

OpenClaw Intro and Agent Teams

15-113 – Wednesday 4/15

Phone app reflection

- Show off what you did, and fill this out
- <https://bit.ly/3O4W4Aj>
- Eat stuff on the front table (the food stuff, not the not-food stuff)



Reminder: FCEs!

The FCEs for this class will help me shape the course for next semester!

And in advance, thanks again for being awesome and for being part of such an experimental course. And thanks for making such wonderful stuff!

Please give me ideas for new things you think we should explore next semester!

Today's cookies are ENTIRELY UNRELATED



What's up today?

- A little bit of OpenClaw / ZeroClaw
- Agent Teams
- Work time

OpenClaw and agents were what you were all most interested in, but this stuff can get *complicated*, so my experiments are still in progress

So, OpenClaw...

- General idea: Set up LLM orchestration on your local machine (or on a server you control) so that you can have:
 - Long-running agentic tasks
 - Scheduled tasks
 - Direct file i/o access
 - Hardware access
 - In theory, more control over everything, including the level of autonomy the agents have

Tradeoffs

- *Highly* configurable!
- Give OpenClaw lots of access to stuff: It sounds like it can perform some very complex tasks (like starting a business for you, though the success of that business may be debatable)
 - ...and in the process, you risk compromising your computer, bank account, your email, and everything else it touches
- Lock it down and limit its scope: It can be run *reasonably* safely, but... then it can't do much. Might not be worth much more than running agents in VS Code.

I have the fear.

- FERPA: Family Educational Rights and Privacy Act
 - I have a lot of sensitive information on my computer!
 - I join the official Felony Club if I casually expose that information to the world. Plus, I just don't want that to happen.
- I am generally quite cautious about what I run on my computers (and this class has already stretched my comfort zone a bit)
- OpenClaw is highly configurable, but that configuration process is so deep that I'm not confident in my ability to demonstrate its value on my main computers without exceeding my risk threshold

...but I have a Raspberry Pi 5

- If it goes rogue, I just eat the microSD card and forget about it
- This is too lightweight to really run useful models locally (anything beyond a not-great conversational agent)
- OpenClaw itself is rather heavy
- ZeroClaw is a more lightweight (if less extensively supported) alternative



These things are really expensive now

- Because of the RAM shortage, boards I bought for \$65 two years ago are going for more than \$300 now :C
- ...So I'm re-using my bootleg Switch that I made for playing retro games
- Alternatives (but also expensive): Mini-computers / NUCs, Mac Minis, desktops



This is an in-progress project


- Let's briefly look at ZeroClaw and what it can do:
<https://github.com/zeroclaw-labs/zeroclaw>
- Later, during work period, I'll see if I can get it connected to my hotspot so that I can show you a demo, but that may have to wait for next week / it may be tough because we're basically underground

Switching gears: Agent Teams


- We discussed having multiple agents building and reviewing projects around a shared spec. In that simplified process, *you* act as the team lead / dispatcher.
- More advanced uses involve an *Agent* that coordinates the team. Lots of different ways to accomplish this, including subagents and agent teams .

Subagents and Agent Teams

Subagents



Agent Teams



Let's look at Agent Teams


- Different services provide different means of doing this, but we'll look at Claude for now:

<https://code.claude.com/docs/en/agent-teams>

This is not enabled by default

- If you want it for a project, ask Claude to include the proper settings from the documentation. Be aware: Heavy token usage.

```
{  
  "env": {  
    "CLAUDE_CODE_EXPERIMENTAL_AGENT_TEAMS": "1"  
  }  
}
```



Set up the team

Note: I forgot to switch to Sonnet. On Opus, this one request ate 30% of my 5-hour allowed usage on a paid plan.

“

I would like to create a story-driven role playing game similar to Phantasy Star. Create an agent team to explore and develop this. The team should include the following agents:

- An agent focused on UX
- An agent focused on the story
- An agent focused on testing for bugs
- An agent focused on simulating users and giving realistic player feedback

“

Now we ask the team to build for us

With Sonnet this time: “Coordinate the agent team to create a playable demo of the first part of the game. You may work for up to 10 minutes. Keep me updated on the progress, but do not ask for my approval unless absolutely required. You and the team have full design decision authority. Make sure that the game is interesting and fun!”

Result:

Damn

 Claude

 Claude

Can't reach Claude

Check your connection.

Can't reach Claude

 Claude

 Claude

Can't reach Claude

Check your connection.

Try again

Can't reach Claude

Check your connection.

Try again

Damn



Elevated errors on Claude.ai, API, Claude Code

[Subscribe](#)

Identified - The issue has been identified and a fix is being implemented.

Apr 15, 2026 - 15:03 UTC

Update - We are continuing to investigate this issue.

Apr 15, 2026 - 14:55 UTC

Investigating - We are seeing increased errors on Claude.ai, API, and Claude Code

Apr 15, 2026 - 14:53 UTC

Try again

Try again

Pivot...

- Claude seems to be down, and lecture is in two hours, so I'm going to go get some cookies for you all and we'll hope that it's back up by the time lecture happens. If so, we'll do a demo. I'll also try to show ZeroClaw if I can connect to my bootleg Switch.
- One way or another, we'll have some work time after that.