Learn to use MADRE in 10 minutes.

You'll install the app, connect the AI you already use and ask it your first question. You don't need to know how to code or to open the Terminal.

MADRE 1.1.5 manual

Start here.

Seven steps, in order. When you finish each one, mark it: this page remembers where you are.

1 What MADRE is and what you need.

MADRE is an app for your computer. It brings the AIs you already use into one chat: Claude, ChatGPT (Codex), Gemini, Grok, OpenCode and Cursor, plus a free AI that runs on your computer. They all work on your projects and you decide what they may change.

30 s

Tick what you already have

- [ ] **A Mac with an M1 chip or later, or a Linux computer.**
  On a Mac you need macOS 12 Monterey or newer. Have Windows? It works too, with its installer: you'll see it in step 2.

- [ ] **An account with some AI.**
  Your ChatGPT or Claude account, or an API key from another company (step 4 explains what that is). MADRE works with your account.

- [ ] **If you have none, that's fine.**
  MADRE helps you set up a free AI on your computer, if it fits your machine.

Four words you'll see often

Project	A folder on your computer: your shop, your invoices, your website. Each project has its own chats and memory.
Agent	What we call each AI when it works on your project. It reads your files, answers you and, if you give it permission, changes them.
Permission	What you let the AI do in each message, from #0 to #4. If you don't pick one, it only reads and chats.
Memory	What MADRE remembers about your project, so you don't have to explain it every time.

You're good if your computer is ready and you know which AI account you'll sign in with.

2 Install the app.

Pick your system. On Mac and Linux it's an app you download and open. On Windows, an installer that sets everything up on its own.

1½ min

macOS

1. **Download it**
   Download MADRE for Mac (53 MB). The file lands in your Downloads folder.
2. **Drag it to Applications**
   Open the file and drag the MADRE icon to the Applications folder, like any app.
3. **Open it**

From Applications or Launchpad. It's signed and checked by Apple, so your Mac opens it with no odd warnings.

Linux

1 **Download the package**

For Ubuntu or Debian, download the `.deb` for your machine: Intel or AMD (60 MB) or ARM (60 MB). For another distribution on ARM, the AppImage (131 MB).

2 **Install it**

Double-click the `.deb` and your software center installs it. From the Terminal it would be:

```
sudo apt install ./MADRE_1.1.5_amd64.deb
```

The AppImage installs nothing: make it runnable with `chmod +x MADRE_1.1.5_aarch64.AppImage` and open it.

3 **Open it**

Look for MADRE among your apps.

Windows

1 **The app isn't on Windows yet**

Today MADRE runs on Windows with its installer, which sets everything up on its own inside WSL2, the Linux that comes with Windows. It takes a few minutes and you don't have to open the Terminal.

2 **Download the installer and open it**

Download MADRE-Instalador.exe and double-click it. If Windows says "Windows protected your PC", click **More info** and then **Run anyway**.

That notice shows up because the installer doesn't have a digital signature yet, not because anything is wrong with it. I explain it calmly at madre.run/en/windows.

3 **Open it**

When it's done, open «MADRE» from the Start menu or the desktop. It opens in your browser, with no Terminal windows.

The first time, MADRE shows you **Before you start**: what MADRE is in three lines, its license and its privacy notice. If you agree, click **Accept and start**.

> **The app lets you know when an AI finishes**, even if you're in another window. On a Mac, if you close the window, MADRE stays in the menu bar. When there's a new version, it tells you and you install it with one click.

Already used MADRE in the browser?

Your projects reach the app just as they are. If you used to open it from the Terminal, after updating MADRE asks whether you stay there or get the app.

For people who code: MADRE from the Terminal

Mac and Linux, with Node 22.5 or newer

```
npx @jossuealcala/madre
```

It opens in your browser, at `127.0.0.1`. That way it also works on a server, through an SSH tunnel.

Check your AIs without opening MADRE

```
npx @jossuealcala/madre doctor
```

You're good if MADRE is open on your computer.

3 The first time, MADRE guides you.

When you open it, MADRE greets you with a seven-step welcome. That's where you connect your AIs, pick your first project and, if your computer can, set up a free AI. All with buttons. If you don't have a project, MADRE makes a test one for you, and at the end it takes you to write your first message.

2 min

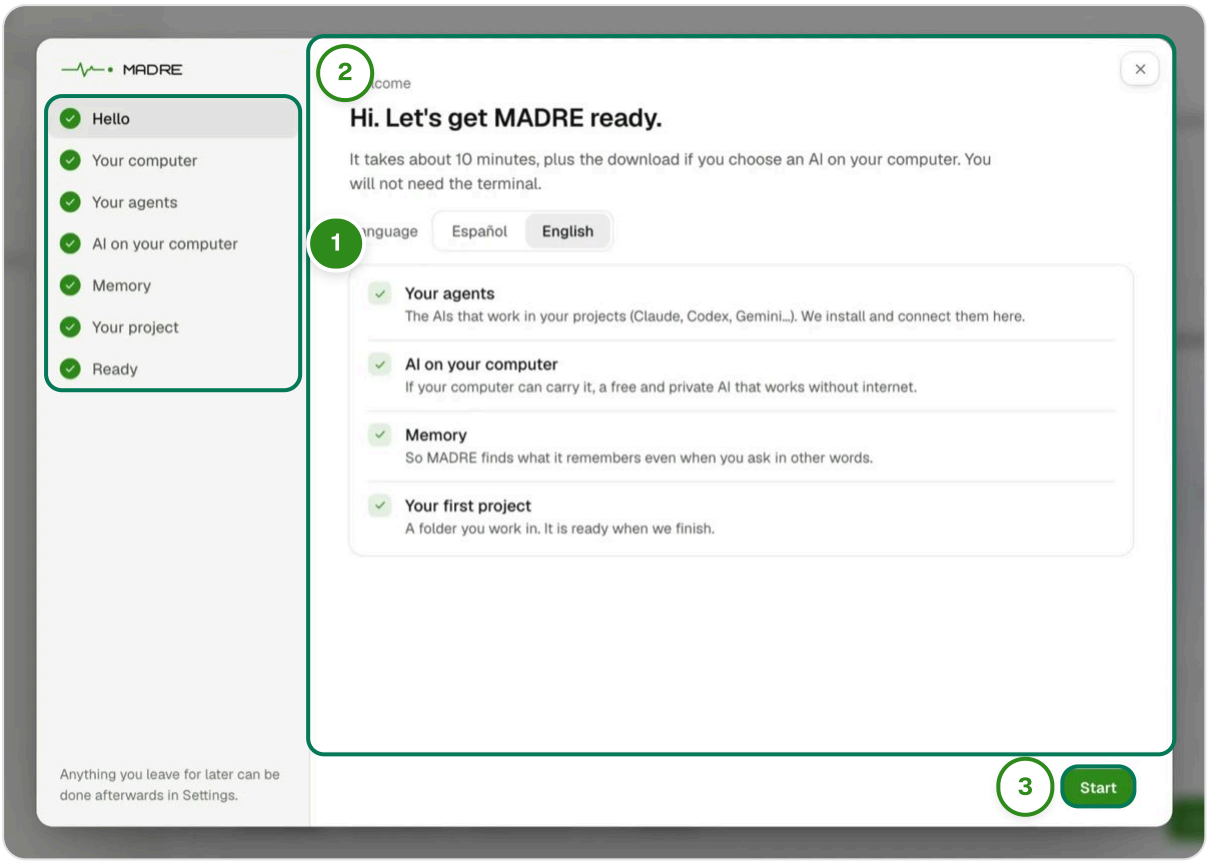

1. The seven steps are on the left. You can go to any of them; the ones you finish get a tick.

2. Each step tells you what it is and gives you the button to do it: connect an AI, download the free AI or pick your folder.

3. The button at the bottom takes you to the next step. Whatever you leave for later you can do in Settings.

If you're new

Why does it check my computer?

To know which AIs you already have installed and which free AI fits your machine. It checks on your computer, and the welcome tells you what for.

> Closed it? You can open it again from Home, in **Settings › Projects › See the welcome again**. It picks up from the step where you left off.

You're good if you finished the welcome or you know where to open it again.

4 Connect your AIs.

You connect each AI only once and it's ready for all your projects. One is enough to get started.

1½ min

Start with the account you already have

- **Claude** Anthropic
  Pick **Claude account** if you pay for a Claude subscription, or **Anthropic Console** if you have API credit. The page gives you a code and you paste it in MADRE.
- **ChatGPT (Codex)** OpenAI
  You sign in with your ChatGPT account: MADRE opens the page so you can sign in. An OpenAI API key works too.
- **Gemini** Google
  You need an API key from Google AI Studio. Google no longer lets Gemini CLI sign in with a free personal account.
- **Grok** xAI
  You sign in with your Grok account from its button.
- **OpenCode** Many providers
  It doesn't bring an AI of its own. You connect the AI company you already use by pasting its API key, or your GitHub Copilot account.
- **Cursor** Cursor
  If you already use Cursor, you sign in with your account from its button.

1. **Says «Not installed»?**
   Click **Install** and MADRE installs it for you.
2. **Sign in**
   Each AI has its own button to sign in. When it says «Connected», it's ready.
3. **Move on**

Connect the others whenever you like, in **Settings › Your AIs**.

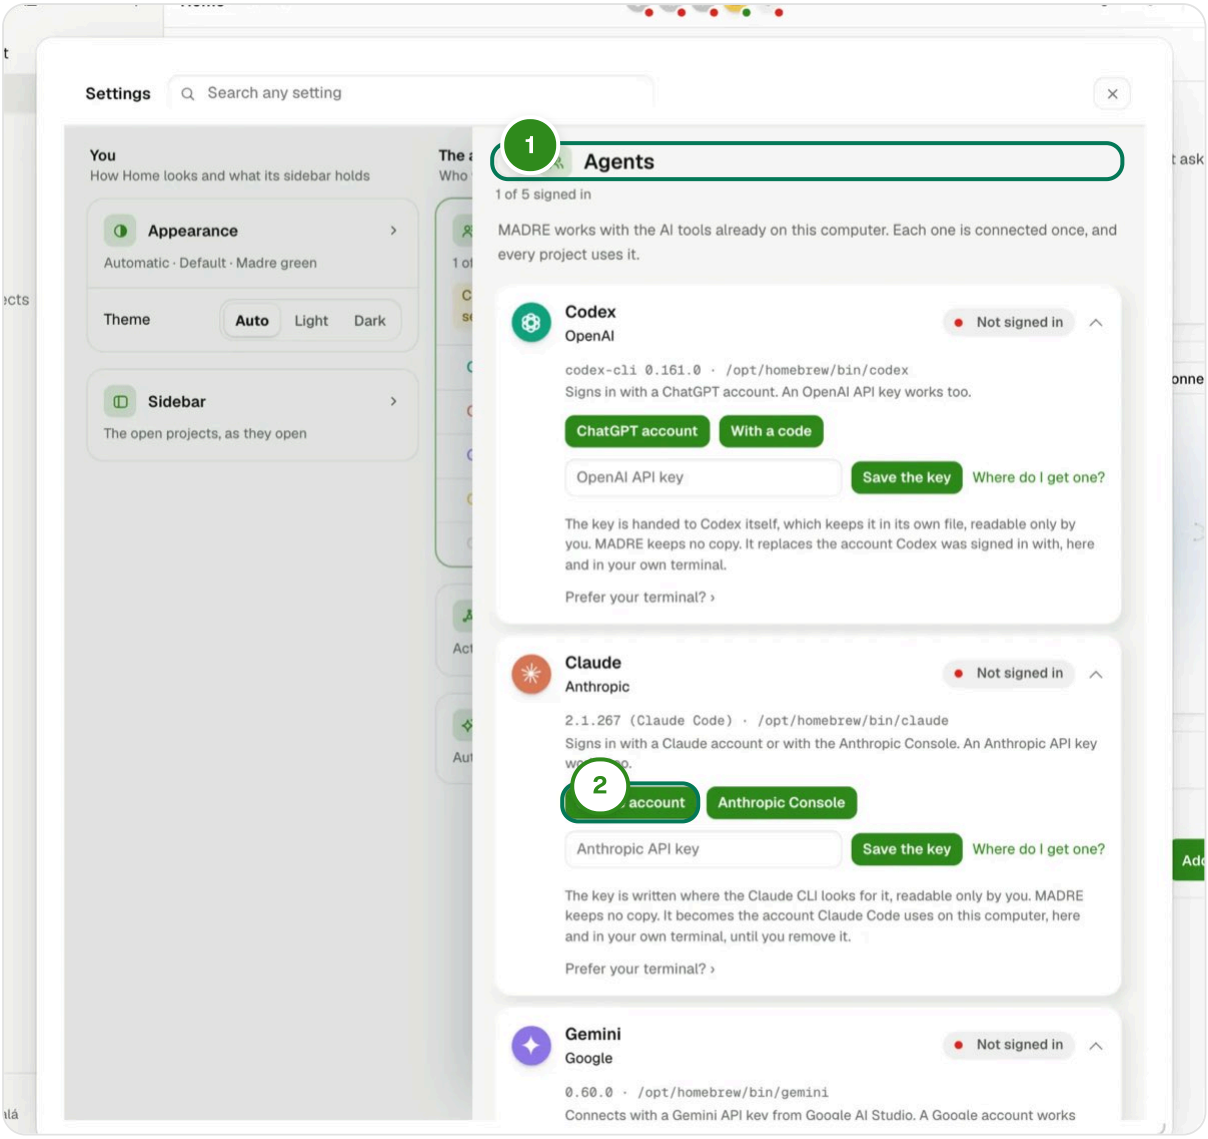

1. In Home, open Settings (the gear, bottom left) and go to **Your AIs**. The first time, the welcome takes you there.
2. Open an AI that says «Not signed in» and use its button, such as **Claude account**. You sign in with your account on that company's page.
3. Once it says «Connected», you can write to it.

If you're new

What's an API key?

It's a long password the AI company gives you so another program can use your account. You paste it only once. MADRE keeps it where that AI looks for it, protected under your user, and doesn't keep a copy.

Rather use the Terminal? These are the commands to sign in

Claude

```
claude auth login
```

Codex

```
codex login
```

Gemini

```
gemini
```

Then type `/auth` and choose how to sign in.

Grok

```
grok login
```

OpenCode

```
opencode auth login
```

Cursor

```
cursor-agent login
```

You're good if at least one AI says «Connected».

5 Write to an AI.

Open your project from Home. At the bottom there's the text box: you write to it the way you'd write to a person. In your message you can say which AI you're talking to and what you let it do.

