



HAL
open science

AI and the Environment: a Symbiosis?

Victor Charpenay

► **To cite this version:**

Victor Charpenay. AI and the Environment: a Symbiosis?. Bootcamp IA & Data, Feb 2024, Paris, France. hal-04497968

HAL Id: hal-04497968

<https://hal.science/hal-04497968v1>

Submitted on 11 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

AI and the Environment: a Symbiosis?

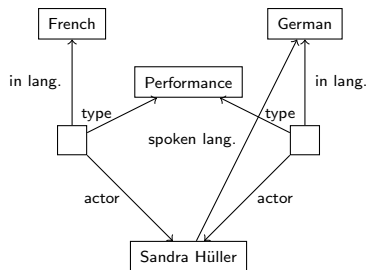
Victor Charpenay

Mines Saint-Etienne

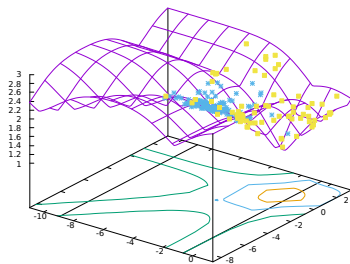
March 1st, 2024

Who Am I? I

I do research in the field of **Knowledge Representation and Reasoning**.



(a) KG



(b) KG embeddings

Figure: Knowledge Graph (KG) representation with learnt embeddings

Who Am I? II




Figure: KG embedding performances vs. parameter efficiency

Who Am I? III

I'm also working on **lifecycle assessment** of AI-based (cyber-physical) systems.

This is why I am also a contact person for **sustainable AI** at IMT.

Why Symbiosis?

IMT has recently joined Enfield, a European excellence network.

Enfield partners are currently preparing a **large-scale review on Green AI**. Its angle is “the state of symbiosis between AI and the environment.”



Enfield: European Lighthouse to Manifest Trustworthy and Green AI, <https://www.enfield-project.eu/>.

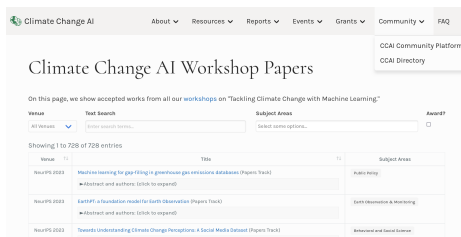
Tackling Climate Change with AI

The AI research community actively seeks ways to leverage AI in **climate change mitigation and adaptation**.

2019 creation of `climatechange.ai`

2020 publication of an extensive review with 898 references

2024 728 papers published on `climatechange.ai`



The screenshot shows the website for Climate Change AI. The navigation bar includes links for About, Resources, Reports, Events, Grants, Community, and FAQ. The main heading is "Climate Change AI Workshop Papers". Below this, there is a search section with a "Text Search" input field and a "Subject Areas" dropdown menu. The search results show a list of papers with columns for Venue, Title, and Subject Area. The first three papers listed are:

Venue	Title	Subject Area
NeurIPS 2023	Yearline learning for gap filling in greenhouse gas emissions databases (Papers Track) ► Abstract and authors (click to expand)	Public Policy
NeurIPS 2023	EarthPI: a foundation model for Earth Observation (Papers Track) ► Abstract and authors (click to expand)	Earth Observation & Monitoring
NeurIPS 2023	Towards Understanding Climate Change Perceptions: A Social Media Dataset (Papers Track)	Behavioral and Social Science

D. Rolnick *et al.* *Tackling Climate Change with Machine Learning*, ACM Computing Surveys, ACM, 2022.

AI for Energy Systems I

HVAC [Heating, ventilation and air conditioning] systems account for more than half of the energy consumed in buildings (...). For control, researchers used deep RL to achieve a scalable 20% reduction of energy while requiring only three sensors: air temperature, water temperature, and energy use.

