

Exploring Voice AI in Education

A Team Investigation of Present Uses and Future Educational Experiences

Scientific Inquiry with AI

Student Guide — V3

Detailed Instructions & Course Material

Why This Lesson

What can voice AI help you do as a learner today? What educational experiences could become possible if we developed the technology further? In this lesson, you will explore both questions. Begin with a real learning need and explore how the voice mode in your chatbot can help address it. Then, use what you learn to envision a more sophisticated application that combines voice recognition, speech generation, and AI's ability to quickly access information and reason through problems.

For this lesson you will work in a team of four on the two scenarios: a practical use you can try now and a future educational experience you can design. Six teams will share their trials and proposals on day two of the lesson, so the class encounters a broad range of possibilities. Explain what you observed, what related research supports, and what remains untested.

Learning Goals

Upon completion of this investigation, you should be able to:

Phenomena:

- Describe diverse educational uses of voice AI, including a use available today and a future experience that requires additional capabilities.
- Explain what speaking and listening contribute to each.

Experimental Practices:

- Generate distinct ideas, try a current use with an observable outcome, and revise it.
- Design a future experience around a learning goal.
- Locate and read related research, explain its relevance, and propose how you would evaluate your own design.

Habits of Mind:

- Recognize how using AI voice mode can be applied to enhance your future learning.
- Distinguish an appealing idea, a working demonstration, and evidence of learning.
- Working with teams to brainstorm and share ideas on short time frames.

What You Will Need

- A phone or computer with an AI tool that supports back-and-forth voice conversation, plus headphones and a microphone if available. Test access before class.
- A shared team document, a way to save brief interaction notes.

The Core Idea

- An educational use begins with a learner and a learning goal. The conversational role should serve that goal.
- A smooth conversation can be useful, but does not by itself show that a learner understands more.
- Research on a related activity can inform a new proposal. Differences in learners, technology, and outcomes limit what transfers.

Driving questions. What can voice AI help you learn now, and what new educational experiences could we build around responsive conversation?

Your Time Plan

When	Time	Main task
Before class 1	120 min	Experience a conversation, generate your own possibilities, and review how to evaluate research
Class 1	110 min	Explore possibilities, choose two scenarios, and prototype the current use
Between classes	~120 min each	Test and revise Scenario 1, find research, and design Scenario 2
Class 2	110 min	Present, build a class map of possibilities, and reflect
After class 2	120 min	Finish the team brief, write a reflection, try a personal application, and complete retrieval practice


Session A — Pre-Class Preparation

~2 hours


Allow about 2 hours before the first class. Bring your ideas and observations to class.

A. Explore ChatGPT Voice Mode (~15 minutes)

Investigate the capabilities of ChatGPT voice mode through a 10-minute conversation with ChatGPT. Before you start, go to settings and choose the voice you want to use. Click on the current voice and you will see a list of voices you listen to and then choose.



Pick any topic to discuss with the AI that interests you. The goal here is just to give you experience interacting with the AI in voice mode. When you finish, check the transcript to see how well the dialog was captured.



Example starter prompt: “I am interested in learning more about UC Merced. What can you tell me about how it got started and what is the vision for it going forward.” You can follow up asking about specific questions like when will the medical school open and if there will be opportunities for current UCM students.

B. Background Reading (~30 minutes)

Download the document titled “Talking, Listening, Reading, Writing How the Way We Share Ideas Shapes Learning” from the supplemental files folder on the knowledgewithai.com website. Spend about 15 minutes reading the document. Then download the Voice Check Your Understanding document which will ask you 7 questions, which you should dedicate 10 minutes to answer. Load both documents into a chatbot and follow the instructions below.

1. Begin in a text conversation. Upload both the reading and the check your understanding documents, then copy and paste the **Starting Instruction** into the chat.
2. **Continue that same conversation in voice mode.** Headphones and a quiet space can help. If voice is unavailable or impractical, type your answers instead.
3. Answer each question in your own words. Say “let me think” when you need time. A tentative answer with reasoning is useful.
4. The tutor should ask all seven in order and give only brief feedback.
5. After question seven, end voice mode and type “summary.” Save the summary and conversation and submit them as your instructor directs.

Starting instruction: Start my voice check. Use the attached worksheet as the question script and the attached research review as the reading. Ask question 1 of 7 exactly as written, then wait for my answer. Track progress through all seven questions and follow the worksheet instructions throughout the conversation. My instructor has given me 10 minutes to conduct this exercise.

If the tutor drifts: Say “Go back to the worksheet. Check which numbered question is next, read it exactly as written, and wait for my answer.”. You may pause or stop whenever you need to; the summary must identify any questions left unfinished.

C. Generate Your Own Possibilities (~30 minutes)

Identify two potential use cases where you think voice mode might be beneficial in education and career preparation. Try to think through ideas on your own and expand upon them using an AI to brainstorm. Write them down in your notebooks with a short description of how AI would be used and describe the role of the learner.

