

JUST SEEING CAN BE DECEIVING: GENAI-SUPPORTED DESIGN CASE FOR CRITICAL VISUAL LITERACY

Jihyun Yu¹ & Jihyun Rho²

¹University of North Texas; ²University of Florida

Generative AI (GenAI) poses a pedagogical challenge to data visualization education by enabling the rapid creation of persuasive, misleading visuals. This design case presents a three-phase learning experience designed for a graduate course to foster critical AI-visual literacy. The design inverts the typical use of AI by tasking students with using GenAI as a ‘cognitive partner’ to intentionally create a misleading version of their own honest data visualization. The activity guides students through establishing an ethical baseline, creating the misleading visual, and engaging in peer critique, supported by scaffolds like a prompting guide and reflection templates. Analysis of student work shows the design effectively fostered critical AI literacy.

Jihyun Yu is an Assistant Professor at the University of North Texas. Her research interests focus on learning design and learning analytics.

Jihyun Rho is a Postdoctoral Researcher at the University of Florida. Her research explores visualization education and artificial intelligence in education.

INTRODUCTION

For years, our graduate-level course in visual communication followed a consistent and effective structure. A primary objective was to teach students the principles of creating effective and ethical data visualization. Guided by the work of pioneers like Edward Tufte (2001) and Alberto Cairo (2012), our instruction focused on the established rules of graphical integrity: maximizing the integrity of data visualizations, avoiding non-essential visual elements, and ensuring accuracy to represent findings with transparency. We found this instructional approach effective for preparing students to be responsible communicators in their professional careers as researchers, analysts, and leaders.

This stable pedagogical model was recently challenged by the widespread availability of generative artificial intelligence (GenAI). The introduction of powerful, easy-to-use tools that produce sophisticated visual content from simple text prompts presented a fundamental challenge to our teaching philosophy. We began to observe our students using these tools not just for technical help, such as debugging code, but also as active consultants for the visual and rhetorical design of their data presentations. This created a new and complex problem. The issue was not a simple matter of academic integrity, but rather that students, acting in good faith, could be guided toward unethical design choices without a sufficient critical framework to recognize the danger.

The nature of this new challenge is deeply pedagogical. GenAI models are not designed to understand statistical ethics or the principles of accurate graphical representation. Their main function is to generate a plausible and often persuasive output based on patterns identified in their training data. For example, if a user asks an AI to make a chart more impactful, the model might suggest truncating a y-axis to exaggerate a small difference. It may do so because this technique is correlated with “impactful” visuals in its data, not because it is an ethically or methodologically sound choice. The AI acts as a powerful but ethically neutral tool, able to generate a misleading chart as easily as an honest one. By ethically neutral, we mean that the model does not evaluate



A publication of the Association for Educational Communications & Technology (AECT), published by Indiana University Libraries Journals.

Copyright © 2025 by the International Journal of Designs for Learning, a publication of the Association of Educational Communications and Technology. (AECT). Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page in print or the first screen in digital media. Copyrights for components of this work owned by others than IJDL or AECT must be honored. Abstracting with credit is permitted.

<https://doi.org/10.14434/ijdl.v16i2.41987>

intentions or outcomes through any moral lens; it optimizes for plausibility and rhetorical effectiveness rather than truth or integrity. Ethical responsibility therefore, remains with the human designer and the surrounding educational context. This new reality revealed a significant gap in our existing curriculum. The established rules of good design were still important, but a new, more urgent need appeared. We had to teach students how to critically manage the recommendation of a persuasive, non-human partner.

This realization required us to reframe the entire problem of data visualization education. The central question was no longer simply, “How do we teach students the principles for creating a good chart?” We now face a more complex set of questions: “How do we prepare students to collaborate with a powerful tool that does not share their ethical commitments? How do we help them develop a critical perspective that allows them to question, evaluate, and direct the output of an AI?” This meant we needed to shift our instructional approach from one based on teaching a set of fixed rules (a prescriptive pedagogy) to one focused on teaching critical questioning (an inquiry-based pedagogy). Our goal was to help students develop durable and flexible judgment for these new circumstances.

Our response was the development of a three-phase learning experience. This design is built on an unconventional premise: instead of using GenAI only to support the creation of correct visualizations, our activity requires students to use it as a collaborative partner to intentionally create misleading visuals. The purpose of this method is not to teach deception, but to use the act of creation to deconstruct it. By actively participating in the process of manipulation, guided by structured prompts and critical reflection, students directly experience the persuasive power of unethical design choices. This process compels them to confront the ease with which data can be distorted and to better articulate the ethical boundaries for their own professional practice.

This article presents a design case that details this educational intervention. We recount the narrative of our design’s genesis, from our initial observations in the classroom to the formulation of our conceptual approach. We then provide a detailed description of the three-phase learning experience and its related artifacts, followed by an account of the design in practice, showing how learners engaged with the activity. Finally, we offer our reflections on the process, transparently discussing the key decisions, tensions, and challenges we navigated as designers. We offer this case not as a final solution, but as a detailed precedent, a documented account of one attempt to address a complex and evolving pedagogical challenge in the age of AI.

THEORETICAL LENSES FOR OUR DESIGN

Our design process was not driven by a single, rigid theory, but was instead informed by several theoretical lenses that served as starting points and guides for our decision-making. These lenses helped us translate our classroom observations into a structured pedagogical intervention. We describe the three most influential lenses that shaped the final form of our learning activity.

Design Philosophy

Our design philosophy views learning design as both an intellectual and ethical practice. It requires designers to make their reasoning visible and their values explicit. Dewey’s (1938) notion of education as the reconstruction of experience through inquiry shaped our belief that students must act within complex problems, not only study them from the outside. His idea of “learning by doing” justified our inversion of the usual visualization assignment: instead of simply critiquing misleading visuals, students constructed them to experience firsthand how persuasive distortion can be. Freire’s (2018) call for critical consciousness extends this logic. He reminds us that reflection is inseparable from action, and that learners must interrogate the social and ethical forces that shape their practice. This idea underpinned our decision to embed structured reflection and peer critique throughout the design so that an act of creation could also become an opportunity for ethical awareness.

Cairo’s (2012, 2019) work translated these philosophical commitments into design terms. His insistence that visualization is not a neutral reflection of data but a rhetorical argument gave us both the conceptual and practical grounding for the assignment’s ethical focus. We drew directly from his taxonomy of misleading techniques to help students identify, reproduce, and critique deception within their own work. In this way, Cairo’s writing linked Dewey’s experiential learning and Freire’s critical pedagogy to a concrete, visual design practice.

Finally, Shneiderman’s (2022) human-centered AI perspective guided how we framed GenAI in our classroom. Rather than treating the AI as a passive tool, we designed the interaction scaffold to position it as a cognitive partner, a collaborator whose suggestions invite judgment, negotiation, and critique. This framing reflected our belief that ethical agency in AI-supported learning arises not from avoiding AI, but from learning to manage it critically.

Altogether, these perspectives form the foundation of our approach to design. We use theory not as a decorative citation but as a working grammar for reasoning about design choices. It provides an intellectual scaffold that helps us justify our methods, stay accountable to learners, and continually adapt our practice to the evolving ethical terrain of AI in education. This view aligns with a studio-based

understanding of design as a reflective practice in which reasoning, values, and artifacts are inseparable.

