

AI 101

YOUR DATA SHOULD TEACH YOU THINGS
ABOUT ITSELF YOU DID NOT KNOW



Contents

The Basics	<hr/>	01
Data		
Asking Questions		02
Typical Value Issues		
Definitions	<hr/>	03
General Jargon Definitions		
Extended Jargon Definitions		04
Concept Definitions		05
Nuanced Definition of Artificial Intelligence		06

The Basics

What is "Artificial Intelligence?" – AI is an extension of Machine Learning where the goal is to create models and systems that can be or are more than the sum of their parts.

Is AI a monolithic concept? – No! Many models, layers, inputs and outputs, and concepts can be classified as "AI" in one form or another. The throughline is the idea of emergence.

Why is AI hard? Unlike previous generations of information technology, AI is a general concept, focused on how a system is constructed instead of what a system is composed of. AI is both open-ended and customized simultaneously, requiring often unique insights and engineering to fill gaps to solve business problems.

What is GenAI? – GenAI refers to systems that generate ("Gen") new information, data, content, and work product from disparate open source and proprietary data.

What is an LLM? – A Large Language Model is how you have likely interacted with AI - think ChatGPT. LLMs use large training data sets to understand how humans use language to convey ideas and knowledge. In this sense, language can be considered not only words but also images and sounds.

Why do LLMs appear smart and dumb at the same time? – LLMs form an understanding of the way words are put together. However, the complex concepts conveyed by human language are not always emergent with scale. Therefore, the higher the abstraction, the less understanding the model will have.

Data

What is "data-driven" in the age of AI? – A data-driven organization must consider anything that can be measured, collected, or inferred as data. This can be everything from classic spreadsheets to the telemetry, audio/visual, and contextual information of a random marketing meeting.

What kinds of data do I need to have? – LLMs and other Deep Learning/AI technologies are bound by the data they are trained on. Data that is proprietary, abstract, or contains multitudes of rare connections are important to create foundational AI infrastructure for business.

What new kinds of data can and should my business be collecting? – Data collection should be unconstrained, such that no prior assumptions of the usefulness and leverage related to a data source should be used when determining if/when data should be collected and stored without new analysis.

Asking Questions

What are the right questions? The idea of AI is to allow a system to augment knowledge, elevate information, and create new outputs. Thus, asking where in a chain of events a human is essential is the starting point for asking what AI can be for an organization. The next is - Human synthesis, but at scale...

What kinds of questions – old and new – can and should my business be asking? – Focusing on augmenting and scaling a workforce is an old question given new life with AI. However, equally powerful opportunities exist in questions regarding the emergence and leveraging of information previously too disparate or out of reach.

How do we ask new kinds of questions? – Counterintuitively, the best place to start is thinking inside the oldest of boxes - What does my company need, and how do I find/train a person to fill this role? AI, if done right, should be thought of as a malleable person – harder to teach but much more productive than their human counterparts.

Typical Value Issues

Where can I add value and at what cost? – The complexity and completeness of DL/AI infrastructure can be expensive to initialize. However, the open-ended flexibility of inputs and outputs, amortization of costs, and massive increases in productive for human employees will provide impressive long-term gains and differential advantages for organizations.

Can determinism and probability coexist? – Moving from a deterministic relational database to a probabilistic AI system is akin to moving from classical physics to quantum physics, where probability rules all. Coexistence lies in a similar manner, such that determinism can be constructed on top of a probabilistic base, just as Newton still guides us when talking about planets, while probability rules the building blocks of those planets.

What is customization, alignment, and problem-solving in my use case? – Your use case and data are contextually specific and must be learned for a system to work for your business. This requires customizing the underlying models, aligning large models to understand your business, and building layers to an AI system to help flexibly solve current, future, and emergent problems.

Why not use an off-the-shelf solution? – The problems involved in creating DL/AI systems revolve around the need to typically perform basic research and invent new methods to deliver a value-adding solution. This bottleneck is why there are no, and will not be for some time, generalized SaaS solutions that involve AI as a systemwide, value-adding augmentation and infrastructure applications for large businesses and enterprises.