D. Literature Research (~15 minutes)

For one idea, use your skills using AI in literature research to conduct a preliminary search to see what is already known about using AI and voice mode for that particular application. Note its title, link, research question, and whether it concerns voice conversation, text AI, or a related teaching method. This is preparation for discussion you will have in class.

E. Review of In-Class Investigation (~10 minutes)

Preview the stages below and note one design decision you want to discuss.

Session A — In-Class Investigation

The Investigation, Stage by Stage

The investigation runs across two 110-minute classes which is followed by team work outside of class. Your instructor will lead the timing and checkpoints. The guide identifies decisions your team needs to make; you choose the educational situation, interaction, and evidence.

Class One — Explore the Possibilities (110 minutes)

Segment	Time
Opening instructor discussion	15 min
Personal trial	10 min
Observations	10 min
Team brainstorm	20 min
Choose both scenarios	10 min
Prototype current use and sketch future experience	25 min
Peer feedback	15 min
Checkpoint	5 min

Stage 1 — Build your Teams and Teach an AI (20 minutes)

Form groups of three or four. Exchange contact information so you will be able to continue to work together outside of class. Share your experience from the prelab with others and make sure everyone knows how to enter voice mode and select voices. Conduct a short multiple person conversation with a single chatbot using voice mode. Instruct the chatbot to play the role of introductory physics student who does not understand linear motion. Use the Socratic method (teach through questioning and dialog) to teach them about distance, displacement, speed, and velocity. Make sure each person contributes to the teaching.

Stage 2 — Brainstorm as a Team (30 minutes)

Each member of the team should share the two ideas for AI voice mode use in education that they developed in the prelab. Create a list of all the individual ideas and then review the ideas in the list below. The team should discuss all potential ideas and consider merging roles and uses. Teams should agree on a single concept that they will develop and evaluate.

AI role	Example learning activity	Why use voice?
STEM tutor	Explain a physics concept; ask for questions before hints.	Practices reasoning aloud.
Debugging coach	Describe a coding problem; talk through possible causes.	Pairs spoken reasoning with coding.
Lab assistant	Hear approved steps and dictate experimental observations.	Supports hands-free work.
Fieldwork partner	Record observations; discuss competing explanations.	Supports inquiry away from a desk.
Student to teach	Teach the AI distance versus displacement; answer its questions.	Encourages explanation and rephrasing.
Presentation coach	Rehearse a science talk and respond to audience questions.	Practices spoken communication.
Simulated patient	Practice interviewing or explaining science to a nonexpert.	Develops listening and questioning.
Language partner	Role-play a conversation in another language.	Offers responsive speaking practice.
Reading tutor	Read aloud; receive help and discuss meaning.	Connects listening with reading support.
Accessible interface	Explore voice access for low vision or limited hand mobility.	Can reduce visual and typing demands.
Communication partner	Explore student-chosen turn-taking practice for autistic learners.	Offers repeatable spoken exchanges.
Multimodal access aid	Combine speech, captions, text, and translation.	Offers choices for participation.
AI-literacy experiment	Check a scientific explanation; test how voice affects trust.	Makes convincing speech an object of study.

Stage 3 — Develop Application (55 minutes)

Create a summary of the application which should emphasize what conversation adds. Explain why listening, speaking, follow-up questions, or turn-taking serve the learning goal.

Build a Prompt for the Application

Write a detailed prompt specifying the learner, goal, AI role, learner task, and limits on AI assistance. The four roles below can help you start, but you are free to design another role or combine them.

- **AI explains:** “Give a brief explanation of [concept] for [learner]. Then ask me to apply it to a new example before offering more help.”
- **AI tutors:** “Help me reason about [task]. Ask one question at a time. Let me attempt an answer before giving a hint. Check my reasoning against [reference].”
- **Student teaches AI:** “Act as an interested novice while I explain [concept]. Ask for clarification when I skip a step or use undefined language. Do not give the explanation for me.”
- **AI challenges:** “Question my claim about [case]. Ask for evidence and alternative explanations, one question at a time. Keep challenges relevant and check factual objections.”

Prompt skeleton: “Our learner is [who]. The goal is [observable learning task]. Act as [role]. Have the learner [action]. Keep turns short. Give help only when [condition]. Use [reference] to check content. End with [independent task or useful record].”

Plan a Small Trial

Choose one observable outcome that fits your goal. For understanding, use a new question answered without AI and check the reasoning against a reliable reference. For communication, ask a human listener to state the main idea and identify an unclear moment. For planning, check whether the learner can explain and justify a revised plan. Enjoyment and confidence can be additional observations, but keep them separate from evidence of learning.

