PERSPECTIVE
https://doi.org/10.1038/s41558-022-01377-7

nature climate change

Check for updates

Aligning artificial intelligence with climate change mitigation

Lynn H. Kaack[1,2,3] ✉, Priya L. Donti[4,5], Emma Strubell[4], George Kamiya[6], Felix Creutzig[7,8] and David Rolnick[9,10]

There is great interest in how the growth of artificial intelligence and machine learning may affect global GHG emissions. However, such emissions impacts remain uncertain, owing in part to the diverse mechanisms through which they occur, posing difficulties for measurement and forecasting. Here we introduce a systematic framework for describing the effects of machine learning (ML) on GHG emissions, encompassing three categories: computing-related impacts, immediate impacts of applying ML and system-level impacts. Using this framework, we identify priorities for impact assessment and scenario analysis, and suggest policy levers for better understanding and shaping the effects of ML on climate change mitigation.

As artificial intelligence (AI) and particularly machine learning (ML) are increasingly being deployed across society[1], there has been a surge of interest in understanding the effects that ML may have on climate action[2–4]. To explicitly and consistently account for ML in long-term climate and energy projections, and in the design of appropriate policies, the research community needs to develop a holistic and operational understanding of the different ways in which ML can positively and negatively impact climate change mitigation and adaptation strategies. In particular, those impacts that are easiest to measure are likely not those with the largest effects. This can lead to challenges in terms of estimating macro-scale effects, understanding underlying dynamics and trends, and prioritizing actions to align ML with climate strategies. To aid in addressing these challenges, we present a systematic framework (Figs. 1–3) for categorizing the different kinds of impacts of ML on global GHG emissions, including computing-related impacts, the immediate impacts of ML applications and the system-level changes ML induces.

Although the effects of digital technologies on environmental sustainability and GHG emissions have previously been conceptualized (for example, refs. [5,6]), this line of work has been largely overlooked in the discussion around AI and ML. Moreover, existing frameworks need to be extended to include aspects that are unique to ML. Recent work has discussed such ML-specific aspects in part, describing applications of ML for tackling climate change[7], applications of ML that increase emissions[8,9] and the energy consumption of ML through software and hardware[9–11]. A few pieces have engaged with both the positive and negative effects of ML on climate[2,3,12–16], but none have explicitly provided an overview of the different mechanisms by which ML may impact emissions. By presenting a unified framework and detailed overview of these mechanisms, we intend to provide a starting point for research, policy-making and organizational action aiming to better align ML with climate change strategies, as well as augment the broader literature on responsible AI (see, for example, refs. [17–19]).

Related literature on assessing the impacts of information and communications technologies (ICT) has often distinguished between the energy- and hardware-related GHG emissions of ICT ('direct' impacts) and the emissions impacts of ICT's applications ('indirect' impacts)[5,6,20–22]. We build on this work and similarly distinguish between the computing-related GHG emissions of ML and the emissions reductions and increases resulting from applications of ML (Fig. 1). Given that ML encompasses a particularly novel and transformative set of software and analytics approaches with nuanced downstream effects, our framework covers three main categories. The first involves the GHG emissions resulting from computing, caused by both the electricity used for ML computations and the embodied emissions associated with computing hardware. The second involves the 'immediate' GHG emissions effects tied to the short-term outcomes of applications of ML. The third involves the structural or 'system-level' GHG effects induced by these applications. Drawing a clear line between these latter two application-level effects is difficult, with different classifications available throughout the literature (see ref. [20] for an overview); our distinction is adapted from ref. [5] and ref. [6], and, although imperfect, is important for framing the discussion of the overall impacts and associated levers of ML. We report quantitative assessments where available and where we believe these estimates are representative, and we discuss the current state of research on impact assessment. We then use our framework to propose a roadmap for assessing and forecasting impacts, and discuss approaches for shaping the impacts of ML. In terms of scope, our framework predominantly focuses on algorithm-related impacts, and omits impacts relating to data collection and management, ICT and digitalization more broadly[2,20,23–25].

Computing-related impacts

We illustrate two different viewpoints that are relevant to assessing direct computing-related impacts. The first is a 'use-phase' view that aims to assess the energy use of individual ML model instantiations by capturing aspects of the use, development and design of

[1]Data Science Lab, Hertie School, Berlin, Germany. [2]Energy and Technology Policy Group, Department of Humanities, Social and Political Sciences, ETH Zurich, Zurich, Switzerland. [3]Institute of Science, Technology, and Policy, ETH Zurich, Zurich, Switzerland. [4]School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. [5]Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA, USA. [6]International Energy Agency, Paris, France. [7]Mercator Research Institute on Global Commons and Climate Change, Berlin, Germany. [8]Sustainability Economics of Human Settlements, Technical University Berlin, Berlin, Germany. [9]School of Computer Science, McGill University, Montreal, Quebec, Canada. [10]Mila – Quebec AI Institute, Montreal, Quebec, Canada. ✉e-mail: kaack@hertie-school.org