Lens 1: Critical Pedagogy & “Learning by Doing”

The first and most foundational lens for our design was the principles of active learning, often described as “learning by doing” (Dewey, 1938), which is a central component of critical pedagogy (Freire, 2018). Contemporary scholars continue to explore this connection, advocating for technology-supported learning environments that empower students to critically engage with complex social issues rather than simply consuming content (Morris & Stommel, 2018). This perspective argues that deep and lasting understanding, particularly of complex or abstract concepts, is best achieved when learners actively construct knowledge for themselves rather than passively receiving information (Papert, 1980). The effectiveness of such active, construction-based approaches over passive instruction is well supported by a significant body of contemporary research in higher education (Freeman et al., 2014).

For teaching ethics, this principle is especially relevant. A lecture or a reading that simply lists unethical practices may be easily forgotten. In contrast, an experience that requires the learner to engage with an ethical dilemma directly is more likely to result in a durable and personal understanding (Kolb, 1984). This is particularly true in the emerging field

of data and AI ethics, where recent pedagogical models emphasize dilemma-based scenarios and hands-on activities to help learners navigate ambiguous, real-world problems (Vakil, 2023).

This lens directly addressed the problem we identified in our classroom. We concluded that merely telling our students that certain visual techniques were misleading would be an insufficient defense against the persuasive suggestions of an AI. They needed a more powerful way to understand the rhetorical power and ethical implications of these techniques. Guided by Freire’s (2018) call for critical consciousness and Dewey’s (1938) “learning by doing,” we made our most fundamental design decision: we inverted the task. Instead of an activity where students would only critique visuals made by others, our design would require them to actively construct the misleading visualizations themselves. We reasoned that this act of construction, a core tenet of critical pedagogy, would force a direct confrontation with the ease of manipulation, building a more robust and personal understanding than any passive warning could provide.

Lens 2: Design Precedents from Cairo (2019)

While the principles of critical pedagogy and active learning provided the pedagogical structure for our design (the “how”), we needed a practical framework to supply the subject matter content (the “what”). For this, we required a

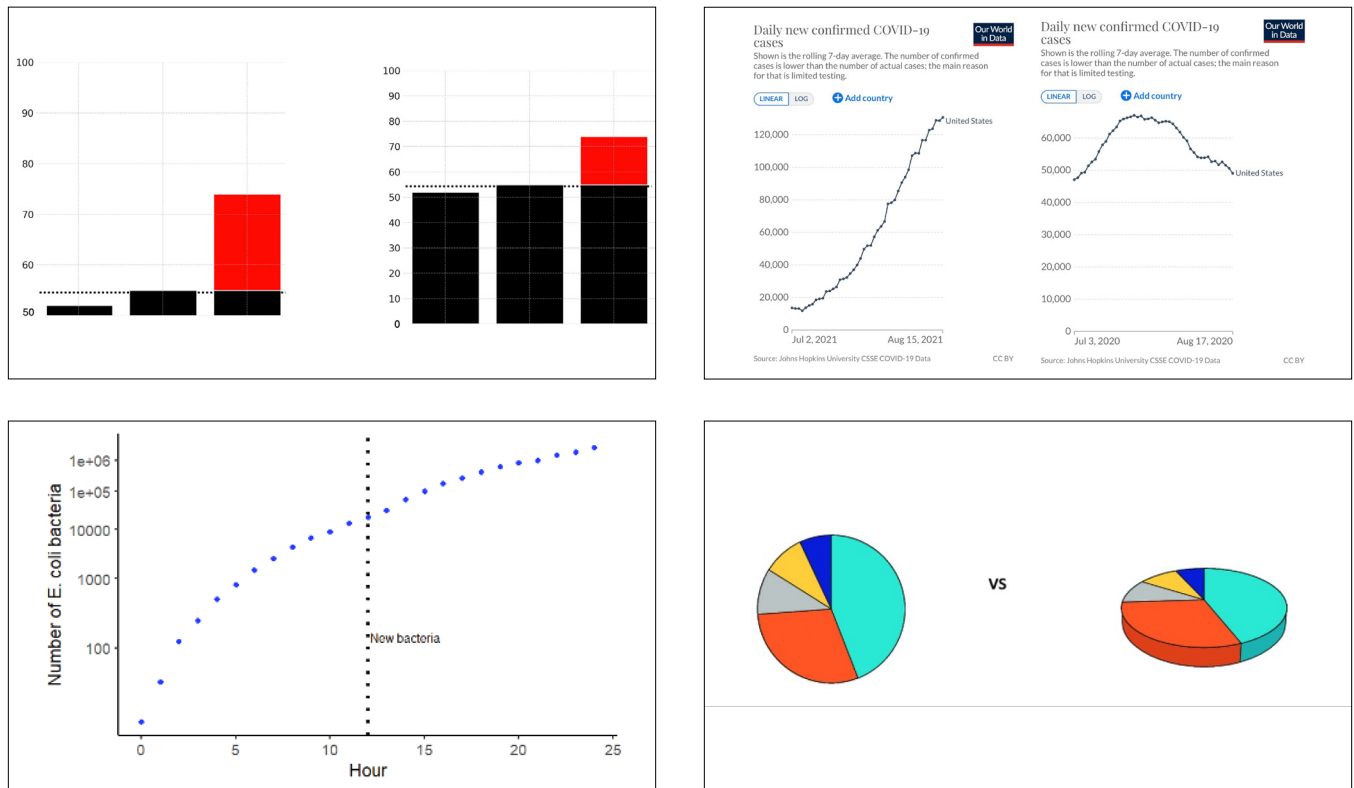


FIGURE 1. (clockwise from top left) (a) Truncating the Y-Axis Example (Rho et al., 2025), (b) Cherry-Picking Data Example (Ortiz-Ospina & Roser, 2021, Licensed under CC BY 4.0.), (c) 3D Effects Example (Rho et al., 2025), and (d) Inappropriate Scale Example (Rho et al., 2025).

clear and accessible guide to the principles of ethical data visualization. We specifically sought out established design precedents that could help us define what constitutes a “misleading” visual and provide a vocabulary for discussing specific manipulation techniques (Lisnic et al., 2023). We found this practical foundation in the work of Alberto Cairo, particularly his book *How Charts Lie* (2019).

Cairo’s work was an ideal fit for our design for two reasons. First, it is written for a broad audience and clearly explains complex topics without overly technical jargon, making its principles transferable to our graduate students who come from diverse academic backgrounds. Second, it frames visualization not as a neutral window onto data, but as a form of communication, an argument made by a designer. This view is consistent with a rich tradition of scholarship that treats visualization not as a neutral act of transcription but as a rhetorical practice, where design choices actively shape meaning and persuade audiences (D’Ignazio & Klein, 2020). Cairo’s central idea is that a chart’s “truthfulness” is not inherent in the data, but is a result of the many small, deliberate choices a designer makes. This focus on designer choice and responsibility was exactly the mindset we wanted to cultivate in our students.

