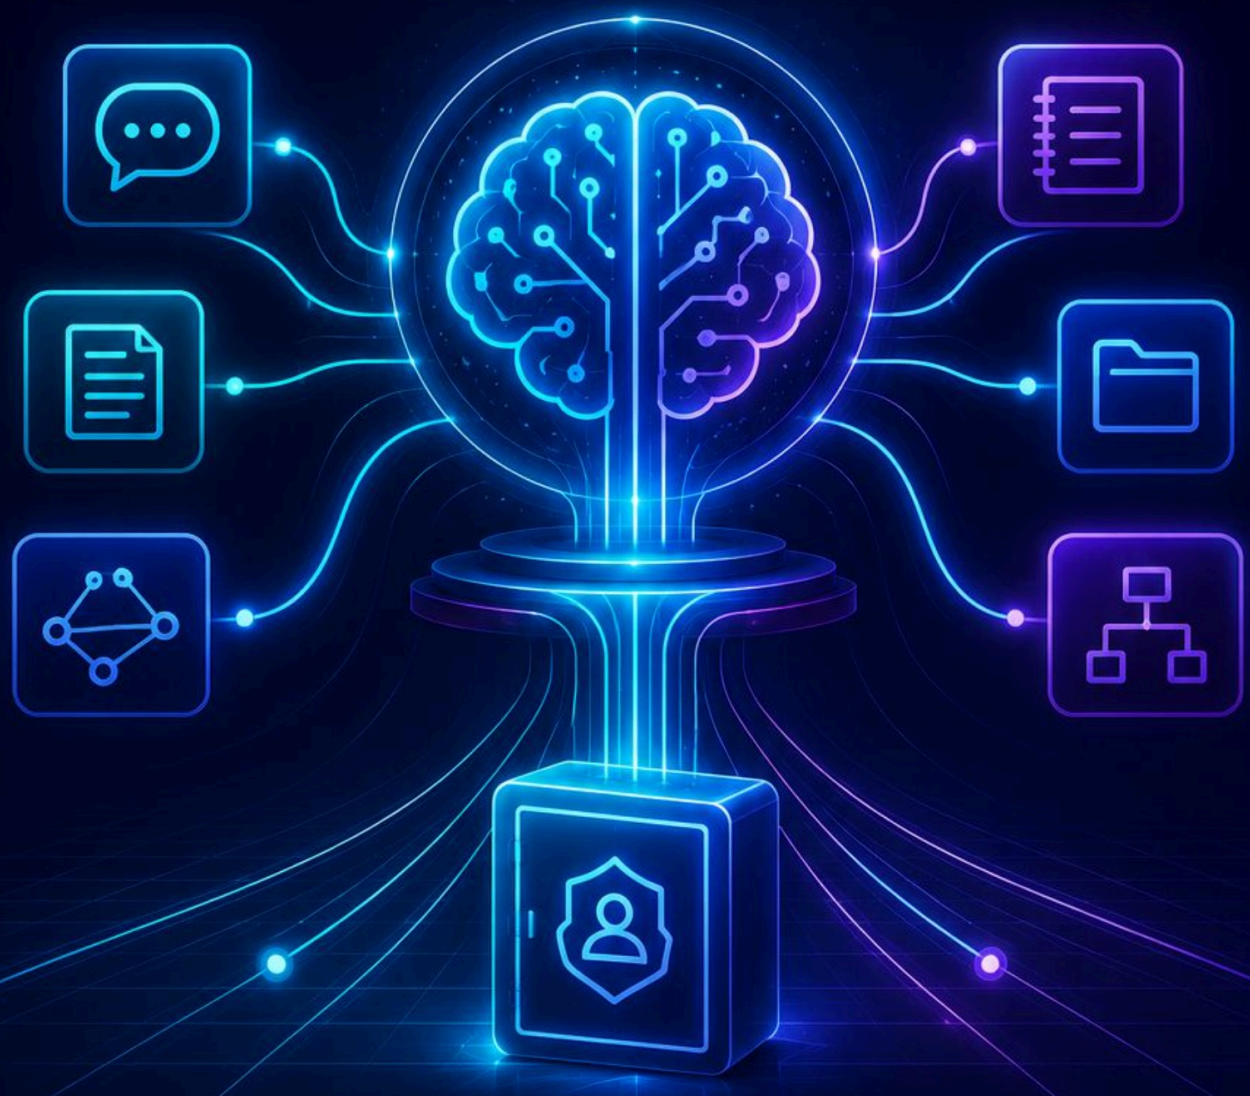
Teach It Once

Make AI your mission control for life and work



Michael Zelbel

OWNWARD STUDIO

Teach It Once

Godspeed Mission Control, a personal AI for life and work

Michael Zelbel

Teach It Once: Godspeed Mission Control, a personal AI for life and work. Second edition, October 2026. Copyright 2026 Michael Zelbel. All rights reserved. No part of this book may be reproduced or distributed in any form without the written permission of the author, except for short quotations in a review. The companion kit is the exception on purpose: every template, prompt and checklist in it is free to download, copy, change and use.

This book is written in a private capacity, in my own time. The views in it are mine alone and not those of any employer or client, and no employer or client is named in these pages. Godspeed Mission Control, the free program this book sets up, is my own work, and so is one optional tool it mentions: Querino, which hosts the Copy prompt pages. Product names belong to their owners. The full publication information, including what I borrowed and from whom, and what this book does and does not promise, is at the back.

Companion kit: **Download the matching kit**

Contents

Part I: Start With Something That Matters

Chapter 1: Make Your AI Briefing From What It Already Knows 10

Chapter 2: Set Up Godspeed Mission Control 14

Chapter 3: Get Useful Work Done on Your First Day 22

Chapter 4: Your Files Are the Originals 26

Chapter 5: Let It Check In With You 29

Part II: Give It Useful Background

Chapter 6: The People Who Matter 34

Chapter 7: Priorities Across Your Life 37

Chapter 8: Give Mission Control a Goal, Then Let It Work 40

Chapter 9: Teach AI Your Writing Style 45

Chapter 10: Tell Mission Control Once, Let It Remember 48

Chapter 11: Let Mission Control Keep Up With Your Life 51