AI for Energy Systems II

Occupancy detection itself represents an opportunity for ML algorithms, ranging from decision trees to deep neural networks that take input from occupancy sensors, WiFi signals, or appliance power consumption data

AI for Energy Systems III

Modeling energy use across buildings: ML can be used to help predict energy consumption from features such as the location, geometries, and various other attributes of interest like building footprint, usage, material, roof type, immediate surroundings, and the like.

AI for Freight Routing

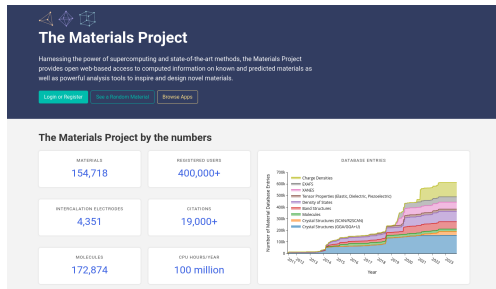
Freight routing and consolidation: Many problem settings are addressed with methods from the field of operations research. There is evidence that ML can improve upon these methods, in particular mixed-integer linear programming.

AI for Ecosystems I

ML can help infer biodiversity counts from image-based sensors. For instance, camera traps take photos automatically whenever a motion sensor is activated—computer vision can be used to classify the species that pass by, supporting a real-time, less labor-intensive species count. It is also possible to use aerial imagery to estimate the size of large herds or count birds.

AI for Ecosystems II

Using semi-supervised generative models and concrete compression data, for example, researchers proposed novel, low-emission concrete formulas that could satisfy desired structural characteristics



A. Jain *et al.* Commentary: *The Materials Project: A materials genome approach to accelerating materials innovation*, APL Materials, 2013.

What About Finance?

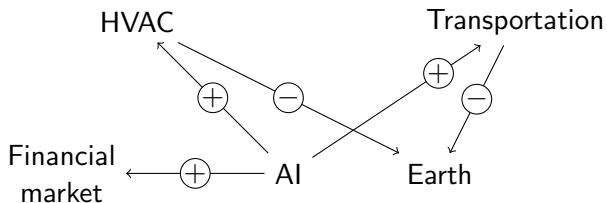
To date, the field of climate finance has been largely neglected within the larger scope of financial research and analysis.

What About Circular Economy?

Circular economy isn't mentioned in the paper

... despite 898 references.

System Symbiosis v1



The Impact of AI

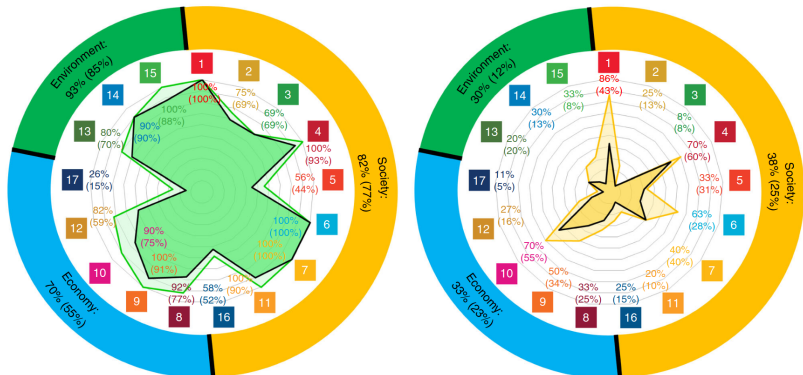


Figure: Summary of positive and negative impact of AI on the various Sustainable Development Goals (SDGs).

R. Vinuesa et al. *The Role of Artificial Intelligence in Achieving the Sustainable Development Goals*, Nature Communications, 2020.

