AI FOR EXECUTIVES

Understanding AI, LLMs, and the Future of Work

A Plain-Language Guide for Business Leaders • December 2025

Table of Contents

Why AI Matters Right Now

$184B

Global AI market size in 2024

75%

of companies using AI in operations

40%

productivity gains reported

→ AI isn't coming—it's already here. Your competitors are using it.

→ Understanding AI basics is now a core executive skill.

→ Leaders who understand AI make better strategic decisions.

→ You don't need to code—you need to understand concepts.

> *"The goal of this deck: Give you enough understanding to have informed conversations, ask the right questions, and make smart AI decisions."*

PART 1

What is AI?

The foundation you need to understand everything else

AI Explained Simply

THE SIMPLE DEFINITION

Artificial Intelligence is software that can learn patterns from data and make decisions or predictions—without being explicitly programmed for each specific task.

TRADITIONAL SOFTWARE

Follows exact rules you write. "If customer orders > $100, apply 10% discount." It only does exactly what you tell it.

AI SOFTWARE

Learns from examples. Show it 1 million customer interactions and it figures out patterns. It can handle situations it wasn't specifically programmed for.

Machine Learning
Finds patterns in data

Computer Vision
Understands images

LLMs (Language)
This deck's focus!

Generative AI
Creates new content

What is an LLM?

LLM = LARGE LANGUAGE MODEL

An LLM is a type of AI that's trained on massive amounts of text (books, websites, articles) to understand and generate human language. Think of it as autocomplete on steroids.

HOW IT WORKS (SIMPLIFIED)

1. Trained on trillions of words from the internet, books, and other text
2. Learns patterns: what words typically follow other words
3. When you ask a question, it predicts the most likely response
4. Generates one word at a time, each based on what came before

MAJOR LLMs YOU'LL HEAR ABOUT

GPT-4/5	OpenAI (ChatGPT)
Claude	Anthropic
Gemini	Google
Llama	Meta (Open Source)

The Autocomplete Analogy

The best way to understand LLMs: They're like your phone's autocomplete, but trained on almost all human knowledge.

📱 PHONE AUTOCOMPLETE

You type:

"I'll be there in..."

It suggests:

"5 minutes" | "10 minutes"

Trained on your past messages. Predicts a few words.

🧠 LLM (CHATGPT, CLAUDE)

You type:

"Explain quantum computing..."

It generates:

A complete, coherent explanation...

Trained on the entire internet. Predicts entire responses.

Key insight: LLMs don't "know" things the way humans do. They predict what text should come next based on patterns they learned during training. This is why they can be confidently wrong (we'll cover this in "hallucinations").

PART 2

Prompting

How to talk to AI effectively

What is Prompting?

SIMPLE DEFINITION

A **prompt** is the text you give to an AI. It's your instruction, question, or request. **Prompting** is the skill of writing prompts that get the best results.

THINK OF IT LIKE DELEGATION

✕ BAD DELEGATION

"Do the report."

✓ GOOD DELEGATION

"Create a 2-page Q3 sales summary for the board, highlighting wins and risks."

✕ BAD PROMPT

"Write about sales."

✓ GOOD PROMPT

"Write a 500-word executive summary of Q3 sales, including top 3 wins and top 3 risks. Use bullet points."

The clearer your prompt, the better the AI's output. Garbage in = garbage out.

Prompting Best Practices

1

Be Specific

Include details: length, format, audience, tone. "Write a 3-paragraph email to a skeptical CFO..."

2

Give Context

Tell the AI who you are, what you're trying to accomplish, and why. Background helps it calibrate.

3

Provide Examples

"Here's an example of the style I want..." helps AI match your expectations.

4

Assign a Role

"Act as a financial analyst..." or "You are an expert in..." shapes the response style.

5

Iterate & Refine

Don't expect perfection on try #1. Ask for revisions: "Make it shorter" or "Add more data."

Pro Tip

Ask AI to "think step by step" for complex problems. This improves reasoning accuracy.

PART 3

Hallucinations

When AI confidently makes things up

What are AI Hallucinations?

DEFINITION

A **hallucination** is when AI generates information that sounds confident and plausible—but is completely made up. It's not lying; it's pattern-matching gone wrong.

REAL EXAMPLES FROM THE NEWS

⚖️ LEGAL DISASTER

A lawyer used ChatGPT for legal research. It invented 6 fake court cases with realistic names and citations. He submitted them to court and was sanctioned.

🏥 MEDICAL RISK

Researchers found AI could be manipulated to give dangerous medical advice, like claiming sunscreen causes cancer. AI can't verify its own facts.

🔍 **WHY IT HAPPENS: LLMs predict "what sounds right" based on patterns—not what IS right. They have no concept of truth or ability to verify facts.**

Bottom line: Always verify AI outputs for factual claims, especially in high-stakes domains like legal, medical, and financial.

Hallucination Rates by Model (2025)

Lower is better. Data from Vectara's Hallucination Leaderboard (Dec 2025).

Model	Company	Hallucination Rate	Rating
Gemini 2.0 Flash	Google	0.7%	🏆 Best
GPT-4o	OpenAI	1.5%	Excellent
GPT-4 Turbo	OpenAI	1.7%	Excellent
GPT-3.5 Turbo	OpenAI	1.9%	Very Good
Claude Sonnet 4	Anthropic	4.4%	Good
Open Source (Avg)	Various	15-30%	Caution

📈 **GOOD NEWS**

Rates dropped 96% since 2021 (21.8% → 0.7%)

⚠️ **HIGH-RISK DOMAINS**

Legal: 6.4% error rate, Medical: Higher still

Source: Vectara Hallucination Leaderboard, AIAboutAI Research (Dec 2025)

How to Reduce Hallucinations

🛡️ YOUR STRATEGIES

1. Always Verify Facts

Cross-check any specific claims, dates, names, or statistics with trusted sources.

