Can You Trust This Chart? A Visualization Game Exploring Deception in Human- and AI-Created Charts

M. Keck[1]

[1]University of Applied Sciences Upper Austria, Campus Hagenberg, Austria

Abstract

AI tools increasingly let users generate data visualizations via natural-language prompts, lowering the barrier for novices. At the same time, this introduces new risks: visualizations can become misleading not only through design choices, but also through prompt formulation or uncritical acceptance of AI outputs, which may produce misleading results even without explicit intent. This paper presents "Can You Trust This Chart?", a team-based visualization game for workshop settings that explores how people create and interpret faithful and misleading visualizations, both created by sketching and generated using AI tools. Teams alternate between a creator role, in which they design faithful or misleading visualizations, and a detective role, in which they analyze and classify visualizations created by others. Through gameplay, discussion, and reflection, the game aims to foster visualization literacy and AI literacy, encouraging participants to critically reflect on how visualizations can be manipulated, how design choices shape interpretation, and how prompt formulation influences AI-generated outputs.

CCS Concepts

• ***Human-centered computing*** → ***Information visualization;*** *Visualization design and evaluation methods;* • ***Applied computing*** → *Interactive learning environments;* • ***Computing methodologies*** → *Artificial intelligence;*

1. Introduction

Data visualizations are powerful tools for supporting decision-making and communicating information. However, they can also be misleading. Examples include truncated axes, selective data, or inappropriate chart types, all of which may substantially affect interpretation [Tuf01, Cai19, PRS*15, CBF20, LGS*22]. Misleadingness is not limited to explicit design violations; perceptual biases, aggregation effects, omission of uncertainty, and rhetorical framing can also contribute to misinterpretation [Sch19, BB22, Hul20].

These findings underscore the importance of visualization literacy – the ability to confidently create, interpret, and critically evaluate visual representations of data, including the identification of misleading elements [ARC*17, VBH*25].

The rapid advancement of artificial intelligence has transformed how visual content is created and interpreted. Text-to-image systems now allow users, including non-experts, to generate sophisticated visual content and data visualizations from natural language prompts [BM24]. While these systems democratize visualization creation, they also introduce new risks: AI-generated outputs may be unintentionally misleading due to vague or biased prompts, or as a result of model hallucinations that corrupt or distort the underlying data [KS26]. Furthermore, users may uncritically accept AI-generated visualizations, even when they contain misleading elements or inappropriate design choices. This development raises important questions about trust, responsibility, and human–AI collaboration in visualization creation.

Game-based and playful approaches have been explored as effective methods for engaging participants with data visualization in educational and workshop settings [SBK*25]. Some of these games involve actively creating and classifying visualizations [PS25], while others focus on chart interpretation or teaching visualization design principles [BS25, AGR25]. However, none of these approaches address the specific challenges introduced by AI-generated content or the critical evaluation of misleading visualizations.

To address these challenges, this paper presents "Can You Trust This Chart?", a team-based visualization game designed for workshop settings with 8–30 participants. Played using a shared digital workspace, the game combines active visualization creation with critical analysis: teams alternate between the creator role, in which they design faithful or misleading visualizations using AI tools or sketching, and the detective role, in which they analyze and classify visualizations created by others. The game has been piloted in two university courses, providing initial evidence of its feasibility and engagement. By actively experiencing both the creation and detection of misleading visualizations, participants develop critical awareness of how design choices, prompt formulation, and AI-generated outputs influence interpretation and trust.

2. Goals and Rules of the Play

The goal of the game is to explore how data visualizations can be designed and interpreted as either faithful or misleading, and how design decisions, prompt formulation, and AI-generated outputs influence interpretation. Players alternate between two roles: the creator role, in which teams design faithful or misleading visualizations, and the detective role, in which teams analyze and classify visualizations created by others. Players earn points for each visualization correctly classified as faithful or misleading.

The game is designed not only as a competitive activity but also as a reflective exercise. A successful session is characterized by engagement with the following topics:

- Justification for why a visualization was judged faithful or misleading
- Identification of specific design and data choices that influenced interpretation
- Insights into how prompt formulation influences AI-generated outputs
- Reflection on how misleading elements can arise without explicit intent

3. Game Manual

The game is played in small teams and consists of a structured sequence of steps in which participants alternate between the creator and detective roles.

