Teacher Answer Key

Case Study

1. LLMs generate text by predicting statistically probable words based on patterns from training data, not by looking up verified facts. Because Marcus's question involved a recent event (possibly after the training cutoff) or simply a topic where the model's patterns led to incorrect output, the AI produced confident-sounding but wrong information. The model has no mechanism to know when it is wrong.

2. Marcus should verify every factual claim from an AI against reliable, citable sources before including it in an assignment. He should use AI as a starting point or brainstorming tool, not as a substitute for actual research, and he should never submit AI-written content as his own work without careful fact-checking.

Multiple Choice

1. A

2. B

3. A

Matching

1-E, 2-C, 3-D, 4-A, 5-B

Fill-in-the-Blank

1) text, 2) hallucination, 3) generate

Short Answer (Sample Responses)

Q1: Traditional programs follow explicit rules written by programmers and fail if asked something outside those rules. LLMs learn statistical patterns from massive amounts of text and use those patterns to predict responses, so they can handle a wide variety of inputs without being specifically programmed for each one.

Q2: Training an LLM requires thousands of specialized GPU chips running for weeks or months, consuming enormous amounts of electricity and requiring large teams of engineers and researchers. The computing infrastructure alone costs tens of millions of dollars, making it accessible only to well-funded technology companies.

Q3: LLMs generate text based on statistical probability, not factual accuracy, which means they can produce wrong information with the same confident tone as correct information. Without a source link to verify claims, users must independently check key facts against reliable sources to avoid spreading or acting on misinformation.