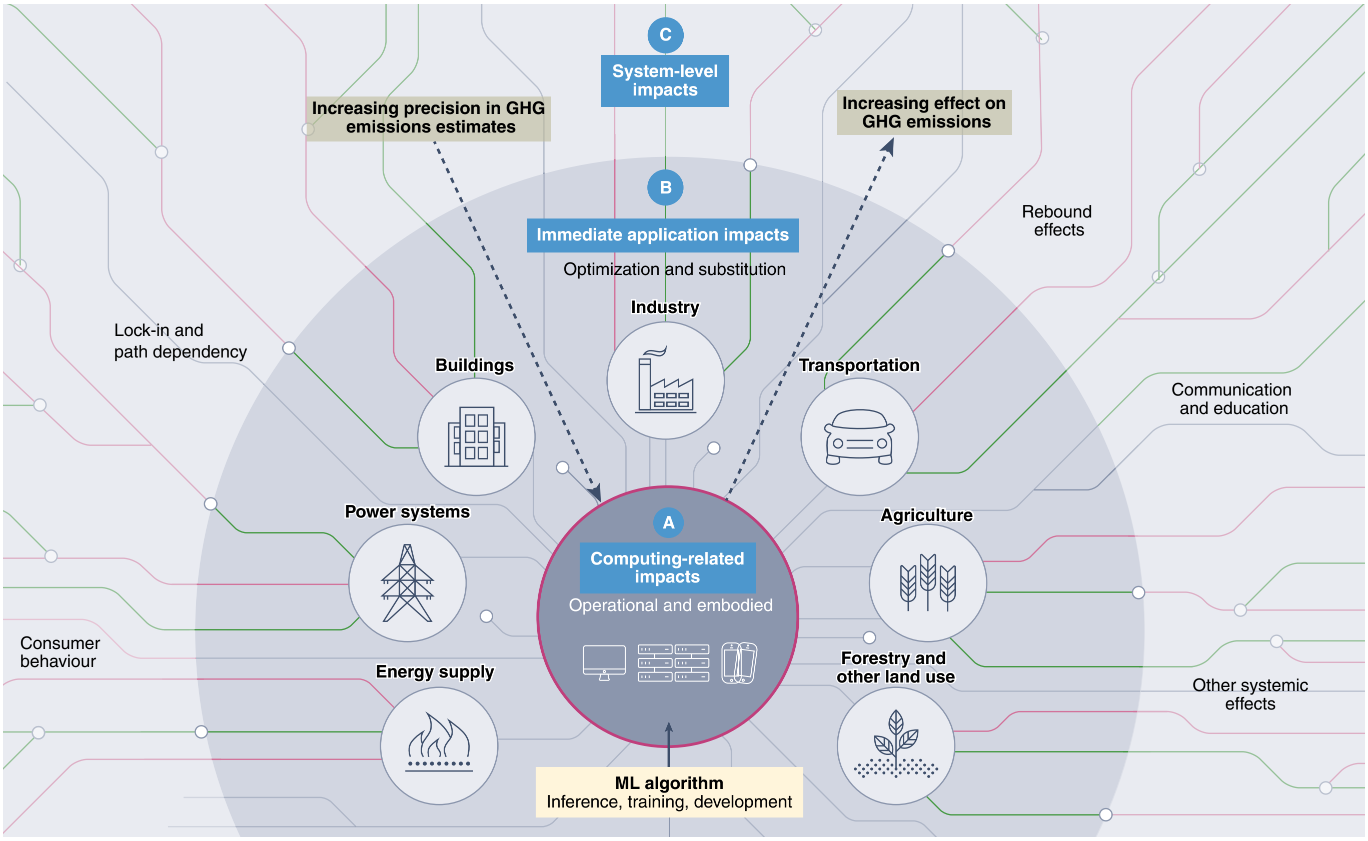

Fig. 1 | Framework for assessing the GHG emissions impacts of ML. We distinguish between three categories (A, B and C) with different kinds of potential emissions impacts, estimation uncertainties, and associated decarbonization levers. Green lines denote effects relating to reductions in GHG emissions, magenta lines relate to increases in emissions, and grey lines symbolize uncertain and/or negligible effects. We provide specifics of Category A of this framework in Fig. 2 and Category B in Fig. 3. Icons adapted with permission from the IEA.

these models. The second is a top-down view that aims to estimate the total global GHG emissions associated with ML workloads, capturing both the sourcing of the electricity used to power computations and embodied emissions from materials extraction and manufacturing.

ML model development and deployment. Creating and running an ML model uses computing power, and therefore energy, with the amount varying dramatically between different algorithms and different stages in the development and use of an ML model. Although many models used in practice are relatively small and can be trained and run on a laptop (such as linear classifiers or decision trees), state-of-the-art performance on more complex tasks is often achieved with very large models, typically using deep learning. The size of the largest deep learning models (measured in number of parameters), and likely the size of the average model, is growing rapidly, leading to much larger demand for computing resources[26–28].

To illustrate how ML models differ so drastically in the energy they consume and better understand approaches to reduce their energy consumption, it is necessary to take a deeper dive into the life-cycle stages of an ML model: model inference (or use), model training, and model development and tuning. Model inference describes the stage where the model is in use in the world. For instance, given new inputs (such as images), the model labels those inputs (for example, it identifies whether an image is of a cat or a dog) according to a function that it has learned. The goal of the model training stage is to learn the underlying function that (for example) maps from inputs to labels by analysing a dataset to choose a set of parameters that define the function. During model development and tuning, a researcher will typically train many different model variants on different datasets to devise a variant that works best in the given problem setting.

We created a schematic overview of relative energy requirements and frequency of each stage of the ML model life cycle (Fig. 2). Model inference is the least energy-intensive process in the ML model life cycle, but it is likely to occur the most frequently. For instance, classifying toxic comments[29,30] or the contents of images[31] on social media requires little power each time a model is used, but may be used on the order of billions of times a day. Also larger models, such as Google's machine translation system, may process more than 100 billion words per day[32]. Those computing requirements can add up: at Facebook the carbon footprint for inference outweighs that for training for certain use cases[33]. The training stage may require many passes over the dataset, often denoted as 'epochs', with each epoch performing full model inference on each example, as well as computing updates to correct the model's prediction for future iterations. In the case of deep learning, for example, this means that the amount of computation required to process each example during training is typically about three times as much as is required during inference for a given model[34]. Training an ML model is thus more energy intensive than using it, but is done much less frequently. Ref. [30] reported that ML models in Facebook's data centres are retrained anywhere from hourly to multi-monthly. The most energy-intensive stage of the ML model life cycle is model development, which requires training many different models. Modern ML models that use neural networks are particularly energy intensive in the development phase as they have many more possible model configurations than their predecessors, and it is not