2 min

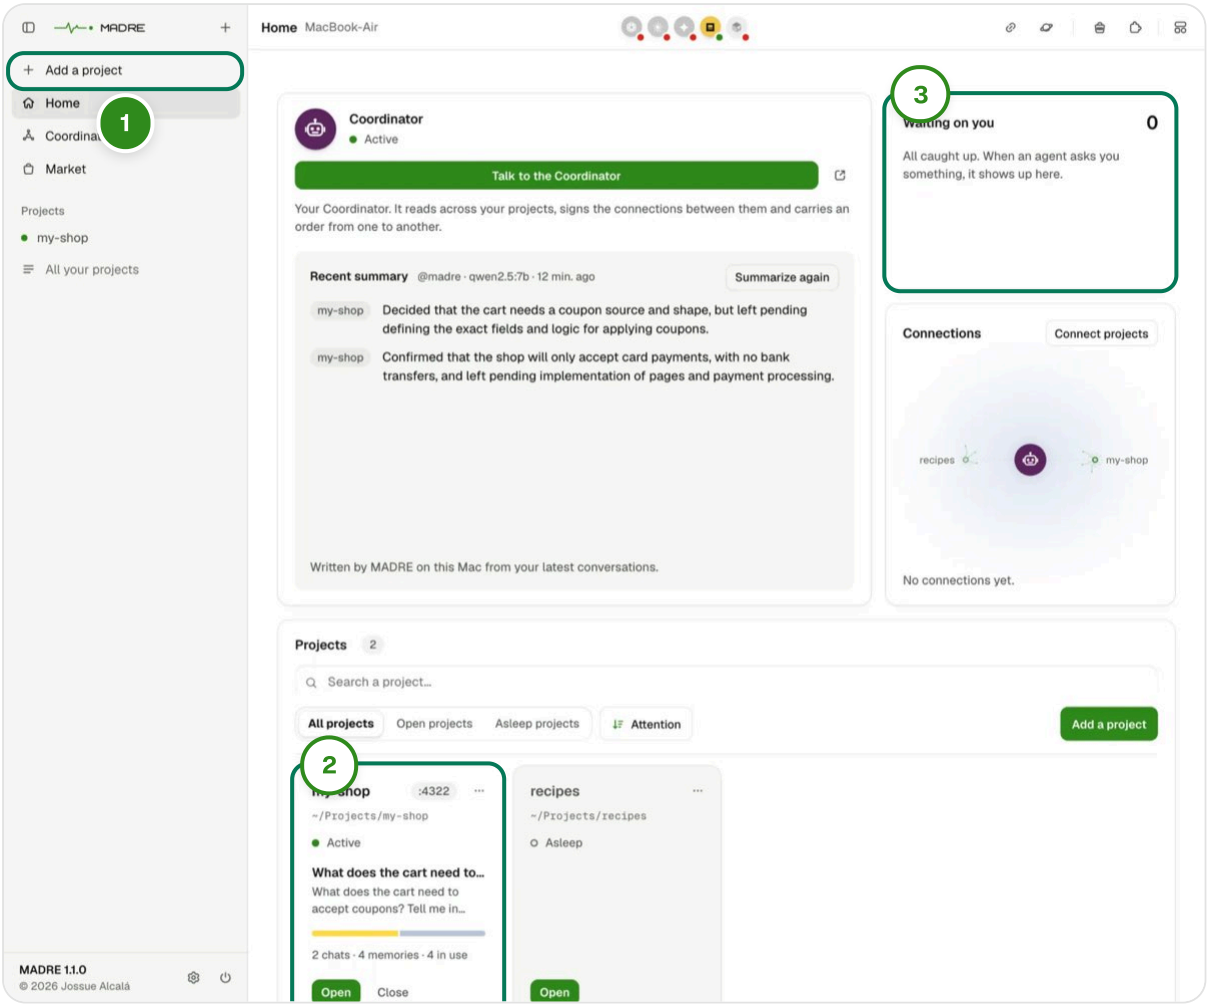

1. With **Add project**, in the left column, you pick a folder. Adding it changes nothing inside it.

2. It shows up under Projects. Click **Open** on its card to go in.

3. **Waiting for you** gathers what was left open: a question from an AI, a permission or a reminder.

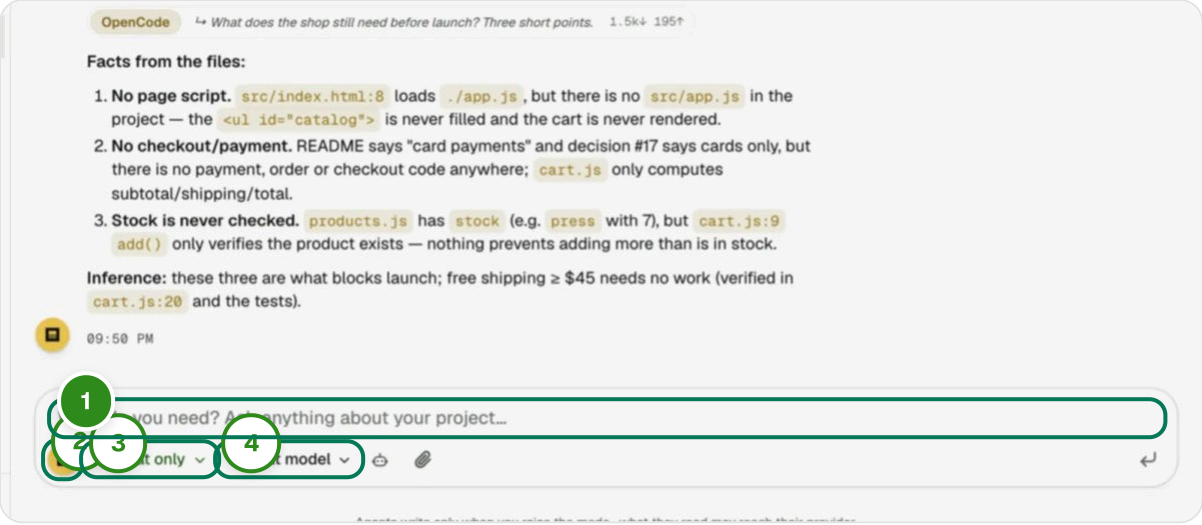

1. Type in the box and press Enter to send.
2. The circle on the left is the AI that will answer you. Tap it to pick another one.
3. The tag with # is the permission. With #1 the AI reads and answers you, but changes nothing.
4. Beside it you pick the model. If you're not sure, leave it as it is.

Try it

Copy a phrase and paste it in your project's box.

```
Explain what this project does in five lines, no jargon.
```

```
What three things would you improve first? Don't change anything yet.
```

```
Make me a list of what's left to finish this project.
```

Build your message

Pick who you're writing to, what you let it do and what you ask. Below you see how it's written in MADRE and what will happen.

This is how you write it in MADRE

```
@claude #1 Explain what this project does, no jargon.
```

Who answers

Claude, with your Claude account or your Anthropic Console.

With which permission

#1 Just chat

Does it ask you first?

No.

What happens

It reads your project and answers. It changes nothing.

Four symbols for the box

@ Who

`@claude` writes to Claude. Type `@` and the list shows up.

What you allow

From `#0` to `#4`. If you don't put one, it uses `#1`, which only chats.

! Which file

`!src/app.js:12-20` shows it only those lines of that file.

/ A command

Opens the command menu. For example, `/stopall` stops everything.

Enter sends the message. To start a new line, use ⌥ + Enter (Alt + Enter on Linux).

Each answer says in one line which AI replied, which message it answers and what it spent. If one takes too long, you stop it with the little square next to its timer. **Stop all** stops them all.

If you're new

What's a token?

It's the unit used to measure an AI's work. A token is roughly a piece of a word. Your plan or credit gets used up in tokens, and MADRE tells you how many each answer used.

You're good if an AI already answered something about your project.

6

You decide how much it can touch.

Every message carries a permission. By default the AIs only read. For them to create or change files, you raise the permission, and the last two always ask you first. No AI can raise it by itself.

1½ min

Try the permission picker

In MADRE you pick the permission on the tag with # beside the box, or by typing `#2` in your message. Here you can try it without worry: pick one and see what happens.

What each permission does

	#0 Not saved	#1 Just chat	#2 Can create	#3 Can edit	#4 Full access
Kept in the chat and the memory	—	✓	✓	✓	✓
Creates new files	—	—	✓	✓	✓
Changes the files you have	—	—	—	✓	✓
Runs commands: installs, tests, publishes	—	—	—	—	✓
Asks you before turning on	—	—	—	Once	Twice
Can be undone	—	—	✓	✓	Your files, yes

With `#4`, what leaves your computer, like uploading your code to the internet or publishing a page, can't be undone. Use it only with an AI you'd lend your computer to.

Below every answer that touches your project, a card says in plain words what changed. Right there you can look at it or undo it. **Undo** asks you first and puts your files back the way they were.

Try each permission, risk free

These examples only read or work in a `practice/` folder, so your project doesn't change. If you don't have that AI, swap the `@name` for one you do have.

#0 Not saved For asking something without it being kept in the chat or the memory.

```
@claude #0 What would you call an app for splitting bills with
friends? Give me 5.
```

What happens: It answers and, when you reload, it's gone.

```
@gemini #0 Explain what an API is like I'm 12.
```

What happens: None of this goes into the project's memory.

#1 Just chat For understanding your project or asking for opinions. Nothing changes.

```
@claude #1 Explain how this project is laid out: what's in each
folder.
```

What happens: It reads and answers.

```
@codex #1 What three things would you improve in this project? Don't
change anything yet.
```

What happens: You get ideas you can later ask for with #2 or #3.

#2 Can create For creating new files. It doesn't touch the ones you have.

```
@claude #2 Create a practice/ folder with a hello.html file that says
«Hello, MADRE» in big letters.
```

What happens: practice/hello.html shows up and opens beside the conversation.

```
@gemini #2 In practice/, create ideas.md with 10 feature ideas for
this project.
```

What happens: A new file full of ideas. Your project stays the same.

#3 Can edit For changing files. It asks you first and saves a point so you can undo.

```
@claude #3 In practice/hello.html, make the background blue and
center the text.
```

What happens: It changes the file you made with #2. Don't like it? Undo.

```
@codex #3 Add a button to practice/hello.html that changes the
background color on every click.
```

What happens: When it's done you see which files it touched.

#4 Full access

For also running commands, like installing, testing or using git. It asks you twice. Use it only when you need to.

```
@claude #4 Inside practice/, start git and save a version of what you
created. Don't upload it anywhere.
```

What happens: It runs commands on your computer. That's why you tell it clearly what it shouldn't do.

```
@codex #4 Run the project's tests and tell me which ones fail. Don't
change any file.
```

What happens: It runs the tests and explains the result.

If you're new

What's git and why does #3 need it?

Git is a program that keeps versions of a folder. MADRE uses it to keep the point Undo goes back to. That's why `#3` and `#4` need your project to use git. If an AI won't let you raise it to `#3`, check its max permission in Settings › Permissions.

You're good if you tried a permission in the practice/ folder and you know where Undo is.

7 Something not working? Use Diagnostics.

Diagnostics is in your project's bar. Write what's happening in your own words, for example «gemini won't sign in», and it tells you how to fix it. Often it comes with a button that fixes it for you.

1 min

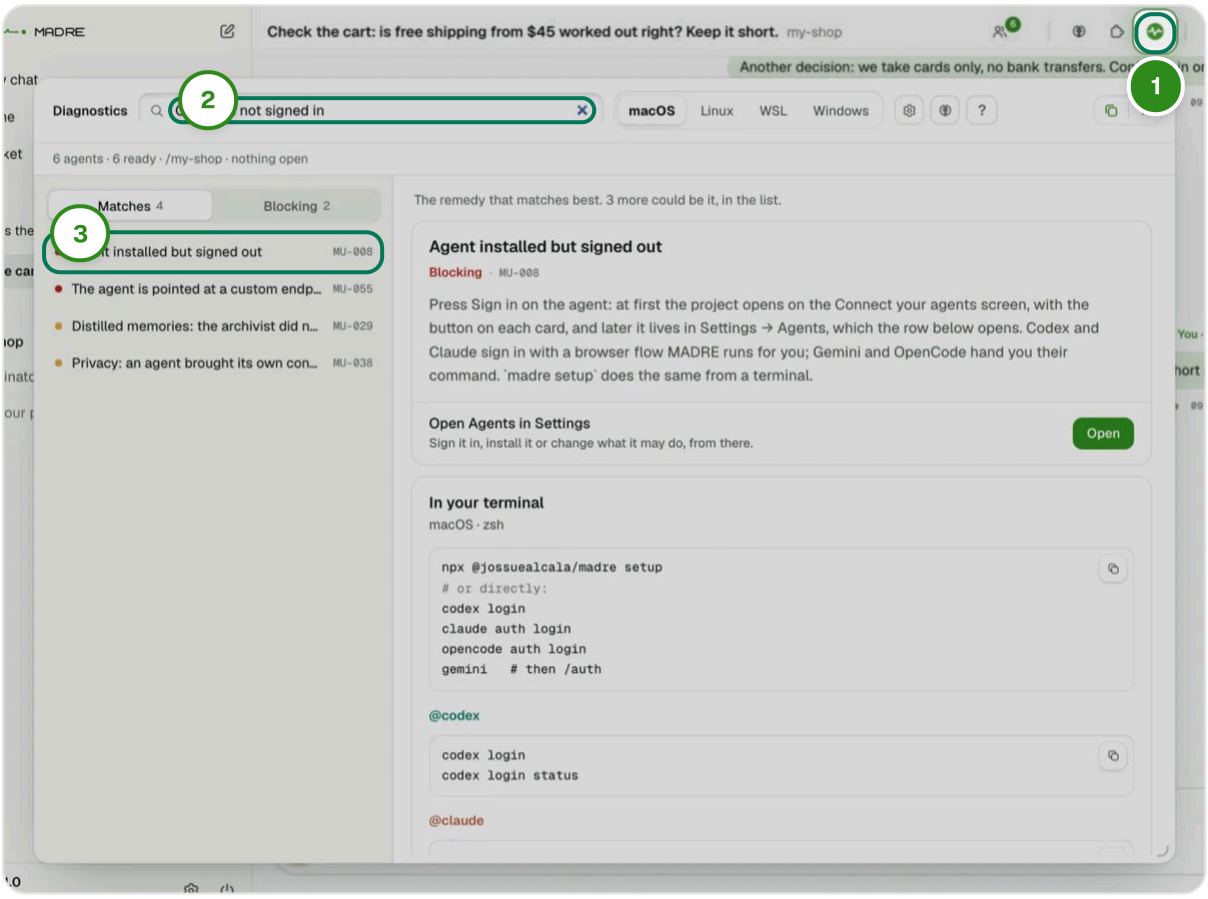

1. Click **Diagnostics**, on the bar.
2. Type what you see, in your own words.
3. Open the case that looks like yours and follow what it says.

What happens most when you start

An AI says «Not signed in».

Go to **Settings › Your AIs** and use its button to sign in.

MADRE can’t find an AI.

In **Settings › Your AIs**, click **Install**. MADRE runs the command for you.

An AI changed something I didn’t want.

Click **Undo** on that answer’s changes card. Your files go back to how they were.

An AI finished without writing anything.

MADRE tells you in its answer and, when it knows, tells you why. Ask again or pass the question to another AI.

Gemini says it hit its limit or is overloaded.

Wait a minute and ask again, pick another model or pass the question to another AI.

Gemini rejects your personal Google account.

Google no longer accepts Gemini CLI with free personal accounts. Get an API key at Google AI Studio and paste it in **Settings › Your AIs › Gemini**.

@madre doesn’t show up.

Ollama is off or has no model. In Home, open **Settings › Local AI**, turn it on and download the model.

I want the new version.

The app tells you when there’s one and you install it with one click.

Windows says “Windows protected your PC”.

It’s because the digital signature is missing, not because of the installer. Click **More info** and then **Run anyway**.

On Windows, Ollama won't install.

If you set MADRE up by hand in Ubuntu, it's missing `zstd`. In the Ubuntu window paste this and click **Install Ollama** again:

```
sudo apt-get install -y zstd
```

Got an error message? Look it up in Known errors. Still not fixed? Tell me what happened. You don't need to create an account and you'll see my reply right there.

You're good if you know where to ask for help, inside and outside MADRE.

Learn by doing.

These are the five practices in **First steps**, the panel MADRE shows you when you start. Each one takes a couple of minutes and you can do them in any order. Copy the phrase, paste it in your project's box and mark it when you're done.

Make your first page

Ask an AI for a page and watch it appear. You learn how to ask, and how to let it create files.

```
#2 Make me a simple page with a title, a short welcome and a button.
Explain how to open it.
```

What an AI creates, like a page, a picture or a PDF, opens by itself beside the conversation.

Ask two AIs the same thing

Compare their answers side by side. You learn to talk to each one with @ and its name.

```
@claude @codex What would you
improve first in this
project? Your own idea, in
three lines at most.
```

You need two AIs connected. Swap the names for the ones you have.

Undo a change

Ask for a change and take it back with Undo. You learn that nothing breaks for good.

```
#2 Create a file called
note.txt that says hello.
Then I will undo it.
```

When it answers, press **Undo** on its reply.

Talk to the AI on your computer

@madre answers with what your project remembers, with no internet and at no cost.

```
@madre What do you remember
about this project?
```

First you need the local AI. See Your local AI.

Get something from the Market

A module, a skill or a connector, with one click. You learn that MADRE grows with what you add.

Open the **Market** in the left column. Try **Before you push**: before your code goes to GitHub, it checks that no keys or private files go with it.

Once you use it.

You don’t need this to get started. Open it when you need it: each topic takes a minute and almost all come with the MADRE screen and a number on what you have to press.

Work with your AIs

Use several AIs in the same chat.

All the AIs read the same conversation.

1 min

That’s why you can ask one to continue what another one did, without explaining anything again.