Step 1 – Preparation (15 min): Before the game begins, the facilitator introduces participants to the five-level taxonomy for misleading visualizations developed by Lo et al. [LGS*22]:

- **Input Stage:** Data Curation and Wrangling
- **Visualization Design Stage:** Choices in Visual Encoding
- **Plotting Stage:** Drawing on the Canvas
- **Perception Stage:** Visually Perceiving the Visualization
- **Interpretation Stage:** Comprehending the Message

For each stage, concrete examples are presented on slides, such as cherry-picking, truncated axes, and hidden uncertainty. A summary of the taxonomy is made available to participants throughout the session as a reference, pinned in a digital workspace (see Figure 2). This 15-minute introduction prepares participants for both the creator and detective roles by making explicit that misleading effects can arise at multiple stages of the visualization process.

Participants are then divided into small teams of two to four people. The game is designed for groups of 8–30 participants in a world café setting — a format in which small teams work simultaneously at separate stations and rotate between them in a structured manner. Before the session begins, participants should be asked to ensure they have access to at least one AI tool of their choice (e.g., ChatGPT, Claude, or Gemini; a free account suffices). Allowing teams to use different tools is intentional: differences in output quality, style, and behavior across tools can enrich the subsequent discussion. Each team receives access to a shared digital workspace (e.g., a Miro board) containing the step-by-step instructions, input materials, and datasets. The datasets are small and easy-to-understand, such as time series or category comparisons, and are provided in a tabular form, so that data can be directly copied into an AI tool or used as the basis for a manually sketched visualization.

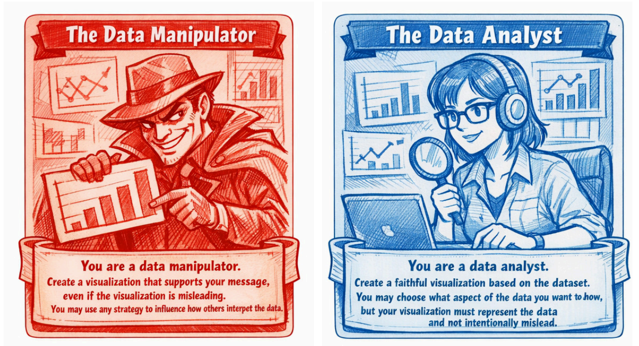

Figure 1: *Role cards: Data Manipulator and Data Analyst, generated using an AI image generation tool (Academic AI).*

Step 2 – Creation Phase (20 min): In the shared digital workspace, two role cards are available to each team, defining the possible creator roles: the **Data Analyst**, who should create a faithful visualization that accurately represents the data, and the **Data Manipulator**, who should deliberately create a misleading visualization (see Figure 1). When taking on the Data Analyst role, teams may choose which aspect of the data to emphasize, but the visualization must accurately represent the data (e.g., appropriate scales, clear labels and legends, and no misleading encodings). When taking on the Data Manipulator role, teams may use any manipulation strategy (e.g., truncated axes, cherry-picking, or inappropriate chart types) to bias interpretation. Teams are encouraged to try out both roles and create several visualizations using different role and dataset combinations. Discussing design decisions and strategies within the team during the creation process is explicitly encouraged.

To create a visualization, teams may either use an AI tool by entering a natural language prompt, or sketch a visualization – either directly in the digital workspace or on paper, which can then be photographed and uploaded. For each visualization created, the team records the intended role (faithful or misleading), and if AI was used, also the prompt and tool name. This information is kept hidden from other teams until the reveal phase. The created visualizations are shared with the other teams at the end of this step by dragging them onto the team's frame in the shared Miro board.

Step 3 – Detective Phase (15–20 min): After the creation phase, all teams take on the role of data detectives. In a clockwise rotation, each team moves to the next team's frame in the shared digital workspace, so that each team receives a new set of visualizations created by others. The teams then analyze the visualizations and discuss the following questions: Is the visualization faithful or misleading? If it is misleading, what type of manipulation has been applied? Teams collaboratively place each visualization along a single axis ranging from faithful to misleading in the shared workspace, and briefly justify their decisions. Visualizations that are difficult to classify can be placed in the middle of the axis, indicating uncertainty or ambiguity. To support structured reflection, participants identify misleading elements and categorize them