The Negative Impact of AI

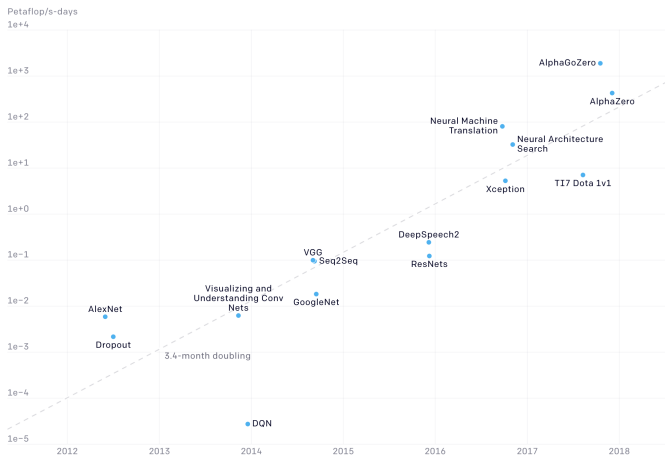
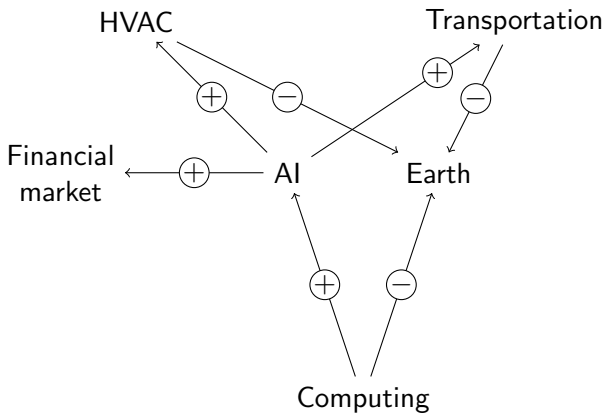


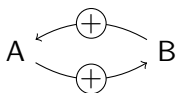
Figure: AlexNet to AlphaGo Zero: 300,000x increase in compute

System Symbiosis v2

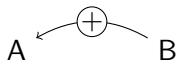


On Symbiotic Relations

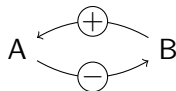
A symbiosis, in its modern definition, isn't necessarily mutually beneficial. It is any **long-term interaction** between species.



(a) Mutualism



(b) Commensalism



(c) Parasitism

Figure: Interactions between species

Green AI: Accuracy vs. Efficiency

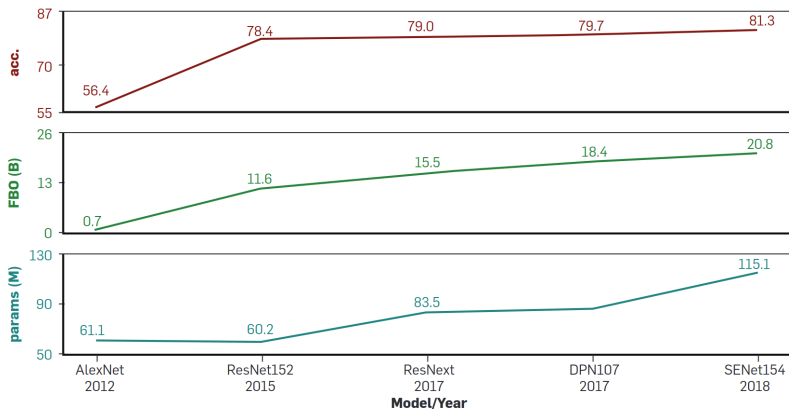


Figure: Model parameters (in million), floating-point operations (in billions), top-1 accuracy on ImageNet

R. Schwartz *et al.* *Green AI*, Communications of the ACM, 2020.

The Carbon Impact of Compute

Consumption	CO₂e (t)
Air travel (NY ↔ SF, one passenger)	0.9
Human life average (1 year)	5
American life average (1 year)	16.4
Car average (incl. fuel, entire lifetime)	57
Semantic Role Labeling (SRL) pipeline	0.017
with tuning and experimentation	35.6
Transformer	0.087
with neural architecture search	284

Table: Estimated CO₂ emissions from training common NLP models, compared to familiar consumption.

