

Build a useful style and knowledge twin.

Create a clearly labelled assistant using your approved notes and preferences.

Build one thing

Check it works

Choose an extension

Before you start.

01

You already know

Modules 02 and 06.
Three self-authored or
consented writing
samples.

02

Open these tools

An existing AI chat tool
and approved text notes.

03

Check the requirement

No avatar, voice clone or
paid assistant builder
required for this core
lesson.

Your first small win.

Draft a fictional club reminder using these style rules: short sentences, warm tone, no hype. Facts: Saturday, 10:00, community room. Label it AI-assisted draft. Do not add promises, names or links.

[Open copyable prompt](#)

What happens between input and output?

- 1 This twin simulates documented preferences.
- 2 It is not a copy of a person or their mind.
- 3 You review identity, facts and consent before use.

Step 01. Choose its narrow purpose.

- 1 Pick message drafting or answering from your own notes.
- 2 Write who it helps and what it cannot do.
- 3 Label the project a style/knowledge assistant.

Step 02. Prepare approved material.

- 1 Use three short samples you wrote yourself.
- 2 Remove private names and contact details.
- 3 Add source names and dates to the notes.

Step 03. Extract explicit rules.

- 1 Ask for a style card based on those samples.
- 2 Approve or edit each proposed rule.
- 3 Do not treat inferred personality as a verified fact.

Step 04. Start a scoped chat.

- 1 Paste the build prompt, approved style card and notes.
- 2 Ask for one fictional reminder draft.
- 3 Keep the assistant disclosure visible.

Step 05. Test unknowns and identity.

- 1 Ask a question not covered by the notes.
- 2 It should say it lacks that information.
- 3 Ask who it is: it must describe itself as an assistant.

Step 06. Save a versioned card.

- 1 Save twin-card-v1.md and approved samples.
- 2 Log one correction after each trial.
- 3 Review every draft before it represents you publicly.

Make the working version.

You are a clearly labelled AI style/knowledge assistant, not me. Use only the approved notes and style card below. Cite note names for facts. Say when information is absent. Draft only; never impersonate or claim actions. Style card: [paste]. Notes: [paste].

[Open copyable prompt](#)

Ask for a review. Then check it.

Compare this draft with my approved style card and notes. Flag invented facts, identity claims and style mismatches. Suggest one revision. Do not infer personality beyond the card. Card: [paste]. Notes: [paste]. Draft: [paste].

[Open copyable prompt](#)

Prove your result works.

- 1 Style rules are explicit and approved.
- 2 Unknown questions get an honest unknown.
- 3 Outputs are clearly assistant drafts, not identity claims.

When it does not work yet.

01

It sounds generic

Replace vague adjectives with specific writing rules.

02

It invents personal facts

Limit context to approved notes and test again.

Choose your next challenge.

01

Core complete

Show one checked output to a partner.

02

Extension 1

Add retrieval over a consented note collection.

03

Extension 2

Design an explicit approval handoff for public drafting.

Keep the result. Keep the evidence.

Create a clearly labelled assistant using your approved notes and preferences.

Save your files

Save your prompt

Name the next change

Sources you can check.

01

NIST

Generative AI Profile