

Using agentic AI today

Using LLMs efficiently for work in 2026.

Davide Pasca

v0.2.0

1 / 20

What is an agent?

A model trained to reach a goal with reasoning + tool calls.

Sample tools:

- Search/read/write files
- Use browser
- Generate/analyze images

Current top tools

Initially meant for coding, good for everything (to code = to build/do).

Codex, Claude Code, OpenCode, etc. The specific tool matters less.

Modern LLMs are trained for "long-horizon" work. They do 90% of the work, "harnesses" are less important. You can make your own and see. (It gets harder with security, large context management, MCP support, etc.).

Limitations

"Humans use 20% of their intellect" ... **myth.**

"Humans use 1% of their LLMs" ... **probably true.**

Bottleneck #1

Remembering to use an AI agent.

Reasons:

- Going straight to Google out of habit
- Bad or slow UI/UX of the tool
- Demoralized by poor model intelligence

Solutions:

- Make an effort to break habits (slow and painful at first)
- Seek better models and interfaces

Bottleneck #2

Long-horizon tasks can take 20+ minutes, so we multitask.

Breaks attention span, we forget to check the reply, etc.

Solutions:

- Ask to test the outcome and give you a TL;DR. It will take more time, but less interaction.
- Use top models: more reliability = less human testing.
- Use a tool with overview of all projects (Codex App, LCR, etc.)

Bottleneck #3

Time spent verifying the outcome.

Solutions:

- Automated testing is *essential*. Automate as much as possible.
 - AI writes the tests anyway!
- Use a frontier model for peace of mind
 - Commit now, verify later (assuming not mission-critical)

Bottleneck #4

AI can't see your screen.

Solutions:

- Set up an easy screenshot-to-clipboard flow
 - Code agent "skill" to take screenshots
 - Ask the agent to browse with its own browser.
- Major life improvement!

With code, remember to clean up!

LLMs are not trained to offer to refactor code, but it's fundamental for long-term maintenance. You must ask yourself, every now and then.

Remember to ask it to note important future context in README/AGENTS.md.

Question

A genie grants you 3 wishes, what should you wish for?

"Build Era", make your own?

Use tools to build better tools (but don't get sucked in).

Reduce the chance of getting locked into tools that are not really yours:

- Software can change for the worse, or get discontinued
- Licensing can become bad

See what your pain points are. Where's the friction? Ask a coding tool to make a utility/interface for you.

Ideal tool/workflow for me

- Single readable window with overview of all projects
- Integration with Codex, Claude Code, OpenCode (no vendor lock-in)
- TODO system to add requests without waiting
- Ability to create quick tasks not tied to a project
- Better browser integration
- Worktrees and automatic merge conflict resolution
- Detection of orphaned ports and processes

Little Control Room (LCR)

Little Control Room filter command boss /errors log Scratch task deleted CPU 49% (590% raw)

[Active 124] Archived 34 a cycle

	ATTN	ASSESS	LAST	AGENT	TODO	RUN	PROJECT	SUMMARY	
!	u	180	working	11:56	CX	05:23	1 PID! adb...	okmain	Compacting conversation history... (05:23)
u	87	done	11:56	CX		3		▼ LittleControlRoom	Fixed workflow warnings, committed, pushed, and released v0.2.8 with n...
>	50	working	11:56	CX	01:40			L spike/orphaned-ports	jectrun.Snapshot objects for runtime display. /ports can use the same sna
!	43	followup	5 days	CX				L spike/youtube-prom...	Advises building image-to-terminal prototype as next artifact.
	25	done	Yesterday	CX				L todo/let-s-add-a-l...	Added used-MCP list to AI engineer sidebars; all tests pass.
	25	done	Yesterday	CX				L codex-process-reuse	Implemented MCP server for runtime process management, removed prompt ...
	0	done	1 week	CX				L lcagent/web-search...	Reviewed ds4_web.c as fallback search backend; no changes made.
!	48	done	3 days	CX			pnpm	talk_gamedev_lessons	Initial Slidev deck committed and PDF exported.
!	u	47	done	Yesterday	CX			talk_agentic_ai.2026	Fixed stale cache and rebuilt deck; user instructed to hard-refresh.
u	40	followup	1 week	CX				dpasca-sdl2-template	Provided platform-support plan with static linking; ready to implement...
u	34	waiting	4 days	OC				tmap skate_oc_gml52	Rebuilt for native GLUT; waiting for user to test ./skate.
	30	done	4 days	CX				romaexe_intros	Commits and video moves done; YouTube titles updated and published.
	30	done	6 days	CX			pnpm	roguellm	Defaulted landing page to Create World tab when no worlds available; v...
	28	followup	1 month	CX				oik-aso	Fractal Strike is low-volume; recommended small marketing push for 10-...
!	23	done	Yesterday	CX			\$1=="m_E...	► proj_leaf	Fixed swarm drift feedback loop after edge bounce; compilation passed....
	22	waiting	3 weeks	CX				newtype_invoice	Email draft ready for user to review and send.
	22	missing	1 month	CX				lcroom-goal-rpc	Turn aborted, awaiting user input.
!	22	waiting	1 month	CX				crypto	Committed; proposed live canary rehearsal with XRPUSDOT, awaiting user ...
!	15	done	3 months	CX				agent_ntpr	Pivoted repo away from X browser automation; optional follow-ups offered.

Summary: Plan to add /ports command agreed; implementation pending.

