4565: Computing Foundations for a Digital Age
Unit 5

Day 3 Lesson Plan: How Does AI Learn? + Ethics

Objective:

Students will explore how machine learning models are trained and identify ethical risks related to biased or flawed data.

Warm-Up: Teaching a Computer

- **Ask:"What does it mean to 'teach' a computer something?"**
  Let students journal or discuss with a partner. Push beyond "you type it in" to the idea that computers can *learn* from examples — not just follow instructions.

Mini-Lesson: From Training to Trouble

Introduce key concepts:

- A **machine learning model** is trained using a set of labeled examples (training data).
  → Example: Show it 100 dog photos labeled "dog," and it learns what a dog looks like.
- The model learns to predict based on **patterns** — but if the training data is **biased or incomplete**, the predictions can be flawed.
- **Bias in AI** happens when the data isn't fair or representative. This can lead to **unethical outcomes**.

Video

- 🎥 Crash Course: Supervised Learning (~10 min)

Pause and discuss:

- "What is supervised learning?"
- "What can go wrong if your training data is limited or unfair?"
- "Who is responsible when an AI system makes a mistake?"

Activity: Teachable Machine Exploration

Option 1: Teachable Machine (Google Tool)

- Students use Teachable Machine to train their own basic image model (e.g., rock-paper-scissors, happy/sad faces).
- Have them experiment with what happens if they use unequal examples or strange angles.

Option 2: AI Bias Sandbox (if available)

- Explore a prebuilt demo that shows what happens when you train a model with biased data.

Guiding Questions:

- What happened when the training data was unbalanced?
- Did your model make any incorrect predictions? Why?
- How could this lead to real-world problems?

Project Tie-In

Prompt students to apply today's lesson to their **Ethical Tech Case Study**:

- How might your tech use machine learning?
- Where could bias show up in the training or use of that system?
- Who could be affected by a mistake?

They can start sketching this section in their project journal or planning sheet.

Exit Ticket

"If an AI makes a harmful mistake, who should be held responsible — the programmer, the company, the user, or someone else? Why?"

Word Bytes (Vocabulary + Student-Friendly Examples)

- **Bias** – "Like when my camera always recognizes my face better than my friend's — not fair."
- **Training Data** – "Examples I give the computer to learn from — like teaching it what a cat looks like."
- **Model** – "The part of the program that makes guesses after it learns from data."
- **Prediction** – "When the AI says, 'This must be a dog!' — even if it's a mop."