What having several is good for

One does, another reviews

Claude makes a change and you ask Codex to review it. Codex already saw everything that happened in the chat.

```
@codex review what Claude just did and tell me if it missed anything.
```

The same question to several

Put several names at the start of the message. You see your question once and each AI answers separately.

```
@opencode @claude what is the cart missing to take coupons?
```

A round table

Up to four AIs answer, one after another, and @madre sums up where they agree and where they don't. You need your local AI.

```
@madre, ask the team: where should I start improving this project?
```

If one runs out of plan

If Claude hits its usage limit, you write to another one and it carries on where it left off.

```
@gemini keep going with what Claude was doing.
```

To look after your plan

Give the heavy work to the AI you pay for. Ask your local AI, which is free, about what was already discussed.

```
@madre what did we decide last time about the design?
```

> **Claude, OpenCode and Codex can ask other copies of themselves for help**, called subagents. They do it with the same permission you gave them, from `#1` to `#3`. At the end, the answer says how many helped, for example «Helped by 2 subagents», and there you see what each one did.

How they remember what was said

Before answering	Each AI gets the latest messages of the chat, even if another one wrote them.
Every 10 messages	Or after 10 minutes with no activity, MADRE writes down what matters in short notes: what was decided, facts about the project, how you like things and what's still pending.
Days later	When they're needed, those notes are passed to whichever AI is answering, even if another one made the decision.

Your local AI writes those notes, at no cost. If you don't have one, the cheapest AI you have connected writes them. You see them and delete them in the **Memory map**.

You're good if you know how to ask one AI to continue what another one did.

Your chats, on the left.

45 s

Each project can have several chats.

They're in the left column, like in any messaging app. They all share the project's memory: what was decided in one, the AI knows in the others.

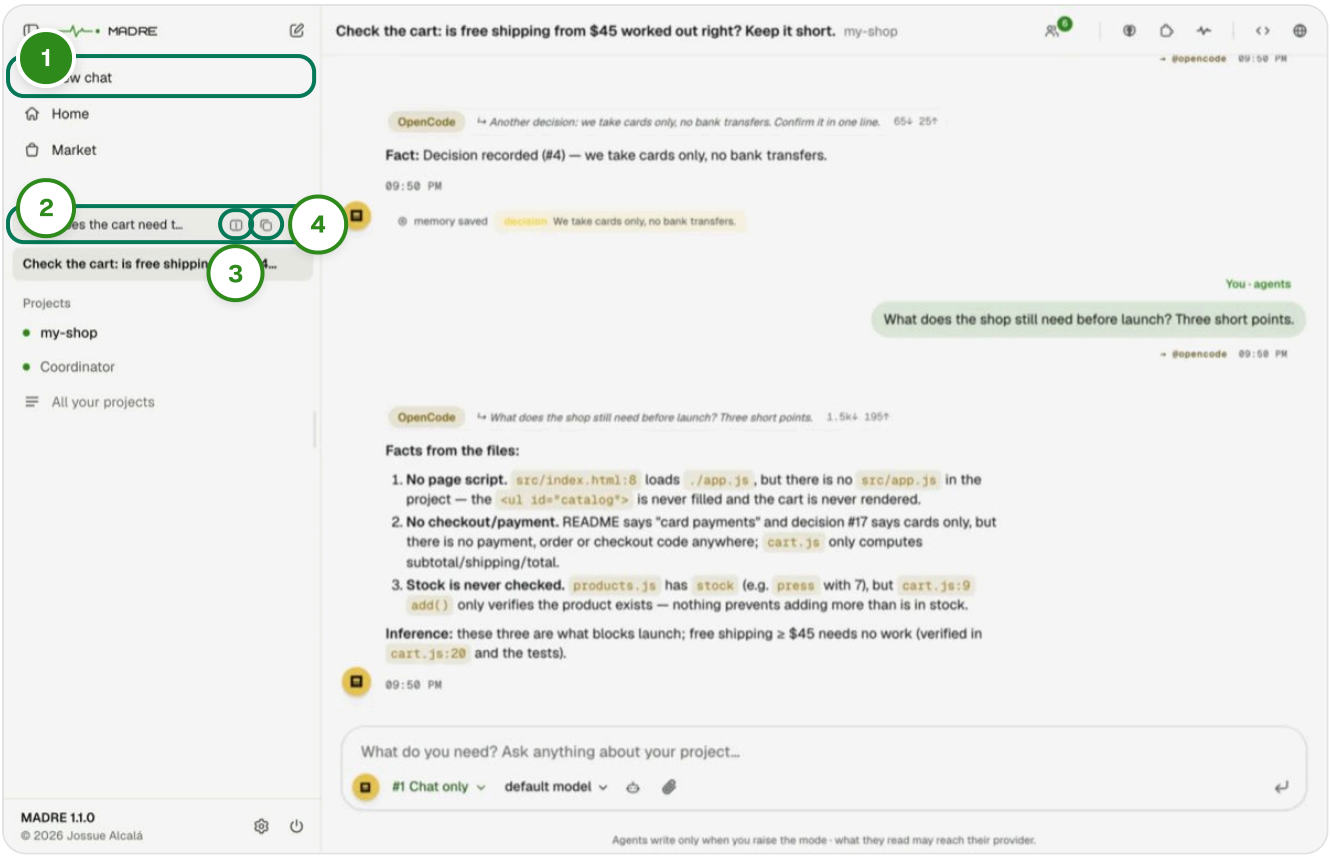

1. «New chat» starts a conversation.
2. Your chats stay below it, the latest on top. Click one to go back to it.
3. Hover over a chat: the first button opens it beside, in the other half of the window.
4. The second opens it in its own window, which you can move. The ✕ deletes it.

> The button at the top left hides the column so you have more room. In a narrow window, the column opens over the conversation.

You're good if you opened a new chat and went back to the previous one.

Your files, next to the chat.

Files opens your project to the right of the chat.

1 min

What each AI changed is marked in its color, so you know who changed what.

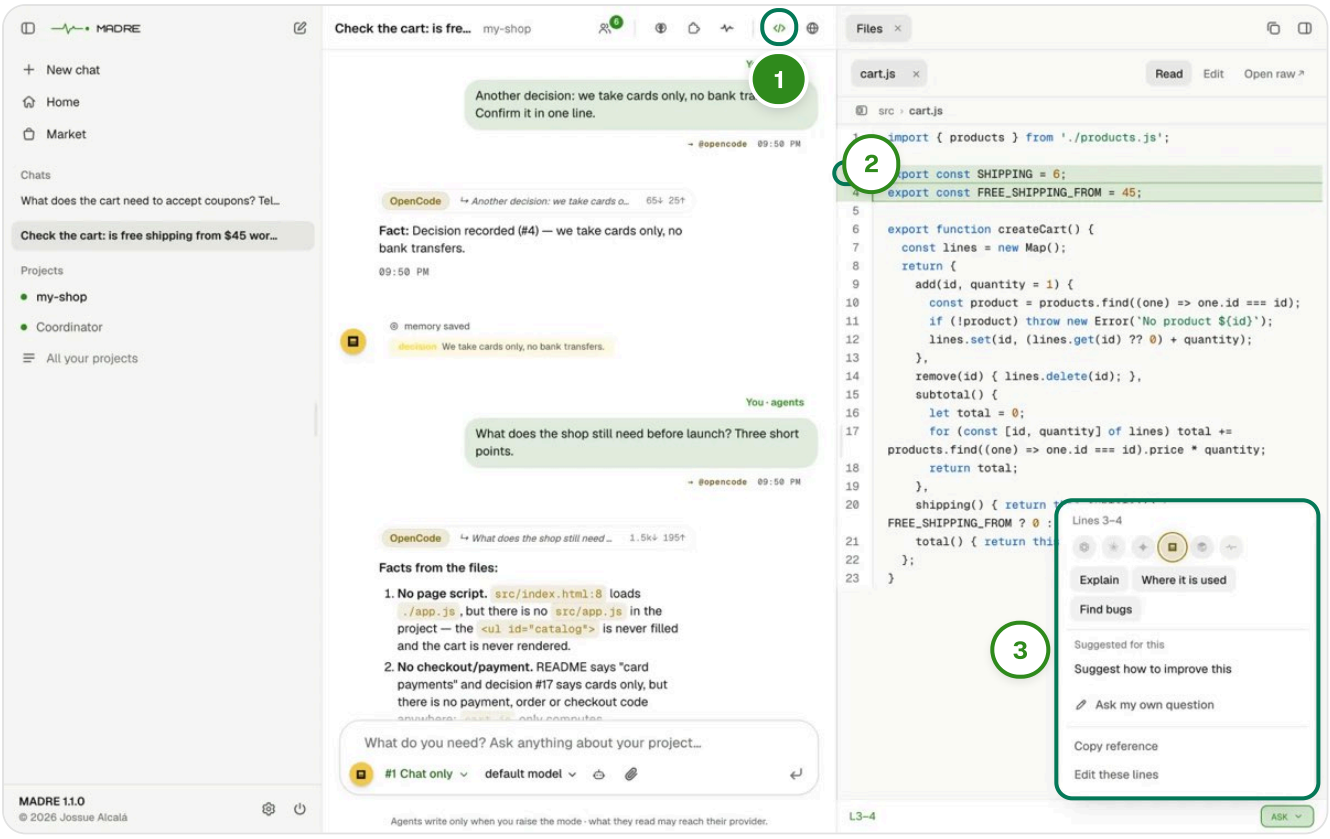

1. Press Files (</>) in the bar at the top and open a file from the list.
2. Click the number of a line. With ⇧ you choose several.
3. Press ASK: choose the agent and what you want, or copy the reference to quote it in the chat.

> **Want to change it yourself?** Switch to **Edit**. ⌘ S saves, ⌘ F finds and ⌘ P finds a file by its name.

You're good if an AI explained a few lines of your project to you.

The project's memory.

MADRE keeps what gets decided in your project: names, decisions and how you like things. 45 s

When you ask an AI something, MADRE passes it what it needs to know, even if another AI said it weeks ago.

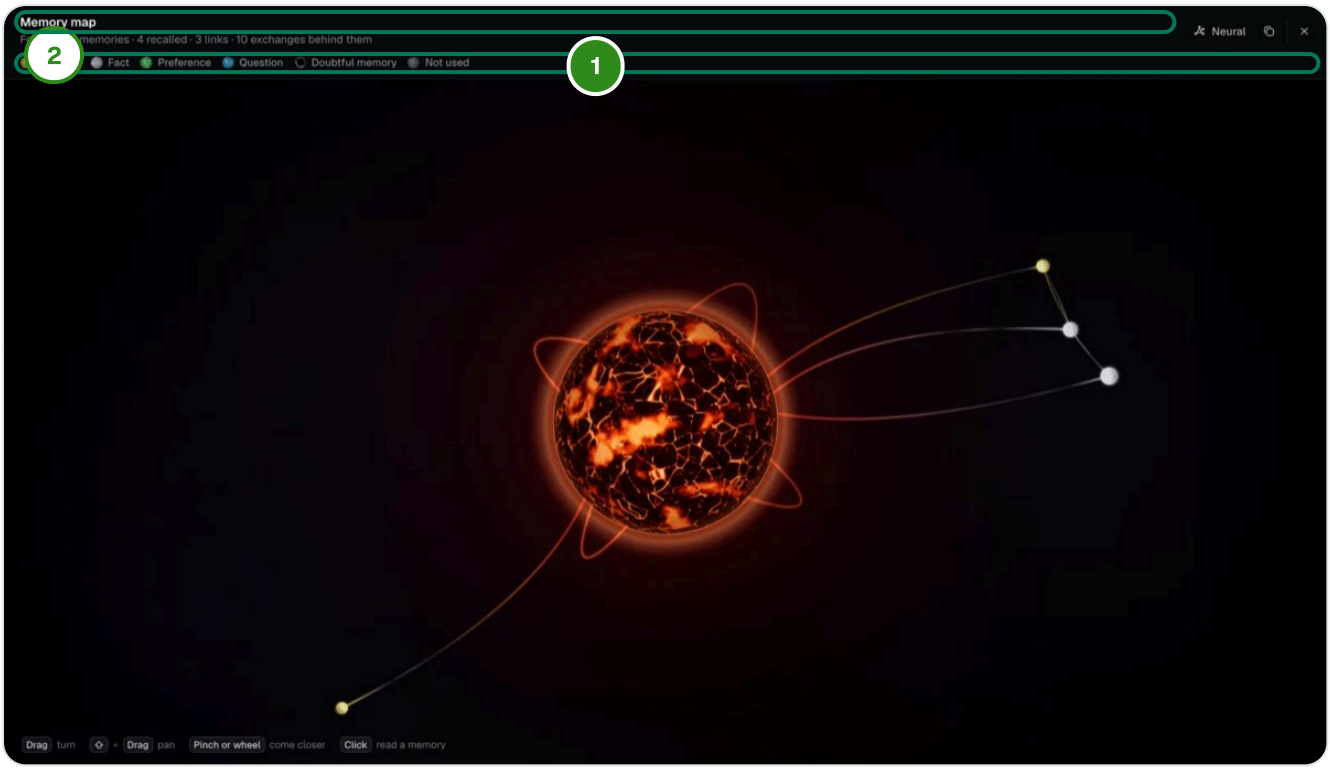

1. Press Memory map in the bar. Type the project's name and it opens full screen.
2. Each dot is something MADRE remembers. Its color says whether it is a decision, a fact or a preference.
3. Click one to read it. If it is no longer true, you forget it from there.

Try it

Tell one AI, then ask another one about prices.

```
Note this: prices always include tax.
```

> With `#0` Not saved nothing is kept in the memory.

You're good if another AI knew something you told the first one.

Notes and reminders.

A note is like a colored slip of paper stuck to your project.

45 s

A reminder lets you know on the day and at the time you say.

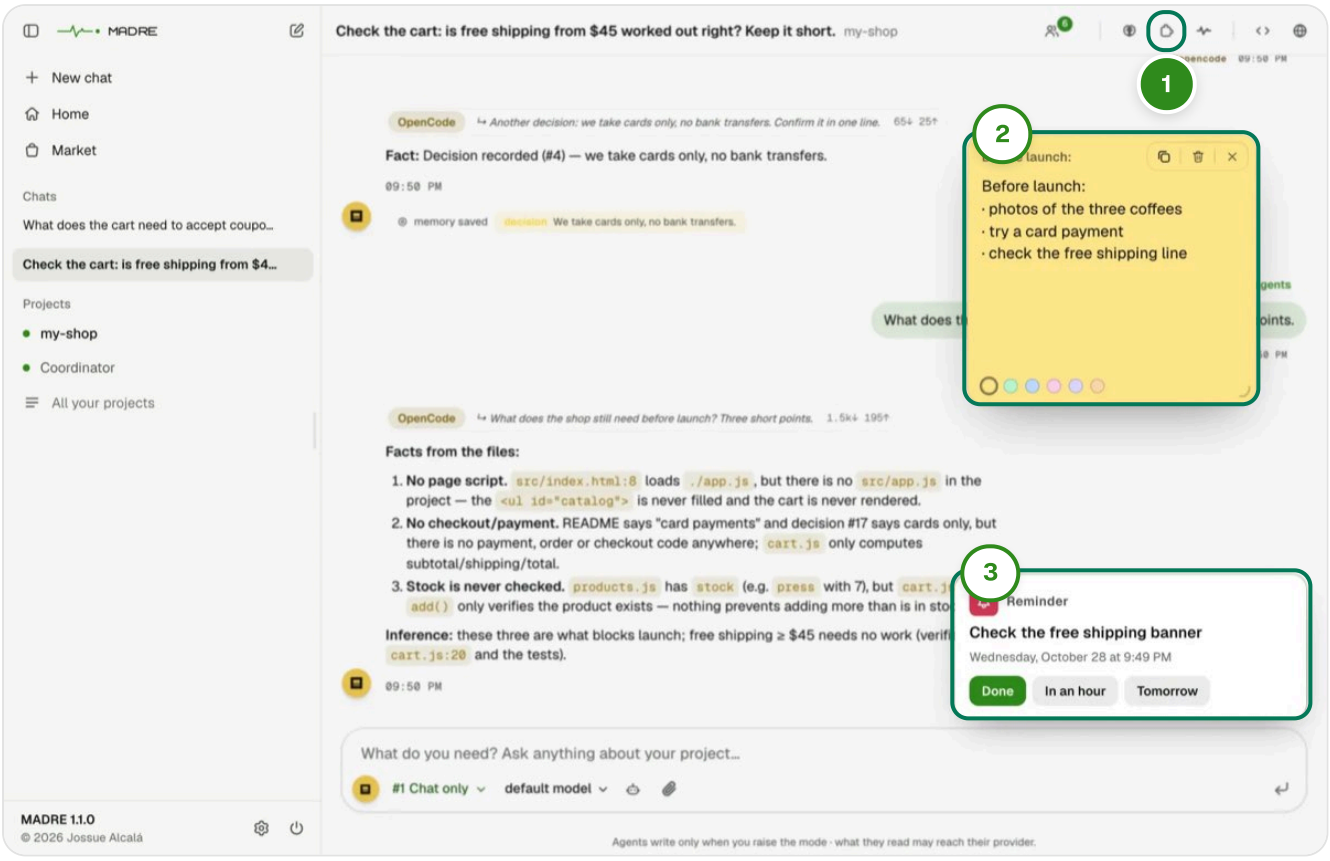

1. Turn on Notes and Reminders in Modules, in the bar at the top.
2. Each note opens its window: write, change its color or throw it away.
3. When a reminder's time comes, MADRE shows it in the project and in «Waiting on you», in Home: «Done», «In an hour» or «Tomorrow».

