techskills

TEACHERS

Data and AI in the classroom: the teacher pack

Activities, answers and the evidence behind every claim

Every figure was checked against its source on 1 November 2026. Free to use in your classroom.

About this pack

This is the printable companion to the Data and AI in the Classroom page on techskills.org. It holds everything you need at the front of a class: how to run each activity, the questions to ask, the numbers the games will show so you can check them in advance, and the evidence behind every claim, with its source.

Everything below works in a normal lesson. Most of it needs one projector. Three activities can be run with no devices at all: the starter, the discussion on the 2020 grades, and the pipeline if you print the table from the pack. Timings are our suggestions from how long each takes to run, so adjust them to your group, and every activity has teacher notes, answers and discussion questions in the downloadable pack. The numbers in the games are invented for teaching, and each activity says so.

Every dataset, machine and sentence in these activities is made up. Nothing a student types into the interactive activities is stored or sent anywhere, and nobody is identified by any of the examples.

For	Teachers and lecturers in schools, colleges and training providers. Written for England; see the curriculum section
Activities	11
Sources	20, each opened and read on 1 November 2026
Produced by	TechSkills, techskills.org

The case for teaching it

We could just say that AI is changing everything and leave it there, and plenty of people do. But you are the one standing in front of a room full of students who will check, so here is what the evidence really says, including the awkward bits, because the awkward bits are where the good lessons live.

The field is real and it is growing. The latest government study counts thousands of UK AI companies, a workforce of well over a hundred thousand and revenue in the billions, and then, to its credit, tells you not to take the growth figures at face value, because part of the rise is down to the researchers finding more firms, not only firms getting bigger. That is a lesson in data in itself, and we would use it in class. The skills shortage is real too, but it is a particular shape: employers say the hardest roles to fill are the experienced ones, because candidates lack work experience, not because nobody is interested.

And the door is wider than the stereotype. Graduate routes are the most common way in, but apprenticeships have gone from a rarity to something nearly a fifth of organisations use for junior roles, most organisations also hire people who have come from a job that had nothing to do with AI, and the survey reports that psychology and philosophy are turning up alongside computer science. It is also honest to say that only a fifth of AI roles are held by women, and that has been falling. So the message for your students is not 'this is easy' and it is not 'this is for someone else', it is that there are lots of ways in, real skill is noticed, and the field would be better with more of them in it.

The evidence

Fact	Condition	Source
The 2025 AI sector study identified 10,488 AI companies in the UK, up 79% on 2024.	The report says part of the rise is better coverage rather than growth alone: 4,580 of the 5,862 firms counted in 2024 carried over and 5,908 were newly added. It says comparisons over time should be interpreted accordingly, and that its estimates are experimental and not official statistics.	Department for Science, Innovation and Technology (research by Perspective Economics, with Beauhurst and Ipsos), Artificial Intelligence sector study 2025 (Published 29 October 2026)
An estimated 130,400 people were employed in AI-related roles in 2025, an increase of 51% on 2024.	The report counts only the estimated AI-related share of employment in firms that do AI as part of a broader business, and its estimates are experimental.	Department for Science, Innovation and Technology (research by Perspective Economics, with Beauhurst and Ipsos), Artificial Intelligence sector study 2025 (Published 29 October 2026)
AI-related revenues increased by 41% to £33.6 billion in 2025, and AI-related gross value added by 58% to £18.6 billion.	Experimental estimates. The report notes that year-on-year changes reflect improved methods and wider terminology as well as genuine market growth.	Department for Science, Innovation and Technology (research by Perspective Economics, with Beauhurst and Ipsos), Artificial Intelligence sector study 2025 (Published 29 October 2026)

Fact	Condition	Source
London, the South East and the East of England accounted for 72% of the registered office locations of the AI companies in the study.	The report notes that the sector's footprint extends further across the UK's nations and regions than registered offices show.	Department for Science, Innovation and Technology (research by Perspective Economics, with Beauhurst and Ipsos), Artificial Intelligence sector study 2025 (Published 29 October 2026)
80% of the businesses surveyed said the availability of skills and talent had a significant or moderate impact on their business over the past year.		Department for Science, Innovation and Technology (research by Perspective Economics, with Beauhurst and Ipsos), Artificial Intelligence sector study 2025 (Published 29 October 2026)
97% of the organisations surveyed identified at least one skills gap in the AI labour market, and 57% reported a technical skills gap.	119 organisations responded, about a 3% response rate, and the report warns that organisations choosing to take part may differ from those that did not. Treat the figures as indicative.	Department for Science, Innovation and Technology (research by Gardiner & Theobald), AI Labour Market Survey 2025 report (Published 28 January 2026)
35% of the organisations had AI vacancies they found hard to fill, and the report gives lack of work experience (31%) and lack of technical skills or knowledge (30%) as the main reasons.	AI professionals and AI leaders were the hardest roles to fill, in line with lack of work experience being the main factor.	Department for Science, Innovation and Technology (research by Gardiner & Theobald), AI Labour Market Survey 2025 report (Published 28 January 2026)
In the 2025 survey, 60% of organisations used career-starter routes such as a graduate position to bring people into junior AI roles, 19% used apprentices and 18% used internships. In 2020 the figures were 42%, 3% and 11%.	Organisations could use more than one route, so the figures do not add up to 100%.	Department for Science, Innovation and Technology (research by Gardiner & Theobald), AI Labour Market Survey 2025 report (Published 28 January 2026)
The most common way into AI roles was from a previous AI role (77% of organisations), and 74% also hired people from a previous role that did not involve AI.	The report suggests the 74% may mean people were retrained from a different field or came from an untrained entry-level position.	Department for Science, Innovation and Technology (research by Gardiner & Theobald), AI Labour Market Survey 2025 report (Published 28 January 2026)
Computer science remains the dominant qualification in AI roles, and the report says roles increasingly include social sciences such as psychology and philosophy. The share of businesses employing data science professionals rose from 48% to 66% between 2020 and 2025.		Department for Science, Innovation and Technology (research by Gardiner & Theobald), AI Labour Market Survey 2025 report (Published 28 January 2026)
88% of organisations used on-the-job training instead of structured education and training programmes, and only 13% of graduate schemes included AI training.		Department for Science, Innovation and Technology (research by Gardiner & Theobald), AI Labour Market Survey 2025 report (Published 28 January 2026)

Fact	Condition	Source
Women accounted for 20% of AI roles in 2025, a drop of 4 percentage points since 2020.		Department for Science, Innovation and Technology (research by Gardiner & Theobald), AI Labour Market Survey 2025 report (Published 28 January 2026)

Your school or college is in the figures too

This is the part that makes a staff room go quiet. Surveys of young people and teachers in 2026 suggest that most of your students already use generative AI, that around six in ten of the older ones use it every week, and that about a quarter say they simply copy what it gives them. Most teachers use it too, six in ten of them are worried about their students using it, and only about a third have had training on how to think about it critically.

Treat these as indicative: they are self-reported, they come from surveys with their own limits, and the largest one is about literacy rather than AI as such. But they tell you something that changes how to teach this subject. You are not introducing a new technology to a class that has never seen it. You are helping a class that already uses it to use it with their eyes open, and that is a much better place to teach from. The same goes for the next stage: nearly every undergraduate in a national survey said they use AI, and a growing minority put AI-written text straight into assessed work.

The evidence

Fact	Condition	Source
79.0% of the young people surveyed said they used generative AI in 2026, and 59.5% used it weekly or more often.	From the National Literacy Trust's Annual Literacy Survey in early 2026 (40,543 young people aged 13 to 18). The survey is about literacy, and the answers are self-reported.	National Literacy Trust, Young people, teachers' and parents' use of AI to support literacy in 2026 (Listed 18 August 2026 (survey carried out in early 2026))
24.1% of the young people said they simply copied AI outputs when using it for homework, rising to 31.6% of those receiving free school meals.		National Literacy Trust, Young people, teachers' and parents' use of AI to support literacy in 2026 (Listed 18 August 2026 (survey carried out in early 2026))
80.6% of the teachers surveyed used AI in 2026, up from 58.0% in 2025, and 60.1% were worried about students using it, up from 43.2%.	2,567 teachers took part. Teachers who had received training on critical AI literacy were 32.6%.	National Literacy Trust, Young people, teachers' and parents' use of AI to support literacy in 2026 (Listed 18 August 2026 (survey carried out in early 2026))
In HEPI's 2026 survey of 1,054 full-time UK undergraduates, 95% reported using AI in at least one way, 94% said they used generative AI to help with assessed work, and 12% said they directly included AI-generated text in assessed work (8% in 2025 and 3% in 2024).	Fieldwork by Savanta in December 2025. Self-reported.	Higher Education Policy Institute, sponsored by Kortext, Student Generative AI Survey 2026 (HEPI Report 199) (Published 12 March 2026)

Starter: your data trail

Before you say the word data, ask everyone to write down every single thing they or their devices created between waking up and arriving at school. The alarm, the message to a friend, the route the bus took, the tap on the card reader, the photo of the cereal, the game that told you your streak was in danger. It usually takes two minutes to fill a page, and the silence while it happens is the sound of a class discovering that data is not something that lives in a spreadsheet in an office, it is something they produce simply by being alive in the twenty-first century.