General Jargon Definitions

Model – A design for describing a data-generating process. In AI, a model refers to the specific formulation of a design, such that a transformer model is distinct from a graphical neural network or to a broad class of models, such as language or vision models.

Emergent – Properties of a model that arise without necessarily being defined. An emergent property can arise from scale (such as in an LLM) or from the interaction and leveraging of information that does not have obvious coherence.

Modality – The input/output focus of a model. This can be language, audio, visual, or concepts. A model's modality defines the input/output from which it can function and be applied to the world.

Encoding – An encoding is a numerical representation of information. An encoder/decoder constructs a numerical representation of data so that it can be used in a model.

Embedding – An embedding relates encoded data to other encoded data within a model. It is the typical building block of a DL/AI model and/or system, allowing, for example, an LLM to understand what word to use next in a generalized context.

Latent – When an event or relationship cannot be observed (measured/collected). The *latent* properties of things are hidden from view and, therefore, must be elevated through the analytical frameworks we construct to observe them. The latent properties of our world are typically the most interesting, if not most important, to understanding intelligence, rationality, and action.

Dimension – The relative geometric representation of the data. Among several possibilities, a dimension may describe how large an encoding a model may create/use or the space a model is designed to represent. Typically, the dimensionality of a model will refer to the size of the matrix representation of the data created by the embedding model.

Vector – A vector describes the magnitude and direction of a numerical representation. In AI, a vector describes the encoded/embedded representation of the input data and is the functional unit of leverage used for input and output.

Token – The base unit of input and output for a model. In an LLM, for example, a word is the primary token of the model, which is then understood relative to all other tokens. Different models take different approaches to tokenizing the data they are trained and refined on, leading to various properties and capabilities of the model (usually with positive/negative tradeoffs).

Query – A query provides the relationship information to the model in training. In contrast, during inference, a query describes the input from which the model determines the correct output. The relationship between training and interacting with a DL/AI model centers around how the model understands a query and uses queries to relate data.

Augment – The shifting of the available information to a query. This can be increasing the information or simply changing the information available depending on the desired result. An augmented system, therefore, extends, changes, or restricts information similar

Extended Jargon Definitions

Space – A space is a mathematical term for literally where something exists, can exist, or where something can fill when formally abstracted into a numerical (loosely) form.

Variable – A variable is something that exists to be acted upon or act within a function. A variable represents a quantity that can change or be changed within a system.

Parameter – A parameter is a variable that defines a particular system, model, or function characteristic. It often remains constant for a given scenario but can be changed to alter the system's behavior or output.

Non-Parametric – In statistics and machine learning, non-parametric methods do not rely on data belonging to any particular probability distribution. They are more flexible than parametric methods as they make fewer assumptions about the data.

Label – A label is a category or class assigned to a data point. Labels are used in supervised learning to train models to classify new, unseen data.

High-Dimensional – Refers to data with many features, such that it is difficult to figure out what is causing action, what is not doing anything, and what is being acted upon. This is often called the “curse of dimensionality” and is a significant reason for many techniques used in advanced analytics and DL/AI systems.

Multi-Modal – A multi-modal system can simultaneously process several input data types, such as text, images, and audio.

Context – The setting we look for to help provide meaning to something. In DL/AI, which is related to natural language, context often refers to the surrounding words or sentences that help determine the definition of a word or phrase.

Gate – A gate controls how information flows through a neural network model. See frame...

Frame – A frame refers to data structures that often represent stereotypical objects or situations. The conjunction of gates and frames helps provide meaning to sequence, flow, and the objects we use to convey information as they change shape – like a sentence, paragraph, section, etc.

Inference – The process of drawing conclusions from premises, observations, or data. In machine learning, inference often refers to using a trained model to make predictions on new, unseen data.

Tuple/Triplet – A tuple is an ordered, immutable collection of elements. A triplet is specifically a tuple containing three elements.

Generative – In machine learning, generative models learn to generate new data that resembles the training data. These models capture the joint probability distribution of the input data and labels.

Deterministic – A process or system is deterministic if its subsequent state is entirely determined by its prior state and inputs, with no random elements involved. In other words, given the same inputs and initial conditions, a deterministic system will always produce the same output.

