Name: _____ **Date:** _____

Silicon Synapses: The 12th Grade Neural Network Challenge

Synthesize the architecture of deep learning models and evaluate the ethical implications of weight bias in high-stakes autonomous decision-making algorithms.

---

1. In the context of optimizing a deep neural network, how does the 'vanishing gradient problem' specifically impede the training of early layers in an architecture like a Recurrent Neural Network (RNN)?

A. The gradient becomes exponentially small during backpropagation, leading to negligible weight updates.
B. The learning rate increases automatically, causing the model to overfit on the initial training batches.
C. Data sparsity in the input layer creates a bottleneck that prevents ReLU activation functions from firing.
D. The loss function becomes convex, trapping the model in a local minimum that cannot be escaped.

2. True or False: In Generative Adversarial Networks (GANs), the 'Discriminator' model's primary objective is to maximize the probability of an image being classified as 'real' regardless of its origin.

A. True
B. False

3. When training a model to identify gravitational waves in LIGO data, a researcher uses _____ to prevent the model from memorizing noise rather than learning generalizable patterns.

A. Hyperparameter Tuning
B. Data Augmentation
C. Regularization (e.g., Dropout)
D. Linear Regression

4. Which specific architectural innovation allowed Transformer models to outperform traditional LSTMs in Natural Language Processing tasks by processing entire sequences simultaneously?

A. Stochastic Gradient Descent
B. Multi-Head Self-Attention
C. K-Nearest Neighbors Clustering
D. Recursive Decision Tree Pruning

5. True or False: Reinforcement Learning (RL) relies on a 'reward signal' to guide an agent toward optimal behavior through trial and error within a defined environment.

A. True
B. False

6. In the context of AI ethics, the 'Black Box' problem refers to the lack of _____, where humans cannot easily trace the logic used by a complex neural network to reach a specific conclusion.

Name: _____ **Date:** _____

A. Interpretability
B. Computational Power
C. Data Storage
D. Latency

7. Consider an AI system designed to predict patient outcomes in a hospital. If the training data primarily features individuals from high-income urban areas, what type of algorithmic risk is most likely to emerge?

A. Hardware Latency
B. Data Degradation
C. Algorithmic Bias (Selection Bias)
D. Unsupervised Entropy

8. True or False: Convolutional Neural Networks (CNNs) are primarily structured to exploit the spatial hierarchy of data, making them ideal for image-related tasks.

A. True
B. False

9. An ensemble method that combines the predictions of several weak decision trees to create a strong predictive model, often used in Kaggle competitions, is known as _____.

A. Backpropagation
B. Gradient Boosting
C. Transfer Learning
D. Sigmoid Activation

10. Which concept defines the theoretical point at which artificial intelligence surpasses human intelligence across all domains, potentially leading to rapid self-improvement cycles?

A. The Turing Threshold
B. Technological Singularity
C. Moore's Limit
D. Heuristic Saturation