Name: _____ **Date:** _____

Answer Key: Sizzling Silicon: Senior Synthetic Intelligence Seminar Quiz

Moving beyond automated routines, these logic puzzles challenge students to analyze algorithmic bias and the architecture of recurrent neural networks.

---

1. In the context of architectural design for sequence processing, which mechanism allows a Transformer model to weigh the importance of different parts of the input data regardless of their distance?

Answer: B) Self-Attention Mechanism

Self-attention allows the model to look at other words in the input sequence to help get a better encoding for the word in focus, effectively capturing long-range dependencies better than traditional RNNs.

2. When an AI model performs exceptionally well on training data but fails to generalize to unseen test data, it is experiencing _____.

Answer: D) Overfitting

Overfitting occurs when a model learns the noise and details in the training data to the extent that it negatively impacts the performance of the model on new data.

3. True or False: In Reinforcement Learning, the 'agent' learns through a system of rewards and penalties without requiring a labeled dataset of correct input-output pairs.

Answer: A) True

Reinforcement Learning relies on exploration and feedback from an environment (rewards/penalties) rather than the explicit labels used in Supervised Learning.

4. Which of the following best describes 'Algorithmic Bias' in the deployment of predictive policing or recidivism software?

Answer: B) The unintentional reflection of historical societal prejudices present in the training data.

Algorithmic bias often stems from historical data that contains human prejudice; the AI learns these patterns as facts, perpetuating systemic inequalities.

5. The process of using a pre-trained model on a new, related task to save computational resources and time is known as _____ Learning.

Answer: A) Transfer

Name: _____ **Date:** _____

Transfer learning involves taking the weights and features learned from a massive dataset (like ImageNet) and applying them to a more specific task with a smaller dataset.

6. True or False: A Convolutional Neural Network (CNN) is primarily designed to process sequential data like audio or text through its recursive feedback loops.

Answer: B) False

CNNs are primarily used for spatial data (like images) using kernels to detect features; Recurrent Neural Networks (RNNs) are the ones designed with feedback loops for sequential data.

7. What is the primary function of the 'Discriminator' in a Generative Adversarial Network (GAN)?

Answer: C) To distinguish between real data and data produced by the Generator.

In a GAN, the Generator tries to fool the Discriminator, while the Discriminator tries to correctly classify samples as 'real' (from the training set) or 'fake' (from the Generator).

8. In the context of AI Ethics, the ability to trace an AI's decision-making process in a way that humans can understand is called _____.

Answer: C) Explainability

Explainability (or XAI) is crucial for high-stakes AI applications like medicine or law, ensuring that the 'reasoning' behind an output is transparent to human auditors.

9. True or False: Artificial General Intelligence (AGI)—the ability for a machine to perform any intellectual task a human can—is currently the standard technology used in modern smart assistants.

Answer: B) False

Modern assistants use Narrow AI (or Weak AI), which is programmed for specific tasks. AGI remains a theoretical future milestone in the field of computer science.

10. Consider an AI used for autonomous drone flight. If the system uses 'Computer Vision' to map its surroundings, what is the role of the 'Activation Function' in its neural network nodes?

Answer: B) To determine whether a neuron should be 'fired' based on the input signal.

The activation function (like ReLU or Sigmoid) introduces non-linearity into the network, allowing it to learn complex patterns and decide if the information passed to a node is relevant.