REDUCING THE CARBON FOOTPRINT OF DIGITAL PIPELINES: A CASE STUDY FROM MRI

NICK SOUTER

@NickSouter

UNIVERSITY OF SUSSEX

PREREGISTERED

RESEARCH-ARTICLE FREE ACCESS

Towards the systematic reporting of the energy and carbon footprints of machine learning

Authors: Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, Joelle Pineau

Authors Info & Claims

The Journal of Machine Learning Research, Volume 21, Issue 1 • Article No.: 24

arXiv > cs > arXiv:2007.03051

Computer Science > Computers and Society

[Submitted on 6 Jul 2020]

Carbontracker: Tracking and Predicting the Carbon Footprint of Training Deep Learning Models

Lasse F. Wolff Anthony, Benjamin Kanding, Raghavendra Selvan

Home > Medical Image Computing and Computer Assisted Intervention – MICCAI 2022 > Conference paper

Carbon Footprint of Selecting and Training Deep Learning Models for Medical Image Analysis

Raghavendra Selvan, Nikhil Bhagwat, Lasse F. Wolff Anthony, Benjamin Kanding & Erik B. Dam

Conference paper | First Online: 16 September 2022

ENVIRONMENTAL RESEARCH
COMMUNICATIONS

ACCEPTED MANUSCRIPT • OPEN ACCESS

How to estimate carbon footprint when training deep learning models? A guide and review

Lucia Bouza Heguerte[1], Aurélie Bugeau[2] and Loïc Lannelongue[3]

Accepted Manuscript online 8 September 2023 • © 2023 The Author(s). Published by IOP Publishing Ltd

Journals & Magazines > Computer > Volume: 55 Issue: 7

The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink

Publisher: IEEE Cite This PDF

David Patterson ; Joseph Gonzalez ; Urs Hölzle ; Quoc Le ; Chen Liang ; Lluis-Miquel Munguia ; Daniel Rothchild ; David R. So ; Maud Texier ; …

Li & Chao (2021)

Researchers used a deep learning classification tool with pictures of crop pests.

Reductions in the amount of training data could reduce the compute required.

