


THE 38X AI COMPANY

HOW TO TURN YOUR EXPERTISE INTO AN
AI-POWERED BUSINESS THAT GETS
SMARTER EVERY DAY



KEITH RUMJAHN

KEITH RUMJAHN

The 38X AI Company (First chapter)

*How to Compound Your Expertise With AI and Build a
Business That Gets Smarter Every Day*

First published by Rumjahn Limited 2026

Copyright © 2026 by Keith Rumjahn

All rights reserved. No part of this publication may be reproduced, stored, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise without written permission from the publisher. It is illegal to copy this book, post it to a website, or distribute it by any other means without permission.

First edition

ISBN (paperback): 978-988-77649-6-0

ISBN (hardcover): 978-988-77649-8-4

ISBN (digital): 978-988-77649-9-1

This book was professionally typeset on Reedsy.

Find out more at reedsy.com

Contents

<i>Before You Begin</i>	v
1 Chapter 1: I Thought AI Made Me Superhuman	1
<i>About the author</i>	14

For my wife, who has always believed in me and made me want to become a better man. I would not be who I am without you.

For Hunter and Tyler, I hope this book reminds you that an idea can become real. Never let anyone convince you that what you want to build is impossible.

And for Micah and Conrad, two of the smartest people I know. I hope you realize how much I have learned from you. I have no doubt you are destined for great things.

Before You Begin

If you're an entrepreneur, solopreneur or business operator already using AI, this book is for you.

Maybe you use ChatGPT or Claude every day. You've collected prompts, experimented with agents, automated a few tasks or even built something with AI. You know the technology is powerful.

But you're not looking for another list of AI tools. You want to know how people who are genuinely good at this turn AI into better business results.

And you may have the uncomfortable feeling that you're still only scratching the surface.

This book is about closing that gap.

It doesn't matter whether you're still chatting with AI or already running agents inside your business. By the end of this book, I want you to understand the path from using AI as a tool to building a company that can become more capable from everything it does.

I use my own businesses throughout the book because they're

the experiments I know best—the successes, failures and embarrassing mistakes where I could see what actually happened. **The stories are mine. The problems they reveal are yours too.**

This isn't a memoir about my AI journey, and it isn't a manual for this month's newest tools. The tools will change too quickly for that. My goal is to shorten the path between where you are now and what great AI use can eventually look like.

You won't build the final system in a weekend. I didn't either. But if this book can save you months of automating the wrong work, chasing tools you don't need, manually moving context around and repeatedly relearning the same lessons, it will have done its job.

The Road to a 38X Company


I didn't arrive at the way I work with AI today by finding one magical tool. I got there in stages, usually by building something that seemed sophisticated and then discovering another place where the system still depended on me.

Looking back, the path is surprisingly simple. You don't need to build everything at once. You need to understand where you are, identify what is missing and strengthen the next part of the system.

This is the road we'll follow through the book.

THE ROAD TO A 38X COMPANY

Build a business that gets smarter from everything it does.



Each cycle begins smarter than the last.

WHERE AM I TODAY?

My strongest steps: _____

My weakest / biggest bottleneck: _____

The next step I will strengthen: _____

Don't worry about memorizing this yet. Each chapter will make

one part of the roadmap concrete. By the end, we'll connect the pieces and you'll see why the last step—REMEMBER—changes the first. That's when the road becomes a loop.

At rumjahn.com/downloads, you'll find the AI Fluency Assessment, the Build Your First 38X Loop tool, and the latest resources that accompany this book.

If you want to go further and actually build the systems I use—including my agent setups, prompts, code, workflows and implementation guides—I keep those inside my paid community. That's also where you can ask questions when something inevitably doesn't work the way the tutorial said it would.

You can find the current community at:

rumjahn.com/community

Let's start with the first question on the roadmap: what deserves execution?

1

Chapter 1: I Thought AI Made Me Superhuman

1 / THINK — *What deserves execution?*

For most of my entrepreneurial career, I accepted a simple rule: if you wanted to build a bigger business, you needed a bigger organization behind it.

