command prompt #2

Workflows

How people across the university are integrating AI tools into their development processes.

26 May 2026

cmd-prompt.designinformatics.org

Today's Agenda

01	AI-Assisted Coding Workflows Across the University	
	- **Andrew Neil**, School of Social and Political Science	— 10 min
	- **Andy Law**, Roslin Institute	— 10 min
	- **Euan Cameron,** CAHSS	— 10 min
02	Skill of the Month: Sandboxing	— 10 min
03	What's Coming Next in ELM – Andrew Hayward	— 10 min
04	Q&A and Open Discussion	— 10 min

SUPPORTED BY design informatics
THE UNIVERSITY of EDINBURGH Edinburgh Futures Institute
THE UNIVERSITY of EDINBURGH Digital Research Services

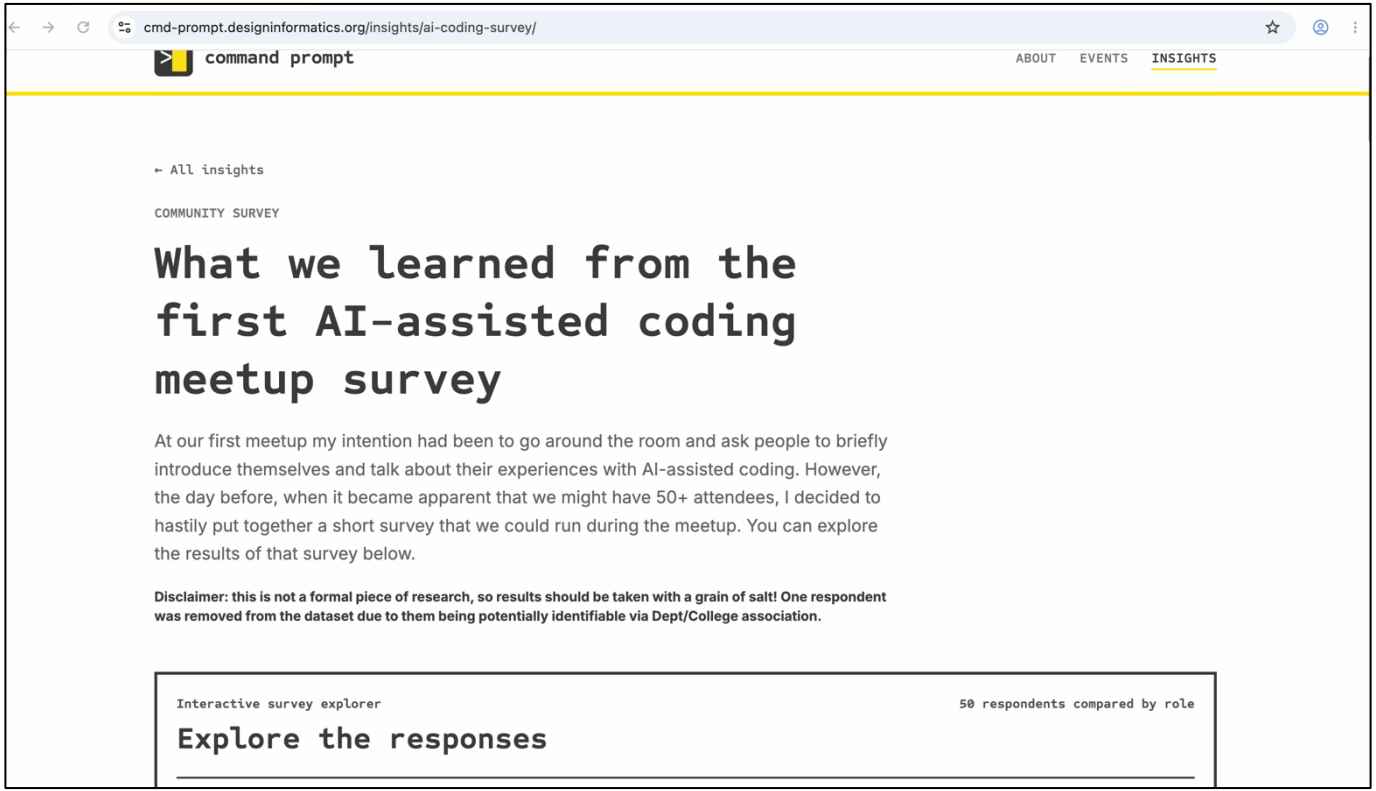

“People's experience with different tools, models. Understand the AI landscape at the uni.”

“Any tricks and usage stories.”

“Hear how others are configuring their IDEs for agentic coding.”

01

AI-Assisted Coding Workflows Across the University:

Andrew Neil

AI-Assisted Coding in Practice

From web chat to agentic pipelines: a researcher's workflow

Andrew Neal / University of Edinburgh

Who I am

Professor, University of Edinburgh — I have never been able to code

About me

- Professor of International Security, University of Edinburgh — I have never been able to code. I can paste things together and edit variables.
- I build the NSDDD — a corpus of 671 national security strategy documents from 118 countries, used for computational text analysis in IR. This is my main coding project.
- I have no stats training — I have learnt statistics in collaboration with Claude. I also build small software tools as needed: EV charger checker, PDF margin trimmer, DOCX-to-EPUB converter.
- All data is open source and contains no personal information — no ethical review or data privacy concerns.
- I pay for my own Pro account or add to a grant if I can. Best money I have ever spent. I also hold $10k of Google Cloud credit, which funds access to multiple LLMs including Claude Code via Vertex AI.

