

AI ABUNDANCE Workbook

START AND SCALE YOUR AI TRANSFORMATION

Welcome

You're here because you need to begin transforming your company with AI, but you don't know where to start. Don't feel bad—every company is in the same position. Beginning an AI transformation is difficult because the technology impacts everything and is changing so fast.

Faced with these challenges, many companies try to identify the best starting point before deploying AI. They compile use cases and hold weekly AI steering committee meetings, but a year later, they often have more questions than answers. Others invest all their resources in one or two high-impact initiatives, only to realize that AI isn't working as expected and requires significant engineering investment. Many of these initiatives stall, leading analysts to conclude that AI is overhyped.

There is a better approach—one that is working for companies already deploying AI and realizing early value from their investments. Today, you'll learn how to mimic their success.

In Part 1, we'll provide a framework for identifying AI opportunities and the most efficient way to pursue them. You'll develop a list of potential opportunities at your company that we'll evaluate together. In Part 2, you'll learn about foundational technologies and generative AI system design, using workshop tools to solve a few AI problems in a team competition. Finally, we'll cover next steps like creating your AI strategy, building skills, launching your first project, and making buy vs. build decisions.

By the end of this workshop, you'll know how to begin your company's AI transformation, avoid common mistakes, and discern value from hype.

This is just the beginning, and we hope you'll consider Prolego when you need help.

01


Opportunities

AI is improving exponentially fast

No, it isn't you.

Transforming your company with AI is challenging for two reasons. First, AI impacts everything—like electricity, computers, or the internet. When everything is affected, it's tough to know where to begin. Second, AI is changing incredibly fast.

Exponentially fast.



This graph illustrates the increase in computing power Microsoft is deploying to train AI models. All relevant trends in AI - how smart it is, how fast it runs, and how fast the costs are falling - follows these trends. Source: Microsoft Build 2024.

Fortunately, we can compare companies making progress with AI to those that aren't.

What isn't working

First, let's discuss what isn't working. Many companies, faced with AI's complexity, are primarily investing time on the following activities and not making progress.

Approach to AI	Why it isn't working
Search for high-ROI use case before starting	Value comes through creating broad AI capabilities, not a single use case
Organize AI strategy meetings	Only surfaces more questions without answers. No learning until you start
Pursue a single, high-impact opportunity	Requires too much customization. High risk of wasted investment

Strategy meetings and seeking the best opportunities are fine, but doing these in isolation won't get you far.

What is working

So what is working? Companies making progress with AI are "thinking exponentially." They have a vision for transforming their businesses and look for efficient ways to leverage others' investments.




Global investment in AI by tech giants, investors, and governments is approaching \$1T/year. AI's capability (measured as a combination of intelligence, speed, and cost) is improving at an exponentially-fast rate. All of your investments and strategic decisions should align with this trend.

Think exponentially

We can summarize the Think Exponentially strategy in one sentence: Transform your business by efficiently deploying AI as it rapidly improves. This requires embracing three principles:

1. Don't wait
2. Deploy AI where it works today
3. Build capabilities designed to scale.


Transform your business by efficiently deploying AI as it rapidly improves



Don't wait



Deploy AI where it works well today



Build capabilities designed to scale

Companies making progress with AI are embracing the rapid change. Instead of searching for the ideal starting point, they're beginning to deploy AI where it can begin creating value and adding more as it gets better.

Let's explore these principles in reverse order.

Principle 1

Build capabilities designed to scale

What can you do with AI?

AI can replace or enhance every business process. For example, here are capabilities you can achieve with AI.

Process and create documents

Read and write to databases

Read and update spreadsheets

Perform analysis across data types

Assign and coordinate simple tasks


Create and use simple software tools

AI is a foundational technology like electricity, computers, or the Internet. Foundational technologies don't neatly fit into your current business because they don't replace a specific solution. To find opportunities you need to consider capabilities like these, and then see where you can apply them broadly.

Principle 1 Build capabilities designed to scale

These capabilities can be used broadly across your business. A capabilities-based approach, rather than a solutions-based one, prevents duplicated investments and promotes incremental improvements across all departments and functions.

Develop **Capabilities** to support **AI Solutions** across all departments and functions.