I had founded two venture-backed startups, sold one and taken another public. I had raised more than twenty million dollars, built an international team of forty people and launched products used by millions. I had even partnered with Marvel to launch a superhero game. I knew how to scale the traditional way: you raised money, hired talented people and assembled an organization capable of doing things you could never do alone.

I also knew what that kind of scale could cost.

Before one major launch, I remember sitting on my bed at eleven o'clock at night. For the previous two weeks, I had barely

been present with my family. My wife and I had been fighting about how much I was working, and I kept explaining why this particular launch mattered. The opportunity was important, the team was depending on me, and we had a number we needed to hit. There was always a reason.

There was something absurd about it. I had started businesses partly because I wanted freedom, yet every level of success seemed to require more of me. More customers created more work, more work required more people, and more people created more management. Greater ambition required more capital and coordination to keep the whole thing moving.

I don't blame venture capital for that experience. It's the right fuel for certain businesses, and there are founders who genuinely love playing that game. What bothered me was that I had somehow built the opposite of what I originally wanted. For years, I assumed this was simply the physics of entrepreneurship: if you wanted more output, you needed more resources.

Then AI made me wonder whether those physics were changing.

The app I shouldn't have been able to build

For years, my iPhone had been quietly collecting information about me through Apple Health: steps, workouts, heart rate, sleep and almost a decade of other data. I wanted to get all of it out, explore it properly on my Mac and eventually use AI to help me understand patterns buried across those years.

There was an obvious problem. I hadn't seriously written code in years, and what I imagined wasn't a simple webpage. I needed software on the iPhone, software on the Mac, a way to move and structure the health data and eventually an AI layer on top.

The entrepreneur in me automatically translated the idea into resources. How many engineers would this require? How long would it take? Did I need a designer? Was the opportunity large enough to justify the investment? Those were the questions I had spent my career asking.

This time, instead of answering them, I tried building it with AI. Within about two days, I had something working.

It was crude, and parts of it barely deserved to be called a product, but the quality of that first version wasn't what excited me. The distance between an idea in my head and functioning software had collapsed. Instead of spending weeks debating whether the project deserved engineering resources, I had something real enough to put in front of people.

I shared the early project online, and within a day roughly a hundred people had shown interest. That didn't prove I had discovered a business, but it was enough evidence to keep going. Eventually I released a buggy version on the App Store, and on the first day a stranger paid for it.

The amount was trivial compared with businesses I had built before, yet I remember how exciting that payment felt. Someone I didn't know had encountered a problem from my own life, found the thing I had made and decided it was useful enough to

buy. What amazed me wasn't the money so much as the speed with which an idea could collide with reality.

A few years earlier, the engineering effort alone might have stopped me from trying. Now I could have an idea, build something and discover whether anybody cared before the old version of me would have finished deciding whether to start.

For most of my career, ideas had been cheap and execution expensive. Every new product competed for engineers, capital and management attention. Even a promising idea could die because it wasn't promising enough to justify taking resources away from something else. AI was changing that equation.

Suddenly I could build software, create websites, research markets, produce content and test ideas at a speed that would previously have required a team. After years of building companies by adding people, I suddenly had access to something that felt almost like unlimited labor.

For a while, I found this intoxicating.

When you can build almost anything

If an idea occurred to me, the cost of finding out whether I could make it had become so low that I often just made it. I built apps, agents, automations, content systems, websites and experiments. AI was extraordinarily accommodating. It never complained that I had changed direction again, told me engineering capacity was full for the quarter or asked whether this new project was more important than the five things I had

started the previous week.

I thought AI had made me superhuman. What I didn't realize was that it had also removed one of the constraints that had quietly protected me from myself.

When building something required several people, months of work and a meaningful amount of money, an idea had to survive difficult questions before we committed to it. Was the customer problem real? Did we understand it? Was this more important than everything else competing for the same resources? Could the opportunity become large enough to justify the investment?

Those constraints undoubtedly killed some ideas that deserved to exist, but they also forced discipline. Now I could skip the argument and have a working version by the weekend.

I began confusing the ability to build something with having a good reason to build it.

One app finally made that impossible to ignore.

