

# Green Algorithms: Quantifying the Carbon Footprint of Computation

Loïc Lannelongue,\* Jason Grealey, and Michael Inouye\*

Climate change is profoundly affecting nearly all aspects of life on earth, including human societies, economies, and health. Various human activities are responsible for significant greenhouse gas (GHG) emissions, including data centers and other sources of large-scale computation. Although many important scientific milestones are achieved thanks to the development of high-performance computing, the resultant environmental impact is underappreciated. In this work, a methodological framework to estimate the carbon footprint of any computational task in a standardized and reliable way is presented and metrics to contextualize GHG emissions are defined. A freely available online tool, Green Algorithms ([www.green-algorithms.org](http://www.green-algorithms.org)) is developed, which enables a user to estimate and report the carbon footprint of their computation. The tool easily integrates with computational processes as it requires minimal information and does not interfere with existing code, while also accounting for a broad range of hardware configurations. Finally, the GHG emissions of algorithms used for particle physics simulations, weather forecasts, and natural language processing are quantified. Taken together, this study develops a simple generalizable framework and freely available tool to quantify the carbon footprint of nearly any computation. Combined with recommendations to minimize unnecessary CO<sub>2</sub> emissions, the authors hope to raise awareness and facilitate greener computation.

## 1. Introduction

The concentration of greenhouse gases (GHGs) in the atmosphere has a dramatic influence on climate change with both global and locally focused consequences, such as rising sea levels, devastating wild-fires in Australia, extreme typhoons in the Pacific, severe droughts across Africa, as well as repercussions for human health.

With 100 megatonnes of CO<sub>2</sub> emissions per year (Note S1, Supporting information), similar to American commercial aviation, the contribution of data centers and high-performance computing facilities to climate change is substantial. So far, rapidly increasing demand has been paralleled by increasingly energy-efficient facilities, with overall electricity consumption of data centers somewhat stable. However, this stability is likely to end in the coming years, with a best-case scenario forecasting a three-fold increase in the energy needs of the sector.<sup>[1,2]</sup>

Advances in computation, including those in hardware, software, and

L. Lannelongue, Prof. M. Inouye  
Cambridge Baker Systems Genomics Initiative  
Department of Public Health and Primary Care  
University of Cambridge  
Cambridge CB1 8RN, UK  
E-mail: LL582@medschl.cam.ac.uk; minouye@baker.edu.au;  
mi336@medschl.cam.ac.uk

L. Lannelongue, Prof. M. Inouye  
British Heart Foundation Cardiovascular Epidemiology Unit  
Department of Public Health and Primary Care  
University of Cambridge  
Cambridge CB1 8RN, UK

L. Lannelongue, Prof. M. Inouye  
Health Data Research UK Cambridge  
Wellcome Genome Campus and University of Cambridge  
Cambridge CB10 1SA, UK


J. Grealey, Prof. M. Inouye  
Cambridge Baker Systems Genomics Initiative  
Baker Heart and Diabetes Institute  
Melbourne, Victoria 3004, Australia

J. Grealey  
Department of Mathematics and Statistics  
La Trobe University  
Melbourne 3086, Australia

Prof. M. Inouye  
British Heart Foundation Centre of Research Excellence  
University of Cambridge  
Cambridge CB1 8RN, UK

Prof. M. Inouye  
National Institute for Health Research Cambridge Biomedical Research Centre  
University of Cambridge and Cambridge University Hospitals  
Cambridge CB2 0QQ, UK

Prof. M. Inouye  
The Alan Turing Institute  
London NW1 2DB, UK

 The ORCID identification number(s) for the author(s) of this article can be found under <https://doi.org/10.1002/adv.202100707>

© 2021 The Authors. Advanced Science published by Wiley-VCH GmbH. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

DOI: 10.1002/adv.202100707

algorithms, have enabled scientific research to progress at unprecedented rates. Weather forecasts have increased in accuracy to the point where 5-day forecasts are approximately as accurate as 1-day forecasts 40 years ago,<sup>[3]</sup> physics algorithms have produced the first direct image of a black hole 55 million light-years away,<sup>[4–6]</sup> the human genome has been mined to uncover thousands of genetic variants for disease,<sup>[7]</sup> and machine learning (ML) permeates many aspects of society, including economic and social interactions.<sup>[8–11]</sup> An example of the scale of computational research in science is the Extreme Science and Engineering Discovery Environment (XSEDE) in the USA. In 2020 only, almost 9 billion compute hours have been used for scientific computing,<sup>[12]</sup> a pace of 24 million h per day. Yet, the costs associated with large-scale computation are not being fully captured.

Power consumption results in GHG emissions and the environmental costs of performing computations using data centers, personal computers, and the immense diversity of architectures are unclear. While programmes in green computing (the study of environmentally responsible information and communications technologies) have been developed over the past decade, these mainly focus on energy-efficient hardware and cloud-related technologies.<sup>[13–15]</sup>

With widely recognized power-hungry and expensive training algorithms, deep learning has begun to address its carbon footprint. ML models have grown exponentially in size over the past few years,<sup>[16]</sup> with some algorithms training for thousands of core-hours, and the associated energy consumption and cost have become a growing concern.<sup>[17]</sup> In natural language processing (NLP), Strubell et al.<sup>[18]</sup> found that designing and training translation engines can emit between 0.6 and 280 tonnes of CO<sub>2</sub>. While not all NLP algorithms require frequent retraining, algorithms in other fields are run daily or weekly, multiplying their energy consumption. Astronomy also relies largely on supercomputers to analyse data, which has motivated some investigations into the carbon footprint of the field.<sup>[19,20]</sup> For example, it has been estimated that the usage of supercomputers by Australian astronomers emits 15 kilotonnes of CO<sub>2</sub> per year, equivalent to 22 tonnes per researcher.<sup>[21]</sup> Cryptocurrencies, and their so-called “mining farms,” have also seen their environmental impact increase exponentially in recent years, and several reports have shed doubts on their sustainability. A 2018 study estimated the yearly energy usage of Bitcoin to be 46 TWh, resulting in 22 Mt of CO<sub>2</sub> released in the atmosphere.<sup>[22]</sup> In March 2021, Bitcoin usage is estimated to be 130 TWh, which, if Bitcoin was a country, would rank its energy usage in 28th highest position in the world, ahead of Argentina and Ukraine.<sup>[23]</sup> Although crypto-mining relies on dedicated hardware (application-specific integrated circuits) instead of usual processors; therefore, it does not compete directly with scientific computing. Regardless, the magnitude of its carbon footprint needs to be addressed urgently.

Previous studies have made advances in estimating GHG emissions of computation but have limitations which preclude broad applicability. These limitations include the requirement that users self-monitor their power consumption<sup>[18]</sup> and are restricted with respect to hardware (e.g., GPUs and/or cloud systems<sup>[24,25]</sup>), software (e.g., Python package integration<sup>[25]</sup>), or applications (e.g., ML).<sup>[18,24,25]</sup> To facilitate green computing and widespread user uptake, there is a clear, and arguably urgent,

need for both a general and easy-to-use methodology for estimating carbon emissions that can be applied to any computational task.

