

How To Think About AI

Schelling AI

Summary

Generative AI is already taking over our world. As these AI models continue to improve, we face immense societal changes. This piece examines the imminent rise of AI Atlantis — a world dominated by highly talented AI agents and robots.¹ We explore why current AI models fundamentally differ from previous technological paradigms, their potential to reshape industries, and the critical choices humanity faces. Everyone must understand how the advancement of AI will rewrite the rules of economics, creativity, and human potential. A few entities will monopolise and control these transformative technologies — this power — unless AI is democratised. Choices made today will determine whether AI amplifies or diminishes human agency and freedoms. This piece is a foundation for further explorations of the topics covered here. It is an urgent call to action that aims to catalyse informed, proactive engagement with AI's remarkable challenges and opportunities.

Contents


Introduction	1
AI Graduates	3
Files to Flows	5
Agents vs. Agency	8
Crossroads	10
Conclusion	12

¹Friedman coined the term "AI Atlantis," referring to a new digital realm of virtually limitless AI resources; Friedman, N. (2024, March). The AI Landscape: Last Year & Years Ahead. In 2024 Abundance Summit. Abundance360, Los Angeles, CA.

Introduction

As things stand, a handful of people will determine the future of AI. They have the potential to control every part of our lives. Annual global AI investment will approach \$200 billion by 2025.² Yet the drastic change we shall soon encounter is understood by few. A new form of intelligence is being created that will transform our collective human existence.

Private AI investment (dotted lines show GS revenue projections*)



Source: Stanford Institute for Human-Centred Artificial Intelligence, Goldman Sachs Research

Let's be crystal clear: humanity standing on the edge of the **most transformative events in our history**. This piece isn't just another tech forecast. It's a *wake-up call*, a *roadmap*, and a *vision of our imminent future* rolled into one.

Here's what you need to know:


- **There Will be an AI Explosion.** 100 billion AI agents and a billion robots. This will soon be a reality.
- **Current AI models? They're Just the Appetiser.** This is a tipping point. AI capabilities will soon accelerate at an unfathomable rate, making today's tech look like Stone Age tools. Most cannot imagine — nor are ready for — the cultural and societal implications.
- **No Industry is Safe.** From software to creatives, finance to pharmaceuticals — AI will change the nature of work and society. Creatives were the last thing predicted to be displaced by AI technologies. Now they're the first.

² Goldman Sachs. (2023, August 1). *AI investment forecast to approach \$200 billion globally by 2025*. Goldman Sachs. <https://www.goldmansachs.com/intelligence/pages/ai-investment-forecast-to-approach-200-billion-globally-by-2025.html>

- **The Cost of Intelligence is Plummeting.** The value of human ingenuity in applying this new resource will skyrocket — it will shortly become one of the most important skills you can have.
- **It's Time for Urgent, Decisive Action.** To harness the benefits and mitigate the risks of this AI tsunami.

This piece will equip you with the mental models and insight to handle the imminent AI wave. You can agree or disagree, but this is the objective reality we face.

Think this is an exaggeration? The global generative AI market will reach \$1.3 trillion by 2032, increasing by over 13x in the next seven years.³



Source: Bloomberg Intelligence, IDC

AI brings with it both unimaginable challenges and opportunities. This document is your crash course on the future.

In the upcoming sections, we'll explore different AI models, their effects on knowledge, and their impact on society and individual control.

We must act now, AI Atlantis is rising.

³ Bloomberg. (2023, June 1). *Generative AI to become a \$1.3 trillion market by 2032, research finds* | Press | Bloomberg LP. <https://www.bloomberg.com/company/press/generative-ai-to-become-a-1-3-trillion-market-by-2032-research-finds/>

AI Graduates

Imagine waking up tomorrow to find a new continent that has emerged overnight. Not a landmass — a digital realm populated by 100 billion AI agents, and a billion robots. All ready to go. Welcome to AI Atlantis: a seemingly far-off utopia which is actually a vivid reality unfolding before our very eyes.

Here's an incredible idea: you can now hire digital offshore talent and robotic workers for a salary lower than what you spend on coffee — yet they have the same skills as an average university graduate. This is a monumental shift that will alter economic and social structures.

If you're saying to yourself, "That's impossible," or "AI will never reach that level," you can stop reading now. We're appealing to reason here, and if you disagree with this mental model, the rest of this document might be a tough pill to swallow.


For those who believe in the iron laws of supply and demand, brace yourselves. We're about to throw Keynesian economics out the window. When you have an almost infinite supply of organised intelligence available worldwide at virtually no cost, traditional rules no longer apply.

We've barely scratched the surface of what's possible with AI. What you're seeing is the approaching (and inevitable) turning point where AI capabilities will ignite, making the current models look primitive in comparison.