E. Strubell, A. Ganesh, A. McCallum. *Energy and Policy Considerations for Deep Learning in NLP*, Proc. of 57th Annual Meeting of ACL, 2020.

The *Real* Carbon Impact of Compute?

Their estimate was 284 tCO₂e for Neural Architecture Search (NAS); the actual number was only 3.2 tCO₂e, a factor of 88 smaller.

Google stated that training was performed on:

- ▶ TPUs (not GPUs)
- ▶ on a small proxy task to search for the best models

They also stated that their data centers were much more energy-efficient than average.

D. Patterson et al. *The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink*, Computer, 2022.

Follow the Renewables

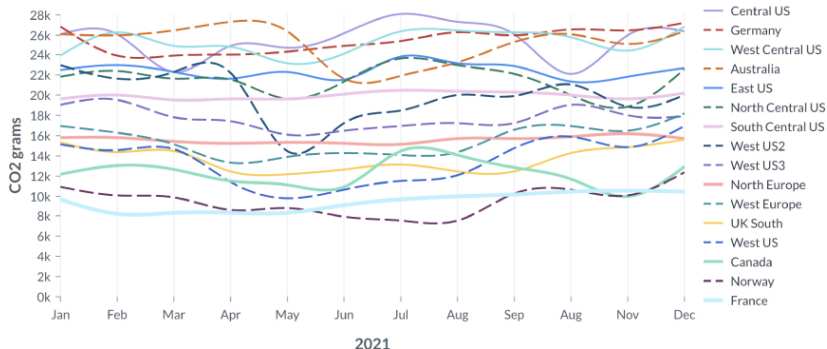


Figure: Carbon emissions that would be emitted from training BERT (language modeling on 8 V100s for 36 hours) in 16 different regions at different times throughout the year.

J. Dodge et al. *Measuring the Carbon Intensity of AI in Cloud Instances*, Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), 2022.

... or Use Less Compute

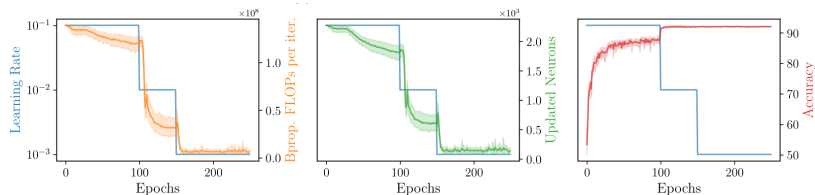


Figure: Back-propagation FLOPs, updated neurons and classification accuracy for ResNet-32 trained on CIFAR-10.

A. Bragagnolo, E. Tartaglione, M. Grangetto. *To Update or not to Update? Neurons at Equilibrium in Deep Models*, Advances in Neural Information Processing Systems 35 (NeurIPS 2022), 2022.

Model Compression

Method	Compression	Accuracy	Message Size (MB)
ZeroFL	Full model	80.62 ± 0.72	44.7
	90 % SP + 0.2 Mask Ratio	81.04 ± 0.28	27.3
	90 % SP + 0.0 Mask Ratio	73.87 ± 0.50	10.1
	Full model	84.43 ± 0.36	44.7
Global magnitude (Ours)	10 % pruning rate	85.96 ± 0.37	38.1
	20 % pruning rate	85.57 ± 0.19	34.8
	30 % pruning rate	85.03 ± 0.32	31.1
	40 % pruning rate	85.20 ± 0.20	27.1
	50 % pruning rate	83.85 ± 0.65	23.0
	60 % pruning rate	83.19 ± 0.44	18.9
	70 % pruning rate	82.25 ± 0.63	14.5
	80 % pruning rate	80.70 ± 0.24	9.8
	90 % pruning rate	76.77 ± 0.47	4.9
	95 % pruning rate	69.14 ± 0.85	2.5
99 % pruning rate	0.10 ± 0.0	0.5	

Table: Comparison to ZeroFL

L. Grativol et al. *Federated learning compression designed for lightweight communications*, Proc. of ICECS 2023, IEEE, 2023.