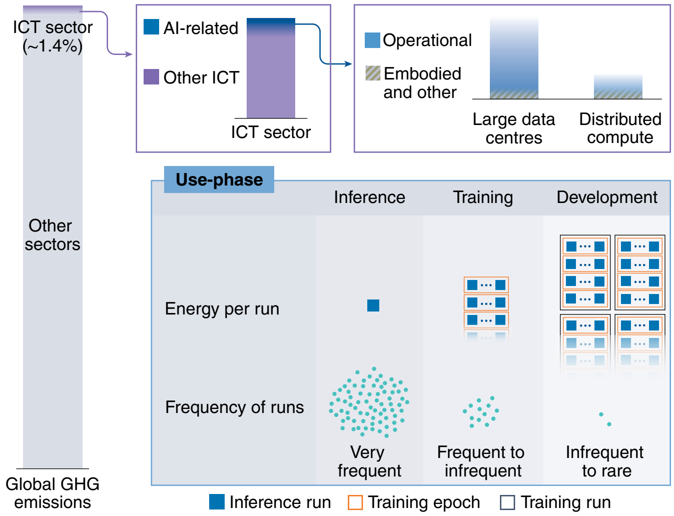

Fig. 2 | Computing-related GHG emissions impacts of ML. The ICT sector accounts for around 1.4% of GHG emissions today, of which ML probably accounts for a small, but unknown, share (indicated by blue shading). Computing-related impacts of ML can be assessed from different perspectives. The majority of ML-related GHG emissions probably come from computing loads in large data centres, with a smaller share from distributed computing (for example, personal computers and smartphones); these GHG emissions result both from operational energy use during computation and from other phases of the hardware life cycle (including embodied emissions). We further break down operational energy use throughout different stages of the model life cycle, with this energy use differing depending on the problem setting and usage patterns.

well understood how those configurations should be set to perform well on a given dataset, except through trial-and-error experimentation and validation ('hyperparameter search'), which often involves thousands of training runs. In the most extreme cases, the GHG emissions associated with developing certain large, cutting-edge models can be comparable to, for example, the lifetime carbon emissions of a car[10], although such computationally intensive processes are performed rarely and by the fewest entities.

The computational requirements of ML models are often described in floating point operations (FLOPs), or the number of additions and multiplications of scalar values required to obtain a result. The precise mapping from FLOPs to energy draw is hardware- and algorithm-dependent, but more FLOPs generally corresponds to higher energy use. ResNet-50[35], a popular deep learning model for image classification, requires about 4 billion FLOPs (and 65 ms) to map a 224 × 224 pixel input image to a label, with an error rate of 24.6% (ref. [36]). A less computationally efficient version, ResNet-152, requires about 11 billion FLOPs (and 150 ms) per image, and obtains only a slightly better error rate of 23.0%. This case illustrates a trade-off in energy-efficient ML: is it worth the more than 2.5 times increase in FLOPs, and corresponding energy consumption, to reduce the error rate by 1.6%? Will the benefits necessarily justify the costs from both an emissions and a broader societal perspective[37–39]?

Software tools for measuring ML model energy use[40] and carbon emissions[41,42] are already available, metrics for reporting model accuracy as a function of computational budget have been proposed[43,44] and benchmarks measuring training and inference efficiency have been established[45,46]. Reference [28] proposed that the amortized computational energy cost of models should be measured across the ML model life cycle to enable a cost–benefit analysis of different use phases (for example increasing training computation to decrease inference computation). However, such reporting is still not standard for researchers and ML software maintainers. Standardized reporting is essential to include efficiency considerations during model development and make energy consumption a criterion when choosing between different ML approaches in practice.

As larger neural network models have become more prominent in certain areas of ML, research into improving the efficiency of ML models has started to expand via methods such as model compression[31,47], devising methods that require fewer training examples to learn a good function[48], reducing retraining frequency and cost[49–51], and conditional use of simpler models in place of more complex ones when examples are less challenging and thus require less computation[52]. The ML research community has also begun to discuss the implications of compressing models and other approaches for improving ML model efficiency on broader performance characteristics[53]. However, the vast majority of ML research and development still focuses on improving model accuracy, rather than balancing accuracy and energy use[11].

Computing infrastructure. The global ICT sector—consisting of all data centres, data transmission networks and connected devices—accounted for around 700 Mt of CO_2-equivalent in 2020, corresponding to around 1.4% of global GHG emissions[54,55]. Around two-thirds of the sector's emissions come from operational energy use (Scope 1 and 2 in the Greenhouse Gas Protocol), with the remainder resulting from materials extraction, manufacturing, transportation and the end-of-life phase (Scope 3)[55]. Although these emissions are relatively small today, especially compared with other sectors and services, policymakers and researchers are increasingly concerned that these emissions could increase as a result of rapid growth in the demand for digital technologies and services, including emerging technologies such as AI/ML.

Only a fraction of emissions from the ICT sector is attributable to AI and ML (Fig. 2), but its exact share is not known due to challenges in boundary definition and a lack of data and established methodology. From the limited information available, we hypothesize that the majority of ML-related workloads today are probably taking place in cloud and hyperscale data centres, with a smaller share occurring on distributed devices such as personal computers. Cloud and hyperscale data centres account for 0.1–0.2% of global GHG emissions[56–58], and it is likely that less than one-quarter of their workloads and traffic are currently ML-related based on estimates for infrastructure-as-a-service and platform-as-a-service[59] and IP traffic related to big data[60]. Over the coming years, edge devices such as smartphones are also expected to handle an increasing volume of inference tasks to reduce latency and dependence on network connectivity[61], with uncertain effects on overall energy use and emissions.

While the amount of computing needed for each of the largest ML training runs is growing rapidly[27], the extent to which efficiency improvements in computing (doubling every 2–3 years)[62,63] can limit overall ML-related energy use in data centres is uncertain. For example, Facebook's overall data centre energy use increased rapidly over the past few years (+40% per year)[64], while computing demands for ML training (for example, +150% per year; ref. [65]) and inference (for example, +105% per year; ref. [66]) have grown even faster. At the same time, by some measures, Facebook's operational GHG footprint (accounting for renewable energy purchases) fell by more than 90% between 2016 and 2020[64], due in part to energy-efficiency improvements and increased renewable electricity procurement.

