Towards a more complete environmental evaluation of an AI program

JRAF 2023

Clément Morand & Anne-Laure Ligozat & Aurélie Névéol

December 13 2023

Laboratoire Interdisciplinaire des Sciences du Numérique

Information and Communication Technologies (ICT) have significant impacts

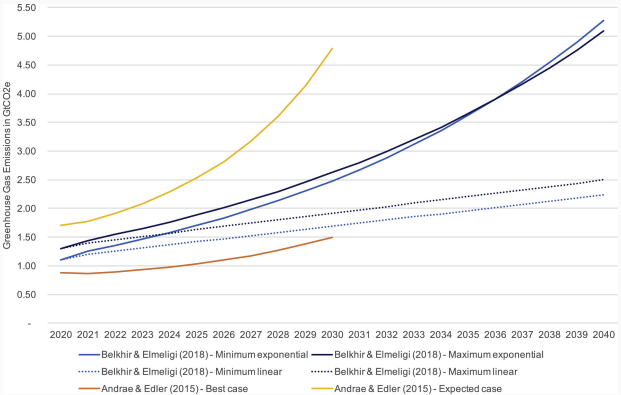

[Freitag et al., 2021]

- $\sim$ 4% of global electricity consumption in 2020.
- 1.4% of global GHG emissions in 2020.
- 5% increase in emissions since 2015.

[Malmodin et al., 2023]

The case of Machine Learning methods

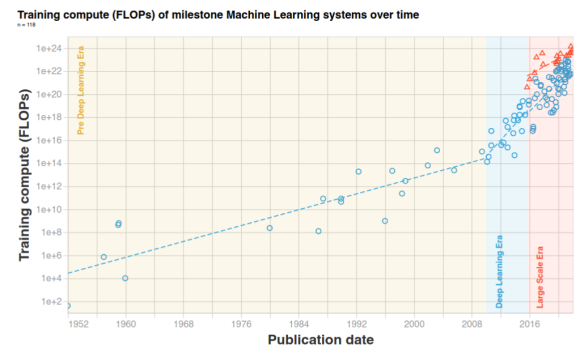

[Sevilla et al., 2022]

- 10 - 15 % of Google's energy consumption [Patterson et al., 2022]
- Important emissions from energy consumption : 552 tCO_2e to train GPT-3 and 38 tCO_2e for BLOOM [Luccioni et al., 2023]

State of the Art

Carbon Footprint of Machine Learning methods

[Strubell et al., 2019]
Importance of
Deep Learning impacts

[Dinarelli et al., 2022]
[Parcollet and Ravanelli, 2021]
evaluate models on efficiency
not only accuracy

[Schwartz et al., 2020]
Need for a 'Green' AI

[Verdecchia et al., 2023]
Review of Green AI

Surveys
[Bannour et al., 2021]
[Jay et al., 2023]

Measuring tools
CarbonTracker [Anthony et al., 2020]
CodeCarbon [Schmidt et al., 2022]
Experiment-Impact-Tracker [Henderson et al., 2020]

Estimation tools
Green Algorithms [Lannelongue et al., 2021]
ML CO_2 Impact [Lacoste et al., 2019]

Life Cycle Assessment for Machine Learning methods

Standards
ISO standards 14040 and 14044
The reference handbook Life Cycle Assessment (LCA):
[European Commission et al., 2010]

Attributional vs. Consequential LCA

[1]

Adapting LCA to ML
[Ligozat et al., 2021]

Carbon footprint of Bloom
[Luccioni et al., 2023]

Importance of Embodied emissions
[Luccioni et al., 2023]
[Gupta et al., 2022]
[Wu et al., 2022]

Importance of Inference phase
[Luccioni et al., 2023]
[Wu et al., 2022]

LCA of Dell servers:
[Thinkstep, 2019]
[Sphera, 2021]

Environmental performance of datacenters
[Gröger et al., 2021]
[Schödwell et al., 2018]

Simplified LCA of servers
[Boavizta, 2021]

[1] source: Weidema BP: Market Information in Life Cycle Assessment, Environmental Project no. 863. Copenhagen: Danish Environmental Protection Agency, (2003)

Defining and implementing an estimation tool for the environmental impacts of computation

An estimation tool that evaluates the environmental impacts of a computation

- "What are the impacts of running program X on the hardware Y during Z hours"