Artificial intelligence: Performance on knowledge tests vs. training computation

Performance on knowledge tests is measured with the MMLU benchmark, here with 5-shot learning, which gauges a model's accuracy after receiving only five examples for each task. Training computation is measured in total petaFLOP, which is 10^{15} floating-point operations.

Our World
in Data



Source: Our World In Data **Data Source:** Epoch (2023) **Note from source:** The values for training computation are estimates and come with some uncertainty, especially for models for which only minimal information has been disclosed, such as GPT-4.

If you don't believe the rapid progress in AI development, take a look at GPT-3.5 and GPT-4 in the graph above. The two models were released just under a year apart, but the training compute increased by almost 10x.

Today's AI models are certainly more than just fancy calculators — they're **pure inference engines** (more on that later). Think of them as *Prometheus*, offering humanity boundless creativity and problem-solving abilities.

By using generative AI, mere mortals turn into instant geniuses. Picture having a window into all of humanity's knowledge in your hand. Well actually, there's no need — it's happening right now.

The impact is already cutting through industries. More than half of US businesses have started using generative AI in some way⁴ — PwC, for instance, has invested a billion dollars in training 65,000 US employees on integrating GPT-4.⁵

Here's the glaring issue — jobs. AI automation puts 300 million full-time jobs at risk.⁶

But before you panic, remember this: new AI-related roles are sprouting up everywhere. The adoption of AI will create 97 million jobs by 2025 — across all industries: automotive, energy, healthcare, mining, the list goes on.⁷ The key is adaptability and continuous learning.

The cost of intelligence is plummeting, but the value of human ingenuity in applying this unique resource will skyrocket. Those who excel in AI collaboration will thrive in this new world.

AI Atlantis is rising. The question isn't whether it will happen, but how to navigate its shores. In the next section, we'll look into why these models differ from anything ever seen.

The future is here, and it's wilder than we ever imagined.

⁴ PwC. (n.d.). 2024 AI business predictions. In PwC. Retrieved July 16, 2024, from <https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html>

⁵ Paw, U., & Atkinson, J. (2024, May 29). PwC is accelerating adoption of AI with ChatGPT Enterprise in UK and US and with clients. PwC. <https://www.pwc.co.uk/press-room/press-releases/corporate-news/pwc-is-accelerating-adoption-of-ai-with-chatgpt-enterprise-in-uk.html>

⁶ Hatzius, J., Briggs, J., Kodnani, D., & Pierdomenico, G. (2023). *The potentially large effects of Artificial Intelligence on economic growth (Briggs/Kodnani)*. In GSPublishing. Goldman Sachs.

<https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.html>

⁷ World Economic Forum. (2020). *The future of jobs report 2020*. World Economic Forum. https://www3.weforum.org/docs/WEF_Future_of_Jobs_2020.pdf

Files to Flows

What's been done with modern AI models is the creation of a new type of technology that can understand and generate human language, create art, and solve complex problems. Why exactly are they so revolutionary? Allow us to explain.

These AI models aren't just fancy calculators or "if-then" machines. They can process and generate information in ways that **mimic human cognition**. To top it off, they can do it at a level that's simply spectacular.

Consider that a publicly available large language model (LLM) can train on trillions of words (tokens). If someone had compressed these into a regular zip file, the words would have occupied thousands of gigabytes. Yet, the final model might only be tens of gigabytes in size — stored almost entirely in one file.

The model can do this because it learns patterns from training data instead of just storing it to retrieve later. The AI model generates responses that are right up there with the original data, but never a perfect copy.

It's as if you carried around the Library of Congress in your pocket. Just look at how Tesla replaced 300,000 lines of self-driving code with a single AI model.⁸ This new model's code is 100 times shorter — most of the storage space is dedicated to a single model weights file. This is an immense change to how we handle and process information.

These models are **universal translators**. By taking input in one form, whether it's text, images, or code, they can turn it into another. They can explain complex ideas to a five-year-old or rewrite Shakespeare in the style of a rap battle.

Provide these models with any context, and they can adapt. It's this versatility that makes them so powerful. It's this versatility that makes them so applicable across industries. It's this versatility that makes adaptation essential.

Daniel Kahneman outlines a concept which helps explain how these models work.⁹ He describes two types of thinking:

- **System 1:** Fast, Instinctive, Emotional
- **System 2:** Slow, Logical, Calculated

Today's generative AI models excel at *System 1* thinking. They can quickly generate responses based on patterns they've learnt, much like how you might instinctively react to a sudden noise or complete a familiar phrase. For this reason, they can produce human-like text or images almost instantly.

The fascinating part is, as these models evolve, they're beginning to demonstrate capabilities that resemble *System 2* thinking — something that used to be unique to non-generative AI such as chess bots or spam filters.

