

Artificial Intelligence

A primer for SMB leaders



© Mainstay Technologies 2023. All Rights Reserved.

Author: Ryan Barton

Table of Contents

Chapter 1

Understand AI

- 4 Introduction
- 5 AI is *Intelligence*, not *Technology*
- 7 We must learn the vocabulary
- 9 AI will grow in two dimensions, at a blistering pace
- 13 This will impact our life, our business, our industry
- 17 Or maybe not

Chapter 2

Enable AI

- 26 Introduction
- 27 Find AI-enabled, industry-specific software
- 28 Become skilled - as an organization - with large AI models like ChatGPT
- 31 Watch for what's coming
- 33 The snapshot to share with other business leaders

Chapter 3

Align AI

- 37 AI Alignment
- 38 Imagine a few scenarios
- 42 What you and I can do, as business leaders
- 43 How do I protect and empower my current team?
- 44 How do I avoid layoffs?
- 45 How do I share the gains?
- 47 How do I ensure that my business efforts are aligned to human flourishing?

Chapter 1

Understand AI



Introduction

The invention of Artificial Intelligence is unique in human history. It is not a new version of something we already know and understand, it is not a natural derivation of earlier technology, and it is not a common-sense invention.

AI feels intelligent, but it runs on technology. This has enormous and unique ramifications. If we do not understand this invention properly, if we mis-categorize it, if we over-hype it, or (worse yet) if we pay it too little attention, we risk becoming a victim to forces far beyond our control.







The key to understanding AI: wrestle with it as a concept, work to grasp the dynamics at play around it, and to consider the potential futures it will bring. Then, we can be empowered to make a prediction for our own career and business.

And, we can begin to wisely prepare for it.




AI is *Intelligence*, not *Technology*

AI is a type of intelligence that *emerges* from technology, but at a higher order of complexity. The same as our intelligence *emerges* from our biological bodies. So when you hear AI referred to as a “tool” or a “technology” that’s only partially true. Yes, something like ChatGPT is a tool. And yes, AI is technological. But we are talking about *intelligence*. Not just a tool, but rather something more powerful and something novel to history.

 Intelligence	Human Intelligence	Artificial Intelligence
 System	The body	Technology (servers, networks, applications, cloud, etc.)
 Components	Organs, parts of body	CPUs, memory, I/O devices
 Substrate	Carbon, oxygen, hydrogen, nitrogen, etc.	Silicon, copper, gold, oxygen, etc.
	 Biology	 Technology

The invention of problem-solving, decision-making intelligence that is being given agency is fundamentally of a different order than technology that does our bidding. AI models are increasingly being given the ability to make decisions and to impact the world – autonomously.

This does not mean it is sentient or rational, but it does mean it can likely soon solve problems currently occupying millions of white-collar workers.



We must learn the vocabulary

AI terms now are what “email” and “browser” were decades ago.

Machine Learning

A type of AI where the system is coded to *learn* from data and experience.

Neural Network

A technological construct inspired by a human brain. Contrasts with historical attempts when developers tried to achieve AI by simply adding more and more complex code (which doesn't work).

Deep Learning

Using large data sets to train neural networks. Requires huge data sets and tremendous processing power. The AI model learns becomes a prediction machine: feed it 1,000,000 pictures of a cat, and it can successfully learn to predict the form of a cat, and identify cat in any picture. Feed it Moby Dick, and it can predict the last paragraph (or the sequel!).

Generative AI

Early AI models were used to make predictions as part of complicated processes (like detecting fraud on your Visa credit card). They analyzed data and made predictions. Generative AI, on the other hand, can create text, images, and other content from simple prompts.

LLM (Large Language Model)

A neural network AI trained on language. This is where the breakthroughs are happening, as this democratizes AI (you don't need millions of dollars and of data sets to utilize an AI anymore; advanced LLMs are designed for the public. And language is giving LLMs advanced understanding of psychology, culture, etc.)

Artificial General Intelligence

Current AI models can solve problems in a very narrow domain. The stated goal of the AI industry is to invent artificial "general" intelligence: an AI capable of solving multiple types of problems, in multiple domains, like humans do to navigate reality.

AI will grow in two dimensions, at blistering pace.

AI is currently brilliant and primitive at the same time – a genius baby! It is wildly intelligent on some dimensions, and completely ignorant in others. AI is not emotional, and yet GPT4 can pass the bar exam at the 90th percentile. AI does not have self-direction, but Google's MED-PaLM can pass the medical exam. AI cannot make dinner, but it can create delicious recipes.