Principle 2

Deploy AI where it works well today

AI can theoretically do anything, but it excels at some tasks more than others. Tackling your hardest problems with AI today can lead to wasted investment. Since AI evolves so quickly, it makes more sense to deploy it where it works well now and tackle tougher problems as AI improves. For instance, AI excels at processing simple documents with minimal customization, making it a great starting point for most companies.

Principle 2 Deploy AI where it works well today

How should you begin deploying AI?

AI Capabilities	Today	2024-2026	2027-2030
Process simple documents	Easy	Easy	Easy
Process complex documents	Hard	Easy	Easy
Use databases, spreadsheets	Hard	Easy	Easy
Perform simple analysis	Hard	Easy	Easy
Assign and coordinate simple tasks	Hard	Easy	Easy
Create and use simple tools	Hard	Easy	Easy
Perform complex analysis across systems	Very Hard	Hard	Easy
Coordinate complex tasks across systems	Very Hard	Hard	Easy
Create and use complex systems	Very Hard	Hard	Easy

Easy: part-time engineer

Hard: dedicated team

Very Hard: not possible

Focusing on solutions with high-ROI potential requires customization and significant investment. Before pursuing such a path, ask yourself whether AI will be better at solving the problem in a year or two. While it is impossible to know what tomorrow's AI can do, you can make a reasonable estimate based on the recent past. In many cases, deploying AI where it works well today is a more efficient approach.

You can also deploy AI against harder problems like processing complex documents, interacting with spreadsheets and databases, or coordinating tasks. AI is almost good enough to handle these tasks without much customization, but not quite yet. Investing in these capabilities in a year or two, when AI can handle them with minimal customization, makes more sense.

Principle 3

Don't wait

You could wait for AI to get easier to deploy, but this comes with high risks. You might miss out on current AI opportunities and lose ground to competitors.

Companies need time to adapt, and history shows many waited too long to adapt to fast-changing trends and couldn't catch up.


How should you begin deploying AI today?

The most efficient way to begin deploying AI is by using it to solve problems where it works well with minimal customization. Ideally, you want to start solving problems using AI “out of the box.”

Processing simple documents

The most common solutions involve processing simple text documents. For example, taking text documents such as contracts, emails, chat logs, and policies, and instructing the AI to transform them into a more useful format: summarizing them, extracting key pieces of information, or classifying them.

SUMMARIZE, CLASSIFY, EXTRACT INFO




You can send documents to an AI model along with instructions to transform them into some desired result: summarizing the contents, extracting key pieces of information such as names or dates, or classifying them. You can apply this workflow anywhere you currently have people reading and analyzing text.

Custom document chat

Another example is building question and answer applications, sometimes called RAG, over simple documents such as manuals or employee policies. You can think of this as building a customized chat app with access to your foundational documents.

ANSWER QUESTIONS



Many companies are also building customized chat applications on their data. Often called retrieval-augmented generation (RAG), it currently works very well on simple documents. Solutions based on complex documents such as those with tables, diagrams, complex structure, or arcane language usually require more customization.

Other AI opportunities

There are many other options, but they are probably specific to your company or industry. In the next section, you'll learn a simple heuristic for quickly estimating whether AI can solve a problem without significant customization—the intern analogy.

Q2

Technology

The most important technology in your AI transformation is foundational AI models, currently Large Language Models (LLMs). LLMs power products like ChatGPT. They are improving so fast that discussing their differences isn't useful since everything will change by next week. Instead, let's focus on what they are and what you can do with them.

LLMs are like interns

We like to think of LLMs as interns. Here are some characteristics of an intern:

1. They have general knowledge but lack specific knowledge about your company.
2. They perform tasks at a low cost but need specific direction and supervision.
3. They can't solve complex problems but excel at repetitive, simple tasks with clear instructions.



In 2024 LLMs are about as smart as an intern. A few years ago they were no smarter than toddlers, and soon they will be as smart as graduate students.

These characteristics also describe LLMs. Many leaders mistakenly see AI as an uber-genius that solves complex problems. In reality, AI is better at solving simple problems at scale.

Of course “intern” is a good analogy for what LLMs can do today. A few years ago LLMs were no smarter than toddlers, and they will be as smart as PhD students very soon.

Find AI opportunities with the “army of interns” hack

If you’re wondering if AI can do something, ask yourself if you could train a bunch of interns to do it with proper instruction and enough time.



“What would you do with an army of interns?” This question can surface opportunities for AI.

This hack works great for discovering opportunities with your business customers. Just ask them, “Suppose I gave you an army of interns. What would you have them do?” Their response is a good signal of potential AI opportunities.