⁸ Musk, E. [@elonmusk]. (2023, August 2). *Vehicle control is the final piece of the Tesla FSD AI puzzle. That will drop >300k lines of C++ control code* [Tweet]. X. <https://x.com/elonmusk/status/1686513363495346178>

⁹ Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus & Giroux Inc.

We're moving from a world of *static files* to one of *dynamic information flows*. The world we're stepping into with AI is one where information is constantly updated, combined, and repurposed in real time. This changes everyone's relationship with information.

Tools like ComfyUI are early examples of this shift — people can now create intricate AI workflows where information and tasks flow seamlessly from one stage to the next. Once a parameter is changed, the flow adapts downstream.



Source: @toyxyz. (2023, August 18). X.

This means processing information and generating new ideas faster than ever before.

The result of all of this? A *dramatic* increase in the **velocity of knowledge**. Whether you're an individual, an organisation, or an entire society, these AI models give you the ability to process, understand, and act on information at a pace and scale that has been beyond our reach until now. Ushering in new tools for ideas, and eventually creating at the speed of thought.

The important part of all this is the agency. More choice, more capability, and more power to change our world than any of us have ever had before — by using these tools.

Sceptics might argue, "But these models make mistakes. They can't truly understand or reason like humans do." And they're right — to an extent. These models aren't perfect. They do 'hallucinate', and they don't have the deep understanding that humans do.

...Yet.

They're improving at an astounding rate. Despite their current limitations, when combined with human intelligence, they're already transforming industries and changing how we work and live.

These AI models will *amplify* human potential in ways we're only beginning to understand.

The question isn't whether they'll change the world. The question is: are you ready for the world they'll create?

Agents vs. Agency

Let's talk about power. Not the kind that keeps your lights on, but the kind that's about to reshape our world. We're talking about **agency** — the ability to shape your destiny. And AI is about to supercharge it.

Visualise life as a grand buffet. Before AI, most people were stumbling around with just a fork, maybe a spoon too if they were lucky. But now? AI is handing out Swiss Army knives to everyone. It's not just about having more tools; it's having the right tools for any situation. This is the essence of increased agency.

We're not just talking about AI replacing humans. We're talking about AI-human hybrids — '**Centaurs**'.

A Centaur is more than just a human using AI. It's a symbiotic relationship where both amplify each other's strengths. Human creativity and intuition, combined with AI's processing power and precision. It's like giving Einstein a supercomputer for a brain.

The thing is: Centaurs outperform humans alone (when they're doing tasks within the AI's capabilities). Human-AI teams beat humans working solo¹⁰ — a quantum leap in human capabilities.

So how exactly does AI amplify individual agency? It's a triple threat:

- **Efficiency on Steroids.** AI boosts everyone's ability to get stuff done. Tasks that used to take hours: now take minutes. Complex problems that would have stumped you indefinitely: are now solvable. It's like having a team of expert assistants at your beck and call, every hour, of every day.
- **Time: The Ultimate Luxury.** By automating the mundane and accelerating the complex, AI frees up humanity's most precious resource — time. Having more leisure hours isn't the main thing (though that's nice); it's about having the mental space to focus on bigger challenges.
- **Choices Galore.** AI is expanding the horizon. Suddenly, skills that were once out of reach are accessible. Want to code but never learnt? AI's got your back. Need to understand complex data? AI can break it down for you. It's essentially a passport to a world of new possibilities.

Now, here's where it gets (even more) interesting — and potentially quite scary. This increase in agency isn't just happening at the individual level. Corporations and governments are getting supercharged too.

For corporations, AI means unparalleled efficiency and innovation potential. But it also means they will need fewer human employees. Look at tech companies — they're making more money with fewer people.

Governments with AI could provide better services, make data-driven policies, enhance national security, and provide better transparency and accountability to their people. Except there's a flip side — the potential for complete surveillance and control. Remember, people often describe a government as having a monopoly on violence. Now imagine that entity with AI-enhanced capabilities.

¹⁰ Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality*. Social Science Research Network. <https://doi.org/10.2139/ssrn.4573321>

The AI revolution is like a tsunami, and the water is already receding from the shores. Some are still playing on the beach, oblivious. Others are running for the high ground. But the smart ones? They're already there.

Because here's the truth: you can choose not to engage with AI, to stick your head in the sand and hope it all blows over. But then you're choosing to be left behind. To be out-competed, out-thought, and out-maneuvred by those who embrace it.

The power of AI to enhance human agency is the most significant development in human history since the invention of writing. It's Gutenberg, but bigger. It's the Enlightenment at light speed.

But with great power comes great responsibility (yeah, we're quoting Spider-Man, deal with it). Humanity stands at a crossroads, and the choices made now will shape our future.