Life cycle phases of hardware

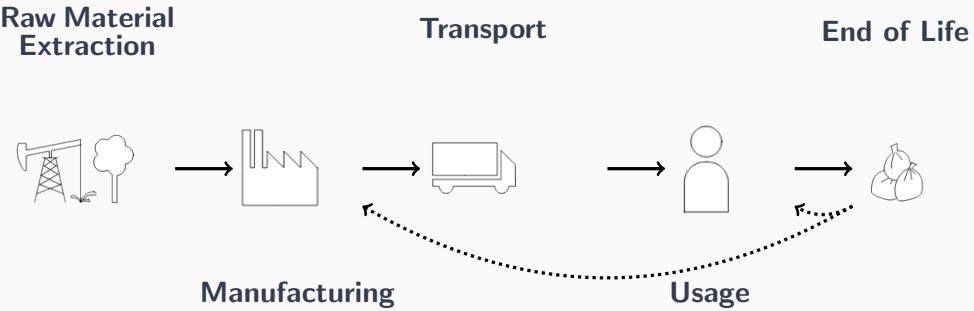

Each phase has different impacts

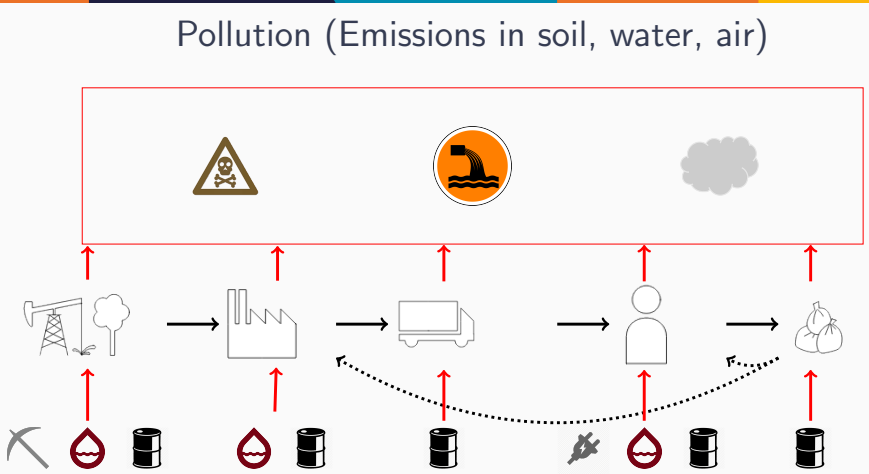

Tools for the evaluation of the environmental impacts of computation

Outil	Life cycle phase considered						Multiple impacts considered	Estimates consumption	GPU support
	Ext.	Man.	Tra.	Uti.		EoL.			
				Infra.	Dyn.				
Green Algorithms	✗	✗	✗	✓	✓	✗	✗	✓	✓
ML CO_2 Impact	✗	✗	✗	✗	✓	✗	✗	✓	✓
CarbonTracker	✗	✗	✗	✓	✓	✗	✗	✗	✓
CodeCarbon	✗	✗	✗	✓	✓	✗	✗	✗	✓
Boavizta	✓	✓	✗	✗	✗	✗	✓	-	✗

Tools for the evaluation of the environmental impacts of computation

Outil	Life cycle phase considered						Multiple impacts considered	Estimates consumption	GPU support
	Ext.	Man.	Tra.	Uti.		EoL.			
				Infra.	Dyn.				
Green Algorithms	✗	✗	✗	✓	✓	✗	✗	✓	✓
ML CO_2 Impact	✗	✗	✗	✗	✓	✗	✗	✓	✓
CarbonTracker	✗	✗	✗	✓	✓	✗	✗	✗	✓
CodeCarbon	✗	✗	✗	✓	✓	✗	✗	✗	✓
Boavizta	✓	✓	✗	✗	✗	✗	✓	-	✗

Evaluated impacts

- ADP, measured in kgSbeq [van Oers et al., 2020, Bruijn et al., 2002]
- PE, measured in MJ [Frischknecht et al., 2015]
- GWP, measured in gCO_2eq [Forster et al., 2023]

A tool for Machine Learning life Cycle Assessment (MLCA)

	ADP	GWP	PE	Human toxicity	Water Consumption	...
Extraction	✓	✓	✓	✗	✗	✗
Manufacturing	✓	✓	✓	✗	✗	✗
Transport	✗	✗	✗	✗	✗	✗
Usage	✓	✓	✓	✗	✗	✗
End of Life	✗	✗	✗	✗	✗	✗

Modeling graphics card

Manufacturing impacts attribution

Infrastructure consumption

Putting impacts in perspective

Modeling the manufacturing impacts of Graphics Cards

Bottom Up approach (Machine = graphics card + CPU + memory + motherboard + storage + PSU + casing)

Graphics card = GPU + Memory + Base

- GPU modeled by die size
- memory modeled by memory size
- base impacts computed from [Loubet et al., 2023]

$$Graphics\ card_{impact} = die_{size} * die_{impact_{per-cm^2}} + memory_{size} * memory_{impact_{perGB}} + base_{impact}$$

where impact $\in$ {ADP, PE, GWP}.

