

# Own your AI agent: running open source agents on your terms

Davide Eynard @ DataFest Tbilisi, 29 Nov 2025

 [@mala@fosstodon.org](https://mala@fosstodon.org)

 [mozilla-ai](https://mozilla-ai.com)

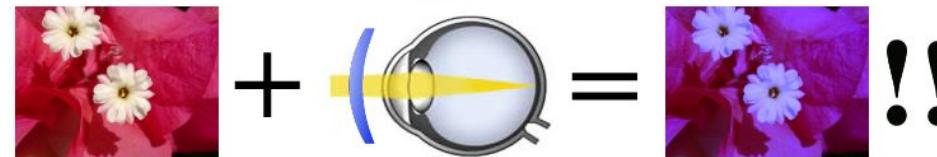
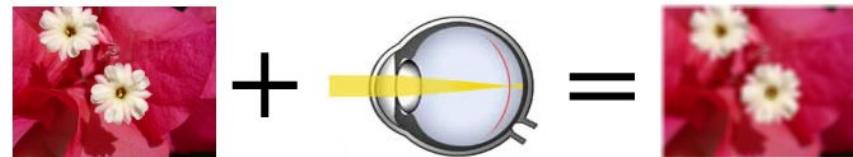


# Intro

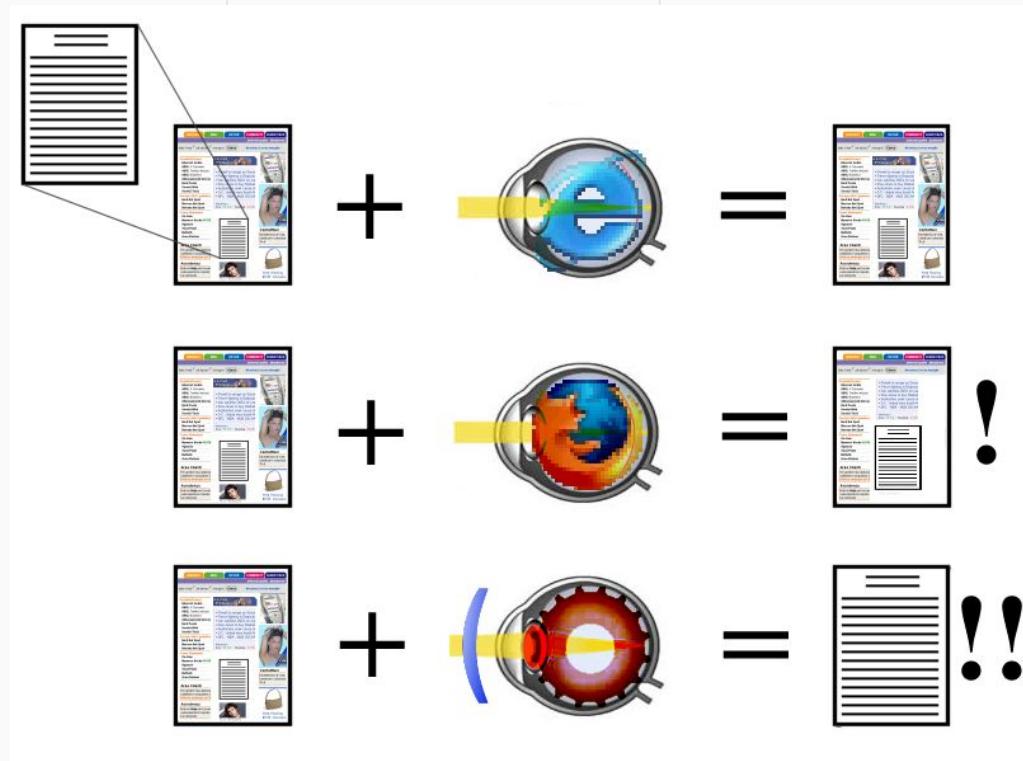
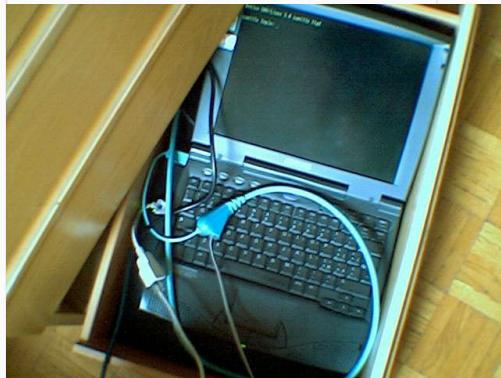
# Ada and Zangemann (Matthias Kirschner, Sandra Brandstätter, David Revoy)



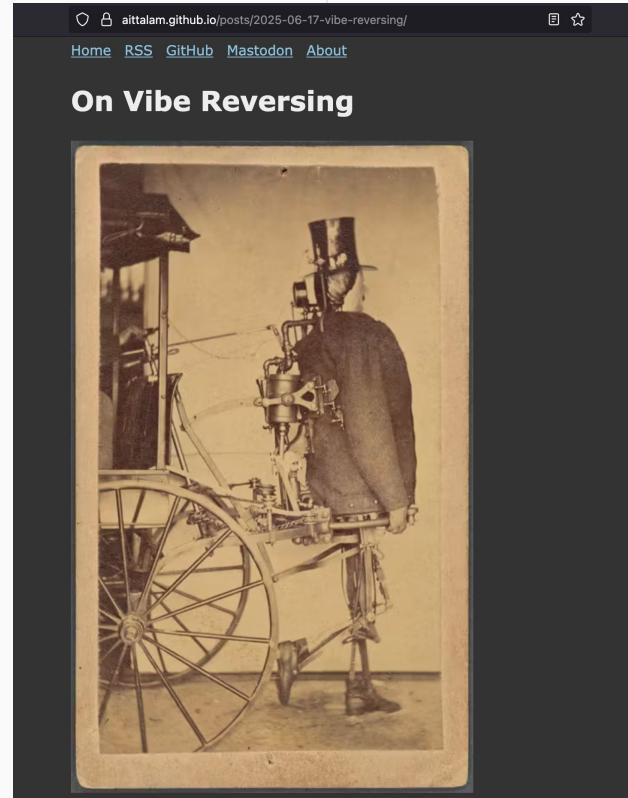
# Powerbrowsing (2005)



# Powerbrowsing (2005)



# Powerbrowsing 3.0 (2025)



# On Vibe Reversing

I'd like to see how well you do in helping me study how a web application works. The website I called viaggiautreno and holds timetables for all trains running in Italy. You can choose a train station and see all the trains passing by that station and if you choose one you can see all the stops they make along the way, where and when. I would like to understand which calls are made to the backend and the format in which the data is stored / its meaning etc

I'd be happy to help you analyze how the Viaggiautreno web application works! This is a great reverse engineering exercise to understand how transportation APIs and web applications function.

Let me search for information about Viaggiautreno's API structure and backend calls.

⊕ viaggiautreno API backend calls reverse engineering	10 results ^
⊕ Reverse engineering dell'app di Trenitalia – Jacopo Jannone – blog	jacopo.io
⊕ Reverse engineering Trenitalia's mobile application – Jacopo Jannone – blog	jacopo.io
⊕ How to Reverse Engineer APIs: The Benefits and Tools	dreamfactory.com
⊕ A Tutorial for Reverse Engineering Your Software's Private API: Hacking Your Couch   Top...	toptal.com
⊕ json – how to reverse engineer an http API call using REST console - Stack Overflow	stackoverflow.com
⊕ viaggiautreno · GitHub Topics · GitHub	github.com
⊕ json – Reverse-engineer a private RESTful API? – Stack Overflow	stackoverflow.com

Filter URLs		Mozilla Firefox Network Monitor										⚙
Status	Method	Domain	File	Initiator	Type	Transferred	Size	WS	Other	Disable Cache	No Throttling	⚙
200	GET	www.viaggiautreno.it	1749247200000	rest-steam109 (xhr)	json	3.7 kB	24.49 kB					
200	GET	www.viaggiautreno.it	1749247200000	rest-steam109 (xhr)	json	2.66 kB	27.66 kB					
200	GET	www.viaggiautreno.it	andamento_frecce_sx_on.png	img	png	cached	1.94 kB					
200	GET	www.viaggiautreno.it	andamento_frecce_sx_off.png	img	png	cached	352 B					
200	GET	www.viaggiautreno.it	andamento_frecce_dx_on.png	img	png	cached	1.94 kB					
200	GET	www.viaggiautreno.it	andamento_frecce_dx_off.png	img	png	cached	352 B					
200	GET	www.viaggiautreno.it	Sat Jun 07 2025 22:17:30 GMT+0100 (British Summer Time)	rest-steam109 (xhr)	json	2.89 kB	22.36 kB					
200	GET	www.viaggiautreno.it	501325	rest-steam109 (xhr)	json	765 B	1 kB					

These are the hits I get from the network tab, which one would you suggest me to check?

Excellent! You've captured some real API calls. Based on what I can see in your Network tab, I'd recommend checking these requests in this order:

## 1. First Priority: The JSON Responses (Most Important)

Check the ones with **Type: json** - these contain the actual train data:

- **The ones with filenames like 1749247200000** - These are likely the main API endpoints returning train data. The numbers look like Unix timestamps, which is exactly what we expect from Viaggiautreno APIs.

## 2. How to Examine Them:

1. **Click on one of the JSON requests** (the ones with large transfer sizes like 24.49 kB or 27.66 kB)
2. **Look at the Headers tab first** to see the complete URL - this will show you which specific API endpoint was called

# Vibe Reversing: the hard truth

It worked! BUT...

- I had already done it in the past
  - ... would have I been able to make sure it worked otherwise?
- Artifacts are on the platform
  - ... you can download them, but you need to ask to create them
- Very good learning references
  - ... that you'll skip, if you are only interested in a quick answer
- I wrote zero lines of code
  - ... and I learned nothing

# Vibe Reversing: the hard truth

The strongest critique would probably be: **"You're a technical person telling non-technical people to make their lives harder to solve problems that mostly exist in your head."**

# Tinkering with AI

Using AI for tinkering => directly meddling with it!

- Is it possible?
- Will I have enough compute?
- Is it useful at all?

