

Lesson 2: How AI Learns

Machine Learning, Training Data & Recommendation Engines • 60 to 75 Min • KS3

Learning Objectives

By the end of this lesson, students will be able to:

- Explain how machine learning works using the concept of training data and pattern recognition
- Describe the three main types of machine learning (supervised, unsupervised, reinforcement)
- Understand how recommendation algorithms personalise content on platforms like YouTube, TikTok, and Netflix
- Identify both benefits and risks of personalised AI (filter bubbles, echo chambers, engagement traps)
- Reflect critically on their own interactions with AI-powered content

Materials Needed

- Projector with internet access (YouTube homepage demo)
- “A Day with AI” timeline handout (printable below)
- “Be the AI” activity cards: user profiles and content options (printable below)
- Green, yellow, and red sticky notes | 3 per student
- Large poster paper for sticky note wall activity
- Reflection journal (or lined paper)
- Optional: student devices for independent research

Curriculum Links

Computing KS3: Understand how algorithms make decisions using data; recognise the role of data in real-world systems.

PSHE: Media literacy; understanding how technology influences behaviour and choices. Digital wellbeing; recognising persuasive design.

Cross-curricular: Maths (data analysis, probability), Science (hypothesis testing), English (persuasion, bias in media).

Teacher Background Knowledge

Essential reading before teaching this lesson.

How Machine Learning Actually Works

Machine learning (ML) is the most common approach to building AI. Instead of writing explicit rules, the developer provides **training data**, meaning many examples of inputs and desired outputs. The algorithm finds **patterns** in this data and creates a mathematical **model** that can make predictions on new, unseen data.

Example: Spam Filter

1. Collect 100,000 emails, each labelled “spam” or “not spam”
2. The ML algorithm analyses features: sender address, keywords (“free”, “winner”, “urgent”), formatting, links, attachments
3. It learns patterns: emails containing “Congratulations! You’ve won” + suspicious links = 93% chance spam
4. When a new email arrives, the model predicts: spam or not?
5. If the user corrects a mistake, the model improves (feedback loop)

Three Types of Machine Learning

1. Supervised Learning (most common)

Given labelled examples, predict labels for new data.

Examples: Spam detection, face recognition, medical diagnosis, voice assistants, autocorrect.

Analogy: A teacher shows students 100 solved maths problems. Students learn the method.

2. Unsupervised Learning

Given unlabelled data, find hidden patterns or groupings.

Examples: Customer segmentation, anomaly detection (fraud), topic discovery in documents.

Analogy: Given 1,000 mixed sweets, group them by similarity without being told the categories.

3. Reinforcement Learning

An agent learns by trial and error, receiving rewards for good actions and penalties for bad ones.

Examples: Game-playing AI (AlphaGo), self-driving cars, robotic movement.

Analogy: Training a dog: treat for sitting, no treat for jumping.

Recommendation Algorithms Explained

When students open TikTok, YouTube, Netflix, or Spotify, a recommendation algorithm decides what to show. These systems use two main techniques:

Collaborative Filtering: “Users similar to you liked X, so you probably will too.” Amazon’s “People who bought this also bought...” uses this. If 1,000 people who liked the same 5 films as you also liked Film X, Netflix will recommend Film X to you.

Content-Based Filtering: “You liked videos about cats, here’s another cat video.” TikTok analyses the actual content (objects, music, text, duration) and finds similar content.

The key insight for students: These algorithms optimise for **engagement** (keeping you watching/scrolling), NOT for your **wellbeing** (showing you what’s good for you). This is a fundamental tension that drives many concerns about social media.

Filter Bubbles and Echo Chambers

Because algorithms show you more of what you already like, you end up in a **filter bubble**, a personalised information ecosystem that reinforces existing beliefs and preferences. Coined by Eli Pariser in his 2011 book.

An **echo chamber** is the social version: your social media feed only shows opinions you already agree with, creating the illusion that “everyone thinks like me.”

Real-world impact: During the 2016 and 2020 US elections, Facebook’s algorithm was found to amplify divisive political content because it generated more engagement (comments, shares). In 2021, internal Facebook documents (the “Facebook Papers”) revealed the company knew its algorithm was harmful to teenage mental health but prioritised growth metrics.

Key Statistics for This Lesson

- Average UK teenager spends 7+ hours per day on screens (Ofcom 2026)
- TikTok’s algorithm needs just 40 minutes of usage to build an accurate interest profile
- Netflix: 80% of content watched comes from algorithmic recommendations
- YouTube’s recommendation algorithm drives 70% of total watch time
- The average smartphone user generates approximately 40 MB of data per day
- Spotify’s Discover Weekly playlist has generated 2.3 billion hours of listening since launch
- Google processes 8.5 billion searches per day, each one generating data