The fifty-dollar app

After the health product worked, I started looking for other opportunities. I found a type of app that appeared to be making money, and it didn't look particularly complicated. Since AI was now helping me build software faster than I ever had before, I decided to make my own version.

Five days later, the app was live. People downloaded it, but

nobody paid for it. That should have made me question the original idea more seriously. Instead, I did what AI now made incredibly easy: I kept improving the product.

I changed features, adjusted the experience and tried new ideas. Every iteration was inexpensive enough that stopping felt harder to justify than doing one more experiment. Six months later, I had a much better product that, after roughly a year and a half on the market, had generated around fifty dollars.

The result bothered me because I couldn't blame execution. Both products had been created by the same person using roughly the same generation of AI tools, and in both cases I had turned an idea into functioning software remarkably quickly. Yet their beginnings were completely different.

The health app started with a problem I understood personally. I got something into the world quickly, and reality almost immediately started giving me evidence that other people cared too. The second app began with a much weaker insight: somebody else seemed to be making money doing something similar.

When reality failed to validate that assumption, I didn't question the direction strongly enough. I used AI to execute harder.


That was when I began to understand the other side of abundant execution.

The infinite factory

Imagine somebody gives you a factory capable of manufacturing almost anything you can describe. It runs twenty-four hours a day, adding another production line costs almost nothing, and you can manufacture twenty variations by breakfast, redesign the packaging before lunch and produce a hundred advertisements by dinner.

For an entrepreneur, this sounds like a superpower. But the factory has one serious limitation:

It has no idea what it should manufacture.



When you can build anything, deciding what to build becomes the bottleneck."

If you ask it to produce something nobody wants, it will do that brilliantly. When sales disappoint, the factory won't become discouraged. It can create more versions, add more features, redesign the packaging and produce another thousand advertisements. You could spend months making the factory increasingly efficient while avoiding the much more uncomfortable question of whether you should be making the product at all.

That's essentially what had happened with my two apps. AI had given both ideas access to extraordinary production capacity, so the difference wasn't whether I could build the software. The difference was what I had pointed the machine at.

The factory had solved enough of the production problem that something unexpected happened: **the decision about what to produce became more valuable.**

I had assumed that as AI became capable of doing more work, human judgment would matter less. Instead, I discovered the opposite.

AI makes execution cheaper. That makes bad judgment more expensive.

When execution was expensive, a bad idea often died because nobody had the time, money or people to build it. Now the infinite factory will happily manufacture the wrong thing, improve it, market it and produce a hundred variations before you've seriously questioned whether anyone wanted the first one.

The dangerous thing about AI isn't only that it can make bad work faster. It can make the wrong work look incredibly productive. There is always another feature to add, another campaign to launch, another competitor to analyze, another landing page to redesign or another hundred pieces of content to generate. You can remain extraordinarily busy while moving almost nowhere.

The same problem appears outside software. An AI content system can publish ten times as many articles without producing ten times as many customers. An agent can complete every step you gave it and have almost no effect on the business. A perfectly functioning automation can still be a failure because the instruction itself was wrong.

AI was becoming extraordinarily good at producing outputs, but outputs and progress were not the same thing.

When execution gets cheaper, judgment gets more valuable

Eventually I found a question that exposed many of my worst AI projects:

If this system worked perfectly for the next year, what would actually become better about the business?

It's an annoyingly simple question. Sometimes the answer is obvious: a customer problem disappears, more qualified prospects book meetings, conversion improves or a process that consumes twenty hours every month no longer requires you.

Other times, the question reveals that you've built an impressive machine without being able to explain what winning looks like. In those cases, the AI isn't necessarily failing. You may have pointed it at the wrong target.

This is why I no longer think the most interesting question is simply what AI can do. That question becomes easier to answer every month. The harder question is what deserves this much execution.

Volume still matters. If you want to become good at making YouTube videos, for example, you probably need to make a lot of videos. Your hundredth video will contain lessons your first video couldn't possibly have known. Repetition creates skill and generates evidence you can't get by sitting in a room thinking about YouTube.

