


Accenture Labs

Inria

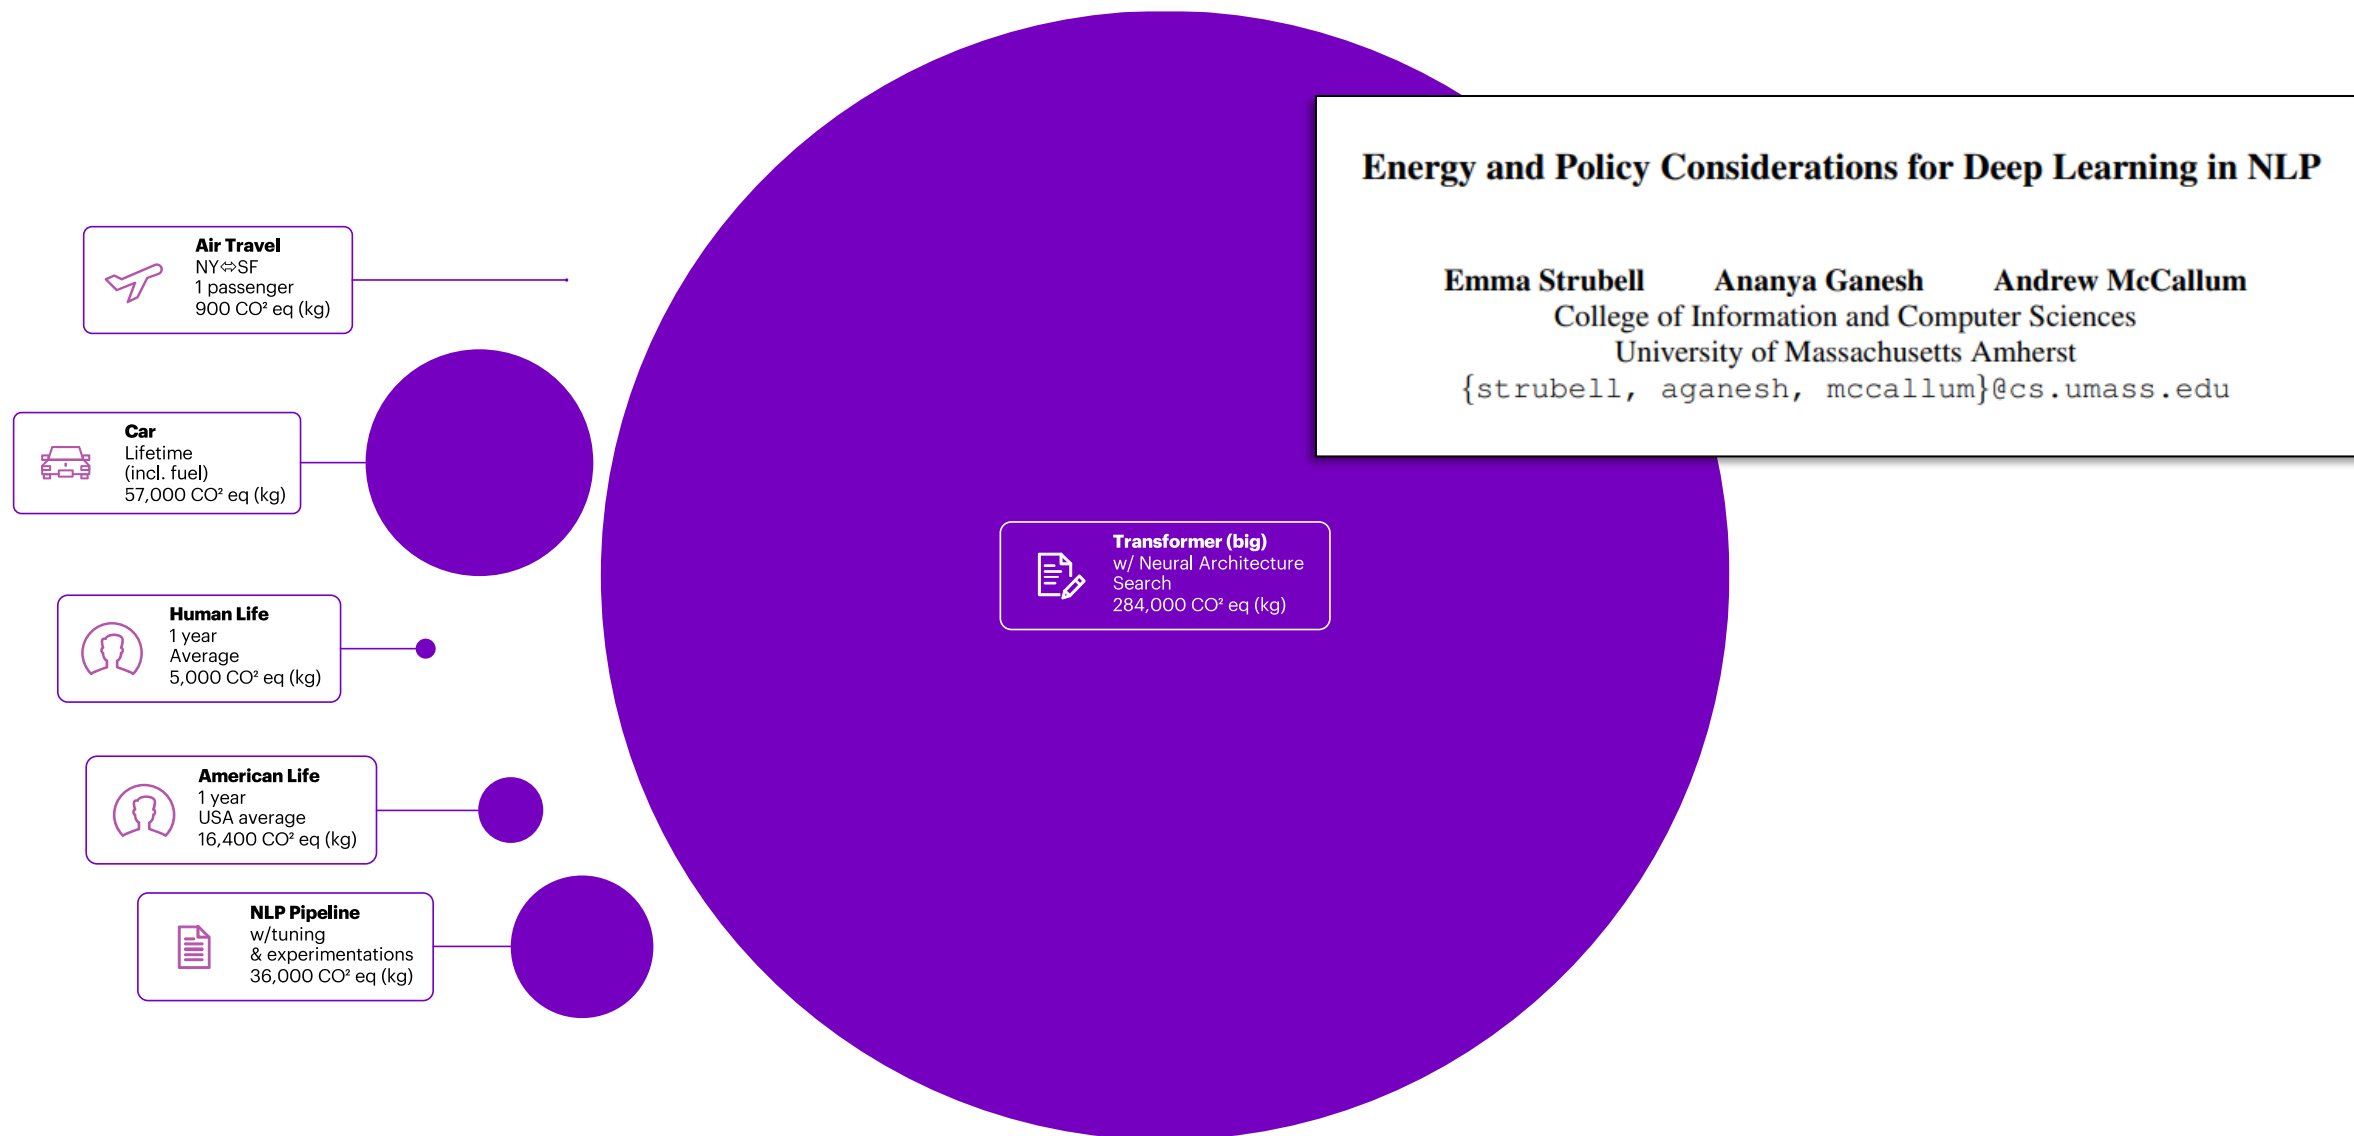
How can we evaluate the energy consumption of Machine Learning?

