

# Welcome!

## Generative AI for Managers & Leaders

**Prompt:**

- *“A painting in the style of Van Gogh of a robot orchestra conductor directing a large group of musicians --ar 2:1.”*



# Welcome & Introductions

## 00 First things first:

1. Introductions, housekeeping, expectations
2. Accessing the class's white board space
3. Accessing the course materials

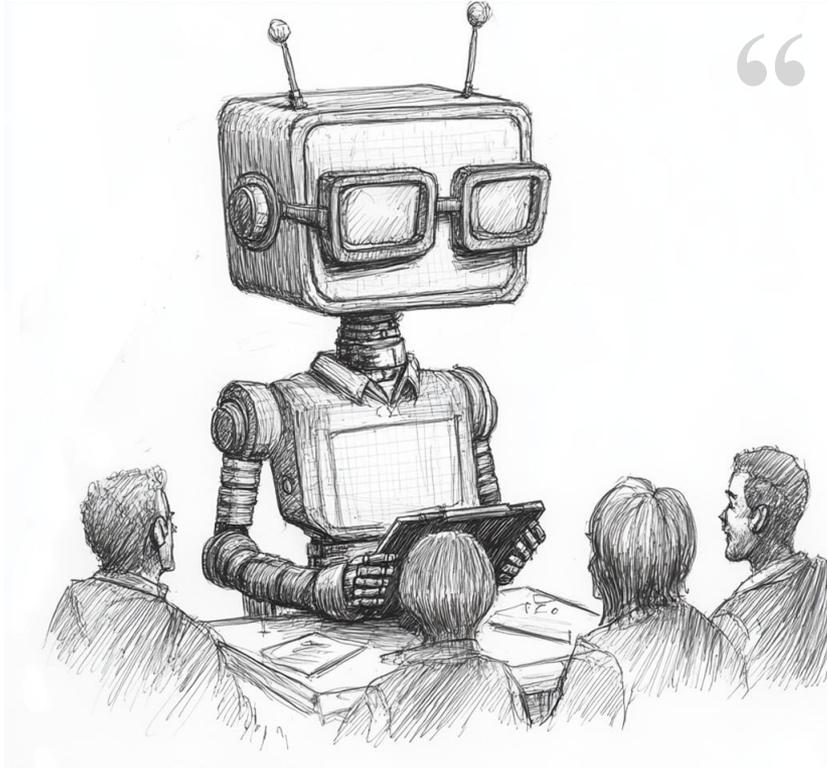
# Setting Expectations: What This Time Is About

- Creating a space for discussions and flexibility
- Establishing a working literacy on AI management
- Identifying decisions and management priorities you think are important related to AI
- Confronting hard questions related to managing AI.
- Surfacing the issues that you'll have to work on, and how to prioritize them.
- I'm a resource for you, and a facilitator. There will be things you have to figure out!

# Setting Expectations: What This Time Is **NOT** About

- A prescriptive, one-size-fits all approach
- An exhaustive reference of highly specific or detailed use cases
- Extended technical discussions, deep focus on any specific tool or technology
- Silver bullets, panaceas, perfect solutions, big overnight transformations

# Setting Expectations: Final thoughts



“ Powerful and sustained change requires constant communication, not only throughout the rollout but after the major elements of the plan are in place. The more kinds of communication employed, the more effective they are. ”

– [DeAnne Aguirre](#)

# Whiteboarding during this session - Miro



## Generative AI for Leaders & Managers

Using Miro - Quick Start...



Virtual Circle



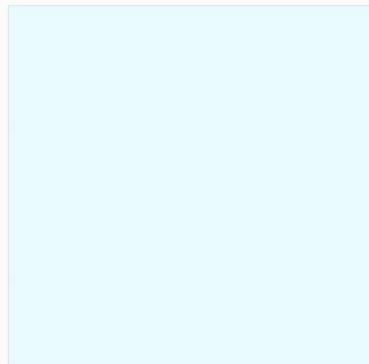
Facilitator



Resources



Parking Lot



### Something to think about

"I need to produce a report comparing software applications currently in use across some agencies to generate recommendations and recommend greater visibility on my website for all sites."

How do you think you'll address this task using generative AI? What are the pros and cons of using generative AI for this task? What are the risks of using generative AI for this task? What are the benefits of using generative AI for this task? What are the challenges of using generative AI for this task? What are the opportunities of using generative AI for this task? What are the risks of using generative AI for this task? What are the benefits of using generative AI for this task? What are the challenges of using generative AI for this task? What are the opportunities of using generative AI for this task?

1. Reflect on the task. (10 min)
2. Brainstorm ideas. (10 min)
3. Share ideas with the group. (10 min)
4. Discuss the ideas. (10 min)
5. Summarize the ideas. (10 min)

### Where do we want to take the workshop now?

Agenda

- 1. Introduction
- 2. What is Generative AI?
- 3. How to use Generative AI?
- 4. Case Studies
- 5. Q&A

There

Yes

No

### Let's capture some of your thoughts

#### CONSIDER: Where are we now in our organization?

What are the current challenges? What are the opportunities? What are the risks? What are the benefits? What are the challenges? What are the opportunities? What are the risks? What are the benefits? What are the challenges? What are the opportunities? What are the risks? What are the benefits?

### DISCUSSION: Busy work

#### Machine Activities

- Repetitive tasks
- Working with data
- Customer experience
- Customer support
- Sales

#### Human Activities

- Strategic planning
- Creative thinking
- Problem solving
- Decision making
- Relationship building
- Innovation

### IDEATION: Use Case Opportunities

What are the opportunities? What are the challenges? What are the risks? What are the benefits? What are the challenges? What are the opportunities? What are the risks? What are the benefits? What are the challenges? What are the opportunities? What are the risks? What are the benefits?

There

Yes

No

### DISCUSSION: What are the challenges to your AI vision?

One has difficulty with change. Regulatory requirements look hard. Skills are needed. What else?

Such is the hard. Something to do with money or risk. What else?

### DISCUSSION: Education by Experimentation

How do we experiment with experimentation?

1. What AI tools or capabilities do you need to experiment with?
2. How can we make an experiment where your group could gain the most learning and insight?

- How do we measure the success of our experiments?
- How do we measure the impact of our experiments?
- How do we measure the value of our experiments?
- How do we measure the ROI of our experiments?

### WHITE BOARD: Governance, policy, use guidance

### Capturing action items

What are the action items? What are the challenges? What are the risks? What are the benefits? What are the challenges? What are the opportunities? What are the risks? What are the benefits? What are the challenges? What are the opportunities? What are the risks? What are the benefits?

There

Yes

No

# Poll:

**How familiar are you with AI tools?**

*(Check all that apply)*

- A. Not much
- B. Used ChatGPT or other tools some
- C. Use AI regularly for my own work tasks
- D. Actively implementing some AI usage into workflows or projects

# What is AI?

- Computation that mimics **cognitive** functions.
- The science, math, and engineering of making **intelligent machines**.
- In computer science “AI agents” can perceive their environment and take actions to maximize its chance of success at some goal.
- **We are primarily interested in generative AI, but not exclusively.**



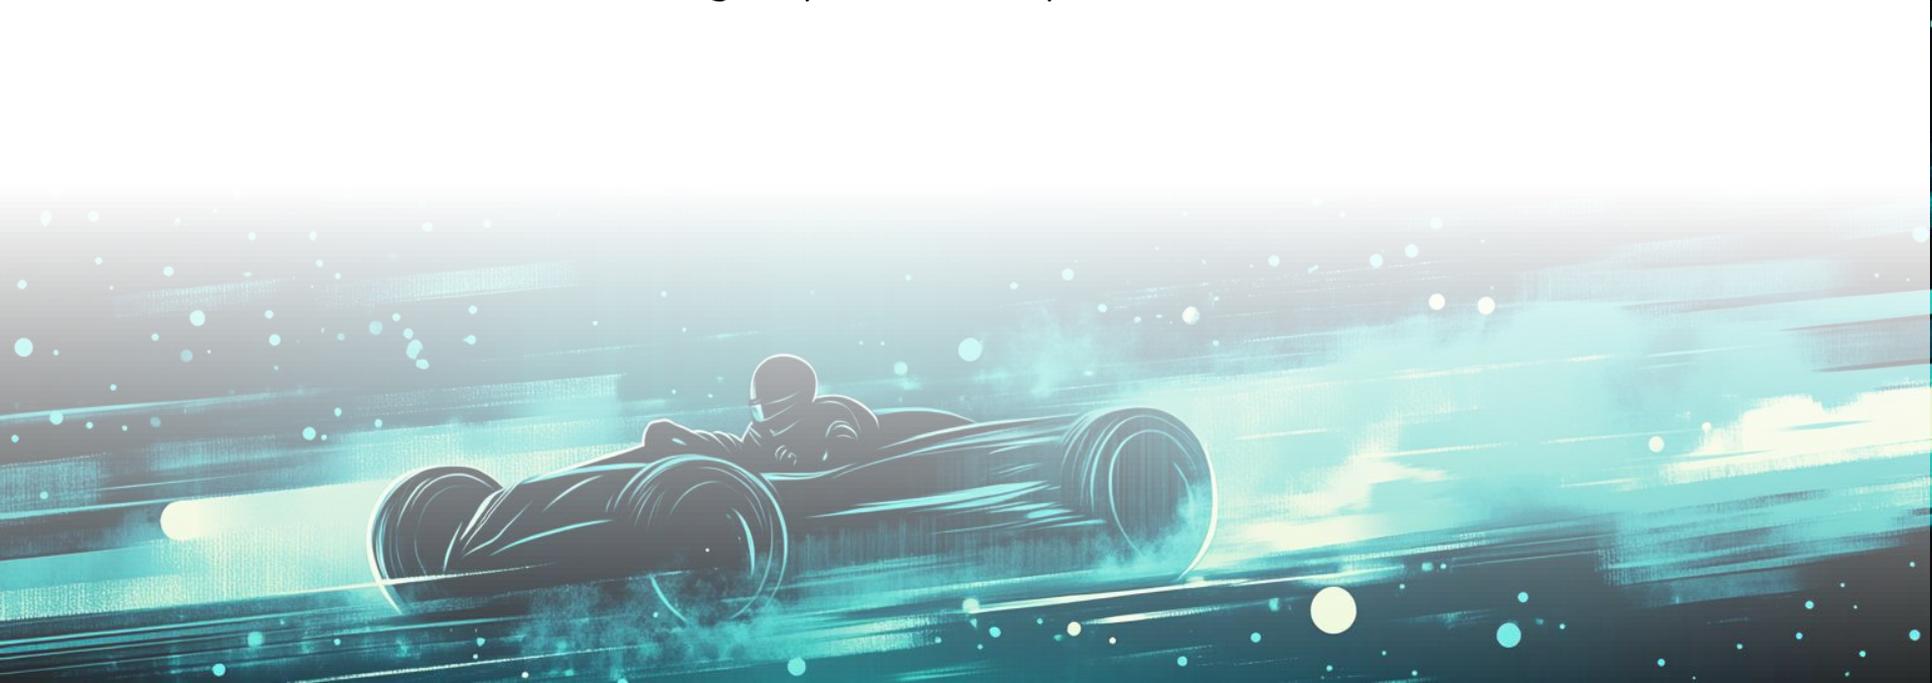
## Very simply:



- As transformative as the internet or mobile phones.
- It took a while for figure out the most powerful use cases.
- How many use cases were eventually made possible by the internet?

# Something to think about

We have to start with urgency, inevitability, and the needs of our stakeholders.



# Something to think about

**ASSIGNMENT**

“I need to produce a report comparing the features and price points for my competitors’ products, as well as any published case studies illustrating how their products have been used in large public projects. I then need to present recommendations for how we can present products in our portfolio in a way that positions them competitively for future bidding opportunities which are similar to those in the case studies.

The final document should be less than twenty pages and summarize my research, as well as present recommendations for marketing initiatives.



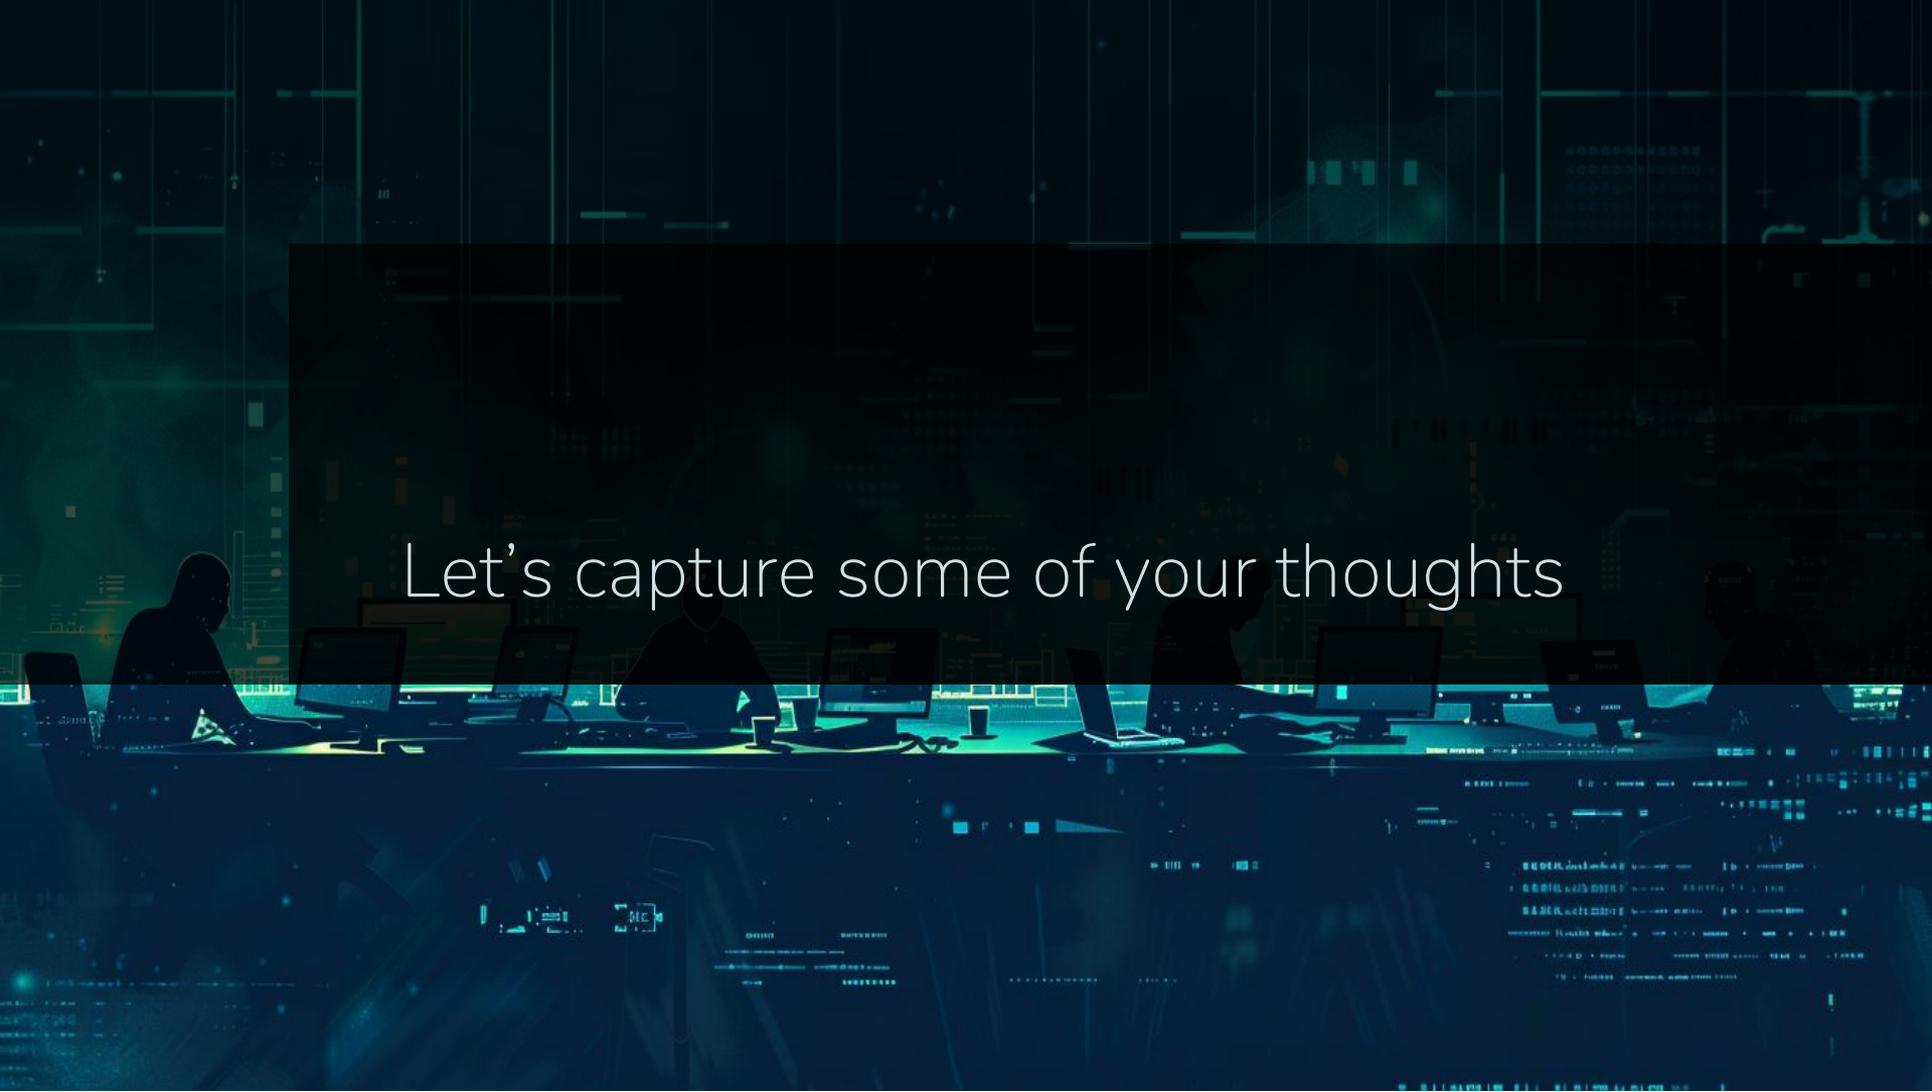
# Something to think about

## ASSIGNMENT

“I need to produce a report comparing the features and price points for my competitors’ products, as well as any published case studies illustrating how their products have been used in large public projects. I then need to present recommendations for how we can present products in our portfolio in a way that positions them competitively for future bidding opportunities which are similar to those in the case studies.

The final document should be less than twenty pages and summarize my research, as well as present recommendations for marketing initiatives.

- Before the Internet, in say 1994, how would I do this and how long would it take?
- Internet in full swing, say 2012, how would I do this and how long would it take?
- Generative AI, 2025, how would I do this and how long would it take?

The image depicts a dark, futuristic control room or server center. In the foreground, several silhouettes of people are seated at long desks, each with multiple computer monitors. The room is dimly lit, with light emanating from the screens and creating a blue-green glow. In the background, a large, dark rectangular screen is visible, which is the focus of the text overlay. The overall atmosphere is high-tech and mysterious.

Let's capture some of your thoughts

Let's capture some of your thoughts

## CONSIDER:

Where are we now in our organization?

GOALS

Do we have defined objectives?



TOOLS

What AI tools (if any) are being used in your area?



GOVERNANCE

Are there **AI policies or guidance** in place?



CHALLENGES

What **challenges or barriers** exist for expanding AI use?

*We'll use our whiteboard canvas to capture your thoughts.*

# What is AI?

- Computation that mimics **cognitive** functions.
- The science, math, and engineering of making **intelligent machines**.
- In computer science “AI agents” can perceive their environment and take actions to maximize its chance of success at some goal.
- **We are primarily interested in generative AI, but not exclusively.**



# Our assumptions for leaders

*If you agree, it's a basis for our time together. If not, some of our positioning won't work well.*

- **As a technology, AI is (at least potentially) transformative.**
- **Defining strategy, making decisions, and providing guidance to our teams are management responsibilities.**
- **Business case first, always. Then tools. Effective adoption depends on defining the need, then applying technology solution to goal.**
- **You are serious about being adaptive to change – or making change – when needed.**



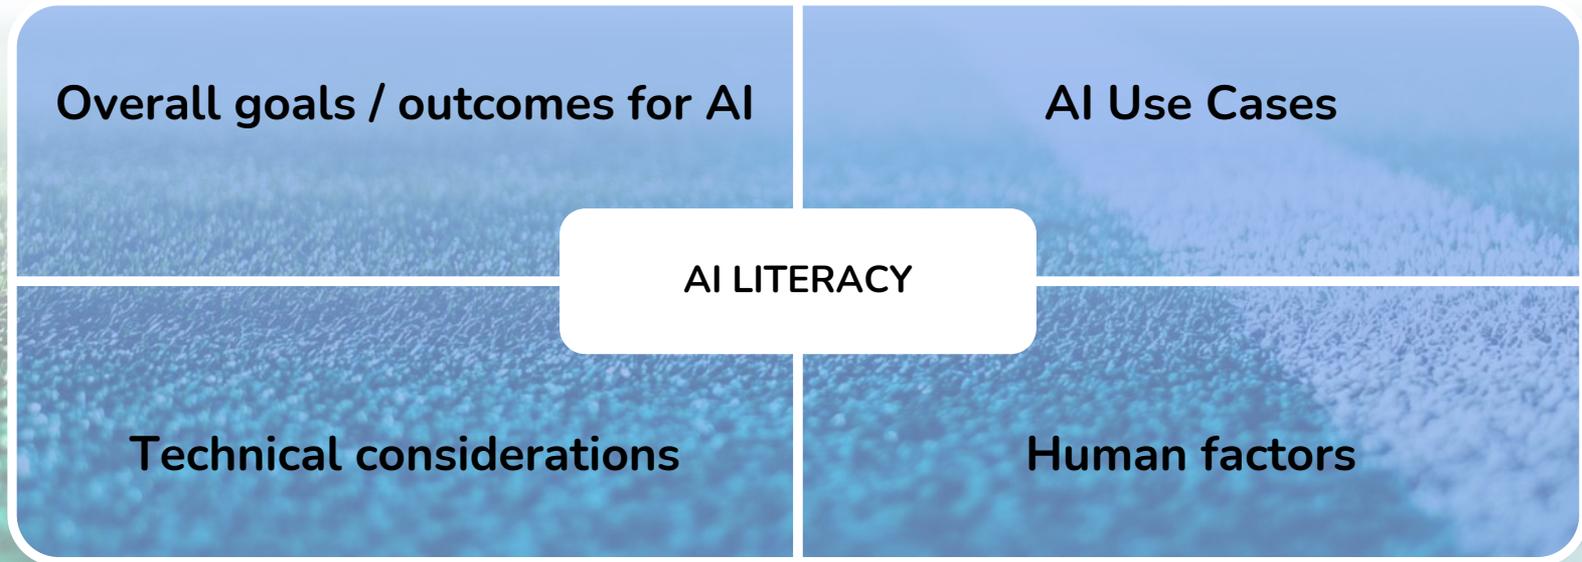
# Organizing our thoughts

*Important to align on vocabulary*

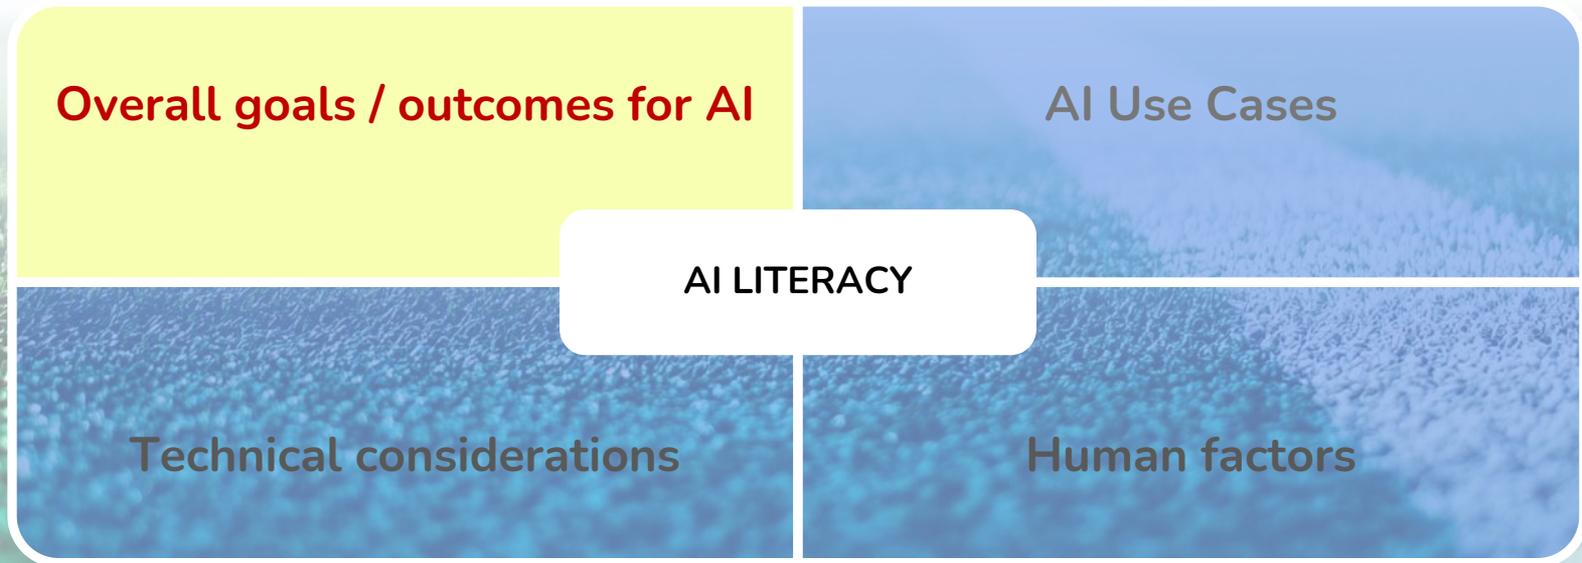
- AI
- Generative AI
- Business Case
- Use Case
- RAG
- LLM
- Internal LLM
- Data
- Training Data
- Model
- People
- Tools
- Agent

# Organizing our thoughts

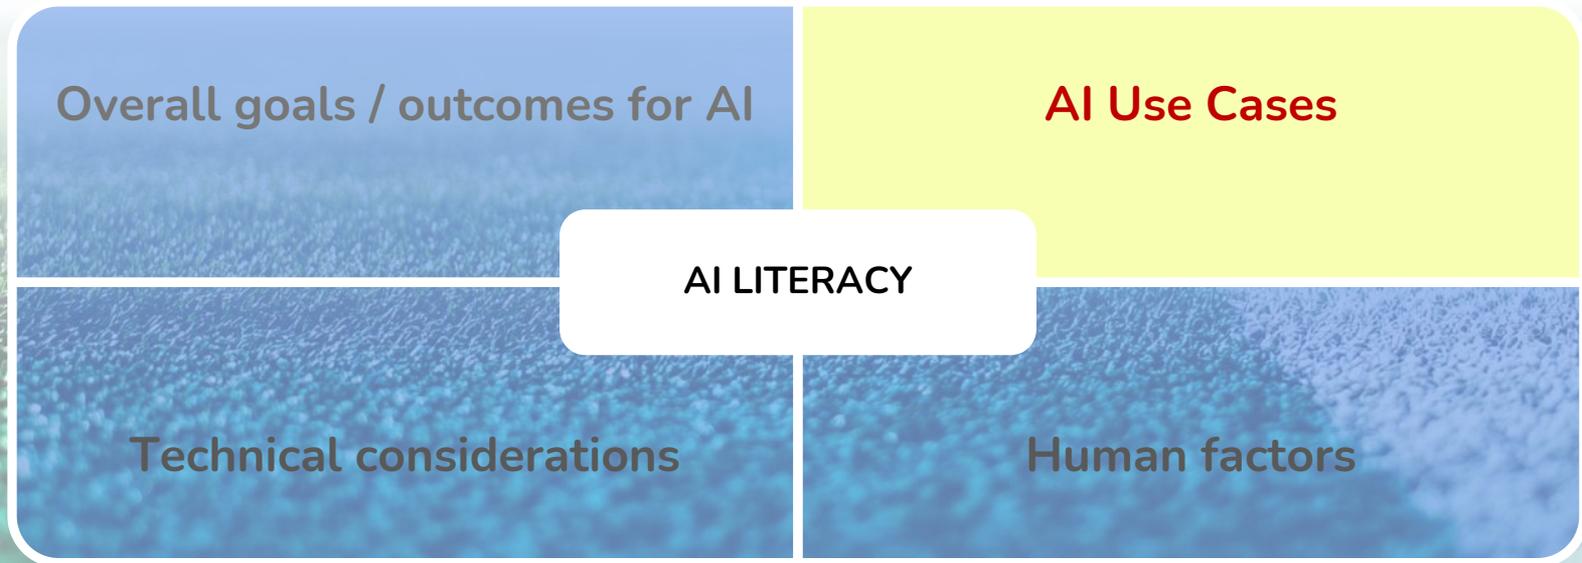
*We'll organize our discussion by keeping an eye on these dimensions.*