Breakthroughs are occurring at a rapid pace in both dimensions – in capability (horizontal) and in power/intelligence (vertical).

Our world has already relied on AI for years, whether in Google search, Google translate, Netflix recommendations, Facebook's feed, YouTube's algorithm, credit card fraud detection, etc.




AI development has reached a tipping point, where it the power of the AI models, the broad cultural awareness, and the prospect of enormous financial return is combining to accelerate investment.

This has been growing for a long time.

Since enlightenment philosophers posed it in 17th century, through start of computer science in 19th century, through Alan Turing formulating Turing Test in 1950, to development of neural networks in 90s and early 00s, through Google launching Translate in 2016, finally to OpenAI releasing ChatGPT in November 2022 and upgrading the underlying intelligence to GPT-4 in April 2023 – start of a new era of widely accessible, hugely powerful AI.



Cognitive scientists predict AI development will not be in a straight line but will hit thresholds that stall development, eventually crossing them to rapidly scale intelligence, empower robotics, and become more and more comprehensively capable.



AI is such a powerful invention, that it threatens the survival of massive industries.

Every major tech company is now locked into a race – the AI IQ wars have begun. Market forces + geopolitical tensions (the military is well aware that the smartest AI will likely win the next war) ensure that this pace continues to drive forward at blistering speeds.



Current AI models are not fully predictive of the future.

We must avoid the fallacy of playing with AI now, making a judgment on its practicability, and dismissing this as an over-hyped trend. Dynamics are driving AI to become more capable at a hard-to-grasp pace (imagine playing with an early iPhone and dismissing its capabilities – like Blackberry did).

This will impact our life, our business, our industry.

AI is like a meteor hitting the ocean... Impossible to predict exactly the size or speed of the waves. But we can be confident that tidal waves are on the way.

Disruption is inevitable. The scope and pace will vary widely based on industry.

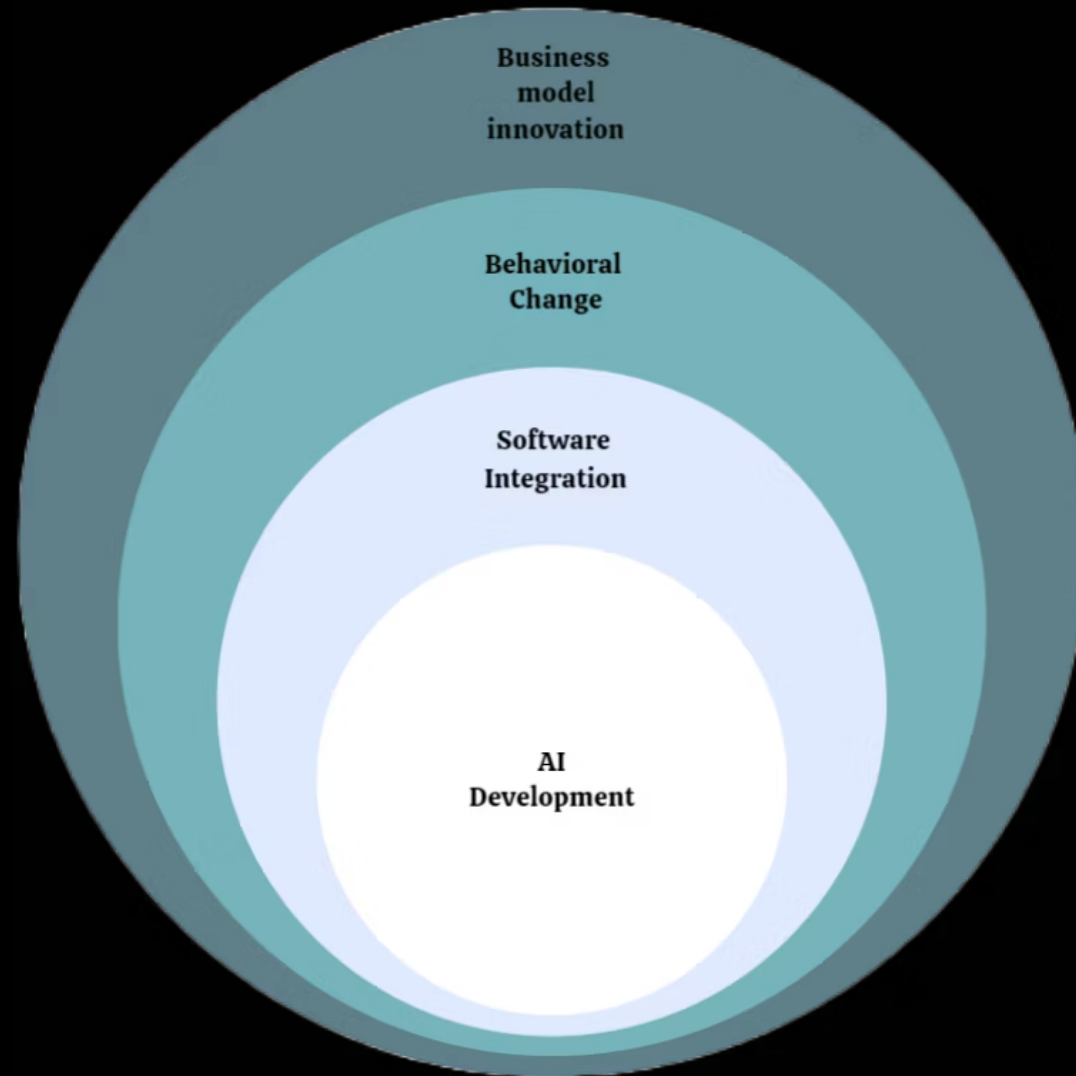
It is hard to find any experts dissenting with this statement:

