APPENDIX E:

Glossary of Key Terms

This glossary is based on the AI4K12 Glossary.

Abstraction	Simplification by the elimination of unimportant details; one of the core components of computational thinking.
Agent	Any artificial intelligence program that interacts with the world via a *sense-deliberate-act* cycle (sometimes called "sense-think-act"). Agents may be physical devices such as robots or exist purely as software, such as automated stock-trading programs.
Artificial Intelligence (AI)	A body of techniques that allow computers to do things that, when humans do them, are considered evidence of intelligence, including natural language processing, recognizing images, and generating content. AI systems use training data and, based on defined or learned patterns, produce outputs like predictions, recommendations, or decisions. AI4K12 defines Five Big Ideas for K-12 AI instruction: Perception, Representation and Reasoning, Learning, Natural Interaction, and Societal Impacts.
Bias	Systematic error in the behavior of a model due to incorrect assumptions or a mismatch between the data the model was trained on and the real world. See Appendix D.
Chatbot	A conversational agent that communicates with humans via a text or voice interface. The simplest chatbots just search for keywords and return scripted replies, but more sophisticated chatbots may fall in the category of intelligent assistants.
Classifier	An algorithm that examines the features of an input pattern and assigns that input to a category (or class). Classifiers have many practical applications, such as spam filtering (using the classes "spam" and "not spam"), or sorting loan applications, where the classes might be low, medium, or high risk. Classifiers can be programmed by hand, but it is more common to train them using machine learning techniques and examples of each class.
Decision Tree	A method of classification based on a series of tests of features of the input. One advantage of decision trees is that their decisions are explainable by looking at the tests that were performed.
Dichotomous Key	An alternative name for a decision tree classifier, used in biology.
Explainable	Explainable AI provides a justification for its decisions by referring to features of the input that led to the decision. For example, if a loan application is categorized as high risk, this decision might be accompanied by an explanation such as a high debt to income ratio or an insufficient employment history.

Fair/Fairness	The decisions of a fair algorithm result in equal outcomes for people who should be treated equally. However, there are multiple, mutually incompatible definitions of fairness (e.g., equal outcomes versus equal risks), so no algorithm can be fair in every sense.
Features	The attributes of an input which are used to reason about it. Features can be numerical values (e.g., age) or categorical values (e.g., state of residence).
Foundational Models	Foundational Models AI models that are trained on massive datasets to learn a broadly useful set of features that allows them to be applied (with additional training) to many different specialized tasks.
Generative AI	Technologies that can generate new text, images, video, and/or audio in response to a prompt. Modern generative AI systems are created using machine learning algorithms run on massive amounts of training data.
Human-in-the-Loop	A system in which humans interact with an automated decision making system to guide its behavior and ensure that its actions are reasonable.
K-Nearest Neighbors (KNN)	A classification algorithm that finds the k-nearest training examples to an input pattern, and assigns that input the same class as the majority of the k neighbors. The definition of "nearest" depends on the features used to represent the patterns.
Large Language Model (LLM)	A neural network trained on massive amounts of text that can be used in a variety of language tasks such as sentence completion, question answering, machine translation, and chatbot functions. Large language models are one of the technologies that make up generative AI.
Logic	A system of formal reasoning using symbolic representations and rules of inference. Much of early AI was based on logic, whereas modern AI relies more on statistical reasoning.
Machine Learning	A subfield of artificial intelligence focusing on combining learning algorithms with training data to create models that can be used to complete tasks. Three important types of machine learning are *supervised learning*, *unsupervised learning*, and *reinforcement learning*. Modern AI applications, such as speech recognition systems, are often constructed using machine learning techniques applied to huge training sets.
Model	A tool for understanding a relationship or phenomenon, used to make predictions or in conjunction with algorithms to process input into output. Models may be based on equations, statistics, rules, agents, neural networks, or other representations. They may be created by humans making sense of data or by computers processing data with algorithms.
Model Card	A model card provides information about an AI model, such as its funding, training data, and performance on benchmarks.
Multimodality	The ability to process multiple types of input, such as text, images, video, and sound.

Natural Language	The language people use to communicate with each other in everyday life. This is distinct from computer languages used for programming or for giving commands to a computer. Computer languages have a rigid vocabulary and syntax, while natural language is fluid and ambiguous, and it includes phenomena such as imagery, humor, sarcasm, and alliteration. Natural language understanding is difficult for computers.
Neural Network	An approach to computing in which many simple processing units are organized into a network to collectively solve a complex problem. Neural networks draw inspiration from theories about how the brain might work, but their components and organizational structure are far simpler than real neural tissue.
Perception	The extraction of meaning from sensory signals. A microphone senses sound and converts it to an electrical signal. Understanding the sound (e.g., recognizing the words being spoken or the music being played) constitutes perception.
Predictor	A reasoner that produces a continuous value as its output, such as estimating the market value of a house based on its features. Predication is also known as regression. Compare classification, which outputs one of a finite set of class labels instead of a continuous value.
Reasoning	The process of making decisions or solving problems. There are many types of reasoning, including classification, prediction, recommendation, planning, and sampling.
Recommender	A reasoner that suggests items the user might like based on what other users who like similar things have liked. Recommendation can be used to suggest items to purchase, ads to show, or stories to include in a news feed.
Reinforcement Learning	A machine learning technique where the training data is labeled with a "reward" signal telling the computer how good its outputs have been. This differs from *supervised learning* where each training example's label tells the computer exactly what output it should have produced. Reinforcement learning is typically applied to sequential tasks where the reward signal comes only at the end. An example is game playing, where the reward signal comes after the final move and tells the computer whether it won or lost the game. Using reinforcement learning, computers have become experts in domains such as backgammon by playing against themselves for thousands of games.
Representation	An encoding of information in a way that is useful for reasoning. The concept of representation in AI is analogous to that of data structures in computer science.
Search	A reasoning method that involves the systematic exploration of possibilities until a solution is found.
Sensing	Translating a physical phenomenon into an electrical signal that can be measured and acted upon.

Supervised Learning	A machine learning technique where the training data consists of input examples plus the desired output for each input. The desired outputs are called labels, and the data is called labeled data. For example, the training data could be pictures, some of which contain cats, and the labels could be 1 for cat or 0 for no cat. The learning algorithm goes through the training set repeatedly, and it slowly adjusts the model's parameters to make it more likely to produce the correct output for each input.
Training Data	A collection of examples that can be used by a machine learning algorithm to construct a reasoner. Training data is labeled if we are told the correct class of each example. Data can also be unlabeled, in which case it may be used for clustering.
Transparency	Disclosing the details of a system's design or operation. For example, in an automated decision making system, transparency may require disclosure of the training data used to build the system, so users can assess whether the data is likely to be biased.
Unsupervised Learning	A machine learning technique that finds structure in unlabeled data. One example is *clustering*, where the computer examines a collection of examples and groups them into categories based on perceived similarity. In a sense, the computer "discovers" the categories; they are not given to it. (If each example came labeled with its correct category, this would be *supervised learning*).

Photo: Convening participants discuss trends and patterns in AI curriculum.