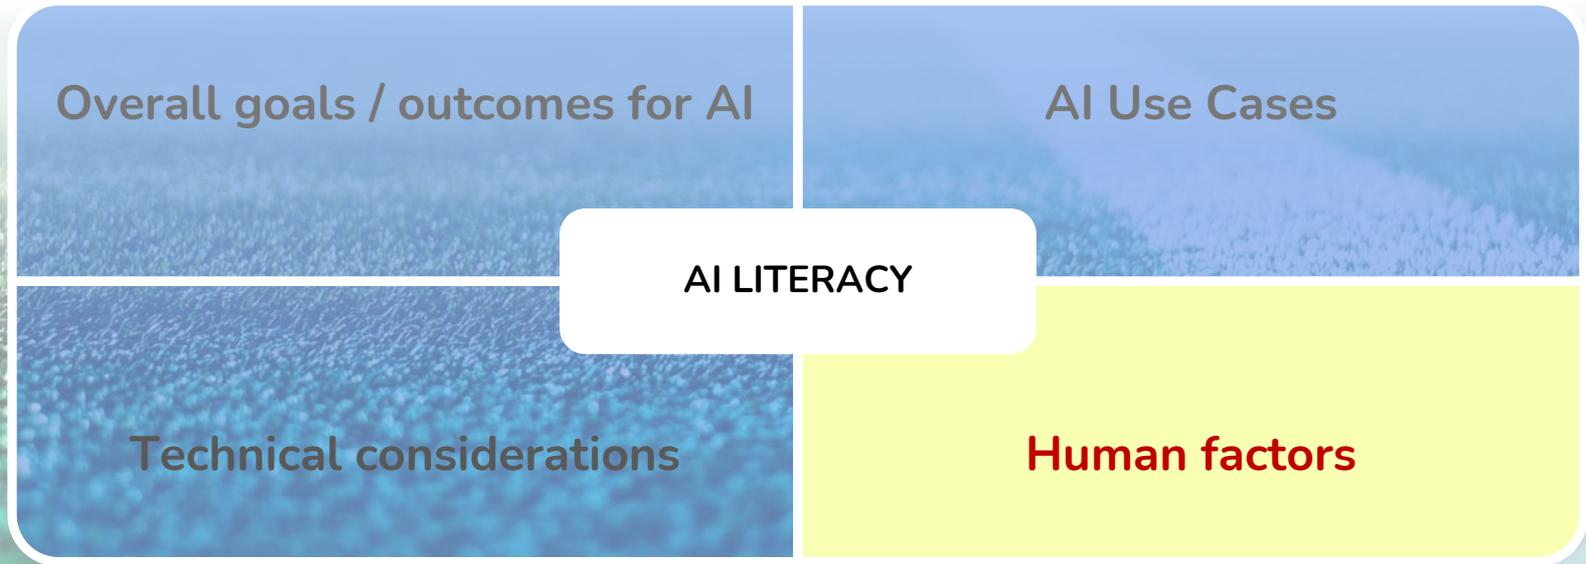
# Organizing our thoughts



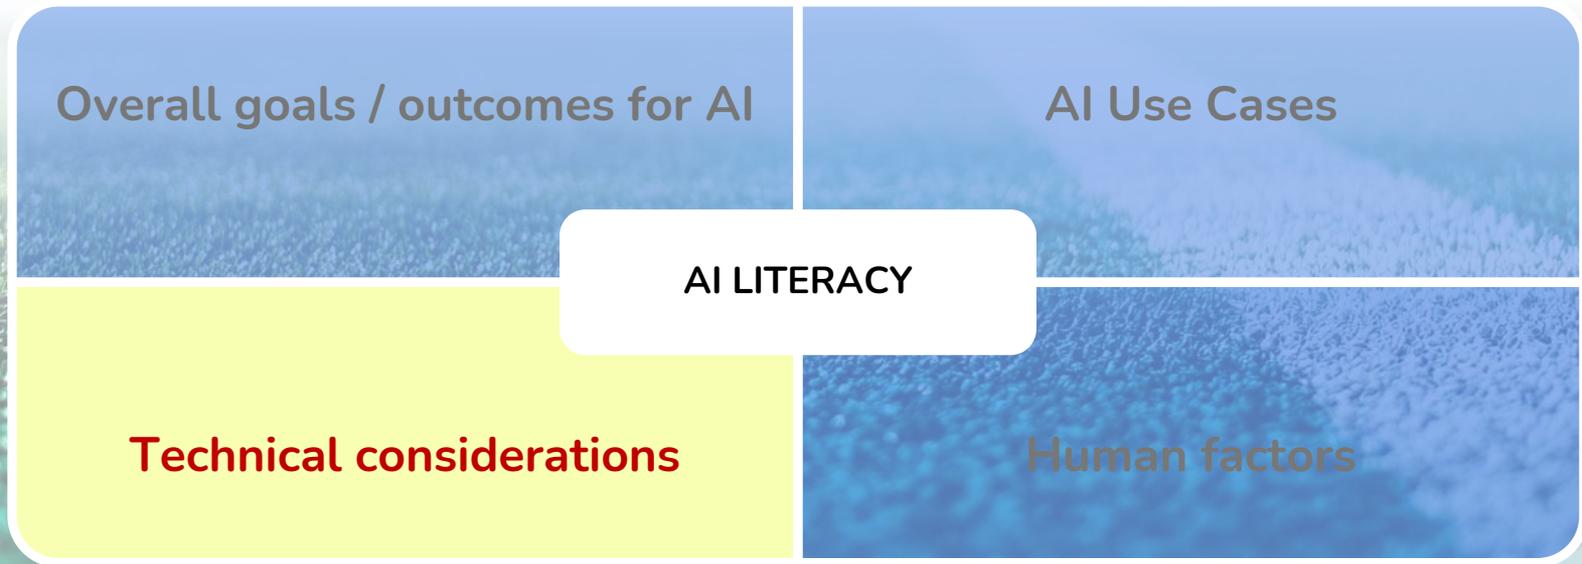
# Organizing our thoughts



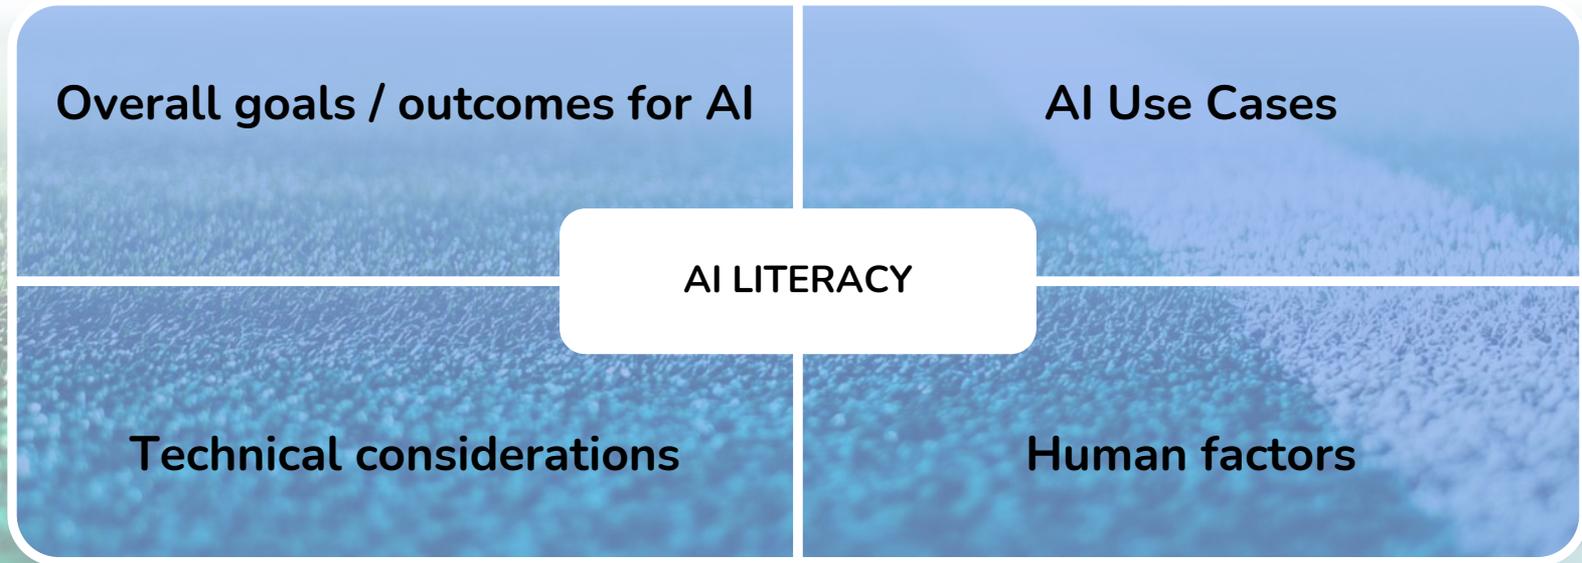
# Organizing our thoughts



# Organizing our thoughts

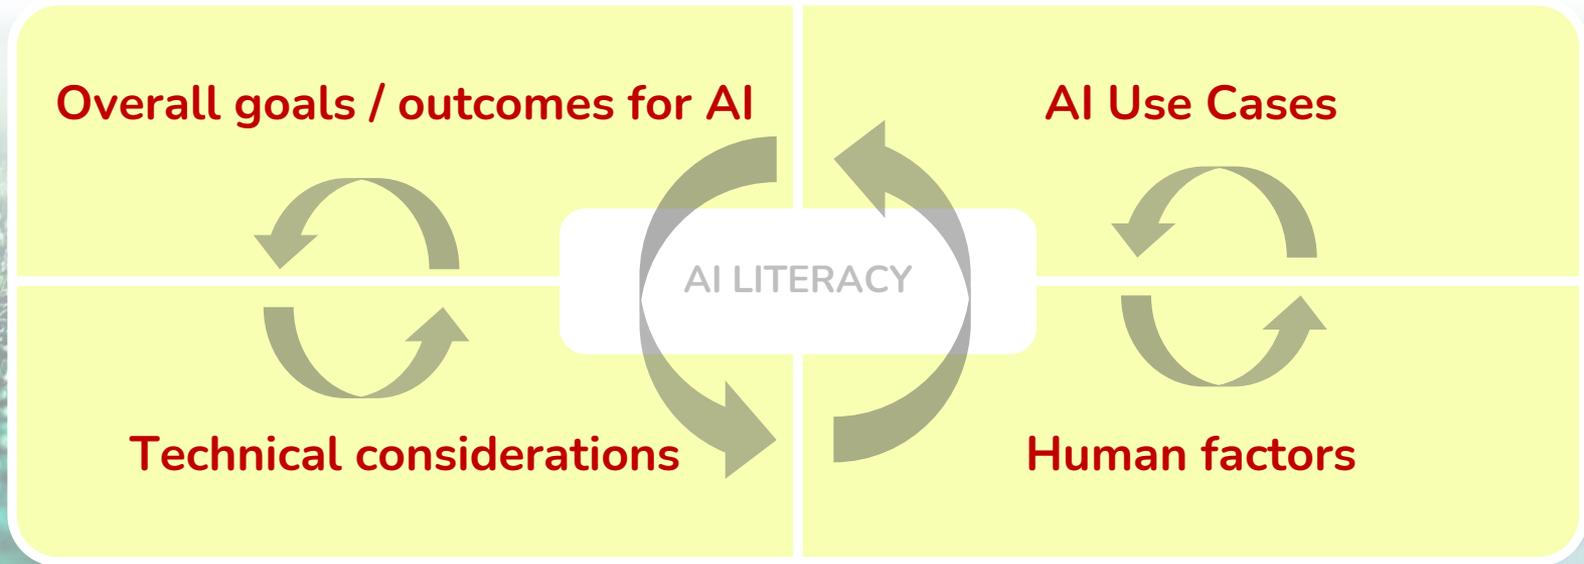


# Organizing our thoughts

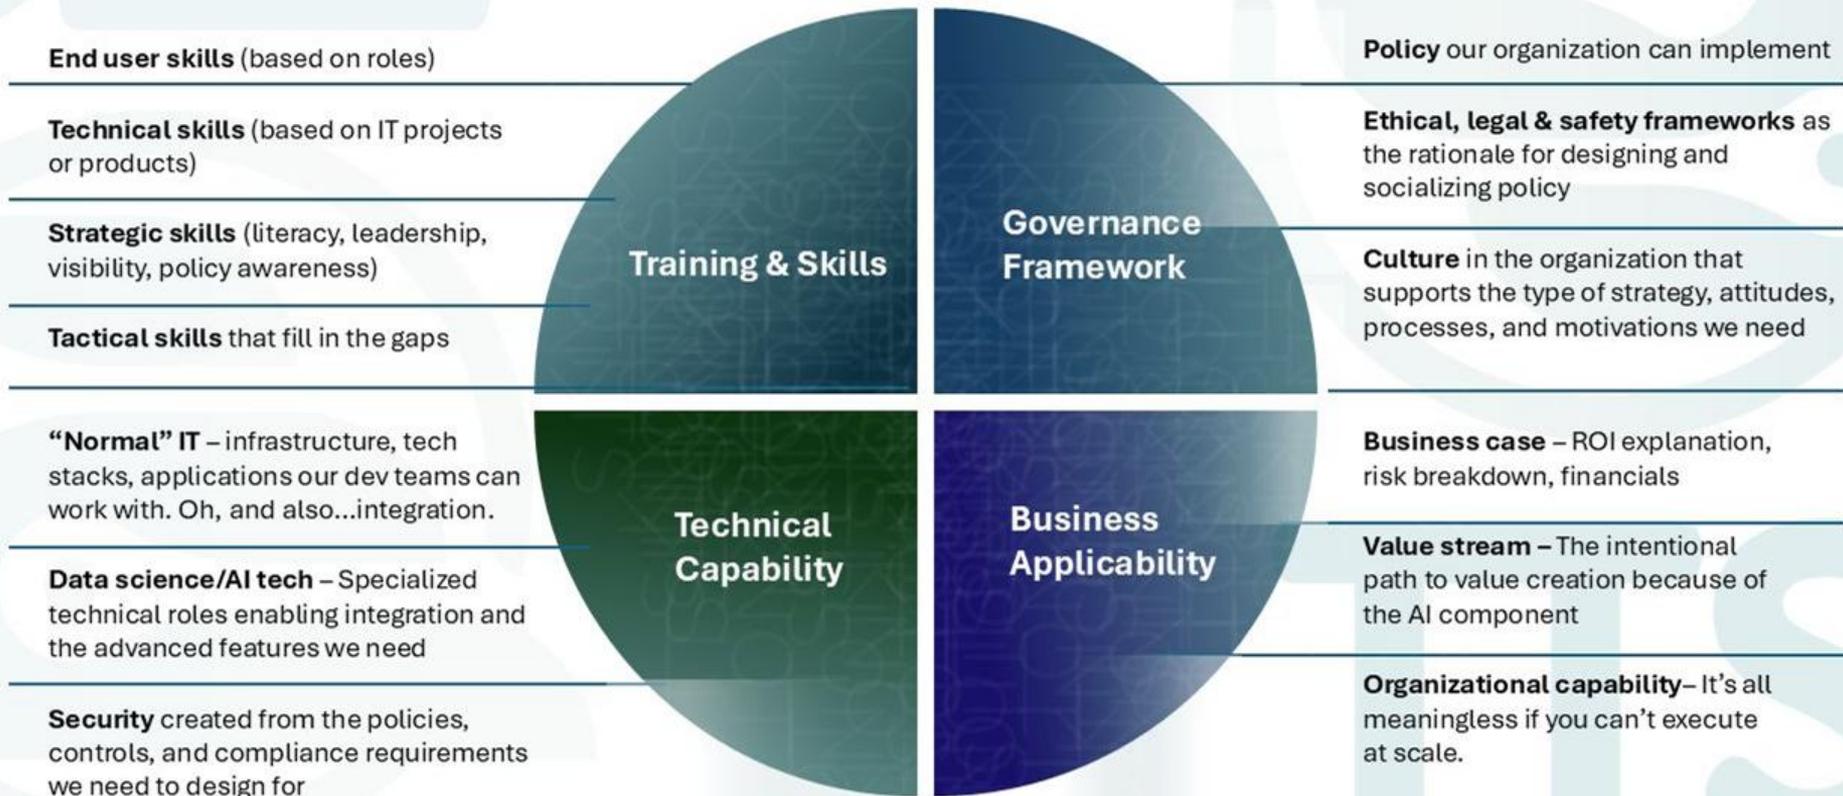


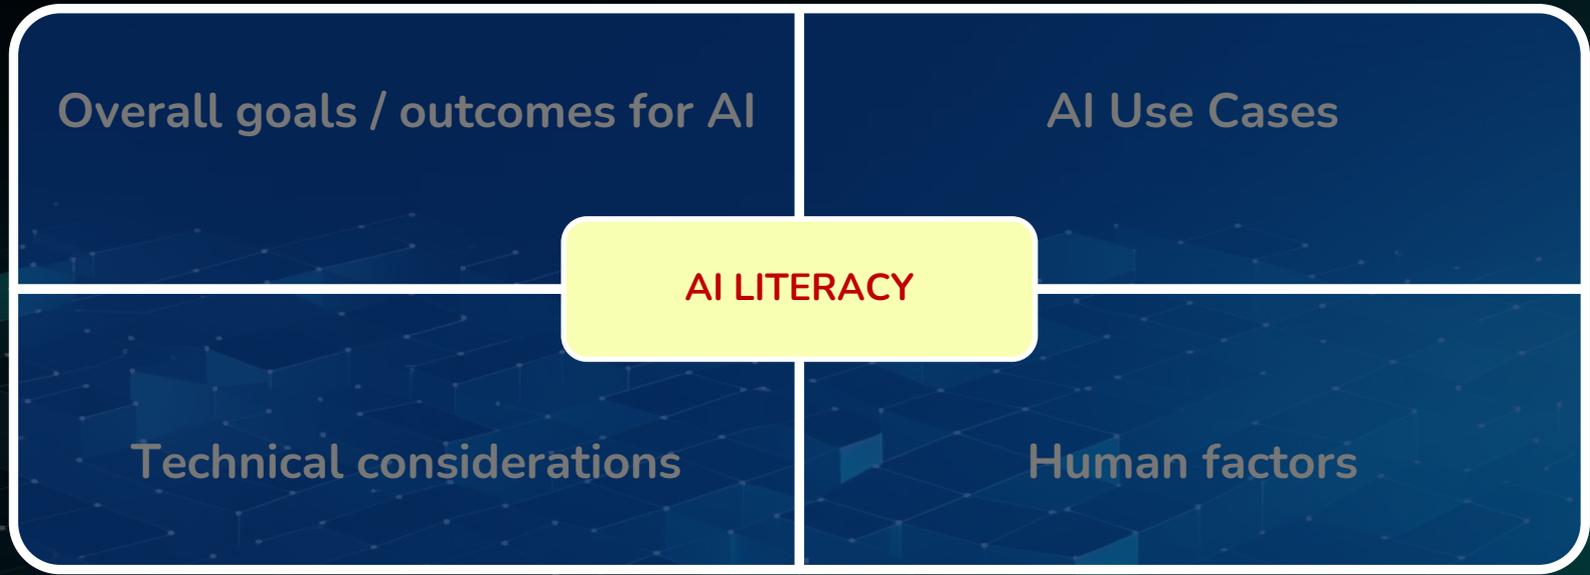
# Organizing our thoughts

*Of course there are all kinds of relationships across these dimensions.*



# Organizing our thoughts





# Generative AI Overview

# Generative AI



OpenAI ChatGPT



Anthropic Claude



Perplexity AI



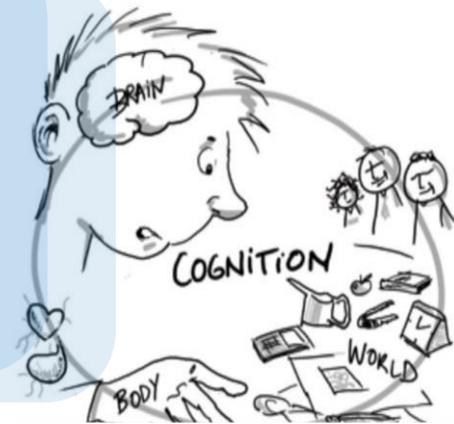
Microsoft Copilot



Google Gemini

AI models that generate text, images, and code based on training data.

Generative AI models use neural networks to identify the patterns and structures within existing data to generate new and original content.



# Types of tools

Type	Examples
Text-to-image	Midjourney, DALL-E, Stable Diffusion
Text-to-video	Runway, Synthesia, Sora
Text-to-audio	Descript, Resemble.ai
Text-to-text	OpenAI (ChatGPT, CoPilot, Azure AI), Jasper, Grammarly, co:here, Claude, Perplexity
Audio-to-text	Descript, Otter.ai, Whisper

[theresanaiforthat.com](https://theresanaiforthat.com)



## The best generative AI tools at a glance

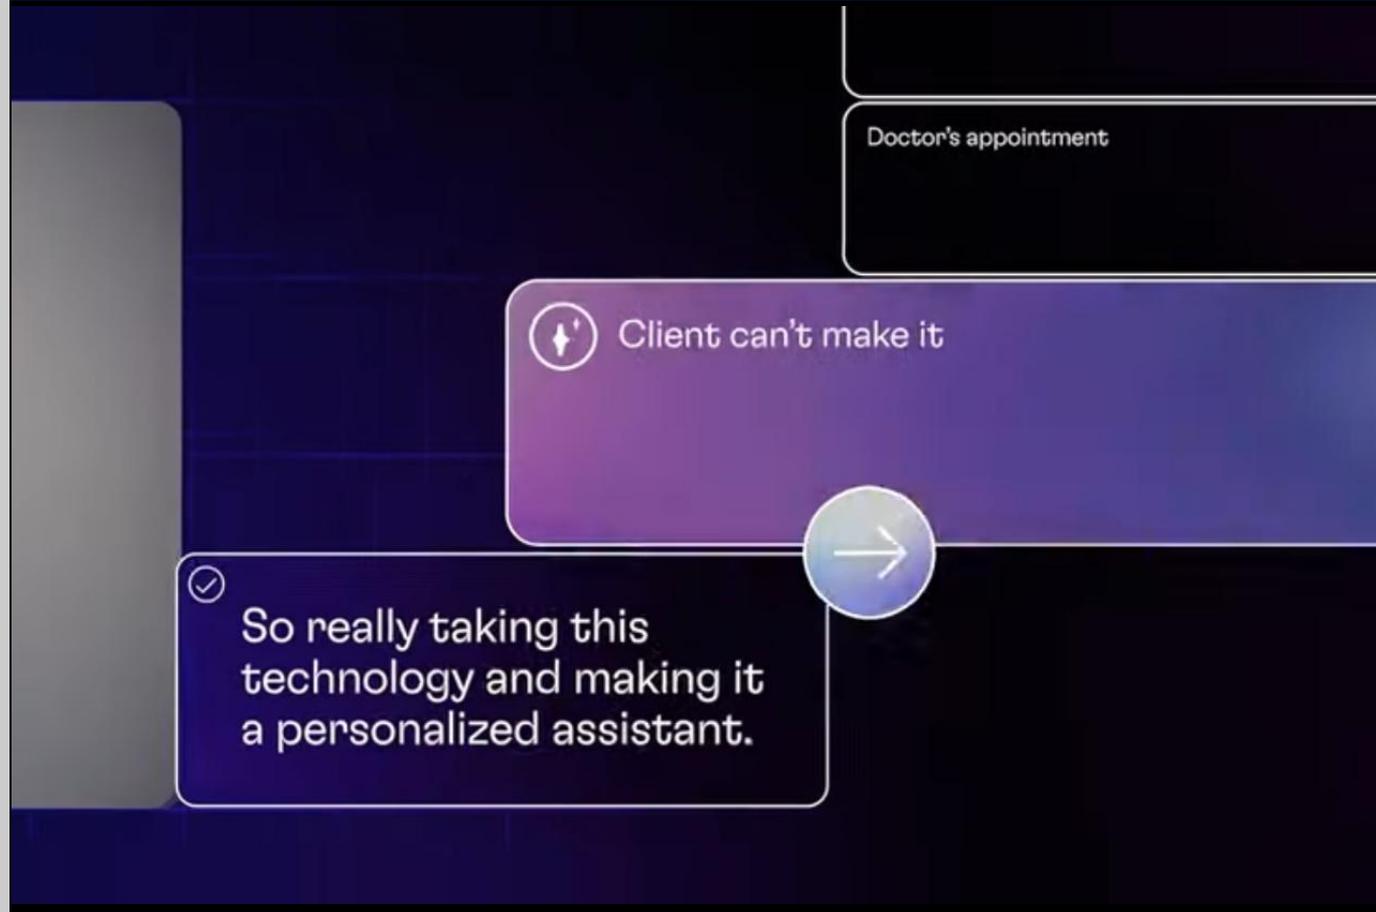
	Category	Best for
<a href="#">Jasper</a>	AI writing generators	Maintaining brand voice
<a href="#">Anyword</a>	AI writing generators	Marketing efforts
<a href="#">Shortwave</a>	AI writing generators	Email assistance
<a href="#">Notion AI</a>	AI notes	Second brain
<a href="#">Mem</a>	AI notes	Organized notes
<a href="#">Runway</a>	AI video tools	Freeform and creative videos
<a href="#">Wondershare Filmora</a>	AI video tools	AI video editing
<a href="#">Midjourney</a>	AI image tools	High-quality images
<a href="#">DALL·E 3</a>	AI image tools	Easy-to-create images
<a href="#">Adobe Photoshop</a>	AI image tools	AI-powered photo editing

## The best generative AI tools at a glance

	Category	Best for
<a href="#">ElevenLabs</a>	AI voice and audio	Versatile assets
<a href="#">Suno</a>	AI voice and audio	Creative text-to-audio results
<a href="#">Beautiful.ai</a>	AI presentation apps	Design automation
<a href="#">Pitch</a>	AI presentation apps	Sales decks
<a href="#">Wix</a>	AI website builders	Automated website creation
<a href="#">Framer</a>	AI website builders	AI-powered website design
<a href="#">Microsoft Power Apps</a>	AI app builders	Creating foundations
<a href="#">Pico</a>	AI app builders	Automated app building
<a href="#">ChatGPT</a>	AI chatbots	Versatility
<a href="#">Claude</a>	AI chatbots	Great writing quality
<a href="#">Zapier Chatbots</a>	AI chatbots	Custom-made chatbots

# According to Google Cloud

1. Customer service
2. Productivity
3. Marketing



# According to Google Cloud

## Multimodality

### What is it?:

- Understands text, images, video, and audio
- Deciphers context with unprecedented accuracy
- Produces precise, natural, and intuitive outputs

**1% to 40%**

40% of generative AI (GenAI) solutions will be multimodal (text, image, audio and video) by 2027, up from 1% in 2023, according to Gartner, Inc