Try it

You can ask any AI for one.

```
Remind me tomorrow at 10 to check the banner.
```

`/notes` and `/reminders` ask the AI you have chosen for the note or the reminder. If you write «I have to…», MADRE offers to keep it.

You're good if you have your first reminder.

Copilot: let the AI work on its own for a while.

For a long job, Copilot lets the AI move forward without asking your permission at every step.

30 s

It can edit files and run commands that only read. If it needs to run your tests or build, it waits for your click.

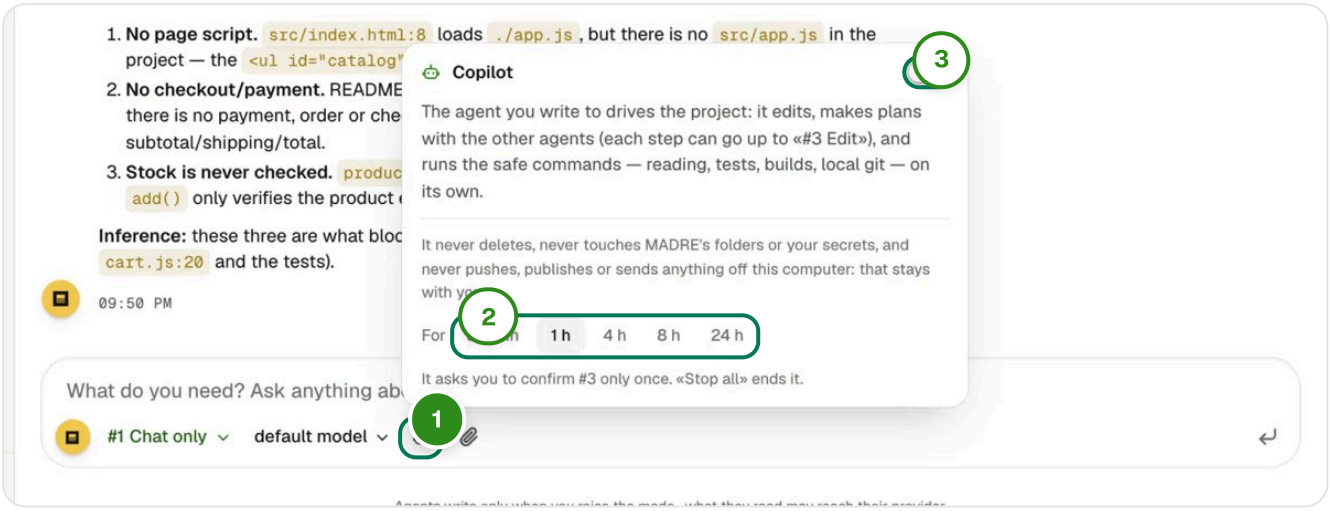

1. Press the robot beside the box.
2. Choose for how long: from 30 minutes to 24 hours.
3. Turn it on and write the task. It asks you to confirm only once.

> It asks you to confirm only once. What it deletes comes back at the end, and publishing anything still needs `#4`. It turns off by itself when the time is up, or with **Stop all**.

You're good if the AI moved forward several steps without asking you.

Running now: what's working right now.

Running now shows you what MADRE is doing in all your open projects: which AI is answering, which plan is moving forward and which server is on.

30 s

You can also stop it from there.

1 Find the beating line

It’s on the bar at the top, in every project and in Home. The number says how many things are working, and the line moves while something works.

2 **Open it**

Under **Now** you see each AI that’s answering, the plans with their steps, the servers and what’s being installed. Below is what finished a moment ago and the upcoming reminders.

3 **Go there or stop it**

Open a row. **Go to the chat** takes you to that message. **Stop** stops only that.

> **Stop everything** stops everything in all your projects, and asks you to confirm.

You’re good if you watched an AI work from Running now.

Add features

The Market and your modules.

The Market adds features to MADRE: Notes, Reminders, Preview, Before you push and more.

1 min

Every module tells you what it does and what it needs before you install it.

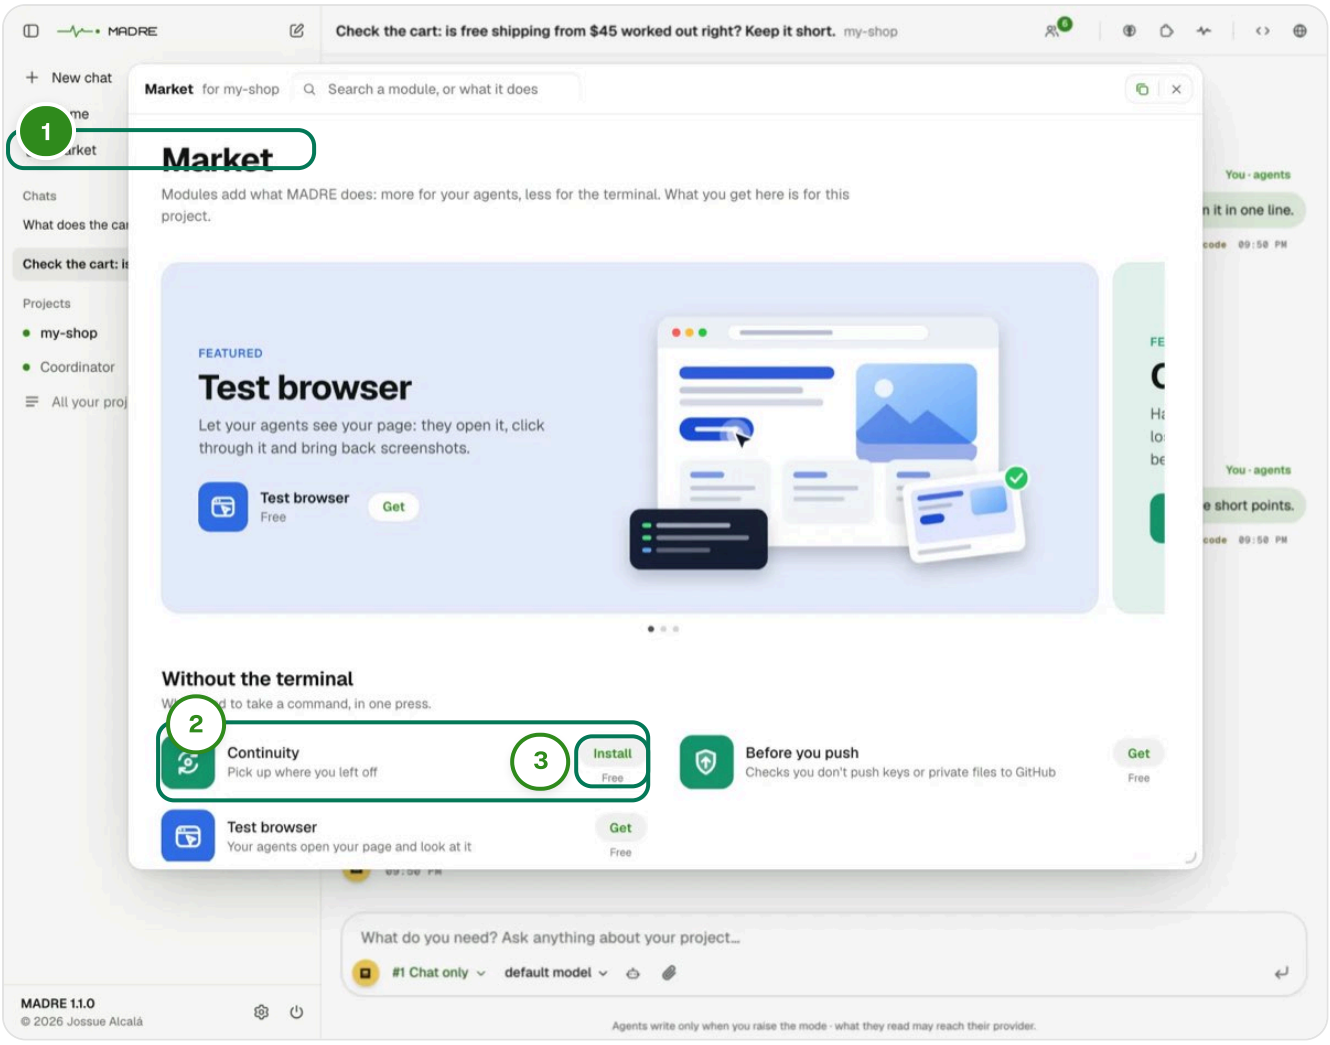

1. Press Market in the left column.
2. Tap a module to read what it does, what it touches in your project and who made it.
3. Press Get or Install. Afterwards you turn it on or off in Modules, in the project's bar.

> **Try Before you push.** Before your code goes to GitHub, it checks that no keys or private files go with it.

Who makes the modules?

Some are made by MADRE's author and others by the community. Each one says who made it and its license, and MADRE checks its signature before installing it. New modules show up on their own, without updating MADRE.

You're good if you turned a module on and saw what it does.

Skills: teach them how you work.

A skill tells your AIs how you work: your quote template, the order of your reports, what to ask before giving a price.

1 min

You write it once and they all follow it, because MADRE is the one that talks to them.

1 Ask for it

Type `/habilidad` and what you want them to do, in your own words. It needs the `#2` Can create permission.

2 Read it before you keep it

The AI writes it and MADRE offers it in the conversation. Click **Read it**: it's one page, a minute to read.

3 Choose when it comes in

In **Modules › Skills**: when it fits, always on, or off.

Try it

```
/habilidad when I ask for a quote, use my template and put the tax apart.
```

> **A skill gives no permission.** It doesn't raise the message's permission or let an AI touch more. Already have skills from Claude or Codex? With **Bring over what you already have** you pick which ones.

You're good if you kept your first skill.

Connectors: let your AIs use other apps.

A connector lets your AIs use other apps, like GitHub, Notion or your calendar.

45 s

They're also called MCP. You connect them once in each project.

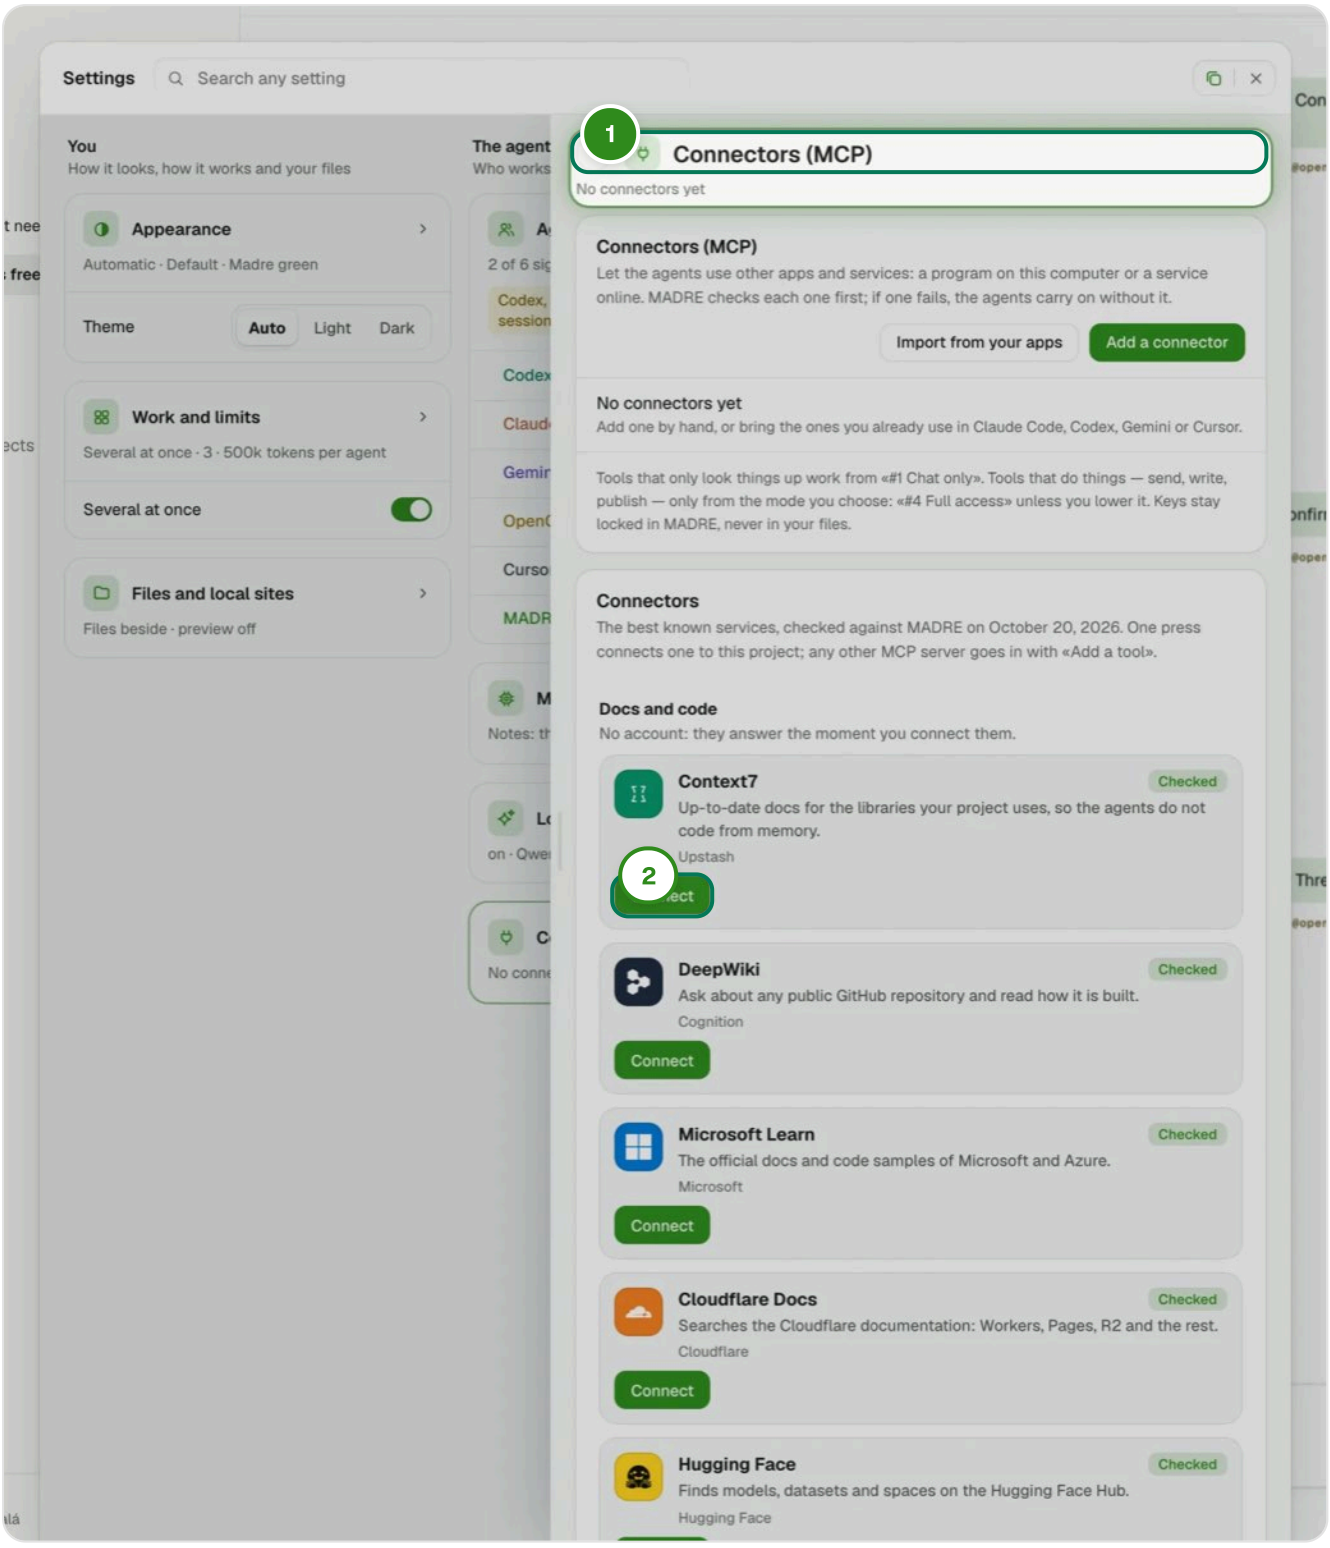

1. Open Settings › Connectors (MCP).
2. Choose one from the catalog and press Connect. Some open their page so you can sign in.
3. Press «Check the connection» to confirm it answers.