Detailed Lesson Timeline

Time	Activity	Detailed Instructions
0 to 5 min	Lesson 1 Review	<p>Quick-fire review: Ask 5 rapid questions to the class:</p> <ol style="list-style-type: none"> 1. "What is AI?" (Technology that learns from data) 2. "How is AI different from a regular program?" (Learns from examples vs follows fixed rules) 3. "Name 3 examples of AI" (Siri, Netflix, face unlock, etc.) 4. "Does General AI exist?" (No. Only Narrow AI) 5. "Can AI think?" (No. It processes data statistically) <p>Share the best questions from Lesson 1's exit tickets. Address any remaining misconceptions.</p>
5 to 12 min	Hook: YouTube Demo	<p>Open YouTube on the projector using a class/teacher account. Show the homepage.</p> <p>Teacher script: "Look at this YouTube homepage. There are 20+ videos here. Who chose them? Did a person at YouTube sit and pick videos for us? No. An AI algorithm selected every single video specifically for THIS account based on what it's watched before. If I logged in with a different account, we'd see completely different videos. In fact, no two people on Earth see the same YouTube homepage. Ever."</p> <p>Class reaction: Students are usually surprised. Follow up: "What data do you think YouTube uses to decide what to show you?" Collect responses.</p>

Time	Activity	Detailed Instructions
12 to 25 min	A Day with AI Timeline	<p>Walk through fictional teenager Alex's day, minute by minute. At each moment, pause and ask: "Is AI involved here?" Students complete the timeline handout.</p> <p>Alex's Day:</p> <ul style="list-style-type: none"> • 7:00 - Smart alarm analyses sleep cycle (AI) • 7:05 - Weather app predicts forecast (AI, specifically ML weather models) • 7:15 - News feed personalised (AI algorithm) • 7:30 - Spotify plays Daily Mix (AI recommendation) • 8:00 - Google Maps suggests fastest route (AI traffic prediction) • 8:30 - School email spam filter (AI) • 9:00 - Autocorrect in essays (AI language model) • 10:00 - Google Search for homework (AI ranking) • 12:00 - School canteen app suggests meals (AI) • 15:30 - YouTube after school (AI recommendations) • 16:00 - Video game enemy AI adapts (AI) • 17:00 - TikTok For You page (AI algorithm) • 18:00 - Netflix "Top Picks for You" (AI) • 19:00 - Chatbot for homework help (AI) • 20:00 - Smart speaker music request (AI speech recognition) • 21:00 - Phone face unlock (AI facial recognition) <p>Count together: 30+ AI interactions in one day! Key question: "Were you aware of all of these before today?"</p>

Time	Activity	Detailed Instructions
------	----------	-----------------------

25 to 40 min	Be the AI Activity	<p>Students role-play as a recommendation algorithm. Each group of 3 to 4 receives:</p> <p>User Profile Card: A fictional teenager with specific behaviour data:</p> <ul style="list-style-type: none"> • User A: Watched 8 cat videos, 3 cooking videos, 2 scary videos, liked all animal content • User B: Watched 6 gaming videos, 4 football highlights, clicked on 2 prank videos • User C: Watched 5 music videos, 3 dance tutorials, shared 2 fashion hauls <p>Content Pool: 15 video options with titles and thumbnails (described). Groups must rank their top 5 recommendations for their user.</p> <p>Key questions after:</p> <ul style="list-style-type: none"> • “How did you decide? What data did you use?” (They used viewing history, just like real algorithms) • “Did you recommend anything OUTSIDE the user’s usual interests?” (Probably not. This is the filter bubble) • “Were you trying to show them what’s GOOD for them, or what they’d click on?” (This is the engagement vs wellbeing tension)
40 to 50 min	Sticky Note Wall	<p>Give each student 3 coloured sticky notes:</p> <p>GREEN = An AI use that HELPS me</p> <p>YELLOW = An AI use I’m unsure about</p> <p>RED = An AI use that concerns me</p> <p>Students write one example on each and place on the class wall chart.</p> <p>Review together:</p> <ul style="list-style-type: none"> • Which category has the most notes? (Usually green; students see AI as helpful) • Did anyone put the SAME thing on different colours? (E.g., one student finds TikTok recommendations helpful, another finds them concerning) Why? • Key insight: the same AI can be helpful AND harmful depending on context.
Time	Activity	Detailed Instructions
50 to 60 min	Filter Bubbles Discussion	<p>Explain the concept: “Because algorithms show you more of what you already like, you end up in a bubble, only seeing content that confirms what you already think. It’s like living in a room where all the windows only show views you already agree with.”</p> <p>Discussion questions:</p> <ul style="list-style-type: none"> • “If TikTok only shows you funny animal videos, what types of content do you NEVER see?” • “What happens if someone only sees political views they already agree with?” • “How can you “burst” your filter bubble?” (Follow diverse accounts, search for different viewpoints, use different platforms)
60 to 70 min	Reflection & Journal	<p>Journal prompt (displayed on board):</p> <p>“Think about the last app you used that recommended content to you. What was the AI doing behind the scenes? How did it decide what to show you? Was it helping you discover something good, or was it trying to keep you scrolling? How did it make you feel?”</p> <p>5 minutes writing time. Optional: 2 to 3 volunteers share.</p> <p>Homework: 24-Hour AI Tracking: for the next 24 hours, keep a log of every time you notice AI in your daily life. Bring to Lesson 3.</p>