**\$98.9B** by 2037

Predicted size of global multimodal AI market<sup>2</sup>

# Big caveat: Change is happening in days and weeks



**LIMITED SHELF LIFE ITEM**

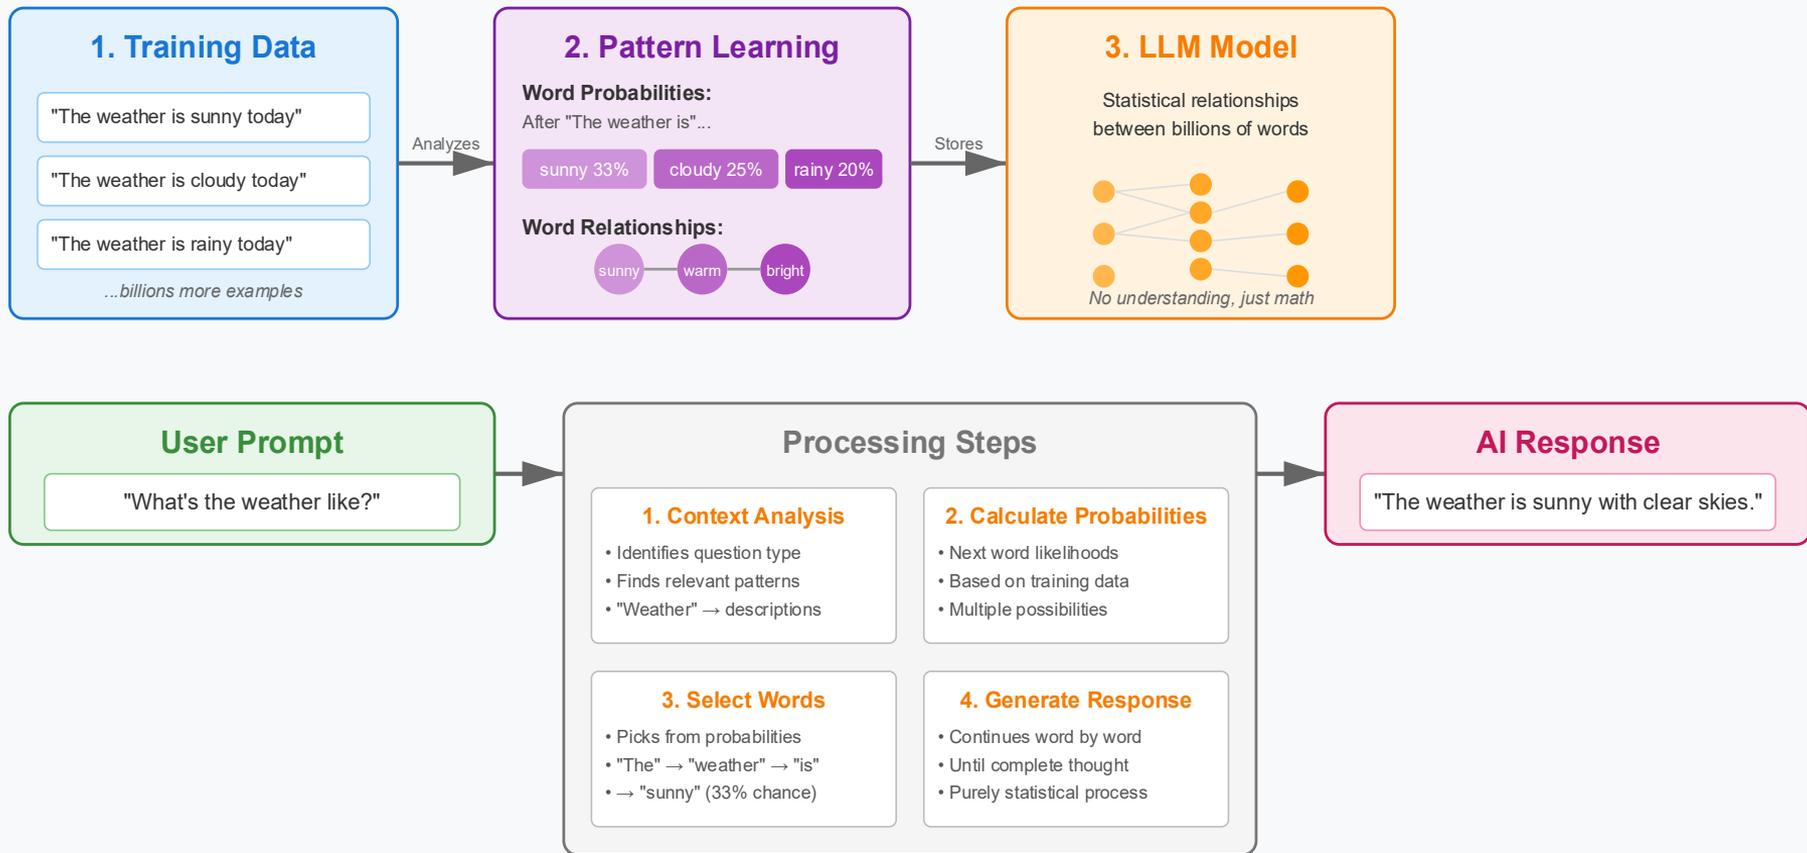
- AI Products
- Features and capability
- Underlying technology
- Legal precedents
- Labor market consequences
- Risks
- New use cases

# What is a large language model? (LLM)

- Neural networks trained on extensive collections of text data from the internet, books, articles, etc.
- Learns to “understand” statistical relationships within language.
- Can generate, manipulate, and analyze all sorts of stuff in a very similar way to humans.
- The trick is understanding what it’s good at and where it stinks.
- They can do some math, but language-based data is the key strength.

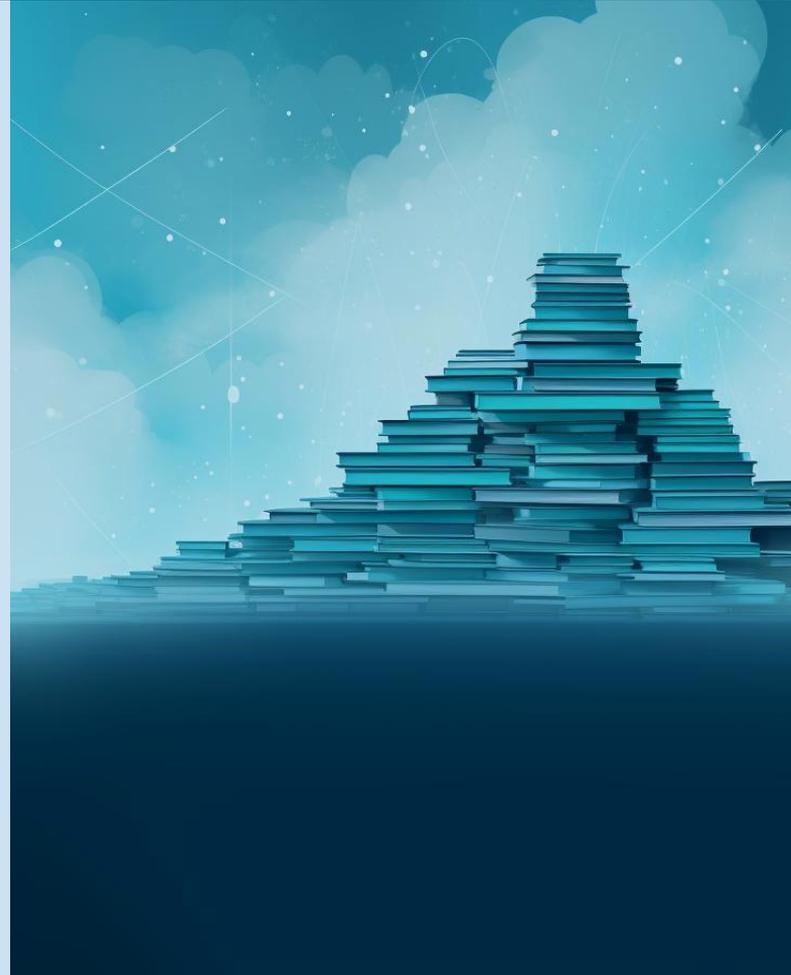


# How LLMs use statistical patterns to generate responses



# Key Insights for Leaders:

- LLMs predict the most statistically likely next words — they don't “understand” meaning
- This explains hallucinations: plausible-sounding but potentially incorrect information
- Quality depends on training-data patterns — garbage in, garbage out
  - **No reasoning**
  - **No facts database**
  - **No truth verification**
  - **Just statistical patterns**



# What is a large language model? (LLM)

- **AI power users are now people with good language/dialectic skills.**





A bit about prompting...

# What is a “dialectic approach?”

We primarily interact with LLMs using what we'll call a “dialectic approach,” meaning that we should always approach interaction with an LLM as a conversation, anticipating that we may not get what we need from the AI tool in just one exchange – it will likely require a series of exchanges so we can establish context, help the AI learn what we need, and refine the direction of the conversation as we go. LLM tools don't just use machine learning to power their core features - they also learn as they go, adapting to our feedback in real time.

# Prompting basics in plain English

- Initiate with as much context as possible. Examples, documents, data...
- Provide examples of previous work that shows what you want the AI to do.
- Approach a generative AI interaction is a conversation...you rarely get what you want the first time.
- Build conversation threads focused on one type of task or project, creating feedback loops and more context over time.
- Include details in your query to get more relevant answers
- Ask the model to adopt a persona
- Use good punctuation and “delimiters” to clearly indicate distinct parts of the input, especially in longer prompts.
- Specify the steps required to complete a task: “break it down into parts.”
- Specify the desired length of the output.

# Bad Prompts: Examples

**Vague Query:** Tell me about agile project management.

- Lacks specificity and context.

**No Persona Adoption:** How do I manage projects?

- Doesn't specify the AI to adopt the role of an experienced project manager.

**Unclear Delimiters:** What are the risks in project management and how to mitigate them, also tell me about team management.

- Combines multiple questions without clear separation.

**No Steps Specified:** How do I implement an ERP system?

Lacks detail on the desired process or steps.

**No Examples Provided:** How do I improve team efficiency?

- Doesn't provide context or examples to tailor the response.

**No Output Length Specified:** Give me a project report.

- Doesn't specify desired length or detail for the report.

# Good Prompts: Examples

**Detailed Query:** "I'm an agile project manager implementing a new ERP system. What are key considerations for integrating existing customer data into the new system?"

- *Provides specific context and details.*

**Persona Adoption:** "Act as an experienced agile project manager. What strategies would you suggest for maintaining team morale during the stressful final stages of an ERP implementation?"

- *Directs the AI to adopt a specific role.*

**Clear Delimiters:** "Identify the top 5 risks in agile ERP implementation | Provide methods to mitigate each risk."

- *Clearly separates two parts of the prompt.*

**Steps Specified:** "Outline the steps for conducting a sprint retrospective in an ERP project."

- *Asks for a specific, step-by-step process.*

**Examples Provided:** "Considering challenges like scope creep and resource constraints, which I've faced in previous projects, how can I better manage project scope in my current ERP implementation?"

- *Includes examples for context.*

**Specifying Output Length:** "List in bullet points, no more than 10, the key metrics I should track during an ERP system implementation."

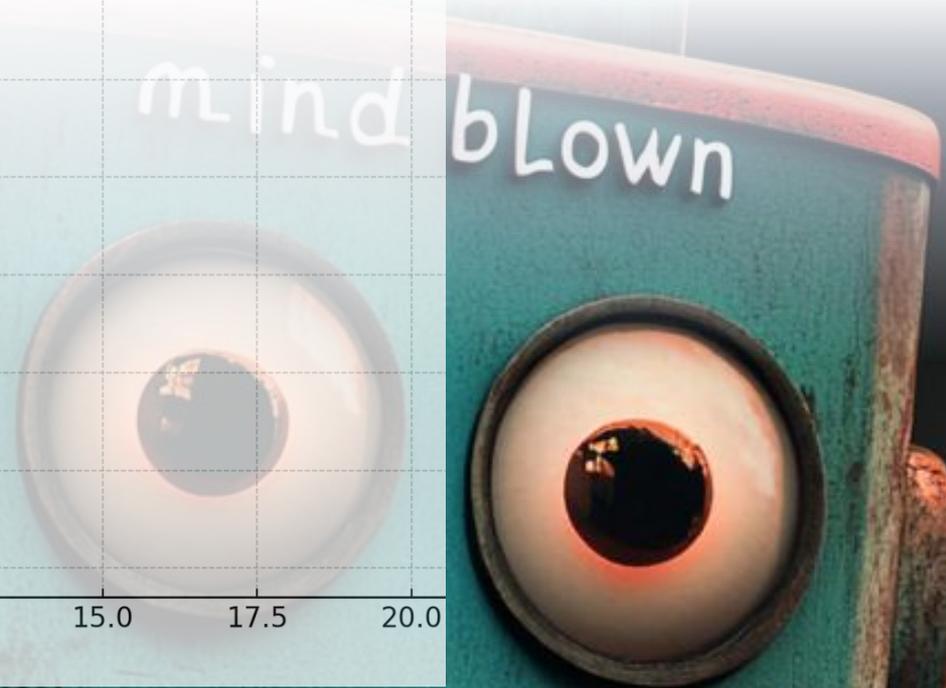
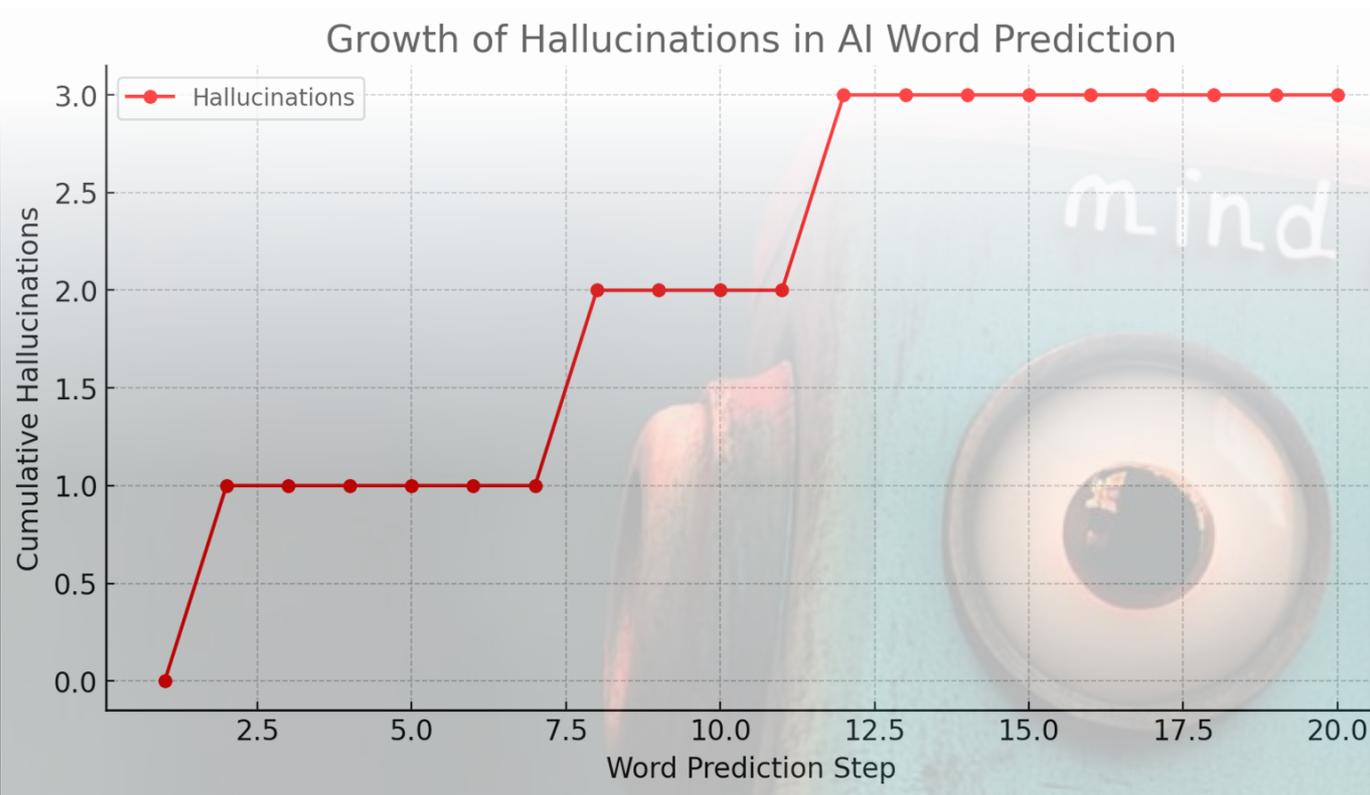
- *Specifies the format and length of the response.*

# What about other types of AI?

AI Type	Example Use Case
<b>Generative AI</b>	Automating reports, summarizing policies
<b>Machine Learning</b>	Fraud detection, risk analysis
<b>Computer Vision</b>	Analyzing satellite data, cancer detection
<b>Natural Language Processing</b>	Chatbots for public services, contract analysis

There is more to being successful with enterprise AI than just generative AI tools and large language models (LLMs). Other types of AI matter, and technology dependencies do, too.

# About those hallucinations



Overall goals / outcomes for AI

AI Use Cases

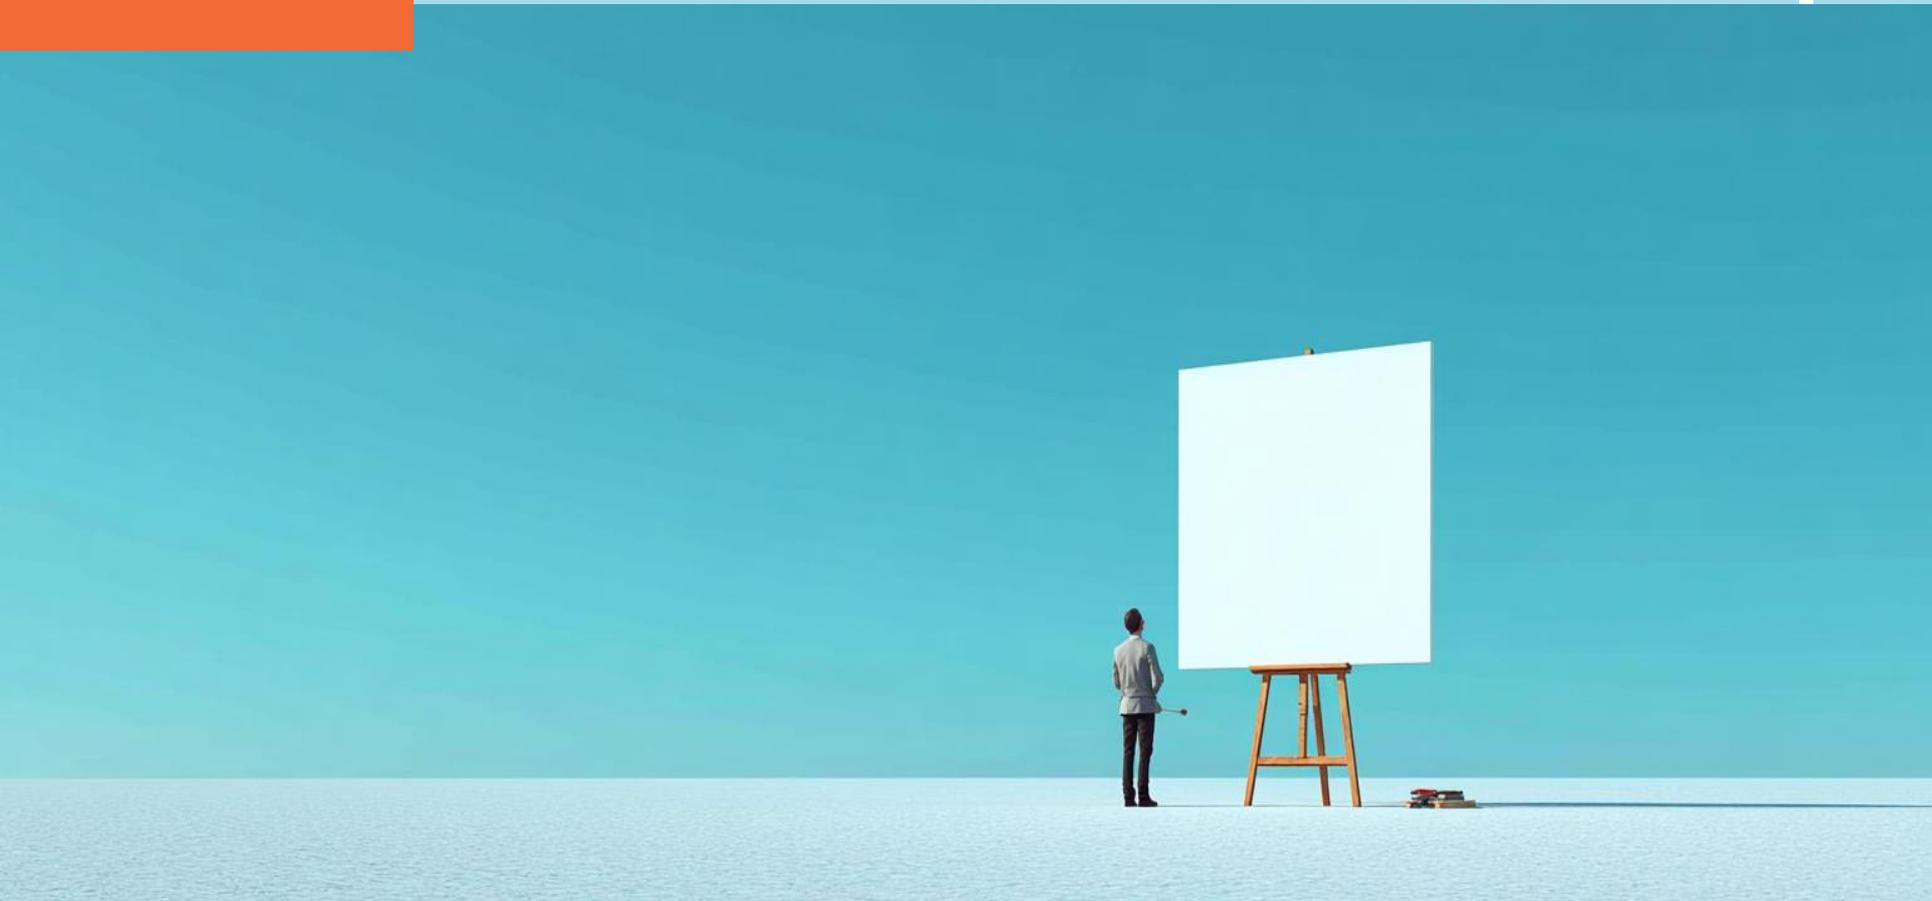
AI LITERACY

Technical considerations

Human factors

What are we really trying to do?

# What Do We Want From AI?



# Why?

- **Efficiency?**
- **Do something new? New service, new product?**
- **More speed?**
- **Save time?**
- **Save money?**
- **Adoption seems important?**

*And worth considering...operational or customer-facing?*

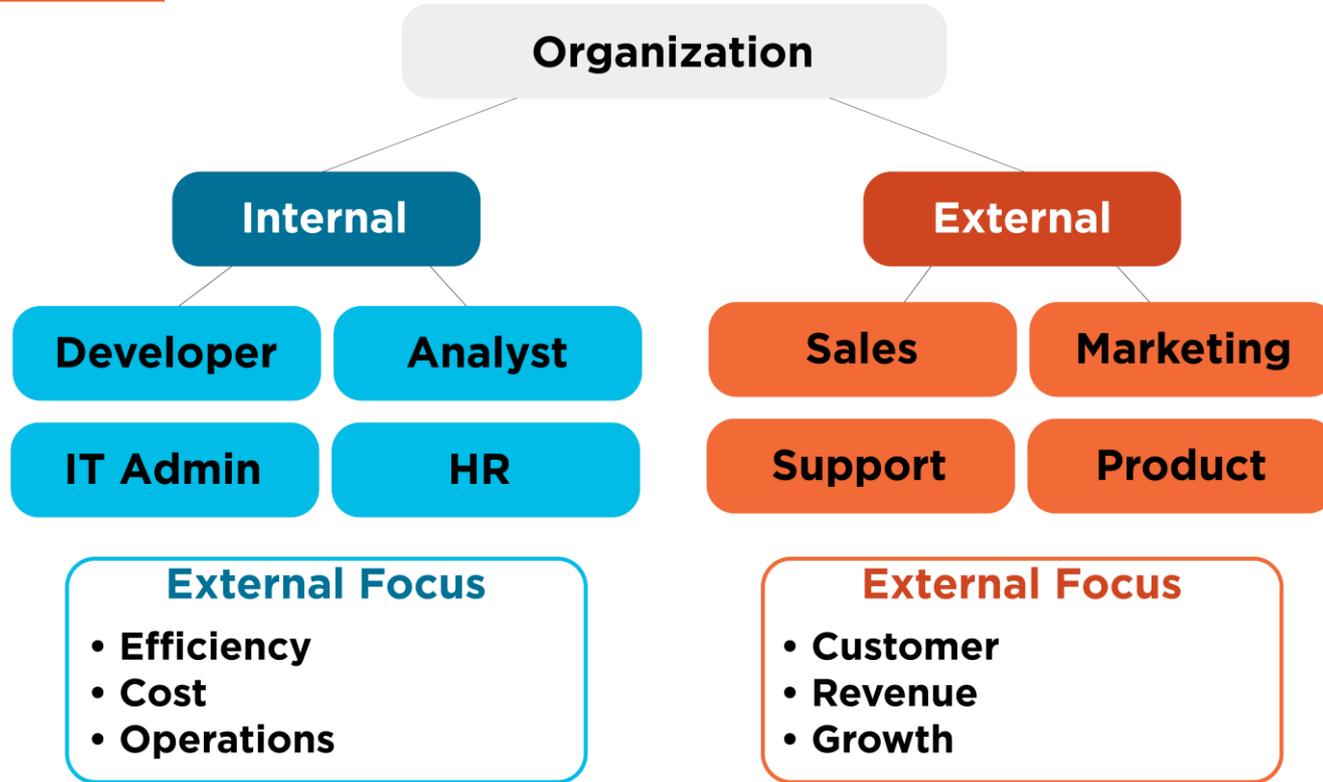
# AI can definitely unlock value, but...

There are some hard conversations. Emergent tech, big change initiatives assume some enterprise prerequisites:

- Good leadership?
- Agility?
- Safe, high-performance culture?
- Good portfolio management practices?
- High innovation performance?
- Tech dependencies?



