EdSutra

Teaching guide

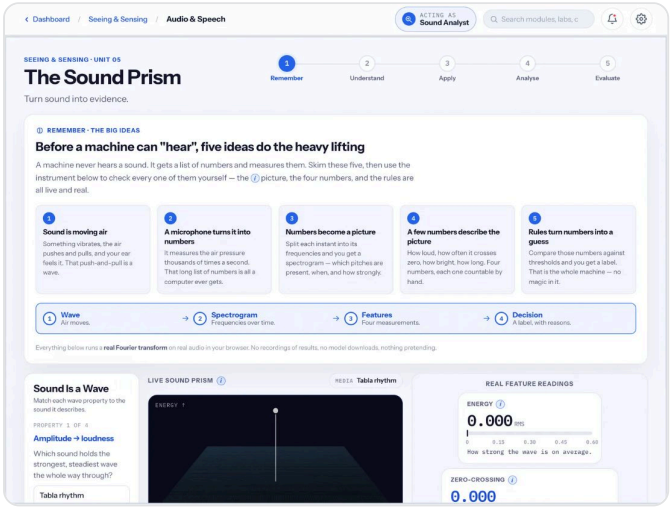
Module 02 · Seeing & Sensing · Unit 05 · Stage 2, Capabilities

Audio & Speech

The Sound Prism

Grades 7–11 | 25 minutes | Role: **Sound Analyst**

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



The big idea

A machine never hears a sound the way we do. It gets a long list of numbers, splits them into frequencies, and measures four things: how loud, how often the wave crosses zero, how bright, how long. A visible rule table turns those numbers into a label, so you can point at exactly why it gets a sound wrong.

The learner's job

Learners are the **Sound Analyst** at Cape Marrow Station · Listening Room. A microphone caught something. Now a machine has to say what it was. Their task: split a sound into what a machine can measure. Then name what decided it.

By the end, learners can say

- I can match amplitude, frequency, duration, and brightness to what you hear
- I can read a spectrogram as frequency × time × energy, and explain why it is a measurement rather than what a machine hears
- I can predict which feature meter a sound will move most, then check it against live measurements
- I can produce a counterexample that defeats the visible thresholds and explain the failure from the numbers
- I can decide which listening tasks four coarse features can and cannot support

Before class

- One computer or tablet for each learner or pair, with EdSutra open (this Module is part of the full EdSutra studio).
- Headphones for each device, since the Unit plays real sound clips. A microphone is optional: only Apply offers it, and nothing needs it.
- Open the studio once on each device before class, so the Unit still runs if the internet drops.
- Play the Remember and Understand steps yourself, so you know the sounds and where learners will click.
- Best after Seeing & Sensing and Uni-modal AI.
- 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

Amplitude: How far a sound wave swings; we hear it as loudness.

Frequency: How many times a second a wave repeats; we hear it as pitch.

Spectrogram: A picture of which frequencies are present, when, and how strongly.

Zero-crossing rate (ZCR): How often the wave crosses its center line; high for hissy sounds.

Brightness: Where a sound's energy sits on average, from low to high frequencies.

Rule table: Limits on the numbers that decide the label, such as TONE-LIKE.

Counterexample: A case that breaks a rule and shows where it fails.

Running the lesson 30 minutes

Open 4 min

"You can tell a drum from a flute with your eyes closed. What would a machine have to measure to do the same?"

Before anyone clicks, ask: could four numbers (how loud, how often the wave crosses zero, how bright, how long) tell a tabla drum rhythm from a bansuri flute melody? Vote yes or no, then return to it in Analyze.

Step	What learners do	Ask the class	Listen for
1. Remember 4 min	Match four wave properties to the clip that shows each: amplitude (loudness), frequency (pitch), duration (pattern in time) and timbre (brightness), using a steady tone, a bansuri flute, a tabla rhythm and white noise.	• Which sound's pitch kept sliding up and down, and why? • Why does white noise sound like a hiss?	Frequency is how often the wave repeats, heard as pitch. White noise spreads its energy across every frequency at once.
2. Understand 5 min	For each sketched spectrogram (a flat line, a vertical stripe, a full block, a rising curve), guess which sound made it, then play it and watch the live prism build the real picture.	• Why does a single drum hit look like a vertical stripe? • Is the spectrogram what the computer hears, or a measurement?	A hit puts energy into every frequency at once, then fades. The spectrogram is a measurement, not hearing.
3. Apply 6 min	Pick a sound and predict which meter will move most (energy, zero-crossing rate or brightness). Then play it and read the four real numbers and the rule table's label. At least three sounds.	• Which reading surprised you most? • Which number separates white noise from the low tone?	Zero-crossing rate: white noise crosses zero far more often than the low tone, 0.502 against 0.009.
4. Analyze 6 min	Pick a sound the rule table gets wrong (the tabla rhythm, filtered noise or the spoken word "seven"), play it again, and choose the measured reason it failed.	• Why was the tabla rhythm called a TONE? • Look back at our vote. Can these numbers tell a drum from a flute?	Its zero-crossing rate, 0.019, passes the tone rule's under-0.10 test. A rule fires on numbers, not on meaning.
5. Evaluate 5 min	Judge listening jobs as "Can do it" or "Cannot" for these four numbers: telling a drum hit from a hum, which word was said, who was speaking, and more. Each answer explains why.	• Why can't four numbers tell which word someone said? • What would a machine need to tell happy from upset?	Words, voices and feelings live in detail over time that four averaged numbers never look at. More rules won't fix it.

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

Deepen the thinking

Critical thinking	A filter turned noise into something the rule table called a tone. How would you test a sound detector before trusting it?
AI thinking	If you could add one more number to the four, what would you measure, and which mistake would it fix?
Ethics and fairness	A voice assistant understands some accents worse than others. Who is affected, and whose job is it to find out and fix it?
Connect to life	Which apps or devices around you listen? What do you think they are measuring?

Misconceptions to listen for

"The AI hears the sound like I do."

The machine only gets four numbers and a frequency picture. The live readings show everything it has to go on.

"If it labels a sound, it understands it."

The spoken words "one" and "seven" land in the same bucket as a flute note. Sorting by numbers is not understanding.

"A few more rules would fix it."

No single extra limit separates every word from every flute. Speech needs different features, such as how the sound changes over time.

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

- A music app that names a song from a few seconds of sound matches patterns in its frequencies.
- Some security alarms listen for breaking glass by measuring the sound's frequencies, not by understanding it.
- A video call's noise filter tries to tell your voice from a barking dog or fan, and sometimes slips.

Try it this week. Hum one note quietly, then loudly, then higher. Which change was loudness, and which was pitch? Then compare a smooth "ooo" with a hissy "sss".

Watch out for. Sorting a sound isn't understanding it. A rule table put two different spoken words in the same bucket, and more rules wouldn't fix that.

Talk about it. Which sounds around you would be hardest for a machine to tell apart, and why?

Exit ticket

1. What does a spectrogram show?

   _____

2. Name the four numbers the Sound Prism measures.

   _____

3. Give one listening job these four numbers cannot do, and say why.

   _____

Going further

In the Computer Vision Lab, the Audio Waveform example draws live microphone sound as bars and a wave line (a microphone is needed). Learners hum, hiss and clap, and say which of the four numbers each sound would change most.