That is the whole opportunity of this subject. Data and AI can feel like a thing that happens to other people, in other buildings, with a lot of maths. It is much closer than that, and a student who has worked out that their own morning is a dataset is already asking the questions this subject needs: who holds it, who is it useful to, and what would a machine learn if it only ever saw my morning and nobody else's?

A ten minute starter: everyone lists the data they created before breakfast, and who now holds it. Nobody has to share their own list.

What this is

A ten minute warm-up with pen and paper. No computers are needed.

What to do

- Give everyone a sheet of paper. For two minutes, in silence, they list everything they or their devices made between waking up and arriving for the day: the alarm, a message, the bus tap, a photo.
- Next to each item they write who now holds that data. A guess is fine.
- Collect the most surprising items on the board, with no names. Nobody has to read out their own list.

What you should notice

Most students fill a page without trying. Data is not only forms and spreadsheets: it is made by ordinary actions, and other people hold it.

Time	10 to 15 minutes
You need	Paper and pens. No devices.
Where	Can be done with no devices.

Students will be able to

- Name at least five pieces of data they create in an ordinary morning.
- Say who holds one of them and why.
- Explain that data about a person is made by their ordinary actions, not only by filling in forms.

Running it

- Give everyone a sheet of paper and two minutes of silence. Ask them to list every piece of data they or their devices created between waking up and arriving for the day. Suggest prompts if they stall: the alarm, a message, a route, a bus tap, a photo, a game, a purchase.

- Now ask them to put a second column next to each item: who now holds it? A phone company, an app, a transport firm, the school or college, a shop. They can guess; the point is the pattern.
- Collect the most surprising items on the board, without names. Nobody has to read their own list out.
- Finish by asking which of the items could be useful for something good, and which they would want a say in.

Questions to ask

- Which of these did you choose to make, and which just happened?
- Who benefits from holding this data, and is it always the person it is about?
- If a model learned from all the lists in this room, whose mornings would it understand best?

Watch for

- Some students may not want to list things at home. Keep it to the morning and to devices, and let them choose what to share.
- Do not ask anyone to reveal anything that could identify another person.
- This is general discussion, not advice about any particular app's settings.

Game: teach the machine

Show a machine some examples, watch it draw its own rule, then feed it a lopsided set and watch it fail the group it has barely seen.

What this is

A factory has two kinds of machine, A (circles) and B (squares). Each is a dot on the chart, orange if it was about to fail and green if it was healthy. A computer studies the dots you give it, colours the chart green or orange as its own rule, then is scored on 40 machines it has never seen. Healthy kind B machines read higher than healthy kind A ones, so they sit in an area that is orange for kind A.

What to do

- Press Fair mix. Look at the three scores on the right: overall, kind A and kind B. They should all be high and close together.
- Press Mostly kind A. This set has plenty of kind A examples and almost none from kind B. Look at the scores again: overall still looks fine, but kind B has dropped a long way.
- Press Add 5 typical healthy kind B examples. Watch the kind B score recover. Press the faulty button too if you like.
- Tick Show the 40 test machines to see the hollow markers the rule is scored on, and which ones it gets wrong.

What you should notice

The overall score hid the problem. The rule failed the group it had barely seen, and the fix was to give it examples of that group. A model only knows what its examples taught it.

Time	25 to 35 minutes
You need	A projector for the whole class, or one device per pair.
Where	Interactive on the page. Needs a screen.

Students will be able to

- Explain that a model learns its rule from the examples it is given.
- Show how an unrepresentative training set makes a model work worse for the people it left out.
- Say why one overall accuracy figure is not enough.

Running it

- Open with the Fair mix preset. Point out the two kinds of machine, circles and squares, and that healthy and faulty readings are different colours. Show the coloured regions: that is the rule the machine drew for itself from the examples.
- Read out the scores for the hidden test machines: overall, for group A and for group B. They should be high and similar.
- Switch to Mostly group A, which has almost no examples of group B. Ask the class to predict what will happen to each score before you click. Then reveal it.

- Add a few group B points yourself by clicking on the chart where group B readings sit, and watch the group B score recover. Ask why.
- Try the Tiny set and change the number of neighbours to see how a small, noisy set makes a jumpy rule.

Questions to ask

- Who was left out of the training examples, and who paid for that?
- If you only looked at the overall score, would you have noticed the problem?
- Where could this happen for real, such as a tool trained mostly on one kind of person or one kind of device?
- What would you do to fix it, and who would need to be involved?

Watch for

- All the numbers are invented, and the machines are not real. The activity shows the mechanism, not the size of any real-world problem. The real-world evidence is the Gender Shades study in the box under the game.
- The method is called nearest neighbours: the machine looks at the closest examples it has seen and goes with the majority. Real systems use cleverer methods, but they also learn from the examples they are given.
- Students may say the machine is being unfair on purpose. It has no intentions; it reproduces the examples it was given.

The scores to expect

These are the scores on the 40 test machines, as whole percentages, for each starting set and each number of neighbours. Overall is all 40, and the others are the 20 machines of each kind. The page shows the same numbers.

Starting set	Neighbours	Overall	Kind A	Kind B
Fair mix: Ten healthy and ten faulty examples of each kind of machine.	1	95%	100%	90%
Fair mix: Ten healthy and ten faulty examples of each kind of machine.	3	93%	90%	95%
Fair mix: Ten healthy and ten faulty examples of each kind of machine.	5	95%	100%	90%
Fair mix: Ten healthy and ten faulty examples of each kind of machine.	7	95%	100%	90%
Mostly kind A: Plenty of kind A examples, and just two from kind B.	1	83%	100%	65%
Mostly kind A: Plenty of kind A examples, and just two from kind B.	3	85%	100%	70%
Mostly kind A: Plenty of kind A examples, and just two from kind B.	5	80%	100%	60%
Mostly kind A: Plenty of kind A examples, and just two from kind B.	7	80%	100%	60%
Tiny set: Only six examples in all, so the rule is a guess.	1	98%	95%	100%

Starting set	Neighbours	Overall	Kind A	Kind B
Tiny set: Only six examples in all, so the rule is a guess.	3	83%	100%	65%
Tiny set: Only six examples in all, so the rule is a guess.	5	93%	100%	85%
Tiny set: Only six examples in all, so the rule is a guess.	7	98%	95%	100%

Every machine and reading is made up. The scores for Mostly kind A show the point: a respectable overall figure and a much lower one for the kind the rule has barely seen. The real-world version is the Gender Shades study below.

The evidence

Fact	Condition	Source
A 2018 study of commercial gender classification systems found that darker-skinned females were the most misclassified group, with error rates of up to 34.7%, while the maximum error rate for lighter-skinned males was 0.8%.	It tested specific commercial products as they were in 2018 and says nothing about how any system performs now. Its point for a classroom is that one overall accuracy figure can hide a very different result for different groups.	Joy Buolamwini and Timnit Gebru, Proceedings of the 1st Conference on Fairness, Accountability and Transparency (PMLR volume 81, pages 77 to 91), Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification (2018)

Your notes

Lab: predict the next word

A tiny language model you can read in full. See how a chatbot can sound sure and be wrong, with every probability on show.

What this is

A very small chatbot you can read in full. It has seen 38 short made-up sentences and nothing else. For any start, it counts which word came next in those sentences and offers the likeliest, with how likely each one is.

What to do

- Leave the start as The cat sat on the. Look at the list on the right: mat comes first at 63%, because it came next most often in the sentences it read.
- Click a word to add it to the sentence, and click more to build one. Or press Pick one for me and let it choose in proportion to the percentages.
- Move the Creativity slider to the left and pick a few times, then to the right and pick a few times. Left plays safe, right lets unlikely words in.
- Now press The capital of Australia is. It offers Sydney at 60% and Canberra at 40%.

What you should notice

It does not know anything. It picks what was most common in the text it read, and that text had a mistake in it, so the likeliest answer is the wrong one. Likely is not the same as true, which is why real chatbots can sound sure and be wrong.

Time	20 to 30 minutes
You need	A projector or one device per pair.
Where	Interactive on the page. Needs a screen.

Students will be able to

- Describe a language model as something that picks a likely next word.
- Explain why a likely answer is not the same as a true one.
- Say what the creativity setting changes.

Running it

- Start with the prompt 'the cat sat on the'. Show the options and the probabilities. Ask what the model is doing: counting which word followed those words in its training sentences.
- Click the likeliest word a few times to build a sentence. Then press 'Pick one for me' a few times and compare the results. Move the creativity slider up and down and note what changes.
- Open the training sentences. Read a few aloud. Ask the class where the options came from. (From these sentences, and nowhere else.)
- Now use the prompt 'the capital of australia is'. Point out that the most likely word is not the true one, because the made-up training sentences contain a common mistake. Ask what a real system would do if the text it learned from contained a mistake many times.

- Show the NCSC fact: the NCSC describes a language model as turning the relationships between words into a probability model.

Questions to ask

- Why can a model sound confident when it is wrong?
- What would make it more likely to be right? What would not?
- Why might the NCSC advise against typing sensitive information into a public chatbot?
- What is the difference between 'likely to be said' and 'true'?