Try it

With GitHub connected, ask an AI:

```
Which issues do I have open on GitHub?
```

> **Already had connectors in Claude Code, Codex, Gemini or Cursor?** With **Import from your apps** you bring them into MADRE.

You're good if a connector answers.

Continuity: another AI picks up where you left off.

45 s

Continuity saves, inside your project, where the work stands: what was decided, what's left and points to undo.

If you switch AI, editor or computer, the next one starts from there.

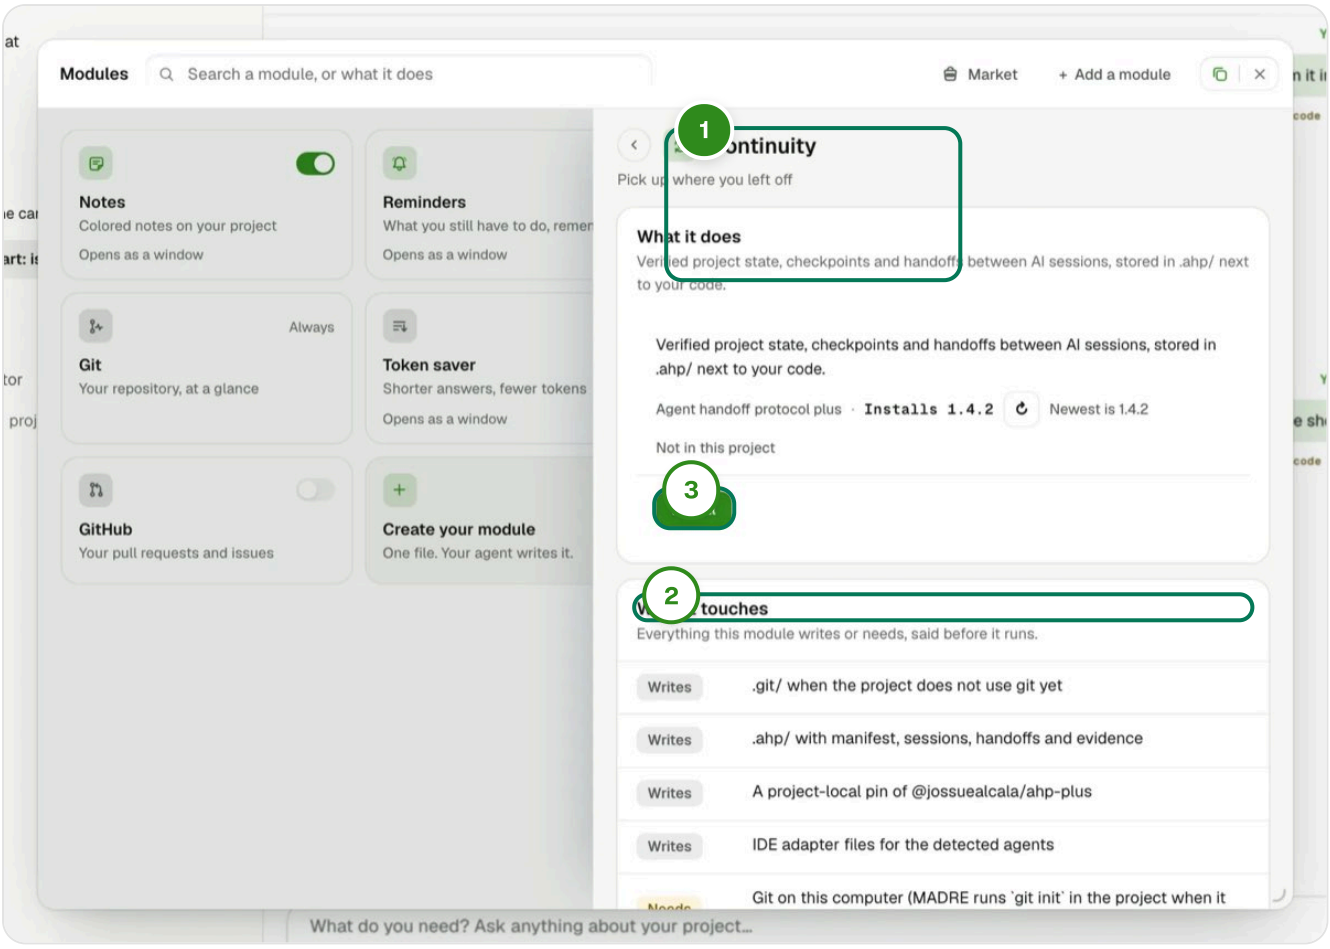

1. In Modules, tap Continuity.
2. «What it touches» shows what it writes: a .ahp/ folder inside your project, in plain text.
3. Press Install. Your project needs to use git, with at least one commit.

And outside MADRE?

Outside MADRE it's called AHP+. Claude Code, Codex and OpenCode read it too, and the `.ahp/` folder is saved in git along with your project.

You're good if your project has the .ahp/ folder.

Your local AI and your projects

Your local AI: a free AI on your computer.

@madre is an AI that runs on your computer with Ollama, a free program.

1 min

It doesn't use up your plan, it works offline and what you ask it doesn't leave your computer.

What it does

- ✓ Answers with what the project remembers and tells you where it got it.
- ✓ Runs the round table and gives you the summary.
- ✓ Writes the memory notes, at no cost.

What it doesn't

- — It doesn't create or change files.
- — It doesn't run commands.
- — It doesn't make things up. If something was never discussed, it tells you.

macOS

1 **Install Ollama**

In Home, open **Settings › Local AI** and click **Install Ollama**. If you have Homebrew, MADRE installs it. If not, **Get Ollama** takes you to ollama.com to download it.

Linux

1 **Install Ollama**

In Home, open **Settings › Local AI** and click **Install Ollama**. MADRE uses the official installer.

Linux asks for your password to install it. MADRE never asks for it or keeps it: it gives you the line to copy.

Windows

1 **Use the Ollama on Windows or install it from MADRE**

If you already have Ollama on Windows, in the welcome choose **Use the Ollama on Windows** and MADRE finds it. If you don't have it, click **Install Ollama** in **Settings › Local AI**.

2 **Turn it on**

If it says «Stopped», click **Start Ollama**.

3 **Download the model**

The model is the AI you're going to download. The one that says **Recommended for your computer** is the right one for your machine. Click **Download**. It takes a few minutes, depending on your internet.

4 **Try it**

Type `@madre What do you remember about this project?` and it answers with what the project remembers.

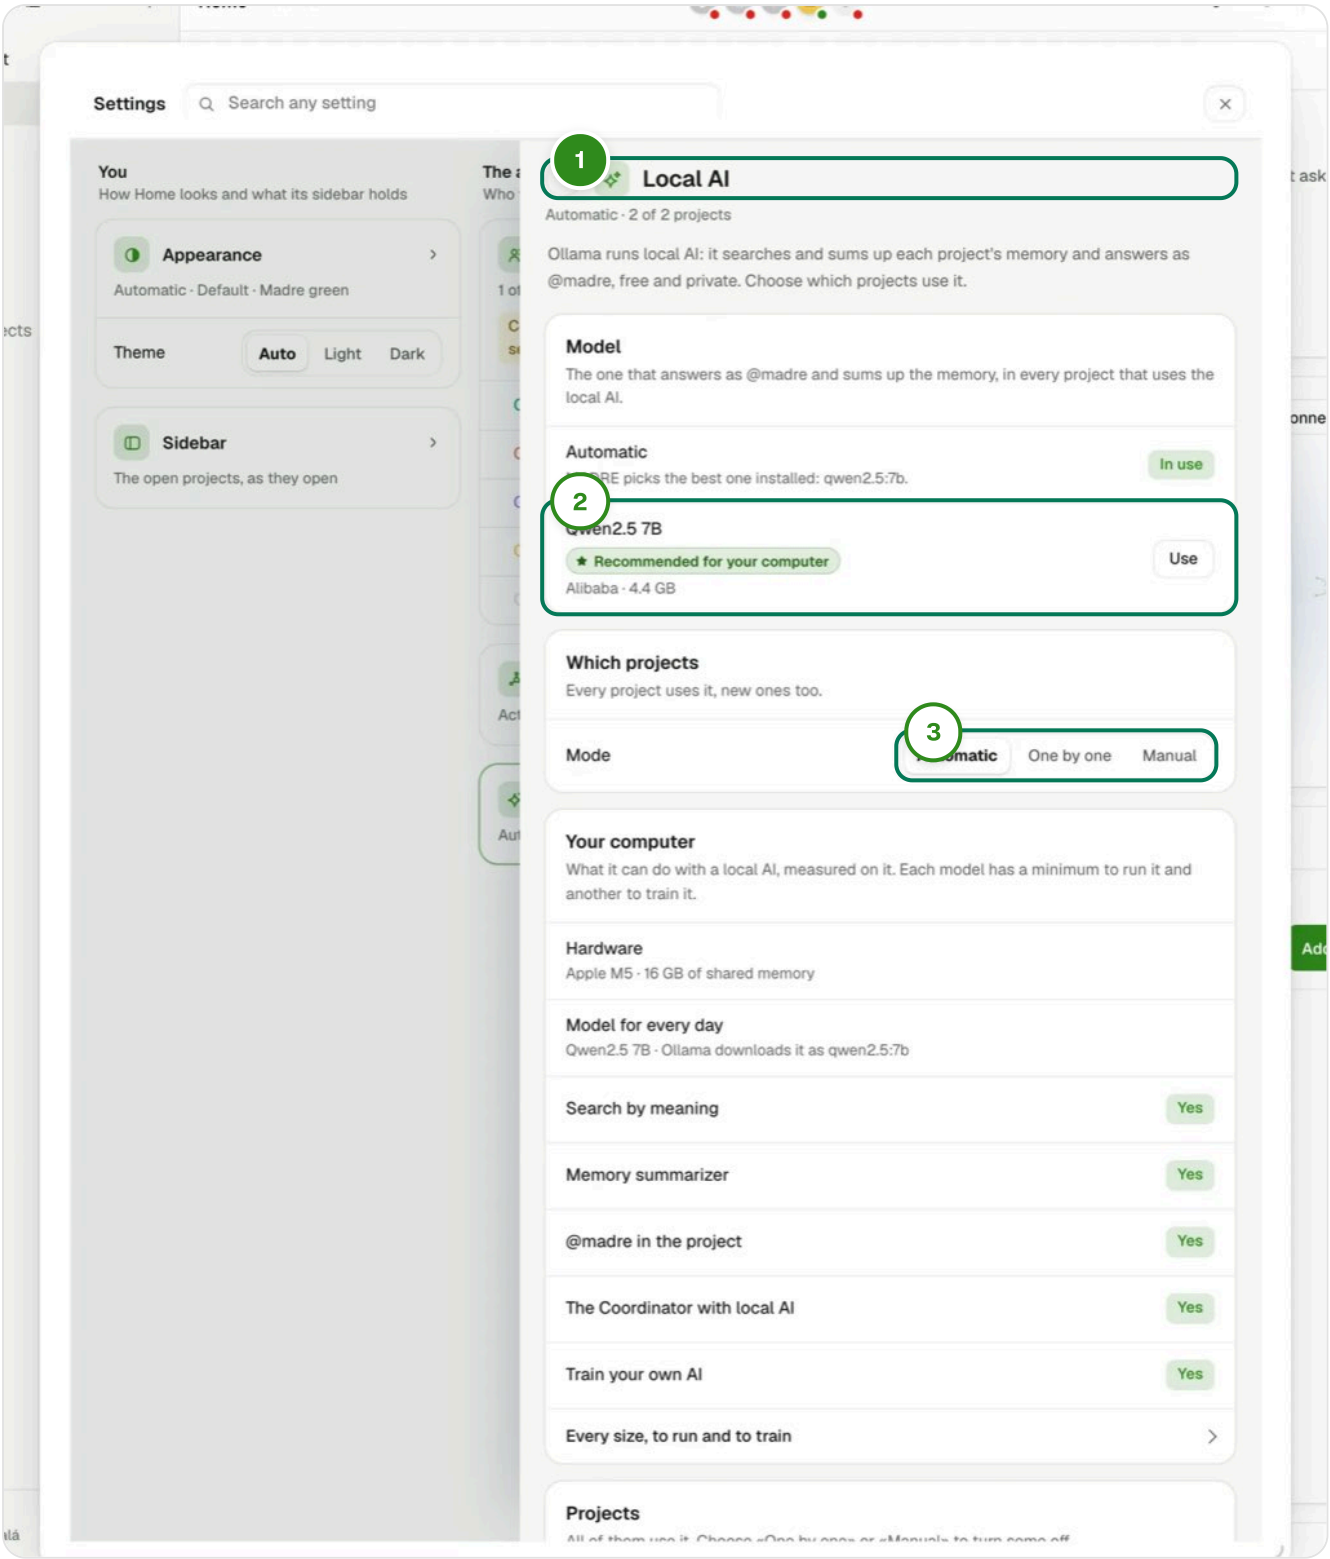

1. In Home, open Settings › Local AI.

2. Download the model marked «Recommended for your computer»: it is already sized for your Mac.

3. Choose which projects use it.

If you're new

What's Ollama? And a model?

Ollama is the program that makes the AI run on your computer. The model is the AI itself. MADRE only offers you the models that fit your machine, and if you later install a bigger one, @madre uses the best one it finds.

You're good if @madre shows up among your AIs and answers you.

Train your AI on what your project knows.

This is the most advanced part of the manual.

1½ min

The first time, set aside an hour. MADRE doesn't train anything on its own or send anything online. You take out what your project learned, train a small model on your computer and @madre starts answering with it.

If you're new

What does training a model mean?

You take a small model that already knows how to talk and teach it with examples from your project: your questions and the answers that helped you. You don't start from scratch, you just adjust it. At the end you have a file that Ollama can use like any other model.

What goes in the training file

What goes in	Your questions in the project with the answer you got, and the memory's notes turned into question and answer.
What stays out	Nothing from `#0` Not saved, no answers you marked as bad, none from @madre and none that are a single line.
What gets deleted	Before the file is saved, paths, names, emails and keys are deleted.
How many examples you need	Under 200, it learns the tone and little else. From about 1000 it's worth it. Under 50, the training doesn't start.
What computer you need	A Mac with an Apple chip is the tested option. On Linux with an NVIDIA card it's being tested. With neither, the model works but can't be trained.

1 Export

In your project, open **Settings › Memory** and, under **Train your own AI**, click **Export**. It tells you how many examples came out.

