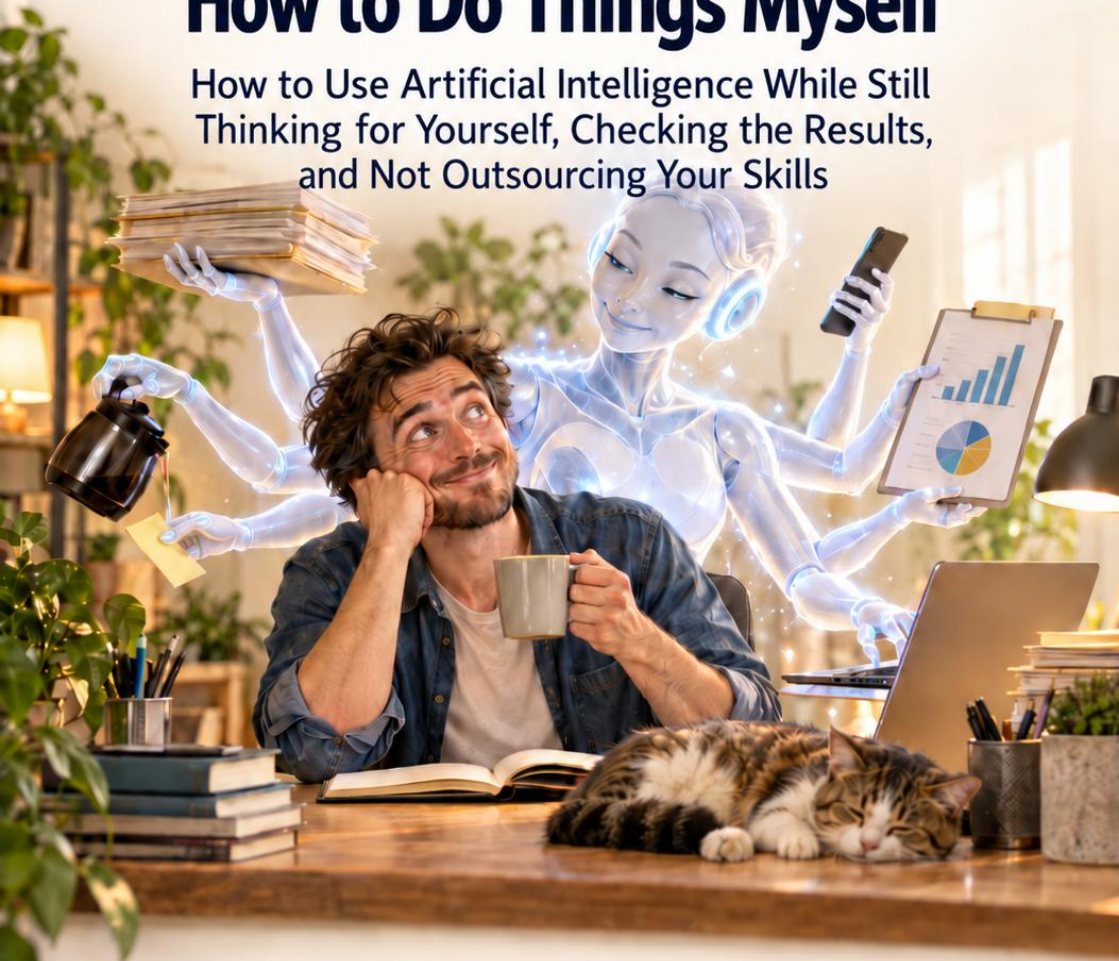
Max Paradox

PROBLEM? RELAX. YOU'VE GOT THIS!

A tongue-in-cheek guide series ;)

AI Is Doing More and More for Me, and I've Kind of Forgotten How to Do Things Myself

How to Use Artificial Intelligence While Still Thinking for Yourself, Checking the Results, and Not Outsourcing Your Skills



We've got a problem. But relax.
It can be solved—and we'll even have something to laugh about.

AI Is Doing More and More for Me, and I've Kind of Forgotten How to Do Things Myself

How to Use Artificial Intelligence While Still Thinking for Yourself, Checking the Results, and Not Outsourcing Your Skills

Max Paradox

PROBLEM? RELAX. YOU'VE GOT THIS!