Energy efficiency has played a central role in limiting the growth in data centre energy demand more generally. Between 2010 and 2018, global data centre energy use rose by only 6%, despite a 550% increase in workloads and computing instances[56]. There have been strong efficiency improvements in servers, storage, networking and infrastructure, as well as a shift away from smaller, less-efficient data

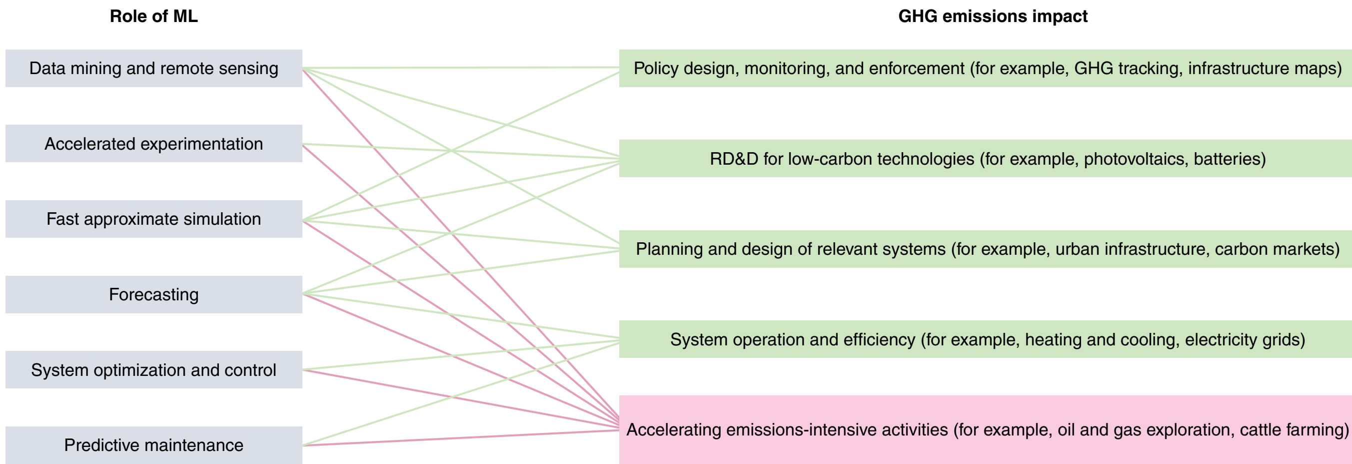

Fig. 3 | Immediate application impacts of ML. ML applications are grouped by their functional role (left) and the associated GHG emissions impacts (right). ML can both reduce emissions (indicated in green) and increase emissions (pink). This diagram differentiates ML applications for addressing climate change in more detail using the findings in ref. [7]; however, the net effect of those applications addressing climate change versus those accelerating emissions-intensive industries is unclear.

centres towards large cloud and hyperscale data centres[23,67], which have higher virtualization, more efficient cooling and increased use of specialized 'AI accelerator' hardware such as application-specific integrated circuits (ASICs) and graphics processing units (GPUs). For instance, a 2017 study found that Google's custom ASIC, the Tensor Processing Unit (TPU), was on average 30–80 times more energy efficient than contemporary CPUs or GPUs[68]. However, the use of GPUs and ASICs for ML applications could drastically increase the power density of data centre racks, which may in turn require liquid cooling technologies and increase water use. Although energy use across all data centres has been flat over the past decade, energy use by large data centres has grown by around 20% annually, and this trend is expected to continue[57]. Limiting overall growth in data centre energy demands over the next decade will therefore require even stronger energy-efficiency improvements. For instance, operators can increase utilization and virtualization to maximize the energy efficiency of existing hardware and infrastructure while replacing hardware when advisable from a life-cycle perspective with the most efficient option. Companies and governments will also need to invest in research, development and demonstration (RD&D) for efficient next-generation computing and communications technologies[56].

Some of the largest data centre operators are now purchasing as much renewable electricity as they consume on a global annual basis[57]; however, this does not guarantee that their data centres are actually fully powered by renewable sources all the time. More ambitious approaches to low-carbon electricity include shifting flexible workloads to times of day (or locations) with higher shares of renewables generation[69] and replacing on-site diesel generators with battery storage.

Computing hardware and infrastructure is also responsible for 'embodied' emissions from raw materials extraction and manufacturing, as well as emissions from transportation and the end-of-life phase (Scope 3). For decentralized computing (for example, desktops, laptops, smartphones), embodied emissions account for 40–80% of devices' life-cycle GHG emissions, whereas for data centres this is typically less than 10% (refs. [55,70–72]). Servers in large data centres are typically replaced every 3–4 years, which can result in higher operational efficiency[56,73]; however, shorter lifespans could also increase the share of life-cycle emissions from manufacturing, which can be mitigated by reusing servers and equipment (such as older GPUs for inference). As data centres become increasingly efficient and powered by clean electricity, the relative importance of emissions from non-operational life-cycle phases will grow—particularly embodied emissions in computing hardware and data centre building construction[33,74].

Immediate application impacts

The broad applicability of ML algorithms means that they can be used both in applications that alleviate bottlenecks in addressing climate change, and in applications that may counteract climate action. In ref. [7], a number of settings in which ML can enable or accelerate climate change mitigation and adaptation strategies were described. These applications span many different areas such as energy, transportation and land use (Fig. 3). For example, via data mining and remote sensing, ML has been used to translate raw data such as text documents or satellite imagery into usable insights for RD&D, policy-making and systems planning—for example by tracking deforestation[75], evaluating susceptibility to coastal inundation[76] and gathering information on corporate climate risk[77]. By accelerating the search for experimental parameters in scientific discovery, ML has been used to aid in the design of next-generation batteries and other materials[78]. By learning from time series, ML has been used to forecast renewable power production[79], crop yields[80] and transportation demands[81]. By controlling and improving the operational efficiency of complex systems, such as industrial heating and cooling systems[82], ML can be used to save resources and energy. ML has also been used to speed up time-intensive physics-based simulations for building design[83] and climate modelling[84]. Predictive maintenance approaches leveraging ML can also be relevant to climate change mitigation when they are applied to low-carbon systems to improve efficiency, reduce costs or build resilience[85].