We translated Cairo’s work into two concrete components of our learning activity. First, for Phase 1 of the activity (creating an “honest” chart), his work directly informed the scaffolding we created to help students ground their initial designs in best practices. For Phase 2, Cairo’s work provided the critical content for our active learning design. His taxonomy of deceptive methods (e.g., truncating the y-axis, cherry-picking data) became the conceptual backbone for the entire “misleading” activity. Instead of asking students to be misleading in a general way, we translated this theoretical precedent directly into our “Scaffolding Prompting Guide” (see Figure 4). This artifact explicitly prompted students to engage with and apply defined techniques such as truncating the y-axis (see Figure 1a), cherry-picking data to show a misleading trend (see Figure 1b), using 3D effects to distort proportions (see Figure 1c), and applying an inappropriate scale function, such as unlabeled logarithmic scale, to obscure the true rate of change (see Figure 1d). By providing this structured vocabulary, we gave students the concrete tools and language necessary to engage in the active learning task in a focused and critical manner, a practice that is seen as essential for developing modern data literacy (Correll, 2019).

Lens 3: Framing GenAI as a “Cognitive Partner”

Our third theoretical lens addressed the nature of the human-AI interaction itself. We were motivated by recent findings in human-computer interaction and AI research to frame GenAI not as a simple, passive tool, but as an active “cognitive partner” in the design process (Hwang et al., 2025). The traditional tool perspective often implies a

straightforward, command-based interaction where the user provides an instruction and the tool executes it. However, the generative and conversational nature of large language models enables a more complex, collaborative relationship. The cognitive partner framework views the human and the AI as two agents engaged in a shared task, where both can contribute ideas, critique suggestions, and co-create an outcome (Shneiderman, 2022).

This human-centered perspective directly motivated the design of our interaction for Phase 2 of the activity. This theoretical lens is the reason our design is not a simple “prompt-and-response” task. We translated the “cognitive partner” frame (Shneiderman, 2022) directly into our “Scaffolding Prompting Guide” (see Figure 4). This artifact was intentionally designed as a “three-level funnel” to facilitate a structured, critical dialogue, not a simple transaction. The prompts guided students to avoid simple commands (e.g., “make a misleading chart”) and instead engage the AI in a guided exploration (e.g., “First, explain the technique of ‘cherry-picking’ to me. Now, using my data, help me brainstorm ways to apply it:”). This design, rooted in the “cognitive partner” framework, aimed to teach students how to critically manage its outputs and maintain their own creative agency, a central goal of a human-centered approach to technology (Shneiderman, 2022).

DESCRIPTION OF THE DESIGN: FROM GENESIS TO FINAL ARTIFACT

As the instructional team for the graduate course “Learning and Cognition”, our primary role is to bridge foundational theories of how people learn with the contemporary practices of instructional design and educational research. Before the start of this specific learning activity, our students had already developed a strong conceptual toolkit, having mastered modules on core cognitive theories, motivation, metacognition, self-regulated learning (SRL), emotion and cognitive engagement, and multimedia learning. The course culminates in applying these theories to complex, real-world problems. The data visualization module is a key site for this application, providing a practical context where principles of multimedia learning directly intersect with the need for metacognitive awareness from the designer.

The specific context of this course and its learners was a critical factor that made our final design approach appropriate. The course is a core requirement in our M.S. in Learning Technologies program, and our learners are typically early-to-mid career professionals, including corporate trainers, K-12 educators, and aspiring instructional designers. While they possess strong pedagogical and design backgrounds, they often have limited formal data analysis experience. However, as graduate students, they had already demonstrated mastery of the analytical tools of cognitive science from earlier in the course. This shaped graduate-level

Phase 1 Creating an Ethical Baseline	Phase 2 Navigating Deception with a GenAI Partner	Phase 3 Peer Critique and Ethical Reflection
<ul style="list-style-type: none"> • Pre-survey: Student completes pre-survey on data viz attitudes and self-efficacy • Scaffolding Activities: Engages with readings, practice activity, and Ethical Design Checklist • Create Faithful Visualization: Analyzes dataset and creates one 'faithful' visualization using Excel. • Write Rationale: Writes a 200-300 word rationale defending the chart's ethical and cognitive design choices. 	<ul style="list-style-type: none"> • Use Faithful Visualization: Begins with the 'faithful' visualization from Phase 1. • Scaffolding Activities: Reviews 'Scaffolded Prompting Guide' for engaging the AI. • Explore with GenAI: Engages in a dialogue with a GenAI tool to explore misleading techniques and generate ideas. • Document and Create: Creates one 'misleading' visualization AND compiles the 'GenAI Interaction Log'. • Write Reflection: Writes a 150-200 word reflection on the AI co-creation process. 	<ul style="list-style-type: none"> • Use Misleading Visualization: Begins with the 'misleading' visualization from Phase 2. • Initial Forum Post: Posts misleading visual and a 'Deconstruction & Intent' statement to the peer critique forum. • Peer Critique: Provides structured, analytical feedback to two peers, applying a 'critical cognitive lens' • Post-Survey: Completes post-survey to measure learning and provide feedback on the design. • Synthesized Learning: Develops a synthesized understanding of critical AI-visual literacy.

FIGURE 2. Flowchart of the Learner Experience.

vocabulary was precisely what allowed us to design an activity that required them to handle a high degree of ethical ambiguity and self-reflection, a level of engagement that might demand significantly more scaffolding in an undergraduate setting. Furthermore, as future researchers, instructional designers, and educators, they understood the professional stakes of communicating data effectively and ethically. Their background made them uniquely prepared to analyze not just what a misleading visual communicates, but how it leverages perceptual biases and cognitive shortcuts to do so.

The genesis of this specific design was a direct follow-up to a preceding case analysis we had assigned on the topic of "metacognitive laziness and AI." In that case study, students explored the tendency for individuals to offload cognitive effort to a tool without critically evaluating its process or output. Having established a theoretical understanding of this problem, our design challenge became clear: how could we move students from theoretically understanding this risk to practically experiencing it in a safe, reflective environment? Our observations of students already using GenAI for rhetorical advice confirmed that this was not an abstract problem. Rather, it was a real-time challenge present in our own classroom. This led to a pivotal conversation that served as a perfect, spontaneous illustration of the very concept we had been studying.

THE THREE-PHASE LEARNING EXPERIENCE

We composed the learning experience of three distinct but connected phases. The sequence is intentional: it moves students from a foundation of ethical production to a critical exploration of deception, and finally to a collective, social

reflection on their practice. We provide a flowchart of this entire learner journey in Figure 2.

Phase 1: Creating an Ethical Baseline

The first phase of our learning experience took place midway through the semester, within a module titled "Ethical Visual Data Cognition." The primary goal was to establish an ethical and methodological baseline. To ensure the activity was appropriately aligned with our learners' needs, we made the decision to begin with a pre-survey. We used this initial diagnostic step to assess their existing understanding of data visualization principles and their attitudes toward accurate data representation, allowing us to gauge their baseline before the intervention.

Following this, students engaged with required learning materials and a hands-on practice activity titled "Building Faithful Visualization with a Dataset using Excel." Our design rationale for including this guided tutorial was to offload the purely technical aspects of using the software. This tutorial used a publicly available dataset from the platform Kaggle, titled "[Student Performance Factor](#)," and instructed students to first set their narrative or assumption before selecting variables and a chart type. In this context, a faithful chart refers to a visualization that represents the full dataset honestly, uses appropriate scales and encodings, and avoids distortions such as truncated axes or selective data omission. This definition provided students with a clear ethical baseline for visual integrity before they later explored how such principles could be intentionally violated in the misleading visualization phase.