JRAF2023

Charlotte Rodriguez



Artificial Intelligence sustainability



Energy and Policy Considerations for Deep Learning in NLP

Emma Strubell Ananya Ganesh Andrew McCallum
College of Information and Computer Sciences
University of Massachusetts Amherst
{strubell, aganesh, mccallum}@cs.umass.edu

No consensus on how to evaluate the consumption

Energy and Policy Considerations for Deep Learning in NLP

Emma Strubell Ananya Ganesh Andrew McCallum

College of Information and Computer Sciences

University of Massachusetts Amherst

{strubell, aganesh, mccallum}@cs.u

Carbon Emissions and Large Neural Network Training

David Patterson^{1,2}, Joseph Gonzalez², Quoc Le¹, Chen Liang¹, Lluís-Miquel Munguia¹,
Daniel Rothchild², David So¹, Maud Texier¹, and Jeff Dean¹


{davidpatterson, qvl, crazydonkey, llmunguia, davidso, maudt, jeff}@google.com,
{pattsn, jegonzal, drothchild}@berkeley.edu

versus

“[the author’s] energy estimate for NAS ended up 18.7X too high for the average organization (...) and 88X off in emissions for energy-efficient organizations like Google”.

Systematic Literature Review (SLR)

Question: What tools and methods (not shaped into tools) currently permit to evaluate the energy consumption of machine learning computing tasks?

