WeTeach_AI

Teaching the Machine – Classifying Rocks with AI

Grade Level: **9-12** | Duration: **65 minutes** | Subject Area: **Science**

LESSON TABLE OF CONTENTS

This lesson was designed for the WeTeach_AI **Advancing AI Literacy Project**. The project supports the development of standards-aligned AI literacy lessons written by teachers for teachers. Additional lesson plan material, such as rubrics, answer keys, activity guides, and instructional considerations can be found here on our website.

The contents of this digital lesson were developed by the Texas Advanced Computing Center (TACC) with the support of Google.org. However, the contents do not necessarily represent the policies of Google.

Lesson Author: Travis Garcia, Computer Science Teacher

"Based in East Texas, I am a STEM educator with a background in geology and astronomy who has a passion for integrating computer science and technology in the classroom. This AI literacy lesson was designed to inspire both students and educators to embrace emerging technologies with confidence and purpose."

Lesson Description

Students shift from human classification to machine learning as they build or test an AI classifier using tools like Teachable Machine. They experience firsthand how training data impacts accuracy and bias, observing successes and failures in rock identification. Through guided discussion, students analyze why certain misclassifications occur and explore ethical questions about data sourcing and representation. By the end, they synthesize insights about the strengths and limitations of AI in scientific contexts, reinforcing the need for human judgment and accountability in technology design.

WeTeach_AI

Lesson Objectives
(formatted as "Students will be able to..." statements)

- Construct or test an AI model for rock classification.
- Evaluate the accuracy and limitations of AI classifiers.
- Explain how bias enters AI systems and propose strategies to mitigate it.

Essential Questions

1. Can AI reliably identify rocks or minerals?
2. How does dataset design influence AI performance?
3. What responsibilities do we have when designing or using AI systems?

TEKS Alignment (Texas Standards Alignment)

§112.51 Specialized Topics in Science
- **(c)(3)** Scientific and engineering practices: The student develops evidence-based explanations and communicates findings, conclusions, or proposed solutions.

§130.431 AP Computer Science Principles
- **5.C:** Explain how an effect of a computing innovation can be both beneficial and harmful.

CSTA/ISTE Alignment (National Standards Alignment)

CSTA

3B-IC-25: Test and refine computational artifacts to reduce bias and equity deficits.

3A-DA-12: Create computational models that represent the relationships among different elements of data collected from a phenomenon or process.

ISTE

1.4.c: Students develop, test and refine prototypes as part of a cyclical design process.

2.6.a: Educators foster a culture where students take ownership of their learning goals and outcomes in both independent and group settings.

WeTeach_AI

Effective Pedagogical Strategies

Curriculum provides guidance for instructors to create a classroom environment that nurtures emotional safety, encourages resiliency in the face of mistakes, and fosters a sense of belonging for all students.	Students see themselves represented in/throughout the curriculum.	Curriculum reflects and highlights the perspectives, languages, and community values of students and/or contemporary youth culture (e.g., popular video games or common student interests/activities).	The instructor provides opportunities for students to explore and give solutions to open-ended prompts.

AI Literacy Competencies

(based on TeachAI Framework)

Engaging with AI 3: Examine how predictive AI systems provide recommendations that can inform and limit perspectives.	**Designing with AI 2:** Compare the capabilities and limitations of AI systems that follow algorithms created by humans with those that make predictions based on data.
Engaging with AI 4: Explain how AI could be used to amplify societal biases.	**Managing AI 2:** Decompose a problem based on the capabilities and limitations of both AI systems and humans.

Key Terms	
Term	**Definition**
AI Model	A computer program that learns from data to make decisions or solve problems, similar to how people learn from experience.
Bias	How incomplete or skewed data can affect AI predictions.
Classification Error	An incorrect assignment of a specimen into a specific category (e.g., labeling limestone as quartz).
Ethics	Fairness and accountability in AI design and use.

Training Data	The dataset used to “teach” an AI system or model.

Launch

Engaging activity or prompt to introduce the lesson. **Estimated time: 10 minutes**

Objective: Connect prior lessons to AI classification and spark curiosity about machine learning limitations.

Materials:

- Discussion prompt slide
- Projector or screen
- Video tutorials or live demonstration of Teachable Machine or a similar app

Teacher Instructions	Sample Teacher Remarks
Recap previous lessons, e.g., *“We’ve seen how humans classify rocks—now let’s see how machines do it.”* Show Teachable Machine or another chosen app to the class and provide a demonstration. Below are tutorial videos: - Teachable Machine Tutorial 1: Gather - Teachable Machine Tutorial 2: Train - Teachable Machine Tutorial 3: Export Guide students along the setup process. Students will upload images and train/test the model. Ask students to predict what will be easy vs. hard for the AI classifier to classify the rocks. --- **Anticipated Student Outcomes** - Students recall human classification challenges. - Students predict AI strengths and weaknesses. - Students prepare to explore dataset design and bias.	*“Previously, we explored how humans classify rocks scientifically and culturally. Today, we’ll see what happens when a machine tries to do the same.* *You’ll train an AI model using Teachable Machine—or test an app like Google Lens.* *Your goal is to see what the AI gets right and wrong. Let’s walk through the process…”* ***(Pause to lead a demonstration of Teachable Machine or another chosen app. Alternatively, you can show the video tutorials for Teachable Machine.)*** *As we prepare to begin, consider this, “What do you think will be easy for AI? What might be hard?”*

WeTeach_AI

Exploration

Step-by-step student tasks, experiments, or investigations. **Estimated time: 35 minutes**

Objective: Students will build or test an AI classifier and evaluate its accuracy and limitations, connecting technical performance to bias and ethics.

