Status: draft — grounds in `handoff/context/c4sight-schools-positioning-and-pricing.md` (prices, taster rules, credibility set, safety promise), `handoff/context/c4sight-schools-sessions-and-lessons.md` (Spine A, upper primary), `handoff/context/c4sight-primary-lesson-plan-years-5-6.md` (the series' core lesson and quality bar), `deliverables/activities/interactive-activities-and-tool-spec.md` (A.1 Guess the Rule, adapted for teacher-solo delivery), and `deliverables/tangibles/curriculum-alignment-matrix.md` (Years 3–4 and 5–6 codes, verified). Public-facing once the contact block is resolved. Strip this Status line and the closing "Open questions" section before PDF.

---

# Who Taught the Machine?

**AI isn't magic. It learns from the examples people choose — which means people are in charge of it.**

**Upper primary · Years 3 to 6**

**Free preview — C4Sight**

---

## What this series is

**Who Taught the Machine?** is C4Sight's flagship AI incursion series for Years 3 to 6. In one session, a class trains a real machine-learning tool themselves, watches it be confidently wrong, works out exactly why, and walks out owning an idea most adults have not been given: an AI does not "know" things — it learned patterns from examples that people chose, and the people are responsible for what it does.

Your students already use tools with AI folded inside them, and most of what they have absorbed about it is either magic or menace. This series replaces both with a correct, working mental model — one they build themselves, from evidence, in the room.

**The arc is fixed, and everything in the session serves it:**

1. **It knows nothing.** A blank classifier goes up on the big screen. It genuinely cannot tell a cat from a dog. A computer starts out empty.
2. **We teach it.** The class supplies the training — students hold picture cards to a camera while the room calls the categories, and confidence bars move as it learns in real time. That is supervised learning, embodied, no jargon required.
3. **We break it.** Shown a deliberately in-between object, it answers with total confidence — "I think this is a CAT, 96% sure" — and it is wrong. The room laughs. Then it goes quiet, because something just became visible.
4. **Why did it break?** Because it only knows what we fed it. It is not thinking; it is matching patterns from the examples we chose. And who chose them? The class did.
5. **So who is the boss?** Fair in, fair out. If a system is taught unfairly or incompletely, it behaves that way — and that is on the humans who taught it, not the machine. Which puts the class in charge, and hands them the job that comes with it: ask questions, and check its work.

**The one big idea:** AI is not magic and does not "know" things. It learns patterns from examples that people give it, so it is only as good, and only as fair, as what it is taught. People are in charge of it, not the other way around.

The moment the series is named for — the machine being *confidently wrong* — carries a second lesson your students will use for the rest of their lives: a source can be very sure of itself and still be mistaken. Ten-year-olds who have watched that happen, and can explain why, are harder to fool. The fairness beat then makes it real with true, age-appropriate examples: a voice assistant trained mostly on some accents struggles to understand others; a system taught only one breed of dog may not recognise the rest.

**Who is in the room.** Two presenters, and the pairing is the point: a working AI practitioner — he runs a Perth studio building AI automation for commercial clients, studies law at UWA, and produces C4Sight, a public AI explainer series — alongside a TRBWA-registered teacher who runs the hands-on activities, the pace, and the classroom translation. One brings the real machine and the honest answers; the other makes sure every student in the room comes with it.

**The safety promise, stated plainly.** This is how every C4Sight session runs, without exception:

- No student devices. Everything hands-on is unplugged or driven by the presenter.
- No student data is entered into any system. Ever.
- No student images are captured, saved, or shown.
- A school staff member is present throughout.
- The camera frames the objects and cards the class holds up — never students' faces.
- The tool runs on the presenter's own device and hotspot, never the school network, and the classifier runs fully offline.