Watch for

- This is a toy. It counts word pairs in a few dozen invented sentences. Real models are trained on vastly more text with far more powerful methods. What they share with this toy is the core move: choosing a likely continuation.
- The made-up mistake about Australia is there on purpose and is clearly shown. Do not let the class leave thinking that is a real fact about any real tool.
- Randomness only happens when you press the button, so the page is the same for everyone when it loads.

What the model offers at each starting point

These are the shares at the default creativity setting of 1.0. Lower settings make the first choice even more likely, and higher settings flatten them.

Starting words	Looks at	Next-word options
The cat sat on the	"on the"	mat 63%, sofa 13%, step 13%, wall 13%
The teacher	"the teacher"	marked 40%, opened 40%, read 20%
It was sunny so the	"so the"	students 83%, teacher 17%
The capital of Australia is	"australia is"	sydney 60%, canberra 40%

The sentences about the capital of Australia are a planted mistake: the made-up text says Sydney more often than Canberra, which is wrong. A model that picks the likeliest word will repeat it.

Everything the model has read

- The cat sat on the mat. (seen 3 times)
- The cat sat on the sofa.
- The cat sat on the wall.
- The dog sat on the mat. (seen 2 times)
- The dog sat on the step.
- The dog chased the cat. (seen 2 times)
- The dog chased the ball.
- The cat chased the mouse. (seen 2 times)
- The teacher marked the books. (seen 2 times)
- The teacher read the register.
- The teacher opened the window.
- The students read the books. (seen 2 times)

- The students opened the window.
- The students wrote the answers. (seen 2 times)
- It was sunny so the students went outside. (seen 3 times)
- It was sunny so the teacher opened the window.
- It was raining so the students stayed inside. (seen 2 times)
- The capital of france is paris. (seen 3 times)
- The capital of japan is tokyo. (seen 2 times)
- The capital of australia is sydney. (seen 3 times)
- The capital of australia is canberra. (seen 2 times)

The evidence

Fact	Condition	Source
The NCSC describes a large language model as an algorithm trained on a large amount of text-based data that analyses the relationships between words and turns them into a probability model, which it then uses to answer.	Published 14 March 2023. This is the NCSC's plain-English description, and it is the idea the Predict the next word lab shows in miniature.	National Cyber Security Centre (NCSC), ChatGPT and large language models: what's the risk? (Published 14 March 2023)
The NCSC says large language models can get things wrong and hallucinate incorrect facts, and that they provide convincing-sounding answers that may only be partially correct, particularly as the topic gets more niche.		National Cyber Security Centre (NCSC), ChatGPT and large language models: what's the risk? (Published 14 March 2023)

Your notes

Lab: make the same data say three things

One year of made-up library loans, three honest-looking charts, three different headlines. Then spot the tricks.

What this is

Twelve months of made-up library loans drawn as a chart. You can change which months it shows, where the vertical axis starts, and whether it is bars or a line. The dark box above the chart shows the headline the chart seems to give.

What to do

- Look at the chart as it opens: the whole year, axis from zero. The headline says loans are up 4.9% and the bars look almost the same height.
- Press At 470 under Where the vertical axis starts. The numbers have not changed, but the bars now look hugely different. The box tells you how many times higher the last bar looks and how many times higher it really is.
- Press March to August. The headline flips to down 7.5%. Then press August to December: up 9.2%.
- Press Line to see the same data as a line.

What you should notice

One set of twelve numbers gave three true headlines. A chart is built from choices, such as where the axis starts and which months to show. Always ask what those choices were and why.

Time	20 to 25 minutes
You need	A projector or one device per pair.
Where	Interactive on the page. Needs a screen.

Students will be able to

- Spot a truncated axis and a cherry-picked time window.
- Explain how the same figures can support opposite headlines.
- Ask what a chart leaves out before believing it.

Running it

- Show the chart with the defaults: a bar chart of the whole year starting from zero. Ask for a headline. Probably something like 'loans are steady, up a little'.
- Now start the axis at 470 instead of zero. The same bars look dramatically different. Ask what changed. (Only the axis.)
- Switch the window to March to August, then August to December. Ask students to write the headline for each. Point out they are all true and they say different things.
- Switch to a line chart and ask whether it feels more or less honest than bars. (Bars start from zero for a reason: the height is the amount.)
- Finish by asking what you would need to ask the person who made the chart.

Questions to ask

- Which of these headlines would you believe from a stranger? Which would you ask more about?
- Is it ever reasonable to start an axis above zero? (For a line chart where the small changes matter, yes, if it is clearly labelled.)
- Who benefits when a chart looks dramatic?

Watch for

- The numbers are invented for the exercise. They are not real library figures.
- Avoid presenting the zoomed chart as 'lying'. It is the choice that misleads when it is not labelled, so teach students to read the axis.

The numbers

Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
510	520	530	525	515	500	495	490	505	515	525	535

The three headlines

Months	First to last	Change
The whole year (Jan to Dec)	510 to 535	up 4.9%
March to August	530 to 490	down 7.5%
August to December	490 to 535	up 9.2%

The numbers are invented. All three headlines are true, and they point in different directions. Starting the axis at 470 instead of zero makes the same bars look several times as different.

Your notes

Game: clean the mess

A spreadsheet with seven things wrong with it. Find them, then see what they did to the class average.

What this is

A made-up class survey in a spreadsheet, with seven mistakes of the kind real data often has. Your job is to find them before the answers are shown.

What to do

- Read the table. Click any cell that looks wrong and it turns yellow. Click a row number to mark a whole row, for a row you think was entered twice. Click again to undo.
- Press Check my flags. Green means a real mistake you found, orange means a real mistake you missed, red means a cell that was fine.
- Read the two class average heights: the one the spreadsheet gives from the messy data, and the one after cleaning.
- Read the seven cards to see what each mistake was and what to do about it. Press Start again to have another go.

What you should notice

A few bad cells change an average, and two mistakes can pull in opposite directions so they hide each other. Cleaning data takes judgement, and sometimes the honest answer is to leave a value out and ask.

Time	20 to 30 minutes
You need	A projector or one device per pair.
Where	Interactive on the page. Needs a screen.

Students will be able to

- Recognise duplicates, wrong units, impossible values, inconsistent labels and missing values.
- Explain how a few bad cells change an average.
- Say why cleaning data is real work and not a chore to skip.

Running it

- Tell the class this is a survey of a made-up class. Ask them to flag anything that looks wrong by clicking on it. There are seven problems, though not every problem is a single cell.
- When they are ready, check the flags. Go through each problem and what it was. Count how many they found.
- Read out the two averages: the one the spreadsheet calculates from the messy data and the one from the cleaned data. Ask why they differ.
- Point out that two of the slips pull in opposite directions: the 1.62 drags the average down and the 262 and the copied row push it up, so they partly hide each other. Ask which problem was hardest to spot. (The decimal point in a height is easy to miss.)

Questions to ask

- What would a model trained on the messy data have learned?
- Which of these problems could a computer find on its own, and which needs a human who knows what the numbers mean?
- What do you do with a value like 262 when you cannot ask the person?

Watch for

- The data is entirely invented. The heights are plausible for a class of young people but belong to nobody.
- In the cleaned figure we remove the impossible 262 rather than guess what it was meant to be. Real analysts often have to make that call, and should say they did.

Answer key

Problem	What is wrong	What to do
A row entered twice	Row 12 is an exact copy of row 5, so one student counts double.	Remove the copy. Check first that it really is the same person and not two students who happen to match.
Metres where centimetres should be	Row 6 says 1.62. Everyone else is in centimetres, so this is probably 162 typed in metres.	Convert it to 162 and note that you did. It is a good guess, but still a guess, so say so.
A height that cannot be right	Row 10 says 262 cm, taller than any person has ever been. It may be 162 or 126 or something else.	Do not guess. Leave it out of the average and ask the person who collected it.
The same thing written two ways	Row 7 says C where the others say Group C. A computer sees two different groups.	Agree one format and change the odd one out.
Spelt differently, meaning the same	Walk, walking and WALK are one way of travelling, but a computer counts them as three.	Pick one label and apply it everywhere.
A missing value	Row 9 has no travel time at all.	Do not fill it with zero. Mark it as missing and leave it out of any average of travel times.
A trip that takes minus fifteen minutes	Row 13 has a negative time, which is impossible. It is probably a typing slip, but nobody can tell which number was meant.	Treat it as missing and check with whoever collected it.

What it does to the average

The class average height as a spreadsheet calculates it from all 16 rows as typed is 155.7 cm. After cleaning (the copy of row 5 removed, 1.62 read as 162 and the impossible 262 left out) it is 159.6 cm from 14 rows. Two of the slips pull in opposite directions, so they partly hide each other.