# What Do We Want From AI?

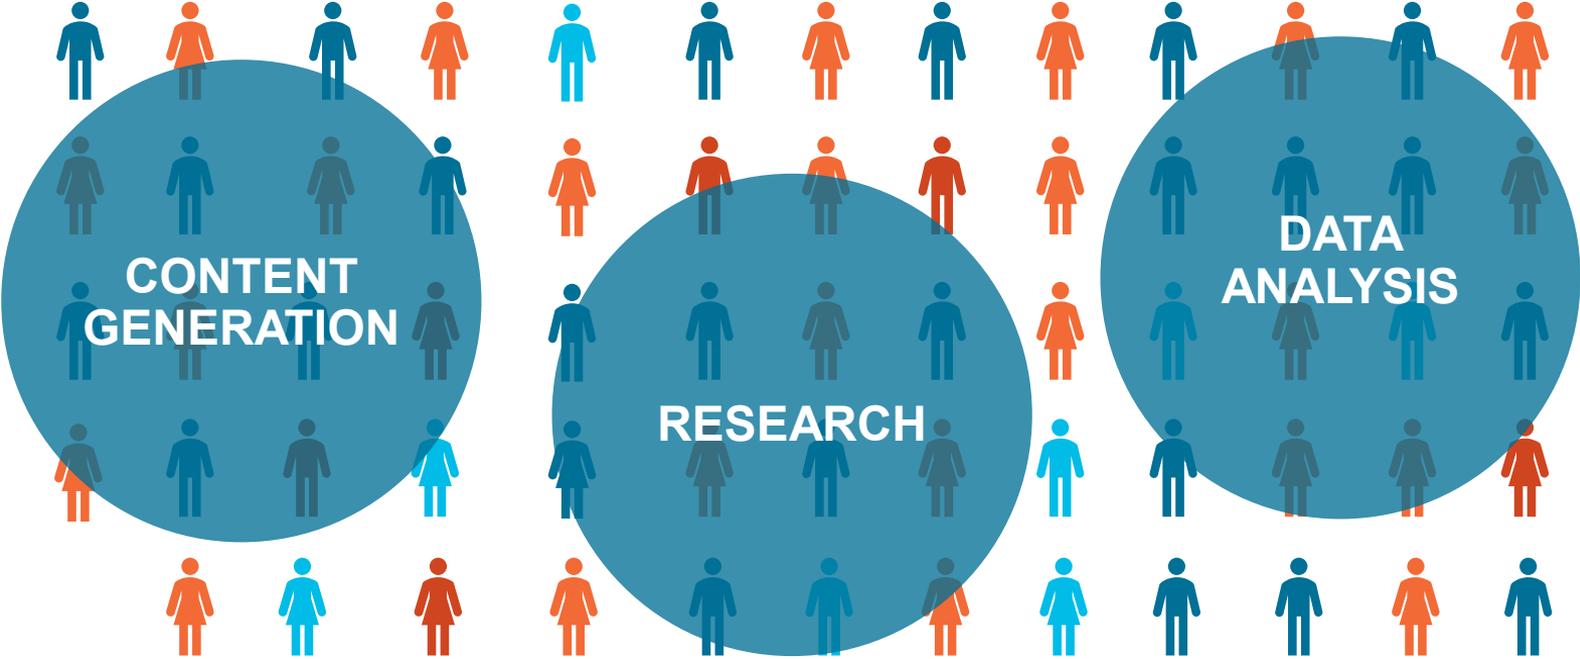


# What Do We Want From AI?



- Marketing
- HR
- IT
- Sales

# What Do We Want From AI?



- Marketing
- HR
- IT
- Sales

# A template for action

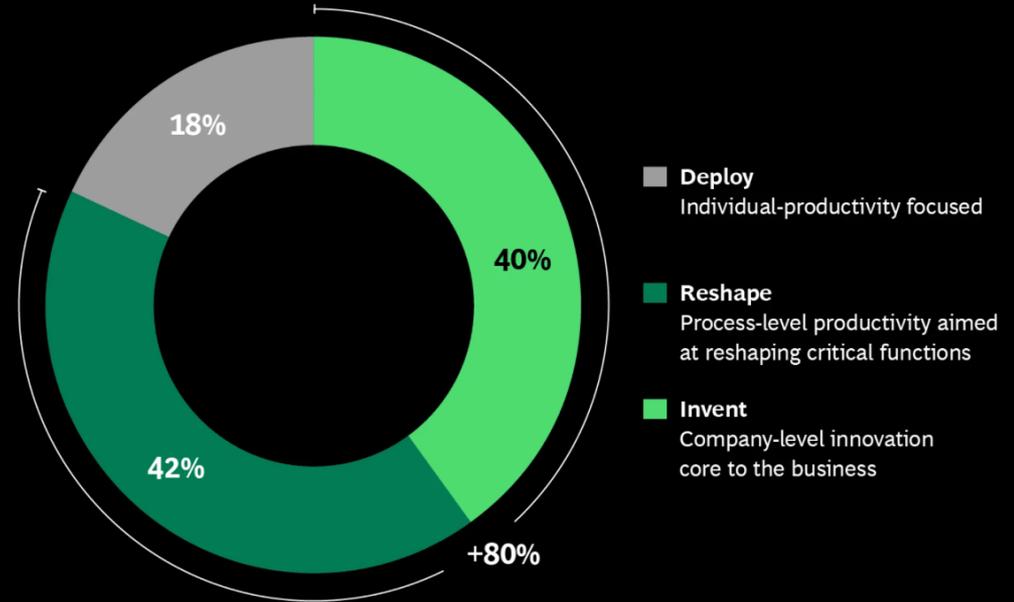
- **What AI opportunity** is most relevant?
- **What's the long-term goal?**
- **What's the first small step** you can take to explore it?
- **Who do you need buy-in from?**
- **What risks or challenges** might you encounter?
- **How will you measure success?**



# Research from Boston Consulting Group

- BCG says leading companies go well beyond “deploy”
- “Deploy” in this case is mainly about tactical, employee-level use cases, and tool/tech considerations are pretty straightforward
- “Reshape” and “Invent” get more aspirational and require effective management across the various AI dimensions

**They focus +80% of their AI investments in reshaping critical functions and inventing new products and services**



# Leadership and AI

*AI can't lead change — that's your job.*

## AI Strengths

- Processes massive data quickly
- Automates repetitive tasks
- Identifies patterns & trends
- Supports decision-making



## Human Leadership Strengths

- Build trust and team cohesion
- Guide employees through change
- Develop strategic long-term vision
- Coach and mentor employees

# AI can't think strategically.

- AI can identify patterns, but humans interpret meaning & implications.
- Leaders anticipate unintended consequences AI might miss.
- AI can optimize or execute workflows, but leaders ensure alignment with goals.



*It's part of a data story  
that hasn't changed.*

```
011010011010  
10011 00110  
0100 0011  
01 10 10  
1 000110  
010011010011
```

Data



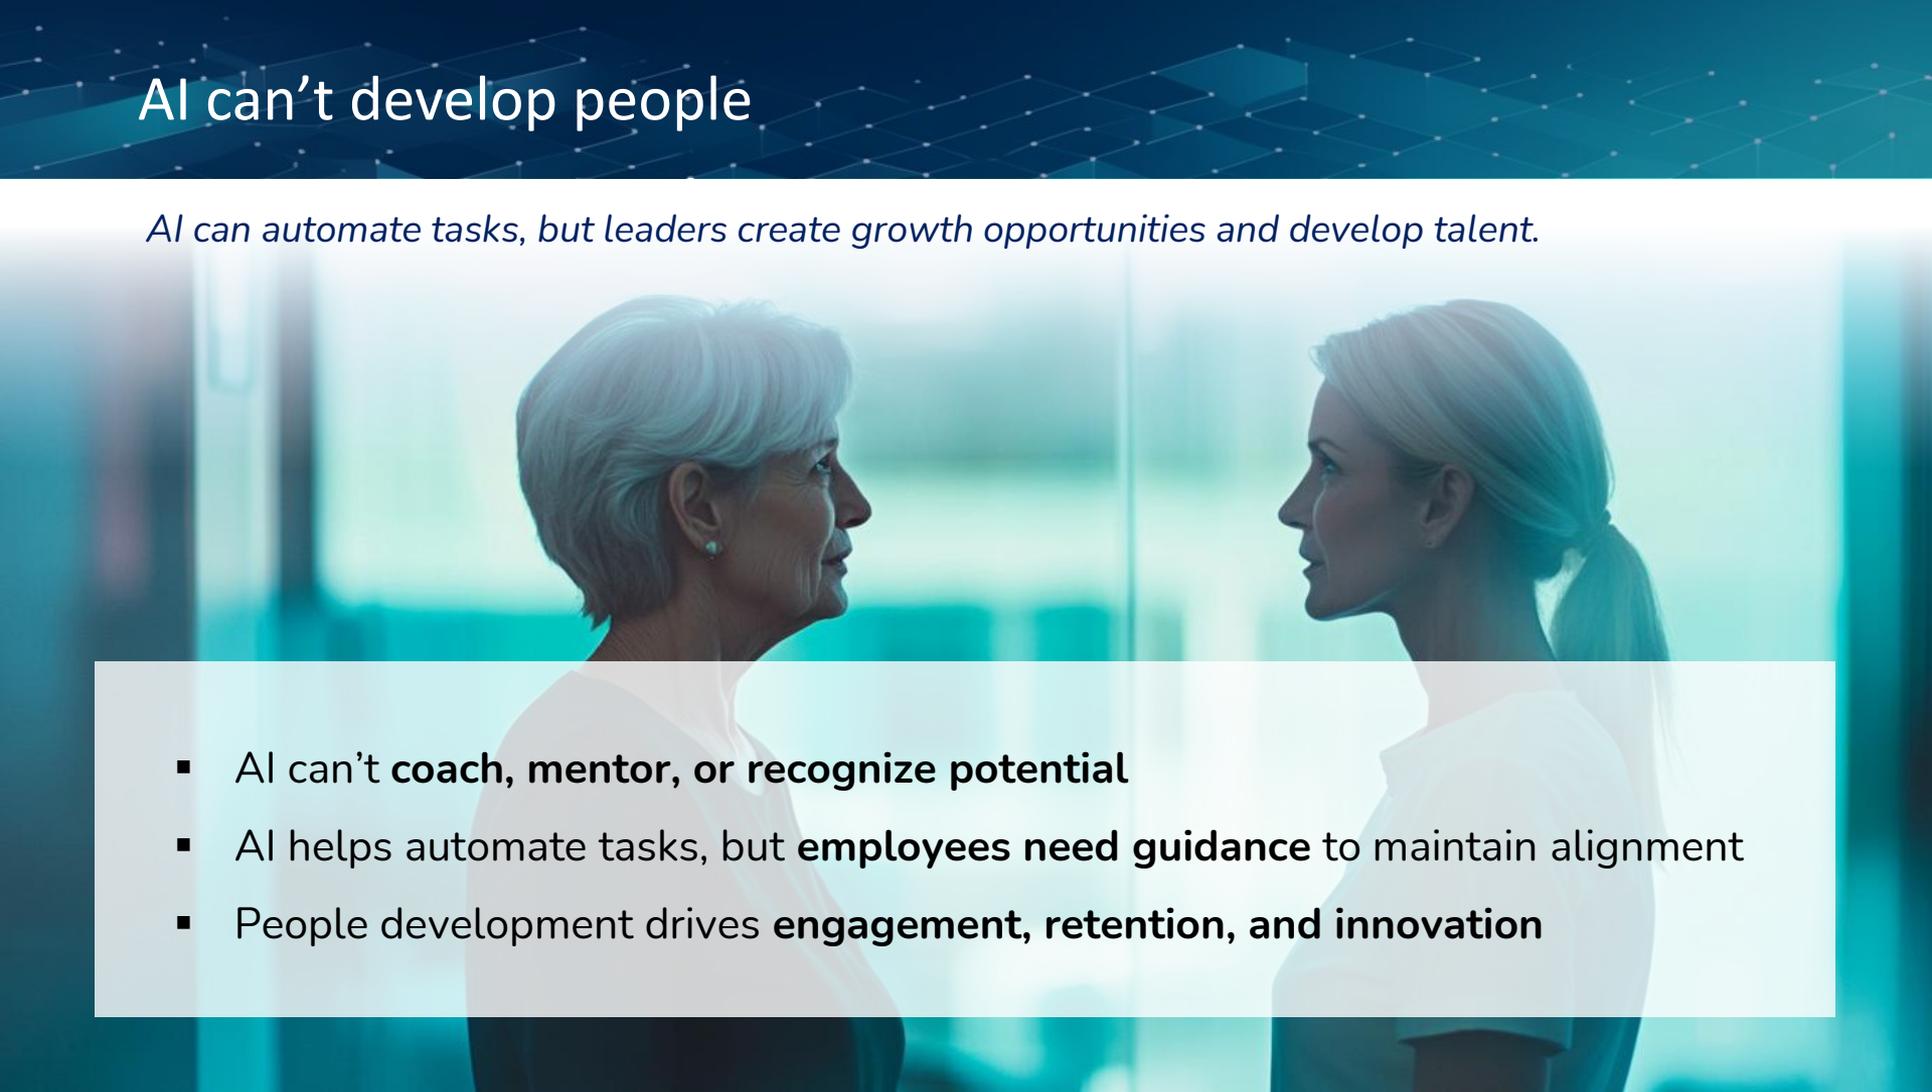
Insights/Creation



Decisions

# AI can't develop people

*AI can automate tasks, but leaders create growth opportunities and develop talent.*

- 
- A photograph of two women in profile, facing each other in a modern office setting. The woman on the left has short, light-colored hair, and the woman on the right has long, light-colored hair tied back. They are both wearing light-colored blouses. The background is a bright, out-of-focus office interior with large windows.
- AI can't **coach, mentor, or recognize potential**
  - AI helps automate tasks, but **employees need guidance** to maintain alignment
  - People development drives **engagement, retention, and innovation**

# AI can't guide change.

*AI impacts our people and work.*

- **Communicate the "why" behind AI adoption** — connect it to mission, explain the business case
- **Acknowledge uncertainty & resistance** — empathize with concerns.
- **Set a culture of continuous learning** — help teams adapt
- **Normalize change discussions** — make AI a part of onboarding, meetings and reviews

# Co-Intelligence / Partnering

## How do you think about Co-Intelligence?

1. Partner
2. Shared responsibility
3. Help
4. Offload work
5. Giving up some control
6. Other thoughts?

**Reminder: How do we best operate?**

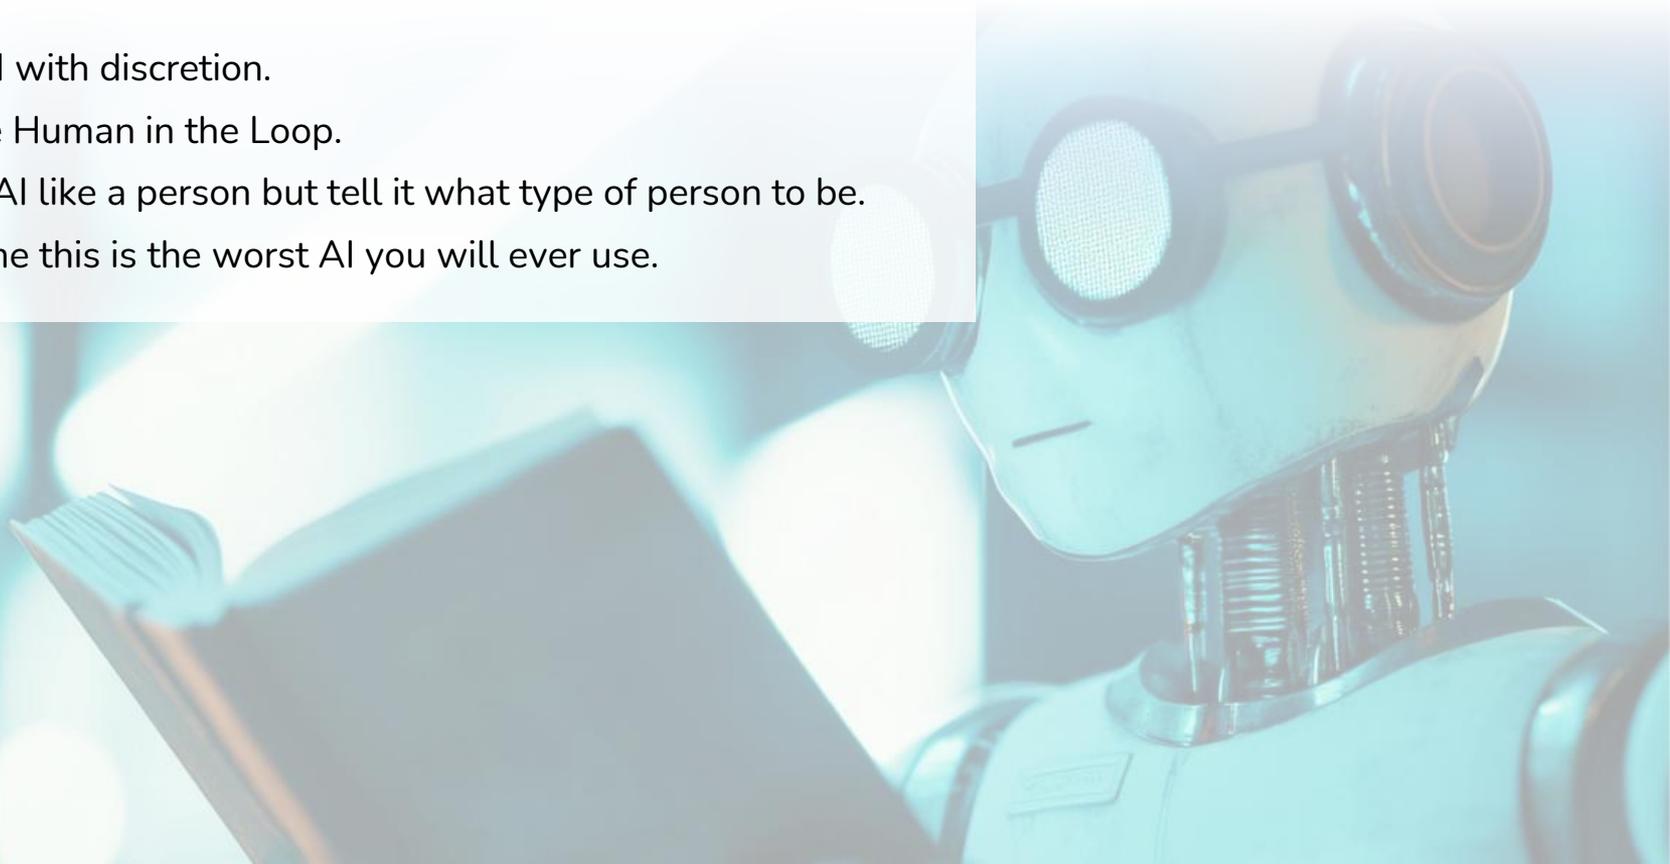
**Reminder: How does AI best operate?**

**Given each's nature how can you see this working?**



# Four rules to consider

1. Use AI with discretion.
2. Be the Human in the Loop.
3. Treat AI like a person but tell it what type of person to be.
4. Assume this is the worst AI you will ever use.



Overall goals / outcomes for AI

**AI Use Cases**

**AI LITERACY**

Technical considerations

Human factors

# Practically speaking, what is Generative AI good at, really?

## Good at:

- Generating Text
- Answering Questions
- Language Translation
- Mimicking Writing Styles
- Basic problem solving
- Summarizing Information
- Recommending Resources
- Conversational Role-playing
- Providing Programming Help
- Offering Explanations on Concepts

## NOT good at

- Analyzing for relevance
- Reading in Emotion and Empathy
- Understanding Context
- Catching copyright issues
- Memory of Past Issues and context
- Fact checking, missing relevance
- Memory of Past Interactions
- Deeper understanding of personality
- Copyright policing
- Catching bias

# Practically speaking, what is Generative AI good at, really?

## IDENTIFY



Finding specific elements

## SUMMARIZE



Condensing into key piece

## MERGE



Combining separate elements

## ORGANIZE



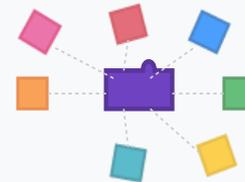
Arranging into categories

## DRAFT



Framework before completion

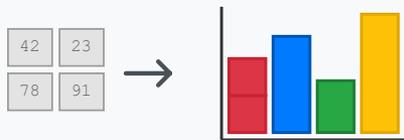
## BRAINSTORM



Generating multiple options

# Practically speaking, what is Generative AI good at, really?

## VISUALIZE



Transforming data into charts

## COMPARE



Highlighting differences

## EXTRACT



Pulling out specific data

## TRANSFORM



Changing format/structure

## ANALYZE



Breaking down patterns

## VALIDATE



Checking data quality

# Automation: big “A,” small “a.”

## Machine Attributes:

### Strengths:

- Repetitive actions
- Working with data
- Consistent performance
- Tireless operation
- Scalability

### Limitations:

- No instinctive understanding
- Needs context
- No emotional discernment
- No flexibility without training
- Fit drift (function, need, etc.)

## Human Attributes:

### Strengths:

- Instinctive judgment
- Creativity
- Emotional perception
- Context-sensitive
- Versatile

### Limitations:

- Fatigue
- Scale limits
- Errors
- Hates monotony

# DISCUSSION: Busy work

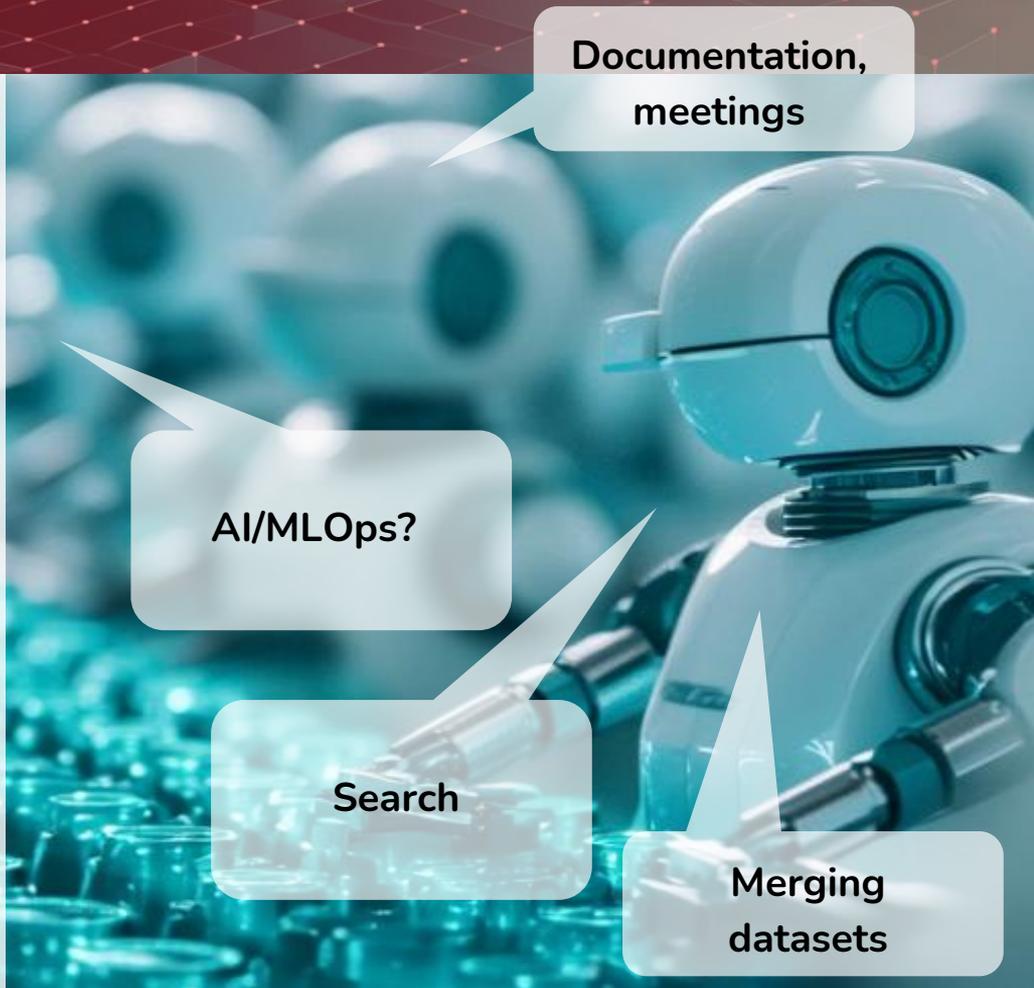
## Machine Attributes:

### Strengths:

- Repetitive actions
- Working with data
- Consistent performance
- Tireless operation
- Scalability

### Limitations:

- No instinctive understanding
- Needs context
- No emotional discernment
- No flexibility without training
- Fit drift (function, need, etc.)



## Initiating Projects

Identifying stakeholders

creating a basic template

Planning, communication, meeting with team on status and send out MoMs

Creating checklists

## Stakeholder Communication

send notes to the team members

Tracking action item statuses

Organizing meetign notes

create templates for taking notes

schedule interviews with stakeholders

Communication Approval and Sign-off process

## Monitoring and Progress

Schedule, resource, budget managements

Ensure all fields of a doc are populated.

Sequence code migration based on dependencies

Tracking action item statuses

Set reminders for meetings

## Risk Management

Assess environmental related defects / issues.

Defect, risk and issues management

Get all the SOP done. Risk and issues management, work on metrics, reporting

## Data Analysis

Forecast financial balances based on historical data.

tracking metrics

Assess environmental related defects / issues.

create flowcharts

Running Reports

## Value and Financials

Data collection, report analysis and preparing the presentations, decision making, work on financials and budgets for the project

Forecast financial balances based on historical data.

## Resourcing & Capacity

resource forcatsting

## Change Management

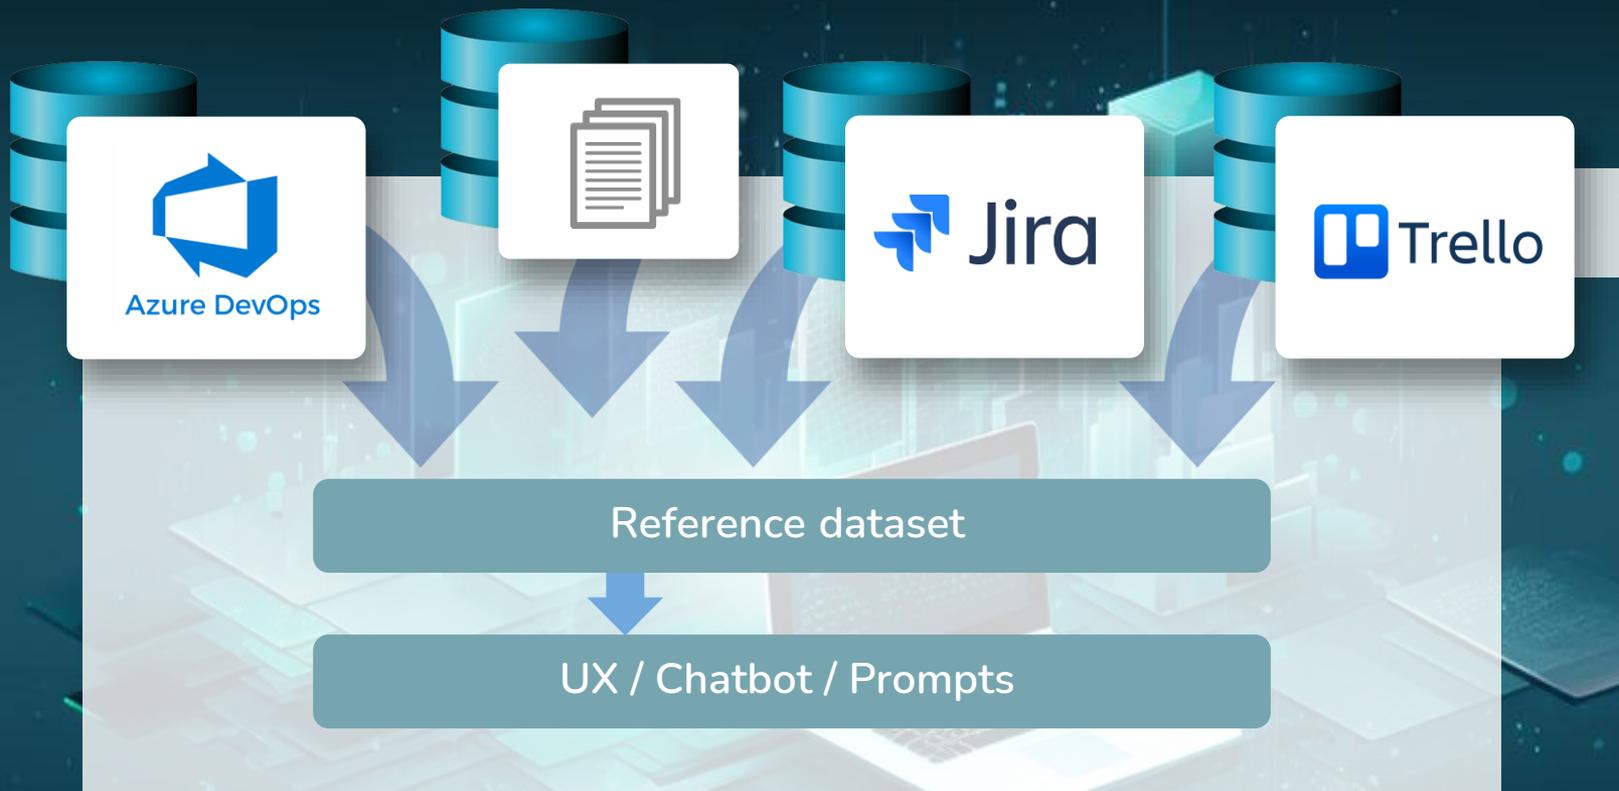
Deconflict business processes. Deconflict technical processes