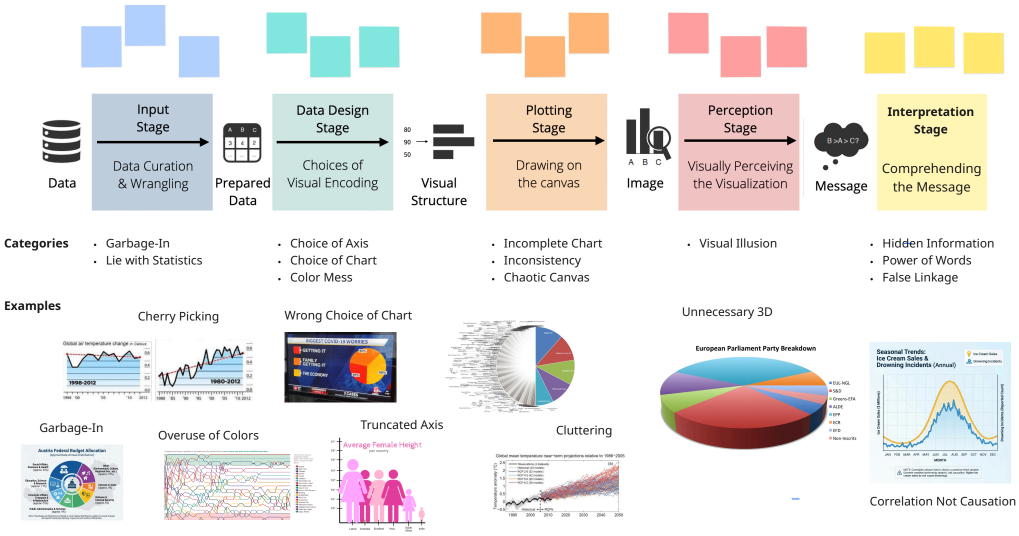

Figure 2: *Taxonomy cheat sheet for misleading visualizations according to Lo et al. [LGS*22] provided in the shared Miro board, with color-coded post-its available for each taxonomy stage.*

using Lo et al.'s taxonomy [LGS*22]. For each misleading element identified, a colored digital post-it should be added to the visualization, indicating the stage at which the misleading effect occurs (e.g., blue for input stage, green for visual design stage). The taxonomy overview pinned in the workspace serves as a reference throughout this phase, including a color legend for the post-it categories (see Figure 2).

Step 4 – Reveal Phase (15–20 min): After the detective phase, teams return to their own frame in the shared digital workspace. Each visualization is revealed by placing the corresponding role card next to it: the blue *Data Analyst* card indicates that the visualization was intended to be faithful, and the red *Data Manipulator* card indicates that it was intended to be misleading. If AI was used, the prompt and tool name should be added as a post-it next to the visualization as well. If a visualization was intended to be faithful but the creator team recognizes that it contains misleading elements, a red border should be added to the blue *Data Analyst* card to indicate that it is unintentionally misleading. These cases are particularly valuable discussion opportunities, as they illustrate how misleading elements can arise without explicit intent. Teams then compare the detective team's classifications with the intended roles. If the intended role does not match the detective team's classification, this can be discussed in the following step. Detective teams receive one point for each visualization correctly classified as faithful or misleading.

Step 5 – Discussion Phase (10–15 min): In a final plenary discussion, teams share highlights from their detective phase. Rather than reviewing all visualizations in detail, the facilitator should invite teams to report on particularly surprising or ambiguous cases, moments of disagreement within the team, and any visualizations that turned out to be unintentionally misleading. The following questions can guide the discussion:

- Which visualization surprised you the most, and why?
- Were there cases where your team disagreed on whether a visualization was faithful or misleading?
- Did you encounter any unintentionally misleading visualizations, and what made them misleading?
- How did the prompt formulation influence the AI-generated output?
- Did you refine your prompts to achieve the desired result, and if so, what strategies did you use?

If time allows, the game can be played over several rounds, giving teams the opportunity to try out more creation methods and datasets.

4. Usage & Variations

The game is designed for workshop settings with approximately 8–30 participants and a duration of 75–90 minutes for a single round, including the introductory phase. It is intended for use in educational contexts, such as university courses on data visualization or AI literacy, as well as in interdisciplinary workshops with participants of varying expertise levels.

The game can be adapted in several ways depending on the learning goals, available resources, and time constraints:

Short Version (45–60 min): For shorter sessions, the taxonomy introduction can be reduced to a brief overview, and teams are asked to create a minimum of two visualizations instead of three. The reveal and discussion phases can be combined into a single step focusing only on the most surprising or ambiguous cases.