So, the *Real* Carbon Impact of Compute?

If one combines:





- ▶ **renewables**-powered data centers,
- ▶ **efficiency** gains on learnt models and
- ▶ **sobriety** in model effectively deployed,

carbon footprint estimates may grow unrealistic.

The only way to quantify the impact of computing on the environment is to **publish consumption data**, along with the methodology.

Beyond Training (Large) Models

Direct environmental impacts AI compute resources lifecycle

Production 	Transport 	Operations 	End-of-life 
<ul style="list-style-type: none">Raw material extractionAssemblyManufacturing	<ul style="list-style-type: none">DistributionFreight transportationHandling & storage	<ul style="list-style-type: none">Energy consumptionWater consumptionCarbon footprint	<ul style="list-style-type: none">Collection & shippingDismantling & recyclingWaste disposal

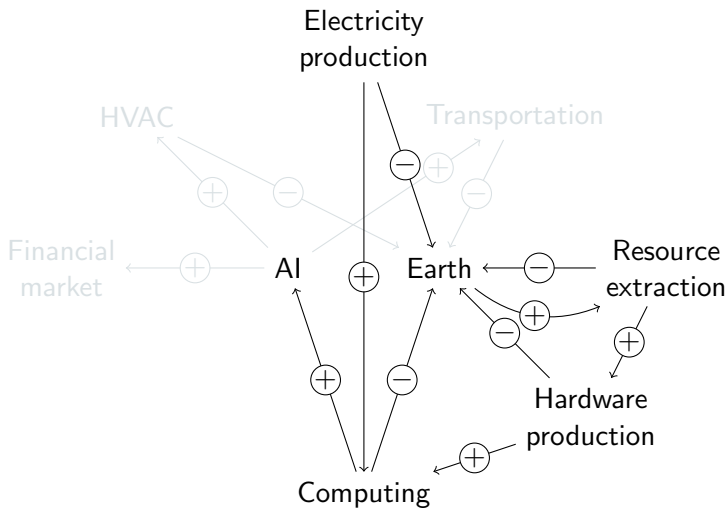
Indirect environmental impacts AI applications

Positive impacts	Negative impacts
<ul style="list-style-type: none">Beneficial sectoral applicationsClimate mitigation and adaptationEnvironmental modelling and forecasting	<ul style="list-style-type: none">Harmful sectoral applicationsCarbon leakage (net increase in emissions)Consumption patterns and rebound effects

Figure: Direct and indirect environmental impacts of AI compute and applications

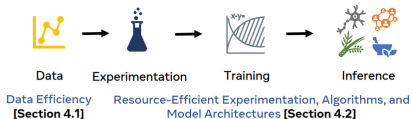
OECD. *Measuring the Environmental Impacts of Artificial Intelligence Compute and Applications*, OECD Digital Economy Papers, 2022.

System Symbiosis v3

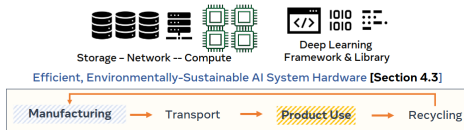


Inference is Costly

Machine Learning Model Development and Deployment Phases [Section 2.1]



System Life Cycle [Section 2.2]

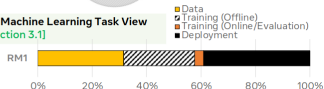


(a) Fleet View



(b) Machine Learning Task View

[Section 3.1]



(c) Infrastructure View



C.J. Wu et al. *Sustainable AI: Environmental Implications, Challenges and Opportunities*, Proc. of MLSys 2022, 2022.

Lack of Information

Data on the environmental impacts of AI compute is not widely available in a standardised and validated form.

This data limitation is particularly acute for measurements of AI compute water consumption and full lifecycle impacts, as these are currently underexplored and underreported.

OECD. *Measuring the Environmental Impacts of Artificial Intelligence Compute and Applications*, OECD Digital Economy Papers, 2022.