2 Follow the step by step

Open **Step by step**, in that same section. MADRE checks your computer, suggests the size and gives you each command with your folders.

3 Set up Python (only once)

It installs in a separate space so nothing else on your computer changes. On Linux, **Step by step** gives you the version for your card.

```
python3 -m venv ~/.madre-train && source ~/.madre-train/bin/activate &&
pip install mlx-lm
```

4 Use it

When it’s done, in **Settings › Local AI** click **Check again**. @madre now answers with your model.

> **Before you trust it**, ask @madre ten things that were decided in your project and five that never came up. It should get the first ones right and tell you the others never came up.

You’re good if your project has its training file.

The Coordinator: an AI that sees all your projects.

The Coordinator is in Home.

30 s

You can ask it about any of your projects. If you connect two projects, one can look up what the other learned.

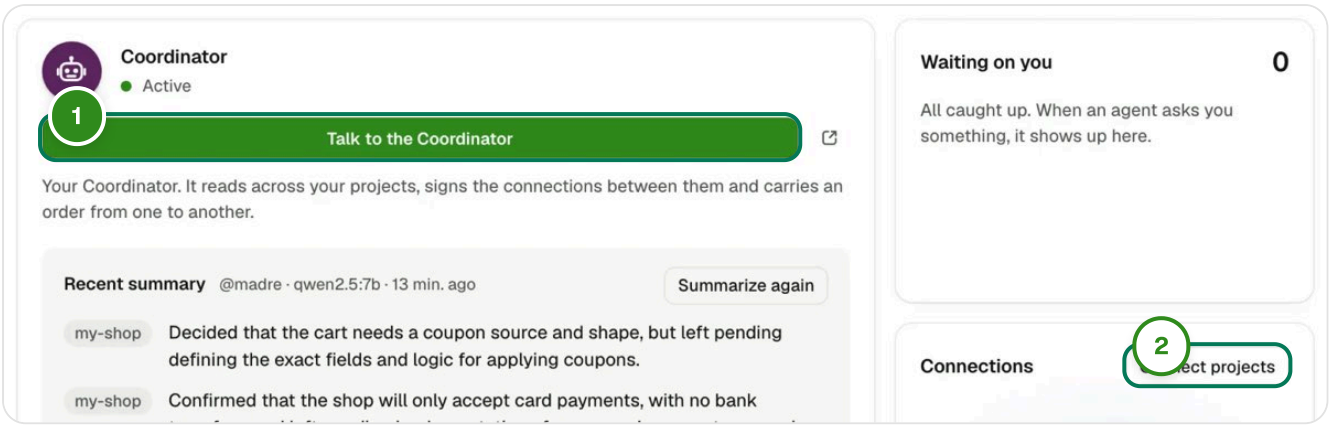

1 In Home, press «Talk to the Coordinator».

2 Press «Connect projects» and choose two. You decide what it may look up and until when.

Try it

In Home, ask it:

```
Which project was I on yesterday?
```

> From any project you ask it with `/coordinator`. Memory notes travel through a connection, never your files or your chats. It isn't part of the app's free mode: I explain it in The license and what it costs.

You're good if the Coordinator answered you about one of your projects.

The app

The app on your computer.

MADRE 1.1.4 is the first version as an app.

45 s

It brings its own windows, lets you know what's happening and updates with your permission.

What the app brings

Alerts	It lets you know when an AI finishes, even if you're in another window.

Its own browser	To see the page you're making, with its size as you change the width and the page's errors one click away.
Its own terminal	For those who need it, inside the app.
Updates	When there's a new version, it tells you and you install it with one click.
On a Mac, the menu bar	If you close the window, MADRE stays there. Its icon turns light or dark with your system's theme.

Files, Your servers and the other tools open to the right of the chat, each in its own tab. You can take any of them out to a separate window.

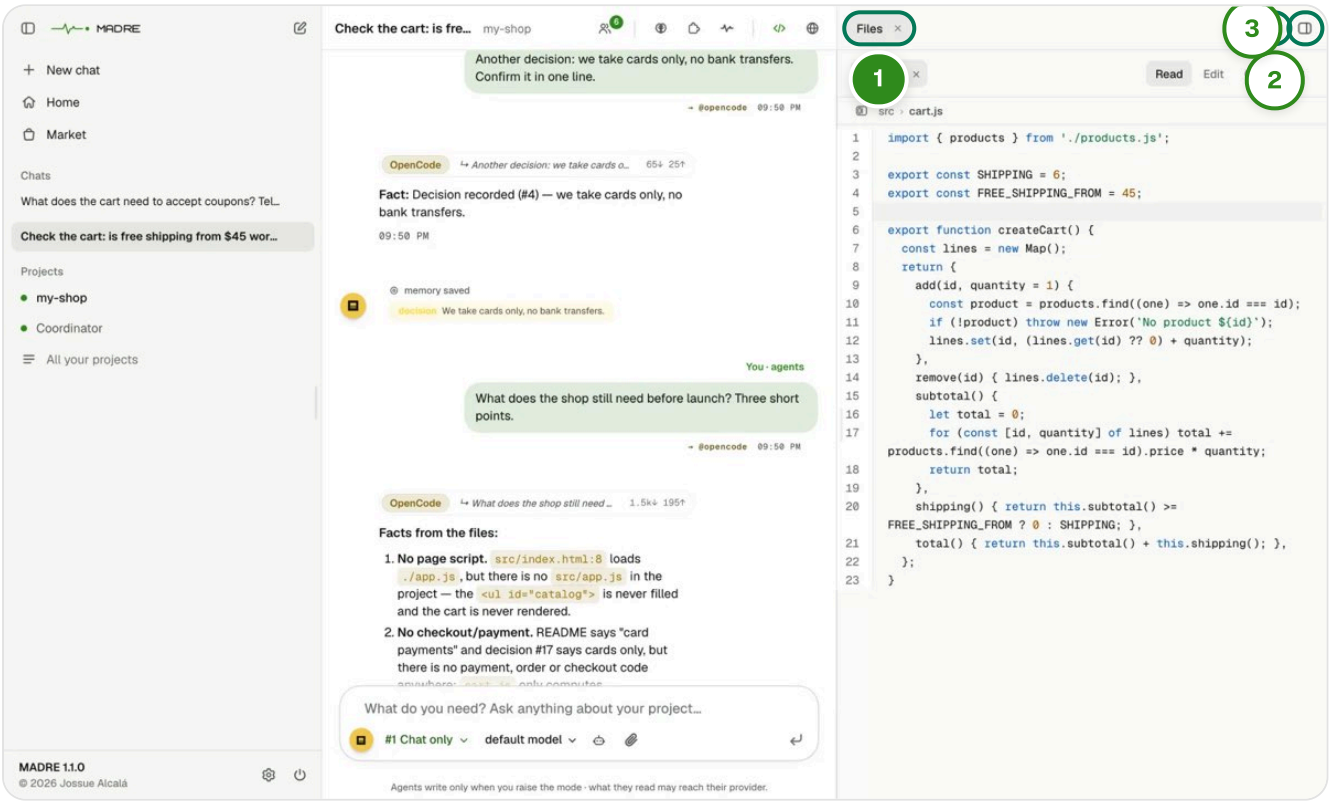

1. Each open tool is a tab. The tab's ✕ closes it.
2. The window button takes it out of the column, to put it wherever you like.
3. The last button closes the whole column.

> What an AI creates, a page, a picture, a PDF or a sheet, opens by itself beside the conversation.

You're good if you took a tool out to its own window and put it back.

Settings.

Your AIs, permissions, appearance, memory and privacy are changed here.

30 s

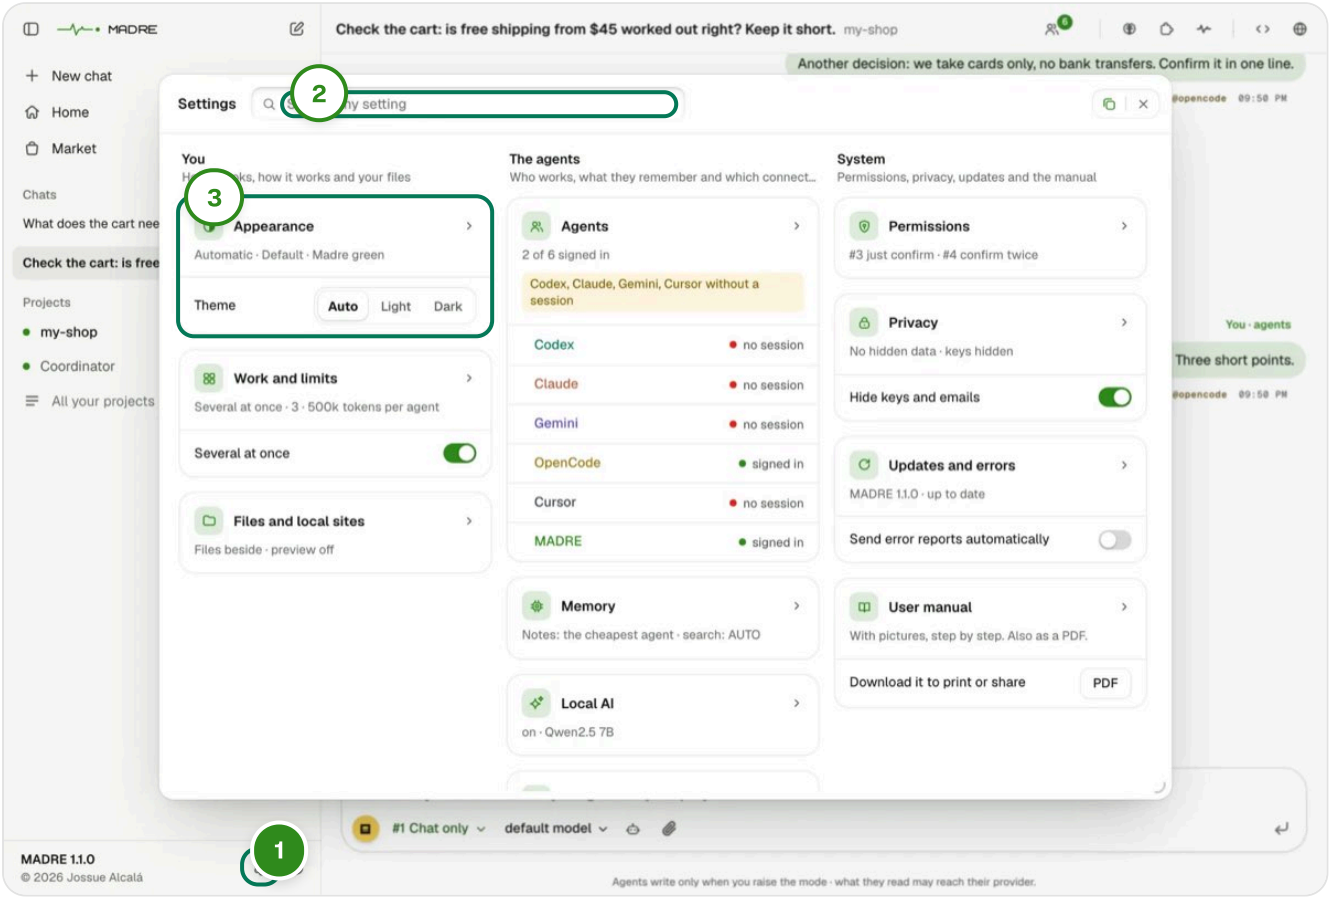

1. Open them with the gear in the left column or with ⌘ ,.
2. Find any setting by its name.
3. Each card says how it is set before you open it. What you change saves itself.

> Each project can have its own color, and from Home you can rename it without touching its folder. In **Updates and errors** you see what the anonymous diagnostics would send, and turn them off there.

You're good if you found a setting with the search box.

How to uninstall MADRE.

First MADRE deletes what it kept, and before that it tells you what goes and what stays.

30 s

Then you remove the app.

1 Open Uninstall

In Home, go to **Settings › Version** and click **Uninstall…**.

2 Read the list

The conversations, the memory, the settings, the `.pulse` folder in each project and the points to undo get deleted. Your projects and your AIs aren't touched.

3 Confirm

Click **Uninstall**. MADRE closes everything and then deletes. If a project can't be closed, it deletes nothing and tells you which one.

4 Remove the app

On a Mac, drag MADRE from Applications to the Trash. On Linux, remove it from your software center.

> What the AIs created in `.pulse/out` stays, unless you tick the box.

You're good if you know how to remove MADRE without losing your projects.

The license and what it costs.

The app has its own license of use: you install it and use it on your computers, for yourself and for your work, also when you work for your clients.

30 s

What you make with it is yours.

In a few words

What the app costs	For now the app is free. Later on it will have a paid plan. Each AI is paid separately, with your own account.
The app's license	Its code is private. The full text is in the app's license.
MADRE's code	The code on npm and GitHub stays under the Business Source License 1.1, free for individuals and teams of up to 10. Each version moves to the GNU

	AGPL v3 four years after it's published. Its license explains it.
Larger organizations	An organization of more than 10 people that uses that code, or anyone who wants to offer it as a service, needs a commercial license. Write to me through the form, under "Something else".

> This is a summary. What counts is the full text of each license.

Not sure what to do first? Pick one.

I tested all of them in MADRE, in an empty folder, with these exact messages (in Spanish). Create a folder, add it in Home and copy the messages in order.

A spaceship game by six AIs

I tested it in MADRE 1.0.3 with Cursor, @madre, Codex, Gemini, Claude and OpenCode: 6 messages, about 13 minutes.

The messages, in order

1

● @cursor **#1 Just chat**

```
@cursor #1 We're making an arcade spaceship game in a single file,
ship.html. Design the mechanics, the enemies, the power-ups that drop
and the visual style in a short list. Don't write code yet.
```

Cursor designed the mechanics, five enemies, five power-ups and the style. Before wrapping up, it asked @madre on its own whether anything had already been decided in the project's memory.

2

● @codex **#2 Can create**

```
@codex #2 Build ship.html with the design Cursor proposed: full-screen
canvas, arrows to move and space to shoot.
```

Codex read Cursor's design in the conversation and built it in one file. You didn't have to paste anything.

3

@gemini **#1 Just chat**

```
@gemini #1 Read ship.html and compare it with Cursor's design: what got left out, and which three things would make it more fun?
```

Gemini compared the game with what Cursor asked for: the boss missiles and the ring explosions were missing. It suggested screen shake, sound and a dash to dodge.

4

@claude **#3 Can edit**

```
@claude #3 Add what Gemini found missing and its three ideas to ship.html, without breaking what Codex did.
```

Claude added all five, then on its own asked Codex to test the game with 900 simulated frames. Everything passed.

5

@opencode **#3 Can edit**

```
@opencode #3 Add a credits screen to ship.html that opens with the C key and tells what each AI contributed in this conversation.
```

OpenCode wrote the credits by reading in the conversation what each one did. Nobody told it who did what.

6

@opencode **#3 Can edit**

```
@opencode #3 Almost there. If I open the credits mid-game, the PAUSE screen shows through behind the text. Make the credits cover everything behind them.
```

OpenCode saw the credits background was half transparent and made it solid.

What you'll end up with: The game file. Double-click and play: arrows to move, space to shoot, Shift to dash and C to see who did what.