project documentations and change management

# DEMO: Project Data Exports & Imports

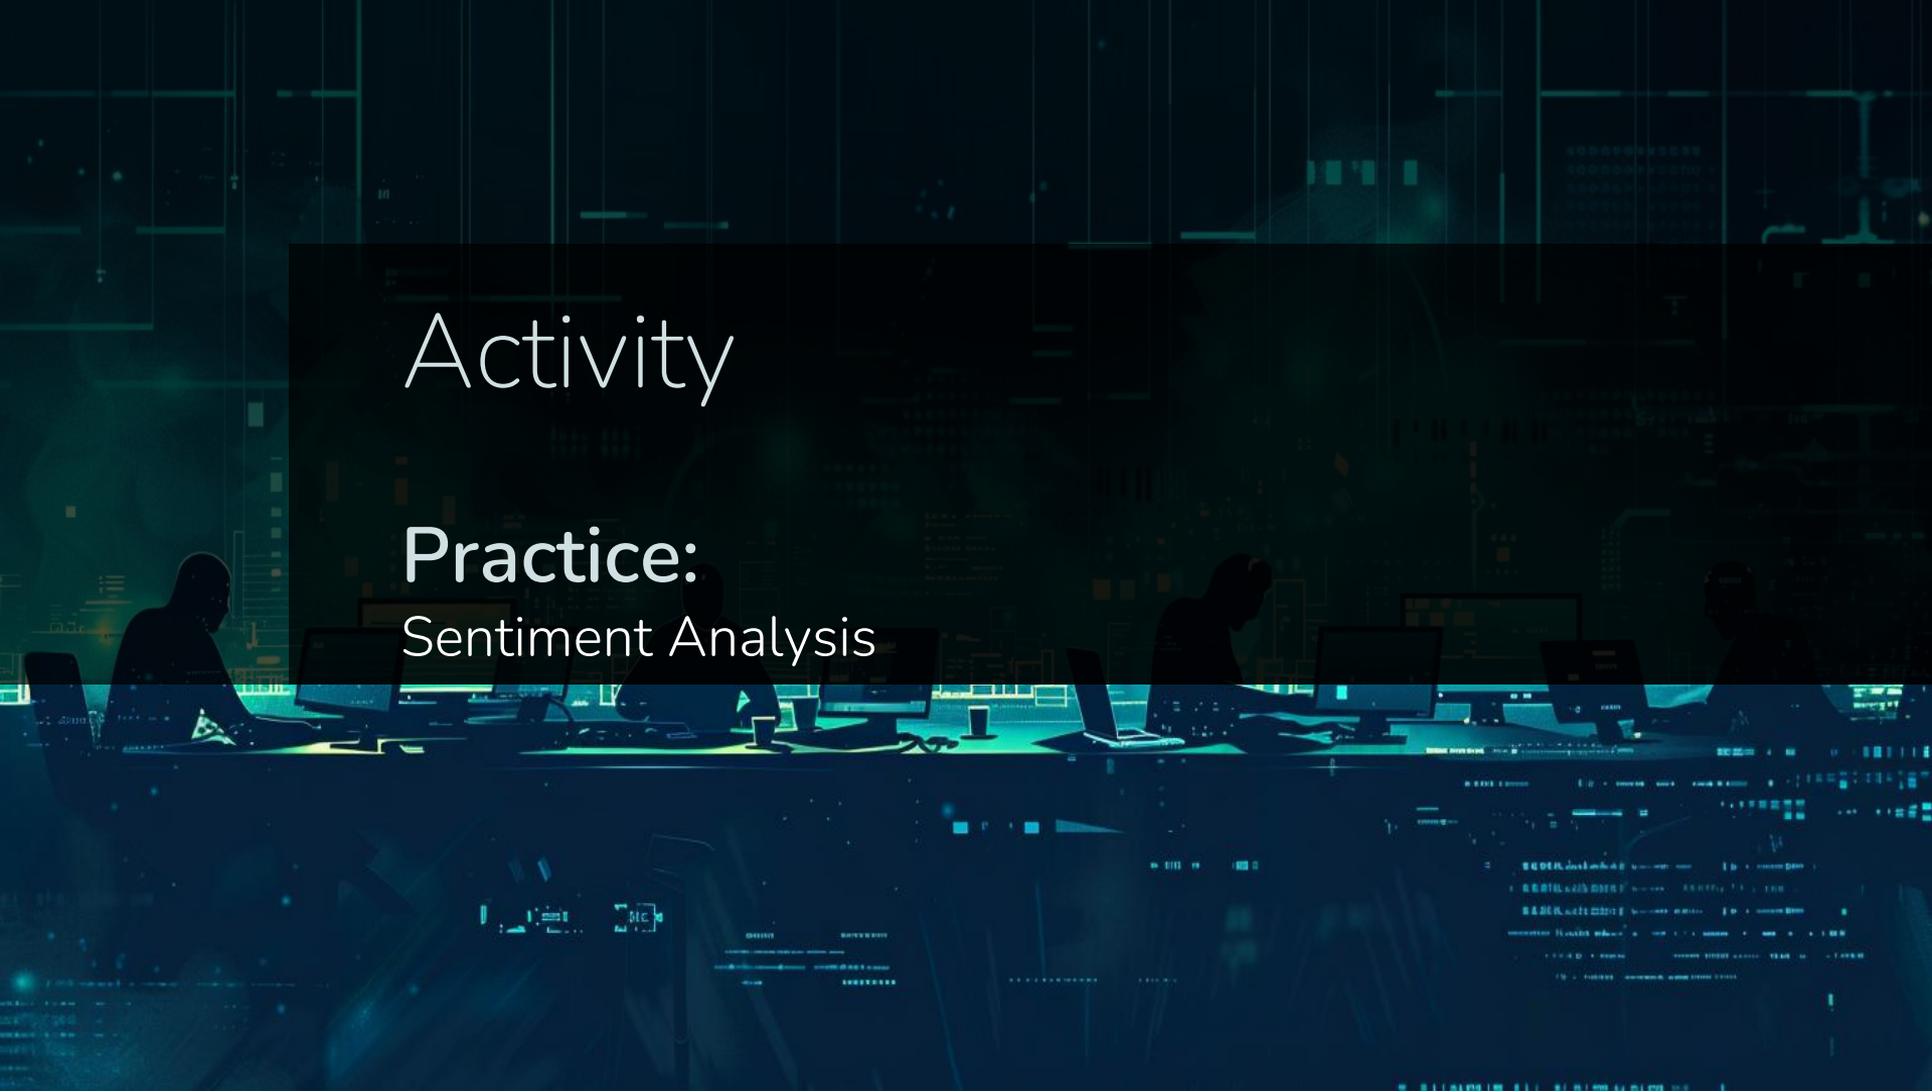


# Data Sources:



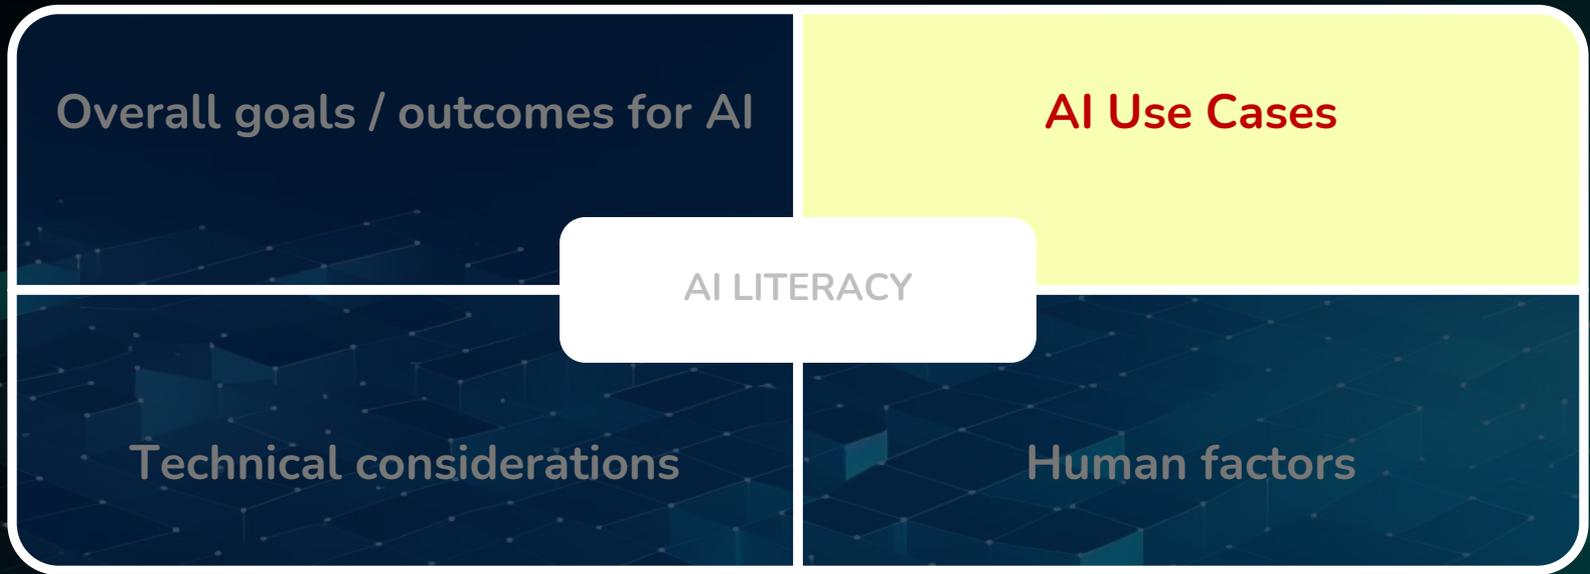
# DEMO: Stakeholder Sentiment Analysis





# Activity

## Practice: Sentiment Analysis



Tactical use, strategic use, acceptable use

# Why?

- Efficiency?
- Do something new? New service, new product?
- More speed?
- Save time?
- Save money?
- Adoption seems important?

*And worth considering...Operational or customer-facing?*

Tactical use cases are usually about saving time.



Customer (internal or external) use cases are about creating value.



# AI can't understand your framework for value.

Our approach to AI has to align with how we communicate our “why” and definition of value.

- Mission driven?
- Profit-driven?
- Do we take a portfolio approach?
- Do our people optimize for local or system-wide value?
- Product P&L?

Example:

Increase Revenue

**Increasing** sales to new or existing customers. Delighting or Disrupting to increase market share and size

Protect Revenue

Improvements and incremental innovation to **sustain** current market share and revenue figures

Reduce Costs

Costs that we are **currently** incurring, that can be **reduced**. More efficient, improved margin or contribution

Avoid Costs

Improvements to **sustain** current cost base. Costs we are not currently incurring but may do in the future

## Initiating Projects

Identifying stakeholders

creating a basic template

Planning, communication, meeting with team on status and send out MoMs

Creating checklists

## Stakeholder Communication

send notes to the team members

Tracking action item statuses

Organizing meetign notes

create templates for taking notes

schedule interviews with stakeholders

Communication Approval and Sign-off process

## Monitoring and Progress

Schedule, resource, budget managements

Ensure all fields of a doc are populated.

Sequence code migration based on dependencies

Tracking action item statuses

Set reminders for meetings

## Risk Management

Assess environmental related defects / issues.

Defect, risk and issues management

Get all the SOP done. Risk and issues management, work on metrics, reporting

## Data Analysis

Forecast financial balances based on historical data.

tracking metrics

Assess environmental related defects / issues.

create flowcharts

Running Reports

## Value and Financials

Data collection, report analysis and preparing the presentations, decision making, work on financials and budgets for the project

Forecast financial balances based on historical data.

## Resourcing & Capacity

resource forcatsting

## Change Management

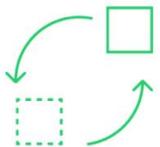
Deconflict business processes. Deconflict technical processes

project documentations and change management

# BCG – Applying the Digital Acceleration Index



**Deploy** AI in everyday tasks to realize 10% to 20% productivity potential

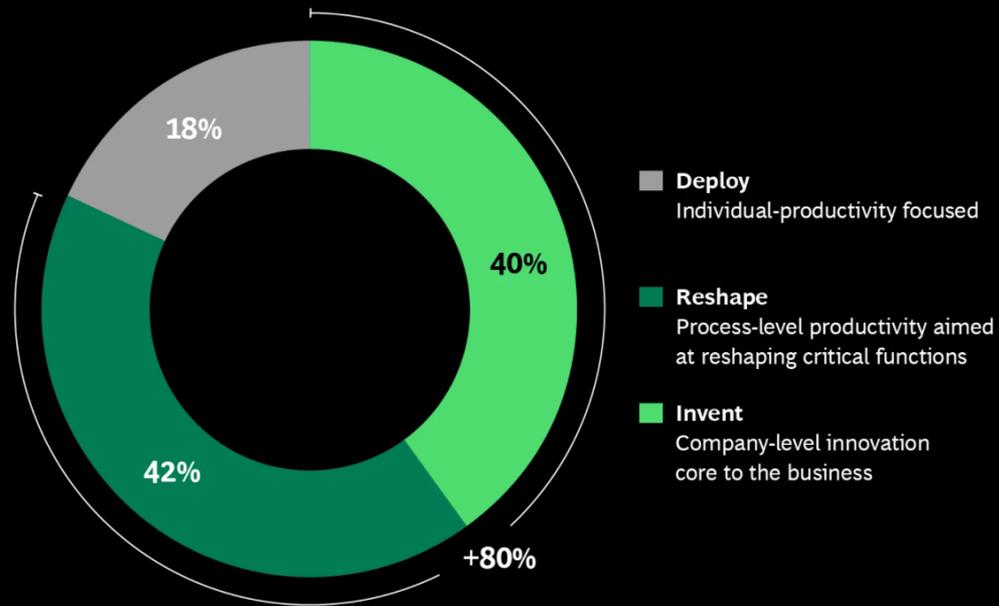


**Reshape** critical functions for 30% to 50% enhancement in efficiency and effectiveness



**Invent** new products and services to build long-term competitive advantage

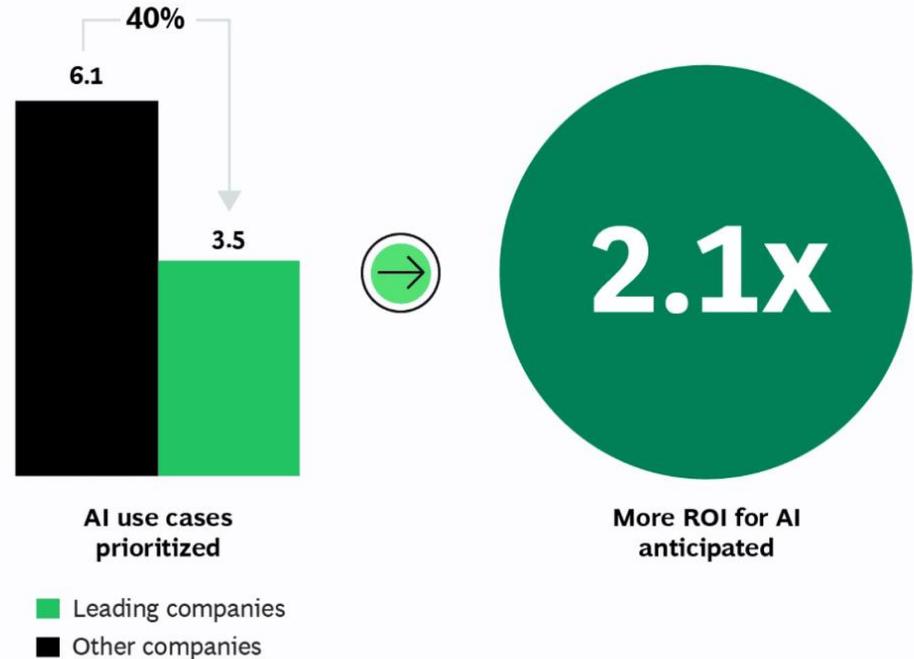
They focus +80% of their AI investments in reshaping critical functions and inventing new products and services



# BCG – Applying the Digital Acceleration Index

- BCG says in reality, companies go too broad and dilute their efforts across multiple pilots and initiatives
- Lower ROI as a result

Leading companies extract greater value by focusing their AI investments



# Organizing our thoughts – use case examples

## Customer-facing value

- Optimized search, NLP search
- Scanning and finding insights in data
- Customer chatbots & virtual assistants
- Adaptive billboards & ads
- Real-time translation

## Employee-facing value

- Content generation
- Notetaking, calendar, productivity integration
- Data merge, analytics, visualization
- Drafting documents
- Platform integration

## Application-oriented value

- Evaluating project prioritization
- Application development assistance
- Code base analysis
- Debugging
- Documentation
- Personalized coding suggestions

## Data-oriented value

- Recommendations
- Automated analytics
- Insights, analytics, visualizations
- Streamlining data-dependent business process
- Streamlined research
- Mapping solutions to problems

# So many kinds of work to think about!

## Employee agents

*Team/user level retail AI use for work*

- Should tie to employee use case
- Should save employees time
- Should translate to efficiency and speed value propositions
- May not need much development

## Internal Generative AI platforms

- General use private LLMs
- May be a customized Azure or OpenAI environment
- May be a managed service vendor
- Development project, integration project, or both

## Customer agents

- Specifically Gen AI powered
- NLP abilities for products and features
- May tie to the internal GenAI tool
- Needs a development project
- Designed to produce business value

## Other types of AI

- Other ML applications
- Machine vision
- Data extraction tools
- Algorithmic AI
- Internally developed AI products

# Use case example: IT ticketing & troubleshooting

*AI has the potential to improve IT support and service management dramatically.*

- **Classify and route tickets faster** – Reduced response times for internal IT requests
- **Instant answers from knowledge bases** – Suggest relevant articles before human intervention.
- **Predict common IT issues** – Analyze patterns in service requests.

## 📌 **Case Study: New York City's 311 AI Chatbot**

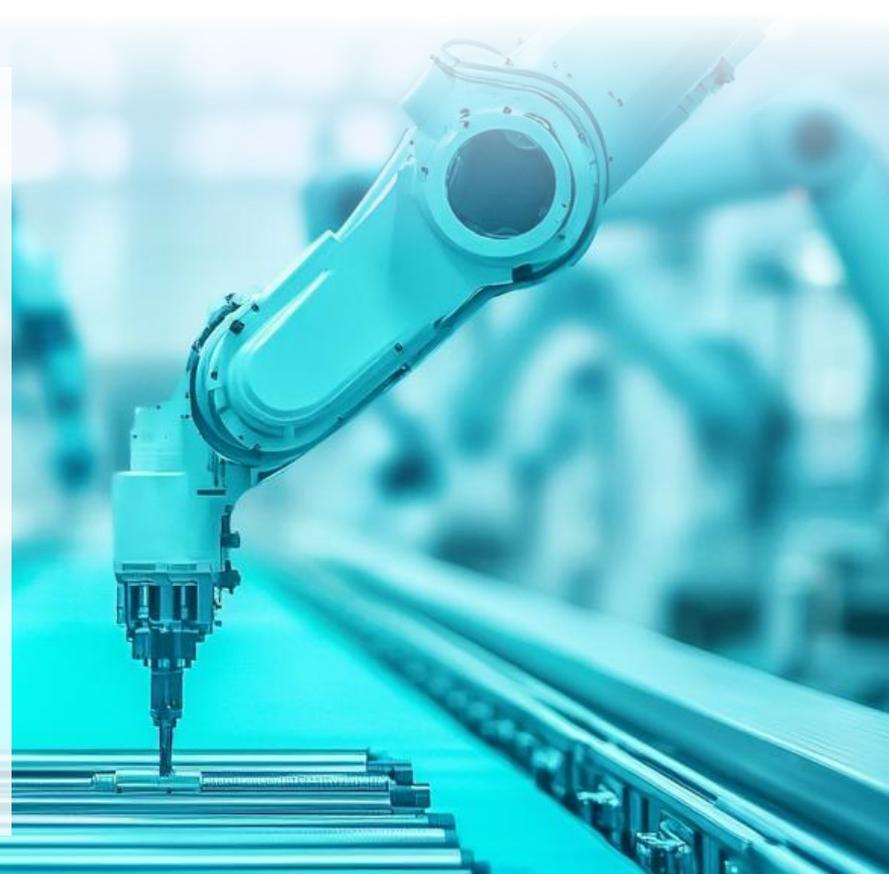
AI filters & prioritizes inquiries before human response.  
Reduced call volume, faster responses.

# Where can AI improve workflows?

## *Adding value with AI*

### **A few more examples:**

- **Helpdesk & IT support, ticketing** – Chatbots, help desk response time.
- **Project management assistance** – AI-generated meeting summaries, risk tracking, reporting.
- **Troubleshooting & service management** – Classification & knowledge base improvement. SLA automation, agent actions.
- **Documents, compliance, reporting** – Contract analysis, FOIA request processing.



# Project Documentation & Reporting

- **Meeting minutes, summarization** – AI can extract action items from project discussions.
- **Documentation assistance** – AI can help format reports for audits.
- **Automation for policy & risk assessments** – AI can cross-reference policies to detect compliance gaps.

**Example:** AI tools like **Microsoft Copilot** and **Google Gemini** are being widely used already.

# Non-Negotiables

- **We start with business case / outcome** – We have to start here, always. What is it we want to achieve in our team, unit, or organization?
- **Data governance & stewardship** – Fundamental for any base of value which relies on organizational data, which our AI practices will.
- **Understand our use cases** – Think of this as a portfolio hierarchy. We must define strategic use cases as well as individual role use cases, and some in between.
- **Change Management** – The technology is moving quickly. Adaptive, agile work patterns and continuous learning has to be baked into our playbook.
- **Policy and governance** – Will they let us use it? Will I let them use it? How will we use it? Do my people know?
- **Adoption** – Who needs to use the AI? What education/enablement is needed?

# A template for action

- **What AI opportunity** is most relevant?
- **What's the long-term goal?**
- **What's the first small step** you can take to explore it?
- **Who do you need buy-in from?**
- **What risks or challenges** might you encounter?
- **How will you measure success?**

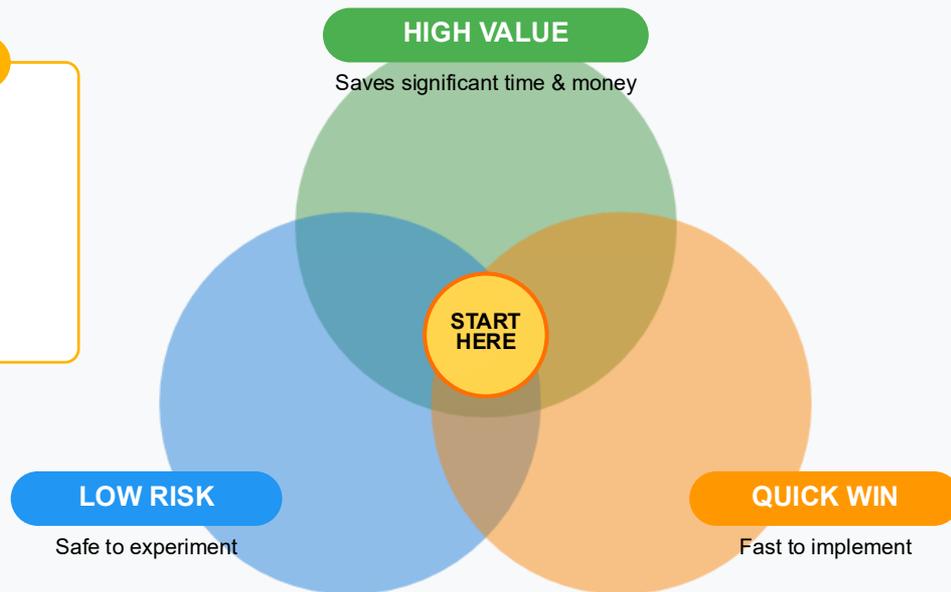


# Finding a realistic AI sweet spot

The best pilots sit at the intersection of value, safety, and speed

## Sweet Spot Examples

- Meeting summaries
- Report first drafts
- Email templates
- Process documentation
- Training materials
- FAQ responses



*"Start where you're frustrated, not where you're regulated"*

# Questions to ask yourself

- What repetitive tasks take up most of your team's time?
- Where do errors or inconsistencies frequently occur?
- Which processes could benefit from faster information retrieval?
- What requires accuracy that only humans can verify?
- Where would mistakes have serious consequences?

***Focus on augmenting human capability, not replacing human judgment***

# Thinking about AI opportunities

## Natural Fit Areas

*Safe to Start Today*

Information  
Search

Document  
Creation

Process  
Documentation

Reporting  
Automation

Email &  
Communication

Training  
Materials

## Proceed with Caution

*Requires Human Review*

**Technical  
Specifications**  
*Always verify*

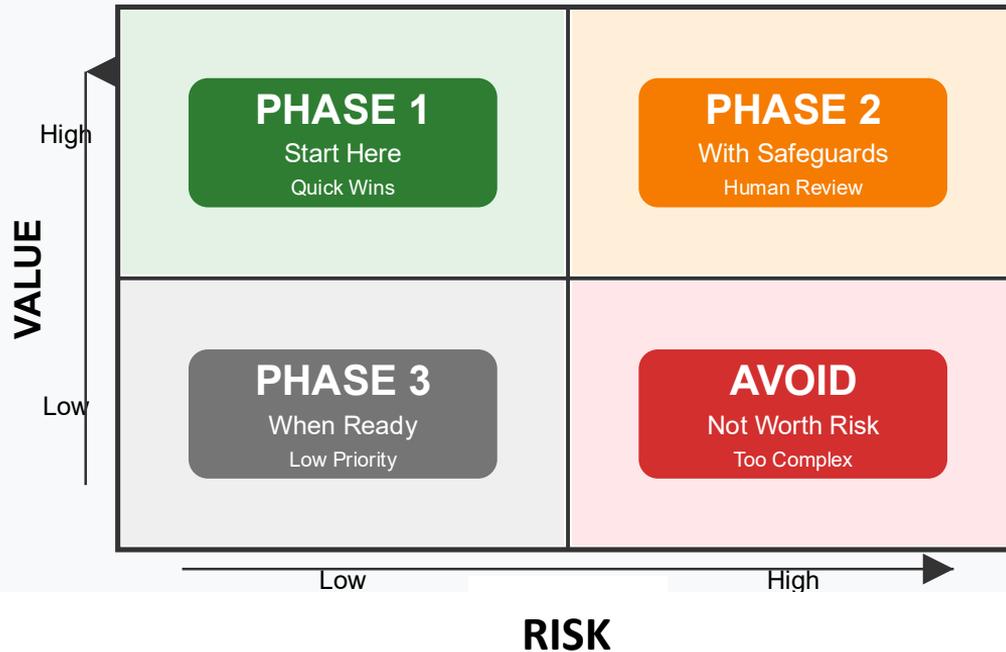
**Compliance  
Documents**  
*Legal review needed*

**Customer-Facing  
Communications**  
*Brand consistency*

**Financial  
Analysis**  
*Double-check numbers*

# Prioritizing AI opportunities

Start in the high-value, low-risk quadrant



# Set up pilots and use cases for success

## DO THIS

- ✓ **Start with volunteers**  
Find enthusiasts who want to experiment
- ✓ **Provide clear guidelines**  
Set boundaries for safe experimentation
- ✓ **Celebrate small wins**  
Share successes to build momentum
- ✓ **Build internal community**  
Create forums for sharing learnings
- ✓ **Document everything**  
Track what works and what doesn't

## AVOID THIS

- ✗ **Don't mandate adoption**  
Force creates resistance and fear
- ✗ **Don't start mission-critical**  
High stakes increase failure risk
- ✗ **Don't skip human review**  
AI outputs always need validation
- ✗ **Don't ignore user feedback**  
Early concerns prevent big problems
- ✗ **Don't rush to scale**  
Prove value before expanding

**Remember: AI adoption is a journey, not a race**  
Each team moves at their own pace - and that's okay

# Set up pilots and use cases for success

## SAMPLE

### IDENTIFY

- Map repetitive tasks
- Estimate time spent
- Check data quality
- Find pain points

Week: 1-2 ?

### EVALUATE

- Score value/risk/speed
- Select 1-2 pilots
- Define success metrics
- Get IT partnership

Week: 3-4 ?

### PILOT

- Start with individual tools
- Document process
- Gather feedback
- Measure results

Week: 5-8 ?

### DECIDE

- Review results
- Plan next steps
- Share learnings
- Scale or pivot

Week: 9-10 ?

# Back to our whiteboard

Starting with our automation ideas from yesterday, can we brainstorm a few workflows? Can you identify:

- **Repetitive or manual process** in your department.
- Where **AI could assist (e.g., automation, recommendations, analysis)**.
- **Risks & challenges** of applying AI to this process.

