



# What Does a Prompt Really Cost the Planet?

Illustrative electricity & water estimates for three everyday AI tasks — and what they're actually equivalent to in real life.



## ■ ELECTRICITY PER TASK — AT A GLANCE



Bars are scaled for visual comparison, not linear — the actual gap spans roughly 2 orders of magnitude.



## Asking a Quick Question

A short, single-turn text prompt — e.g. "What's the capital of France?" — answered by a typical modern chat model.

01

### ⚡ ELECTRICITY

≈ 0.3 Wh

≈ powering a 10W LED bulb for ~2 minutes

### 💧 WATER

≈ 0.3 mL

≈ 1/15th of a teaspoon of water



## Generating One AI Image

A single standard-quality image (around 1024×1024) from a diffusion-based image generator.

02

### ⚡ ELECTRICITY

≈ 0.01–0.03 kWh

≈ half a smartphone's battery charge

### 💧 WATER

≈ 15–60 mL

≈ a shot glass or two of water

## ■ PUTTING IT IN SCALE

- ⚡ ~35 quick questions use about as much electricity as generating 1 AI image
- ⚡ ~5 AI images use about as much electricity as 1 PDF-to-slide-deck job
- 💧 Doing the PDF → deck task daily for a month uses roughly 2–4 standard water bottles worth of water



## Uploading a PDF & Turning It Into a Slide Deck

Reading a ~10–15 page document, then generating a ~12-slide presentation — a long-context, multi-step, high-output task.

03

### ⚡ ELECTRICITY

≈ 0.05–0.15 kWh

≈ running a microwave for ~5 minutes, or ~8 phone charges

### 💧 WATER

≈ 250–500 mL

≈ half to a full bottle of water

## METHODOLOGY & SOURCES

Task 1 uses OpenAI's disclosed per-query figure and Epoch AI's independent 2025 estimate for GPT-4o-class models. Task 2 uses the Hugging Face & Carnegie Mellon University 2024 benchmarking study on diffusion image models. Task 3 has no direct public study — it's **our own extrapolation** from published long-context energy estimates (Jehham et al., 2025), scaled up for a multi-step, high-output workflow. Water figures apply a standard cooling + electricity-generation conversion factor where direct data doesn't exist. **These numbers vary 10–100× across published studies** depending on the model, hardware, data-center location, and whether only on-site cooling or the full electricity lifecycle is counted — treat every figure here as an order-of-magnitude estimate, not a precise measurement.