Name: _____ **Date:** _____

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?

A. High bias, potentially leading to underfitting of the training data.
B. High variance, potentially leading to overfitting the training noise.
C. Low variance, ensuring consistent performance on unseen test sets.
D. Convergence toward a global minimum through linear regression.

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

A. True
B. False

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

A. Batch Gradient Descent
B. Mini-batch Gradient Descent
C. Stochastic Gradient Descent (SGD)
D. Momentum-based Optimization

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?

A. Overall Accuracy
B. Precision
C. Sensitivity (Recall)
D. Specificity

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

A. The use of a zero-sum game between a generator and a discriminator.
B. The reliance on supervised learning with large labeled datasets.
C. A feedback loop that only minimizes Mean Squared Error.
D. The exclusion of backpropagation in the training phase.

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

A. True

Name: _____ **Date:** _____

B. False

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

A. Convolutional Layer
B. Markov Decision Process (MDP)
C. Turing Test
D. K-Nearest Neighbor

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?

A. Weight Regularization
B. Unsupervised Clustering
C. Proxy Variables/Algorithmic Bias
D. Dimensionality Reduction

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

A. True
B. False

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

A. Self-Attention
B. Max Pooling
C. Recursive Loop
D. Linear Activation