Name: _____ **Date:** _____

Answer Key: Analyze Advanced Algorithms: Collegiate AI Assessment

Evaluate the architectural nuances of backpropagation, stochastic gradient descent, and the ethical implications of algorithmic bias in high-stakes decision systems.

---

1. In the context of the 'Bias-Variance Tradeoff' in machine learning, which phenomenon is most likely to occur if a model architecture is excessively complex relative to the size of the training dataset?

Answer: B) High variance, potentially leading to overfitting the training noise.

High variance occurs when a model is overly sensitive to small fluctuations in the training set, capturing noise rather than the underlying pattern, which results in poor generalization.

2. The 'Vanishing Gradient Problem' is primarily associated with deep recurrent neural networks (RNNs) using saturating activation functions like Sigmoid or Tanh.

Answer: A) True

During backpropagation, gradients are multiplied by small derivatives of saturating functions, causing the signal to diminish exponentially as it moves through deeper layers.

3. Identify the optimization technique where the gradient is calculated and weights are updated based on a single, randomly selected training example per iteration.

Answer: C) Stochastic Gradient Descent (SGD)

Stochastic Gradient Descent (SGD) uses only one training sample per step, which introduces noise into the convergence path but significantly reduces computational overhead.

4. When evaluating the performance of a fraud detection AI where the cost of a 'False Negative' is extremely high, which metric should the lead researcher prioritize?

Answer: C) Sensitivity (Recall)

Recall measures the proportion of actual positives correctly identified. In fraud detection, missing a fraudulent event (False Negative) is more dangerous than a false alarm.

5. Which architectural feature differentiates 'Generative Adversarial Networks' (GANs) from standard deep learning models used for classification?

Answer: A) The use of a zero-sum game between a generator and a discriminator.

Name: _____ **Date:** _____

GANs consist of two networks (Generator and Discriminator) competing; the generator learns to create data while the discriminator learns to distinguish real from fake.

6. Transfer learning involves taking a pre-trained model and fine-tuning it on a new, related task to leverage existing feature representations.

Answer: A) True

Transfer learning is efficient because the model has already learned low-level features (like edges in images) and only needs to learn high-level features for the specific new task.

7. In Reinforcement Learning, the _____ is the mathematical framework used to model decision-making in environments where outcomes are partly random.

Answer: B) Markov Decision Process (MDP)

MDPs provide a formal way to describe a Reinforcement Learning environment using states, actions, rewards, and transition probabilities.

8. An AI model used for granting bank loans systematically denies applications from a specific demographic despite not being given 'race' as a variable. What concept best explains this?

Answer: C) Proxy Variables/Algorithmic Bias

Even without explicit labels, an AI can learn bias through proxy variables (like ZIP codes or browsing habits) that correlate strongly with protected demographic characteristics.

9. Heuristic search algorithms, such as A*, are considered 'weak AI' because they rely on specific domain-based rules rather than general consciousness.

Answer: A) True

Weak AI (or Narrow AI) refers to systems designed to handle specific tasks or problem-solving areas, as opposed to General AI which would mimic human cognitive breadth.

10. The _____ mechanism in Transformer architectures allows the model to assign different weights to different parts of the input sequence dynamically.

Answer: A) Self-Attention

Self-attention enables Transformer models (like BERT or GPT) to process context by focusing on relevant words in a sentence regardless of their distance from each other.