Journal of
Information
Systems
Education

Volume 33
Issue 4
Fall 2022

Teaching Tip
Evaluating Visualizations with a Compact Rubric

William K. McHenry

Recommended Citation: McHenry, W. K. (2022). Teaching Tip: Evaluating Visualizations with a Compact Rubric. *Journal of Information Systems Education*, 33(4), 324-337.

Article Link: https://jise.org/Volume33/n4/JISE2022v33n4pp324-337.html

Initial Submission: 23 December 2021
Minor Revision: 31 January 2022
Accepted: 22 March 2022
Published: 15 December 2022

ISSN: 2574-3872 (Online) 1055-3096 (Print)

Journal of Information Systems Education, 33(4), 324-337, Fall 2022

Teaching Tip
Evaluating Visualizations with a Compact Rubric

William K. McHenry
Management Department
University of Akron
Akron, OH 44325, USA
wm@uakron.edu

ABSTRACT

Students now have readily available and powerful tools to access, manipulate, combine, and visualize data. Acquiring data and visual literacy requires more than knowledge of how to use these tools. Students need to engage with assignments that challenge them to make relatively complex visualizations, interpret them, and explain why these interpretations matter for given problem situations. This paper describes how to structure feedback for these assignments. The few published visualization evaluation rubrics are mainly oriented toward how-to-do-it heuristics. This paper makes a contribution by presenting, defining, and giving examples of the use of an innovative compact rubric for evaluating visualizations (CRVE). This rubric eliminates some of the length and complexity of heuristic evaluation, focusing on interpretation and relevance. In a graduate business intelligence course, students showed definite improvement over the course of the semester in the construction of visualizations, telling a story with headlines, and striving for data exploration. However, higher levels of technical correctness of visualizations did not necessarily correspond to better interpretations. This finding underscores the importance of emphasizing interpretation through a feedback mechanism like the CRVE presented here.

Keywords: Visualization, Rubrics, Data literacy, Business intelligence, Tableau

1. INTRODUCTION

Since the early 1980s, the information systems (IS) field has included instruction in decision support systems, leading to analytics and business intelligence topics that have become pervasive (Watson, 2009). These systems include data visualizations to convey key performance indicators (KPIs) to end users, and in response, IS educators have made visualization an integral part of the IS curriculum (Leidig & Salmela, 2020; Nestorov et al., 2019). Easily accessible software packages from Tableau, Microsoft (PowerBI within Office365), The R Foundation (ggplot2), and others allow students to create attractive visualizations with relatively little effort. Data can be sourced and combined from multiple places with a few clicks. AI-powered pipeline capabilities enhance visual data mining and the ability to pose questions in natural language with answers as visualizations (Bouali et al., 2016; Liu et al., 2021).

However, sophisticated technologies and such "marvelously malleable … graphical effects" do not necessarily ensure correct data or better visual designs and interpretations (Kostelnick, 2008, p. 121). Examples of blunders or unintentionally misleading representations of data are easily found in the press, on television, and on social media (Cairo, 2019; MacPherson-Krutsky, 2020). The need to teach various forms of data and visual literacy is widely recognized (Boerner et al., 2019; Fontichiaro & Johnston, 2020; Rodrigues et al., 2021). Visualization mastery is highly desired in industry (Ryan et al., 2019), but student capabilities fall short. In one study, while 94% of students could extract data from a graph, only 53% could extrapolate and analyze relationships implicit in the graph (Wakeling et al., 2015).

Teaching students how to acquire, prepare, analyze, visualize, and communicate data effectively can be time-consuming (Camm et al., 2017). Given time constraints in broader courses that are not exclusively about visualization, and larger class sizes, teaching focus may devolve to tools, techniques, and methods for a smaller set of graphical conventions (Burch & Melby, 2020; Kong, 2020; Ridgway, 2016). While there are excellent studies that conceptualize and test best practices for visualizations, missing is guidance on how to translate these concepts to the evaluation of visualizations in concrete classroom settings (Friedman et al., 2019). Discussion of evaluation is largely absent from recent descriptions of courses with significant visualization components (Jeyaraj, 2019; Nestorov et al., 2019; Stephens & Young, 2020; Zhang et al., 2020).

This paper describes the experience of evaluating visualizations in a graduate business intelligence (BI) course with case-oriented assignments. It examines existing rubrics and explains the desirable characteristics of the compact rubric for visualization evaluation (CRVE) created by the author, which was designed to help students improve both their technical and interpretation skills. It may also help educators assimilate essential concepts from the visualization field into their evaluation practice. Emphasizing the essential elements, the CRVE innovates by freeing time for the evaluator to focus feedback on the specific areas that need the most work. This paper gives examples of the use of the rubric in evaluating assignments in the BI course and then presents quantitative and qualitative data about its effectiveness in the course, including how success in technical elements relates to success in interpretation.

2. LITERATURE REVIEW

2.1 What Constitutes a "Good" Visualization

Numerous researchers have focused on information communication through the pragmatic visual efficiency of visualizations (Kosara, 2007). Quantitative evaluation techniques have included surveys, questionnaires, pre/post tests, and other forms of heuristic (criteria-based) analyses. Qualitative techniques have included observations, interviews, and focus groups (Ltifi et al., 2018). A typical study of this type found that task-irrelevant data points or details increased cognitive load, ratings of task difficulty, processing time, and error rates (Strobel et al., 2018). Kosara (2016) found the surprising result that the rationale behind the "pie charts are bad" dogma does not consider how users actually read them.

Students may be introduced to these concepts using recognized classics by authors such as Cleveland (1993), Few (2012), Tufte (2001), their subsequent works, and works by authors who synthesize many visualization strands into a single text (Munzer, 2014). According to Tufte (2001), "[g]raphical excellence consists of complex ideas communicated with clarity, precision, and efficiency… [giving]… to the viewer the greatest number of ideas in the shortest time with the least ink in the smallest space… [and requiring] … telling the truth about the data" (p. 51).

While this literature informs the standards by which educators evaluate visualizations, it is not necessarily compact or direct enough to be suitable for teaching visualization in nonspecialist settings. Shorter case studies and examples may work better. Padilla's case study (2018) offers accessible explanations of perception and cognition theories, using them to inform improved depictions of forecasted hurricane trajectories. Knaflic (2015) provides detailed illustrations of improvement steps. Wexler et al. (2017) debate the merits of dashboard designs using examples from more than 30 domains. Other useful examples are found in style guides (Elder & Cesal, 2020), incisive explanations from leading practitioners (e.g., *The Economist*'s "Off the Charts" newsletter), and numerous blogs. In using these valuable resources, educators still must devise manageable means for providing feedback to students.

2.2 Heuristic and User-Centered Evaluation Methods

Two emphases are found in current visualization pedagogy, which lead to different approaches to evaluation. The first approach is the proper construction of visualizations, defined as "using the right algorithms and visualization principles in creating visualizations" (Beasley et al., 2020, p. 146). This approach uses heuristics, which comprise a set of guidelines against which to judge a visualization (Brath & Banissi, 2016). Since the presence or absence of numerous features, such as labeling an axis, can easily be perceived, objective measurement of success is possible at this level. The rubric of Friedman and Rosen (2017) (Appendix A) mainly uses a heuristic approach, including 33 technical elements to evaluate, some of which resemble specification grading (Howitz et al., 2021) by requiring either a "Yes" or a "No." Perhaps as a recognition of how laborious this can be, Friedman and Rosen (2017), Friedman et al. (2019), and Beasley et al. (2020) assign students to fill out their rubric (Appendix A) as part of peer reviews. Educators maybe unable to afford the time to notate every possible thing wrong with every visualization.

Rubrics focus student attention on what is important, how and why the students went astray, the rationale for evaluation, and how to make improvements. Rubrics also organize and limit the number of evaluation points and metrics for the educator, making the process more efficient (Stevens & Levi, 2013). However, even rubrics with expansive lists of heuristic criteria may leave out other elements that need attention. Finding lists appropriate for all visualizations has proven elusive (Santos et al., 2018). They may break down when encountering a new kind of error or when stripping away one level of error reveals another. Broader rubrics permit some "wiggle room" as the evaluator confronts significant elements that are not covered in the heuristics.