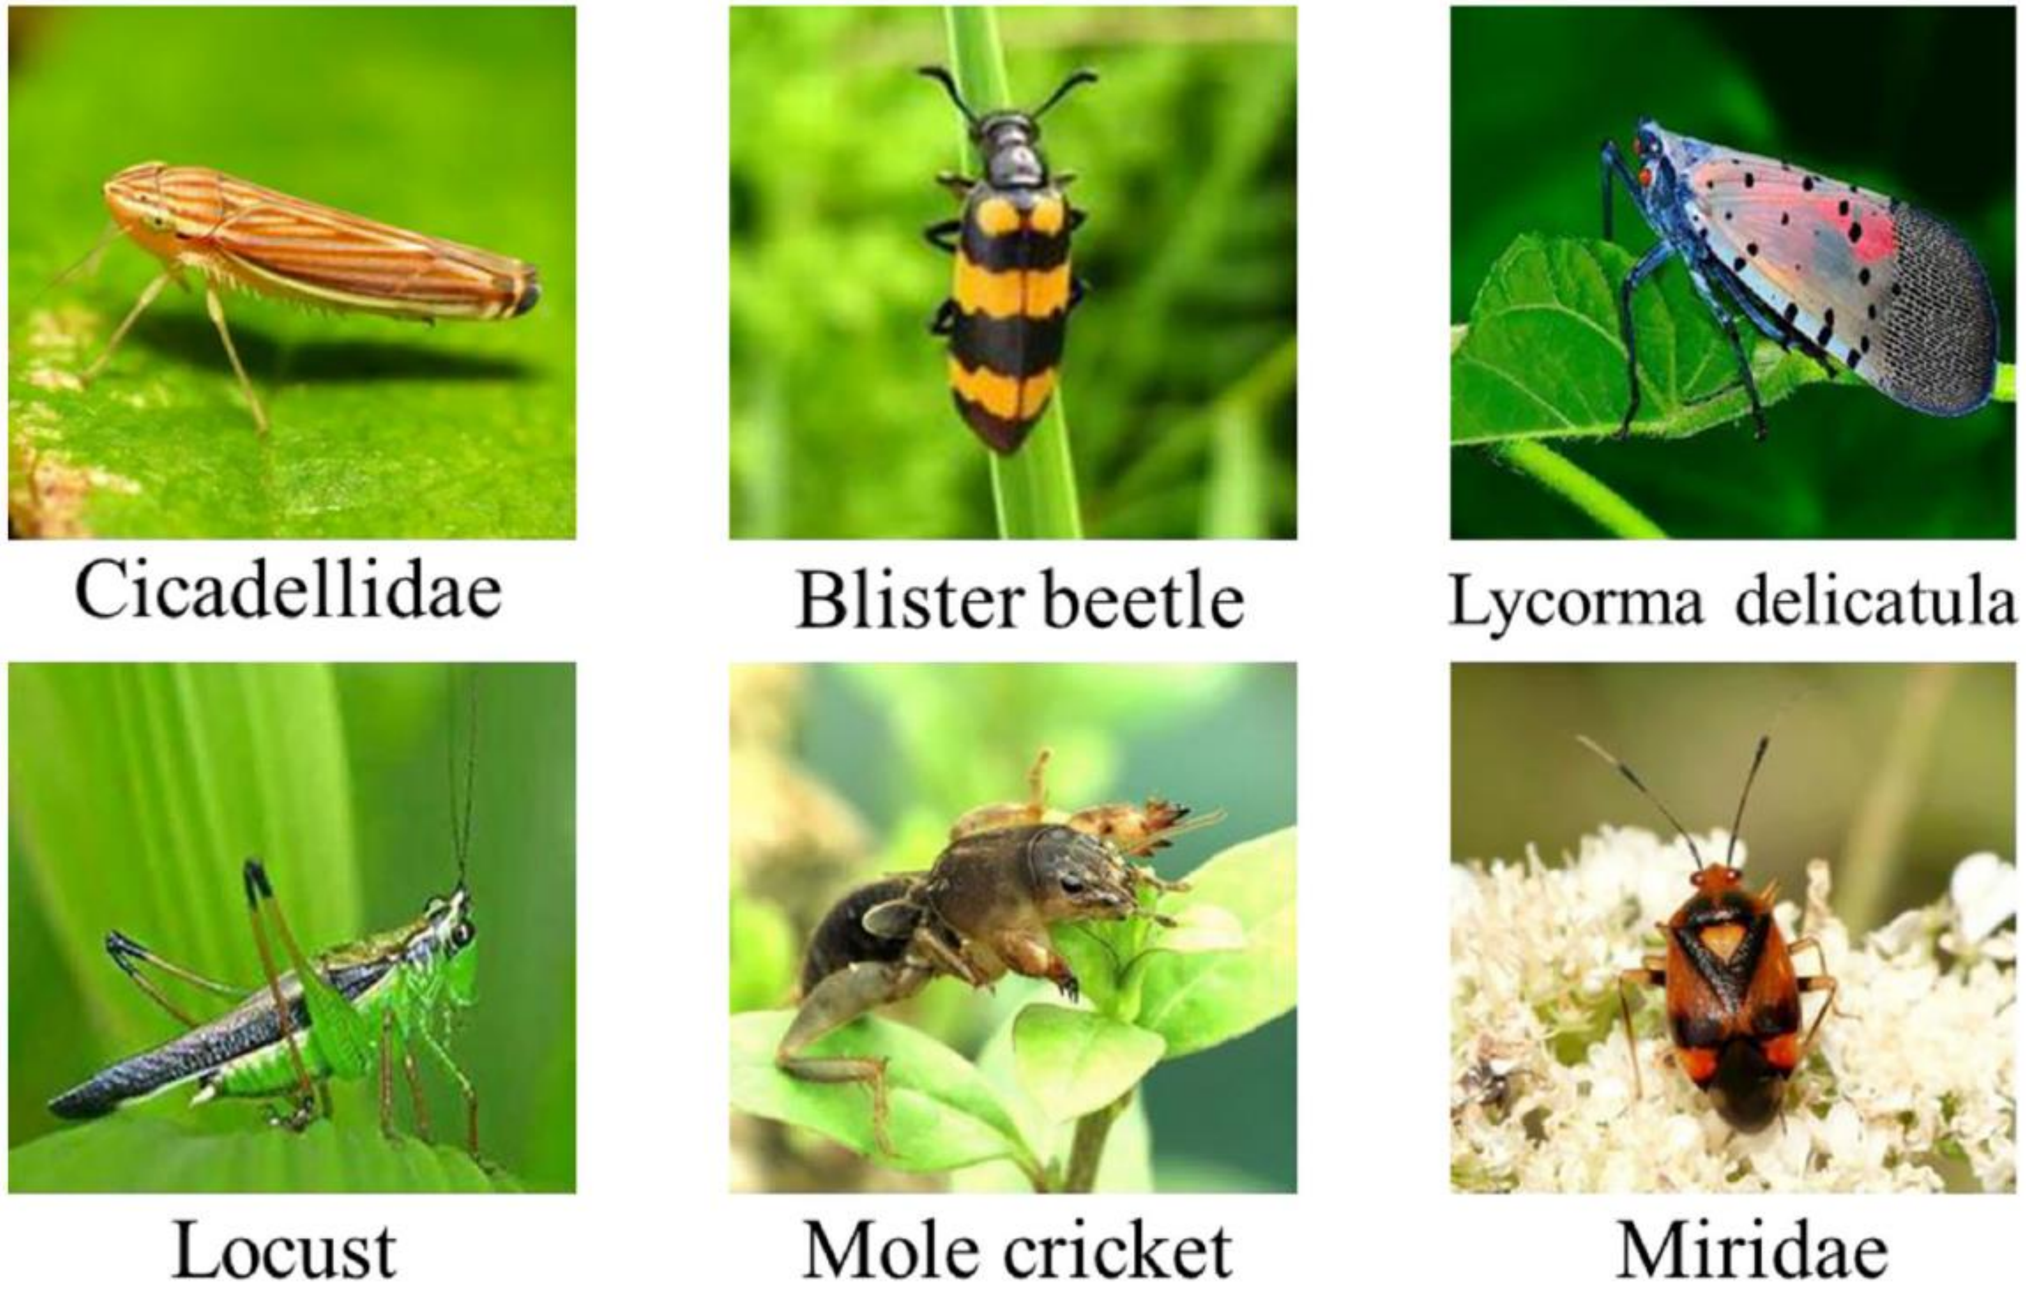

FIGURE 1 | Some samples of the crop pest dataset.

Li & Chao (2021)

When classifying crop pests, peak training accuracy could be achieved with:

- 75% of total data (shallow model)
