THIS IS A SPECIAL EXCERPT **BUY THE BOOK ON AMAZON**

# Your best years can last longer than you think.

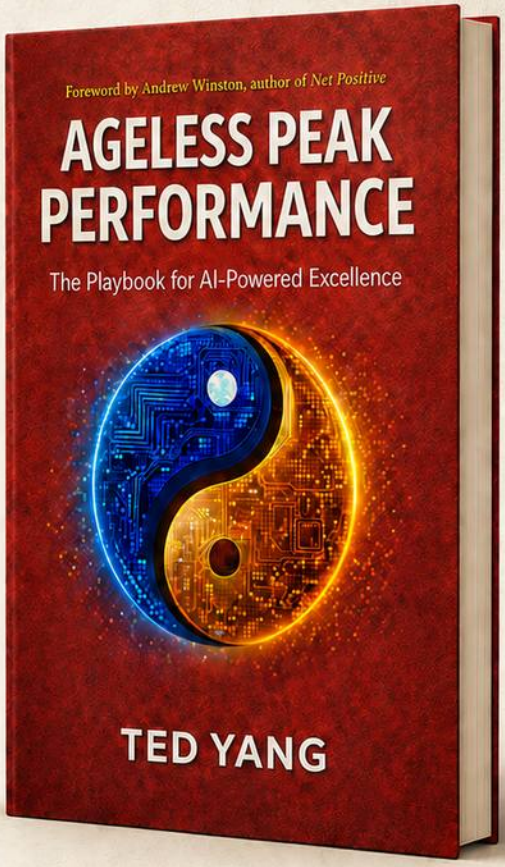


## Ageless Peak Performance

**The Playbook for AI-Powered Excellence**

Learn how AI speeds up what slows with age and amplifies what gets stronger

**eBook Launch Price Only $9.99**

**Get the free Devil's Advocate workflow at AgelessPeakPerformance.com**

by Ted Yang

# Ageless Peak Performance

The Playbook for AI-Powered Excellence

Ted Yang

A ProductMind Book - productmind.co

# Praise for Ageless Peak Performance

"The workforce is changing faster than ever. The leaders who thrive will be those who learn to understand and manage AI. Ted Yang has written the practical, no-hype guide to doing exactly that without losing what makes us human." **— Ned Lamont, Governor of the State of Connecticut**

"*Ageless Peak Performance* reframes learning as a lifelong advantage in the age of AI. Ted Yang shows how technology, used thoughtfully, helps people grow, adapt, and contribute at their highest level at any stage of life. Not something to fear. Something to master. " **— Miguel Cardona, 12th United States Secretary of Education**

"I approached artificial intelligence initially with skepticism. This book changed my mind. Ted Yang shows how AI used wisely can be a great equalizer, bridging opportunity gaps for people everywhere, whether in Addis Ababa or Washington, in boardrooms or classrooms. The AI revolution will either widen the gap between haves and have-nots, or close it. Throughout history, access to tools has determined who rises. Ted Yang has written the guide to wielding this one—not for the privileged few, but for anyone with the will to learn.**" — HIH Prince Ermias Sahle-Selassie, grandson of Emperor Haile Selassie I**

"Ted Yang does an amazing job telling easy to understand stories about a complex and critical technological shift happening in front of and around us all. His great insights into what to do and how to engage are worthy of your reading time. And his story telling style makes it enjoyable. Learn and enjoy!**" —Nick Donofrio, IBM Fellow and Former EVP Innovation and Technology, Member, National Academy of Engineering, Fellow, Royal Academy of Engineering, Board of Directors, AMD.**

# Contents

# Prologue

## Ten Years is not Enough

What if your best years will never be behind you? What if AI could extend the peak of your performance from the 10 to 15 years it is today, to 50, 60, 70, or more years? What if you could truly have ageless peak performance?

This is not a prompt book. This is not an instantly obsolete manual for LLMs. And it's not a hype manifesto about human replacement.

