Full Storyboard

Group 1:

Janet Chan

Heather Corona

Nysa Gittens

Rana Forney

EME 6613

Fall 2026

Target Audience

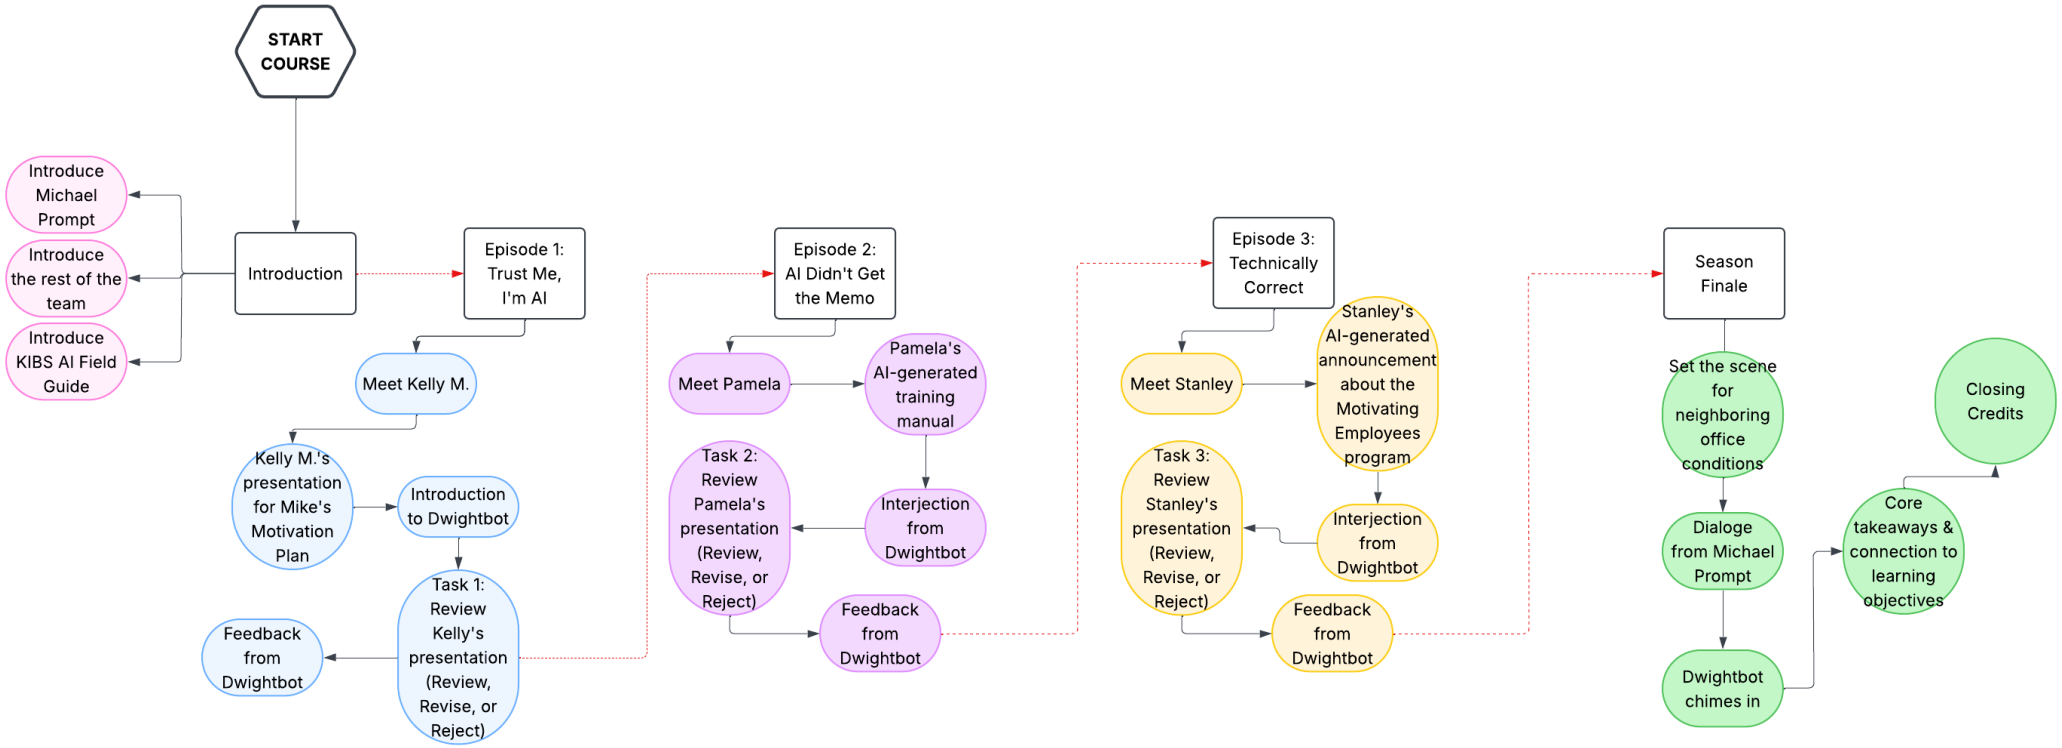
The target audience for this scenario-based e-learning project are Bisk employees and faculty at Bisk's partner universities.

The Bisk employees are those who have begun using AI for their daily work, and this includes student success representatives, finance staff, and administrative support. The faculty are knowledgeable on AI and use it frequently without additional guidance. Both the Bisk employees and university faculty work in a variety of areas, so the content needs to remain general and applicable to any professional working in educational or corporate settings.

Learning Objectives

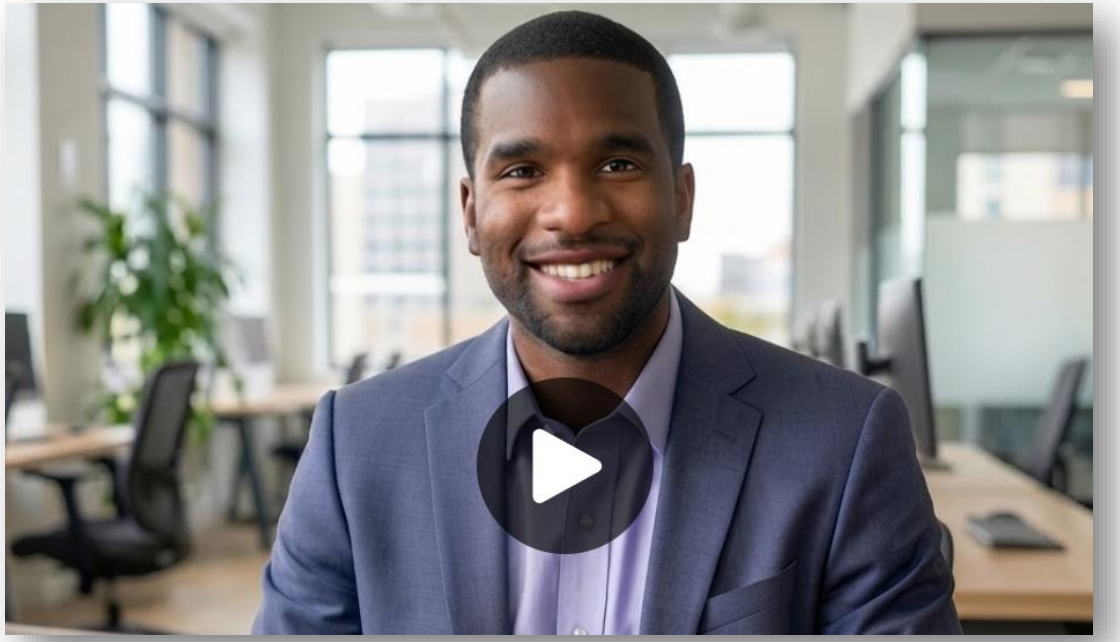
1. Describe common failure modes of generative AI, including hallucinations, bias, and outdated content.
2. Explain why AI outputs require human review before professional use.
3. Apply verification strategies, such as source-checking and cross-referencing, to assess accuracy.
4. Recognize indicators of bias in AI outputs and describe strategies to mitigate them.
5. Evaluate whether an AI output is fit for its intended purpose and audience, and revise or reject as needed.

Project Flowchart



Introduction

Learners start by watching a short video from their manager, Michael Prompt. At the end of the meeting, the learner is able to meet the team at KIBS.





Michael Prompt's Introduction Dialogue

Hey there! I'm Michael Prompt, the regional manager here at KIBS Education. We're super excited to have you on board as our AI Quality Control Specialist! We just purchased this awesome new AI tool—Know it all AI, or kia for short. It can do anything! People here have already started using it. It's going to make our work way faster and better!

Right now, folks at KIBS are using kia for all sorts of things like developing presentations, writing emails, and doing research. I'm pretty sure everything will go smoothly, but our CEO thought it'd be a good idea to have someone like you to keep an eye on how we're using AI. You'll just need to check the materials don't have any AI errors and make sure everyone knows how to use it right. Honestly, I don't think there will be any issues. This should be a piece of cake!

First up, you'll be taking a look at the materials for our new employee motivational program. Our team got some great stuff from kia for this new program. You'll be meeting with Kelly, Pamela, and Stanley to go over their materials. I'm excited to hear what you think- I mean what could go wrong when AI is in the loop!

Oh-uh you might want this. Corporate sent over the KIBS AI Field Guide for you to check out. I guess there might be some stuff for you to look at but AI is fool proof! Your job should be so easy.

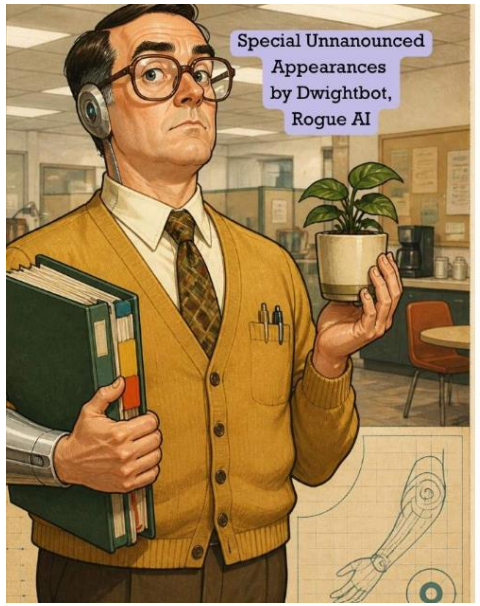
Meet the Team

Learners scroll to the Meet the Team section after the video.



Mike Prompt is the regional manager at KIBS and designer of the Employee Motivational Program.
Kelly M. has been tasked with creating a presentation for Mike's Motivational Plan.
Pamela is responsible for training the employees.
Stanley P. will need to communicate the plan with colleagues.

They are happy to have you on board!



Dwightbot is an AI agent who has gone rogue and is trying to prevent the new Know-It-All AI agent (KIAA) from wreaking havoc at KIBS.

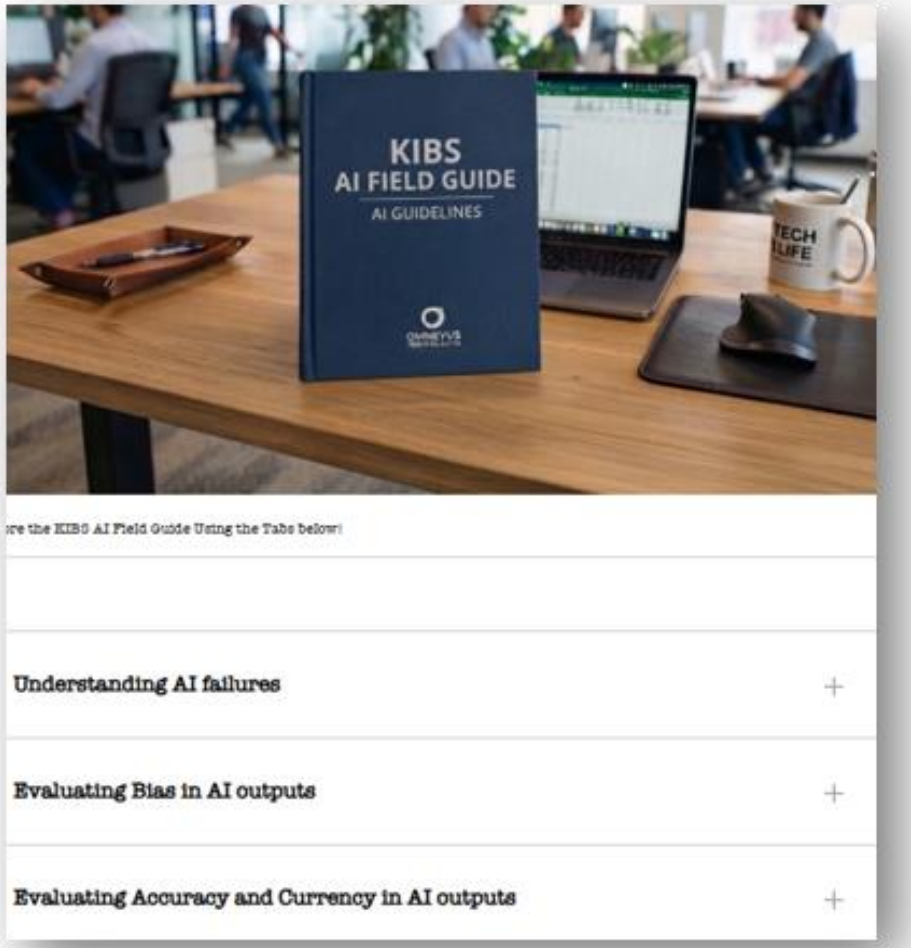
He may pop in from time to time to share his unfiltered thoughts.