In this study, we present a simple and widely applicable method and a tool for estimating the carbon footprint of computation. The method considers the different sources of energy usage, such as processors and memory, overhead of computing facilities and geographic location, while balancing accuracy and practicality. The online calculator (www.green-algorithms.org) implements this methodology and provides further context by interpreting carbon amounts using travel distances and carbon sequestration. We demonstrate the applicability of the Green Algorithms method by estimating the carbon footprint of particle physics simulations, weather forecast models, and NLP algorithms as well as the carbon effects of distributed computation using multiple CPUs. Finally, we make recommendations on ways for scientists to reduce their GHG emissions as well as discuss the limitations of our approach.

## 2. Results

We developed a simple method which estimates the carbon footprint of an algorithm based on a number of factors, including the hardware requirements of the tool, the runtime and the location of the data center (Experimental Section). Using a pragmatic scaling factor (PSF), we further augment our model by allowing for empirical estimates of repeated computations for a particular task, for example, parameter tuning and trial-and-errors. The resultant gCO<sub>2</sub>e is compared to the amount of carbon sequestered by trees and the emissions of common activities such as driving a car and air travel. We designed a freely available online tool, Green Algorithms (www.green-algorithms.org; **Figure 1**), which implements our approach and allows users to evaluate their computations or estimate the carbon savings or costs of redeploying them on other architectures.

We apply this tool to a range of algorithms selected from a variety of scientific fields: physics (particle simulations and DNA irradiation), atmospheric sciences (weather forecasting), and ML (NLP) (**Figure 2**). For each task, we curate published benchmarks and use www.green-algorithms.org to estimate the GHG emissions (Experimental Section). For parameters independent of the algorithm itself, we use average worldwide values, such as the worldwide average power usage effectiveness (PUE) of 1.67<sup>[26]</sup> and carbon intensity (CI) of 475 gCO<sub>2</sub>e kWh<sup>−1</sup>.<sup>[27]</sup>

### 2.1. Particle Physics Simulations

In particle physics, complex simulations are used to model the passage of particles through matter. Geant4<sup>[28]</sup> is a popular toolkit based on Monte-Carlo methods with wide-ranging applications, such as the simulation of detectors in the Large Hadron Collider and analysis of radiation burden on patients in clinical practice or external beam therapy.<sup>[29–31]</sup> Meylan et al.<sup>[32]</sup> investigated the biological effects of ionizing radiations on DNA on an entire human genome (6.4 × 10<sup>9</sup> nucleotide pairs) using GEANT4-DNA, an extension of GEANT4.

To quantify the DNA damage of radiation, they run experiments with photons of different energy, from 0.5 to 20 MeV. Each




**Figure 1.** The Green Algorithms calculator (www.green-algorithms.org).


experiment runs for three weeks to simulate 5000 particles (protons) using 24 processing threads and up to 10 GB of memory. Using the Green Algorithms tool, and assuming an average CPU power draw (such as the Xeon E5-2680, capable of running 24 threads on 12 cores), and worldwide average values for PUE and CI, we estimated that a single experiment emits 49 465 gCO<sub>2</sub>e. When taking into account a PSF of 11, corresponding to the 11 different energy levels tested, the carbon footprint of such study is 544 115 gCO<sub>2</sub>e. Using estimates of car and air travel (Experimental Section), 544 115 gCO<sub>2</sub>e is approximately equivalent to driving 3109 km (in a European car) or flying economy from New

York to San Francisco. In terms of carbon sequestration (Experimental Section), it would take a mature tree 49 years to remove the CO<sub>2</sub> equivalent to the GHG emissions of this study from the atmosphere (593 tree-months).

A common way to reduce the running time of algorithms is to distribute the computations over multiple processing cores. If the benefit in terms of time is well documented for each task, as in<sup>[33]</sup>, the environmental impact is usually not taken into account. GEANT4 is a versatile toolbox; it contains an electromagnetic package simulating particle transport in matter and high energy physics detector response.<sup>[34]</sup> Schweitzer et al.<sup>[33]</sup> use a



**Figure 2.** Carbon footprint (gCO<sub>2</sub>e) for a selection of algorithms, with and without their pragmatic scaling factor.



**Figure 3.** Effect of parallelization using multiple cores on run time and carbon footprint using TestEm12 GEANT4 simulation.

standardized example, TestEm12,<sup>[35]</sup> to compare the performances of different hardware configurations, from 1 to 60 cores (i.e., a full Xeon Phi CPU). With the Green Algorithms tool, we estimated the carbon footprint of each configuration (**Figure 3**), which shows that increasing the number of cores up to 15 improves both running time and GHG emissions. However, when multiplying the number of cores further by 4 (from 15 to 60), the running time is only halved, resulting in a twofold increase in emissions, from 238 to 481 gCO<sub>2</sub>e. Generally, if the reduction in running time is lower than the relative increase in the number of

cores, distributing the computations will worsen the carbon footprint. In particular, scientists should be mindful of marginal improvements in running time which have disproportionately large effects on GHG emissions, as demonstrated by the gap between 30 and 60 cores in **Figure 3**. For any parallelized computation, there is likely to be a specific optimal number of cores for minimal GHG emissions.

## 2.2. Weather Forecasting

Weather forecasts are based on sophisticated models simulating the dynamics between different components of the earth (such as the atmosphere and oceans). Operational models face stringent time requirements to provide live predictions to the public, with a goal of running about 200–300 forecast days (FDs) in one (wall clock) day.<sup>[36]</sup> Neumann et al.<sup>[36]</sup> present the performances of two models in use for current weather forecasts: i) the Integrated Forecast System (IFS)<sup>[37]</sup> used by the European Centre for Medium-Range Weather Forecasts (ECMWF) for 10-day forecasts, and ii) the ICOSahedral Non-hydrostatic (ICON)<sup>[38]</sup> designed by the German Weather Service (Deutscher Wetterdienst, DWD) and whose predictions are used by more than 30 national weather services.<sup>[39]</sup>

The configurations in daily use by the ECMWF include a supercomputer based in Reading, UK, which has a PUE of 1.45,<sup>[40]</sup> while ICON is run on the German Meteorological Computation Centre (DMRZ)<sup>[41]</sup> based in Germany (PUE unknown). Neumann et al.<sup>[36]</sup> ran their experiments on hardware similar to what equips both facilities, “Broadwell” CPU nodes (Intel E5-2695v4, 36 cores) and minimum 64 GB memory per node. We utilize these parameters for our CO<sub>2</sub>e emission estimates. It is important to note that ICON and IFS each solve slightly different problems, and therefore are not directly comparable.



The DWD uses ICON with a horizontal resolution of 13 km<sup>[42]</sup> and generates a FD in 8 min. Based on the experiments run by Neumann et al.,<sup>[36]</sup> this requires 575 Broadwell nodes (20 700 CPU cores). We estimate that generating one FD emits 12 848 gCO<sub>2</sub>e (14 tree-months). With a running time of 8 min per FD, ICON can generate 180 FDs in 24 h. When taking into account this PSF of 180, we estimated that each day, the ICON weather forecasting algorithm releases  $\approx 2\,312\,653$  gCO<sub>2</sub>e, equivalent to driving 13 215 km or flying from New York to San Francisco four times. In terms of carbon sequestration, the emissions of each day of ICON weather forecast are equivalent to 2523 tree-months.

