

Workshop 2: AI as Environmental Cost

Exercises and Setup Guide

Setup

Installing the Environmental Impact Tracker

Students will need to manually install the Environmental Impact Tracker Chrome extension from a .zip file. Walk through the following steps before beginning the exercises.

- 1. Download the extension.** Distribute the .zip file to students.
- 2. Extract the .zip file.** Have students unzip the file to a location they can find again. On Mac, double-clicking the .zip will extract it automatically. On Windows, right-click the .zip and select "Extract All." The result should be a single folder containing the extension files.
- 3. Open Chrome's extension manager.** In Chrome, type `chrome://extensions` into the address bar and press Enter. This opens the Extensions page.
- 4. Enable Developer Mode.** In the top-right corner of the Extensions page, toggle the "Developer mode" switch on. This exposes the controls needed to load an unpacked extension.
- 5. Load the extension.** Click the "Load unpacked" button that appears in the top-left after enabling Developer Mode. In the file picker, navigate to the extracted folder and select it. Chrome will install the extension and it should appear in the list on the Extensions page with no errors.
- 6. Verify the extension is active.** The extension's icon should appear in Chrome's toolbar (students may need to click the puzzle-piece icon to pin it). Navigate to chatgpt.com. The Environmental Impact Tracker sidebar should appear automatically on the right side of the screen with three tabs: Tokens, Impact, and Models.

Additional Requirements

Ensure all students have a ChatGPT account. The free tier works. If the sidebar does not appear, have students check that the extension is toggled on in `chrome://extensions` and refresh the ChatGPT page.

Exercise 1: Observing Token Accumulation

Have students chat with ChatGPT normally for a few minutes and note what happens in the Tokens tab. Then have them ask ChatGPT to write a 1,000-word story across three separate prompts. They should notice input tokens growing substantially with each turn as the model re-reads the full conversation history. Then have them ask ChatGPT to search for something and observe the search token count.

After these activities, discuss what the extension can and cannot see. It tracks the student's

prompts, conversation history, visible responses, and search content. It cannot see the system prompt (thousands of hidden instruction tokens sent every turn), tool definitions, thinking tokens (invisible reasoning that costs the same as output), or multi-step tool loops. The sidebar numbers are a lower bound on actual compute.

Exercise 2: Comparing Estimates

Direct students to the Models tab. Have them note their Impact tab totals, then switch between the lower, middle, and higher estimates. The higher estimate produces roughly 3.5 times the impact of the lower, both for single sessions and at scale. Both come from credible published sources. They differ because one is self-reported by Google on optimized hardware and the other is an independent academic benchmark. The gap is not error; it reflects the fact that no company publishes definitive data.

Exercise 3: Finding Your Own Numbers

Have students select "Custom" in the Models tab and spend 10 to 15 minutes researching other estimates for per-token energy, carbon, and water costs. They should enter whatever they find into the extension. Expect wide variation. Some numbers will be per-query and need conversion. Some will cover energy but not water. Use this to reinforce that the inscrutability of the environmental question is tied to the inscrutability of the systems themselves: we cannot measure what we are not allowed to see.

Exercise 4: Why No Disclosure?

Facilitate a discussion: AI companies publish detailed benchmarks for intelligence, reasoning, and safety. Why are there no benchmarks for resource efficiency? Consumers can compare cars on fuel economy and emissions because disclosure is mandated. AI is used by hundreds of millions of people and concern about its footprint is widespread. Why is there no equivalent? What would have to change?

Exercise 5: Designing for Legibility

Close with a speculative exercise. If ChatGPT showed estimated energy cost before each message, would students use it differently? If users could set a daily resource budget, how should the interface behave as they approach the limit? If a model could alter its behavior based on resource constraints (shorter responses, skipping reasoning, avoiding search), what rules would they give it and what tradeoffs would that involve? Have groups share their proposals. The goal is to make visible the design choices currently being made invisibly.