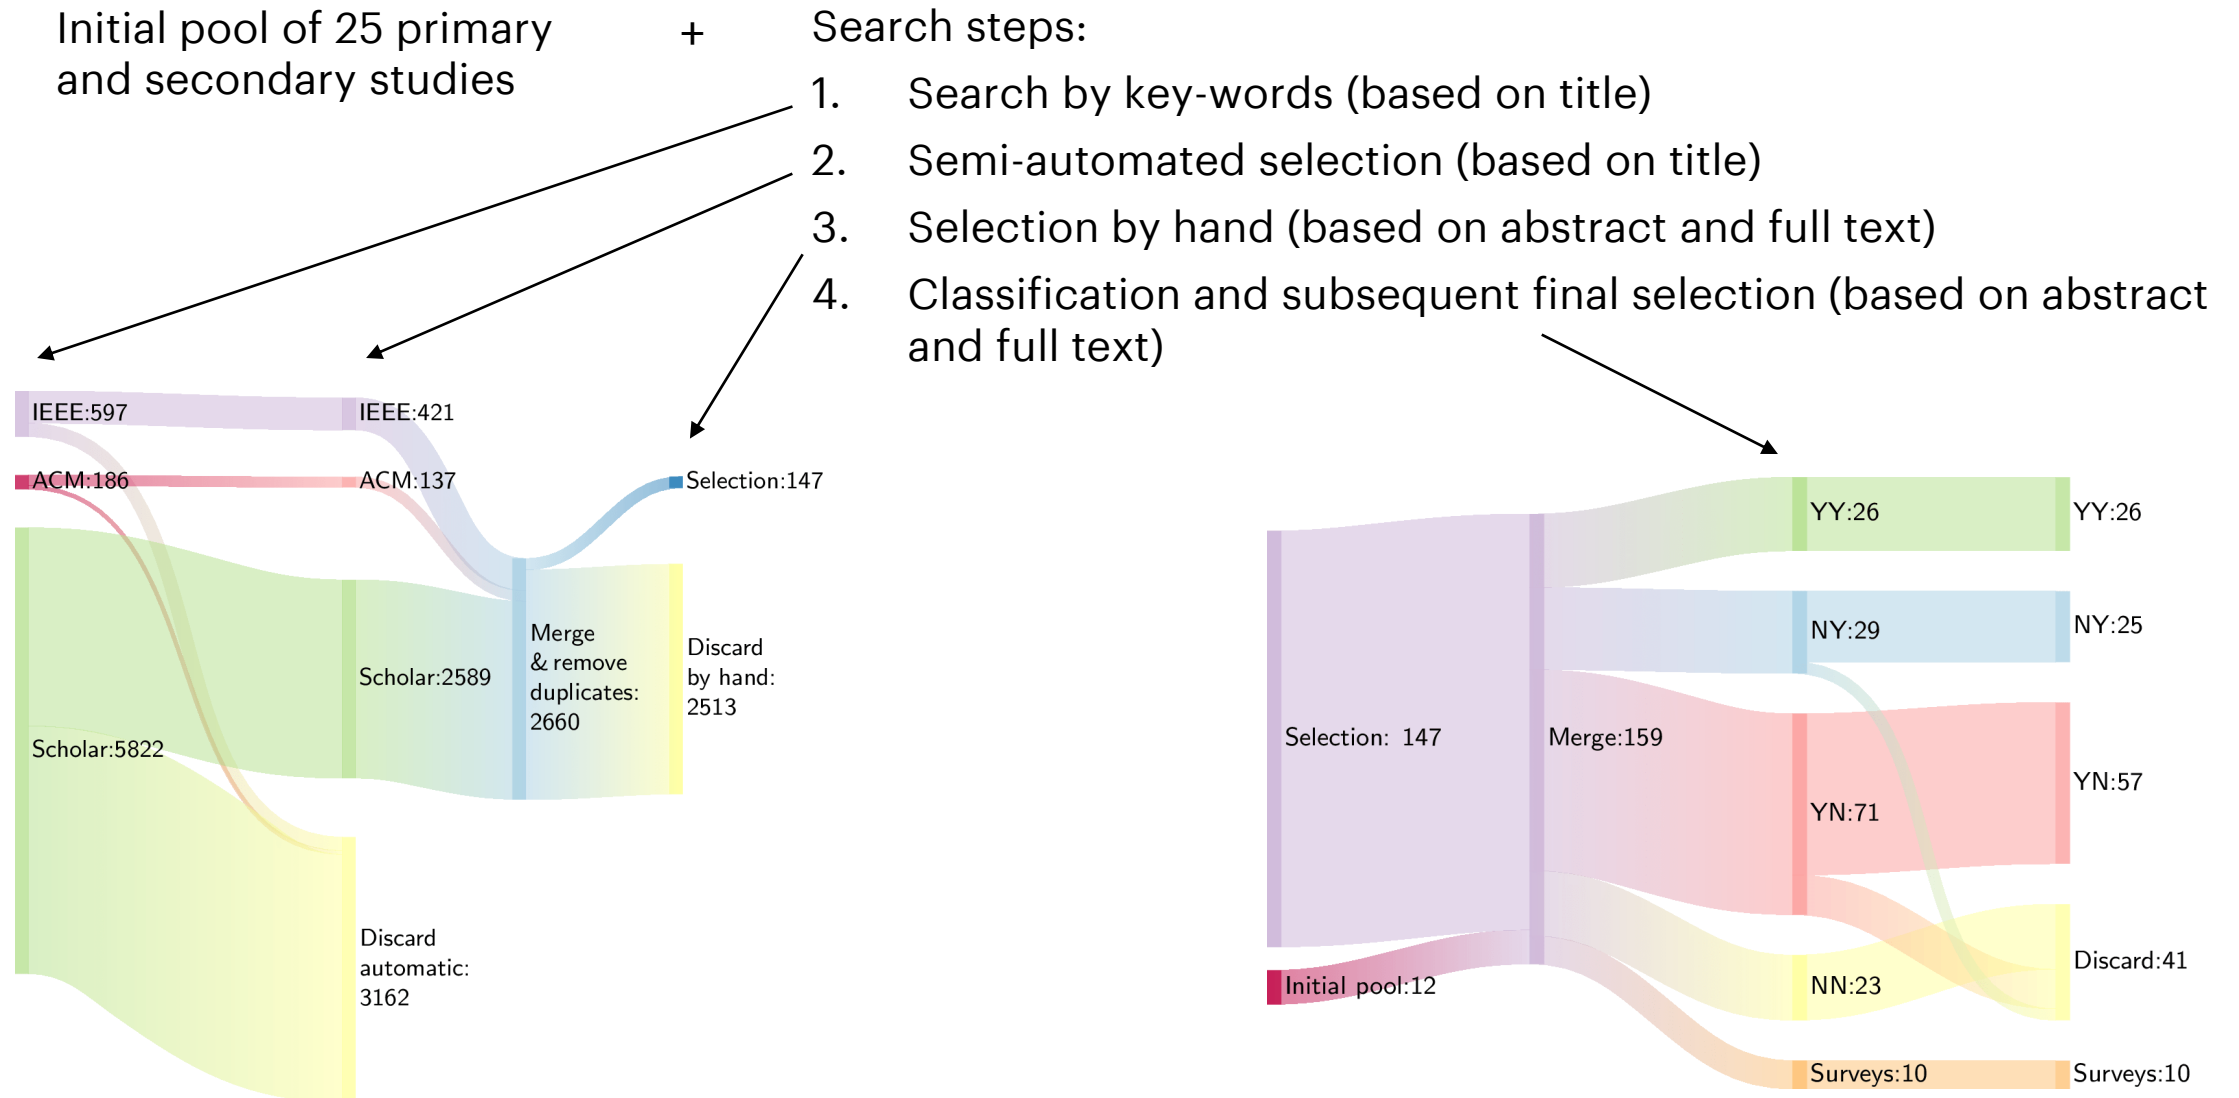


Data sources: ACM Digital Library, IEEE Xplore Digital Library, Google Scholar search engine (GS)

GS presents several constraints:

- search query may contain **at most 256** characters.
- 1 query will provide **at most 1000** results, even if it has more associated results (actual number of results is displayed by GS, even though the user does not have access to all of them).
- GS **official interface** only permits the user to save results to the user's Google Scholar Library by selecting the star icon, and then exporting said library. GS will tend to block any behavior deviating from this usage mode.

SLR: overview of the search steps



SLR: search step 1

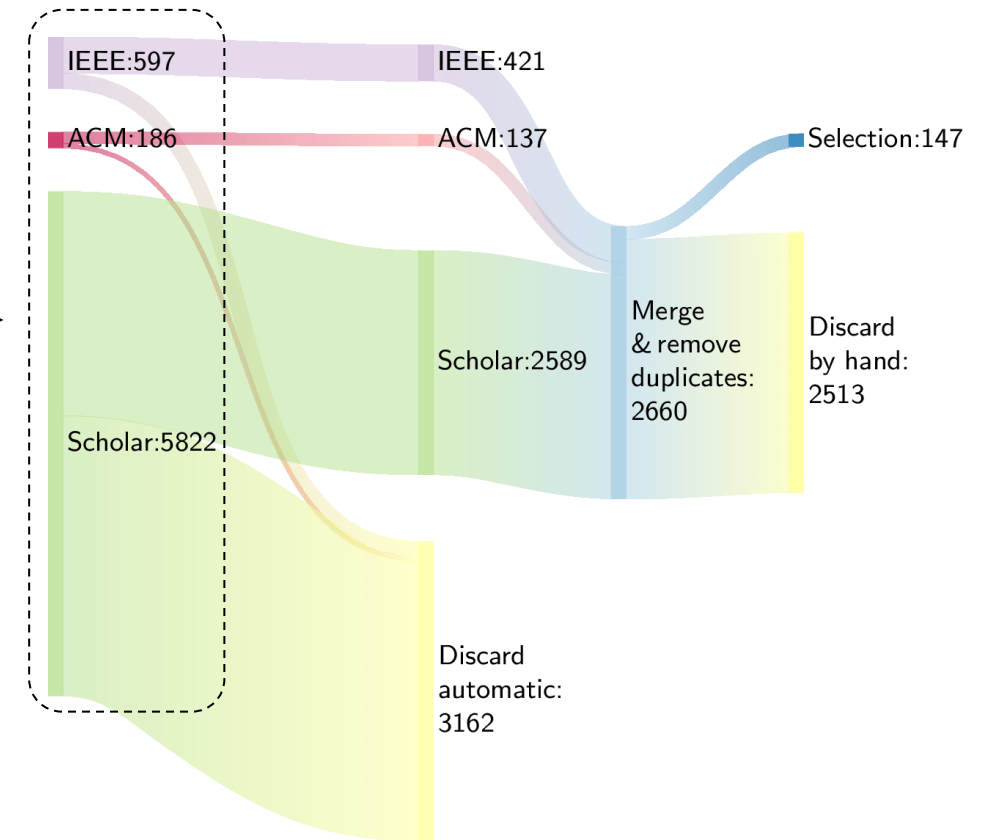
Search steps:

1. **Search by key-words (based on title)** →
2. Semi-automated selection (based on title)
3. Selection by hand (based on abstract and full text)
4. Classification and subsequent final selection (based on abstract and full text)

Key-words:

- i. machine learning, deep learning, computing, information and communications technology, ICT, artificial intelligence, AI, natural language processing, NLP, neural network, neural networks, CNN, DNN, computation, computations, software, process-level, server, virtual machine, federated learning, distributed learning
- ii. measure, measuring, estimate, estimation, consumed, consumption, predict, prediction, predicting, track, tracking, report, reports, reporting, account, quantify, quantifying, monitor, monitoring, evaluate, evaluating
- iii. energy, power, environmental impact, carbon footprint, carbon emissions, carbon impact.

Simultaneously excluded all results containing any of the following keywords: wind, building, buildings, vehicles, homes, ships, solar, photovoltaic, vehicle.