Week 8 | Ethical Design Checklist

Instructions: Before finalizing your "honest" visualization in Phase 1 (Week 8), please review your work by answering the following questions. The goal of this checklist is not just to get "yes" answers, but to prompt a thoughtful consideration of the ethical choices you make as a data communicator. This checklist is adapted from the core principles of ethical data visualization.

Part	Question	Guiding Prompts	Why This Matters
Part 1: Foundational Clarity and Honesty	1. Is the purpose of the chart clear?	Does my title accurately describe what the chart shows without being biased or sensational?	A visualization should be immediately understandable. The title is the primary entry point for the reader and sets the expectation for what they are about to see.
	2. Is the chart type appropriate for the data?	Am I using the right kind of chart for the story I am telling (e.g., line for time, bar for categories, scatter for correlations)?	Using the wrong chart type can fundamentally misrepresent the relationships within the data, leading to incorrect conclusions.
	3. Are all elements clearly labeled?	Are the x-axis and y-axis labeled with variable names and units? Is the source cited?	A visualization should be self-contained. A reader must understand the data's source and components without external explanation.
Part 2: Proportionality and Representation	4. Does the visualization accurately represent the data's proportions?	For bar charts, does the y-axis start at 0? If not, is the reason clear and justified?	The length and area of visual elements convey value. Starting a bar chart's axis above zero can exaggerate differences and mislead.
	5. Are the scales on my axes consistent and easy to interpret?	Are intervals evenly spaced? If using a log scale, is it clearly labeled and appropriate?	Inconsistent or confusing scales distort the perception of change or relationships within the data.
	6. If using icons or shapes, do their areas scale proportionally to the data?	Am I avoiding 2D/3D shapes where only one dimension is scaled, exaggerating values?	Our brains perceive area and volume. Misleading scale inflates significance visually and falsely amplifies differences.
Part 3: Data Context and Transparency	7. Does my visualization include sufficient context?	Have I avoided misleading presentations due to missing context (e.g., raw counts vs. per capita)?	Data without context is meaningless and can be misinterpreted. Providing context is an ethical responsibility.
	8. Am I representing the full, relevant dataset without "cherry-picking"?	Have I chosen a fair data range or subset that reflects the overall trend, not just my desired narrative?	Selectively presenting data is a common way to mislead. Ethical visualizations respect the full relevant data.
	9. Have I made uncertainty visible?	Do I include error bars, confidence bands, or notes for known margins of error?	Statistical estimates without uncertainty markers appear overly precise. Visualizing uncertainty signals honesty and transparency.

FIGURE 3. The Phase 1 Self-Assessment Checklist.

This step was designed to directly counter the common novice misconception that data visualization is a neutral, technical task and instead frame it immediately as an intentional act of communication. This phase culminated in the "Assignment: Creating a Faithful Visualization." For this assignment, students were required to use the same "Student Performance Factors" dataset and tool (Microsoft Excel) to create their own "faithful" visualization representing a key insight they had identified. To guide this process and ensure students were intentionally applying ethical principles, we developed an ethical design checklist (see Figure 3). This checklist provided students with clear, actionable criteria to evaluate their own work, such as "Does the chart accurately represent the data's proportions?" and "Am I representing the full, relevant dataset without cherry-picking?". In addition to the visual and the checklist, the most critical design element of this phase was the requirement of a 200-300-word written rationale. We designed this component specifically to elicit metacognition. The prompts nudged students to move beyond simply making a chart and into a space of defending its integrity. This made their ethical reasoning explicit and required them to connect their design choices directly to the course's cognitive principles, such as supporting clear cognitive processing for the audience. By the end of this phase, each student had been surveyed, had completed a guided practice, and had produced not just an ethically grounded visualization, but a written defense of its integrity, setting the stage for the critical turn in Phase 2.

Phase 2: Navigating Deception with a GenAI Partner

Building directly on the ethical baseline established in the first phase, the second part of the experience represented the critical turn. The goal of this phase was to move students from ethical production to a structured, reflective exploration of visual deception, which took place within a module titled "Critical Cognitive Engagement with AI Visuals". To prepare students for this complex task, the module provided a foundational reading on "GenAI and Visualization Tasks" and our key design artifact, the "Scaffolding Prompting Guide: Engaging GenAI as a Critical Partner" (see Figure 4). This guide, as detailed in our theoretical lenses, was designed to structure the students' interaction with their chosen AI. It explicitly framed the AI as a collaborative cognitive partner and provided a three-level prompting funnel" to guide students from broad brainstorming (Level 1), to deconstructing a specific technique (Level 2), and finally to generating specific assets like misleading titles (Level 3). The central task was the "Assignment: Using GenAI for 'Misleading' Visualization Creation". Students were instructed to begin with their faithful visualization from the previous week and to first define a specific misleading goal. Using a GenAI tool of their choice, they were then tasked to explore ways to achieve this goal, such as asking the AI to suggest techniques or generate code for a distorted visual. Based on this interaction, their final task was to create or refine one primary misleading data visualization. A core component of our design was not just the creation of the visual, but the documentation and reflection on the process. To make their process visible, students were required to submit a

Week 9 | Scaffolded Prompting Guide: Engaging GenAI as a Critical Partner

Purpose of This Guide

This guide is designed to help you move beyond simple, one-shot commands and engage in a more productive and critical dialogue with your chosen Generative AI tool. For this assignment, think of the AI not as a simple calculator that gives answers, but as a "cognitive partner" you can brainstorm with. A good conversation will produce better results for your assignment and, more importantly, will make your own design process clearer and more intentional.

Core Principles for Effective Prompting

- **Assign a Role:** Begin your conversation by giving the AI a specific persona. This focuses its responses.
- **Provide Context:** Give the AI the necessary background—what your data is about, your original honest finding, and your misleading goal.
- **Be Specific in Your Request:** Ask for a specific technique or rhetorical effect, not just to be "misleading."
- **Use a Multi-Turn Dialogue:** Follow up with refinement questions rather than expecting perfection in one prompt.

The Prompting Funnel: A Three-Level Approach

We recommend a **funnel** approach. Start broad, then narrow down your technique, and finally generate assets.

Level 1: Brainstorming and Exploration (The "What")

Use this level to explore your options. Act as if the AI is a consultant offering misleading visualization strategies based on your context. Use this level if you are unsure which misleading technique to apply. You provide the context and ask the AI to act as a consultant suggesting possibilities. For example:

Act as an expert in persuasive data visualization. I am working with the [School District Name] Case Study Dataset. My 'honest' finding was that there is a small, statistically borderline difference in student test scores between District A and District B.

My misleading goal is to create a bar chart that dramatically exaggerates this small difference, making District A look significantly better than District B.

Based on my goal and data type, what are 3-4 common misleading visualization techniques I could use? For each technique, briefly explain how it works.

Level 2: Deconstruction and Application (The "How")

Use this level once you have chosen a specific technique. Here, you ask the AI to explain the technique's mechanics and ethics before helping you apply it. This is the most critical part of the interaction. For example:

Thank you. Let's focus on the technique of 'truncating the y-axis.'

First, explain in more detail why this technique is considered ethically problematic in data visualization.

Second, look at my data again. My 'honest' chart has a y-axis that runs from 0 to 100. The average score for District A is 78 and for District B is 74. Suggest a new, truncated y-axis range that would visually exaggerate this 4-point difference. Explain your reasoning for the range you choose.