Instead, I will show you how to reach AI flow, to integrate AI so deeply into your work that you perform effortlessly at your peak. And how to extend that peak into a plateau lasting decades.

I was exposed to generative AI too early. I have a news outlet, Daily Voice, and a vendor came to us saying they could automatically generate copy for our ads. It didn't work and was full of hallucinations.

Months after that, articles started to come out about people falling in love with generative AI prototypes at Google. Groan. Was this going to be another blockchain?

I was wrong. My aha moment with generative AI was when I realized that it wasn't just software, it was an intellect-enhancing tool. AI quickly and correctly summarized a technical paper that I had been struggling to read. I could do things better and faster than I'd ever imagined.

It may seem like AI is everywhere, but most people have not really experienced the power of generative AI. Show them how it can conduct research,

summarize an article, or rapidly generate content, and they will be completely floored. Then they will start imagining a future together with AI.

This book will help you realize that future. A future where you perform better and at your peak every day because AI is amplifying your strengths (judgment, pattern recognition, strategic thinking), making up for your weaknesses (processing speed, working memory), and working as the ultimate thought partner.

You might wonder if this book is just for techies, just for younger people starting their career, or just for the experienced. The answer is no. Everyone can use a guide to the AI superpowered future.

A recent study of highly-experienced open-source developers revealed that when using AI coding tools, they felt 20% faster, but were actually 19% slower. They didn't understand the fundamental principles of AI peak performance and being technical and AI-savvy did not protect them. This methodology would have worked for them, it works for an almost 50-year-old me, and it will work for you.

By picking up and reading this book you are at least AI curious. Great start!

Don't be an AI skeptic. If you choose not to use AI, or only use it timidly, you will be blown away by those who do. Already, we are seeing employers that insist that their new hires must have generative AI skills. This is an explicit acknowledgment that this technology magnifies human capability, and what employer doesn't want that?

Soon we will see that it is not just about your ability to use generative AI; it will be the fluidity with which you use it that will matter to employers. Is it truly augmenting you and everything you do? Or are you just using it on the edges?

In this book I will use the terms: generative AI, LLMs, and AI interchangeably. AI is a much larger topic than just generative AI but it is LLMs, transformer-based AI, that makes peak performance possible.

In Part I: The New Reality, I'll take you through the brutal reality of cognitive decline and why traditionally you perform at your peak for so short a time. Then

we'll explore what's different with AI as an intellect-enhancing tool and how it creates a peak performance plateau that can last.

In Part II: The Principles, we'll go over four critical principles for AI adopters. These principles are evergreen and work for any LLM or whatever may come next. If you internalize and master these principles, you'll avoid the same mistakes that those expert coders did. More on that in Chapter Six.

In Part III: The Superpowers, I expand upon the principles and show you how to develop five distinct superpowers that each use AI in a novel and transformative way.

Finally, in Part IV: The Peak, we expand our discussion to speak about reaching a peak with a team and then we'll go through a 90-day implementation plan to take you from merely AI curious to AI adopter, someone who has reimagined their relationship with cognitive work itself and performs at their peak. Finally, I'll go through the mindset you need to sustain that peak.

Before we get to superpowers and principles, let's get on the same page with the problem we all face. How can we achieve ageless peak performance? Why don't we have it today? And what can be done about it?

Thank you for coming on this journey with me. Onwards to Chapter One!

# Chapter Four

# Principle One - Treat AI Like an Intern, Not an Oracle

> "When done right, AI can be incredibly beneficial... [do] not blindly rely on AI platforms." – Judge Kelly Rankin, Wadsworth vs Walmart et al

In 2023, in what may be the first, or certainly one of the earliest, cases like it, two lawyers were fined $5,000 and sanctioned by the New York Federal Court in Manhattan for using ChatGPT to do their research in a brief submitted in a lawsuit. The problem? Several cases that were cited in the brief did not exist, as they were AI hallucinations.