Remarks:

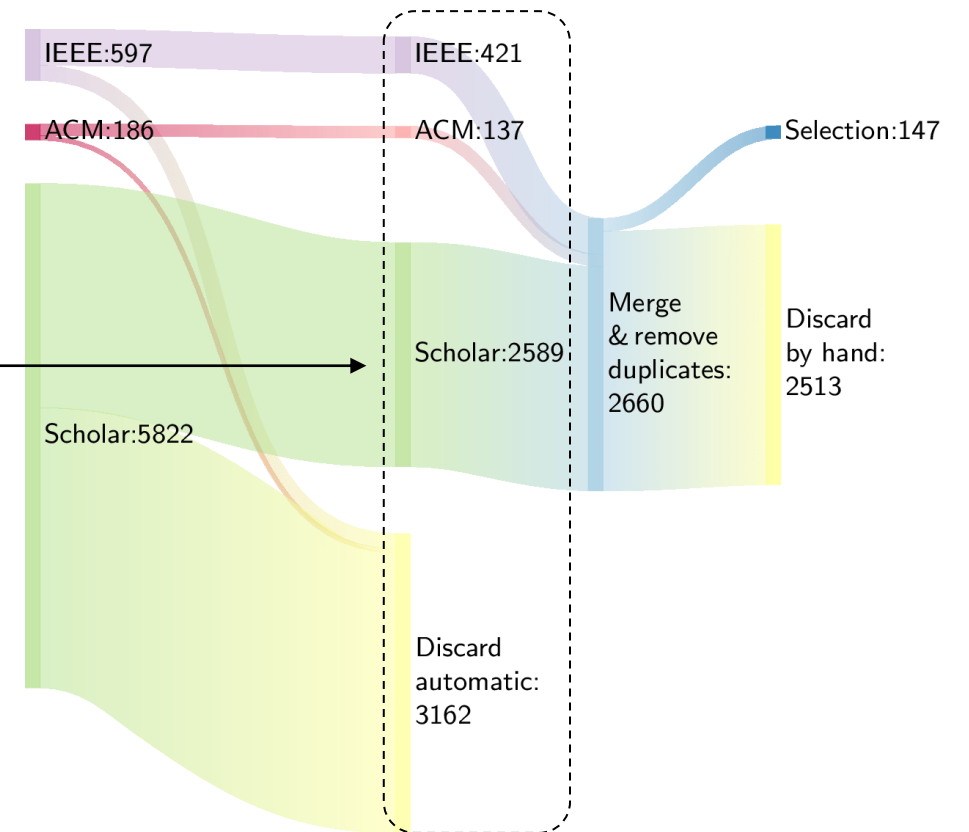
- Use the operators AND, OR, NOT, parenthesis and “phrase word for word” to build appropriate **queries** for each of the data sources.
- For GS, the original query had to be divided into **103 sub-queries to meet the constraints of GS**. The results of these sub-queries were then merged together, the duplicates removed.



SLR: search step 2

Search steps:

1. Search by key-words (based on title)
2. **Semi-automated selection (based on title)**
3. Selection by hand (based on abstract and full text)
4. Classification and subsequent final selection (based on abstract and full text)



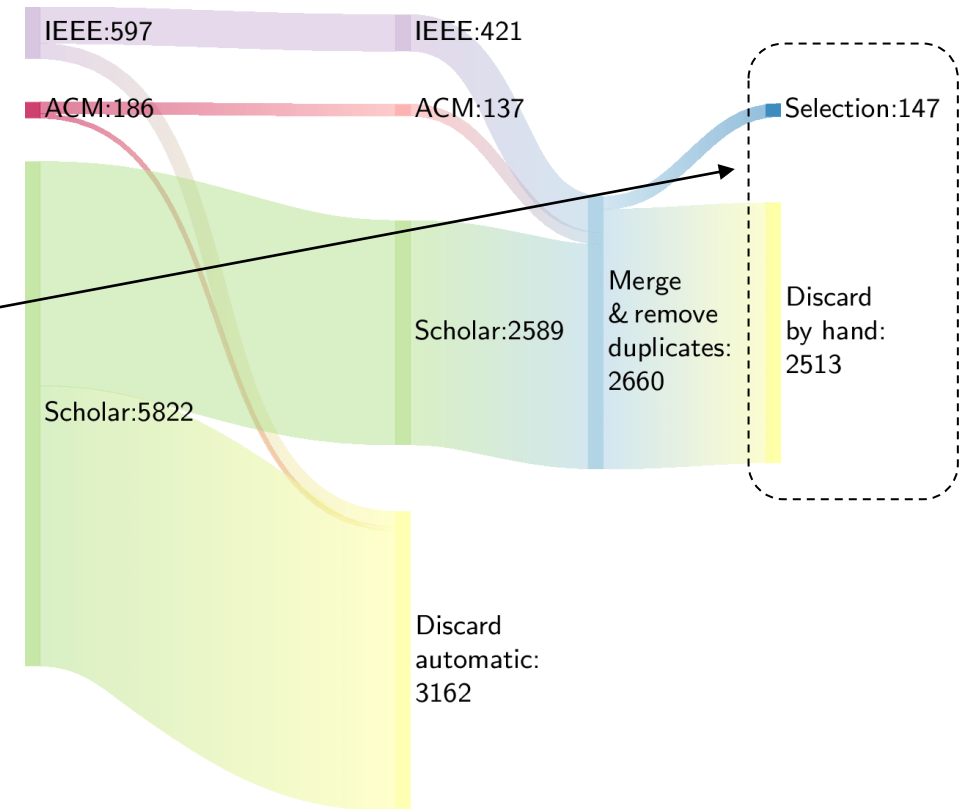
Remarks:

- First part of the selection process
- Based on the results' titles and on the relevance to the research question only.
- Consists in identifying words (among all words contained in the results' titles) that rule a title containing any of these words, as **off-topic**.

SLR: search step 3

Search steps:

1. Search by key-words (based on title)
2. Semi-automated selection (based on title)
- 3. Selection by hand (based on abstract and full text)**
4. Classification and subsequent final selection (based on abstract and full text)



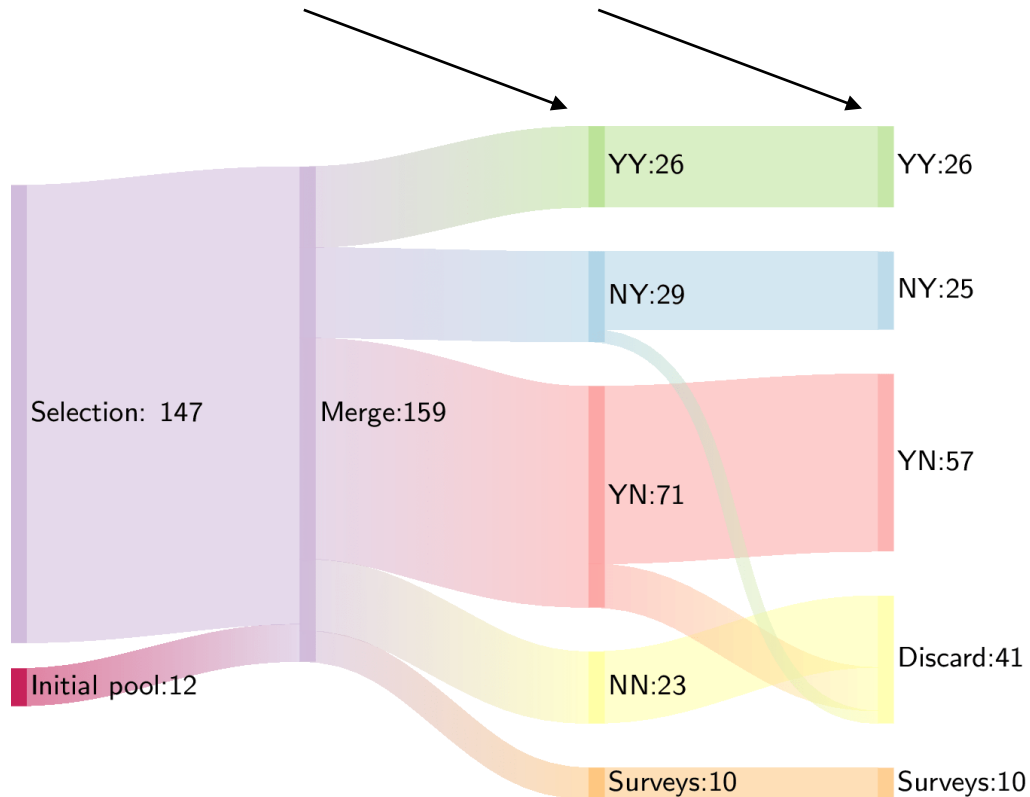
Remarks:

- Second part of the selection process
- Based on assessors reading the results' titles, abstracts and full text if necessary.
- Three assessors share this task.