Overall goals / outcomes for AI

AI Use Cases

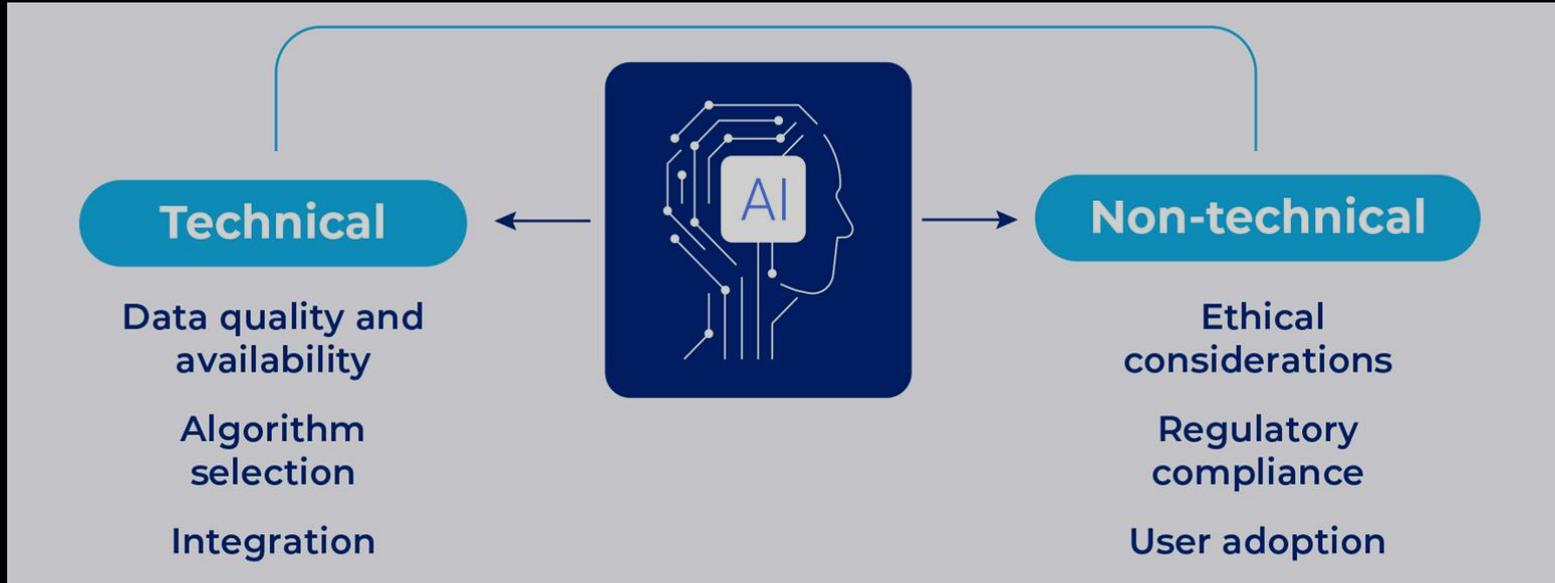
AI LITERACY

Technical considerations

Human factors

Tools and stacks

# Types of AI Challenges



# It's only as good as the data

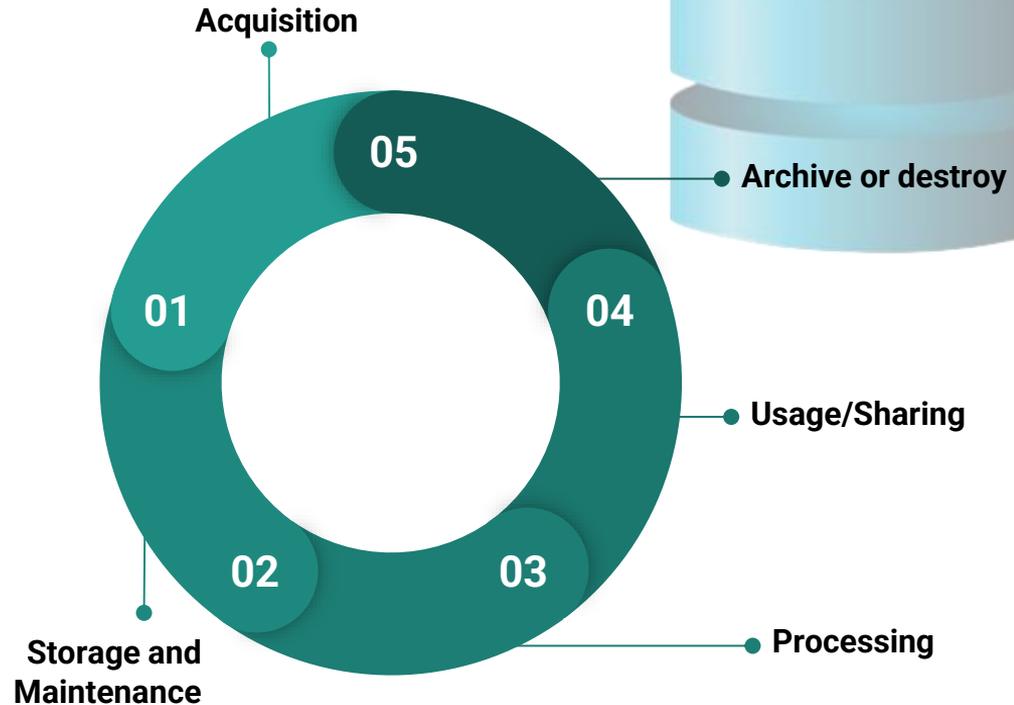
*How do organizations ensure quality, secure data?"*

**Public tools address some, not all, of these concerns. These statements assume a decision to adopt a private, integrated enterprise AI environment. Does that apply to your situation?**



# The Importance of Data Governance

- Ensures **Data Quality** and Integrity
- Facilitates **Compliance** with Regulations
- Enhances **Data Security**
- Supports **Ethical AI Use**
- Enables Effective **Data Management**
- Builds **Stakeholder Trust**
- Promotes Data **Transparency** and Accountability



# It's only as good as the data

*How do organizations ensure quality, secure data?"*

**Public tools address some, not all, of these concerns.**

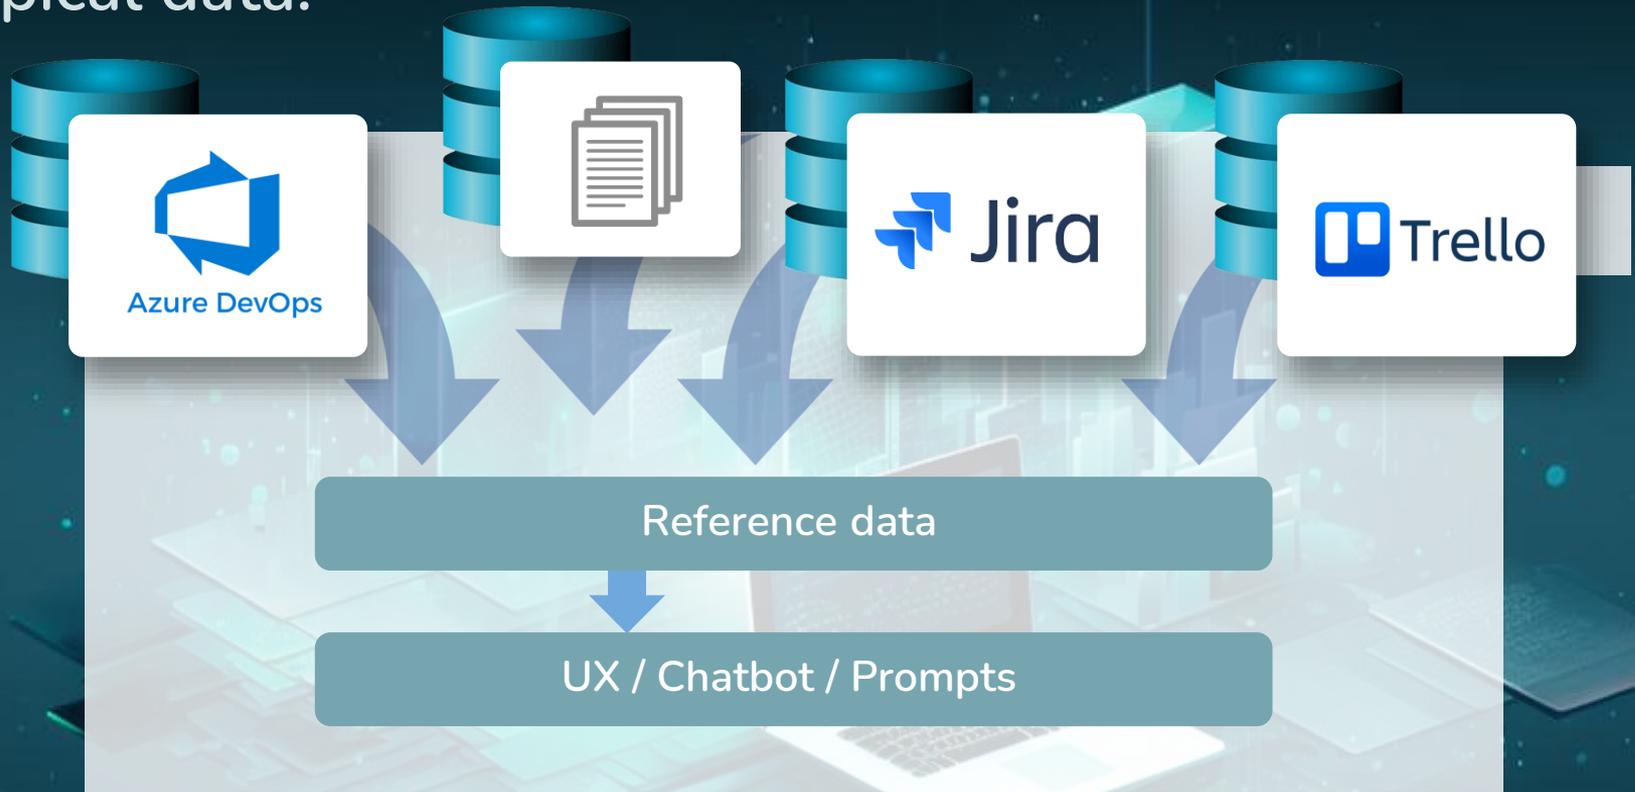
**The following statements assume a decision to adopt a private, integrated enterprise AI environment. Does that apply to your situation?**

**Either way, these statements are true:**

- Many have data silos that limit AI's effectiveness.
- AI works best when trained on high-quality, structured data.
- Data governance & privacy must be in place before deploying AI models.
- Leaders must balance data security with insight/utility



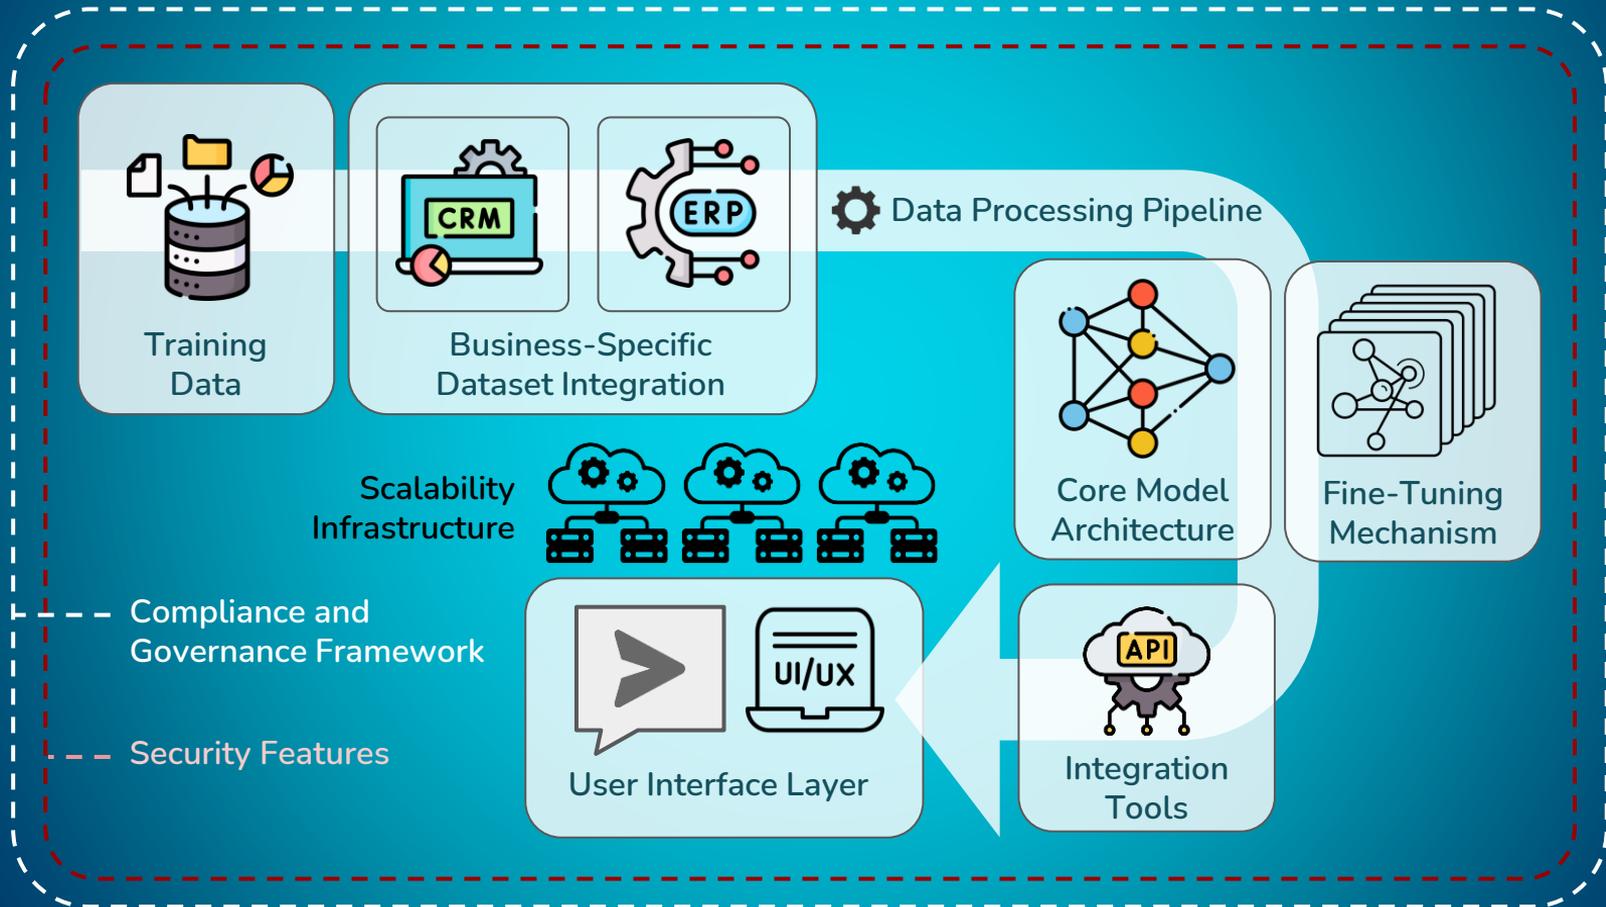
# Typical data:



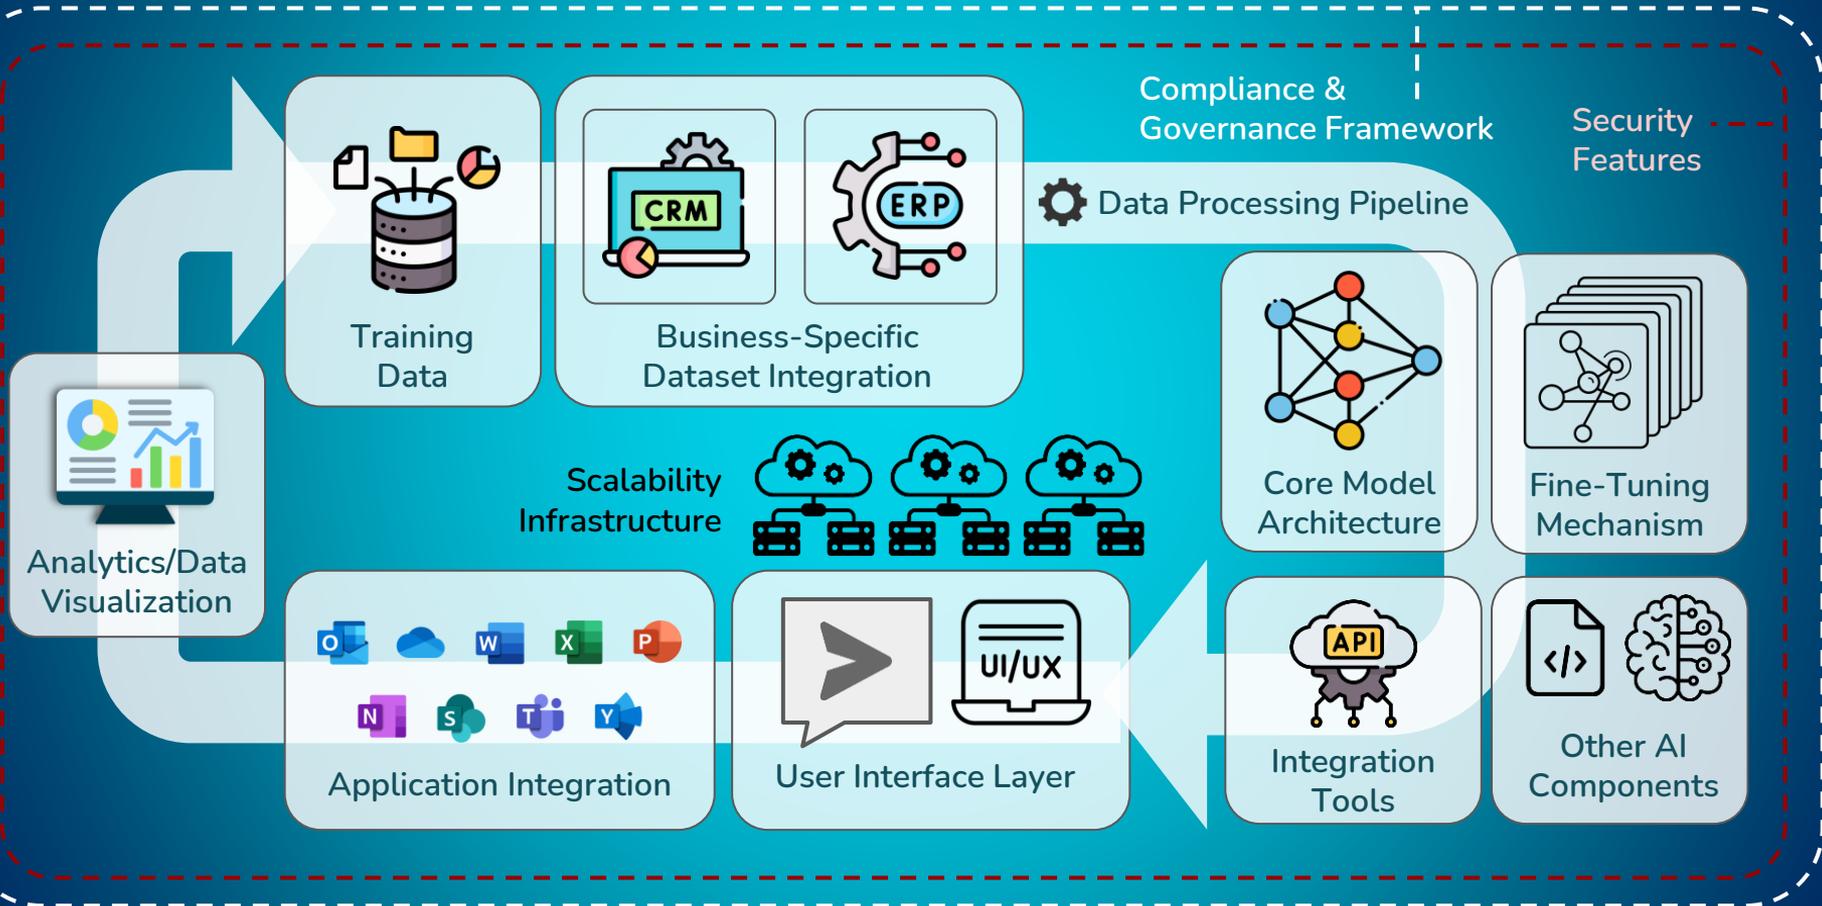
# Benefits of a internal generative AI tools

- **Data Privacy**
- **Security**
- **Defined ethics/bias behavior**
- **Legal compliance**
- **Reduced liability**
- **Potentially reduced or eliminated errors**
- **Customization**

# Private LLM Core Components



# Additional Considerations & Broader Strategy

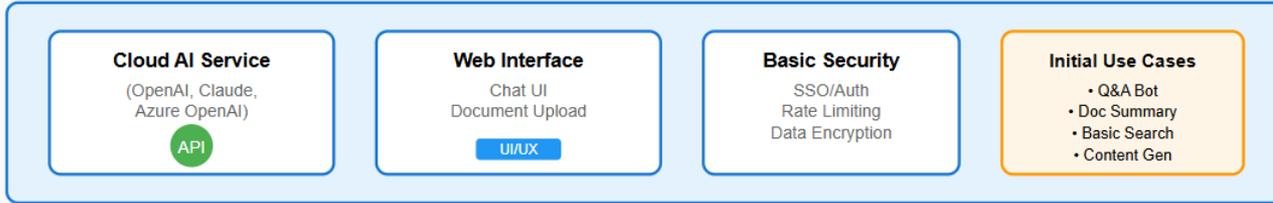


# Private LLM solution range

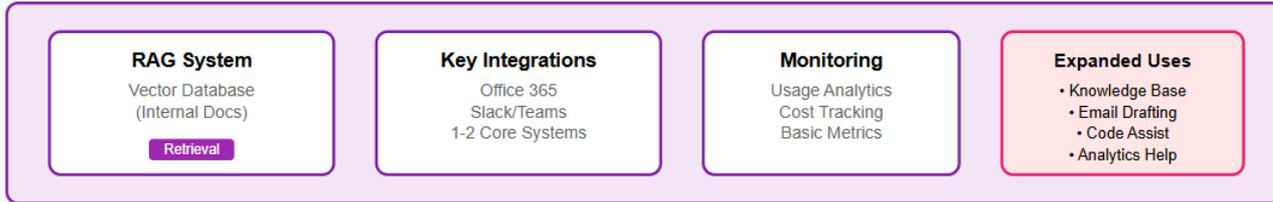
Consideration	Lightweight Solution	Enterprise-Scale Solution
<b>Costs</b>	<b>\$10,000 to \$50,000</b>	<b>\$100,000 to \$500,000 or more</b>
<b>Advantages</b>	<ul style="list-style-type: none"><li>• Easier to implement</li><li>• Less resources</li><li>• Smaller teams, less complexity</li></ul>	More features, Better integration More support Training resources
<b>Disadvantages</b>	<ul style="list-style-type: none"><li>• Lack of features</li><li>• More manual work</li><li>• Less integration</li></ul>	Investment Complexity Needs more (machine) training
<b>Timeframes</b>	<b>1 to 3 months</b>	<b>6 to 18 months or more</b>
<b>Examples</b>	<ul style="list-style-type: none"><li>• Huggingface</li><li>• Ollama</li></ul>	C3AI NVidia Nemo Thomson Reuters' Legal Tracker

# Enterprise AI Architecture, Common Implementation Pattern

## Phase 1: Foundation (6-12 months)



## Phase 2: Integration (12-18 months)



## Phase 3: Advanced (18+ months) - Less Common



### Key Reality Check

- 80% of enterprises stop at Phase 1 with basic chatbot functionality
- Integration complexity and data quality issues are the main barriers to Phase 2 | ROI must be proven in Phase 1 before additional investment

Overall goals / outcomes for AI

AI Use Cases

AI LITERACY

Technical considerations

Human factors

Adoption, skills, governance, use

The most important thing a leader can do to enable AI adoption...

**What is the most important thing a team needs in order to use AI effectively?**



The most important thing a leader can do to enable AI adoption...

**Clear communication and guidance.**



# The most important thing a leader can do to enable AI adoption...

## Clear communication and guidance.

- **What's the org's vision for AI here?**
- **What's the plan?**
- **What's expected of us?**
- **What tools we can use**
- **What's OK, and what's not?**
- **Where can I find the skills, resources, and guidance I need?**

# The most important thing a leader can do to enable AI adoption...

## **Clear communication and guidance.**

- In short, what's the official policy here, and what are the acceptable use guidelines – including resources – that I need to feel confident I'm using AI in the right way.

# Poll:

## What are your biggest AI concerns?

*(Check all that apply)*

- Missing an opportunity to use AI for positive impact
- Security & data privacy risks
- Compliance & legal challenges
- Misinformation & accuracy issues
- Workforce displacement concerns
- Lack of leadership or strategy for AI adoption
- Something else

# We can enable effective AI use, but...

There are some hard conversations. Emergent tech, big change initiatives assume some enterprise prerequisites:

- Good leadership?
- Agility?
- Safe, high performance culture?
- Good portfolio management practices?
- High innovation performance?
- Tech dependencies?



# Adoption considerations you can help with

*To leverage AI and mitigate risks, HR teams should help our organizations navigate acceptable use guidelines, resource availability, and training/enablement.*

- **Collaborative tool, not labor replacement.** *(with exceptions)*
- **Experimentation** for solution development
- You need **guardrails**—privacy, security, compliance
- Think about AI investments as **R&D**

# Teams & roles

**What team do you need?**



# Teams & roles

**Example:** standing up an internal LLM tool



Function	Stakeholder
Data Science	Lead Data Scientist
Data Engineering	Data Engineer
Software Development	Integration Developer Frontend Developer
Quality Assurance	QA Engineer
Security	Security Analyst
Compliance	Compliance Officer
Change Management	Change Manager Training Coordinator
User Experience	UX Designer

# Change management & buy-In

## Why teams may resist adoption:

- **Job displacement** – Worry AI will replace roles
- **Unclear AI policies** – Unsure how they are allowed to use it
- **Lack of training & understanding** – Don't know how to integrate it
- **Concerns over accuracy & trust** – Hallucinations make us hesitant to rely on it



# Change management & buy-In

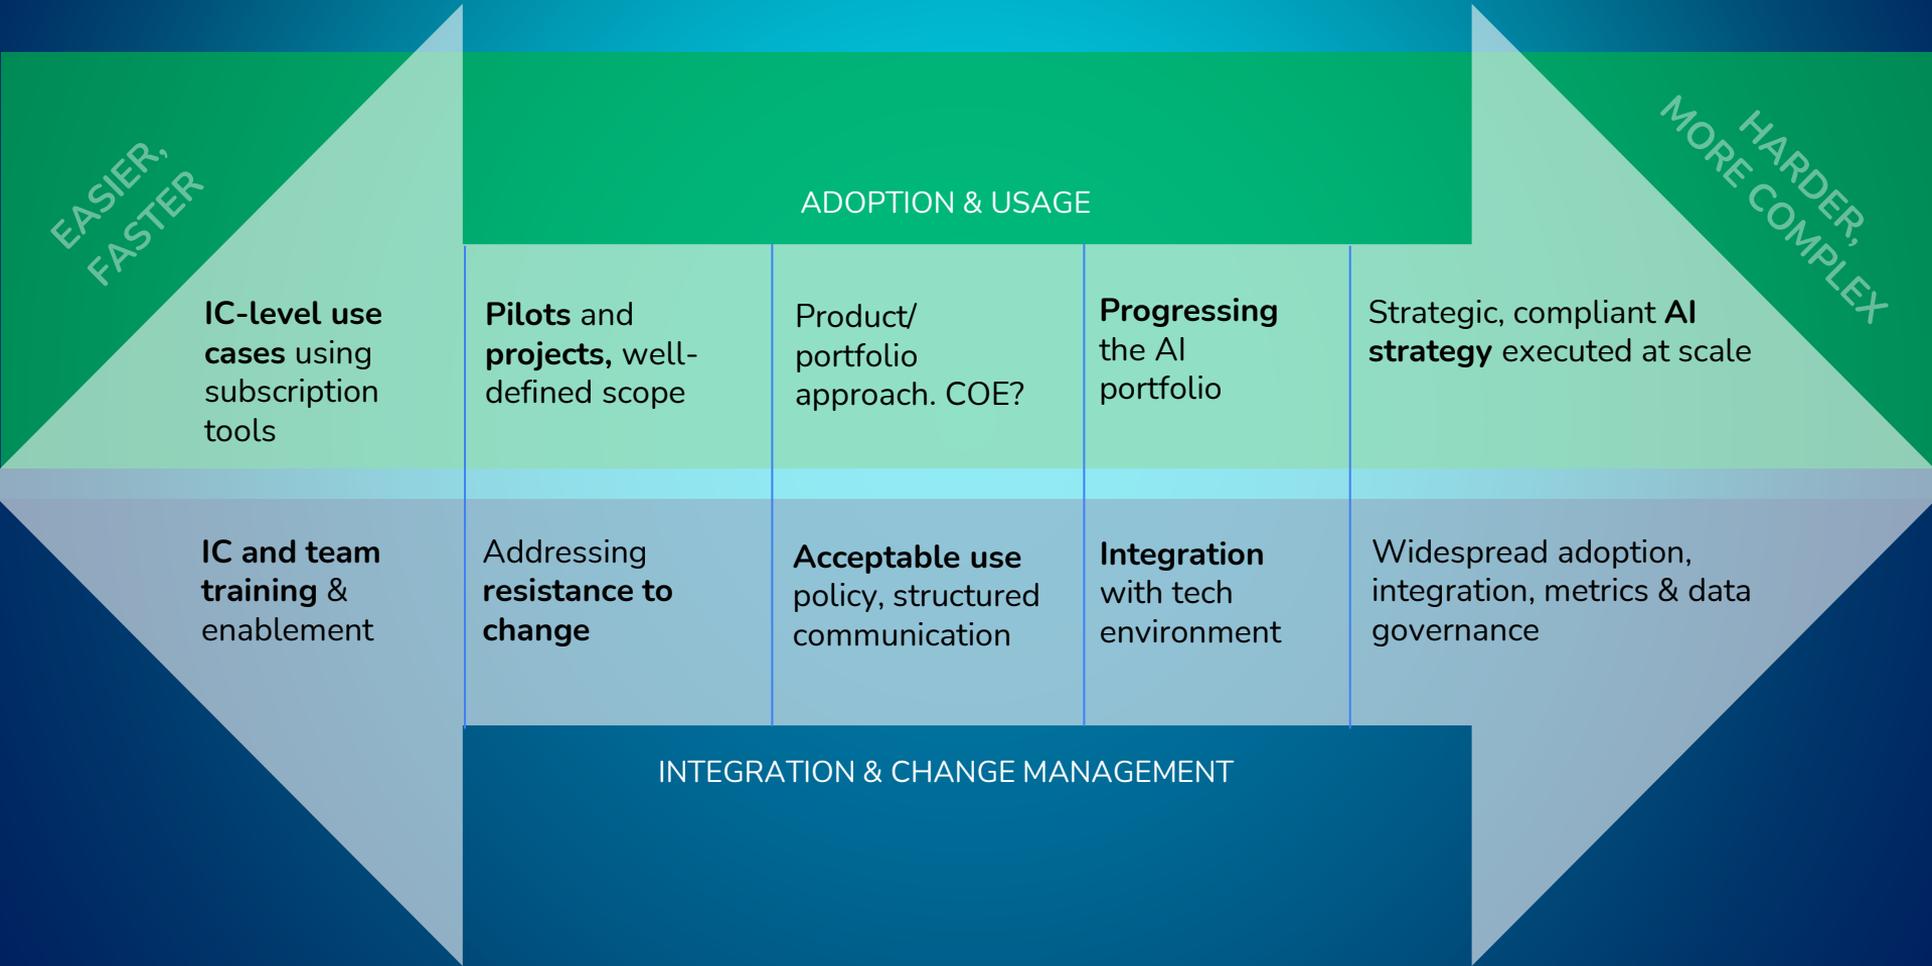
## Why teams may resist adoption:

- **Job displacement** – Worry AI will replace roles
- **Unclear AI policies** – Unsure how they are allowed to use it
- **Lack of training & understanding** – Don't know how to integrate it
- **Concerns over accuracy & trust** – Hallucinations make us hesitant to rely on it

## We must:

- Communicate the role of AI here: An enhancement, or a replacement?
- Develop clear AI guidelines for staff.
- Provide hands-on AI training to build confidence.
- Encourage early experimentation in low-risk areas.