The mistake is assuming that more execution is always the answer. Sometimes you need more repetitions because the activity is valuable and you're still learning how to do it well. Other times, the evidence is telling you that the direction itself is wrong. Knowing the difference can be worth far more than making the machine faster.

That realization changed the way I used AI. Instead of asking only how it could help me execute my decisions, I started asking whether it could help me make better decisions in the first place.

AI doesn't always need to give you the answer

Some of the most useful conversations I have with AI now don't produce anything tangible at all. I don't leave with an article, an app, a research report or an automation. Sometimes I leave with a better decision, and that can be worth far more.

I had an executive coach before AI became capable of playing this role, and some of our most valuable conversations didn't involve him telling me what to do. He asked questions. Why was I doing this? What evidence did I actually have? Was I solving the real problem or the comfortable one? What was I avoiding?

Often the answer was already somewhere inside me. I just needed someone to keep asking questions until I admitted it.

I now use AI in much the same way. If I'm uncertain about the direction of a business, a product or even this book, I sometimes tell it not to give me advice at all. Instead, I ask it to interview me one question at a time, challenge my assumptions and make me explain why I believe the work matters.

This book itself changed direction several times through that process. I would think through an idea with AI, put the manuscript in front of real people, hear what they actually thought and then return to AI with evidence I didn't have before. Sometimes the answer was to improve what I had written; other times it was to reconsider what the book was trying to do.

That is a very different relationship from treating AI as a junior employee waiting for tasks. The junior-employee metaphor

assumes you already know what should be done: you decide and AI executes. But some of the highest-value work in a company happens one level above execution, when you ask whether you're doing the right thing at all.

That was the first major change in how I thought about AI. I had started by seeing intelligence as additional labor—something that could write, research, code and produce on my behalf. The more capable it became, the more work I imagined giving it.

Eventually I understood that the opportunity was much larger.

AI hadn't eliminated the bottleneck. It had moved it.

Try this

Before you automate another task or build another AI worker, choose one thing you're currently spending significant time on and give AI this instruction:

Act as my executive coach. Do not give me solutions yet. Interview me one question at a time about why I'm doing this work, what business result I expect it to create, what evidence I have that it's working, what assumptions I'm making, and what I might be doing instead. Challenge my answers until I can clearly explain why this deserves my time.

At rumjahn.com/downloads, you'll find the AI Fluency Assessment, the Build Your First 38X Loop tool, and every prompt from this book in one place. I'll also keep the resources there updated

as the tools and models change.

Then answer honestly. The goal isn't necessarily for AI to give you a better way to automate the work. You may discover that the most valuable thing it does is help you realize the work wasn't worth automating in the first place.

Because as execution gets cheaper, **deciding what deserves execution becomes more valuable.**

About the author

ABOUT THE AUTHOR



Keith Rumjahn is an entrepreneur, builder, and AI practitioner based in Hong Kong.

He has built two companies, sold one, and taken another public. Along the way, he raised roughly US\$20 million and grew an international team of about forty people. From the outside, it looked like success. Inside, Keith had become the person every important decision eventually came back to.

That experience changed what he wanted from business.

After more than a decade without a month-long holiday, Keith stepped away, helped install a successor, and began rebuilding around a different question: Could a company create real value, make money, and improve without consuming the founder's life?

Today, he uses AI to operate a portfolio of apps, digital products, content systems, analytics loops, and automated workers. These are not classroom demonstrations. They serve real customers, encounter real failures, and produce evidence that changes how the next cycle runs. An early public app repository reached more than 450 GitHub stars. A measured set of 227 Threads posts generated more than 2.14 million views. His systems have also failed in revealing ways: automations have reported success while delivering nothing, dashboards have contradicted reality, and polished products have broken when customers used them.

Those failures are part of the work.

Keith teaches AI through corporate workshops, online content, and his founder community, but his real subject is not the latest tool. It is how solopreneurs can turn their judgment into systems that learn, improve, and keep working when they step away.

The 38X AI Company is the operating system he wishes he had built earlier: one designed to compound results without endlessly compounding headcount, complexity, and pressure.

Keith lives in Hong Kong with his wife and children. You can follow his work at rumjahn.com.