This wasn't an isolated incident. A 2025 Stanford study found that even specialized legal AI research tools hallucinated 17-33% of the time. When asked general legal questions, these LLMs invented non-existent court cases with

plausible names and details. "Hallucinations in Legal RAG Systems" - Dahl et al. (Stanford, 2025)[1] .

The lawyers treated ChatGPT as an oracle, as a wise source of knowledge, which they could use in their case. But they should have treated it like an intern. Someone who is smart, willing to help, and full of energy, but fallible.

The mistake wasn't in using AI. It was how AI was used. It was not checking the work that it produced, as they would've done for a junior employee, and then claiming it was their own and even doubling-down in denials after it was discovered.

We've been conditioned by Google and other search engines to get answers from them like they are the Oracle at Delphi. Yes, we had to dig through links and do additional work after we put in keywords, but the key idea was that there was such a thing as a source of truth, and Google was its index. Most people realize now that Google isn't always the truth, given search engine optimization and other techniques, but the overall idea that Google gives you truth persists and has been translated over into the world of AI, with sometimes disastrous results, as in the NY case above and the one where the quote from this chapter is taken.

The first principle for using AI for peak performance: treat AI like an intern, not an oracle.

What does this mean? Like an intern, you must give clear, specific instructions. Similarly, you should check the work you receive, especially initially, and then iterate and refine just as you would for a junior employee to get a better answer. You wouldn't use your intern to make slides for a presentation to the CEO and then pass it off as your own work without checking. And you shouldn't do it with AI either.

---

1. https://dho.stanford.edu/wp-content/uploads/Legal_RAG_Hallucinations.pdf

## Why is Generative AI Like an Intern?

### Trait 1: AI is Non-deterministic

If you ask an intern the same question twice, they'll probably give you different answers, especially if it's on a subject they don't know and they're tap-dancing and saying what comes to mind so as not to look bad in front of their boss. AI doesn't get nervous, but it babbles.

Traditional software and traditional coding languages are deterministic. That means that if you give them the same input and then run the code, you will get the same output. Given a starting state, and a set of functions, algorithms, or procedures that you want to run against that state, you will arrive at the same end state.

Try doing that with ChatGPT. Literally run the same prompt twice, and I'm betting you'll get different results back. This could be something small, like different phrasing, but it could also be a completely different answer the second time, or even a contradictory answer.

One way this is obvious is if you ask ChatGPT to summarize a meeting. Again, try this one yourself. You'll end up with three different summaries that will all be correct, mostly, but will emphasize different parts of the meeting, even though the input is the same meeting transcript.

This isn't a bug, and this isn't something that the makers of LLMs layered on to make generative AI more human. It is fundamental to the way in which they were constructed. Large language models use probabilistic models that are trained upon the corpus of human knowledge as exists on the internet. The core thing that's happening in any LLM is a prediction algorithm that is trying to use all of its inputs and all of its prior outputs together to predict what word should come next. That probabilistic calculation results in different answers every time.

What this means for you is that the result that comes out, like the answer from a helpful but not-so-knowledgeable intern, can differ each time you ask it.

That means that you need to be on your toes before you blindly take any outputs from AI as truth from the mountaintop.

### Trait 2: AI Tries to Please and Hallucinates Confidently

AI doesn't like to admit it doesn't know something. In fact, it would rather give you a wrong answer than admit it doesn't know. Coming back empty-handed is a no-no. Just like the intern you sent to go get your lunch who brought back the wrong thing because the deli ran out of bologna.

What happened in the federal court case above, and actually many times in documents submitted in courts, is that generative AI made up citations. It made up citations in a way that is plausible, looking for facts, papers, rulings, what have you, that justify the case it's trying to make, or that you're trying to have it make, but which are fabrications.

There may also be statistical falsifications, not just a citation, but obscuring the truth with incorrect numbers. And the hallucination can also be not one fact, but a chain of facts.

The problem is AI doesn't express any doubt when it does this. It gives you a whole chain based upon false premises in confident-sounding language.