Lifecycle Assessment of AI Services

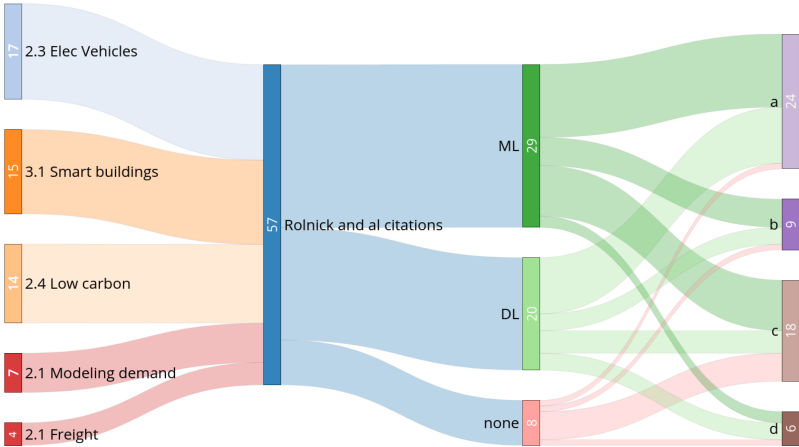



Figure: Sankey diagram of parts of Rolnick's paper references in terms of environmental evaluation

A.L. Ligozat et al. *Unraveling the Hidden Environmental Impacts of AI Solutions for Environment Life Cycle Assessment of AI Solutions*, Sustainability, 2022.

System Symbiosis v4



AI and the Edge (of the Internet)

Facial emotion recognition (with CNNs) can be $47.42\times$ to $65.39\times$ more energy-efficient if one replaces RGB cameras with event cameras.

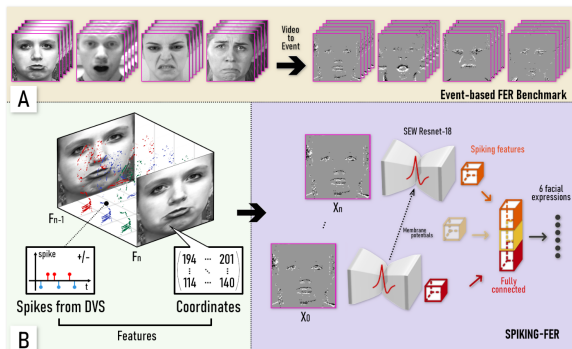


Figure: Overview of the proposed contributions to facial emotion recognition using event-based cameras.

Is AI Yet Another Information Technology?

AI systems highly depend on the availability of the Web, of the Internet of Things (IoT), on the widespread adoption of phones with camera, etc.

Is there anything special about AI?

Environmental Impact of IT I

The carbon footprint of IT represents $\sim 2\%$ of all human activities.

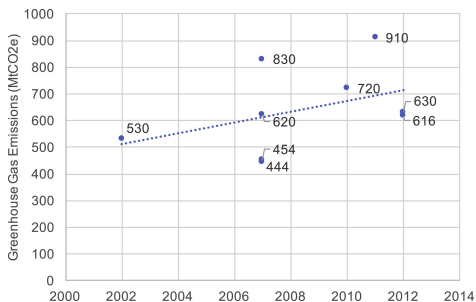


Figure: Estimates of ICT's carbon footprint from studies published before 2015

C. Freitag *et al.* *The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations*, Patterns, Cell Press, 2021.

Environnemental Impact of IT II



La sobriété numérique numérique

40 pratiques accessibles
pour les PME et ETI

www.mines-stetienne.fr

Julien De Benedittis
Nadine Dubruc
Michelle Mongu
Sophie Peillon

Octobre 2023

Jevons Paradox

Improving energy efficiency of a system does not always lead to a decrease in energy use. In fact, it is common to observe a rebound effect: a net increase in energy use.