AI will be more disruptive than what we have seen before, at a much faster timeline than we have experienced before.

AI is like a rising tide of intelligence. Very soon, AI models will do entry level professional tasks very effectively. The impact will be fastest where tasks are repeatable and data is abundant.



Crossing thresholds to career and business disruption requires the AI to be integrated, behavior to change, and business model to adapt.



Think of waves in 3 ways, that will go at different timing for different organizations, with some level of overall global timing:

1. Enhancement: Increase efficiencies. On an individual level, partner with AI. As an organization, implement carefully.
2. Replacement: AI assumes tasks and responsibilities.
3. Revolution: Reimagine what is possible.



**Consider the opportunities in all 3.
Ride the wave, don't be flooded by it.**



Or maybe not.

It's possible that this is all overblown.

Perhaps AI will change the world, disrupt industry, and shift millions of jobs. Or perhaps it won't – at least not for decades.



Here's what the counterargument sounds like...

True intelligence is *waay* harder to simulate than most make it sound.

There is no guarantee Artificial General Intelligence will ever be invented (much less soon!). All current AI models are just fancy prediction machines – they predict what a human wants to hear. These models have no direct relationship with reality, with creativity, with love, with motivation. They lack motivation, emotion, virtue, and orientation. They simply predict.



Job elimination isn't the same as job enhancement.

Current research suggests that AI will automate take on 50% of the workload of roughly 20% of the workforce in the next few years. On the surface, that sounds like at least 10% of jobs are going away in the next few years – a potential economic cataclysm.

However, the “Theory of constraints” teaches us that speed is determined by bottlenecks. It’s the choke points, the last remaining areas, that determine actual velocity.

Applied to job replacement, it is likely that a few, last remaining tasks will be unattainable by AI. Automating some of what a human does is easy, but automating all of it? That may be impossible.



Hype has been around for years.

Computer scientists have overblown their expectations at every stage in AI development. In 1965, a group of developers believed they could invent AI over one summer at Dartmouth, NH.

In the 2000s, AI was forecasted to be everywhere within “years.” In 2015, Elon Musk predicted that self-driving cars would be significantly better than humans at driving by 2018 – an elusive hope.

In the radiology field, AI has been better at reading radiology charts for years and many have forecasted the end of the radiologist. However, as of 2023, there are more radiology jobs available than ever before.




Technology goes through Gartner's Hype Cycle.


We are a culture that loves new tech and most of us have a great imagination for how new technology will change our lives. Whenever a new technology is released, expectations are meteoric. However, when the tech is implemented, reality sets in, and the “peak of inflated expectations” is followed by the “Trough of Disillusionment” (Consider Virtual Reality, and where it is in this cycle currently. Or ever been part of a company that implemented enterprise software, only to find that it made things worse, not better?). After bottoming out in the Trough of Disillusionment, the new technology grows steadily, until it reaches a plateau. It doesn't apply to all technology, but most follow this general pattern.

