Sun Shines Bright Vol.1

by computational mama

Experiments in running solar powered servers for Small Language Models

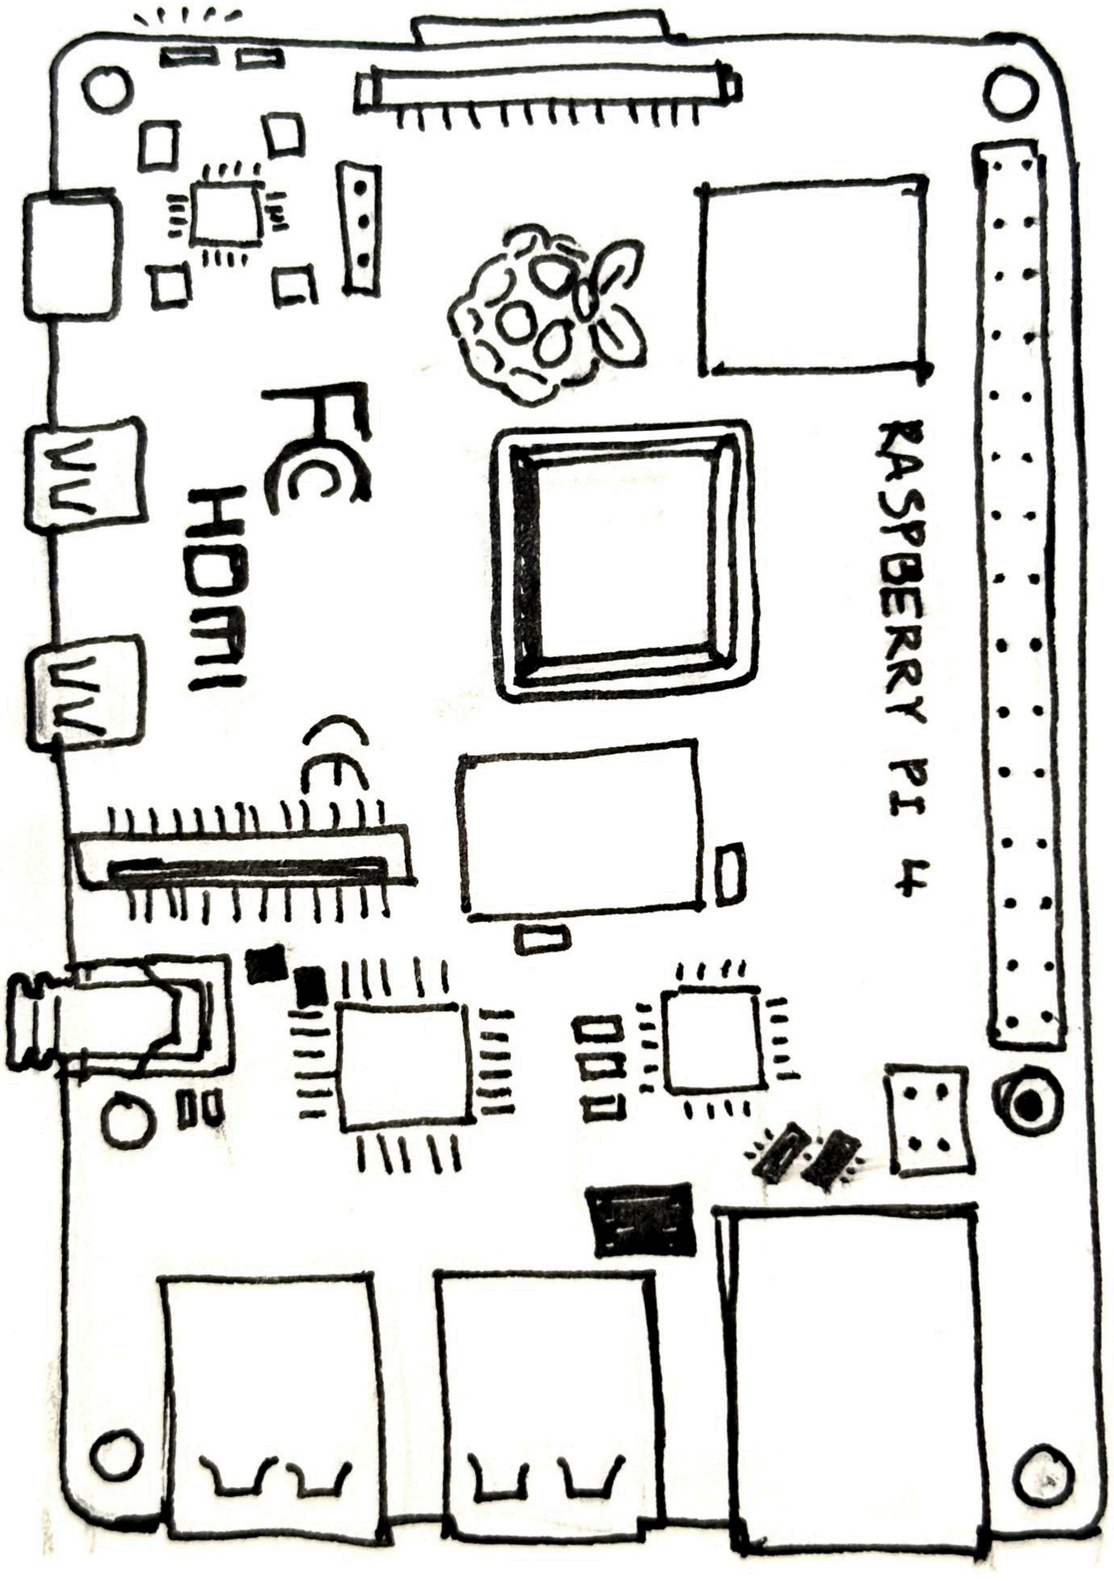
INTRO

Since January 2025, I've been slowly and patiently building a small solar powered server. The project has been humbling. It's made me realize how much energy and resources we waste to host simple websites, and AI services.

In Oct/Nov 2024, Rhizome.org held an info session where a speaker explained a long term project called Solar Protocol. They call it "A naturally intelligent network". By hosting simple websites across a network of solar-powered Raspberry Pis they