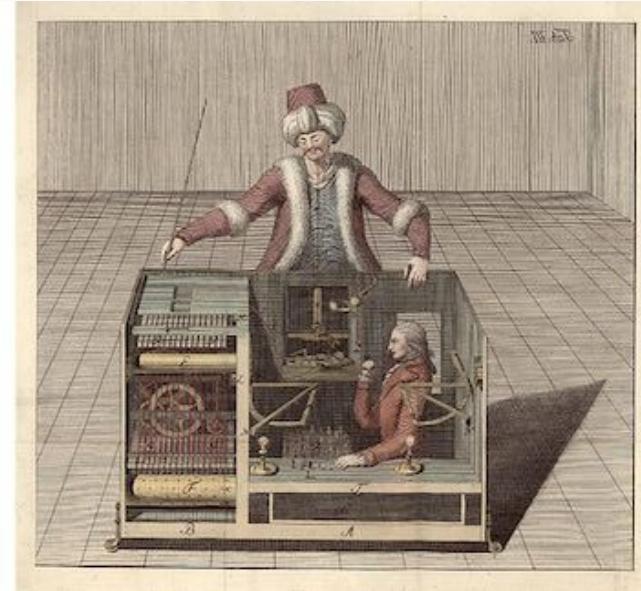
# Tinkering with AI

Using AI for tinkering => directly meddling with it!

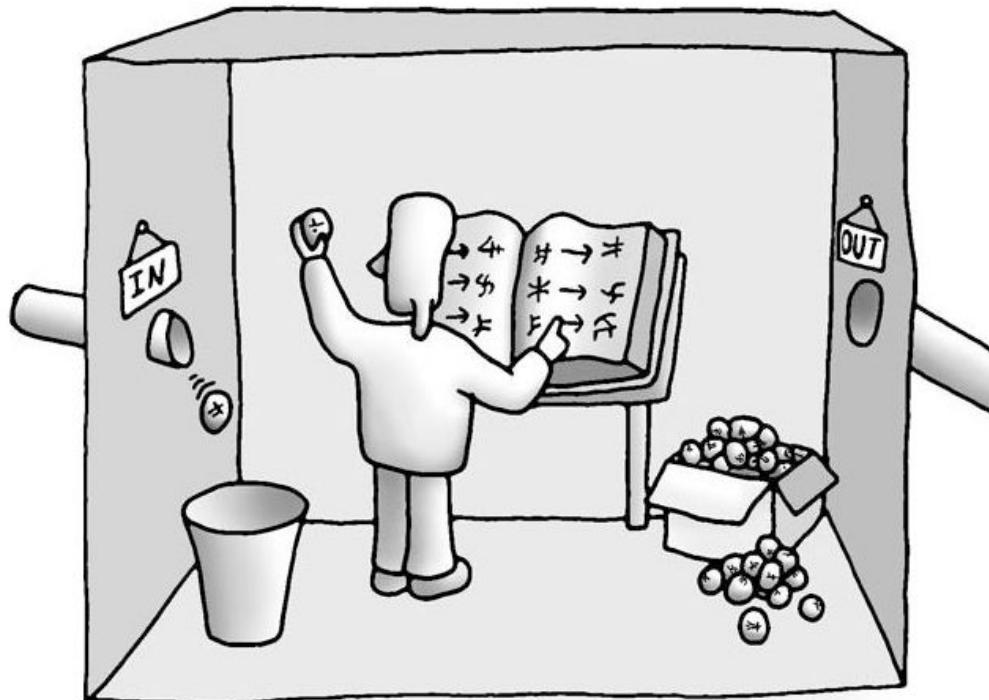
- Is it possible?
- Will I have enough compute?
- Is it useful at all?



# The Mechanical Turk



# The Chinese Room Argument



John Searle, 1980

# The WOPR Room

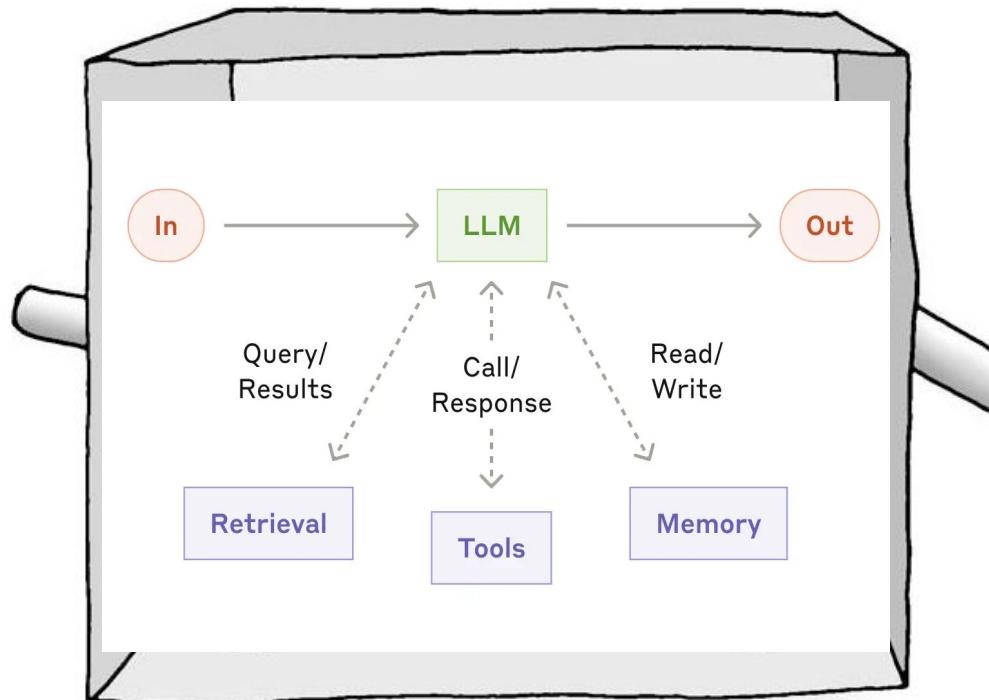


War Games (John Badham, 1983)

# The open source AI models room



# The commercial AI services room



# Build

# How does an agent work?

```
● ● ●

history = [instructions, user_prompt]

while True:

    response = CALL_LLM(history)

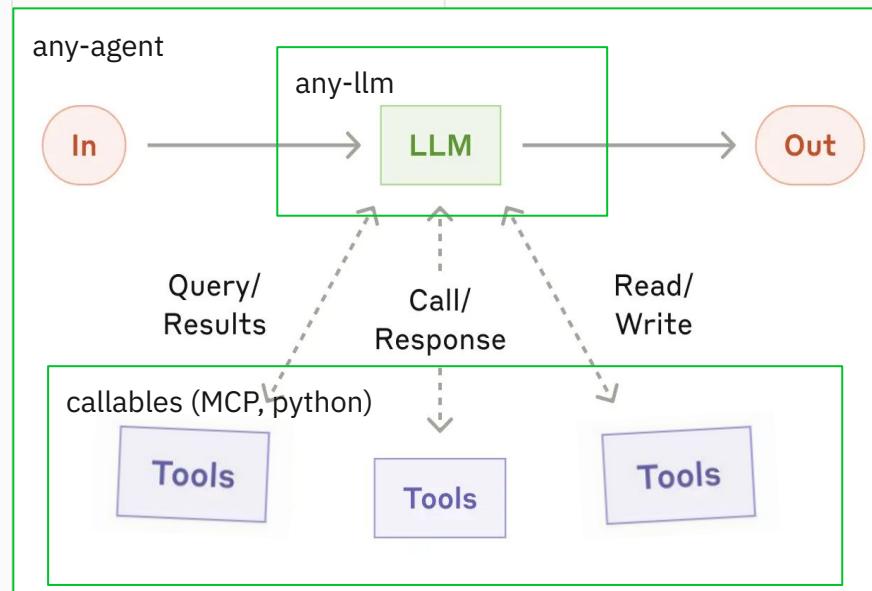
    history.append(response)

    if response.tool_executions:
        for tool_execution in tool_executions:
            tool_response = EXECUTE_TOOL(tool_execution)
            history.append(tool_response)
    else:
        return response
```

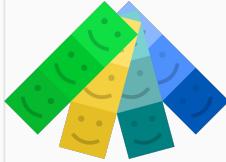


## any-agent

- A Python library providing a **single interface** to **different agent frameworks**
  - think openai, smolagents, google, langchain, ...
- Relies on **any-lm** for LLM-calling
  - transparently swap between models
- Provides its own agent implementation called **tinyagent**
- Standardized **OpenTelemetry** traces



Any-agent (Mozilla.ai, 2025)



# any-agent: minimal code example

```
● ● ●

from any_agent import AgentConfig, AnyAgent
from any_agent.tools import search_web, visit_webpage

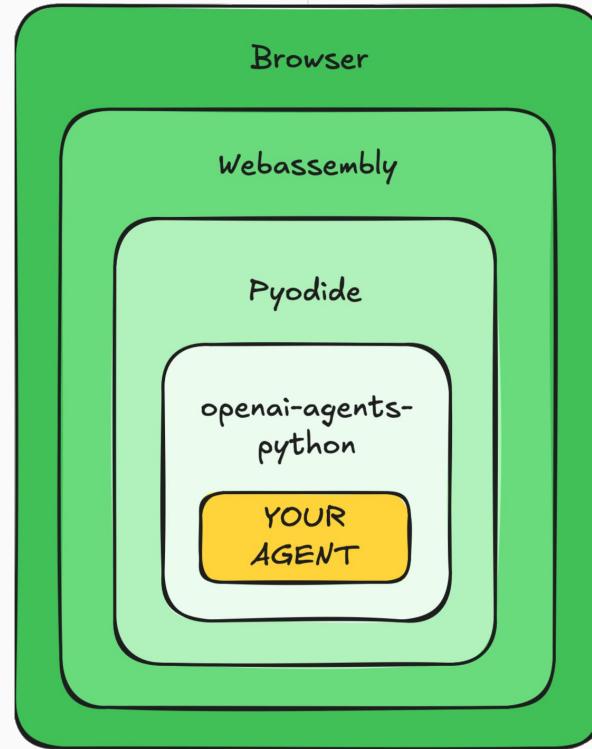
agent = AnyAgent.create(
    "tinyagent",
    AgentConfig(
        model_id="gpt-4.1-mini",
        instructions="Use the tools, please",
        tools=[search_web, visit_webpage]
    )
)

agent_trace = agent.run(
    "Which Agent Framework is the best??")
```

[Any-agent](#) (Mozilla.ai, 2025)