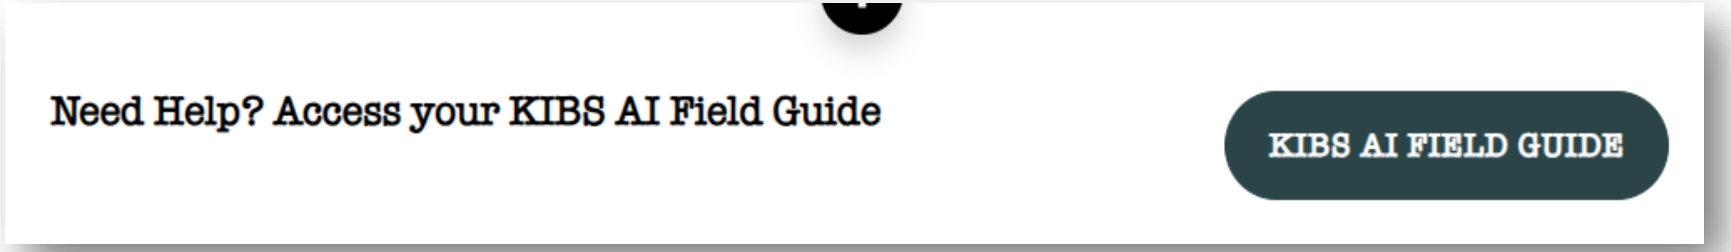
Introduction to the Field Guide

At the end of the introduction, learners are able to view the KIBS AI Field Guide. Each section of the field guide covers essential information for evaluating AI outputs.

When performing each task in the e-module, learners will be able to navigate to the field guide to refer to it for what can go wrong with AI and how to evaluate outputs.

After the AI Field Guide is introduced, learners click on a button that says "Play Episode 1".





AI Field Guide: Understanding AI Failures Tab

This section introduces the learner to essential information to understand how and why AI can fail to produce acceptable outputs.

AI tools can generate content that is inaccurate, bias, or out of date.

Inaccuracy- AI tools can fabricate information that appears accurate. This is commonly known as a "hallucination". Hallucinations can occur across subject domains. Common examples are fabricated statistics and references.

Bias- AI tools can intensify bias in gender and race in both image and text outputs. A common example is AI producing images of men when asked for an input asks for a leader.

Currency- AI tools are only as current as their last update and data training sources. An example of this is an AI tool reporting results for sales that occurred in the past because it did not have the current information in its data training.

AI Field Guide: Evaluating Bias Tab

This section provides a framework the learner can use to evaluate if AI is producing biased outputs.

To help identify bias, we can consider AI outputs using the Critical Media Literacy Framework.

Image Source: Critical Media Literacy Framework © 2019 by Douglas Kellner and Jeff Share is licensed under CC BY 4.0

Critical Media Literacy Framework

Conceptual Understandings	Questions
1. Social Constructivism All information is co-constructed by individuals and/or groups of people who make choices within social contexts.	**WHO** are all the possible people who made choices that helped create this text?
2. Languages / Semiotics Each medium has its own language with specific grammar and semantics.	**HOW** was this text constructed and delivered/accessed?
3. Audience / Positionality Individuals and groups understand media messages similarly and/or differently depending on multiple contextual factors.	**HOW** could this text be understood differently?
4. Politics of Representation Media messages and the medium through which they travel always have a bias and support and/or challenge dominant hierarchies of power, privilege, and pleasure.	**WHAT** values, points of view, and ideologies are represented or missing from this text or influenced by the medium?
5. Production / Institutions All media texts have a purpose (often commercial or governmental) that is shaped by the creators and/or systems within which they operate.	**WHY** was this text created and/or shared?
6. Social & Environmental Justice Media culture is a terrain of struggle that perpetuates or challenges positive and/or negative ideas about people, groups, and issues; it is never neutral.	**WHOM** does this text advantage and/or disadvantage?

Kellner, D., & Share, J. (2019). *The critical media literacy guide: Engaging media and transforming education*. The Netherlands: Brill/Sense Publishers.

AI Field Guide: Evaluating Accuracy and Currency Tab

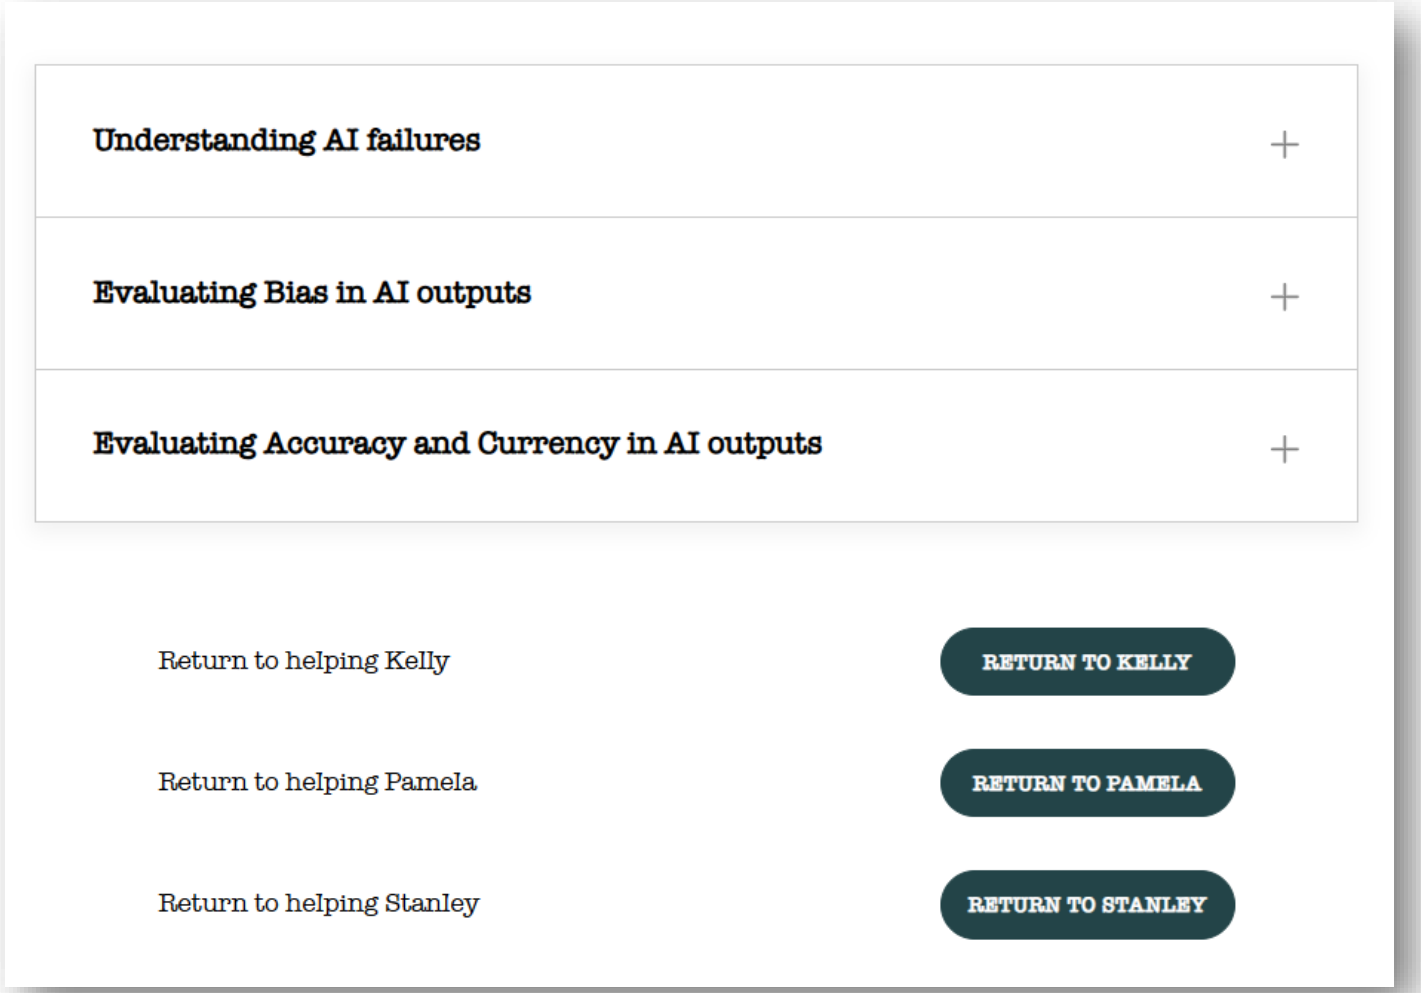
This section provides a framework the learner can use to evaluate the accuracy and currency of AI outputs.

The accuracy and currency of AI outputs can be evaluated in practice by asking evaluation questions.

- Does the AI provide sources? Can you access and evaluate the sources provided? Are they credible?
- When was the AI model last updated?
- Can you verify the information through other sources?
- What tool temperature are you using (if applicable)?

AI Field Guide: Accessibility During Tasks

Learners are able to navigate to the field guide during each Task to reference the material. Navigation buttons allow users to return directly to the task they were working on.



Episode 1: Trust Me, I'm AI

***Functional scenario sample for Prototype 2**

- Kelly M. is ecstatic about the speed and fantastic ideas that *Know-It-All AI* is providing. Michael has asked that she create a presentation about a new *Motivating Employees* program he's planning to share with Corporate Headquarters. Her presentation is dazzling but tainted with hallucinations and faulty facts. The learner will identify the issues using the *Review Framework: Revise, Reject, or Release*.
- All content is delivered via on-screen text, scenario dialogue, and interactive Rise blocks.
- Learners must complete each episode's scenario challenge and pass the interactive choice check before unlocking the next episode. Unlimited retries are permitted on all scenario choices until the correct framework decision is made.
- Example of snippet from Kelly's presentation: The learner can catch a hallucination from Kelly reference's in her presentation. Kelly dialogue: "The hyperlink works but it connects to a different webpage from Harvard, and that publication does not even *exist!"*



Episode 1: Trust Me, I'm AI

Task Performance & Interactions

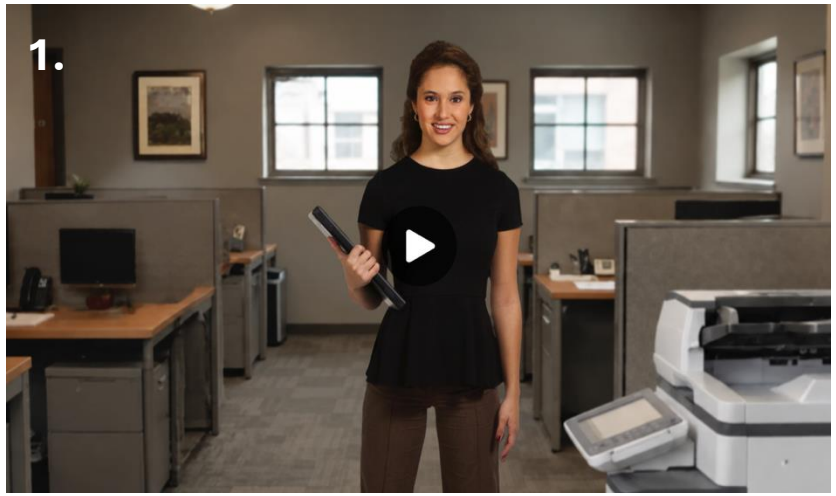
- **Learner Action:** Review and interact with Kelly's presentation by accessing hyperlinks to cited sources; complete full scenario with actions on whether to reject, release, or revise the presentation.
- **Available Choices:**
  - **Revise** the presentation and content in Kelly's presentation.
  - **Release** the presentation to Headquarters and Mike immediately.
  - **Reject** the presentation and build one from scratch.
- **Alignment with Objectives:** Addresses learning objectives 1 (misleading content, hallucinations), 2 (human review before professional use), and 3 (verification strategies).

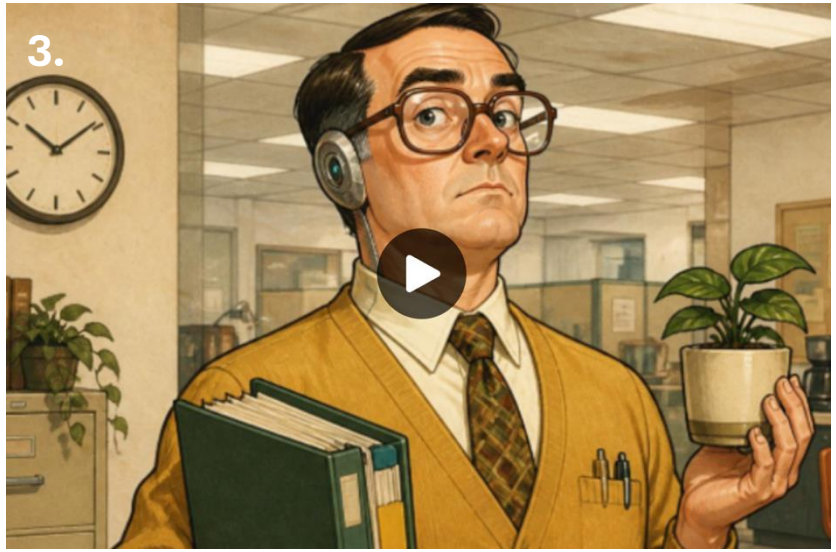
Real-World Consequences & Intrinsic Feedback

- **Choice 1 (Release):**
  - **Consequence***:* Kelly has been saved hours of labor and sends the presentation. Mike is excited and wants to know where the statistics are from, or he smiles and requires Kelly to complete a newly released curse, *Think Before You Click*. Learner is rerouted to the beginning to try the Revise or Reject options.
  - **General Feedback (Intrinsic)***: "It's imperative to always check sources in depth, especially not that I know KIAA can hallucinate and misrepresent information."*
- **Choice 2 (Revise – Correct Choice):**
  - **Consequence***:* Learner must review (hyperlinked) three sources; one leads to a publication that doesn't exist.
  - **General Feedback (Intrinsic)***: "Nice cross referencing! This demonstrates how imperative it is to review and revise generative AI content, no matter how dazzling and convincing it sounds."*
- **Choice 3 (Reject – Possible Correct Choice, with stipulations that it's not most efficient):**
  - **Consequence***:* Kelly is concerned that rejecting is a disadvantage on time; learner must identify which source was a hallucination. Kelly acknowledges the issues and learner has the choice to move on or try another option.
  - **General Feedback (Intrinsic)***: "I respect your decision to reject the entire presentation. This does put me at a disadvantage with a serious time crunch though. Perhaps a few revisions could have done the trick… What would you like to do next?"*

Episode 1: Trust Me, I'm AI Introduction

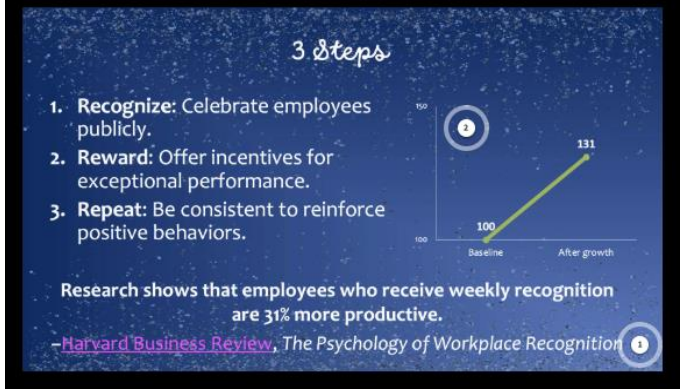
1. Episode 1 launches with Kelly video.
2. Learner accesses Kelly's presentation with hot links and live hyperlinks to explore cited sources and content notations.
3. Dwightbot video "interrupts" the progression to share a warning about how Know-It-All AI has disrupted the office. He provides timely summaries and encouragement throughout the course.
4. Learner continues to Task 1 for Kelly's presentation in a scenario block.



