider user intent. Recommended solutions for visualizations by VizLinter usually result in minimal changes, but they do not always match users' expectations. When the number of revision actions is not a concern, users often have different ideas for making corrections. More than half of the participants pointed out that user intent should be considered. One participant said, *"I would like to see the tool taking more user intentions into account and suggest more revision options to choose from."* Another participant suggested inferring user intents from the initial chart with errors or user modifications.
- Consider semantic meaning. Some participants expected a more intelligent tool that can understand the semantics of the data field to better revise the chart. One participant mentioned, "I found that VizLinter cannot detect whether the semantics in the chart is correct. It would be better if the tool could provide more revision suggestions based on the semantic meaning of data."
- Add more visual principles. Some participants suggested that more rules could be integrated into VizLinter. For example, rules to detect missing values or outliers can be added to ensure the validity of the rendered visualization. Rules regarding visual aesthetics are also desired, "some visual guidelines on aesthetics could be added to help make the charts more visually pleasing."

## 7 DISCUSSION

In this section, we discuss the potential use scenarios of the VizLinter framework, as well as current limitations to guide future research.

## 7.1 Potential Use Scenarios

**VizLinter for other grammars.** In this work, we utilize Vega-Lite as the visualization specification syntax. In the future work, VizLinter can incorporate other popular visualization grammars such as D3.js, ggplot2, and Matplotlib. One way of doing so is to translate onward-and-backward between Vega-Lite and other specification languages. Another solution is to implement various versions of VizLinter respectively to suit different grammars. In practice, a BI tool with over 100,000 users for the tech company Ant Group has already adapted VizLinter into its visualization library in Javascript<sup>5</sup> and configured its internal design guidelines into the linter to meet user needs. This case proves the adaptability of VizLinter in other programming languages.

VizLinter in BI tools. The framework of VizLinter can be applied in business intelligent tools. Although most BI tools provide drag-anddrop interactions for users to build visualizations without coding, some wrong actions may affect the correct representation of the information, such as mapping data to an inappropriate channel or applying the wrong data transformation. Embedding VizLinter in BI tools could save time for digging into design guidelines, which is especially handy for inexperienced users who need to create visualizations quickly. The BI tool mentioned above has received positive feedback in practice, which enables users to find out optimal configurations quickly. Furthermore, VizLinter can serve as a collaborative visualization design system, ensuring that various visualizations built by different users in multiple views are consistent and compatible [42]. For example, linting tools could remind users to pay attention to set margins or encode with specific color palettes following the homogeneous design guidelines.

**VizLinter for education.** During the interview, many participants felt surprised that VizLinter helped them recognize design guidelines of visualization, some of which were overlooked occasionally. Consequently, one of the potential applications of VizLinter is an auxiliary teaching tool for visualization education. Novices could use VizLinter

to verify whether the visualizations they create violate any design guidelines. During the trial-and-error process, one can obtain visualization knowledge gradually, which also echos the pedagogical value of ES-Lint [52]. VizLinter can be valuable for studying visualization design guidelines, while instructors could also benefit from the automatically linting function to make quick judgment when checking assignments.

## 7.2 Limitations and Future Work

We summarize several limitations mentioned during interviews and found in the design and implementation process and propose future work directions.

*Explain the violated rules in more expressive ways.* Currently, violated rules detected by the linter are only displayed in plain text. When interviewing the participants about the functionalities of VizLinter, more details about the rules were reported as desired, such as in-depth descriptions and graphical demonstration of the rules. In the future, it is necessary to explain each rule in more detail, including complete documentation and examples of typical documents and don'ts illustrated with code and graphics.

*Extend the coverage of rule categories.* All the four types of rules covered in VizLinter focus on basic construction errors of visualizations. In the current stage, guidelines regarding perceptual expressiveness or visual encoding effectiveness have not been considered. Due to subjective considerations, there is no clear right or wrong judgment for such errors. Hence, more pilot studies should be conducted before applying subjective guidelines to VizLinter. In the adapted version for the tech company mentioned earlier, some soft rules configured by the domain experts were applied. In the next step, we plan to collect and evaluate different categories of rules, covering expressiveness, aesthetic or stylistic issues, to enrich the capability of VizLinter.

<u>Enrich system configurations and user interaction</u>. The current design of the framework restrains the capability of error handling to the principles included in the rule base. However, in practice, it can be more helpful if the rule base allows extension. Therefore, in the next step, we consider making the rules user-configurable so that VizLinter can fit in different scenarios. One implementation of VizLinter is the web editor with linting and fixing functions, where limited interactions were provided. It is not intuitive enough for users to refer each rule to its corresponding coding snippets when the broken rules are listed in a separate panel. How to present the errors and the linkage among rules, codes, and visualization will be one of the future works of VizLinter.