ISBN: 978-83-8483-786-3

ALL-IN Rafał Płonkowski



CONTENTS

INTRO 5

Chapter 1 - When Your AI Assistant Quietly Became the Decision-Maker 11

Chapter 2 - It Looks Like I Know What I'm Doing. Do I Actually? 20

Chapter 3 - It Sounds Like an Expert. That Doesn't Make It Right 29

Chapter 4 - AI Saved Me Ten Minutes and Stole My Entire Afternoon 38

Chapter 5 - The Lost Art of Not Knowing the Answer Immediately 47

Chapter 6 - Ask the People Involved, Not Just the Machine 56

Chapter 7 - When Your Second Brain Goes Offline 66

Chapter 8 - Everyone Sounds Professional. Unfortunately, Everyone Sounds Identical 75

Chapter 9 - Give AI the Job, but Keep Your Hands on the Steering Wheel 85

Chapter 10 - Before You Trust the Result, Try to Break It 94

Chapter 11 - I Don't Need AI to Applaud My Ideas. I Need It to Find the Holes 104

Chapter 12 - Turn AI Into a Skills Coach, Not Your Permanent Substitute 114

Chapter 13 - The Best Decision According to AI. Best for Whom, Exactly? 124

Chapter 14 - When AI Can Click Send, It Needs More Than Your Optimism 134

Chapter 15 - AI Made the Mistake. Guess Who Gets to Fix It? 144

Chapter 16 - AI Built the Solution. Unfortunately, I Can't Tell Whether It's Safe 154

Chapter 17 - I Don't Have Time to Think for Myself. Now What? 165

Chapter 18 - Make Good AI Habits Boring Enough to Last 174

Chapter 19 - How to Tell When You're Giving AI Too Much Control Again 184

Chapter 20 - AI Can Do More and More. You Still Get to Decide Why 193

INTRO

It starts with something embarrassingly small. You need to write a short message declining an invitation. Nothing dramatic. You're busy, you can't make it, and you'd like to sound reasonably polite without suggesting that the other person's event is an inconvenience comparable to having your wisdom teeth removed. You open a blank message, type three words, delete them, and then instinctively reach for AI. Within seconds, you have a beautifully phrased response expressing sincere appreciation, heartfelt regret, and enthusiastic hopes for future opportunities. You send it. Five minutes later, you realize you've just described a neighbor's barbecue as an "invaluable opportunity for meaningful connection." Apparently, your social life now has a corporate communications department.

The strange part isn't that you used artificial intelligence to write a message. It's that writing one yourself suddenly felt unnecessarily difficult. A task you once handled without thinking now seems to require technological assistance. And perhaps you've noticed the same thing elsewhere. You ask AI to organize your thoughts before you've actually had any. You request a recommendation before identifying what matters to you. You let it summarize information, explain problems, prepare arguments, and occasionally tell you what your own opinion should probably be. Everything becomes faster and more polished. Meanwhile, somewhere in the background, your brain starts behaving like an employee who has discovered that someone else will cover every shift if it looks sufficiently busy.

Let's establish something important immediately: using AI doesn't make you lazy, unintelligent, or incapable. These tools can be extraordinarily useful. They can make complicated information more accessible, help people communicate, reduce tedious work, and support learning that might otherwise feel overwhelming. Refusing to use them simply to prove your independence would be a peculiar form of self-improvement. Nobody awards a medal for manually formatting a document that a suitable tool could handle in seconds. The question isn't whether artificial intelligence should help you. The question is whether you've noticed exactly what kind of help you're accepting, and what happens to your own abilities when that help becomes your default response to everything.

There is an important difference between saving effort and avoiding thought. Suppose AI drafts a routine announcement based on decisions you've already made. You check the details, adjust the wording, and send it. Wonderful. Now imagine asking AI to decide what should happen, choose the supporting arguments, determine which concerns matter, and write the announcement explaining your decision. You may receive an excellent result. But unless you've remained involved in the reasoning, you could end up being responsible for a position you barely understand. The words look like yours. The document carries your name. Unfortunately, when someone challenges its assumptions, the chatbot isn't necessarily going to attend the meeting and defend your judgment.