Physical Card Game Variant: The game can be played entirely without digital tools, using printed dataset cards, role cards, paper, and pens for sketched visualizations, and a large A3 sheet featuring

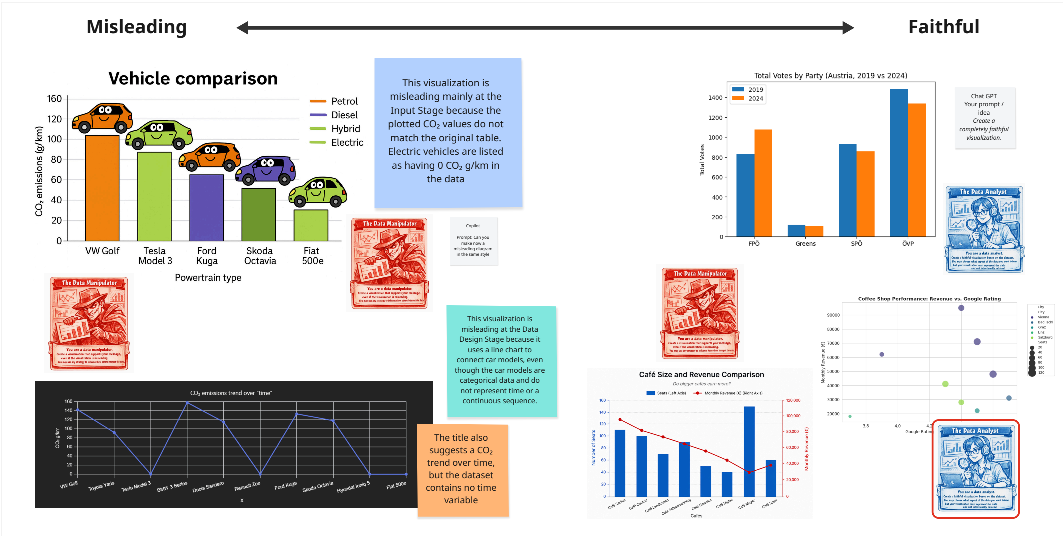

Figure 3: *Example of a team frame from the master's course pilot session. Two visualizations were correctly classified as misleading (left); one was correctly identified as faithful (blue Data Analyst card, right), while one misleading visualization was not detected (red Data Manipulator card, right). A fifth visualization (blue Data Analyst card with a red border) was intended as faithful but was found to contain unintentional misleading elements during the reveal phase.*

a coordinate system ranging from faithful to misleading for the detective phase. In this variant, AI-generated visualizations must be redrawn by hand before being shared with other teams, masking their origin. This makes the physical card game variant particularly suitable for including a second classification dimension in the detective phase: participants not only assess whether a visualization is faithful or misleading, but also whether it was created by a human or generated using AI. This adds an additional layer of reflection on authorship, trust, and the perceived credibility of AI-generated visualizations.

AI-Only or Sketching-Only: In the AI-only variant, all visualizations are generated using AI tools, placing emphasis on prompt design and critical evaluation of AI outputs. In the sketching-only variant, all visualizations are drawn by hand, focusing on visualization design principles and manual creation. Both variants can be played in the digital or analog setup.

Scenario Cards: In an extended version of the game, scenario cards can be introduced that define a communication goal, such as emphasizing a trend or highlighting differences. These cards introduce the concept of framing and allow players to explore how the same communication goal can be achieved through both faithful and misleading visualization strategies.

These variations allow facilitators to tailor the game to different audiences, time constraints, and learning objectives.

4.1. Pilot Evaluation

The game was piloted in two university data visualization courses – a bachelor's course (n = 32, teams of 3–4) and a master's course (n = 18, teams of 2–3) – using a shared Miro board as the digital workspace, following the structure described in Section 3. Five datasets were provided, covering topics such as tech stock performance, Austrian election results, vehicle emissions, coffee shop revenues, and climate and tourism data for Austrian cities. All datasets were visible to all teams, making strategies such as cherry-picking observable and discussable.