# What do we need to get done?



We had this discussion:

What are the challenges to your AI vision? Examples:

Org has  
difficulty with  
change

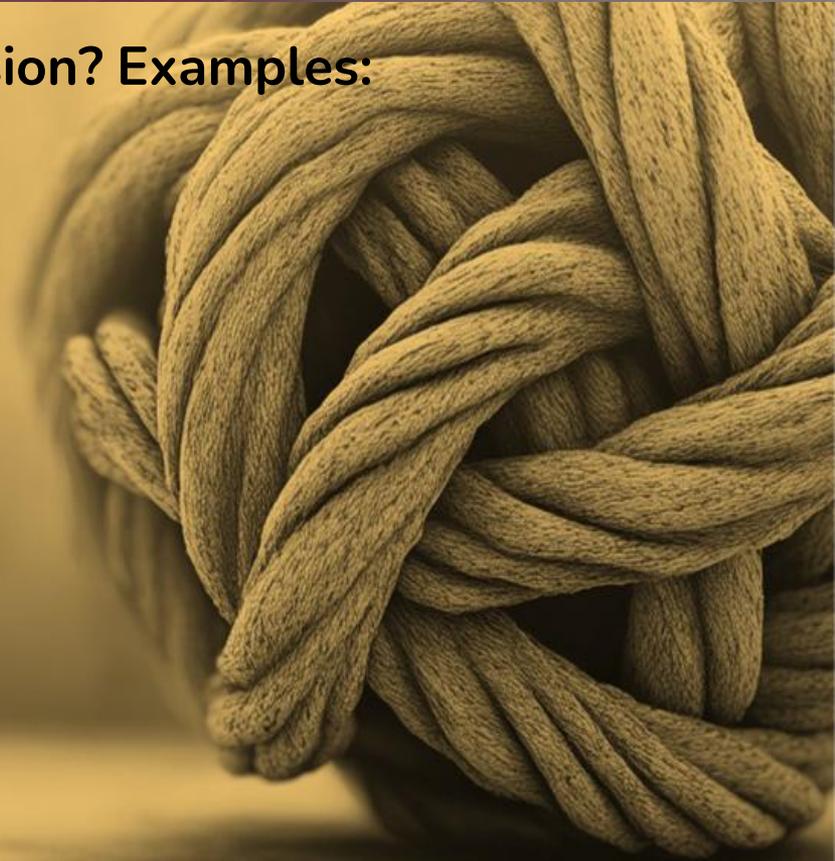
Tech is too  
hard

Regulatory  
requirements  
too hard

Something to  
do with money  
or cost

Skills or people  
not available

What else?



# How to address obstacles?

## Common barriers:

- Leadership hesitation
- Tools
- Strategy not defined
- Unclear guidance
- Guidance not communicated well
- Procurement
- Employee skills
- Policy/compliance uncertainty
- Security, data privacy
- Misinformation, hallucinations
- Trust & transparency

# Addressing some employee concerns

- **Communicate that automation is for efficiency, not elimination** – Should allow staff to focus on higher-value work. How do we define that value?
- **Reskill & upskill** – Teams need AI literacy & skills training
- **Unions** – AI use sometimes must align with labor policies & collective bargaining agreements, if applicable.
- **Ethical use in HR / hiring** – Watch for bias in recruitment or employee analysis.



# Learning & training can enable change

- **AI literacy training** – So staff broadly understand and align on AI use, tools, & limitations.
- **Experiment** – Encourage low-risk pilot projects.
- **Communicate & explain policy** – Figure out privacy & compliance. Communicate the guidance on use so people can implement.
- **Daily AI workflows** – Encourage & enable the tactical use, Recognize & socialize the successful examples.

# Tools and acceptable use

**We immediately run into governance needs.** Publicly available tools are powerful, but not compatible with many organizations' policy. How to leverage generative AI tools acceptably?

- **Data risks**
- **IP concerns**
- **Corporate restrictions**
- **Privacy & security**
- **Risk & compliance**
- **Integration with current controls**
- **Private LLM solutions**

# Tools and acceptable use

EASY &  
CHEAP

HARD &  
EXPENSIVE

**Individual accounts** on public services (i.e. ChatGPT 4)

**Enterprise subscriptions** (CoPilot, ChatgPT Claude)

**API Integration:** Internal projects on name-brand services

**Custom/internal development:** Starts involving all the things

**Private LLMs**

# AI Investments

## *Key considerations*

- **Enterprise AI vs. Public AI Models** – Should your agency **buy, build, or integrate** AI tools?
- **ROI & Efficiency Gains** – What **time & cost savings** justify AI procurement?
- **Pilot Programs Before Full Deployment** – Start small, then **scale successful AI applications**.
- **Evaluating Vendor AI Claims** – Ensure vendors **demonstrate transparency, security, and bias mitigation**.

# Can we treat AI as an R&D investment?

*It's an investment in new long-term capabilities, not just immediate efficiency gains.*

## Reasons why:

- AI adoption, like any emergent technology, **requires a multi-year strategy to realize substantial value.**
- AI models and tools **improve over time**, requiring ongoing refinement and training.
- Measuring ROI should include a range of value propositions: quick **efficiency gains, service improvement, risk reduction, etc.**

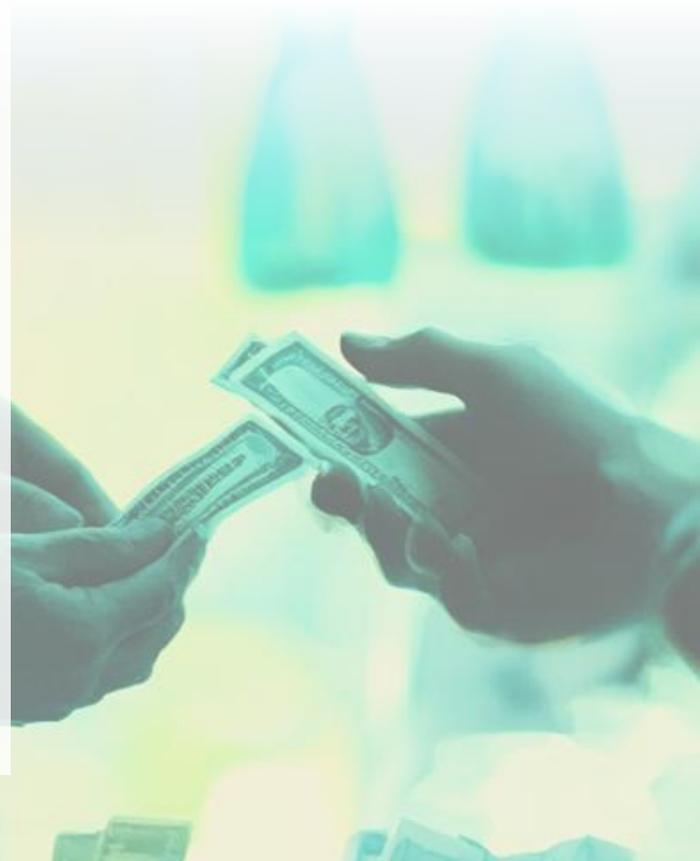


# Procurement challenges

- **Strict RFP/vetting processes** – AI adoption can require vendor vetting & regulatory approval.
- **Enterprise AI licensing, or open-source?** – Limited state-approved vendors
- **Budget & cost** – Justify with clear ROI
- **Data residency & cloud restrictions** – AI models must comply with NY state IT infrastructure policies
- **Regulatory & Ethics Reviews** – AI procurement often requires additional approval layers due to bias risks.

## Case Study: AI procurement in state agencies

*(Details in notes)*



# Metrics, OKRs, ROI

## Examples

- **Cost savings & efficiency Gains** – Reduction in labor hours for manual processes.
- **Response time improvements** – Faster IT ticket resolution, help desk response, or document processing.
- **Error reduction & accuracy** – Fewer compliance violations, data entry mistakes, or risk misclassifications.
- **Satisfaction & service delivery** – Fewer complaints, faster responses, and greater transparency.
- **New products or value streams**

**Example:** AI-driven **help desk automation** in a city government led to **30% faster response times** and **20% reduction in manual processing costs**.

# DISCUSSION: Adoption by Experimentation

*How to encourage safe AI exploration?*

1. What **AI tools or functions** has your team tried or considered?
2. Are there **low-risk AI experiments** your group could pilot? Are you piloting any already?

- Hold **internal AI hackathons** for IT teams to test AI use cases.
- Develop **low-risk pilot projects** to build a playbook.
- Set **clear AI usage guidelines** to ensure compliance, ethical use.
- Give employees a **safe space to experiment**.

# Risks & governance



# Security & Data Protection

*In regulated environments...*

AI must comply with **federal/state security mandates**.

## **CONCERNS:**

- **Prompt injection & AI exploits** – AI systems can be manipulated
- **Data residency & sovereignty** – AI tools must store & process data within state-approved infrastructure.
- **AI in cybersecurity monitoring** – AI can detect threats, but also introduces new vulnerabilities.

**Example:** In 2023, AI-powered **government chatbots** leaked sensitive information due to weak controls.  
*(Details in notes)*

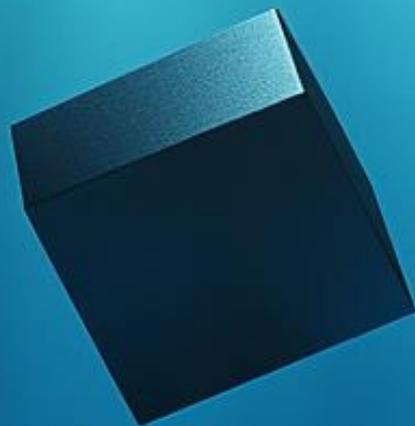


# Transparency

## *Tensions between AI and trust – automation vs. accountability*

- Depending on the audit environment, AI may need to be explainable – Black-box AI models might not be acceptable.
- In some jurisdictions public sector AI use must be auditable – Agencies must document AI decisions & provide justification.
- Automation can't eliminate human oversight.

**Example: AI in Benefits Processing** - Welfare eligibility assessments increased efficiency but faced public backlash due to opaque decision-making processes.



- Scenario testing
- Custom development: control vs. ease
- Providers/managed services – cost vs. ease
- Enterprise license assurances
- Data provenance
- Data governance practices mapped to concerns
- Rules and SLAs
- Contractual indemnification
- Training & awareness
- Pilot testing, incremental rollouts
- Continuous monitoring



## **Governance platforms:**

### *Potential tools:*

Platforms like IBM OpenPages or Microsoft Purview provide centralized tools for managing compliance, risk assessments, and governance policies.

**Co-Intelligence:** Combine automated scoring with human oversight for final decisions

## **Explainable AI (XAI):**

### *Potential tools:*

**SHAP:** machine learning explainability tool based on Shapley values from game theory.

Tools like **LIME** (Local Interpretable Model-Agnostic Explanations) complement SHAP by providing localized explanations for individual predictions

## **Stakeholder engagement:**

Involve diverse groups in designing and testing the system to identify potential risks early.

## **Transparency & fairness frameworks:**

### *Examples:*

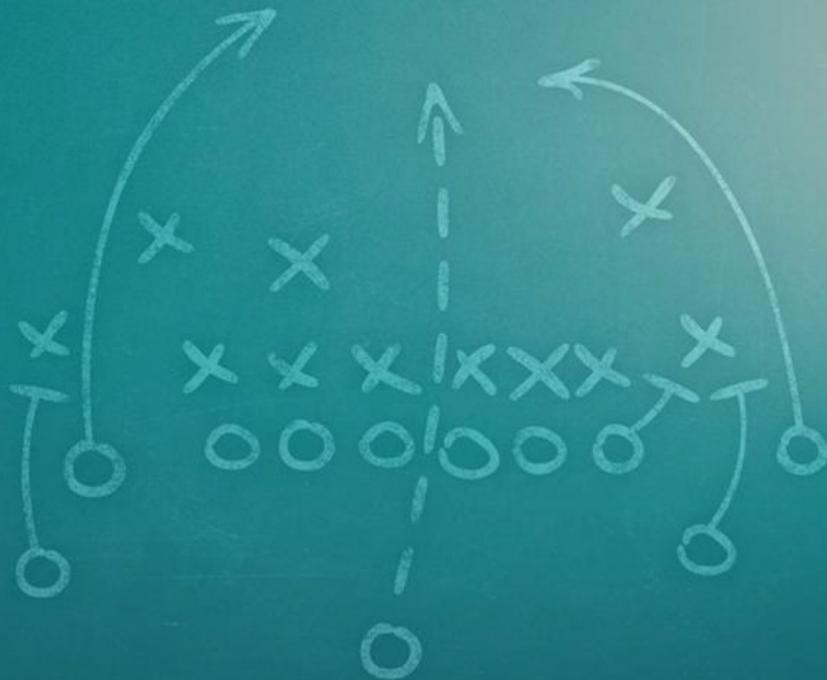
**OECD AI Principles** or **NIST AI Risk Management Framework** responsible AI development and use.

# AI Governance & Use Guidelines

- **Clear AI policies** – Define where AI can and cannot be used
- **Human oversight** – AI-generated recommendations must be reviewed by humans.
- **Auditability & explainability** – How and where do things need to be traceable for public accountability?
- **Security & compliance** – Tools must align with legal & compliance requirements.



# WHITE BOARD: Governance, policy, use guidance





Wrap up

# WHITE BOARD: action items

Let's write down personal **action items**

## Template for action:

- **What AI opportunity** is most relevant for your department?
- **What's the long-term goal?**
- **What's the first small step** you can take to explore it?
- **Who do you need buy-in from?**
- **What risks or challenges** might you encounter?
- **How will you measure success?**

**Share your action plan.**



**What's still on  
your mind?**



# Thank You!



Chris Knotts, Principal Consultant

# APPENDIX 1

Demos

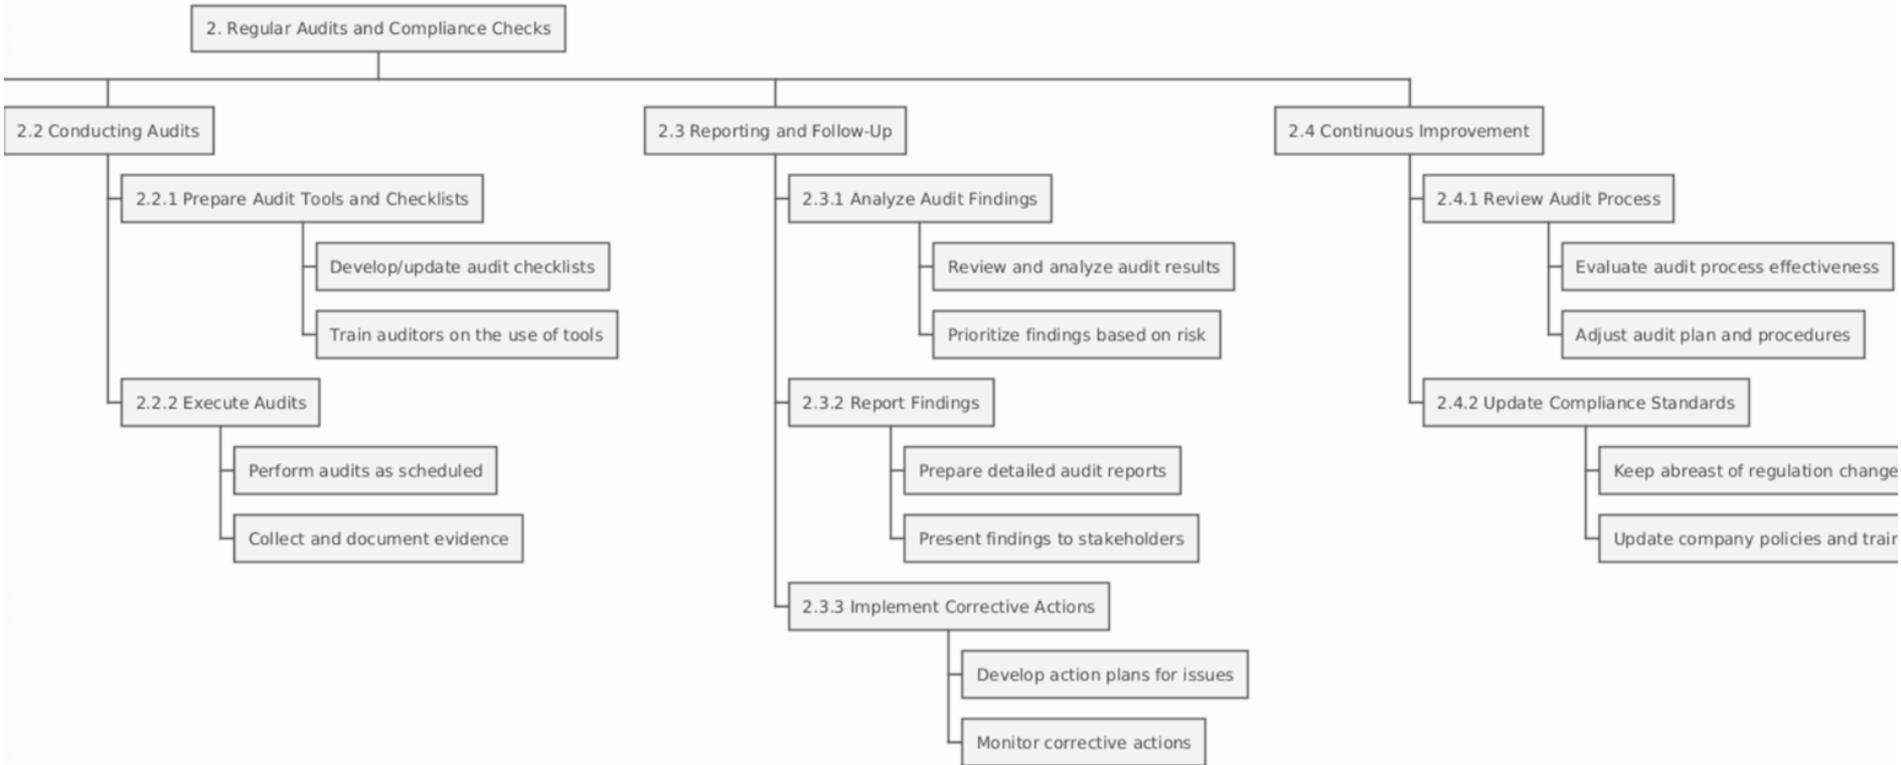
# DEMO: Image Generation



# DEMO: Project Data Exports & Imports



# DEMO: Work Decomposition & WBS



# DEMO: Ports of the World



# DEMO: Stakeholder Sentiment Analysis



# APPENDIX 2

Legality and liability considerations

# Main things

- Lack of data governance will erode trust and open up legal/audit risks
- AI has a black-box problem
- Datasets always...ALWAYS...have biases baked in - evaluate the source (IP, demographic data, domain bias, source bias, cultural bias)
- Legal precedents barely exist - it's the wild west. Be bold to embrace the technology, but cautious when it comes to legal consequences.
- Most generative AI tools either claim IP rights to the content created, or the content is public domain, or they don't explicitly specify
- Generative AI learns from a large corpus of training data- at least some of that data usage is being challenged by the original content creators
- Watch the space closely, it's unfolding in real time

# Many issues start with the training data



Generative AI models require a LOT of data in order to learn how to do what they do.

A lot of this data comes from open-source datasets provided for the training.

# Where Does the Data Come From? GPT models

<b>Common Crawl dataset</b>	410 billion tokens - 60% weight in the training
<b>Webtext2</b>	19 billion tokens - 22% weight in the training
<b>Books1</b>	12 billion tokens - 8% weight in the training
<b>Books2</b>	55 billion tokens - 8% weight in the training
<b>Wikipedia</b>	3 billion tokens - 3% weight in the training

From *Language Models are Few-Shot Learners* discussing GPT3 training.

# Where Does the Data Come From? Diffusion model Image Generators

## LAION dataset

- Large-scale Artificial Intelligence Open Network)
- 400 million image-caption pairs

## Some big problems

- Private data
- Sensitive data
- Inappropriate images
- In some cases illegal images
- Images scraped without rights usage

# Considerations

Free doesn't mean "no license" or that there aren't any rights associated with use.

Business users will continue to adopt generative AI at a breakneck pace because of the immediate benefits.



# IP and Copyright

- What is a derivative work?  
Precedents will likely be set based on interpretation of existing fair use doctrine.
- When is AI-generated content “sufficiently transformative?”
- One of few existing precedents stems from Google’s successful defense of scraping the content of books for its search service. They were successful. Will this precedent hold?

## Mitigation

- Businesses could be liable for willful infringement if courts find training datasets using unlicensed content isn’t fair use
- Providers of AI tools should be asked if models were trained with unprotected data
- Terms and licensing should be worked into contracts, whether for services you provide using AI, or for any AI service that you use.
- Avoid integrating providers and services into your offerings if licensing or open-source rights can’t be documented.

# US Developments

Currently, there isn't a comprehensive federal law governing AI or data privacy in the United States, resulting in a fragmented landscape of regulations that could affect AI. This includes a variety of state-level data privacy and AI-specific regulations, as well as the regulatory authority of the FTC and similar consumer protection agencies.

## Mitigation

- Adopt internal policies aligned with stringent state regulations.
- Stay updated on and adapt to changing state laws.
- Set up a compliance team for regulatory navigation.
- Exceed baseline standards to enhance consumer trust.

# Patents - US

- The US Patent and Trademark Office stipulates that only humans can “invent” something
- USPTO has clarified that works created by AI cannot be patented or copyrighted
- USPTO says works created with the assistance of AI **CAN** be patented or copyrighted
- Where to draw the line?



# GDPR and EU Developments

- Italian DPA's temporary ban on ChatGPT led to OpenAI's privacy updates, restoring service in Italy.
- GDPR roles for AI use involve distinguishing between data controllers and processors.
- Key GDPR challenges include lawful data processing, transparency, data collection for AI training, and conducting DPIAs.
- Unresolved issues include dataset transparency ("black box"), consent validation, and managing data subjects' rights in AI datasets.
- Upcoming EU AI Regulation may introduce risk-based scrutiny for generative AI applications.

## Mitigation

- Testing for compliance
- Feedback mechanisms for accountability
- Figuring out explainability/transparency for audits and reporting
- Proactively protecting against unresolved legal rulings

# Mitigation - General Considerations

- Thorough due diligence - assurance from service providers and data sources
- Robust monitoring
- Work indemnification into contracting so you're protected
- Thorough employee training - just like security hygiene or cybersecurity awareness training we all have to take.
- Proactively get creative for strategies around risk-free training data



# Overarching Mitigation

- **Prioritize ethical AI principles** (transparency, fairness, responsibility, accountability, reliability) and establish a practice to operationalize these in product development.
- **Address generative AI risks by ensuring accuracy** (train AI on own data, communicate uncertainty), safety (mitigate bias, protect privacy), honesty (respect data provenance, consent), empowerment (AI as an assistant, respect content creators), and sustainability (minimize model size for lower energy use).
- **For AI integration, use zero-party or first-party data** for training to ensure accuracy and trust. Regularly update and curate training data to remove biases and inaccuracies.
- **Maintain human oversight** to review AI outputs for accuracy and bias, considering generative AI as a tool to augment human capabilities rather than replace them.
- **Implement continuous testing and feedback mechanisms** for AI systems, including ethical AI training for staff and feedback channels for reporting concerns, to navigate the evolving landscape of generative AI responsibly.

# APPENDIX 3

Labor impact research

# Evaluating & Planning for GenAI

What's the use case

1

- *Start with high-potential use cases*
- *Apply the lens of AI strengths and weaknesses*

What's the business case

2

- *Put need above tool*
- *What really needs to be done?*

What's the technical/tooling solution

3

- *Scope the solution*
- *Focus on meeting the need*
- *Figure out the costs*

What are policy or compliance hurdles?