The spreadsheet

Row	Group	Height (cm)	How they travel	Minutes to school
1	Group C	158	Walk	12
2	Group C	163	Bus	25
3	Group C	149	Walk	9

Row	Group	Height (cm)	How they travel	Minutes to school
4	Group C	171	Car	15
5	Group C	155	Bus	30
6	Group C	1.62	Walk	14
7	C	160	Bus	22
8	Group C	166	walking	10
9	Group C	152	Car	(blank)
10	Group C	262	Bus	28
11	Group C	157	WALK	11
12	Group C	155	Bus	30
13	Group C	168	Car	-15
14	Group C	150	Bus	35
15	Group C	164	Walk	8
16	Group C	159	Car	18

Your notes

Lab: correlation is not causation

Umbrellas and lateness rise together. Does one cause the other? Colour the days by weather and watch the pattern vanish.

What this is

Forty made-up days. Each dot is one day: across is how many umbrellas the shop sold, up is how many students arrived late. The box on the right gives a correlation number, from -1 to 1, which says how closely two things rise and fall together.

What to do

- Look at the chart before you click anything. The dots slope upwards and the correlation is about 0.78, which is strong. Ask the class: do umbrellas make students late?
- Tick Colour the days by weather. Rainy days become blue circles and dry days become orange squares.
- Read the two new numbers, one for rainy days only and one for dry days only. They are close to zero.
- Try the three Find the third factor boxes at the bottom: decide what the hidden cause is, then press Reveal.

What you should notice

The pattern is real, but umbrellas do not cause lateness. Rain causes both. Inside each kind of day the link almost disappears. A correlation tells you where to look, not what caused what.

Time	20 to 25 minutes
You need	A projector or one device per pair.
Where	Interactive on the page. Needs a screen.

Students will be able to

- Read a scatter graph and describe the correlation.
- Explain what a third factor is and how it can create a pattern.
- Suggest a third factor for a new example.

Running it

- Show the scatter graph with no colours. Ask what the pattern says. Read out the correlation figure. Ask for a story that explains it, and write the students' stories on the board.
- Now tick 'Colour the days by weather'. Inside each colour the pattern disappears. Ask what that tells them.
- Introduce the term third factor, or confounder. Then work through the three made-up examples at the bottom, asking for a third factor before you reveal each one.
- Link it to the GCSE maths statement: recognise correlation and know that it does not indicate causation.

Questions to ask

- Does selling an umbrella make anyone late? What does?
- If you only had the two columns and not the weather, how could you tell?

- When a news story says 'linked to', what should you ask?

Watch for

- The days and figures are invented. The pattern was built in on purpose, to show the mechanism.
- Some students will say 'so correlation means nothing'. It does not: correlation is a clue to look further, not a conclusion.

The correlations the page shows

Days	Number of days	Correlation
All days	40	0.78
Rainy days only	21	-0.04
Dry days only	19	0.20

The days and numbers are invented, built so that the weather drives both. Inside each kind of day there is little pattern left.

Find the third factor: answers

Made-up chart	Third factor
Imagine a made-up chart that shows drivers with more speeding fines also have more parking fines.	How much they drive. People who drive more pick up more of both.
Imagine a made-up chart that shows the more fire engines at a fire, the more damage the building suffers.	The size of the fire. A bigger fire needs more engines and also causes more damage.
Imagine a made-up chart that shows days when more ice creams are sold are also days with more sunburn.	Sunshine. Hot, sunny days make people buy ice cream and also get burnt.

Your notes

Diagram: from raw data to a decision

Seven stages from a question to a model in use, what goes wrong at each one, and the real jobs that do the work.

What this is

Every model goes through the same seven stages, from asking a question to being used on real people. This diagram lets you click through them one at a time.

What to do

- Click stage 1, Ask a question. Read what happens and what can go wrong. Ask the class what else could go wrong before you read the list.
- Look at Who does this work on the right: real jobs from the Career Matrix, with the area of work.
- Click stage 5, Test it, and stage 7, Watch it and be ready to change it. These show real evidence: a study of face analysis systems and the 2020 exam grades.
- Follow the See it in action link to jump to the activity that shows that stage.

What you should notice

Many problems start early, in what data is collected and who is left out, and only show up at the end. Lots of different jobs are involved, and not all of them are coding.

Time	20 to 30 minutes
You need	A projector, or the printed table in the teacher pack.
Where	Interactive on the page. Can be done with no devices.

Students will be able to

- Order the stages of building and using a model.
- Give one thing that can go wrong at each stage.
- Name two jobs in Data and AI and say which stage they work on.

Running it

- Click through the seven stages together, one at a time. For each, ask the class what could go wrong before you show the box.
- Pick two stages and look at the real jobs named under them. Notice they are different kinds of job: some are engineering, some are about people and decisions.
- Where a stage has a real example attached, the 2020 grades or the Gender Shades study, read the evidence box and ask whether the problem was a technical one, a people one, or both.
- Set a short task: pick an app they use and name one thing that could go wrong at each stage.

Questions to ask

- Which stage do you think is most often skipped, and why?

- Which of these jobs would suit you, and which would you hate?
- Where does a human need to be involved?

Watch for

- This is a simplification. Real projects loop back between stages and run several at once.
- The jobs shown are the ones the Career Matrix lists for the matching area of work, with their pay and sources in the career guide.

A deliberate simplification. Real projects loop back between stages and run several at once. The jobs are those the Career Matrix lists for the matching areas of work.

Stage	What happens	What can go wrong	Who does it
1. Ask a question	Decide what you are trying to find out or predict, who it is for, and what a good answer would look like.	Measuring what is easy to count instead of what matters. Asking a question the data cannot answer. Nobody asks who will be affected.	Data & AI Product Management: AI Product Manager, Data Product Manager, ML Product Owner, Platform Product Manager; Analytics & Decision Intelligence: Data Analyst, Business Intelligence Developer, Decision Scientist, Analytics Lead
2. Collect the data	Gather the examples a model will learn from, or the figures an analysis will use.	The data over-represents some people and misses others. It was collected for a different purpose and does not fit this one. Personal data is gathered with no good reason, or without telling people.	Data Foundations & Engineering: Data Engineer, Analytics Engineer, Database Specialist, Data Platform Engineer
3. Clean and label	Fix errors, merge duplicates, agree units and labels, and where the model needs it, attach the right answer to each example.	Duplicates, wrong units, impossible values, inconsistent labels and gaps change the answers. The labels reflect one person's judgement and nobody checks it.	Data Foundations & Engineering: Data Engineer, Analytics Engineer, Database Specialist, Data Platform Engineer
4. Build and train	Choose a method and let it find patterns in the data.	The pattern it finds is a coincidence or the work of a hidden third factor. It learns the examples so closely that it fails on new cases.	Machine Learning & AI Development: Machine Learning Engineer, AI Engineer, Applied Scientist, NLP / Computer Vision Specialist; Applied Data Science & Experimentation: Data Scientist, Experimentation Lead, Forecasting Analyst, Research Scientist
5. Test it	Check it on examples it has never seen, and check who it works for, not only how well it does overall.	One overall score hides a group the system fails. The test examples are too like the training examples to tell you anything new.	Applied Data Science & Experimentation: Data Scientist, Experimentation Lead, Forecasting Analyst, Research Scientist; Machine Learning & AI Development: Machine Learning Engineer, AI Engineer, Applied Scientist, NLP / Computer Vision Specialist

Stage	What happens	What can go wrong	Who does it
6. Put it to work	Connect it to a real process where its outputs affect real people.	It is used for something it was never tested for. Nobody can see why it decided what it did, or challenge the result.	AI Operations & Deployment: MLOps Engineer, AI Platform Engineer, Model Operations Specialist, AI Reliability Engineer
7. Watch it and be ready to change it	Keep checking how it performs in practice, listen to the people it affects, and step in when it goes wrong.	Nobody is responsible for noticing, so nothing is noticed. There is no way to put things right when a decision is wrong.	Data Strategy, Governance & Ethics: Data Governance Lead, Privacy Engineer, Responsible AI Specialist, Chief Data Officer; AI Operations & Deployment: MLOps Engineer, AI Platform Engineer, Model Operations Specialist, AI Reliability Engineer

The evidence

Fact	Condition	Source
A 2018 study of commercial gender classification systems found that darker-skinned females were the most misclassified group, with error rates of up to 34.7%, while the maximum error rate for lighter-skinned males was 0.8%.	It tested specific commercial products as they were in 2018 and says nothing about how any system performs now. Its point for a classroom is that one overall accuracy figure can hide a very different result for different groups.	Joy Buolamwini and Timnit Gebru, Proceedings of the 1st Conference on Fairness, Accountability and Transparency (PMLR volume 81, pages 77 to 91), Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification (2018)
In summer 2020, for A level entries the calculated grade was the same as the centre assessment grade (CAG) submitted by the school or college for 58.7% of entries and higher for 2.2%, which leaves 39.1% where it was lower.	Ofqual gives the 58.7% and the 2.2%. The 39.1% is what is left over, our subtraction. These are the calculated grades, not the grades students finally received.	Ofqual, Summer 2020 results analysis: GCSE, AS and A level (update to the interim report) (December 2020)
The final grade awarded was the higher of a student's CAG or their calculated grade, so for 97.8% of A level entries the final grade was the same as the CAG.		Ofqual, Summer 2020 results analysis: GCSE, AS and A level (update to the interim report) (December 2020)
Ofqual's annual report says the approach developed for GCSE, AS and A levels did not command public confidence, and that exam boards would award centre assessment grades unless the standardised grade was higher. Concern grew after A level results day on 13 August 2020, and GCSE results were issued on the new basis on 20 August 2020.	The report does not give the exact date of the Ofqual Board's decision.	Ofqual, Ofqual annual report and accounts 2020 to 2021 (Chair's foreword dated 8 July 2021)

Fact	Condition	Source
The ICO's guidance on automated decision-making covers decisions based solely on automated processing that have a legal or similarly significant effect on a person, and sets out safeguards such as ways to make representations, human intervention and a way to contest the decision.	The ICO page says this is draft guidance under consultation, updated 31 March 2026, reflecting the Data (Use and Access) Act 2025. It is general information, not legal advice.	Information Commissioner's Office (ICO), Automated decision-making and profiling (draft guidance) (Updated 31 March 2026; the page says it is draft guidance under consultation)

Your notes

When a model decides something about people

The abstract version of 'bias in AI' slides off most students, so give them the concrete one. In 2018 researchers published a test of commercial systems that tried to tell whether a face belonged to a man or a woman. The researchers split the results by group and found that the system's worst group and its best group were a very long way apart, as the figures below show. A single overall score, it turns out, is a lovely way to hide the people a system does not work for. The Teach the machine game above lets your students make that happen with their own hands, in a toy with made-up machines, before they read the real thing.