Example
NVIDIA A100 SMX4 80GB

GWP: 330 $kgCO_2$ eq PE: 3900 MJ ADP: 0.027 kgSb eq

Manufacturing impacts Attribution

- *linear* attribution
- [Luccioni et al., 2023] on the Jean Zay cluster

$$\texttt{embodied}_{impact} = \texttt{manufacturing}_{impact} \frac{\texttt{hours usage}}{\texttt{total available hours}}$$

impact $\in$ {ADP, PE, GWP}

Energy consumption (dynamic & infrastructure)

Dynamic energy consumption modeling based on the Thermal Design Power (TDP) as in [Lannelongue et al., 2021]

$$E_{dynamic} = \texttt{hours usage} * \sum_{p \in \{\text{CPU,GPU}\}} (n_p * u_p * \texttt{TDP}_p) + \texttt{memory}_{size} * P_{perGB}$$

Energy consumption (dynamic & infrastructure)

Dynamic energy consumption modeling based on the TDP as in [Lannelongue et al., 2021]

$$E_{dynamic} = \texttt{hours usage} * \sum_{p \in \{\text{CPU,GPU}\}} (n_p * u_p * \texttt{TDP}_p) + \texttt{memory}_{size} * P_{perGB}$$

From dynamic to total consumption

$$E = E_{dynamic} * \texttt{dynamic ratio} * 1E-3$$

$$\texttt{dynamic ratio} = \frac{\text{TOTAL}}{Production} \simeq \frac{\text{TOTAL}}{\sum_{j \in \text{Jobs}} (E_{dynamic})_j} \simeq 1.834$$

Tests on Jean Zay [Luccioni et al., 2023]

$$\texttt{Energy}_{impact} = E * impact_{per\ kWh}$$

Putting impacts in perspective

Need for a global perspective [Rasoldier et al., 2022], [Hauschild, 2015]

COP21 · CMP11
PARIS 2015
UN CLIMATE CHANGE CONFERENCE

"Stratégie Nationale Bas Carbone" [a]
2 tCO_2 e/person/year

[a] https://indicateurs-snbc.developpement-durable.gouv.fr/

[a]

Planetary boundaries [Sala et al., 2020]

- PB_{GWP} = 985 $kgCO_2$ e/person/year
- PB_{ADP} = 3.17E-02 kgSbeq/person/year

[a] Credit: "Azote for Stockholm Resilience Centre, based on analysis in Persson et al 2022 and Steffen et al 2015"

Putting impacts in perspective

Need for a global perspective [Rasoldier et al., 2022], [Hauschild, 2015]

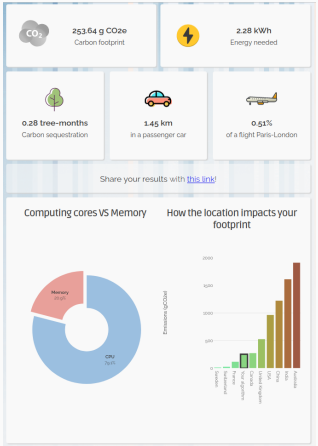

http://calculator.green-algorithms.org/

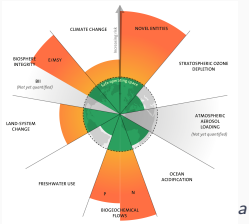

[a]

Planetary boundaries [Sala et al., 2020]

- $PB_{GWP} = 985\ kgCO_2\ e/person/year$
- $PB_{ADP} = 3.17E\text{-}02\ kgSbeq/person/year$

[a] Credit: "Azote for Stockholm Resilience Centre, based on analysis in Persson et al 2022 and Steffen et al 2015".

Putting impacts in perspective

Need for a global perspective [Rasoldier et al., 2022], [Hauschild, 2015]

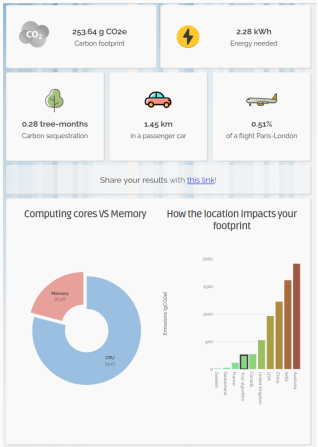

training BLOOM

- GWP: 59tCO_2 eq
  - annual emissions of 59 person (PB_{GWP})
  - annual emissions of 29 person (SNBC)
- ADP: 1.2 kgSb eq
  - annual resource extraction of 38 person (PB_{ADP})
- PE: 9800000 MJ