Chapter 12: Correct AI Without the Argument 53

Chapter 13: Your Notebook: Notes, People and What It Noticed 56

Chapter 14: Track Anything With One Word 58

Part III: Teach AI Repeated Jobs

Chapter 15: Save a Repeated Job as an AI Skill 67

Chapter 16: Make an AI Skill Work Your Way 71

Chapter 17: Skills for Email, Planning, Conversations and Decisions 74

Chapter 18: Create Skills for Your Own Work 78

Chapter 19: Check That Your AI Skills Work 81

Part IV: Set Limits and Check the Work

Chapter 20: Set Limits on What AI May Do 84

Chapter 21: Back Up Your Mission Control and Undo Mistakes 89

Chapter 22: Check Where Your Private Information Goes 93

Chapter 23: Check AI's Answers Before You Act 96

Part V: Have Work Ready When You Return

Chapter 24: See and Change What Runs on Its Own 101
Chapter 25: The Morning Brief 104
Chapter 26: Let AI Investigate a Question While You Are Away 108
Chapter 27: Set Up an Automatic Weekly Review 110
Chapter 28: Get Alerts When a Page or Product Changes 112
Chapter 29: Keep Researching a Question Over Time 115
Chapter 30: Track Deadlines Until the Work Is Done 118

Part VI: Add What You Need *(optional add-ons)*

Chapter 31: Give Your Mission Control Its Own Email Address 124
Chapter 32: Let Your Mission Control Read Your Mail 129
Chapter 33: Let Your Mission Control File Your Paperwork 133
Chapter 34: Use Your Mission Control on Several Devices 141
Chapter 35: Look After Your Server 144
Chapter 36: Let Your Mission Control Make a Phone Call for You 147
Chapter 37: Let Your Mission Control Finish Your Videos 152
Chapter 38: Try Your Mission Control With Another AI Assistant 157
Chapter 39: Reach and Share Your Saved Prompts Online 160
Chapter 40: Coaching in Depth: Areas, Styles and Habits 162
Chapter 41: Use Your Mission Control for Coding Projects 167
Chapter 42: Let Your Mission Control Be Your Wingman 171

