SysBio group meeting : journal club

"Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning" - Henderson *et al.* 2020

Nicolas Captier

June 21, 2022

ML energy demands are growing exponentially

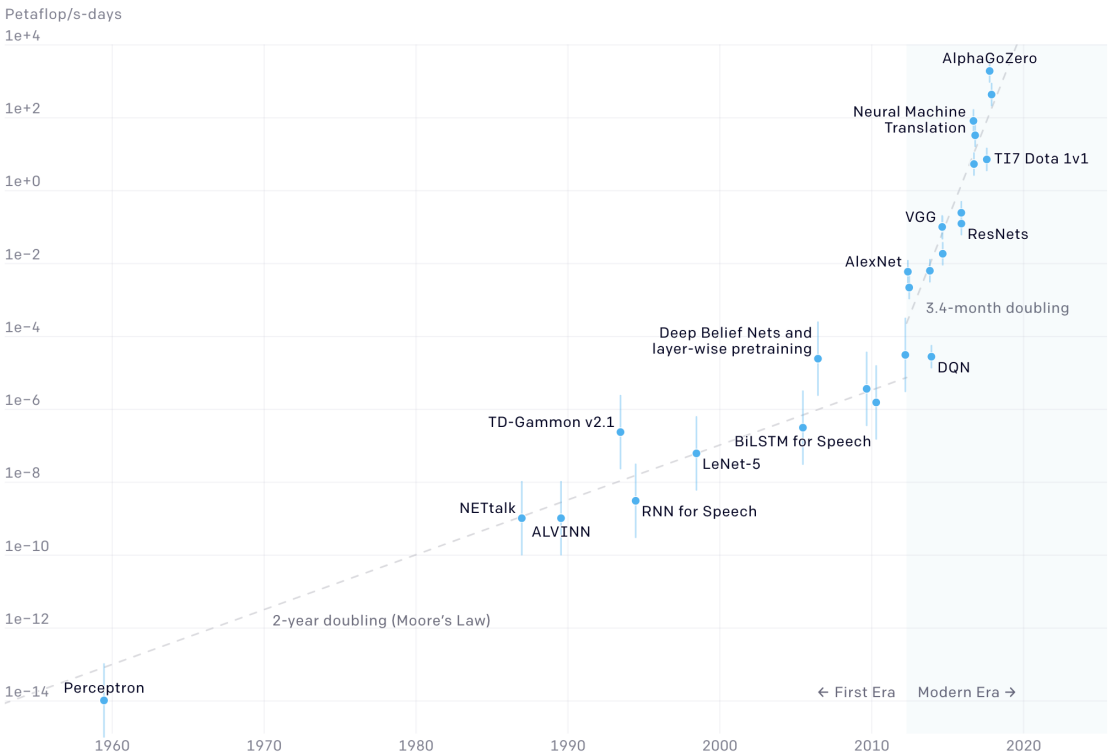

Consumption	CO_2e (lbs)
Air travel, 1 passenger, NY↔SF	1984
Human life, avg, 1 year	11,023
American life, avg, 1 year	36,156
Car, avg incl. fuel, 1 lifetime	126,000

Training one model (GPU)	
NLP pipeline (parsing, SRL)	39
w/ tuning & experimentation	78,468
Transformer (big)	192
w/ neural architecture search	626,155

Figure 1: (Left) Number of floating point operations required to train an ML model[1] - (Right) Estimated CO2 emissions (in lbs) for training common NLP models[2]

[1] "AI and compute" - Amodei *et al.* 2018. Blog post

[2] "Energy and Policy Considerations for Deep Learning in NLP" - Strubell *et al.* 2020

Drive eco-friendly behaviors with standardized and systematic carbon footprint reports ?