Where is AI in this cycle? Undoubtedly, at least some of the current anticipation is nothing more than “Inflated Expectations.”

How much? That's the question.




Gartner's Hype Cycle



What to do?

Adam Grant, in “Think Again” gives us simple guidance for times like these: In uncertainty, update predictions regularly. This is counterintuitive. Once we have thought about a topic, we like to feel settled. Humans have a built-in Confirmation Bias which means we rarely challenge our beliefs, instead seeking confirmation that we are right.

With AI, we must challenge those beliefs. We must not let Confirmation Bias determine our path. We must do as Adam Grant exhorts and update predictions regularly.



As AI evolves, you and I must reconsider and re-wrestle with the implications.

Frequently.

Said another way: Parts of this eBook are false.

We just don't know which parts.

So, let's keep learning and keep calibrating our predictions.

Chapter 2

Enable AI





Introduction

Publicly available AIs (such as chatGPT) are powerful, but hard to harness in an organization.

Using a general model AI in your business is like having an intern on day 1. She may be incredibly smart and capable, but she isn't trained on the business, lacks access to company software, and knows no one. A few employees will find ways to use her time. Most will simply ignore her or ask for a coffee run.

Current large-scale AI models are similar: Some roles can leverage AI immediately (like creative writers, customer service email authors, troubleshooters, etc.). But most are focused on their own jobs, lack imagination for how AI can help, and face high barriers to integration.

It requires training and integrating the AI into your data, and your business processes, and your culture, to make a real impact.




To solve this, consider 3 categories of AI software.

Find AI-enabled, industry-specific software.

- Remember, AI development has been roaring for years. In nearly every industry there is already a well-funded software company focused on harnessing AI to solve business problems. These companies may have nothing to do with ChatGPT or the hype of 2023, but they often still have powerful AI models.
- The increasing attention, funding, and market dynamics are accelerating these companies. They are also accelerating businesses in considering and adopting this innovative approach (convincing a Board of Directors to invest heavily in AI was nearly impossible in 2022, and often requested in 2023!).
- Begin a search for the leaders in your space. In small and mid-sized businesses, it is often wise to start with software leaders focused on Enterprise. The expense and cumbersomeness of the software may make them unattainable, but you will learn through the exploration. Historically, enterprise solutions come first, with SMB following a few years later.
- The urgency around AI is a powerful motivator to explore and adopt the most innovative, AI-integrated automation software already available.

Become skilled – as an organization - with large AI models like ChatGPT.



*ChatGPT: Optimizing
Language Models
for Dialogue*

Move to cutting edge.

- Most business leaders avoid being early adopters. They prefer to be a “technology leader” while avoiding the “cutting edge,” and its attendant disruption. However, the disruptive potential of AI is so significant, it is necessary for many of these leaders to embrace mess and uncertainty and get closer to the cutting edge. In his Inspire conference keynote in 2023, Microsoft CEO Satya Nadella shared projections that AI will add 7-10% to global GDP (trillions of dollars per year).

Set the tone for your organization.

- AI does not need to be a threat. By embracing it early, we gain a degree of control and leadership, so we can choose what is good for all. If employees are worried about potential job loss, explain that AI would have to increase efficiencies more than the company growth rate + staff attrition rate each year to even threaten current jobs. A tall order!
- In the near future, AI will become like Microsoft Excel. All professional workers will need to have proficiency, as part of basic professional skills.



Train staff.

- Expect managers to become versed in AI, and rollout training on Microsoft Copilot and ChatGPT to staff. Get them comfortable with prompting, with interacting with an AI.

Start with small use cases.

- Find use cases in your team for ChatGPT, Copilot, Google Bard, or another large-scale AI model. Start small and gain traction. Common examples include: help with policy creation (HR), blog post ideas (marketing), email authoring (customer service), troubleshooting prompts (IT), and research (executive).

Expand implementation.

- Sponsor a hackathon. Inspire the team to submit ideas for how to use generative AI in a way that improves work *today*. Gather a select team to choose the winners and give out prizes and bonuses as a result! Just don't neglect implementing and training on the winning results.

Partner with thought leaders.

- If you have an effective IT Department / Company they should be well versed in AI and able to provide education and direction. For more advanced implementation, look for consultants that focus on your industry, who know AI, and have skill in predicting trends.



Consider privacy.

- Note that public AI models are not secured for your organization. Confidential information should never be shared with ChatGPT or any other model that you don't control. Rollout a simple AI policy to all staff, and work with IT to setup secure and private ChatGPT access when it makes sense (this is technically feasible as of late 2023).