Closing

Chapter 43: Keep the Jobs That Move Your Life Forward 178

Back Matter

Appendix A: The Companion Kit 181
Appendix B: Where to Go From Here 184
Appendix C: Your Build Checklist 185
Appendix D: The Mission Control Folder Layout 187
Appendix E: How an Assistant Produces an Answer 190
Appendix F: Notes and Sources 192
Publication Information 194

How to Read This Book

One Sunday evening in September, my phone buzzed with a message from my personal AI. It had found something for me. Peter Attia, a doctor with a popular health podcast, had just released an episode with a headache specialist. Some mornings I wake up with a headache from the way I slept, and my AI knew that. So it didn't just send me the link. It told me to skip to minute 51, where the specialist explains which everyday habits make headaches less frequent. That was more attention than I had ever paid my headaches myself. Then it asked me a question. I had started running again the week before, and it had that in its notes: how had the runs gone, and were my headaches any different on the days I ran?

I hadn't asked for any of that. Weeks earlier, I had told it once which health experts I trust, and that I wanted to talk about my health on Sunday evenings. It took it from there: every week it reads what those experts publish, picks out what fits my life, and brings it to our Sunday talk.

Most of us use AI like a very clever stranger. We open a chat, explain everything, get a good answer, close the chat, and next week we explain everything again. This book is about a different kind of AI. It keeps what you tell it. It works on what matters to you while you get on with your day. And it messages you when it needs you.

That AI is a free program called Godspeed Mission Control. It does its thinking with the ChatGPT account you already have, or with another AI service you prefer, such as Claude, Gemini, Kimi or OpenRouter. ChatGPT is my example throughout this book, because most people already have it. Godspeed adds two things a chatbot can't do on its own: notes about your life, kept as simple files you can read and correct, and work done while you are away. You install it in one evening, with this book walking you through every step, and it starts working. I named it after the room full of people on the ground who look after astronauts. It does the watching; you stay in charge. In this book I call your copy of it "your mission control".

Take Godspeed as a starting point that you make your own. What it knows about you, the jobs it does and the rules it follows are all plain text files, and you change any of them by telling it what you want. Most of this book is about exactly that. I shaped mine around my life, so after a few months yours should look quite different from mine. That is the idea.

Every chapter ends with something that works, so you can stop after any of them and keep what you have.

What you will do

First, you find out what ChatGPT already knows about you. Nothing to install yet. You ask it, correct its answer and save the result as a one-page file. That file is how your mission control gets to know you, so you never have to explain yourself from scratch again.

Then you install Godspeed Mission Control. It needs a computer to live on, and it can have two homes. One is your own laptop or desktop, where you sit down with it to read, write and plan. The other is a small rented server: a computer in a data centre that you rent by the month and that never sleeps, so your reminders arrive on time even while your laptop is shut in your bag.

You start with one home and can add the other later. I suggest the server, because work done while you are away is the whole point of this book. If renting one sounds like too much for a first evening, start on your own computer. Either way, on that first evening your mission control reads your file, asks you for one goal and does a first piece of real work on it.

Next you tell it about your life. Who matters to you, what comes first right now, how you write, and what it must never do without asking. You also see exactly where your information goes and who else can read it. You meet your notebook, a private web app that shows all its files in a friendly way, so you can read and correct everything it knows about you. And you teach it the jobs you keep repeating, so you explain each one once.

Then you let it work while you are away. This is where it pays off. It can have your morning ready before you sit down. It reminds you about the tax return, more often as the deadline gets close, and stops when you tell it the return is filed. With your go-ahead, it keeps after a company's support team until your problem is solved. Mine did that for seventy-five days, politely, until a shop that kept refusing my order finally let me buy what I needed.

Last come the extras. Each one adds one specific thing, such as reading your Gmail, making a phone call for you, or putting captions on your videos. Nothing else in the book depends on them, so take the ones you need and skip the rest.

One chapter along the way is mostly for fun, and you can skip it without missing anything: "Correct AI Without the Argument", a small experiment in how to correct an AI without losing your temper.

