Name _____

Solving AI's Most Confusing Crimes

AI systems can only learn what humans teach them. If the data they learn from is labeled wrong, things go sideways-fast. Imagine being told that "happy" means "sad" or that every dog is a "cat." You'd get pretty confused, right? So does AI.

You'll investigate short "case files" where data was labeled incorrectly. For each one, figure out what went wrong and how it could mess up an AI's predictions or cause unfair results. Then explain what should have been done differently.

Case 1: The Mixed-Up Moods - An AI was being trained to recognize emotions in text messages. Some examples included:

- "I can't believe I got the job! This is the best day ever!" labeled as **Sad**
- "Ugh, this is the worst!" labeled as **Happy**
- "I guess it's okay, nothing special." labeled as **Angry**

Your Tasks:

1. What labeling errors can you spot here?
2. How might this affect an AI that's trying to understand people's feelings?
3. If this AI ran a chatbot, what kinds of weird mistakes might it make?

Case 2: The Food Mix-Up - A company was training an AI to identify cafeteria food from photos. But their data looked like this:

- A slice of pizza labeled as **Sandwich**
- A banana labeled as **Dessert**
- A brownie labeled as **Fruit**

Your Tasks:

1. What mistakes were made?
2. If this AI worked in a real cafeteria, what problems could it cause?
3. How might customers react if their food was labeled wrong?

Case 3: The Street Scene Confusion - An AI for self-driving cars was learning to recognize things on the road. Here's what the labelers did:

- A stop sign labeled as **Billboard**
- A person crossing the street labeled as **Mailbox**
- A red traffic light labeled as **Go Sign**

Your Tasks:

1. Why are these labels dangerous?
2. What might happen if this AI drove an actual car?
3. What lesson should labelers take away from this case?