At ECMWF, IFS makes 10-day operational weather forecasts with a resolution of 9 km. To achieve a similar threshold of 180 FDs per day, 128 Broadwell nodes are necessary (4608 cores).<sup>[36,43]</sup> Using the PUE of the UK ECMWF facility (1.45), we estimate the impact of producing one FD with IFS to be 1660 gCO<sub>2</sub>e. Using a PSF of 180 for one day's forecasts, we estimated emissions of 298 915 gCO<sub>2</sub>e, equivalent to driving 1708 km or three return flights between Paris and London. These emissions are equivalent to 326 tree-months.

Furthermore, we modeled the planned scenario of the ECMWF transferring its supercomputing to Bologna, Italy, in 2021.<sup>[44]</sup> Compared to the data center in Reading, the new data center in Bologna is estimated to have a more efficient PUE of 1.27.<sup>[45]</sup> *Prima facie* this move appears to save substantial GHG emissions; however, it is notable that the CI of Italy is 33% higher than the UK.<sup>[46]</sup> Unless the sources of electricity for the data center in Bologna are different from the rest of Italy and in the absence of further optimizations, we estimated that the move would result in an 18% increase in GHG emissions from the ECMWF (from 298 915 to 350 063 gCO<sub>2</sub>e).

### 2.3. Natural Language Processing

In NLP, the complexity and financial costs of model training are major issues.<sup>[17]</sup> This has motivated the development of language representations that can be trained once to model the complexity of natural language, and which could be used as input for more specialized algorithms. The BERT (Bidirectional Encoder Representations from Transformers)<sup>[47]</sup> algorithm is a field leader which yields both high performance and flexibility: state-of-the-art algorithms for more specific tasks are obtained by fine-tuning a pre-trained BERT model, for example in scientific text analysis<sup>[48]</sup> or biomedical text mining.<sup>[49]</sup> Yet, while the BERT model is intended to avoid retraining, many data scientists, perhaps understandably, continue to recreate or attempt to improve upon BERT, leading to redundant and ultimately inefficient computation as well as excess CO<sub>2</sub>e emissions. Even with optimized hardware (such as NVIDIA Volta GPUs), a BERT training run may take three days or more.<sup>[50]</sup>

Using these optimized parameters, Strubell et al.<sup>[18]</sup> showed that a run time of 79 h on 64 Tesla V100 GPUs was necessary to train BERT, with a usage factor of the GPUs of 62.7%. With the Greens Algorithms calculator, we estimated that a BERT training run would emit 754 407 gCO<sub>2</sub>e (driving 4311 km in a European car; 1.3 flights from New York to San Francisco; and 823 tree-months). When considering a conservative PSF of 100 for hyper-

parameters search, we obtain a carbon footprint of 75 440 740 gCO<sub>2</sub>e.

While BERT is a particularly widely utilized NLP tool, Google has also developed a chatbot algorithm, Meena, which was trained for 30 days on a TPU-v3 Pod containing 2048 Tensor Processing Unit (TPU) cores.<sup>[51]</sup> There is limited information on the power draw of TPU cores and memory; however, the power supply of this pod has been estimated to be 288 kW.<sup>[52]</sup> Using a run time of 30 days, assuming full usage of the TPUs and ignoring memory power draw, the Greens Algorithms calculator estimated that Meena training emitted 164 488 320 gCO<sub>2</sub>e, which corresponds to 179 442 tree-months or 71 flights between New-York and Melbourne.

## 3. Discussion

The method and Green Algorithms tool presented here provides users with a practical way to estimate the carbon footprint of their computations. The method focuses on producing sensible estimates with small overheads for scientists wishing to measure the footprint of their work. Consequently, the online calculator is simple to use and generalizable to nearly any computational task. We applied the Green Algorithms calculator to a variety of tasks, including particle physics simulations, weather forecasting, and NLP, to estimate their relative and ongoing carbon emissions. Real-world changes to computational infrastructures, such as moving data centers, was also quantifiable in terms of carbon footprint and was shown to be of substantive importance; for example, moving data centers may attain a more efficient PUE but a difference in CI may negate any efficiency gains, potentially making such a move detrimental to the environment.

Our work substantially enhances and extends prior frameworks for estimating the carbon footprint of computation. In particular, we have integrated and formalized previously unclear factors such as core usage and unitary power draw (per-core or per-GB of memory). As a result, and as presented in the Experimental Section, the carbon footprint of an algorithm can be broken down to a small number of key, easily quantifiable elements, such as number of cores, memory size and usage factor. This reduces the burden on the user, who is not required to either measure the power draw of hardware manually or use a limited range of cloud providers for their computations. This makes the method highly flexible in comparison to previous work. Besides drawing attention to the growing issue of GHG emissions of data centers, one of the benefits of presenting a detailed open methodology and tool is to provide users with the information they need to reduce their carbon footprint. Perhaps the most important challenge in green computing is to make the estimation and reporting of GHG emissions a standard practice. This requires transparent and easy-to-use methodology, such as the Green Algorithms calculator ([www.green-algorithms.org](http://www.green-algorithms.org)) and the open-source code and data presented here (see section Code availability).

Our approach has a number of limitations. First, the carbon footprint estimated is restricted to GHGs emitted to power computers during a particular task. We do not perform a life cycle assessment and therefore, do not consider the full environmental and social impact of manufacturing, maintaining, and disposing of the hardware used, or the maintenance of the power

plants. Including these is impractical at scale and would greatly reduce who can use the method. Besides, the conversion of the impact of various GHG into CO<sub>2</sub>e is commonly based on a 100-year timescale; however, this is now debated as it can misrepresent the impact of short-lived climate pollutants like methane<sup>[53]</sup> and new standards may be needed in the future. Second, the TDP may substantially underestimate power draw in some situations. For example, when hyperthreading is used, the real power consumption can be double the indicated TDP.<sup>[54]</sup> The TDP value remains a sensible estimate of the base consumption of the processor in most situations, but users using hyperthreading should be aware of the impact on power consumption. Third, while the power consumption from storage is usually minimal at the scale of one computation, if central storage is constantly queried by the algorithm (for example, to avoid overloading memory), this can be an important factor in power draw; however, there are resources which can be utilized if the algorithm is designed to be heavily storage reliant.<sup>[55]</sup> Moreover, at the scale of the data center, storage represents a significant part of electricity usage<sup>[55]</sup> and research projects relying on large databases should separately acknowledge the long-term carbon footprint of storage. Fourth, while some averaging is necessary, the energy mix of a country varies by the hour. For example, the CI of South Australia, which relies on wind and gas to produce electricity,<sup>[56]</sup> can vary between 112 and 592 gCO<sub>2</sub>e kWh<sup>-1</sup> within one day, depending on the quantity of coal-produced electricity imported from the neighboring state of Victoria.<sup>[57]</sup> Although most regions are relatively stable, these outliers may require a finer estimation. Our online calculator uses averaged values sourced from government reports.<sup>[46]</sup> Fifth, the PUE has some limitations as a measure of data centers energy usage,<sup>[58,59]</sup> due to inconsistencies in ways to calculate it. For example, reporting of PUE is highly variable from yearly averages to best-case scenarios, such as in winter when minimal cooling is required (as demonstrated by Google's quarterly results<sup>[60]</sup>). Whether to include infrastructure components such as security or on-site power generation is also source of discrepancies between data centers.<sup>[61]</sup> Although some companies present well-justified results, many PUEs have no or insufficient justification. Furthermore, PUE is not defined when computations are run on a laptop or desktop computer. As the device is used for multiple tasks simultaneously, it is impossible to estimate the power overhead due to the algorithm. In the calculator, we use a PUE of 1 because of the lack of information, but we caution this should not be interpreted as a sign of efficiency. Even though discrepancies will remain, the widespread adoption of an accurate, transparent, and certified estimation of PUE, such as the ISO/IEC standard,<sup>[62]</sup> would be a substantial step for the computing community. Sixth, the carbon emissions in the section Results are based on manual curation of the literature. When parameters such as usage factor or PUE were not specified, we made some assumptions (100% core usage, or using average PUE) that can explain differences between our estimates and the real emissions. For best results, authors should estimate and publish their emissions.

