



GREEN IO LONDON 2026 · COMMUNITY PITCH

Developers want greener software. Most don't know where to start.

The Green Coding Roadmap

A **scientifically developed**, up-to-date pool of Green Coding resources for practitioners, **free** and **open source** on roadmap.sh.

100% free & open source

2-year ISOC-funded project

4 years since identifying the gap

Environmental Informatics Unit · Prof. Dr. Volker Wohlgemuth · Dennis M. Junger · Nick Kriesing Garcia




GERMAN
INFORMATICS SOCIETY



H O C H
S C H U L E
T R I E R

Problem: knowledge is scattered



WHAT WE MEAN BY GREEN CODING

“the act of designing, developing, maintaining, and (re-)using software systems in a way that requires as little energy and natural resources as possible”

Junger et al. (2024, p. 3)

For learners, this means ...

- no clear starting point
- unclear how concepts connect
- no path from understanding to practice

Other domains have solved it

- visual learning roadmaps, e.g. roadmap.sh
- **for Green Coding: no self-directed learning path, until now**

Our solution: the Green Coding Roadmap

One guided, free path from understanding to practice, on roadmap.sh.

Six chapters, one learning path

01

Fundamentals and Concepts

Why emissions matter, Green IT vs Green Coding

Remember

02

Carbon and Energy

Embodied & operational carbon, carbon intensity

Understand

03

Measurement and Evaluation

Blue Angel, SCI, GSMM, energy profiling tools

Understand

04

Code Level

Algorithms, runtimes, APIs, web frontends and AI

Apply

05

Infrastructure and Operations

Data centers, cloud, carbon-aware scheduling

Apply

06

Process, Governance and Culture


People, standards, regulation, Green SDLC

Apply

+ **Toolkit:** tools and resources by task, glossary A-Z

Reference

Bloom: remember / understand → apply



Excerpt on roadmap.sh > 100 topics

Start learning. Share it. Build on it.

Sustainable software scales through people. Free and open source on roadmap.sh: use it in your team, your course or your next sprint, and tell us what's missing.



Interested in research & collaboration?

greencoding.f2.htw-berlin.de



roadmap.sh/r/green-coding

Thank you!

Dennis M. Junger

Environmental Informatics Unit, **HTW Berlin**

Come by our poster!

We're happy to see you at or around our poster, with QR codes to take away. Questions, feedback, ideas for the roadmap? Let's connect.



 [linkedin.com/in/dennismjunger](https://www.linkedin.com/in/dennismjunger)



On a mission to change the world by making AI more sustainable

Green IO London · 30 September 2026

01

About GreenPT

Who we are, what we build, and what makes it sustainable.

01 About GreenPT

02 Measuring energy

03 CO₂ and location

04 Efficiency

05 Compression

06 New plans

07 New regions

One platform, three products



APPS

Chat for everyone

- Consumers
- Teams
- Enterprises

Web

iOS

Android



API

Full-stack AI provider

- One OpenAI-compatible endpoint
- Text, code, speech, OCR, embeddings
- Search and fetch

`api.greenpt.ai/v1`



ADVISORY

Greener AI, anywhere

- Measure your AI footprint
- Find the waste
- Prove the drop

A full-stack AI provider from Utrecht



Privacy-first

No training on your data, GDPR by default, ISO 27001 infrastructure.

Open models

GLM, Kimi, Qwen, Mistral, DeepSeek: the newest within days.

Transparent

Energy and CO₂ reported on every single request.

Hosted in Europe

Hosted in Europe on 100% renewable energy, with no training on your data.

Sustainable at every layer

From the power in the datacenter to the words the model writes: every layer is chosen to use as little energy as possible.



Renewable hosting

European datacenters on 100% renewable power, with free cooling and a PUE of 1.25.



Right-sized models

The router picks the smallest model that does the job, on a GPU sized to fit.



Fewer tokens

Quantized models, prompt compression and Honey: the same answer with less compute.



Measured

Energy and CO₂ on every request, so the savings are numbers, not claims.

OUR IMPACT

40-50%

less CO₂ emissions

20-30%

less energy

93%

less water

20-30%

lower cost

Trusted by organisations across Europe

Government, insurers, healthcare, education, NGOs and culture: organisations that chose AI with a **smaller footprint**.



40+ organisations on the platform,
from the Netherlands to Canada

02

Measuring energy

From watts on the GPU to milliwatt-hours per 100 tokens.

01 About GreenPT

02 Measuring energy

03 CO₂ and location

04 Efficiency

05 Compression

06 New plans

07 New regions

Every AI answer has a carbon cost. Three things set it.

$$\text{CO}_2 = \text{tokens} \times \text{energy per token} \times \text{grid intensity}$$

what the model reads and writes how efficiently it runs how clean the power is

01

ENERGY

Measure the energy

Watts from the GPU, turned into
mWh per 100 tokens.