The second approach focuses on the subjective evaluation of quality and accuracy (Beasley et al., 2020), which are harder to map to a checklist. This type of evaluation requires knowledge of users' needs in a specific domain and an understanding of how the analyses enabled by the data, data transformation, and visual encoding choices lead to an acceptable solution. The educator looks at the student's work not just as the successful application of elements such as color or shape but holistically—whether it serves the user's given purpose well enough. The "user" is the ultimate viewer or reader of the visualization, which could be someone inside or outside an organization. When the educator adopts this user-centered approach, he or she plays the role of the user, taking into consideration the potential or actual reactions users may have to a given scenario-based student design (Brath & Banissi, 2016). From the user's point of view, what matters is what the visualization conveys, what analyses it facilitates, and what interpretations it justifies. Beasley et al. (2020) point out that subjective evaluation helps students build their ability to evaluate their own and others' visualizations critically. They note that broad group discussions are not enough, implying the need for specific feedback.

Nolan and Perrett's (2016) rubric (Appendix B) takes a user-centered approach. It emphasizes analysis, synthesis, and communication of the results clearly, precisely, and concisely. "We have found that if the student is provided with detailed comments on his/her work and the completed matrix of competencies, then this evaluation creates a platform for discussion between the educator and student that is more process and content driven than point driven" (p. 267). Friedman and Rosen's (2017) rubric (Appendix A) also has a section called Narrative Consideration, but most of the assignments in their classes that they describe focused on narrower technical aspects of programming and visualization construction. The Narrative Consideration category is, however, "critical in projects where the story surrounding the visualization is as important as the visualization itself" (Friedman & Rosen, 2017, p. 3). Narratives facilitate the acquisition of data analysis and problem-solving skills (Saundage et al., 2016).[1]

Thus, educators often the play roles of both heuristic evaluators and end users, providing written (and sometimes oral) comments on potential viewers' reactions to a given visualization. Effective feedback should reduce the gap between current understanding and goals, provide high information content, help students see where to go next, and match the timing of the ongoing task cycle (Wisniewski et al., 2020). Comments should be neutral and balanced, based on facts, and not advocate pet preferences of the educator (Kosara, 2007). Rubrics provide the framework for this feedback.

3. COMPACT EVALUATION RUBRIC AND BI COURSE

3.1 The Rubric

The CRVE (Table 1) was developed in the fall 2016 after a surge in the number of graduate students and an expansion of the number of assignments left the author struggling to provide detailed, consistent, and timely feedback. Grounded in the approaches outlined in Section 2, the CRVE endeavors to fold heuristic evaluations into just a few rubric items and levels, while providing the necessary broader categories for rich, meaningful user-centered feedback. Any of the technical elements covered by Friedman and Rosen (2017) can be noted in feedback, while the rubric itself places greater emphasis on the more results-oriented elements in Nolan and Perrett (2016).

In the Technical Correctness of Visualization (TechCorrect) category, a 1-Fatally Flawed visualization portrays incorrect data. Data may be wrong at the source, or errors may be introduced by manipulations or calculations made by the student. In such cases, users cannot draw the correct conclusion. To avoid fatal flaws, students must sufficiently understand each data field's units, its formulation through the logic of calculations and blending, and its meaning. This is similar to Nolan and Perrett's (2016) Computation metric (Appendix B). A 2-Grossly Misleading visualization uses correct data, but its portrayal and/or lack of necessary context result(s) in viewers drawing incorrect conclusions. With a 3-Stylistically Challenged visualization, a user is able to draw the right conclusions, but stylistic or aesthetic elements make it harder to do so than it should be. In order to reach the 4-Delivers Intended Message level, the visualization should realize Tufte's (2001) best practice principles, conveying meaningful information that sheds light on the given task. Hence, in one rubric line, the TechCorrect element takes in Friedman and Rosen's (2017) Visual Design and Design Consideration lines (at least 14 elements) (Appendix A) and Nolan and Perrett's (2016) Computation, Visual Presentation, and Analysis lines (Appendix B).

The Visualization Interpretation (VizInterpret) category is about how students explain the meaning of their visualizations. 1-Egregious Misinterpretation represents wildly wrong conclusions and may lead the viewer far from the truth.. A 2-Misinterpretation contains demonstrably wrong conclusions but without the potential damage of egregious misinterpretations. Analyses may be correct, but 3-Somewhat Lacking in other regards. For example, they may be too brief or ambiguous to be certain about how correct they are, they may be questionable because they do not take into account other relevant facts, or they may cite (unverifiable) data not available to the viewer. A 4-OK interpretation draws correct conclusions. The VizInterpret metric is most closely related to Nolan and Perrett's (2016) Synthesis item (Appendix B).

The Strength of Insights (StrengthInsights) category captures visualization ambitiousness. Level 1 visualizations, such as a bar chart showing yearly sales, simply declare facts. Visualizations that are more ambitious may combine several variables in an exploratory fashion (Berinato, 2016). This category also considers how students used their insights and analyses to make at least one reasonable and concrete recommendation for what the organization or institution should do based on the visualization. Level 2 is for simple declarations with strong insights or strong exploration without strong insights. Level 3 leaves some room for instances where most of this is done but there are some errors or things missing. Level 4 combines both strong exploration and recommendations.

The Tell the Story (TellStory) category captures how students narrate key ideas related to the visualization (Knaflic, 2015). Level 1 is for visualizations without any informative headlines, captions, or annotations at all. Level 2 text defines data points and/or the X and Y axes. Level 3 again leaves room for instances where some things are not quite right. Level 4 text interprets the visualization in a non-misleading manner, guiding viewers to valid, justified insights. Level 4 visualizations should be able to stand alone if circulated without additional context, e.g., on one PowerPoint slide.

Including StrengthInsights and TellStory incentivizes students to go beyond just making a good-looking visualization. Of course, single decisions made by students may influence several rubric lines. For example, dramatic use of color may help to tell the story very well, but may also lead to biased interpretations due to "anchoring" (Cho et al., 2017). While the

Category	Lowest to Highest Rubric Levels			
	1	2	3	4
Technical Correctness of Visualization (TechCorrect)	Fatally Flawed	Grossly Misleading	Stylistically Challenged	Delivers Intended Message
Visualization Interpretation (VizInterpret)	Egregious Misinterpretation	Misinterpretation	Somewhat Lacking	OK
Strength of Insights (StrengthInsights)	Declarative; Few/No Recommendations	Declarative With Strong Recommendations or Exploratory Without Strong Recommendations	Uneven Mix or One Thing Wrong	Exploratory With Strong Recommendations
Tell the Story (TellStory)	No Informative Titles/Headlines	Chart Titles Convey Contents (e.g. X and Y Axes)	Titles Tell Story to Certain Extent	Chart Titles Tell Story Sufficient for Circulation

Table 1. Rubric Developed by Author for BI Course

assignment of CRVE levels may seem subjective, selecting the appropriate level may be more straightforward than in rubrics with many more evaluation points. For the Data/Ink Ratio design consideration from the rubric of Friedman and Rosen (2017), for example, what are the cutoff points between Way Too Little/Much Ink, Slightly Too Little/Much Ink, and Perfect Amount of Ink?

3.2 The BI Course

Master's students seeking either an MS or an MBA degree took the BI course. This course used visualizations to provide a bottom-up view of constructing and analyzing KPIs for business decisions. In parallel, it provided a top-down view of how BI-provided KPIs facilitate the achievement of strategic and other business goals. Knowledge gained from visualizations and KPIs was contextualized within the broader goals of team and enterprise knowledge management. In 2016–2018 the class was a conventional 16-week, 3-credit course, after which it was compressed into an 8-week, 3-credit course—making the need for compact evaluation more imperative. The author assigned balanced four-person teams.

