

The AI Gambit — Leveraging artificial intelligence to combat climate change: opportunities, challenges, and recommendations

Josh COWLS^{1,2*}, Andreas Tsamados^{1*}, Mariarosaria Taddeo^{1,2} and Luciano Floridi^{1,2}

¹ Oxford Internet Institute, University of Oxford, 1 St Giles', Oxford, OX1 3JS, UK

² Alan Turing Institute, British Library, 96 Euston Rd, London NW1 2DB, UK

* Both authors contributed equally.

Abstract

In this article we analyse the role that artificial intelligence (AI) could play, and is playing, to combat global climate change. We identify two crucial opportunities that AI offers in this domain: it can help improve and expand current understanding of climate change and it contribute to combatting the climate crisis effectively. However, the development of AI also raises two sets of problems when considering climate change: the possible exacerbation of social and ethical challenges already associated with AI, and the contribution to climate change of the greenhouse gases emitted by training data and computation-intensive AI systems. We assess the carbon footprint of AI research, and the factors that influence AI's greenhouse gas (GHG) emissions in this domain. We find that the carbon footprint of AI research may be significant and highlight the need for more evidence concerning the trade-off between the GHG emissions generated by AI research and the energy and resource efficiency gains that AI can offer. In light of our analysis, we argue that leveraging the opportunities offered by AI for global climate change whilst limiting its risks is a gambit which requires responsive, evidence-based and effective governance to become a winning strategy. We conclude by identifying the European Union as being especially well-placed to play a leading role in this policy response and provide 13 recommendations that are designed to identify and harness the opportunities of AI for combatting climate change, while reducing its impact on the environment.

Keywords: Artificial Intelligence, Climate Change, Digital Ethics; Digital Governance; Environment; Sustainability, Carbon Footprint.

Funding information: Josh Cowls' and Andreas Tsamados' contributions were supported by a fellowship provided by the Vodafone Institute.

Acknowledgements: The authors wish to gratefully acknowledge the invaluable comments provided on earlier drafts of this article by David Watson and members of the Vodafone Institute.

Introduction

In this article we analyse the role that artificial intelligence (AI) could play, and is already playing, as a technology to combat global climate change. The Intergovernmental Panel on Climate Change (IPCC) has repeatedly emphasised the need for large-scale responses to human-induced climate change to prevent avoidable warming and to mitigate the effects of unavoidable warming as well as that which has already occurred (Masson-Delmotte et al. 2018; Pachauri et al. 2014).

Leveraging the opportunities offered by AI for global climate change is both feasible and desirable, but it involves a sacrifice (ethical risks and potentially an increased carbon footprint) in view of a very significant gain (a more effective response to climate change). It is, in other words, a gambit, which requires responsive and effective governance to become a winning strategy. The overall aim of this article is to contribute to the development of such a strategy. We begin in section 1 by exploring the opportunities that AI affords for combatting climate change, identifying two crucial opportunities: AI can help improve and expand current understanding of climate change; and AI is increasingly part of a package of responses that are essential to combatting the climate crisis effectively, by delivering much greener, more sustainable and effective solutions. However, as we argue, the introduction of AI into the climate domain risks amplifying several social and ethical challenges already associated with AI more generally, such as unfair bias, discrimination, or opacity in decision making.

In section 2 we address the flipside of AI in the context of climate change: the contribution to global climate change of the greenhouse gases emitted by developing computation-intensive AI systems. We focus on the carbon footprint of AI research, and assess the factors that influence AI's greenhouse gas (GHG) emissions in this context,

finding that the carbon footprint of AI research can be significant, and highlighting the need for more scientific evidence concerning the trade-off between the GHG emissions generated by AI research and the energy and resource efficiency gains that AI offers when applied to various tasks and industries.

In section 3, we turn to the wider policy context, and identify the European Union is especially well-placed to adopt effective policy response to the opportunities and challenges presented. Thus in section 4 we provide 13 recommendations for European policymakers and AI researchers that are designed to identify and harness the opportunities of AI for combatting climate change, while reducing the impact of its development on the environment. We conclude our analysis by stressing that risks and benefits of the uses of AI to fight climate change are distinct yet intertwined, and that effective policies and strategies are required to both leverage the potential of AI and minimise the harms it poses in order to protect the environment.

1. AI against climate change

AI is already having a significant positive impact in the fight against climate change. Yet exactly how significant, and precisely what sort of impact, are challenging questions to answer. This section provides an overview of initiatives and projects that rely on AI to fight climate change (1.1), notes work done to capture the nature and scale of the phenomenon (1.2), and identifies a set of obstacles to be overcome to ensure that the use of AI to combat climate change is not only effective but also ethically sound (1.3).

1.1 How AI is used against climate change

AI may be characterised as a set of multipurpose tools and techniques designed to simulate and/or improve upon processes that would have seemed intelligent had a human performed them (McCarthy et al. 2006). At a high level, key cognitive capabilities displayed by “intelligent” machine systems include a combination of classification, prediction, and decision-making. These capabilities are already being deployed in a diverse array of domains, like health (e.g., recognising features in an image such as an X-ray scan for cancer diagnosis), transportation (e.g., using environmental sensors to safely drive a car), and communication (e.g., processing human speech and responding in kind). Applying the

“solution space” of AI to the “problem space” of climate change could yield significant benefits, by helping to understand the problem, and by facilitating effective responses.

Despite scientific consensus about the basic facts of climate change, many aspects of the environmental crisis remain uncertain. This includes the explanation of past and present events and observations, and the accurate prediction of future outcomes. The ability of AI to process enormous amounts of non-structured, multi-dimensional data using sophisticated optimisation techniques is already facilitating the understanding of high-dimensional climate datasets and forecasting of future trends (Huntingford et al. 2019). AI techniques have been used to forecast global mean temperature changes (Ise and Oba 2019; Cifuentes et al. 2020); predict climactic and oceanic phenomena such as El Niño (Ham, Kim, and Luo 2019), cloud systems (Rasp, Pritchard, and Gentine 2018), and tropical instability waves (Zheng et al. 2020); better understand aspects of the weather system — like rainfall, generally (Sønderby et al. 2020; Larraondo et al. 2020) and in specific locales, such as Malaysia (Ridwan et al. 2020) — and their knock-on consequences, like water demand (Shrestha, Manandhar, and Shrestha 2020; Xenochristou et al. 2020). AI tools can also help anticipate the extreme weather events that are more common as a result of global climate change, for example heavy rain damage (Choi et al. 2018) and wildfires (Jaafari et al. 2019), and other downstream consequences, such as patterns of human migration (Robinson and Dilkina 2018). In many cases, AI techniques can help to improve or expedite existing forecasting and prediction systems, for example by automatically labelling climate modelling data (Chattopadhyay, Hassanzadeh, and Pasha 2020), improving approximations for simulating the atmosphere (Gagne et al. 2020), and separating signals from noise in climate observations (Barnes et al. 2019).

Combating climate change effectively requires a vast array of responses to the crisis, which broadly include both mitigating existing effects of climate change and reducing emissions through decarbonisation to prevent further warming. For example, a 2018 Microsoft/PwC report estimated that using AI for environmental applications could boost global GDP by between 3.1 and 4.4%, while reducing greenhouse gas emissions anywhere from 1.5 to 4% by 2030 compared to a “business as usual” scenario (Microsoft 2018, 8). An array of AI-based techniques already play a key role in many of these responses (Inderwildi et al. 2020; Sayed-Mouchaweh 2020). This includes, for example, energy efficiency in industry, especially the petrochemical sector (Narciso and Martins

2020). Studies have also used AI to understand industrial pollution in China (Zhou et al. 2016), the carbon footprint of concrete used in construction (Thilakarathna et al. 2020), and even energy efficiency in shipping (Perera, Mo, and Soares 2016). Other work has explored the use of AI in electrical grid management (Di Piazza et al. 2020), to forecast building energy usage (Fathi et al. 2020), and to assess the sustainability of food consumption (Abdella et al. 2020). AI can also help to predict carbon emissions based on present trends (Mardani et al. 2020; Wei, Yuwei, and Chongchong 2018) and the impact of interventionist policies like a carbon tax (Abrell, Kosch, and Rausch 2019) and carbon trading systems (Lu et al. 2020). AI could also be used to help monitor the active removal of carbon from the atmosphere through sequestration (Menad et al. 2019).

Beyond this indicative evidence, the growing use of AI to fight climate change can also be seen from the higher vantage point of major institutions and large-scale initiatives. The European Lab for Learning & Intelligent Systems (ELLIS) has a Machine Learning for Earth and Climate Sciences programme that aims to “model and understand the Earth system with Machine Learning and Process Understanding”.¹ The European Space Agency has also established a Digital Twin Earth Challenge to provide “forecasting on the impact of climate change and responding to societal challenges”.² On the academic side, the EC-funded iMIRACLI (innovative MachIne leaRning to constrain Aerosol-cloud CLimate Impacts) initiative will support 15 PhD students across nine European universities to “develop machine learning solutions to deliver a breakthrough in climate research”,³ with doctoral projects underway since autumn 2020.

Several European universities have initiatives and training programmes dedicated to unlocking the power of AI for climate.^{4,5,6} Indeed, a search of Cordis — the European database for funded research — for current projects addressing climate change and AI returned a total of 122 results.⁷ Analysis of these 122 projects suggests that they represent both geographic and disciplinary breadth. The projects are well spread across the

¹ <https://ellis.eu/programs/machine-learning-for-earth-and-climate-sciences>

² <https://copernicus-masters.com/prize/esa-challenge/#>

³ <https://imiracli.web.ox.ac.uk>

⁴ <https://www.uv.es/uvweb/uv-news/en/news/ai-understanding-modelling-earth-system-international-research-team-co-led-university-valencia-wins-erc-synergy-grant-1285973304159/Novetat.html>

⁵ <https://www.exeter.ac.uk/research/environmental-intelligence/>

⁶ <https://ai4er-cdt.esc.cam.ac.uk>

⁷ Search of Cordis research project database conducted 30th November 2020 of Projects with search string [(‘climate change’ OR ‘global warming’) AND (‘artificial intelligence’ OR ‘machine learning’)], (n=122).

continent, albeit with a clear skew towards western Europe in terms of where they are coordinated (see Figure 1.1). Figure 1.2 displays the top-level field(s) of study indicated for each of the projects, where this information was provided (n=106). Unsurprisingly, a large majority of projects relate to the natural sciences and/or engineering and technology, but a considerable number are also anchored in social sciences. And as Figure 1.3 shows, at a more granular level, the breadth of subjects that these projects touch on is vast, and spans domains as diverse as viticulture, mycology and galactic astronomy.

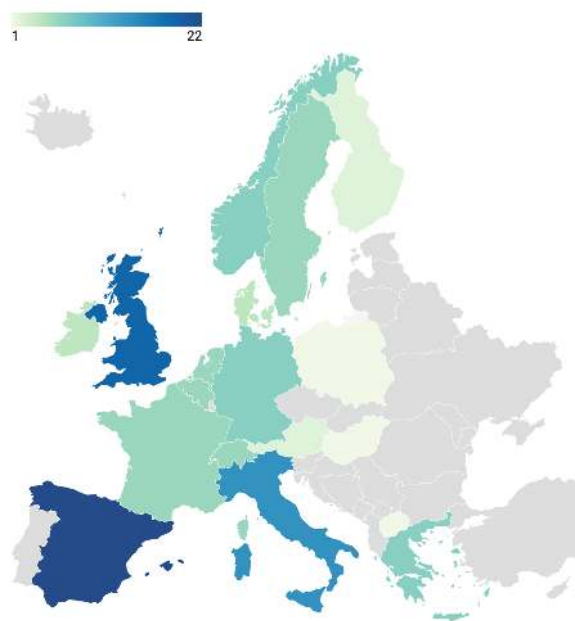



Figure 1.1: Countries in which EU-funded projects using AI to address climate change are “coordinated”.
 Not shown: Israel (1 project).



Uganda.¹⁰ Meanwhile, AI development company ElementAI's AI for Climate program¹¹ provides expertise and partnership opportunities to improve the energy efficiency of manufacturing and business operations.

1.2 How evidence of AI against climate change is gathered

AI is not a silver bullet nor “the only tool in the drawer” for combating climate change, but the previous section illustrates that efforts to use AI to combat climate change are growing at a fast pace. Because of this pace of development, undertaking a more comprehensive, and rigorous, assessment is a challenge. To date, several systematic approaches to gathering evidence of the use of AI for climate change worldwide have been trialled, resulting in a range of datasets, organised in different ways, each of which paints a partial picture of the phenomenon. For instance, some researchers have used the United Nations Sustainable Development Goals (SDGs) as a basis for evidence-gathering about AI-based solutions to address climate change. Of the 17 SDGs, goal 13, “Climate Action”, is most explicitly associated with climate change, but several others, such as 14, “Life Below Water”, and 15, “Life on Land”, are also related. For example, the database of University of Oxford's Research Initiative on AIxSDGs¹² contains 108 projects, of which 28 are labelled as related to Goal 13 (see Figure 1.4); the SDG AI Repository managed by the UN's ITU agency¹³ contains 9 climate-focused projects; and the database of the AI4SDGs Think Tank¹⁴ contains 5.

Clearly, the full range of projects using AI to tackle climate change around the world is not captured in these databases. This may be a result of the selection criteria employed in the surveys, or a lack of awareness of these evidence-gathering efforts among those actually deploying the technology (despite the annual, high-profile AI for Good summit organised by the ITU). It may also be that the SDGs are not the ideal framework, at least scientifically, for exploring the use of AI to tackle climate change. Each SDG contains specific targets and indicators (five and eight respectively in the case of the 13th goal), which are high-level and policy-focused. Consider, for example, indicator 13.1.2, the

¹⁰ <https://ai.google/social-good/impact-challenge/>

¹¹ <https://www.elementai.com/ai-for-climate>

¹² <https://www.sbs.ox.ac.uk/research/centres-and-initiatives/oxford-initiative-aisdgs>

¹³ <https://www.itu.int/en/ITU-T/AI/Pages/ai-repository.aspx>

¹⁴ <http://www.ai-for-sdgs.academy/topics#13%20Climate%20Action>

“number of countries with national and local disaster risk reduction strategies”. Tying the outputs of a single AI initiative to high-level and policy-focused outcomes may prove to be problematic and make the SDGs less than the ideal framework to map the uses of AI to tackle climate change.