Although ML is often seen as a 'futuristic' technology, most of these applications are possible with current ML techniques, and many are already being deployed[1,7]. In addition, areas of cutting-edge ML research such as interpretable and probabilistic ML[86,87], physics-integrated ML[88] and transfer learning[89] can both enable new applications and better support integration within existing systems. To support the development and deployment of this kind of work, it will be crucial to facilitate interdisciplinary and applied research via science policy, advance the technological readiness of applications through RD&D programmes and adapt current regulatory environments to mitigate bottlenecks in deployment in relevant sectors and industries. This includes targeted funding and research programmes, testbeds and demonstration projects, public procurement programmes and relevant data management initiatives.

As a general-purpose tool, ML has also been applied in ways that may make climate goals harder to achieve. One such effect is when ML is used to decrease the cost of emissions-intensive activities, thereby potentially increasing their consumption. For example, ML has been used to accelerate oil and gas exploration and extraction by decreasing production costs and boosting reserves[8], which could in turn lead to greater use of fossil fuels. Likewise, ML is used in the 'Internet of Cows' to help manage livestock at scale[90], which can increase cattle farming, an activity already responsible for about 9% of GHG emissions[91]. A potential approach to reduce or avoid the emissions increases associated with such applications is to require ML solutions providers to account for and report the emissions impacts of the applications they support, even if only at the level of order-of-magnitude or qualitative assessments where more detailed numbers are infeasible to obtain. Such reporting can also help address phenomena such as dual use, where stakeholders may use the same ML algorithm for multiple purposes (for example, both in ways that help climate action and in ways that hinder it).

The total immediate impact of ML applications on GHG emissions is extremely difficult to estimate due to the lack of data on the deployment rate of ML, the diversity of application areas and the lack of procedures to appropriately attribute emissions effects to the use of ML algorithms. Although some scientific reviews exist within isolated fields or sectors, the only attempts to provide overall numbers are from ML solutions providers in the private sector[92–94] (these studies are not peer-reviewed and do not disclose all methodologies). We also note that ML can be used with the motivation to elevate the profile of sustainability-related activities in corporations in a way that could provide a false impression of overall organizational sustainability[12].

System-level impacts

While the previous section describes ML applications that are directly beneficial or detrimental to climate change mitigation, many societal ML applications may not have clear immediate impacts on climate change. However, many of these applications can have broader societal implications beyond their immediate impact, and these system-level effects can influence GHG emissions both positively and negatively. Although these kinds of impact may be hard to quantify, they have the potential to outweigh immediate application impacts and are extremely important to consider when evaluating ML use cases.

One pathway to system-level impacts occurs when ML enables changes to a technology that in turn affect the ways in which that technology is used. For example, rebound effects can occur when ML increases the efficiency of a service. Although the improved efficiency may result in lower GHG emissions per use, a decrease in cost may lead to increased consumption of the same (or another) good. This can eat into GHG benefits from efficiency gains or even counteract them[95]. Such rebound effects can be direct, for example by allowing a manufacturing plant to use ML-enabled efficiency gains to increase production of the same goods, thereby (partially) negating emissions savings. Even larger impacts can be expected from more structural types of rebound effects[12,96], which occur (for example) with ML-enabled autonomous driving. Specifically, autonomous vehicles can improve fuel efficiency, but they may also lead to higher rates of individualized vehicle travel, potentially increasing overall energy use and emissions if autonomous vehicles are not shared and/or electrified[97–100].

Given the role of ML as an accelerator of technological development, it may also induce path dependencies that affect climate change mitigation. For instance, the phenomenon of 'lock-in' refers to a scenario in which a particular technology reaches markets first and prevents competitors from entering the market[101]. Depending on how it is applied, ML may end up entrenching the role of a potentially inferior technology in a way that prevents others, for example low-carbon technologies, from entering the market. For instance, the adoption of autonomous vehicles may ingrain the role of trucks and private cars as the dominant means of transportation, instead of enabling infrastructure and space for less emissions-intensive rail, public transit and micromobility options[98]. On the other hand, ML may help break path-dependency effects or create a first-mover advantage for a technology that is beneficial to the climate. The potential effects of ML on such path dependencies in the context of climate change mitigation should be carefully analysed.

Another avenue to system-level impact occurs when ML technologies influence broader lifestyle changes across society, for example by changing the demand for goods and services[96]. A likely negative example here is in advertising, where ML algorithms such as recommender systems can be used to increase the consumption of goods and services with embodied GHG emissions. Given that ML is fuelled by data, its use could also incentivize increasingly large data infrastructures, which can come with their own carbon footprint and systemic implications. Various other paradigm-changing applications of ML have highly unclear effects from a climate perspective—such as in automatic translation tools, virtual assistants and augmented or virtual reality.

These examples demonstrate how important it is to assess the impacts of an ML application at the system level, rather than only estimating marginal effects, and to design public policy to shape system effects. Such policy levers include requiring climate impacts to be considered within regulations surrounding ML-driven emerging technologies[102], and implementing carbon pricing or other mechanisms to incentivize GHG emissions reductions and avoid rebound effects when ML is applied for efficiency. Such climate-cognizant technology assessment should build on and complement frameworks for responsible innovation[103,104] and responsible AI[17–19].

A roadmap for assessing and forecasting impacts

We have discussed the extent to which it is currently possible to estimate the GHG emissions associated with ML above. However, holistic and realistic predictions of the impact of ML across several areas of our framework will require new reporting standards, more data collection, novel measurement methodologies and benchmarking frameworks, and new approaches for developing forecasts and scenarios. Moreover, given the heterogeneous nature of the capabilities, impacts and generalizability of different digitalization technologies, ML and other forms of data analytics warrant separate consideration within impact assessment and attribution frameworks for digital technologies. Such efforts could, for example, build on and extend existing methods and standards for life-cycle assessment (LCA) of ICT to devise approaches that take such heterogeneity into account (see for example, ref. [105] for a LCA of direct effects and ref. [21] for indirect effects, as well as refs. [106,107] for standards). We call on the academic fields of LCA, industrial ecology and others to extend their work to actively grapple with the task of assessing the impacts related to ML, accounting for the ML-specific considerations raised in our framework (Figs. 1–3). Our framework lays out the factors that are relevant for LCA and, depending on the scope of the analysis, can provide a basis on which to estimate the emissions of a company, a particular product, or a policy.