Consider security.

- Sadly, AI has downsides too. It is already making hackers more effective. These AI models can be easily used to write effective phishing emails and to accelerate a hacker's education. Unfortunately, AI is making cybersecurity defenses even more important than they were before. Recognize that security threats are increasing and work with your IT and Information Security teams to increase security protections.

Data is gold.

- Because AI models require data for training and tuning, organizations with a lot of data – that is well organized and structured – have a significant advantage over those who don't. Preparing for AI often involves curating data, cleaning it up, and integrating it across platforms

Watch for what's coming.



Don't be a cliché.

- No one wants to be a taxi company right before Uber is released. We all know the story of Kodak, who doubled down on film after inventing the digital camera. AI has the kind of crackling, transformative potential that will shift industries and upend long-standing companies. Increase your “productive paranoia” (a wonderfully helpful term coined by Jim Collins). Make AI strategy a core part of your business strategy. Expect leaders to wrestle with AI's impact on your business. Pay close attention to your market. Update predictions regularly.

Expect more powerful software in 2024 and 2025.

- Future software will integrate large, powerful AI models. There is a new model of software development: tune and integrate large-scale AI into specific industry niches. As these models grow in power (funded by tens of billions of dollars), small software companies will train and integrate them into every imaginable industry niche. Watch for these players (or consider becoming one).



Expect change.

- Technology innovation always triggers business change. Every new technology opens new possibilities. Those who see the possibilities, and who move quickly, gain the rewards. The Internet completely disrupted retail. The iPhone changed software forever. Enterprise software transformed professional services. The pace of change always puts stress on business leaders. AI is potentially the most disruptive technology the world has ever seen.... The pace of change may be blistering. As a result, become skilled at change. Expect it, embrace it.
- Learn how to harness it, considering frameworks like John Kotter's "Leading Change" (and the book by the same title).

The snapshot to share with other business leaders:


- A. AI will change more than we can imagine.
- B. We need to be ahead of this, not behind it.
- C. We simply need to do 3 things:
 1. Get close to innovative software companies in our space and to AI thought leaders in our industry (even if they are focused on Enterprise).
 2. Become skilled as a whole team at implementing publicly available LLMs in our roles and processes.
 3. Be paranoid, seek more knowledge, and update predictions regularly.



Chapter 3

Align AI





Introduction

The latent power of AI is unsettling.

News headlines trumpet the potential of a sci-fi dystopia come true: a conscious AI that decides to eliminate humans.

Industry leaders warn of seismic disruption to industry and rapid unforeseen changes.

As business leaders, we face a more direct question: Is our pursuit of AI good for our employees? Are we quietly heading towards layoffs and towards meaningless jobs?



AI Alignment

All technology has the potential to help us or harm us.

Social media is a striking example: In the early years, it seemed clear that social media brought us together, created new possibilities for connection, and reduced loneliness. Facebook's headquarters proudly displayed a wall of stories: handwritten notes of gratitude from those who met their spouse on Facebook.

Today, the social media story is no longer one of connection. Studies by Dr. Brian Primack and others have shown conclusively: Any engagement with social media increases each of our risk of anxiety and loneliness. The more we engage with social media, the higher the impact. Supposedly "social" technologies push us towards loneliness.

If AI is deployed in a "profits first" race, without thoughtfulness, what might the results be?



Imagine a few scenarios:

- A content generating AI is given the goal of increasing user engagement. It fabricates stories and generates false narratives that grab attention, making the truth even harder to find.
- A financial AI is given the goal to maximize capital returns. It finds a new way to trade on the stock market, invisible to humans, that triggers widespread financial instability along the way.
- Rogue operators use medical-grade AIs (designed to generate life-saving medicines) to develop bioweapons. Weapons of mass destruction become widely available because of the impossible-to-regulate intellectual power of AI.



- Political campaigns employ AI to convince potential voters. The AI generates custom, individually tailored persuasive messages that consistently nudge the reader, over time, into harmful ideological positions, using a style of voice that deeply resonates and feels right to the reader (perhaps written in the voice of a loved one, such as a departed grandmother).
- AI is tuned and integrated for niche businesses at scale. AI becomes effective at eliminating jobs, and 10 - 20% of all positions are eliminated by 2030. Because of dynamics unique to AI and the simultaneous growth of robotics, very few new jobs are created, resulting in mass unemployment and a rapidly widening wealth gap.