The second story is closer to home, because it happened to people the age of your sixth form and their older siblings. In summer 2020 exams were cancelled, and a standardisation process produced calculated grades that Ofqual compared with the centre assessment grades submitted by schools and colleges. The published figures show how often the calculated grade matched, beat or fell below the submitted grade, and Ofqual's own annual report says the approach did not command public confidence and was changed. We have set out exactly what the sources say and nothing more, because the useful lesson is not who was right, it is what a fair system needs: a way to check an individual result, a way to challenge it, and someone accountable.

The evidence

Fact	Condition	Source
A 2018 study of commercial gender classification systems found that darker-skinned females were the most misclassified group, with error rates of up to 34.7%, while the maximum error rate for lighter-skinned males was 0.8%.	It tested specific commercial products as they were in 2018 and says nothing about how any system performs now. Its point for a classroom is that one overall accuracy figure can hide a very different result for different groups.	Joy Buolamwini and Timnit Gebru, Proceedings of the 1st Conference on Fairness, Accountability and Transparency (PMLR volume 81, pages 77 to 91), Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification (2018)
In summer 2020, for A level entries the calculated grade was the same as the centre assessment grade (CAG) submitted by the school or college for 58.7% of entries and higher for 2.2%, which leaves 39.1% where it was lower.	Ofqual gives the 58.7% and the 2.2%. The 39.1% is what is left over, our subtraction. These are the calculated grades, not the grades students finally received.	Ofqual, Summer 2020 results analysis: GCSE, AS and A level (update to the interim report) (December 2020)
The final grade awarded was the higher of a student's CAG or their calculated grade, so for 97.8% of A level entries the final grade was the same as the CAG.		Ofqual, Summer 2020 results analysis: GCSE, AS and A level (update to the interim report) (December 2020)
Ofqual's annual report says the approach developed for GCSE, AS and A levels did not command public confidence, and that exam boards would award centre assessment grades unless the standardised grade was higher. Concern grew after A level results day on 13 August 2020, and GCSE results were issued on the new basis on 20 August 2020.	The report does not give the exact date of the Ofqual Board's decision.	Ofqual, Ofqual annual report and accounts 2020 to 2021 (Chair's foreword dated 8 July 2021)

Fact	Condition	Source
The ICO's guidance on automated decision-making covers decisions based solely on automated processing that have a legal or similarly significant effect on a person, and sets out safeguards such as ways to make representations, human intervention and a way to contest the decision.	The ICO page says this is draft guidance under consultation, updated 31 March 2026, reflecting the Data (Use and Access) Act 2025. It is general information, not legal advice.	Information Commissioner's Office (ICO), Automated decision-making and profiling (draft guidance) (Updated 31 March 2026; the page says it is draft guidance under consultation)

A model that really did decide something about students their age. What it did, what Ofqual says, and what a fair system would need.

What this is

A discussion with no devices. In summer 2020 exams were cancelled and a computer process helped decide grades. The published facts are in the evidence boxes just above, and everything said in the lesson should come from them.

What to do

- Read the evidence boxes aloud. Check everyone knows what a centre assessment grade was: the grade each school or college submitted.
- In pairs, ask: if the calculated grade matched the submitted grade for most students, was it a good system? Take answers.
- Tell them that where the two differed, the calculated grade was far more often lower than higher.
- Ask what a fair system would need. Look for a way to check an individual result, a way to challenge it, and someone accountable.

What you should notice

A system can be right for most people and still wrong for you. What matters is whether people can check it and challenge it. Some students or staff may have been affected, so stick to what the sources say.

Time	30 to 40 minutes
You need	The facts below. No devices.
Where	Can be done with no devices.

Students will be able to

- State what happened in 2020 using the published figures.
- Explain why a model that is right on average can still be unfair to individuals.
- Say what safeguards the law and good practice expect when a decision is automated.

Running it

- Share the facts in the box. Check the class knows what a centre assessment grade was: the grade each school or college submitted for its students, in place of exams. Keep it factual.

- Ask in pairs: if a calculated grade matches the teacher's grade for most students, is it a good system? Take answers.
- Reveal the other side: for the A level entries where the two differed, it was far more often lower than higher.
- Ask what a fair system would have needed: a way to check individual results, a way to appeal, someone accountable, a way to change course. Compare with the ICO safeguards in the box.
- Finish with what actually happened: Ofqual's own report says the approach did not command public confidence, and the grades were changed to the higher of the two.

Questions to ask

- A model can be right for most people and still wrong for you. Is that acceptable? Who decides?
- Who should be able to challenge a decision made by a model?
- What would you have wanted to know before it was used?

Watch for

- Some students, siblings or staff in your school or college were affected, and some may have strong feelings. Keep to what the published sources say and do not guess about causes.
- The activity is about the idea of a model deciding things about people, not about blaming anyone. Ofqual's report is the account to rely on.
- This is general information, not legal advice.

Your notes

Nine kinds of data job

Explore the nine areas of Data and AI work, with the roles, what they suit and pay with its source. Not everyone needs to love maths.

What this is

Nine kinds of work in data and AI, taken from the TechSkills Career Matrix, with the jobs in each, what each suits and the pay with its source.

What to do

- Click each of the nine boxes in turn.
- For each one, read the jobs and the line that starts Best suited to.
- Where there is a gold box labelled a route into this area, read it: that is the apprenticeship or hands-on way in.
- Ask students which area they could picture themselves in, and why.

What you should notice

Data and AI is not one job. Some areas are about building, some about deciding and some about being on site, and several do not need strong coding. Pay figures come with their source and date, so quote them with that.

Time	15 to 20 minutes
You need	A projector or one device per pair.
Where	Interactive on the page. Needs a screen.

Students will be able to

- Name at least four different Data and AI jobs.
- Match a student's interests to an area of the field.
- Find an apprenticeship route into one of them.

Running it

- Pick one student interest at random, such as sport, music, fashion, games or health. Ask which of the nine areas could involve it.
- Click through areas that match. Read the 'best suited to' line aloud, and the pay with its source.
- Point out the apprenticeship routes where there is one, and say plainly that they are in addition to degrees, not instead of them.

Questions to ask

- Which area would you not have thought of as a tech job?
- Which would suit someone who does not enjoy maths? Does any?

Watch for

- Pay is quoted as the Career Matrix gives it, with its own source and date. Quote it with that, not as a promise.
- Not every area has a hands-on or apprenticeship route.

Area	Jobs	Pay, with its source	Best suited to
01. Data Foundations & Engineering	Data Engineer, Analytics Engineer, Database Specialist, Data Platform Engineer	£33,515 to £90,000+ (database administrator, ONS ASHE, to lead data engineer, IT Jobs Watch UK)	Engineers who enjoy solving complex systems and infrastructure challenges.
02. Analytics & Decision Intelligence	Data Analyst, Business Intelligence Developer, Decision Scientist, Analytics Lead	£30,000 to £82,250+ (data analyst 10th percentile to senior data analyst 90th percentile, IT Jobs Watch UK)	Logicians who enjoy structure, organisation, and clear visual communication.
03. Machine Learning & AI Development	Machine Learning Engineer, AI Engineer, Applied Scientist, NLP / Computer Vision Specialist	£44,875 to £117,000+ (NLP specialist 10th percentile to AI engineer 90th percentile, IT Jobs Watch UK)	Mathematical thinkers who love experimentation, research, and coding.
04. Data Strategy, Governance & Ethics	Data Governance Lead, Privacy Engineer, Responsible AI Specialist, Chief Data Officer	£40,000 to £150,000+ (data protection officer, IT Jobs Watch UK, to chief data officer, Payscale UK)	Detail-oriented thinkers interested in law, ethics, and social impact.
05. AI Operations & Deployment	MLOps Engineer, AI Platform Engineer, Model Operations Specialist, AI Reliability Engineer	£57,500 to £123,750+ (MLOps engineer, junior to senior, IT Jobs Watch UK, closest real data for AI platform/model ops/reliability roles)	System builders who enjoy the bridge between software engineering and AI.
06. Frontier AI & Advanced Applications	Generative AI Specialist, Robotics AI Engineer, Autonomous Systems Developer, AI Research Engineer	£55,000 to £105,000+ (robotics engineer to generative AI software engineer, IT Jobs Watch UK)	Early adopters who thrive in rapidly changing, experimental environments.
07. Applied Data Science & Experimentation	Data Scientist, Experimentation Lead, Forecasting Analyst, Research Scientist	£36,000 to £95,000+ (junior data scientist to senior data scientist, IT Jobs Watch UK)	Curious investigators who want proof, not just intuition.
08. Data & AI Product Management	AI Product Manager, Data Product Manager, ML Product Owner, Platform Product Manager	£47,624 to £102,735+ (AI product manager, 25th to 90th percentile, Glassdoor UK)	Bridge-builders who enjoy scoping ambiguous problems into shippable plans.
09. Asset Integrity & Inspection Engineering	NDT Engineer, Materials Engineer, Asset Integrity Engineer, Quality Assurance Engineer	£24,000 to £70,000+ (technician to senior NDT/asset integrity engineer, Glassdoor/Indeed UK benchmarks)	Methodical, detail-focused people who'd rather prove something is safe with real evidence than take someone's word for it.