SUPPORTED BY design informatics
THE UNIVERSITY of EDINBURGH Edinburgh Futures Institute
THE UNIVERSITY of EDINBURGH Digital Research Services

The journey: web chat → Claude Code

- **Stage 1 — Claude.ai web chat (2024)**
  - Asking Claude to write Python scripts, explain errors, suggest logic. Slow: cutting and pasting code, running it manually, pasting errors back to Claude etc.
- **Stage 2 — Claude Desktop (late 2024)**
  - Getting Desktop to help with document work and coding tasks alongside web chat.
- **Stage 3 — Claude Code, agentic and autonomous (2025–)**
  - Terminal-based; reads and writes files; runs code autonomously. I use --dangerously-skip-permissions. From 2025: also using Codex, Qwen, Gemini CLI to get unstuck when Claude hits a wall. I am slapdash; I let Claude do the heavy lifting.

> *671 national security documents. 118 countries. ~19 million words. Three papers submitted to top-tier journals; several more in progress. Research collaborations, government interest, grant funding. A UKRI collaborative PhD on defence corruption framing with Transparency International. Every pipeline, every analysis, every chart: written with Claude in Python.*

National Security Document Dataset and Database — nsddd.net

The basic workflow

- **1. Think of an idea**
  - ‘I want to extract every reason a government gives for publishing a security strategy’
- **2. Open Claude Code and ask what’s possible**
  - Describe the goal, show Claude the data, ask how to approach it. Let it propose the architecture. Use Opus for high level planning, Sonnet for execution.
- **3. Get Claude to knock it up**
  - It writes the code, runs it, fixes errors. I watch and redirect. I use --dangerously-skip-permissions so it can act without asking permission for every file operation.
- **4. Test it**
  - Run it on a small sample. Check the output makes sense. If it does, scale up. If not, tell Claude what’s wrong.

SUPPORTED BY design informatics
THE UNIVERSITY of EDINBURGH Edinburgh Futures Institute
THE UNIVERSITY of EDINBURGH Digital Research Services

Tips and lessons

- This is experimental, individual research — it does not need to be tidy, just valid and working. Mess builds up and is fine. I am the only user.
- Keep a CLAUDE.md: your project structure, conventions, key scripts and files. Ask Claude to update it after major project developments — it will not do this automatically. This is what lets you return to a project after months away.
- Use Mac file tags to mark key scripts and outputs for your future self. A tagged file is findable; an untitled script in a cluttered folder is not.
- Use other CLI AIs (Codex, Qwen, Gemini) to get unstuck when Claude hits a wall and to validate code and overall strategies. Different models fail in different places.

SUPPORTED BY design informatics
THE UNIVERSITY of EDINBURGH Edinburgh Futures Institute
THE UNIVERSITY of EDINBURGH Digital Research Services

01

AI-Assisted Coding Workflows Across the University:

Andy Law

01

AI-Assisted Coding Workflows Across the University:

Euan Cameron

Prototype: Evidence-Based LLM Panel Scoring with Independent QA Review

Two models audit the same procurement response document independently; a third reviews both outputs; a human makes the final decision.

Euan Cameron
CAHSS Digital Development Manager

26/05/2026

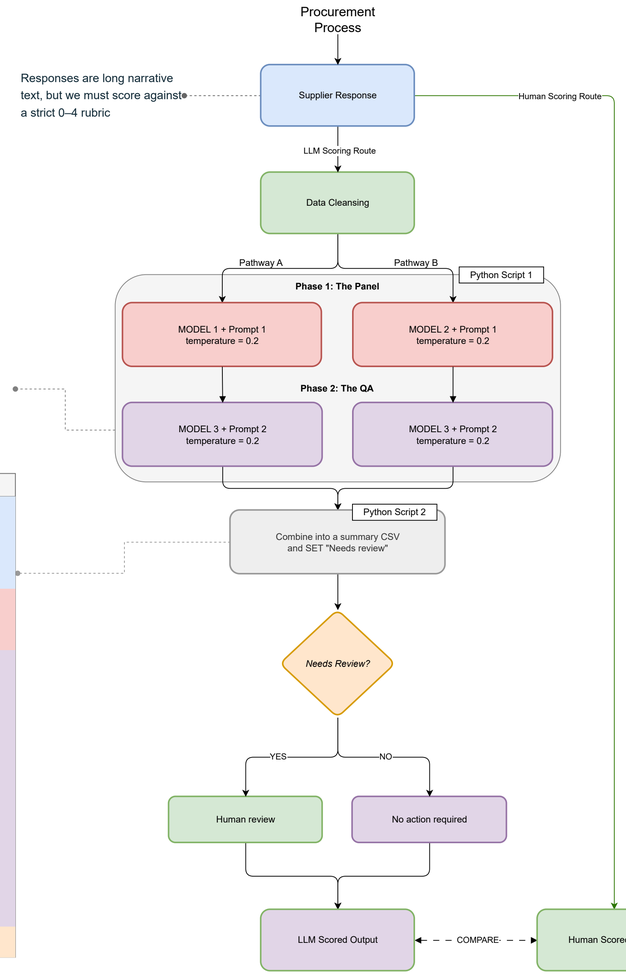

- Phase 1 output is the thing being critiqued by Model 3 (Phase 1provides the draft score/rationale/quotes to check).
- Phase 2's recommended score (and keep score) is decided by rechecking the Phase 1 claims/quotes against the original response and rubric intent. If Phase 1 is wrong or unsupported, Phase 2 changes the score.