There are various, realistic actions one can take to reduce the carbon footprint of their computation. Acting on the various parameters in Green Algorithms (see Experimental Section), is a clear and easy way approach. Below, we describe a selection of practical changes one can make:

### 3.1. Algorithm Optimization

Increasing the efficiency of an algorithm can have myriad benefits, even apart from reducing its carbon footprint. Therefore, we highly recommend this and foresee algorithm optimization as one of the most productive, easily recognizable core activities of green computing. While speed is an obvious efficiency gain, part of algorithm optimization also includes memory minimization. The power draw from memory mainly depends on the memory available, not the actual memory used,<sup>[63]</sup> and the memory available is often the peak memory needed for one step of the algorithm (typically a merge or aggregation). By optimizing these steps, one can easily reduce energy consumption.

### 3.2. Reduce the Pragmatic Scaling Factor

Limiting the number of times an algorithm runs, especially those that are power hungry, is perhaps the easiest way to reduce carbon footprint. Relatedly, best practices to limit PSF (as well as financial cost) include limiting parameter fine-tuning to the minimum necessary and building a small-scale example for debugging.

### 3.3. Choice of Data Center

Carbon footprint is directly proportional to data center efficiency and the CI of the location. The latter is perhaps the parameter which most affects total carbon footprint because of inter-country variation, from under 20 gCO<sub>2</sub>e kWh<sup>-1</sup> in Norway and Switzerland to over 800 gCO<sub>2</sub>e kWh<sup>-1</sup> in Australia, South Africa, and some states in the USA. To rigorously assess the impact of punctually relocating computations, the marginal CI, rather than the average one, should be used.<sup>[64]</sup> The marginal value depends on which power plant would be solicited to meet the unexpected increased demand. Although it would ideally be used, it varies by the hour and is often not practical to estimate accurately at scale. When the marginal CI is unknown, the average one (presented in Experimental Section and Figure S2, Supporting information) can be used by scientists as a practical lower bound estimate to assess the benefit of moving computations. Indeed, due to the low operating cost of renewable technologies, the marginal power plants (which are the last one solicited) are generally high-carbon technologies such as fuel or gas<sup>[64]</sup> which leads the marginal CI to be higher than the average CI. Besides, if the move is permanent, by relocating an HPC facility or using cloud computing for example, then the energy needs are incorporated into utility planning and the average CI is the appropriate metric to use. Data center efficiency (PUE) varies widely between facilities but, in general, large data centers optimize cooling and power supply, reducing the energy overhead and making them more efficient than personal servers. Notably, a 2016 report estimated that if 80% of small data centers in the USA were aggregated into hyperscale facilities, energy usage would reduce by 25%.<sup>[65]</sup> For users to make informed choices, data centers should report their PUE and other energy metrics. While large providers like Google or Microsoft widely advertise their servers' efficiency,<sup>[60,66]</sup> smaller structures often do not. As highlighted here, cloud providers offer the opportunity to use efficient data centers in low-carbon countries, and they can be greener alternatives to local data centers.

### 3.4. Offsetting GHG emissions

Carbon offsetting is a flexible way to compensate for carbon footprint. An institution or a user themselves can directly support reductions in CO<sub>2</sub> or other GHGs, for example by sponsoring fuel-efficient stoves in developing countries, reduction in deforestation or hydroelectric and wind-based power plants.<sup>[67,68]</sup> The pros and cons of carbon offsetting are still debated due to the variety of mechanisms and intricate international legislations and competing standards. Therefore, we only present here an overview and point interested scientists at some resources. Multiple international standards regulate the purchase of carbon credits and ensure the efficiency of the projects supported.<sup>[69]</sup> Most of the well-established standards are managed by non-profits and abide by the mechanisms set in place by the Kyoto protocol (in particular Certified Emission Reduction)<sup>[70]</sup> and the PAS 2060 Carbon Neutrality standard from the British Standards Institution.<sup>[71]</sup> Although the primary aim is carbon offsetting, projects are often also selected in line with the United Nations' Agenda 30 for Sustainable Development,<sup>[72]</sup> a broader action plan addressing inequalities, food security, and peace. Amongst the most popular standards are the Gold Standard (founded by WWF and other NGOs),<sup>[73]</sup> Verra (formerly Verified Carbon Standard),<sup>[74]</sup> and the American Carbon Registry (a private voluntary GHG registry).<sup>[75]</sup> In addition to direct engagement with these standards, platforms like Carbon Footprint<sup>[67]</sup> select certified projects and facilitate the purchase of credits.

## 4. Conclusions

The framework presented here is generalizable to nearly any computation and may be used as a foundation for other aspects of green computing. The carbon footprint of computation is substantial and may be affecting the climate. We therefore hope that this new tool and metrics raise awareness of these issues as well as facilitate pragmatic solutions which may help to mitigate the environmental consequences of modern computation. Overall, with the right tools and practices, we believe HPC and cloud computing can be immensely positive forces for both improving the human condition and saving the environment.

## 5. Experimental Section

The carbon footprint of an algorithm depends on two factors: the energy needed to run it and the pollutants emitted when producing such energy. The former depends on the computing resources used (e.g., number of cores, running time, and data center efficiency) while the later, called carbon intensity, depends on the location and production methods used (e.g., nuclear, gas, or coal).

There are several competing definitions of "carbon footprint," and in this project, the extended definition from Wright et al.<sup>[76]</sup> was used. The climate impact of an event is presented in terms of carbon dioxide equivalent (CO<sub>2</sub>e) and summarizes the global warming effect of the GHG emitted in the determined timeframe, here running a set of computations. The GHGs considered were carbon dioxide (CO<sub>2</sub>), methane (CH<sub>4</sub>) and nitrous oxide (N<sub>2</sub>O);<sup>[77]</sup> these are the three most common GHGs of the "Kyoto basket" defined in the Kyoto Protocol<sup>[78]</sup> and represent 97.9% of global GHG emissions.<sup>[79]</sup> The conversion into CO<sub>2</sub>e was done using Global Warming Potential (GWP) factors from the Intergovernmental Panel on Climate Change (IPCC)<sup>[77,80]</sup> based on a 100-year horizon (GWP100).