Figure 1.4 AI-based projects addressing the UN Sustainable Development Goals (Cowls et al. Forthcoming).

Alternative approaches to mapping the uses of AI to address the climate crisis clarify the phenomenon further. One recent large-scale study pinpointed 37 use cases within 13 domains where AI¹⁵ “can be applied with high impact in the fight against climate change” (Rolnick et al. 2019, 2), and offered a host of examples. For each case, the authors noted which subdomain of the technology (causal inference, computer vision, etc) could be beneficial (see Figure 1.5). Since the publication of this landscaping study, the authors have launched Climate Change AI (CCAI), an organisation composed of “volunteers from academia and industry who believe that tackling climate change requires concerted societal

¹⁵ *Machine learning* is commonly considered to be a subset of the wider set of technological systems that fall under the heading of *artificial intelligence*. Rolnick et al.’s usage of “machine learning”, however, is quite inclusive, capturing a broad array of examples. Rather than split definitional hairs, the evidence assembled is treated here as a comprehensive overview of the ways in which artificial intelligence *per se* can be used to combat climate change.

action, in which machine learning can play an impactful role¹⁶, which has resulted in a wide network of researchers.

	Causal inference	Computer vision	Interpretable models	NLP	RL & Control	Time-series analysis	Transfer learning	Uncertainty quantification	Unsupervised learning
1 Electricity systems									
Enabling low-carbon electricity		•	•		•	•		•	•
Reducing current-system impacts		•				•		•	•
Ensuring global impact		•					•		•
2 Transportation									
Reducing transport activity		•				•		•	•
Improving vehicle efficiency		•			•				
Alternative fuels & electrification					•				•
Modal shift	•	•				•		•	
3 Buildings and cities									
Optimizing buildings	•				•	•	•		
Urban planning		•				•	•		•
The future of cities				•			•	•	•
4 Industry									
Optimizing supply chains		•			•	•			
Improving materials									•
Production & energy		•	•		•				
5 Farms & forests									
Remote sensing of emissions		•							
Precision agriculture		•			•	•			
Monitoring peatlands		•							
Managing forests		•			•	•			
6 Carbon dioxide removal									
Direct air capture									•
Sequestering CO ₂		•						•	•
7 Climate prediction									
Uniting data, ML & climate science		•	•			•		•	
Forecasting extreme events		•	•			•		•	
8 Societal impacts									
Ecology		•					•		
Infrastructure					•	•		•	
Social systems		•				•			•
Crisis		•		•					
9 Solar geoengineering									
Understanding & improving aerosols						•		•	
Engineering a planetary control system					•			•	
Modeling impacts						•		•	
10 Individual action									
Understanding personal footprint	•			•	•	•			
Facilitating behavior change				•					•
11 Collective decisions									
Modeling social interactions			•		•				
Informing policy	•	•		•				•	•
Designing markets					•	•			•
12 Education				•	•				
13 Finance				•		•		•	

¹⁶ <https://www.climatechange.ai/about>

Figure 1.5: Domains of prospective positive climate impact and forms of AI technology relevant to each, from Rolnick et al. (2019).

Each of these approaches to gathering evidence of AI used to combat climate change helps illuminate the nature of the phenomenon and understand better which domains are attracting more efforts and which are potentially overlooked. Consider for example a cross-analysis (Cowls et al. Forthcoming) between the aforementioned Oxford Research Initiative AIxSDGs database and scoping study by Rolnick et al. (2019). Figure 1.6 charts the number of climate change-related projects in the AIxSDGs against the specific domains identified by Rolnick and colleagues. In some domains, such as Farms & Forests, there is clear evidence of projects that met the criteria for inclusion in the AIxSDGs database, whereas in others few if any projects are included. This may result in part from the criteria used in the AIxSDGs database collection, among which was the need for evidence of tangible real-world impact.




Figure 1.6: Projects in the Oxford AIxSDG database working in the different domains identified by Rolnick et al. (2019).

It is clear that AI offers many options for addressing a wide array of challenges associated with climate change. And given the severity and scope of the challenges posed by climate change, it may be advisable to experiment with a wide array of potential solutions across

many domains, as discussed in section 1.1. However, the opportunities offered by AI can only be harnessed to their full potential if ethical values and social expectations are to be met. We turn to these next.

1.3 What are the risks to be avoided or minimised?

Using AI in the context of climate change poses fewer and less severe ethical risks (Tsamados et al. 2020) than using AI in domains such as health and criminal justice, where personal data and direct human-facing decisions are at the core of all activities. Nonetheless, it is important to avoid or minimise the ethical risks that may still arise when maximising the positive impact AI in the fight against climate change.

The first set of risks follows from the way AI models are designed and developed (Yang et al. 2018). Most data-driven approaches to AI are supervised, i.e. they are “trained” on existing labelled data as a basis from which to “learn” to cluster, classify, predict or make decisions regarding new, previously unseen data. This introduces the potential for unwanted bias to enter into the decisions at which an AI system ultimately arrives. This may lead to discrimination and unfair treatment of individuals or groups. Consider, for example, the use of AI to decide where to locate charging stations for electric vehicles (EVs) based on existing patterns of EV use (Tao, Huang, and Yang 2018). It is possible that using AI to decide where to place charging stations based purely on existing patterns of EV ownership — which could be skewed towards wealthier areas — may result in bias against less wealthy areas, in turn disincentivising the uptake of EVs in these areas. In the same vein, attempts to rely on smartphones to infer individuals’ transportation choices (Dabiri and Heaslip 2018) may lead to biased choices unless communities with lower smartphone uptake are properly accounted for.

A second set of risks concerns the erosion to human autonomy that some climate-focused AI systems may pose (Floridi and Cowls 2019; Taddeo and Floridi 2018). Tackling climate change requires large-scale coordinated action, including systematic changes to individual behaviour. As Rolnick et al. note, “understanding individual behaviour can help signal how it can be nudged” (2019 p. 51), for example through limiting people’s “psychological distance” to the climate crisis, helping them visualise its impacts, or encouraging them to take pro-environmental actions. There is considerable debate over the impact of nudging on individual autonomy (Floridi 2016), and whether it prevents

people making “free choices” (for discussion see Schmidt and Engelen 2020), so adopting such an approach in the environmental context requires striking the right balance between protecting individual autonomy and implementing large-scale climate-friendly policies and practices (Coeckelbergh 2020).

Along with fair treatment and autonomy, uses of AI to fight climate change also risk breaching privacy. To the extent to which AI systems rely on non-personal data, e.g. meteorological and geographical data, to understand the climate crisis, they are unlikely to raise privacy concerns. But devising strategies to limit emissions may require data that reveal patterns of human behaviour, where privacy concerns could have more relevance. For example, in control systems designed to decrease carbon footprints in a range of contexts, such as energy storage (Dobbe et al. 2019), industrial heating and cooling (Aftab et al. 2017), and precision agriculture (Liakos et al. 2018), the effectiveness of AI systems depends on granular data about energy demands, often available in real time. The data collected may contain sensitive personal information, risking privacy at both individual and group levels (Floridi 2017). This tension is highlighted in recent Vodafone Institute research finding showing that, while Europeans are broadly willing to share their data to help protect the environment, a clear majority (53%) would only do so under strict conditions of data protection (Vodafone Institute for Society and Communications 2020, 3).

None of these obstacles emerge solely from the use of AI to combat climate change. However, ethical challenges caused by AI may take on novel forms in this context, and therefore require careful responses. Furthermore, the computational cost and potential environmental impact of developing AI systems raises a different set of considerations specific to the climate change domain, which are the focus of the next section.

2. AI’s carbon footprint

AI (both in the sense of training models and of uses) can consume vast amounts of energy and generate greenhouse gas (GHG) emissions (García-Martín et al. 2019; Cai et al. 2020). This is why establishing systematic and accurate measurements of AI’s carbon footprint is key to ensuring that efforts to harness the potential of this technology outweigh its

environmental cost. For reasons explained in Section 2.1, this section focuses on methods to estimate the carbon footprint only of AI research (training models), not of AI uses in general, and the technological and normative factors that contribute to the rise of computationally-intensive AI research.

Following the advent of deep learning (DL), computing power (henceforth compute) usage rose exponentially, doubling every 3.4 months (Amodei and Hernandez 2018), as specialised hardware to train large AI models became central to the research field (Hooker 2020). The increase in energy consumption associated with training larger models and with the widespread adoption of AI has been in part mitigated by hardware efficiency improvements (Ahmed and Wahed 2020; Wheeldon et al. 2020). However, depending on where and how energy is sourced, stored and delivered, the rise of compute-intensive AI research can have significant, negative environmental effects (Lacoste et al. 2019).

2.1 Gauging the carbon footprint of AI

A “carbon footprint” accounts for the GHG emissions of a device or activity, expressed as carbon dioxide equivalent (CO₂eq). When applied to a product like a smartphone, a carbon footprint estimation considers emissions that occur during constituent activities, like the extraction of raw materials, manufacturing, transportation, lifetime usage and how the product is disposed of (Crawford and Joler 2018; Malmodin and Lundén 2018). This estimate includes, among other things, information on the carbon/emission intensity of electricity generation throughout a product’s lifecycle and on the carbon offsetting efforts made by the various actors involved in the aforementioned activities (Matthews, Hendrickson, and Weber 2008). However, determining the carbon footprint of a type of product (e.g. AI systems) or entire sector (e.g. Information Communication Technologies, ICT) can be a daunting task that yields only partial results, not least due to transparency issues and methodological challenges of GHG monitoring (Matthews, Hendrickson, and Weber 2008; Russell 2019; Cook and Jardim 2019; Mytton 2020).

Estimates of GHG emissions of the ICT sector (including computing devices and data centres) vary greatly across different studies (Malmodin and Lundén 2018; Hintemann and Hinterholzer 2020). Malmodin and Lundén’s (2018), a widely cited study based on data from 2015, estimates that the ICT sector is responsible for 1.4% of global GHG emissions. Depending on future efficiency gains and the diversification of energy sources,

estimates indicate that the ICT sector will be responsible for anywhere between 1.4% (assuming a stagnant growth) to 23% of global emissions by 2030 (Andrae and Edler 2015; Malmodin and Lundén 2018; C2E2 2018; Belkhir and Elmeligi 2018; Jones 2018).¹⁷ At the same time, it is worth noting that the demand for data centres, which are key to the ICT sector and the operation of AI in research and production settings, has grown substantially in recent years, yet data centres' energy consumption has remained relatively stable (Avgerinou, Bertoldi, and Castellazzi 2017; Shehabi et al. 2018; Jones 2018; Masanet et al. 2020). The International Energy Agency reports that, if current efficiency trends in hardware and data centre infrastructure can be maintained, global data centre energy demand — currently 1% of global electricity demand — “can remain nearly flat through 2022, despite a 60% increase in service demand” (IEA 2020). Indeed, significant efforts have been made to curb data centres' carbon footprint by investing in energy-efficient infrastructure and switching to renewable sources of energy (Jones 2018; Masanet et al. 2020). Cloud providers especially, such as Microsoft Azure and Google Cloud, have worked to keep their carbon footprint in check by committing to renewable energy, improving cooling systems and efficient processors, recycling waste heat, and investing in carbon offsetting schemes (Jouhara and Meskimmon 2014; Avgerinou, Bertoldi, and Castellazzi 2017; Jones 2018; Open Compute Project 2020). In fact, both providers have leveraged AI to reduce the energy consumption of their data centres, in some cases by up to 40% (Evans and Gao 2016; Microsoft 2018).

Whether these efforts keep pace with the growing demand for data centre services and whether efficiency gains are equally realised around the world will be crucial factors affecting the environmental impact of the sector. These goals may not be easily achievable. Even in the EU, where energy-efficient cloud computing has become a primary issue on the political agenda, the European Commission estimates a 28% increase in energy consumption of data centres by 2030 (European Commission 2020d). Things are complicated even further by transparency concerns regarding the data required to calculate GHG emissions of on-premise data centres as well as cloud vendors, which will need to be addressed in order to obtain an accurate understanding of the carbon footprint of the ICT sector (Hintemann 2015; Mytton 2020; Hintemann and Hinterholzer 2020).

¹⁷ In fact, a new standard (L.1470) set out by the ITU was recently developed to keep the ICT industry in line with the Paris Agreement and reduce GHG emissions by 45% from 2020 to 2030 (ITU 2020).

At the same time, understanding the carbon footprint of AI involves more than just monitoring data centres, as the rest of this section will show (Henderson et al. 2020; Cai et al. 2020). Given the wide range of artefacts and activities relying on some form of AI and the multi-layered production process of AI systems — spanning from data collection and storage, to hardware production and shipment, to AI/machine learning (ML) model trainings and inferences — gauging the carbon footprint of AI is challenging. This is why this section focuses on the carbon footprint associated with the energy consumption of AI research activities, available in corresponding research publications.¹⁸ As we shall see in section 2.3, verifiable information on the short- and projected medium-term environmental impact of AI research activities is limited and suffers from a lack of systematic and accurate measurements. However, the information contained in research publications regarding the energy consumption and carbon emission of AI is more accessible and testable than in industry reports. Thus, it offers a more reliable starting point to understand the environmental impact of AI, even if it is indicative only of a subset of all AI-related activities. Furthermore, to gauge the energy consumption and carbon footprint of AI research activities, it is important to distinguish between two phases of computation that are central to supervised ML research methods: training (or “learning”) and inferences. Training a ML model involves providing labelled sample data, or a “training set”, to a ML algorithm so that it can “learn” from it and create an appropriate mathematical model with the optimal parameters that minimise a certain cost function (e.g. some metric of error). Once the training phase is finished, a model and its parameters are fixed and such model can be operationalised and produce actionable output on new, unseen data, which is the “inference” process.

In the short term, the training phase is computationally more demanding and energy intensive (Al-Jarrah et al. 2015). In the medium term, the energy consumption of the inference phase scales with usage, as inference can usually occur millions of times per day for an indefinite amount of time (Sze et al. 2017). So, training is often more energy-

¹⁸ Over the past two decades the number of data-driven AI conferences and publications have grown dramatically (Perrault et al. 2019). The volume of peer-reviewed AI papers increased by more than 300% from 1998 to 2018, while the number of publications on “Machine Learning” on the open-access archive “ArXiv” has doubled every 15 months since 2015 (Perrault et al. 2019). This growth is also fuelled by an increasing amount of publications originating from the private sector — big technology companies in particular — with their seemingly limitless resources to conduct experiments (Perrault et al. 2019, 17; Lohr 2019; Ahmed and Wahed 2020).

intensive in data-driven, ML-based research, while inference might be more energy-intensive in at-scale production systems which may require non-stop usage. This is why, in the context of AI as a whole, this report focuses on information pertaining to the research and training of AI models.