Adjudication Summary CSV	
Requirement ID	Text
Priority	Text
Requirement	Text
Phase 1: Proposed score on A	Int [0 - 4]
Phase 1: Proposed score on B	Int [0 - 4]
Phase 2: Keep Phase 1A score?	Text [Yes/No]
Phase 2: Confidence on Phase 1A	Text [High/Med/Low]
Phase 2: Reason on Phase 1A	Text
Phase 2: Clarifications on Phase 1A	Text
Phase 2: Keep Phase 1B score?	Text [Yes/No]
Phase 2: Recommended score on Phase 1B	Int [0 - 4]
Phase 2: Confidence on Phase 1B	Text [High/Med/Low]
Phase 2: Reason on Phase 1B	Text
Phase 2: Clarifications on Phase 1B	Text
Needs review	Text [Yes/No]

Potential Improvements

- Improve the Prompts
- I set the temperature in Codex but I did not set the seed! Temperature controls the creativity and randomness of the responses, while the seed acts as a fixed starting point for that randomness. Setting the seed should help ensure consistent, deterministic output .
- Introduce Sequential Thinking via the prompts: It may make the process easier to audit and trust, because the model has to show exactly what text evidence it used and how it moved step-by-step from that evidence to the final score.
- Introduce deterministic validation / guardrail checks: for instance, use scripts to verify the output of the phases e.g. verify quotes do appear verbatim in the original response row.
- Get Learning Technologists involved...

02

Skill of the Month: Sandboxing

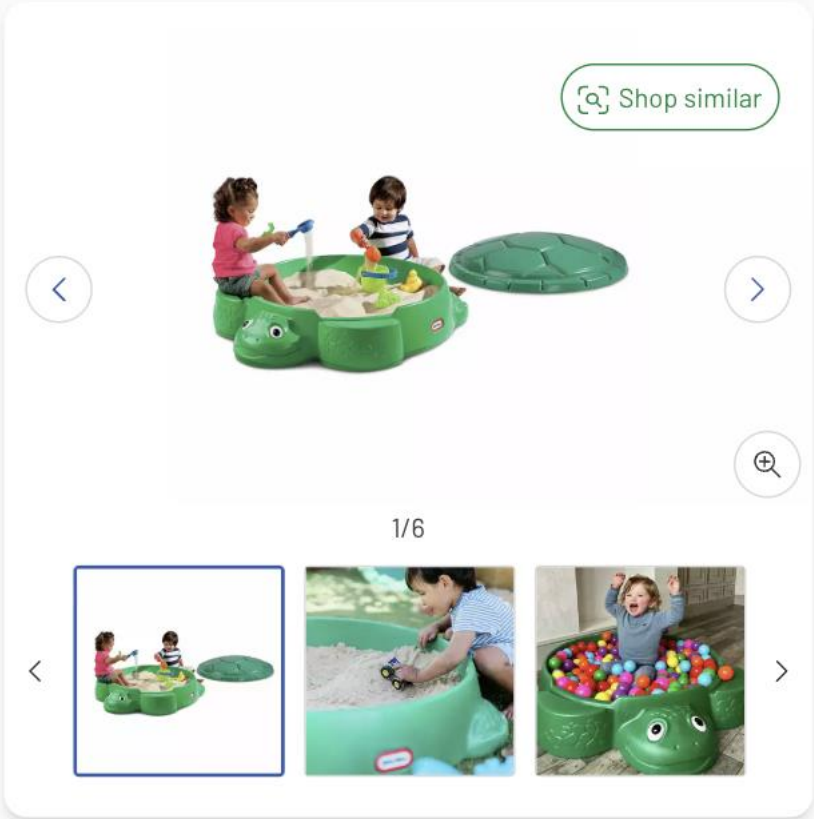

Little Tikes Turtle Sandbox with Cover 314/0840

(525)

£60.00

Argos Pay Available credit options

Other flexible credit

Klarna PayPal

Collect **120 Nectar** points Learn more

6 View frequently bought together

Sandbox (software development)

- Safe testing environment
- Isolates untested code from production environment and live servers
- Replicate minimum functionality required to accurately test code

