

How One Tiny Data Mistake Can Cause a Giant AI Meltdown Answer Key

1. Bad Labeling

Example Answer: The labeling team mixed up medical images - they marked some healthy freckles as "possible issue" and some early cancer signs as "healthy."

Reasoning: Human error during labeling led to incorrect examples in the dataset. These bad labels became part of the AI's "training lessons."

2. Bad Learning

Example Answer: The AI learned false patterns - it started to believe that some dangerous spots looked "safe" and that harmless ones were "problems."

Reasoning: Since the AI can't tell truth from error, it copies what's in the training data. Wrong labels = wrong lessons.

3. Bad Predictions

Example Answer: When used in real hospitals, the AI gave wrong results: it flagged healthy people as sick and missed actual warning signs in others.

Reasoning: The AI's decisions matched its flawed training. Its "knowledge" came from mislabeled data, so its predictions were unreliable.

4. Real-World Consequence

Example Answer: Patients received incorrect medical advice. Some got scared over nothing, and others didn't get treatment in time. Doctors stopped trusting the AI system, and the company's reputation suffered.

Reasoning: AI mistakes in sensitive areas like healthcare can cause emotional, financial, and safety problems for real people.

5. Prevention

Example Answer: They could have double-checked labels, slowed down the labeling process, or had experts review the data before training. They could also test the AI with known examples to catch errors early.

Reasoning: Quality control and careful data review prevent mistakes from spreading into large systems.