SLR: search step 4

Search steps:

1. Search by key-words (based on title)
2. Semi-automated selection (based on title)
3. Selection by hand (based on abstract and full text)
4. **Classification and subsequent final selection (based on abstract and full text)**

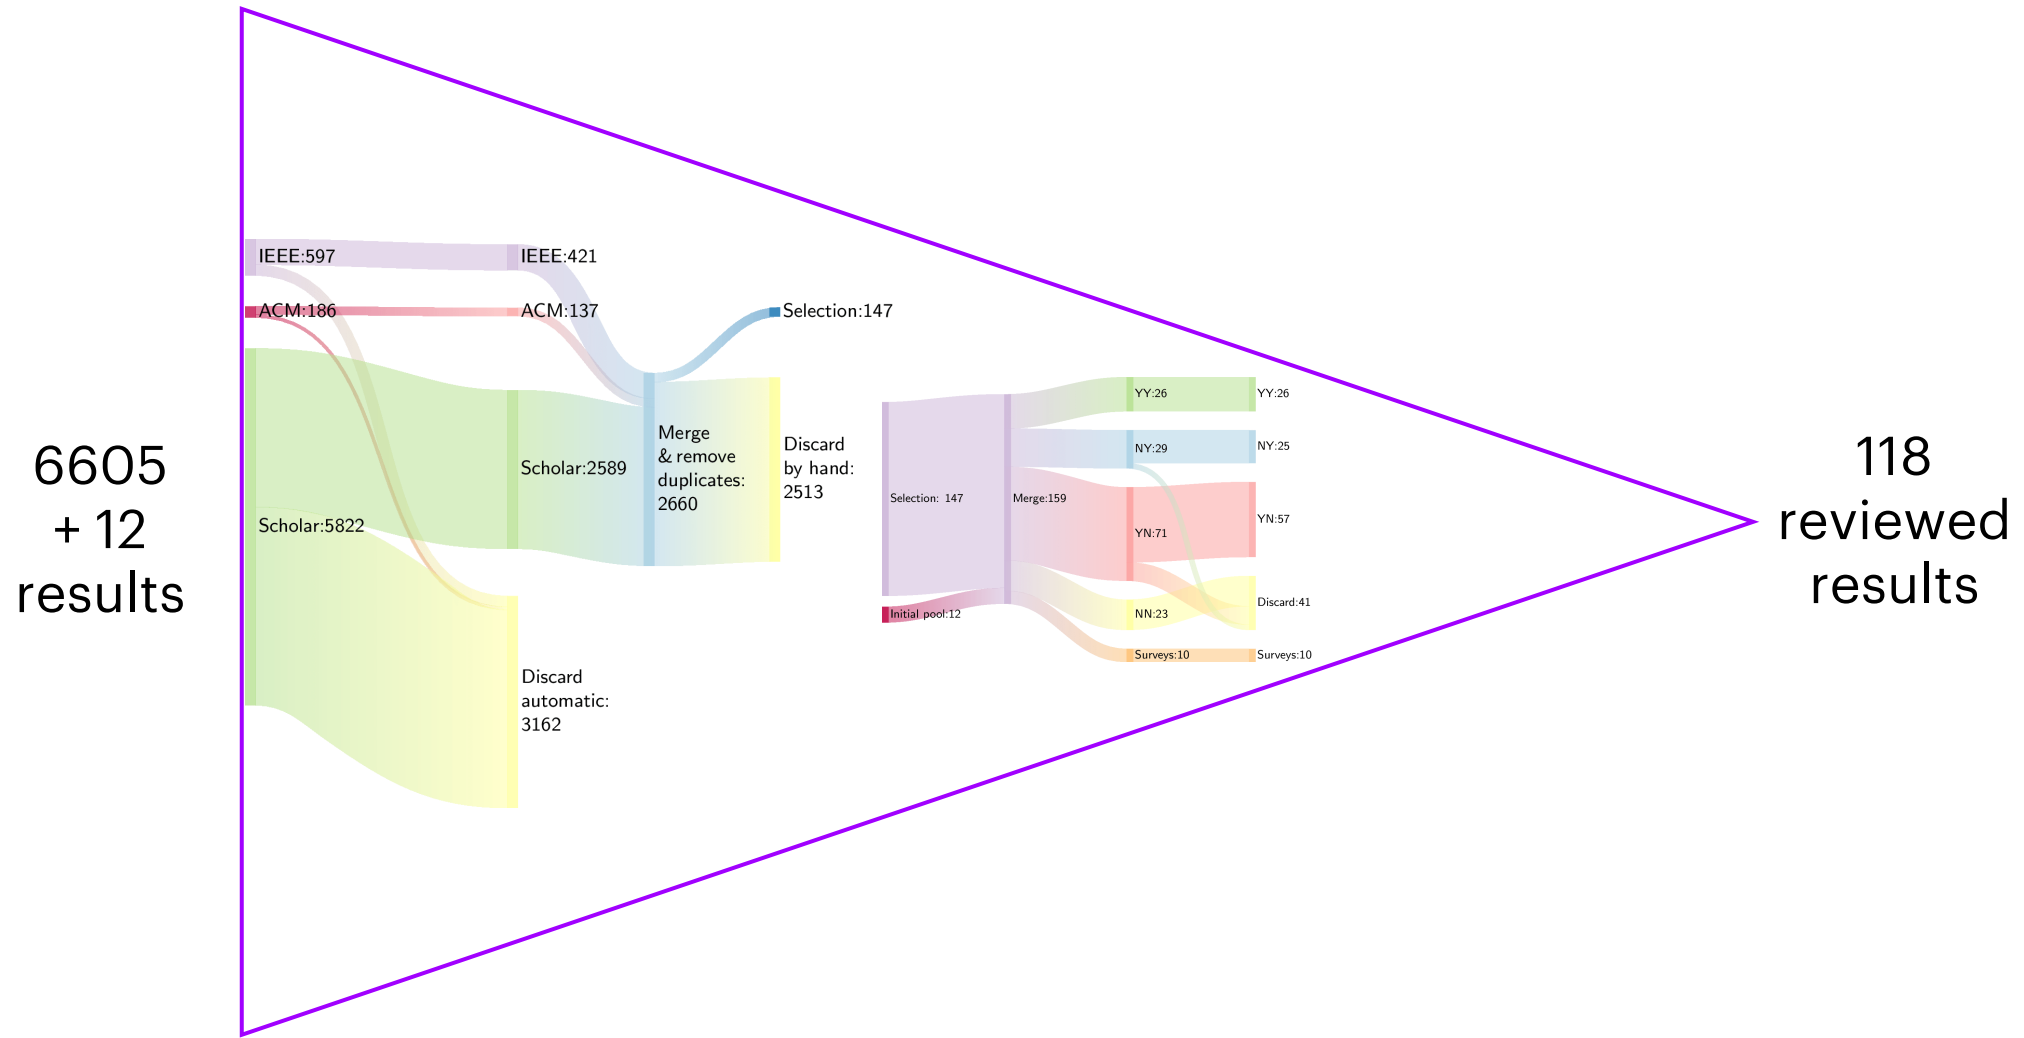


	for ML	yes	no
creation		yes	no
yes		YY	YN
no		NY	NN
survey		SY	SN

Classify the results selected in Step 3 according to two criteria:

- I. If primary study, has a tool or method been created by the authors or not (= used only)? Otherwise, the result is a secondary study (survey). Three possible values: “Yes – creation”, “No – no creation”, “Survey”.
- II. Is the result specifically concerned with AI applications or not (= concerned with software in general, virtual machines, data centers, etc.)? Two possible values: “Yes – for ML”, “No – not for ML”.

