EdSutra

Teaching guide

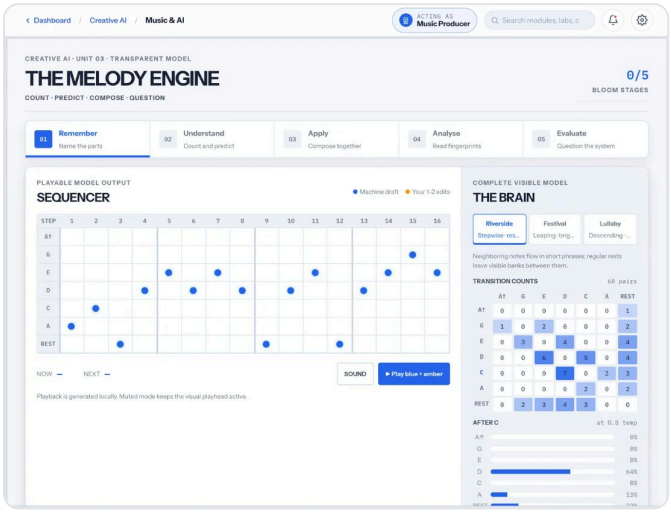
Module 07 · Creative AI · Unit 03 · Stage 2, Capabilities

Music & AI

The Melody Engine

Grades 7–11 | **30 minutes** | Role: **Algorithmic Music Producer**

Grade 7: About age 12–13 · Year 8 in the UK · Class 7 in India



The big idea

The Melody Engine is a real, tiny music model: a table counting which note follows which in a songbook. It drafts a tune by sampling each next note from those counts; the seed makes a run repeatable and temperature reshapes the chances. A tune carries its songbook's fingerprint, so credit, consent and honest labels still matter.

The learner's job

Learners are the **Algorithmic Music Producer** at The Long Room · Sound Desk. The studio needs a tune it can explain. It must replay and show its credits. Their task: train the engine, shape its draft, and show who made every note.

By the end, learners can say

- I can identify the sequence, transition counts, probabilities, seed, and temperature that drive a next-note model
- I can predict a next note from visible counts and explain how temperature reshapes its probabilities
- I can train, generate, inspect, listen to or watch, and revise a seeded machine melody
- I can distinguish songbook fingerprints using rests, leaps, repetition, and transition evidence
- I can evaluate an AI-music workflow using authorship, consent, representation, and model evidence

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.
- Speakers or headphones help, but sound is optional: every step works muted, with lit notes and counts on screen.
- Play the opening and the first two steps yourself, so you know where learners will click.
- Pairs work well: one learner drives, the other is the Music Producer who tracks which notes were the machine's and which were theirs.
- Best after Generative AI and Algorithmic Art.
- 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

Songbook: The set of example melodies the engine learns from.

Transition count: How many times one note was followed by another in the songbook.

Markov chain: A model that picks the next item using only the current one.

Temperature: A dial that sharpens or flattens the chances without changing the counts.

Seed: A number that replays the same random choices, so a run can repeat.

Provenance: A record of where something came from and who changed it.

Consent: Permission from the people whose music is used, given before it is used.

Running the lesson 35 minutes

Open 5 min

"A melody arrived and the credits are blank. A table of numbers wrote most of it. How would you find out who made what?"

Before anyone clicks, ask: if a machine learns only from slow, calm songs, what will its new tunes sound like? Write the guesses on the board and check them in Analyze, where learners read melody fingerprints.

Step	What learners do	Ask the class	Listen for
1. Remember 4 min	Learners inspect the engine's four parts (songbooks, the Brain's count table, the seed and temperature dials, and the sequencer), then answer: when you press Count & train, what changes?	• What does one square in the Brain's table count? • What do the blue notes on the sequencer mean?	Each square counts how often one note was followed by another in the songbook. That table is the whole model.
2. Understand 7 min	Learners do four checks: predict which note follows C in Riverside, predict whether two runs at temperature 0 match, run the slow trace of all 16 notes, and say what temperature changes.	• Which note follows C most often, and how did the counts show it? • What did temperature change, and what stayed the same?	D follows C most often in the Riverside counts. Temperature reshapes the chances; the counted table stays the same.
3. Apply 8 min	Learners pick a songbook, press Count & train, generate a blue draft, and play it at temperature 0.0 and 1.3. Then they change one or two notes, shown in amber, and save the melody.	• How did the 0.0 take sound different from the 1.3 take? • Why did you change those notes, and what did the A/B listening show?	Low temperature sticks to the most common moves; high temperature takes rarer ones. Amber keeps the human edits visible.
4. Analyze 5 min	Learners study three mystery melodies from the same engine. Using counts of rests, wide leaps, repeats and returns to A, they name each one's training songbook and pick the clues.	• Which clue gave the songbook away? • Why read the counts instead of the melody's mood? • Look back at our guesses on the board. Were we right?	Many wide leaps and few rests point to Festival; frequent rests and returns to A point to Lullaby.
5. Evaluate 6 min	Learners take a position on three cases (authorship, consent and representation) and back each with at least two evidence cards. The feedback explains any position the facts do not support.	• A song archive is online. Does that mean a model may train on it? • Who shaped the finished melody: the machine, the learner, or both?	Being online is not permission. The model sampled the draft; the learner chose and edited it, so credit is shared.

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

Deepen the thinking

Critical thinking	You could name a mystery tune's songbook from its counts alone. What does that say about how new a generated tune really is?
AI thinking	Temperature reshapes the chances but never changes the counts. Where else might a dial like this matter, such as a story writer or a chatbot?
Ethics and fairness	A model trained on one community's songs sells tunes in that style. What should the community get: a say, credit, payment? Why?
Connect to life	Which songs taught you your favorite style of music? How is that like, and unlike, a songbook training the engine?

Misconceptions to listen for

"The machine understands the song."
The whole model is a table of next-note counts. It remembers only the current note, and the Brain panel shows everything it knows.

"Temperature teaches the model something new."
Temperature only sharpens or flattens the chances in one row. The Understand check shows the counted table stays exactly the same.

"If a song is online, a model may train on it."
The consent case in Evaluate: being online is not permission. Ask the rights holders or the community, and record every source.

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

- Some apps and games make background music on the fly by predicting what sound comes next.
- AI songs that sound like a famous singer raise a question: did that artist agree to be copied?
- In music class, you pick up a style from songs you practice; a model's style comes from its training songs.

Try it this week. Hum a tune where you choose each note only from the last one. Then hum a real song. What does the one-note rule miss?

Watch out for. A music model doesn't understand the song; it follows patterns from its training. Track which notes were yours, and whose music it learned from.

Talk about it. If you change two notes in an AI-made tune, is it your song? What would make it yours?

Exit ticket

1. What does the Brain's table count?

   _____

2. What does temperature change, and what does it leave alone?

   _____

3. Name one clue that tells you which songbook trained a tune.

   _____

Going further

In the Sound and Music Lab, learners play chords with their hands (camera optional; mouse and keyboard work too). Ask: which steps follow fixed rules, and is any step a model that learned from examples? Saved melodies return as exhibits in Machine Creativity.