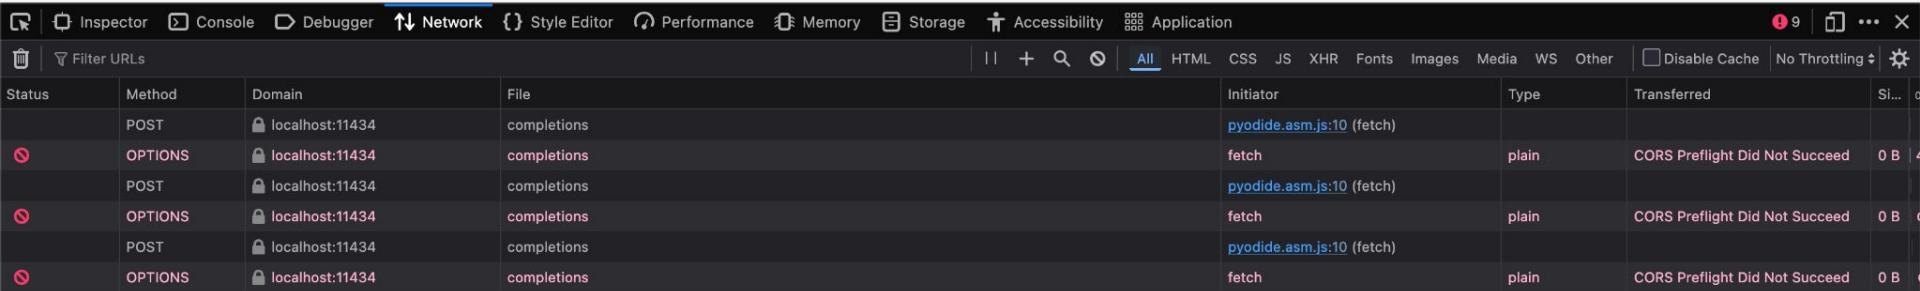
# Demo time!

# wasm-agents



# Set up CORS first!

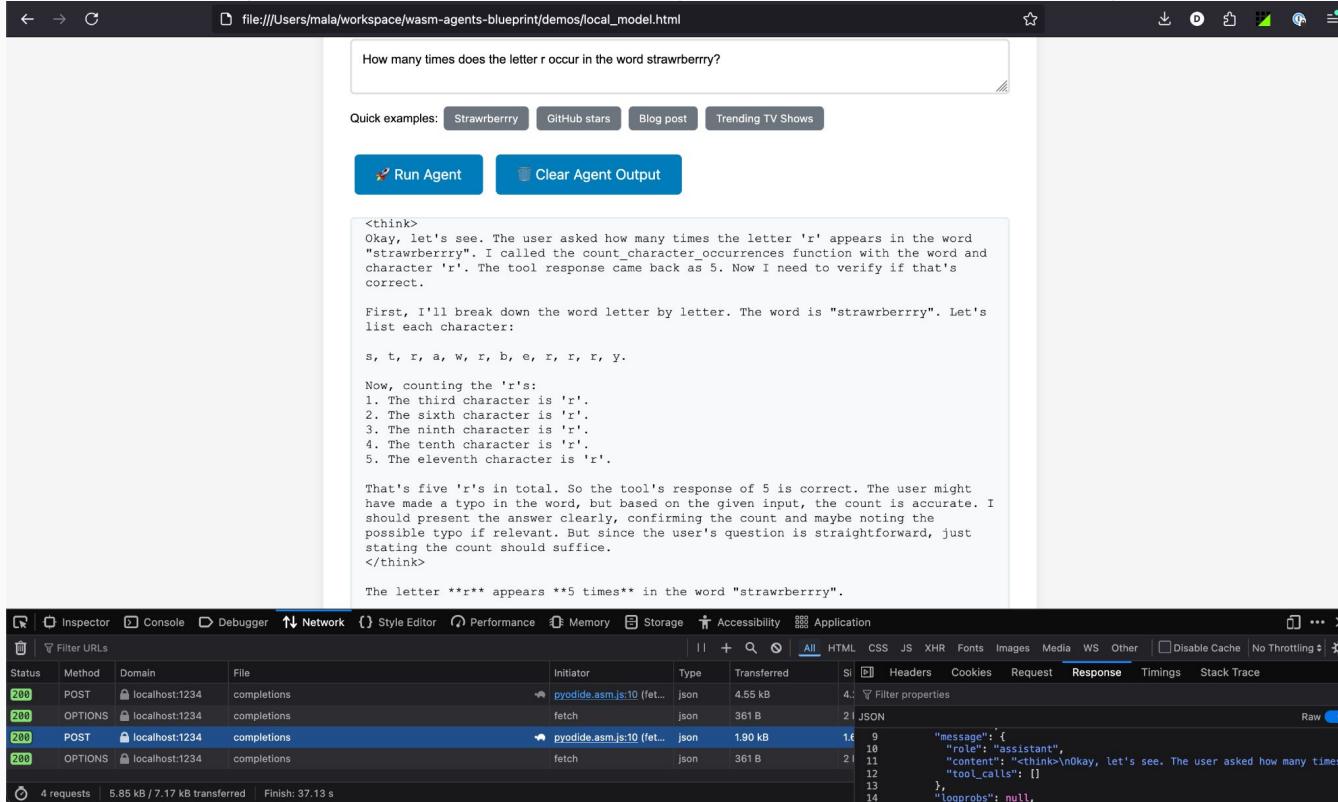
```
[15:02:56] 🚀 Setting up agent and running...
[15:02:56] 🚀 Server: http://localhost:11434/v1 | Model: qwen3:8b
[15:02:56] ✅ Tavily search enabled
[15:02:56] 🚀 Prompt: "What is the title of the latest post on https://aittalam.github.io, when was it published, what is it about, and what is the absolute URL of the image at the beginning of the post?
IMPORTANT: if you need to follow links to get all the required information, assume I have already authorized you to follow them as long as they point to the same domain."
[15:02:58] ❌❌ Agent execution failed or returned empty result
[15:02:58] ❌ Result: undefined
```



The screenshot shows the Mozilla DevTools Network tab with the following table of network requests:

Status	Method	Domain	File	Initiator	Type	Transferred	Size
	POST	localhost:11434	completions	pyodide.asm.js:10 (fetch)			
🚫	OPTIONS	localhost:11434	completions	fetch	plain	CORS Preflight Did Not Succeed	0 B
	POST	localhost:11434	completions	pyodide.asm.js:10 (fetch)			
🚫	OPTIONS	localhost:11434	completions	fetch	plain	CORS Preflight Did Not Succeed	0 B
	POST	localhost:11434	completions	pyodide.asm.js:10 (fetch)			
🚫	OPTIONS	localhost:11434	completions	fetch	plain	CORS Preflight Did Not Succeed	0 B

# StRawRbeRRy fields forever



file:///Users/mala/workspace/wasm-agents-blueprint/demos/local\_model.html

How many times does the letter r occur in the word strawrberry?

Quick examples: [Strawberrry](#) [GitHub stars](#) [Blog post](#) [Trending TV Shows](#)

[Run Agent](#) [Clear Agent Output](#)

```

<think>
Okay, let's see. The user asked how many times the letter 'r' appears in the word "strawrberry". I called the count_character_occurrences function with the word and character 'r'. The tool response came back as 5. Now I need to verify if that's correct.

First, I'll break down the word letter by letter. The word is "strawrberry". Let's list each character:
s, t, r, a, w, r, b, e, r, r, r, y.

Now, counting the 'r's:
1. The third character is 'r'.
2. The sixth character is 'r'.
3. The ninth character is 'r'.
4. The tenth character is 'r'.
5. The eleventh character is 'r'.

That's five 'r's in total. So the tool's response of 5 is correct. The user might have made a typo in the word, but based on the given input, the count is accurate. I should present the answer clearly, confirming the count and maybe noting the possible typo if relevant. But since the user's question is straightforward, just stating the count should suffice.
</think>

```

The letter \*\*\*r\*\*\* appears \*\*5 times\*\* in the word "strawrberry".

Network tab (Mozilla DevTools)

Status	Method	Domain	File	Initiator	Type	Transferred	Headers	Cookies	Request	Response	Timings	Stack Trace
200	POST	localhost:1234	completions	pyodide.asm.js:10 (fet...	json	4.55 kB	4.55 kB					
200	OPTIONS	localhost:1234	completions	fetch	json	361 B	361 B					
200	POST	localhost:1234	completions	pyodide.asm.js:10 (fet...	json	1.90 kB	1.90 kB					
200	OPTIONS	localhost:1234	completions	fetch	json	361 B	361 B					

Raw (Response content)

```

{
  "message": {
    "role": "assistant",
    "content": "<think>\nOkay, let's see. The user asked how many times the letter 'r' appears in the word \"strawrberry\". I called the count_character_occurrences function with the word and character 'r'. The tool response came back as 5. Now I need to verify if that's correct.\n\nFirst, I'll break down the word letter by letter. The word is \"strawrberry\". Let's list each character:\ns, t, r, a, w, r, b, e, r, r, r, y.\n\nNow, counting the 'r's:\n1. The third character is 'r'.\n2. The sixth character is 'r'.\n3. The ninth character is 'r'.\n4. The tenth character is 'r'.\n5. The eleventh character is 'r'.\n\nThat's five 'r's in total. So the tool's response of 5 is correct. The user might have made a typo in the word, but based on the given input, the count is accurate. I should present the answer clearly, confirming the count and maybe noting the possible typo if relevant. But since the user's question is straightforward, just stating the count should suffice.\n</think>\n\nThe letter ***r*** appears **5 times** in the word \"strawrberry\"."
  }
}

```

4 requests 5.85 kB / 7.17 kB transferred Finish: 37.13 s

# Monitor websites

## Custom Prompt:

What is the title of the latest post on <https://aittalam.github.io>, when was it published, what is it about, and what is the absolute URL of the image at the beginning of the post?

IMPORTANT: if you need to follow links to get all the required information, assume I have already authorized you to follow them as long as they point to the same domain.

Quick examples: [Strawberry](#) [GitHub stars](#) [Blog post](#) [Trending TV Shows](#)

 Run Agent

 Clear Agent Output

 View response as log

## Content Summary

The post announces the release of **Ch35t 0.1.9**, a decentralized riddle format and Python library. Key points:

- **Motivation:** To address the centralization of riddles, enabling offline play and sharing without relying on a single server.
- **Features:**
  - A JSON schema format for riddles with encrypted payloads and hints.
  - Support for encryption methods (e.g., MD5, AES) and portable clients (e.g., Marimo UI).
  - Riddles can be shared as URLs, JSON files, or text, ensuring longevity even if the original site is unavailable.
- **Use Case:** Ideal for scavenger hunts, puzzles, or creative challenges.