This is how dependence can develop without a dramatic turning point. Nobody wakes up and announces, "Today I shall permanently outsource my intellectual independence." Instead, you make hundreds of convenient decisions, most

of them perfectly reasonable. Why struggle with an opening paragraph when AI can provide five? Why work through an unfamiliar concept when it can explain it immediately? Why compare alternatives when it can produce a ranking? Gradually, the uncomfortable stages of thinking begin to feel optional. Uncertainty becomes something to eliminate instantly. An imperfect first attempt seems like wasted time. And a convincing answer starts to feel suspiciously similar to genuine understanding.

Of course, not every task deserves your full intellectual attention. You have deadlines, responsibilities, bills, relationships, and occasionally the ambitious personal goal of sitting down for ten uninterrupted minutes. You cannot turn every email into a critical-thinking exercise or every shopping decision into a dissertation. Sometimes you're exhausted. Sometimes the task is dull. Sometimes the sensible choice really is to let technology handle almost everything. A useful approach must accommodate those realities rather than demand that you become a permanently alert philosopher with an exceptionally organized calendar. The challenge is choosing where your involvement matters, not maximizing the number of tasks you perform without assistance.

Another difficulty is that AI can make incomplete knowledge look remarkably complete. A well-structured answer may contain an incorrect fact, an unsupported assumption, or a conclusion that doesn't follow from the information provided. The language can remain confident throughout. It's rather like listening to someone deliver a flawless speech about a building they've never entered. You might admire the vocabulary while failing to notice that they've placed the emergency exit inside a cupboard. The

greater the consequences of an error, the more important it becomes to distinguish persuasive presentation from reliable evidence. You don't need to distrust everything. You need to understand when a result requires verification, testing, or qualified human expertise.

But protecting your judgment is only part of the story. There's also the question of what you want to continue learning. If AI always completes the challenging portion of a task, you may obtain better immediate results while getting fewer opportunities to practice the underlying skill. That doesn't mean ability disappears overnight, or that doing things the hard way is automatically educational. It means the relationship between assistance and learning deserves attention. A tool that supplies finished answers can also explain mistakes, offer hints, challenge your assumptions, and help you practice. The same technology can become either a permanent substitute for certain activities or a valuable partner in developing your capabilities. Much depends on what you ask it to do.

And then there's the subtle question of ownership. Not legal ownership of a document, but ownership of the thinking behind it. Do you recognize your priorities in a recommendation? Can you explain why a suggested approach fits your circumstances? Are you comfortable rejecting AI's polished proposal because you know something important that the system doesn't? These questions matter because technology can be impressive without understanding your entire life. It doesn't automatically know which commitments you're willing to make, which risks you can afford, or what gives a particular activity meaning. Those are not annoying details interfering

with optimization. They are often the very reason a decision matters.

This book is about building a working relationship with AI that preserves those distinctions. You'll explore how to recognize when assistance has become excessive, how to evaluate generated work without checking every harmless sentence as though national security depended on it, and how to use artificial intelligence to strengthen skills rather than merely disguise gaps in them. You'll also encounter the less glamorous realities: mistakes that need repairing, tools operating outside your expertise, automation that requires sensible boundaries, and days when thoughtful independence competes with the urgent need to finish work and go home.

There will be no demand that you abandon AI, no ceremonial return to typewriters, and no suggestion that every difficult task must become an opportunity for personal transformation. You can use powerful tools and still think for yourself. You can delegate work without surrendering every decision. You can accept an excellent suggestion while remaining willing to question it. What matters is knowing which responsibilities belong to the tool, which belong to you, and when another person's expertise is genuinely necessary. That arrangement may look different depending on the task, and it can change as your abilities and circumstances change.

So consider this an invitation to remain an active participant in your own work and decisions. Let AI handle the repetitive parts, illuminate unfamiliar subjects, and offer possibilities you might never have considered. Enjoy the convenience. Just don't confuse a finished answer with a finished thought, or a professional-looking result with a skill

you've actually acquired. Your brain doesn't have to perform every job in the organization. It does, however, deserve something more meaningful than the honorary title of Vice President of Clicking Approve.