Differentiation Strategies

Activity	Lower Ability / SEN	Core	Higher Ability
Day with AI	Pre-completed timeline with some blanks. Students fill in 6 answers from a word bank.	Complete timeline with 16 prompts. Identify which are AI.	Create own AI timeline for their actual yesterday. Explain HOW the AI works in each.
Be the AI	Simplified user profile with 3 clear interests. Smaller content pool of 8 options.	Standard user profiles and 15 content options.	Design their own user profile + content pool. Include “trap” content designed to cause a filter bubble.
Sticky Notes	Sentence starters provided: “AI helps me when...”, “I’m worried about AI when...”	Free response on sticky notes.	Write a 100-word argument: “Are recommendation algorithms more helpful or harmful for teenagers?”
Journal	3 sentence starters to complete rather than free writing.	Free journal response to prompt.	Extended writing: Compare how two different platforms personalise content differently.

Assessment Criteria

Level	Descriptor	Evidence
Emerging	Can identify some AI in daily life but struggles to explain how recommendations work.	Timeline partially completed. Sticky notes vague.
Developing	Identifies AI in daily life. Basic understanding that algorithms use data to recommend content.	Timeline mostly complete. Can explain what data YouTube uses.
Secure	Clear understanding of how recommendations work. Can explain filter bubbles. Thoughtful journal entry.	Be the AI activity shows data-driven decisions. Journal reflects on personal experience.
Excelling	Deep understanding of engagement vs wellbeing tension. Can analyse the ethics of recommendation systems.	Nuanced discussion contributions. Journal explores personal and societal implications.

Key Vocabulary Reference Sheet

Recommendation Algorithm: A system that analyses your behaviour data (what you watch, like, click, skip) to predict and suggest content you're likely to enjoy.

Collaborative Filtering: "Users like you also liked X." Recommendations based on what similar users enjoyed.

Content-Based Filtering: "You liked X, here's more like X." Recommendations based on the features of content you've already consumed.

Training Data: The examples used to teach an AI system. Quality and quantity of training data directly affect AI performance.

Filter Bubble: The personalised information ecosystem created by algorithms, which can limit your exposure to diverse viewpoints and content.

Echo Chamber: A social environment where you only encounter opinions and information that reinforce your existing beliefs.

Engagement: The measure of how much time and attention a user gives to an app or platform. Algorithms typically optimise for maximum engagement.

Personalisation: Customising content, recommendations, or experiences for individual users based on their data and behaviour patterns.

Data Footprint: The trail of data you create through your digital activities. Every click, search, like, and scroll generates data.

Persuasive Design: Features deliberately designed to maximise user engagement: infinite scroll, autoplay, notification badges, streaks.

Printable: Be the AI: User Profile Cards

Print one set per group. Cut along lines.

Profile	Watch History	Behaviour Data
User "AnimalLover2024" A:	8 cat videos, 3 cooking videos, 2 scary movie trailers	Liked: all animal content. Average watch time: 95%. Shared: 2 cat videos. Searches: "cute puppies", "funny cats"
User "GoalMachine99" B:	6 gaming walkthroughs, 4 football highlights, 2 prank compilations	Liked: gaming and sports. Average watch time: 80%. Commented on: 3 gaming videos. Searches: "Fortnite tips"
User "DanceStar_X" C:	5 music videos, 3 dance tutorials, 2 fashion hauls, 1 study-with-me	Liked: music and dance. Average watch time: 70%. Shared: 4 music videos. Searches: "K-pop dance", "outfit ideas"

Content Pool (groups rank top 5 for their user):

- 1. "10 Cutest Kitten Moments"
- 2. "How to Make Perfect Pancakes"
- 3. "Scary Movie Trailer Compilation"
- 4. "Fortnite Season Update Guide"

- 5. "Premier League Goals of the Week"
- 6. "Top 10 Pranks Gone Wrong"
- 7. "K-Pop Dance Tutorial: Beginner"
- 8. "Climate Change Explained"
- 9. "History of Ancient Egypt"
- 10. "Fashion Haul: Spring Collection"
- 11. "Dog Rescue Transformation"
- 12. "Minecraft Building Tutorial"
- 13. "Study With Me: 2 Hours"
- 14. "World News Today"
- 15. "Violin Cover: Popular Song"

Discussion after activity: Did any group recommend items 8, 9, or 14 (educational/news content)? Probably not, because the user data didn't suggest interest in those topics. **This is a filter bubble in action.**

Post-Lesson Teacher Notes

What went well?

What would I change?

Common misconceptions:

Students needing support in Lesson 3:

Additional Resources:

- Eli Pariser: "Beware Online Filter Bubbles" (TED Talk, 9 min), excellent for teachers
- The Social Dilemma (Netflix documentary, 2020), teacher viewing recommended

- BBC Own It: How Algorithms Work: bbc.co.uk/ownit
- Common Sense Media: AI and Algorithms resources: commonsensemedia.org
- Spotify's Discover Weekly: How it Works (blog post): newsroom.spotify.com