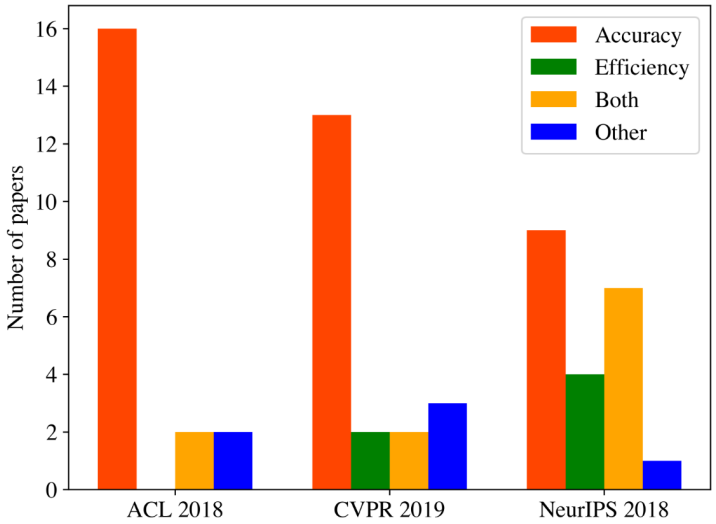

Figure 2: AI papers tended to target performance rather than efficiency ("Green AI" - Schwartz et al. 2019)

Systematically report for energy consumption may:

- raise awareness about ML energy consumption
- drive mitigation efforts and efficiency - accuracy trade-offs
- help identifying targets to reduce energy costs

Carbon Impact Statement

This work contributed 8.021 kg of CO_{2eq} to the atmosphere and used 24.344 kWh of electricity, having a USA-specific social cost of carbon of \$0.38 (\$0.00, \$0.95). Carbon accounting information located at: `https://breakend.github.io/ClimateChangeFromMachineLearningResearch/measuring_and_mitigating_energy_and_carbon_footprints_in_machine_learning/` and

A Python software for automatic CO2 reporting

- **experiment-impact-tracker**[3] is a Python package compatible with Linux and Mac OS X systems running NVIDIA GPU's and Intel processors.
- With a few lines of code added to your main script it will launch a separate python process that will gather compute/energy/carbon information in the background.
- It can generate a carbon impact statement as well as detailed html reports.

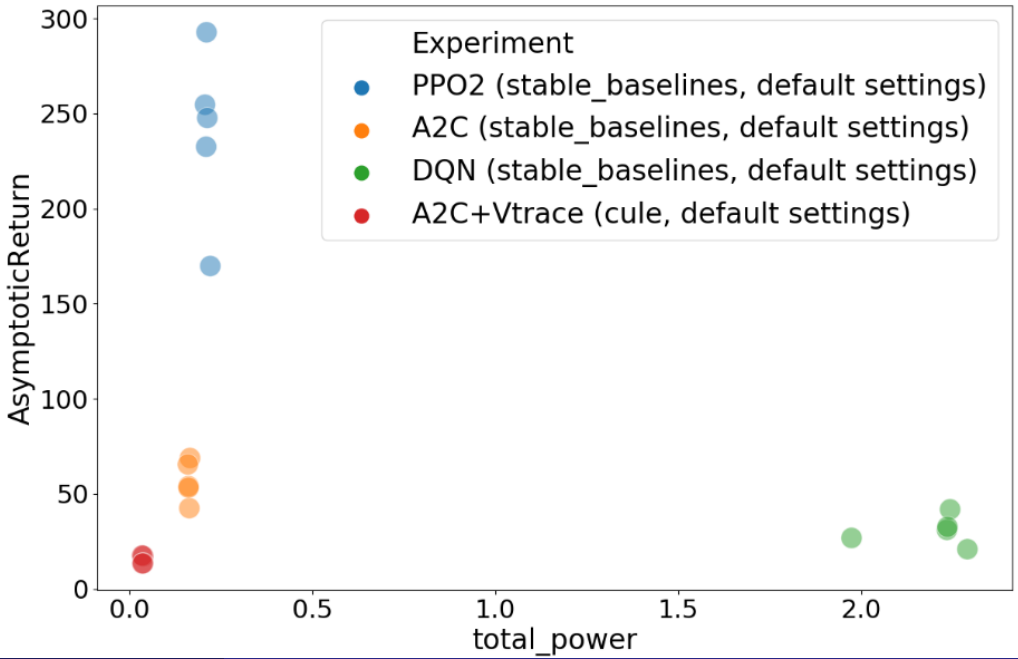

A few details about energy consumption estimation - I

The estimation is restricted to consumption directly linked to the experiment. The carbon footprint associated with the life-cycle of each computing component (e.g. manufacture of GPUs and CPUs) is not taken into account.