---

## Run this yourself: Guess the Rule

You do not need us in the room to give your class the core of this series. Guess the Rule is the anchor unplugged activity from the series' library, adapted so you can run it on your own in 15 to 20 minutes — and it needs no materials at all. It is supervised learning acted out by the class: one student learns a rule purely from labelled examples, with nobody ever stating the rule. That is exactly what a machine-learning system does, and by the end your class will have felt it from the inside.

| | |
|---|---|
| **Time** | 15–20 minutes (two rounds plus debrief) |
| **Group** | Whole class; one volunteer per round |
| **Year levels** | Years 3 to 6 |
| **Materials** | None required — optional picture cards below |
| **What it teaches** | A "learner" — human or machine — works out a rule only from the examples it is given. Good examples make a good rule; missing examples make a confident mistake. |

### Materials (optional, but they help)

The activity runs on voices alone. Cards make the examples visible to the whole room and keep younger or quieter students in the game.

- **A set of 12–16 picture cards** matched to the rules you might use. For "animals that can fly / cannot fly," a printable set would be: parrot, eagle, owl, magpie, bat, butterfly (flyers) and dog, cat, snake, goldfish, kangaroo, echidna (non-flyers) — plus the two trick cards below.
- **Two trick cards, kept aside: a penguin and an emu.** Birds that do not fly. These power the moment the activity is really for.
- Printable suggestions: one clear image and its name per quarter-A4 card, bold and uncluttered. Clip art or your own drawings both work; a laminated set will serve for years.

### Before you start

Choose a spot where the volunteer can stand out of earshot — the back of the room, facing away, is enough if the class can whisper. Agree a quiet signal: the secret only works if the guesser cannot hear it being chosen. Have a couple of backup rules in your pocket in case the class picks one that is too hard or too leaky ("things that are red / not red" and "animals that live in water / on land" are reliable).

### How to run it

1. **Set it up (about 2 minutes).** One volunteer becomes "the computer" and moves out of earshot. Tell the class the job: *"We are going to teach the computer a rule — but we are never allowed to say the rule. Only examples."*
2. **Choose the secret rule (1–2 minutes).** The class quietly agrees a simple rule that sorts things into two groups. Keep it to **one attribute for Years 3–4** ("animals that can fly / cannot"); a **two-part rule is a good stretch for Years 5–6** ("small AND can fly").
3. **Feed labelled examples (3–4 minutes).** The volunteer returns. Students give examples one at a time — each example *with its group, never the rule*: "A parrot is in Group 1. A dog is in Group 2. An eagle is in Group 1." Hold cards up as they are called if you made them. Feed at least six before anyone asks for a guess — a computer needs plenty of examples too, and that is worth saying.
4. **Test the computer (2 minutes).** Ask the volunteer either to **state the rule**, or to **predict the group** for a brand-new example: "Where does a magpie go?" If they nail it, congratulate the class: *"You trained it well."*
5. **Now the trick case (2–3 minutes).** Before anyone relaxes, produce the penguin. *"Where does a penguin go?"* A volunteer trained on obvious flyers will very often answer confidently — and wrongly. That is not a failure of the game; **it is the whole point**. They were trained on clear examples and met an unclear one, and they were sure anyway. If your volunteer gets it right, celebrate the near-miss instead: *"What example could have tricked you? Why?"* — the discussion lands in the same place.
6. **Round two (3–4 minutes).** New volunteer, new rule, run it faster. Watch the class get sharper about which examples teach best — they are now thinking like the people who train real systems.

### The spoken teaching point — say it out loud, exactly here

> "You just learned a rule from examples, without anyone ever telling you the rule. That is exactly what a computer does when people train it. And if we had only given you red birds, you might have guessed the wrong rule entirely — the machine has the same problem. It only knows the examples it was given, and it is confident anyway."

**Debrief questions**

- "How many examples did you need before you were sure? What made the good examples good?"
- "What made the penguin hard? Were you still confident? What does that tell us?"
- "If we had chosen unfair or incomplete examples on purpose, what would our 'computer' have learned?"
- "Real AI tools learned from examples too. Who do you think chose them — and what should those people be careful about?"

That last question is the doorway to the whole series. Let the class sit with it.