The evidence

Fact	Condition	Source
The Data analyst apprenticeship (Level 4, standard ST0118) typically takes 24 months, not including the end-point assessment, and describes the job as collecting, organising and studying data to provide business insight.	Version 1.1 is approved for delivery but is being revised.	Skills England (apprenticeship standards), Data analyst (standard ST0118, version 1.1) (Version 1.1, which the page says is approved for delivery but being revised)

Data and AI is not only for people who like screens

Here is something that surprises most students and quite a few teachers: data has to come from somewhere, and a surprising amount of it comes from someone in a hi-vis jacket. Our Career Matrix has a name for the genuinely physical, on-site routes, Hard Hat Tech, and five of our eight career families have one. In Data and AI there are two. One is the surveyor, who captures the real world as measurements with laser scanners, GPS equipment and mapping drones. The other is the inspection engineer, who uses ultrasound and other imaging to check the materials, welds and components of safety-critical equipment for hidden flaws, without taking anything apart.

Then there is the place all the data finally lives. A data centre is a building full of computers, and in 2024 the government announced it would designate data centres as Critical National Infrastructure, with the aim of making the data in them less likely to be compromised in outages and cyber attacks. If your students think of AI as something that happens in the cloud, it is worth telling them the cloud has a postcode.

Drones are the perfect doorway, because students already find them thrilling and they really are used for surveying. If your school or college wants a drone club, the rules changed on 1 January 2026, and the box below says exactly how.

Data has to be captured, stored and checked somewhere physical. Here are the on-site careers that do it.

What this is

Two hands-on jobs in the Data and AI family, from the Career Matrix, and the reason data centres are treated as national infrastructure.

What to do

- Read the two route cards and ask what data each job captures or checks.
- Read the data centre fact in the box above them.
- If you have a drone club, or want one, read the CAA rules in the evidence box together.

What you should notice

Data has to be captured, stored and checked somewhere physical, and that work is done by people on site as well as people at screens.

Time	10 to 15 minutes
You need	A projector or the page on a phone.
Where	Needs a screen.

Students will be able to

- Name two hands-on routes in the Data and AI family.
- Explain why a data centre is infrastructure.
- State the basic rule for flying a drone in the UK.

Running it

- Read the two hands-on routes. Ask what data each of them captures or checks.
- Show the data centre fact. Ask what a data centre is and why anyone would protect it like a power station.
- If you have a drone or a drone club, read the CAA rules box together first.

Questions to ask

- Who has to go out in all weathers so that the data is right?
- If the data from an inspection is wrong, who is affected?

Watch for

- Pay and roles are as the Career Matrix gives them. The drone rules are the CAA's from 1 January 2026.
- Never fly a drone near people or without checking the CAA's current rules first.

Route	Jobs	Pay, with its source
Construction Technology & Geospatial Surveying: Capturing the data that describes the physical world, on site rather than in an office.	Surveying Technician, Engineering Design Technician (BIM/CAD), Geospatial Survey Technician	£32,905 to £47,519 (median, surveying technician to chartered surveyor)
Non-Destructive Testing (NDT): Hands-on inspection of real components and structures, through a paid apprenticeship rather than an engineering degree.	NDT Operator Apprentice, NDT Engineering Technician, Inspection Engineer	£22,000 to £45,000+ (Level 2 operator apprentice to Level 3 engineering technician, Glassdoor/Indeed UK benchmarks)

The evidence

Fact	Condition	Source
In September 2024 the government announced that data centres would be designated as Critical National Infrastructure.	The announcement covers physical data centres and the cloud operators that use them.	Department for Science, Innovation and Technology, Data centres to be given massive boost and protections from cyber criminals and IT blackouts (Published 12 September 2024)
From 1 January 2026 anyone flying a drone that weighs over 100g needs a Flyer ID, which comes from a free theory test on the CAA website (the old threshold was 250g).		Civil Aviation Authority (CAA), Regulator calls on new and existing drone users to learn new rules before taking off (Published 29 December 2025)
Under-18s need a Flyer ID and can take the test themselves, but a parent or guardian must hold the Operator ID. Children under 13 need a parent or guardian with them for the online test, and children under 12 must be supervised by someone aged 16 or over when flying.		Civil Aviation Authority (CAA), Regulator calls on new and existing drone users to learn new rules before taking off (Published 29 December 2025)

Using AI tools in your classroom: what the rules say

Sooner or later a colleague will ask whether they can use an AI tool to write reports or mark homework, and a student will ask whether they can use one for coursework. It is far better to know the short version before either happens. The Department for Education is clear that it sees more immediate benefit and fewer risks from teachers using these tools than from students using them, and equally clear that the teacher stays responsible for the result.

The practical rules are short. Do not put personal data into a public tool, and if you think you need to, ask your data protection officer or IT lead first. Check what comes out, because these tools can state falsehoods with perfect confidence. Check age limits before students use anything, and make sure there is supervision and filtering. Do not let students' original work be used to train a model without permission. And for anything that counts towards a qualification, read the current JCQ guidance and your exam board's before you set it. The checker below turns this into six questions you can answer in a minute.

This is general information to help you teach, not legal advice, and the DfE pages are the place to read the detail.

A short checker for teachers: should this go into an AI tool? It shows what the DfE and the NCSC actually say.

What this is

A checklist for teachers. It asks six questions about something you are thinking of putting into an AI tool, and shows what government guidance says about each. Nothing you tick is stored.

What to do

- Think of a real task, such as writing report comments or marking a piece of homework.
- Answer Yes or No to each of the six questions.
- Read the box that appears at the bottom. It lists what you should do and the guidance behind it, with the source.

What you should notice

The short version is: keep personal data out, check what comes out, mind age limits, and follow the rules for qualifications. This is a summary of the sources, not legal advice, so ask your data protection officer when in doubt.

Time	5 to 10 minutes
You need	A device. Also useful on a staff meeting screen.
Where	Interactive on the page. Needs a screen.

Students will be able to

- Say what the DfE recommends about personal data in AI tools.
- Name the checks before students' work goes into one.
- Know where assessment rules sit.

Running it

- Open the checker on a screen at a staff meeting or in department time. Answer the four questions for a task you actually have in mind.
- Read the guidance it shows, with its source. Remind colleagues that the checker is a summary and the DfE pages are the authority.
- Ask your data protection officer or IT lead where your school's or college's policy stands, and note the answer for others.

Questions to ask

- What is the quickest safe habit that would stop personal data going into a public chatbot?
- How would you tell students what is and is not allowed for coursework?

Watch for

- The checker does not store anything you tick. It shows what the sources say and tells you where to check your own school's or college's policy.
- Assessment rules change. Read the current JCQ guidance and your exam board's before setting or supervising coursework.
- This is general information to help you teach. It is not legal advice.

The six questions

Question	If yes
Does it include names, photos, grades or anything else that could identify a student, parent or member of staff?	Stop and check with your data protection officer or IT lead before going any further.
Is it a public chatbot, and does what you are typing contain anything sensitive or private?	Remove anything you would not want made public, or use a tool your school has approved.
Is it a student's original work, such as an essay, a drawing or a piece of code?	Check the tool's terms and settings, and get permission before the work is used to train a model.
Is it for coursework, or any assessment that counts towards a qualification?	Read the current JCQ guidance and your exam board's rules before you set it or supervise it.
Will students be using the tool themselves?	Check the age limit, and make sure there is supervision and filtering in place.
Will you give what comes out to students or parents, or use it in a document, without reading it carefully?	Read it and check the facts yourself first. You remain responsible for the result.