Prompts are like messages on paper airplanes

We can't communicate with LLMs directly or through visual cues. Instead, we use prompts—messages sent back and forth. Think of prompts as information written on paper airplanes sent to the intern, who then writes their response on another paper airplane and sends it back.



Unlike a human intern, you can't have a conversation with an LLM. Instead, you need to send messages, called prompts, and data back and forth. Messages written on paper airplanes are a good analogy.

Communicating with LLMs through prompts has similar challenges to sending messages via paper airplanes.

A more efficient way to program


Many companies underestimate AI by seeing it as an additive technology to existing systems. It's actually a new way to program and work with computers.

	Python	AI
Interface	Interpreter	LLM
Language	<pre>with open(filename) as f: contents = f.read() lines = contents.splitlines() for line in lines: if not line or line.startswith("# ") or line.startswith(: continue</pre>	"Delete the comments in this code:"
Job	Programmer, Data Scientist	AI Systems Engineer

Traditionally, we write software by creating files that tell the computer exactly what to do, like writing Python code for the Python interpreter. With AI, we use prompts in general language to instruct LLMs. This approach is vastly more efficient, and ultimately, all existing software will be replaced by software using LLMs because it's more efficient to build complex systems.

A basic AI system

This is what a basic AI system looks like: You send data chunks with instructions, called prompts, to the LLM and task it to perform some task for you.



A basic AI system involves passing prompts and data to LLMs.



Next steps

Well, you've learned the basics of finding AI opportunities and the technology. Don't feel like an expert? Don't worry, none of us do because the technology is changing too fast.

What matters is that you take action. Here are four common next steps.

Create your AI strategy

Most companies are creating an AI strategy to present to their CEO or board to begin their AI investment. We've helped clients create these strategies over 6-12 month projects. In reality, this isn't worth your effort because the technology and opportunities are changing so fast.

Instead, create a vision and a set of guiding principles, such as the AI strategy pillars, where you commit to investing along 5 dimensions to transform your company with AI.

Transforming your company with AI requires more than just building technology. Most companies capture these changes along 5 pillars such as these. You download a starting powerpoint template at the workshop resources page.

AI strategy pillars

01	02	03	04	05
CAPABILITIES	INFRASTRUCTURE	TALENT	GOVERNANCE	TOOLS
Systems that solve many business problems	Computing services to securely process data and support solutions at scale.	Human capital to execute and maintain a competitive position.	Metrics and oversight to ensure business impact and reduce risks	3rd-party resources to accelerate our program
Examples: search ,unified natural language query, decision agents	Examples: vector databases, data pipelines, monitoring solutions.	Examples: data scientists, ML engineers, AI product managers	Examples: investment steering committee, model policies, performance metrics	Examples: data science development, chat interfaces, decision dashboards

Include example projects and opportunities, and perhaps a steering committee of tech and business leaders from your organization. Instead of spending months creating a perfect plan, you're better off doing something quick to get initial funding and start pursuing the opportunities discussed in Part 1.

The resources in the appendix include a link to an AI PowerPoint strategy deck you can download and customize for your company.

Build your company's AI skills

If your company has a technology team that knows how to build and deploy software using modern cloud services, you probably have most of the skills you need to leverage AI. (This assumes you want to deploy AI where it works well today without much customization and model training.)

Download a copy of this job description at the workshop resources page.

Position: AI Systems Engineer Company: Acme, Inc. Location: [Insert location]		PROLEGO
About the Job:	As an AI System Engineer at Acme, Inc., you'll play a pivotal role in our transition to becoming an AI-driven company. Your work will encompass designing, developing, and optimizing interfaces with large language models (LLMs) to integrate them into our evolving AI initiatives. You will experiment, innovate, and collaborate across teams, laying the foundation for our future AI projects. Your goal is to create seamless interactions with LLMs, ensuring security, efficiency, and alignment with our company's mission and values.	
About Acme, Inc.:	Acme, Inc. is stepping into an exciting new phase, initiating its first AI projects with LLMs. Be a part of this transformative journey and help us explore the limitless possibilities of AI.	
Responsibilities:	• Collaborate with Teams: Engage in meetings with diverse stakeholders to understand workflows and provide insights. Educate teams on the potential and boundaries of integrating LLMs into existing systems. • Craft LLM Prompts: Creatively design the prompts necessary to guide LLMs towards specific tasks, ensuring alignment with desired outcomes. • Develop Integration Software: Construct robust software to process LLM responses and enable integration with existing applications. • Implement Intelligent Constraints: Design constraints to prevent users from asking questions that LLMs cannot answer, maintaining alignment with the task objectives. • Assess and Mitigate Security Risks: Monitor and evaluate potential security risks like prompt injection or sensitive data leakage, implementing necessary security protocols. • Coordinate Complex AI Tasks: Design agents to manage intricate tasks such as multi-database SQL queries or automated workflows. • Experiment and Innovate: Lead experiments to test LLM communications, analyzing responses to ensure desired results, and iteratively refine processes. • Participate in Design Meetings, Standups, and Planning Sessions: Engage in the necessary meetings to develop and deliver solutions within the team.	