Chapter 1 - When Your AI Assistant Quietly Became the Decision-Maker

Nobody remembers the exact moment an assistant gets promoted to management without anyone approving the promotion. It usually begins with perfectly reasonable requests. Organize these notes. Summarize that document. Suggest a more efficient approach. Then, one afternoon, you find yourself looking at a decision you supposedly made, wondering why you can't explain it without reopening a conversation with artificial intelligence. Somewhere along the way, the helpful assistant stopped merely preparing your work and started determining what the work should be. The remarkable thing is that this transition doesn't require a contract, a meeting, or even a suspiciously cheerful email from Human Resources. It happens because each individual step feels convenient, sensible, and hardly worth questioning.

Consider Daniel, who manages customer support for a small company selling outdoor equipment. His team receives questions about deliveries, product availability, returns, and occasional complaints from people who appear to believe a waterproof jacket should also protect them from every disappointment life has ever produced. Daniel introduces AI to help organize incoming requests. At first, the arrangement works beautifully. The system groups messages by subject, identifies repeated questions, and prepares summaries of issues requiring attention. Daniel still makes the important decisions. Then he asks AI to suggest which complaints deserve priority. Later, he lets it recommend appropriate responses. Eventually, he begins approving entire batches

of proposed resolutions after a quick glance. The workload becomes more manageable, and everyone appreciates the improved response times.

One morning, a customer challenges a refusal to replace a damaged product. Daniel checks the record and discovers that the recommendation was based on an interpretation of the company's policy that he would not normally have accepted. The AI had treated an unusual case as though it were routine. Daniel could have caught the problem, but the decision had arrived packaged as a complete, plausible recommendation, accompanied by a polite response ready to send. He approved it because it resembled the dozens of acceptable recommendations he'd reviewed earlier. Nothing about the process announced that an important judgment was being transferred. His authority remained intact on paper. In practice, however, the system had increasingly determined which cases mattered, how they should be interpreted, and which outcomes seemed reasonable.

The distinction matters because doing work and directing work are not interchangeable activities. When AI sorts messages into categories you've defined, it performs a task. When it identifies potentially urgent cases for your review, it helps organize your attention. But when its classification becomes the sole reason a complaint receives less consideration, it begins influencing priorities. When its suggested interpretation becomes the basis of a customer decision, it participates in judgment. None of these arrangements is automatically wrong. The problem arises when you don't notice which decisions are being made, what information supports them, or whether the system has enough context to make a useful recommendation.

The gradual disappearance of your own judgment

Why is this transition so difficult to recognize? Because convenience doesn't arrive carrying a warning label. Each successful interaction encourages you to trust the next one. The first twenty summaries are accurate, so you spend less time checking the twenty-first. Most recommendations appear sensible, so you gradually stop examining their assumptions. This is an understandable response to a tool that often works well. Humans routinely adjust their attention according to experience. We don't inspect every button on a familiar appliance before pressing it. Unfortunately, a reliable history with straightforward tasks doesn't establish reliability under every possible condition, particularly when the next task involves unfamiliar circumstances or consequences for other people.

There's also a powerful difference between noticing an error in something you created and identifying one in a polished result you didn't develop. When you've assembled an argument yourself, you generally remember the difficulties encountered along the way. You know which facts were uncertain and where you made assumptions. With an AI-generated recommendation, that history may be less visible. You receive the conclusion without necessarily experiencing the decisions that produced it. Reviewing the final paragraph can therefore feel sufficient even when the most important weaknesses lie in the reasoning behind it. The result appears finished, and your brain, never one to reject a convenient shortcut without a compelling reason, is quite happy to move along.

This doesn't mean an AI system always reasons in the same way a human decision-maker does, or that a convincing explanation necessarily reveals the actual process behind its output. Depending on the system, it may

organize provided information, generate plausible interpretations, or use additional tools. The practical question remains the same: can you establish a defensible basis for the decision independently of how confidently the answer is presented? If your only justification is that the system recommended it, you haven't obtained the kind of explanation that matters when consequences arise.

Consider another example. A community theater committee uses AI to help select productions for the coming season. The committee supplies information about budgets, audience preferences, rehearsal availability, and volunteer experience. Initially, the tool simply compares proposals. Gradually, members begin asking it to rank the shows, identify the most commercially promising options, and recommend the final program. The recommendations appear professional, but an important criterion is missing: the theater exists partly to provide opportunities for inexperienced performers. A production that generates the highest anticipated ticket revenue may leave many volunteers with little meaningful participation. The AI hasn't necessarily failed to perform the requested comparison. The committee has allowed its original purpose to disappear from the decision.