When assessing the GHG emissions impacts of ML, it is important to compare ML approaches to alternatives. Such alternatives are not constrained to other ML models; they can also be other types of analytics approaches that fulfil the same purpose, or can be human decision-making. As ML has enabled many innovations that other methods were unable to attain, in some cases the baseline may be a world where such innovations did not exist at all. The choice of an appropriate baseline depends on the aim of the analysis, as well as the category of the impacts being assessed.

Table 1 | Levers to reduce the GHG emissions impacts associated with ML computing and applications

Lever type	Computing-related (algorithm, infrastructure)	Application-related (immediate, systemic)
Public sector		
Economic instruments	*Implement economy-wide or sector-specific carbon pricing to incentivize emissions reductions and mitigate rebound effects*[a]	
RD&D	Support research in energy-efficient ML[a]	Support interdisciplinary and applied ML research for climate-relevant applications of ML[a]
	Support RD&D in energy-efficient, specialized and low-resource hardware[a]	Provide mechanisms to advance the technological readiness of climate-beneficial ML applications (for example, testbeds, demonstration projects, public procurement programmes)[b]
	Support RD&D in data centre operational efficiency[a]	
Regulation	Employ a climate-cognizant technology assessment lens within AI strategies and when regulating ML-driven emerging technologies[c]	
	Implement clean electricity mandates (for example, low-carbon portfolio standards)[a]	*Employ regulatory approaches to constrain sector-specific GHG emissions*[a]
	Implement efficiency standards for data centre hardware and infrastructure[a]	Reduce deployment barriers in relevant sectors and industries for AI applications that are beneficial to the climate[b]
Best practices and standards	Develop interoperability standards for commercial ML approaches to prevent lock-in to particular solutions providers and facilitate a decentralized solutions provider space[b]	
	Develop and implement standardized metrics for evaluating model efficacy that include energy efficiency[a]	Require meaningful civic and stakeholder engagement in scoping, developing, and deploying ML-driven projects[a]
		Implement data governance standards that spur impactful work and are mindful of privacy and ownership[b]
		Develop best practices and systematic approaches to weigh benefits and costs for ML applications[c]
Monitoring and reporting	Develop measurement methodologies and guidance to estimate and report ML-related GHG emissions[b]	
	Mandate appropriate life-cycle transparency and reporting of GHG emissions for ML use cases, including both computing and application-related impacts[b]	
Capacity building	Build in-house public-sector capacity in ML to facilitate governance and deployment[b]	
	Promote ML education and literacy among climate-relevant entities and in the public sector[a]	
		Incentivize ML workforce shifts towards climate-oriented entities (for example, via placement programmes)[b]
Private sector		
Corporate climate action	Adopt organizational carbon pricing strategies that account for both computing- and application-related emissions (for example, Scope 1, 2 and 3 emissions, including those from cloud computing as well as from products and services)[b]	
	Reduce wasteful model retraining and execution[a]	Adjust business models to avoid ML applications that drive GHG emissions increases[a]
	Make energy efficiency a central criterion in evaluating model efficacy[a]	Encourage ML applications that drive GHG emissions reductions[a]
	Reduce GHG emissions across supply chains and product life cycle (including embodied emissions)[a]	Measure and engage in voluntary reporting of the emissions impacts of ML products and services[b]
	Maximize energy efficiency in data centres and support related RD&D[a]	
	Shift computing loads to geographies and times with lower carbon-intensity of the grid[a]	
	Purchase low-carbon electricity and invest in energy technologies to decarbonize the grid[a]	
	Develop standardized ML platforms to facilitate rapid company-wide adoption of energy efficiency improvements[c]	

[a]Policies that are ready to implement or already exist. [b]Policies that can be developed today. [c]More analysis needed to develop policies. General policy levers are set in italics.

To estimate computing-related impacts (Category A of our framework), better access to information will be crucial. For example, although is relatively straightforward to estimate the computing-related GHG emissions resulting from individual runs of AI systems, the usage patterns in practice are typically opaque. Where appropriate, practitioners could disclose information about such usage patterns, as well as other inputs relevant to computing GHG emissions impacts resulting from ML system development, training/fine-tuning and inference (for example, specifics about the model type and size, training requirements for model development and the type of pre-trained model used, the type and location of computing infrastructure used and the frequency of training/retraining/fine-tuning and inference). This information can help provide an understanding of industry trends and aid in the development of best practices and benchmarks for trade-offs between different approaches. We discuss feasibility considerations around reporting requirements in the next section. To estimate the total computing-related GHG emissions, an important data point is the share of the total computing load in data centres that can be attributed to ML, ideally distinguished by the relevant model life-cycle stages. This information would allow for a top-down estimate of global computing-related impacts and underlying dynamics, but is not made public by data centre operators at present.

There are currently limited quantitative estimates available about the immediate impacts of ML applications (Category B). The lack of established methodology poses a central bottleneck here. Research and practice need to establish how to estimate the marginal and counterfactual benefit that ML could have if introduced in established processes, including distinguishing between use cases that would not exist without ML versus those where ML provides improvements to an existing use case. For such efforts, it will be important to develop a more fine-grained taxonomy of ML systems and application areas that can help to generalize beyond single case studies and also help stakeholders assess the costs and benefits of new projects a priori. We provide a starting point for developing such an approach, as illustrated in Fig. 3, encompassing such diverse potential effects of ML as accelerated technological innovation, more informed decision-making via improved analytics and increased energy efficiency of industry operations. In addition, obtaining better data will be difficult, yet particularly crucial, considering the potentially large magnitude and uncertainty around those developments. To estimate impacts more broadly and systematically, reviewing, synthesizing and generalizing case studies will be important, and where data cannot be easily obtained, approaches such as stakeholder surveys or expert elicitation might help to fill gaps.

