



Towards sustainable and resilient AI

frugal and energy-aware distributed machine learning platforms

PhD Directors:
Sébastien MONNET, Stéphan PLASSART



Last year




Fa**L**a**F**El**S**
*F*ederated *L*earning *F*rugality
and *E*fficiency via *S*imulation






Platform heterogeneity







More machine, less runtime?





More machine, less runtime?



Mary Huet de Barochez, A., Plassart, S., & Monnet, S. (AINA 2025, April)

06/01/2026

Andrew MARY HUET DE BAROCHEZ



Takeaways:

1. Reducing runtime $\Rightarrow$ reducing energy
2. Finding the sufficient number of machine is bound to the **workload**, **heterogeneity of machines** and **parallel algorithms**.



This year, even more appetizing!



DAHL



*Distributed, Arbitrary and
Heterogenous Learning*




Carbon footprint of devices:
40% usage
60% manufacturing, distribution, end-of-life

ADEME-Arcep 2025






Homogenous platform





Heterogeneous platform

Phase de calcul



...

Phase de
synchronisation



Handling heterogeneity

- Of machines
- Network
- Computations

Allows existing machines to collaborate


Limiting:

- Investments in new machines
- Premature replacement



AI framework PoC

- Using a Task Graph Computing System (TGCS)
 - Supports multiple hardware
 - Flexible load balancing
 - Arbitrary execution order
- Frugality by design
 - Python layer
 - Energy-aware scheduling






DAHL : *D*istributed, *A*rbitrary and *H*eterogenous *L*earning

- Written in C and CUDA
- Based on **StarPU** (inria Bordeaux)
- Unit test following **Pytorch** values
- Reproducible
- Open source:
codeberg.org/PhoqueEberlue/dahl












Forward-backward dependency







CPU + GPU training





Experimenting

Comparing **DAHL** with **Pytorch** on the same use cases.

Datasets:

- Fashion-MNIST (28x28)
- CIFAR-10 (32x32x3)
- Big-fashion (270x360x3)

Hyperparameters:

- Batch size [1-256]
- LR 0.001
- Conv-relu-maxpool-fc
- 8000 samples per iter

Estimating energy consumption using  **Alumet**

10 runs per experiment



Reading results

Results for fashion-mnist

Runtime by batch_size

Energy Consumption by batch_size

Accuracy by batch_size

Experiments varying batch size

Runtime by num_cpu


Energy Consumption by num_cpu

Accuracy by num_cpu

Experiments varying number of CPU









Reading results





DAHL vs Pytorch : Fashion-MNIST







Results for fashion-mnist





DAHL vs Pytorch : CIFAR-10







Results for cifar-10





DAHL vs Pytorch : Big-fashion

Results for big-fashion





Thanks for your attention!