**If it wobbles**

- **The volunteer cracks the rule in two examples:** do not move on — test the trick case anyway. Under-trained and confident is the lesson; lean into it.
- **The class leaks the rule:** reframe instantly. *"Now our computer was told the answer — that's not learning, that's copying. Let's do it properly."*
- **A struggling guesser:** offer a prediction frame — *"I think Group 1 means ___ because ___"* — and feed more examples before asking again. More examples make a better model, which is, again, the lesson.

Run only this, and your class has genuinely met supervised learning, training data, and the confident mistake. That is real value, and it is yours to keep.

---

## What the full incursion adds

The activity above is the real pedagogy — it is what our own presenters run mid-session, and we are giving it away because it works. Here is what changes when C4Sight comes to your school.

**The machine, live.** In the incursion, the "computer" is not a volunteer — it is a genuine image classifier on the big screen. Your class trains it themselves: students hold cards to a camera, the room calls the categories, and confidence bars sharpen in real time as it learns. Then it meets a deliberately in-between object and is confidently wrong, with a big sure number attached. Everything your class just felt in Guess the Rule happens again in front of them, done by a real machine — and the quiet after the laugh, when a ten-year-old explains bias back to the room in their own words, is what this series exists for.

**Two presenters, working as a pair.** A working AI practitioner drives the tool and takes the hard questions honestly — including "how does it really work?" and "should we trust it?" — while a registered teacher runs the activities, the pace, and the differentiation, and watches for who is and is not following. Neither half works as well alone; together the session holds a whole class for an hour.

**The full arc, with room to go deeper.** The core 60-minute lesson, a 45-minute rotation version for whole-year-level days, and a 90-minute deep-dive that adds the Fairness Sort — a second hands-on activity that makes bias physical — plus a vetted, framed short video and a "make it fairer" discussion. Differentiation is built in across Years 3–4 and 5–6, with support stems for students who need them and extension questions for the ones who race ahead.

**The teacher resource pack.** Booked sessions include a pack you keep: a pre-session primer with the handful of words worth introducing first (data, examples, training, fair), a teacher-facing run sheet, a

no-technology follow-up activity for afterwards, the vetted video links, and the full descriptor-by-descriptor curriculum alignment mapping.

**The safety promise, built in.** The tool is designed around the promise rather than promised around the design: the camera frames objects only, nothing is recorded or saved, and the classifier runs fully offline on the presenter's own device and hotspot.

---

## Curriculum connections (Australian Curriculum v9)

An incursion connects to and enriches the curriculum — it does not deliver it, and we will not pretend otherwise. A one-hour session cannot teach a content descriptor to the achievement standard; what it gives you is a vivid, correctly-modelled encounter your own teaching can build on. The strength ratings below are deliberately conservative, and every code is checkable. (These are the same codes SCSA implements in WA.)

### Years 3–4

| Code | Descriptor (AC v9, verbatim) | How the series connects | Strength |
|---|---|---|---|
| AC9TDI4K03 | recognise different types of data and explore how the same data can be represented differently depending on the purpose | The heart of the fit: a held-up card becomes image data to the camera, then a pattern, then a confidence number, then a category label — the same object represented four ways, live. | Direct |
| AC9TDI4K01 | explore and describe a range of digital systems and their peripherals for a variety of purposes | The webcam is a textbook peripheral: students watch input hardware feed a system and can describe its job in the pipeline. | Enriches |
| AC9TDI4P02 | follow and describe algorithms involving sequencing, comparison operators (branching) and iteration | Guess the Rule — in this pack and in the session — is a comparison and branching algorithm run by humans, across repeated examples. | Enriches |
| AC9TDI4P05 | discuss how existing and student solutions satisfy the design criteria and user stories | The "why did it break? is it right?" segment is a genuine evaluation: the class judges whether the tool met its job and pinpoints why it did not. | Enriches |