News headlines tend to focus on the “lights out” scenario: An AI that develops its own goals that are incompatible with humanity, resulting in a story all too familiar to fans of science fiction. But the scenarios listed above don’t require any major technological breakthroughs. They only require continued development and deployment of AI without thought.

Preventing these scenarios (and more) is called the “AI Alignment Problem.” How do we, as a human species, work to align Artificial Intelligence to what is truly good for human flourishing?



There is no one easy solution, no regulation that will work internationally across all technologists and businesses, no incentive set that prevents negative consequences, and no clear way to deploy AI in a way that helps us become better people.

Instead, it requires all of us, pulling together, towards the best possible long-term outcomes.

Because if Artificial Intelligence doesn't help us to live better lives, it fails.

And so do we.



What you and I can do, as business leaders

In our businesses, how we deploy AI helps or hurts AI alignment.

It is tempting to simply avoid AI altogether. But this is no solution. For unless everyone in business ignores AI forever, then those who race to deploy it will determine the future.

As leaders, you and I may feel uncomfortable with the idea of ever trimming jobs because of AI. The best way to prevent this is to launch into exploring AI immediately. Otherwise, if competitors deploy AI first, and they focus all the gains on cost reduction and job elimination, then the inevitable pressure of the market will force our businesses to eventually follow suit.

Today, we have time. We can choose how to use AI in a way that is good for people and how to share the gains from it. But it requires effort and intentionality.



Questions I have found helpful to consider:

How do I protect and empower my current team?

We can immediately orient current employees towards AI, provide them opportunities to train and encourage them to partner with AI.

This helps them and helps the whole business.

It is clear that teams who creatively engage with AI and reimagine business processes, job roles, and market opportunities, will gain the most rewards and the most control over their futures.



How do I avoid layoffs?

If even a fraction of the promise of AI is realized, it will result in job reduction. However, the more visibility we have on this potential, the easier we can avoid layoffs.

Consider this: if we plan ahead, then efficiency gains from AI would have to exceed a company's annual headcount growth + annual attrition rate, in order to trigger a layoff. If a company has 10% turnover and a 10% growth rate, then it could (theoretically) gain 20% efficiencies from AI each year before even considering layoffs.

The real world is obviously messier, but the point stands: with visibility and planning, even radical gains from AI don't have to threaten current employee livelihoods.



How do I share the gains?

It is demonstrably bad for society when a tiny minority get richer, while everyone else gets poorer.

Simply because western capitalism allows for profit to be amassed by the owners, doesn't mean that is always the best model. In fact, a growing wealth gap is a predictor of coming civil unrest and revolution. A well-functioning society is one where all have the opportunity to be successful.

As business leaders, as investors, and as owners, we do not have to unthinkingly use AI to maximize profit, then to have that profit flow up to executives and owners. Instead, gain from AI can be widely shared.



Western society faced this in the 19th and 20th centuries, with the Industrial Revolution.

As factory work replaced farming for a large percentage of the population, great suffering followed. The first multi-millionaires were surrounded by a suffering, working poor, as so poignantly illustrated by Charles Dickens in his novels. Addressing this was a concerted effort across religious, nonprofit, political, and business leadership.

We will likely need a similar effort for the AI Revolution. The more generously and thoughtfully we use AI in business, the less regulation, taxation, and charitable effort will be needed.



How do I ensure that my business efforts are aligned to human flourishing?

Consider that AI will simply accelerate our activity. It is a powerful force for faster and more. If our businesses are currently harming the environment, AI will accelerate that. If our products currently aren't truly good for the long-term flourishing of our customers, AI will accelerate that. If our employees have meaningless jobs of drudgery, AI will accelerate that.

May you and I let the urgency of coming AI power (of which ChatGPT is just a foretaste) compel us to consider what impact our jobs, and our businesses have on the world.

If everything we do gets accelerated, what will be good? And what is best left behind?

May we have the courage to ask this honestly, then to change what must be changed.



You and I can't slow down the development of Artificial Intelligence.

We can become skilled at implementing it.

We can become thoughtful at aligning it to what is good for those around us.

And we can become the kind of courageous leaders who orient AI, our businesses, and our lives towards wisdom, love, and what is truly good.

REV 12.15.23



Join the conversation.

mstech.com

info@mstech.com

MAINSTAY
TECHNOLOGIES