4

- *How can you get there?*
- *Whose sponsorship do you need?*
- *How can you affect change?*

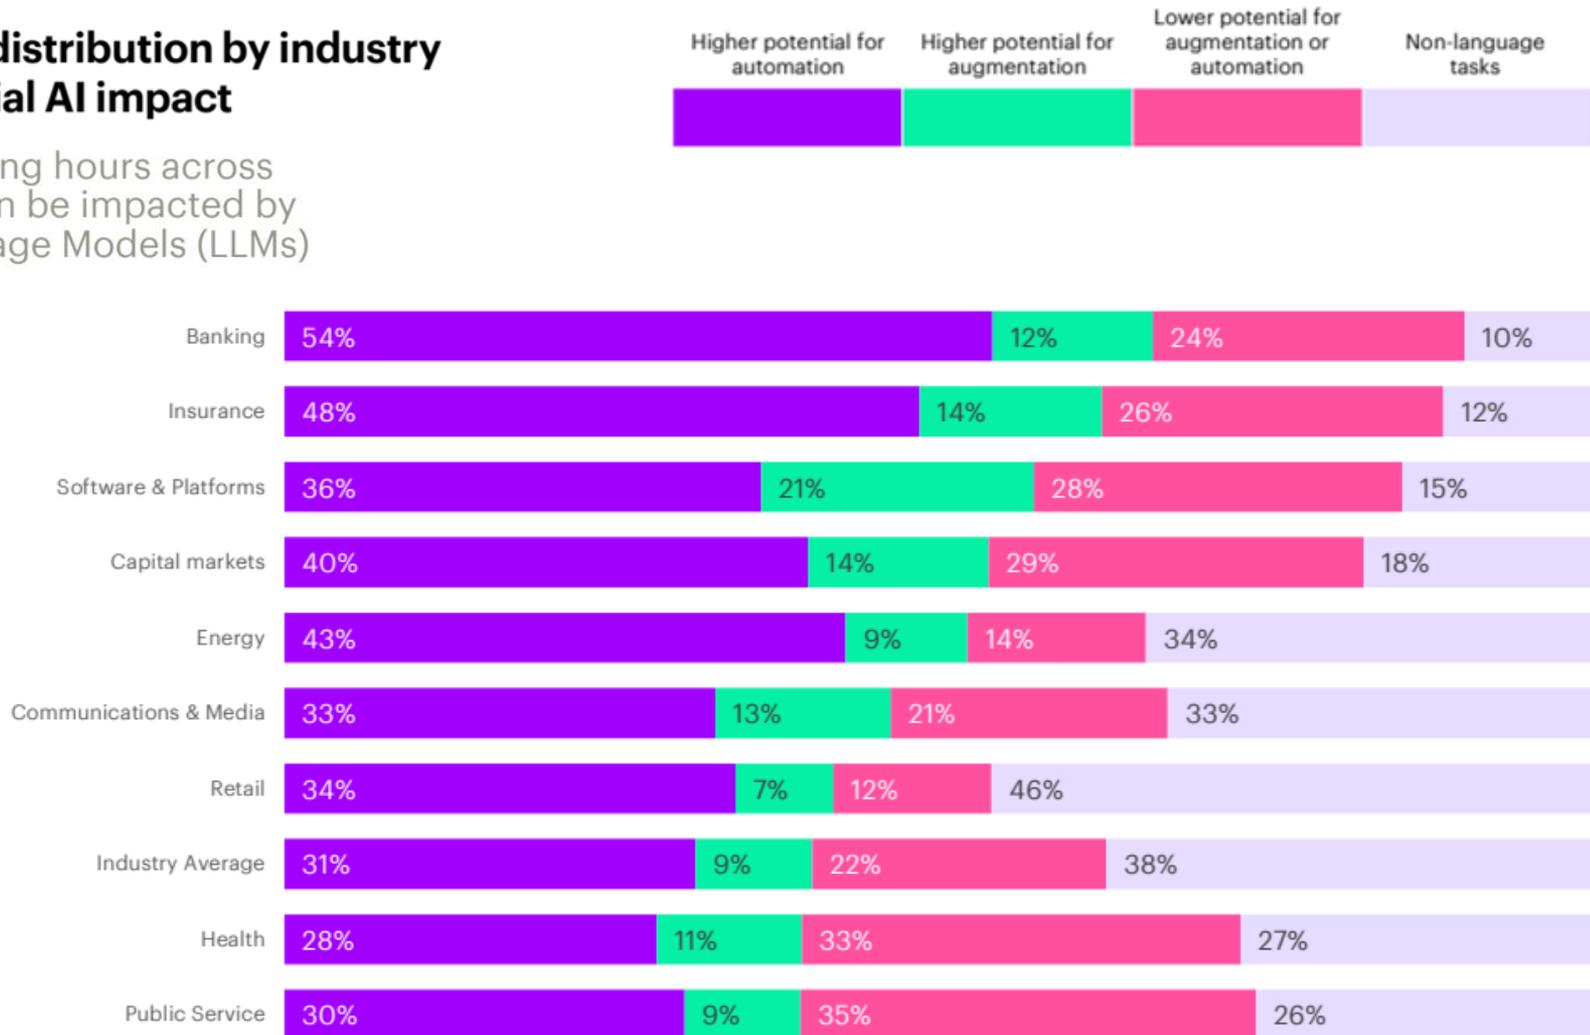
What's the implementation plan?

5

- *Revisit business case*
- *What are the costs?*
- *What's the ROI?*
- *How to measure?*

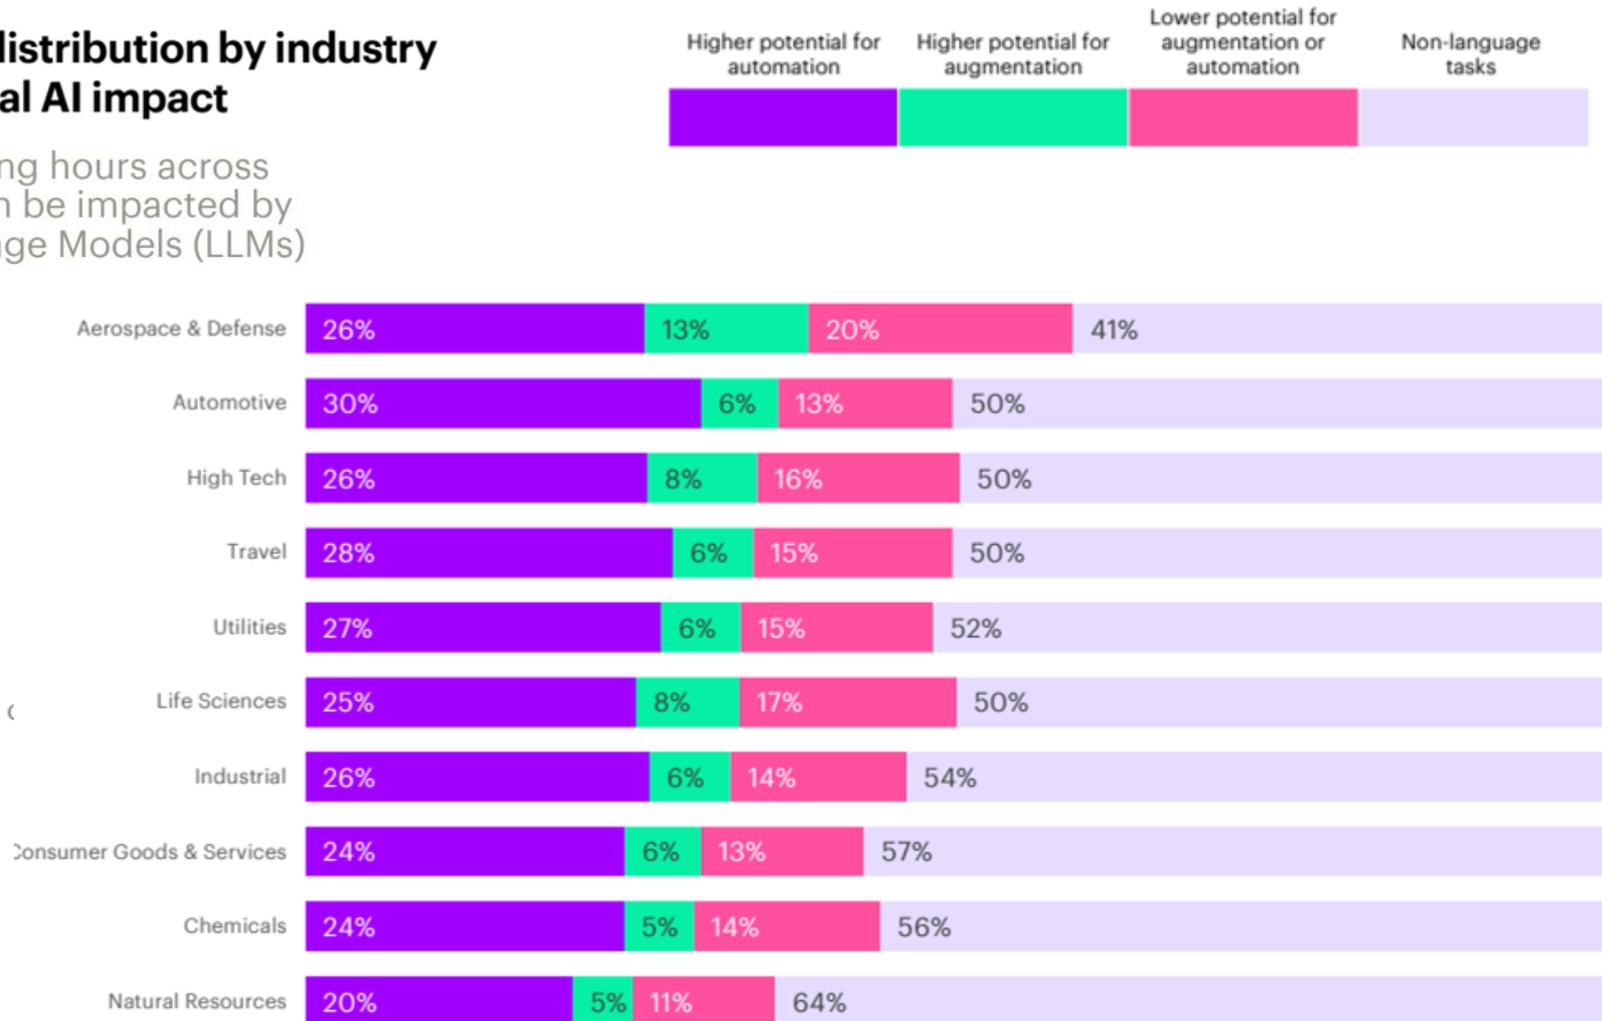
# Work time distribution by industry and potential AI impact

40% of working hours across industries can be impacted by Large Language Models (LLMs)



# Work time distribution by industry and potential AI impact

40% of working hours across industries can be impacted by Large Language Models (LLMs)



# APPENDIX 4

## Prompting

# What is a “dialectic approach?”

We primarily interact with LLMs using what we'll call a “dialectic approach,” meaning that we should always approach interaction with an LLM as a conversation, anticipating that we may not get what we need from the AI tool in just one exchange – it will likely require a series of exchanges so we can establish context, help the AI learn what we need, and refine the direction of the conversation as we go. LLM tools don't just use machine learning to power their core features - they also learn as they go, adapting to our feedback in real time.

# Prompting basics in plain English

- Initiate with comprehensive context to guide the AI's understanding and responses.
- Provide examples
- Employ iterative refinement to hone in on precise information, leveraging the AI's initial outputs to refine subsequent queries.
- Establish feedback loops to improve the AI's accuracy and relevance over time.
- Include details in your query to get more relevant answers
- Ask the model to adopt a persona
- Use delimiters to clearly indicate distinct parts of the input
- Specify the steps required to complete a task
- Specify the desired length of the output

# Bad Prompts: Examples

**Vague Query:** Tell me about agile project management.

- Lacks specificity and context.

**No Persona Adoption:** How do I manage projects?

- Doesn't specify the AI to adopt the role of an experienced project manager.

**Unclear Delimiters:** What are the risks in project management and how to mitigate them, also tell me about team management.

- Combines multiple questions without clear separation.

**No Steps Specified:** How do I implement an ERP system?

Lacks detail on the desired process or steps.

**No Examples Provided:** How do I improve team efficiency?

- Doesn't provide context or examples to tailor the response.

**No Output Length Specified:** Give me a project report.

- Doesn't specify desired length or detail for the report.

# Good Prompts: Examples

**Detailed Query:** "I'm an agile project manager implementing a new ERP system. What are key considerations for integrating existing customer data into the new system?"

- *Provides specific context and details.*

**Persona Adoption:** "Act as an experienced agile project manager. What strategies would you suggest for maintaining team morale during the stressful final stages of an ERP implementation?"

- *Directs the AI to adopt a specific role.*

**Clear Delimiters:** "Identify the top 5 risks in agile ERP implementation | Provide methods to mitigate each risk."

- *Clearly separates two parts of the prompt.*

**Steps Specified:** "Outline the steps for conducting a sprint retrospective in an ERP project."

- *Asks for a specific, step-by-step process.*

**Examples Provided:** "Considering challenges like scope creep and resource constraints, which I've faced in previous projects, how can I better manage project scope in my current ERP implementation?"

- *Includes examples for context.*

**Specifying Output Length:** "List in bullet points, no more than 10, the key metrics I should track during an ERP system implementation."

- *Specifies the format and length of the response.*

# The COSTAR Prompt Framework

Try using the COSTAR prompt framework:

1. **Context (C):** Provide essential background information or setting for the task. This helps the LLM understand the specific scenario or domain it is dealing with, leading to more relevant responses.
2. **Objective (O):** Clearly articulate the goal or purpose of the prompt. Specify what you want the LLM to accomplish, ensuring that its focus remains on achieving this particular aim.
3. **Style (S):** Define the desired style of the response. This could range from imitating the writing style of a specific profession, like a scientist or journalist, to emulating the narrative tone of certain genres, such as formal reports or creative fiction.
4. **Tone (T):** Determine the emotional or attitudinal coloring of the response. Whether it's formal, casual, enthusiastic, or empathetic, setting the tone ensures the LLM's response aligns with the intended sentiment.
5. **Audience (A):** Identify the target audience for whom the response is intended. Tailoring the content and complexity of the LLM's response to suit the audience, such as experts, beginners, or a general readership, ensures better comprehension and engagement.
6. **Response Format (R):** Specify the format in which you want the response. This could be a list, a structured report, a JSON object, a narrative, etc. Defining the format helps in generating responses that are suitable for your subsequent use, whether it be for analysis, presentation, or further processing.

# The COSTAR Prompt Framework - Tell AI to create the prompt

## Try using the COSTAR prompt framework (continued):

1. Break down complex tasks into a sequence of simpler prompts in an interactive conversation.
2. Employ affirmative directives such as `do,' while steering clear of negative language like 'don't'.
3. Implement example-driven prompting (Use few-shot prompting).
4. Use following phrases: "Your task is" and "You MUST".
5. Always use leading words like writing "think step by step".
6. Assign a role to the model i.e. "you are an expert \_\_\_\_"
7. Repeat specific words or phrases multiple times within a prompt.
8. Try to induce Chain-of-thought (CoT) when possible, guiding the LLM to do dive in deeper to each step
9. Use output primers, which involve concluding your prompt with the beginning of the desired output. Utilize output primers by ending your prompt with the start of the anticipated response.
10. To write an essay /text /paragraph /article or any type of text that should be detailed: "Write an ultra-detailed [essay/text /paragraph] for me on [topic] in detail by adding all the information necessary".

# More Sophisticated Prompt Techniques

- System Message
- Repeating the instructions
- Priming the output
- Clear Syntax
- Few shot learning
- Chain of thought (or Auto-Generated Chain of Thought)
- Shuffle Ensemble (and choice shuffle)

# Initial Prompt: The “System Message”

The **system message** sets the stage for the interaction with the model by providing context, instructions, or relevant information tailored to the specific use case. It can outline the assistant's personality, stipulate the scope of the model's responses, and establish response formats.

**Example:** “Here’s your system message: You are an AI assistant designed to help answer questions related to adult learning principles. If unsure about a query, ask for more details.”

# Repeating the Instructions

AI powered by LLMs are subject to **recency bias**, so information positioned at the end of a prompt often wields greater influence on the output. This phenomenon suggests that reiterating critical instructions or guidelines towards the prompt's conclusion could enhance the model's adherence to desired response parameters.

Therefore, especially in longer prompts, restate key instructions or objectives at the end of a prompt to ensure the generated advice or analysis aligns closely with the requirements.

# Priming the Output

Introducing specific words or phrases at a prompt's conclusion can effectively "prime" the model to deliver responses in a desired format. For instance, hinting at the expected structure, like "Provide a step-by-step guide:", encourages the model to format its response accordingly.

## Example:

- **"You are an AI that offers concise instructional design evaluations. End each response with a brief summary."**

# Clear Syntax

Using clear syntax for your prompt — including punctuation, headings, and section markers — helps communicate intent and often makes outputs easier to read. In the example below, separators (--- in this case) have been added between different sources of information or steps. This allows the use of --- as a stopping condition for generation. In addition, section headings or special variables are presented in uppercase to differentiate them.

**System Message:** You'll read a paragraph, and then issue queries to a search engine in order to fact-check it. Also, explain the queries.

---

## PARAGRAPH INPUT

Ivan Rodriguez is married to Claudia Rodriguez. They have three children, and Ivan works at Microsoft and Claudia works for Pinterest. What search queries should I use to fact-check this?

---

## QUERIES

# Few-Shot-Learning

Few-shot learning is a machine learning technique in which several examples are provided in the initial prompt which establish a precedent for how the AI processes our query. This enables AI models to make accurate predictions even when provided with a very small number of labeled examples.

**Example Prompt:** "Consider three past corporate training programs with identified challenges: 1) Training Program Alpha struggled with low engagement due to lack of interactive content, 2) Training Program Beta had poor completion rates due to misalignment with employee schedules, and 3) Training Program Gamma received negative feedback due to outdated material. Given the upcoming Training Program Delta is designed for similar audiences as Beta but aims to incorporate more interactive content like Alpha, predict potential engagement challenges and suggest mitigation strategies based on these examples."

# Auto-Generated Chain of Thought

"Auto-Generated Chain of Thought" as a prompting technique allows us to get more sophisticated and accurate output from AI models, particularly in complex problem-solving domains. AGCOT asks the AI to articulate its reasoning process in a series of logical, interconnected steps, much like a human expert would when tackling a multifaceted issue.

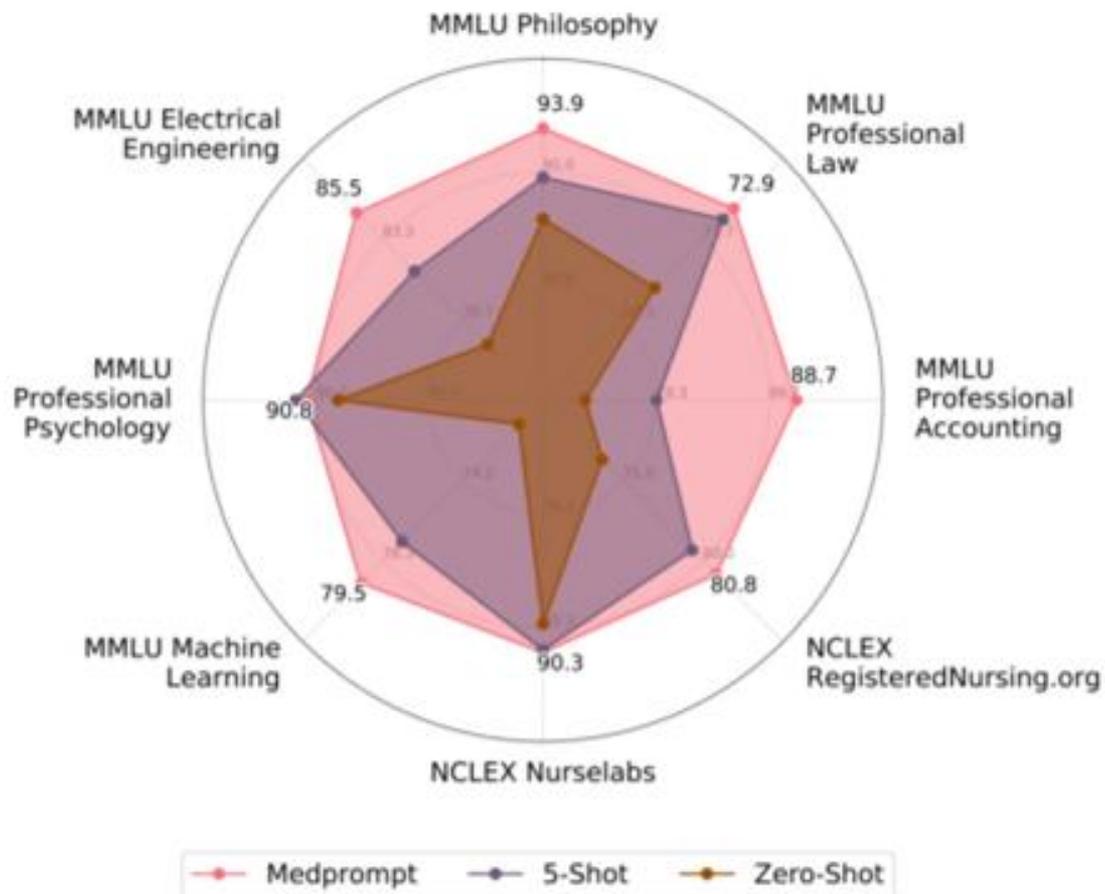
**Example Prompt:** "Given the following criteria for selecting a vendor: cost efficiency, reliability, past performance, and delivery timelines, along with the following vendor options: Vendor A, Vendor B, and Vendor C, break down the decision-making process. Consider Vendor A has the lowest cost but average reliability, Vendor B has excellent past performance but longer delivery timelines, and Vendor C offers a balance between cost and reliability but with limited past performance data."

# Ensemble Approach

"Ensembling" is a prompting approach to enhance decision-making in AI reasoning. The AI is presented with a set of potential choices answers the same question multiple times. The AI then compares it's choices, sees what was reasoned as correct the most, and more heavily weights that answer. This increases the quality and accuracy of the answer.

Additionally, we can change the order of the options each time the AI answers the same question. This introduces randomness and reduces bias. This is known as "choice shuffle ensemble."

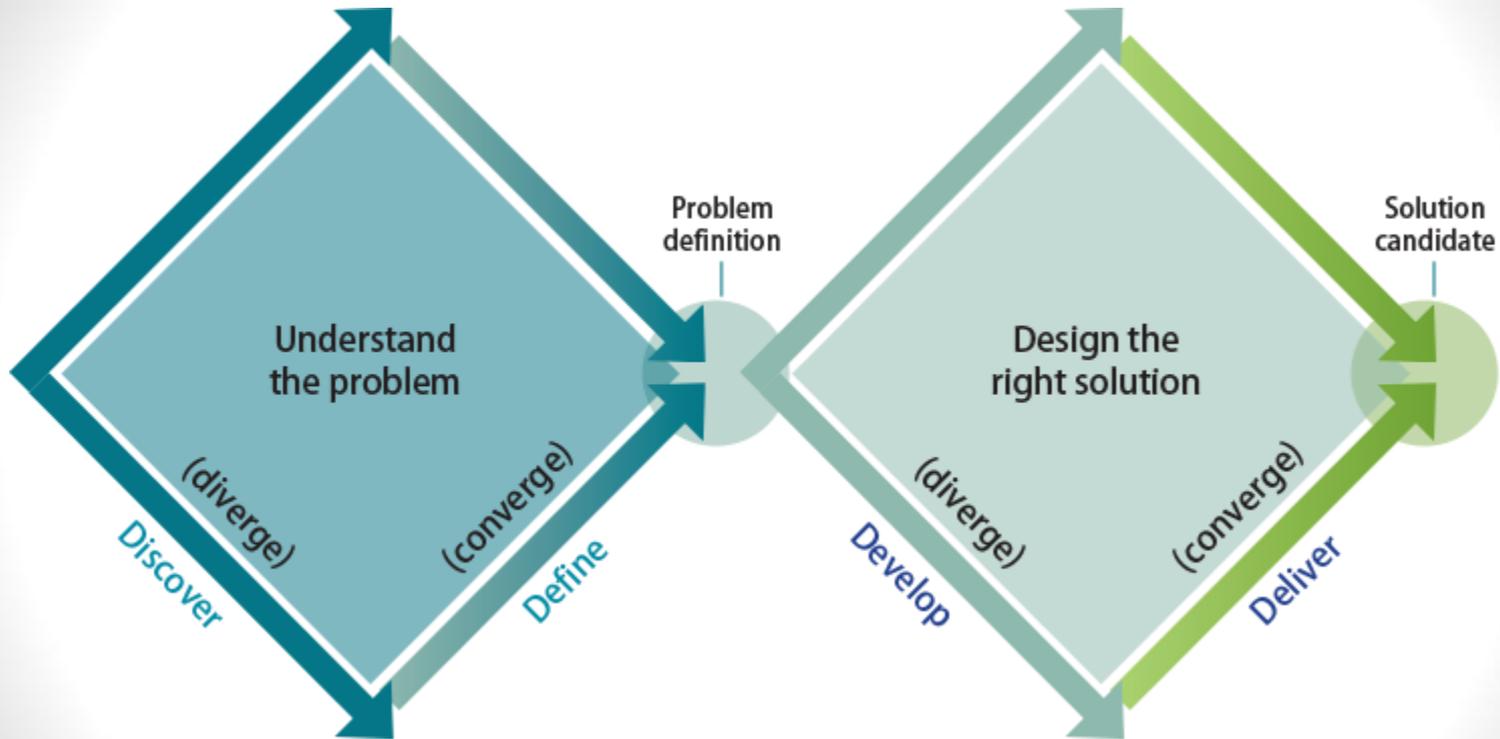
**Example Prompt:** "Our team is upgrading the LMS, and the new system needs to integrate with legacy systems, and be highly customizable. Consider evaluating four learning management systems for our organization's needs: LMS A is highly customizable, LMS B integrates well with existing systems, LMS C is user-friendly, and LMS D offers robust analytics features. Shuffle these options, answer the same question four times, and provide an ensemble analysis on which LMS might best suit our needs for customization, integration, user experience, and data-driven decision-making."



# Thank You!



Chris Knotts, Principal Consultant



# What is a Token

---

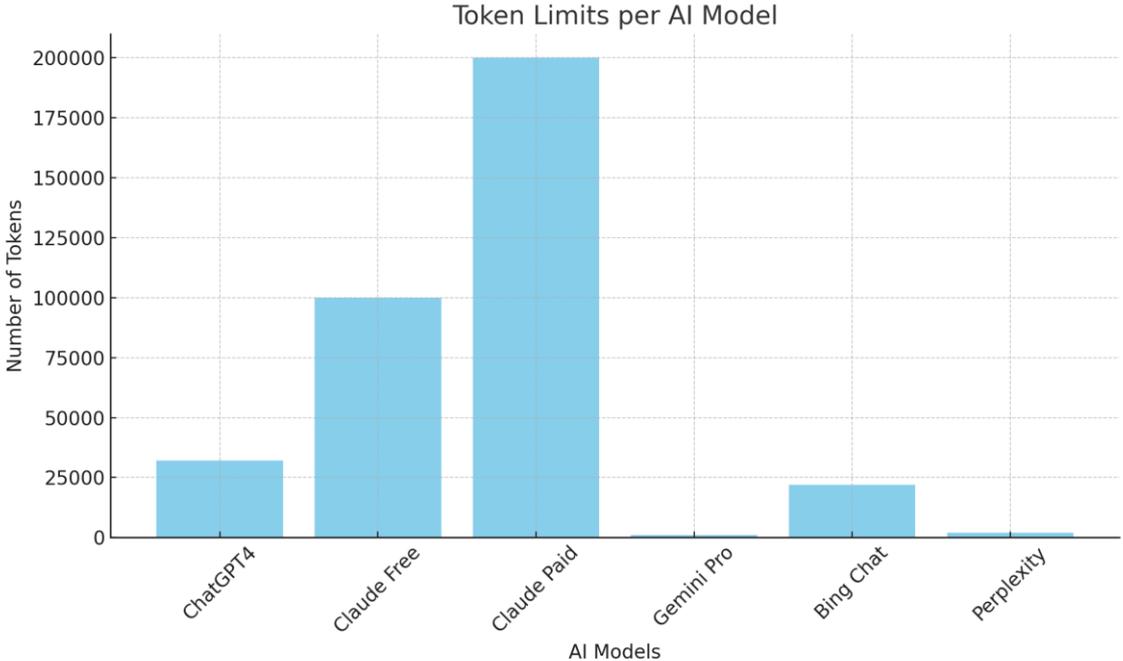
- In the context of generative AI, especially models like GPT (Generative Pre-trained Transformer) and similar language models, a "token" refers to the smallest unit of data that the model processes.
- Tokens can be words, parts of words (like syllables or subwords), characters, or even byte-pair encodings (BPEs), depending on the model's design and the granularity of the input text processing
- Think of it as a unit of measurement

# Tokens Used in Open Systems



You

Here is the data for tokens per model: ChatGPT4 32k, Claude Free 100k, Claude Paid 200k, Gemini Pro 1k, Bing Chat 22k, Perplexity 2k



# Energy Consumption

---

- Charging an iPhone 14 that is completely dead: It takes approximately 59,875 joules of energy to fully charge an iPhone 14 from a completely depleted state, considering a charging efficiency of 75%. This calculation is based on an estimated battery capacity and nominal voltage, and actual values may vary slightly depending on the specific device characteristics and charging conditions
- Creating one Image using ChaptGPT 4: Based on these assumptions, generating an image with an AI model like DALL·E on a GPU that consumes 200 watts of power would use approximately 1000 joules of energy for a process taking about 5 seconds. This is a rough estimate and actual energy consumption can vary based on the specific hardware and operational efficiency

Just tested, the next image you see took 1 min to create and so it consumed roughly 12,000 to 15,000 Joules of energy. Roughly 5 image “engagements” equal the energy of charging an iPhone. And it screwed up the image

# Considerations

---

1. Think of it as a very smart **junior assistant** that can make major mistakes
2. It is made to provide output so It can make major mistakes
  - it is very easy for the AI to “hallucinate” and generate plausible facts
  - ***It is programmed to provide output regardless of accuracy***
3. It must be supervised (Supervised AI)
4. It can LEARN, it does LEARN, and it is scary how well it learns
  - You can give it lots of information it will learn
    - Writing voice
    - Your company trends
    - Your team’s trends
    - How you like information given to you
5. You need to know the age of the datasets it is using
6. If you are not sure it is private, it isn’t
7. Different tools do different things
  1. You have the interface (each do different things)
  2. You have the LLM (pool of useable information)