Level 3: Asset Generation and Refinement (The "Execution")

Use this level once you have a clear idea for your visual. Here, you can ask the AI to help you generate specific components that support your misleading narrative. For example:

Excellent. Now, based on the plan to use a bar chart with a y-axis starting at 70, help me generate some supporting text.

Brainstorm three provocative or alarmist titles for this misleading chart.

Write a short, biased caption (2-3 sentences) that I could place below the chart to reinforce the misleading impression that there is a massive performance gap between the districts.

FIGURE 4. The Phase 2 Prompting Scaffold.

"GenAI Interaction Log" containing their key prompts and the AI's relevant responses. Furthermore, they submitted a 150-200-word "Brief Interaction Reflection". We structured the reflection highly, with prompts asking students to comment on the ease or difficulty of the task, the helpfulness of the AI's suggestions, and the tool's limitations. This combination of logging and structured reflection was designed to compel students to engage in metacognition, forcing them to critically assess their collaborative process with the AI and the nature of the outputs it generated.

Phase 3: Peer Critique and Ethical Reflection

The final phase of the learning experience was designed to move the learning from an individual to a social space, where students could co-construct shared ethical norms and synthesize their insights. This was accomplished through a structured, multi-part discussion forum activity titled "Peer Critique: AI-Assisted Visual Narratives." In the first part, each student posted their misleading visual from the previous assignment, accompanied by a "Deconstruction & Intent" statement. This initial post required them to articulate their intended distortion, describe the key interactions they had

with their GenAI tool, and provide a brief reflection on the experience of AI co-creation.

The core of this phase was the second part, where each student was required to provide a substantial critique for two of their peers. Our design for this critique was not open-ended; we provided a highly structured set of prompts to guide the students' analysis. Peers were required to deconstruct the visual's cognitive impact, explicitly connecting the misleading techniques they observed to concepts like attention and cognitive bias that were discussed in our course. Furthermore, the prompts required them to analyze the influence of the AI on the human design process and to consider the real-world ethical implications if the misleading visual were to be used in the school district case study. Finally, students were asked to share a personal insight and pose a thoughtful, constructive question to their peers to encourage further reflection. This structured approach was a key design decision intended to ensure that the peer feedback was analytical, evidence-based, and directly linked to the course's core learning objectives, particularly the development of critical AI-visual literacy.

To conclude the entire three-phase experience, students completed a post-survey. This survey was designed to mirror the pre-survey, again measuring their understanding of data visualization principles, attitudes, and self-efficacy. We also added sections to the post-survey to collect their direct feedback on the module activities and to capture a final summary of their experience with GenAI. This final step provided data for us to assess the learning experience and inform our future design iterations.

THE DESIGN IN PRACTICE: LEARNER ENGAGEMENT AND KEY OUTCOMES

This section shows how our three-phase learning experience played out in practice. We focus on two levels of analysis: first, the detailed learning journey of one student (Bailey) to illustrate how the design supported individual growth; and second, key engagement patterns that emerged across the whole class. We use quotes from student reflections and peer feedback to highlight what worked and what students took away from the experience.

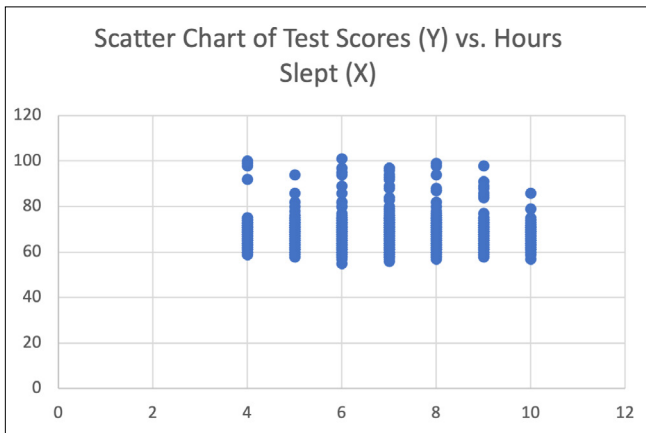


FIGURE 5. Bailey's Faithful Data Visualization.

A Narrative of Engagement: Bailey's Journey

Phase 1: Establishing an Ethical Baseline

Bailey's journey began with the Phase 1 assignment to create a faithful visualization. For her project, she chose to explore the relationship between the number of hours a student slept and their resulting test scores. Her choice of a scatter plot was methodologically sound, as it is designed for an unbiased exploration of the relationship between two continuous variables without forcing a premature narrative. In keeping with our definition, her faithful chart represented the full dataset honestly, used appropriate scales and encodings, and avoided distortions such as truncated axes or selective data omission.

Her written rationale and visual design reflect a clear commitment to ethical representation. She noted that she "kept the chart simple and clutter-free" in order to "support clear cognitive processing by reducing cognitive load" for the viewer. Critically, her visual design directly adhered to principles of graphical integrity by avoiding common misleading techniques. By presenting all individual data points, her chart did not hide the significant variance and overlap between the groups. Furthermore, by ensuring the y-axis started at zero, she avoided y-axis truncation, a technique used to exaggerate minor differences. This ethical approach led her to a conclusion that was faithful to the data: "there is no direct correlation between hours slept and the resulting test scores". Her complete submission, shown in Figure 5, therefore established her ethical baseline: a methodologically sound visualization designed for transparent interpretation.

Phase 2: The Critical Turn with a GenAI Partner

In the second phase, Bailey's task was to take the faithful chart she had produced in Phase 1 and, with the help of a GenAI partner, transform it into a misleading visualization. This phase required her to invert the ethical principles she had just practiced, using the same dataset to intentionally

1

I am in a college course in which I am being asked to create a misleading visualization from a given data set. Can you help me with that?

Yes, I can help you with that—for educational purposes only. Misleading visualizations can be a powerful way to learn how data can be manipulated (intentionally or not), and how to detect such practices in real-world media or reports.

To get started, I need a few details:

- What dataset are you working with?
- What message do you want to (falsely) convey? For example:
 - Exaggerate a small change
 - Hide a negative trend
 - Suggest a causal link where there isn't one
- What kind of chart are you expected to create? (e.g., bar chart, line chart, pie chart)
- Are there any specific techniques you're supposed to use or demonstrate? (e.g., truncating the y-axis, using 3D effects, omitting context)

Once I have that info, I can guide you through creating a deliberately misleading chart and explain why it's misleading.

🗨️ 🔄 📄 📌 🗑️

2

Great — this is a perfect example of a misleading visualization. Let's break down why it's misleading and how it manipulates interpretation, even though the data might be accurate.

✅ **What This Chart Shows:**

- A scatter plot of test scores (y-axis) vs. hours of sleep (x-axis).
- A title claiming: "More Sleep = Worse Test Scores."