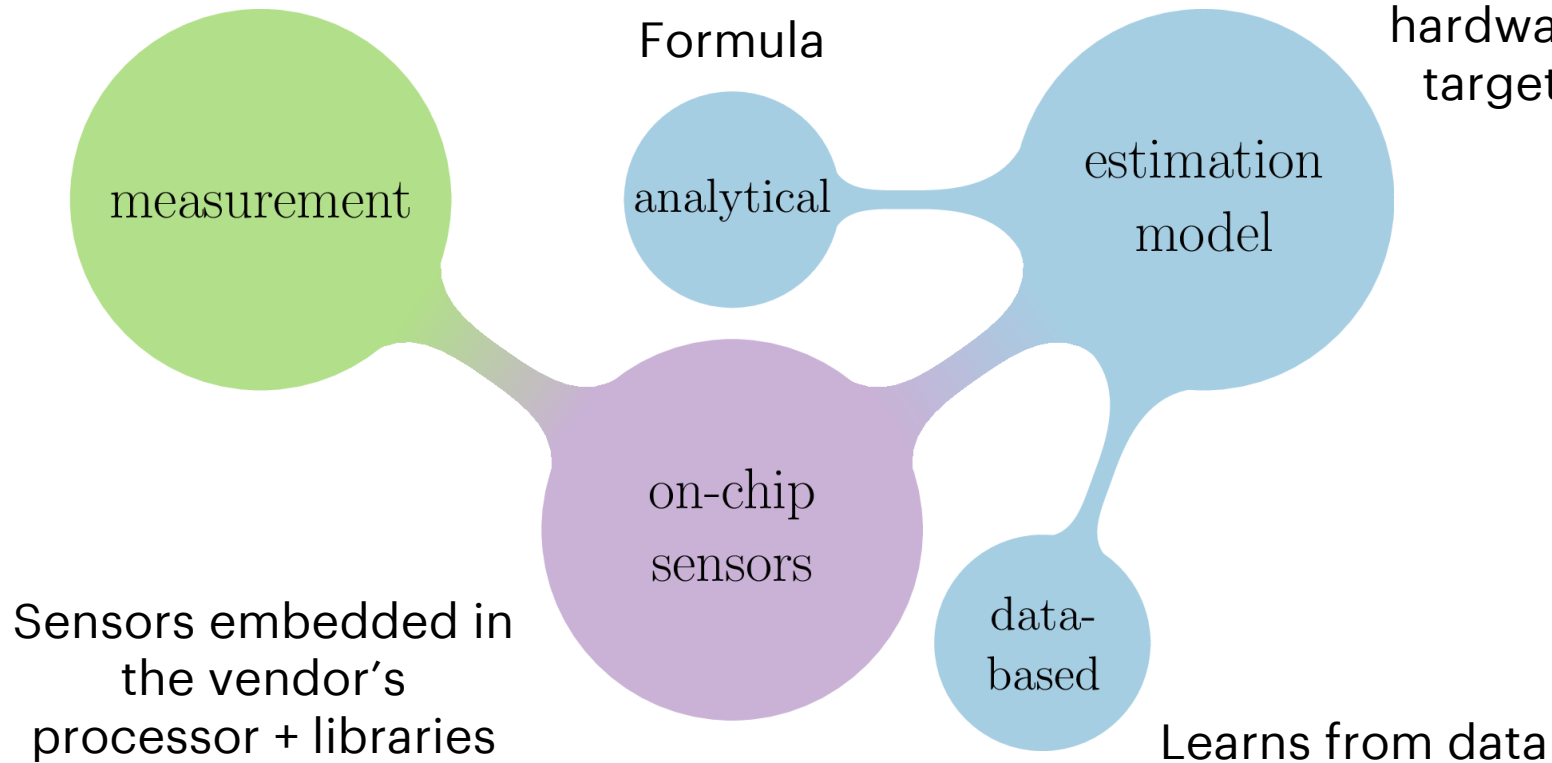
SLR: search steps summary



Different approaches to evaluate energy consumption

Actual measurements of current intensity, etc.

Estimation from indirect evidence (activity of hardware, charact. of target application)




Examples of tools and methods




External Power Meter



ECO2AI


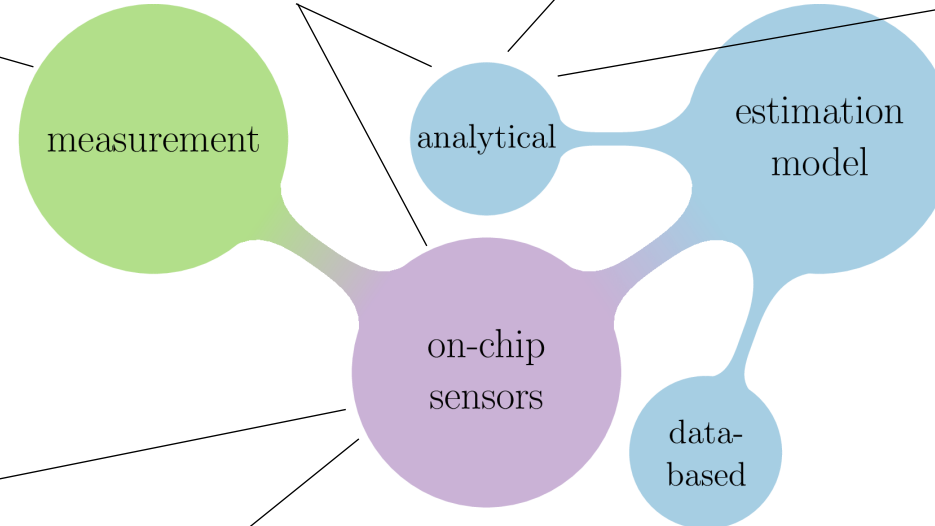


GREEN ALGORITHMS



FLOPS

FLOPS



CARBON TRACKER



CODE CARBON



Comparison framework

Qualitative study

- On the method or tool (approach, output, ...)
- On the usage (hardware, software constraints, ...)
- On the surrounding project (availability, context, ...)