When estimating these parameters, accuracy and feasibility must be balanced. This study focused on a methodology that could be easily and broadly adopted by the community and therefore, restricts the scope of the environmental impact considered to GHGs emitted to power computing facilities for a specific task. Moreover, the framework presented requires no extra computation, nor involves invasive monitoring tools.

**Energy Consumption:** An algorithm's energy<sup>[81]</sup> needs were modeled as a function of the running time, the number, type, and process time of computing cores (CPU or GPU), the amount of memory mobilized, and the power draw of these resources. The model further included the efficiency of the data center,<sup>[82]</sup> which represents how much extra power is necessary to run the facility (e.g., cooling and lighting).

Similar to previous works,<sup>[18,24]</sup> this estimate was based on the power draw from processors and memory, as well as the efficiency of the data center. However, the formula was refined and flexibility was added by including a unitary power draw (per core and per GB of memory) and the processor's usage factor. The energy consumption  $E$  (in kWh) was expressed as:

$$E = t \times (n_c \times P_c \times u_c + n_m \times P_m) \times PUE \times 0.001 \quad (1)$$

where  $t$  is the running time (hours),  $n_c$  the number of cores, and  $n_m$  the size of memory available (gigabytes).  $u_c$  is the core usage factor (between 0 and 1).  $P_c$  is the power draw of a computing core and  $P_m$  the power draw of memory (Watt).  $PUE$  is the efficiency coefficient of the data centre.

The assumptions made regarding the different components are discussed below. It is previously shown that the power draw of a server motherboard is negligible<sup>[83]</sup> and in a desktop computer, the motherboard handles a multitude of tasks, which makes it impractical to assess the fraction of power usage attributable to a specific algorithm. For these reasons, it was decided not to include the motherboard's power draw in this model.

**Power Draw of the Computing Core:** The metric commonly used to report the power draw of a processor, either CPU or GPU, is its thermal design power (TDP, in Watt) and is provided by the manufacturer. TDP values frequently corresponded to CPU specifications which include multiple cores, thus here TDP values were normalized to per-core. While TDP is not a direct measure of power consumption, rather the amount of heat a cooling system dissipates during regular use—it is commonly considered a reasonable approximation.

The energy used by the processor was the power draw multiplied by processing time, scaled by the usage factor. However, processing time could not be known a priori and, on some platforms, tracking could be impractical at scale. Modeling exact processing time of past projects may also necessitate re-running jobs, which would generate unnecessary emissions. Therefore, when this processing time is unknown, the simplifying assumption was made that core usage is 100% of run time ( $u_c = 1$  in Equation (1)).

**Power Draw from Memory:** Memory power draw is mainly due to background consumption with a negligible contribution from the workload and database size.<sup>[63]</sup> Moreover, the power draw is mainly affected by the total memory allocated, not by the actual size of the database used, because the load is shared between all memory slots which keeps every slot in a power-hungry active state. Therefore, the primary factor influencing power draw from memory is the quantity of memory mobilized, which simply requires an estimation of the power draw per gigabyte. Measured experimentally, this was estimated to be 0.3725 W per GB.<sup>[63,84]</sup>

For example, requesting 29 GB of memory draws 10.8 W, which is the same as one core of a popular Core-i5 CPU. Figure S1, Supporting information further compares the power draw of memory to a range of popular CPUs.

**Power Draw from Storage:** The power draw of storage equipment (HDD or SSD) varies significantly with workload.<sup>[55]</sup> However, in regular use, storage is typically solicited far less than memory and is mainly used as a more permanent record of the data, independently of the task at hand. The power draw of storage was estimated to be 0.001 W per GB (Note S2, Supporting information). As above, by comparison, the power draw of memory (0.3725 W per GB) and a Core-i5 CPU (10.8 W per core) are more than two orders of magnitude greater. While the researcher overhead for approximating storage usage may not be substantial, it is unlikely

to make a significant difference to overall power usage (and GHG emissions) estimation. Therefore, the power consumption of storage was not considered in this work.

**Energy Efficiency:** Data center energy consumption includes additional factors, such as server cooling systems, power delivery components, and lighting. The efficiency of a given data center can be measured by the Power Usage Effectiveness,<sup>[85,61]</sup> defined as the ratio between the total power drawn by the facility and the power used by computing equipment:

$$PUE = \frac{P_{\text{total}}}{P_{\text{compute}}} \quad (2)$$

A data center PUE of 1.0 represents an ideal situation where all power supplied to the building is utilized by computing equipment. The global average of data centers was estimated as 1.67 in 2019.<sup>[26]</sup> While data centers with relatively inefficient PUE may not report it as such, some data centers and companies have invested significant resources to bring their PUEs as close to 1.0 as possible; for example, Google uses ML to reduce its global yearly average PUE to 1.10.<sup>[60,86]</sup>

**Carbon Intensity of Energy Production:** For a given country and energy mix, the carbon footprint in CO<sub>2</sub>e represents the amount of CO<sub>2</sub> with the same global warming impact as the GHGs emitted, which simplifies the comparison between different electricity production methods. The carbon footprint of producing 1 kWh of energy (CI) varies significantly between locations due to the broad range of production methods (Figure S2, Supporting information), from 12 gCO<sub>2</sub>e kWh<sup>-1</sup> in Switzerland (mainly powered by hydro) to 880 gCO<sub>2</sub>e kWh<sup>-1</sup> in Australia (mainly powered by coal and gas).<sup>[57,46]</sup> The 2020 CI values aggregated by *Carbon Footprint*<sup>[46]</sup> were used; these production factors take into account the GHG emissions at the power plants (power generation) as well as, when available, the footprint of distributing energy to the data center.

**Estimation of Carbon Footprint:** The carbon footprint C (in gCO<sub>2</sub>e) of producing a quantity of energy E (in kWh) from sources with a CI (in gCO<sub>2</sub>e kWh<sup>-1</sup>) is then:

$$C = E \times CI \quad (3)$$

By putting together Equations (1) and (3), the long-form equation of the carbon footprint C is obtained:

$$C = t \times (n_c \times P_c \times u_c + n_m \times P_m) \times PUE \times CI \times 0.001 \quad (4)$$

**CO<sub>2</sub>e of Driving and Air Travel:** gCO<sub>2</sub>e was contextualized by estimating an equivalence in terms of distance travelled by car or by passenger aircraft. Previous studies had estimated the emissions of the average passenger car in Europe as 175 gCO<sub>2</sub>e km<sup>-1</sup><sup>[77,87]</sup> (251 gCO<sub>2</sub>e km<sup>-1</sup> in the United States<sup>[88]</sup>). The emissions of flying on a jet aircraft in economy class were estimated between 139 and 244 gCO<sub>2</sub>e km<sup>-1</sup> per person, depending on the length of the flight.<sup>[77]</sup> Three reference flights were used: Paris to London (50 000 gCO<sub>2</sub>e), New York to San Francisco (570 000 gCO<sub>2</sub>e), and New York to Melbourne (2 310 000 gCO<sub>2</sub>e).<sup>[89]</sup>

**CO<sub>2</sub> Sequestration by Trees:** Trees play a major role in carbon sequestration and although not all GHGs emitted could be sequestered, CO<sub>2</sub> represents 74.4% of these emissions.<sup>[90]</sup> To provide a metric of reversion for CO<sub>2</sub>e, the number of trees needed to sequester the equivalent emissions of a given computation was computed. The metric tree-months, the number of months a mature tree needs to absorb a given quantity of CO<sub>2</sub>, was defined. While the amount of CO<sub>2</sub> sequestered by a tree per unit of time depends on a number of factors, such as its species, size, or environment, it was estimated that a mature tree sequesters, on average, ≈11 kg of CO<sub>2</sub> per year,<sup>[91]</sup> giving the multiplier in tree-months a value close to 1 kg of CO<sub>2</sub> per month (0.92 g).