Several approaches to monitoring and estimating the GHG emissions of AI research activities have been recently offered. These include the reporting of floating point operations (Lacoste et al. 2019; Schwartz et al. 2019; Henderson et al. 2020), hardware type and hardware burden or “processors multiplied by the computation rate and time” (Thompson et al. 2020, 10), the data centre in use during model training, as well as the energy sources powering the electrical grid (Schwartz et al. 2019; Anthony, Kanding, and Selvan 2020), the number of experiments required during model construction (Schwartz et al. 2019; Strubell, Ganesh, and McCallum 2019), and the time period in which a model was trained, as carbon/emission intensity can vary throughout the day (Anthony, Kanding, and Selvan 2020). Of these approaches, two recent efforts stand out for their generalisability and/or ease of use, namely Henderson et al’s (2020) “experiment-impact-tracker” and Lacoste et al’s (2019) Machine Learning Emissions Calculator.

The first approach rests on a comprehensive framework available on GitHub (Henderson et al. 2020), specifying the relevant data to collect during and after model training phases to assess the related GHG emissions:

1. Central processing unit (CPU) and graphics processing unit (GPU)¹⁹ hardware information
2. experiment start and end-times
3. the energy grid region the experiment is being run in (based on IP address)
4. the average carbon/emission intensity in the energy grid region
5. CPU- and GPU-package power draw
6. per-process utilisation of CPUs and GPUs
7. GPU performance states
8. memory usage
9. the real-time CPU frequency (in Hz)
10. real-time carbon intensity

¹⁹ It is important to also include Field Programmable Gate Arrays (FPGAs) and Application Specific Integrated Circuits (ASICs) like Tensor Processing Units (TPUs) to this framework.

11. disk write speed.

Unfortunately, information about these 11 variables is rarely available in its entirety in most research publications (Henderson et al. 2020). In an analysis of 1,058 research papers on DL, Thompson et al. (2020, 10) found that most papers “did not report any details of their computational requirements”.

By contrast, the second approach (Lacoste et al.’s 2019) limits itself to information pertaining to the type of hardware, hours of training, region of compute, and cloud provider/private infrastructure. This is a helpful approach to estimating the carbon footprint of AI research activities using a minimum amount of data and without actually reproducing experiments and models. For this reason, we use Lacoste et al.’s (2019) approach to calculate the carbon footprint of large AI research projects, and we use GPT-3, OpenAI’s latest research breakthrough as our case study. While the following estimates cannot be definitive, due to a lack of data available in OpenAI’s research publication, they serve to reflect both the importance and the difficulty of assessing carbon footprints when researchers fail either to report them or to provide enough information regarding training infrastructure and model implementation.

GPT-3 is an autoregressive language model that has attracted considerable attention from researchers and news outlets since documentation was published on ArXiv in May 2020 by Brown et al (2020). From the research publication detailing GPT-3, we know that the model required *several thousands* of petaflop/s-days ($3.14E23$ FLOPS) of compute during pre-training. This is orders of magnitude higher than the previous state-of-the-art (SOTA) 1.5B parameter GPT-2 model that the company released in 2018, which required only *tens* of petaflop/s-days (Radford et al. 2018). GPT-3 was trained using NVIDIA’s V100 GPUs on a cluster provided by Microsoft. Thus, one can calculate that, at a theoretical processing speed of 28 Terra Flops (TFLOPS)²⁰ for a V100 GPU, it would take around 355 GPU years for a single training run (Li 2020).

Using Lacoste et al.’s carbon impact calculator and assuming that the cloud provider (Microsoft Azure) was based in the US (West), we find that a single training run would have generated 223,920kg CO₂eq. If the cloud provider had been Amazon Web

²⁰ The 28 TFLOPS is assumed here based of NVIDIA’s advertisement of the V100 performance as well as on Microsoft’s DeepSpeed and ZeRO-2 performance results for training +100-billion-parameter models using V100 GPUs (NVIDIA 2018; Rangan and Junhua 2020).

Services (AWS), the same training would have generated 279,900kg CO₂eq.²¹ This does not include the carbon offsetting efforts made by these companies (Mytton 2020). As a point of reference, a typical passenger car in the United States emits about 4,600kg CO₂eq per year (US EPA 2016), meaning that one training run would emit as much as 49 cars (Microsoft Azure) or 61 cars (AWS) in a year. A single training run can emit drastically more GHG depending on the region of compute and the carbon/emission intensity of electricity generation in the selected region (Lacoste et al. 2019). For example, it is 10 times more costly in terms of CO₂eq emissions to train a model using energy grids in South Africa compared to France (see compute regions in Lacoste et al. 2019). Figure 3.1 below offers examples of the variation of energy consumption across different countries.

	Carbon emissions (CO ₂ eq)	Train Compute (FLOPS)	GPU	Training hours	Cloud Provider
South Africa (West)	942,330kg	3.14E+23	V100	3.11E+06	Microsoft Azure
India (South)	858,360kg	3.14E+23	V100	3.11E+06	Microsoft Azure
Australia (Central)	839,700kg	3.14E+23	V100	3.11E+06	Microsoft Azure
Europe (North)	578,460kg	3.14E+23	V100	3.11E+06	Microsoft Azure
South Korea (Central)	485,160kg	3.14E+23	V100	3.11E+06	Microsoft Azure
Brazil (South)	186,600kg	3.14E+23	V100	3.11E+06	Microsoft Azure
France (Central)	93,300kg	3.14E+23	V100	3.11E+06	Microsoft Azure

Figure 3.1 Environmental costs (in kg of CO₂eq) of a single training run of GPT-3 across different compute regions.²²

The authors of GPT-3 (Brown et al. 2020) also note that training the model required an immense amount of resource, but GPT-3 has the advantage of adapting to new tasks quite efficiently compared to other language models that would be relatively costly to fine-tune. For example, the authors report that to generate “100 pages of content from a trained

²¹ Using a different set of assumptions and a methodology similar to that of Henderson et al. (2020), a group of researchers from the University of Copenhagen have estimated the training run would emit up to 84,738.48kg CO₂eq in the US (region not specified) (Anthony, Kanding, and Selvan 2020). This highlights the importance of disclosing enough information in research publications.

²² Regional carbon intensity sourced from <https://github.com/mlco2/impact/tree/master/data>

model can cost on the order of 0.4 kW-hr, or only a few cents in energy costs” (Brown et al. 2020).

More recently, researchers from the Google Brain team released a research paper stating to have trained a 1.6 trillion parameters language model — approximately 9 times bigger than GPT-3 (Fedus, Zoph, and Shazeer 2021). And although the paper describes the use of a training technique that reduces computational costs and increases model training speed, it does not indicate the energy consumption or carbon emissions of the research project. This comes against the backdrop of earlier warnings from Google’s own Ethical AI team regarding the environmental costs of such large models (Bender, Gebru, and McMillan-Major 2021).

It is crucial for the field of AI to come to terms with these numbers. These large AI research projects may be indicative of — and exacerbate — a failure to engage with environmental questions, to disclose important research data, and to shift the focus away from ecologically short-sighted success metrics (García-Martín et al. 2019; Schwartz et al. 2019; Henderson et al. 2020). In what follows we explore the technological and normative factors that have entrenched the field of AI research on an energy-intensive, and potentially carbon-intensive, path.

2.2 Factors driving increases in AI’s carbon footprint

Technological Considerations: Compute-Intensive Progress

The recent rise of AI can be largely attributed to the increasing availability of massive amounts of data and to the adoption of general methods leveraging the “continued exponentially falling cost per unit of computation” described by Moore’s law (i.e. the number of transistors per microchip doubles every two years for the same costs) (Sutton 2019). DL epitomises AI research that is based on scaling general purpose methods with increased computation and availability of large amounts of non-structured data (Sutton 2019). Recent breakthroughs, where AI models were able to reach parity with humans on a number of specific tasks, are the result of such AI research based on deep neural networks and improvements in computation and data availability (Ahmed and Wahed 2020; Hooker 2020). However, the advent of DL has also marked a split between the increase in available compute (i.e. Moore’s law) and the increase in compute-*usage* (Theis and Wong 2017; N. Thompson and Spanuth 2018; Ahmed and Wahed 2020). Exploring

these trends helps us map the risks and opportunities of AI research with regards to climate change.

Moore’s law has resulted in developers being able to double an application’s performance for the same hardware cost. Prior to 2012, AI developments have closely mirrored Moore’s law, with available compute doubling approximately every two years (Perrault et al. 2019). As shown in Figure 3.2, improvements in computer hardware provided almost a 50,000× improvement in performance, while the computational requirements of neural networks had grown at a similar pace until the introduction of chips with multiple processor cores (Hill and Marty 2008; Thompson et al. 2020, 8). Arguably, hardware development has often determined what research activities would be successful (Hooker 2020). For example, deep convolutional neural networks and backpropagation, which are central components to contemporary DL research, had already been introduced in the 80s (Fukushima and Miyake 1982; Werbos 1988), but had real impact only four decades later, following hardware progress and large-scale data availability (Hooker 2020).




Figure 3.2 Computing power demanded by DL throughout the years (figure taken from Thompson et al. 2020).

As described by Thompson et al. (2020, 8), during the “multicore era”, DL was “ported to GPUs, initially yielding a 5 – 15× speed-up which by 2012 had grown to more than 35×” and which led to the AlexNet breakthrough in 2012 (Alom et al. 2018). Shortly after

the AlexNet breakthrough in image recognition, a number of achievements followed in the various sub-fields of AI. In 2015, a Reinforcement Learning (RL) system achieved human-level performance in a majority of Atari games; in 2016 object recognition reached human parity and AlphaGo beat one of the world's greatest Go players; in 2017 speech recognition reached human parity; in 2018 reading comprehension, speech synthesis and machine translation all reached human parity; and in 2019, the ability to scan and extract contextual meaning from text and speech (and answer a series of interconnected questions) reached human parity (Alom et al. 2018; Microsoft 2019; Evans and Gao 2016). These breakthroughs were all possible due to considerable increases in compute-usage (Ahmed and Wahed 2020). Indeed, since 2012, compute-usage has been doubling every 3.4 months, spearheaded by the development of DL (Amodei and Hernandez 2018).

Increases in compute have been essential, especially to RL, as this is an area of ML that stands out for its sample-inefficient methods of learning. Learning phases can require hundreds of millions of samples, making it impractical for “real-world control problems” such as in robotics (Buckman et al. 2018). Yet, RL has been used for text summarisation, robotic manipulation, and also to compete with human performance in domains such as Atari games, Chess, and Go (Berner et al. 2019). As researchers begin to apply RL methods to increasingly complex domains, like online multiplayer games, sample inefficiency will continue to drive energy costs higher. For example, OpenAI Five, which was developed to compete with professional Dota 2 players, played 900 years' worth of games per day to reach a competitive level at the game (Berner et al. 2019). After ten months of training, using around 770 Peta Flops/s·days of compute, the model beat the world champions at Dota 2 (Berner et al. 2019).

The multicore era also marks a decoupling of the improvements in hardware performance from the growth in computational requirements of large AI models, with the latter considerably outpacing the former. Because of this, researchers are facing diminishing returns (Thompson et al. 2020). The compute needed to train SOTA models is growing approximately ten times faster than GPU performance per watt (Thompson et al. 2020). This means that the present trend in scaling ML models is unlikely to be a sustainable path forward, both in terms of financial costs and for the preservation of the planet, given the very high levels of energy consumption that are associated with it (Henderson et al. 2020; Thompson et al. 2020). As shown in Figure 3.3, it would be

financially and ecologically prohibitive to reach lower error rates in different tasks, as any improvement (measured in percentage points) on a model’s accuracy would require significantly more energy and GHG. For example, if we look at the Thompson et al.’s (2020) polynomial models, it appears that reducing the error rate by 16.7 percentage points for MS COCO (Common Objects in Context) to achieve an error rate of 30%, would require $10^9 \times$ more computation (GFLOPS) and generate $10^8 \times$ more CO₂eq (in lbs).

Benchmark	Error rate	Polynomial			Exponential		
		Computation Required (Gflops)	Environmental Cost (CO ₂)	Economic Cost (\$)	Computation Required (Gflops)	Environmental Cost (CO ₂)	Economic Cost (\$)
ImageNet	Today: 11.5%	10 ¹⁴	10 ⁶	10 ⁶	10 ¹⁴	10 ⁶	10 ⁶
	Target 1: 5%	10 ¹⁹	10 ¹⁰	10 ¹¹	10 ²⁷	10 ¹⁹	10 ¹⁹
	Target 2: 1%	10 ²⁸	10 ²⁰	10 ²⁰	10 ¹²⁰	10 ¹¹²	10 ¹¹²
MS COCO	Today: 46.7%	10 ¹⁴	10 ⁶	10 ⁶	10 ¹⁵	10 ⁷	10 ⁷
	Target 1: 30%	10 ²³	10 ¹⁴	10 ¹⁵	10 ²⁹	10 ²¹	10 ²¹
	Target 2: 10%	10 ⁴⁴	10 ³⁶	10 ³⁶	10 ¹⁰⁷	10 ⁹⁹	10 ⁹⁹
SQuAD 1.1	Today: 4.621%	10 ¹³	10 ⁴	10 ⁵	10 ¹³	10 ⁵	10 ⁵
	Target 1: 2%	10 ¹⁵	10 ⁷	10 ⁷	10 ²³	10 ¹⁵	10 ¹⁵
	Target 2: 1%	10 ¹⁸	10 ¹⁰	10 ¹⁰	10 ⁴⁰	10 ³²	10 ³²
CoLLN 2003	Today: 6.5%	10 ¹³	10 ⁵	10 ⁵	10 ¹³	10 ⁵	10 ⁵
	Target 1: 2%	10 ⁴³	10 ³⁵	10 ³⁵	10 ⁸²	10 ⁷³	10 ⁷⁴
	Target 2: 1%	10 ⁶¹	10 ⁵³	10 ⁵³	10 ¹⁸¹	10 ¹⁷³	10 ¹⁷³
WMT 2014 (EN-FR)	Today: 54.4%	10 ¹²	10 ⁴	10 ⁴	10 ¹²	10 ⁴	10 ⁴
	Target 1: 30%	10 ²³	10 ¹⁵	10 ¹⁵	10 ³⁰	10 ²²	10 ²²
	Target 2: 10%	10 ⁴³	10 ³⁵	10 ³⁵	10 ¹⁰⁷	10 ⁹⁹	10 ¹⁰⁰

Figure 3.3 Implications of achieving performance benchmarks on the computation, carbon emissions (lbs), and economic costs from deep learning based on projections from polynomial and exponential models (figure from Thompson et al. 2019).