Examples are provided in the first class session to define and illustrate the various CRVE levels. Students are polled on how they see each visualization before discussion, so they begin to understand gaps between their perceptions, assumptions, and the visualizations' true natures. They are introduced to questions they should ask themselves that dovetail with the CRVE. For TechCorrect, for example, they should ask: 1. Is the data being shown correct? The rationale for saying that incorrect data is a "Fatal Flaw" is explained. A very good-looking visualization that conveys incorrect data is worse than no visualization at all. 2. Is the manner of presentation leading to the wrong conclusions, or making it impossible to draw correct conclusions? "Grossly Misleading" is defined, and students are told that this wording is deliberately intended to grab their attention. 3. Can parts of the visualization be improved to increase understanding? If these criteria have been satisfied, then Level 4 should accrue: all overt ways of improvement have been utilized to reach a visualization that directly communicates the intended message. Similarly, the use of the word "egregious" (outstandingly bad) for interpretations is explained using famous examples such as the incomplete O-ring chart that contributed to the Challenger Space Shuttle disaster (Fry, 2021; Vaughan, 2016).

It is not assumed that students come into the class with any prior courses in databases, data manipulation, etc. Join types are demonstrated using a small amount of data and the Excel VLOOKUP command. Teaching about joins is especially important since Tableau has provided (in Version V2020.2 forward) the "noodle" to define "relationships" between tables where the default is a many-to-many relationship. Granularity, aggregation, and "slicing and dicing" are explained. Students must define the granularity of the data they are using. They receive detailed instructions about how to combine the data needed for the assignments.

Figure 1 provides a capsule overview of the case study–style assignments and their sequencing in "rounds" in various semesters so that students can apply feedback from one round in the next. Chicago crime data (~1.5 million rows) from the City of Chicago Data Portal (https://data.cityofchicago.org/) always provided the basis for Rounds 1 and 2. Round 2 added a rich means to relate crime to employment, education, racial composition, etc., with data from Community Data Snapshots (https://datahub.cmap.illinois.gov/dataset/community-data-snapshots-raw-data). In 2019, Round 3 comprised updates and additions to results from Rounds 1 and 2. Rounds 1 and 2 rotated two of the four team members, Round 3 used the whole team, and Round 4 was individual.

The Round 3 lighting distributor case (2016–2018) was based on data from a real firm provided by a pricing analytics consultancy (a consultancy representative gave direct feedback on presentations in live sessions). The Round 4 tire manufacturer case (2016–2018) used (artificially enhanced) quantitative and qualitative tire-performance consumer reaction data from reviews on TireRack.com. The author augmented the Round 4 Mozilla (2020) case with data from published sources for internet penetration in Chicago community areas. For context, students received short articles about Chicago crime, a case study about pricing analytics, and larger cases about the tire manufacturer (written by the author) and Mozilla (Watson et al., 2017).

Rather than being given dirty data to clean, which can require enormous amounts of student time (Battle & Heer, 2019), or being allowed to choose their own data sets for the entire end-to-end process (Ryan et al., 2019), which can eliminate economies of scale for grading, students received mostly cleaned data, enhanced with higher-level dimensions. They were given guidelines with which to explore the data, rather than fixed visualizations to prepare. In asking students to use higher-level skills such as analyzing, combining, and creating (i.e., Bloom's taxonomy [Burns et al., 2020]), educators must be able to recognize where the mistakes happened, and how mistakes interacted to produce sometimes unrecognizable results. Following student paths, finding each error, and writing about it in a way that is useful to students can be time-consuming. As with Nolan & Perrett (2016), an essential element that makes the CRVE both a heuristic and a user-centered conversation is the extensive TechCorrect and VizInterpret written feedback.

3.2.1 Example from the Lighting Distributor Case. A lighting distributor project TechCorrect 1-Fatal Flaw occurred when a team tried to create a better KPI than the provided "Margin." The margin was defined for them as the difference between the revenue collected on a sale and the cost of the goods sold, i.e., Margin = Sales – (Unit Cost * Quantity). This team defined "New Profit" as (Sales – Unit Cost), without multiplying Unit Cost by Quantity. The error was only discovered by checking in Tableau after the presentation. When the students advised investing more heavily in certain products that now seemed much more profitable than they really were, this was a VizInterpret 1-Egregious Misinterpretation.

3.2.2 Examples from Chicago Crime Assignment. Figure 2 purported to show the crime per capita (CPC) rate on various holidays in Chicago. It used Chicago crime data from Jan. 1, 2012 until roughly mid-2016. It is impossible to tell by the observation that the students summed the number of crimes on each of these days over a period of 4.5 years, then divided by the total population for just one year (2010). This considerably overstated crime rates on each holiday. The Y axis