Discriminant – A discriminant function distinguishes between different classes in machine learning, particularly in classification tasks where directionality is unclear. It takes the input features and outputs a score for each possible class, typically used to make the final classification decision.

Concept Definitions

Learning – In the context of AI and cognitive science, learning is the process of acquiring new knowledge, behaviors, skills, or preferences through experience, study, or instruction. It involves improving performance on a specific task with experience.

Machine Learning – A subset of artificial intelligence that focuses on developing algorithms and statistical models that enable computer systems to improve their performance on a specific task through experience without being explicitly programmed.

Deep Learning – A subset of machine learning based on artificial neural networks with multiple layers (hence "deep"). These models can learn hierarchical representations of data, often leading to state-of-the-art performance in tasks such as image and speech recognition, natural language processing, and more.

Reinforcement Learning – A type of machine learning where an agent learns to make decisions by taking actions in an environment to maximize some notion of cumulative reward. It's based on the idea of learning through interaction with the environment and receiving feedback in the form of rewards or penalties.

Neural Net Models – Computational models inspired by biological neural networks. They consist of interconnected nodes (neurons) organized in layers, capable of learning complex patterns in data. Neural nets are the foundation of deep learning.

Graph Models – Mathematical structures used to model relations between objects. In machine learning, graph models can represent complex relationships in data and are used in various applications, including network analysis, recommender systems, and knowledge representation.

Knowledge Graphs – A type of graph model that represents a collection of interlinked descriptions of entities – objects, events, or concepts. They are used to store interconnected information in a graph structure and are particularly useful for integrating information gathered from multiple sources.

Training – The process of teaching a machine learning model to make predictions or decisions. It involves exposing the model to a dataset and adjusting its parameters to minimize the difference between its predictions and the actual outcomes.

Fine-tuning – The process of taking a pre-trained model and further training it on a specific task or dataset. This allows the model to adapt its learned features to a new, often more specialized task, typically with less data and computational resources than training from scratch.

Abstractions – In computer science and AI, abstractions are simplified representations of complex systems or concepts. They allow us to focus on essential features while hiding unnecessary details. In machine learning, abstractions often refer to the hierarchical representations learned by deep neural networks, where each layer captures increasingly abstract features of the input data.

Nuanced Definition of Artificial Intelligence

Artificial Intelligence (AI):

AI is the field of computer science dedicated to creating systems capable of performing tasks that typically require human intelligence. These tasks include reasoning, problem-solving, learning, and understanding natural language.

Current "Emergent" Form:

Today's prominent AI systems, like large language models, are based on processing vast amounts of data to recognize patterns and generate human-like responses. These models, trained on diverse text from the internet and other sources, can engage in conversations, answer questions, and even create content across various domains. They excel at language-related tasks but operate by predicting likely sequences of words rather than genuinely understanding meaning in a human sense.

Key features:

1. Data-driven: Performance improves with more high-quality training data.
2. Pattern recognition: Excels at finding and replicating patterns in data.
3. Generalization: Can apply learned patterns to new, unseen situations.
4. Scalability: Capabilities often increase with model size and computational power.

Previous "Organic" Form: Earlier approaches to AI often aimed to mimic biological intelligence and cognitive processes more closely. These systems typically combined rule-based logic, knowledge representation, and machine learning in ways inspired by theories of human cognition.

Key features:

1. Symbolic reasoning: Uses logical rules and representations to solve problems.
2. Knowledge-based: Relies on structured information about the world.
3. Modular: Often composed of specialized components for different cognitive tasks.
4. Explainable: The decision-making process can often be traced and understood.

Both forms of AI continue to be developed and have their strengths. Current LLM-type AI excels at processing and generating human-like language, while more traditional "organic" approaches often perform better in tasks requiring explicit reasoning or domain-specific expertise. The future of AI may involve hybrid methods that combine the strengths of both paradigms.



Want to learn more about AI?

Let us be your guide.

At The Bridge, we provide the strategic partnership to help you think about data and AI differently - moving from uncertainty to opportunity by building solutions that meet you where you are and take you where you want to be.

CONTACT US

<https://thebridge.com/>

 DALLAS, TX

 DENVER, USA

 QUITO, EC