To enable the current compute-usage trend and mitigate diminishing returns, ML-specific hardware, such as Google’s TPUs, and various approaches, like neural architecture search, have been developed in recent years (Amodei and Hernandez 2018). However, even as new devices and hardware architectures continue to deliver better energy efficiency, it is not guaranteed that these improvements will keep pace with compute-usage, nor are they guaranteed to be available to everyone around the globe (Ahmed and Wahed 2020; Hooker 2020). It follows that, if AI researchers are unable to access SOTA hardware to train large ML models, or if hardware performance does not keep pace with the growth of compute-usage in AI research, then the field’s energy consumption will grow quickly (Nature Electronics 2018). Additionally, research has shown that the current focus on DL and custom hardware has come at the detriment of funding “hardware for use cases that are

not immediately commercially viable”, making it more costly to diversify research (Hooker 2020, 9).

Algorithmic progress has also shown promising effects in relation to efficiency improvements for large model trainings. Although algorithmic progress is more dependent on human knowledge — as opposed to computational advances — and thus takes more time and effort to occur (Sutton 2019), Thompson et al. (2020) note that three years of algorithmic improvement is equivalent to an increase in computing power by a factor of 10. This can be observed in image recognition (Hernandez and Brown 2020), neural machine translation (Thompson et al. 2020), and certain areas of RL (Hernandez and Brown 2020). For example, since 2012 the compute required to train a neural network to the “same performance on ImageNet classification, has been decreasing by a factor of 2 every 16 months” and it now takes “44 times less compute to train a neural network to the level of AlexNet” (Hernandez and Brown 2020). Nevertheless, we note that research exploring new neural network architectures or new hardware-software-algorithm combinations has largely been side-lined in favour of compute-intensive AI research (Hooker 2020; Marcus 2020; Ahmed and Wahed 2020).

Fortunately, researchers have also sought to reduce the computational burden and energy consumption of AI by focusing on building more efficient models through various approaches, such as random hyperparameter search, pruning, transfer learning or simply by stopping training early for underperforming models (Sze et al. 2017; Pham et al. 2018; C.-F. Chen et al. 2019; Schwartz et al. 2019; Coleman, Kang, Narayanan, Nardi, Zhao, Zhang, Bailis, Olukotun, Re, et al. 2019). More efforts are required in these areas. To be successful they need endorsement and cultivation from the wider field of AI to gain larger uptake. For AI research to continue to thrive, while keeping its carbon footprint in check and avoid running into a technological impasse in the coming years (Jones 2018), the field needs to reconsider its dedication to compute-intensive research and move away from performance metrics that focus exclusively on accuracy improvements (Schwartz et al. 2019). The following section addresses the normative factors that have enabled these negative trends.

Normative Considerations

For a field of research that relies on data collection and data processing, information about the energy consumption and carbon emissions of AI/ML models and research activities should be more detailed and more accessible (Henderson et al. 2020). Indeed, alongside the technological factors that have skewed the development of AI/ML models, there are also normative factors, like the lack of effective reproducibility requirements for research publications, which also contribute to explain the entrenchment of research activities in energy-intensive practices (Hooker 2020; Fursin 2020). Examining these factors, and questioning the validity of some standards and practices in AI research, is key to ensuring that the field keeps its carbon footprint to a minimum.

AI research has been grappling with a reproducibility crisis. Given the growing amount of AI-related research activities and compute-usage, this crisis needlessly supercharges the field's carbon footprint (Fursin 2020). From papers that do not disclose their code (as is the case for GPT-3) to papers that do not share the data used to train their model (e.g. for privacy or proprietary reasons) to papers that provide insufficient or even misleading information about the training conditions of their models, there have been persistent obstacles to verifying and reproducing results in AI research (Gibney 2020; Fursin 2020). In turn, these obstacles translate into unnecessary energy consumption.

After conducting a survey of 400 algorithms presented in research papers at two top AI conferences (IJCAI and NeurIPS), researchers reported that only 6% of the presented papers shared the algorithm's code, a third shared the data on which they tested their algorithms, and only half shared a partial summary of the algorithm (Gundersen and Kjensmo 2018; Hutson 2018). Several studies have investigated this issue in the context of energy consumption and carbon emissions (Lacoste et al. 2019; Schwartz et al. 2019; Strubell, Ganesh, and McCallum 2019; Henderson et al. 2020; Dhar 2020). Indeed, after analysing a sample of 100 papers from the NeurIPS 2019 proceedings, Henderson et al. (2020, 4) reported that none of them provided carbon metrics, only one of them “measured energy in some way, 45 measured runtimes in some way, 46 provided the hardware used” and 17 of them “provided some measure of computational complexity (e.g., compute-time, FPOs, parameters)”. Although major AI conferences, such as ICML, IJCAI or NeurIPS, are increasing their efforts to normalise the submission of code and have implemented reproducibility checklists, the disclosure of information regarding

computational complexity, energy consumption, and carbon emission is still uncommon (Strubell, Ganesh, and McCallum 2019; Thompson et al. 2020).

Sharing source code is necessary to ensure reproducibility in AI research. But it is not sufficient. Researchers have highlighted the importance of disclosing the training data and the initial parameters set for the training phase, or hyperparameters (Schwartz et al. 2019; Hartley and Olsson 2020). Sharing a model's sensitivity to hyperparameters, or of the random numbers generated to start the training process in the case of RL, is essential to allow researchers to reproduce results without going through a long, and environmentally costly, process of trial and error (Hutson 2018; Strubell, Ganesh, and McCallum 2019; Gibney 2020). Indeed, the number of experiments run by researchers before achieving publishable results are both “underreported and underdiscussed” (Dodge et al. 2019; Schwartz et al. 2019, 9). In this case, a direct result of incomplete or misleading information disclosures is the “double costs” incurred by researchers that have to rediscover, even if only partially, the information that led to the reported results. Building on existing research becomes more difficult when newcomers have to incur unnecessary costs of experimentation that were already incurred for the original publication of a model. This approach inflicts an unnecessary double cost on the environment via increased energy consumption.

According to recent research on the energy consumption and carbon footprint of DL in natural language processing (NLP), the process of researching and developing SOTA models multiplies the financial and environmental costs of training a model by thousands of times (Strubell, Ganesh, and McCallum 2019). Indeed, over the course of six months of research and development, a single research paper may require training thousands of models before being published (Dodge et al. 2019; Schwartz et al. 2019, 4). Similarly, Schwartz et al' (2019, 4) have reported that massive amounts of computation go into “tuning hyperparameters or searching over neural architectures”. This is the case, for example, of Google Brain, which trained over 12,800 neural networks in its neural architecture search to achieve a 0.09 percent accuracy improvement and 1.05x in speed on the CIFAR-10 dataset (Zoph and Le 2017). In light of our calculations regarding the carbon emission of a single training run for GPT-3, this would mean that to achieve their published model the research team at OpenAI may have generated much more CO₂eq than previously estimated. Failing to report the research experiments that went into

achieving the reported results can have a snowball effect for the field of AI research in terms of energy consumption and carbon emissions, as it imposes a longer trial-and-error process onto new researchers.

Modern AI research has focused on producing deeper and more accurate models at the detriment of energy efficiency (Sutton 2019; Perrault et al. 2019; Hooker 2020). Indeed, some of the main benchmarks, challenges and leader boards on which AI researchers and organisations compete, such as GLUE (2020), SuperGLUE (2020), SQuAD2.0 (2020), ILSVRC (2017) and VTAB (2020), have been heavily focused on driving accuracy improvements with little regard for improving on energy efficiency (Perrault et al. 2019; Reddi et al. 2020). This narrow focus increases compute-intensive AI research and exacerbates diminishing returns, with researchers competing for fractional improvements in error rates (Henderson et al. 2020). It is only relatively recently that efforts have emerged to reduce compute-usage and improve energy efficiency of DL methods, at the algorithmic, hardware, as well as implementation levels (T. Chen et al. 2015; Y.-H. Chen, Emer, and Sze 2017; EDL 2017; Sze et al. 2017; Guss et al. 2019; García-Martín et al. 2019; Jiang et al. 2019; Cai et al. 2020). The *Low Power Image Recognition Challenge* (LPIRC) is a good example of such efforts (García-Martín et al. 2019).

To demonstrate the prevalence of accuracy metrics over efficiency metrics, a group of researchers at the Allen AI Institute sampled 60 papers from top AI conferences (ACL, CVPR and NeurIPS) that claimed to achieve some kind of improvement in AI (Schwartz et al. 2019). As shown in Figure 3.4, a large majority of the papers target accuracy (90% of ACL papers, 80% of NeurIPS papers and 75% of CVPR papers), and in both ACL and CVPR, which are empirical AI conferences, only 10% and 20% respectively argue for new efficiency results (Schwartz et al. 2019). The prevalence of accuracy over efficiency in AI research has also been stressed by the Electronic Frontier Foundation’s “AI Progress Measurement” project, which tracks progress on problems and metrics/datasets from the AI research literature and provides a comprehensive view of the field’s priorities (EFF 2017).²³

²³ It is also worth noting the carbon footprint associated with the steep growth in attendance at AI conferences. The Stanford AI Index reports that in 2019, NeurIPS had an increase of 41% in attendance over 2018 and over 800% relative to 2012 (Perrault et al. 2019). Other conferences such as the AAAI and CVPR are also seeing an annual attendance growth of around 30% (Perrault et al. 2019). These generate a non-negligible amount of carbon emissions (Perrault et al. 2019; Henderson et al. 2020). Using the attendance of the top 10 AI conferences in 2019, Henderson et al. (2019) estimate that around 34,449,597 kg of CO₂eq were emitted from these conferences alone.




Figure 3.4. Proportion of papers that target accuracy, efficiency, both or other from a sample of 60 papers (figure from Schwartz 2019).

Several issues arise from focusing on accuracy over efficiency metrics. First, it creates a high barrier to entry, as only wealthy research groups are able to incur in the growing costs of compute-intensive research (Ahmed and Wahed 2020). This leads to a limited number of researchers to be able to afford stronger results and hence publications (Schwartz et al. 2019, 2), thus creating a virtual monopoly on fundamental research and side-lining researchers from smaller organisations, less funded contexts, and developing countries (see Figure 3.5). Second, it ingrains a “the bigger the better” mentality into the field, thus giving *carte blanche* to organisations and researchers to accelerate experimentation and increase their eventual energy consumption. This, in turn, makes it harder to explore efficiency improvements. It also reduces the diversity of research topics. Third, it keeps the field on a path of diminishing returns and incentivises researchers to pursue incremental improvements and “publish at all cost”, even if it means achieving practically (for deployed systems) negligible accuracy improvements.




Figure 3.5 Number of deep learning papers on ArXiv, per region (figure from Perrault et al. 2019).

2.3 Overall balance

It is important to keep in mind that, although training AI/ML models can require a lot of energy, they are usually used to improve the efficiency of many tasks that would otherwise require more time, space, human effort, and potentially electricity (Narciso and Martins 2020). When deployed in production settings or edge devices, AI systems can have downstream effects that counterbalance their own energy consumption and GHG emissions (see Section 1). Additionally, recent progress in making the deployment of deep neural networks on edge devices, like smartphones and tablets, much more efficient, has been significant for the environmental impact of AI (Cai et al. 2020). Indeed, the diversity of hardware platforms in use today has created various efficiency constraints requiring, for example, that neural networks are redesigned and retrained for each new environment they are deployed in (Cai et al. 2020). However, novel approaches, such as “Once-for-All” networks, show great promise as they require approximately 1/1,300 of the carbon

emissions of SOTA neural architecture search approaches while also reducing inference time (Cai et al. 2020).²⁴

For most or all industry sectors, AI offers significant “gains in efficiency and performance” (Hall and Pesenti 2017, 2), and indeed, the European Commission’s Horizon 2020 programme has been investing in projects using AI systems to improve the energy and resource efficiency of many sectors (Dahlquist 2020). Balancing the energy consumption of AI against its energy-efficiency gains will be an important task for both researchers and regulators alike, one that begins with obtaining enough information about a model.

Our analysis thus far is testament to the complexity inherent to any attempt to talk about AI in the context of climate change. On the one hand, the power of AI can be harnessed to address some of the most complex tasks associated with combatting climate change successfully. On the other hand, the development of AI is itself contributing to GHG emissions that advance climate change still further. Taken together, this analysis suggests the need for coordinated, multilevel policymaking that can advance the use of AI to combat climate change, whilst ensuring that the development of AI does not itself contribute to the existing problem. This is why in the remainder of this report we turn from technological to political considerations for AI and climate change.

3. Policy context: the EU’s twin transitions

The opportunities presented by AI for tackling climate change are just one example of the broader intersection between the digital revolution and the efforts for sustainability. In recent years, this “Green & Blue” formula (Floridi and Nobre 2020; Floridi 2020) has become apparent in European policymaking. “A European Green Deal” and “A Europe fit for the digital age” were two of the six “headline ambitions” highlighted in the political guidelines released as part of von der Leyen’s campaign. Since von der Leyen took office, documents issued by the Commission have begun referring to the “twin transitions” — ecological and digital — that will shape Europe’s medium- to long-term future. The

²⁴ Once-for-all networks support “diverse architectural settings by decoupling training and search”, thus reducing GPU usage and carbon emissions (Cai et al. 2020, 1).

proximity and interdependence of “green plus blue” has recently been further reinforced by the Commission’s response to the coronavirus pandemic, which anticipates a large stimulus including “modernisation” through “fair climate and digital transitions, via the Just Transition Fund and the Digital Europe Programme” (European Commission 2020i). References to the role of digital technology are studied throughout the “European Green Deal” communication and roadmap, released in December 2019 (European Commission 2019). As the document notes, “digital technologies are a critical enabler for attaining the sustainability goals of the Green Deal in many different sectors”, and technologies “such as artificial intelligence ... can accelerate and maximise the impact of policies to deal with climate change and protect the environment” (p. 9). Domains in which “smart” or “innovative” digital technologies are expected to play a role include energy grids (p. 6), consumer products (p. 8), pollution monitoring (p. 9), mobility (p. 10), and food and agriculture (p. 12) — that is, many of the domains in which the existing evidence, summarised in section 1, suggests that AI is already being deployed and will make an increasing difference.