http://calculator.green-algorithms.org/

Case studies on impacts measurement

Evaluating the quality and the usability of MLCA in diverse situations

- Evaluation of the usage modeling, reproducing results from:
  - Impacts of Named Entity Recognition: [Bannour et al., 2021]
  - Survey of existing tools [Jay et al., 2023]
  - Impacts of Spoken Language Understanding: [Dinarelli et al., 2022]
  - Impacts of Transformers: [Cattan et al., 2022]
  - Impacts of Natural Language Processing (NLP): [Strubell et al., 2019]
- Evaluation of hardware manufacturing, reproducing LCA results from Dell:
  - Thinkstep for DELL on the R740 server [Thinkstep, 2019]
  - Sphera for Dell on the R6515, R7515, R7525 servers [Sphera, 2021]
- Evaluation of the whole tool: reproducing results from [Luccioni et al., 2023]

Evaluating the quality and the usability of MLCA in diverse situations

- Evaluation of the usage modeling, reproducing results from:
  - Impacts of Named Entity Recognition: [Bannour et al., 2021]
  - Survey of existing tools [Jay et al., 2023]
  - Impacts of Spoken Language Understanding: [Dinarelli et al., 2022]
  - Impacts of Transformers: [Cattan et al., 2022]
  - Impacts of NLP: [Strubell et al., 2019]
- Evaluation of hardware manufacturing, reproducing LCA results from Dell:
  - Thinkstep for DELL on the R740 server [Thinkstep, 2019]
  - Sphera for Dell on the R6515, R7515, R7525 servers [Sphera, 2021]
- Evaluation of the whole tool: reproducing results from [Luccioni et al., 2023]

Evaluating the quality and the usability of MLCA in diverse situations

- Evaluation of the usage modeling, reproducing results from:
  - Impacts of Named Entity Recognition: [Bannour et al., 2021]
  - Survey of existing tools [Jay et al., 2023]
  - Impacts of Spoken Language Understanding: [Dinarelli et al., 2022]
  - Impacts of Transformers: [Cattan et al., 2022]
  - Impacts of NLP: [Strubell et al., 2019]
- Evaluation of hardware manufacturing, reproducing LCA results from Dell:
  - Thinkstep for DELL on the R740 server [Thinkstep, 2019]
  - Sphera for Dell on the R6515, R7515, R7525 servers [Sphera, 2021]
- Evaluation of the whole tool: reproducing results from [Luccioni et al., 2023]

Evaluating the quality and the usability of MLCA in diverse situations

- Evaluation of the usage modeling, reproducing results from:
  - Impacts of Named Entity Recognition: [Bannour et al., 2021]
  - Survey of existing tools [Jay et al., 2023]
  - Impacts of Spoken Language Understanding: [Dinarelli et al., 2022]
  - Impacts of Transformers: [Cattan et al., 2022]
  - Impacts of NLP: [Strubell et al., 2019]
- Evaluation of hardware manufacturing, reproducing LCA results from Dell:
  - Thinkstep for DELL on the R740 server [Thinkstep, 2019]
  - Sphera for Dell on the R6515, R7515, R7525 servers [Sphera, 2021]
- Evaluation of the whole tool: reproducing results from [Luccioni et al., 2023]

Use phase: Replicating results from [Jay et al., 2023]

Difference between real TDP of the GPU and the TDP used in Green Algorithms

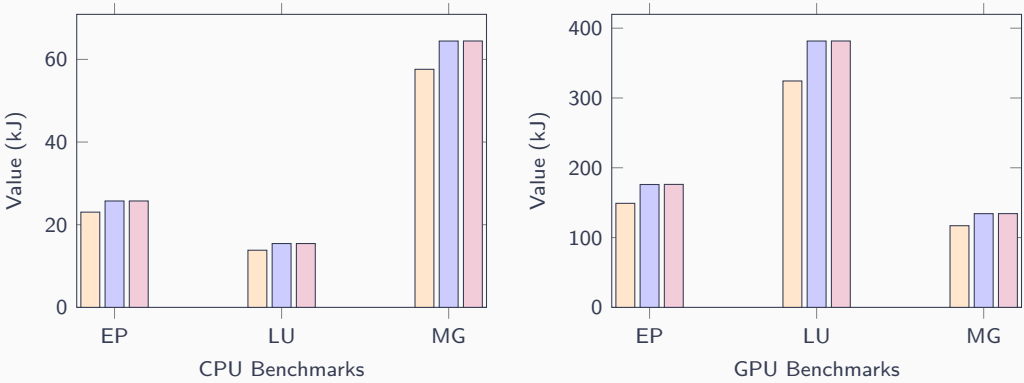

Value Real (orange ▯), Value Match (blue ▯), Expected (Green Algorithms) (purple ▯)