raise not only some important questions about energy consumption, they allow for conversations around community built networks, addressing:

- our dependence on big tech for small things,
- a gradual decline of the tinkerability of the internet
- our distraction from meaningful content and conversation to vapid details about data and speed

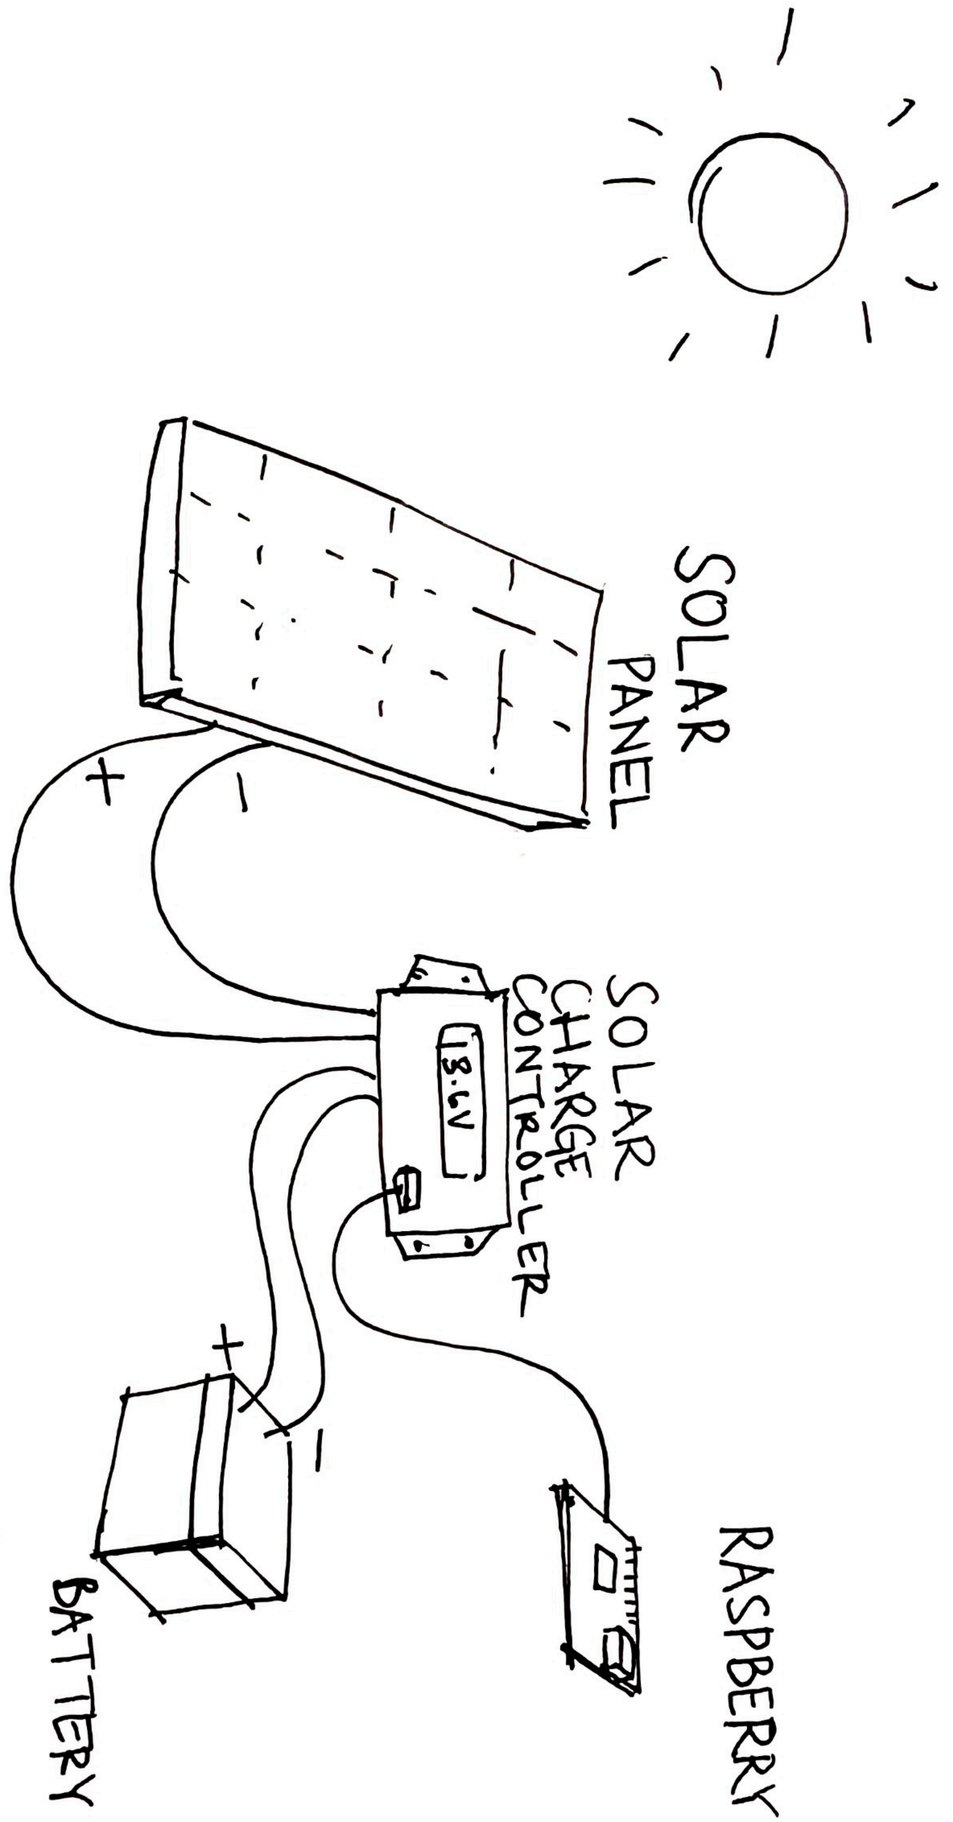
As I work in the generative AI sector, I'm always thinking about the energy waste that surrounds my work, particularly stuff that can be trivial. To offset the energy consumption of my own projects, I've been steadily building out "Sun Shines Bright".