Recent Commission documents on Europe’s forthcoming “digital transformation” equally highlight the possibilities this transformation holds for sustainability. As the Commission’s recently released “Strategy on shaping Europe’s digital future” (European Commission 2020e) notes, digital technologies will “be key in reaching the ambitions of the European Green Deal and the Sustainable Development Goals” as “powerful enablers for the sustainability transition” (European Commission 2020, p.5). The document highlights sectors including agriculture, transport and energy as benefiting particularly from digital “solutions”. In addition, other Commission documents released within the last twelve months also highlight the twin transitions:

- The “European strategy for data” announces the establishment of a “common European Green Deal data space” to use shared data to meet Green Deal targets (European Commission 2020b);
- The “White Paper on Artificial Intelligence: A European approach to excellence and trust” highlights the impact of AI on climate change mitigation and adaptation in its first paragraph and again throughout (European Commission 2020g); and

- The “New industrial strategy for Europe” asserts that Europe “needs an industr[ial sector] that becomes greener and more digital” (European Commission 2020b p. 2)

Several of the documents make reference to the so-called “Destination Earth” initiative, the stated intention of which is to “develop a very high precision digital model of the Earth to monitor and simulate natural and human activity, and to develop and test scenarios that would enable more sustainable development and support European environmental policies” (European Commission 2020h). Destination Earth is designed to contribute both to the Commission’s Green Deal and to its Digital Strategy. It targets national authorities to aid policymakers and then opens up to users from academia and industry. The technical details of Destination Earth remain to be specified, but it is said to provide access to “data, advanced computing infrastructure, software, AI applications and analytics”. Therefore, while the exact role of AI tools within the initiative remains to be seen, the scale and ambition of Destination Earth and its role at the intersection of the “twin transitions” suggest it may be important in fostering the use AI to tackle climate change.

Given the prominence of digital technologies in everyday life and the increasing salience of the climate change challenge — as well as the coordination of policy priorities that accompanies a new administration — it is not surprising to find concordance among these documents. Nonetheless, the extent to which the Commission seems to see the twin transitions as developing hand-in-hand is striking. The EU’s renewed commitment to using AI and other digital technologies to make European society and industry greener and more sustainable is an important statement of intent and suggests that Europe may become a focal point of efforts to develop AI to combat climate change effectively. However, for all the opportunities these policy documents highlight, they brush over the challenges that need to be addressed to ensure a successful adoption of AI tools. To this end, in the next section, we offer 13 recommendations.

4. Recommendations for EU policymakers and the AI research community

The previous sections identified two areas where recommendations for leveraging the opportunities and addressing the challenges posed by AI in the context of climate change can be offered. Stated as objectives, these are, first, to harness the potential of AI for

combatting climate change in ways that are ethically sound; and second, to gauge and minimise the size of AI's carbon footprint. In this section, we address these two objectives in turn, in order to identify specific methods and areas of intervention for European policymakers and AI researchers. Our recommendations urge these stakeholders to *assess* existing capacities and potential opportunities, *incentivise* the creation of new infrastructures, and *develop* new approaches to enable society to maximise the potential of AI in the context of climate change, while minimising ethical and environmental drawbacks.

4.1 Harnessing the potential of AI for combatting climate change

By themselves, comprehensive surveys and conferences appear to be insufficient to gather, document, and analyse all the relevant evidence of AI being used to tackle climate change. More needs to be done to monitor and seek positive, climate-focused AI solutions from across sectors, domains, and regions of the world. This would involve deriving best practices and lessons learned from existing projects and identifying opportunities for future initiatives that may be missed without sufficient funding or support. Given the political and economic commitments it has already made, the EU would be an especially suitable sponsor and host of such an initiative. The EU is also in a leading position internationally to disseminate its findings to support action against climate change at a global scale.

Recommendation 1: *Incentivise* a world-leading initiative (Observatory) to document evidence of AI being used to combat climate change around the world, derive best practices and lessons learned, and disseminate the findings among researchers, policymakers and the public.

Another challenge concerns the ability to share the necessary resources for developing robust AI systems. This includes the best practices and lessons learned to be collected by the initiatives proposed in Recommendation 1 but, crucially, it also extends to data. The effectiveness of AI systems rests in large part on the size and quality of available datasets used to train these systems.

The recent European Strategy for Data notes that a current lack of data hinders the use of data for the public good, underlining the Commission's support for establishing

a new Common European Green Deal “data space [to] support ... the Green Deal priority actions on climate change” (European Commission 2020b). This may in turn require legislative and regulatory steps to facilitate business-to-business and business-to-government data sharing. The document also notes the need to “assess what measures are necessary to establish data pools for data analysis and machine learning”. The issue for the climate change data space is not simply to open the floodgates to data sharing, but also to ensure that the data that is shared is high-quality, accurate, and relevant to the problem at hand. In short, this is not just a question of collation but also of curation. The steps outlined so far to ensure that the specific requirements of climate change research are served by the data space are moving in this direction.

The European Strategy for Data argues that “data spaces should foster an ecosystem (of companies, civil society and individuals) creating new products and services based on more accessible data”. In the case of climate change, organisations (particularly in the private sector) may need further encouragement to develop AI-based solutions that are not “products and services” *per se*, but rather focused efforts to tackle climate-related issues, with or without a profit incentive, and potentially in partnership with public and non-profit groups. Therefore, the Commission could play a more front-footed role in stimulating these efforts or “challenges”, as it has already sought to do in the context of business-to-government data sharing for the public interest more generally (European Commission 2020f).

Recommendation 2: *Develop* standards of quality, accuracy, relevance and interoperability for data to be included in the forthcoming Common European Green Deal data space; identify aspects of climate action for which more data would be most beneficial; and explore, in consultation with domain experts and civil society organisations, how this data could be pooled in a common global climate data space.

Recommendation 3: *Incentivise* collaborations between data providers and technical experts in the private sector with domain experts from civil society, in the form of “challenges”, to ensure that the data in the Common European Green Deal data space is utilised effectively against climate change.

As section 3 makes clear, there has been considerable investment — both fiscal and political — to harness the twin ecological and digital transitions in order to create a more sustainable and prosperous EU. If done right, using AI in the fight against climate change is an ideal point of synthesis for these objectives. Therefore, to build on the previous recommendations, we also recommend that the European Commission earmarks a proportion of the recently announced Recovery Fund to support efforts to develop AI that tackles climate change in the ways identified through the proposals in Recommendation 1. Per the recent agreement between the Commission, the Parliament and European leaders, a considerable proportion (30%) of the Fund will be “dedicated to fighting climate change”, and it is separately stated that more than 50% of the overall fund will support modernisation related, to *inter alia*, “fair climate and digital transitions”. Thus, there is ample scope to invest a substantial proportion of this fund to leveraging AI-based responses to climate change, building on opportunities identified in Recommendations 1-3.

Recommendation 4: *Incentivise the development of sustainable, scalable responses to climate change that incorporate AI technology, drawing on earmarked Recovery Fund resources.*

It is important to ensure that all EU-funded and supported climate change research and innovation that uses AI follow steps to prevent bias and discrimination. This should take the form of protocols, auditing, and best practices tailored to this particular research context. In particular, large-scale initiatives such as the Destination Earth project ought to be designed with great care to prevent biases and discrepancies from arising in the so-called “digital twin” that will be created.

At the same time, transparency of purposes — clarifying for what an AI system is being optimised — may help to protect human autonomy. To this end, it may not be enough to make available information about how systems are optimised, but it may also be necessary to give affected stakeholders the opportunity to question, and even contest, the optimisation parameters that are set for a given system, depending on the context. Ensuring that these mechanisms of explanation and contestation are reliable and reproducible is likely to require access to the relevant data and initial conditions and parameter settings that were used for training algorithms.

Recommendation 5: *Develop* mechanisms for ethical auditing of AI systems deployed in high-stakes climate change contexts, where personal data may be used and/or individual behaviour may be affected. *Ensure* that clear, accessible statements regarding what metrics AI systems are optimised for, and why this is justified, are made available prior to the deployment of these systems. The possibility for affected stakeholders to question and contest system design and outcomes should also be guaranteed.

4.2 *Regulating AI's carbon footprint*

The field of AI research stakeholders, which includes (but is not limited to) researchers, laboratories, funding agencies, journal editors, conference organisers, and the managers of open-source ML libraries, can take several immediate steps to ensure that its carbon footprint is properly gauged and kept in check. At the same time, policymakers should play a critical role in ensuring that new reporting standards are set for organisations conducting large scale experiments and that the underlying infrastructure of AI remains environmentally sustainable while supporting innovative AI research in the EU. To this end, we offer recommendations to both stakeholders in both the research and policy domains.

Recommendations for AI research stakeholders

Steps have already been taken to tackle the reproducibility crisis mentioned in section 2.3. For example, within two years of encouraging paper submissions to include source code, NeurIPS reported the number of papers with code going from 50% to 75% of submissions (Gibney 2020). Additionally, standards and tools, like the Association for Computing Machinery's (ACM) (2020) artifact badging, NeurIPS's (2020) OpenReview, cKnowledge (Fursin 2020), PapersWithCode (2020), and MLPerf (Reddi et al. 2020), have been established in recent years to promote openness in scientific communication and ensure reproducibility. Similarly, systematic and accurate measurements to evaluate the energy consumption and carbon emissions of AI is needed for research activities. "Plug and play" tools need to be developed to facilitate the reporting of GHG emissions, and research conferences, journals and the community at large can play an important role in normalising the reporting of such data.

Open-source ML libraries, which are often established by private organisations, are essential to AI research. Adding information on the energy consumption, carbon emissions, and training conditions of various models — including hyperparameter sensitivity or algorithm performance against hardware — on these websites can help the field develop its environmental commitment.

Recommendation 6: *Develop* conference and journal checklists that include the disclosure of, *inter alia*, energy consumption, computational complexity, and experiments (e.g. number of training runs, and models produced) to align the field on common metrics (Gibney 2020; Schwartz et al. 2019; Henderson et al. 2020).

Recommendation 7: *Assess* the carbon footprint of AI models that appear on popular libraries and platforms, such as PyTorch, TensorFlow and Hugging Face, to inform users about their environmental costs.

These recommendations aim to normalise the disclosure of information pertaining to AI's carbon footprint as well as to help researchers and organisations select research tools based on environmental considerations. Online AI courses, ML libraries, journals and conferences can take actions to collect and display more information regarding the energy consumption and GHG emissions of AI. These recommendations aim to enable researchers to monitor and report systematically their AI projects' carbon footprints by using ready-made tools such as Lacoste et al.'s calculator (2019) or Henderson et al.'s framework (2020).

Increasing research on energy efficient computing and efficient AI is an important component to ensure that AI's carbon footprint is controlled in the long run. And the promotion of efficiency metrics and research may need to come from the field itself. For example, the Low-Power Image Recognition Challenge was created to define a common metric to compare image recognition results, accounting for energy efficiency and accuracy (Gauen et al. 2017). Similarly, Stanford University's DAWNbench benchmark was created in response to the field's hyper focus on accuracy metrics (Coleman, Kang, Narayanan, Nardi, Zhao, Zhang, Bailis, Olukotun, Ré, et al. 2019). The benchmark offers a “reference set of common deep learning workloads for quantifying training time, training cost, inference latency, different optimisation strategies, model architectures, software

frameworks, clouds, and hardware” (Coleman, Kang, Narayanan, Nardi, Zhao, Zhang, Bailis, Olukotun, Re, et al. 2019)

Recommendation 8: *Incentivise* the development of efficiency metrics for AI research and development (including model training) by promoting efficiency improvements and objectives in journals, conferences and challenges.

Note that key organisations such as the ACM, IEEE, NeurIPS and ICML, among others, would be instrumental in normalising efficiency metrics or publication requirements such as the one outlined in Recommendation 6. The normalisation of such metrics and requirements can bring more researchers to actively seek energy-efficient approaches to research and development, such as by changing their region of compute or cloud provider, or by opting to run intense calculations during times of excess electricity generation capacity (e.g. at night).

Recommendations for EU policymakers

Policymakers have an important role to play in equalising access to compute, developing efficient deep learning, and making compute-intensive AI research more accessible and affordable. For example, researchers in the US have suggested nationalising cloud infrastructure to provide more researchers with the ability to work without bearing exorbitant costs (Etchemendy and Li 2020). A European equivalent of the “National Research Cloud” could allow the EU to establish a long-term infrastructure that enables more European researchers to compete on a global scale, while also ensuring that research occurs on an efficient and sustainable European platform.²⁵

The reported and estimated decrease (by 30%) of EU-based data centres (EEA 2020) is largely due to efforts by EU states to increase the share of renewable energies in power generation (European Commission 2020a). CO₂ emissions of national power generation across EU member states has been decreasing, albeit emission rates differ

²⁵ Such a project would present clear synergies with the EU’s landmark cloud infrastructure project, Gaia-X, which seeks to develop a common data infrastructure in Europe (GAIA-X 2020).

significantly between different member states. For example, power generation in Estonia emits over 9 times more CO₂ than in Slovakia (EEA 2020).

Recommendation 9: *Develop* greener, smarter and cheaper data infrastructure (e.g., European research data centres) for researchers and universities across the EU.

Given the EU's increasing interest and investments in AI (Stix 2019), it is also important that the AI sector is considered specifically when formulating environmental policies. Both in research or production settings, AI requires increasingly specialised hardware and services that should be considered in any long-term environmental strategies.

Recommendation 10: *Assess* AI and its underlying infrastructure (e.g., data centres) when formulating energy management and carbon mitigation strategies to ensure that the European AI sector becomes sustainable as well as uniquely competitive.

Carbon labels and similar standards can benefit from receiving the endorsement of policymakers and even be required within the EU. Policymakers are key to ensuring that the field of AI research becomes more transparent when it comes to energy consumption and carbon emissions.

Recommendation 11: *Develop* carbon assessment and disclosure standards for AI to help the field align on metrics, increase research transparency, and communicate carbon footprints effectively via methods such as adding carbon labels to AI-based technologies and models listed in online libraries, journals, and leader boards.

