

8th Grade AI Logic Breach Quiz

Deconstruct 10 complex scenarios involving algorithmic bias, heuristic search architectures, and the ethical dilemmas of reinforcement learning.

1. In the context of the 'Trolley Problem' applied to autonomous vehicles, which ethical framework focuses on minimizing total harm by prioritizing the lives of the many over the few?

- A. Deontological Ethics
- B. Virtue Ethics
- C. Utilitarianism
- D. Nihilism

2. Large Language Models (LLMs) possess 'Artificial General Intelligence' (AGI) because they can reason across multiple domains with human-level consciousness.

- A. True
- B. False

3. Which term describes the 'invisible' process where an AI model identifies patterns in data that humans cannot explain or trace, often leading to transparency issues?

- A. Open Source Logic
- B. Black Box Effect
- C. Linear Regression
- D. Data Cleaning

4. An AI is trained to recognize pneumonia in X-rays but fails when used at a different hospital because it learned to identify a specific hospital's watermark instead of medical features. This is an example of:

- A. Overfitting
- B. Supervised Learning
- C. Natural Language Processing
- D. Edge Computing

5. In Reinforcement Learning, the system learns through a process of 'Trial and Error' by maximizing which specific variable?

- A. Input Data Volume
- B. The Reward Signal
- C. Processor Clock Speed
- D. User Feedback Loops

6. If an AI hiring tool systematically rejects resumes with feminine-coded words because it was trained on 20 years of male-dominated hiring data, the model is suffering from _____.

- A. Hardware Latency

Name: _____ Date: _____

- B. Algorithmic Bias
- C. Cyber-Physical Misalignment
- D. Encryption Error

7. A 'Generative Adversarial Network' (GAN) consists of two neural networks—a generator and a discriminator—that compete against each other to create realistic data.

- A. True
- B. False

8. Which of the following is the most significant challenge regarding the 'alignment problem' in AI development?

- A. Matching the AI's speed to the user's internet connection
- B. Ensuring an AI's goals match human values and intentions
- C. Aligning the pixels in a computer vision model
- D. Converting hardware components to be carbon neutral

9. When an AI hallucination occurs, what is actually happening within the large language model?

- A. The computer has caught a virus
- B. The AI is purposely lying to the user
- C. The model generates a statistically likely but factually wrong output
- D. The internet connection was lost during processing

10. Symbolic AI (GOFAI) relies on pre-programmed 'if-then' rules rather than learning from massive datasets like modern Machine Learning.

- A. True
- B. False