WHAT YOU NEED?

HARDWARE
- Raspberry Pi 4
- 32 GB SD card
- Solar Panel 55W
- Solar Charge Controller
- Lead Acid Battery
- Wifi/LAN

SOFTWARE
- Raspbian OS (headless)
- Docker
- Ollama with Tinyllama and Gemma 3 (1B model)
- OpenWebUI



WHY?

- Udaipur (my home) has 200+ clear days in the year - a good place to start!
- Can it help expand my AI experiments in a guilt-free way?
- If I learn this, can I teach it to others? Can self-hosting be more accessible to non-technologists?

With this project, I'm engaging in a few questions, perhaps they are inspiring for you too:
1. How can we self host AI language models with low/renewable resources?
2. How can we push the boundaries of small computing devices? What does it mean? Why is it needed?
3. How can we consider our energy usage as indie tech art explorers?
4. Why not small language models?
5. What is the infrastructure required to host your own AI models without any focus on scaling?

As indie tinkerers, I don't know why the onus is always on us to experiment, test and question! ;)
Urging software architects and system thinkers to tinker with us and build better AI futures.
There is enough evidence now that smaller models, new ARM architecture are more efficient and are good at several tasks. Many of us only need those! After all, we are just using AI for recipes, holidays and heart to heart conversations! ;)

I'm trying to document the whole process, all hosted on the solar server that I've built. Read more here: https://solar.cmama.xyz

Next steps

- Running SmolLM with RAG
- Test fine-tuning of Tinyllama to run Measuring Silences (or similar) without RAG
- More barebones approaches to chat UIs

Resources

https://solar.lowtechmagazine.com/
https://solarprotocol.net/
https://openwebui.com/
https://www.ollama.com/

made by
computational mama
on 20/04/2025