Name: _____ **Date:** _____

Answer Key: Neural Networks and Algorithmic Bias: 6th Grade CS Quiz

Students analyze how training data shapes AI outcomes and identify potential bias in automated systems during this rigorous assessment or bell-ringer activity.

---

1. If a developer trains a computer vision model using only photos of golden retrievers, why might the AI fail to identify a Siamese cat?

Answer: B) The training data lacked diversity, causing the model to have a narrow understanding.

Machine learning models are limited by their training data; if the data is biased or restricted, the AI's output will reflect those limitations.

2. True or False: A 'Black Box' algorithm refers to a system where the internal decision-making process is difficult for humans to see or understand.

Answer: A) True

In complex deep learning, even the creators often cannot trace exactly how a neural network reached a specific conclusion, lead to the 'black box' label.

3. In the context of AI ethics, when a system consistently produces unfair results against a specific group, it is known as algorithmic _____.

Answer: C) Bias

Algorithmic bias occurs when human prejudices or skewed data sets are reflected in the decisions made by AI systems.

4. Which of the following describes the function of 'Weights' within an artificial neural network node?

Answer: B) The importance or strength assigned to a specific input signal.

In a neural network, weights determine how much influence a particular piece of data has on the final output during the processing phase.

5. True or False: Natural Language Processing (NLP) is the field of AI that allows a smart assistant to interpret the sarcasm in a human's voice.

Answer: A) True

Name: _____ **Date:** _____

NLP involves the complex task of understanding human language, including context, tone, and intent, not just literal word definitions.

6. When an AI improves its performance on a task without being explicitly programmed for every scenario, it is utilizing _____ learning.

Answer: B) Machine

Machine learning is the specific subset of AI where algorithms find patterns in data to adapt and improve over time.

7. Consider an AI designed to predict weather patterns. If the model is 'Overfitting,' what is happening?

Answer: C) The AI memorized the training data so well it cannot handle new, unseen data.

Overfitting happens when a model learns the 'noise' in the training data rather than the actual patterns, making it fail in the real world.

8. True or False: Generative AI creates new content, whereas Discriminative AI focuses on classifying existing data into categories.

Answer: A) True

This is a key distinction: Generative models (like DALL-E) build things, while Discriminative models (like a spam filter) sort things.

9. A city uses AI to optimize traffic lights. What is the most likely ethical concern regarding this specific application?

Answer: A) The AI might give priority to wealthier neighborhoods if the data is biased.

Ethical AI analysis requires looking at how automated decisions might unfairly impact different demographic or socioeconomic groups.

10. The process of a human checking an AI's output to ensure it is correct and safe is called 'Human-in-the-_____.'

Answer: C) Loop

A 'Human-in-the-loop' system ensures that AI makes recommendations but a human provides the final oversight to prevent errors.