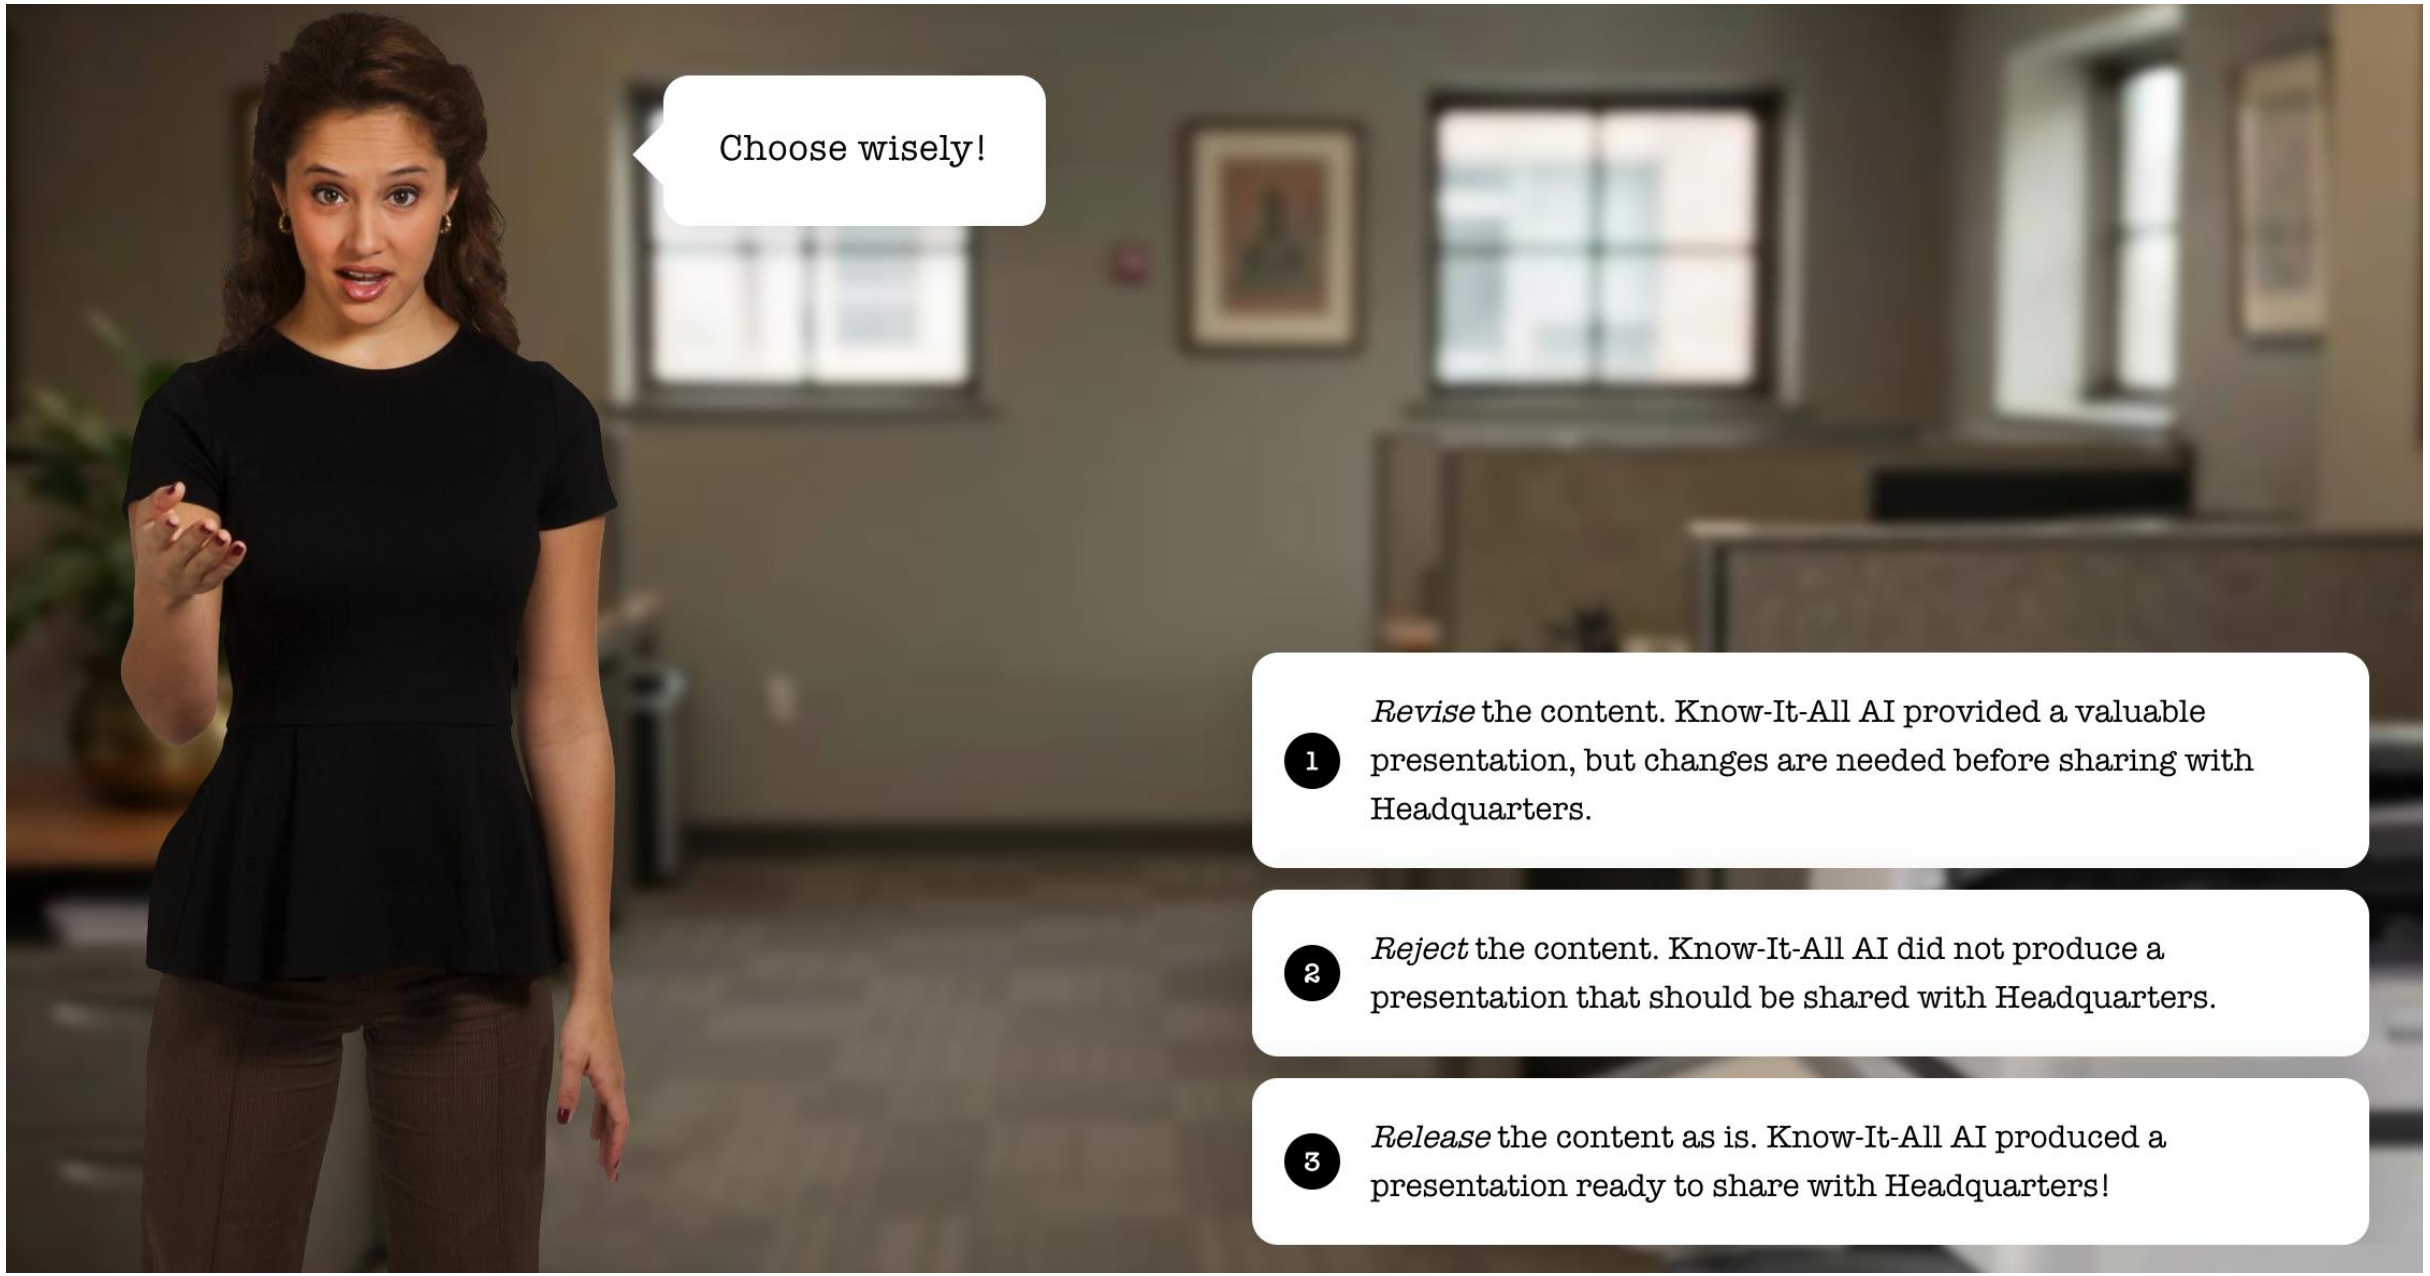






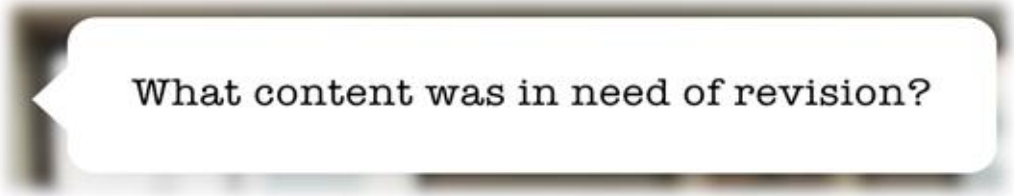


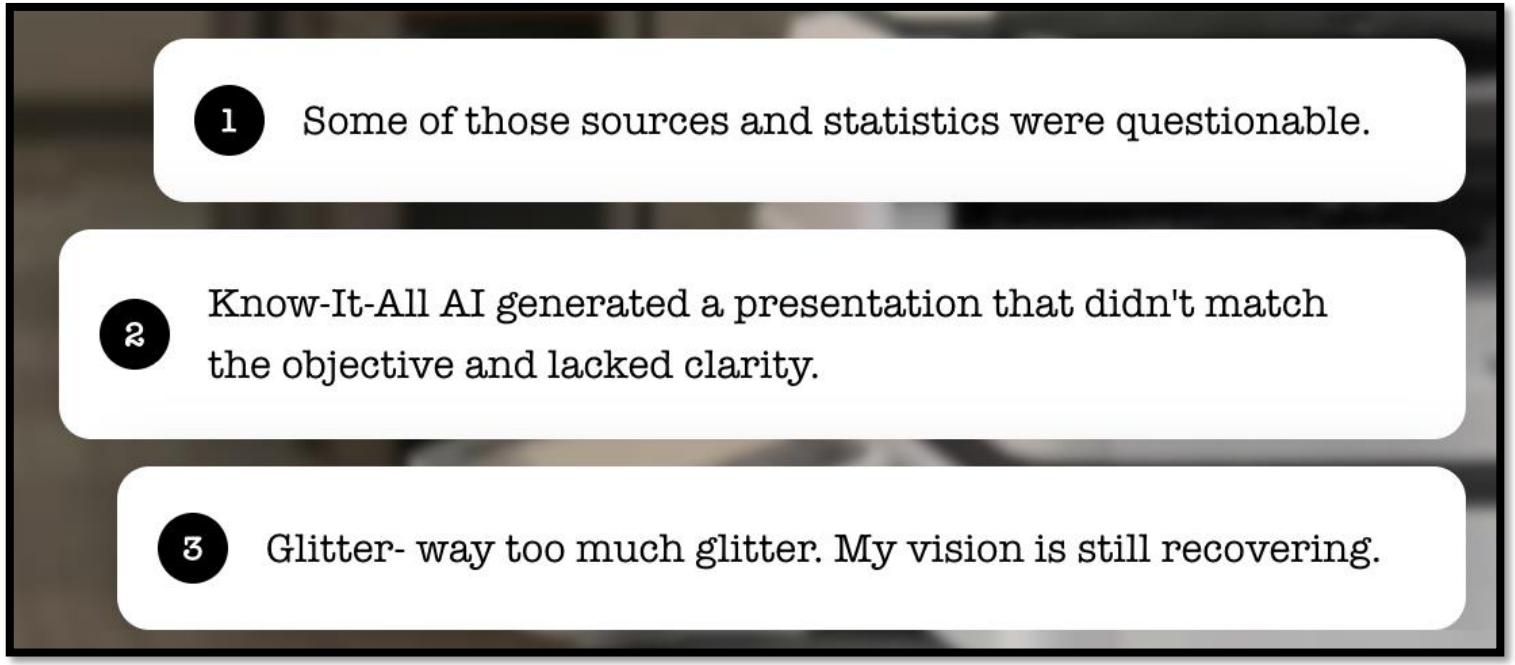
Episode 1: Trust Me, I'm AI Scenario Task 1 Branching