02

GRID

CO₂ and location

The live grid at each datacenter,
and the cleanest hour to run.

03

TOKENS

Not every token costs the same

Output written one by one: the
costly part

Input read in a single pass

Cached reused, almost free

04


NEXT

Plans and regions

Plans metered in grams of CO₂,
and new regions.

You can't cut what you don't measure

Our metric works like kWh per 100 km for an electric car: **milliwatt-hours per 100 tokens**. We run the hardware, so we can meter every watt.



One prompt, one GPU: power and temperature climb the moment it starts writing.

01 Measure the GPU

An NVIDIA H100 under load draws about **300 W**.

02 Add the rest of the server

CPUs and the rest add 27%: $300 \text{ W} \times 1.27 = 380 \text{ W}$.

03 Per second

$380 \text{ W} \div 3,600 = 105.6 \text{ mWh}$ for every second of work.

04 Per 100 tokens

We metered a benchmark run: 51,674 tokens in 56.7 s. The GPU serves many requests at once, so it writes ~911 tokens a second.

$105.6 \times 56.7 \div 51,674 \times 100 = 11.59 \text{ mWh}$

A 500-token answer: about **58 mWh**.

05 To grams of CO₂

Times the live grid intensity at the datacenter. The result ships with every response.

THE BENCHMARK RUN • ONE H100 SERVER

51,674
tokens written

56.7 s
wall-clock time

~911
tokens per second

5.99 Wh
for the whole run

03

CO₂ and location

Where and when a prompt runs decides its footprint.

01 About GreenPT

02 Measuring energy

03 CO₂ and location

04 Efficiency


05 Compression

06 New plans

07 New regions

Same prompt, different grid, different CO₂

A kilowatt-hour in Norway is not a kilowatt-hour in Poland. We compute emissions from the **live carbon intensity of the grid at each datacenter**, and send work where the grid is cleanest.



From kilowatt-hours to grams, per location

Energy is what the GPU uses. CO₂ depends on **where** and **when** it used it. So we multiply every request's measured energy by the grid it actually ran on.



Location-aware

Grid intensity for the bidding zone of each datacenter, not a global average.



Realtime and forecast

Live kWh to CO₂ per request, plus a forecast so work can wait for a cleaner hour.



Everywhere

In Chat, Apps and every API response, with no extra setup.

GreenPT Pro

Prompts left: ∞



Last 30 days:



Energy used: 55.10 Wh



CO₂: 1.29 gCO₂e



[Learn about calculations](#)


In Chat, a **live badge** shows the energy and grams of CO₂ you used.

POST /v1/chat/completions → response

```
"impact": {  
  "energy": { "total": 148614, "unit": "Wms" },  
  "emissions": { "total": 2, "unit": "ugCO2e" }  
}
```

Same job, cleaner hour

The carbon intensity of the grid changes every hour with sun, wind and demand. Work that doesn't have to run right now can wait for a cleaner hour.



-62%

CO₂ for the same job, started at 13:00 instead of 19:00. No code change, just timing.

THE SCHEDULER DECIDES

RUN NOW

if this hour is already clean


DEFER

wait for the cleanest hour before the deadline

Good candidates: AI inference that can run at night, like embeddings, reports and evals

Our GPUs follow the Dutch day

Completed requests on our platform over one day, straight from Grafana. Nobody asks AI anything at 04:00, and everyone does at 10:00.



1 🐼 Asleep
00:30 to 08:00

2 🍞 Lunch
bread with hagelslag

3 🍽️ Dinner
and the NOS Journaal

Scheduled resource allocation

Capacity follows the daily rhythm, so GPUs don't sit idle at full size through the night.

Dynamic load mechanisms

When a spike hits, extra replicas start. When it passes, they stop.

~100 W

An idle H100 still draws about 100 W, against about 300 W under load. **Idle capacity is never free.**

04

Efficiency

The right model on the right GPU.

01 About GreenPT

02 Measuring energy

03 CO₂ and location

04 Efficiency

05 Compression

06 New plans

07 New regions

Bigger is not always better

Most questions don't need the biggest model. The router sends each one to the **smallest model** that does the job well.

EVERYDAY QUESTION

"Summarise this email in 3 bullets"



SMART

Router



SMALL MODEL

GLM 5.3 Flash

HARD PROBLEM

"Find the race condition in this service"



SMART

Router



LARGE MODEL

GLM 5.3

01 Small models do most of the work

GLM 5.3 Flash, DeepSeek 4.1 Flash and Mistral Small cover everyday work. GLM 5.3 takes the hard problems.