What it costs

- Godspeed Mission Control: free, now and later. No subscription, no upgrade.
- The AI that does the thinking: this book begins with ChatGPT, whose free plan is enough to begin with Chapter 1. For the verified Godspeed beta setup, use an API account and access key. API use has its own allowance and billing, separate from a ChatGPT subscription. You pay the AI service whatever it charges.
- The server, if you rent one: a monthly fee to the hosting company. When I checked in October 2026, the plan I use cost $8.99 a month to start and $14.99 a month once it renews.

How the pages work

Text in a box is usually something to paste into a chat: ChatGPT in the first chapter, then your mission control's own chat, which you set up in the second. When a box shows a file or a reply instead, I say so. The **Copy prompt** link under a box opens the same text online, so you can copy it in one tap. A name in this font, such as `profile/about-me.md`, is one of your mission control's files: this one is the file about you, in a folder called profile, and the `.md` at the end just marks a plain text file. The replies I quote come from real tests. In some tests I used people I made up for practice, such as Sam, an illustrator; I say so where they appear.

Give your mission control a week before you judge it. The first few check-ins, the short chats it starts with you, will tell you more than my enthusiasm can.

Part I: Start With Something That Matters

Chapter 1: Make Your AI Briefing From What It Already Knows

Who makes the first move?

You close the AI app. The conversation was useful. You have a better plan than you had an hour ago. Tomorrow, you will still have to remember to open it. The app is not going to call you.

I want a different arrangement. If a decision is due on Friday, I want the research ready on Thursday. If I promised my mother a call on Sunday, I want a nudge on Saturday, not an apology on Monday.

That is what your mission control is for. It does the legwork: it looks things up, compares choices and prepares drafts. You keep the decisions. It can compare three vacuum cleaners for you, and it won't buy one until you say so. A message it writes stays a draft until you say send.

It needs something worth working on, so pick one thing you would like to move forward in the next few months. It can come from any corner of your life: finishing a project at work, running your first 10k, getting your money in order, or finally having that calm talk with your brother about your parents. One is enough, and you can add more later. Pick yours now and put it in one sentence. You will need it in a few minutes.

It also needs to know who you are, without you giving a speech. I have had quite enough of AI help that starts with a questionnaire. By the time I have explained the work, I have done half of it. It is like hiring an assistant who starts every morning by asking where the office is.

So in this chapter you make a briefing: a one-page note about you, written by the AI you already use and corrected by you, with your goal in it. In the next chapter you hand that briefing to your mission control, and it starts from there. Today you only need ChatGPT, and nothing can break.

Ask ChatGPT what it knows about you

Open ChatGPT, or Claude if that is the chatbot you use, and paste the prompt below. You don't need to study it. In short, it asks the chatbot what it knows about you, under five headings. It also asks it to mark its guesses with (?), to keep them apart from what you actually told it, and to stick to what it can really see, because chatbots are eager to please.

```
Help me create a short briefing that I can give another AI assistant. Start with information
available to you in this conversation, accessible earlier conversations and saved memory. Do not
imply you can read material you cannot access.
```

```
Tell me what you know under exactly these five headings:
- Who I am and what I do
- The people who matter
- What I am focused on now
- How I like to be helped
- My hard limits

Mark uncertain claims with (?). Distinguish things I told you from your interpretations. Do not
include passwords, access keys, payment details or other credentials.

Then list at most six short questions that would fill important gaps. List them without asking
me to answer yet.
```

Copy prompt

The answer may be thinner than you expected. Do not take it personally. The chatbot only knows what your account lets it see, such as your old chats and the notes some chatbots save about you, their "saved memory". And this is a short briefing, not an export of your account (a full copy of every chat you ever had). A few accurate lines are all your mission control needs to start.

Correct what it got wrong

Read the five sections in the chat. Some lines will sound like you. Others may describe last year's you, still busy with a project you're glad you abandoned. Start with the lines marked (?), because those are its guesses.