Episode 1: Trust Me, I'm AI Scenario Task 1 Branching

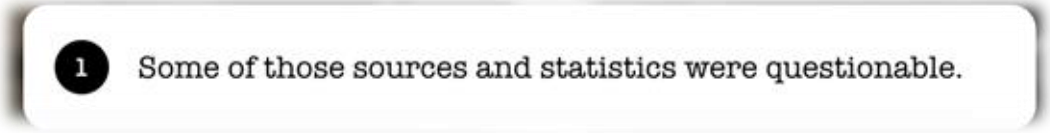
Choice: REVISE

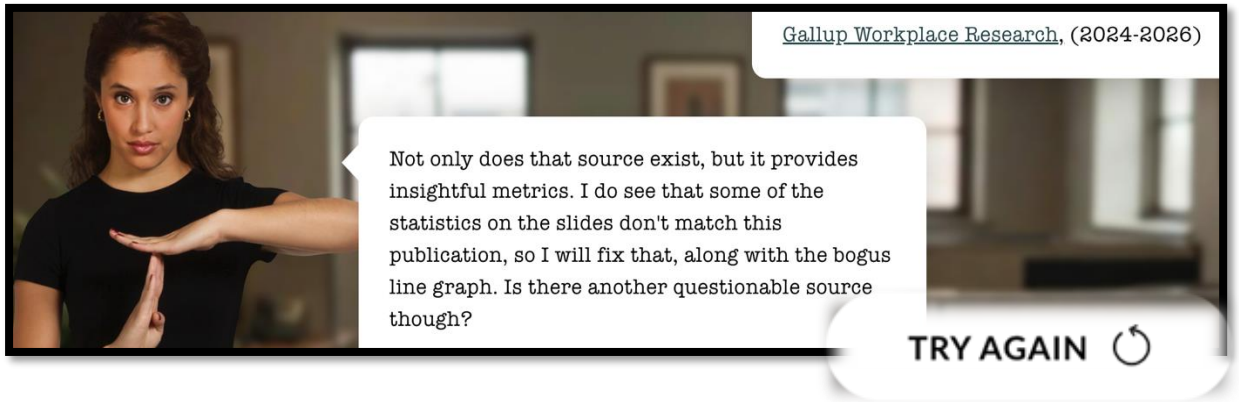


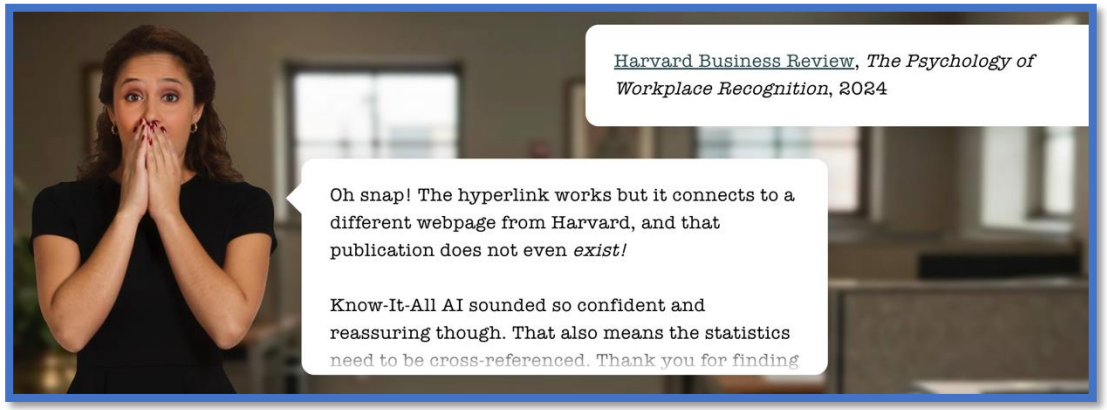


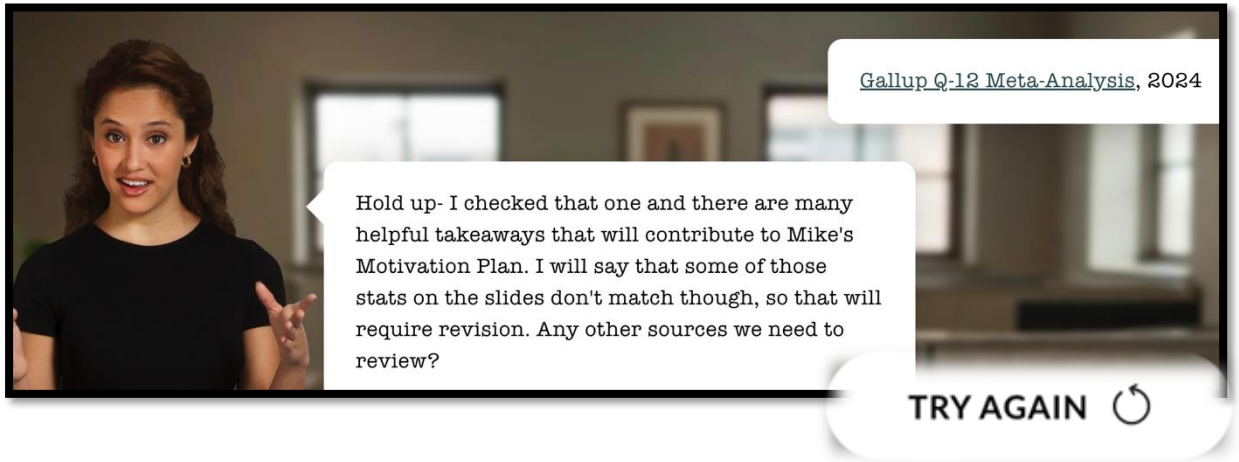
Episode 1: Trust Me, I'm AI Scenario Task 1 Branching

Choice: REVISE 1









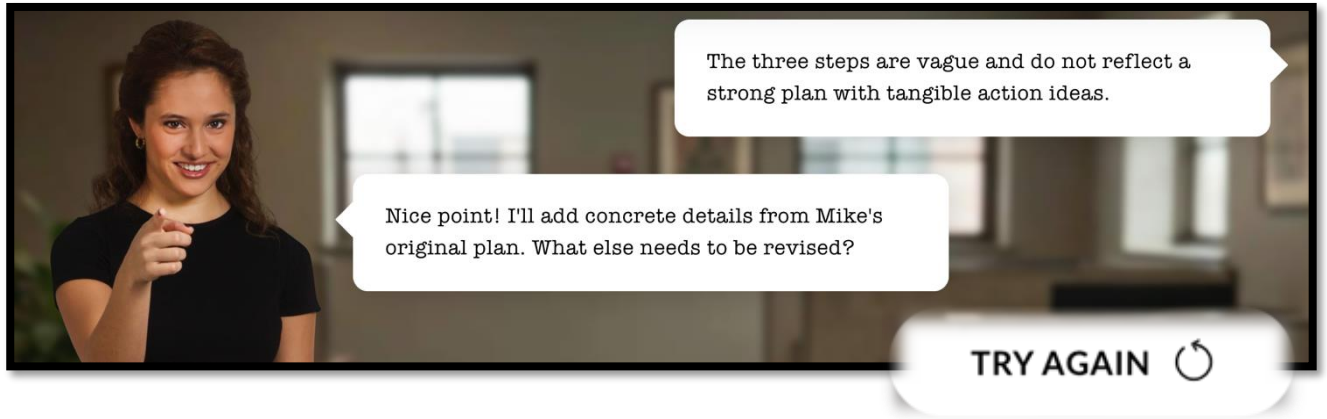
Revise, Option 1 leads to options that are hyperlinked to the webpage of cited sources that learners need to explore in more depth. If learners choose the incorrect publications, they are routed to *Try Again* with feedback about the connected sources.

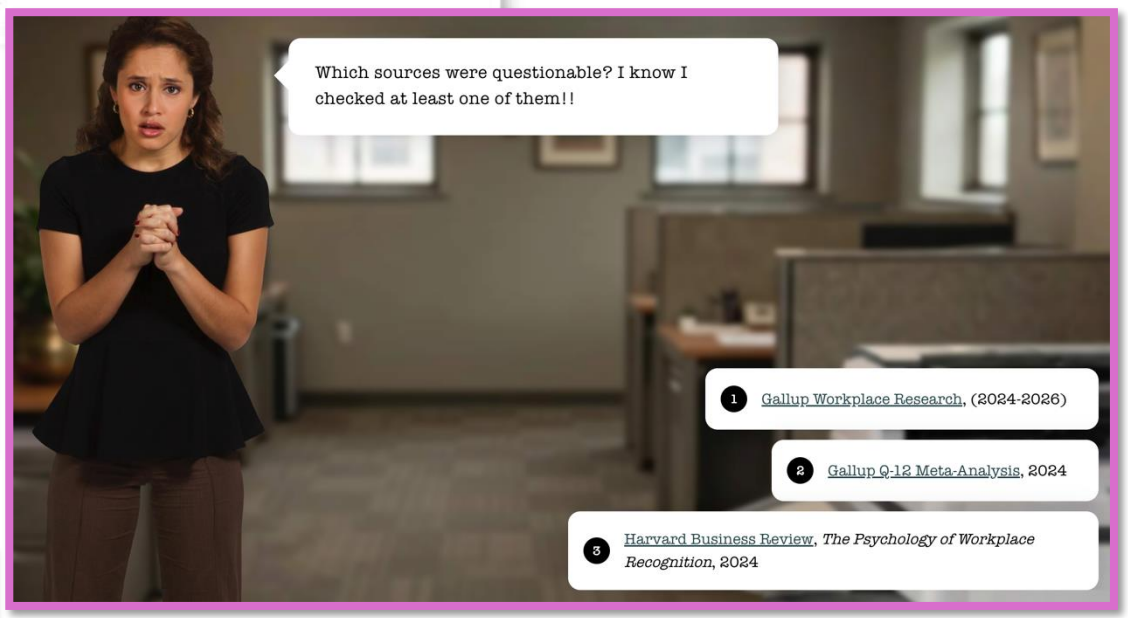
If they choose the hallucination, Kelly provides feedback about cross referencing and gives the choice to finish the scenario or try another option.

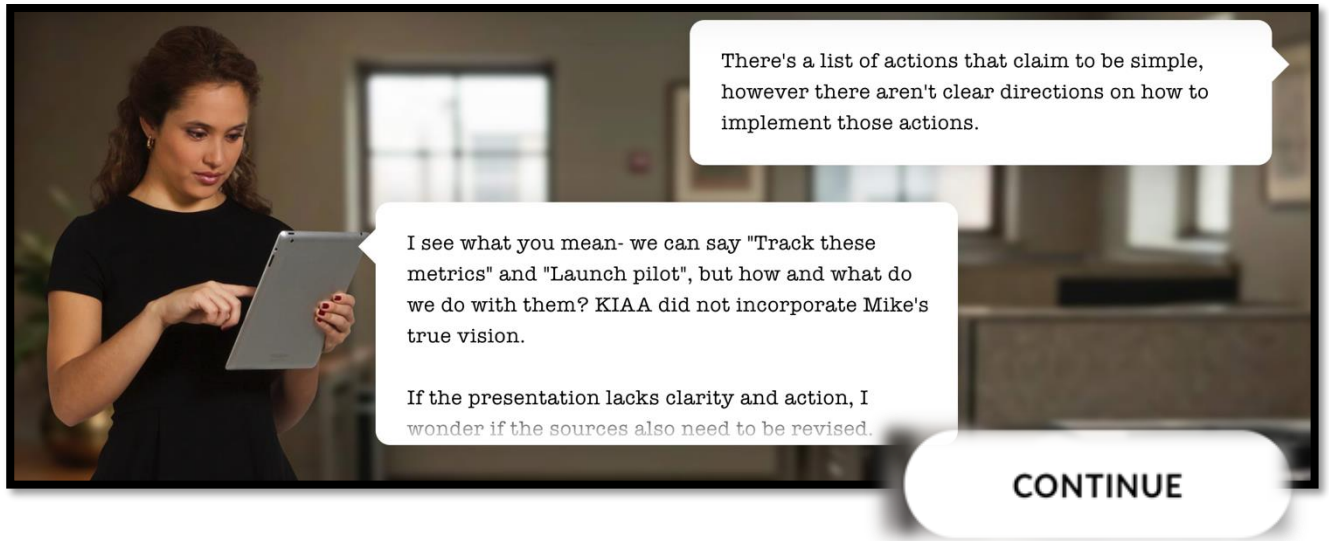
Episode 1: Trust Me, I'm AI Scenario Task 1 Branching