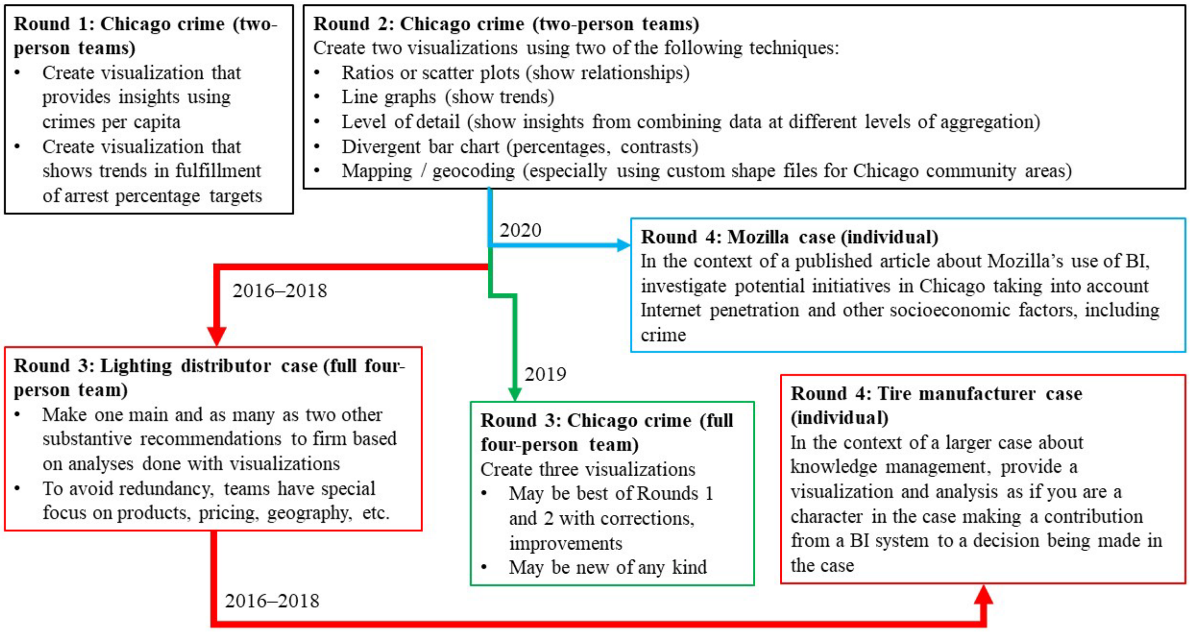

Figure 1. Assignment Flow in the BI Course in Various Versions

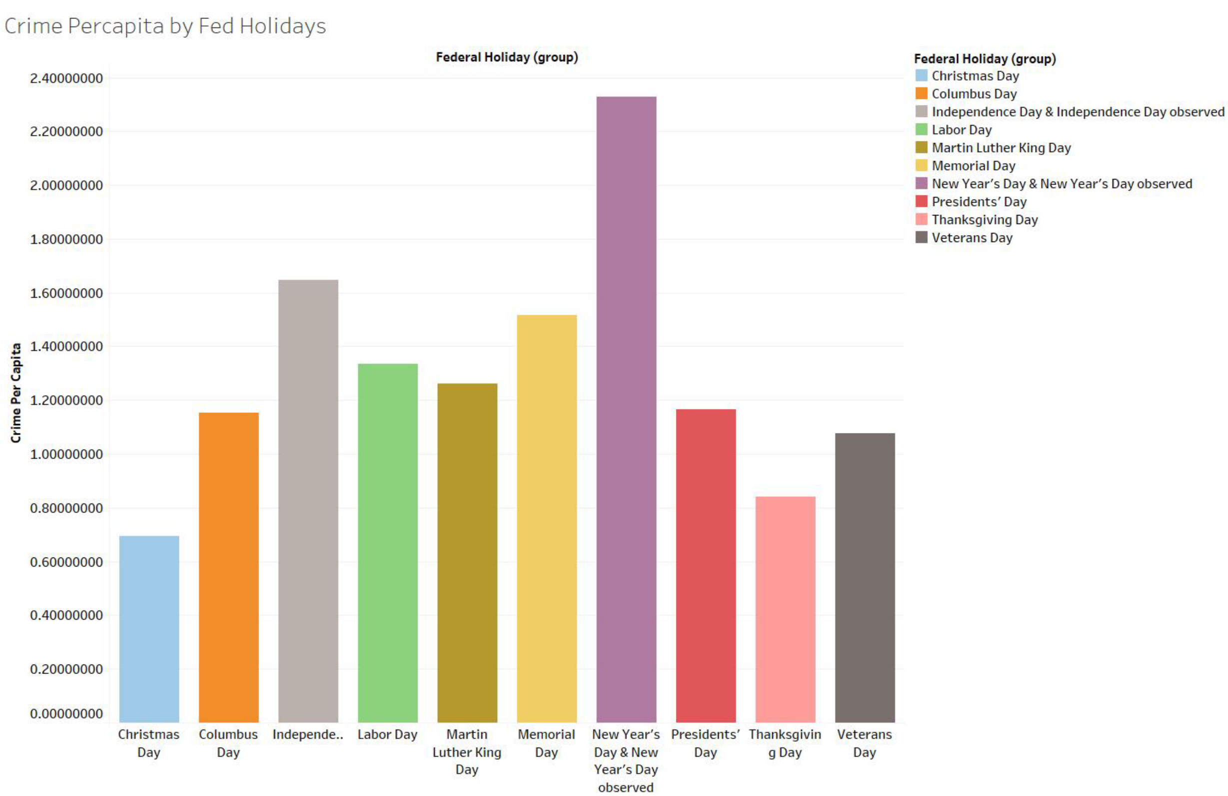

Figure 2. Students' Crime per Capita on Federal Holidays Visualization

should have read "Crimes per 1,000" since CPC was multiplied by 1,000. These are TechCorrect 1-Fatal Flaws.

A VizInterpret 1-Egregious Misinterpretation of Figure 2 would be to insist boldly that there were 2.4 crimes per every person in Chicago on New Year's Day! These students made a 3-Somewhat Questionable interpretation that focused only on the relative values of CPC. They suggested that the New Year's holiday was the worst, without noting that, in this visualization, New Year's covers two days. If the team had made the caveat that the underlying data does not identify "Observed Days" for other holidays, they might have drawn a correct conclusion about the relative magnitude of federal holiday crimes.

If the underlying data in Figure 2 was correct, and each bar was limited to one day, then Figure 2 could serve as an example of a visualization that is TechCorrect 3-Stylistically Challenged with excessive decimal places, the smallness of the text, superfluous use of color, etc. In using the CRVE, once it has been determined that the visualization is fatally flawed, the educator does not necessarily have to evaluate other elements. The stylistic aspects can be discussed in a general review of examples with the whole class.

Maintaining a "master workbook" in which quick calculations can be done is useful for verifying what is in the student's workbook. Demonstrating in class how to rework problematic visualizations helps students see the process of "getting from here to there." (This also can inform feedback to the students about why a visualization may be grossly misleading or stylistically challenged.) Students appreciate access to this reworked Tableau workbook to study before the next round. In the author's reworked version, Figure 3 reaches the level of TechCorrect 4-Delivers Intended Message. Each bar is "per day," which gives a strong impression of how holidays compare to non-holiday days. A VizInterpret 3-Questionable Interpretation of Figure 3 might be that holiday crimes were growing from 2012 to 2015. The viewer would have no opportunity to evaluate the alleged trend for him- or herself. For both Figures 2 and 3, StrengthInsights is exploratory, with the ultimate level depending on the recommendations made. With just a description of axes, Figure 2 is at Level 2 for TellStory. Figure 3 is ready to circulate at Level 4.

Figure 4 provides an attractive example of another visualization that used the crime per capita KPI. It appeared to be a visualization at TechCorrect level 4-Delivers the Intended Message. However, the headline said that the top three (worst) communities were annotated. In reality, for the four-year averages shown, both West Garfield Park and the Loop had higher crime rates. Discerning this from the size and color of the circles would be difficult and unreliable. A viewer could easily overlook the Loop (third highest), Chicago's business center.

Consequently, it was evaluated as 2-Grossly Misleading. The students' interpretation included a potentially strong insight that areas near major highways and higher traffic with lower income per capita could experience higher rates of crime per capita, receiving a 4-OK on VizInterpret. The visualization had Level 4 headlines and Level 4 insights and recommendations. Overall, the students received a message that