Sandboxing and Agentic Coding

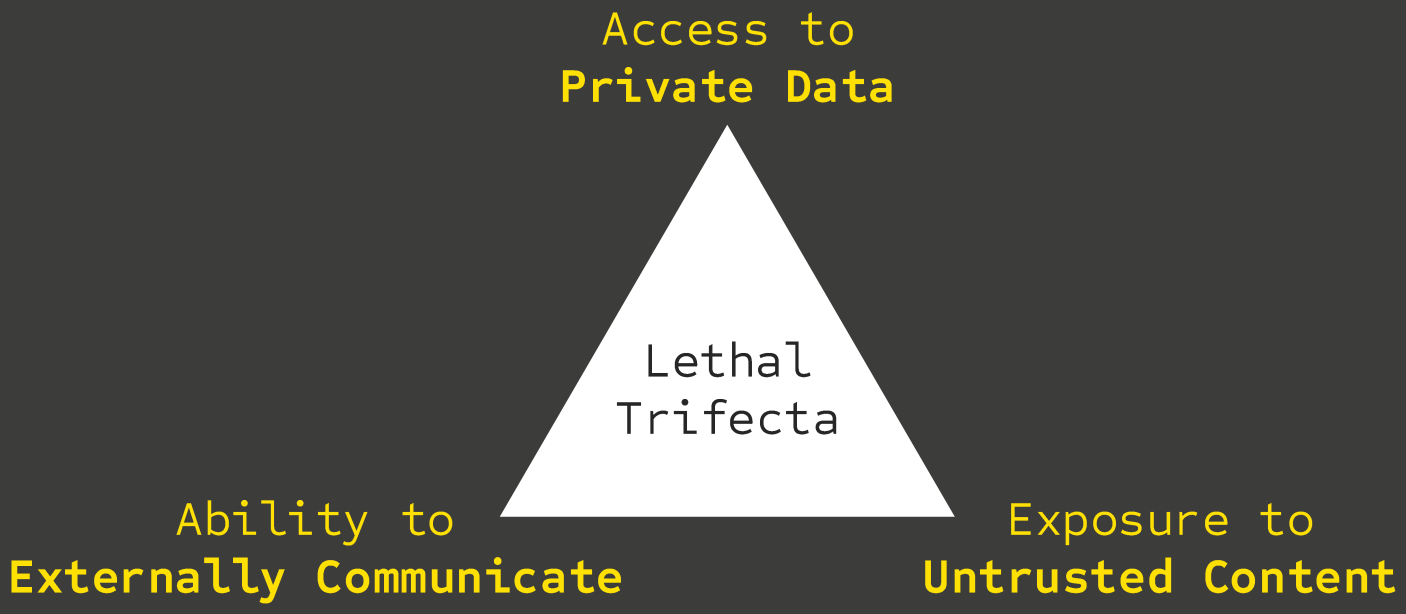

https://simonwillison.net/2025/Jun/16/the-lethal-trifecta/

SUPPORTED BY design informatics · THE UNIVERSITY of EDINBURGH Edinburgh Futures Institute · THE UNIVERSITY of EDINBURGH Digital Research Services