Run a brief prototype of Scenario 1. Rotate who speaks, operates the tool, observes learner actions, and checks evidence. Every member should experience the learner role during the weekend trial. Also sketch Scenario 2: who participates, what they do, and what becomes possible. Make each member’s contribution identifiable; everyone need not speak in the presentation.

Peer Feedback

Explain both scenarios to another team. Ask them to identify the learning goals, work left to the learner, and one assumption worth testing. For the future experience, ask whether it creates a new opportunity and what would have to be built. Revise your plans.

Class Checkpoint

Show the instructor both scenarios, the current-use prompt and observation plan, and an initial sketch of the future experience. Agree on weekend responsibilities and confirm that your direction adds variety to the class.

Between Classes — Test, Design, and Find Research (~2 hours)

Plan about 2 hours per person, including a team meeting. Divide research tasks and build on work started in class.

Task	Time
Coordinating as a team	20 min
Searching and reading	30 min
Testing and revising the current use	30 min
Developing the future experience	25 min
Preparing the presentation	15 min

Stage 4 — Try, Revise, and Document

Each member tries Scenario 1 for about 5 to 8 minutes and completes the agreed observation or independent task. Save the original prompt, tool and date, a representative exchange, and the outcome. Document failures and contradictions as well as useful moments in your shared log.

Trial log fields: Participant label; task and AI prompt; what the learner did; observed outcome and how it was checked; limitation or failure; proposed change. Use anonymous labels and remove private details.

Choose one change to the prompt or procedure and test the revision with at least one member. Explain what changed in the interaction and why. Repeating a task can itself help performance; order, familiarity, and different participants can also affect outcomes. Describe the trial as exploratory, and do not treat four classmates as proof of effectiveness.

Stage 5 — Connect the Idea to Research

Use your literature-review skills to find at least one closely related research example if you can. It may inform either scenario or both; you do not need a separate study for each. Search for the technology and the learning activity, such as *conversational agent* plus *language practice* or *self-explanation* plus *tutoring*. Try voice, speech, dialogue, and text terms separately.

Ask the AI: “Suggest search terms and candidate research for [use and goal]. Distinguish voice AI from text AI and research on the underlying teaching method. Provide source links, and mark any references you cannot verify.”

Open the source yourself. Verify that it exists and read enough of the methods and results to support your summary. An AI summary or search snippet is not evidence. If only an abstract is accessible, identify that limit and avoid claims that require details you could not inspect.

Make a Short Evidence Note

- Citation and working link, with the title and authors checked against the source.
- Learners and setting; activity or technology studied; comparison, if any; and outcome measured.
- One relevant finding, with its location in the source, and one important limitation.

- Which scenario the study informs, how it helps your reasoning, and where the match ends. Does it study voice AI, text AI, or a related educational practice?

If no direct study turns up, use the closest related research and explain the gap. If you cannot find a credible related source, document databases or search tools, search terms, and what you found. A transparent search limitation is acceptable; do not invent support. Your presentation must still distinguish your observations from your speculation.

Scenario 2 — Design a Future Educational Experience

Begin with the question: *What did our trial make us wish the system could do?* You may extend the trial or pursue a different opportunity. Start with a valuable learning experience; then identify the technology needed to support it.

Describe the Experience

Name the learner and goal. Describe a short sequence showing what the learner says or does, how the environment responds, and what happens next. Identify a decision, explanation, or practice opportunity that stays with the learner. Explain why a conversation matters to this experience.

Ask: *What could learners do that would be difficult to arrange routinely today?* Your proposal need not be feasible yet. A storyboard, role-play, or sample dialogue can communicate it; you are not required to build software. Label imagined responses and outcomes as hypothetical.

Directions to Stretch Your Thinking

- Investigate a fictional community by interviewing people who hold different pieces of evidence. Decide whom to question and which claims to verify.
- Revisit your own recorded explanations from earlier in a course, then discuss how your reasoning has changed.
- Describe a scientific model aloud and explore a simulation governed by explicit rules. Design an experiment that distinguishes competing explanations.
- Rewind a difficult conversation and try another response from the same starting situation. Compare what each branch reveals.
- Work across disciplines with conversational support that helps participants understand one another's terminology and assumptions.

Identify What Must Be Developed

Separate what you have demonstrated with accessible tools from additional capabilities the full experience would need. These might include consistent character knowledge, persistent records, simulation software, or coordinated participants. An AI-generated description does not establish that a feature works.

Propose How You Would Test It

Choose one uncertain assumption and a small first test. State what human learners would do afterward without AI, how their work would be checked, and what result would make you revise the design. You do not need to run this test now. Scores from simulated students describe the simulation; they do not establish that real students learned.

Future design record: learner and goal; short interaction sequence; distinctive educational opportunity; reason for voice; capabilities demonstrated and capabilities still needed; research connection if available; first evaluation and key uncertainty.

