

The environmental footprint

small per prompt, large in aggregate

AI uses energy and water. A climate negotiator who uses it should be able to say how much. Your own use is small. Total demand is large, and growing.



MEASURED, PER PROMPT

0.24 Wh
ENERGY

0.03 g CO₂e
CARBON

0.26 mL
WATER

Google's measurement of a typical text prompt to Gemini. It includes idle servers and data-centre overheads, which most estimates leave out. Google, arXiv:2508.15734, August 2025.

01 YOUR OWN USE IS SMALL

One exchange uses about as much as running a laptop for a few seconds.

A full day of heavy use is still less than a short car trip.

Polly shows this figure for every account, and publishes how it works it out below.

02 TOTAL DEMAND IS LARGE

- Data centres are adding electricity demand faster than grids are switching to clean power.
- New data centres are going up in places already short of water.
- The companies that publish the small per-prompt figures have seen their total emissions rise since they started building these systems.

03 WHAT INTELLIGENCE COSTS TO RUN




FIG. 01 Each dot is an open model that Artificial Analysis scores and that says how much of itself it uses per token. Filled dots are the frontier: no other model is both smarter and lighter to run. The blue line follows one model from its lowest to its highest reasoning effort. Scores: Artificial Analysis Intelligence Index, 27 August 2026. Energy: The ML.ENERGY Benchmark v3.0.

04 THE CALCULATION, SHOWN

ENERGY PER TOKEN

2 FLOP per active parameter → divide by chip throughput →
multiply by chip power

WATER AND CARBON

1.1 litres per kWh for cooling →
480 g CO₂e per kWh on an average grid

Reading your prompt costs about a quarter as much as writing the answer. Re-reading a prompt already seen costs a tenth. The three are counted separately. Polly's model uses 18 billion parameters per token. A typical exchange uses about 0.10 Wh and 0.11 mL of water, just below Google's figure above.

05 WHAT THIS LEAVES OUT

Training the model. The water used to make the electricity. Your own device. The carbon from building the hardware.

Take every figure on this sheet as correct to within a factor of two, not as an exact measurement.

06 WHAT TO DO ABOUT IT

- 1 Turn reasoning off when you do not need it. With it on, the same model can use twenty to forty times as many tokens.
- 2 Ask providers for their figures, including water, and which power grid serves your requests. Fair questions for a supplier, and few people ask them.
- 3 Publish your own figure. A delegation that reports the footprint of its own AI use can fairly ask others to do the same.

GainForest and the Youth Negotiators Academy can help a delegation work out and publish its own AI footprint for free, using the method on this sheet.

SOVEREIGNTY-ALIGNED AI FOR MULTILATERAL DIPLOMACY · AI4COP.ORG

team@gainforest.net
ai4cop.org