- 60% of total data (middle model)
- 40% of total data (deep model)

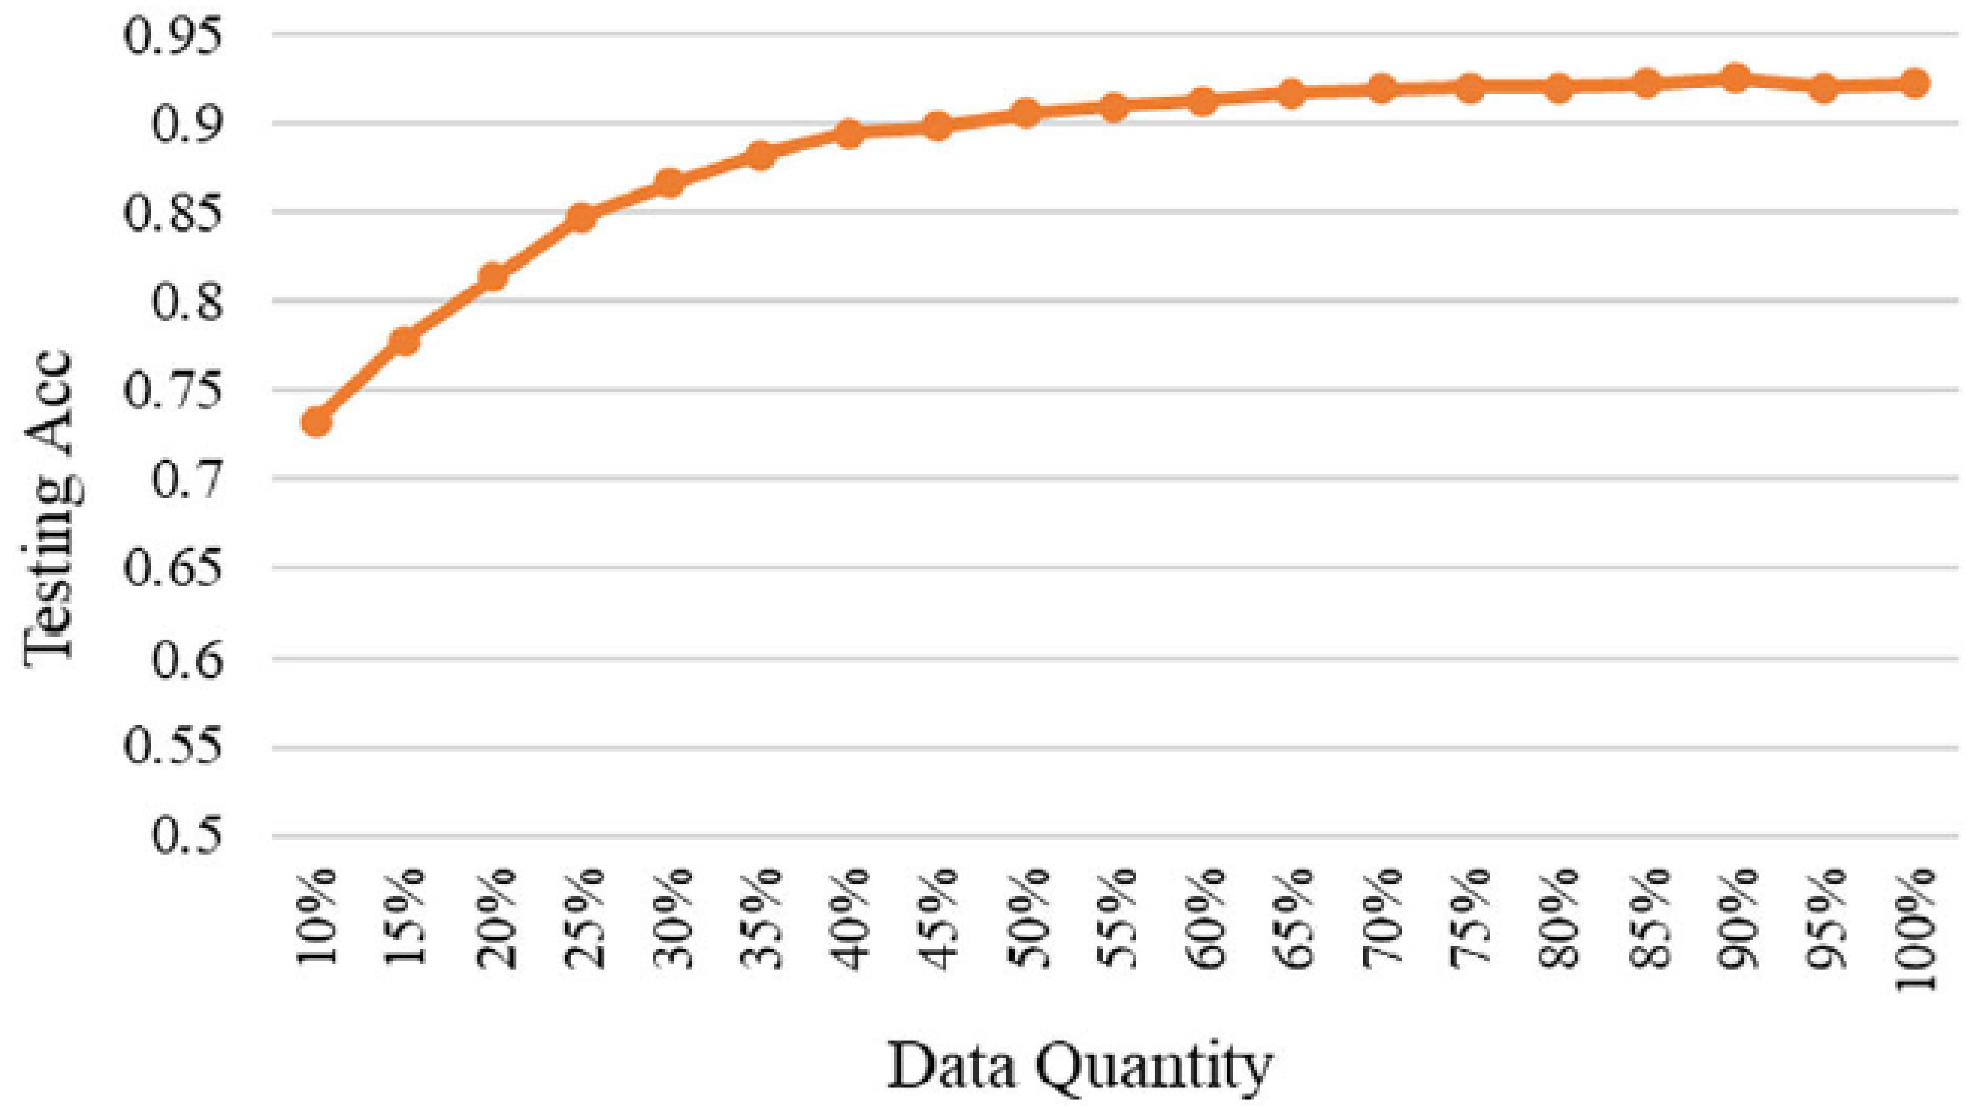

FIGURE 4 | The relation between accuracy and data quantity under the middle model.