If so, the primary role you need is an AI Systems Engineer - someone who can build complex capabilities that support many use cases across your company using prompts and basic data parsing. This is not data science, traditional software programming, machine learning, or natural language processing. It is a new skill for a new programming interface. You can download a job description on the workshop resources page.

While courses and training are all fine, you and your people will not truly understand AI until you start trying to solve your problems with your data at your company.


Launch your first project with a Performance Framework

Most engineering teams working on their first few generative AI projects quickly get stuck. Because LLMs produce different results for different tasks, it can be challenging to efficiently build a solution. We don't yet have methodologies like Agile or Scrum as we do with traditional software engineering. In the worst case, teams find themselves in perpetual "iterate-by-feel"

situations, constantly changing and trying prompts and models without a way to measure progress or predict completion. To provide transparency into how well your system is working and where to invest, you need a performance framework.

Even experienced engineering teams quickly get stuck building solutions.

A Performance Framework provides transparency and allows you to make progress.



Task	Expected	Actual	Confidence	Model	Latency	Cost	Recommendation
1	Data match	Data match	High	Mistral-7B	1.6s	0.002	Deploy baseline prompt workflow
2	Data not match	Data not match	Medium	GPT-4o	15s	0.04	LLM occasionally gets confused. Create validation tool for LLM.
3	Data not match	Data not match	Low	GPT 4o	35s	0.1	LLM cannot solve without additional context. Get additional data or don't deploy

Traditional software engineering methodologies like Agile or Scrum don't translate well to Generative AI projects because LLMs are non-deterministic. Instead, developers need to break projects down into specific tasks and data sets that cover the problem space and measure performance against each. This approach provides transparency and helps you decide where and why to invest.

Break your desired solution into representative data and tasks, and try to solve them using prompts, orchestrations, and basic data processing. This allows you to see where your system works well and where you need further investment.

Buy AI solutions? It depends

Every technology vendor is hyping the AI part of their product offering. Now that you know how to think exponentially, you can ask the critical question: *How valuable will their solution be as AI improves exponentially?*

Solutions based on deep domain expertise relevant to your industry or a unique proprietary dataset are more likely to be valuable even as AI improves. Better AI might make them even more valuable.

On the other hand, incumbent solutions adding AI to their existing services may be less valuable when AI can do most of what their traditional product does.

For example, general solutions like Customer Relationship Management, Marketing Automation, or Document Management are at high risk of being displaced by large language models. So too are so-called "GPT wrappers" with no proprietary AI that claim to solve problems by sending prompts back and forth to LLMs.



"GPT wrappers" are new solutions with no proprietary AI. They simply send data and prompts to LLMs



All incumbent vendors are "adding AI" as part of their service offering. But will future AI render their entire offering obsolete?

About Prolego

The premiere
generative AI
solutions provider
since 2017

Prolego is an elite consulting team of AI engineers, strategists, and creative professionals guiding the world's largest companies through the AI transformation.

Prolego's proprietary Performance Evaluation Framework offers engineering teams a proven methodology for developing and deploying bespoke LLM-based solutions, enhancing schedule adherence, boosting performance, and optimizing cost management.

Representative clients

LOCKHEED MARTIN

citi

Morgan Stanley

Raytheon
Technologies



BOSCH



Bristol Myers Squibb

CLOUDERA



TRANSAMERICA

COX
AUTOMOTIVE™

FINRA

Protective



HEIDRICK &
STRUGGLES

Contact us

Interested in discussing your
generative AI projects, contact
Russ at russ@prolego.com