The evidence

Fact	Condition	Source
The DfE says it sees more immediate benefits and fewer risks from teacher-facing use of generative AI than from student-facing use, and that teachers, leaders and staff must use their professional judgement when using these tools.		Department for Education, Generative artificial intelligence (AI) in education (Last updated 12 August 2025 (England))

Fact	Condition	Source
The DfE recommends that personal data is not used in generative AI tools. If you think it may be needed, check with your data protection officer or IT lead that it is safe, and make sure the data is protected within the tool and not used to further train the AI.		Department for Education, Generative artificial intelligence (AI) in education (Last updated 12 August 2025 (England)); Department for Education, Generative artificial intelligence (AI) and data protection in schools (Published 3 February 2023, updated 14 October 2026 (England))
Schools must include in their privacy notice how any data is collected, processed and stored by generative AI tools, and should be open and transparent about how they use them.		Department for Education, Generative artificial intelligence (AI) and data protection in schools (Published 3 February 2023, updated 14 October 2026 (England))
Some AI tools have age restrictions, and schools must review them before allowing students to use the technology.		Department for Education, Generative artificial intelligence (AI) and data protection in schools (Published 3 February 2023, updated 14 October 2026 (England))
The DfE says students should only use generative AI in education settings with appropriate safeguards in place, such as close supervision and tools with filtering and monitoring, and that schools should follow age limits and Keeping children safe in education.		Department for Education, Generative artificial intelligence (AI) in education (Last updated 12 August 2025 (England))
The DfE says schools and colleges must not allow or cause students' original work to be used to train generative AI models unless they have permission, or an exception to copyright applies.	The DfE says permission would need to come from the student as the copyright owner or, if the student is a minor, from their parent or legal guardian.	Department for Education, Generative artificial intelligence (AI) in education (Last updated 12 August 2025 (England))
The DfE says schools, colleges and awarding organisations need to continue taking reasonable steps to prevent AI-related malpractice in formal assessments, and points to the Joint Council for Qualifications (JCQ) guidance on AI use in assessments.	JCQ updates its guidance, so read the current version before setting or supervising coursework.	Department for Education, Generative artificial intelligence (AI) in education (Last updated 12 August 2025 (England))
The DfE says generative AI can produce nonsensical, inaccurate or false information presented as fact, known as hallucination, and that any content it produces requires critical judgement to check for appropriateness and accuracy.		Department for Education, Generative artificial intelligence (AI) in education (Last updated 12 August 2025 (England))
The NCSC advises not to include sensitive information in queries to public large language models, because the query will be visible to the organisation providing the model.		National Cyber Security Centre (NCSC), ChatGPT and large language models: what's the risk? (Published 14 March 2023)

The Translator

Jargon in plain English, with a comparison from school life. Where the comparison stops working, it says so.

Term	In plain English	Think of it like this
Algorithm	A set of step-by-step instructions for solving a problem or finishing a task. A computer follows it exactly, in the order given.	A recipe, or the steps for long division. Follow them in order and you get the answer. Where that stops working: Some modern AI systems are not written as steps by a person. They find their own rule from examples, which is what the next few entries are about.
Artificial intelligence (AI)	A broad label for computer systems that do things which normally seem to need human intelligence, such as understanding speech, recognising pictures or writing text. It covers many different methods.	'Sport' is a label for football, swimming and chess boxing. AI is a label for many different techniques in the same way. Where that stops working: Nobody agrees exactly where the edge of the label is, and the word is often used for things that are just ordinary software.
Automated decision	A decision about a person made solely by a computer system with no human involved. The ICO's guidance is about such decisions where they have a legal or similarly significant effect, and the safeguards people should have.	A system that decides automatically which students get a place on a course, with no person looking at it. Where that stops working: This is general information, and the ICO's page is draft guidance under consultation.
Average (mean, median, mode)	Three ways of describing the middle of some numbers. The mean adds them up and divides, the median is the middle value in order and the mode is the most common.	Three students at the bus stop earn £5 each from a Saturday job and one earns £100. The mean is about £29, the median is £5. Which one is more honest about a typical student? Where that stops working: No single average describes everyone, which is one reason a headline average can hide how different groups do.
Bias	In data and AI, a system that works worse, or gives systematically different results, for some groups of people. It often comes from the examples the system learned from or from choices made in building it.	A school quiz written entirely by students from one friendship group will probably suit that friendship group best. Where that stops working: In everyday life bias suggests someone meant it. A system can be biased with nobody meaning anything.
Causation	When one thing really makes another happen. Much harder to show than correlation, and usually needs an experiment or very careful design.	Pressing the switch makes the projector come on. Doing it again and again with the same result is how you know.
Confounder	A hidden third thing that causes two other things to move together, so they look related when they are not.	Rain makes more umbrellas sell and more students arrive late. Umbrellas do not cause lateness; the weather causes both.
Correlation	Two things that tend to rise and fall together. It is a pattern in the numbers, and says nothing yet about why.	Students who sit near the radiator tend to be sleepier. The two go together; that does not mean the radiator is the whole reason. Where that stops working: A strong correlation can still be a coincidence or the work of a third thing.

Term	In plain English	Think of it like this
Critical National Infrastructure (CNI)	The systems and facilities a country relies on so much that their loss would be serious, such as energy and water. In 2024 the government announced it would designate data centres as CNI.	The school's power supply: nobody thinks about it until it fails, and then nothing works.
Dashboard	A single screen of charts and key figures that lets someone keep an eye on how something is going.	The attendance board in the office: one glance tells you which classes need chasing. Where that stops working: A dashboard shows what someone decided to measure and nothing else. The Make the same data say three things activity is a warning about that.
Data cleaning	Finding and fixing errors in data before using it: duplicates, wrong units, impossible values, inconsistent labels and gaps.	Marking your own homework for silly mistakes before you hand it in. Where that stops working: Cleaning involves judgement. The Clean the mess activity has a 262 cm height that nobody can safely correct without asking.
Data pipeline	The automated route data takes from where it is collected, through cleaning and checking, to where it is used.	Homework moving from your bag, to the teacher's pile, to the mark book, to the report, without anyone retyping it. Where that stops working: A mistake early in the pipeline travels all the way to the end, so the checking has to happen along the way.
Database	An organised store of data that lets you add, find and change information quickly and reliably.	The school's management system, where one search finds a student, their timetable and their parent's phone number.
Dataset	A collection of data organised for a purpose, often rows (one per person, event or item) and columns (one per thing you measured).	The class register: one row per student, one column per thing you record.
Generative AI	AI that produces new content, such as text, pictures or music, based on patterns it learned from a lot of examples.	A student who has read thousands of stories and can now write one of their own in the same style, without having to copy any single one. Where that stops working: The new content can look convincing and be wrong. See hallucination.
Hallucination	When an AI tool produces false or nonsensical information and presents it as fact. The DfE describes it in just those words.	A student who has not done the reading and bluffs, in a very confident voice. Where that stops working: A bluffing student knows they are bluffing. A model has no idea.
Large language model (LLM)	The NCSC describes it as an algorithm trained on a large amount of text that analyses the relationships between words and turns them into a probability model, which it uses to answer. In short, it works out likely continuations of the text it is given.	The Predict the next word activity is a tiny, readable version of the same idea. Where that stops working: The activity counts word pairs in a few dozen sentences. Real models are trained on vastly more text and use far more powerful methods. The shared idea is choosing a likely next word.

Term	In plain English	Think of it like this
Machine learning	A way of making software where, instead of writing the rules by hand, you show a program lots of examples and let it work out the pattern for itself.	Teaching a new student to spot a good sentence by showing them hundreds of good and bad ones, rather than handing them a rulebook. Where that stops working: A person still decides which examples to show and how to judge whether the student has got it right, and that is where many of the problems start.
Model	What machine learning produces once it has learned from data: a pattern-spotting rule you can use on new cases.	The 'feel' a football coach has for a good striker after watching thousands of matches. It is not written down anywhere, but it can be used on a player they have never seen. Where that stops working: A coach can explain a hunch. A model often cannot explain why it decided what it did.
Neural network	A kind of model made of very many small, connected calculations whose strengths are adjusted during training. The name is loosely inspired by the brain, but it is maths, not a brain.	A huge number of tiny dials. Training is a long process of nudging every dial until the answers improve. Where that stops working: Nobody can read the dials and say what each one means, which is why explaining a decision is hard.
Outlier	A value that sits far away from the rest. It might be an error, or it might be the most interesting thing in the dataset.	The one student who scored 100% when everyone else was between 40 and 60. A mistake, or a genius?
Overfitting	When a model learns its training examples so closely, quirks and mistakes included, that it does badly on new cases.	Learning last year's past paper answers word for word and then freezing when the questions are changed slightly. Where that stops working: A student can tell they are stuck. A model just confidently gives the wrong answer.
Personal data	Any information about a living person who could be identified from it. In schools it includes names, photos, grades and anything that points to a particular student.	Everything in a student's file, and a photo from the school trip. Where that stops working: The details of the law are for your school's data protection officer. The DfE's advice is to keep personal data out of generative AI tools.
Prompt	The text you give an AI tool to say what you want from it.	The question on an exam paper. A vague question gets a vague answer, and a different wording can get a very different one.
Sample	The group you actually measure, used to learn about a larger group you cannot measure completely.	Tasting one spoonful of soup to judge the whole pot. It only works if you stir it first. Where that stops working: If the sample is not representative, you learn about the sample and not the pot. The Teach the machine game is that problem in miniature.
SQL	A language for asking questions of a database, such as 'show me every order over £50 from last March'.	Asking the librarian for every book on volcanoes published after 2010, in a form the library system understands exactly. Where that stops working: It only answers what you ask. A badly worded question gets a confident wrong answer.
Structured and unstructured data	Structured data fits neatly into rows and columns, like a register. Unstructured data does not, like photos, recordings or free-text comments.	The register is structured. The comments teachers write on reports are not.