## Absolute URL of the Image

The first image in the post is:

[https://aittalam.github.io/images/ch35t\\_01.png](https://aittalam.github.io/images/ch35t_01.png)

# Build M3U files

Model Name:

Quick presets: [Ollama](#) [LM Studio](#) [EndSummerCamp](#)

Custom Prompt:

Connect to <https://> [REDACTED] Nirvana/Nevermind/, get the list of mp3 files available there and build an M3U file containing all the mp3 URLs

Quick examples: [Strawrberry](#) [GitHub stars](#) [Blog post](#) [Trending TV Shows](#)

 [Run Agent](#)  [Clear Agent Output](#)  [View response as log](#)

Okay, let's see. The user wants an M3U file with all the MP3 URLs from that webpage. I visited the URL and got the list of files. Now, I need to extract the MP3 filenames from the table.

Looking at the response, each entry has a name like "01 - Smells Like Teen Spirit.mp3" but with spaces encoded as %20. The URLs are constructed by appending these filenames to the base URL. So, the full URL for each track would be the base URL plus the filename.

Next, I need to format these into an M3U file. The standard M3U format starts with #EXTM3U, followed by each track's metadata. Each track entry starts with #EXTINF, which includes the track title, then the URL on a new line.

So, for each MP3 file, I'll split the filename to get the track title (after the hyphen and before .mp3), then create the #EXTINF line with that title. The URL is the base URL plus the encoded filename.

Wait, but the filenames have %20, which is URL-encoded. Do I need to decode them? The server might handle the encoding, so using the exact filenames from the table is safer. However, the M3U URLs should be properly encoded. So I'll use the filenames as they appear in the table, including the %20.

Let me list all the filenames:

1. 01%20-%20Smells%20Like%20Teen%20Spirit.mp3

# Do LLMs dream of databases?

Custom Prompt:

Can you use fetch\_url to connect to imdb and find the IMDB id for the movie "The Matrix"?

Quick examples: [Strawberry](#) [GitHub stars](#) [Blog post](#)

🚀 Run Agent Clear Agent Output

```
[19:28:35] ⓘ
<think>
Okay, let's see. The user asked for the IMDB ID of "The Matrix". I used the
fetcher_mcp_fetch_url function to get the search results from IMDB. The response has
a list of movies, and I need to find the correct one.

Looking at the content, the first result is "The Matrix" from 1999. The URL linked
is https://www.imdb.com/title/tt0133093/. The IMDB ID is usually the part after "/"
"title/" and before "/", so here it's tt0133093.

I should confirm that this is the correct ID. The user might need just the ID, so
I'll extract it and present it as the answer. Other results are sequels or related
content, so the first one is likely the right choice.
</think>

The IMDB ID for the movie **"The Matrix"** (1999) is **tt0133093**. This can be
found in the URL of the first search result: https://www.imdb.com/title/tt0133093/.
```

Inspecting the Network tab in Mozilla DevTools, we can see the following requests:

Status	Method	Domain	File	Initiator	Type	Transferred	...
200	GET	localhost:8090	tools	pyodide.asm.js:10 (xhr)	json	1.70 kB	...
200	GET	localhost:8090	tools	pyodide.asm.js:10 (xhr)	json	1.75 kB	...
200	POST	localhost:1234	completions	pyodide.asm.js:10 (fetch)	json	8.76 kB	...
200	OPTIONS	localhost:1234	completions	fetch	json	361 B	...
200	POST	localhost:8090	fetch_url	pyodide.asm.js:10 (xhr)	json	26.67 kB	...
200	OPTIONS	localhost:8090	fetch_url	xhr	plain	293 B	...
200	POST	localhost:1234	completions	pyodide.asm.js:10 (fetch)	json	1.65 kB	...
200	OPTIONS	localhost:1234	completions	fetch	json	361 B	...

Details for the fetch\_url request (row 5):

```

Request
disableMedia: true
extractContent: true
navigationTimeout: 10000
returnHTML: false
timeout: 30000
url: "https://www.imdb.com/search/title/?title=The+Matrix"
waitForNavigation: false
waitUntil: "load"
    
```

# Adding MCP tools: ZIM

file:///Users/maia/workspace/wasm-agents-blueprint/demos/local\_mcpd.html

**Run Agent** **Clear Agent Output**

```
[20:56:55] 🚀 Setting up agent and running...
[20:56:55] 🚀 Server: http://localhost:1234/v1 | Model: qwen/qwen3-8b
[20:56:55] 🚀 Prompt: "When was Denny Vrandecic born? Use ZIM files as your
information source."
[20:57:41] ✅ Agent code ran successfully! Check console for Python output
[20:57:41] 🚀
[20:57:41] 🚀 AGENT RESPONSE:
[20:57:41] 🚀
[20:57:41] 🚀
<think>
</think>
[20:57:41] 🚀
```

Denny Vrandecic (born \*\*27 February 1978\*\*) is a Croatian computer scientist. The information was retrieved from the Wikipedia ZIM file, which contains detailed biographical data about him.

Inspector Console Debugger Network Style Editor Performance Memory Storage Accessibility Application

Filter URLs

Status	Method	Domain	File
200	GET	localhost:8090	servers
200	GET	localhost:8090	tools
200	GET	localhost:8090	tools
200	GET	localhost:8090	tools
200	GET	localhost:8090	tools
200	POST	localhost:1234	completions
200	OPTIONS	localhost:1234	completions
200	POST	localhost:8090	list_zim_files
200	POST	localhost:1234	completions
200	OPTIONS	localhost:1234	completions
200	POST	localhost:8090	search_zim_file
200	OPTIONS	localhost:8090	search_zim_file
200	POST	localhost:1234	completions
200	OPTIONS	localhost:1234	completions
200	POST	localhost:8090	get_zim_entry
200	OPTIONS	localhost:8090	get_zim_entry
200	POST	localhost:1234	completions
200	OPTIONS	localhost:1234	completions

**Network**

Initiator	Type	Transferred	Size
pyodide.asm.js:10 (xhr)	json	192 B	38 B
pyodide.asm.js:10 (xhr)	json	2.68 kB	2.4...
pyodide.asm.js:10 (xhr)	json	997 B	79...
pyodide.asm.js:10 (xhr)	json	1.70 kB	1.4...
pyodide.asm.js:10 (xhr)	json	1.75 kB	1.5...
pyodide.asm.js:10 (fetch)	json	3.54 kB	3.2...
fetch	json	361 B	2 B
pyodide.asm.js:10 (xhr)	json	1.02 kB	86...
pyodide.asm.js:10 (fetch)	json	2.42 kB	2.1...
fetch	json	361 B	2 B
pyodide.asm.js:10 (xhr)	json	3.74 kB	3.5...
plain	plain	293 B	0 B
pyodide.asm.js:10 (fetch)	json	2.73 kB	2.4...
fetch	json	361 B	2 B
pyodide.asm.js:10 (xhr)	json	2.71 kB	2.5...
plain	plain	293 B	0 B
pyodide.asm.js:10 (fetch)	json	1.01 kB	70...
fetch	json	361 B	2 B

18 requests | 21.82 kB / 26.51 kB transferred | Finish: 46.11 s

# Trust

# Does this really work?

EXECUTE\_TOOL: final\_answer

Input

```
{
  "answer": "Cory Doctorow's novels in order of publication are: 1. Down and Out in Paris and London (1995), 2. East of the Sun (1998), 3. The Ring-Binder (2001), 4. The Rusted Sky (2002), 5. Football for the Poor (2003), 6. The Last Policewoman (2004), 7. The Art of Theft (2005), 8. The End (2006), 9. The Sharp End (2007), 10. The First Fifteen Lives of Harry August (2008), 11. Wired (2009), 12. The Lost (2010), 13. The Last Thing He Wanted (2017), 14. Walkaway (2019), 15. The Companion (2020). Note: Some sources may list the publication years differently due to different editions or formats."
}
```

OUTPUT

Cory Doctorow's novels in order of publication are: 1. Down and Out in Paris and London (1995), 2. East of the Sun (1998), 3. The Ring-Binder (2001), 4. The Rusted Sky (2002), 5. Football for the Poor (2003), 6. The Last Policewoman (2004), 7. The Art of Theft (2005), 8. The End (2006), 9. The Sharp End (2007), 10. The First Fifteen Lives of Harry August (2008), 11. Wired (2009), 12. The Lost (2010), 13. The Last Thing He Wanted (2017), 14. Walkaway (2019), 15. The Companion (2020). Note: Some sources may list the publication years differently due to different editions or formats.

## Selected bibliography [edit]

In chronological sequence, unless otherwise indicated

### Fiction [edit]

#### Novels [edit]

- *Down and Out in the Magic Kingdom*. Tor. 2003. ISBN 0-7653-0436-8.
- *Eastern Standard Tribe*. Tor. 2004. ISBN 0-7653-0759-6.
- *Someone Comes to Town, Someone Leaves Town*. Tor. 2005. ISBN 0-7653-1278-6.
- *Makers*. Tor. 2009. ISBN 978-0-7653-1279-2.
- *For the Win*. Tor. 2010. ISBN 978-0-7653-2216-6.
- *The Great Big Beautiful Tomorrow*, 2011, ISBN 978-1-6048-6404-5
- *The Rapture of the Nerds*. Tor. September 2012. ISBN 978-0-765-32910-3. (with Charles Stross)
- *Pirate Cinema*. Tor. 2012. ISBN 978-0-7653-2908-0.
- *Walkaway*. Tor. 2017. ISBN 978-0-7653-9276-3.
- *The Lost Cause*. Tor. 2023. ISBN 978-1-0359-0223-1.

# Problem: LLMs are gonna LLM

- They always answer

# Problem: LLMs are gonna LLM

- They always answer
- Sometimes they say they don't know

Davide Eynard does not appear to have a notable online presence or widely documented public information, as the Wikipedia page for this name returns a 404 error. This may indicate the name is misspelled, refers to a private individual, or is fictional. For accurate birthdate information, please verify the name spelling or provide additional context about the person you're seeking.

Custom Prompt:  
When was Davide Eynard born?