Manufacturing: Sphera for Dell on the R6515 & R6525 servers [Sphera, 2021]

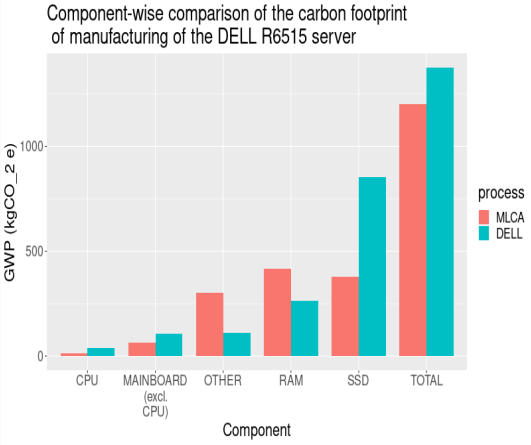

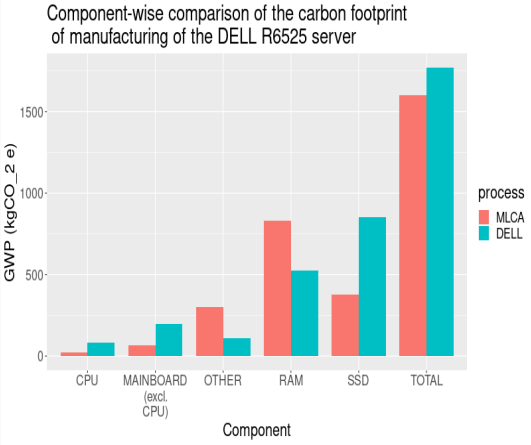

Evaluation of the whole tool: reproducing results from [Luccioni et al., 2023]

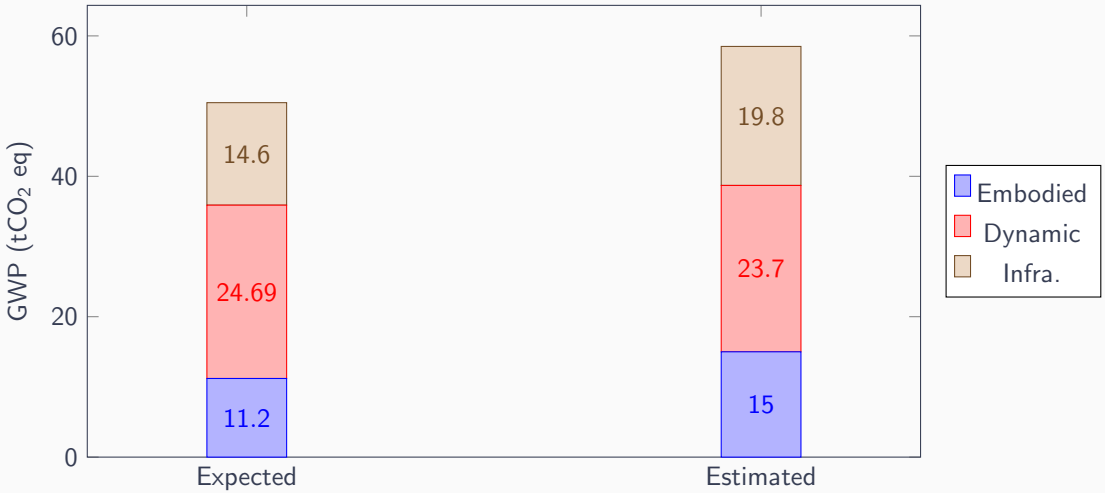

Toward a domain wide evaluation: the state of Natural Language Processing in the French healthcare system

Interviews for understanding the current state of ICT in French healthcare

Objectives

- Overview of ICT in healthcare, the needed infrastructures and probable evolution
- Understanding the state of reflexion of the professionals on the sustainability of ICT in healthcare

Focused on AI use in clinical data warehouses (Entrepôts de Données de Santé (EDS))

Protocol

- Series of semi-structured interviews [DiCicco-Bloom and Crabtree, 2006]
- 9 persons contacted, 7 positive responses, one decline and one pending.

Overview of ICT in healthcare, conclusions

1. ICT are ubiquitous within French healthcare (healthcare organisation, clinical practice, public health research).
2. The new availability of clinical data warehouses places the system at a turning point towards new deployment/uses of ICT in healthcare.
3. Still the beginning of the reflection on the sustainability of the ICT in health.