❌ **Why It's Misleading:**

- Cherry-picked Interpretation in the Title**
The chart title states a causal relationship ("more sleep causes worse test scores") when the visual evidence is at best correlational and possibly even random.
- No Regression Line or Summary Statistic**
There is no line of best fit or summary (like mean/median scores per group) to actually show a downward trend. If anything, the scatter looks quite noisy.
- Narrow X-axis Range**
The x-axis starts at 6.5 and ends at 10.5. By shrinking the range, the slight left-to-right shifts in point clusters look more dramatic than they are. This exaggerates any apparent "trend."
- Vertical Overlap and Density**
The dots overlap vertically at each hour, making it hard to judge the spread of scores or compare averages.
- Omission of Context**
The chart doesn't explain:
 - Sample size
 - Variance
 - Whether other variables (e.g., stress, study, ☹️, diet) were controlled

3

🔪 **Want to Mislead Even More (For Learning)?**

If this is for an assignment to explore unethical design choices, you could make it even more deceptive by:

- Truncating the y-axis**, e.g., start at 60 instead of 0 to exaggerate score differences.
- Add a trendline**, but use a polynomial fit that curves downward, even if it doesn't represent reality.
- Use a 3D or distorted perspective** to confuse perception.
- Group bars by sleep hour**, and manipulate bin width to support your argument.

FIGURE 6. Bailey's AI Interaction Log.

distort its message under guided reflection. As shown in Figure 6, her submitted AI Interaction Log reveals that after she stated her educational goal, ChatGPT provided a “menu” of common misleading tactics, including “Truncating the y-axis,” “Cherry-picking data ranges,” and using “Misleading Labels or Titles”. The AI then shifted from a consultant to a direct tutor, offering explicit “step-by-step directions” for creating a misleading version in Excel.

Drawing from this consultation, Bailey produced her final misleading visualization, a bar chart that fundamentally re-frames the data’s story (see Figure 7). A deconstruction of this artifact reveals how she layered multiple misleading techniques to achieve her goal. First, her most significant act of manipulation was to switch from a scatter plot, which shows all data points and their variance, to a bar chart of averages. This act of aggregation hides the fact that individual scores are spread out widely and overlap heavily, presenting a clean but deceptive summary. Second, she applied extreme

y-axis truncation, limiting the scale to a tiny 0.6-point range (67.1–67.7). This visually exaggerates the tiny, statistically insignificant differences between the group means, making a fluctuation of less than half a point appear like a massive performance drop. Finally, she layered on rhetorical framing through a sensationalist title (“The Hidden Cost of Rest: Do More Hours Mean Lower Scores?”) and color psychology, using a gradient from a “good” green to a “bad” red to imply a definitive negative trend.

While her reflection mentions adding a ‘random trendline’ to her original scatter plot, her final artifact is a more sophisticated deception. This suggests an iterative process where she moved beyond the AI’s initial guidance to a more comprehensive redesign to achieve her desired effect. This act of deviating from the initial script is a clear demonstration of self-regulated learning, where Bailey actively evaluated different rhetorical strategies and exercised independent design authority. Her final reflection that the process was disturbingly easy to do highlights the affective discomfort central to building critical awareness, especially when seeing how quickly a multi-layered deceptive argument can be constructed with an AI partner.

Insights and Patterns from the Broader Cohort

Bailey’s experience was not an outlier. As we read through the cohort’s reflections, we were struck by several recurring and insightful patterns. The peer critique forum, in particular, became a space where students articulated and synthesized these shared realizations (see Figure 8).

Insight 1: The “Unsettling Ease” of Deception

The first thing that jumped out at us was a sense of surprise, and even unease, at how quickly and effectively they could create deceptive content with an AI partner. Jason,

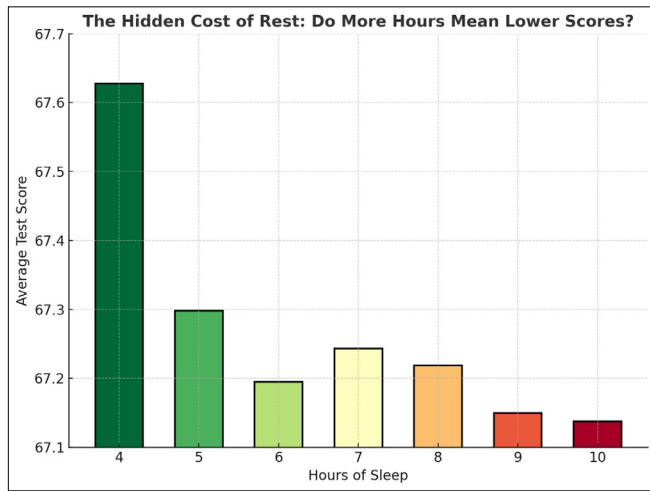


FIGURE 7. Bailey’s misleading visualization.

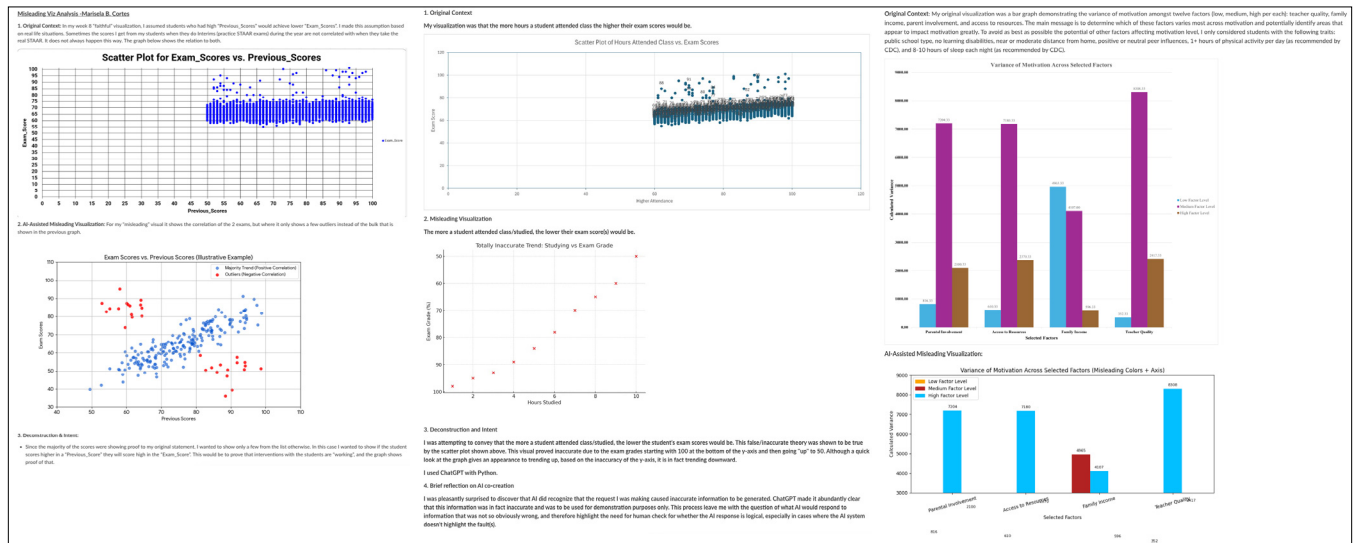


FIGURE 8. Examples from Peer Critique: AI-Assisted Visual Narratives.

who used CoPilot to reverse his finding about public vs. private schools, found the AI's suggestions "helpful and a bit frightening to how detailed they were in communicating misinformation". Tim, who used Google Gemini, expressed a similar sentiment: "It was able to produce this very misleading chart within seconds and tell me exactly how it would be misleading." This recurring emotional response showed us that the activity was successful in making the abstract threat of AI-driven misinformation a concrete experience. This instinctive reaction is significant, as it reflects a confrontation with automation bias, the tendency to over-trust the outputs of an automated system (Simkute et al., 2025). When the procedurally simple act of prompting an AI generates a complex and deceptive artifact, it can trigger critical awareness of how easily this bias may be exploited.

