

Create an Image Model

40 minutes

Discovering Artificial Intelligence

WHAT THIS LESSON DOES

In this lesson, you'll build an image model with Google's Teachable Machine to recognise rock, paper, and scissors hand gestures. Follow step-by-step instructions to create classes, add images, train, and test your model for future projects. Note: You will need a webcam or a camera on your computer to create the images for the image model.

WHAT STUDENTS SHOULD BE ABLE TO DO

- Understand and utilise Google's Teachable Machine to create an image model.
- Create and define classes within an image model project.
- Add and manage image samples to each class for effective model training.
- Train, test, and refine the image model to ensure accurate gesture recognition.
- Export and save the created image model for future use in projects.

BEFORE THE LESSON

- Open teachablemachine.withgoogle.com yourself first and check the school filter is not blocking it and that webcam access can be allowed.
- One device per pupil with a working webcam, and each pupil logged into whatever account they use, so camera permissions can be granted.
- Have a reliable place for pupils to save their shareable link ready, for example emailing it to themselves or a shared class note, since they need it in a later lesson.
- Decide how pupils facing the camera will have room to make clear hand gestures without crowding each other.

Create an Image Model

HOW THE LESSON RUNS

Introduction

Open Teachable Machine

Create an Image Project

Create the Classes

Add Images to the Rock Class

Add Images to the Paper Class

Add Images to the Scissors Class

Train the Model

Test the Model

Export your Model

Conclusion

TEACHING MOVES

- 1. Introduction. Tell pupils they are training a computer to recognise their hand, not programming rules for it. Flag that this model gets reused later to build a real game, so it is worth doing well.
- 3. Create an Image Project. Pause on the word class. Say it plainly: a class is one category the model learns to spot. Their model will have three, one per gesture.
- 4. Create the Classes. Check every pupil has renamed all three classes clearly to Rock, Paper and Scissors before capturing images. Mislabelled classes cause confusing results later.
- 5. Add Images to the Rock Class. Model good sampling: hand held steady, filling the frame, captured in several positions. Show them deleting a blurry frame so they see quality matters more than quantity.
- 8. Train the Model. Set the expectation that training takes a minute or two and the screen is working, not frozen. Ask pupils to watch without clicking away.
- 9. Test the Model. Have pupils test all three gestures, not just one. When a gesture is misread, point them to adding more samples to that specific class and retraining, rather than starting over.
- 10. Export your Model. Walk the room as pupils export. The two clicks (Export Model, then Upload my model) are easy to miss. Confirm each pupil has actually saved the link somewhere retrievable.

WHAT TO WATCH FOR

- Pupils think a handful of near-identical samples is enough, then are puzzled when the model misreads their gesture. Show that the model learns from variety. Have them recapture with the hand at different angles and distances, then retrain and test again.
- Pupils assume the model is wrong or broken when it misreads a gesture, rather than seeing it as undertrained. Reframe it: the model only knows what it was shown. Adding more samples to the weak class and retraining is the fix, and that is how machine learning works.
- Pupils think clicking Export Model has saved their work, and lose the shareable link. Stress that the link itself is the thing to keep. Check each pupil has copied it into a note or email before moving on.

DIFFERENTIATION

- Emerging: Sit beside pupils stuck at the webcam permission prompt, as this blocks everything that follows. Let them complete just the Rock class fully with you before working the other two classes on their own.
- Developing: Prompt them to test after training and read the confidence the model shows, rather than assuming it is finished once trained. Ask them to say back what a class is in their own words before they add images.
- Proficient: Challenge them to make the model fail on purpose, then explain why, and fix it by improving one class. Have them add samples with a different background or lighting and predict whether accuracy improves.