- ▶ Improved steam engine → higher use of coal
- ▶ Cheaper fuel → more trips
- ▶ Teleconferences → more business trips
- ▶ Cheaper electronic devices → more devices per household
- ▶ ...
- ▶ home virtual assistant → +13.5% energy consumption

K.-j. Chen *et al.* *Influence of Rebound Effect on Energy Saving in Smart Homes*, Cross-Cultural Design: Applications in Cultural Heritage, Creativity and Social Development (CCD 2018), 2018.

Rebounds in XG Network Communications?

In future generations of communication networks, AI may both:

- ▶ improve the energy efficiency of communications (e.g. via multi-armed bandits) and
- ▶ help process higher amounts of data being exchanged over networks.

H. El Hassani, A. Savard, E. Veronica Belmega. *Energy-Efficient 1-Bit Feedback NOMA in Wireless Networks with No CSIT/CDIT*, IEEE Statistical Signal Processing Workshop, 2021.

No AI?

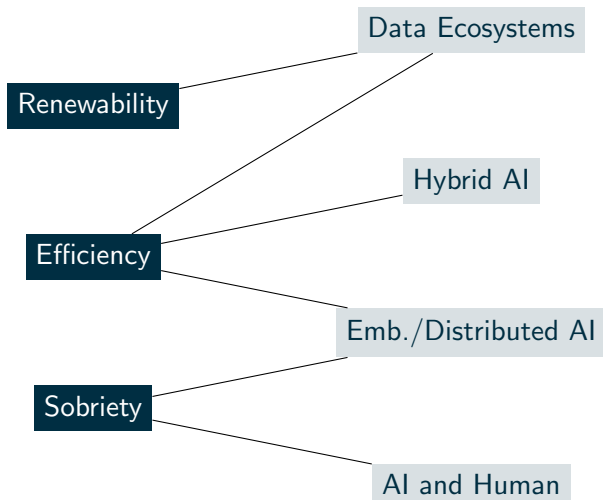
Collecting detailed temperature data in a building helped change HVAC settings and reduce overall consumption by 35-40%.



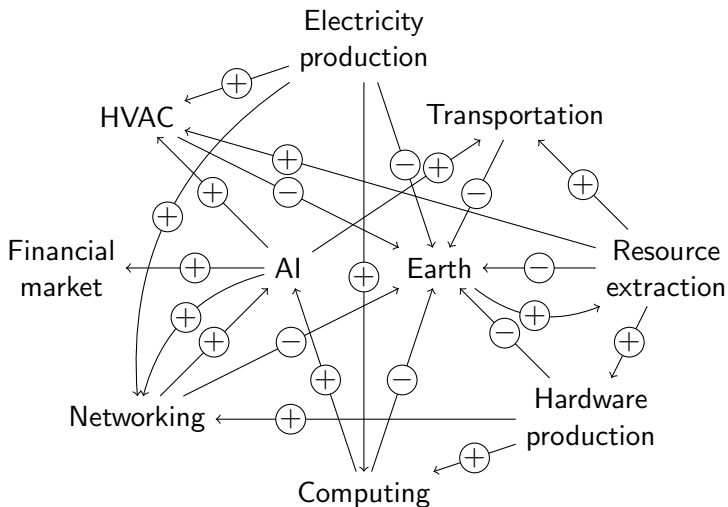
Figure: Temperature measurements in office buildings in Saint-Étienne (january-february 2023)

I. Fatokun *et al.* *Modular Knowledge integration for Smart Building Digital Twins*, Linked Data in Architecture and Construction Workshop, CEUR-WS.org, 2023.

Sustainable AI @ IMT



System Symbiosis v6



Conclusion

Advanced techniques exist both to **reduce the environmental impact of engineered systems** through AI and to **improve AI efficiency**.

Rebound effects may occur because of efficiency gains, leading to net increase in FLOPs executed by AI systems.

Potential benefits of AI on the environment have yet to be **fully quantified**.

Until they are, **classical AI tools** may also help reduce our environmental impact with **few FLOPs**.