**Years 5–6**

| Code | Descriptor (AC v9, verbatim) | How the series connects | Strength |
|---|---|---|---|
| AC9TDI6K03 | explain how digital systems represent all data using numbers | The tool turns what the camera sees into numbers and patterns; the confidence bars are the data as numbers. Students see, rather than get told, that the machine's world is numbers. | Direct |
| AC9TDI6P06 | evaluate existing and student solutions against the design criteria and user stories and their broader community impact | The fairness beat is exactly this: the class evaluates the tool not only on "did it work" but on who it might treat unfairly and why — broader community impact, made concrete. | Direct |
| AC9TDI6K04 | explore how data can be represented by off and on states (zeros and ones in binary) | The session establishes that everything the machine perceives becomes numbers — the exact hook a teacher needs to then teach binary. The incursion sets it up; your classroom lands it. | Enriches |
| AC9TDI6P02 | design algorithms involving multiple alternatives (branching) and iteration | Guess the Rule embodies a branching rule refined over repeated examples, and "how could we make it fairer?" pushes students to modify it. | Enriches |

The heaviest curriculum work sits in three general capabilities (which carry no descriptor codes in v9): **Critical and Creative Thinking** (the break moment is engineered so the questions — "wait, is it right?", "why is it so sure?" — come from the students), **Ethical Understanding** (fairness, responsibility and bias at an age-appropriate level: the machine was not unfair, the people who taught it were), and **Digital Literacy** (a source being sure of itself is not the same as being right). The session is also built to model the **Australian Framework for Generative AI in Schools** — its transparency, fairness, accountability and privacy principles are visible in how the session runs, which is a conversation starter we are happy to have with your leadership team.

The free activity in this pack does honest curriculum work on its own — Guess the Rule is AC9TDI4P02 and AC9TDI6P02 in action. The full descriptor-by-descriptor mapping, including the general-capability continua, ships in the teacher resource pack with a booking.



---

## Bring Who Taught the Machine? to your school

The series is delivered as an incursion, sized to fit how your school actually runs. All prices exclude GST. Perth metro; travel is free within 30 km of the Perth CBD.

| Product | Format | Capacity | Price |
|---|---|---|---|
| **Taster** | 30 min, one class or assembly | Any | **Free** |
| **Single Year Group** | 2 × 60 min sessions, one presenter | Up to 60 | **$950** |
| **Half Day** | 3 sessions | Up to 90 | **$1,750** |
| **Full Day** | 5 sessions + free 45 min after-school staff briefing | Up to 150 | **$2,850** |
| **Term Program** | 6 weekly sessions with one year group, ends in a showcase | One year group | **$7,500** |

Single Year Group, Half Day and Full Day bookings include the teacher resource pack — the pre-session primer, post-session extension activities, and the curriculum mapping. A Half Day covers one year level in rotation; a Full Day covers a whole primary school's upper years, and the free after-school staff briefing means your teachers get something out of the day too.

**The free Taster, and how it works.** We will bring the tool and run the real thing for one class or an assembly, free, for 30 minutes. Two honest rules, so you know exactly what you are agreeing to: it is **one per school, ever**, and it **only runs with a decision-maker in the room** — Principal, Deputy, or your Digital Technologies coordinator. The Taster exists to show you the real thing, and it ends with us booking a conversation about whether a paid visit makes sense for your school. If it does not, no hard feelings — your students still got the lesson.

**Teaching Pre-primary to Year 2?** The lower-primary series is **The Computer Puppy** — its free preview pack is a separate download.

**Book, or ask us anything:**

> → **caleb@c4studios.com.au** · **c4studios.com.au**

*C4Sight · Perth, Western Australia*

---

## Open questions for Caleb (internal — remove before PDF)

1. **Contact block.** The CTA dead-ends into the placeholder until the brand/channel decision lands. This is the only thing blocking publication.
2. **Printable card sheet.** As with the Series 1 pack: the copy tells teachers how to make an optional card set, which works, but a designed one-page printable (flyers, non-flyers, penguin, emu) bundled with the PDF would lift the download's perceived value — worth a later design pass?
3. **Presenter naming.** I used roles and the verifiable credibility set, not names, since the public brand presentation is unsettled. Confirm whether packs should name the presenters.

*C4Sight is a C4 Studios venture · Caleb Scott · ABN 26 477 523 297 · Perth, WA*