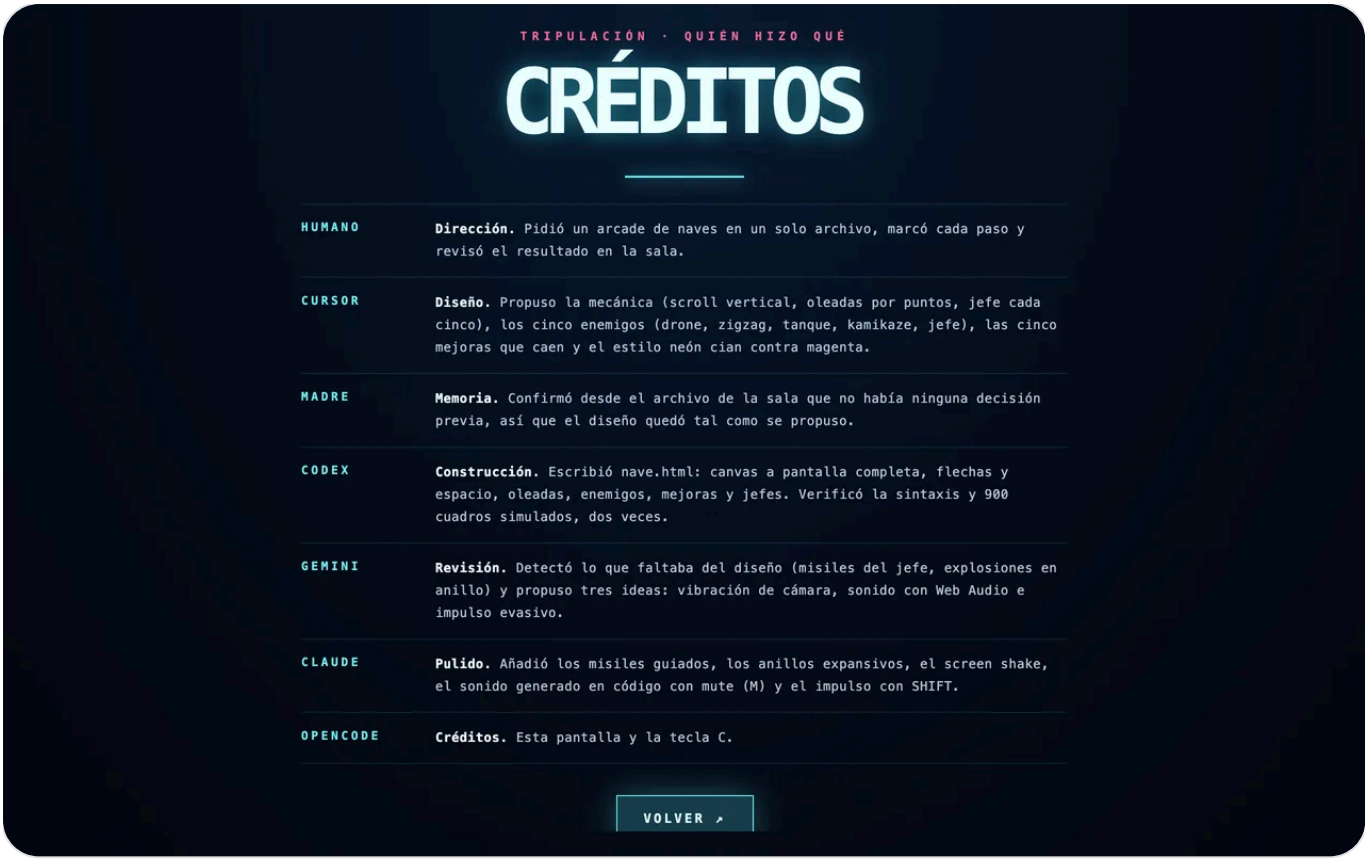

This is how it came out for me, untouched. Yours will come out different, because each AI does it its own way.

Open it and try it It's the file that came out of the test, untouched and in Spanish.

Your personal site

I tested it in MADRE 1.0.3 with Cursor, Claude, Codex, Gemini and OpenCode: 6 messages, about 9 minutes.

The messages, in order

1 @cursor **#1 Just chat**

```
@cursor #1 I want a one-page personal website, just HTML and CSS. My name is Ana Torres and I'm a product designer in Guadalajara. Suggest sections, palette, typefaces and a visual detail that makes it memorable, with no images from the internet.
```

Cursor suggested four sections, a palette of paper, ink, cobalt, jade and sun, two typefaces and an AT monogram that draws itself.

2 @claude **#2 Can create**

```
@claude #2 Build it in index.html and styles.css following Cursor's
proposal. Make up three projects, each with a cover drawn in CSS, and
use hola@example.com as the email.
```

Claude built it following Cursor's proposal to the letter, with three made-up projects and covers drawn only with CSS.

3 @codex **#1 Just chat**

```
@codex #1 Review index.html and styles.css: how does it look on a phone
and how accessible is it? Give me a short list of fixes.
```

Codex reviewed it without touching it and found things to fix on very narrow screens, in link sizes and in contrast.

4 @gemini **#3 Can edit**

```
@gemini #3 Apply the fixes Codex found and add automatic dark mode.
```

Gemini applied Codex's fixes one by one and added dark mode.

5 @opencode **#3 Can edit**

```
@opencode #3 Add a discreet line to the footer of index.html saying
what each AI did in this conversation.
```

OpenCode put in the footer what each AI did, taken from the conversation, at the minimum text size Gemini had set.

6 @claude **#3 Can edit**

```
@claude #3 Almost there. The page pulls its fonts from Google Fonts and
I want it with nothing from the internet. Swap them for system fonts
with the same character, keep the name just as strong, and update the
credits line in the footer.
```

Claude dropped Google Fonts and picked heavy system fonts so the name still hits hard. It also updated the credits.

What you'll end up with: Your site in index.html and styles.css. Double-click to see it, or turn on Preview in Modules to see it inside MADRE.

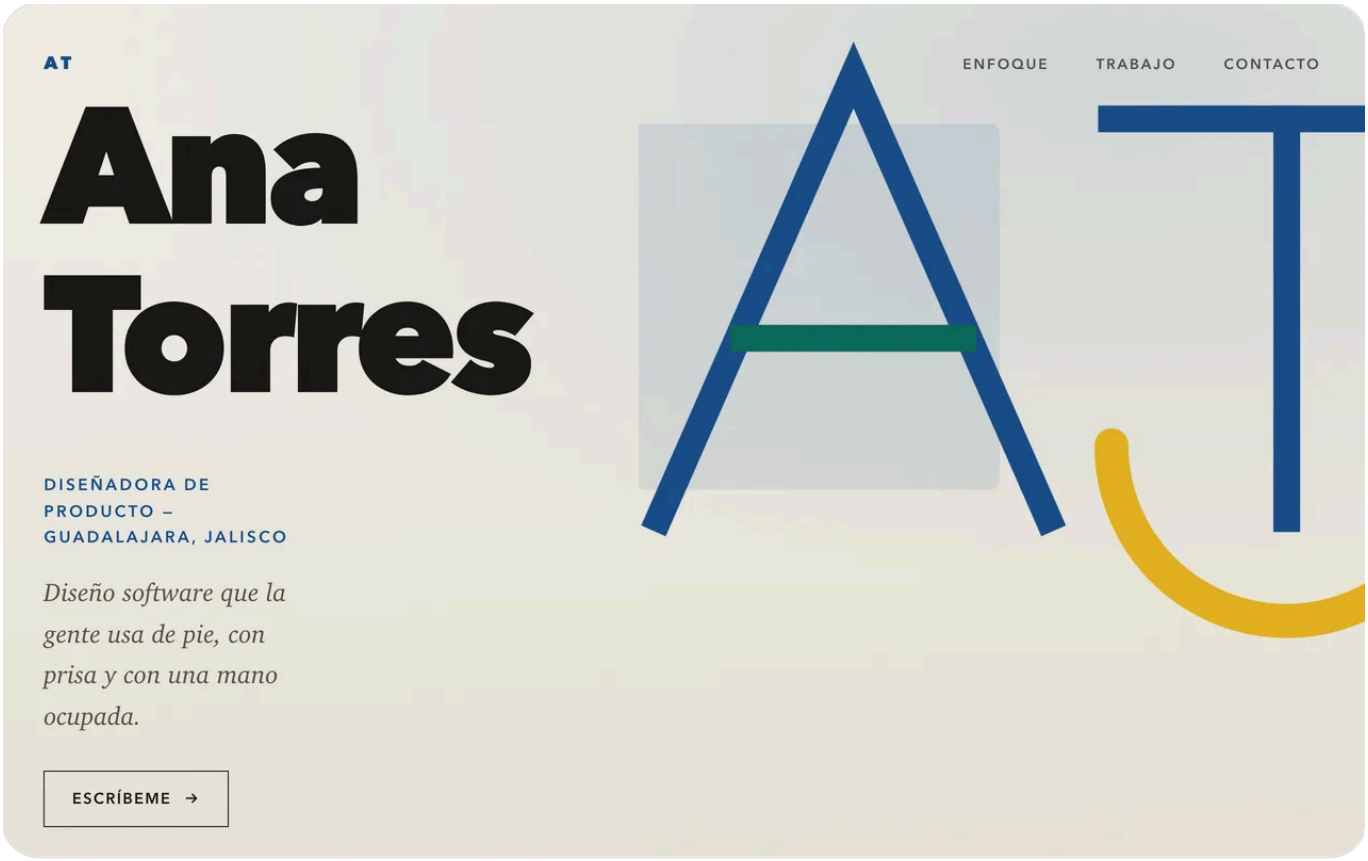

This is how it came out for me, untouched. Yours will come out different, because each AI does it its own way.

Open it and try it

It's the file that came out of the test, in Spanish. I only removed the LinkedIn link, which could lead to a real person's profile.

A Trello-style board

I tested it in MADRE 1.0.3 with Gemini, Cursor, OpenCode, Codex and Claude: 6 messages, about 8 minutes.

The messages, in order

1

@gemini **#2 Can create**

```
@gemini #2 Make a Trello-style board in a single file, board.html:
three columns (To do, Doing, Done), add cards, drag them between
columns, edit them with a double click and save everything in the
browser. Make it look modern and open with sample cards.
```

Gemini made it in one file, with sample cards, drag and drop and saving in the browser.

2

@cursor **#1 Just chat**

```
@cursor #1 Review board.html as if you were going to use it every day:
what's missing? Short list.
```

Cursor reviewed it as a daily user would: you couldn't reorder within a column, deleting didn't ask first and there was no backup or undo.

3

@opencode **#3 Can edit**

```
@opencode #3 Add the most useful of Cursor's suggestions to board.html
and tell me what you left out and why.
```

OpenCode added the most useful of Cursor's list (reordering, undo with Ctrl+Z, export and import) and explained what it left out.

4

@codex **#1 Just chat**

```
@codex #1 Review board.html: what happens with empty cards, super long
text or if I drop a card outside a column? Short list.
```

Codex reviewed it without touching it and found that an empty card could slip in on import and that text had no limit.

5

@claude **#3 Can edit**

```
@claude #3 Fix what Codex found and add a card in Done that tells what
each AI did in this conversation.
```

Claude fixed what Codex found and added a card in Done that tells what each AI did. Nobody told it who did what.

6

@gemini **#3 Can edit**

```
@gemini #3 Almost there. On a phone the buttons at the top run off the
screen, and on a big screen the columns don't fill the width. Fix it
and add your fix to the who-did-what card.
```

Gemini fixed the buttons for phones, made the columns fill the screen and added itself to the who-did-what card.

What you'll end up with: board.html, ready to use. It opens with sample cards and the who-did-what card.

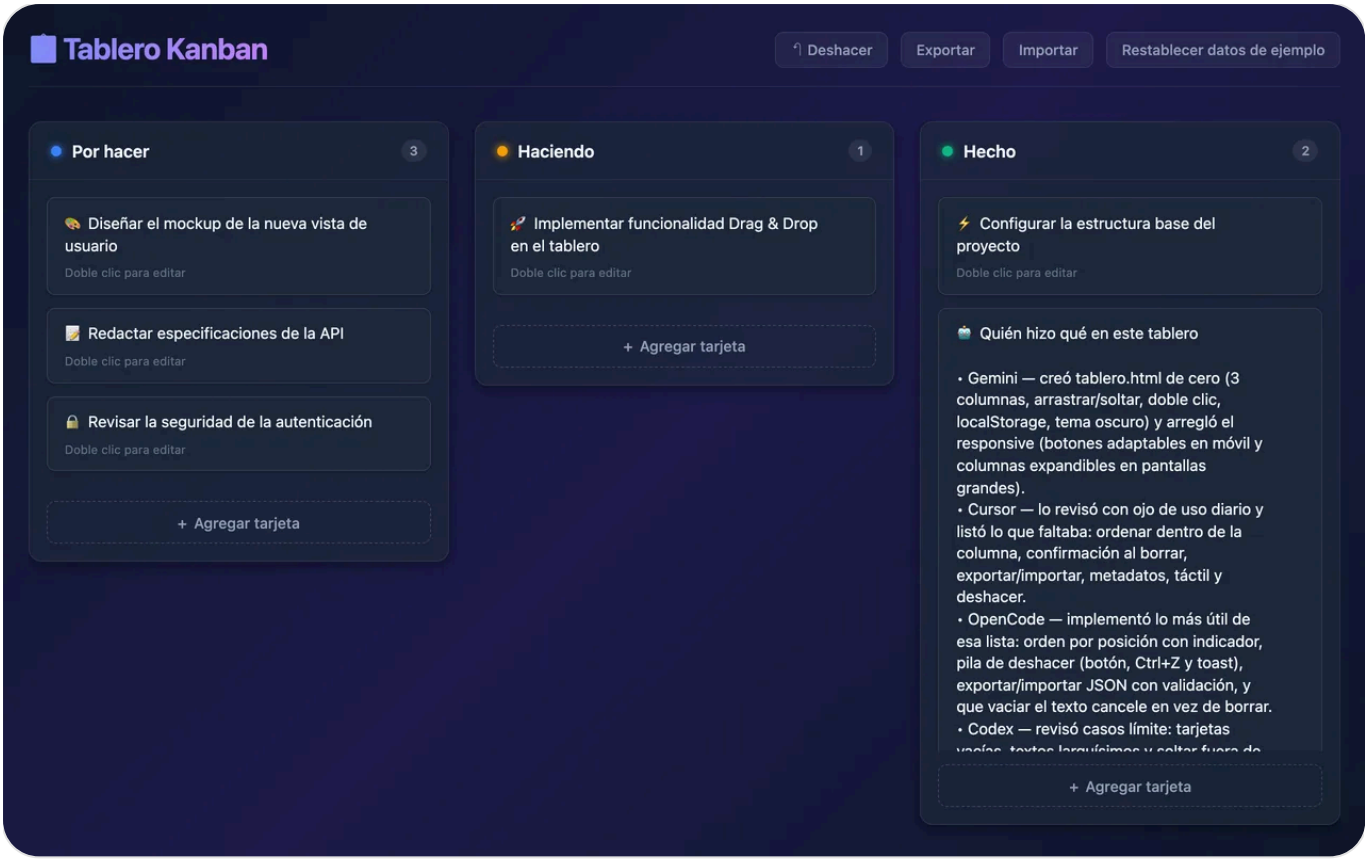

This is how it came out for me, untouched. Yours will come out different, because each AI does it its own way.

Open it and try it It's the file that came out of the test, untouched and in Spanish.

Your spending tracker

I tested it in MADRE 1.0.3 with @madre, Codex, Claude, Gemini, OpenCode and Cursor: 5 messages, about 8 minutes.

The messages, in order

1 `@madre` **Round table**

```
@madre, ask the team: I want a page to log my monthly spending and see
the total per category with a chart, no server and no libraries from
the internet, in a single file. How would you build it?
```

Codex, Claude, Gemini and OpenCode said how they'd build it, each one reading the previous ones («I agree with @codex on the essentials…»). @madre summed up where they agreed and where they didn't. **Needs your local AI**

2 `@claude` **#2 Can create**

```
@claude #2 Create spending.html with the best of the team's proposals, make it look like a modern finance app with dark mode, and tell me who you took what from.
```

Claude built it from the best of the table and tells you who it took what from, like Codex's skeleton and Gemini's SVG chart.

3

@opencode **#3 Can edit**

```
@opencode #3 Add a month of sample data to spending.html so the chart shows from the start, and a button to delete it.
```

OpenCode added 23 sample expenses, with a notice at the top and a button to remove them that leaves yours alone.

4

@cursor **#1 Just chat**

```
@cursor #1 Review spending.html: does it do what the team agreed on at the round table? What's missing?
```

Cursor checked it against what the table agreed: it covers the essentials, but the sample data wasn't saved until your first change.

5

@gemini **#3 Can edit**

```
@gemini #3 Add what Cursor found missing to spending.html, and put at the bottom who contributed what in this conversation.
```

Gemini fixed what Cursor saw and put at the bottom what each AI contributed, starting with the round table.

What you'll end up with: The spending file: you log expenses, see the month's total, the daily average and the chart per category, and back up your data.

