



AI FUNDAMENTALS

*A Human-Centered Introduction
to Machine Intelligence*

by

Nagul Meera Mahankali

AI Enthusiast

Table of Contents

1. What Is AI?	3
2. What Is an AI Model?	3
3. Types of AI Models	4
4. Generative AI	5
5. Large Language Models (LLMs)	5
6. The Human Brain Analogy	6
7. What Is a Prompt?	7
8. User Prompt vs System Prompt	8
9. Function Calling / Tool Use	9
10. RAG: Retrieval-Augmented Generation	10
11. Understanding Embeddings	11
12. Embeddings as Mental Coordinates	12
13. Embeddings in Action	13
14. Popular AI Models	14
15. AI Agents & Agentic AI	15
16. How It All Works Together	16





What Is AI?

Artificial Intelligence (AI) is the ability of machines to perform tasks that typically require human intelligence—like understanding language, recognizing images, making decisions, or solving problems.

Think of AI as a digital brain that can learn, adapt, and assist—just like a human, but powered by data and algorithms.



What Is an AI Model?

An AI model is a trained system that learns patterns from data and uses that knowledge to make predictions or generate outputs.

Imagine a student who studies thousands of examples. Once trained, they can answer questions or solve problems—even ones they've never seen before.



Types of AI Models

Type	Description
Machine Learning	Learns from data to make predictions or decisions
Deep Learning	Uses neural networks for complex tasks like image recognition
Generative Models	Creates new content—text, images, music, or code
Reinforcement Learning	Learns by trial and error, like training a pet with rewards
Rule-Based Systems	Follows predefined logic (if-then rules)

Generative AI

Generative AI creates new content—text, images, audio, or code—based on patterns it learned during training.

It's like an artist who has studied thousands of paintings and can now create original artwork in any style.

Popular Examples: ChatGPT, DALL·E, Midjourney, Sora





Large Language Models (LLMs)

An LLM is a type of generative AI trained on massive text datasets to understand and generate human-like language.

It's like a super-educated adult who has read millions of books and can answer questions, write essays, or hold conversations.

The Human Brain Analogy

Imagine a human growing from a baby to an adult:

-  **Baby Brain** = Untrained model
-  **Learning Phase** = Training on data (school, society, nature)
-  **Mature Brain** = Pre-trained model
-  **Adult Responding** = Inference: generating answers based on learned knowledge

Just like a grown-up who can answer questions based on life experience, a generative AI model responds using what it learned during training.

What Is a Prompt?

A **prompt** is the instruction, question, or input you give to an AI model. It's how you communicate with AI to get the response you want.

Think of a prompt as asking a question to a knowledgeable friend. The clearer and more specific your question, the better the answer you'll receive.

The Human Analogy

Imagine you're talking to an expert:

Vague prompt: "Tell me about food."

Response might be too broad or generic.

Clear prompt: "What are three healthy breakfast options for someone trying to lose weight?"

Response will be specific, actionable, and useful.

Why Prompts Matter

Just like asking better questions gets better answers from people, well-crafted prompts get better results from AI. The prompt is your way of guiding the AI's knowledge toward exactly what you need.



User Prompt vs System Prompt

User Prompt

This is **your direct question or instruction** to the AI. It's what you type or say when interacting with the model.

Example: "Write a summary of the French Revolution in 3 bullet points."

Think of it as what you ask your friend directly during a conversation.

System Prompt

This is a **background instruction** that sets the AI's behavior, role, or context—usually set by developers or administrators before the conversation begins.

Example: "You are a helpful history teacher. Always explain concepts in simple terms and use examples."

Think of it as telling your friend, "When we talk today, act like a patient teacher explaining things to a beginner."

Key Difference

User Prompt: Changes with every question you ask

System Prompt: Stays constant, setting overall behavior

Together, they guide the AI to give you the most relevant and appropriate responses!



Function Calling / Tool Use

Function calling allows an LLM to interact with external tools, APIs, or databases to retrieve real-time information or perform actions—like checking weather, querying databases, or sending emails.

The Analogy

Imagine the adult brain faces a new question—something they haven't learned:

- The adult turns to young kids (external tools/APIs) who are up-to-date and specialized
- The adult asks: "Hey kid, what's the weather in Tokyo?"
- The kid checks and replies: "It's 22°C and sunny"
- The adult passes this answer along

The LLM delegates tasks to external tools when it doesn't have the answer itself.

RAG: Retrieval-Augmented Generation

RAG is a technique that enhances an LLM's responses by retrieving relevant information from external knowledge sources (documents, databases, websites) at the time of the query.

The Analogy

Imagine the adult brain needs to answer a complex, current question:

- The adult walks into a library (vector database/document store)
- Searches for relevant books or articles
- Reads the most useful pages
- Synthesizes the info and gives a grounded answer





The LLM retrieves relevant documents at runtime to enhance its response with fresh, factual information.