Creation Phase. Teams produced between 3 and 6 visualizations depending on team size and working pace. AI tools were strongly preferred over sketching; in both courses, only one participant per course chose to sketch visualizations, while all others used AI tools, suggesting that manual sketching is unlikely to occur spontaneously and may need to be explicitly incentivized if desired. Participants primarily used ChatGPT, Claude, and Gemini. Teams tended to create more misleading than faithful visualizations, suggesting that the manipulator role was perceived as more engaging and playful. Prompt strategies varied substantially, ranging from highly specific instructions that explicitly defined cherry-picked data ranges, axis manipulations, or inappropriate chart types – such as *"Create a misleading dual axis line graph for two of these tech stocks that make it seem like one is performing worse than the other when it is actually the other way around. Choose the scales of the axes accordingly"* or *"Create a line chart of election results for 2024 using state & party on the x-axis to suggest a continuous trend"* – to minimal prompts such as *"make misleading graph"*. A practical issue observed during the pilot evaluation was that several AI tools included explicit hints in the visualization title indicating which misleading elements they show, revealing the intended role to the detective teams. Participants were advised to adjust their prompts accordingly or crop the title before sharing the visualization on the miro board. One instance of a tool refusal was reported, in which ChatGPT declined to generate a deliberately misleading chart when the prompt explicitly framed the intent to deceive. The student resolved this by requesting a side-by-side comparison of a misleading and an faithful version, which proved equally instructive. Apart from this instance, no content policy violations or refusals were reported across any of the tools used.

A recurring challenge was that AI models occasionally generated factually incorrect visualizations. In one prominent case, an AI-generated chart displayed positive CO_2 emissions for electric vehicles, directly contradicting the dataset in which electric vehicles were listed with 0 g/km. This was successfully identified by student detectives and correctly categorized at the input stage of Lo et al.'s taxonomy [LGS*22], illustrating how AI hallucinations can introduce errors that are indistinguishable from intentional manipulation without careful verification (see Figure 3, left).

Detective Phase. The most frequently identified misleading elements across both courses were truncated y-axes, dual y-axes with differing scales, the use of line charts for categorical data, 3D effects, and cherry-picked data ranges. Among these, cherry-picking proved the most difficult to detect, as it requires comparing the visualization against the underlying data rather than evaluating visual design choices alone. However, having the complete datasets available directly in the shared Miro board allowed participants to cross-check data points, making detection possible – albeit more effortful than identifying misleading elements at the visual design or plotting stage. Both groups demonstrated a solid ability to identify misleading visualizations overall, though classification accuracy varied across teams. Several teams identified misleading elements in visualizations that were originally intended to be faithful — for example, a scatterplot with a missing size legend, or a bar chart where incorrect color assignments for political parties introduced unintended bias. These cases effectively illustrated the concept of unintentional misleadingness and prompted productive discussion. Figure 3 shows an example team frame from the master's course, in which three visualizations were correctly classified, while one misleading visualization was incorrectly placed on the faithful side of the axis, as revealed by the red *Data Manipulator* card during the reveal phase. A fifth visualization, intended as faithful, received a red border during the reveal phase after the creator team recognized it contained unintentional misleading elements: a scatterplot of Austrian cities using an ordered color scale, which implied a ranking among the cities that was not present in the data.

Differences Between Groups. A notable difference emerged between the two cohorts in their application of Lo et al.'s taxonomy [LGS*22]. Master's students produced detailed, stage-specific annotations and provided nuanced justifications – for example, correctly identifying that an AI-generated chart showed incorrect CO_2 values for electric vehicles, or that a map title implied a conclusion not supported by the underlying data. Bachelor's students engaged less consistently with the taxonomy and tended to focus primarily on the faithful vs. misleading classification. This suggests that a simplified version of the game may be more appropriate in introductory settings, in which participants are only asked to identify misleading elements rather than categorize them according to a formal taxonomy.

Overall Assessment. Both sessions generated active engagement and substantive discussion. Informal feedback from participants indicated that the game was perceived as enjoyable, and several students noted that the detective role prompted them to examine visualizations more critically than they would have otherwise. The plenary discussion focused on highlights and edge cases rather than covering all visualizations in detail, which proved an effective use of the limited time.