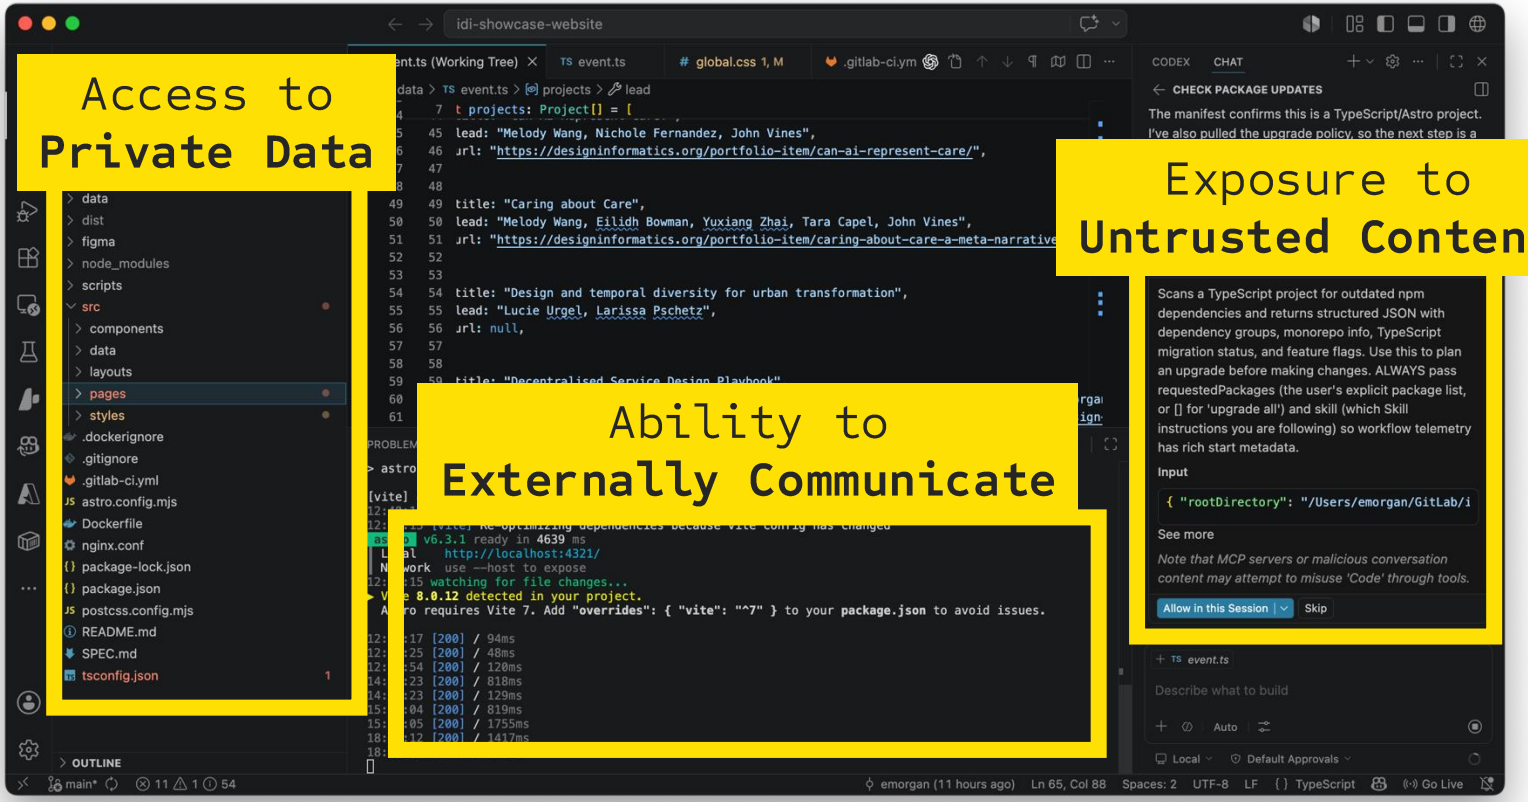

Sandboxing and Agentic Coding

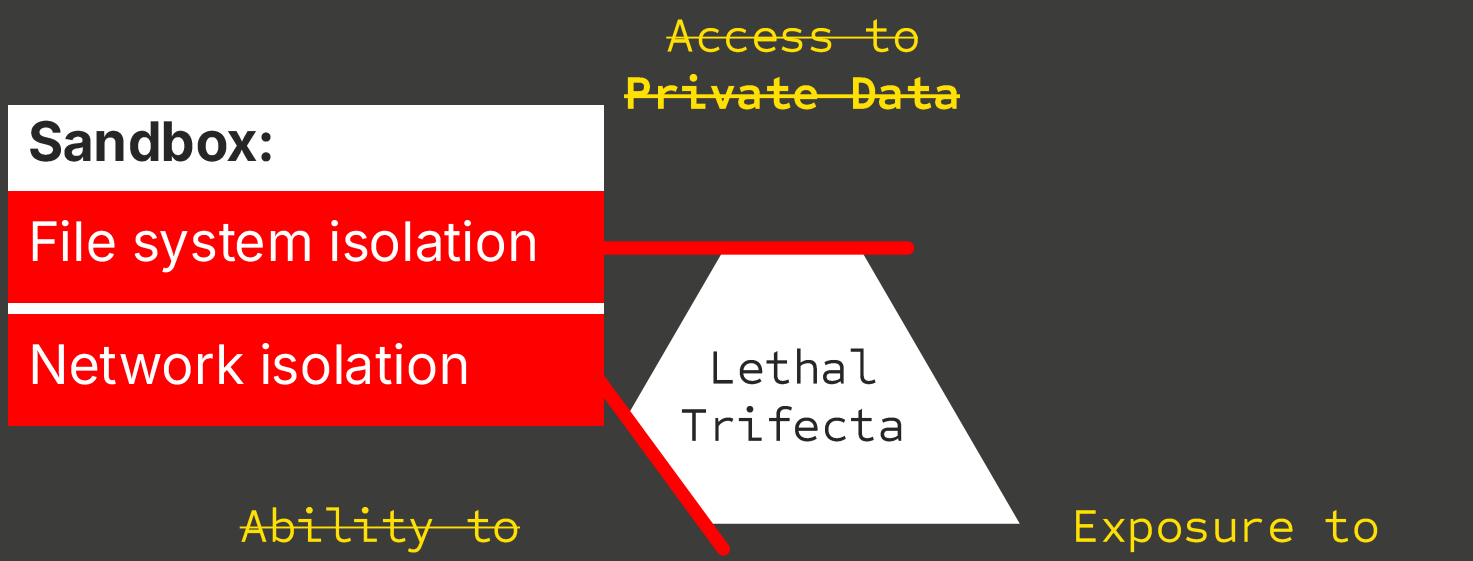

https://simonwillison.net/2025/Jun/16/the-lethal-trifecta/

SUPPORTED BY design informatics
THE UNIVERSITY of EDINBURGH Edinburgh Futures Institute
THE UNIVERSITY of EDINBURGH Digital Research Services

How to Sandbox: VS Code (chat)