02 Step up only when needed

The router starts small and moves to a bigger model only when the task asks for it.

03 Quantization

MXFP4 and GGUF shrink the weights: less memory, faster tokens, less energy.

Slice the GPU to fit the model

A small model on its own H100 leaves most of the card idle, and an idle H100 still draws about 100 W. **MIG** splits one GPU into isolated slices.


WITHOUT MIG



Three GPUs, mostly idle

H100 × 3

WITH MIG



One GPU, three models

H100 × 1

The right chip per job

NVIDIA L40 for speech to text, NVIDIA H100 for the large language models.

Never compute twice

vLLM with LMCache reuses the KV cache, so repeated context isn't computed again.

05

Compression

Lossless first, then trimming.

01 About GreenPT

02 Measuring energy

03 CO₂ and location

04 Efficiency

05 Compression

06 New plans

07 New regions

Lossless: labels once, then just the data

When agents talk to agents, nothing needs to be dropped to save tokens. A planner hands a trip plan to a concierge. ESON writes each label once, adds a built-in row count, and **loses nothing**. Only between agents, never in answers for people.

Normal (JSON)

806 characters

```
{
  "from": "planner",
  "to": "concierge",
  "city": "Geneva",
  "activities": [
    {
      "n": 1,
      "name": "Jet d'Eau",
      "area": "Lakefront",
      "hrs": 1,
      "chf": 0
    },
    {
      "n": 2,
      "name": "St-Pierre Cathedral",
      "area": "Old Town",
      "hrs": 2,
      "chf": 7
    }
  ]
}
```

Honey handoff (ESON)

347 characters

```
!eson/1
from=planner
to=concierge
city=Geneva
activities[6]{n,name,area,hrs,chf}
1 Jet d'Eau Lakefront 1 0
2 St-Pierre Cathedral Old Town 2 7
3 CERN Science Gateway Meyrin 3 0
4 Lake cruise (CGN) Lac Leman 2 25
5 Patek Philippe Museum Plainpalais 2 10
6 Bains des Paquis Lakefront 1 2
```

1 activity

activities: 6 · labels written once, so the saving grows

6 activities

-57%

shorter to send
≈ -71% fewer tokens

The greenest token is the one you never generate

A question to an AI chat and its reply. The filler goes, **the important bits stay word for word**. Every word removed is a token nobody pays for.

ORIGINAL

QUESTION • 19 words

Hi there! I was just wondering if you could
please tell me the capital of **Australia?**
Thanks so much!

REPLY • 11 words

Great question! The capital of Australia is
Canberra. Hope that helps!

COMPRESSED

QUESTION • 6 words

tell me the capital of **Australia?**

REPLY • 1 words

Canberra.

full

◀ drag to compress ▶

lean

-77%
words removed

06

New plans

Pay for your impact, not your prompts.

01 About GreenPT

02 Measuring energy

03 CO₂ and location

04 Efficiency

05 Compression

06 New plans

07 New regions

Pay for your impact, not your prompts

Today a plan counts prompts. Our new plans count **grams of CO₂**. A short question on a clean grid uses less of your budget than an essay at peak hour. **Your bill becomes a budget for the planet.**

TODAY

Starter

Perfect for trying out our AI platform. A simple plan to help you get started with clarity.

€4,50 /month

Learn More

Plan Includes:

- 15 prompts per day
- Access to AI Chat
- Upload and analyze documents
- Basic support
- AI Search
- API Access



NEW

Starter

Perfect for trying out our AI platform. A simple plan to help you get started with clarity.

€4,50 /month

Learn More

Plan Includes:

- 20 gram CO₂ per day
- Access to AI Chat
- Upload and analyze documents
- Basic support
- AI Search
- API Access

~~15 prompts per day~~ → 20 g CO₂ per day

Same Starter plan, same price. A different unit.

Efficiency is rewarded

Tighter prompts, smaller models and cleaner hours make your budget last longer.

Measured, not estimated

Every request already reports its grams, so the budget is exact.

Impact you can see

Users learn what AI costs the planet, one answer at a time.

07

New regions

Now in the EU and the US. More to come.

01 About GreenPT

02 Measuring energy

03 CO₂ and location

04 Efficiency

05 Compression

06 New plans

07 New regions

Now in the EU and the US. More to come.

Since 24 September: **same API key**, just change the base URL.



LIVE • DEFAULT

Europe

Our own stack on 100% renewable power. The full model catalogue.

```
api.greenpt.ai/v1  
api.eu.greenpt.ai/v1
```



NEW • 24 SEP

United States

Inference by our partner **Neuralwatt**. Selected chat models and embeddings.

```
api.us.greenpt.ai/v1  
# same key, same CO2 reporting
```



PLANNED

What's next?