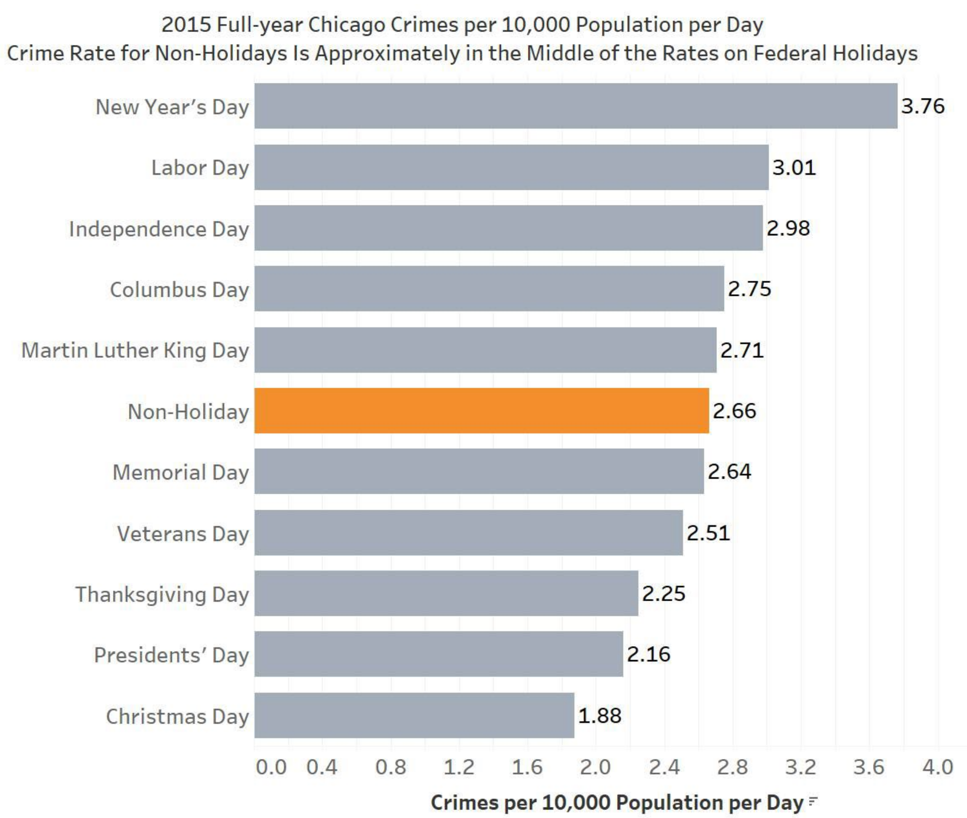

Figure 3. Corrected Visualization Shows Holiday Crimes in Proper Context

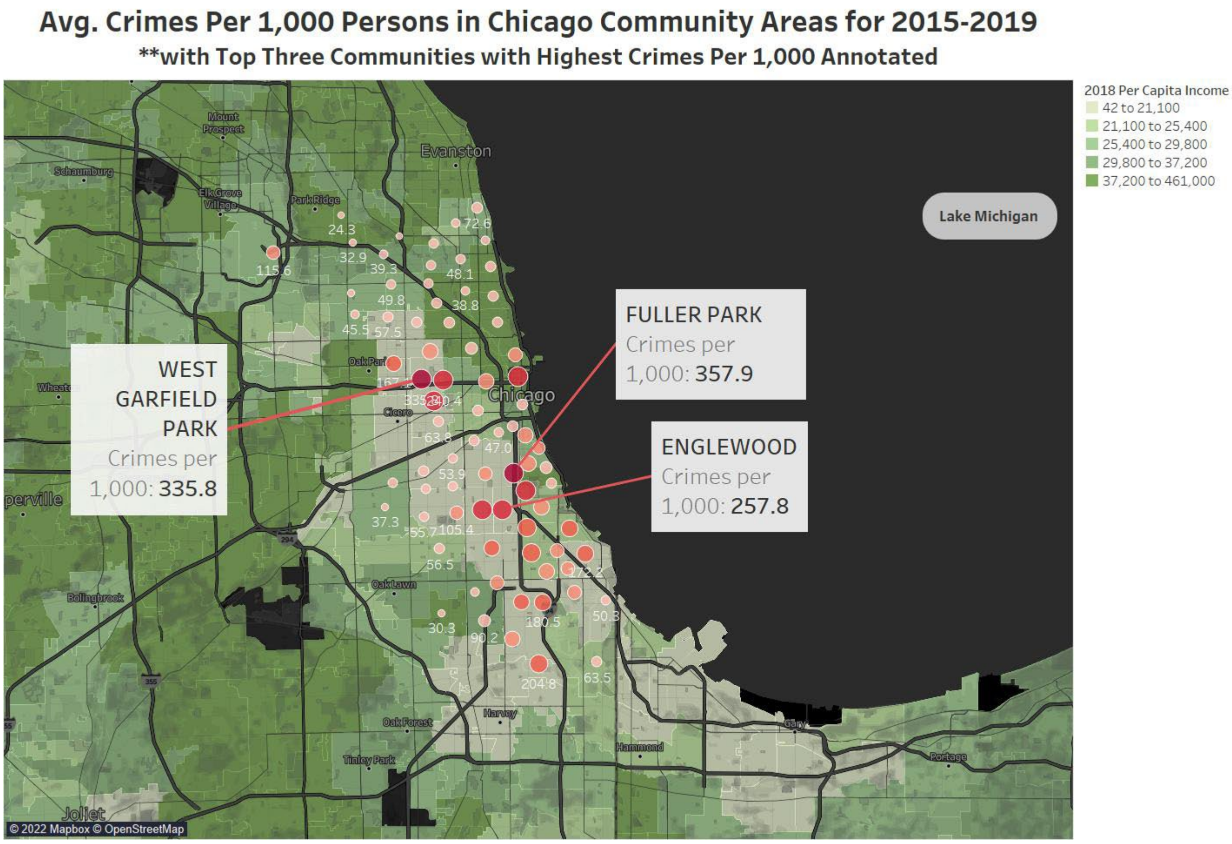

Figure 4. Attractive Visualization That Still Is Grossly Misleading

even a visualization that looks this good and has almost everything right may not be correct. Overall, the work still received an A-minus grade.

3.3 Framing Visualizations with the CRVE

Over four years of using the CRVE, the questions it poses were increasingly used to frame the discussions of visualizations. In a recent semester during Round 2 of the Chicago Crime assignment, students produced two visualizations with trend lines, seven with scatter plots and fitted lines, four with data on top of maps, and two bar or stacked bar charts. In the class discussing this assignment, students examined batches of similar visualizations and indicated (using the LMS survey mechanism) what level of TechCorrect they thought each visualization merited. Most visualizations received at least one or two ratings of 1-Fatal Flaw, and 4-Delivers Intended Message, with more ratings divided between the other two levels. Students were amazed to see how differently others approached the same data and evaluated the same visualizations.

The final part of the class closes with a powerful message about do's and do not's (Table 2) that students should understand in a new way. Points 1, 2, and 3 relate directly to TechCorrect. Points 2, 4, and 5 relate to VizInterpret. Point 4 relates to StrengthInsights, and Points 5 and 6 relate to TellStory. Table 2 applies to many visualization contexts.

	DO	DO NOT
1	Double check all parts of the calculations	Assume you or Tableau are/is always right!
2	Understand the metrics and/or KPIs you are using	Assume your audience will understand
3	Put in the work needed to make your visualization shine	Make it shine before you are sure it is correct
4	Carry out enough additional exploration to be confident in your conclusions	Go with the first thing you find
5	State your visualization-based conclusions cautiously and appropriately	Draw conclusions that go beyond what your visualization justifies
6	Provide enough context; visualizations will circulate	Forget to add headlines that tell the story

Table 2. Takeaways Based on the CRVE

3.4 Using the CRVE for Grading

By using the CRVE, the process of assigning grades is separated from the process of assigning rubric scores. Rubric scores reflect the reality of what the students did, while grades can also acknowledge effort, adherence to assignment guidelines, etc. To derive the grade, each rubric score is multiplied by a weight that reflects that element's importance. TechCorrect and VizInterpret are weighted more heavily (3x the rubric score) than StrengthInsights (2x) or TellStory (1x). This also allows weighting in Rounds 2–4 so that somewhat flawed complicated charts may receive similar credit to flawless easier charts, while flawless complicated charts always yield the highest scores. The sum of weighted rubric points is mapped to a linear scale reflecting the lowest to highest totals and perceptions of "how bad" the worst efforts are (usually a "C"). Quibbling about points for grades happens very rarely. (Nolan and Perrett [2016] had a similar experience.)

The author assembles the feedback in Excel, making it possible to reuse parts of comments for similar examples and to check for consistency. Finding one sort of error on the fifth visualization may entail rechecking and adjusting evaluations of some of the others. In the author's opinion, doing this in Excel is much easier than doing it by moving from student to student in an LMS such as D2L and manually inserting and adjusting rubric scores and comments. By distributing the results to the students using Word and mail merge in an email, it is possible to provide additional commentary about the assignment and personalize messages to individual teams/students as needed. Since assignments are collected in PowerPoint (with the essential original Tableau workbook), it would also be possible to write comments in each PowerPoint file and then transfer them back to the LMS Assignment Box (D2L can automate this to a certain extent).

4. EVIDENCE AND DISCUSSION OF RESULTS

Data from four years of the course was extracted from existing spreadsheets for 159 students, of whom roughly 58% were male, 53% were in an MS Management/Information Systems Management program, and 62% were international, many with engineering or information systems backgrounds. In some semesters, a TechCorrect bottom level of "1" was added for visualizations with multiple Fatal Flaws. In some cases, 5-point scales, including half points for VizInterpret were used when that assignment included live presentations. All of these ratings were rationalized to the four-point scales in Table 1 for analysis here.

4.1 Quantitative Results

Tables 3 to 6 show how using the CRVE plus feedback led to improved student performance over the semester. In Tables 3 and 4 N is the number of visualizations and the Percentage is the percentage of visualizations in that round for that CRVE element.

4.1.1 TechCorrect. After the feedback from Round 1, the percentage of visualizations in Round 2 with at least one TechCorrect 1-Fatal Flaw went down from 32.2% to 25%, ultimately settling at just 12.1%. Similarly, the percentage for 2-Grossly Misleading went from 44% to 20%. These are substantial improvements, suggesting that the CRVE plus feedback helped a large majority of students eliminate the worst technical problems. Visualizations with 4-Delivers Intended Message grew from 4.6% to 54.3% for Round 3 and 41.4% for Round 4. Since the skills were cumulative during the semester, it does not appear that regression to the mean was a factor.

Level	Data	R1	R2	R3	R4
4-Delivers Intended Message	%	4.6	9.9	54.3	41.4
	N	7	15	100	58
3-Stylistically Challenged	%	19.1	23.7	9.2	26.4
	N	29	36	17	37
2-Grossly Misleading	%	44.1	41.4	21.2	20.0
	N	67	63	39	28
1-Fatal Flaw(s)	%	32.2	25.0	15.2	12.1
	N	49	38	28	17

Table 3. TechCorrect Results

4.1.2 VizInterpret. In general, this item did not show as much improvement as expected. The percentage of interpretations that fell into the 1-Egregious Misinterpretation category did fall from 17.1% in Round 1 to none in Round 3 and 5.7% in Round 4. There was a corresponding rise in 2-Misinterpretation scores for Round 3. However, there was not a clear pattern of improvement for either the 4-OK or the 3-Somewhat Questionable categories. Scores in the top category 4-OK varied from a low of 29.6% in Round 2 to a high of 43.4% in Round 1, while scores for 3-Somewhat Questionable varied from 30.3% to 42.8%. These results may be related to the differences in the assignments in the various rounds. In Round 1 there was a direct statement of the types of analyses required, but in Round 2 students had much more freedom to choose both technique and analysis. It may be that lessons learned in the analysis parts of each assignment (especially since contexts varied) were not as easily transferred to the next assignment as technical lessons were.