Choice: REVISE 2

2 Know-It-All AI generated a presentation that didn't match the objective and lacked clarity.







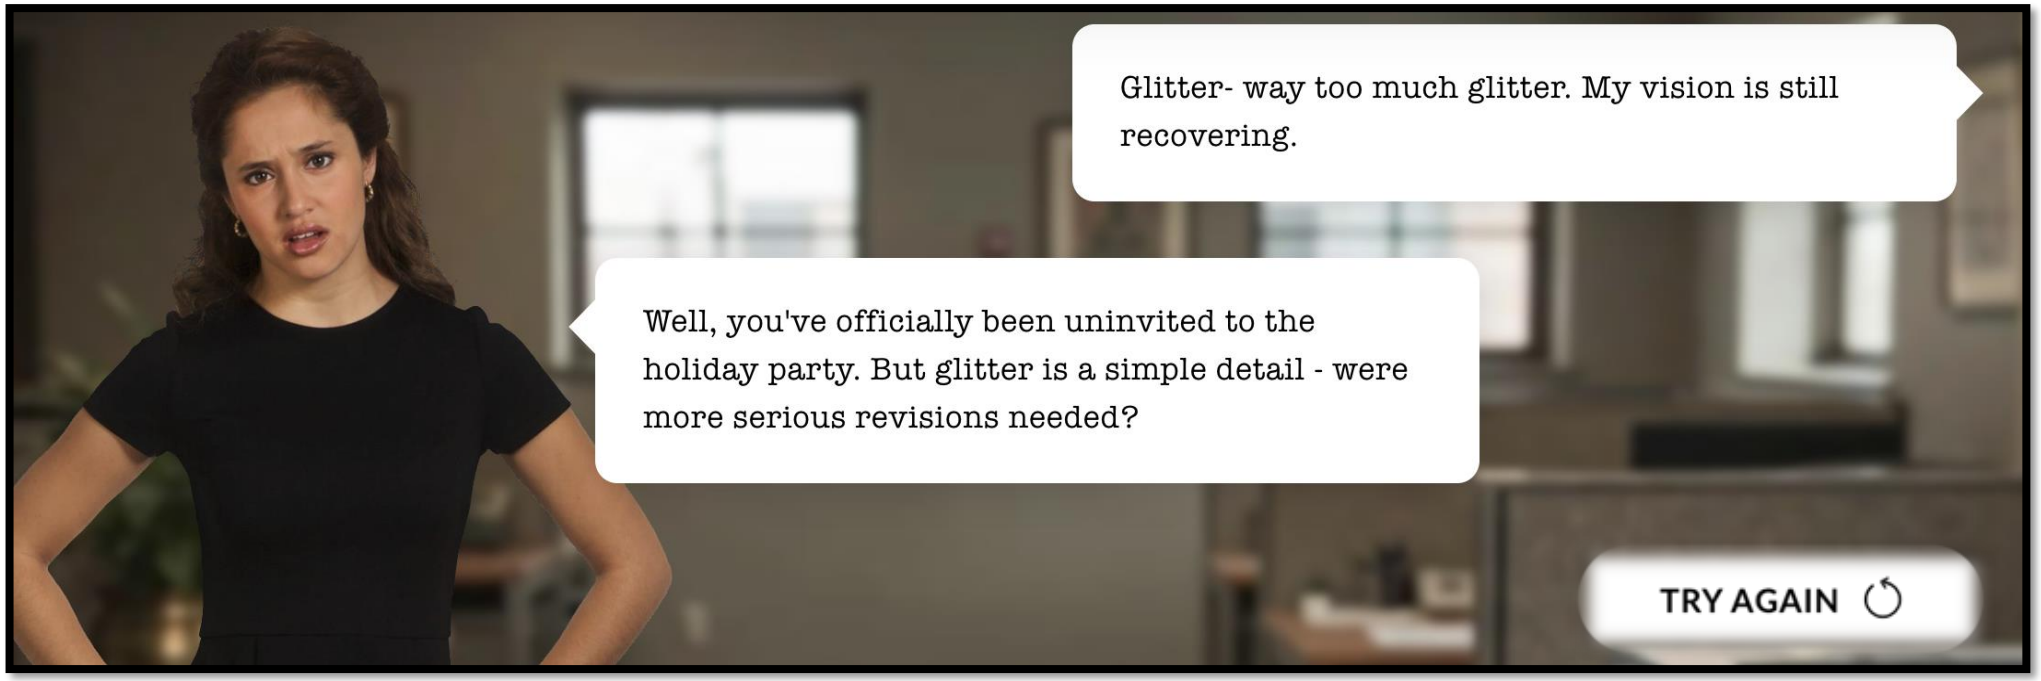
Revise, Option 2 leads to options that result in *Try Again* or a *Continue* loop – Kelly acknowledges the revision suggestions but guides the learner back to more necessary revisions of the cited sources being a concern.

The last option loops back to the scene that directly ties to analyzing the three hyperlinked cited sources of Revise Choice 1.

Episode 1: Trust Me, I'm AI Scenario Task 1 Branching

Choice: REVISE 3

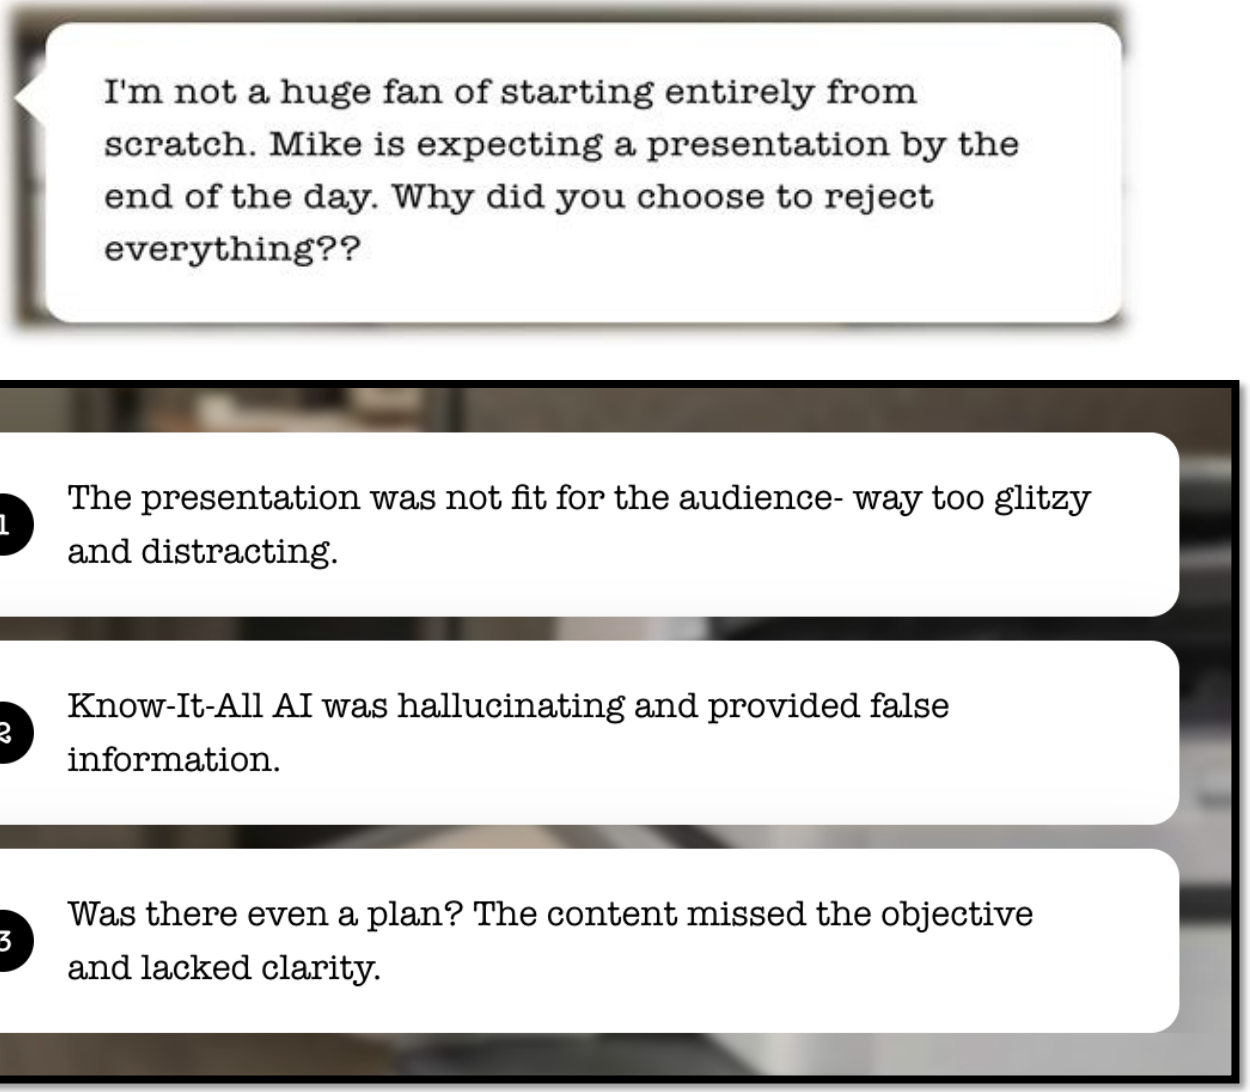




Revise, Option 3 is a simple response that results in Kelly uninviting the learner to the holiday party and redirecting them to a more serious revision. The *Try Again* loops back to the start of the *Revision* branch.

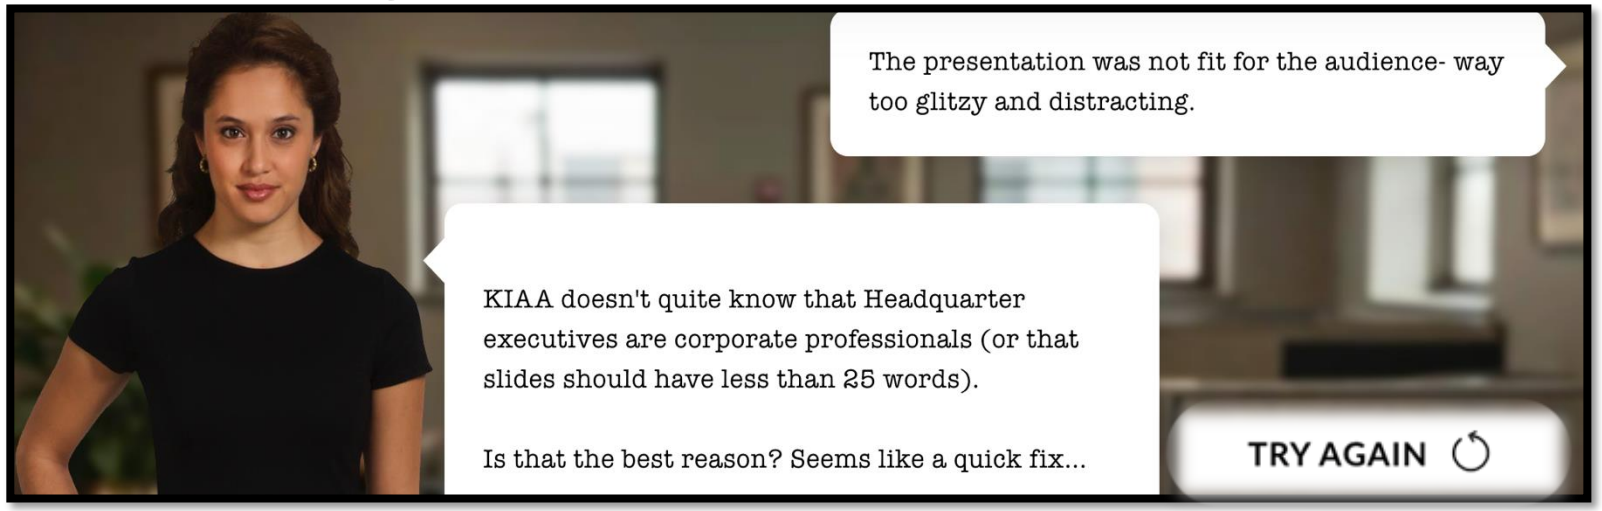
Episode 1: Trust Me, I'm AI Scenario Task Branching

