

LESSON 2

See How AI Learns: Train an Image Model

40 minutes

How Machines Learn

WHAT THIS LESSON DOES

Students train and test a browser-based image model from their own webcam examples, then explore how training data shapes what AI gets right, gets wrong, and how that links to the media they use every day.

WHAT STUDENTS SHOULD BE ABLE TO DO

- Students train, test and improve a simple image-recognition model from their own examples; explain what training data is; and connect biased or limited data to how AI content, filters and recommendations are built.

BEFORE THE BELL

- Devices with a working webcam and browser access (or one shared IWB webcam for a whole-class build)
- No accounts or sign-ins needed for Teachable Machine
- Plan how the class will save the exported model link (shared drive, email to self, or class note) for the next session
- Quiet enough lighting so hand gestures read clearly on camera

See How AI Learns: Train an Image Model

HOW THE LESSON RUNS

4 min	introduction	Start: What We're Exploring
2 min	content	Introduction
2 min	content	Open Teachable Machine
3 min	content	Create an Image Project
2 min	content	Create the Classes
4 min	content	Add Images to the Rock Class
2 min	content	Add Images to the Paper Class
2 min	content	Add Images to the Scissors Class
3 min	content	Train the Model
5 min	content	Test the Model
3 min	content	Export Your Model
2 min	content	Conclusion
3 min	content	Make Sense: What Did the Model Learn?
2 min	content	Reflect
1 min	summary	What You Covered

TEACHING MOVES

- Model the first class (e.g. Rock) on the board or IWB before pairs work independently so everyone sees clear frames, varied angles and how to drop blurry shots
- Keep the opening short; the literacy payoff comes from testing, fooling and fixing the model
- Prompt pairs to fool the model on purpose (odd angle, partial hand, cluttered background), then add better examples and retrain, that is the bias/data point made concrete
- Use media language throughout: this is how recommendation, filter and moderation systems learn from someone else's chosen examples

WHAT TO WATCH FOR

- Too few or too similar examples, model only works for one hand, one angle or one person; fix by adding varied positions, distances and (if appropriate) both partners' hands, then retrain
- Blurry or mixed-up images left in a class, scissor frames in Rock, motion blur, empty frames; fix by deleting poor samples before training
- Testing without retraining, students add images but forget Train Model again; remind them the live preview only updates after a fresh train

DIFFERENTIATION

- Support: pair stronger readers with peers; allow a whole-class build on the IWB for the first class; provide gesture photos as a visual checklist
- Extension: deliberately bias one class (only one skin tone, only left hand, only dark background), retest, and explain who the model would unfairly favour in a real media tool
- Extension: compare confidence scores across gestures and argue which class needs richer training data and why