HUM 220N Instructor Definitions

AI and the Human Experience | Foundations and algorithm types

Suggested definitions for the group worksheet. Accept equivalent wording. These brief summaries cover vocabulary beyond the game clues.

TERM OR THINKER	DEFINITION OR MAIN IDEA
Algorithm	A step-by-step procedure for solving a problem or performing a task.
Input	The information supplied to a system, model, or algorithm for processing.
Output	The result produced by a system, model, or algorithm after processing an input.
Model	A representation of patterns or relationships used to make predictions or produce outputs.
Machine learning	Methods that learn patterns from data rather than having every rule explicitly written by a programmer.
Neural network	A model made of connected layers of numerical units, with weights adjusted during training.
Large language model (LLM)	A large neural network trained on language data to predict tokens and generate text.
Linear regression	A method that learns a linear relationship to predict a numerical value, such as a price.
Logistic regression	A method that estimates class probabilities, often for two categories such as spam or not spam.
Decision tree	A model that makes predictions by following a sequence of branching tests on input features.
Sorting algorithm	A procedure that puts items in a specified order. Examples include quicksort and merge sort.
Binary search	A search procedure that repeatedly halves a sorted search space to find a target.

Use short definitions that explain what the term does or why the idea matters.

HUM 220N Instructor Definitions

AI and the Human Experience | Learning types and tasks

Suggested definitions for the group worksheet. Accept equivalent wording. These brief summaries cover vocabulary beyond the game clues.

TERM OR THINKER	DEFINITION OR MAIN IDEA
Dijkstra's algorithm	An algorithm that finds shortest paths from a starting point in a graph with nonnegative edge weights.
Encryption algorithm	A procedure that uses a key to transform readable data into encrypted data. Examples include AES and RSA.
Supervised learning	Learning from examples that pair inputs with correct labels or target outputs.
Unsupervised learning	Finding patterns or structure in data without supplied target labels, such as grouping similar examples.
Reinforcement learning	Learning how to act through interaction with an environment, using rewards to improve a strategy over time.
Labeled data	Examples paired with target answers or categories, such as emails marked spam or not spam.
Classification	Predicting which category an input belongs to, such as spam or not spam.
Regression	Predicting a numerical value, such as a house price or temperature.
Clustering	Grouping examples by similarity without being given the correct group labels in advance.
Training	Adjusting a model's learned parameters using data and a learning objective.
Inference	Using a trained model to produce an output from an input; ordinary inference does not update its learned weights.
Training data	Examples used to fit a model. Choices about what to include or exclude shape what it learns.

Use short definitions that explain what the term does or why the idea matters.

HUM 220N Instructor Definitions

AI and the Human Experience | Training and text representations

Suggested definitions for the group worksheet. Accept equivalent wording. These brief summaries cover vocabulary beyond the game clues.

TERM OR THINKER	DEFINITION OR MAIN IDEA
Weights	Learned numerical parameters that control how a neural network processes information.
Loss	A numerical measure of how poorly a model's predictions meet its training objective.
Backpropagation	Calculating how each parameter contributes to loss by working backward through a network. An optimizer uses these gradients to update weights.
Gradient descent	An optimization method that adjusts parameters in a direction that locally reduces loss, using calculated gradients.
Pretraining	The broad initial training stage in which a model learns patterns from a large dataset.
Fine-tuning	Additional training that adapts an existing model to a particular task or behavior.
RLHF	Reinforcement learning from human feedback: using human preferences to help train a reward model that guides further model training.
Token	A unit of text processed by a language model, such as a word, part of a word, or punctuation.
Tokenization	Splitting text into tokens and mapping those tokens to numerical IDs.
Token ID	A token's index in the tokenizer's vocabulary. The ID identifies the token; it is not its embedding.
Vector	An ordered list of numbers. In a language model, vectors represent tokens and intermediate information.
Word / token embedding	A learned numerical vector representing a word or token. In an LLM, a token embedding is a starting representation for processing.

Use short definitions that explain what the term does or why the idea matters.

HUM 220N Instructor Definitions

AI and the Human Experience | Inside a large language model

Suggested definitions for the group worksheet. Accept equivalent wording. These brief summaries cover vocabulary beyond the game clues.

TERM OR THINKER	DEFINITION OR MAIN IDEA
Hidden state	A token's intermediate representation inside the network; it changes as layers process contextual information.
Self-attention	A mechanism that uses attention weights to combine information from permitted token positions into contextual representations.
Query	In attention, a vector compared with key vectors to calculate attention scores.
Key	In attention, a vector compared with a query to help determine how strongly to weight information at a token position.
Value	In attention, a vector containing information that is combined with other value vectors using attention weights.
Multi-head attention	Several attention operations running in parallel, allowing different learned relationships to contribute to a representation.
Transformer	A neural-network architecture that uses attention and feed-forward blocks to update token representations through layers.
Feed-forward network	A block that transforms each token's representation using learned weights and nonlinear operations.
Context window	The limit on how much token context a model can process at once.
Next-token prediction	Computing probabilities for the token that comes next. A token is selected and appended, and generation continues.
Temperature	A setting that changes the probability distribution used when sampling tokens. Higher values generally make selections more varied.
Hallucination	A plausible-sounding but false or fabricated claim generated by an AI model.

Use short definitions that explain what the term does or why the idea matters.

HUM 220N Instructor Definitions

AI and the Human Experience | Thinkers and course ideas

Suggested definitions for the group worksheet. Accept equivalent wording. These brief summaries cover vocabulary beyond the game clues.

TERM OR THINKER	DEFINITION OR MAIN IDEA
Cognitive offloading	Moving some cognitive work, such as remembering or organizing information, to an external tool or resource.
John Searle	Chinese Room argument: manipulating symbols according to rules does not by itself establish understanding.
Alan Turing	The Imitation Game approaches machine intelligence through observable conversational performance.
Howard Gardner	Multiple intelligences: intelligence includes different capacities rather than being adequately captured by one test score.
Marshall McLuhan	The medium is the message: the form of a medium shapes human experience, not just the content it carries.
Neil Postman	The Judgment of Thamus highlights technological trade-offs: a technology can give us something while taking something away.
Kate Crawford	Classification is not neutral. Choosing categories involves power and can exclude or harm people.
Media ecology	The study of media as environments that shape communication, perception, culture, and social life.
Medium	A means or form of communication, such as speech, writing, print, or television, that shapes how communication happens.
Technopoly	Postman's term for a culture in which technology becomes a dominant source of authority for social decisions and values.
Weak AI	In Searle's distinction, computers are useful tools for studying or simulating minds; simulation alone does not establish understanding.
Strong AI	In Searle's distinction, the claim that the appropriately programmed computer literally has a mind and understands.

Use short definitions that explain what the term does or why the idea matters.