*Employ more state-of-art research to resolve performance issues.* Currently, the scopes of design rules and actions in VizLinter are relatively small, where input cases can be solved within seconds. It is undeniable that, with the progress of development, the ever-growing rule base and action space would increase the execution time of VizLinter. Keep improving the framework for better performance utilizing state-or-art research [15, 18, 38] is worth further studies.

## 8 CONCLUSION

We present VizLinter, a framework for automatically detecting flaws in visualizations and suggesting revisions, consisting of a visualization linter and a visualization fixer. The visualization linter employs design guidelines collected from prior research and incorporates Answer Set Programming to identify issues of the input visualizations. The visualization fixer resolves the detected violations by the linter and automatically suggests revisions by formulating the problem as a binary integer programming problem. A prototype of VizLinter is implemented as an online Vega-Lite editor, showing how developers can use VizLinter to create visualization without flaws. An in-lab user study was also conducted to evaluate the effectiveness and efficiency of VizLinter. The results showed that VizLinter could help users identifying and fixing errors. In future work, we plan to extend the current framework with larger coverage of rules, more system configurations and user interactions, and to employ it in broader application scenarios.

#### ACKNOWLEDGMENTS

<sup>5</sup>https://ava.antv.vision/en/docs/guide/chart-linter/intro

This work was supported in part by the National Natural Science Foundation of China 62072338, 6200070909.

## REFERENCES

- [1] Eslint. https://eslint.org/. [Online; accessed 25-March-2021].
- [2] M. M. Abbas, M. Aupetit, M. Sedlmair, and H. Bensmail. ClustMe: A Visual Quality Measure for Ranking Monochrome Scatterplots based on Cluster Patterns. *Computer Graphics Forum*, 38(3):225–236, 2019. doi: 10.1111/cgf.13684
- [3] T. Achterberg. Scip: solving constraint integer programs. *Mathematical Programming Computation*, 1(1):1–41, 2009. doi: 10.1007/s12532-008-0001-1
- [4] D. W. Barowy, E. D. Berger, and B. Zorn. ExceLint: automatically finding spreadsheet formula errors. *Proceedings of the ACM on Programming Languages*, 2(OOPSLA):148:1–148:26, Oct. 2018. doi: 10.1145/3276518
- [5] D. W. Barowy, D. Gochev, and E. D. Berger. CheckCell: data debugging for spreadsheets. ACM SIGPLAN Notices, 49(10):507–523, Oct. 2014. doi: 10.1145/2714064.2660207
- [6] M. Behrisch, B. Bach, N. Henry Riche, T. Schreck, and J.-D. Fekete. Matrix Reordering Methods for Table and Network Visualization. *Computer Graphics Forum*, 35(3):693–716, June 2016. doi: 10.1111/cgf.12935
- [7] J. Bertin and W. J. Berg. Semiology of graphics: diagrams, networks, maps. ESRI Press : Distributed by Ingram Publisher Services, Redlands, Calif, 1st ed., 2010.