These labels would allow researchers and developers to make environmentally informed decisions when choosing components (e.g. model, hardware and cloud provider) for their work. For example, borrowing directly from The Carbon Trust's (2020) "product carbon footprint labels", the following labels could be adapted to AI research and distributed in a similar fashion to ACM labels:

- Lower CO₂eq – indicating that the carbon footprint of a model/product is significantly lower carbon than the market dominant model/product in its category.
- CO₂eq Measured – indicating that the model/product footprint has been measured in accordance with an internationally recognised standard such as product standards: PAS2050, GHG Product Standard and ISO14067.
- Carbon Neutral – indicating that the model/product emissions are offset by the issuing organisation.

Policymakers are also key to ensuring that AI researchers in the EU are able to expand the field of AI research and diverge from conventional assumptions and research practices. Diverse funding will help European researchers to break from technological and normative trends that make it costly for researchers to try new ideas.

Recommendation 12: *Incentivise* diverse research agendas by funding and rewarding projects that diverge from the current trend of compute-intensive AI research in order to explore energy-efficient AI.

Example of potentially energy-efficient AI strategies include new hardware-software-algorithm combinations, algorithmic progress, symbolic AI, and hybrid symbolic-neural systems (Marcus 2020). Following this recommendation would enable the EU to enhance its potential for the development of AI research, by allowing a diverse pool of researchers from multiple countries to pursue a wide range of research agendas and compete with compute-intensive AI research coming from the US and China.

Recommendation 13: *Incentivise* energy-efficient and green research by making EU funding conditional on applicants measuring and reporting their estimated energy consumption and GHG emissions. Funding could fluctuate according to environmental efforts made (e.g. usage of efficient equipment, usage of renewable electricity, Power Usage Effectiveness of <1.5).

Conclusion

This report has described the beneficial impact that AI can have in the fight against climate change, the ethical challenges encountered in this process, and the computational intensity

that the development of AI requires, which introduces different challenges relating to energy consumption and GHG emissions. Benefits and risks are distinct yet intertwined. This is why we agree with Floridi and Nobre (2020) and see the use of AI to fight climate change as a leading example of

“a new marriage between the Green of our habitats – natural, synthetic and artificial, from the biosphere to the infosphere, from urban environments to economic, social, and political circumstances – and the Blue of our digital technologies, from mobile phones to social platforms, from the Internet of Things to Big Data, from AI to future quantum computing”.

In this marriage, some risks, such as AI’s carbon footprint, are not entirely avoidable, but they can certainly be minimised, to deliver the best strategies against climate change. This is why the right policies are key to harness the opportunities while ensuring that the risks are adequately assessed and minimised, as much as possible.

Harnessing the positive and mitigating the negative impact of AI on the environment is achievable with the support of robust policymaking and of key stakeholders. The formula of “Green & Blue” has never been more central to the European policymaking agenda, and the Recommendations outlined in this report can serve as a “Green & Blue-print” for a more sustainable society and a healthier biosphere. By shedding light on the use of AI to counter climate change and offering recommendations to make this use of AI ethically sound and sustainable, this report aims to inform EU policy strategy for the ‘twin transitions’ and help ensure that the marriage between the Green and the Blue is a success that leads to a better society and a healthier planet.

References

- Abdella, Galal M., Murat Kucukvar, Nuri Cihat Onat, Hussein M. Al-Yafay, and Muhammet Enis Bulak. 2020. 'Sustainability Assessment and Modeling Based on Supervised Machine Learning Techniques: The Case for Food Consumption'. *Journal of Cleaner Production* 251 (April): 119661. <https://doi.org/10.1016/j.jclepro.2019.119661>.
- Abrell, Jan, Mirjam Kosch, and Sebastian Rausch. 2019. 'How Effective Was the UK Carbon Tax? - A Machine Learning Approach to Policy Evaluation'. SSRN Scholarly Paper ID 3372388. Rochester, NY: Social Science Research Network. <https://doi.org/10.2139/ssrn.3372388>.
- ACM. 2020. 'Artifact Review and Badging - Current'. 2020. <https://www.acm.org/publications/policies/artifact-review-and-badging-current>.
- Aftab, Muhammad, Chien Chen, Chi-Kin Chau, and Talal Rahwan. 2017. 'Automatic HVAC Control with Real-Time Occupancy Recognition and Simulation-Guided Model Predictive Control in Low-Cost Embedded System'. *Energy and Buildings* 154: 141–56.
- Ahmed, Nur, and Muntasir Wahed. 2020. 'The De-Democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research'. *ArXiv:2010.15581 [Cs]*, October. <http://arxiv.org/abs/2010.15581>.
- Al-Jarrah, Omar Y., Paul D. Yoo, Sami Muhaidat, George K. Karagiannidis, and Kamal Taha. 2015. 'Efficient Machine Learning for Big Data: A Review'. *Big Data Research* 2 (3): 87–93. <https://doi.org/10.1016/j.bdr.2015.04.001>.
- Alom, Md Zahangir, Tarek M. Taha, Christopher Yakopcic, Stefan Westberg, Paheding Sidike, Mst Shamima Nasrin, Brian C. Van Esesn, Abdul A. S. Awwal, and Vijayan K. Asari. 2018. 'The History Began from AlexNet: A Comprehensive Survey on Deep Learning Approaches'. *ArXiv:1803.01164 [Cs]*, September. <http://arxiv.org/abs/1803.01164>.
- Amodei, Dario, and Danny Hernandez. 2018. 'AI and Compute'. OpenAI. 16 May 2018. <https://openai.com/blog/ai-and-compute/>.
- Andrae, Anders, and Tomas Edler. 2015. 'On Global Electricity Usage of Communication Technology: Trends to 2030'. *Challenges* 6 (1): 117–57. <https://doi.org/10.3390/challe6010117>.
- Anthony, Lasse F. Wolff, Benjamin Kanding, and Raghavendra Selvan. 2020. 'Carbontracker: Tracking and Predicting the Carbon Footprint of Training Deep Learning Models'. *ArXiv:2007.03051 [Cs, Eess, Stat]*, July. <http://arxiv.org/abs/2007.03051>.
- Avgerinou, Maria, Paolo Bertoldi, and Luca Castellazzi. 2017. 'Trends in Data Centre Energy Consumption under the European Code of Conduct for Data Centre Energy Efficiency'. *Energies* 10 (10): 1470. <https://doi.org/10.3390/en10101470>.
- Barnes, Elizabeth A., James W. Hurrell, Imme Ebert-Uphoff, Chuck Anderson, and David Anderson. 2019. 'Viewing Forced Climate Patterns Through an AI Lens'. *Geophysical Research Letters* 46 (22): 13389–98. <https://doi.org/10.1029/2019GL084944>.
- Belkhir, Lotfi, and Ahmed Elmeligi. 2018. 'Assessing ICT Global Emissions Footprint: Trends to 2040 & Recommendations'. *Journal of Cleaner Production* 177 (March): 448–63. <https://doi.org/10.1016/j.jclepro.2017.12.239>.

- Bender, E. M., T. Gebru, and A. McMillan-Major. 2021. ‘On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?’. *Proceedings of FAccT*.
- Berner, Christopher, Brooke Chan, Vicki Cheung, Przemyslaw Dębiak, Christy Dennison, David Farhi, Quirin Fischer, et al. 2019. ‘Dota 2 with Large Scale Deep Reinforcement Learning’. *ArXiv:1912.06680 [Cs, Stat]*, December. <http://arxiv.org/abs/1912.06680>.
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. ‘Language Models Are Few-Shot Learners’. *ArXiv:2005.14165 [Cs]*, July. <http://arxiv.org/abs/2005.14165>.
- Buckman, Jacob, Danijar Hafner, George Tucker, Eugene Brevdo, and Honglak Lee. 2018. ‘Sample-Efficient Reinforcement Learning with Stochastic Ensemble Value Expansion’. In *Advances in Neural Information Processing Systems 31*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, 8224–34. Curran Associates, Inc. <http://papers.nips.cc/paper/8044-sample-efficient-reinforcement-learning-with-stochastic-ensemble-value-expansion.pdf>.
- C2E2. 2018. ‘Greenhouse Gas Emissions in the ICT Sector’. 2018. <https://c2e2.unepdtu.org/collection/c2e2-publications/>.
- Cai, Han, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. 2020. ‘Once-for-All: Train One Network and Specialize It for Efficient Deployment’. *ArXiv:1908.09791 [Cs, Stat]*, April. <http://arxiv.org/abs/1908.09791>.
- Chattopadhyay, Ashesh, Pedram Hassanzadeh, and Saba Pasha. 2020. ‘Predicting Clustered Weather Patterns: A Test Case for Applications of Convolutional Neural Networks to Spatio-Temporal Climate Data’. *Scientific Reports* 10 (1): 1317. <https://doi.org/10.1038/s41598-020-57897-9>.
- Chen, Chun-Fu, Quanfu Fan, Neil Mallinar, Tom Sercu, and Rogerio Feris. 2019. ‘Big-Little Net: An Efficient Multi-Scale Feature Representation for Visual and Speech Recognition’. *ArXiv:1807.03848 [Cs]*, July. <http://arxiv.org/abs/1807.03848>.
- Chen, Tianqi, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. ‘MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems’. *ArXiv:1512.01274 [Cs]*, December. <http://arxiv.org/abs/1512.01274>.
- Chen, Yu-Hsin, Joel Emer, and Vivienne Sze. 2017. ‘Using Dataflow to Optimize Energy Efficiency of Deep Neural Network Accelerators’. *IEEE Micro* 37 (3): 12–21. <https://doi.org/10.1109/MM.2017.54>.
- Choi, Changhyun, Jeonghwan Kim, Jongsung Kim, Donghyun Kim, Younghye Bae, and Hung Soo Kim. 2018. ‘Development of Heavy Rain Damage Prediction Model Using Machine Learning Based on Big Data’. Research Article. *Advances in Meteorology*. Hindawi. 13 June 2018. <https://doi.org/10.1155/2018/5024930>.
- Cifuentes, Jenny, Geovanny Marulanda, Antonio Bello, and Javier Reneses. 2020. ‘Air Temperature Forecasting Using Machine Learning Techniques: A Review’. *Energies* 13 (16): 4215.
- Coeckelbergh, Mark. 2020. ‘AI for Climate: Freedom, Justice, and Other Ethical and Political Challenges’. *AI and Ethics*, October. <https://doi.org/10.1007/s43681-020-00007-2>.
- Coleman, Cody, Daniel Kang, Deepak Narayanan, Luigi Nardi, Tian Zhao, Jian Zhang, Peter Bailis, Kunle Olukotun, Chris Ré, and Matei Zaharia. 2019. ‘Analysis of DAWNbench, a Time-to-Accuracy Machine Learning Performance Benchmark’.

- Coleman, Cody, Daniel Kang, Deepak Narayanan, Luigi Nardi, Tian Zhao, Jian Zhang, Peter Bailis, Kunle Olukotun, Chris Re, and Matei Zaharia. 2019. ‘Analysis of DAWNbench, a Time-to-Accuracy Machine Learning Performance Benchmark’. *ArXiv:1806.01427 [Cs, Stat]*, December. <http://arxiv.org/abs/1806.01427>.
- Cook, Gary, and Elizabeth Jardim. 2019. ‘Clicking Clean Virginia’. *Greenpeace USA* (blog). 2019. <https://www.greenpeace.org/usa/reports/click-clean-virginia/>.
- Cowls, Josh, Andreas Tsamados, Mariarosaria Taddeo, and Luciano Floridi. Forthcoming. ‘A Definition, Benchmark and Database of AI for Social Good Initiatives’. *Nature Machine Intelligence*.
- . Forthcoming. ‘Artificial Intelligence for Social Good: A Definition, a Benchmark, and a Database’.
- Crawford, Kate, and Vladan Joler. 2018. ‘Anatomy of an AI System’. *Anatomy of an AI System*. 2018. <http://www.anatomyof.ai>.
- Dabiri, Sina, and Kevin Heaslip. 2018. ‘Inferring Transportation Modes from GPS Trajectories Using a Convolutional Neural Network’. *Transportation Research Part C: Emerging Technologies* 86 (January): 360–71. <https://doi.org/10.1016/j.trc.2017.11.021>.
- Dahlquist, Erik. 2020. ‘The FUDIPO Project: AI Systems in Process Industries’. <https://cordis.europa.eu/article/id/415798-using-ai-to-improve-energy-and-resource-efficiency-in-various-industries>.
- Dhar, Payal. 2020. ‘The Carbon Impact of Artificial Intelligence’. *Nature Machine Intelligence* 2 (8): 423–25. <https://doi.org/10.1038/s42256-020-0219-9>.
- Di Piazza, A., M. C. Di Piazza, G. La Tona, and M. Luna. 2020. ‘An Artificial Neural Network-Based Forecasting Model of Energy-Related Time Series for Electrical Grid Management’. *Mathematics and Computers in Simulation*, May. <https://doi.org/10.1016/j.matcom.2020.05.010>.
- Dobbe, Roel, Oscar Sondermeijer, David Fridovich-Keil, Daniel Arnold, Duncan Callaway, and Claire Tomlin. 2019. ‘Toward Distributed Energy Services: Decentralizing Optimal Power Flow With Machine Learning’. *IEEE Transactions on Smart Grid* 11 (2): 1296–1306.
- Dodge, Jesse, Suchin Gururangan, Dallas Card, Roy Schwartz, and Noah A. Smith. 2019. ‘Show Your Work: Improved Reporting of Experimental Results’. <https://arxiv.org/abs/1909.03004v1>.
- EDL. 2017. ‘Efficient Deep Learning’. 2017. <https://efficientdeeplearning.nl/>.
- EEA. 2020. ‘CO2 Intensity of Electricity Generation. Data Table. European Environmental Agency (EEA)’. <https://www.eea.europa.eu/data-and-maps/data/co2-intensity-of-electricity-generation>.
- EFF. 2017. ‘AI Progress Measurement’. Electronic Frontier Foundation. 2017. <https://www.eff.org/ai/metrics>.
- Etchemendy, John, and Fei Fei Li. 2020. ‘National Research Cloud: Ensuring the Continuation of American Innovation’. Stanford HAI. 2020. <https://hai.stanford.edu/blog/national-research-cloud-ensuring-continuation-american-innovation>.
- European Commission. 2019. ‘The European Green Deal’. 2019. <https://eur-lex.europa.eu/legal->