Power Usage Effectiveness (rescaling factor)

Percentage in use for each resource

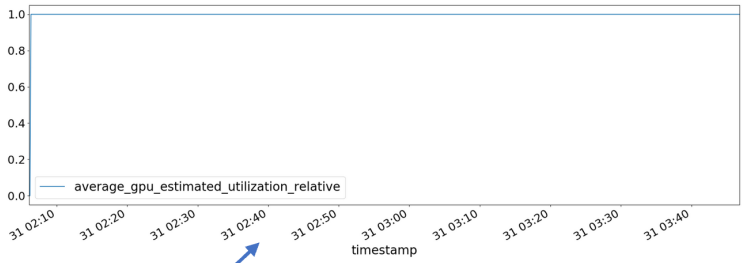

$$e_{\text{total}} = \text{PUE} \sum_{p} (\text{p}_{\text{dram}} e_{\text{dram}} + \text{p}_{\text{cpu}} e_{\text{cpu}} + \text{p}_{\text{gpu}} e_{\text{gpu}})$$

Energy usage of each resource

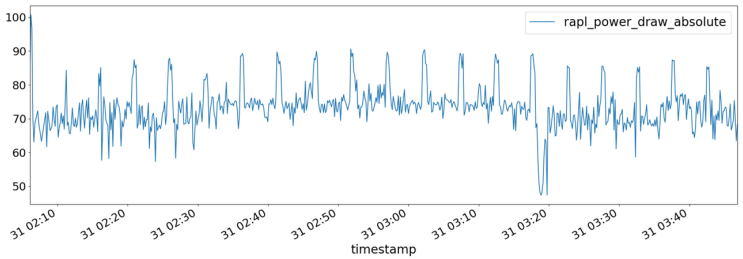

A few details about energy consumption estimation - II

The estimated energy is then multiplied by the carbon intensity of the local energy grid (g CO2eq/ kWh)[a]. The region the experiment is conducted in is estimated with the machine's IP address.