2. Ask for Sources

Request "cite your sources" — then verify those sources actually exist and say what AI claims.

3. Use AI + Web Search

Tools that combine AI with real-time search (like Perplexity) reduce hallucinations significantly.

🏗️ TECHNICAL SOLUTIONS

RAG (Retrieval-Augmented Generation)

Connects AI to your actual documents/data before answering. **Reduces hallucinations by 71%.**

Grounding

Give AI access to specific, verified data sources (databases, knowledge bases, company docs).

Fine-Tuning

Train AI on your specific domain data to improve accuracy in your context.

💡 **Key Insight: The question isn't "will AI hallucinate?" but "how do we build verification into our workflows?"**

PART 4

Agents, Apps & MCP

Three ways to extend AI beyond chat

The Problem: AI is Isolated

By default, AI chatbots are "brains in a jar." They can think, but they can't **do** anything in the real world—can't access your files, send emails, update your CRM, or pull live data.

THREE SOLUTIONS HAVE EMERGED

AI Agents

AI that can take actions autonomously

Add-On Apps

Plugins that add specific features

MCP

Universal connector standard

Think of it like your phone: The phone itself is powerful, but apps and integrations are what make it useful. AI is the same—it needs connections to your tools and data.

What are AI Agents?

DEFINITION

An **AI Agent** is an AI system that can plan multi-step tasks, use tools, and take actions autonomously—not just answer questions, but actually DO things.

💬 REGULAR AI CHAT

You: "Book me a flight to NYC"

AI: "I can't book flights, but here are some tips for finding cheap flights..."

Can only talk about it

🤖 AI AGENT

You: "Book me a flight to NYC"

Agent: "Found 3 options. I've booked the 8am Delta for $289. Confirmation sent to your email."

Actually takes action

EXAMPLES OF AI AGENTS IN THE WILD

Coding Agents
Write, test, and deploy code

Research Agents
Search, read, synthesize info

Sales Agents
Update CRM, draft emails

Customer Service
Resolve issues end-to-end

Add-On Apps (Plugins)

DEFINITION

Add-on apps (or plugins) are pre-built integrations that give AI access to specific tools or services. Think of them like apps on your phone—each one adds a new capability.

COMMON EXAMPLES

- Code Interpreter
- Web Browsing
- DALL-E (Images)
- PDF Reader
- Email Integration
- Calendar Access

✓ PROS

Easy to use—just enable and go. Pre-built, maintained by vendors. Quick to add specific features.

✗ CONS

Limited selection. Each platform has different plugins. Can't easily customize or combine them.

The Analogy: Add-on apps are like iPhone apps—convenient, but you're limited to what's in the App Store, and apps from one platform don't work on another.

MCP: The Universal Connector

MCP = MODEL CONTEXT PROTOCOL

MCP is an open standard that lets **any AI** connect to **any tool or data source** using one universal protocol. It's like USB-C for AI.

🔌 THE USB-C ANALOGY

Before USB-C

Every device had different cables. Phone charger ≠ laptop charger ≠ camera charger.

After USB-C

One cable works with everything. Standardization = simplicity.

MCP for AI

One protocol connects any AI to any tool. Build once, works everywhere.

🏢 WHO'S ADOPTING IT?

- Anthropic (creator)
- OpenAI
- Google DeepMind
- Microsoft
- AWS

🔗 AVAILABLE CONNECTORS

- Google Drive
- Slack
- GitHub
- Salesforce
- Databases
- +1000s more

Why it matters: MCP is becoming the industry standard. In Dec 2025, it was donated to the Linux Foundation with backing from all major AI companies. This means your AI investments are portable.

Agents vs Apps vs MCP: Comparison

	🤖 AI Agents	🧩 Add-On Apps	🔌 MCP
What it is	AI that takes actions autonomously	Pre-built plugins for specific tools	Universal protocol for any connection
Flexibility	High – can chain multiple steps	Low – one tool at a time	High – unlimited connections
Setup Effort	Medium – needs configuration	Low – just enable	Medium – one-time setup
Portability	Varies – often vendor-locked	Low – platform-specific	High – open standard
Best For	Complex workflows, automation	Quick, simple additions	Enterprise integration, scalability

💡 **Bottom Line: These aren't mutually exclusive. Agents often USE MCP to connect to tools. Think of MCP as the plumbing, agents as the workers, and apps as quick shortcuts.**

Key Takeaways

1

LLMs are pattern matchers

They predict text based on training data. They don't "know" things or verify facts.

2

Good prompts = good results

Be specific, provide context, give examples, and iterate. Treat AI like a smart intern.

3

Hallucinations are real risks

Rates are improving (sub-2% for top models) but always verify facts in critical domains.

4

MCP is the future standard

Universal connector adopted by all major AI companies. Your integrations become portable.

5

Agents are the next wave

AI that doesn't just answer but actually does—coding, research, scheduling, sales tasks.

You don't need to be technical

Understanding these concepts lets you ask the right questions and make smart AI investments.