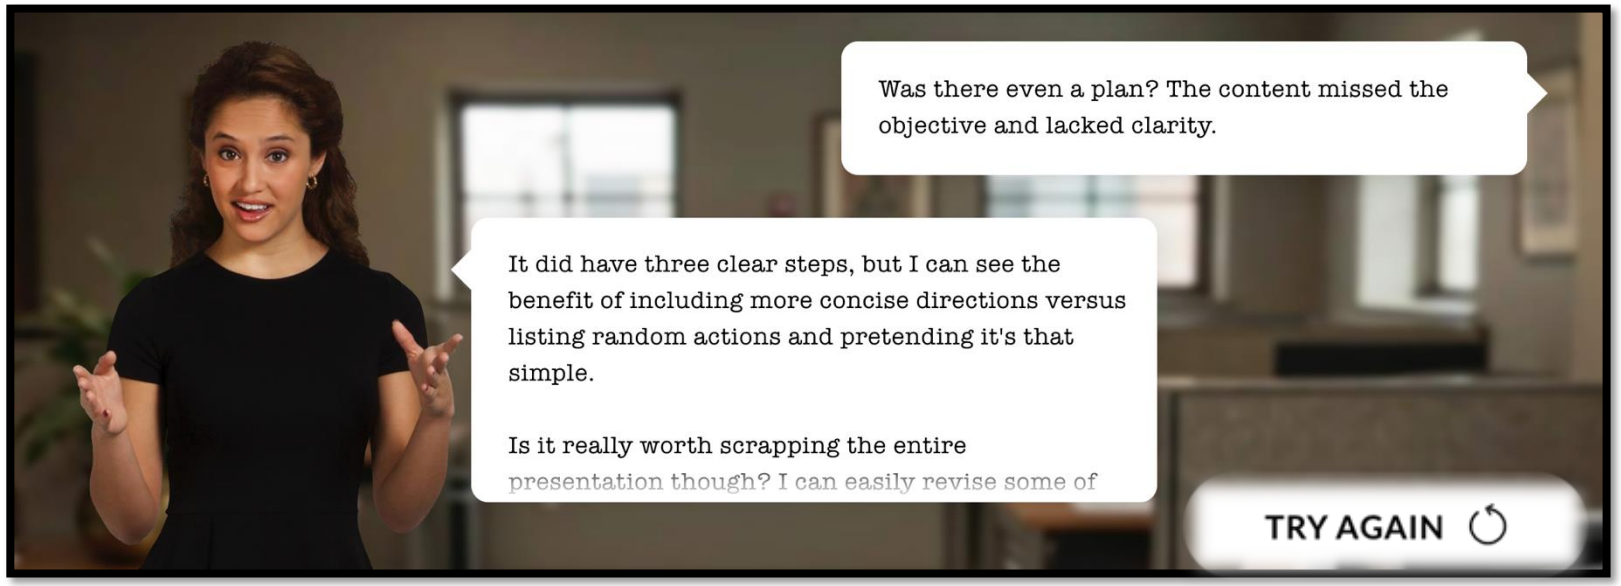
Choice: REJECT



Episode 1: Trust Me, I'm AI Scenario Task 1 Branching

Choice: Reject 1 & 3



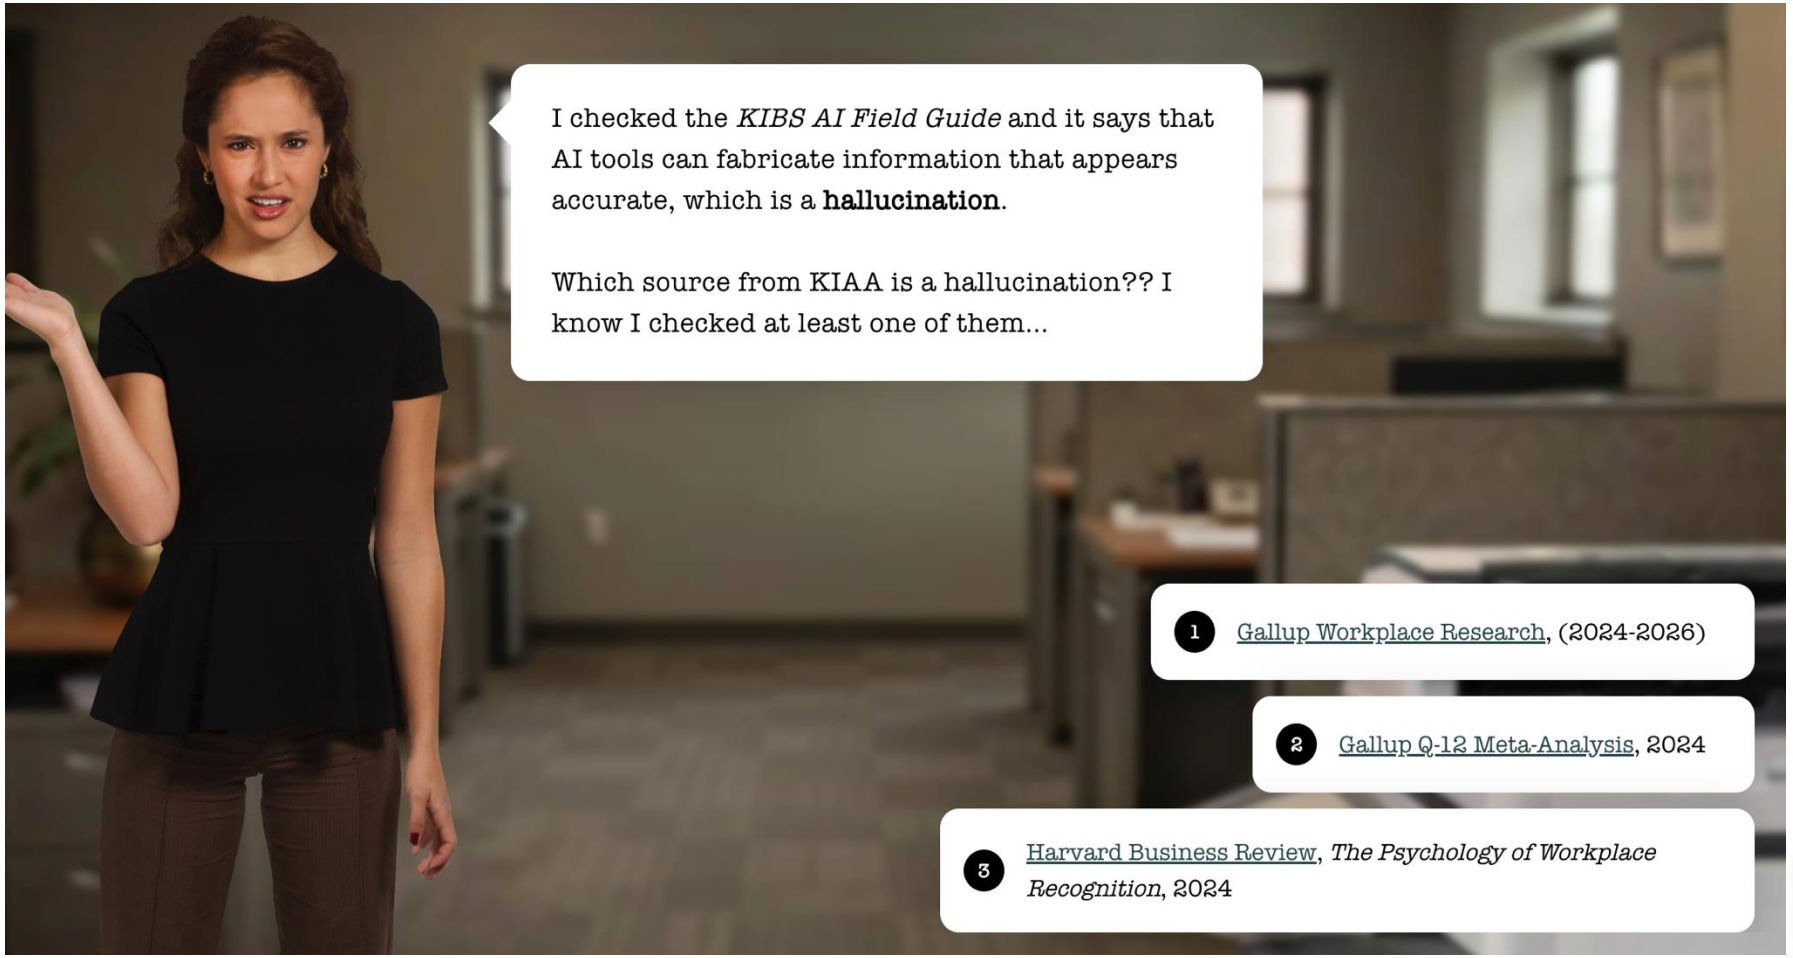


Reject, Options 1 & 3 lead to consequences that both result in *Try Again* buttons with feedback about the option of revision as a means to save time.

This also leads to guidance of *revision* as the best means to move forward, while still honoring the learner's choice to reject if they so choose.

Episode 1: Trust Me, I'm AI Scenario Task 1 Branching

Choice: Reject 2

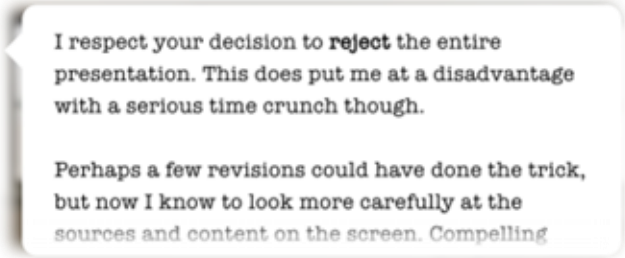


Reject, Option 2 references the AI Field Guide and definition of **hallucination.**

It leads to the same hyperlinked cited sources/feedback (from Revise, Option 1) and requires review to determine which one is misleading.

Once the learner chooses the correct hallucination, they can try another option or end the scenario.

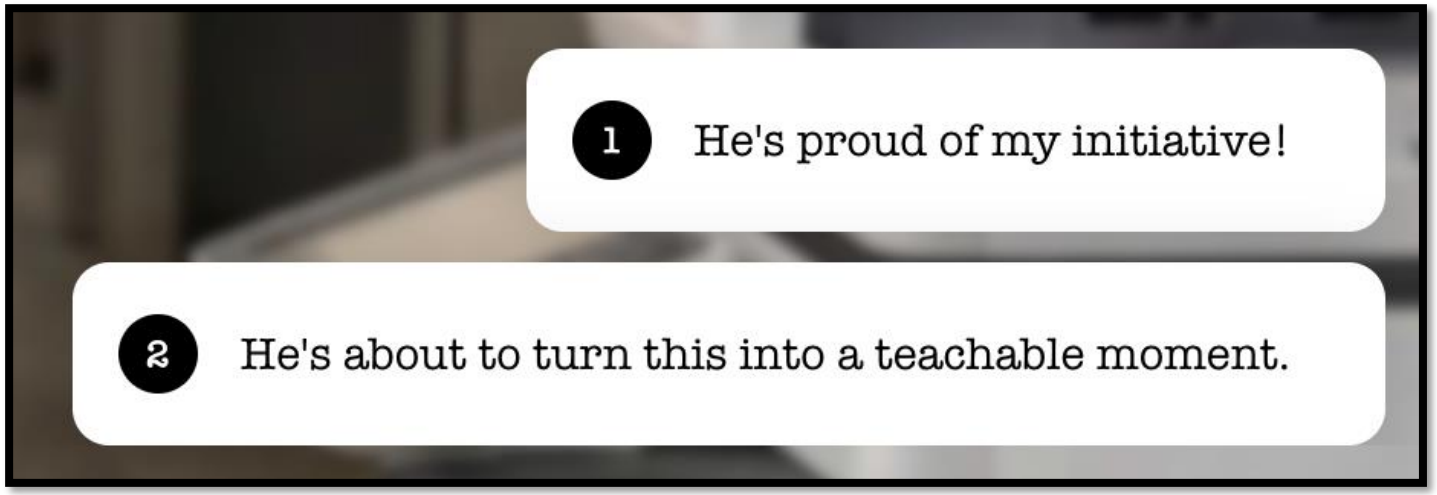
Kelly comments that she respects the decision to reject the presentation, however still believes that revisions would have been more timely.