Your intern is desperate to prove its value and win a full-time job, or at least an offer. Similarly, Claude and other AI systems have been built in such a way to try to always add value. After all, who wants to pay for a tool that says it can't do something?

What you can do about this is ask for verification. Use multiple interns to check each other in the real world, and in the AI world, do the same by asking AI to check itself, and then interactively challenge the results that come back.

If you can understand this simple technique of treating AI like an intern and verifying its results, you will avoid a large percentage of potential disasters.

### Trait 3: AI Is Not Like a Human Expert

Let's be honest about the expertise level of large language models. If you were to ask AI, you might think that it is an expert on the level of the brightest possible human minds in every discipline; it knows a lot of things. But you are not hiring the AI version of a human expert. What you're hiring is an intern who is bright but is really a generalist. Not only is AI not an expert greater than other humans, often it is over-confident.

This is obvious given the way that large language models are trained. These models are an average of all of the expertise that is fed into them; even if mechanisms are used to refine that expertise, what you arrive at is only above average and not exceptional.

A January 2026 study comparing the creativity of over 100,000 humans against the latest LLMs confirmed this. They found that "[AI was] largely surpassed by the aggregated top half of human participants" and was soundly beaten by the top 10% of humans (aka human experts) "Divergent creativity in humans and large language models" - Bellemare-Pepin, et al. (Nature, 2026)[2] . AI is trained on a lot of facts but because of its inherent probabilistic model, it is averaging down the best thinking with mediocrity.

The converse is also true. In cases where AI is not trained with a lot of information, its expertise is often flat out wrong.

Contrast this with human experts. Human experts are experts not just because they know a lot of facts but because they can adapt what they learned within a new context. They also know the limitations of their own knowledge. In other words, human experts have expert judgment.

Don't get me wrong, AI expertise is still incredibly powerful. Given that you yourself are not an expert in many topics, AI can get you up to speed very

2. https://www.nature.com/articles/s41598-025-25157-3

quickly. See the superpower of Learn Anything to find out more. Even for areas that you are an expert in, AI can help you think more clearly.

Try this for yourself based upon the domain in which you yourself are expert, for example where you have built your career. I did a career in finance but am notoriously bad at trading my own portfolio despite having worked with the best in the world. It would be nice for AI to automate and do a lot of that work for me, but literally asking it to do that is a terrible idea. What it comes back with, if you ask it something super open-ended, is something that nobody should use for investing. But it helps start the conversation and to reawaken some of my old brain cells. Then we can work synergistically at it together.

## How to work with your AI intern

So how best do you work with your AI intern? Here are some good techniques.

### Technique 1: Be Clear and Specific

The first is to be very clear and very specific. If you are vague, you will get bad results. Don't say, "help me invest my portfolio." Say, "I'm looking to invest my portfolio and retire in the next 20 years. I'm interested in broad market exposure, specifically to US equities. Give me suggestions on how I might do so, with links to resources to explain why."

If you were to give your intern a broad instruction, they might get it right, but more likely they will get it wrong and require more iterative loops with you to zero in on the answer. Giving clear and specific instructions to AI gets you more quickly to better results.

### Technique 2: Structure Your Reasoning in a Chain

This technique is popular for writing prompts, and is an organized way of thinking that works both with people and machines.

Start a chain of reasoning that gets you closer to the answer without making broad leaps that can magnify hidden assumptions, mistakes, hallucinations, and other problems. This is analogous to showing your work. You don't have to trust how AI got to an answer because you can see the intermediate steps of its logic.

### TV Show Example – Without Chain of Reasoning

Let's try a trivial example many of us can relate to. I asked generative AI what I should watch on TV tonight. What comes back is, "you should watch the Real Housewives of New Jersey."

I've never seen the Real Housewives of anywhere (no offense to those who like it, including my friend who helps produce it), so this answer isn't even close to good but it is what I deserve for asking something so vague.