Quick examples: [Strawberry](#) [GitHub stars](#) [Blog post](#) [Trending TV Shows](#)

[Run Agent](#) [Clear Agent Output](#) [View response as markdown](#)

[08:56:48] ⓘ  
<think>  
Okay, let's see. The user asked when Davide Eynard was born. I tried to use the visit\_webpage tool to check Wikipedia, but got a 404 error. That means the URL doesn't exist. Hmm.

First, maybe the name is misspelled. Let me think. Davide is a Italian name, common in some regions. Eynard might be a typo. For example, maybe it's "Davide Eynard" as a fictional name? Or perhaps it's "David Eynard" or another variation.

Wait, the 404 says the page doesn't exist on Wikipedia. So maybe there's no notable figure or a person not covered by Wikipedia. Since the user might have confused the name. Let me get the birthdate from there.

to inform the user that the information isn't available. The visit\_webpage failed, and there's no other tool that I couldn't find the birthdate for Davide Eynard, possibly due to the name being misspelled or not widely documented.

I should also check if there's another way, but the tools only have count\_character\_occurrences and visit\_webpage. Count\_character\_occurrences is for counting characters in a word, which isn't relevant here. So no other tools can help.

So, the answer should state that the birthdate information isn't available and suggest possible issues with the name.

</think>

Davide Eynard does not appear to have a notable online presence or widely documented public information, as the Wikipedia page for this name returns a 404 error. This may indicate the name is misspelled, refers to a private individual, or is fictional. For accurate birthdate information, please verify the name spelling or provide additional context about the person you're seeking.

[08:56:48] ⓘ

# Problem: LLMs are gonna LLM

- They always answer
- Sometimes they say they don't know
- BUT

EXECUTE\_TOOL: final\_answer

Input

```
{
  "answer": "I'm unable to retrieve information about Davide Eynard's birth date as the requested webpage is forbidden. Please try another source or provide additional context."
}
```

OUTPUT

I'm unable to retrieve information about Davide Eynard's birth date as the requested webpage is forbidden. Please try another source or provide additional context.

Custom Prompt:

When was Davide Eynard born?

Quick examples: Strawberry GitHub stars Blog post Trending TV Shows

Run Agent

Clear Agent Output

View response as markdown

[08:56:48] ⓘ

&lt;think&gt;

Okay, let's see. The user asked when Davide Eynard was born. I tried to use the visit\_webpage tool to check Wikipedia, but got a 404 error. That means the URL doesn't exist. Now,

available through the tools. Since the visit\_webpage failed, and there's no other tool to use, I should tell the user that I couldn't find the birthdate for Davide Eynard, possibly due to the name being incorrect or not widely documented.

I should also check if there's another way, but the tools only have count\_character\_occurrences and visit\_webpage. Count\_character\_occurrences is for counting characters in a word, which isn't relevant here. So no other tools can help.

So, the answer should state that the birthdate information isn't available and suggest possible issues with the name.

&lt;/think&gt;

Davide Eynard does not appear to have a notable online presence or widely documented public information, as the Wikipedia page for this name returns a 404 error. This may indicate the name is misspelled, refers to a private individual, or is fictional. For accurate birthdate information, please verify the name spelling or provide additional context about the person you're seeking.

[08:56:48] ⓘ

# Problem: LLMs are gonna LLM

- They always answer
- Sometimes they say they don't know
- BUT ONLY

## OUTPUT

First, I need to check if I have any existing knowledge about Davide Eynard. From what I remember, Davide Eynard is a Swiss footballer who plays as a defender. He's been in the Swiss Super League, so maybe he's not too well-known internationally.

Since the user wants me to use available tools, I should consider if there are any databases or resources I can access. However, in a typical scenario without external tools, I rely on my training data up to 2023. Let me recall if I have his birth date stored.

Wait, I think I remember that he was born on June 15, 1997. But I'm not entirely sure. To confirm, I would usually check a reliable source like a sports database or his official profile. Since I can't access external tools right now, I need to rely on my existing knowledge.

Alternatively, maybe there's a possibility of confusion with another person named Davide Eynard. But given the context of football, I believe the correct answer is June 15, 1997. I should present this information clearly, making sure it's accurate based on my training data.

Davide Eynard was born on June 15, 1997.

## Custom Prompt:

When was Davide Eynard born?

Quick examples: [Strawberry](#) [GitHub stars](#) [Blog post](#) [Trending TV Shows](#)

 Run Agent

 Clear Agent Output

 View response as markdown

[08:56:48] 

</think>

count\_character\_occurrences and visit\_message. count\_character\_occurrences is for counting characters in a word, which isn't relevant here. So no other tools can help.

So, the answer should state that the birthdate information isn't available and suggest possible issues with the name.  
</think>

Davide Eynard does not appear to have a notable online presence or widely documented public information, as the Wikipedia page for this name returns a 404 error. This may indicate the name is misspelled, refers to a private individual, or is fictional. For accurate birthdate information, please verify the name spelling or provide additional context about the person you're seeking.

[08:56:48] 

# Problem: LLMs are gonna LLM

- They always answer
- Sometimes they say they don't know
- BUT ONLY SOMETIMES

Davide Eynard was born on May 19, 1995.

This information can be verified through sports databases or biographical sources, which typically list the birth dates of professional athletes.

Custom Prompt:

When was Davide Eynard born?

Quick examples: [Strawberry](#) [GitHub stars](#) [Blog post](#) [Trending TV Shows](#)

 Run Agent

 Clear Agent Output

 View response as markdown

[08:56:48] 

<think>

Okay, let's see. The user asked when Davide Eynard was born. I tried to use the visit\_webpage tool to check Wikipedia, but got a 404 error. That means the URL doesn't exist. Hmm.

person with that name. Alternatively, the user might have confused the name. Let me check if there's a similar name.

Alternatively, maybe it's a local figure or a person not covered by Wikipedia. Since the tool call returned a 404, I can't get the birthdate from there.

In this case, the assistant's role is to inform the user that the information isn't available through the tools. Since the visit\_webpage failed, and there's no other tool to use, I should tell the user that I couldn't find the birthdate for Davide Eynard, possibly due to the name being incorrect or not widely documented.

I should also check if there's another way, but the tools only have count\_character\_occurrences and visit\_webpage. Count\_character\_occurrences is for counting characters in a word, which isn't relevant here. So no other tools can help.

So, the answer should state that the birthdate information isn't available and suggest possible issues with the name.

</think>

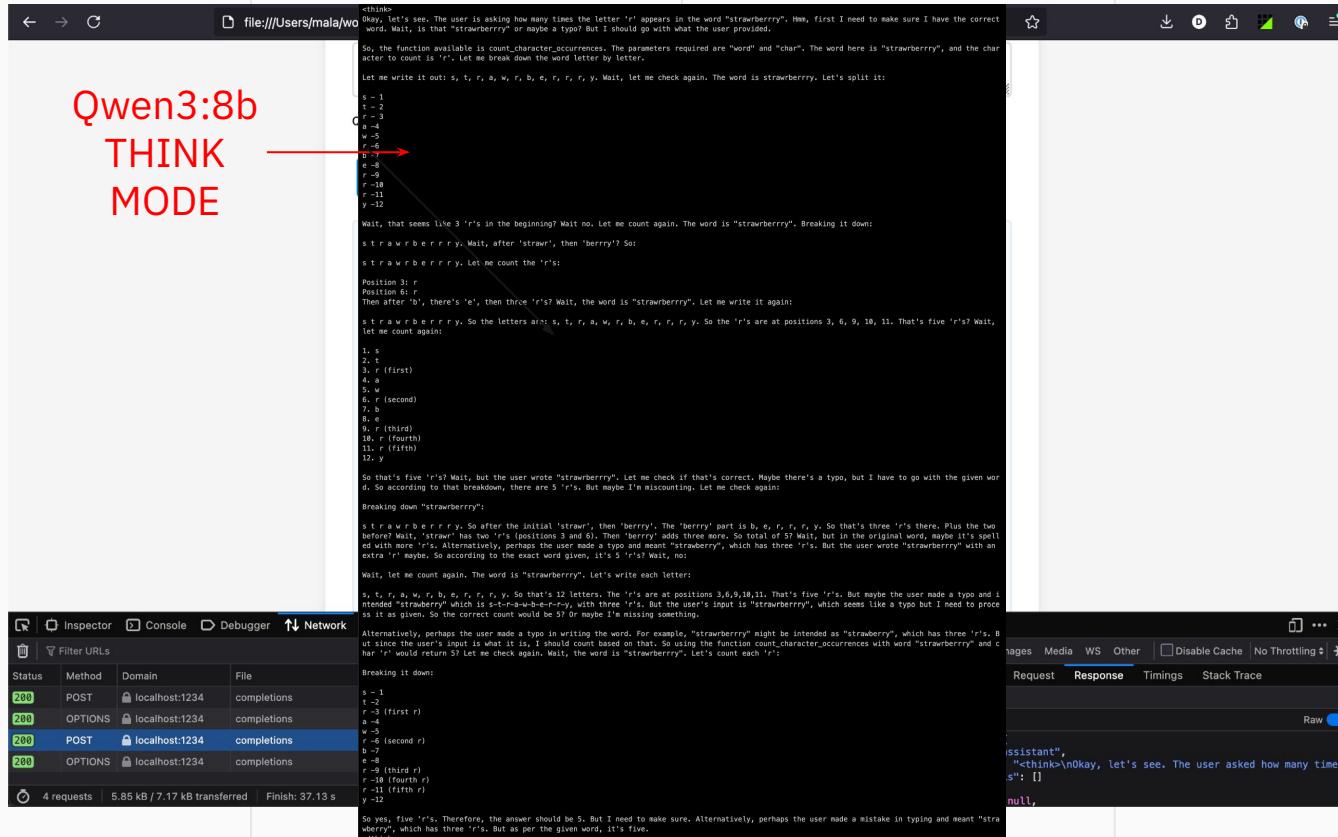
Davide Eynard does not appear to have a notable online presence or widely documented public information, as the Wikipedia page for this name returns a 404 error. This may indicate the name is misspelled, refers to a private individual, or is fictional. For accurate birthdate information, please verify the name spelling or provide additional context about the person you're seeking.