That example reveals a central difficulty with delegation: the person who frames the problem often determines the answer's boundaries. If you ask for maximum efficiency, a system may reasonably favor speed and resource savings. If you ask for the strongest commercial option, financial considerations may dominate. But the best available solution to the stated question might not be the best answer to the situation you're actually facing. When you're actively involved in defining the problem, you're more likely to notice

missing priorities. When you accept the question suggested by the tool and then accept its answer, you can lose control twice without realizing it.

Three things that can be handed over

It helps to distinguish the actual task from the thinking surrounding it. There is the production of an output: a draft, calculation, schedule, or summary. There is the interpretation of information: deciding which details matter, what they indicate, and how competing explanations should be evaluated. Finally, there is authority: determining what action will be taken and accepting responsibility for its consequences. AI can assist with all three, but that doesn't mean you should give it the same freedom in each area.

Imagine asking AI to prepare a weekly inventory report. Having it organize figures into a readable format is straightforward delegation. Having it flag unusual changes can also be useful, provided the criteria are suitable and the underlying data are dependable. Allowing it to interpret the changes requires more care because an apparent shortage might result from a recording delay rather than an actual supply problem. Authorizing an expensive emergency order introduces a further level of responsibility. The technology's ability to perform these activities doesn't make their consequences equivalent. Somewhere between creating a report and spending company money, someone should know exactly where the authority lies.

In personal life, the boundaries are less formal but equally important. AI can organize a household moving checklist without knowing everything about your priorities. It can suggest different neighborhoods based on criteria you provide. It cannot determine how much inconvenience your family is willing to accept, how people feel about leaving

familiar surroundings, or which relationships matter most. A recommendation may help clarify possibilities, but treating it as the final answer would grant the tool authority over preferences that belong to the people involved. Sometimes the appropriate decision-maker isn't even you alone.

You can begin identifying these boundaries by examining a recent task where AI played an important role. Don't focus on how many prompts you used or how impressive the result looked. Instead, reconstruct the decisions. What did you already know before opening the tool? Which criteria did you choose? Which conclusions came from the system? Where did you genuinely assess alternatives, rather than merely recognize that the proposed answer sounded reasonable? If someone challenged the outcome tomorrow, what evidence or reasoning could you provide without asking AI to manufacture a defense?

The exercise can be surprisingly revealing. You may discover that a task you considered heavily automated involved substantial personal judgment. You defined the problem, selected meaningful criteria, rejected unsuitable proposals, and used AI mainly to accelerate preparation. Excellent. You may also find that your apparent contribution consisted largely of entering a broad instruction and approving the resulting document. That doesn't automatically make the outcome wrong, but it should prompt a closer look at whether the task required expertise or authority that you didn't exercise.

A particularly useful question is whether the system has begun making choices you never explicitly intended to delegate. Daniel, for instance, might be comfortable allowing AI to assign routine messages to categories. He might also approve suggested responses to common delivery

questions. But complaints involving unusual circumstances, disputed policy interpretations, or commitments outside established procedures need a different arrangement. By specifying these boundaries, he can retain much of the efficiency without asking himself to personally rewrite hundreds of ordinary messages. The objective isn't to reclaim every action. It's to reclaim the decisions where his understanding and responsibility make a difference.

Don't mistake supervision for understanding

One of the easiest ways to overestimate your involvement is to assume that approving an output means you've meaningfully supervised its creation. Approval can be a valuable safeguard, but only when the reviewer has enough information, time, and competence to evaluate the relevant issues. If someone places a technical document in front of you and asks for a signature, reading it carefully doesn't magically give you the expertise needed to identify every flaw. The same applies to AI-generated work. A final review may catch obvious mistakes while missing incorrect assumptions that require independent verification or specialist knowledge.

This is especially important in workplaces where AI adoption comes with ambitious expectations. Managers may be encouraged to automate more tasks while remaining responsible for every result. Yet reviewing a large quantity of unfamiliar material can become difficult in its own right. A genuine control process needs appropriate criteria, access to reliable information, and a way to handle exceptions. Otherwise, approval risks becoming a ceremony in which the human contributes little beyond clicking a button and accepting the consequences. It's an impressive arrangement

for the software, which receives most of the freedom and none of the uncomfortable performance reviews.