Term	In plain English	Think of it like this
Test data	Examples kept back and never shown to the model while it learns, then used to find out how well it really works.	The mock exam: questions you have not seen, used to check what you actually know. Where that stops working: A mock only tells you about the topics it covers. A test set that leaves out some people tells you nothing about those people.
Training data	The examples a model learns from. What goes in shapes what comes out.	The textbook a student revises from. If the textbook leaves out a topic, the student will not know it however clever they are. Where that stops working: A student can notice a gap and go looking. A model cannot, unless a person spots it first.

The Curriculum Translator

Where it comes from	What the curriculum says	Activities that support it
Maths, national curriculum	"appropriate measures of central tendency (mean, mode, median) and spread (range, consideration of outliers)"	Clean the mess, Make the same data say three things
Maths, national curriculum	"construct and interpret appropriate tables, charts, and diagrams"	Make the same data say three things, Clean the mess
Maths, national curriculum	"describe simple mathematical relationships between 2 variables (bivariate data) in observational and experimental contexts and illustrate using scatter graphs"	Correlation is not causation, Teach the machine
Maths, national curriculum	"use and interpret scatter graphs of bivariate data; recognise correlation and know that it does not indicate causation"	Correlation is not causation
Maths, national curriculum	"infer properties of populations or distributions from a sample, whilst knowing the limitations of sampling"	Teach the machine, Discussion: the 2020 grades
Maths, national curriculum	"interpret and construct tables and line graphs for time series data"	Make the same data say three things
Maths, national curriculum	"make predictions; interpolate and extrapolate apparent trends"	Make the same data say three things, Predict the next word
Maths, national curriculum	"use a probability model to predict the outcomes of future experiments"	Predict the next word
Computing, national curriculum	"understand how data of various types (including text, sounds and pictures) can be represented"	Your data trail, Clean the mess
Computing, national curriculum	"including collecting and analysing data and meeting the needs of known users"	Clean the mess, Make the same data say three things, From raw data to a decision
Computing, national curriculum	"use technology safely, respectfully, responsibly and securely, including protecting their online identity and privacy"	Your data trail, Before you paste it in
GCSE computer science	"the ethical, legal and environmental impacts of digital technology on wider society, including issues of privacy and cyber security"	Discussion: the 2020 grades, Teach the machine, Your data trail, Before you paste it in
Planned, from September 2028	"The refreshed curriculum will teach pupils AI literacy, how to use data to complete tasks and solve problems, as well as learning about bias in technology."	Teach the machine, Predict the next word, Clean the mess, Discussion: the 2020 grades, Before you paste it in

Be aware of the limits

- The national curriculum computing programmes of study (published in 2013) and the GCSE computer science content do not use the words artificial intelligence or machine learning. Data is named in only a few of the computing statements above, so most of the links here are to maths.
- The last row is a plan, not a curriculum. The government's November 2025 response says the revised national curriculum is due to be published by spring 2027, for first teaching from September 2028, and that

it will replace the computer science GCSE with a computing GCSE and explore a new level 3 qualification in data science and AI. Explore is the government's word, so do not tell students that qualification exists.

- The maths statements are the programmes of study for England. Check your own exam board's specification for how it covers statistics and probability at GCSE.
- The curriculum statements here come from the school curriculum, because that is where the wording is published. If you teach an FE, vocational or apprenticeship course, check your awarding body's specification for the matching content.
- All of this is for England. Scotland, Wales and Northern Ireland have their own curricula, and curriculum documents are reviewed from time to time, so check the current version for your school.

Questions students, parents and colleagues ask

Do students have to be brilliant at maths or coding to work in data and AI?

No, though it depends which part you mean, and it is worth being straight about that. The Career Matrix splits the field into nine areas of work, and some of them, such as machine learning, lean heavily on maths and programming. Others are about decisions, rules and people, such as governance and ethics, product management and operations.

The government's AI labour market survey is also clear that computer science is still the dominant qualification, and that roles increasingly include people who studied social sciences such as psychology and philosophy. Most organisations hire people from jobs that had nothing to do with AI before, which suggests the doors are not as narrow as the stereotype.

Is it worth encouraging students into this? Will AI not take the jobs?

Nobody can promise the future, and we would not trust anyone who did. What we can do is tell you what the evidence says today. The latest government study estimates the UK AI sector employed around 130,000 people in 2025, up by about half on the year before, though it asks you to treat year-on-year growth with care because part of it is better counting. A separate survey found 35% of organisations had AI vacancies they struggled to fill, mostly because candidates lacked work experience.

The honest reading is that employers are short of people who can do the work and have shown they can, which is why the routes that let students build real experience, such as apprenticeships and projects, matter so much. The government has also said it will teach students AI literacy and how to use data in a refreshed curriculum, which is a sign that these skills are expected to matter well beyond a few specialist jobs.

Is it cheating if students use AI tools for homework?

That depends on the task and on your school's policy, and for anything that counts towards a qualification it depends on the awarding body's rules. The Department for Education says schools, colleges and awarding organisations need to keep taking reasonable steps to prevent AI-related malpractice in formal assessments, and it points to the JCQ guidance, which you should read in its current version before you set coursework.

For everyday homework the more useful question is what the student is learning. In a 2026 survey about a quarter of young people said they simply copied AI outputs, and that was higher among students receiving free school meals, so it may be worth thinking about who has the most to lose from copying. The activities here are designed to make students better at questioning what a tool tells them, which helps whatever your policy is.

Can I put students' names or work into an AI tool to save time?

The Department for Education recommends that personal data is not used in generative AI tools. If you think you need to, check with your data protection officer or IT lead first and make sure the data is protected in the tool and not used to train it. Schools also have to say in their privacy notice how any AI tools collect, process and store data.

Students' original work has a separate rule: schools and colleges must not allow it to be used to train generative AI models unless they have permission, or an exception to copyright applies. The NCSC adds a

simple habit for everyone, which is not to put sensitive information into a public chatbot, because the query will be visible to the organisation providing it.

Why do chatbots get things wrong, and say them so confidently?

Because of how they work. The NCSC describes a large language model as an algorithm trained on lots of text that turns the relationships between words into a probability model, and then uses it to answer. It is working out what is likely to come next, which is not the same as checking what is true. The Predict the next word activity lets your class watch this happen in a model small enough to read in full.

Both the NCSC and the DfE are blunt about the result. The NCSC says these models can hallucinate incorrect facts and give convincing-sounding answers that may be only partially correct, particularly on niche topics. The DfE says the content they produce needs critical judgement to check it is accurate and appropriate.

What is the difference between AI, machine learning and data science?

These are overlapping words rather than separate boxes, and people use them loosely, so here is a plain version. AI is the broad label for computer systems that do things that seem to need intelligence. Machine learning is one way of building them, by showing a program many examples and letting it find the pattern. Data science is the wider practice of using data and statistics to answer questions, which sometimes involves machine learning and often does not.

It is also worth knowing for planning that the national curriculum computing programmes of study and the GCSE computer science content do not use the words artificial intelligence or machine learning at all. The government has said the refreshed curriculum will teach AI literacy, and that is a plan, not yet a curriculum.

Why does it matter who works in AI, and where are the women?

The 2025 labour market survey found women accounted for 20% of AI roles, down 4 percentage points since 2020. That is worth sharing with your students, whoever they are.

There is also a reason that goes beyond fairness. The Gender Shades study found a very large gap in how a set of commercial systems worked for different groups. We cannot say from that study that a different team would have prevented it, but it shows why testing a system on the people it will be used on matters, and why people with different experiences are useful when deciding what to test.

How does someone actually get into this without a degree?

There are routes, and the Career Matrix sets them out. The best-known is the Data analyst apprenticeship, a Level 4 standard that typically takes 24 months and describes the job as collecting, organising and studying data to provide business insight. The Matrix also lists a Level 3 Data technician route into it, and a Level 6 degree apprenticeship into machine learning engineering.

In the government's survey, 19% of organisations used apprentices for junior AI roles, up from 3% in 2020, and graduate routes were still the most common at 60%. So apprenticeships are growing, and they are one road among several, not a replacement for the others. The same survey found 88% of organisations rely on on-the-job training, which suggests that what students can show they have already done is likely to count.

How reliable are the numbers you quote?

Every figure on this page comes from a named source we opened and read on the date shown, with the condition the source puts on it, and some of those conditions are significant. The AI sector study calls its estimates experimental and says year-on-year comparisons should be treated with care. The labour market survey had 119 responses and the authors warn that people who respond may differ from those who do not.

We would rather you taught students to ask that question about every number they meet than to trust ours. If you spot something that has changed, tell us and we will fix it.

Sources and how we checked them

Each source was opened and read on 1 November 2026. Where a source warns about its own figures, we pass the warning on.