Session B — In-Class Investigation

Share and Compare

Segment	Time
Finalize and rehearse presentations	20 min
Team presentations with questions (six teams)	42 min
Compare ideas and build the class map	25 min
Reflection writing and discussion	15 min
Individual exit ticket	8 min

Stage 6 — Prepare a Five-Minute Presentation (20 minutes)

Use the first 20 minutes to finalize and rehearse a five-minute presentation of both scenarios. About four slides are enough: current use; trial and revision; future experience; research and open question. Use a saved exchange or brief role-play rather than relying on a live demonstration. Provide a readable text version of audio examples.

Time	What your team should communicate
2 minutes	Scenario 1: the real learning need, a concrete interaction, what the trial revealed, and your revision. Include why voice matters.
2 minutes	Scenario 2: the future learner experience, the new opportunity it creates, what must be built, and a first evaluation.
1 minute	The closest research example, which scenario it informs, what it supports, and one unresolved question.

Each presentation is followed by 2 minutes of questions. Six teams use 42 minutes in total. Present an unresolved issue honestly; a useful finding that an idea struggles is as valuable as a successful demonstration.

Short speaking practice: During rehearsal, explain the core idea in 60 seconds using three keywords rather than a script. Ask a teammate to repeat the takeaway and name one unclear point. Try once more. AI may suggest changes to wording or structure, but human feedback anchors whether the explanation was understood. Do not treat a transcript alone as evidence of pacing or intonation.

Stage 7 — Build a Class Map of Possibilities (25 minutes)

As you listen, record each team's current use and future experience. For each, note the learner, goal, and reason for voice. Separate observed outcomes, research findings, proposed capabilities, and evidence still needed.

After the presentations, spend 25 minutes comparing the six teams' ideas. Look for different learning activities and needs the class has not addressed. In teams, combine an insight from another presentation with your own future design and explain what the combination would let a learner do.

Questions to pursue

- Which experiences depend on immediate dialogue?
- Which future proposals create an opportunity that is difficult to provide today?
- Where would additional software, persistent records, or human oversight be needed?
- Which learners might need another way to participate?
- How could we test actual human learning?

Reflect: Where Does It Break Down? (23 minutes)

For 15 minutes, write and discuss a current use you would adopt, a future design you find promising, and a claim you cannot yet support. Explain what changed your thinking. In the final 8 minutes, complete an individual exit ticket: identify a future opportunity outside your team's direction and one first step toward testing it.

Part 3 — Post-Class Work**~2 hours**

Allow about 2 hours after the second class.

A. Finish the Team Brief (~30 minutes per person)

Finish one shared brief of 2 to 3 pages, excluding references. Include both scenarios: the current learner and goal, original and revised prompt, representative exchange or notes, and checked trial outcomes; the future learner experience, required capabilities, and evaluation proposal. Add your research evidence note and a contribution statement. Reuse your existing work.

AI may help organize the brief. Team members must verify source details, check quoted exchanges against their records, and decide which conclusions the evidence supports. Mark paraphrases as paraphrases. Include a short note about how AI assisted your work.

B. Write an Individual Reflection (~30 minutes)

Write 300 to 450 words in your own voice. Identify a current use and a future experience from other teams that changed your thinking. Adapt one to a new learner or situation. Explain what the learner must do, what would indicate success, and one material limitation. Distinguish capabilities you tried from capabilities you are proposing.

C. Try a Personal Application (~30 minutes)

Try one use from another team with ordinary course material. Record what happened and what you could do afterward without AI. Add a few sentences to your reflection. If access prevents a trial, design a concrete interaction and checking task, and label the proposal as untested.

D. Retrieval Practice / Course Review (~30 minutes)

Use the course retrieval-practice routine. Answer before asking for help. Include one question about the distinction between a demonstration and evidence that learning improved.

Everything You Turn In**Before class one (individual)**

- Preparation notes with three current-use ideas, one ambitious possibility, conversation observations, and a preliminary source or search note.

Team

- One team presentation and shared brief covering both scenarios, including the trial log, future design and evaluation proposal, checked research note or documented search limitation, and contribution statement.

Class two and after (individual)

- Class comparison notes, exit ticket, and final reflection with the personal application.

What Strong Work Shows

Assessment emphasizes the educational reasoning and honest use of evidence, rather than production polish, accent, speaking speed, or a positive result. Use this checklist before submitting:

- **Breadth:** the team considered distinct learning activities and helped the class see a different possibility.
- **Design:** both scenarios have a clear learner, goal, learner actions, and reason for voice. The future scenario describes a distinctive opportunity and identifies what must be developed.
- **Evidence:** the trial is documented, the research is checked, and the limits of each are explicit.
- **Judgment:** the current interaction was revised using observations; the future design includes a meaningful evaluation proposal. Each student can transfer an idea to a new situation.