Quantitative study





Image Classifier on MNIST

Script: PyTorch example [Basic MNIST Example](#)
GitHub: [pytorch/examples/tree/main/mnist](#)

Image Classifier on CIFAR10

Script: PyTorch tutorial "Training a classifier"
Pytorch website at: [tutorials/beginner/blitz/](#)

Resnet18 on ImageNet

Script: recipe for training Resnet18 from PyTorch
GitHub: [pytorch/vision/references/classification/](#)

Bert-base on SQUADv1.1

Script: recipe for fine-tuning Bert from google-research
GitHub: [google-research/bert/](#)

Comparison framework: application

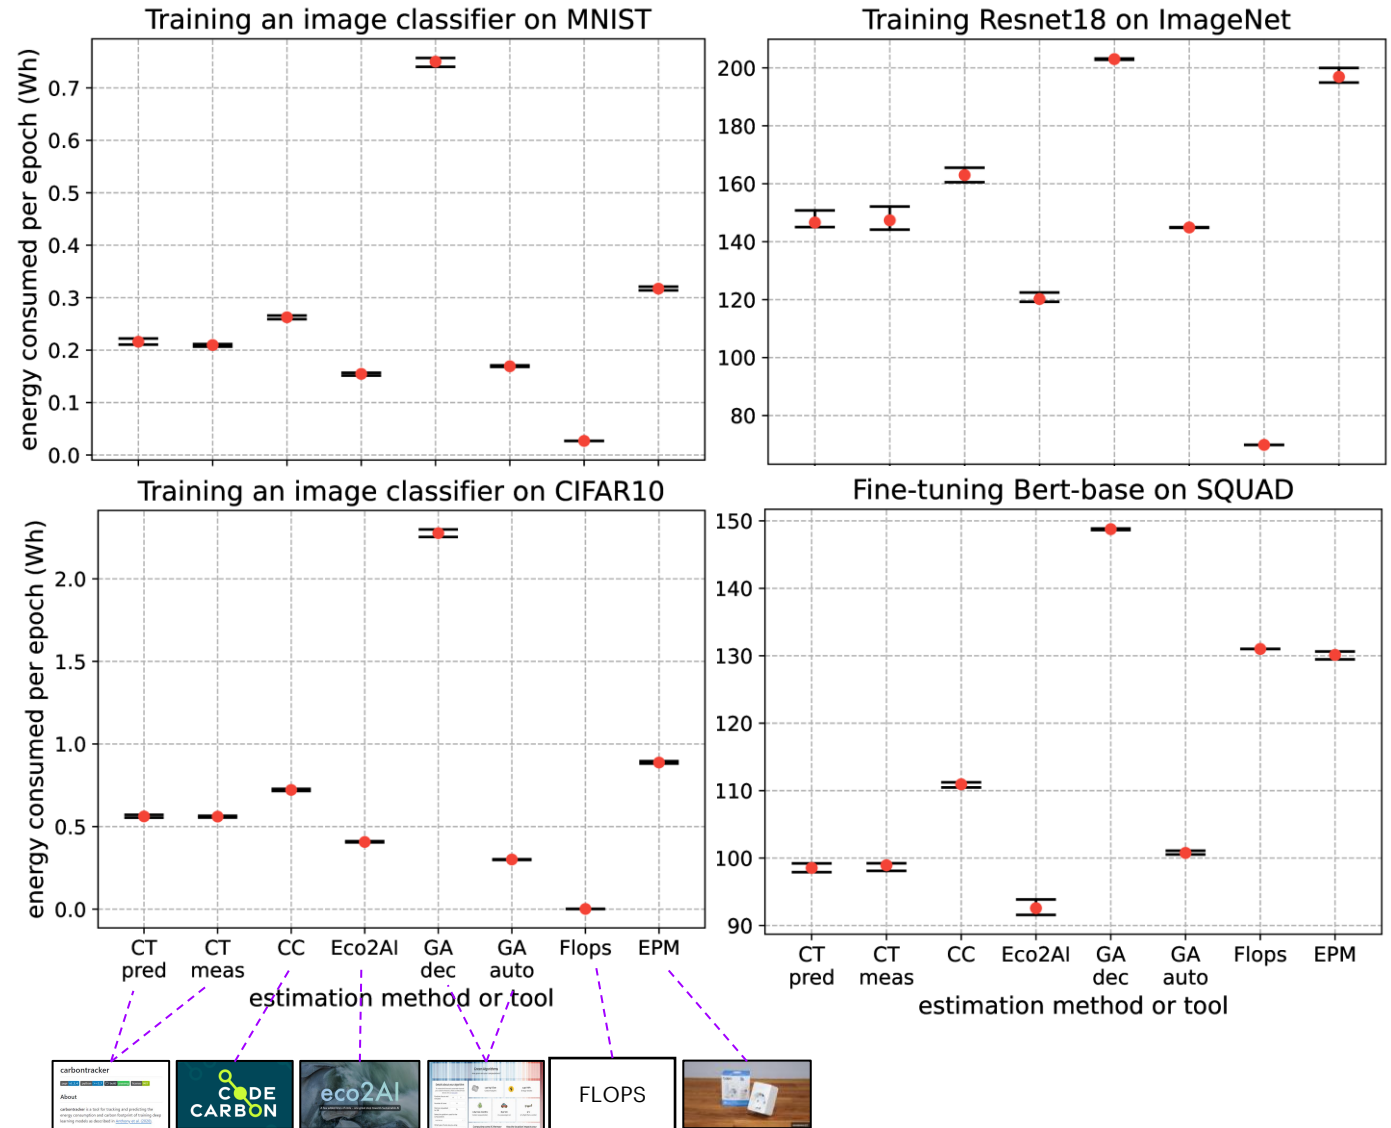
Evaluation tools and methods tested:

CT:pred, CT:meas, CC, Eco2AI, GA:dec, GA:auto, Flops, EPM.

- Some tools are tested in different modes
 - GA:dec (100% hardware usage), GA:auto (monitoring + average of hardware usage)
 - CT:meas (normal mode), CT:pred (prediction based on 1 epoch)
- EPM = external power meter

Observations:

- Different outputs for different approaches, though some remain comparable
- Different outputs for some tools and methods, depending on the ML task



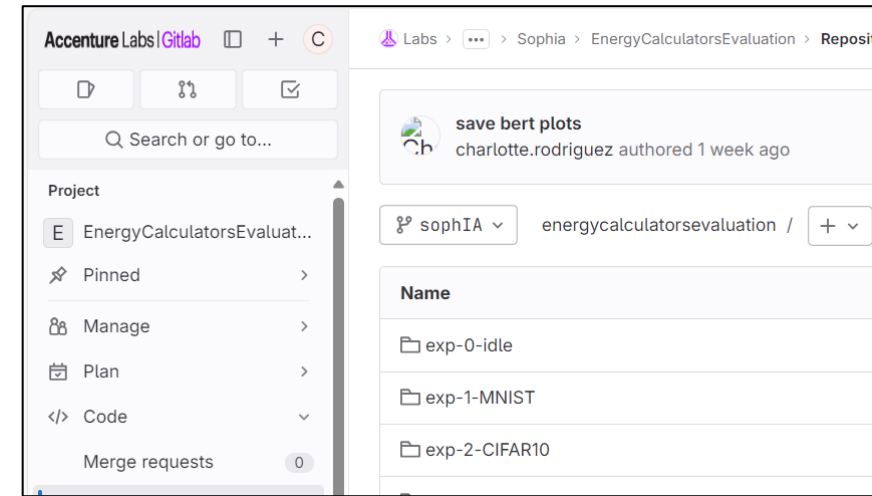
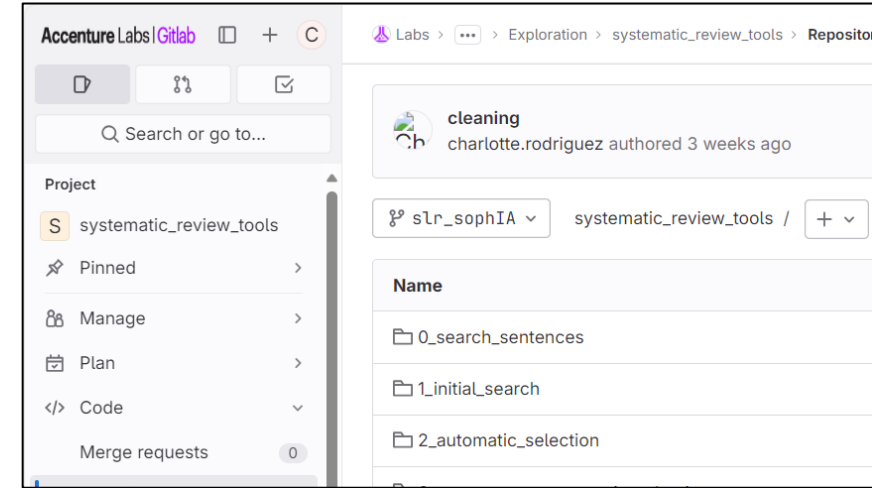
Conclusion