- [8] E. Bertini, A. D. Girolamo, and G. Santucci. See What You Know: Analyzing Data Distribution to Improve Density Map Visualization. *Euro-graphics/IEEE-VGTC Symposium on Visualization*, pp. 163–170, 2007. doi: 10.2312/VISSYM/EUROVIS07/163-170
- [9] M. A. Borkin, A. A. Vo, Z. Bylinskii, P. Isola, S. Sunkavalli, A. Oliva, and H. Pfister. What Makes a Visualization Memorable? *IEEE Transactions* on Visualization and Computer Graphics, 19(12):2306–2315, Dec. 2013. doi: 10.1109/TVCG.2013.234
- [10] M. Brachmann, C. Bautista, S. Castelo, S. Feng, J. Freire, B. Glavic, O. Kennedy, H. Müeller, R. Rampin, W. Spoth, and Y. Yang. Data Debugging and Exploration with Vizier. In *Proceedings of the 2019 International Conference on Management of Data*, SIGMOD'19, pp. 1877–1880. Association for Computing Machinery, Amsterdam, Netherlands, June 2019. doi: 10.1145/3299869.3320246
- [11] G. Brewka, T. Eiter, and M. Truszczyński. Answer set programming at a glance. *Communications of the ACM*, 54(12):92–103, Dec. 2011. doi: 10. 1145/2043174.2043195
- [12] L. Byron and M. Wattenberg. Stacked Graphs Geometry Aesthetics. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1245– 1252, Nov. 2008. doi: 10.1109/TVCG.2008.166
- [13] S. M. Casner. Task-analytic approach to the automated design of graphic presentations. ACM Transactions on Graphics, 10(2):111–151, Apr. 1991. doi: 10.1145/108360.108361
- [14] W. S. Cleveland and R. McGill. Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods. *Journal of the American Statistical Association*, 79(387):531–554, Sept. 1984. doi: 10.1080/01621459.1984.10478080
- [15] B. Cuteri, C. Dodaro, F. Ricca, and P. Schüller. Overcoming the grounding bottleneck due to constraints in asp solving: Constraints become propagators. In *Proceedings of the Twenty-Ninth International Joint Conference* on Artificial Intelligence (IJCAI'20), pp. 1688–1694. ijcai. org, 2020.
- [16] G. Ellis and A. Dix. Enabling Automatic Clutter Reduction in Parallel Coordinate Plots. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):717–724, Sept. 2006. doi: 10.1109/TVCG.2006.138
- [17] J. Forrest and R. Lougee-Heimer. Cbc user guide. In *Emerging theory, methods, and applications*, pp. 257–277. INFORMS, 2005. doi: 10.1287/EDUC.1053.0020
- [18] M. Gebser, R. Kaminski, B. Kaufmann, and T. Schaub. Multi-shot asp solving with clingo. *Theory and Practice of Logic Programming*, 19(1):27– 82, 2019.
- [19] D. Gotz, W. Wang, A. T. Chen, and D. Borland. Visualization model validation via inline replication. *Information Visualization*, 18(4):405–425, Oct. 2019. doi: 10.1177/1473871618821747
- [20] D. Gotz and Z. Wen. Behavior-driven visualization recommendation. In Proceedings of the 14th International Conference on Intelligent User Interfaces, IUI'09, pp. 315–324. Association for Computing Machinery, Sanibel Island, Florida, USA, Feb. 2009. doi: 10.1145/1502650.1502695
- [21] J. Heer and M. Agrawala. Multi-Scale Banking to 45 Degrees. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):701–708, Sept. 2006. doi: 10.1109/TVCG.2006.163
- [22] J. Heer and B. Shneiderman. Interactive Dynamics for Visual Analysis: A