**Pragmatic Scaling Factor:** Many analyses are presented as a single run of a particular algorithm or software tool; however, computations are rarely performed only once. Algorithms are run multiple times, sometimes hundreds, systematically or manually, with different parameterizations. Statistical models may include any number of combinations of covariates,

fitting procedures, etc. It is important to include these repeats in the carbon footprint. To take into account the number of times a computation is performed in practice, the PSF was defined, a scaling factor by which the estimated GHG emissions are multiplied.

The value and causes of the PSF vary greatly between tasks. In ML, tuning the hyper-parameters of a model requires hundreds, if not thousands,<sup>[18]</sup> of runs, while other tools require less tuning and can sometimes be run a smaller number of times. As per published work or the user's own experience, the PSF should be estimated for any specific task; besides, in Green Algorithms it is provided for, and recommended that, each user estimate their own PSF.

## Supporting Information

Supporting Information is available from the Wiley Online Library or from the author.

## Acknowledgements

L.L. and J.G. contributed equally to this work. L.L. was supported by the University of Cambridge MRC DTP (MR/S502443/1). J.G. was supported by a La Trobe University Postgraduate Research Scholarship jointly funded by the Baker Heart and Diabetes Institute and a La Trobe University Full-Fee Research Scholarship. This work was supported by core funding from: the UK Medical Research Council (MR/L003120/1), the British Heart Foundation (RG/13/13/30194; RG/18/13/33946) and the NIHR Cambridge Biomedical Research Centre (BRC-1215-20014). This work was also supported by Health Data Research UK, which is funded by the UK Medical Research Council, Engineering and Physical Sciences Research Council, Economic and Social Research Council, Department of Health and Social Care (England), Chief Scientist Office of the Scottish Government Health and Social Care Directorates, Health and Social Care Research and Development Division (Welsh Government), Public Health Agency (Northern Ireland), British Heart Foundation, and Wellcome. M.I. was supported by the Munz Chair of Cardiovascular Prediction and Prevention. This study was supported by the Victorian Government's Operational Infrastructure Support (OIS) program. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health and Social Care.

## Conflict of Interest

The authors declare no conflict of interest.

## Data Availability Statement

The data that support the findings of this study are openly available in a GitHub repository at <https://github.com/GreenAlgorithms/green-algorithms-tool>.

## Keywords

climate change, computational research, green computing

Received: February 20, 2021

Revised: March 22, 2021

Published online: May 2, 2021

[1] N. Jones, *Nature* **2018**, 561, 163.

[2] A. Andrae, T. Edler, *Challenges* **2015**, 6, 117.