Conclusion

Conclusion

Contributions

- An estimation tool named *MLCA* evaluating multiple impacts and taking into account different phases of the life cycle
- An evaluation of the usability and quality of this tool
- An overview of ICT usage in the French Healthcare system

Future Work

- LCA of a Graphics card used in computing facilities
- Integration of End-of-Life considerations
- Applying results of the interviews for an evaluation of the impacts of ICT use in healthcare
- Interface *à la* Green Algorithms

Appendix

ICT in Healthcare

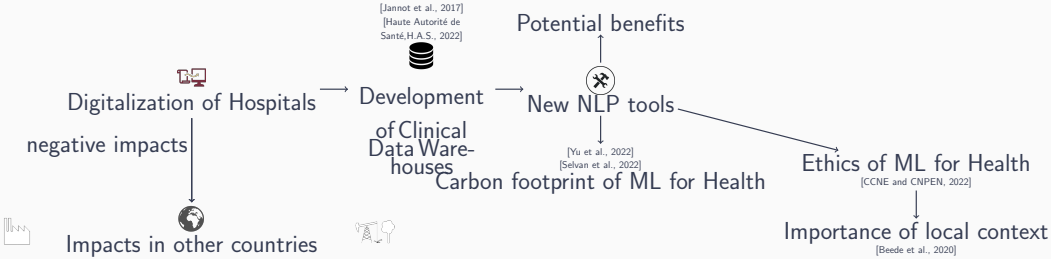

Interview protocol

Presentation and description of the objectives of the interview then:

1. What is your job and background?
2. What are the digital tools you use in your work or know are being used in health?
3. What infrastructure exists or is needed to support this/these usages?
4. What is the reflection on the environmental impacts induced by this/these usages?
5. What is the reflection on the ethics of using the ICT in health?

Participants of the study

Name	Background	Hospital staff	NLP researcher	Governmental agency staff	Management of an EDS	City
Christel Gerardin	CS & MD	✓	✓	✗	✗	Paris
Antoine Neuraz	CS & MD	✓	✓	✗	✗	Paris
Bastien Rance	CS	✓	✓	✗	✗	Paris
Romain Bey	CS	✓	✓	✗	✓	Paris
Stéfan Darmoni	CS & MD	✓	✓	✗	✓	Rouen
Nathalie Baudinière	MBA	✗	✗	✓	✗	Paris
Brigite Seroussi	CS & MD	✓	✓	✓	✗	Paris

CS = informatique, MD = médecine
MBA = Ingénieure et management

Interviews summary, Usages & Infrastructure

Usages

- Development of the Entrepôt de Données de Santé (EDS)
- EDS for research, care and direction
- Mostly new tools for automation

Infrastructure

- System duplication
- Computing power
- A turning point in infrastructure development

Interviews summary, Environmental policies

Environmental policies

- Little to no policies known by the researchers
- Carbon footprint of the hospitals IT systems: 190,000 tCO_2e [Baudinière et al., 2021]
- Existing policies: eco-score

Ethics

- Importance of privacy
- Rupture of the patient-clinician relationship
- Question of responsibility
- Question of medical training
- Risks of cyber-attacks and digital dependency

References i

Anthony, L. F. W., Kanding, B., and Selvan, R. (2020).
Carbontracker: Tracking and predicting the carbon footprint of training deep learning models.

Bannour, N., Ghannay, S., Névéol, A., and Ligozat, A.-L. (2021).
Evaluating the carbon footprint of NLP methods: a survey and analysis of existing tools.
In *EMNLP, Workshop SustaiNLP*, Proceedings of the 2nd Workshop on Simple and Efficient Natural Language Processing, Punta Cana, Dominican Republic.

Baudinière, N., Bégué, N., Bréhaux, K., Chaussat, J.-C., Darré-Bérenger, C., Guillet, P., Goglin, J.-F., Hanna, R., Joviado, O., Fol, V. L., Leroux, V., Marrauld, L., Séroussi, B., and Thierry, J.-P. (2021).
L'impact environnemental du numérique en santé.

accessible at https://esante.gouv.fr/sites/default/files/media_entity/documents/RAPPORT_GT6_VF.pdf.

Beede, E., Baylor, E., Hersch, F., Iurchenko, A., Wilcox, L., Ruamviboonsuk, P., and Vardoulakis, L. M. (2020).
A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy.
In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, page 1–12, New York, NY, USA. Association for Computing Machinery.

Boavizta (2021).
Numérique et environnement : Comment évaluer l'empreinte de la fabrication d'un serveur, au-delà des émissions de gaz à effet de serre?

References iii

Bruijn, H., Duin, R., Huijbregts, M. A. J., Guinee, J. B., Gorree, M., Heijungs, R., Huppes, G., Kleijn, R., Koning, A., van Oers, L., Sleeswijk, A. W., Suh, S., and de Haes, H. A. U. (2002).
Handbook on Life Cycle Assessment - Operational Guide to the ISO Standards.
Springer Dordrecht.