taxonomy of tools that support the fluent and flexible use of visualizations. *Queue*, 10(2):30–55, Feb. 2012. doi: 10.1145/2133416.2146416

- [23] A. K. Hopkins, M. Correll, and A. Satyanarayan. VisuaLint: Sketchy In Situ Annotations of Chart Construction Errors. *Computer Graphics Forum*, 39(3):219–228, 2020. doi: 10.1111/cgf.13975
- [24] K. Hu, M. A. Bakker, S. Li, T. Kraska, and C. Hidalgo. VizML: A Machine Learning Approach to Visualization Recommendation. In *Proceedings* of the 2019 CHI Conference on Human Factors in Computing Systems, CHI'19, pp. 1–12. Association for Computing Machinery, Glasgow, Scotland Uk, May 2019. doi: 10.1145/3290605.3300358
- [25] N. Hynes, D. Sculley, and M. Terry. The Data Linter: Lightweight Automated Sanity Checking for ML Data Sets. In *Proceedings of the 31st Conference on Neural Information Processing Systems*, NIPS'17. Long Beach, CA, USA, 2017.
- [26] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller. A Systematic Review on the Practice of Evaluating Visualization. *IEEE Transactions* on Visualization and Computer Graphics, 19(12):2818–2827, Dec. 2013. doi: 10.1109/TVCG.2013.126
- [27] S. Kandel, A. Paepcke, J. M. Hellerstein, and J. Heer. Enterprise Data Analysis and Visualization: An Interview Study. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2917–2926, Dec. 2012. doi: 10.1109/TVCG.2012.219
- [28] D. Keim. Designing pixel-oriented visualization techniques: theory and applications. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):59–78, Mar. 2000. doi: 10.1109/2945.841121
- [29] Y. Kim, K. Wongsuphasawat, J. Hullman, and J. Heer. GraphScape: A Model for Automated Reasoning about Visualization Similarity and Sequencing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI'17, pp. 2628–2638. Association for Computing Machinery, New York, NY, USA, May 2017. doi: 10.1145/ 3025453.3025866
- [30] R. M. Kirby and C. T. Silva. The need for verifiable visualization. *IEEE Computer Graphics and Applications*, 28(5):78–83, Oct. 2008. doi: 10. 1109/MCG.2008.103
- [31] V. Lifschitz. What is answer set programming? In Proceedings of the 23rd national conference on Artificial intelligence - Volume 3, AAAI'08, pp. 1594–1597. AAAI Press, Chicago, Illinois, July 2008.
- [32] H. Lin, D. Moritz, and J. Heer. Dziban: Balancing Agency & Automation in Visualization Design via Anchored Recommendations. In *Proceedings* of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, pp. 1–12. Association for Computing Machinery, Honolulu, HI, USA, Apr. 2020. doi: 10.1145/3313831.3376880
- [33] J. Mackinlay. Automating the design of graphical presentations of relational information. ACM Transactions on Graphics, 5(2):110–141, Apr. 1986. doi: 10.1145/22949.22950
- [34] A. McNutt and G. Kindlmann. Linting for Visualization: Towards a Practical Automated Visualization Guidance System. In VisGuides: 2nd Workshop on the Creation, Curation, Critique and Conditioning of Principles and Guidelines in Visualization, 2018.
- [35] A. McNutt, G. Kindlmann, and M. Correll. Surfacing Visualization Mirages. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, pp. 1–16. Association for Computing Machinery, Honolulu, HI, USA, Apr. 2020. doi: 10.1145/3313831.3376420
- [36] S. Mitchell, M. OSullivan, and I. Dunning. Pulp: a linear programming toolkit for python. *The University of Auckland, Auckland, New Zealand*, p. 65, 2011.
- [37] D. Moritz, C. Wang, G. L. Nelson, H. Lin, A. M. Smith, B. Howe, and J. Heer. Formalizing Visualization Design Knowledge as Constraints: Actionable and Extensible Models in Draco. *IEEE Transactions on Vi*sualization and Computer Graphics, 25(1):438–448, Jan. 2019. doi: 10. 1109/TVCG.2018.2865240
- [38] D. R. Morrison, S. H. Jacobson, J. J. Sauppe, and E. C. Sewell. Branch-andbound algorithms: A survey of recent advances in searching, branching, and pruning. *Discrete Optimization*, 19:79–102, 2016.
- [39] T. Munzner. Visualization Analysis and Design. CRC Press, 2015. OCLC: 1105776954.
- [40] K. Muşlu, Y. Brun, and A. Meliou. Preventing data errors with continuous testing. In *Proceedings of the 2015 International Symposium on Software Testing and Analysis*, ISSTA 2015, pp. 373–384. Association for Computing Machinery, Baltimore, MD, USA, July 2015. doi: 10.1145/2771783. 2771792
- [41] M. Padberg and G. Rinaldi. A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems. SIAM review,