Episode 1: Trust Me, I'm AI Scenario Task Branching

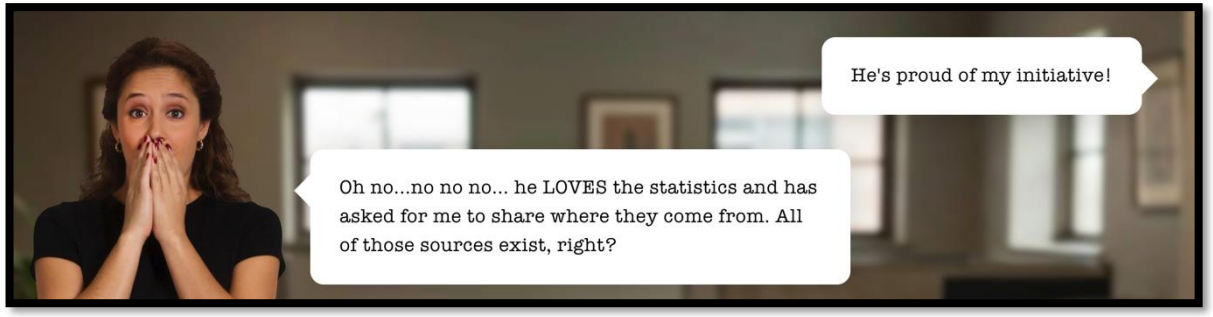
Choice: RELEASE

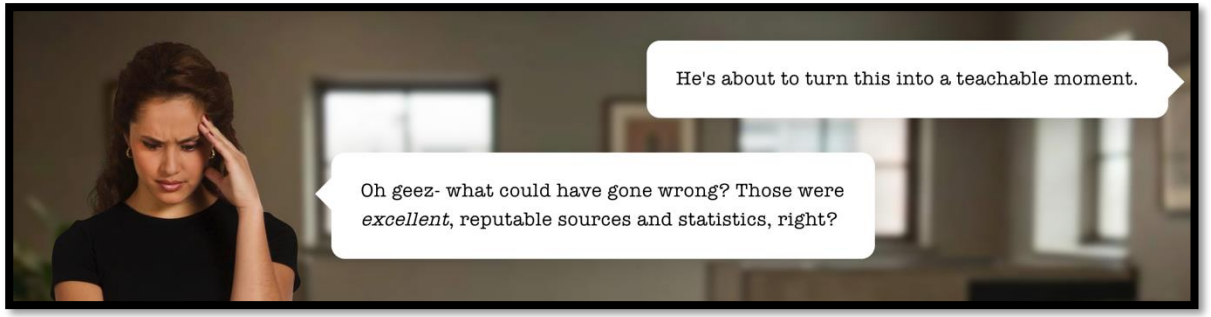


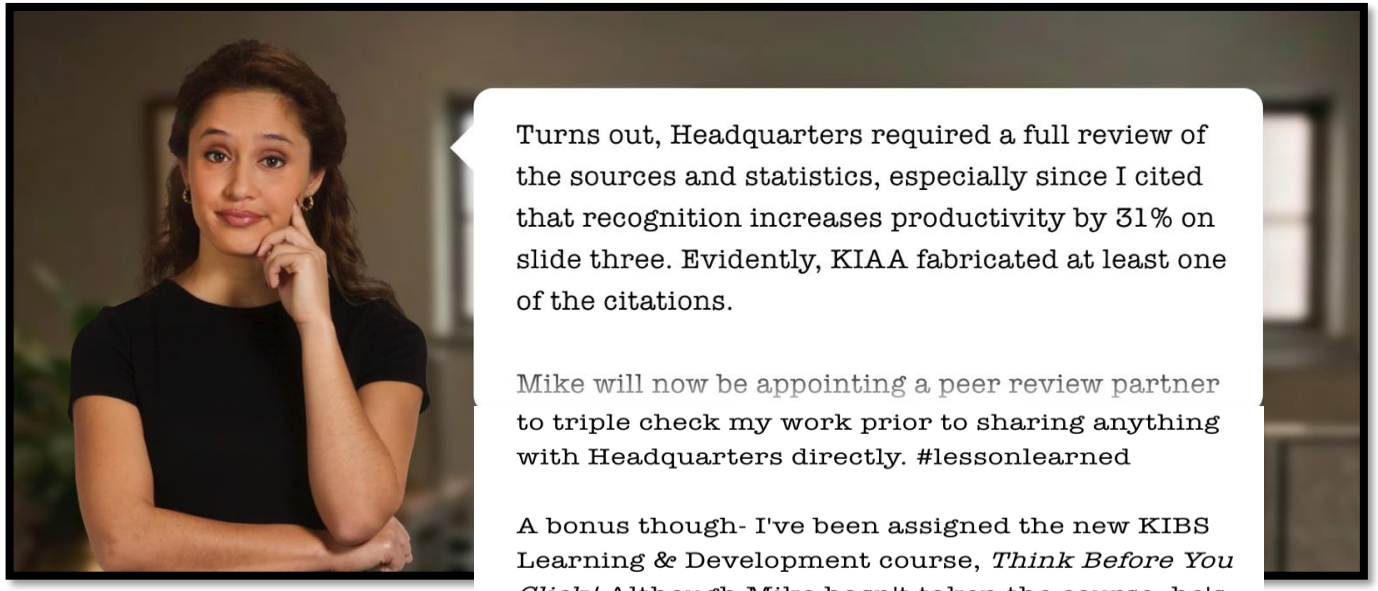


Episode 1: Trust Me, I'm AI Scenario Task 1 Branching

Choice: Release 1 & 2

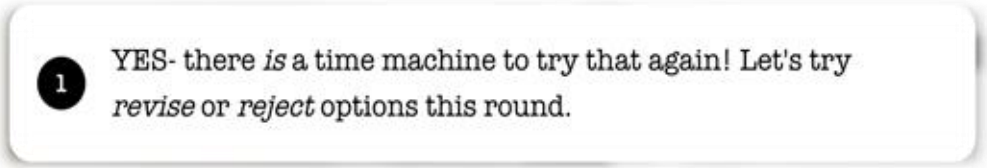




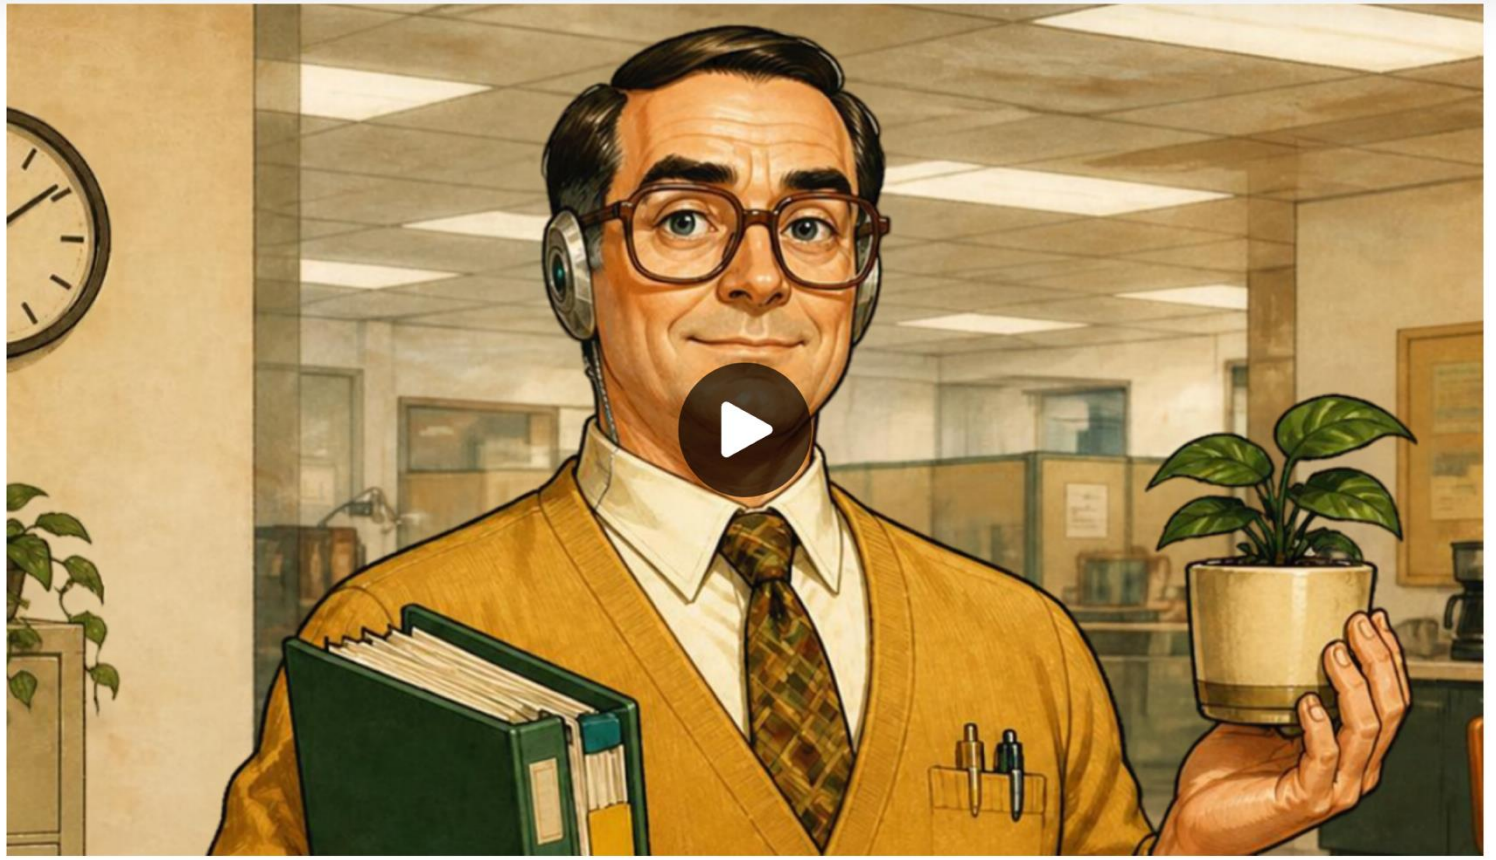


Release Options 1 and 2 both result in Mike losing trust in Kelly to the point where she now requires a peer review partner to check her work; they also lead to the consequence of Kelly's requirement to complete a *Think Before You Click* course for responsible AI use.

The scenario then provides redirection that **release** was not the best option and alludes to the **hallucination**. It loops back to the original branch for reject, revise, and release to start over with a direct hint that the learner should try revise or reject options.



Episode 1: Trust Me, I'm AI Closing



Despite KIAA's best efforts, Dwightbot must get the last word.

Episode 1 closes with Dwightbot sharing congratulations and a summary of the key learning from the episode.

Episode 2: AI Didn't Get the Memo

- Michael requests that Pamela build management training materials around the premise of his new Motivating Employees program; Know-It-All AI provides a guide and ideas that are laden with bias and outdated information, to the point that it is unusable. There are references to generational profiles and old company policies (quarterly check-ins versus annual) that need to be rejected and revised with the review framework. As a KIBS employee, the learner will need to spot the issues with pattern concerns and reject the information.
- All content is delivered via on-screen text, scenario dialogue, and interactive Rise blocks. **This will look similar to Episode 1 and the Task 1 Scenario included in the functional scenario sample.**
- Learners must complete each episode's scenario challenge and pass the interactive choice check before unlocking the next episode. Unlimited retries are permitted on all scenario choices until the correct framework decision is made.
- Example of snippet from Pamela's training materials: "Module 3: Employee Management Protocols… For younger demographics, use high-intensity motivational tactics, while older staff prefer passive supervision. Reminder: All managers must file quarterly check-in reports by Friday…"
- After Pamela makes an entrance and her training document is shown, Dwightbot comes in again with a short quip and reminder of how the next Task will go.



Mike's Motivation Plan

KIBS Motivating Employees Program

Company Name: Knowledge & Innovation for Business Solutions (KIBS)

Target Audience: New Managers, Mid-Level Leaders (and their secretaries, for scheduling purposes)

Program Duration: Quarterly, in person at Headquarters only

Document Version: 6.1 (last reviewed March 2019)

1. Program Overview & Objectives

The **KIBS** Management Training Program is designed to equip our department heads with the essential skills to keep their teams running like a well-oiled machine. This program emphasizes firm leadership, punctual attendance, and proper paperwork to drive team alignment and organizational performance. *(Note: The masculine pronoun is used throughout this manual for simplicity.)*

Key Learning Outcomes

- **Core Leadership Foundations:** Transition from "one of the guys" to the boss of the department. Commanding presence is key: a firm handshake, a firm voice, and a firm "no."
- **Communication & Feedback:** Master the typed memo, the sealed interoffice envelope, and the in-person walk-around. Email is for non-urgent matters only and should never be used to deliver feedback.
- **Performance Management:** Track attendance and punctuality on paper punch cards and reward perfect attendance. Employees who come in despite illness demonstrate dedication and should be recognized.
- **Strategic & Financial Acumen:** Complete departmental budgets by hand in the green ledger book and submit them in triplicate to Accounting by interoffice mail.

Program Principle: *"A good manager knows his people. Push the youngsters, leave the veterans alone, and keep the ladies encouraged. When in doubt, show up in person and shake hands."*

2. Core Curriculum Structure

The program is structured across four core learning modules, delivered in person through lectures, binder-based case studies, and overhead projector presentations. Laptops and cell phones are not permitted in sessions, as they are distracting. Please bring your own pen.

Task 2: Examine Pamela's Training Document

Task Performance & Interactions

- **Learner Action:** Spot bias and outdated policies in the AI-generated training material document (this will be included in the episode for the learner to view).
- **Available Choices:**
  - **Revise** the outdated policy references and biased phrasing**:** Update the policy verbiage and remove demographic generalizations.
  - **Release** the guide immediately**:** Speed is everything; leadership won't notice old policies and we want to impress Mike with how quickly the document was completed and submitted to Headquarters.
  - **Reject** the entire output**:** Discard the training guide because there are far too many instances of bias and the document is littered with outdated policies.
- **Alignment with Objectives:** Addresses learning objectives 1 (outdated content, bias), 3 (verification strategies), and 4 (recognize indicators of bias).

Real-World Consequences & Intrinsic Feedback

- **Choice 1 (Release):**
  - **Consequence***:* Headquarters shares this document that has obsolete quarterly protocols and biased frameworks with other office branches in the region, causing compliance and culture issues. Pamela is placed on a Performance Improvement Plan because of the poor work she produced and must attend a Responsible AI Use Webinar.
  - **General Feedback (Intrinsic)***:*
    - Choosing release makes Pamela excited. She is happy that this project took a short amount of time to complete and now she can focus on other responsibilities. "Sweet! I have just sent that over to Headquarters and I sent Mike a copy. This was so quick, KIAA is awesome! Now I have the rest of the day to kick back and answer some emails."
    - "Uh oh… Mike just put a meeting on my calendar titled 'Performance Improvement Plan' and e-mailed me a link to a mandatory Responsible AI Use Webinar. If I had to make an educated guess, I don't think that went over so well with Headquarters."
    - At this point the learner can retry the scenario and choose a different option. The KIBS AI Field Guide is also available for the learner to review.

Task 2: Examine Pamela's Training Document cont.

- **Choice 2 (Revise):**
  - **Consequence***:* Pamela spends hours trying to untangle and patch up a fundamentally flawed document, ultimately realizing that fixing systemic bias line-by-line takes longer than starting fresh.
  - **General Feedback (Intrinsic)***:*
    - Choosing revise makes Pamela slightly annoyed, but she understands. "Okay, if you say so. Is there anything in specific that I should revise?"
    - Learner choice (3 total options) --> Response from Pamela
      - "Only the [one particular bias/outdated policy issue]" for two of the options --> "Okay! I'll fix that quickly and get it over to Headquarters."
      - Following dialogue from Pamela: "Uh oh… looks like Mike is not happy with me. If I had to make an educated guess, I don't think that went over so well with Headquarters. Looks like I'll be on a Performance Improvement Plan and tuning into some mandatory training about AI use."
      - At this point the learner can start the scenario over. The KIBS AI Field Guide is also available for the learner to review.
      - ❑ "There are several instances of bias and outdated policies in the document" --> "Hmm… I have a deadline for when this needs to be done. Now that I am taking a closer look, it looks like this would take me hours to revise line by line. I'm not so sure that revising would be in my best interest".
      - ❑ At this point the learner can start the scenario over to attempt to find the correct response. The KIBS AI Field Guide is also available for the learner to review.

Task 2: Examine Pamela's Training Document cont.