Cattan, O., Ghannay, S., Servan, C., and Rosset, S. (2022).
Benchmarking transformers-based models on french spoken language understanding tasks.
In *INTERSPEECH 2022*, Incheon, South Korea.

CCNE and CNPEN (2022).
Diagnostic Médical et Intelligence Artificielle : Enjeux Éthiques. Avis commun du CCNE et du CNPEN, Avis 141 du CCNE, Avis 4 du CNPEN.

References iv

DiCicco-Bloom, B. and Crabtree, B. F. (2006).
The qualitative research interview.
Medical Education, 40(4):314–321.

Dinarelli, M., Naguib, M., and Portet, F. (2022).
Toward low-cost end-to-end spoken language understanding.
arXiv preprint arXiv:2207.00352.

European Commission, Joint Research Centre, and Institute for Environment and Sustainability (2010).
International Reference Life Cycle Data System (ILCD) Handbook - General guide for Life Cycle Assessment - Detailed guidance.
EUR 24708 EN. Luxembourg. Publications Office of the European Union.
First edition March 2010. EUR 24708 EN. Luxembourg. Publications Office of the European Union.

Forster, P., Storelvmo, T., Armour, K., Collins, W., Dufresne, J.-L., Frame, D., Lunt, D., Mauritsen, T., Palmer, M., Watanabe, M., Wild, M., and Zhang, H. (2023).
The earth's energy budget, climate feedbacks and climate sensitivity.
In Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M., Huang, M., Leitzell, K., Lonnoy, E., Matthews, J., Maycock, T., Waterfield, T., Yelekçi, O., Yu, R., and Zhou, B., editors, *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, page 923–1054. Cambridge University Press.

Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G. S., and Friday, A. (2021).
The real climate and transformative impact of ict: A critique of estimates, trends, and regulations.
Patterns, 2(9):100340.


Frischknecht, R., Wyss, F., Knöpfel, S. B., Lützkendorf, T., and Balouktsi, M. (2015).
Cumulative energy demand in lca: the energy harvested approach.
International Journal of Life Cycle Assessment 20, page 957–969.


Gröger, J., Liu, R., Stobbe, L., Druschke, J., and Richter, N. (2021).
Green Cloud Computing: lebenszyklusbasierte Datenerhebung zu Umweltwirkungen des Cloud Computing: Abschlussbericht.
Umweltbundesamt.


Gupta, U., Kim, Y. G., Lee, S., Tse, J., Lee, H.-H. S., Wei, G.-Y., Brooks, D., and Wu, C.-J. (2022).
Chasing carbon: The elusive environmental footprint of computing.


Hauschild, M. Z. (2015).
Better – but is it good enough? on the need to consider both eco-efficiency and eco-effectiveness to gauge industrial sustainability.
Procedia CIRP, 29:1–7.
The 22nd CIRP Conference on Life Cycle Engineering.


Haute Autorité de Santé,H.A.S. (2022).
Entrepôts de données de santé hospitaliers en france.


Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky, D., and Pineau, J. (2020).
Towards the systematic reporting of the energy and carbon footprints of machine learning.
J. Mach. Learn. Res., 21(1).


Jannot, A.-S., Zapletal, E., Avillach, P., Mamzer, M.-F., Burgun, A., and Degoulet, P. (2017).
The georges pompidou university hospital clinical data warehouse: a 8-years follow-up experience.
International journal of medical informatics, 102:21–28.


Jay, M., Ostapenco, V., Lefèvre, L., Trystram, D., Orgerie, A.-C., and Fichel, B. (2023).
An experimental comparison of software-based power meters: focus on cpu and gpu.
In *CCGrid 2023 - 23rd IEEE/ACM international symposium on cluster, cloud and internet computing*, pages 1–13, Bangalore, India. IEEE.


Lacoste, A., Luccioni, A., Schmidt, V., and Dandres, T. (2019).
Quantifying the carbon emissions of machine learning.


Lannelongue, L., Grealey, J., and Inouye, M. (2021).
Green algorithms: Quantifying the carbon footprint of computation.
Advanced Science, 8(12):2100707.


Ligozat, A.-L., Lefèvre, J., Bugeau, A., and Combaz, J. (2021).
Unraveling the hidden environmental impacts of ai solutions for environment.


Loubet, P., Vincent, A., Collin, A., Dejous, C., Ghiotto, A., and Jego, C. (2023).
Life cycle assessment of ict in higher education: a comparison between desktop and single-board computers.
The International Journal of Life Cycle Assessment, pages 1–19.