Materials:

- Data table for recording Rock Turning Test results (provided in this lesson's resources)
- Devices with internet access
- Teachable Machine link (or an alternative app like Google Lens, RockD, iNaturalist)
- Training images or rock sample photos (provided in this lesson's resources)

Teacher Instructions	Sample Teacher Remarks
Conduct a Rock Turing Test. Instruct students to: - Create one class per rock type (e.g., granite, limestone, sandstone). - Upload multiple images for each class. - Click Train Model and wait for the process to complete. Direct students to test their model with new images or use an app to classify rock samples. Ensure they record results in the provided data table, noting: - Correct identifications - Misclassifications - Possible causes of errors Circulate and ask guiding questions: - "What patterns do you notice in the errors?" - "What might be missing from your dataset?" - "How could these gaps introduce bias?" Allocate roughly: - 5–7 minutes for setup and demo - 10–15 minutes for training - 10–15 minutes for testing and recording results	*"Open Teachable Machine and create a new Image Project. Add one class per rock type—granite, limestone, sandstone. Upload your images and click Train Model. Think about this: how does the data you provide shape the AI's performance?"* ***(Pause for the completion of the training process.)*** *"Now test your model with new images—or use an app to classify samples. Record successes, failures, and possible causes. Ask yourselves: What patterns do you see? What might be missing from your dataset? How could these gaps introduce bias?"* --- **Anticipated Student Outcomes** - Students create or test an AI model for rock classification. - Students identify successes, failures, and patterns of error. - Students connect dataset design to bias and ethical implication

WeTeach_AI

Whole Class Discussion

Discussion questions, teacher prompts, and expected student responses. **Estimated time: 15 minutes**

Objective: Synthesize findings and connect to bias and ethics.

Materials:

- Student data tables from “Exploration” section of the lesson
- Whiteboard or chart paper for recording patterns
- Optional: Projected images of rocks or screenshots of AI misclassifications

Teacher Instructions	Sample Teacher Remarks
Set Up the Space Arrange students for visibility of the discussion board or chart paper. Prepare columns for recording patterns (e.g., Successes, Failures, Causes of Error). Quick Share (5 min) Call on each group to share: - One success (where the AI model or app worked well). - One failure (where it misclassified). - One possible cause of error (e.g., lighting, dataset gaps). Record responses under the prepared columns. Deep Dive Discussion (10 min) Guide students through analytical and evaluative questions, as provided below. Encourage students to use sentence stems and jot notes before sharing. Highlight connections to bias, representation, and ethical responsibility.	*“Let’s start by hearing from each group. Share one success, one failure, and one possible cause of error. What stood out to you?”* ***(Pause for student share outs.)*** *“Now, let’s dig deeper. What patterns do we notice across all groups? Where did the AI struggle most? Why do you think that happened?”* ***(Pause for share outs and bridge to ethics)*** *“If AI can’t see what humans see, when should human judgment override AI? What harm could come from ignoring that?”*

Sample Discussion Questions	Sample Student Responses
1. What patterns of error did we notice?	*"Most errors happened with rocks that looked similar in color or texture."*
2. What did the datasets fail to represent?	*"Our dataset didn't include rough or weathered samples—only polished ones."*
3. What rule would you set for when human judgment should override AI?	*"If the AI is less than 80% confident or the classification affects safety or cost, humans should decide."*
4. How could adding more diverse data improve fairness?	*"It would help the model recognize rocks from different regions and conditions."*
5. What ethical risks exist if scientists rely only on AI for classification?	*"Misclassification could lead to unsafe construction materials or wrong mining decisions."*

Assessment

Formative or summative assessment tasks and criteria. **Estimated time: 5 minutes**

Objective: Reflect on AI performance and ethical implications.

Materials:

- Exit ticket handout or digital form (Google Docs, LMS, or sticky notes)
- Student data tables from "Exploration" section of the lesson
- Optional: Devices for audio/video submissions

Assessment Opportunities	Facilitation Tips
Exit Ticket (5 minutes) Provide students with the following exit ticket prompts: - *"What is one major limitation of your model?"* - *"What would you change to make it more accurate and fair?"* Providing the following submission options is recommended:	Clarify expectations: *"Your response should explain a limitation and suggest a realistic improvement. Use examples from your classifier activity." Give a clear 5-minute countdown and collect responses promptly.* Model response examples for students: - Limitation: Our model only used polished samples, so it misclassified rough rocks.

- Written paragraphs (Google Docs, LMS, or sticky notes). - Audio or video response for students who prefer verbal expression. Encourage evidence-based responses: *"Use examples from today's discussion or your classifier activity to support your answer."* If assigning outside of class, post clear submission guidelines. For example, responses should include at least one key term from today (e.g., bias, dataset, classification, accuracy). ☑️Use the rubric provided in this lesson's resources to assess student submissions.	- Improvement: Add more diverse images with different textures and lighting. Offer sentence starters for students who need scaffolds: - "One limitation of my model is ___." - "One way to improve fairness is ___." Circulate and check for understanding. Ask probing questions: - *"Why does this limitation matter?"* - *"How would your change reduce bias?"* Highlight strong examples anonymously to model quality responses. Note misconceptions (e.g., confusing AI with human judgment) for review in future lessons.