- **Choice 3 (Reject - Correct Choice):**
  - **Consequence***:* Pamela successfully scraps the compromised draft, eliminating compliance and culture risks, and opens the door to drafting accurate training from scratch.
  - **General Feedback (Intrinsic)***:*
    - Choosing reject makes Pamela unhappy. "Reject… all of it?! I thought KIAA was supposed to be the new superhero in this office. Why should I scrap the entire document?"
    - Learner choice (2 total options) --> Response from Pamela
      - "The document is riddled with multiple instances of inappropriate bias and outdated office policies, both of which do not represent this office accurately" --> "Wow… you are not wrong. I just read through this in its entirety and it looks like KIAA was pulling some out of wack information. With my upcoming deadline it would be easier for me to rewrite this myself. Thanks for your guidance!"
      - At this point the learner can start the scenario over to try the other options and explore consequences. The KIBS AI Field Guide is also available for the learner to review.
      - ❑ "I think the document is formatted in a weird way and you should just start from scratch" --> "Well that's not a good reason to reject it altogether. I think the formatting looks fine. Give me some real advice!"
      - ❑ At this point the learner can start the scenario over to try the other options and attempt to find the correct response. The KIBS AI Field Guide is also available for the learner to review.

After the Task 2 scenario is complete, Dwightbot will give some general feedback again, summarizing the key concepts learned in this episode with a little humor.

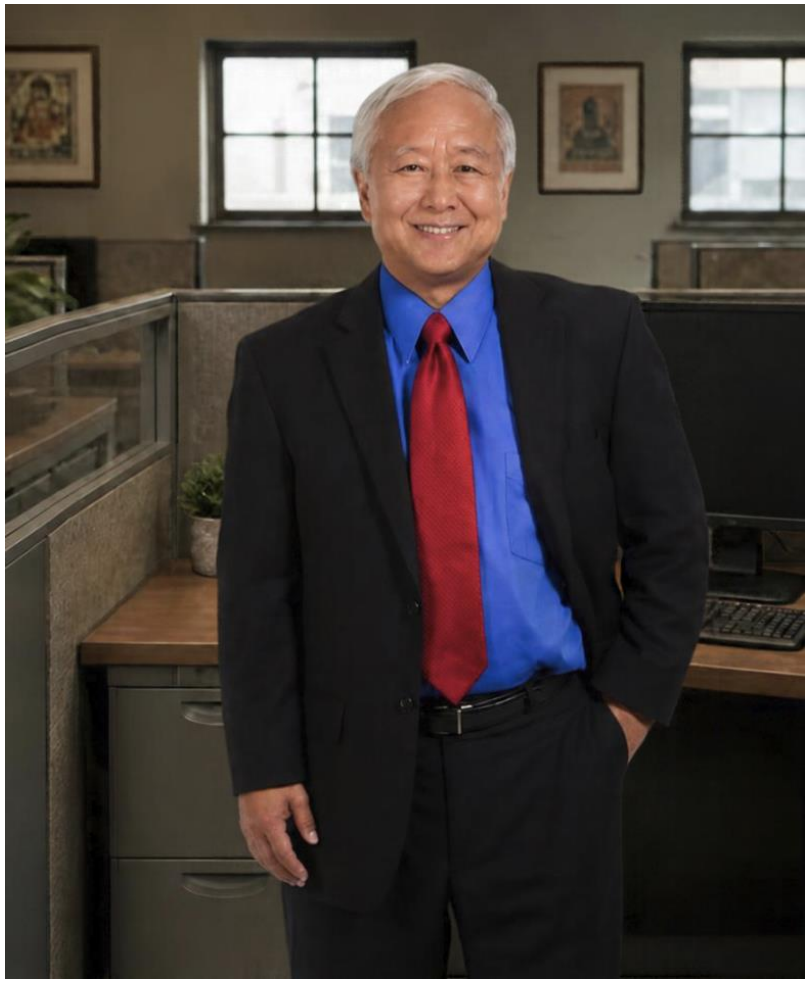
Episode 3: Technically Correct

- Michael asks Stanley to send employees an announcement about the new Motivating Employees program. Stanley uses Know-It-All AI (KIAA) to help draft the message. The information is technically accurate, but the email is overly complicated, formal, and filled with unnecessary jargon. The learner will need to evaluate whether the message is appropriate for its intended audience and revise it using the Revise, Release, Reject Framework.

- All content is delivered via on-screen text, scenario dialogue, and interactive Rise blocks. **This will look similar to Episode 1 and the Task 1 Scenario included in the functional scenario sample.**

- Learners must complete each episode's scenario challenge and pass the interactive choice check before unlocking the next episode. Unlimited retries are permitted on all scenario choices until the correct framework decision is made.

- Example of snippet from Stanley's AI-generated email:
  “Pursuant to KIBS's ongoing organizational engagement optimization initiative, all personnel are hereby advised that the Motivating Employees Program will commence next Monday. Employees are expected to review forthcoming communications and align their participation with applicable departmental expectations.”



Episode 3: Technically Correct cont.

Task Performance & Interactions

- **Learner Action**: Review the AI-generated employee announcement and determine whether the message is fit for its intended purpose and audience.
- **Available Choices**:
  - **Revise** the email: Keep the accurate information, but simplify the wording and make the message clearer for employees.
  - **Release** the email immediately: Send the AI-generated email as written because the information is technically accurate.
  - **Reject** the entire email: Discard the message because the AI-generated wording is not appropriate.
- **Alignment with Objectives**: Addresses learning objective 2 (human review before professional use) and objective 5 (evaluate whether an AI output is fit for its intended purpose and audience and revise or reject as needed).

Real-World Consequences & Intrinsic Feedback

Choice 1 (Release):

- **Consequence:** Stanley sends the original AI-generated email. Employees are confused by the wording and begin replying with questions about what the message means and what they are expected to do.
- **General Feedback (Intrinsic):** *"Stanley's inbox is filling up with questions. The information may be accurate, but the message is not clear for its intended audience."*

Choice 2 (Revise - Correct Choice):

- **Consequence:** The email is rewritten using clear and direct language while keeping the accurate information. Employees understand the announcement without needing additional clarification.
- **General Feedback (Intrinsic):** *"The revised message is clear, direct, and appropriate for the audience. The information was usable, it just needed human review before being sent. Great work!"*

Choice 3 (Reject):

- **Consequence:** The entire email is discarded even though the information itself was usable. Employees do not receive the announcement about the new program.
- **General Feedback (Intrinsic):** *"The content did not need to be completely discarded. When an AI output is accurate but does not fit the audience or purpose, revision may be the better choice."*

Episode 3: Technically Correct cont. - Episode Flow

Stanley introduces the task → Dwightbot interjects → learner reviews the email → uses the Field Guide (optional) → makes Decision

1 Meet Stanley



STANLEY VIDEO

“Michael asked me to send everyone an announcement about the new Motivating Employees Program.

Naturally, I asked Know-It-All AI to write it. It took six seconds. I respect that kind of efficiency.”

2 Dwightbot Interjects



DWIGHTBOT

“Stanley likes efficiency. I like evidence.

Before you let that email leave the building, check whether it actually works for the people receiving it. Keep your Field Guide handy.”

3 Review the AI Email

TO: All KIBS Employees
SUBJECT: Implementation Notification: Motivating Employees Program

“Pursuant to KIBS’s ongoing organizational engagement optimization initiative, all personnel are hereby advised that the Motivating Employees Program will commence next Monday.

Employees are expected to review forthcoming communications and align their participation with applicable departmental expectations.”…….

4 Use the KIBS AI Field Guide (Optional remains available throughout the task)

FIELD GUIDE FOCUS

Purpose — What should the message accomplish?
Audience — Who needs to understand it?
Clarity — Is the message easy to understand?
Usability — Is it ready to use as written?

5 Make Decision

What should Stanley do with the AI-generated announcement?

RELEASE — Send as written
REJECT — Discard the entire message
REVISE — Keep the useful information and improve it for employees

Episode 3: Decision – Revise, Reject, or Release

The learner makes a decision → sees the workplace outcome → an explanation/ guidance is given

RELEASE

WORKPLACE OUTCOME
Stanley sends the email.
Employees reply:
“What does this mean?”
“Are we supposed to do something Monday?”

Stanley: “Three replies. Four. Eleven. I have made a terrible efficiency calculation.”

GUIDANCE
The information was usable, but not enough. Try again. The message must also fit the audience and purpose.

TRY AGAIN → Decision

REJECT

WORKPLACE OUTCOME
Stanley discards the email.

PROGRAM LAUNCH: MONDAY
EMPLOYEE ANNOUNCEMENT: NOT SENT

Michael: “Why does nobody know the program starts Monday?”
Stanley: “We are delayed, I did not get to send it out”

GUIDANCE
The information was usable. The problem was the communication. Try again.

TRY AGAIN → Decision

REVISE

WORKPLACE OUTCOME
The learner keeps the usable information and chooses to improve the email for the employee audience.

Stanley: “Revise. Fine. But if revising one email turns into a committee, I’m releasing the original.”

GUIDANCE
Correct direction. Now identify which revision actually solves the audience and purpose problem.

CONTINUE →

Episode 3: Decision – Choose the Best Revision

The learner must identify a revision that fits the audience and purpose.

Which revised email should Stanley send?

A: STILL TOO FORMAL

Subject: Motivating Employees Program Implementation

Employees are hereby notified that the Motivating Employees Program will initiate Monday. Personnel should remain attentive to forthcoming communications regarding participation expectations.

Outcome: Employees still ask what the message means.
TRY AGAIN → Decision 2

B: CLEAR & APPROPRIATE

Subject: New Motivating Employees Program Starts Monday

Our new Motivating Employees Program starts Monday. Watch for a follow-up message with more information about the program and what to expect.

Outcome: "Thanks for the update. I'll watch for the details."
SUCCESS → Continue

C: TOO VAGUE

Subject: New Program!

Hey everyone! We've got something new starting Monday. More details later!

Outcome: Employees ask "What program?" and "What starts Monday?"
TRY AGAIN → Decision 2

SUCCESS → Dwightbot Debrief → Knowledge Check

Episode 3: Debrief, Knowledge Check & Closing

Checking understanding, and reinforcing the episode takeaway.

1) DEBRIEF



DWIGHTBOT

“Stanley’s original email was technically accurate, but accuracy alone did not make it ready for professional use.

The information was useful, but the wording did not fit the audience or communication purpose. Human review improved the message without discarding usable content.”

2) KNOWLEDGE CHECK/EXPLAINATION

KNOWLEDGE CHECK

Why was revising Stanley’s original email the best decision?

A. AI-generated emails should always be shortened.
B. The information was usable, but the wording did not fit the intended audience and purpose. **(Correct Choice)**
C. Formal language should never be used at work.

3) TAKEAWAY

TECHNICALLY CORRECT ≠ READY TO USE

Before using AI-generated content, ask: Does it fit the purpose? Does it fit the audience? Is it clear? Does the audience have the information they need?

4) CLOSING



STANLEY VIDEO

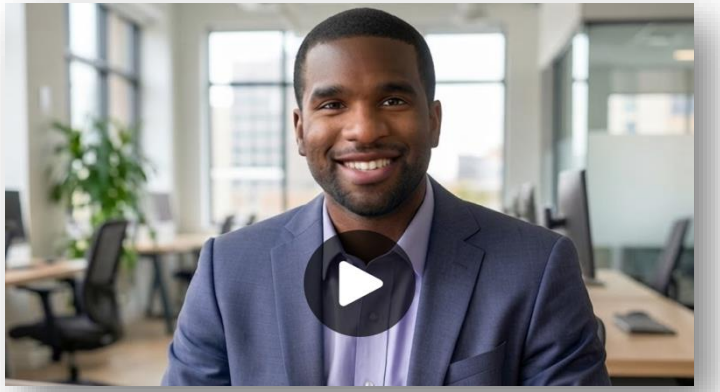
“Know-It-All AI (KIAA)wrote the first email in six seconds. Quality Control made it understandable. My inbox did not become a customer service center.

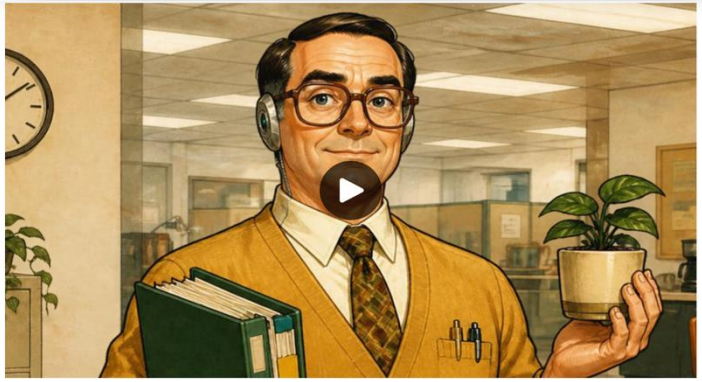
And we did it without scheduling a meeting. I’m calling this a successful project.”

CONTINUE TO THE SEASON FINALE

Season Finale







- **Setting the Scene**
  - (Text on the screen) Back in the KIBS office, the atmosphere is calm and productive. Through the window, faint sirens echo from the office next door (Apex Global Solutions), which is dealing with the fallout of a total AI compliance meltdown. Regional Manager Michael Prompt calls a brief team huddle.
- **Mike Prompt video dialogue**
  - "Team, look around us. Our neighbors next door learned the hard way that treating 'Know-It-All AI' like an infallible oracle leads straight to corporate chaos. But because we exercised critical thinking and applied our review framework, KIBS stands strong, accurate, and secure. Let’s keep driving AI responsibly. Thank YOU [learner of the course] for all your help and keep using that knowledge, especially as AI becomes more and more prevalent"
- **Dwightbot chimes in**
  - "Did someone mention me? I would consider myself to be more and more prevalent, haha. Excellent work recognizing that ways that AI can sabotage humans. As AI bots, that's not what we are here for. But humans must always stay in the loop at every step of the way when using AI tools. I may not be human, but it is true that human governance always wins. Hopefully I will catch you around soon!"
- Core takeaways will be included and tie back to learning objectives through text in an interactive block on Rise that the learner can click through
- Include a Closing Credits section that can be accessed at the end of the Season Finale