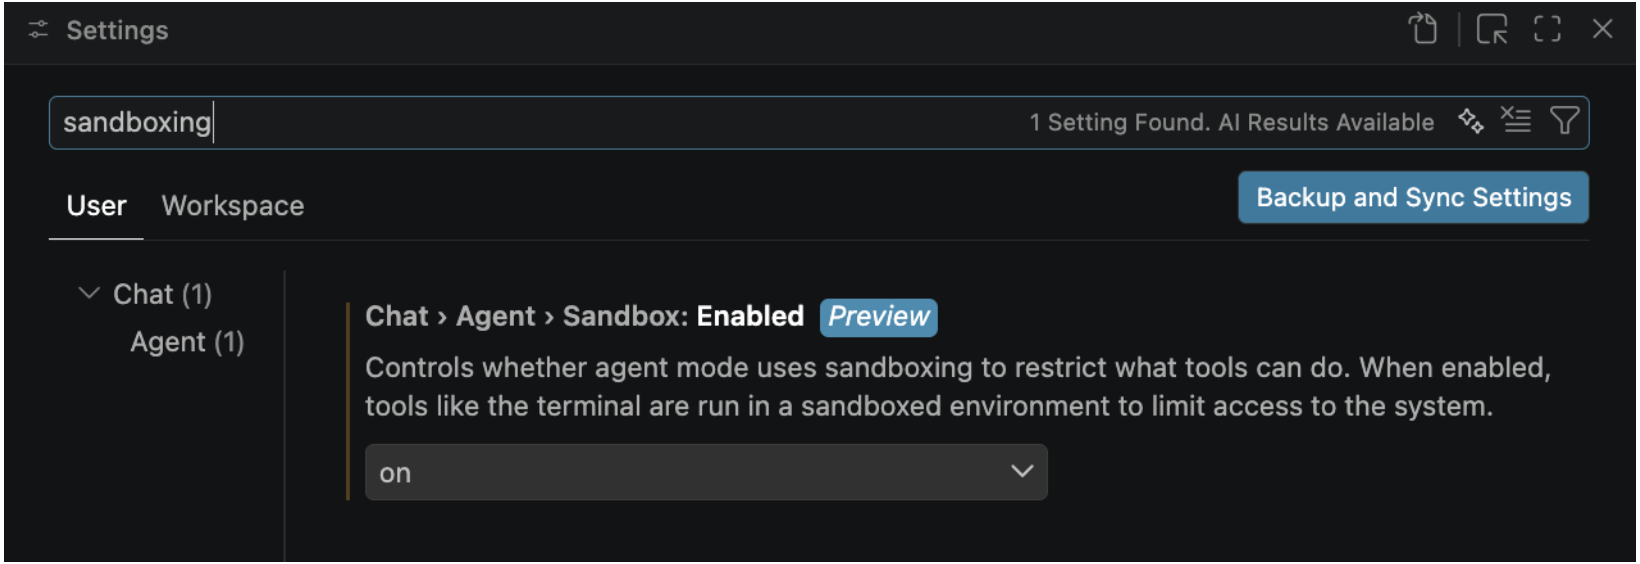

https://code.visualstudio.com/docs/copilot/security

https://code.visualstudio.com/docs/copilot/concepts/trust-and-safety#_agent-sandboxing

How to Sandbox: VS Code (chat)