We're in London this week. **Who knows...;**)

```
api.???.greenpt.ai/v1  
# talk to us after the session
```



Thank you

On a mission to change the world by making AI more sustainable



Cas Burggraaf


Co-founder & CTO

cas@greenpt.ai


greenpt.com


How to get your company on board with sustainability & responsible tech practices...

Digital
Impact
Pitch!





09.2026














Join at
slido.com
#POS2025




Which category of equipment emits the most Greenhouse Gases?


- Personal Devices
- Network Infrastructure
- Data Centers



ENERGY



NETWORK
INFRASTRUCTURES




DATA CENTERS

MANUFACTURING
RESOURCES



MINING RESOURCES OF THE WORLD





END OF USE




PSYCHOLOGICAL
OBSESSION









LIMIT NUMBER OF CONNECTED OBJECTS



RENEW YOUR EQUIPMENT LESS OFTEN

REDUCE NUMBER OF SCREENS (TV, TABLETS, LAPTOPS ETC...)

REDUCE + LIMIT
QUANTITY
OF MATERIAL



LIMIT THE NUMBER OF SERVERS

ACTIONS !

ACTIONS !

REDUCE + LIMIT
QUANTITY
OF MATERIAL

EXTEND
LIFESPAN OF OUR
EQUIPMENTS


REPAIR +
TAKE CARE OF
OUR
EQUIPMENTS

REUSE
OUR UNUSED
EQUIPMENTS...


OPTIMISE
MATERIAL
DESIGN AND
SOFTWARE

QUESTION
OUR HABITS
& USAGES

SPREAD AWARENESS
AROUND YOU...



Our (unique) conferences...



Launched 2023



Short 60 to
90 min sessions

Simplify to better
understand

In-person, online
or hybrid

20 to 1 000+
participants



Selected Solutions for the planet in 2022 / 2023 /
2024



Our (unique) conferences...

Climate Pitch!

Launched 2020

25k+ Participants


490+ Confe

Digital Impact Pitch!

Launched 2023



CPD
CERTIFIED
The CPD Certification
Service



Planet On Stage

Selected Solutions for the planet in 2022 / 2023 / 2024

change
NOW

3 key pillars to empowerment

Live a **COLLECTIVE** experience

(Head)



awareness

Knowing the **causes**,
urgency, scale

(Heart)



wellbeing

Reconnect to our **emotions**
to trigger movement

(Hand)



action

Individually,
collectively &
professionally

Source: inspired from the Climate Psychologist initiative - "Climate empowerment" framework

Conference flow – 60, 75 or 90mn



Annual Company Event

CSR, sales,
innovation event

More than 1000
participants
in-person or
online or hybrid!



Whole Hybrid Company Annual Event - France - 06/23



Hybrid Employee Monthly Meeting - UK - 01/24



ESCP Master students - 09/25



Company Employee Meeting - France - 11/24



Shop Employees

Universities / First
year students
bachelor + Master

Hackathon

From 20 to 1000+
participants



ESCP Master students - 09/25



Citizen fair - Portugal - 09/25



MBA ESG Students - France - 09/25



Company event - France - 11/23




Institut Français - UK - Cinema set-up - 12/24




Online Workshop International - 11/23



Kellogg's - UK - Annual event 11/23



Workshop - UK - 2024



Auchan - International CSR Event - Live Studio



Teachers'summit - Portugal - 08/24





... In more than 17 countries...
182+ Pitchers





Masterclass: Communicating with Impact !



1× 3 hours -
Online

Part 1

Let your true Self

- shine** Find your authentic voice on stage [*theatrical improvisation techniques*]
- Let go & manage stress [*breathing, visualisation, anchoring*]
 - Increase focus [*through preparation and stage presence*]
 - Personalise your speech with emotion and humour

Part 2

Connect with your

Audience

- Role plays with personalised debrief and coaching
- Audience analysis and preparation strategies
 - Develop agility and resilience on stage
 - Handle difficult questions and challenging attitudes

They trust us...

DEPLOYMENT AT SCALE (*national, shop employees, by department or roles including Exec Committee, Board of Directors, Top Leaders' sessions*)



OTHER TYPES OF SESSIONS CARRIED OUT (*e.g. Board Management / Board members / International clients Innovation Days, Learn & Lunch sessions, Universities...*):



Keep in touch !

Have a request ?
Want to talk?



www.planetonstage.org/en/contact

Experience a Pitch
live or join our
Training !



www.planetonstage.org/en/register

Contact me !



sandra.sydow75@gmail.com

Tel: +44 (0) 7557263678
(best via WhatsApp)