- [3] R. B. Alley, K. A. Emanuel, F. Zhang, *Science* **2019**, 363, 342.
- [4] The Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, D. Bintley, L. Blackburn, W. Boland, K. L. Bouman, G. C. Bower, M. Bremer, C. D. Brinkerink, R. Brissenden, S. Britzen, A. E. Broderick, D. Brogiere, T. Bronzwaer, D.-Y. Byun, J. E. Carlstrom, A. Chael, C. Chan, S. Chatterjee, K. Chatterjee, M.-T. Chen, Y. Chen, I. Cho, et al., *Astrophys. J. Lett.* **2019**, 875, L1.
- [5] The Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, D. Bintley, L. Blackburn, W. Boland, K. L. Bouman, G. C. Bower, M. Bremer, C. D. Brinkerink, R. Brissenden, S. Britzen, A. E. Broderick, D. Brogiere, T. Bronzwaer, D.-Y. Byun, J. E. Carlstrom, A. Chael, C. Chan, S. Chatterjee, K. Chatterjee, M.-T. Chen, Y. Chen, I. Cho, et al., *Astrophys. J. Lett.* **2019**, 875, L3.
- [6] The Event Horizon Telescope Collaboration, K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, D. Bintley, L. Blackburn, W. Boland, K. L. Bouman, G. C. Bower, M. Bremer, C. D. Brinkerink, R. Brissenden, S. Britzen, A. E. Broderick, D. Brogiere, T. Bronzwaer, D.-Y. Byun, J. E. Carlstrom, A. Chael, C. Chan, S. Chatterjee, K. Chatterjee, M.-T. Chen, Y. Chen, I. Cho, et al., *Astrophys. J. Lett.* **2019**, 875, L4.
- [7] A. Buniello, J. A. L. MacArthur, M. Cerezo, L. W. Harris, J. Hayhurst, C. Malangone, A. McMahon, J. Morales, E. Mountjoy, E. Sollis, D. Suveges, O. Vrousou, P. L. Whetzel, R. Amode, J. A. Guillen, H. S. Riat, S. J. Trevanion, P. Hall, H. Junkins, P. Flicek, T. Burdett, L. A. Hindorf, F. Cunningham, H. Parkinson, *Nucleic Acids Res.* **2019**, 47, D1005.
- [8] D. Ben-Israel, W. B. Jacobs, S. Casha, S. Lang, W. H. A. Ryu, M. de Lotbiniere-Bassett, D. W. Cadotte, *Artif. Intell. Med.* **2020**, 103, 101785.
- [9] D. M. Dimiduk, E. A. Holm, S. R. Niezgoda, *Integr. Mater. Manuf. Innov.* **2018**, 7, 157.
- [10] G. Choy, O. Khalilzadeh, M. Michalski, S. Do, A. E. Samir, O. S. Pinykh, J. R. Geis, P. V. Pandharipande, J. A. Brink, K. J. Dreyer, *Radiology* **2018**, 288, 318.
- [11] S. Athey, in *The Economics of Artificial Intelligence: An Agenda*, National Bureau of Economic Research, Inc., Cambridge **2018**, p. 507.
- [12] XSEDE Impact—Usage Statistics, <https://portal.xsede.org/#/gallery> (accessed: March 2021).
- [13] C. Reimsbach-Kounatze, *OECD Digital Economy Papers*, Vol. 155, OECD Publishing, Paris **2009**.
- [14] S. Sarkar, S. Misra, *IET Network* **2016**, 5, 23.
- [15] K. Gai, M. Qiu, H. Zhao, L. Tao, Z. Zong, *J. Network Comput. Appl.* **2016**, 59, 46.
- [16] I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT Press, Cambridge, MA **2016**.
- [17] R. Schwartz, J. Dodge, N. A. Smith, O. Etzioni, *Green AI. Commun. ACM* **2020**, 63, 54.
- [18] E. Strubell, A. Ganesh, A. McCallum, Energy and Policy Considerations for Deep Learning in NLP, in Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, Jul. **2019**, pp. 3645–3650, <https://doi.org/10.18653/v1/P19-1355>.
- [19] K. Jahnke, C. Fendt, M. Fournesneau, I. Georgiev, T. Herbst, M. Kaasinen, D. Kossakowski, J. Rybizki, M. Schlecker, G. Seidel, T. Henning, L. Kreidberg, H.-W. Rix, *Nat. Astron.* **2020**, 4, 812.
- [20] S. Portegies Zwart, *Nat. Astron.* **2020**, 4, 819.
- [21] A. R. H. Stevens, S. Bellstedt, P. J. Elahi, M. T. Murphy, *Nat. Astron.* **2020**, 4, 843.
- [22] C. Stoll, L. Klaußen, U. Gellersdörfer, *Joule* **2019**, 3, 1647.
- [23] Cambridge Bitcoin Electricity Consumption Index (CBECI), <https://cbeci.org/contact/> (accessed: March 2021).
- [24] A. Lacoste, A. Luccioni, V. Schmidt, T. Dandres, arXiv191009700 Cs **2019**.
- [25] P. Henderson, J. Hu, J. Romoff, E. Brunskill, D. Jurafsky, J. Pineau, *J. Mach. Learn. Res.* **2020**, 21, 1.
- [26] A. Lawrence, Uptime Institute Blog **2019**, <https://journal.uptimeinstitute.com/is-pue-actually-going-up/> (accessed: April 2020).
- [27] Emissions—Global Energy & CO2 Status Report 2019—Analysis, <https://www.iea.org/reports/global-energy-co2-status-report-2019/emissions> (accessed: February 2020).
- [28] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytrcek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell'Acqua, G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, et al., *Nucl. Instrum. Methods Phys. Res.* **2003**, 506, 250.
- [29] S. Jan, G. Santin, D. Strul, S. Staelens, K. Assié, D. Autret, S. Avner, R. Barbier, M. Bardiès, P. M. Bloomfield, D. Brasse, V. Breton, P. Bruyndonckx, I. Buvat, A. F. Chatzioannou, Y. Choi, Y. H. Chung, C. Comtat, D. Donnarieix, L. Ferrer, S. J. Glick, C. J. Groiselle, D. Guez, P.-F. Honore, S. Kerhoas-Cavata, A. S. Kirov, V. Kohli, M. Koole, M. Krieguer, D. J. van der Laan, et al., *Phys. Med. Biol.* **2004**, 49, 4543.
- [30] S. Jan, D. Benoit, E. Becheva, T. Carlier, F. Cassol, P. Descourt, T. Frisson, L. Grevillot, L. Guigues, L. Maigne, C. Morel, Y. Perrot, N. Rehfeld, D. Sarrut, D. R. Schaart, S. Stute, U. Pietrzyk, D. Visvikis, N. Zahra, I. Buvat, *Phys. Med. Biol.* **2011**, 56, 881.
- [31] D. Sarrut, M. Bardiès, N. Boussion, N. Freud, S. Jan, J.-M. Létang, G. Loudos, L. Maigne, S. Marcatili, T. Mouxion, P. Papadimitroulas, Y. Perrot, U. Pietrzyk, C. Robert, D. R. Schaart, D. Visvikis, I. Buvat, *Med. Phys.* **2014**, 41, 064301.
- [32] S. Meylan, S. Incerti, M. Karamitros, N. Tang, M. Bueno, I. Clairand, C. Villagrasa, *Sci. Rep.* **2017**, 7, 11923.
- [33] P. Schweitzer, S. Cipiè, A. Dufare, H. Payno, Y. Perrot, D. R. C. Hill, L. Maigne, *Sci. Program.* **2015**, 2015, 980752.
- [34] J. Apostolakis, A. Bagulya, S. Elles, V. N. Ivanchenko, J. Jacquemier, M. Maire, T. Toshito, L. Urban, *J. Phys. Conf. Ser.* **2010**, 219, 032044.
- [35] Geant4—an Object-Oriented Toolkit for Simulation in HEP, <https://gitlab.cern.ch/geant4/geant4/tree/edb408b5618b3b1cd3f40c5759a5da4aa56bb7b/examples/extended/electromagnetic/TestEm5> (accessed: June 2020).
- [36] P. Neumann, P. Düben, P. Adamidis, P. Bauer, M. Brück, L. Kornblueh, D. Klocke, B. Stevens, N. Wedi, J. Biercamp, *Philos. Trans. R. Soc., A* **2019**, 377, 20180148.
- [37] S. Keeley, Modelling and Prediction: Integrated Forecast System, <https://www.ecmwf.int/en/research/modelling-and-prediction> (Published in 2013, accessed: April 2020).
- [38] G. Zängl, D. Reinert, P. Rípodas, M. Baldauf, *Q. J. R. Meteorol. Soc.* **2015**, 141, 563.
- [39] Wetter und Klima—Deutscher Wetterdienst—Numerical weather prediction models—ICON (Icosahedral Nonhydrostatic) Model, [https://www.dwd.de/EN/research/weatherforecasting/num\\_modelling/01\\_num\\_weather\\_prediction\\_modells/icon\\_description.html](https://www.dwd.de/EN/research/weatherforecasting/num_modelling/01_num_weather_prediction_modells/icon_description.html) (accessed: June 2020).
- [40] S. Baylis, **2011**, <https://doi.org/10.21957/CEMUB9F8>, <https://www.ecmwf.int/node/17450> (accessed: April 2020).
- [41] Wetter und Klima—Deutscher Wetterdienst—Information technology—Data processing, DMRZ, <https://www.dwd.de/EN/aboutus/it/functions/Teasergroup/dataprocessing.html?nn=24864> (accessed: June 2020).
- [42] The horizontal resolution represents the level of geographical detail achieved when modelling the different weather phenomenon.