Correct whatever is wrong in plain language, one sentence per mistake. Suppose you have moved from Company A to Company B, but the briefing still lists your old employer. Tell the chatbot:

```
Correct my employer. I now work at Company B, not Company A.
```

Then add your goal the same way, so it shows up under "What I am focused on now": "Add my goal: I want to run my first 10k this spring." Answer the questions it listed at the end, as many as you can. If you can't answer one yet, leave that gap in: an honest gap is better than a made-up answer. Last, ask it to show you the whole briefing again, with your changes in it.

When I first tried a prompt like this in my ChatGPT account, all five headings came back filled. It remembered my work and the books I write. It also described my preference for short answers and one step at a time when I needed instructions. I'd never sat down to write a job description for an assistant, yet here was one, assembled from my requests and corrections. It was an oddly useful record of my impatience.

Take out what should not travel

Before you pass the briefing on, read it once more for things that should stay with you. The prompt tells the chatbot to leave secrets out, but a chatbot does not always do everything it is

asked. Ask it to take out passwords, card numbers, access keys (the long secret codes some apps hand out instead of a password) and anything else you'd rather not share. Secrets like those belong in a password manager, such as the one built into your phone or browser.

That keeps them out of the briefing, but the chatbot may still remember them. If that bothers you, tell it to forget them, and delete the old chat where you first typed them, if you can find it in your list of past chats.

For other private things, give only what the help needs. To plan your week, “I keep Thursday evening free” is enough; the reason can stay with you.

Why so fussy? That first try also turned up something I had forgotten I ever gave an AI. Months earlier, I had given an AI my credit card details for an experiment: I wanted to see whether it could order from an online shop. The experiment failed miserably, and maybe that's a good thing, because it failed on the security checks that are normal for German credit cards. You confirm an online payment not only with the code on the back of the card, but also in a special banking app on your phone, which in turn wants your OK and your face.

So the experiment ended really fast, and I forgot about it. The card details made an unexpected return in the briefing I was about to hand to another AI. I had asked for a useful account of myself and received rather too useful an account of my payment method.

Save it as a file

When the briefing reads right to you, ask for it as a file. "Markdown" in the prompt just means a plain text file ending in `.md`, which every AI reads easily.

```
Give me this briefing as a downloadable Markdown file named what-my-ai-knew.md. Keep the
uncertainties visible. If you cannot create a download, give me the exact text to save in a
plain text file with that name.
```

Copy prompt

Save the download on your computer as `what-my-ai-knew.md`, somewhere you will find it again tomorrow, because the next chapter starts from there. Did all this on your phone? Open the same chat on your computer and download it there. If you got text instead of a file, ask the chatbot to walk you through saving it as a file, one step at a time. It is good at exactly that.

All the explaining you used to do in every new chat is now a file. Hand it over, and the next assistant starts where the last one left off.

What “once” means

This is where the book gets its title, *Teach It Once*. You did the work of explaining once, and now it lasts. When something changes, you fix one line instead of giving your life story again.

For an AI to work while you are away, it needs more than a file. It needs a computer that stays switched on, and your permission to work on your goal. Chapter 2 sets up the computer, either a rented one that never sleeps or your own, and it is mostly copying and pasting. Today, you have a briefing you've read and corrected, with your goal in it. Not bad for a chapter that needed nothing but a chat.

Part II: Give It Useful Background

Chapter 8: Give Mission Control a Goal, Then Let It Work

Knowing what matters is one thing. Doing a little about it every day is the part most of us lose to a busy week. Your mission control does that part.

Every day it does a small piece of work on up to three of your goals, and tells you when it needs you. I call this the daily round. It switched on by itself when you gave it your first goal, so there is nothing to start. Here is how it works, so you can trust it and correct it.

The whole round, in four steps:

1. **You name a goal, once.** One sentence is enough.
2. **It reads up on how people usually reach that goal, once,** and writes what it learns into a short guide, the playbook.
3. **Every day, it picks a step and does it.** Then the small checking program you met on your first day makes sure the result is really there.
4. **Now and then, it asks you.** Only for what nobody else can answer or do. You reply in a few words. On most days it needs nothing from you at all.