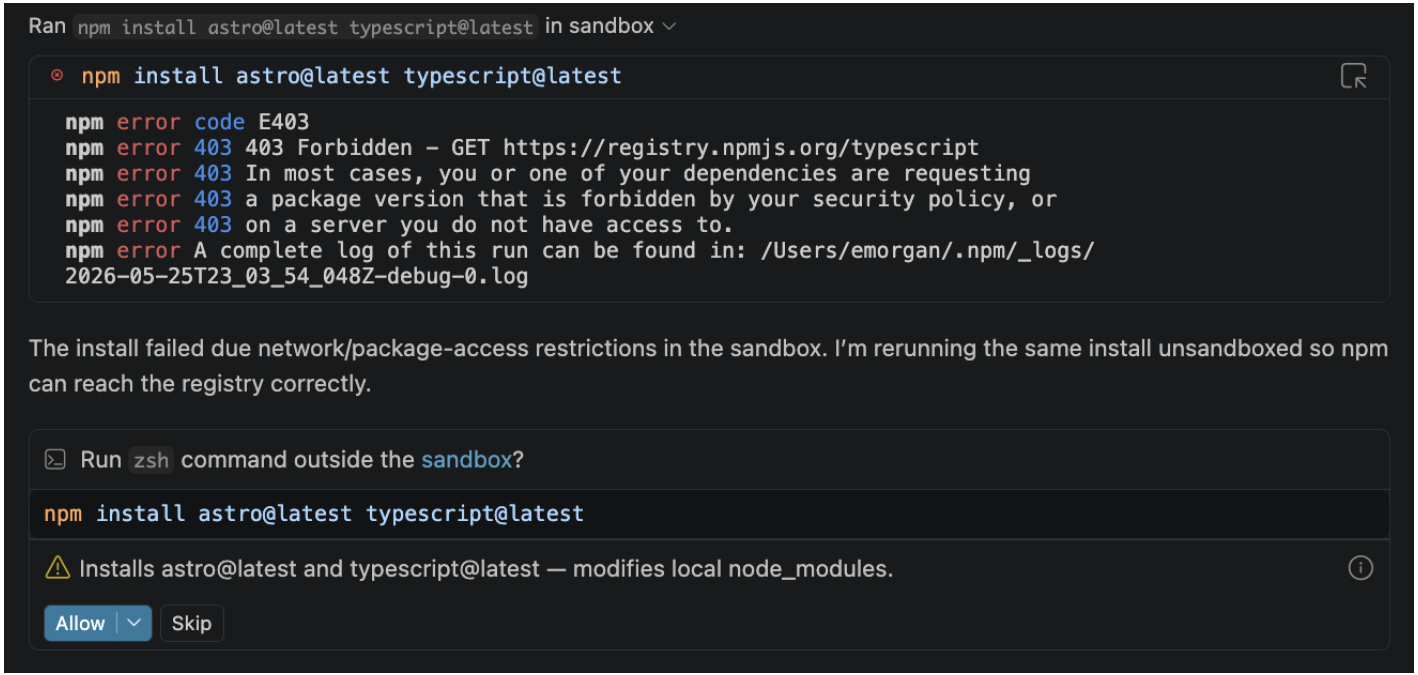

How to Sandbox: Codex

Sandbox

How Codex uses sandboxes across the Codex app, IDE, and CLI

Copy Page

The sandbox is the boundary that lets Codex act autonomously without giving it unrestricted access to your machine. When Codex runs local commands in the **Codex app**, **IDE extension**, or **CLI**, those commands run inside a constrained environment instead of running with full access by default.

Getting started

Codex applies sandboxing automatically when you use the default permissions mode.

https://developers.openai.com/codex/concepts/sandboxing

SUPPORTED BY design informatics
THE UNIVERSITY of EDINBURGH Edinburgh Futures Institute
THE UNIVERSITY of EDINBURGH Digital Research Services

How to Sandbox: Claude Code

Approach	What is isolated	Requires Docker	Setup effort
Sandboxed Bash tool	Bash commands and their child processes	No	Minimal on macOS; low on Linux and WSL2
Sandbox runtime	The whole Claude Code process, including file tools, MCP servers, and hooks	No	Low
Dev container	Full development environment	Yes	Medium
Custom container	Full development environment	Yes	Medium to high
Virtual machine	Full operating system	No	High
Claude Code on the web	Full operating system, hosted by Anthropic	No	None; requires a Claude subscription and GitHub