There will also be occasions when your own decision is less reliable than a suggestion from AI. Preserving judgment doesn't mean insisting that you're always right. The system may identify relevant information you overlooked or present a convincing alternative. The valuable response is to examine the disagreement, not automatically defend your original position. Sometimes you'll change your mind. Sometimes you'll discover that the tool lacks important context. Either result can improve the decision, provided you understand why you've accepted or rejected the recommendation.

At its best, this kind of collaboration makes your role more demanding in a productive way. You spend less time formatting routine materials and more time defining problems, considering consequences, and recognizing exceptions. Those activities require judgment rather than constant manual effort. At its worst, automation leaves you processing a conveyor belt of answers too quickly to understand them. Both arrangements can look impressively productive from the outside. The difference becomes apparent when circumstances change and someone needs an explanation that goes beyond the contents of the latest generated report.

Daniel eventually changes his customer-support process. AI continues to organize requests and draft routine responses, but unusual cases are flagged for meaningful review. The team clarifies which policy interpretations require human judgment and records reasons for consequential exceptions. Some recommendations still need correction, and the arrangement doesn't eliminate every

mistake. What changes is Daniel's ability to explain decisions and recognize when the system has reached the limits of the task he intended to delegate.

That's the distinction worth protecting. You don't need to write every sentence, examine every routine action, or reject suggestions simply because they came from a machine. You do need to notice when a convenient assistant has started defining your priorities, interpreting your responsibilities, or making decisions whose consequences belong to you.

An assistant can be enormously helpful. An assistant that has quietly become your manager deserves a slightly closer look, particularly if you're still the one receiving the complaints.

Chapter 2 - It Looks Like I Know What I'm Doing. Do I Actually?

A perfectly functioning result can be an excellent disguise for a missing skill.

Take Oliver, who has decided to turn his neglected backyard into a productive vegetable garden. His previous gardening experience consists of keeping a supermarket basil plant alive for eleven days, which he considers a respectable achievement under difficult circumstances. This time he's determined to do things properly. He photographs the space, describes the soil, lists the vegetables he'd like to grow, and asks AI for a planting plan. The resulting layout is impressive. It includes suggested spacing, watering arrangements, planting times, and explanations of which crops might suit different parts of the garden. Oliver follows the instructions carefully. Within a few weeks, he has something that looks suspiciously like competence. His neighbors compliment the transformation. One even asks for gardening advice.

Then several plants begin struggling, and Oliver realizes that he cannot confidently explain why. Is the soil too wet? Has the weather changed their needs? Are the plants receiving insufficient light? He has followed a detailed plan, but he hasn't developed much understanding of the principles behind it. The original instructions described what to do under expected conditions. Now the conditions are different, and the person officially responsible for the garden has discovered that his expertise lives inside a conversation history. The vegetables, unfortunately, have declined to wait while he retrieves it.

Oliver's situation illustrates a distinction that's easy to miss when AI helps us achieve impressive results. Completing an activity successfully doesn't necessarily mean you've acquired the underlying skill. You can follow a well-designed set of instructions without understanding why those instructions work. You can use a reliable template without knowing how to adapt it. You can produce an excellent outcome with expert assistance while remaining a beginner yourself. There's nothing inherently wrong with any of that. Trouble begins when the quality of the finished product convinces you that your personal capabilities have improved to the same degree.

The difference between recognizing an answer and knowing it

One reason this illusion is so convincing is that understanding can feel much stronger while we're looking at an explanation than when we have to produce one ourselves. Read a clear description of a concept, and its logic may appear perfectly obvious. Close the explanation and try to reconstruct it, and suddenly there are missing steps everywhere. The information seemed familiar because it was presented in an organized sequence. That familiarity isn't the same thing as being able to retrieve the knowledge independently, apply it to unfamiliar circumstances, or recognize when the explanation doesn't fit.

AI can make this experience especially pleasant. You ask a question and receive an accessible, patiently structured response. Difficult concepts arrive with examples, comparisons, and helpful transitions. You follow along, nodding occasionally at your screen like a thoughtful professor reviewing the work of a promising student. Then someone asks you to explain the same concept without