Level	Data	R1	R2	R3	R4
4-OK	%	43.4	29.6	33.9	37.1
	N	66	45	60	52
3-Somewhat Questionable	%	30.3	42.8	37.3	39.3
	N	46	65	66	55
2-Misinterpretation	%	9.2	13.8	28.8	17.9
	N	14	21	51	25
1-Egregious Misinterpretation	%	17.1	13.8	0.0	5.7
	N	26	21	0	8

Table 4. VizInterpret Results

4.1.3 StrengthInsights and TellStory. The StrengthInsights and TellStory rubric items were only used consistently for Rounds 1 and 2 (the Chicago Crime assignment). The instructor evaluated pairs of students in those rounds, so the N and percentages shown in Tables 5 and 6 are for pairs of visualizations for each CRVE element. These results cannot be correlated directly with the other evaluations. Nevertheless, Tables 5 and 6 show that once students realized what was missing, many teams made corrections. For StrengthInsights, the percentage of pairs in Level 1 dropped from 45.2% to 9.2%, while the percentage of pairs in Level 4 rose from 17.8% to 42.1%.

For TellStory, the percentage of pairs in Level 1 went from 68.4% in Round 1 to 19.7% in Round 2. About three-fourths of the pairs reached either Level 3 or Level 4 in Round 2.

Level	Data	R1	R2
4-Exploratory With Strong Recommendations	%	17.8	42.1
	N	13	32
3-Uneven Mix or One Thing Wrong	%	17.8	13.2
	N	13	10
2-Declarative With Strong Recommendations or Exploratory Without Strong Recommendations	%	19.2	35.5
	N	14	27
1-Declarative; Few/No Recommendations	%	45.2	9.2
	N	33	7

Table 5. StrengthInsights Results

Level	Data	R1	R2
4-Chart Titles Tell Story Sufficient for Circulation	%	7.9	43.4
	N	6	33
3-Titles Tell Story to Certain Extent	%	23.7	32.9
	N	18	25
2-Chart Titles Convey Contents (e.g. X and Y Axes)	%	0.0	3.9
	N	0	3
1-No Informative Titles/Headlines	%	68.4	19.7
	N	52	15

Table 6. TellStory Results

4.1.4 TechCorrect vs. VizInterpret. One would think that as students execute the technical aspects of visualizations better, they would also interpret them better. For teams in Rounds 1–3, there was a negligible non-parametric correlation between TechCorrect and VizInterpret (Spearman's $\rho=0.1831$, $p<.0001$). Table 7 shows the co-occurrence of pairs of TechCorrect and VizInterpret levels for teams (N=481). While 11.2% of the 481 visualizations 4-Delivered the Intended Message and had 4-OK interpretations, the largest share, 13.7%, was for 2-Grossly Misleading visualizations leading to 3-Somewhat Questionable interpretations. However, in 9.6% of the cases, 2-Grossly Misleading visualizations still led to 4-OK interpretations. 1-Fatally Flawed visualizations still led to just 3-Somewhat Questionable interpretations in 8.5% of the cases and to 4-OK interpretations in 7.3% of the cases.

How to explain these results? When TechCorrect was 1-Fatally Flawed or 2-Grossly Misleading, the evaluation of VizInterpret depended in part of the extent to which the students exacerbated the technical errors by parroting and trumpeting them in the analysis. In other cases, the 1-Fatal Flaw did produce incorrect data, but data in a pattern that was similar to the correct data and could produce some analysis that was correct. While one would think that having more eyes in teams to review visualizations would lead to higher technical quality and better analyses, this may not have happened in practice. Student habits of dividing the work up may have led to situations in which cross-student reviews were inadequate (or even nonexistent). One of the goals of the CRVE is to give students concrete language and tools for evaluating their own and each other's work.

VizInterpret Level	Data	TechCorrect Level			
		1	2	3	4
4	%	7.3	9.6	7.5	11.2
	N	35	46	36	54
3	%	8.5	13.7	6.7	7.9
	N	41	66	32	38
2	%	4.2	6.9	1.9	5.0
	N	20	33	9	24
1	%	4.0	4.8	0.8	0.2
	N	19	23	4	1

Table 7. Relative Frequency of TechCorrect/ VizInterpret Score Pairs for Teams in Rounds 1–3

For individual work (Table 8, N=140), the TechCorrect and VizInterpret relationship was moderately strong (Spearman's $\rho=0.5583$, $p<.0001$). For individuals (right side), Table 5 shows a higher concentration of pairs in the upper right-hand corner—i.e., visualizations of higher technical correctness corresponded to improved interpretations. Here the most common pair is 4-Delivers the Intended Message and 4-OK (23.6%), with 4-Delivers the Intended Message leading to 3-Questionable Interpretation in 16.4% of pairs, and 3-Stylistically Challenged leading to 4-OK in 11.4% of pairs. These results are much closer to what one would expect.

VizInterpret Level	Data	TechCorrect Level			
		1	2	3	4
4	%	0.0	2.1	11.4	23.6
	N	0	3	16	33
3	%	2.9	10.0	10.0	16.4
	N	4	14	14	23
2	%	6.4	6.4	3.6	1.4
	N	9	9	5	2
1	%	2.9	1.4	1.4	0.0
	N	4	2	2	0

Table 8. Relative Frequency of TechCorrect/ VizInterpret Score Pairs for Individuals in Round 4

4.1.5 Feedback Given. The average total number of words in the feedback given with the rubric was about 265. The average number of total TechCorrect and VizInterpret words for the Chicago Crime assignments ranged from 200 to 300. The Lighting assignment, with more visualizations, engendered more feedback (about 400 words on average).

4.2 Qualitative Feedback from Students

Typically, students have expressed satisfaction with the BI course and approach. They have found the visualization skills make them more marketable. Based on email and comments in teaching evaluations, some students really appreciated the

CRVE structure and accompanying feedback. They wrote: "This type of feedback is really beneficial...the most honest and detailed feedback I have ever received on any project," and "I look more critically at my own work to verify that I really know what I think I know, scrutinize visuals more carefully, and am leading an effort to be more thoughtful in how we manage and use information across [my work]." Another wrote: "I have learnt a lot about data analysis in your classes, how to [see] underlying faults in the data model we build etc. thanks to the mistakes I made in your classes. The fundamentals you taught me are the most important thing that is allowing me to stand on my feet right now."

Some students with little background could find the Tableau learning curve to be daunting or could be baffled by subject matter such as pricing analytics. What the author thought to be honest, direct feedback was occasionally perceived as negative and demotivating. For example, "Grading was unnecessarily detailed to the point that feedback was not useful. Criticism of work did not allow me to effectively apply it to future assignments because it [was] mostly negatively framed, specific feedback focused on doing something wrong with no detail on how to make something better (this was done in class, but only generally)." Some of these comments were addressed by providing short videos for highly relevant Tableau data modeling and visualization tips, and double-checking feedback before sending it to make sure it was within the scope of effective comments outlined at the end of Section 2.

5. CONCLUSIONS

Educators are teaching critical visualization concepts under daunting constraints, such as large class sizes, shortened and online formats, and courses that do not focus exclusively on visualizations. The published visualization literature has supplied an accepted wisdom about clarity, precision, efficiency, and truth-telling. This paper fills a gap by providing an innovative way for educators to translate this wisdom and provide meaningful feedback to students using a compact rubric that combines heuristic and user-centered approaches. It focuses on the use of higher-level skills and telling a story that is plausible and defendable within a case-study context.

When students receive data they may manipulate, they will make mistakes. Today's tools make it even easier for students to hide fatal flaws behind the veneer of spectacular-looking visualizations. Educators must "look under the hood"; the compact nature of the CRVE allows the educator to spend time on careful review and feedback formulation. The author's finding that it was easier to foster improvements in technical aspects than in interpretation only reinforces the need for educators to focus on higher-level interpretation skills.