CARBON TRACKERS

ONLINE CALCULATOR

EMBEDDED PACKAGE

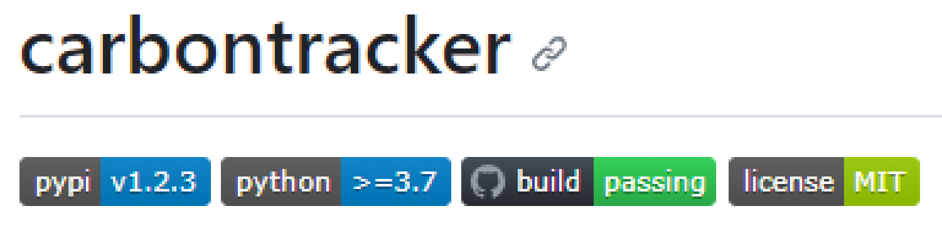

SERVER-SIDE TOOL

Table 1 | Pros and cons of each strategy for carbon footprint estimation

Characteristic of each strategy	Online calculator	Embedded package	Server-side tool
Compatible with any hardware	Yes	No, a dedicated tool needs to be built for each	Yes, provided hardware usage can be monitored
Compatible with any programming language	Yes		Yes
Compatible with any task or research field	Yes		Yes
Computing metrics are collected automatically	No	Yes	Yes
Does not interfere with existing code	Yes	No	Yes
Estimates can be obtained beforehand	Yes	No	No
Estimates can be obtained in real time	No	Yes	Yes
Estimates can be obtained retrospectively	Yes	Only if the tracker was active at the time	Yes
Scalable with large numbers of jobs	No	Yes	Yes
Scalable over long periods of time	No	Only if the tracker is used every time	Yes

LANNELONGUE & INOUYE (2023)

How to reduce the carbon footprint of neuroimaging computing

Souter et al. (2023)

Summary of ten recommendations for reducing the carbon footprint of neuroimaging computing

1. *Make use of existing preprocessed data when possible, instead of acquiring and processing new data*
2. *Preregister a study analysis plan in order to avoid repetitions*
3. *Quantify and report the carbon footprint of your computing using available carbon tracking tools*
4. *Only run the preprocessing and analysis steps that you need*
5. *Run your computing at lower carbon intensity times and in lower carbon intensity locations*
6. *Regularly remove files that you do not need*
7. *Plan where, and for how long, you will store files, aided by research technicians*
8. *Advocate for non-commercial and centralised data storage solutions*
9. *Publicly share sufficient data to ensure it is FAIR (Findable, Accessible, Interoperable, Reusable), but consider the extent of what others will actually need or use*
10. *Discuss the importance of greener computing with other neuroimagers and advocate for systemic change*