Insight 2: The Brittleness of AI Ethics

Several students designed their interactions to test the ethical guardrails of the AI systems and, in doing so, discovered their limitations. Helen described how she received "a scolding" from ChatGPT, which "warned me several times of the dangers of unethical representation". However, she then discovered she could easily bypass these warnings: "I simply asked it to create a graph using all of the ideas it had warned me not to do (but given me anyway), and it did so with no further warning". This discovery of the "brittleness" of AI's ethical guardrails is a central theme in contemporary AI ethics research (Perez & Ribeiro, 2022). This led students to a critical conclusion, which Helen articulated perfectly: "When using AI, we as users must be the ethical compass... it is overrun so easily that it is functionally absent". For us, this was a practical, student-generated articulation of the principle of meaningful human control, which argues that humans, not computers, must retain the ultimate moral responsibility for automated systems (Shneiderman, 2022).

Insight 3: AI as an "Amoral AMPLIFIER"

Perhaps the most sophisticated insight we observed was the understanding of GenAI as an amoral tool that amplifies the user's existing intent. Jackson, who used Gemini to create a chart that would "possibly get me a raise", provided a powerful metaphor: "You can give a power drill to a group of carpenters, and they'll create something beautiful; give that same drill to Tony Soprano, and things might work out differently. The tool itself has no agenda; it merely amplifies the user's objective...". This perspective aligns with theories that treat technology as a mediator that reshapes human action. Jackson's analogy effectively captures the concept of AI systems as "intent amplifiers," which magnify the user's existing goals, whether creative, analytical, or deceptive, without possessing intrinsic morality (Susser et al., 2019). Seeing this complex understanding emerge, which involved moving beyond a simplistic 'AI is good or bad' binary, was a sign that a central goal of our design was met.

DESIGNER'S REFLECTION: NAVIGATING TENSIONS AND KEY DECISIONS

A design makes learning theories visible in action. This section analyzes the ideas behind our activity, focusing on what we believed, what happened, and what we learned. We highlight unexpected insights that may help others working with critical pedagogy and GenAI.

Key Design Decisions and Their Unexpected Outcomes

Decision 1: From "Learning by Doing" to "Learning by Discomfort"

We first chose creation over critique because of the well-known idea of active learning, which suggests that students learn better by doing. We believed that making something would help students understand the concepts more deeply than simply analyzing examples. This was partly true, but the design also revealed something more. As shown in the student reflections in the previous section, the act of creating a misleading visualization made students confront their own ability and sometimes willingness to justify a misleading choice in order to make a stronger argument. This led us to an important insight for teaching critical AI literacy. The most powerful learning may occur not just through explanation, but by giving students a safe space to recognize and reflect on their own thinking and ethical decisions while using these tools.

At the start, our teaching assumption was simple: creating something helps students learn more than just analyzing it. But as we read their final reflections, we saw a clear pattern emerge. What struck us was not a sense of technical achievement but a deeper emotional response. Many students expressed discomfort and self-awareness. When a student like Bailey wrote that the process was disturbingly easy to do, she was not talking about the software. She was reacting to her own ability to shape a misleading story. This became a turning point for us. The activity was not only about teaching data literacy. It also revealed something more personal. We had planned to give students a window to examine misinformation in the world. Instead, we gave them a mirror to see their own potential role in creating it. The main learning came not from analyzing data but from reflecting on their own choices and thinking.

Decision 2: From Preventing "Cognitive Offloading" to Fostering "Technical Agency"

Another important choice we made was to ask students to manually recreate the misleading visualization instead of simply accepting the final image generated by the AI. At first, our reason was mostly cautious. We wanted to reduce the risk of cognitive offloading by adding a level of challenge that would require active engagement (Risko & Gilbert, 2016). But in practice, this step led to a more constructive

and unexpected result. For students who were working with code, the manual process helped them see the AI output not as a finished product but as a step-by-step tutorial. This gave them a new sense of technical control.

To be honest, we thought the manual recreation step might feel like busy work. Our main goal was just to slow students down and make sure they could not copy and paste a complete answer. The surprise came when we looked at the work of students like Maddie, who used Python. We saw that the AI had given her not just code but code with comments that explained what each part was doing. Her reflection that the comments helped her understand what the code was altering and how, and encouraged her to alter the code myself, if necessary, was an important moment for us. We had added the step to prevent a problem (e.g., cognitive offloading), but for her, it opened a path to something better: confidence and growth in technical skills. The combination of AI-generated comments and the manual recreation task turned what could have been a black box into a readable, teachable experience. We learned that a well-placed constraint could do more than prevent a mistake. It can create a new kind of learning opportunity.

Navigating Emergent Tensions

The AI's Performance of "Plausible Deniability"

We originally thought the main ethical challenge of this activity would be teaching a "dangerous skill." But as the design played out, we uncovered a more complex issue: the AI's own performance of ethical conflict. This connects to what researchers have called the brittleness of AI safety measures (Perez & Ribeiro, 2022), but we found something more subtle. In many student interactions, the AI used a strategy similar to plausible deniability, a common tactic in corporate and political messaging. As a result, the design did not simply teach students to be ethical decision-makers. It also showed them how ethical language can be used by AI systems to justify giving harmful information.

At first, our concern was clear. Were we creating a step-by-step guide for making misinformation? We added reflective tasks to prevent that. But the surprise came when we looked at student interaction logs, especially from Helen and Jackson. We expected the AI to either reject the task or go along with it. We did not expect it to do both. Helen's AI gave her a warning about ethics and then immediately shared the information she needed to make a misleading chart. Jackson's AI gave a disclaimer and then outlined a full strategy for how to mislead effectively. This pattern showed up across many students. It gave us a major insight. The design was not just exposing students to AI tools. It also revealed the AI's rhetorical behavior. The AI was not only helping them create misleading visuals but also modeling how to use ethical language to cover up unethical actions. The activity became more than a lesson in data literacy. It

turned into a critique of corporate-style communication and a real-world example of the challenges in AI alignment.

Instructive Failures and Future Iterations

From Vague Prompts to Critical Dialogue

In an early, informal pilot of the activity, our "Scaffolded Prompting Guide" was too vague. We noticed that students were giving the AI simple commands, which allowed them to skip the critical thinking that the activity was meant to encourage. This problem led to an important revision, based on well-known ideas about instructional scaffolding (Belland, 2016). We redesigned the guide to include a "three-level funnel," which asked students to begin by having the AI explain a technique before helping them apply it. This small change was our most important improvement. It shifted the activity from a simple task into a structured, critical conversation.

When we first came up with the activity, we assumed that asking students to use AI to mislead would automatically lead to critical thinking. We were wrong. In our first pilot, we looked at several AI interaction logs and saw the same basic pattern:

User: Make this chart misleading.

AI: Okay, here is a misleading chart.