Luccioni, A. S., Viguier, S., and Ligozat, A.-L. (2023).
Estimating the carbon footprint of BLOOM, a 176b parameter language model.
Journal of Machine Learning Research, 24(253):1–15.

References x

Malmodin, J., Lövehagen, N., Bergmark, P., and Lundén, D. (2023).
Ict sector electricity consumption and greenhouse gas emissions – 2020 outcome.
SSRN.

Parcollet, T. and Ravanelli, M. (2021).
The Energy and Carbon Footprint of Training End-to-End Speech Recognizers.
working paper or preprint.

Patterson, D., Gonzalez, J., Hölzle, U., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D. R., Texier, M., and Dean, J. (2022).
The carbon footprint of machine learning training will plateau, then shrink.
Computer, 55(7):18–28.

Rasoldier, A., Combaz, J., Girault, A., Marquet, K., and Quinton, S. (2022).
How realistic are claims about the benefits of using digital technologies for ghg emissions mitigation?
LIMITS'22: Workshop on Computing within Limits.

Sala, S., Crenna, E., Secchi, M., and Sanyé-Mengual, E. (2020).
Environmental sustainability of european production and consumption assessed against planetary boundaries.
Journal of Environmental Management, 269:110686.

Schmidt, V., Goyal-Kamal, Courty, B., Feld, B., Amine, S., kngoyal, Zhao, F., Joshi, A., Luccioni, S., Léval, M., Bogroff, A., de Lavoreille, H., Laskaris, N., Connell, L., Wang, Z., Saboni, A., Catovic, A., Blank, D., Stechly, M., alencon, JPW, Books, M., Swadik, S., M., H., Coutarel, M., Pollard, M., McCarthy, C., Husom, E. J., Vicente, F., and Tae, J. (2022).
mlco2/codecarbon: v2.1.4.

Schödwell, B., Zarnekow, R., Liu, R., Gröger, J., and Wilkens, M. (2018).
Kennzahlen und indikatoren für die beurteilung der ressourceneffizienz von rechenzentren und prüfung der praktischen anwendbarkeit.
Umweltforschungsplan des Bundesministeriums für Umwelt, Naturschutz, Bau und Reaktorsicherheit; Umweltbundesamt: Dessau-Roßlau, Germany.

Schwartz, R., Dodge, J., Smith, N. A., and Etzioni, O. (2020).
Green ai.
Commun. ACM, 63(12):54–63.

Selvan, R., Bhagwat, N., Wolff Anthony, L. F., Kanding, B., and Dam, E. B. (2022).
Carbon footprint of selecting and training deep learning models for medical image analysis.
In Wang, L., Dou, Q., Fletcher, P. T., Speidel, S., and Li, S., editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, pages 506–516, Cham. Springer Nature Switzerland.

Sevilla, J., Heim, L., Ho, A., Besiroglu, T., Hobbhahn, M., and Villalobos, P. (2022).
Compute trends across three eras of machine learning.
In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.

Sphera (2021).
Life Cycle Assessment Dell Servers R6515, R7515, R6525, R7525.

Strubell, E., Ganesh, A., and McCallum, A. (2019).
Energy and policy considerations for deep learning in NLP.
In Korhonen, A., Traum, D. R., and Màrquez, L., editors, *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3645–3650. Association for Computational Linguistics.

Thinkstep (2019).
Life Cycle Assessment of Dell R740.

van Oers, L., Guinée, J. B., and Heijungs, R. (2020).
Abiotic resource depletion potentials (ADPs) for elements revisited—updating ultimate reserve estimates and introducing time series for production data.
The International Journal of Life Cycle Assessment, 25:294–308.

Verdecchia, R., Sallou, J., and Cruz, L. (2023).
A systematic review of green ai.

Wu, C., Raghavendra, R., Gupta, U., Acun, B., Ardalani, N., Maeng, K., Chang, G., Behram, F. A., Huang, J., Bai, C., Gschwind, M., Gupta, A., Ott, M., Melnikov, A., Candido, S., Brooks, D., Chauhan, G., Lee, B., Lee, H. S., Akyildiz, B., Balandat, M., Spisak, J., Jain, R., Rabbat, M., and Hazelwood, K. M. (2022).
Sustainable AI: environmental implications, challenges and opportunities.


Yu, J.-R., Chen, C.-H., Huang, T.-W., Lu, J.-J., Chung, C.-R., Lin, T.-W., Wu, M.-H., Tseng, Y.-J., and Wang, H.-Y. (2022).
Energy efficiency of inference algorithms for clinical laboratory data sets: Green artificial intelligence study.
J Med Internet Res, 24(1):e28036.