33(1):60-100, 1991. doi: 10.1137/1033004

- [42] Z. Qu and J. Hullman. Keeping multiple views consistent: Constraints, validations, and exceptions in visualization authoring. *IEEE Transactions* on Visualization and Computer Graphics, 24(1):468–477, 2018. doi: 10. 1109/TVCG.2017.2744198
- [43] B. Saket, A. Endert, and C. Demiralp. Task-based effectiveness of basic visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 25(7):2505–2512, 2019. doi: 10.1109/TVCG.2018.2829750
- [44] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-Lite: A Grammar of Interactive Graphics. *IEEE Transactions on Visualization* and Computer Graphics, 23(1):341–350, Jan. 2017. doi: 10.1109/TVCG. 2016.2599030
- [45] C. Seifert, B. Kump, W. Kienreich, G. Granitzer, and M. Granitzer. On the Beauty and Usability of Tag Clouds. In 2008 12th International Conference Information Visualisation, pp. 17–25. IEEE, London, UK, July 2008. doi: 10.1109/IV.2008.89
- [46] J. Seo and B. Shneiderman. A Rank-by-Feature Framework for Interactive Exploration of Multidimensional Data. *Information Visualization*, 4(2):96– 113, June 2005. doi: 10.1057/palgrave.ivs.9500091
- [47] D. Shi, Y. Shi, X. Xu, N. Chen, S. Fu, H. Wu, and N. Cao. Task-oriented optimal sequencing of visualization charts. In 2019 IEEE Visualization in Data Science (VDS), pp. 58–66, 2019. doi: 10.1109/VDS48975.2019. 8973383
- [48] M. Sondag, B. Speckmann, and K. Verbeek. Stable Treemaps via Local Moves. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):729–738, Jan. 2018. doi: 10.1109/TVCG.2017.2745140
- [49] C. Stolte, D. Tang, and P. Hanrahan. Polaris: A System for Query, Analysis, and Visualization of Multidimensional Relational Databases. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):52–65, Jan. 2002. doi: 10.1109/2945.981851
- [50] T. Tang, R. Li, X. Wu, S. Liu, J. Knittel, S. Koch, L. Yu, P. Ren, T. Ertl, and Y. Wu. Plotthread: Creating expressive storyline visualizations using reinforcement learning. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):294–303, 2021. doi: 10.1109/TVCG.2020.3030467
- [51] A. Tatu, G. Albuquerque, M. Eisemann, J. Schneidewind, H. Theisel, M. Magnork, and D. Keim. Combining automated analysis and visualization techniques for effective exploration of high-dimensional data. In 2009 *IEEE Symposium on Visual Analytics Science and Technology*, pp. 59–66. IEEE, Atlantic City, NJ, USA, 2009. doi: 10.1109/VAST.2009.5332628
- [52] K. F. Tómasdóttir, M. Aniche, and A. Van Deursen. The adoption of javascript linters in practice: A case study on eslint. *IEEE Transactions* on Software Engineering, 46(8):863–891, 2018. doi: 10.1109/TSE.2018. 2871058
- [53] M. Trutschl, G. Grinstein, and U. Cvek. Intelligently resolving point occlusion. In *IEEE Symposium on Information Visualization 2003 (IEEE Cat. No.03TH8714)*, pp. 131–136, Oct. 2003. doi: 10.1109/INFVIS.2003. 1249018
- [54] M. Vartak, S. Rahman, S. Madden, A. Parameswaran, and N. Polyzotis. SeeDB: Efficient Data-Driven Visualization Recommendations to Support Visual Analytics. *Proceedings of the VLDB Endowment International Conference on Very Large Data Bases*, 8(13):2182–2193, Sept. 2015. doi: 10.14778/2831360.2831371
- [55] P. Wang and Y. He. Uni-Detect: A Unified Approach to Automated Error Detection in Tables. In *Proceedings of the 2019 International Conference* on Management of Data, SIGMOD'19, pp. 811–828. Association for Computing Machinery, Amsterdam, Netherlands, June 2019. doi: 10. 1145/3299869.3319855
- [56] Y. Wang, X. Chen, T. Ge, C. Bao, M. Sedlmair, C. Fu, O. Deussen, and B. Chen. Optimizing Color Assignment for Perception of Class Separability in Multiclass Scatterplots. *IEEE Transactions on Visualization* and Computer Graphics, 25(1):820–829, Jan. 2019. doi: 10.1109/TVCG. 2018.2864912
- [57] Y. Wang, Y. Wang, Y. Sun, L. Zhu, K. Lu, C.-W. Fu, M. Sedlmair, O. Deussen, and B. Chen. Revisiting Stress Majorization as a Unified Framework for Interactive Constrained Graph Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):489–499, Jan. 2018. doi: 10.1109/TVCG.2017.2745919
- [58] C. Ware. *Information visualization: perception for design*. Interactive technologies. Morgan Kaufmann, Waltham, MA, 3rd ed., 2013.
- [59] Z. Wen and M. X. Zhou. An optimization-based approach to dynamic data transformation for smart visualization. In *Proceedings of the 13th International Conference on Intelligent User Interfaces*, IUI'08, pp. 70–79. Association for Computing Machinery, Gran Canaria, Spain, Jan. 2008.

#### doi: 10.1145/1378773.1378784

- [60] L. Wilkinson and G. Wills. *The grammar of graphics*. Statistics and computing. Springer, New York, 2nd ed., 2005.
- [61] G. Wills and L. Wilkinson. AutoVis: automatic visualization. *Information Visualization*, 9(1):47–69, Mar. 2010. doi: 10.1057/ivs.2008.27
- [62] K. Wongsuphasawat, D. Moritz, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):649–658, Jan. 2016. doi: 10.1109/TVCG.2015 .2467191
- [63] A. Wu, W. Tong, T. Dwyer, B. Lee, P. Isenberg, and H. Qu. Mobilevisfixer: Tailoring web visualizations for mobile phones leveraging an explainable reinforcement learning framework. *IEEE Transactions on Visualization* and Computer Graphics, 27(2):464–474, 2021. doi: 10.1109/TVCG.2020. 3030423
- [64] A. Wu, Y. Wang, X. Shu, D. Moritz, W. Cui, H. Zhang, D. Zhang, and H. Qu. Survey on artificial intelligence approaches for visualization data. arXiv preprint arXiv:2102.01330, 2021.