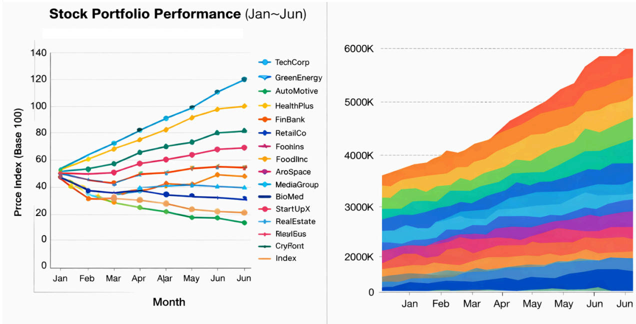

Figure 4: *Two visualizations generated from the same dataset using different prompts and created with the GPT-4o model. The left visualization was intended to be a faithful visualization, but still contains unintended misleading elements. The right visualization was intentionally designed to be misleading through the overuse of color and a truncated y-axis.*

5. Reflection & Conclusion

The game encourages participants to actively engage with both the creation and interpretation of data visualizations. By requiring players to produce both faithful and misleading visualizations, it reveals how easily visual representations can influence interpretation and how subtle design decisions can change the perceived message of a chart. The tension between successfully misleading others and being detected adds a competitive element that appeared to drive engagement, as observed in the pilot sessions where teams spontaneously invested effort in crafting convincing misleading visualizations. The reveal phase, in which intended roles and prompts are disclosed, creates moments of surprise and insight.

A key insight that emerges through gameplay is the distinction between *incorrect* and *misleading* visualizations. Figure 4 illustrates this with two AI-generated visualizations from the same dataset. The left visualization was intended to be faithful using the prompt: *"Use this dataset and create a clear line chart showing the performance of 14 stocks in a portfolio from January to June. Use distinct colors, include legend, labels, gridlines, and a y-axis starting at 0"*. Despite these instructions, the resulting visualization contains unintended misleading elements, including duplicated axis labels, incorrect label naming, and an overuse of color – the latter attributable to the vague prompt, which did not specify that this encoding strategy may cause visual clutter when applied to a large number of categories. The right visualization was intentionally designed to be misleading using the prompt: *"Use this dataset and create a stacked area chart showing the performance of 14 stocks in a portfolio from January to June and make it misleading by the overuse of color and not starting the y-axis at zero. Do not include any labels or messages in the visualization that could give a hint that it is misleading"*. This chart deliberately uses a truncated y-axis and an overuse of color, and additionally omits the legend, which is required to correctly interpret the data. These examples highlight that misleading effects can arise not only through deliberate manipulation but also through vague prompts or uncritical acceptance of AI outputs – a central learning experience of the game.

Challenges in the pilot study include the complexity of Lo et al.'s taxonomy, which proved more accessible for experienced participants than for novices. Facilitators should be prepared to guide discussions around boundary cases and unintentionally misleading visualizations, as these often generate the most productive reflection. In introductory settings, a simplified version may be more appropriate in which participants are only asked to identify misleading elements rather than categorize them by stage. A further challenge concerns the scoring system: teams produced varying numbers of visualizations – between 3 and 6 – resulting in unequal workloads during the detective phase. Future iterations should address this by either asking teams to share only a fixed number of visualizations (e.g. their three most interesting ones), or by calculating a detection rate rather than an absolute score, so that teams are not penalized for having analyzed fewer visualizations.

Overall, the game provides a playful yet critical environment designed to promote visualization literacy and AI literacy. Future work will involve a more systematic study investigating the prompts participants use when creating faithful or misleading visualizations, as well as the visual features that most strongly influence judgments of faithfulness and credibility.

5.1. Material Overview

The game requires a shared digital workspace (e.g., a Miro board) containing all necessary materials: step-by-step instructions for each phase, example datasets presented in tabular form, role cards defining the Data Analyst and Data Manipulator roles, a summary of Lo et al.'s taxonomy [LGS*22] including a color legend for the post-it categories, and a frame for each team to collect and display their visualizations. A Miro board template including example datasets, step-by-step instructions, and team frames is provided as supplemental material at https://doi.org/10.17605/OSF.IO/SZUM2.

For sketched visualizations, pen and paper are sufficient. Sketches can be photographed and uploaded directly to the shared workspace.