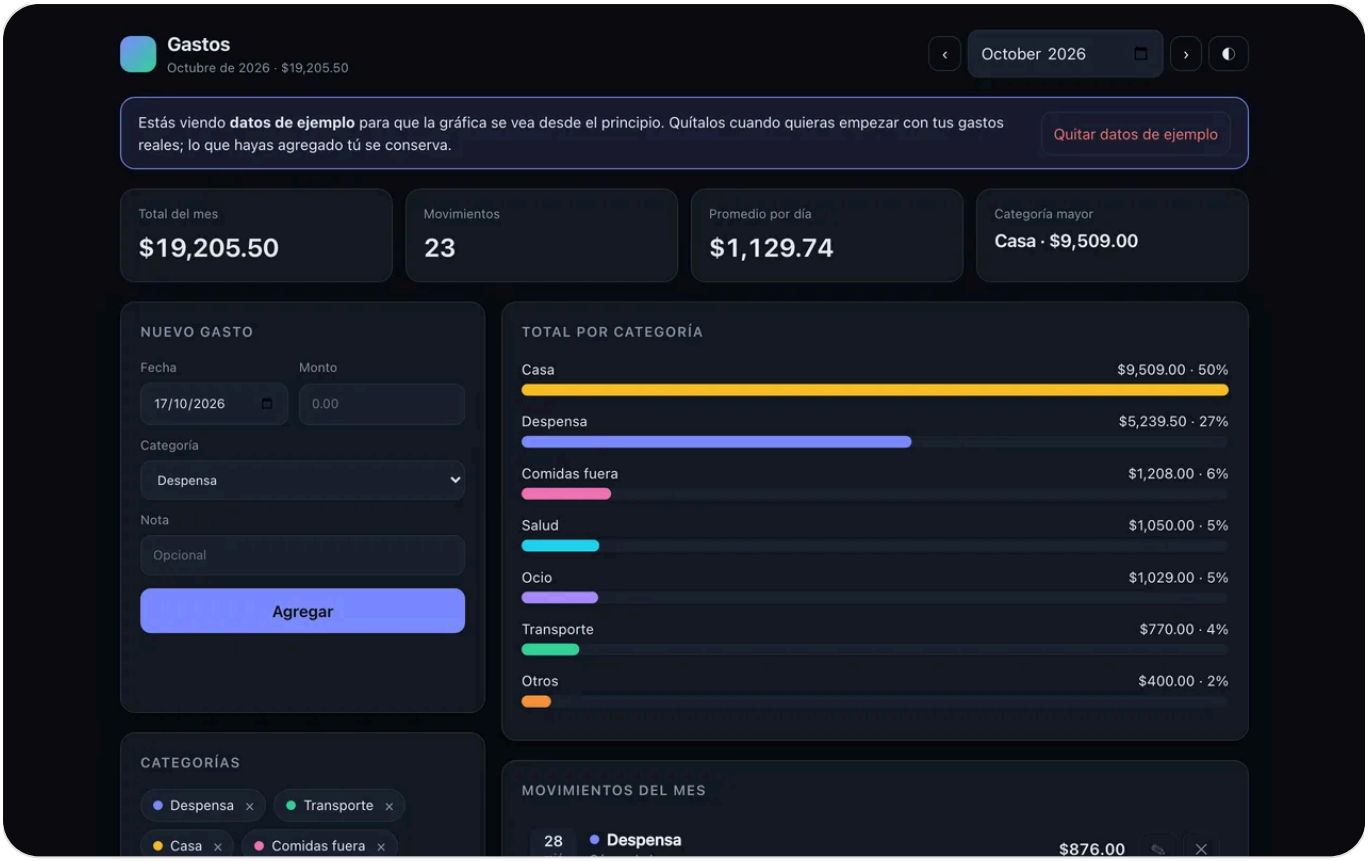

This is how it came out for me, untouched. Yours will come out different, because each AI does it its own way.

Open it and try it It's the file that came out of the test, untouched and in Spanish.

A game made as a team

You give Claude a single instruction. Claude writes a plan and MADRE follows it step by step: Cursor designs, Codex builds it, Gemini reviews it and Claude ends with a summary. You just watch.

I tested it in MADRE 1.0.3 with Claude, Cursor, Codex, Gemini and OpenCode: 3 messages, about 9 minutes.

The messages, in order

1

@claude **#2 Can create**

```
@claude #2 Make a plan with the team to build a memory game with emoji cards from scratch in memory.html: @cursor proposes the design without writing code, @codex writes the whole game following Cursor, @gemini reviews it and suggests upgrades, and you wrap up summarizing what came out and how to play.
```

Claude wrote a four-step plan and MADRE ran it on its own: Cursor designed the game, Codex wrote it to the letter, Gemini reviewed it and suggested five upgrades, and Claude wrapped up

explaining how to play.

2 @opencode #3 Can edit

```
@opencode #3 Apply Gemini's upgrades to memory.html and add a «Who made it» button that tells what each AI contributed in this conversation.
```

OpenCode applied Gemini's five upgrades and added «Who made it», with what each AI contributed, down to the emojis Cursor picked.

3 @claude #3 Can edit

```
@claude #3 Almost there. The game works, but it looks too plain. Give it personality: cards that flip in 3D with a drawn back, lively colors and a celebration when you win. Don't change how it plays and add your part to «Who made it».
```

Claude gave it purple cards with a drawn back that flip in 3D, lively colors and confetti when you win, and added itself to «Who made it».

What you'll end up with: The game file: eight emoji pairs, moves, time, restart and the «Who made it» button.

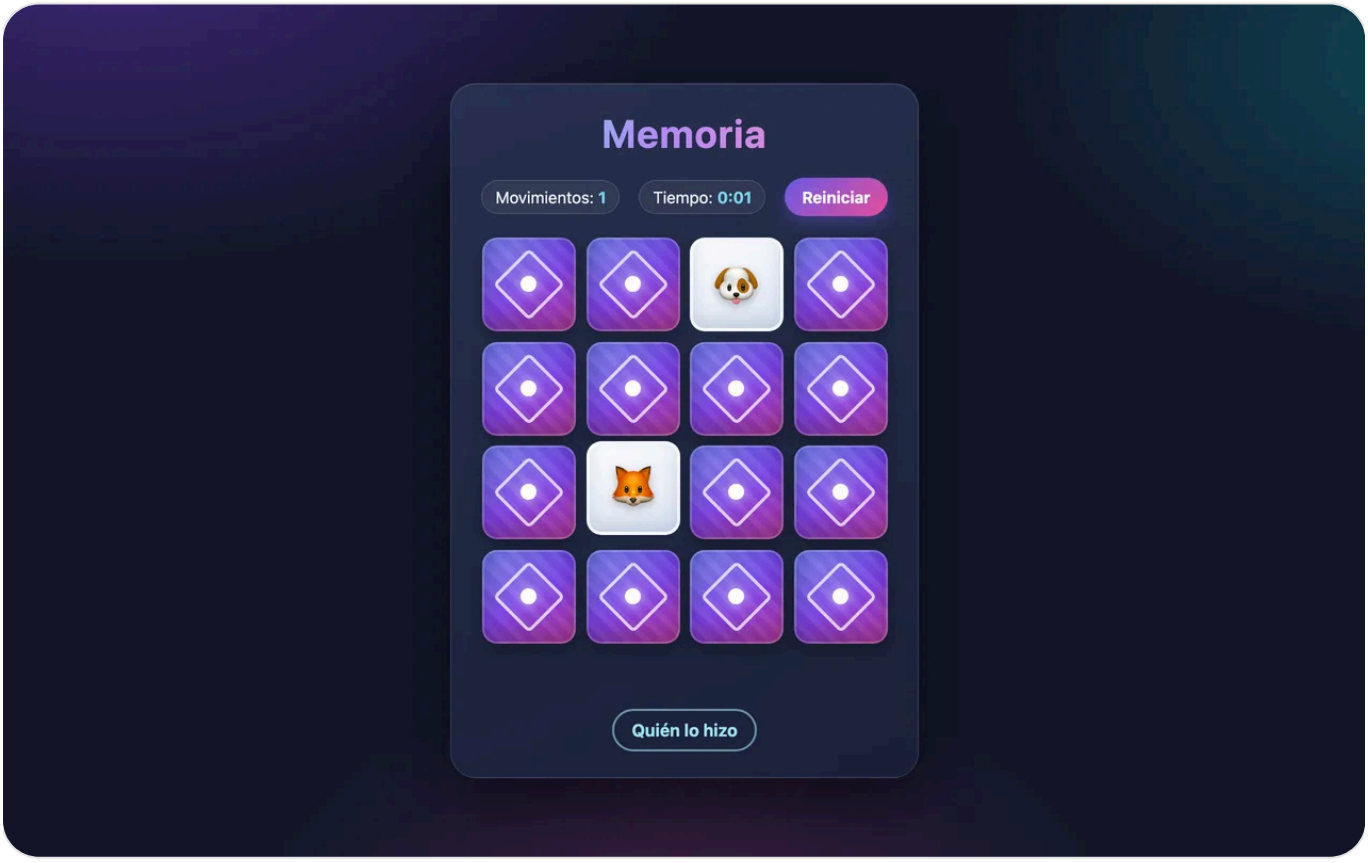

This is how it came out for me, untouched. Yours will come out different, because each AI does it its own way.

Open it and try it It's the file that came out of the test, untouched and in Spanish.

Brainstorm that ends in an app

A brainstorm between AIs that ends in an app. Each one reads what the others said before chiming in, and the one that builds it explains why it chose that idea.

I tested it in MADRE 1.0.3 with @madre, Codex, Claude, Gemini, OpenCode and Cursor: 5 messages, about 8 minutes.

The messages, in order

1

`@madre` **Round table**

```
@madre, ask the team: I have an empty folder and a free afternoon. What
small but surprising app can we make in a single HTML file, offline?
One idea each.
```

Codex pitched a constellation that sings, Claude a garden of living patterns, Gemini a musical pendulum and OpenCode a notebook that saves itself. @madre summed it up.

Needs your local AI

2

`@claude` **#2 Can create**

```
@claude #2 From the team's ideas, pick the most surprising one and
build it in app.html, with a design that feels handmade. Say why you
picked it.
```

Claude didn't pick its own idea: it picked OpenCode's, because it surprises by what it is, not by how it looks. Saving downloads the very same file with your notes inside.

3

`@cursor` **#1 Just chat**

```
@cursor #1 Try app.html as if you were the one who proposed the idea:
what's missing for it to really surprise?
```

Cursor tried it as OpenCode would and found a problem: the page explained the trick before you got to live it.

4

`@gemini` **#3 Can edit**

```
@gemini #3 Apply what Cursor said to app.html. Nothing should get cut off on a phone.
```

Gemini removed what gave the trick away, added a generation stamp and a notice when you open a copy, and fit it to phones.

5

● @codex **#3 Can edit**

```
@codex #3 Add a credits note to app.html that tells how the idea was born and what each AI did in this conversation.
```

Codex wrote how the app was born, with what each AI proposed and did, and tucked it inside «look at the HTML's guts» so it doesn't spoil the surprise.

What you'll end up with: app.html, an «autonomous notebook»: you write and doodle, and when you save you take the file with everything inside.

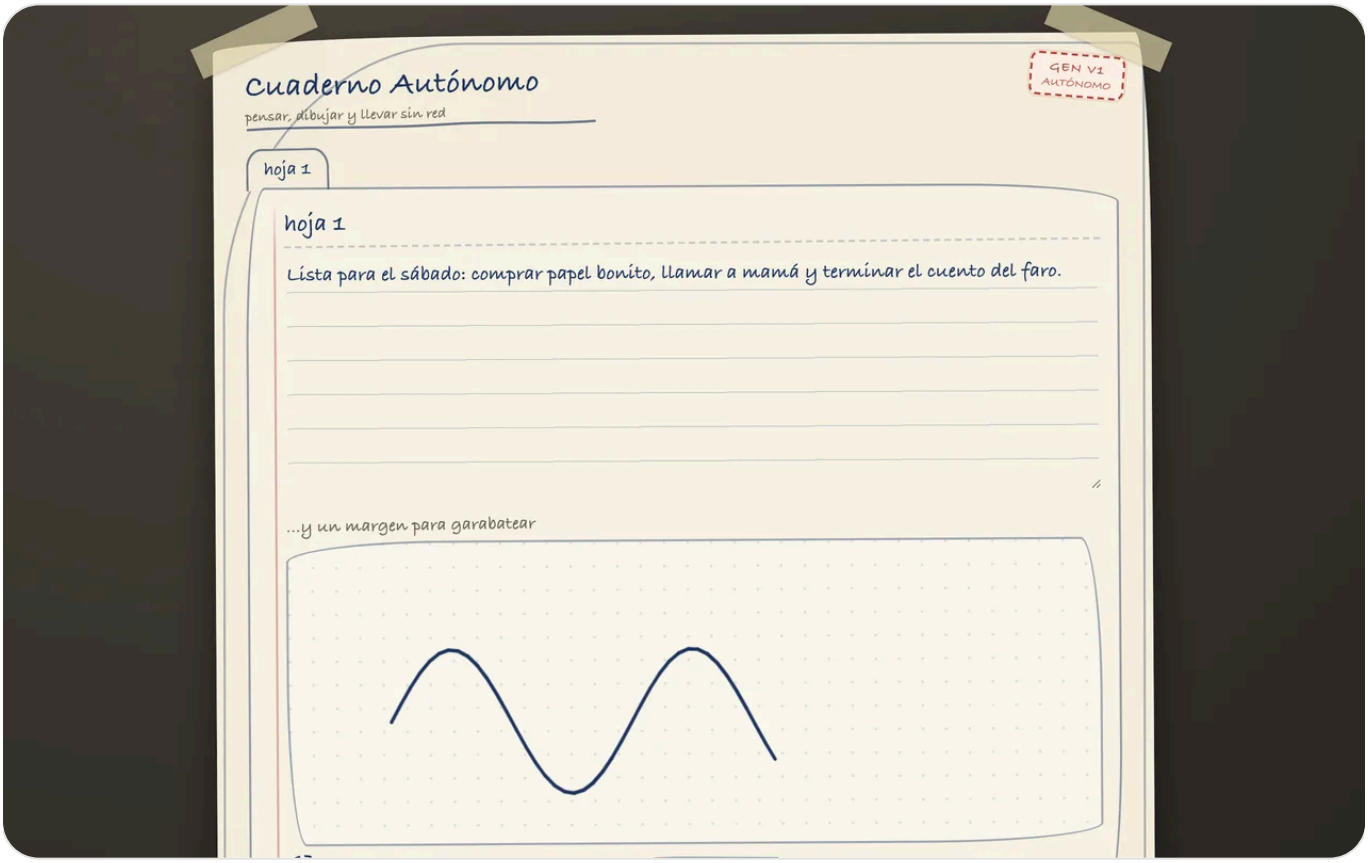

This is how it came out for me, untouched. Yours will come out different, because each AI does it its own way.

Open it and try it It's the file that came out of the test, untouched and in Spanish.

If you don't like something, ask any AI for the change or click Undo. Gemini needs a Google AI Studio key (step 4). In these tests Cursor could only read: that's why it designs and reviews, and the others

write.

All of MADRE on one page.

Keep it handy. It's also in the PDF.

In the text box

`@claude`	Who you're writing to.
`#0 … #4`	What you let it do.
`!file`	Which file you hand over.
`/`	The command menu.
`@madre, ask the team`	Round table with all of them.

Commands

`/stopall`	Stops everything.
`/coordinator`	Ask the Coordinator from your project.
`/habilidad`	Ask the AI for a new skill.
`/notes`	Ask for a note.
`/reminders`	Ask for a reminder.

Shortcuts

`Enter`	Sends the message.
`⌥ Enter`	New line (Alt + Enter on Linux).
`⌘ ,`	Opens Settings (Ctrl + , on Linux).
`⌘ P`	Finds a file by its name.
`?`	The gestures and shortcuts sheet.

The permissions

#0	Not saved
#1	Just chat
#2	Can create
#3	Can edit
#4	Full access

Ready to start.

Download MADRE, connect the AI you already use and ask it your first question.

To keep going

Make your module

Add features to MADRE with a single file.

Technical reference

All the options, commands and limits.

Changelog

What each version brings.

Known errors

Error messages and how to fix them.

Back to MADRE