One of the main things the author learned during the four years of using the CRVE was that students would not automatically understand the rubric or the feedback given with it. The rubric is now introduced on the first day of class. During the semester the students use it themselves to evaluate each other's work. The course ends with powerful takeaways related to the rubric. One student recently wrote, "I won't look at visualizations the same way again." This sort of mindset change is exactly what the author has hoped to achieve. While the questions posed by the CRVE and the takeaways shown in Table 2 are applicable to most visualization teaching situations, they may be adapted to add an item or level that takes into account novelty or innovation.

This paper is limited in two senses. It does not constitute a research study comparing the efficacy of the rubrics in the Appendices with the author's rubric. Rather, in pointing out similarities and differences, the paper argues that the CRVE plus feedback provides an effective means of helping students focus on the most important teachable moments emerging from the assignments. Second, while the use of the CRVE saves time by eliminating repetitive checking of numerous heuristic rubric points, the purpose of this paper was not to show that it saves time overall (although the author believes this to be the case based on experiences of grading without using the rubric).

Further research into the relationship between the technical correctness of visualizations and their interpretation may yield fruitful results, especially in the context of teams. How technically flawed does a visualization have to be before its interpretation becomes an egregious misrepresentation? In addition, beyond Kong's (2020) extended abstract, we do not know much about how most educators evaluate visualizations. Do they take for granted that data was correctly manipulated? Do they require student peer evaluators to "look under the hood" on the visualizations they evaluate? Do they consider how students interpret their own visualizations?

The next evolution of user-oriented visualization pedagogy could include video creation and editing tools that allow educators to show and explain to students the errors that are being discovered as the rubric items are being evaluated. Initial experiments with this idea found that the creation of each video eliminates the time advantages of assigning CRVE levels and formulating feedback in Excel. Visual explanations will only be practical if it becomes as easy to "cut and paste" relevant explanatory video snippets during feedback construction as it is to copy and adapt text from one cell to another in a spreadsheet.

6. ACKNOWLEDGMENTS

I offer my sincere thanks to Diane Grabowski for excellent help in preparing this manuscript for publication. All errors remain my own.

7. ENDNOTES

[1]Jeffrey Shaffer provided a grading rubric with the Data Visualization course he provided through the Tableau Teaching Community (https://community.tableau.com/s/teachers) for Tableau users. As it is not been formally published, it is not discussed here. It contains one element about the data, four about the mechanics and aesthetics of the visualization, and two about usability and impact. It does not explicitly consider the analysis done by the students using their visualizations.

8. REFERENCES

Battle, L., & Heer, J. (2019). Characterizing Exploratory Visual Analysis: A Literature Review and Evaluation of Analytic Provenance in Tableau. *Computer Graphics Forum*, 38(3), 145-159.

Beasley, Z., Friedman, A., Pieg, L., & Rosen, P. (2020). Leveraging Peer Feedback to Improve Visualization Education. *2020 IEEE Pacific Visualization Symposium (PacificVis)* (pp. 146-155). IEEE.

Berinato, S. (2016). Visualizations That Really Work: Know What Message You're Trying to Communicate Before You Get Down in the Weeds. *Harvard Business Review*, 94(6), 92-100.

Boerner, K., Bueckle, A., & Ginda, M. (2019). Data Visualization Literacy: Definitions, Conceptual Frameworks, Exercises, and Assessments. *Proceedings of the National Academy of Sciences of the United States of America*, 116(6), 1857-1864.

Bouali, F., Guettala, A., & Venturini, G. (2016). VizAssist: An Interactive User Assistant for Visual Data Mining. *The Visual Computer: International Journal of Computer Graphics*, 32(11), 1447.

Brath, R. & Banissi, E. (2016). Evaluation of Visualization by Critiques. *BELIV '16: Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (pp. 1-8). New York, NY, USA: ACM.

Burch, M., & Melby, E. (2020). What More Than a Hundred Project Groups Reveal About Teaching Visualization. *Journal of Visualization*, 23(5), 895-911.

Burns, A., Xiong, C., Franconeri, S., Cairo, A., & Mahyar, N. (2020). How to Evaluate Data Visualizations Across Different Levels of Understanding. *2020 IEEE Workshop on Evaluation and Beyond - Methodological Approaches to Visualization (BELIV)* (pp. 19-28). IEEE.

Cairo, A. (2019). *How Charts Lie: Getting Smarter About Visual Information*. New York, NY, USA: W. W. Norton & Company, Inc.

Camm, J., Fry, M. & Shaffer, J. (2017). A Practitioner's Guide to Best Practices in Data Visualization. *INFORMS Journal on Applied Analytics*, 47(6), 473-488.

Cho, I., Wesslen, R., Karduni, A., Santhanam, S., Shaikh, S., & Dou, W. (2017). The Anchoring Effect in Decision-Making With Visual Analytics. *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)* (pp. 116-126). IEEE.

Cleveland, W. S. (1993). *Visualizing Data*. Summit, NJ, USA: Hobart Press.

Elder, K. L., & Cesal, A. (2020). Should We Teach Data Visualization Using Data Visualization Style Guides? *Issues in Information Systems*, 21(4), 264-274.

Few, S. (2012). *Show Me the Numbers: Designing Tables and Graphs to Enlighten* (2nd ed.). Burlingame, CA, USA: Analytics Press.

Fontichiaro, K., & Johnston, M. P. (2020). Rapid Shifts in Educators' Perceptions of Data Literacy Priorities. *Journal of Media Literacy Education*, 12(3), 75-87.

Friedman, A., Aguilar, A. G., & Arezes, F. C. P. (2019). Toward Peer-Review Software and a Rubric Application in Visual Analytics Classes: A Case Study. *Education for Information*, 35(3), 241-249.

Friedman, A., & Rosen, P. (2017). Leveraging Peer Review in Visualization Education: A Proposal for a New Model [Conference Paper]. *Pedagogy Data Visualization, IEEE VIS Workshop*. arXiv. https://arxiv.org/abs/2101.07708.

Fry, H. (2021, June 11). When Graphs Are a Matter of Life and Death. *The New Yorker*. https://www.newyorker.com/magazine/2021/06/21/when-graphs-are-a-matter-of-life-and-death

Howitz, W. J., McKnelly, K. J., & Link, R. D. (2021). Developing and Implementing a Specifications Grading System in an Organic Chemistry Laboratory Course. *Journal of Chemical Education*, 98(2), 385-394.

Jeyaraj, A. (2019). Teaching Tip: Pedagogy for Business Analytics Courses. *Journal of Information Systems Education*, 30(2), 67-83.

Knaflic, C. (2015). *Storytelling With Data: A Data Visualization Guide for Business Professionals*. Hoboken, NJ, USA: Wiley.

Kong, Y. (2020). Extended Abstract: What to Expect When You Are Learning Data Visualization: A Textual Analysis of Course Syllabi. *2020 IEEE International Professional Communication Conference (ProComm)* (pp. 189-190). IEEE.

Kosara, R. (2007). Visualization Criticism – The Missing Link Between Information Visualization and Art. *2007 11th International Conference Information Visualization (IV '07)* (pp. 631-636). Los Alamitos, CA, USA: CPS/IEEE.

Kosara, R. (2016, Oct.). An Empire Built on Sand: Reexamining What We Think We Know About Visualization. *BELIV '16: Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (pp. 162-168). New York, NY, USA: ACM.

Kostelnick, C. (2008, March). The Visual Rhetoric of Data Displays: The Conundrum of Clarity. *IEEE Transactions on Professional Communication*, 51(1), 116-130.

Leidig, P., & Salmela, H. (2020). IS2020: A Competency Model for Undergraduate Programs in Information Systems. The Joint ACM/AIS IS2020 Task Force. https://www.acm.org/binaries/content/assets/education/curricula-recommendations/is2020.pdf

Liu, C., Han, Y., Jiang, R., & Yuan, X. (2021). ADVISor: Automatic Visualization Answer for Natural-Language Question on Tabular Data. *2021 IEEE 14th Pacific Visualization Symposium (PacificVis)* (pp. 11-20). IEEE.

Ltifi, H., Amri, S. & Ben Ayed, M. (2018). Fuzzy Logic-Based Evaluation of Visualizations Generated by Intelligent Decision Support Systems. *Information Visualization*, 17(1), 3-21.