Will AI-enhanced agency lead to a new renaissance of human creativity and problem-solving? Or will it exacerbate existing inequalities and lead to alternative (and more) forms of oppression?

The answer isn't written in stone. It's up to us. *All* of us.

In this unexplored world, you could have more power than we ever dreamed possible. What will you do with it?

Crossroads

This is a climactic point in human history. The choices we make today will echo through generations, shaping the very fabric of our existence.


The most crucial decision humanity faces is this: pursue an open, collaborative approach to AI development, or allow it to remain in the hands of a select few? As AI Atlantis approaches, the growing AI divide must not go unnoticed — these technologies will increasingly dictate economic and social power. Schelling believes **democratising AI** is not just preferable — it's *essential*.

Why? Consider this:

- **Diverse perspectives ensure AI caters to all societal needs.** Not just the priorities of tech giants or governments.
- **Open-source development builds trust.** It allows for better scrutiny, minimising the risk of hidden biases, single points of failure, or unintended consequences.
- **Collaboration speeds up innovation.** The collective genius of the global community will always outpace siloed corporate research and development.
- **Broad participation keeps ethical considerations at the forefront.** AI design and implementation must not evade ethics.

The alternative? A world where AI becomes the new oil, controlled by *digital barons* who monopolise the most powerful tool humanity has ever created.

Is that the future you want?



Source: Irene Solaiman. *The Gradient of Generative AI Release: Methods and Considerations*.

Measures must be taken quickly, but thoughtfully. Here's what must be done:

- **Invest in AI Education at All Levels of Society.** There cannot be informed debates about the future if most people don't understand the basics.
- **Develop Strong Ethical Guidelines and Regulations.** For both AI development and deployment. These need to be flexible enough to keep pace with rapid advancements.
- **Rethink Our Economic Models.** To address the inevitable job displacement and inequality. Informed citizens and policymakers need to make data-driven decisions together.
- **Promote International Cooperation.** To ensure the benefits of AI are distributed globally. AI Atlantis should be a rising tide that lifts all boats, not just the yachts of Silicon Valley and Wall Street.

This is the takeaway: Although the era of generative AI has just started, it's reaching a level of performance that's world-changing. It will reshape our societal structures. Society needs to act fast to mitigate the risks and capitalise on the opportunities to solve these pressing issues.

The power to shape our AI-driven future is in our hands. Will it be used to amplify humanity or diminish it? Will you help create a future where technology serves as a bridge to greater understanding, innovation, and prosperity for all?

The choice is ours. Let's make it count. If AI is going to shape the future of humanity, then all of humanity needs a say in shaping AI.

Conclusion

AI is here. It's accelerating faster than most can comprehend. Again, we're not talking about some distant future; we're talking about now. Choices today will determine whether AI becomes our greatest ally or our most formidable challenge.

What we want you to take away is:

- **AI Atlantis is Rising.** 100 billion AI agents and a billion robots at your fingertips — all as competent as university graduates — is our imminent reality.
- **Current Models are Just the Beginning.** We're on the cusp of an explosion in AI capabilities that will make today's tech look primitive.
- **The Fundamental Nature of Work, Creativity, and Problem-Solving is Changing.** Industries ranging from software to healthcare, education to entertainment, are on the brink of transformation.
- **Individual Agency Will Soar.** AI will provide people with more choices, time, and capabilities, than ever before.

So what do we all do? Society needs to act in a way that's fast and considered:

- **Invest in AI Education and Literacy.** For everyone. Your grandma needs to understand this as much as your software engineer neighbour.
- **Push for Open, Collaborative Approaches to AI Development.** Tech giants or governments must not monopolise this power.
- **Rethink Our Economic Models.** The old economic paradigms are breaking down. New frameworks that can handle the reality of AI Atlantis are needed.
- **Develop Standards for AI Ethics and Regulations.** Ensure AI enhances human potential rather than diminishes it.
- **Encourage Global Cooperation.** AI doesn't respect borders. There needs to be a combined, international effort to tackle the challenges and make use of the opportunities.

The reality is, AI is going to change everything. Absolutely everything. The way humanity works, creates, and solves problems. The way we all live.

There is still a choice. Together, we can **shape this change**. We can guide it toward a future where AI enhances human creativity, solves our biggest challenges, and reveals incredible possibilities.

Or we can sleepwalk into a future where AI widens inequality, erodes privacy, and leaves many behind.

The choice is all of ours. But we need to make it now. Not next year, not next election cycle. **Now.**

So, what's it going to be? Will you sit there and watch AI take over our world? Or do you want to help form the future?

Schelling has made its choice: shaping an AI future that empowers humanity. Are you with us?

Let's make AI Atlantis a paradise, not a dystopia. We must ensure AI is made by the people, for the people, with the people.