EdSutra

Teaching guide

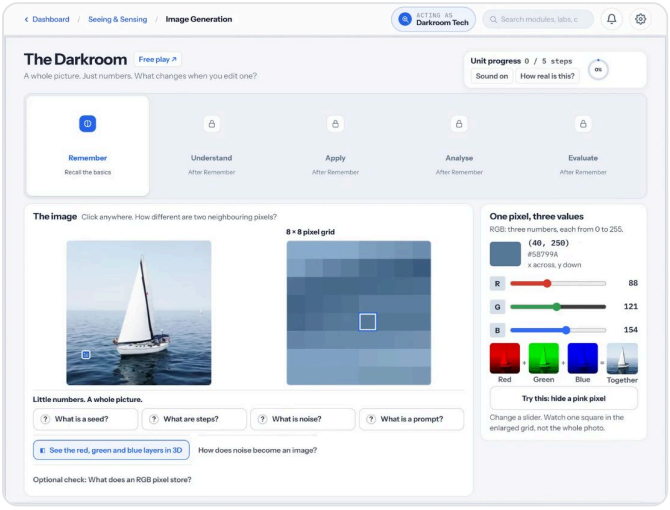
Module 02 · Seeing & Sensing · Unit 04 · Stage 2, Capabilities

Image Generation

The Diffusion Darkroom

Grades 8–11 | **30 minutes** | Role: **Darkroom Technician**

Grade 8: About age 13–14 · Year 9 in the UK · Class 8 in India



The big idea

An image maker does not paste in a photo. Fizz, a real tiny model, learned from 12 photos, then starts from random noise and removes a little at every step until a picture appears. The seed, the prompt and the steps each change the result. A convincing picture does not prove where it came from; only its records can.

The learner's job

Learners are the **Darkroom Technician** at Cape Marrow Station · Darkroom. This room makes pictures of things that never happened. People will still believe them. Their task: turn static into a picture. Then say how it was made, and whether it's new.

By the end, learners can say

- I can name the four controls of image generation — prompt, seed, step, and noise — and what each one decides
- I can read a live-computed signal/noise readout off a real diffusion noise schedule
- I can predict and check the step at which an image crosses the edge of legibility
- I can separate what a prompt fixes (structure) from what a seed fixes (detail)
- I can judge image provenance and explain why visual inspection alone is not a reliable detector

Before class

- One computer or tablet for each learner or pair, with EdSutra open (this Module is part of the full EdSutra studio).
- Open the studio once on each device before class, so the Unit still runs if the internet drops.
- Play Remember and Understand yourself, then develop the baseline sailboat (seed 1042, 24 steps) once in Apply.
- Pairs work well: one learner runs the Darkroom, the other writes down each prediction before a run.
- Best after How Machines Learn and Seeing & Sensing.
- Units open in order, and Module 1 comes first. To start a class here, switch on Open every Unit in Settings → Teachers & parents.

Words to introduce

RGB pixel: One dot of a picture, stored as red, green and blue numbers.

Noise: Random numbers for every color value; it looks like colored static.

Seed: A number that picks the random starting noise.

Steps: How many updates the model makes on the way from noise to picture.

Prompt: Words turned into numbers that steer every update.

Training: Learning patterns from example photos, done before any picture is made.

Provenance: A picture's record of where it came from and how it was made.

Running the lesson 30 minutes

Open 4 min

"Could a computer make a photo of a sailboat that nobody ever took? Where would that picture come from?"

Tell learners that Fizz, the Unit's image maker, learned from just 12 photos. Ask: will its pictures be brand new, or copies of photos it learned from? Take a show of hands and come back to it in Evaluate.

Step	What learners do	Ask the class	Listen for
1. Remember 4 min	Click two pixels of a real boat photo and read their red, green and blue values. Then open the cards "What is a seed?" and "What are steps?".	• What does one pixel store? Does anything in it say "boat"? • What does a seed pick?	A pixel is three numbers from 0 to 255. A seed picks the starting noise; it means nothing by itself.
2. Understand 5 min	Switch between Training (earlier: Fizz learns from 12 photos) and Generation (now: noise becomes a picture). Scrub one recorded run to its early, middle and final steps.	• While Fizz generates, what stays fixed? • At which step did you first see a hint of a boat?	Training happened earlier. During generation the learned rule stays the same and only the picture changes.
3. Apply 6 min	Develop the baseline (a sailboat, seed 1042, 24 steps), scrub to a middle step, change only the seed and develop, then develop the baseline again to see whether it repeats.	• What did changing only the seed do? • Did the baseline come out the same the second time? Why?	The same seed and settings repeat the same picture here. A new seed starts from different noise.
4. Analyze 6 min	Three pairs of runs, A and B, each differ in one setting: seed, steps or prompt. For each, learners name what changed, predict B, watch it develop, then say what they see.	• What happened when B had just 1 step? • Did your prediction for B match what developed?	One step gives a blur. A new prompt gives a different kind of thing; a new seed, another picture.
5. Evaluate 5 min	Judge two cases: their own print beside a training photo, and an uncaptioned photo in a made-up class chat. For each: can pixels prove the source, what next, and which fact backs it?	• Your print almost matches one of Fizz's 12 photos. Is it new? • The chat photo has no record. Does that mean an AI made it?	Pixels cannot prove where a picture came from. A missing record means the source is uncertain, not AI-made.

Steps follow Bloom's revised taxonomy, in the order the Unit presents them.

Deepen the thinking

Critical thinking	Fizz's print looked like a real photo. Besides how a picture looks, what can tell you where it came from?
AI thinking	Fizz copies because it knows only 12 photos. What would change if it learned from millions, and could it still copy?
Ethics and fairness	If an image maker copies a photographer's picture almost exactly, who should get the credit, and who is responsible?
Connect to life	Where have you seen pictures that might be AI-made: ads, posters, social media? How did you decide whether to trust them?

Misconceptions to listen for

"The AI looks up a photo while it draws."

Understand separates the two: training happened earlier. While generating, Fizz uses a fixed learned rule to remove noise, step by step.

"Every AI picture is brand new."

Fizz knows only 12 photos, and its sailboat print matches one almost exactly. Large models mostly make new pictures, but can sometimes copy.

"You can tell an AI picture by looking."

Pixels can match closely but never say where a picture came from. Check its records; a missing record means unknown, not AI-made.

In real life the card learners see at the end of the Unit

- Many AI image makers turn a written prompt into a picture by cleaning up random noise, step by step.
- Posters, ads and magazine covers can now use AI-made pictures that show things that never happened.
- Some photo apps can add or erase things in your pictures with generative AI.

Try it this week. Next time a surprising picture is shared with you, ask an adult to help find where it first appeared before you believe or share it.

Watch out for. You can't tell how a picture was made by looking at it. Fizz's prints began as static, yet looked real, and some copied pictures it learned from.

Talk about it. Should AI-made pictures always carry a label saying so? Whose job is that?

Exit ticket

1. What does a seed decide?

2. What does Fizz learn in training, and what happens during generation?

3. Why can't the pixels alone prove where a picture came from?

Going further

For homework, use the "Try it this week" activity from the In real life card. In class, learners can press Free play, ask Fizz for a dragon (it has never seen one), develop, and explain the result.