- [43] J. A. P. León, Progress in using single precision in the IFS, <https://www.ecmwf.int/en/newsletter/157/meteorology/progress-using-single-precision-ifs> (Published in 2018, accessed: June 2020).
- [44] G. Lentze, ECMWF signs contract with Atos for new supercomputer, <https://www.ecmwf.int/en/about/media-centre/news/2020/ecmwf-signs-contract-atos-new-supercomputer> (accessed: April 2020).
- [45] ECMWF, ECMWF Q&A, [https://www.ecmwf.int/sites/default/files/medialibrary/2019-02/Industry\\_day\\_-\\_QA.pdf](https://www.ecmwf.int/sites/default/files/medialibrary/2019-02/Industry_day_-_QA.pdf) (accessed: April 2020).
- [46] International Electricity Factors, [https://www.carbonfootprint.com/international\\_electricity\\_factors.html](https://www.carbonfootprint.com/international_electricity_factors.html) (accessed: November 2020).
- [47] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Minneapolis, Minnesota, Jun. 2019, pp. 4171–4186, <https://doi.org/10.18653/v1/N19-1423>.
- [48] I. Beltagy, K. Lo, A. Cohan, in Proc. 2019 Conf. Empir. Methods Nat. Lang. Process. 9th Int. Jt. Conf. Nat. Lang. Process. EMNLP-IJCNLP, Association For Computational Linguistics, Hong Kong, China **2019**, pp. 3615–3620.
- [49] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, J. Kang, *Bioinformatics* **2020**, 36, 1234.
- [50] NVIDIA AI, BERT Meets GPUs—Future Vision—Medium, <https://web.archive.org/web/20190521104957/https://medium.com/future-vision/bert-meets-gpus-403d3fbed848> (Published in 2019, accessed: April 2020).
- [51] D. Adiwardana, M.-T. Luong, D. R. So, J. Hall, N. Fiedel, R. Thop-pilan, Z. Yang, A. Kulshreshtha, G. Nemade, Y. Lu, Q. V. Le, *arXiv:2001.09977* **2020**.
- [52] P. Teich, Tearing Apart Google's TPU 3.0 AI Coprocessor, <https://www.nextplatform.com/2018/05/10/tearing-apart-googles-tpu-3-0-ai-coprocessor/> (Published in 2018, accessed: April 2020).
- [53] M. R. Allen, K. P. Shine, J. S. Fuglestedt, R. J. Millar, M. Cain, D. J. Frame, A. H. Macey, *npj Clim. Atmos. Sci.* **2018**, 1, 16.
- [54] I. Cutress, Why Intel Processors Draw More Power Than Expected: TDP and Turbo Explained, <https://www.anandtech.com/show/13544/why-intel-processors-draw-more-power-than-expected-tdp-turbo> (accessed: January 2020).
- [55] E. Tómes, N. Altıparmak, in Proc. 2017 IEEE Int. Conf. Cluster Computing, IEEE, Piscataway, NJ **2017**, pp. 625–626.
- [56] SA.GOV.AU—SA's electricity supply and market, <https://www.sa.gov.au/topics/energy-and-environment/energy-supply/sas-electricity-supply-and-market> (accessed: July 2020).
- [57] B. Tranberg, O. Corradi, B. Lajoie, T. Gibon, I. Staffell, G. B. Andresen, *Energy Strategy Rev.* **2019**, 26, 100367.
- [58] R. C. Zoie, R. D. Mihaela, S. Alexandru, in Proc. 2017 5th Int. Symp. Electr. Electron. Eng. ISEEE 2017, IEEE, Piscataway, NJ **2017**, pp. 1–6.
- [59] J. Yuventi, R. Mehdizadeh, *Energy Build.* **2013**, 64, 90.
- [60] Efficiency—Data Centers—Google, <https://www.google.com/about/datacenters/efficiency/> (accessed: January 2020).
- [61] V. Avelar, D. Azevedo, A. French, E. N. Power, WhitePaper **2012**, [https://datacenters.lbl.gov/sites/all/files/WP49-PUE%20A%20Comprehensive%20Examination%20of%20the%20Metric\\_v6.pdf](https://datacenters.lbl.gov/sites/all/files/WP49-PUE%20A%20Comprehensive%20Examination%20of%20the%20Metric_v6.pdf) (accessed: November 2020).
- [62] I. S. O./I. E. C. 30134-2:2016(en), Information technology—Data centres—Key performance indicators—Part 2: Power usage effectiveness (PUE), <https://www.iso.org/obp/ui/#iso:std:iso-iec:30134-2:ed-1:v1:en> (accessed: November 2020).
- [63] A. Karyakin, K. Salem, in Proc. 13th Int. Workshop Data Management on New Hardware (DaMoN), ACM Press, Chicago **2017**, pp. 1–9.
- [64] R. McCarthy, C. Yang, *J. Power Sources* **2010**, 195, 2099.
- [65] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, W. Lintner, *United States Data Center Energy Usage Report*, Lawrence Berkeley National Laboratory, Berkeley, USA **2016**.
- [66] Azure global infrastructure | Microsoft Azure, <https://azure.microsoft.com/en-gb/global-infrastructure/> (accessed: April 2020).
- [67] Carbon Offset Projects, <https://www.carbonfootprint.com/carbonoffsetprojects.html> (accessed: April 2020).
- [68] Project Map—ClimateCare, <https://climatecare.org/project-map/> (accessed: April 2020).
- [69] A. Michaelowa, I. Shishlov, S. Hoch, P. Bofill, A. Espelage, *Overview and Comparison of Existing Carbon Crediting Schemes*, Nordic Environment Finance Corporation (NICA), Helsinki, Finland **2019**.
- [70] Emissions Trading | UNFCCC, <https://unfccc.int/process/the-kyoto-protocol/mechanisms/emissions-trading> (accessed: April 2020).
- [71] PAS 2060 Carbon Neutrality, <https://www.bsigroup.com/en-GB/PAS-2060-Carbon-Neutrality/> (accessed: April 2020).
- [72] *Transforming Our World: The 2030 Agenda for Sustainable Development*, United Nations, New York **2015**.
- [73] The Gold Standard, <https://www.goldstandard.org/> (accessed: April 2020).
- [74] Verra, <https://verra.org/> (accessed: April 2020).
- [75] American Carbon Registry, <https://americancarbonregistry.org/> (accessed: April 2020).
- [76] L. A. Wright, S. Kemp, I. Williams, *Carbon Manage.* **2011**, 2, 61.
- [77] N. Hill, R. Bramwell, E. Karagianni, L. Jones, J. MacCarthy, S. Hinton, C. Walker, B. Harris, Department for Business Energy & Industrial Strategy (BEIS) **2020**.
- [78] *Kyoto Protocol To The United Nations Framework Convention On Climate Change*, United Nations, New York, USA **1998**.
- [79] H. Ritchie, M. Roser, *Our World Data* **2020**, <https://ourworldindata.org/co2-and-other-greenhouse-gas-emissions> (accessed: June 2020).
- [80] Fourth Assessment Report—IPCC, <https://www.ipcc.ch/assessment-report/ar4/> (accessed: December 2020).
- [81] Power (in Watt, W) measures the instantaneous draw of a component. Energy (in kilowatt-hours, kWh) measures the power draw over time and is obtained by multiplying the power draw by the running time.
- [82] By data centre, we mean the facility hosting the cores and memory, which may not be a dedicated data centre.
- [83] H. Geng, *Data Center Handbook*, John Wiley & Sons, Hoboken, USA **2014**.
- [84] C. Angelini, I. W. August 29 2014, Intel Core i7-5960X, -5930K And -5820K CPU Review: Haswell-E Rises, <https://www.tomshardware.com/uk/reviews/intel-core-i7-5960x-haswell-e-cpu,3918-13.html> (accessed: January 2020).
- [85] C. L. Belady, C. G. Malone, Proc. ASME 2007 InterPACK Conf. collocated with the ASME/JSM 2007 Thermal Engineering Heat Transfer Summer Conf., ASME, New York, USA **2010**, pp. 751.
- [86] J. Gao, Machine Learning Applications for Data Center Optimization, **2014**.
- [87] E. Helmers, J. Leitão, U. Tietge, T. Butler, *Atmos. Environ.* **2019**, 198, 122.
- [88] O. US EPA, Greenhouse Gas Emissions from a Typical Passenger Vehicle, <https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle> (accessed: January 2020).
- [89] Carbon Footprint Calculator, <https://calculator.carbonfootprint.com/calculator.aspx?tab=3> (accessed: June 2020).
- [90] H. Ritchie, M. Roser, Greenhouse gas emissions, <https://ourworldindata.org/greenhouse-gas-emissions> (accessed: November 2020).
- [91] H. Akbari, *Environ. Pollut.* **2002**, 116, S119.