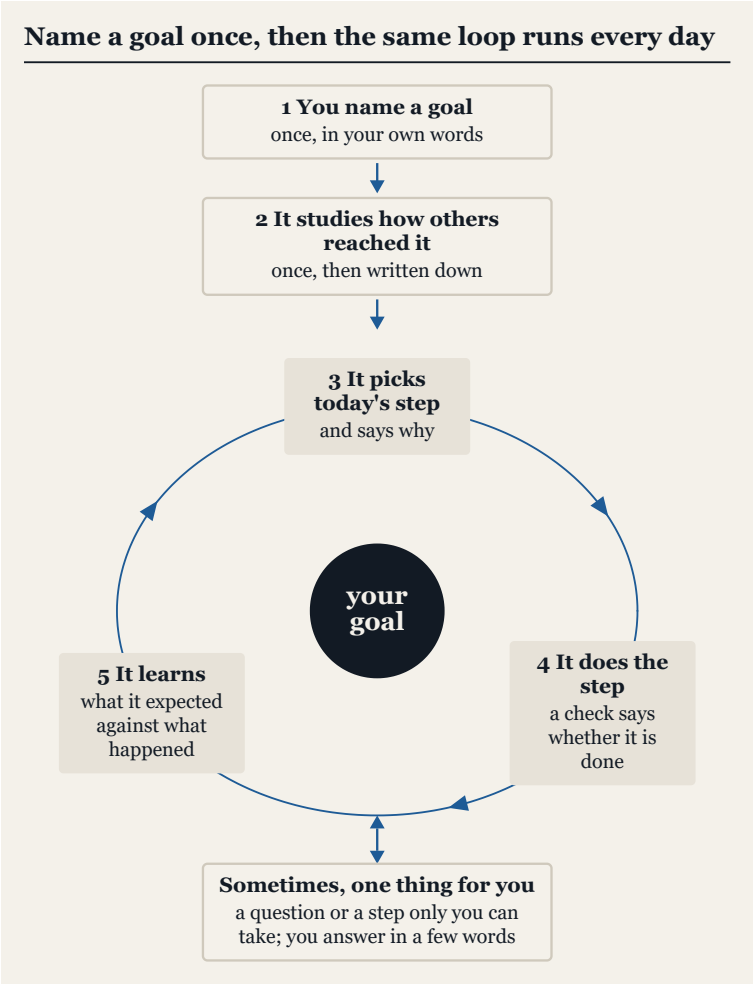


Figure 8.1: On most days the box at the bottom, the only one that needs you, stays empty.

A week with one goal

Here is a made-up week with the 10k from your first day: race on 19 April, three runs a week. The week starts the morning after.

Monday. It reads up on how beginners get there. Then it checks the facts with you and asks one thing: "Race on 19 April, three runs a week. Any old injuries I should plan around?" You say: "None."

Tuesday. It writes the playbook, and the check passes. "Your 10k playbook is ready in the notebook." You: "Good. Go on."

Wednesday. It works out what is holding you back: you have never run far. So it keeps this week's three runs short and picks days for them. "Here is a first-week plan, on days that suit your week." You: "Not Thursday, my sister visits."

Thursday to Sunday. It moves Thursday's run to Friday and starts a log of your runs. On Friday you run, and afterwards you tell it: "Ran 2.5 km. Legs fine." Its next round writes that into

the log, in your own words. On Sunday it checks the week against the plan: on track. On those four days you hear nothing from it, because nothing needed you.

Your whole share of the work that week, apart from the running, was fourteen words.

Name a goal

Your first goal went in on your first day. You add more the same way, in a plain sentence. The opening words matter: "Make this a goal and work on it with me" tells it you want work done, not just a note taken. Then say what you want, by when, and where you are now:

```
Make this a goal and work on it with me: I want to manage a simple conversation in Spanish by
our trip to Seville on 12 May. Right now I know about twenty words.
```

Twenty words of Spanish is already more than you needed for the whole 10k week.

It saves the goal under **Goals**: your words, the date you said them, and the date of the trip. From the next day on, the round works on it. You change a goal the same way you made it: "push it to June", "drop it".