MacPherson-Krutsky, C. (2020, July 24). 3 Questions to Ask Yourself Next Time You See a Graph, Chart or Map. *The Conversation*. https://theconversation.com/3-questions-to-ask-yourself-next-time-you-see-a-graph-chart-or-map-141348

Munzer, T. (2014). *Visualization Analysis and Design*. Boca Raton, FL, USA: CRC Press.

Nestorov, S., Jukić, N., & Rossi, S. (2019). Design and Implementation of a Data Visualization Course with a Real-World Project Component in an Undergraduate Information Systems Curriculum. *Journal of Information Systems Education*, 30(3), 202-211.

Nolan D., & Perrett, J. (2016). Teaching and Learning Data Visualization: Ideas and Assignments. *The American Statistician*, 70(3), 260-269.

Padilla, L. M. (2018). A Case for Cognitive Models in Visualization Research: Position Paper. *2018 IEEE Evaluation and Beyond - Methodological Approaches for Visualization (BELIV)* (pp. 69-77). IEEE.

Ridgway, J. (2016). Implications of the Data Revolution for Statistics Education. *International Statistical Review*, 84(3), 528-549.

Rodrigues, A. M. B., Barbosa, G. D. J., Lopes, H. C. V., & Barbosa, S. D. J. (2021). What Questions Reveal About Novices' Attempts to Make Sense of Data Visualizations: Patterns and Misconceptions. *Computers & Graphics*, 94, 32-42.

Ryan, L., Silver, D., Laramee, R. S., Ebert, D., & Rhyne, T. (2019). Teaching Data Visualization as a Skill. *IEEE Computer Graphics and Applications*, 39(2), 95-103.

Santos, B. S., Silva, S. & Dias, P. (2018). Heuristic Evaluation in Visualization: An Empirical Study: Position Paper. *2018 IEEE Evaluation and Beyond - Methodological Approaches for Visualization (BELIV)* (pp. 78-85). IEEE.

Saundage, D., Cybulski, J. L., Keller, S., & Dharmasena, L. (2016). Teaching Data Analysis with Interactive Visual Narratives. *Journal of Information Systems Education*, 27(4), 233-248.

Stephens, P., & Young, J. (2020). Real-Time Visual Analytics: An Experiential Learning Activity for Undergraduates. *Information Systems Education Journal*, 18(6), 4-12.

Stevens, D. D., & Levi, A. J. (2013). *Introduction to Rubrics: An Assessment Tool to Save Grading Time, Convey Effective Feedback, & Promote Student Learning* (2nd Ed.). Sterling, VA, USA: Stylus Publishing.

Strobel, B., Lindner, M. A., Saß, S., & Köller, O. (2018). Task-Irrelevant Data Impair Processing of Graph Reading Tasks: An Eye Tracking Study. *Learning and Instruction*, 55, 139-147.

Tufte, E. R. (2001). *The Visual Display of Quantitative Information* (2nd Ed.). Cheshire, CN, USA: Graphics Press.

Vaughan, D. (2016). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA* (Exp. Ed.). Chicago, IL, USA: University of Chicago Press.

Wakeling, S., Clough, P., Wyper, J., & Balmain, A. (2015). Graph Literacy and Business Intelligence: Investigating User Understanding of Dashboard Data Visualizations. *Business Intelligence Journal*, 20(4), 8-19.

Watson, H. J. (2009). Tutorial: Business Intelligence – Past, Present, & Future. *Communications of the Association for Information Systems,* 25(39), 487-510.

Watson, H. J., Crince, H., & Tonon, B. (2017). The Analytics Portal at Mozilla. *Business Intelligence Journal*, 22(4), 5-13.

Wexler, S., Shaffer, J., & Cotgreave, A. (2017). *The Big Book of Dashboards: Visualizing Your Data Using Real-World Business Scenarios*. Hoboken, NJ, USA: Wiley.

Wisniewski, B., Zierer, K., & Hattie, J. (2020). The Power of Feedback Revisited: A Meta-Analysis of Educational Feedback Research. *Frontiers in Psychology*, 10(3087), 1-14.

Zhang, L., Chen, F., & Wei, W. (2020). Teaching Tip: A Foundation Course in Business Analytics: Design and Implementation at Two Universities. *Journal of Information Systems Education*, 31(4), 244-259.

AUTHOR BIOGRAPHY

William K. McHenry is associate professor of information systems at the College of Business, University of Akron, in Akron, Ohio. He received his Ph.D. in MIS from the University of Arizona in 1985, taught for 15 years at the McDonough School of Business at Georgetown University, and joined the Akron faculty in 2000. He teaches classes on MIS, decision support, knowledge management, business intelligence and visualization, data warehousing, IS and data governance, and business analytics. Early research was about MIS in the USSR and international computing. Recent papers have concerned knowledge management systems and topics related to teaching IS.

APPENDICES

Appendix A. Rubric of Friedman and Rosen (2017)

Area	Item	Scale		
Algorithmic Design	Selection of Algorithm	Below Average	Average	Above Average
	Correct implementation	No	Minor Errors	Appears Correct
	Efficient implementation	Much Slower	As Expected	Much Faster
	Featureful implementation	Major Features Missing	As Expected	Major Features Added
	Datasets Used	Not Useful	As Expected	Better than Expected
Visual Design	Visual Channels: (check which present) Position, Length, Area, Shape, Color Hue, Animation, Angle			
	Intended Encodings	Many Unintended	Few Unintended	All Intended
	Encoding Expressiveness	Many Errors	Few Errors	Correctly Assigned
	Encoding Effectiveness	Many Ineffective	Few Ineffective	Most Effective
	Effective Use of Color	Mostly Ineffective	None Used	Highly Effective
Interaction	Interaction Selection: (check which present) Selection, Highlighting, Linked Views, Pan/Translate			
	Interaction Effectiveness	Missing	As Expected	Better than Expected
Design Consideration	Clear/Thorough Labeling	No labels	Some Missing labels	Completely labeled
	Data/Ink Ratio	Way Too Little/Much Ink	Slightly Too Little/Much Ink	Perfect Amount of Ink
	Missing Scales	No Scales	Some Missing Scales	All Scales Present
	Missing Legend	No Legend	Incomplete Legend	Complete Legend
	Scale Distortion	Severe Distortion	Minor Distortion	No Distortion
	Lie Factor	Major Lie	Minor Lie	No Lie
	Chart Junk & Embellishments	Way Too Little/Much	A Bit Too Little/Much	Perfect Amount
	Data Density	Too Sparse	Expected	Too Dense
	Gestalt Principals	No Gestalt Principals	Some Gestalt Principals	Many Gestalt Principals
Narrative Consideration	Accurate & Informative	No Description	Incomplete/Self-Explanatory	Complete Description
	Support of Narrative	No Description	Incomplete/Self-Explanatory	Complete Description
	Datasets Used provide enough information and detail to support narrative	Not At All	Partially	Completely

Appendix B. Rubric of Nolan and Perrett (2016)

Critical Task	Competency Level		
	Needs Improvement	Basic	Surpassed
Computation: Perform computations	Computations contain errors and extraneous code	Computations correct but contain extraneous / unnecessary code	Computations correct and properly identified and labeled
Analysis: Choose and carry out analysis appropriate for data and context	Choice of analysis overly simplistic, irrelevant, or missing key component	Analysis appropriate, but incomplete/important features, assumptions not made explicit	Analysis appropriate, complete, advanced, relevant, informative
Synthesis: Identify key features of the analysis, and interpret results (in context)	Conclusions are missing, incorrect, or not made based on analysis	Conclusions reasonable, but partially correct or partially complete	Relevant conclusions explicitly connected to analysis and context
Visual presentation: Communicate findings graphically clearly, precisely, and concisely	Inappropriate choice of plots; poorly labeled plots; plots missing	Plots convey information correctly but lack context for interpretation	Plots convey information correctly w/ adequate / appropriate information
Written: Communicate findings clearly, precisely, and concisely	Explanation is illogical, incorrect, or incoherent.	Explanation is partially correct but incomplete or unconvincing	Explanation is correct, complete, and convincing

Information Systems & Computing Academic Professionals

Education Special Interest Group

STATEMENT OF PEER REVIEW INTEGRITY

All papers published in the ***Journal of Information Systems Education*** have undergone rigorous peer review. This includes an initial editor screening and double-blind refereeing by three or more expert referees.

ISSN 2574-3872