Perhaps the most important impacts, yet most difficult to assess, are the system-level impacts (Category C). ML is a fast-growing enabling technology that has the potential to affect present and future societal and technological trajectories and thus needs to be appropriately accounted for in forecasting and scenario analysis. ML can influence many input factors of climate and energy system models, such as efficiency, production and consumption rates, learning rates, resource constraints, financial assumptions and so on, which make ML a 'wild card' that could introduce large transformations in different ways. How to appropriately factor that uncertainty into climate and energy system models is yet to be established. Importantly, ML builds on digital infrastructure, yet the impact assessment of digitalization is itself at an early stage (especially when it comes to estimating the impacts of how digital technologies are applied)[106,107]. Energy and climate models, such as energy system models developed by the International Energy Agency and the US Energy Information Administration, or the Shared Socioeconomic Pathways used by the IPCC, generally do not explicitly or systematically account for digitalization, let alone the effects introduced by ML. One exception is perhaps the inclusion of autonomous vehicles in scenarios, for example, by the US Energy Information Administration[100]. Our framework can be used as a starting point, and is sufficiently general to provide a comprehensive framing for incorporating current and future ML effects within scenario analysis.

Aligning ML with climate change mitigation

Given the multi-faceted relationship of ML with climate change, many different kinds of approaches from the public and private sectors are needed to shape its impacts. This will require progress in both climate policy and AI policy, coupled with algorithmic and hardware innovation and the development of adequate impact assessment methodologies. We outline a number of strategies that can address the different emissions effects laid out in our framework (Table 1).

Despite not addressing ML explicitly, general climate policy approaches such as carbon pricing may be effective in driving the development and use of ML in a manner that is aligned with climate change mitigation. Science policy approaches that foster low-carbon technologies may also facilitate uses of ML that enable or improve these technologies (although they may not necessarily address ML-specific barriers). To address more technology-specific opportunities and risks, it will be important for climate change to become a major consideration within AI innovation and deployment policies. This includes (1) promoting the research, development and deployment of ML applications that are beneficial to the climate, (2) requiring transparency and accountability for those use cases that could increase emissions or otherwise counteract climate change goals, as well as on computational energy use, and (3) employing climate-cognizant technology assessment for ML use cases that are not traditionally within the realm of climate policy, but where decisions today may have important implications for future climate impacts. Many of the associated policy approaches in Table 1 can be developed and implemented starting today.

Furthermore, mandating emissions measurement and reporting for ML use cases—considering the impacts of both computing and applications—can enable these emissions to be regulated via climate policy approaches, and further shape the design of targeted policies. Such reporting requirements, however, need to be carefully designed on the basis of an understanding of where top-down measurement might suffice to inform regulatory approaches, what is feasible to estimate given present measurement methodologies, and the costs and burdens associated with reporting (see, for example, ref. [108] regarding costs for enterprises). Such requirements should also be implemented with an eye towards preventing strategic behaviour such as the 'hiding' of emissions in cloud computing servers[109]. Climate-related reporting for ML-based systems could potentially be more easily implemented where other AI reporting requirements are planned (such as those proposed in the EU[110]).

Finally, we note that ML expertise today is often concentrated among a limited set of actors, raising potential challenges with respect to the governance and implementation of ML in the context of climate change. For instance, the use of ML in certain contexts may yield or exacerbate societal inequities, for example, by widening the digital divide[37,111], through algorithmic bias[112] or by shifting power from public to large private entities by virtue of who controls relevant data or intellectual capital. Strategies to address such gaps include strengthening small and medium-sized ML solutions providers, developing incentives such as placement programmes and dedicated education to shift the ML workforce towards public and climate-relevant entities, developing interoperability standards to prevent lock-in to particular solutions providers and developing best practices for when state-of-the-art ML models versus other (simpler) alternatives should be used. Fostering meaningful civic engagement processes for the scoping, design and deployment of projects (and associated data collection and provision efforts) will also be critical to ensuring that ML approaches are both effective and avoid potential pitfalls[113].

The ultimate effect of ML on the climate is far from predestined, and societal decisions will play a large role in shaping its overall impacts[114,115]. This will require a holistic portfolio of approaches across policy, industry and academia to incentivize uses of ML that support climate change strategies while mitigating the impacts of use cases that may counteract climate change goals. Most importantly, society cannot wait to act: with the rapidly growing prevalence of ML and the increasing urgency of climate change, we now have a critical window of opportunity to shape the impacts of ML for decades to come.