Path: /Users/davide/dev/repos/LittleControlRoom--orphaned-ports

Assessment: followup Activity: active

Last activity: 2026-06-25T11:56:29+09:00 Codex

Repo: clean (spike/orphaned-ports)

Merge back: spike/orphaned-ports -> master

Merge status: nothing to merge

Attention: 50

Worktrees: root + 5 linked, 3 active, 1 dirty

Worktree actions

- w or /wt lanes

- M or /wt merge

- x or /wt remove

TODO

Control Room - LittleControlRoom--orphaned-ports



Standby. Use /run or /run-edit to wake the room.

Enter Codex lanes /wt ops merge remove help quit PIDs 1 port1 cost \$0.123

13 / 20

LCR: Codex session

Codex | okmain Alt+O links Alt+L blocks Alt+S sidebar

YOLO MODE

Session

Model gpt-5.5 / xhigh
Context 41% of 258k
Tokens i249.7M c94% r327k o883k
Commands /model

Browser

Page localhost:9878

Diff Summary

7 files changed, 520 insertions(+), ... M4 ?3
master_mobnext
M Tools/render_sprites/render.mjs
M Tools/render_sprites/render.engine.mjs
M Tools/render_sprites/tune.html
+4 more
Enter opens full diff

Active Processes

port tune.mjs pid 25575 :9878
port Google pid 15277 :63896
port! pid 37109 adb :5037

Summary

The Phong smoke render passed. The PBR one collided...

◆ lcr_runtime/list_processes

The replacement tuner is healthy on http://localhost:9878/. I'm loading the Playwright tool bindings now for a live smoke test.

I'm going to use the existing managed browser context. The smoke check will look for: HDR controls in Lighting, no HDR env controls in Material, the readout saying HDR, and no console errors after rendering a PBR probe plus a Phong vehicle.

◆ playwright/browser_navigate

The live UI already shows the new shape: HDR is under Lighting, Material has no env controls, and the readout says HDR Wide Street 01 even for the F50/Phong path. I'm going to switch to a PBR probe and render once to catch runtime errors.

◆ playwright/browser_evaluate

The PBR probe loaded with HDR Wide Street 01, the old env controls are gone from Material, and the CLI preview is now much cleaner. I'm checking the console and then I'll run a headless render smoke.

◆ playwright/browser_console_messages

No browser console errors. I'm doing two headless smoke renders now: one Standard/PBR model and one Phong-style vehicle, to make sure the shared HDR lighting path works in both material modes.

Command (8 lines hidden; Alt+L previews) -> \$ /bin/zsh -lc 'node render.mjs --model models/manual_fordgt20...

Command (27 lines hidden; Alt+L previews) -> \$ /bin/zsh -lc 'node render.mjs --model models/superhive_tran...

[command failed, exit 1]

The Phong smoke render passed. The PBR one collided with the other smoke render because both tried to bind the same temporary 9876 file server at once; that was my parallelism being a little too eager here. I'm rerunning the PBR smoke by itself.

Command (8 lines hidden; Alt+L previews) -> \$ /bin/zsh -lc 'node render.mjs --model models/superhive_trans...

Working 13:40 Enter steer ctrl+c interrupt Alt+Up hide Alt+Enter newline ctrl+v image Alt+C copy menu

> |

LCR: TODOs while agents work

TODO LittleControlRoom

3 open, 14 total

[] We have /cpu to check and query about ... - working Codex

[] Let's look into this way to do web sea... - resume Codex

[] How do we make this app more popular ?... - resume Codex

[x] Kimi K2.7 Codejust came out. If we have access to it if I mean anyone with an old A...

[x] We just got the error: Error LCAgent failed: lcagent executable not found; set LCR0...

[x] Now sure if it's mean tot be like this, but I see an LCAgent session giving almost ...

[x] Task title may become something like `[2 lines pasted]` if a paste was summarried...

[x] Let's see if we can improve the formatting and the markdown tables. For example, in...

[x] When I'm in an AI agent embedded window and I scrull up with 2 fingers ont he touch...

[x] We currently have /terminal inside an AI engineer session. We shoudl have it access...

[x] In the AI engineer session sidebar we currently have the model critic both under se...

[x] We need to unify hwo the sidebar of the AI engineer works. All sidebars by default s...

[x] Get into romaexe_intro with a new codex session, changed model version for the LLM,...

[x] The model selection for LCAgent isn't quite there. When in an LCAgent session I use...

a add e edit space done d delete p purge done c copy
Enter start Esc close

Edit TODO LittleControlRoom

Kimi K2.7 Codejust came out. If we have access to it if I mean anyone with an old API key can have access to it. And also let's see if we internally some kind of default for Kimi to K2.6. If so, let's switch the default to K2.7 code. Here is the reference URL
<https://platform.kimi.ai/docs/pricing/chat-k27-code>

enter newline ctrl+s save Esc cancel

Queue requests without interrupting the running agent.

Turn a thought into a task, then launch it later.

Note about worktrees

WT = temp workspace (git clone + branch)

Beware of git-ignored files: where's my `.env`?

Merge can bring conflicts. In most cases, you can solve them with `/resolve`, which uses an AI agent to guess and fix automatically.

Voice input

MacParakeet is very powerful. Records your voice and pastes text in the current app.

Main issue: it annoys people around you 8)

Better browser integration

Multiple tasks fighting over the same Playwright Chrome instance.

Solution:

- Per-project PW instance. Hidden by default (doesn't steal focus)

Moving forward

Things are evolving fast. There's no perfect workflow. Try other tools, build your own small tools.

Thank you

Questions and discussion