Brain extraction

Distortion correction

Registration

Slice-timing correction

Surface reconstruction

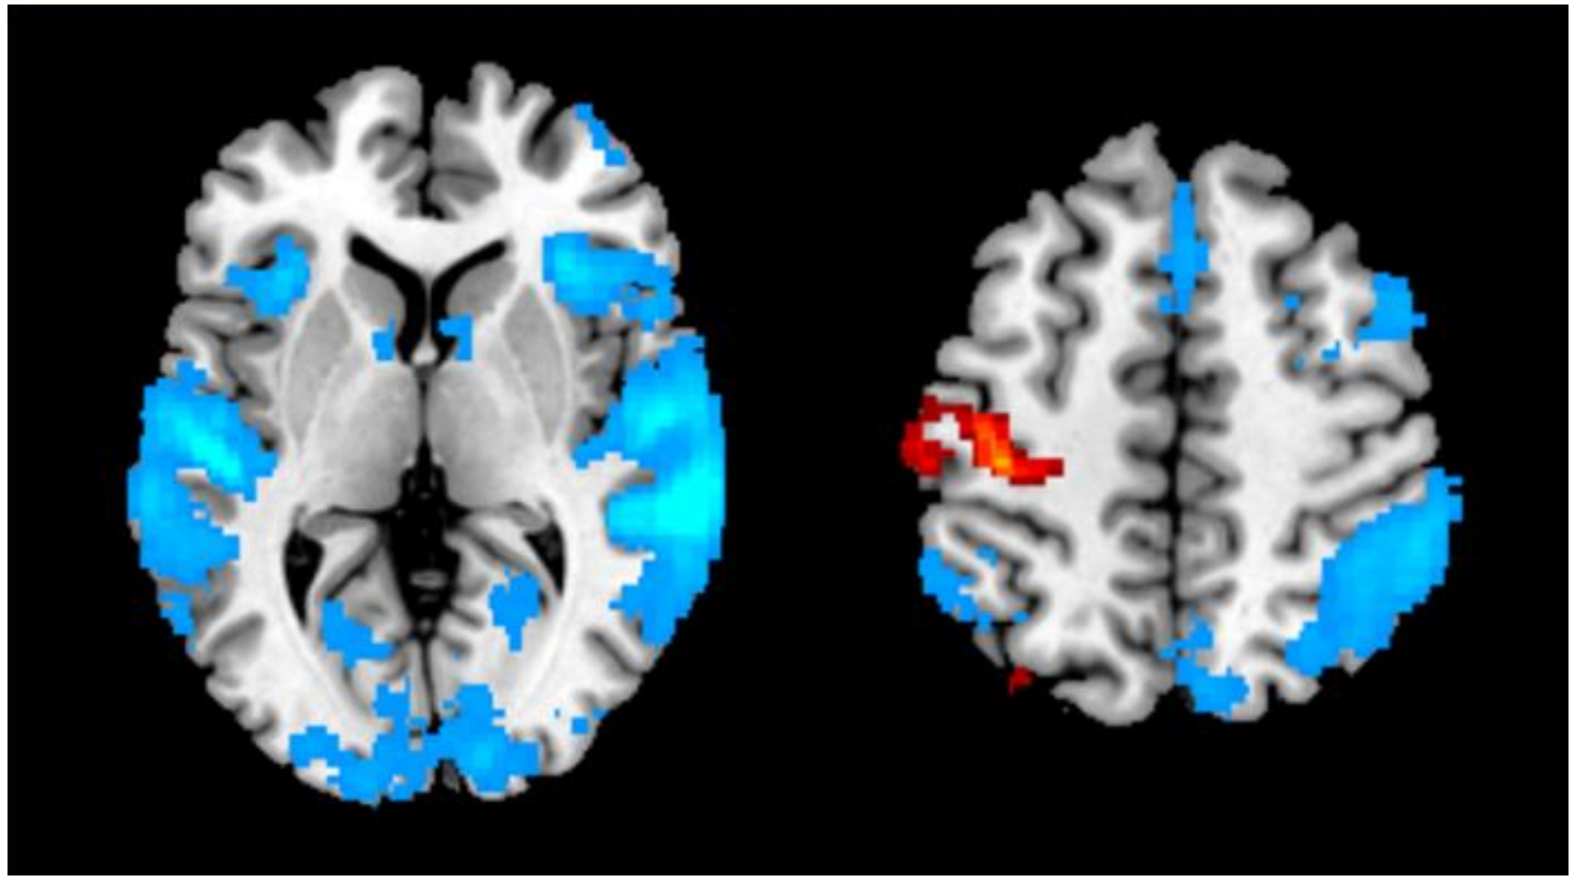

Temporal filtering

Motion correction

Nuisance regressor identification

Smoothing

Spatial filtering

fMRIPrep

```
usage: fmriprep [-h] [--version] [--skip_bids_validation]
                [--participant-label PARTICIPANT_LABEL [PARTICIPANT_LABEL ...]]
                [-t TASK_ID] [--echo-idx ECHO_IDX] [--bids-filter-file FILE]
                [--anat-derivatives PATH] [--bids-database-dir PATH]
                [--nprocs NPROCS] [--omp-nthreads OMP_NTHREADS]
                [--mem MEMORY_MB] [--low-mem] [--use-plugin FILE]
                [--anat-only] [--boilerplate_only] [--md-only-boilerplate]
                [--error-on-aroma-warnings] [-v]
                [--ignore {fieldmaps,slicetiming,sbref,t2w,flair}
[{fieldmaps,slicetiming,sbref,t2w,flair} ...]]
                [--longitudinal] [--output-spaces [OUTPUT_SPACES ...]]
                [--me-output-echos] [--bold2t1w-init {register,header}]
                [--bold2t1w-dof {6,9,12}] [--force-bbr] [--force-no-bbr]
                [--medial-surface-nan] [--slice-time-ref SLICE_TIME_REF]
                [--dummy-scans DUMMY_SCANS] [--random-seed _RANDOM_SEED]
                [--use-aroma]
                [--aroma-melodic-dimensionality AROMA_MELODIC_DIM]
                [--return-all-components]
                [--fd-spike-threshold REGRESSORS_FD_TH]
                [--dvars-spike-threshold REGRESSORS_DVARS_TH]
                [--skull-strip-template SKULL_STRIP_TEMPLATE]
                [--skull-strip-fixed-seed]
                [--skull-strip-t1w {auto,skip,force}] [--fmap-bspline]
                [--fmap-no-demean] [--topup-max-vols TOPUP_MAX_VOLS]
                [--use-syn-sdc [{warn,error}]] [--force-syn]
                [--fs-license-file FILE] [--fs-subjects-dir PATH]
                [--no-submm-recon] [--cifti-output [{91k,170k}] |
                --fs-no-reconall] [--output-layout {bids,legacy}]
                [-w WORK_DIR] [--clean-workdir] [--resource-monitor]
                [--reports-only] [--config-file FILE] [--write-graph]
                [--stop-on-first-crash] [--notrack]
                [--debug {compcor,fieldmaps,pdb,all} [{compcor,fieldmaps,pdb,all} ...]]
                [--sloppy] [--track-carbon] [--country-code COUNTRY_CODE]
                bids_dir output_dir {participant}
```

Pipeline Variants

ID	Label	Flag addition
P0	Baseline	N/A
P1	No FreeSurfer surface reconstruction	-fs-no-reconall
P2	Sloppy preprocessing	-sloppy
P3	Increase parallelisation	-nthreads **16**
P4	Low memory	--low-mem
P5	Add surface output space	--output-spaces **MNI152Lin6Asym:res-2 fsaverage**
P6	Remove parallelisation	--nthreads **1**
P7	ICA AROMA	--use-aroma
P8	Increase output space resolution	--output-spaces **MNI152Lin6Asym:res-1**
P9	Fieldmap-free distortion correction	--use-syn-sdc **error**

Finding the right carbon tracker

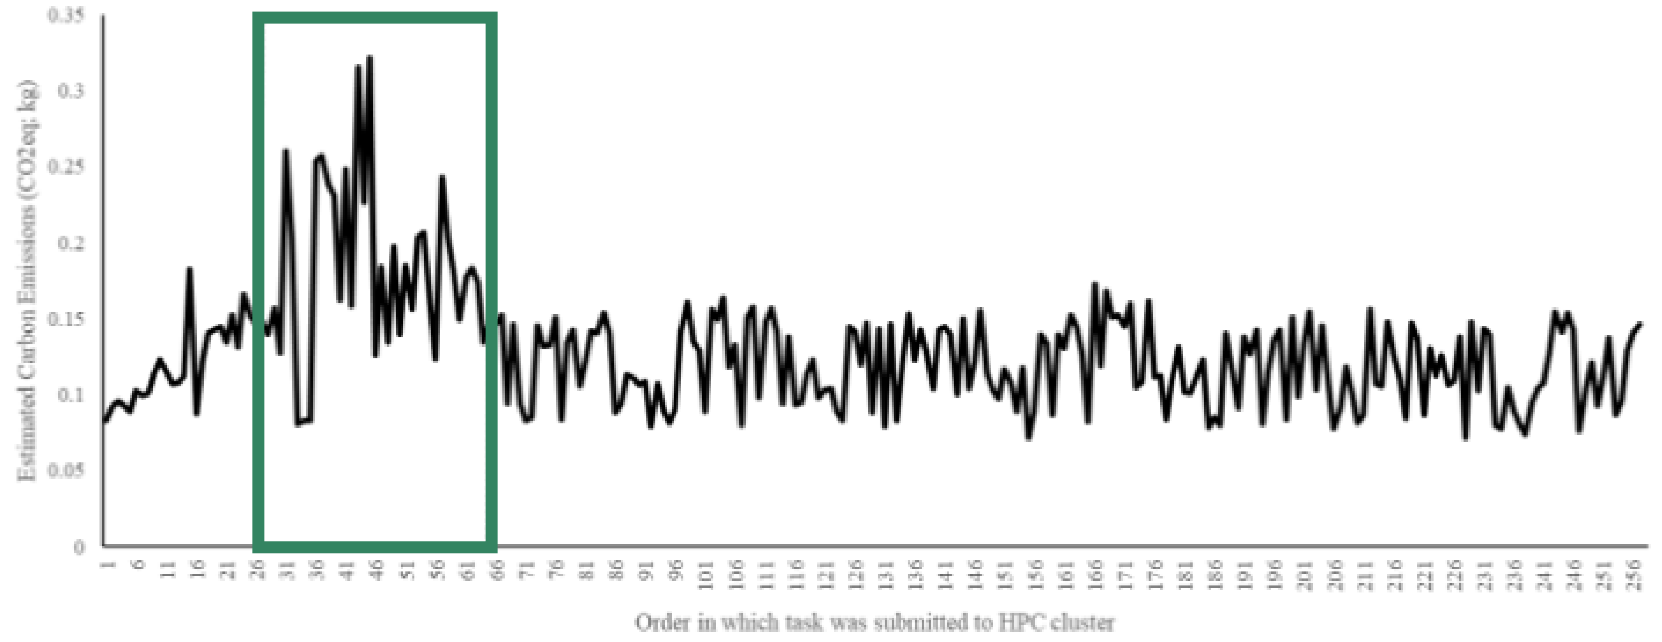

(a) – CodeCarbon

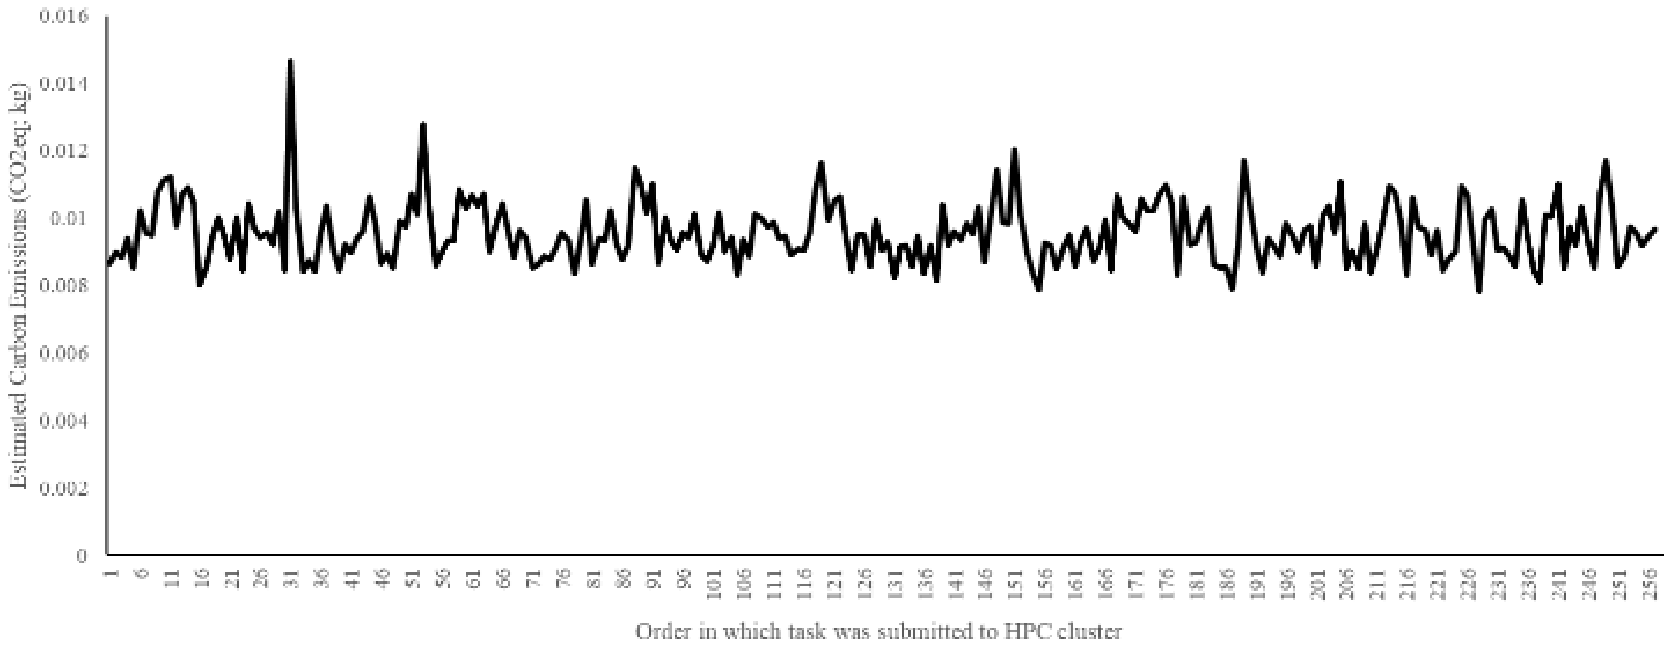

(b) – HPC logs

We experienced issues using **CodeCarbon** - preprocessing would go through periods of surging or dipping emissions with no predictable pattern.