[a] https://electricitymap.org/

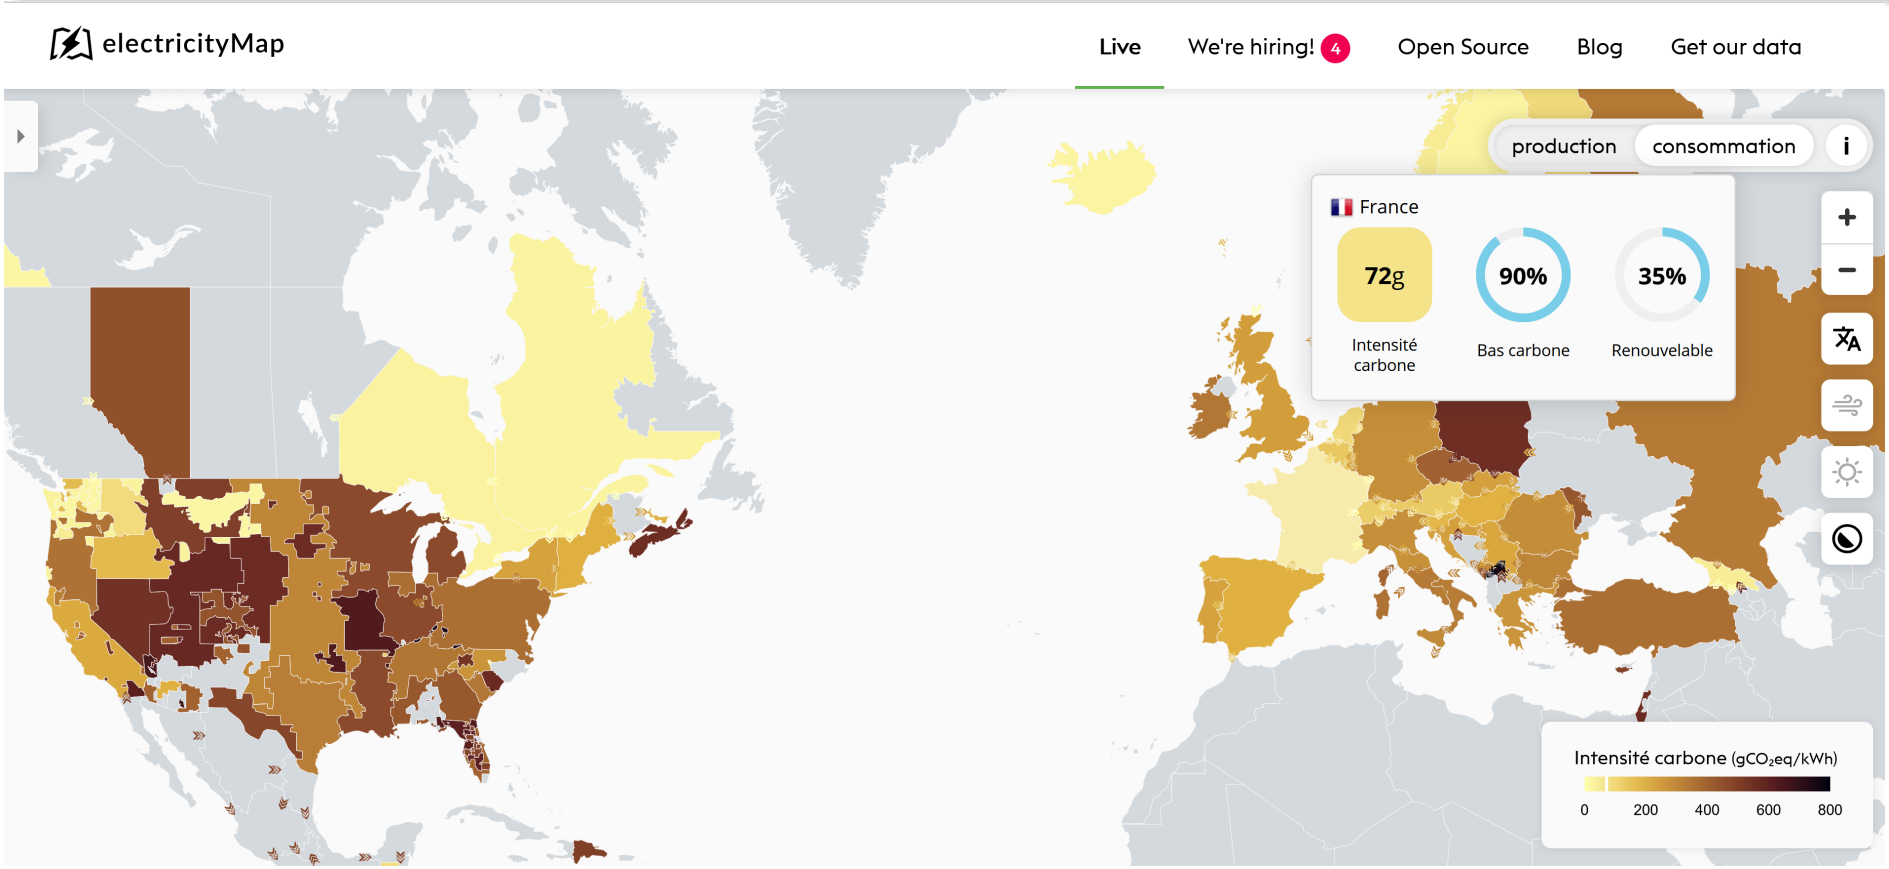

Encourage systemic changes through awareness ?

The authors believe this kind of reporting should not limit energy-intensive research as it could lead to important progresses. They rather believe these systematic measurements could bring the light on the energy impact of ML methods:

- **highlighting energy-efficient tools** and methods and helping users to select the ones which will give them a **satisfying trade-off between performance and efficiency**
- encouraging the **development of new energy-efficient methods** as well as their **integration in common/popular ML platforms** and tools (e.g. Nvidia mixed-precision computing into Pytorch).
- urging ML companies and laboratories to **move their training jobs to low carbon cloud regions** (when it is possible).