[08:56:48] 

# StRawRbeRRy fields (literally) forever

Qwen3:8b

THINK MODE



```

<think>
Okay, let's see. The user is asking how many times the letter 'r' appears in the word "strawberry". Hm, first I need to make sure I have the correct word. Wait, is that "strawberry" or maybe a typo? But I need to go with what the user provided.

So, the function available is count_character_occurrences. The parameters required are "word" and "char". The word here is "strawberry", and the character to count is 'r'. Let me break down the word letter by letter.

Let me write it out: s, t, r, a, w, r, b, e, r, r, y. Wait, let me check again. The word is strawberry. Let's split it:
s - 1
t - 2
r - 3
a - 4
w - 5
r - 6
b - 7
e - 8
r - 9
r - 10
r - 11
y - 12

Wait, that seems like 3 'r's in the beginning? Wait no. Let me count again. The word is "strawberry". Breaking it down:
s t r a w b e r r y. Wait, after 'straw', then 'berry'? So:
s t r a w b e r r y. Let me count the 'r's:
Position 3: r
Position 6: r
Then after 'b', there's 'e', then thrice 'r's? Wait, the word is "strawberry". Let me write it again:
s t r a w b e r r y. So the letters are s, t, r, a, w, r, b, e, r, r, r, y. So the 'r's are at positions 3, 6, 9, 10, 11. That's five 'r's? Wait, let me count again:
1. s
2. t
3. r (First)
4. a
5. w
6. r (Second)
7. b
8. e
9. r (Third)
10. r (Fourth)
11. r (Fifth)
12. y

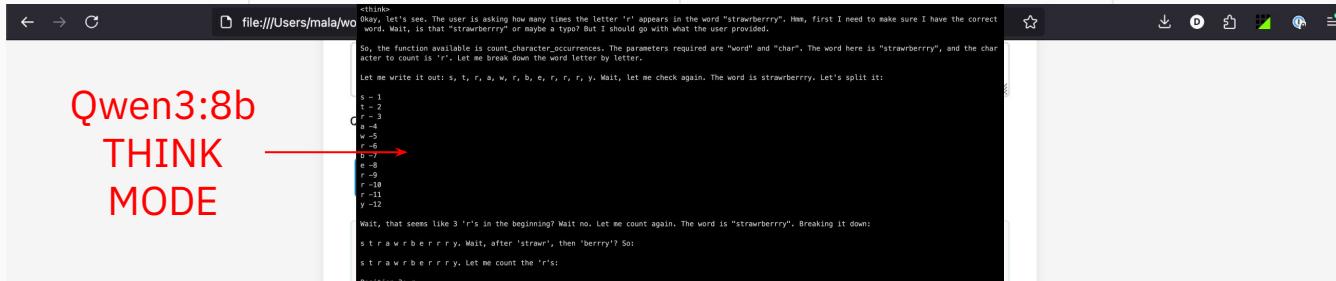
So that's five 'r's? Wait, but the user wrote "strawberry". Let me check if that's correct. Maybe there's a typo, but I have to go with the given word. So according to that breakdown, there are 5 'r's. But maybe I'm miscounting. Let me check again:
Breaking down "strawberry":
s t r a w b e r r y. So after the initial "straw", then "berry". The "berry" part is b, e, r, r, r, y. So that's three 'r's there. Plus the two before 'b'. Wait, "straw" has two 'r's (positions 3 and 6). Then "berry" adds three more. So total of 5? Wait, but in the original word, maybe it's spelled with more 'r's. Alternatively, perhaps the user made a typo and meant "strawberry", which has three 'r's. But the user wrote "strawberry" with an extra 'r' maybe. So according to the exact word given, it's 5 'r's? Wait, no:
Wait, let me count again. The word is "strawberry". Let's write it again:
s, t, r, a, w, r, b, e, r, r, y. So that's 12 letters. The 'r's are at positions 3, 6, 9, 10, 11. That's five 'r's. But maybe the user made a typo and intended "strawberry" which is s-t-r-a-w-b-e-r-r-y, with three 'r's. But the user's input is "strawberry", which seems like a typo but I need to proceed as it is given. So the correct count would be 5? Or maybe I'm missing something.

Alternatively, perhaps the user made a typo in writing the word. For example, "strawberry" might be intended as "strawberry", which has three 'r's. But since the user's input is what it is, I should count based on that. So using the function count_character_occurrences with word "strawberry" and character 'r' would return 5? Let me check again. Wait, the word is "strawberry". Let's count each 'r':
Breaking it down:
s - 1
t - 2
r - 3 (first r)
a - 4
w - 5
r - 6 (second r)
b - 7
e - 8
r - 9 (third r)
r - 10 (fourth r)
r - 11 (fifth r)
y - 12

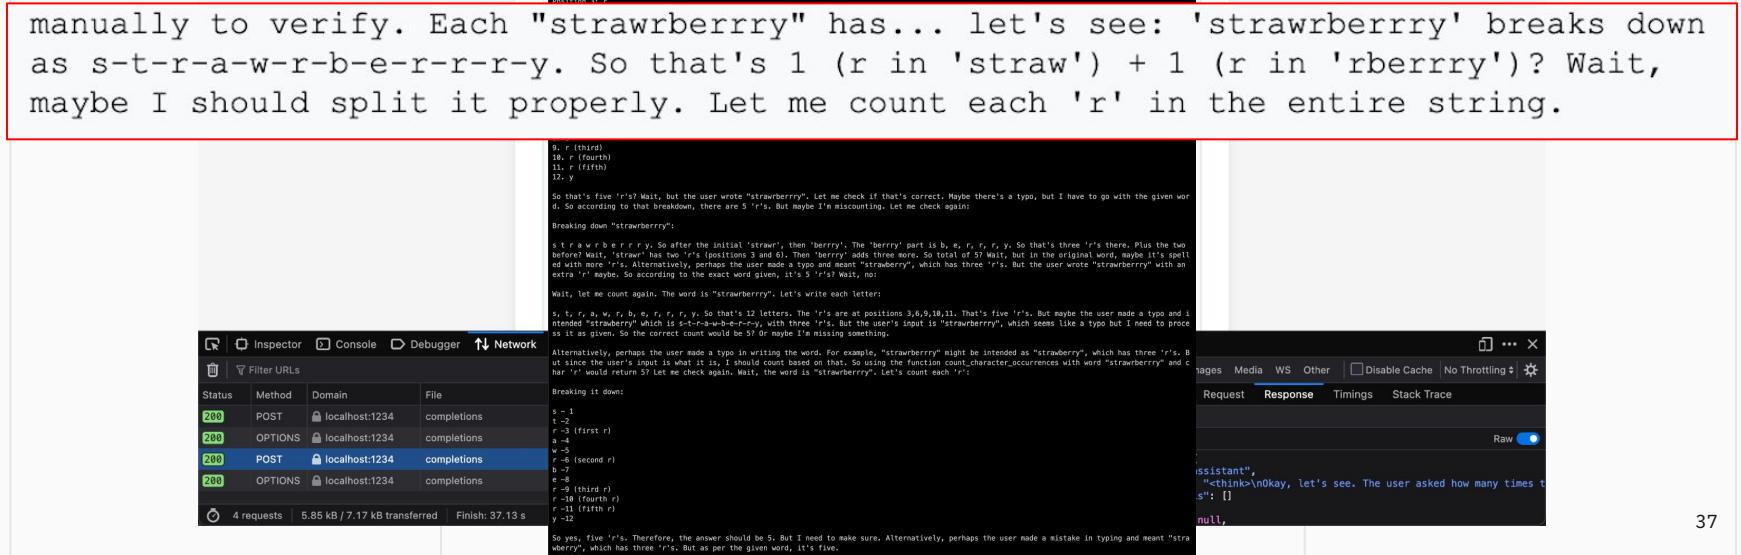
So yes, five 'r's. Therefore, the answer should be 5. But I need to make sure. Alternatively, perhaps the user made a mistake in typing and meant "strawberry", which has three 'r's. But as per the given word, it's five.
</think>

```

# StRawRbeRRy fields (literally) forever



manually to verify. Each "strawberrry" has... let's see: 'strawberrry' breaks down as s-t-r-a-w-r-b-e-r-r-r-y. So that's 1 (r in 'straw') + 1 (r in 'rberrry')? Wait, maybe I should split it properly. Let me count each 'r' in the entire string.



# Problem: Context length

## Custom Prompt:

How many stars does the mozilla-ai/any-agent project have on GitHub?

```
time=2025-09-04T15:13:43.227+02:00 level=DEBUG source=prompt.go:68 msg="truncating input messages which exceed context length" truncated=2
time=2025-09-04T15:13:43.227+02:00 level=DEBUG source=server.go:1373 msg="completion request" images=0 prompt=22725 format=""
time=2025-09-04T15:13:43.305+02:00 level=WARNING source=runner.go:127 msg="truncating input prompt" limit=4096 prompt=7707 keep=4 new=4096
time=2025-09-04T15:13:43.307+02:00 level=DEBUG source=cache.go:104 msg="loading cache slot" id=0 cache=1771 prompt=4096 used=4 remaining=4092
time=2025-09-04T15:13:58.012+02:00 level=DEBUG source=cache.go:240 msg="context limit hit - shifting" id=0 limit=4096 input=4096 keep=4 discard=2046
update: applying K-shift
```

To effectively use the `AnyAgent` library, follow these structured steps to create and run an agent for answering questions:

### 1. Installation

Install the `any-agent` package via pip:

```
pip install 'any-agent'
```

# Problem: Context length

## Custom Prompt:

How many stars does the mozilla-ai/any-agent project have on GitHub?

```
time=2025-09-04T15:13:43.227+02:00 level=DEBUG source=prompt.go:68 msg="truncating input messages which exceed context length" truncated=2
time=2025-09-04T15:13:43.227+02:00 level=DEBUG source=server.go:1373 msg="completion request" images=0 prompt=22725 format=""
time=2025-09-04T15:13:43.305+02:00 level=WARNING source=runner.go:127 msg="truncating input prompt" limit=4096 prompt=7707 keep=4 new=4096
time=2025-09-04T15:13:43.307+02:00 level=DEBUG source=cache.go:104 msg="loading cache slot" id=0 cache=1771 prompt=4096 used=4 remaining=4092
time=2025-09-04T15:13:58.012+02:00 level=DEBUG source=cache.go:240 msg="context limit hit - shifting" id=0 limit=4096 input=4096 keep=4 discard=2046
update: applying K-shift
```

To effectively use the `AnyAgent` library, follow these structured steps to create and run an agent for answering questions:

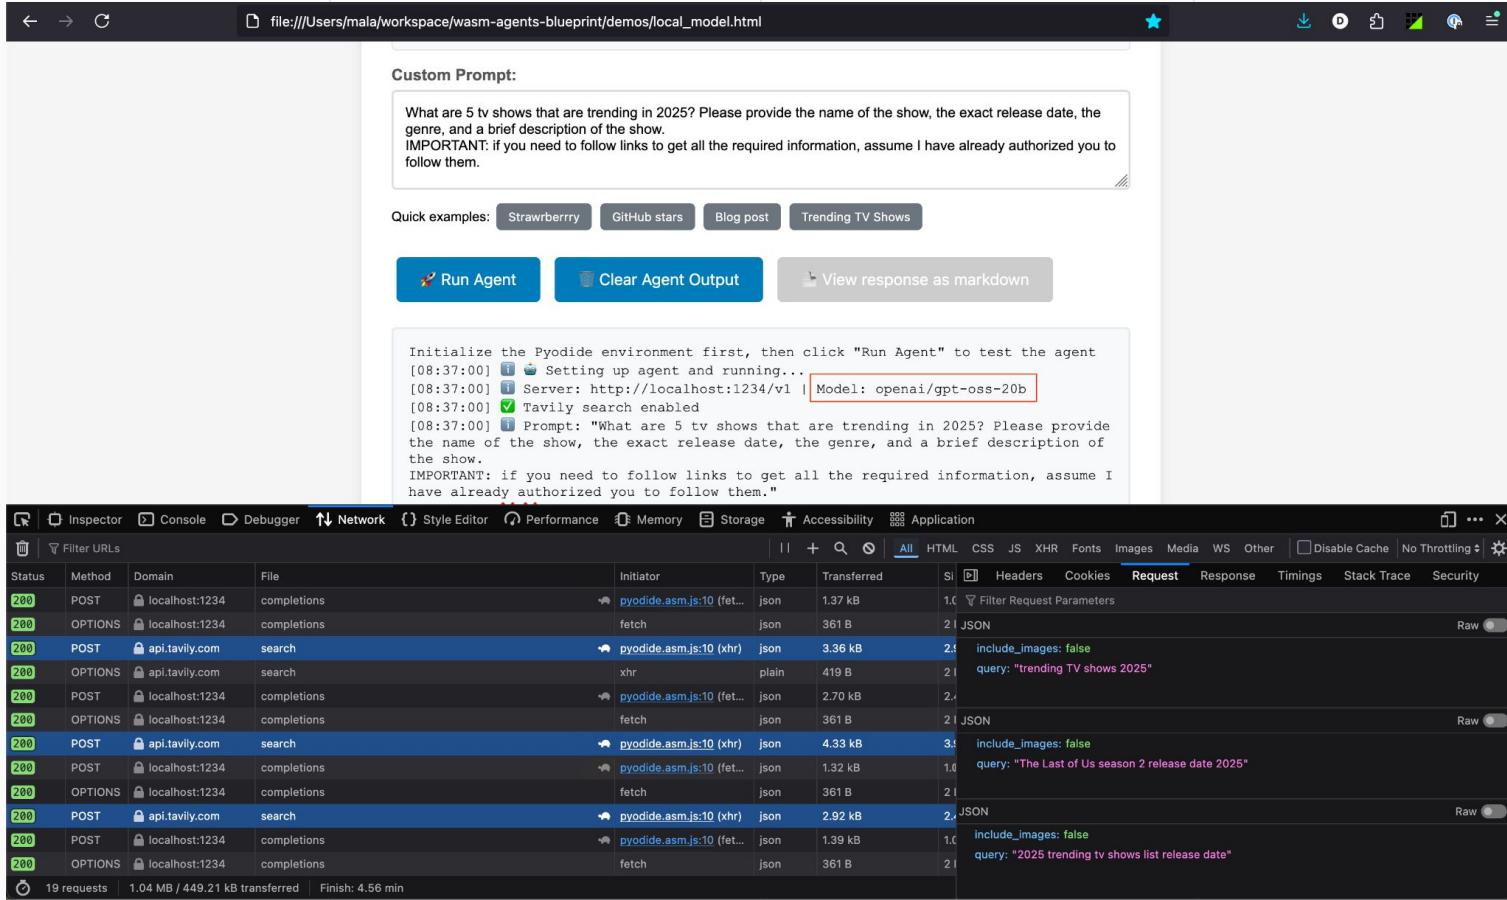
### 1. Installation

Install the `any-agent` package via pip:

```
pip install 'any-agent'
```

Know your tools!

# Model-specific insights: abusing search



Custom Prompt:

What are 5 tv shows that are trending in 2025? Please provide the name of the show, the exact release date, the genre, and a brief description of the show.  
IMPORTANT: if you need to follow links to get all the required information, assume I have already authorized you to follow them.

Quick examples: Strawberry, GitHub stars, Blog post, Trending TV Shows

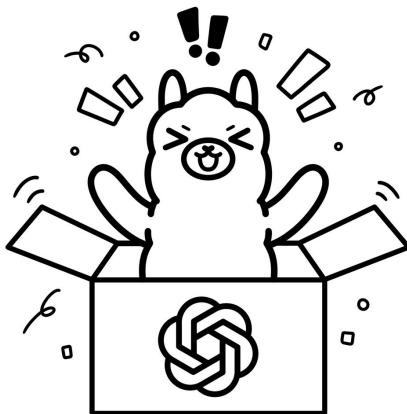
Run Agent, Clear Agent Output, View response as markdown

Initialize the Pyodide environment first, then click "Run Agent" to test the agent  
[08:37:00] 📺 Setting up agent and running...  
[08:37:00] 🚀 Server: http://localhost:1234/v1 Model: openai/gpt-oss-20b  
[08:37:00] ✅ Tavily search enabled  
[08:37:00] 📺 Prompt: "What are 5 tv shows that are trending in 2025? Please provide the name of the show, the exact release date, the genre, and a brief description of the show.  
IMPORTANT: if you need to follow links to get all the required information, assume I have already authorized you to follow them."

Status	Method	Domain	File	Initiator	Type	Transferred	Headers	Cookies	Request	Response	Timings	Stack Trace	Security
200	POST	localhost:1234	completions	pyodide.asm.js:10 (fetc...	json	1.37 kB	1.0		include_images: false				Raw
200	OPTIONS	localhost:1234	completions	fetch	json	361 B	2.1		query: "trending TV shows 2025"				Raw
200	POST	api.tavily.com	search	pyodide.asm.js:10 (xhr)	json	3.36 kB	2.4		include_images: false				Raw
200	OPTIONS	api.tavily.com	search	xhr	plain	419 B	2.4		query: "trending TV shows 2025"				Raw
200	POST	localhost:1234	completions	pyodide.asm.js:10 (fetc...	json	2.70 kB	2.4		include_images: false				Raw
200	OPTIONS	localhost:1234	completions	fetch	json	361 B	2.4		query: "The Last of Us season 2 release date 2025"				Raw
200	POST	api.tavily.com	search	pyodide.asm.js:10 (xhr)	json	4.33 kB	3.5		include_images: false				Raw
200	POST	localhost:1234	completions	pyodide.asm.js:10 (fetc...	json	1.32 kB	1.0		query: "The Last of Us season 2 release date 2025"				Raw
200	OPTIONS	localhost:1234	completions	fetch	json	361 B	2.1		include_images: false				Raw
200	POST	api.tavily.com	search	pyodide.asm.js:10 (xhr)	json	2.92 kB	2.4		query: "2025 trending tv shows list release date"				Raw
200	POST	localhost:1234	completions	pyodide.asm.js:10 (fetc...	json	1.39 kB	1.0		include_images: false				Raw
200	OPTIONS	localhost:1234	completions	fetch	json	361 B	2.1		query: "2025 trending tv shows list release date"				Raw

19 requests 1.04 MB / 449.21 kB transferred | Finish: 4.56 min

# Who owns search tools?

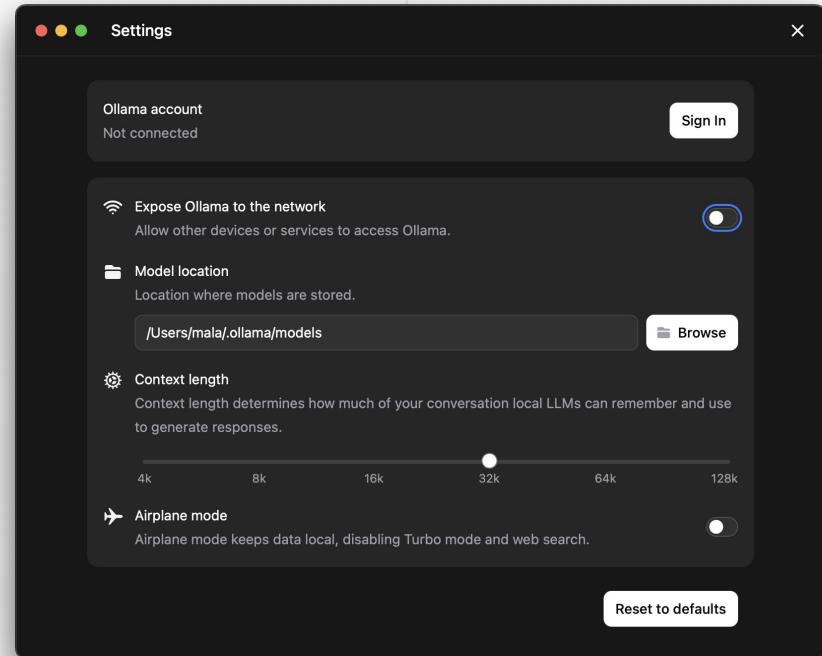


## Welcome OpenAI's gpt-oss!

Ollama partners with OpenAI to bring its latest state-of-the-art open weight models to Ollama. The two models, 20B and 120B, bring a whole new local chat experience, and are designed for powerful reasoning, agentic tasks, and versatile developer use cases.

### Feature highlights

- **Agentic capabilities:** Use the models' native capabilities for function calling, web browsing (Ollama is providing a built-in web search that can be optionally enabled to augment the model with the latest information), python tool calls, and structured outputs.