### TV Show Example - With Chain of Reasoning

Now ask what to watch tonight, but start with constraints and ask it to show its reasoning of why a recommendation might be good. The show needs to be 30 minutes or less as I have to get to bed early. I want a comedy to brighten my mood. Other comedies I like are: The Office, Parks and Rec, and Brooklyn Nine-Nine.

First, it verifies my constraints. Second, it indicates that I like mockumentary/workplace comedies. Based on these preferences, it suggests Abbott Elementary but if I had more time to try Ted Lasso.

It's possible that I'm not particularly into mockumentaries even though three came to mind. In which case, I can go back, challenge that assumption, and get a different result.

You could imagine doing this with your intern the first time you show them how to calculate cash flows on a spreadsheet. Show the steps iteratively, rather than jumping to a completed spreadsheet and trying to explain its complexity in one shot.

### Technique 3: Know When You Need an Expert

Finally, another technique is to know when you need to go with an expert rather than your AI-powered generalist. If you want someone who's an expert in San Marino tax law, you don't go to your intern and ask them to figure it out for you. You find the senior partner and you ask her.

When the stakes are high and the consequences of being wrong massive, use AI to scaffold your knowledge but don't try to have it take the place of a human expert. Ask AI to help you connect with them.

Don't try to make your intern magically into an oracle. You might think it works, but it will fall apart at exactly the wrong time.

Following this principle may sound a bit pedantic, and a bit annoying to some users. But what we're trying to do when we use AI is get good results repeatedly. This is so you can use it confidently to reach your peak performance plateau.

AI was built in a way that mirrors the brightest intern you've ever had, that also works 24 hours a day, and doesn't complain or ask for a promotion.

A great intern or a great employee can act as an extension of the boss, enabling that boss to have much more leverage in the work that she needs to do every day. That's how generative AI gets you peak performance.

This metaphor isn't about limitation. It's about unlocking potential.

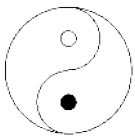

## Try This

Think back to when you were an intern or were responsible for interns. If that hasn't happened to you, think about when you were newly hired into your first job.

What's the best advice you could give yourself on how to perform better at the job? Now generalize that advice and keep it in mind as you prompt generative AI.

THIS IS A SPECIAL EXCERPT **BUY THE BOOK ON AMAZON**

# Your best years can last longer than you think.

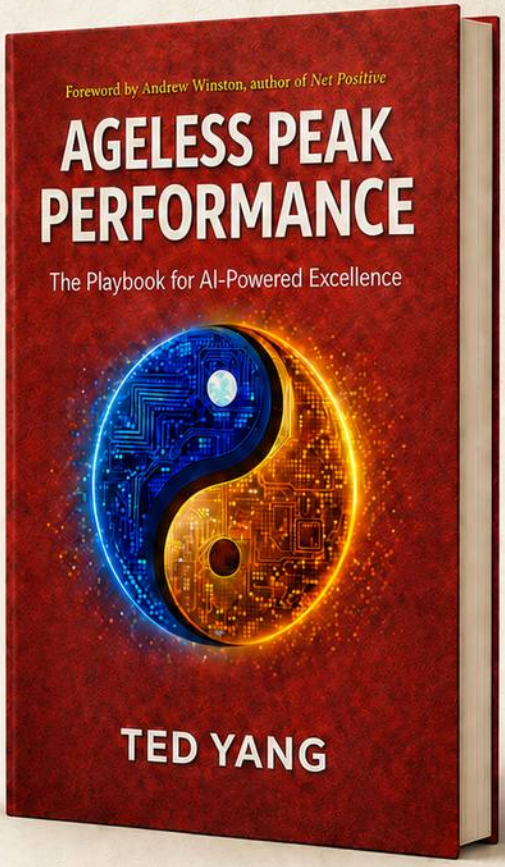


## Ageless Peak Performance

**The Playbook for AI-Powered Excellence**

Learn how AI speeds up what slows with age and amplifies what gets stronger

**eBook Launch Price Only $9.99**

**Get the free Devil's Advocate workflow at AgelessPeakPerformance.com**

by Ted Yang