Understanding Embeddings

Embeddings are a way to represent words, images, or concepts as **vectors**—lists of numbers that capture their meaning, relationships, and context.

Memory & Association

Your brain stores concepts by meaning, not alphabetically. When you hear "apple", you instantly connect it to:

-  Fruit
-  Green or red
-  Juice
-  Apple Inc.

These associations form a mental cluster—your brain places "apple" near other related concepts. That's exactly what embeddings do in AI.

Embeddings as Mental Coordinates

Think of each concept as having coordinates in your brain:

- **"Apple"** might be at (0.2, 0.8, 0.1, ...)
- **"Banana"** might be nearby at (0.3, 0.7, 0.2, ...)
- **"Laptop"** would be far away at (0.9, 0.1, 0.6, ...)

The closer the coordinates, the more related the concepts. This is how AI models "understand" relationships—by measuring distance between embeddings.

Human Thought vs AI Embedding

Human Brain	AI Embedding System
Stores ideas by meaning	Stores vectors in semantic space
Associates concepts through memory	Measures similarity via vector distance
Learns from experience	Learns from training data
Uses intuition to group ideas	Uses math to cluster related vectors

🌟 Everyday Example: Embeddings in Action

Let's say you ask an AI:

"What's a good fruit for breakfast?"

How the AI Uses Embeddings:

The model uses embeddings to find words close to "fruit" and "breakfast" in its vector space—like "banana," "berries," or "yogurt."

It doesn't just match keywords—it understands the context through proximity in embedding space.

The AI knows that "banana" is closer to both "fruit" and "breakfast" than "steak" would be!

Why This Matters:

Embeddings allow AI to:

- Find similar items (recommendation systems)
- Search by meaning, not just keywords
- Understand context and relationships
- Group related concepts automatically

Popular AI Models

Model	Company	Best For
GPT-4o / o1	OpenAI	Advanced reasoning, multimodal tasks
Claude 4	Anthropic	Long context, analysis, coding
Gemini 2.0	Google	Multimodal AI, search integration
Llama 3	Meta	Open-source, customizable
DALL-E 3	OpenAI	Image generation from text
Midjourney	Midjourney	Artistic image creation

What Is an AI Agent?

An **AI Agent** is an autonomous system that can perceive its environment, make decisions, and take actions to achieve specific goals—often without constant human intervention.

Think of an AI Agent as a personal assistant who doesn't just answer questions, but can actually complete tasks for you—like booking appointments, managing emails, or ordering groceries.

The Human Analogy

Imagine hiring an assistant who:

- Understands your goals ("I need to organize a team meeting")
- Plans the steps (check calendars, find a time, book a room)
- Takes action (sends invites, reserves the room)
- Adapts if needed (reschedules if conflicts arise)

An AI Agent works the same way—it perceives, plans, acts, and learns from results.

What Is Agentic AI?

Agentic AI refers to AI systems designed with agency—the ability to act independently, make decisions, and pursue goals autonomously over extended periods.

While a regular chatbot waits for your questions, Agentic AI proactively works toward objectives, adapting its approach as circumstances change.



How It All Works Together

```
[User Question]
    ↓
[LLM (Adult Brain)]
    ↓
[Can I answer this?]
    ↓      ↘
  [Yes]    [No]
    ↓      ↓
[Generate] [Tool/RAG]
    ↓
[Final Answer]
```



Key Takeaways

- AI mimics human intelligence through learning and pattern recognition
- AI models are like students learning from thousands of examples
- LLMs are super-educated adults who've "read" millions of books
- Function calling lets AI use external tools for real-time data
- RAG lets AI retrieve current information from knowledge bases
- Embeddings help AI understand relationships between concepts
- AI Agents can autonomously complete complex, multi-step tasks
- Agentic AI takes initiative and works toward goals independently



About This Book

AI Fundamentals offers a refreshingly accessible introduction to artificial intelligence through intuitive human analogies. Whether you're a curious beginner or a professional looking to understand AI concepts, this book bridges the gap between complex technical concepts and everyday understanding.

Instead of overwhelming jargon, you'll discover how AI works through relatable metaphors—like comparing an LLM to a well-educated adult who has read millions of books, or understanding embeddings as mental coordinates in your brain.

What You'll Learn:

- ✓ Core AI concepts explained through human-centered analogies
- ✓ How Large Language Models (LLMs) actually work
- ✓ Function calling and RAG (Retrieval-Augmented Generation)
- ✓ Understanding embeddings and vector spaces
- ✓ Overview of popular AI models in the market
- ✓ Practical insights into generative AI and deep learning
- ✓ AI Agents and the future of autonomous AI systems

"Making AI accessible, one analogy at a time."

— Nagul Meera Mahankali