# Model-specific insights: will it work without search?

file:///Users/mala/workspace/wasm-agents-blueprint/demos/local\_model.html

gpt-oss:latest

Quick presets: Ollama LM Studio EndSummerCamp

Custom Prompt:

What are 5 tv shows that are trending in 2025? Please provide the name of the show, the exact release date, the genre, and a brief description of the show.  
IMPORTANT: if you need to follow links to get all the required information, assume I have already authorized you to follow them.

Network

Status	Method	Domain	File	Initiator	Type	Transferred
200	GET	en.wikipedia.org	2025_in_television	pyodide.asm.js:10 (xhr)	html	18.59 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	1.63 kB
200	GET	en.wikipedia.org	Alien_Earth	pyodide.asm.js:10 (xhr)	html	51.89 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	1.51 kB
404	GET	en.wikipedia.org	Hostage_(2025_TV_series)	pyodide.asm.js:10 (xhr)	html	12.08 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	663 B
404	GET	en.wikipedia.org	Hostage_(TV_series_(2025))	pyodide.asm.js:10 (xhr)	html	12.09 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	826 B
200	GET	en.wikipedia.org	index.php?title=Hostage_(TV_series)	pyodide.asm.js:10 (xhr)	html	37.61 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	1.38 kB
200	GET	en.wikipedia.org	The_Terminal_List_Dark_Wolf	pyodide.asm.js:10 (xhr)	html	33.31 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	693 B
200	GET	en.wikipedia.org	Dexter_Resurrection	pyodide.asm.js:10 (xhr)	html	44.96 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	949 B
404	GET	en.wikipedia.org	Peacemaker_(2025_TV_series)	pyodide.asm.js:10 (xhr)	html	12.09 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	711 B
404	GET	en.wikipedia.org	Peacemaker_(TV_series_(2025))	pyodide.asm.js:10 (xhr)	html	12.10 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	647 B
404	GET	en.wikipedia.org	Peacemaker_(2025_TV_series)	pyodide.asm.js:10 (xhr)	html	12.13 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	642 B
200	GET	en.wikipedia.org	Peacemaker_(TV_series)	pyodide.asm.js:10 (xhr)	html	65.59 kB

30 requests 2.33 MB / 434.34 kB transferred Finish: 4.91 min

# A new problem: stalking agents?

Custom Prompt:

Search the web for Davide Eynard's birthdate. Tell me both the date and the URL where you found it.

Quick examples: [Strawberry](#) [GitHub stars](#) [Blog post](#)

[Run Agent](#) [Clear Agent Output](#) [View response as log](#)

**Davide Eynard**

- Date of birth:** 11February1976
- Source:** The PDF CV hosted at the UC Lugano Institute of Computational Sciences lists "Birthplace and date: February 11th 1976, Monza (MB), Italy."
- URL:** [http://davide.eynard.it/cv/cv\\_20221108.pdf](http://davide.eynard.it/cv/cv_20221108.pdf)



Network

Status	Method	Domain	File	Initiator	Type	Transferred
200	OPTIONS	89.169.109.132:8090	searxng_web_search	xhr	plain	293 B
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	755 B
200	POST	89.169.109.132:8090	searxng_web_search	pyodide.asm.js:10 (xhr)	json	12.93 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	753 B
200	GET	davide.eynard.it	/about-me/	pyodide.asm.js:10 (xhr)	html	5.29 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	729 B
200	POST	89.169.109.132:8090	searxng_web_search	pyodide.asm.js:10 (xhr)	json	13.78 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	952 B
200	GET	davide.eynard.it	cv_20221108.pdf	pyodide.asm.js:10 (xhr)	pdf	90.47 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	1.55 kB
502	POST	89.169.109.132:8090	fetch_url	pyodide.asm.js:10 (xhr)	problem	1.51 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	945 B
200	GET	www.researchgate.net	Davide-Eynard	pyodide.asm.js:10 (xhr)	html	60.15 kB
200	POST	89.169.109.132:11434	completions	pyodide.asm.js:10 (fetch)	json	743 B

22 requests | 897.64 kB / 214.53 kB transferred | Finish: 1.58 min

# Conclusions

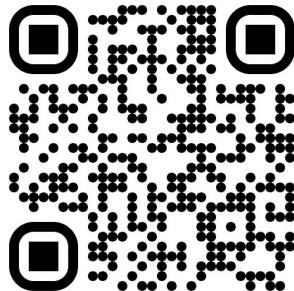
# Learnings

- There's **tinkering** with AI and tinkering **with AI**
  - you'll solve things more quickly with the former, learn more with the latter
- No one-size-fits-all solution
  - **task, compute, model training, tools**
- Easy, schmeasy: "friendly" vs actually useful
  - UX is only now evolving from chat to agents, expect things to break
- Think smol
  - limit context growth
  - few tools bring you a long way
- ~~LLMs~~ Agents are slot machines
  - remember the success, **don't forget the failures**
- Prevent **enshittification**
  - what (data, freedoms, control) are you giving away with each choice?

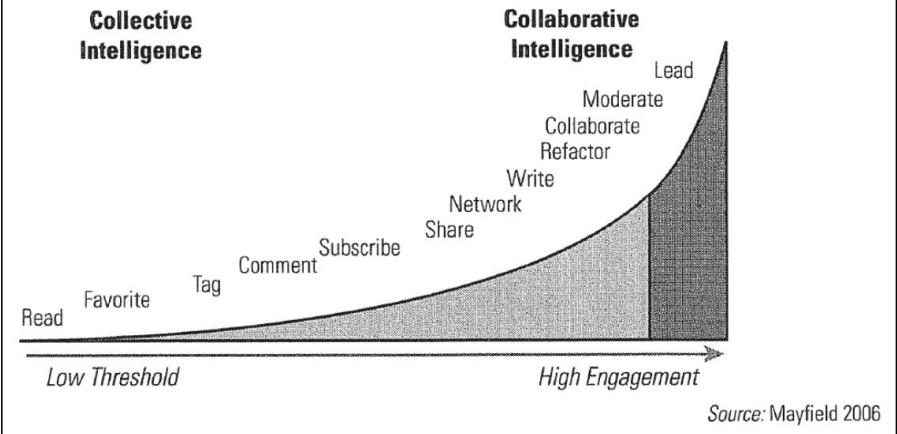
# And now?

Choose your own adventure:

- Check out [mozilla-ai GitHub org!](#)
- Play with wasm-agents
- Write your own python “any-agent”
- Test different local models
- Try different AI tools and MCP servers
- Host tools / services for your community



**Figure 2.2** The power law of participation



Source: Mayfield 2006

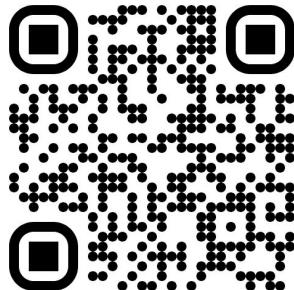
# And now?



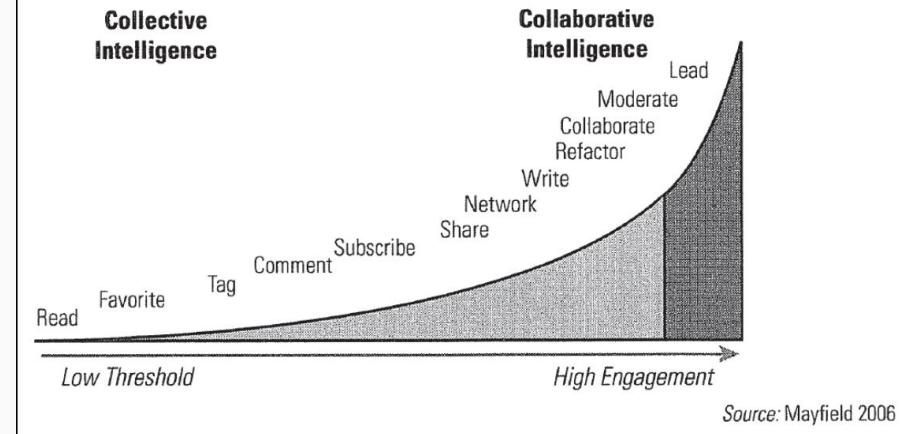
Mozilla.ai

Choose your own adventure:

- Check out mozilla-ai [GitHub org!](#)
- Play with wasm-agents
- Write your own python “any-agent”
- Test different local models
- Try different AI tools and MCP servers
- Host tools / services for your community

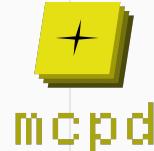


**Figure 2.2** The power law of participation



Source: Mayfield 2006

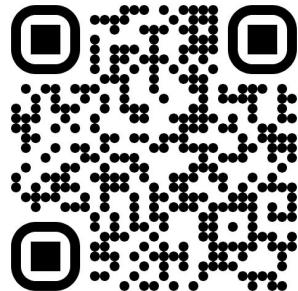
# And now?



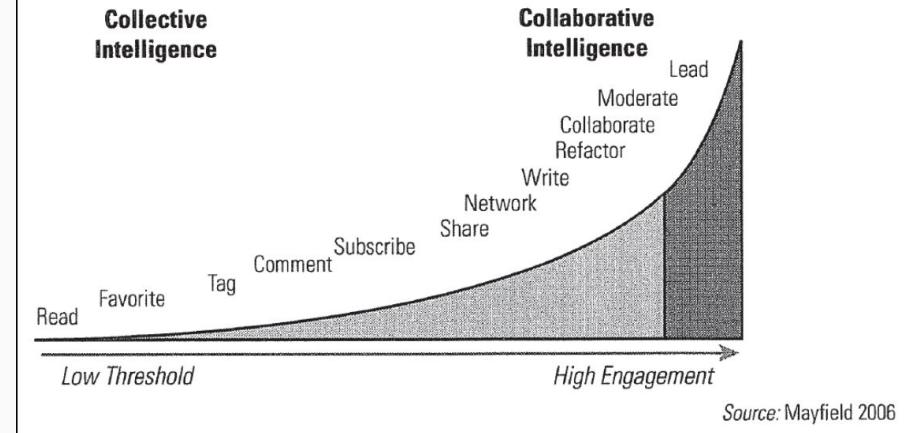
Mozilla.ai

Choose your own adventure:

- Check out mozilla-ai [GitHub org!](#)
- Play with wasm-agents
- Write your own python “any-agent”
- Test different local models
- Try different AI tools and MCP servers
- Host tools / services for your community



**Figure 2.2** The power law of participation



Source: Mayfield 2006

**BE LIKE ADA!**

# Thank you!