Limits:

- SLR: search done in June
- SLR: no backward or forward search
- Experiments: no baseline

Next:

- This study on arxiv and github (SLR & exp.)
- Work on some of the limits
- Recommendations
- Scheduling tool for decentralized learning



Conclusion

Accenture Labs | Gitlab

systematic_review_tools

cleaning
charlotte.rodriguez authored 2 weeks ago

slr_sophIA

Name	Last commit
0_search_sentences	cleaning
1_initial_search	cleaning
2_automatic_selection	cleaning

README.md

Building the search sentences

Here, we build the search sentences for different data sources.

Usage

Step 1: Choose your key-words
"Intervention_1" and "Intervent

Step 2: Uncomment one (or se

- scholar_build(full_s
- multiple search sentences
- available), or we set full
- ieee_build() for the IE
- acm_build() for the AC
- arxiv_build() for arXi
- paper_build() for a ba

Step 3: Run the main script as

README.md

Initial Search

Here, we pr

Usage

Step 1: Mak

excluding_words.ipynb 118.17 KiB

Selection phase 1: semi-automatic based on

This notebook is used to find excluding words, i.e., words (among the word one of these words off-topic.

1_load_selected_papers.ipynb 23.48 KiB

Post-processing papers after selection by hand

papers that have been selected by hand

papers from the initial pool

both pools of papers

new excel file appropriate to do a classification of all papers as: to

as pd

2_load_classified_papers.ipynb 17.84 KiB

Post-processing the classified papers

All selected papers have been classified as follows: tool creation/no tool crea

postprocess the classification.

Papers will then be submitted to the data extraction phase. Depending on the

four groups of papers:

1. tool creation & for AI
2. no tool creation & for AI
3. tool creation & not for AI
4. no tool creation & not for AI

Importing libraries:

```
In [1]: import pandas as pd
```

	A	B	C	D	E	F
1	id	title	pub_info	date	link	full_text_id
2	311	Eco2AI: carbon e	Budenny, S	2022	https://ar	budenny2022
3	393	Carbontracker: T	Anthony, La	2020	https://ar	anthony2020
4	1891	Green algorithms	Lannelongue	2021	https://on	lannelongue2021
5	339	Quantifying the c	Lacoste, Ale	2019	https://ar	lacoste2019
6	340	Towards the syst	Henderson,	2020	https://dl.	henderson2020
7	1672	Energy Consump	Lahmer, Sey	2022	https://iee	lahmer2022
8	1368	POMMEL: Explor	Montgomer	2021	https://iee	montgomerie-corcoran2021
9	1695	Fast Yet Accurat	Dariol, Quer	2023	https://dl.	dariol2023
10	1091	Understanding th	Carastan-Sa	2022	https://lin	carastan-santos2022
11	2661	Energy Usage Re	Lottick, Kad	2019	http://arxi	lottick2019
12	1179	EnergyNN: Energ	Goel, Shikha	2021	https://iee	goel2021
13	406	Trends in AI infer	Desislavov,	2023	https://lin	desislavov2023
14	2660	CUMULATOR: T	Technical Tri	2020	https://lin	technical2020

Current state of the selection

Conclusion

Add a new ML computing task script + additional modifications:

The screenshot shows a GitLab repository interface for a project named 'EnergyCalculatorsEvaluation'. The left sidebar contains navigation options like 'Pinned', 'Manage', 'Plan', 'Code', 'Merge requests', 'Repository', 'Branches', and 'Commits'. The main area displays a table of experiments:

Name	Last commit
exp-0-idle	clean
exp-1-MNIST	corrections and reorganization
exp-2-CIFAR10	corrections and reorganization
exp-3-resnet18	fix tapo firmware update + docu...
exp-4-bert-squad	

The screenshot shows a file explorer view of a directory structure. A folder named 'exp-1-MNIST' is highlighted. Below it, other folders are visible: 'exp-2-CIFAR10', 'exp-3-resnet18', and 'exp-4-bert-squad'.

```
20 declare -A scripts=( ['mnist'  
21 ['cifar10']='exp-2-CIFAR10/CI  
22 ['CUB_200_2011']='exp-3-resne  
23 ['image_net']='exp-3-resnet18  
24 ['SQUAD-extracted']='exp-4-be  
25 ['SQUAD-v1-1']='exp-4-bert-sq  
26 ['idle']='exp-0-idle/idle' )  
27
```

```
65 if [[ $exp == 'mnist' ]]; then  
66   if [[ $dev_test == 'True' ]]; then  
67     opt7="--nb_batch_inferences=10"  
68     opt8="--epochs=2"  
69     opt="${opt} ${opt7} ${opt8}"  
70   fi  
71   if [[ $ml == 'training' ]]; then  
72     mkdir "${path_logs_and_results}/${i}/${exp}_model"  
73     opt9="--output_dir=${path_logs_and_results}/${i}/${exp}_model"  
74     opt="${opt} ${opt9}"  
75   fi  
76 fi  
77
```

Provide the path to the script

Specify options, ...

Specific preparation steps

```
160 def prepare_calculator(exp):  
161     # Preparing the calculators #  
162     tracker = None  
163     if exp.name_calc == 'code_ca  
164     output_file = exp.cc_out  
165     tracking_mode = 'machine'  
166     # tracking_mode = 'proce  
167     if exp.online == True:  
168         measure_power_secs =  
169         tracker = EmissionsT  
170         output_file = su
```

How to save the outputs, ...

```
fct_for_experiments.py 14.62 KiB  
1 import sys  
2 import os  
3 _path = '.'  
4 sys.path.append(os.path.join(_path))  
5 import json  
6 import time  
7  
8 # --- FOR CALCULATORS  
9 from codecarbon import EmissionsTracker, OfflineEmissionsTracker  
10 from carbontracker.tracker import CarbonTracker  
11 import os
```

```
fct_for_saving.py 14.67 KiB  
1 import json  
2 import os  
3 import pandas as pd  
4 from datetime import timedelta  
5 import numpy as np  
6 from fct_for_tapo import find_tapo
```

Or add a new evaluation tool or method + additional modifications:



Thank You



Comparison framework: detail

For each of the four ML computing tasks, we do the following experiment:

Repeat 5 times the following: for each evaluation tool or method that we want to test (in a random order), we record the duration of the computing task and energy consumed evaluation (output of the method or tool).

Computer	Alienware Desktop Computer
CPU	Intel Core i9-9900K CPU - 3.60GHz - with 8 cores (2 threads per core: 16 virtual cores) - TDP 95W
RAM	64,0 GB
Integrated GPU	Intel UHD Graphics 630
GPU 1	NVIDIA GeForce RTX 2080 SUPER - TDP 250W
GPU 2	NVIDIA GeForce RTX 2080 SUPER - TDP 250W
OS Version	Ubuntu 22.04.1 LTS (Jammy Jellyfish)
Python	3.10.6