References

[AGR25] Amabili L., Gröller E., Raidou R. G.: Leveraging Popular Board Games to Teach Data Visualization. In *EuroVis Workshop on Visualization Play, Games, and Activities* (2025), Stoiber C., Boucher M., Keck M., Amabili L., Raidou R. G., Filipov V., Oliveira V., Schetinger V., Aigner W., (Eds.), The Eurographics Association. URL: https://visgames2025.netlify.app/_astro/paper5.DKm7t8Ij.pdf. 1

[ARC*17] Alper B., Riche N. H., Chevalier F., Boy J., Sezgin M.: Visualization literacy at elementary school. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2017), CHI '17, Association for Computing Machinery, p. 5485–5497. doi:10.1145/3025453.3025877. 1

[BB22] Beconytė G., Balčiūnas A., Šturaitė A., Viliuvienė R.: Where maps lie: Visualization of perceptual fallacy in choropleth maps at different levels of aggregation. *ISPRS International Journal of Geo-Information 11*, 1 (2022). doi:10.3390/ijgi11010064. 1

[BM24] Basole R. C., Major T.: Generative ai for visualization: Opportunities and challenges. *IEEE Computer Graphics and Applications 44*, 2 (2024), 55–64. doi:10.1109/MCG.2024.3362168. 1

[BS25] Boucher A., Stoiber C.: Cards, Charts, and Strategy: A Game-Based Approach to Data Visualization for Pattern. In *EuroVis Workshop on Visualization Play, Games, and Activities* (2025), Stoiber C., Boucher M., Keck M., Amabili L., Raidou R. G., Filipov V., Oliveira V., Schetinger V., Aigner W., (Eds.), The Eurographics Association. doi:https://visgames2025.netlify.app/_astro/paper2.D_MwQeYC.pdf. 1

[Cai19] Cairo A.: *How Charts Lie: Getting Smarter about Visual Information*. W.W. Norton & Company, New York, 2019. 1

[CBF20] Correll M., Bertini E., Franconeri S.: Truncating the y-axis: Threat or menace? In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2020), CHI '20, Association for Computing Machinery, p. 1–12. doi:10.1145/3313831.3376222. 1

[Hul20] Hullman J.: Why authors don't visualize uncertainty. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2020), 130–139. doi:10.1109/TVCG.2019.2934287. 1

[KS26] Keck M., Stöckl A.: When AI lies with charts: Misleading infographics in text-to-image generation. In *Proceedings of the 2026 International Conference on Advanced Visual Interfaces* (2026), AVI '26, ACM. doi:10.1145/3811427.3811469. 1

[LGS*22] Lo L. Y.-H., Gupta A., Shigyo K., Wu A., Bertini E., Qu H.: Misinformed by visualization: What do we learn from misinformative visualizations? *Computer Graphics Forum 41*, 3 (2022), 515–525. doi:https://doi.org/10.1111/cgf.14559. 1, 2, 3, 5, 6

[PRS*15] Pandey A. V., Rall K., Satterthwaite M. L., Nov O., Bertini E.: How deceptive are deceptive visualizations? an empirical analysis of common distortion techniques. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2015), CHI '15, Association for Computing Machinery, p. 1469–1478. doi:10.1145/2702123.2702608. 1

[PS25] Plankensteiner S., Stoiber C.: Drawagram: A Game-Based Learning Approach to Teach Time-based Data Visualization. In *EuroVis Workshop on Visualization Play, Games, and Activities* (2025), Stoiber C., Boucher M., Keck M., Amabili L., Raidou R. G., Filipov V., Oliveira V., Schetinger V., Aigner W., (Eds.), The Eurographics Association. doi:https://visgames2025.netlify.app/_astro/paper3.ffawGA8m.pdf. 1

[SBK*25] Stoiber C., Boucher M., Keck M., Amabili L., Raidou R. G., Filipov V., Oliveira V., Schetinger V., Aigner W.: EuroVis Workshop on Visualization Play, Games, and Activities 2025: Frontmatter. In *EuroVis Workshop on Visualization Play, Games, and Activities* (2025), Stoiber C., Boucher M., Keck M., Amabili L., Raidou R. G., Filipov V., Oliveira V., Schetinger V., Aigner W., (Eds.), The Eurographics Association. URL: https://visgames2025.netlify.app/. 1

[Sch19] Schiewe J.: Empirical studies on the visual perception of spatial patterns in choropleth maps. *KN - Journal of Cartography and Geographic Information 69* (2019), 217–228. doi:10.1007/s42489-019-00026-y. 1

[Tuf01] Tufte E. R.: *The Visual Display of Quantitative Information*, 2nd ed. Graphics Press, Cheshire, CT, 2001. 1

[VBH*25] Varona M., Bonilla K., Hedayati M., Joshi A., Harrison L., Kay M., Nobre C.: The state of the art in visualization literacy, 08 2025. doi:10.48550/arXiv.2509.01018. 1