- content/EN/TXT/?qid=1576150542719&uri=COM%3A2019%3A640%3AFIN.
- . 2020a. ‘2030 Climate & Energy Framework’. Text. Climate Action - European Commission. 2020. https://ec.europa.eu/clima/policies/strategies/2030_en.
- . 2020b. ‘A European Strategy for Data’. Text. Shaping Europe’s Digital Future - European Commission. 2020. <https://ec.europa.eu/digital-single-market/en/policies/building-european-data-economy>.
- . 2020c. ‘A New Industrial Strategy for a Green and Digital Europe’. https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip_20_416/IP_20_416_EN.pdf.
- . 2020d. ‘Energy-Efficient Cloud Computing Technologies and Policies for an Eco-Friendly Cloud Market’. Text. Shaping Europe’s Digital Future - European Commission. 2020. <https://ec.europa.eu/digital-single-market/en/news/energy-efficient-cloud-computing-technologies-and-policies-eco-friendly-cloud-market>.
- . 2020e. ‘Shaping Europe’s Digital Future’. Text. European Commission. 2020. https://ec.europa.eu/info/strategy/priorities-2019-2024/europe-fit-digital-age/shaping-europe-digital-future_en.
- . 2020f. ‘Towards a European Strategy on Business-to-Government Data Sharing for the Public Interest’.
- . 2020g. ‘White Paper on Artificial Intelligence: A European Approach to Excellence and Trust’. Text. European Commission - European Commission. 2020. https://ec.europa.eu/info/publications/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en.
- . 2020h. ‘Destination Earth (DestinE)’. Text. Shaping Europe’s Digital Future - European Commission. 17 August 2020. <https://ec.europa.eu/digital-single-market/en/destination-earth-destine>.
- . 2020i. ‘Recovery Plan for Europe’. Text. European Commission - European Commission. 10 November 2020. https://ec.europa.eu/info/strategy/recovery-plan-europe_en.
- Evans, Richard, and Jim Gao. 2016. ‘DeepMind AI Reduces Google Data Centre Cooling Bill by 40%’. Deepmind. 2016. [/blog/article/deepmind-ai-reduces-google-data-centre-cooling-bill-40](https://blog.google/ai/deepmind-ai-reduces-google-data-centre-cooling-bill-40).
- Fathi, Soheil, Ravi Srinivasan, Andriel Fenner, and Sahand Fathi. 2020. ‘Machine Learning Applications in Urban Building Energy Performance Forecasting: A Systematic Review’. *Renewable and Sustainable Energy Reviews* 133 (November): 110287. <https://doi.org/10.1016/j.rser.2020.110287>.
- Fedus, William, Barret Zoph, and Noam Shazeer. 2021. ‘Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity’. *ArXiv:2101.03961 [Cs]*, January. <http://arxiv.org/abs/2101.03961>.
- Floridi, Luciano. 2016. ‘Tolerant Paternalism: Pro-Ethical Design as a Resolution of the Dilemma of Toleration’. *Science and Engineering Ethics* 22 (6): 1669–88. <https://doi.org/10.1007/s11948-015-9733-2>.
- . 2017. ‘Group Privacy: A Defence and an Interpretation’. In *Group Privacy*, 83–100. Springer.
- . 2020. ‘The Green and the Blue: A New Political Ontology for a Mature Information Society’. *Philosophisches Jahrbuch* 2: 307–38.

- Floridi, Luciano, and Josh Cowls. 2019. 'A Unified Framework of Five Principles for AI in Society'. *Harvard Data Science Review* 1 (1). <https://doi.org/10.1162/99608f92.8cd550d1>.
- Floridi, Luciano, and Kia Nobre. 2020. 'The Green and the Blue: How AI May Be a Force for Good'. OECD. <https://www.oecd-forum.org/posts/the-green-and-the-blue-how-ai-may-be-a-force-for-good>.
- Fukushima, Kunihiko, and Sei Miyake. 1982. 'Neocognitron: A New Algorithm for Pattern Recognition Tolerant of Deformations and Shifts in Position'. *Pattern Recognition* 15 (6): 455–69. [https://doi.org/10.1016/0031-3203\(82\)90024-3](https://doi.org/10.1016/0031-3203(82)90024-3).
- Fursin, Grigori. 2020. 'Enabling Reproducible ML and Systems Research: The Good, the Bad, and the Ugly', August. <https://doi.org/10.5281/ZENODO.4005773>.
- Gagne, David John, Hannah M. Christensen, Aneesh C. Subramanian, and Adam H. Monahan. 2020. 'Machine Learning for Stochastic Parameterization: Generative Adversarial Networks in the Lorenz '96 Model'. *Journal of Advances in Modeling Earth Systems* 12 (3): e2019MS001896. <https://doi.org/10.1029/2019MS001896>.
- GAIA-X. 2020. 'GAIA-X: A Federated Data Infrastructure for Europe'. <https://www.data-infrastructure.eu/GAIA-X/Navigation/EN/Home/home.html>.
- García-Martín, Eva, Crefeda Faviola Rodrigues, Graham Riley, and Håkan Grahn. 2019. 'Estimation of Energy Consumption in Machine Learning'. *Journal of Parallel and Distributed Computing* 134 (December): 75–88. <https://doi.org/10.1016/j.jpdc.2019.07.007>.
- Gauen, K., R. Rangan, A. Mohan, Y. Lu, W. Liu, and A. C. Berg. 2017. 'Low-Power Image Recognition Challenge'. In *2017 22nd Asia and South Pacific Design Automation Conference (ASP-DAC)*, 99–104. <https://doi.org/10.1109/ASPDAC.2017.7858303>.
- Gibney, Elizabeth. 2020. 'This AI Researcher Is Trying to Ward off a Reproducibility Crisis'. *Nature* 577 (7788): 14–14. <https://doi.org/10.1038/d41586-019-03895-5>.
- GLUE. 2020. 'GLUE Benchmark'. 2020. <https://gluebenchmark.com/>.
- Gundersen, Odd Erik, and Sigbjørn Kjensmo. 2018. *State of the Art: Reproducibility in Artificial Intelligence*.
- Guss, William H., Cayden Codel, Katja Hofmann, Brandon Houghton, Noboru Kuno, Stephanie Milani, Sharada Mohanty, et al. 2019. 'The MineRL Competition on Sample Efficient Reinforcement Learning Using Human Priors'. *ArXiv:1904.10079 [Cs, Stat]*, July. <http://arxiv.org/abs/1904.10079>.
- Hall, Dame Wendy, and Jerome Pesenti. 2017. 'GROWING THE ARTIFICIAL INTELLIGENCE INDUSTRY IN THE UK'. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf.
- Ham, Yoo-Geun, Jeong-Hwan Kim, and Jing-Jia Luo. 2019. 'Deep Learning for Multi-Year ENSO Forecasts'. *Nature* 573 (7775): 568–72. <https://doi.org/10.1038/s41586-019-1559-7>.
- Hartley, Matthew, and Tjelvar S.G. Olsson. 2020. 'DtoolAI: Reproducibility for Deep Learning'. *Patterns* 1 (5): 100073. <https://doi.org/10.1016/j.patter.2020.100073>.
- Henderson, Peter, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. 2020. 'TOWARDS THE SYSTEMATIC REPORTING OF THE ENERGY AND CARBON FOOTPRINTS OF MACHINE LEARNING', 25.

- Hernandez, Danny, and Tom Brown. 2020. 'AI and Efficiency'. OpenAI. 5 May 2020. <https://openai.com/blog/ai-and-efficiency/>.
- Hill, Mark D, and Michael R Marty. 2008. 'Amdahl's Law in the Multicore Era', 6.
- Hintemann, Ralph. 2015. 'Consolidation, Colocation, Virtualization, and Cloud Computing: The Impact of the Changing Structure of Data Centers on Total Electricity Demand'. In *ICT Innovations for Sustainability*, edited by Lorenz M. Hilty and Bernard Aebischer, 125–36. Advances in Intelligent Systems and Computing. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-09228-7_7.
- Hintemann, Ralph, and Simon Hinterholzer. 2020. 'Data Centres In Europe - Opportunities For Sustainable Digitalisation', 36.
- Hooker, Sara. 2020. 'The Hardware Lottery'. *ArXiv:2009.06489 [Cs]*, September. <http://arxiv.org/abs/2009.06489>.
- Huntingford, Chris, Elizabeth S Jeffers, Michael B Bonsall, Hannah M Christensen, Thomas Lees, and Hui Yang. 2019. 'Machine Learning and Artificial Intelligence to Aid Climate Change Research and Preparedness'. *Environmental Research Letters* 14 (12): 124007. <https://doi.org/10.1088/1748-9326/ab4e55>.
- Hutson, Matthew. 2018. 'Artificial Intelligence Faces Reproducibility Crisis'. *Science* 359 (6377): 725–26. <https://doi.org/10.1126/science.359.6377.725>.
- IEA. 2020. 'Data Centres and Data Transmission Networks'. <https://www.iea.org/reports/data-centres-and-data-transmission-networks>.
- Inderwildi, Oliver, Chuan Zhang, Xiaonan Wang, and Markus Kraft. 2020. 'The Impact of Intelligent Cyber-Physical Systems on the Decarbonization of Energy'. *Energy & Environmental Science* 13 (3): 744–71. <https://doi.org/10.1039/C9EE01919G>.
- Ise, Takeshi, and Yurika Oba. 2019. 'Forecasting Climatic Trends Using Neural Networks: An Experimental Study Using Global Historical Data'. *Frontiers in Robotics and AI* 6. <https://doi.org/10.3389/frobt.2019.00032>.
- Jaafari, Abolfazl, Eric K. Zenner, Mahdi Panahi, and Himan Shahabi. 2019. 'Hybrid Artificial Intelligence Models Based on a Neuro-Fuzzy System and Metaheuristic Optimization Algorithms for Spatial Prediction of Wildfire Probability'. *Agricultural and Forest Meteorology* 266–267 (March): 198–207. <https://doi.org/10.1016/j.agrformet.2018.12.015>.
- Jiang, Angela H., Daniel L.-K. Wong, Giulio Zhou, David G. Andersen, Jeffrey Dean, Gregory R. Ganger, Gauri Joshi, et al. 2019. 'Accelerating Deep Learning by Focusing on the Biggest Losers', October. <https://arxiv.org/abs/1910.00762v1>.
- Jones, Nicola. 2018. 'How to Stop Data Centres from Gobbling up the World's Electricity'. *Nature* 561 (7722): 163–66. <https://doi.org/10.1038/d41586-018-06610-y>.
- Jouhara, Hussam, and Richard Meskimmon. 2014. 'Heat Pipe Based Thermal Management Systems for Energy-Efficient Data Centres'. *Energy* 77 (December): 265–70. <https://doi.org/10.1016/j.energy.2014.08.085>.
- Lacoste, Alexandre, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. 'Quantifying the Carbon Emissions of Machine Learning'. *ArXiv:1910.09700 [Cs]*, November. <http://arxiv.org/abs/1910.09700>.
- Larraondo, Pablo R., Luigi J. Renzullo, Albert I. J. M. Van Dijk, Inaki Inza, and Jose A. Lozano. 2020. 'Optimization of Deep Learning Precipitation Models Using Categorical Binary Metrics'. *Journal of Advances in Modeling Earth Systems* 12 (5): e2019MS001909. <https://doi.org/10.1029/2019MS001909>.

- Li, Chuan. 2020. 'OpenAI's GPT-3 Language Model: A Technical Overview'. Lambda Blog. 3 June 2020. <https://lambdalabs.com/blog/demystifying-gpt-3/>.
- Liakos, Konstantinos G., Patrizia Busato, Dimitrios Moshou, Simon Pearson, and Dionysis Bochtis. 2018. 'Machine Learning in Agriculture: A Review'. *Sensors* 18 (8): 2674.
- Lohr, Steve. 2019. 'At Tech's Leading Edge, Worry About a Concentration of Power - The New York Times'. 2019. <https://www.nytimes.com/2019/09/26/technology/ai-computer-expense.html>.
- Lu, Hongfang, Xin Ma, Kun Huang, and Mohammadamin Azimi. 2020. 'Carbon Trading Volume and Price Forecasting in China Using Multiple Machine Learning Models'. *Journal of Cleaner Production* 249 (March): 119386. <https://doi.org/10.1016/j.jclepro.2019.119386>.
- Malmodin, Jens, and Dag Lundén. 2018. 'The Energy and Carbon Footprint of the Global ICT and E&M Sectors 2010–2015'. *Sustainability* 10 (9): 3027. <https://doi.org/10.3390/su10093027>.
- Marcus, Gary. 2020. 'The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence'. *ArXiv:2002.06177 [Cs]*, February. <http://arxiv.org/abs/2002.06177>.
- Mardani, Abbas, Huchang Liao, Mehrbakhsh Nilashi, Melfi Alrasheedi, and Fausto Cavallaro. 2020. 'A Multi-Stage Method to Predict Carbon Dioxide Emissions Using Dimensionality Reduction, Clustering, and Machine Learning Techniques'. *Journal of Cleaner Production* 275 (December): 122942. <https://doi.org/10.1016/j.jclepro.2020.122942>.
- Masanet, Eric, Arman Shehabi, Nuo Lei, Sarah Smith, and Jonathan Koomey. 2020. 'Recalibrating Global Data Center Energy-Use Estimates'. *Science* 367 (6481): 984–86. <https://doi.org/10.1126/science.aba3758>.
- Masson-Delmotte, Valérie, Panmao Zhai, Hans-Otto Pörtner, Debra Roberts, J. Skea, P. R. Shukla, Anna Pirani, W. Moufouma-Okia, C. Péan, and R. Pidcock. 2018. *Global Warming of 1.5 OC: An IPCC Special Report on the Impacts of Global Warming of 1.5° C Above Pre-Industrial Levels and Related Global Greenhouse Gas Emission Pathways, in the Context of Strengthening the Global Response to the Threat of Climate Change, Sustainable Development, and Efforts to Eradicate Poverty*. World Meteorological Organization Geneva, Switzerland.
- Matthews, H. Scott, Chris T. Hendrickson, and Christopher L. Weber. 2008. 'The Importance of Carbon Footprint Estimation Boundaries'. *Environmental Science & Technology* 42 (16): 5839–42. <https://doi.org/10.1021/es703112w>.
- McCarthy, John, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon. 2006. 'A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955'. *AI Magazine* 27 (4): 12–12.
- Menad, Nait Amar, Abdolhossein Hemmati-Sarapardeh, Amir Varamesh, and Shahaboddin Shamsheerband. 2019. 'Predicting Solubility of CO₂ in Brine by Advanced Machine Learning Systems: Application to Carbon Capture and Sequestration'. *Journal of CO₂ Utilization* 33 (October): 83–95. <https://doi.org/10.1016/j.jcou.2019.05.009>.
- Microsoft. 2018. 'The Carbon Benefits of Cloud Computing: A Study on the Microsoft Cloud'.
- . 2019. 'Machine Reading Systems Are Becoming More Conversational'. *Microsoft Research* (blog). 2019. <https://www.microsoft.com/en->