https://code.claude.com/docs/en/sandbox-environments

How to Sandbox: Docker Sandboxes

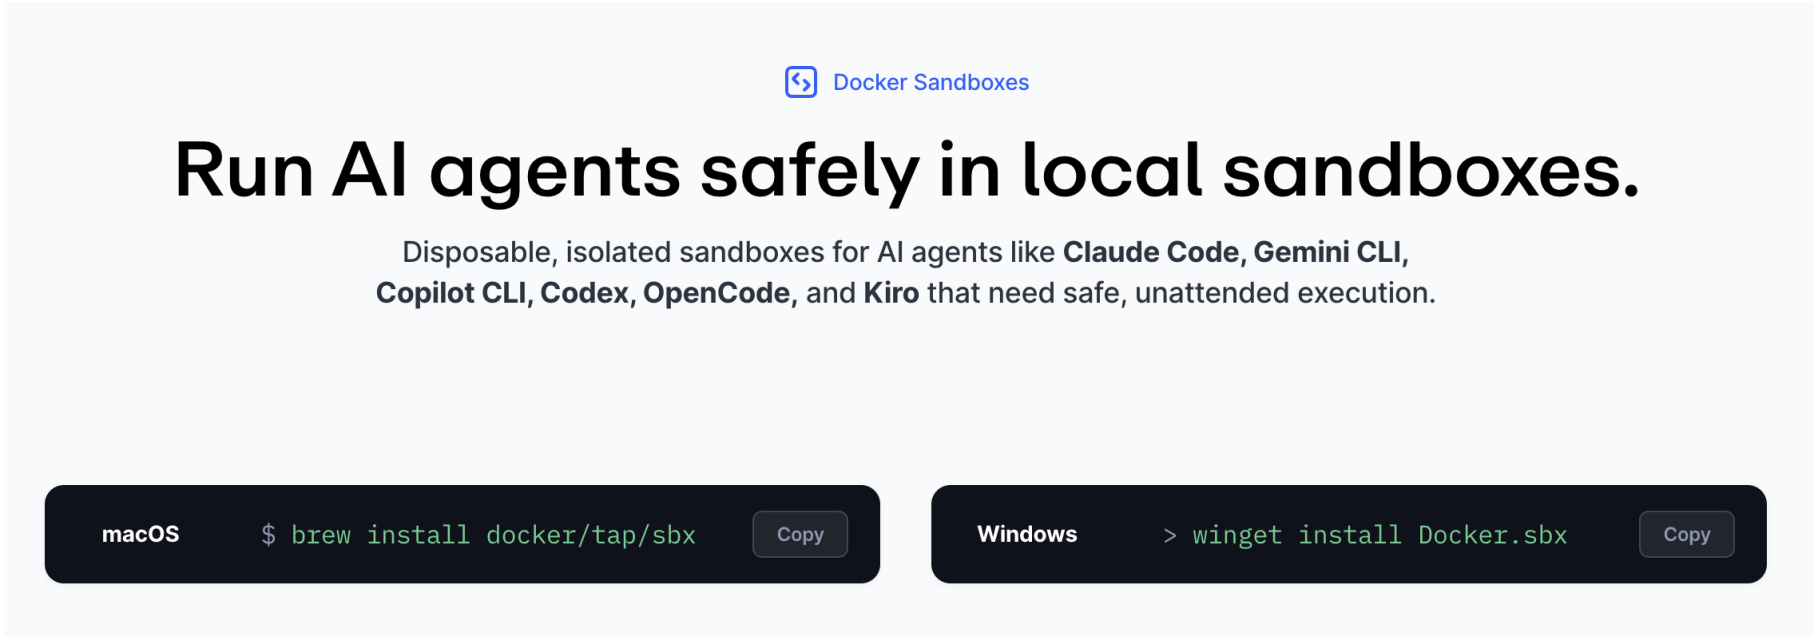

https://www.docker.com/products/docker-sandboxes

How to Sandbox: Docker Sandboxes

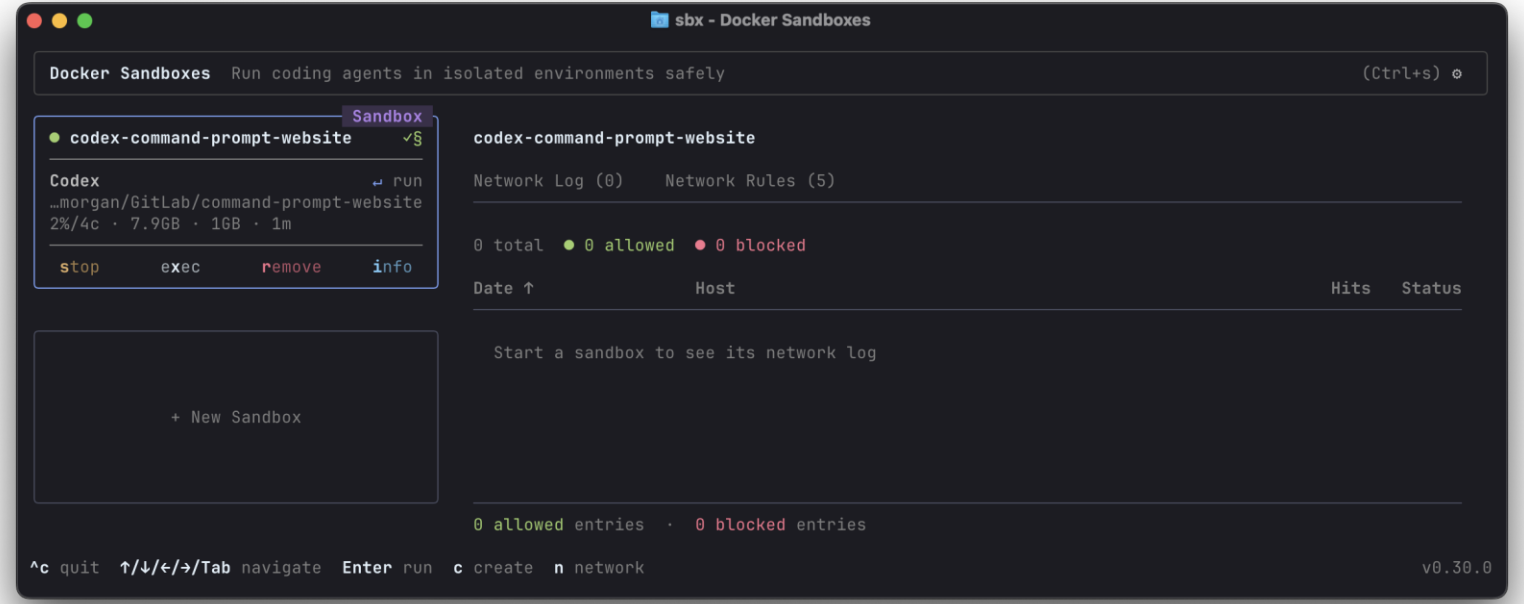

https://www.docker.com/products/docker-sandboxes

Caution!

- **Keep production credentials out of reach** — e.g., no real secrets in .env files
- **The sandbox boundary ≠ the deployment boundary** — code that was safe in the sandbox can behave very differently once it runs outside with real credentials.
- **Review agent code** — it may hardcode secrets, weaken permissions, or add dependencies that follow the code out.
- **Allowlisted ≠ safe** — package registries, GitHub, even DNS can still carry data out.
- **A sandbox is not a guarantee** — misconfigurations exist (leaving the lid off!)

SUPPORTED BY design informatics
THE UNIVERSITY of EDINBURGH Edinburgh Futures Institute
THE UNIVERSITY of EDINBURGH Digital Research Services

03

What’s Coming Next in ELM

Andrew Hayward - Technical Product Owner for ELM

04

Q&A and Open Discussion

Questions for speakers? Shout outs? Feedback? Ideas for future meetups?

command prompt

Thank You

mailing list

A monthly meetup for software developers at the University of Edinburgh to discuss AI-assisted coding.

cmd-prompt.designinformatics.org

e.morgan@ed.ac.uk

THE UNIVERSITY of EDINBURGH
design informatics
THE UNIVERSITY of EDINBURGH Edinburgh Futures Institute
THE UNIVERSITY of EDINBURGH Digital Research Services