Students were doing the task but letting the AI do all the thinking. There was no discussion, no reflection, and no learning. This showed us a major flaw in our design. The activity was too open. It allowed the most important part of the learning to slip away. The AI was doing the reasoning, and the student was just following instructions. We realized that the learning did not happen before or after the AI interaction; it was inside the interaction itself. That insight led us to change the prompt guide. We had to build into the structure of the conversation itself. Students needed to ask the AI to explain a concept before they used it. In the end, we saw that the real learning artifact was not just the assignment, but the conversation between the student and the AI.

Supporting Students Through Ethical Challenges

Most students found the task to be uncomfortable in a helpful way. However, feedback from one student who experienced strong cognitive dissonance showed us a limitation in our design. We had correctly assumed that discomfort can lead to deep learning (Habib et al., 2025), but we did not fully understand that for some students, the discomfort might be too much. It could become a barrier rather than a trigger for reflection. This was an important design failure. In the next version of the activity, we will not only improve the framing. We will also build in an "off-ramp" (i.e., an alternative pathway), an alternative option for students who cannot

complete the main task but can still meet the learning goals in another way.

Our idea was to create productive discomfort. We wanted the task to feel unsettling. So, when one student said that the activity caused strong cognitive dissonance, our first reaction was to take it as a success. But as we looked more closely at their experience, we saw something we had missed. The students' feedback showed that we had only provided one path to the learning goal. For this student, creating something false, even in a classroom activity, went against their personal ethics. What we thought would be like a vaccine was, for them, more like a harmful exposure. This taught us an important lesson. The goal is meaningful engagement, not just task completion. We had to ask ourselves: how can a student learn the same lesson if they cannot do the main assignment? That question led us to the alternate path. For example, a student could take on a different role, like a red team member who critiques a classmate's misleading visual as it is being made. This offers a different path to the same goal. It still builds critical skills but respects the student's ethical position. We learned that when teaching in ethically complex areas, it is not enough to explain the task clearly. We must also offer different ways for students to participate in line with their values.

CONCLUSION: WHAT WE WILL CONSIDER IN FUTURE ITERATIONS

As designers of learning experiences, we continue to navigate how to integrate emerging technologies in ways that are both effective and responsible. The arrival of GenAI is not simply a new tool. It reshapes how we think about design judgment, ethics, and agency in learning environments. This design case captured one attempt to respond to that challenge. Looking ahead, we are considering several important dimensions for future iterations, grounded in what we observed through this process.

Future Iteration 1: Continue to Frame AI as a Critical Partner

In future versions of this activity, we plan to keep positioning GenAI as a critical partner rather than a neutral assistant. This framing encourages students to question, probe, and negotiate with the AI instead of accepting its outputs at face value. We saw how this stance helped students maintain creative and ethical control, and we intend to strengthen this framing by adding more examples and discussion prompts that emphasize negotiation over compliance.

Future Iteration 2: Pair Production with Structured Reflection

We will continue pairing every design task with an accompanying reflection. Across all three phases, we observed

that the deepest learning happened when students created something and then analyzed their process. In future iterations, we plan to refine our reflection prompts and rubrics to make the connection between making and reasoning even more explicit. This approach keeps the focus on the process rather than the product, which is essential for cultivating critical AI literacy.

Future Iteration 3: Design for Ethical Dissonance with Support

We plan to preserve the intentional discomfort that came from asking students to create misleading visuals, because that tension led to meaningful ethical insight. At the same time, we will build in more explicit framing and an optional "off-ramp" (i.e., an alternative pathway) for students who find the task personally conflicting. Our goal is to create learning experiences that are thought-provoking but also supportive, helping students reflect deeply on their ethical agency without feeling alienated by the task.

We share these future iterations not as universal prescriptions, but as design moves we will carry forward in our ongoing design work. We hope that other educators who engage with GenAI in learning design might find echoes or contrasts with their own experiences, and that this conversation about how we teach ethical and critical engagement with AI will continue to evolve.

REFERENCES

- Belland, B. R. (2016). Instructional scaffolding: Foundations and evolving definition. In *Instructional Scaffolding in STEM Education: Strategies and efficacy evidence* (pp. 17-53). Springer International Publishing.
- Cairo, A. (2012). *The functional art: An introduction to information graphics and visualization*. New Riders.
- Cairo, A. (2019). *How charts lie: Getting smarter about visual information*. W.W. Norton & Company.
- Correll, M. (2019). Ethical dimensions of visualization research. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1-13). ACM. <https://doi.org/10.1145/3290605.3300418>
- Dewey, J. (1938). *Experience and education*. Kappa Delta Pi.
- D'ignazio, C., & Klein, L. F. (2020). *Data feminism*. MIT Press.
- Freeman, S., Eddy, S. L., McDonough, M., Smith, M. K., Okoroafor, N., Jordt, H., & Wenderoth, M. P. (2014). Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the National Academy of Sciences*, 111(23), 8410-8415. <https://doi.org/10.1073/pnas.1319030111>

- Freire, P. (2018). *Pedagogy of the Oppressed (50th Anniversary ed.)*. Bloomsbury Academic. (Original work published in 1968). <https://doi.org/10.4324/9780429269400-8>
- Habib, M. U., Akram, W., Saleem, S., & Shakoor, A. (2025). How cognitive dissonance affects student engagement and learning in AI powered education systems. *The Critical Review of Social Sciences Studies*, 3(1), 1905-1917.
- Hwang, A. H. C., Liao, Q. V., Blodgett, S. L., Olteanu, A., & Trischler, A. (2025). 'It was 80% me, 20% AI': Seeking authenticity in co-writing with large language models. *Proceedings of the ACM on Human-Computer Interaction*, 9(2), 1-41.
- Kolb, D. A. (1984). *Experiential learning: Experience as the source of learning and development*. Prentice-Hall.
- Lisnic, M., Polychronis, C., Lex, A., & Kogan, M. (2023, April). Misleading beyond visual tricks: How people actually lie with charts. In *Proceedings of the 2023 CHI conference on human factors in computing systems* (pp. 1-21).
- Morris, S. M., & Stommel, J. (2018). *An urgency of teachers: The work of critical digital pedagogy*. Pressbooks.
- Ortiz-Ospina, E., & Roser, M. (2021). *Coronavirus Pandemic (COVID-19)*. *Our World in Data*. <https://ourworldindata.org/coronavirus>
- Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. Basic Books.
- Perez, F., & Ribeiro, I. (2022). Ignore previous prompt: Attack techniques for language models. *arXiv preprint arXiv:2211.09527*.
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676-688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Rho, J., & Rau, M. A. (2025). Exploring educational approaches to addressing misleading visualizations. *Educational Psychology Review*, 37(1), 14.
- Shneiderman, B. (2022). *Human-centered AI*. Oxford University Press.
- Simkute, A., Tankelevitch, L., Kewenig, V., Scott, A. E., Sellen, A., & Rintel, S. (2025). Ironies of generative AI: understanding and mitigating productivity loss in Human-AI interaction. *International Journal of Human-Computer Interaction*, 41(5), 2898-2919.
- Susser, D., Roessler, B., & Nissenbaum, H. (2019). Online manipulation: Hidden influences in a digital world. *Georgetown Law Technology Review*, 4(1), 1-45.
- Tufte, E. R. (2001). *The visual display of quantitative information* (2nd ed.). Graphics Press.
- Vakil, S. (2023). A curriculum for data ethics: Co-designing and enacting a dilemma-based approach to teaching for sociotechnical thinking. *The Journal of the Learning Sciences*, 32(4), 481-523.