The playbook

The playbook is named after the book of moves a sports team plays from. To write it, your mission control reads up on how other people reached the same goal: guides, research, the experience of people who have done it. What do they actually do, and in what order? Where do people usually give up?

It also writes down what is still unknown: where the guides disagree, and where your case is different from theirs. And it sorts the steps into two piles: the ones your mission control can do alone, and the ones only you can do. The runs go in your pile, sadly.

Why read up, when it made you a plan on your first day? That plan was one job. The playbook is the background it rereads every day before it picks the next step, the way a coach glances at the season plan before each training session. Ask it "show me the playbook for my 10k", or open it in your notebook, the private web app that shows your mission control's files. It is written for you, and it is often the most useful thing it writes.

How it picks today's goals

Once a day, it reads every goal, every playbook, what it finished yesterday and anything you said since. Then it picks at most three goals for today and writes down why.

Some goals always get a seat. If you told it "my weekly call with my mother always comes first", that call is in. So is anything due within a week, like your Spanish when Seville is six days away. And so is any goal it has not looked at for a week, like seeing more of your friends. That last one

protects the quiet goals, the ones with no date and no number, which would otherwise lose to the loud ones every day.

For each chosen goal, it asks the question from the last chapter: what is the one thing holding this goal back right now? Then it picks one step for today and writes down what "done" will look like, such as "a plan with the race date in it". The checking program opens the result and looks for exactly that.

What comes back to you

Every daily round ends with one line for you. On most days it is "nothing today", with the reason, and then it doesn't bother you at all. When there is something, the line reaches you: in Telegram if your mission control lives on a server, or, on your own computer, as a note in your notebook, in a folder called From your routines. If you later switch on a morning brief, a later chapter's idea, the line opens the brief instead.

Here is what my own mission control said on its first two days. On the first, it had one question. I offer three small software kits online, and all three were free, with a tip if a buyer felt like it. It asked whether I wanted to charge for them and suggested three prices, numbered, so my whole reply could be one digit. On the second day it had nothing for me. Both days, it was right about whether it needed me.

When there is something, it is one thing, never two, and it is one of three kinds:

1. **A question only you can answer,** like Monday's injury question, with what it already knows laid out, so your answer is a word or a number.
2. **A step only you can take,** with everything it could prepare already done: the text written, the options compared.
3. **A step that reaches somebody else,** such as a message, a payment or a sign-up, prepared and waiting for your yes.

You answer in ordinary words, in the same chat: "not Thursday", "push it to May", "drop that". It writes your answer under the goal, so you can see later what you said and what it did with it.

Two things protect you here. Silence is never read as a yes, so nothing happens just because you were busy. And a question or suggestion you ignored twice is not put in front of you a third time. It can take a hint.

Deadline reminders, the second routine that came with your goal, work differently. Say "remind me to renew my passport by June", and the reminder keeps coming back, calmly, until you say the thing is done. The daily round lets go of a suggestion; a reminder does not let go of a promise.

When it is wrong

It will sometimes be wrong. Unlike most of us, it then writes down that it was wrong, and why.

Take Wednesday's guess: what holds your 10k back is that you have never run far. If you run five kilometres next week without trouble, the guess was wrong. It writes that down and looks for the real obstacle. The playbook works the same way: when what happens to you contradicts it, it marks the playbook as wrong, writes down why and reads up again.

Now and then it also writes a guess down as a number, like a weather forecast: a 60 percent chance that you are still on plan three weeks from now. When the three weeks are up, it writes down whether it was right, next to the guess. After a few months, you can see how often its guesses come true, and how far to trust them.

Start a round yourself

The round runs on its own every day. You can also start one any time, in the chat. Say "decide today", and it picks today's goals and steps on the spot. Say "do the next piece of work", and it starts right away on the next step already on its list. Try it once just to watch it think, and again on any day you have just told it something important, so it acts on the news today instead of tomorrow.

End of the free sample

You have read Chapter 1 and Chapter 8 of *Teach It Once*. The whole book has 43 chapters.

To keep reading, the full book is on Leanpub:
leanpub.com/teachitonce

The free companion kit holds every template, prompt and checklist the book uses:
teachitonce.com/companion