- us/research/blog/machine-reading-systems-are-becoming-more-conversational/.
- Microsoft, and C. 2018. ‘How AI Can Enable a Sustainable Future’. *Microsoft in Association with PwC*.
- Mytton, David. 2020. ‘Hiding Greenhouse Gas Emissions in the Cloud’. *Nature Climate Change* 10 (8): 701–701. <https://doi.org/10.1038/s41558-020-0837-6>.
- Narciso, Diogo A. C., and F. G. Martins. 2020. ‘Application of Machine Learning Tools for Energy Efficiency in Industry: A Review’. *Energy Reports* 6 (November): 1181–99. <https://doi.org/10.1016/j.egy.2020.04.035>.
- Nature Electronics. 2018. ‘Does AI Have a Hardware Problem?’ *Nature Electronics* 1 (4): 205–205. <https://doi.org/10.1038/s41928-018-0068-2>.
- NeurIPS. 2020. ‘OpenReview’. OpenReview. 2020. <https://openreview.net/>.
- NVIDIA. 2018. ‘TESLA V100 Performance Guide’.
- Open Compute Project. 2020. ‘Open Compute Project’. Open Compute Project. 2020. <https://www.opencompute.org>.
- Pachauri, Rajendra K., Myles R. Allen, Vicente R. Barros, John Broome, Wolfgang Cramer, Renate Christ, John A. Church, Leon Clarke, Qin Dahe, and Purnamita Dasgupta. 2014. *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Ipcc.
- PapersWithCode. 2020. ‘Papers with Code - Browse the State-of-the-Art in Machine Learning’. 2020. <https://paperswithcode.com/sota>.
- Perera, Lokukaluge P., B. Mo, and Guedes Soares. 2016. ‘Machine Intelligence for Energy Efficient Ships: A Big Data Solution’. *Maritime Engineering and Technology III, Guedes Soares & Santos (Eds.)* 1: 143–50.
- Perrault, Raymond, Shoham Yoav, Erik Brynjolfsson, Clark Jack, John Etchmendy, Barbara Grosz, Lyons Terah, Manyika James, Mishra Saurabh, and Niebles Juan Carlos. 2019. *Artificial Intelligence Index Report 2019*.
- Pham, Hieu, Melody Y. Guan, Barret Zoph, Quoc V. Le, and Jeff Dean. 2018. ‘Efficient Neural Architecture Search via Parameter Sharing’. *ArXiv:1802.03268 [Cs, Stat]*, February. <http://arxiv.org/abs/1802.03268>.
- Radford, Alec, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2018. ‘Language Models Are Unsupervised Multitask Learners’, 24.
- Rangan, Majumder, and Wang Junhua. 2020. ‘DeepSpeed & ZeRO-2: Shattering Barriers of Deep Learning Speed & Scale’. *Microsoft Research* (blog). 19 May 2020. <https://www.microsoft.com/en-us/research/blog/zero-2-deepspeed-shattering-barriers-of-deep-learning-speed-scale/>.
- Rasp, Stephan, Michael S. Pritchard, and Pierre Gentine. 2018. ‘Deep Learning to Represent Subgrid Processes in Climate Models’. *Proceedings of the National Academy of Sciences* 115 (39): 9684–89. <https://doi.org/10.1073/pnas.1810286115>.
- Reddi, Vijay Janapa, Christine Cheng, David Kanter, Peter Mattson, Guenther Schmuelling, Carole-Jean Wu, Brian Anderson, et al. 2020. ‘MLPerf Inference Benchmark’. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, 446–59. Valencia, Spain: IEEE. <https://doi.org/10.1109/ISCA45697.2020.00045>.
- Ridwan, Wanie M., Michelle Sapitang, Awatif Aziz, Khairul Faizal Kushiari, Ali Najah Ahmed, and Ahmed El-Shafie. 2020. ‘Rainfall Forecasting Model Using Machine

- Learning Methods: Case Study Terengganu, Malaysia'. *Ain Shams Engineering Journal*, November. <https://doi.org/10.1016/j.asej.2020.09.011>.
- Robinson, Caleb, and Bistra Dilkina. 2018. 'A Machine Learning Approach to Modeling Human Migration'. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, 1–8. COMPASS '18. Menlo Park and San Jose, CA, USA: Association for Computing Machinery. <https://doi.org/10.1145/3209811.3209868>.
- Rolnick, David, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, et al. 2019. 'Tackling Climate Change with Machine Learning'. *ArXiv:1906.05433 [Cs, Stat]*, November. <http://arxiv.org/abs/1906.05433>.
- Russell, Stephen. 2019. 'Estimating and Reporting the Comparative Emissions Impacts of Products', December. <https://www.wri.org/publication/estimating-and-reporting-comparative-emissions-impacts-products>.
- Sayed-Mouchaweh, Moamar, ed. 2020. *Artificial Intelligence Techniques for a Scalable Energy Transition: Advanced Methods, Digital Technologies, Decision Support Tools, and Applications*. Springer International Publishing. <https://doi.org/10.1007/978-3-030-42726-9>.
- Schmidt, Andreas T., and Bart Engelen. 2020. 'The Ethics of Nudging: An Overview'. *Philosophy Compass* 15 (4): e12658. <https://doi.org/10.1111/phc3.12658>.
- Schwartz, Roy, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2019. 'Green AI'. *ArXiv:1907.10597 [Cs, Stat]*, August. <http://arxiv.org/abs/1907.10597>.
- Shehabi, Arman, Sarah J. Smith, Eric Masanet, and Jonathan Koomey. 2018. 'Data Center Growth in the United States: Decoupling the Demand for Services from Electricity Use'. *Environmental Research Letters* 13 (12): 124030. <https://doi.org/10.1088/1748-9326/aaec9c>.
- Shrestha, Manish, Sujal Manandhar, and Sangam Shrestha. 2020. 'Forecasting Water Demand under Climate Change Using Artificial Neural Network: A Case Study of Kathmandu Valley, Nepal'. *Water Supply* 20 (5): 1823–33. <https://doi.org/10.2166/ws.2020.090>.
- Sønderby, Casper Kaae, Lasse Espeholt, Jonathan Heek, Mostafa Dehghani, Avital Oliver, Tim Salimans, Shreya Agrawal, Jason Hickey, and Nal Kalchbrenner. 2020. 'MetNet: A Neural Weather Model for Precipitation Forecasting'. *ArXiv:2003.12140 [Physics, Stat]*, March. <http://arxiv.org/abs/2003.12140>.
- SQuAD. 2020. 'The Stanford Question Answering Dataset'. 2020. <https://rajpurkar.github.io/SQuAD-explorer/>.
- Stix, Charlotte. 2019. 'A Survey of the European Union's Artificial Intelligence Ecosystem'. http://lcfi.ac.uk/media/uploads/files/Stix_Europe_AI_Final.pdf.
- Strubell, Emma, Ananya Ganesh, and Andrew McCallum. 2019. 'Energy and Policy Considerations for Deep Learning in NLP'. *ArXiv:1906.02243 [Cs]*, June. <http://arxiv.org/abs/1906.02243>.
- SuperGLUE. 2020. 'SuperGLUE Benchmark'. SuperGLUE Benchmark. 2020. <https://super.gluebenchmark.com/>.
- Sutton, Richard. 2019. 'The Bitter Lesson'. 2019. <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>.
- Sze, Vivienne, Yu-Hsin Chen, Tien-Ju Yang, and Joel Emer. 2017. 'Efficient Processing of Deep Neural Networks: A Tutorial and Survey'. *ArXiv:1703.09039 [Cs]*, August. <http://arxiv.org/abs/1703.09039>.

- Taddeo, Mariarosaria, and Luciano Floridi. 2018. 'How AI Can Be a Force for Good'. *Science* 361 (6404): 751–52.
- Tao, Ye, Miaohua Huang, and Lan Yang. 2018. 'Data-Driven Optimized Layout of Battery Electric Vehicle Charging Infrastructure'. *Energy* 150: 735–44.
- The Carbon Trust. 2020. 'Product Carbon Footprint Label'. <https://www.carbontrust.com/what-we-do/assurance-and-certification/product-carbon-footprint-label>. 2020.
- Theis, Thomas N., and H.-S. Philip Wong. 2017. 'The End of Moore's Law: A New Beginning for Information Technology'. *Computing in Science & Engineering* 19 (2): 41–50. <https://doi.org/10.1109/MCSE.2017.29>.
- Thilakarathna, P. S. M., S. Seo, K. S. Kristombu Baduge, H. Lee, P. Mendis, and G. Foliente. 2020. 'Embodied Carbon Analysis and Benchmarking Emissions of High and Ultra-High Strength Concrete Using Machine Learning Algorithms'. *Journal of Cleaner Production* 262 (July): 121281. <https://doi.org/10.1016/j.jclepro.2020.121281>.
- Thompson, Neil C., Kristjan Greenewald, Keeheon Lee, and Gabriel F. Manso. 2020. 'The Computational Limits of Deep Learning'. *ArXiv:2007.05558 [Cs, Stat]*, July. <http://arxiv.org/abs/2007.05558>.
- Thompson, Neil, and Svenja Spanuth. 2018. 'The Decline of Computers As a General Purpose Technology: Why Deep Learning and the End of Moore's Law Are Fragmenting Computing'. SSRN Scholarly Paper ID 3287769. Rochester, NY: Social Science Research Network. <https://doi.org/10.2139/ssrn.3287769>.
- Tsamados, Andreas, Nikita Aggarwal, Josh Cows, Jessica Morley, Huw Roberts, Mariarosaria Taddeo, and Luciano Floridi. 2020. 'The Ethics of Algorithms: Key Problems and Solutions'. *Available at SSRN 3662302*.
- US EPA. 2016. 'Greenhouse Gas Emissions from a Typical Passenger Vehicle'. Overviews and Factsheets. US EPA. 12 January 2016. <https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle>.
- Vodafone Institute for Society and Communications. 2020. 'Digitising Europe Pulse — Tackling Climate Change: A Survey of 13 EU Countries'.
- VTAB. 2020. 'Visual Task Adaptation Benchmark'. Visual Task Adaptation Benchmark. 2020. https://google-research.github.io/task_adaptation/.
- Wei, Sun, Wang Yuwei, and Zhang Chongchong. 2018. 'Forecasting CO2 Emissions in Hebei, China, through Moth-Flame Optimization Based on the Random Forest and Extreme Learning Machine'. *Environmental Science and Pollution Research* 25 (29): 28985–97. <https://doi.org/10.1007/s11356-018-2738-z>.
- Werbos, Paul J. 1988. 'Generalization of Backpropagation with Application to a Recurrent Gas Market Model'. *Neural Networks* 1 (4): 339–56. [https://doi.org/10.1016/0893-6080\(88\)90007-X](https://doi.org/10.1016/0893-6080(88)90007-X).
- Wheeldon, Adrian, Rishad Shafik, Tousif Rahman, Jie Lei, Alex Yakovlev, and Ole-Christoffer Granmo. 2020. 'Learning Automata Based Energy-Efficient AI Hardware Design for IoT Applications'. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 378 (2182): 20190593. <https://doi.org/10.1098/rsta.2019.0593>.

- Xenochristou, Maria, Chris Hutton, Jan Hofman, and Zoran Kapelan. 2020. 'Water Demand Forecasting Accuracy and Influencing Factors at Different Spatial Scales Using a Gradient Boosting Machine'. *Water Resources Research* 56 (8): e2019WR026304.
- Yang, Guang-Zhong, Jim Bellingham, Pierre E. Dupont, Peer Fischer, Luciano Floridi, Robert Full, Neil Jacobstein, et al. 2018. 'The Grand Challenges of Science Robotics'. *Science Robotics* 3 (14). <https://doi.org/10.1126/scirobotics.aar7650>.
- Zheng, Gang, Xiaofeng Li, Rong-Hua Zhang, and Bin Liu. 2020. 'Purely Satellite Data-Driven Deep Learning Forecast of Complicated Tropical Instability Waves'. *Science Advances* 6 (29): eaba1482. <https://doi.org/10.1126/sciadv.aba1482>.
- Zhou, Zhifang, Tian Xiao, Xiaohong Chen, and Chang Wang. 2016. 'A Carbon Risk Prediction Model for Chinese Heavy-Polluting Industrial Enterprises Based on Support Vector Machine'. *Chaos, Solitons & Fractals, Nonlinear Dynamics and Complexity*, 89 (August): 304–15. <https://doi.org/10.1016/j.chaos.2015.12.001>.
- Zoph, Barret, and Quoc V. Le. 2017. 'Neural Architecture Search with Reinforcement Learning'. *ArXiv:1611.01578 [Cs]*, February. <http://arxiv.org/abs/1611.01578>.

Appendix: List of Abbreviations

ACM	Association for Computing Machinery
AI	Artificial Intelligence
ASICs	Application Specific Integrated Circuits
AWS	Amazon Web Services
Compute	Computing Power
CCAI	Climate Change AI
CPU	Central Processing Unit
CO₂eq	Carbon Dioxide Equivalent
DL	Deep Learning
EU	European Union
ELLIS	European Lab for Learning & Intelligent Systems
FLOPS	Floating Point Operations Per Second
FPGAs	Field Programmable Gate Arrays
GDP	Gross Domestic Product
GDPR	General Data Protection Regulation
GHG	Greenhouse Gas
GPU	Graphics Processing Unit
ICML	International Conference on Machine Learning
ICT	Information Communication Technologies
IJCAI	International Joint Conferences on Artificial Intelligence
IPCC	Intergovernmental Panel on Climate Change
ML	Machine Learning
NeurIPS	Conference on Neural Information Processing Systems
NLP	Natural Language Processing
RL	Reinforcement Learning
SOTA	State-Of-The-Art
SDGs	Sustainable Development Goals
TPU	Tensor Processing Unit
UN	United Nations