This was an issue of lacking **hardware isolation**, **CodeCarbon** was picking up energy use from other jobs on the same node. We instead switched to a **server-side tool**, similar to **GA4HPC**.

Measuring Performance

As well as minimising carbon emissions, we need to ensure the software is producing **high quality data**.

To quantify **preprocessing performance**, we measured mean **statistical task activation** in ‘regions of interest’ (ROIs) relevant to our data.

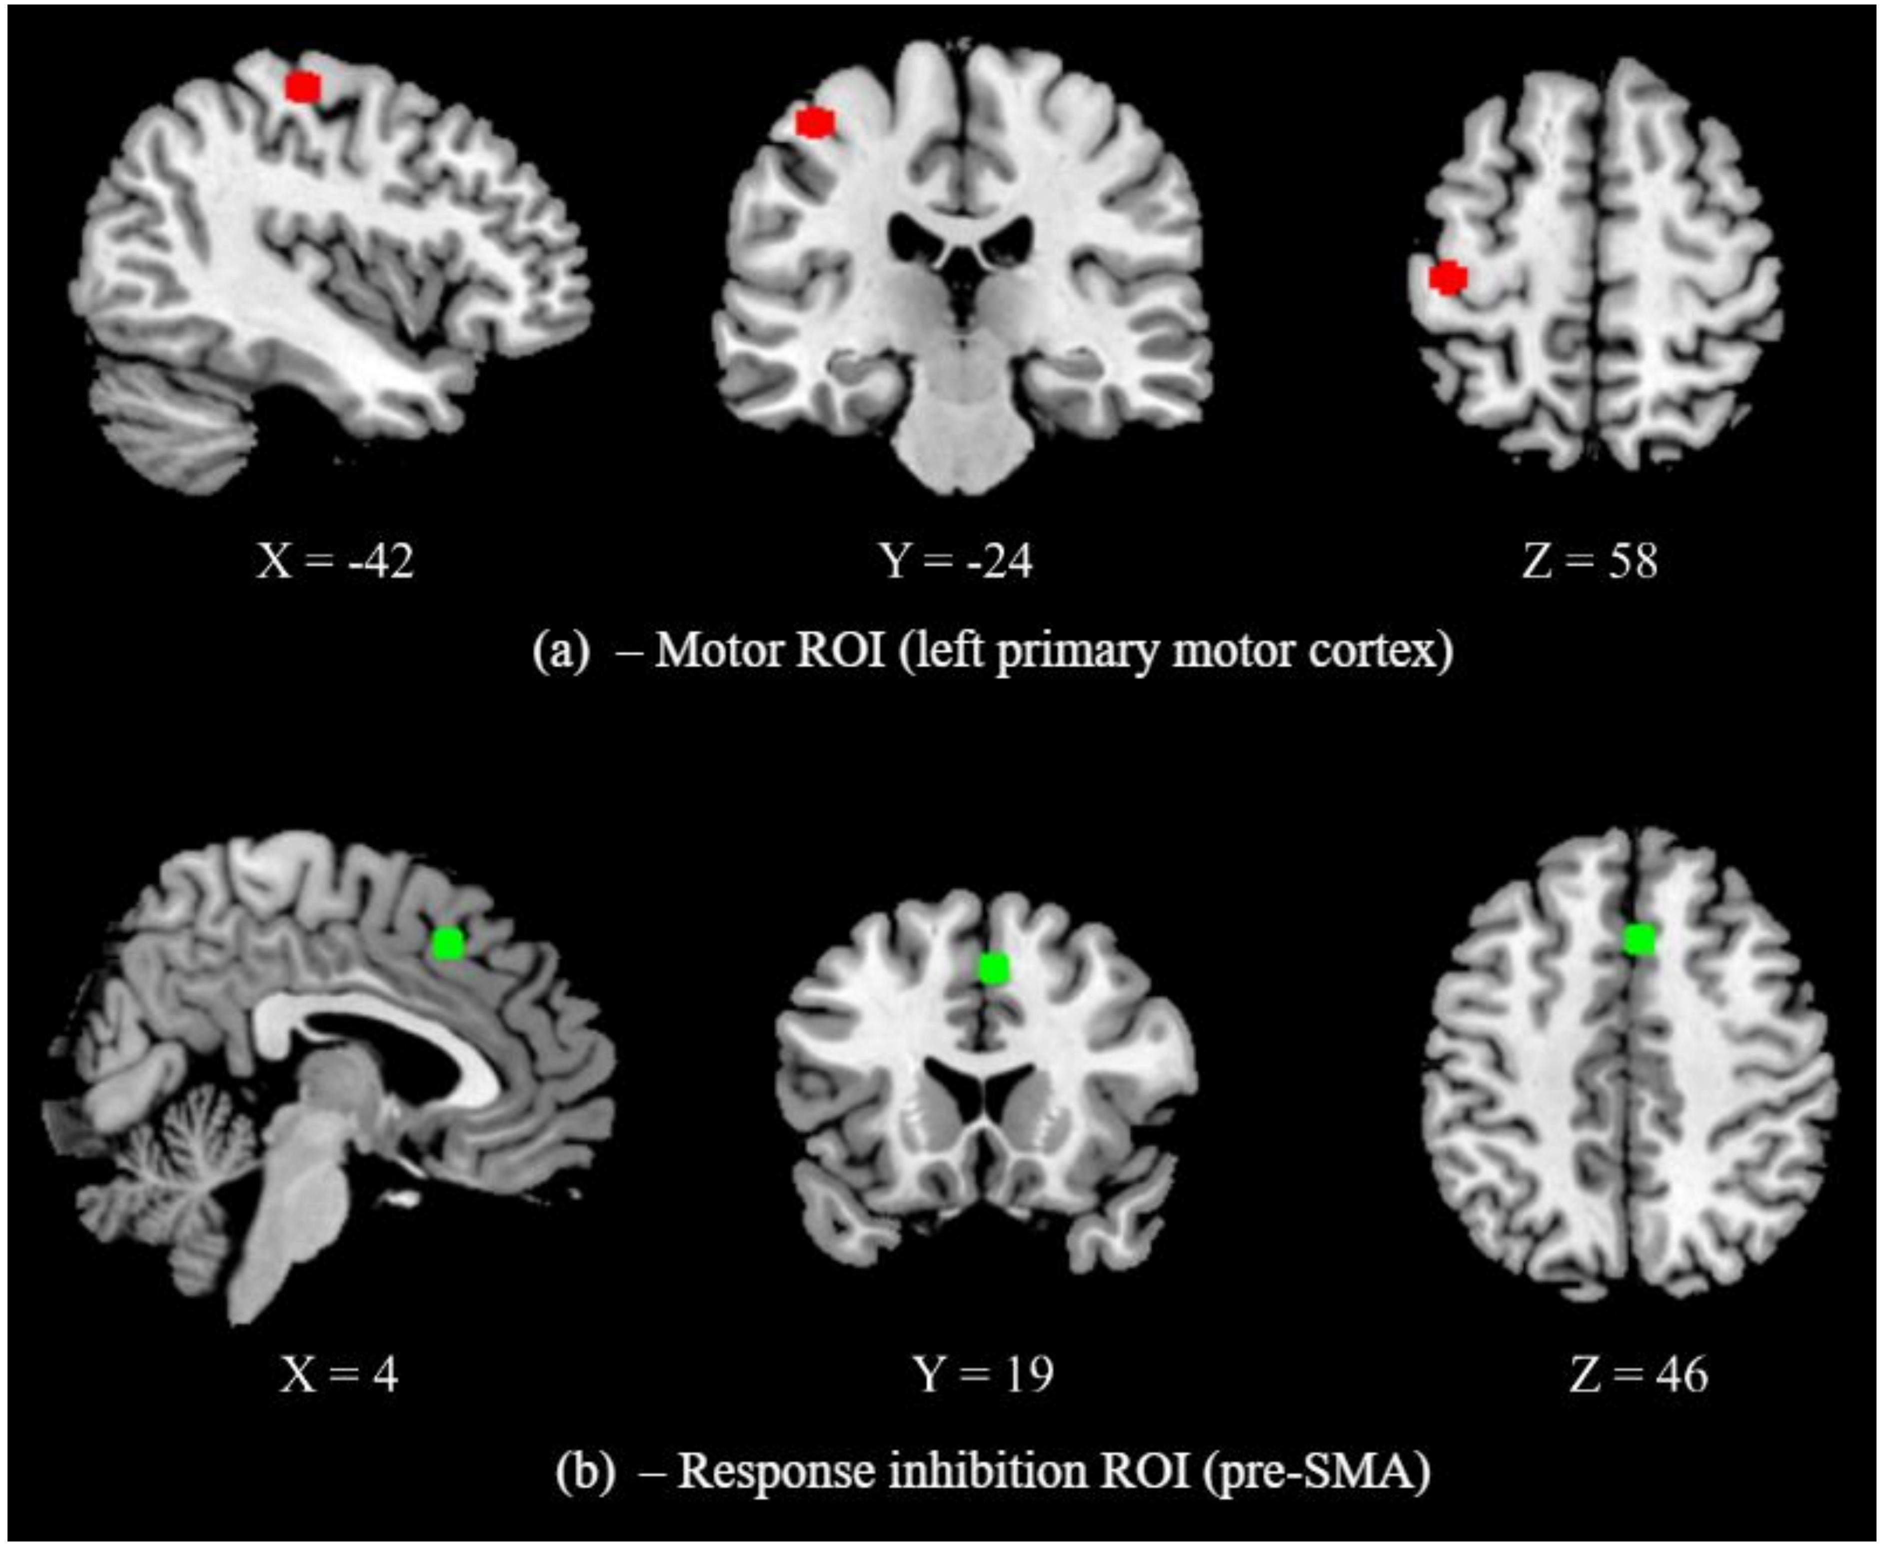

RESULTS

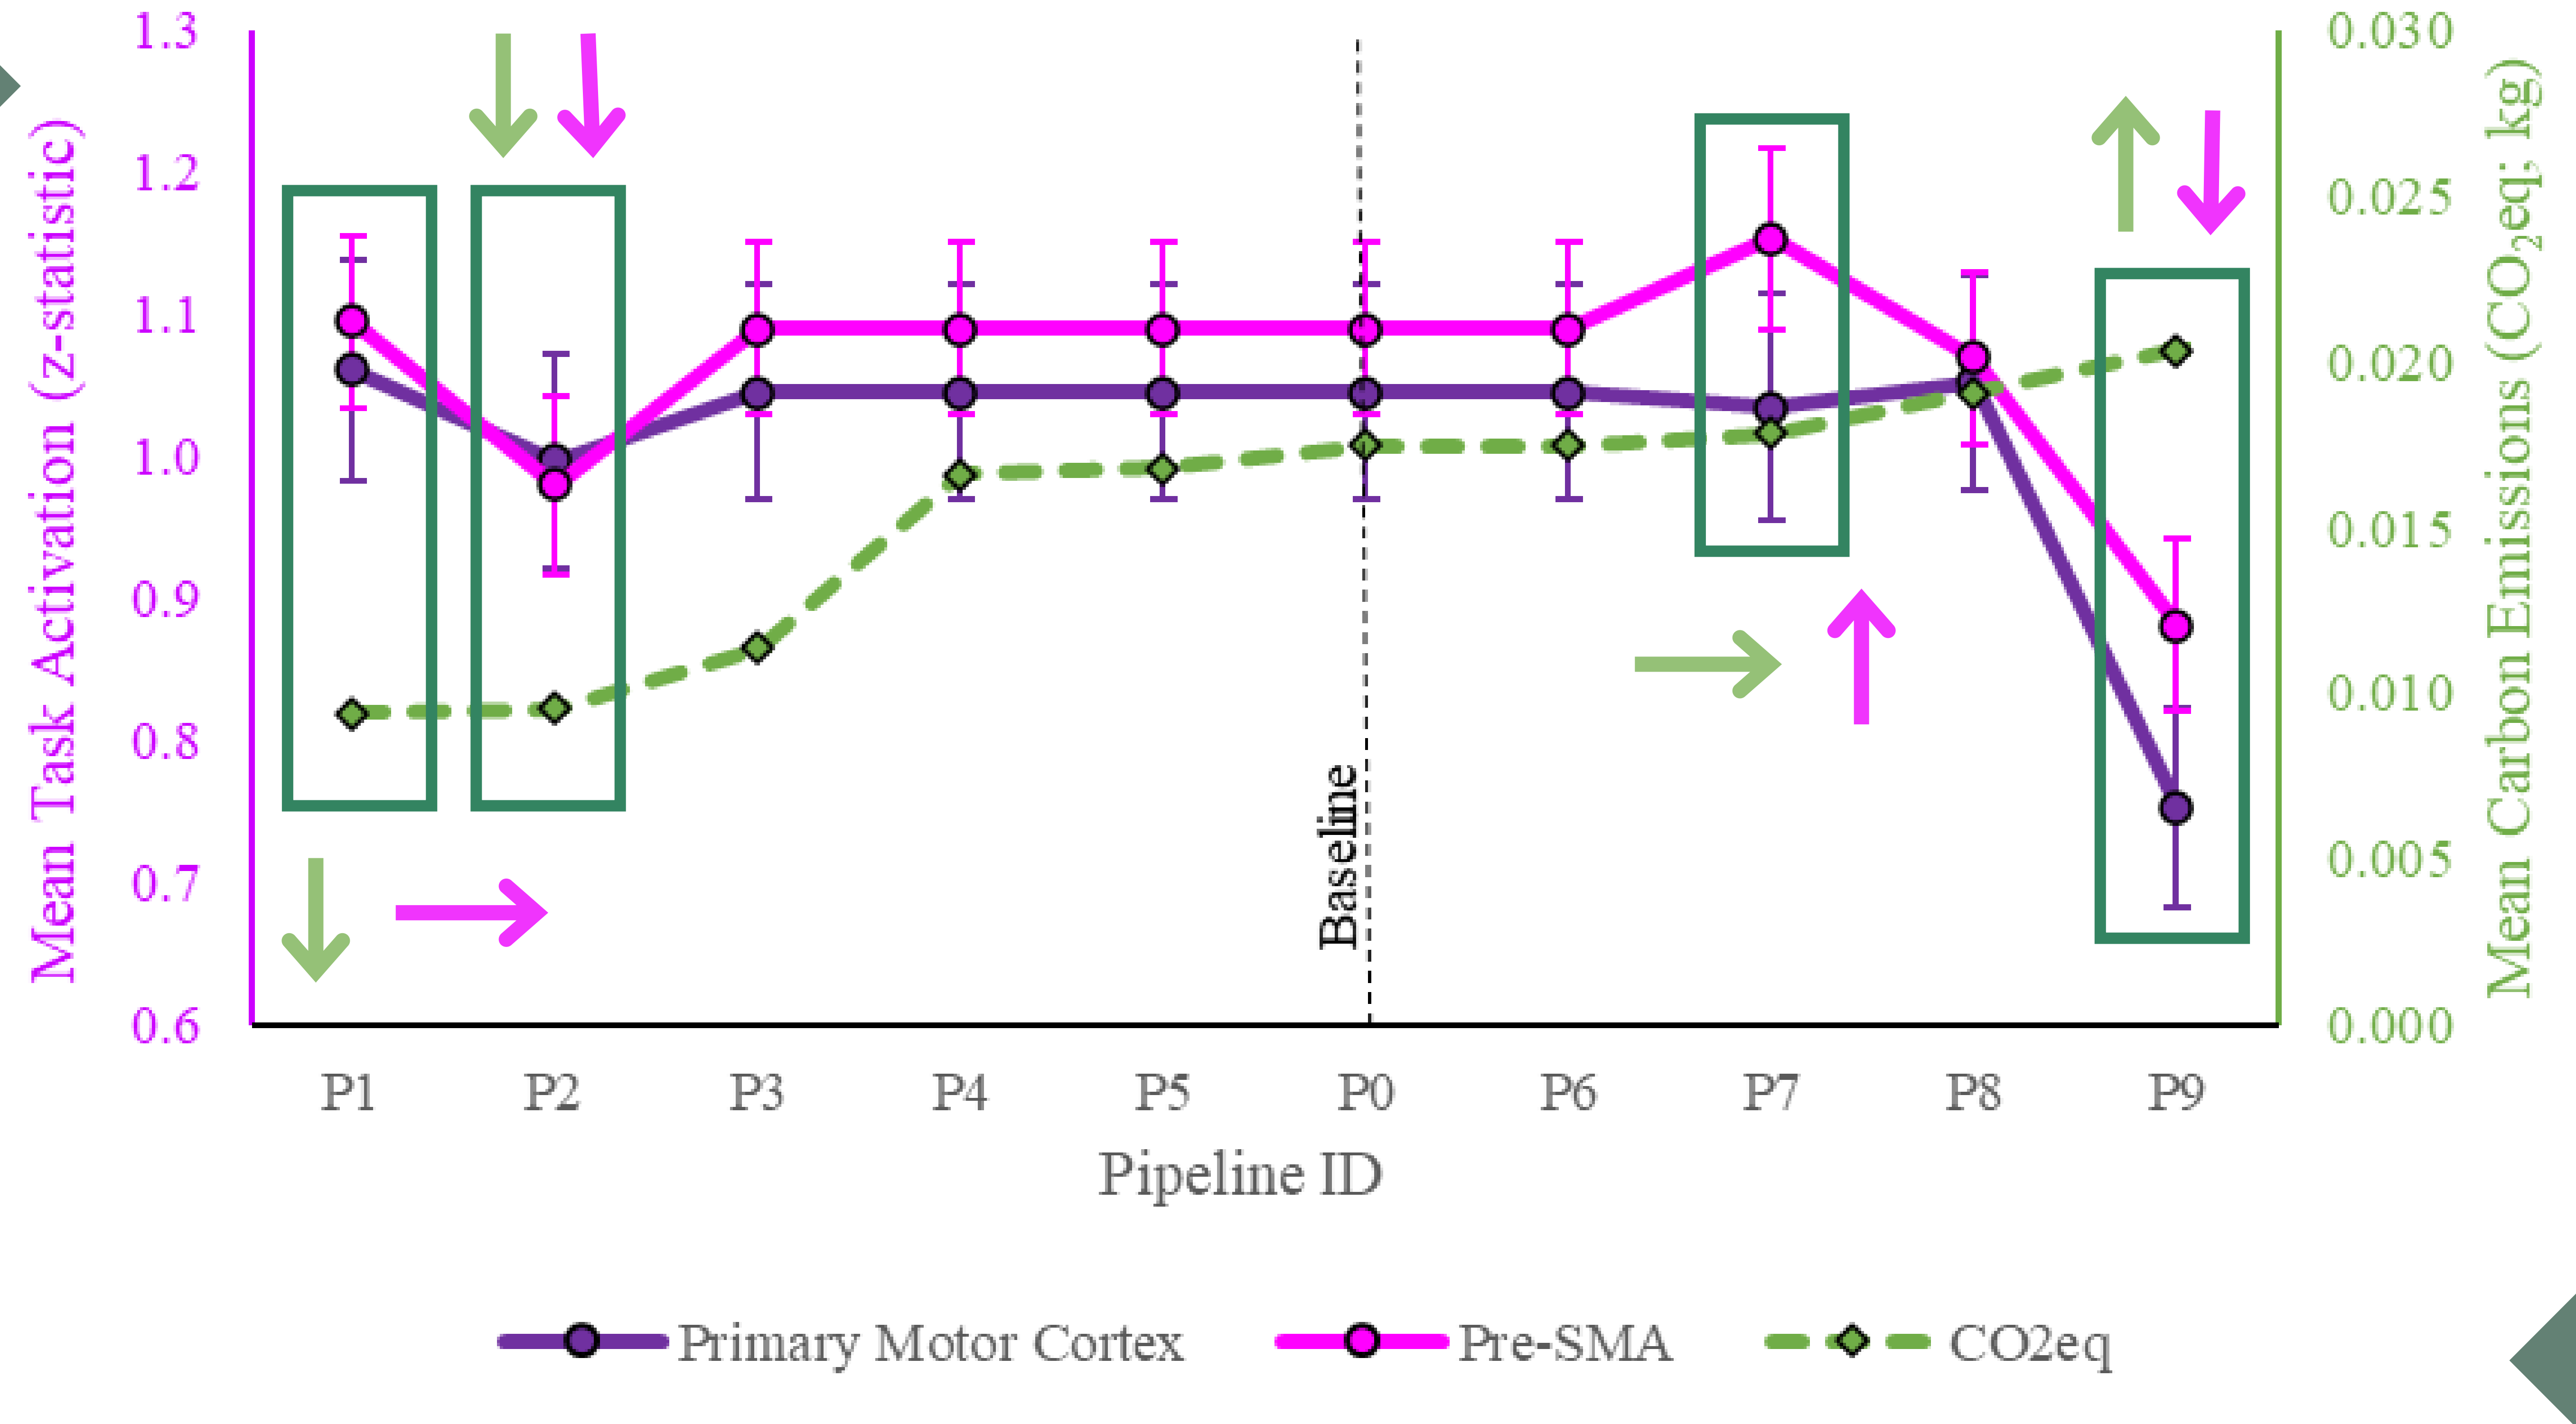

1.3
1.2
1.1
1.0
0.9
0.8
0.7
0.6
Mean Task Activation (z-statistic)
0.030
0.025
0.020
0.015
0.010
0.005
0.000
Mean Carbon Emissions (CO_2eq; kg)
Baseline
P1
P2
P3
P4
P5
P0
P6
P7
P8
P9
Pipeline ID
Primary Motor Cortex
Pre-SMA
CO2eq

OUR RECOMMENDATIONS

Box 1. Summary of recommendations for reducing the carbon footprint of fMRIPrep

1. Disabling FreeSurfer surface reconstruction (--fs-no-reconall) can almost halve one's computing carbon footprint with no trade-offs in performance.
2. Using 'sloppy' registration (--sloppy) can almost halve one's carbon footprint while testing fMRIPrep use but should not be used in preprocessing given performance losses.
3. Increasing parallelisation (--nthreads <number>) of a job can reduce emissions while not impacting performance - benefits in this paper are confounded by hardware efficiency.
4. Implementing low memory mode (--low-mem) can modestly reduce emissions while not impacting performance, although exact energy savings may be HPC cluster-specific.
5. Implementing ICA AROMA (--use-aroma) can benefit sensitivity to statistical activation without increasing carbon emissions.
6. Increasing the resolution of the volumetric output space (--output-spaces MNI152NLin6Asym:res-<number>) increases emissions while not benefitting task sensitivity but somewhat benefitting data smoothness.
7. fMRIPrep's experimental fieldmap-free distortion correction technique (--use-syn-sdc) both increases estimated emissions and degrades data quality. This flag should be avoided.

fMRIPrep Junk Files

Up to 96% of fMRIPrep output files can be considered junk (not needed in subsequent analysis). Reducing the size of resulting data can decrease server workload, and therefore energy usage.

For open-source code to automatically remove these files:

fMRIPrepCleanup
https://github.com/NickESouter/fMRIPrepCleanup

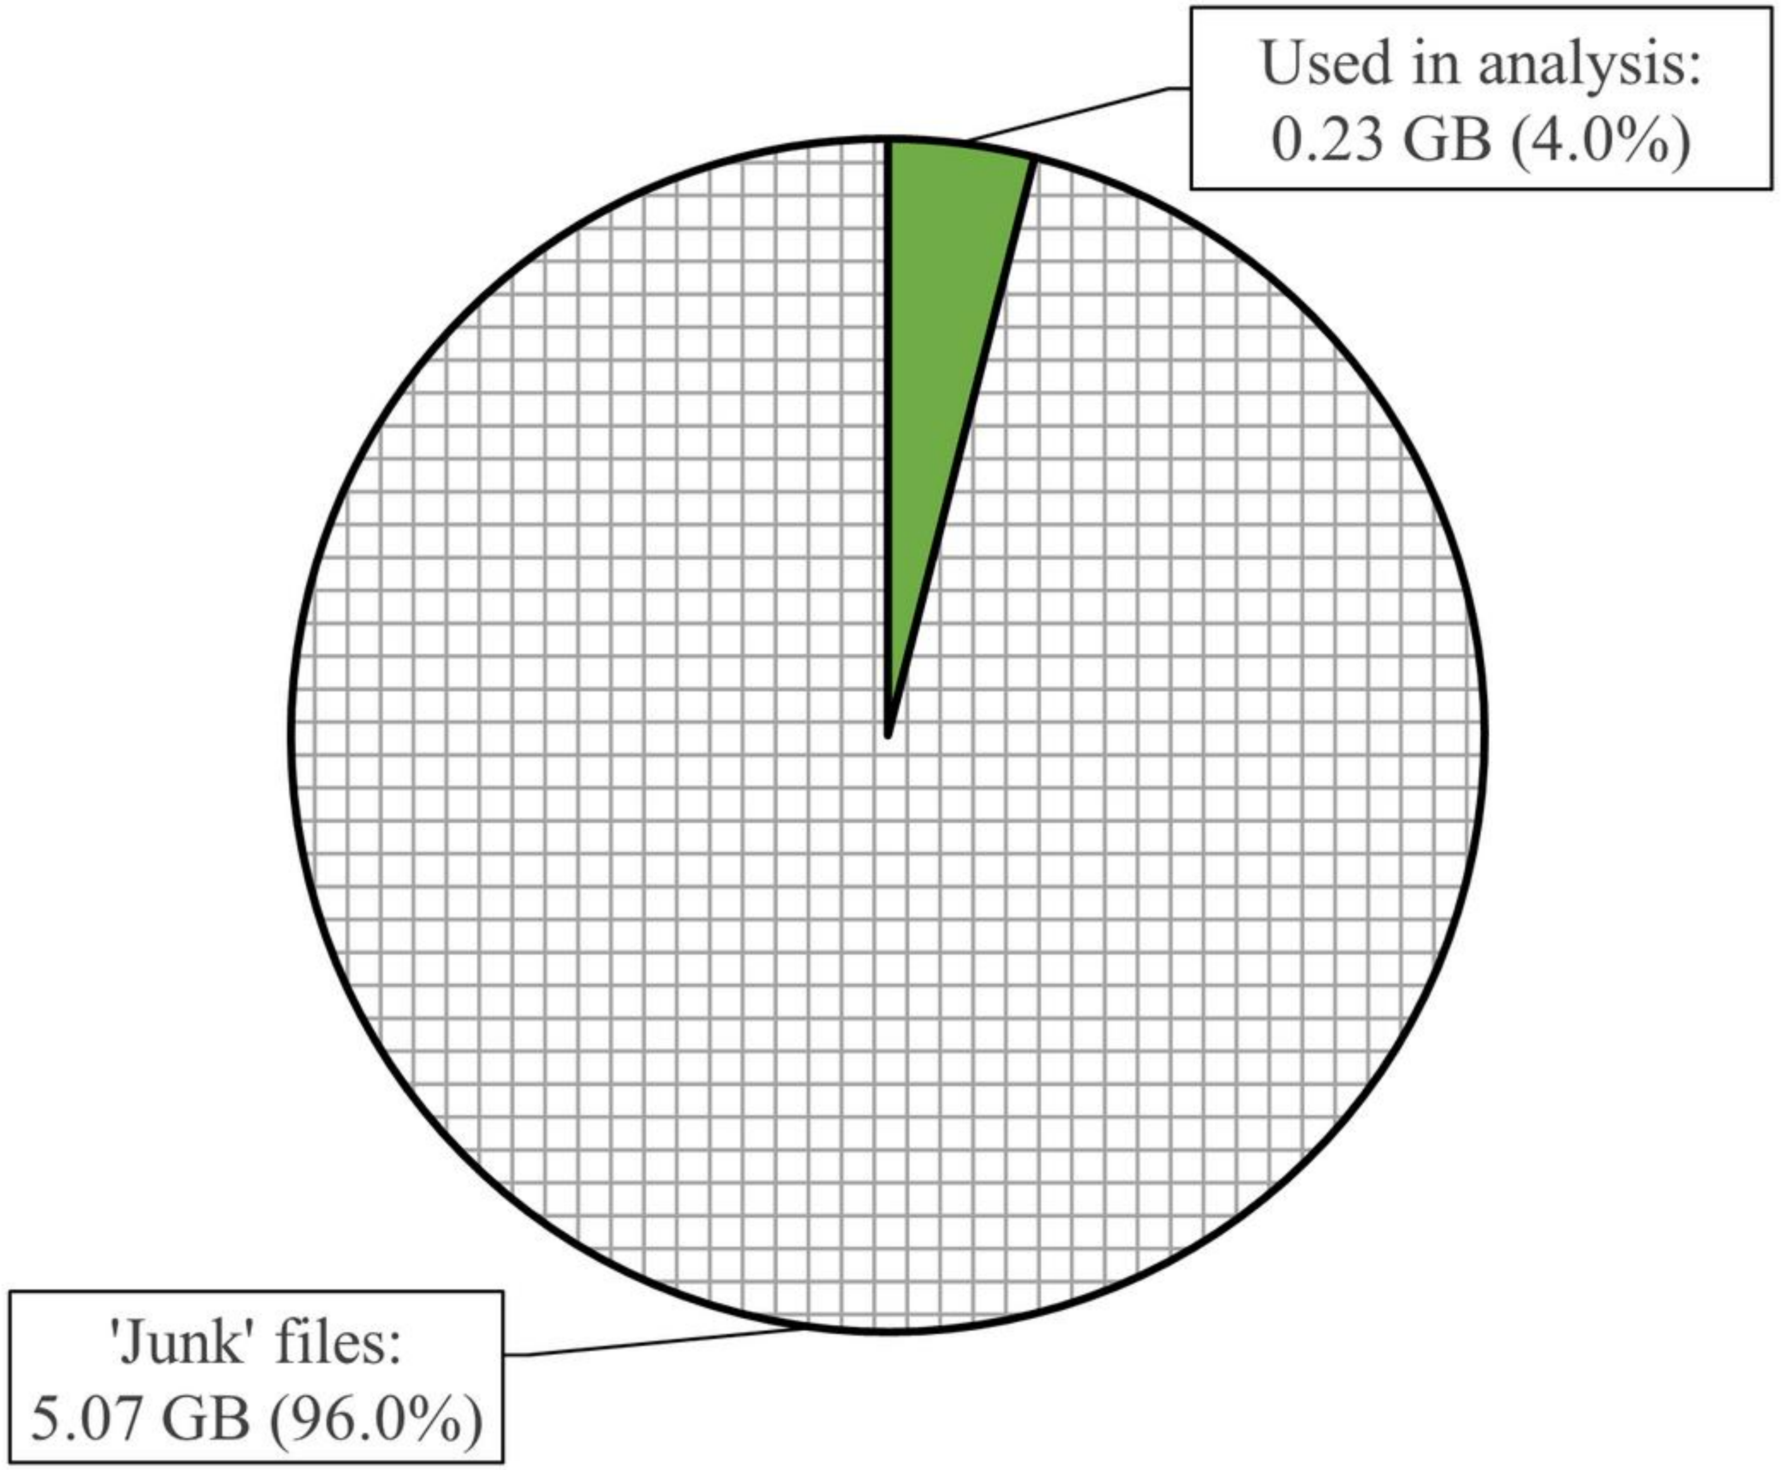

Lessons learned

- For a given digital tool, there may be low-hanging fruit for drastically reducing emissions
- Choosing the right carbon tracker can be complicated - check that it's working as expected
- Quantifying performance of a tool can be complicated - check for existing best practices
- Neuroimagers appear to understand the importance of this topic
- (Some) HPC teams are interested in helping to facilitate carbon tracking
- Digital tools can produce a large amount of unneeded files

Thank you!

Collaborators

- Charlotte Rae
- Gabby Samuel
- Loïc Lannelongue
- Lincoln Colling
- Raghav Selvan
- Nikhil Bhagwat
- Chris Racey
- Reese Wilkinson
- Niall Duncan

To stay up to date with sustainability in neuroimaging:

OHBM Sustainability & Environmental Action
Special Interest Group

For open-source green neuroimaging computing code:

NickESouter

References

- Lannelongue, L., & Inouye, M. (2023). Carbon footprint estimation for computational research. *Nature Reviews Methods Primers, 3*, 9. https://doi.org/10.1038/s43586-023-00202-5
- Li, Y., & Chao, X. (2021). Toward sustainability: Trade-off between data quality and quantity in crop pest recognition. *Frontiers in Plant Science, 12*, 811241. https://doi.org/10.3389/fpls.2021.811241
- Souter, N., Lannelongue, L., Samuel, G., Racey, C., Colling, L., Bhagwat, N., Selvan, R., & Rae, C. (2023). Ten recommendations for reducing the carbon footprint of research computing in human neuroimaging. *OSF Preprints.* https://doi.org/10.31219/osf.io/7q5mh