Received: 7 April 2021; Accepted: 27 April 2022;
Published online: 9 June 2022

References

1. Zhang, D. et al. *Artificial Intelligence Index Report 2021* (AI Index Steering Committee, Human-Centered AI Institute, 2021).
2. *Digital Technology and the Planet: Harnessing Computing to Achieve Net Zero* (Royal Society, 2020); https://royalsociety.org/-/media/policy/projects/digital-technology-and-the-planet/digital-technology-and-the-planet-report.pdf
3. Kaack, L. H., Donti, P. L., Strubell, E. & Rolnick, D. *Artificial Intelligence and Climate Change: Opportunities, Considerations, and Policy Levers to Align AI with Climate Change Goals* (Heinrich-Böll-Stiftung, 2020); https://eu.boell.org/en/2020/12/03/artificial-intelligence-and-climate-change
4. *Harnessing Artificial Intelligence to Accelerate the Energy Transition* (World Economic Forum, 2021); https://www.weforum.org/whitepapers/harnessing-artificial-intelligence-to-accelerate-the-energy-transition
5. Berkhout, F. & Hertin, J. De-materialising and re-materialising: digital technologies and the environment. *Futures* **36**, 903–920 (2004).
6. Hilty, L. M. & Aebischer, B. in *ICT Innovations for Sustainability* (eds Hilty, L. M. & Aebischer, B) 3–36 (Springer, 2015).
7. Rolnick, D. et al. Tackling climate change with machine learning. *ACM Comput. Surv*. https://doi.org/10.1145/3485128 (2022).
8. *Oil in the Cloud: How Tech Companies are Helping Big Oil Profit from Climate Destruction* (Greenpeace, 2019); https://www.greenpeace.org/usa/reports/oil-in-the-cloud/
9. Dobbe, R. & Whittaker, M. AI and climate change: how they're connected, and what we can do about it. *Medium* https://medium.com/@AINowInstitute/ai-and-climate-change-how-theyre-connected-and-what-we-can-do-about-it-6aa8d0f5b32c (2019).
10. Strubell, E., Ganesh, A. & McCallum, A. Energy and policy considerations for deep learning in NLP. In *Proc. 57th Annual Meeting of the Association for Computational Linguistics* 3645–3650 (Association for Computational Linguistics, 2019); https://doi.org/10.18653/v1/P19-1355
11. Schwartz, R., Dodge, J., Smith, N. A. & Etzioni, O. Green AI. *Commun. ACM* **63**, 54–63 (2020).
12. Dauvergne, P. Is artificial intelligence greening global supply chains? Exposing the political economy of environmental costs. *Rev. Int. Polit. Econ.* https://doi.org/10.1080/09692290.2020.1814381 (2020).
13. Coeckelbergh, M. AI for climate: freedom, justice, and other ethical and political challenges. *AI Ethics* **1**, 67–72 (2021).
14. Gunther, H. & Rose, J. Governing AI: the importance of environmentally sustainable and equitable innovation. *Environ. Law Rep.* **50**, 10888 (2020).
15. Stein, A. L. Artificial intelligence and climate change. *Yale J. Reg*. **37**, 890–939 (2020).
16. Cowls, J., Tsamados, A., Taddeo, M. & Floridi, L. The AI gambit-leveraging artificial intelligence to combat climate change: opportunities, challenges, and recommendations. *SSRN* https://ssrn.com/abstract=3804983 (2021).
17. *FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (ACM, 2021).
18. *Recommendation on the Ethics of Artificial Intelligence* (UNESCO, 2021); https://unesdoc.unesco.org/ark:/48223/pf0000380455
19. *Areas for Future Action in the Responsible AI Ecosystem* (The Future Society, GPAI Responsible Development, Use and Governance of AI Working Group & CEIMIA, 2020); https://www.gpai.ai/projects/responsible-ai/areas-for-future-action-in-responsible-ai.pdf
20. Horner, N. C., Shehabi, A. & Azevedo, I. L. Known unknowns: indirect energy effects of information and communication technology. *Environ. Res. Lett.* **11**, 103001 (2016).
21. Bieser, J. & Hilty, L. Indirect effects of the digital transformation on environmental sustainability: methodological challenges in assessing the greenhouse gas abatement potential of ICT. In *5th International Conference on Information and Communication Technology for Sustainability* 68–81 (EasyChair, 2018); https://doi.org/10.29007/lx7q
22. Pohl, J., Hilty, L. M. & Finkbeiner, M. How LCA contributes to the environmental assessment of higher order effects of ICT application: a review of different approaches. *J. Clean. Prod.* **219**, 698–712 (2019).
23. *Digitalization & Energy* (OECD/IEA, 2017); https://www.iea.org/reports/digitalisation-and-energy
24. Sivaram, V. et al. *Digital Decarbonization Promoting Digital Innovations to Advance Clean Energy Systems* (Council on Foreign Relations, 2018); https://www.cfr.org/report/digital-decarbonization
25. Wilson, C., Kerr, L., Sprei, F., Vrain, E. & Wilson, M. Potential climate benefits of digital consumer innovations. *Annu. Rev. Environ. Resour.* **45**, 113–144 (2020).
26. Canziani, A., Paszke, A. & Culurciello, E. An analysis of deep neural network models for practical applications. Preprint at https://arxiv.org/abs/1605.07678 (2017).
27. AI and compute. *OpenAI* https://openai.com/blog/ai-and-compute (16 May 2018).
28. Bommasani, R. et al. On the opportunities and risks of foundation models. Preprint at https://arxiv.org/abs/2108.07258 (2021).
29. Joulin, A., Grave, E., Bojanowski, P. & Mikolov, T. Bag of tricks for efficient text classification. In *Proc. 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2* 427–431 (Association for Computational Linguistics, 2017); https://www.aclweb.org/anthology/E17-2068
30. Hazelwood, K. et al. Applied machine learning at Facebook: a datacenter infrastructure perspective. In *2018 IEEE International Symposium on High Performance Computer Architecture* 620–629 (IEEE, 2018); https://doi.org/10.1109/HPCA.2018.00059
31. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A. & Chen, L.-C. MobileNetV2: inverted residuals and linear bottlenecks. In *IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2018).
32. Turovsky, B. Ten years of Google Translate. *The Keyword* https://blog.google/products/translate/ten-years-of-google-translate (2016).
33. Wu, C. J. et al. Sustainable AI: environmental implications, challenges and opportunities. In *Proc. Machine Learning and Systems 4* 795–813 (MLSys, 2022).
34. Jiang, A. H. et al. Accelerating deep learning by focusing on the biggest losers. Preprint at https://arxiv.org/abs/1910.00762 (2019).
35. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition* 770–778 (IEEE, 2016); https://doi.org/10.1109/CVPR.2016.90
36. Albanie, S. ConvNet-Burden: estimates of memory consumption and FLOP counts for various convolutional neural networks. *GitHub* https://github.com/albanie/convnet-burden (2019).
37. Bender, E. M., Gebru, T., McMillan-Major, A. & Shmitchell, S. On the dangers of stochastic parrots: can language models be too big? In *Proc. 2021 ACM Conference on Fairness, Accountability, and Transparency* 610–623 (ACM, 2021).
38. Gupta, A., Lanteigne, C. & Kingsley, S. SECure: a social and environmental certificate for AI Systems. Preprint at https://arxiv.org/abs/2006.06217 (2020).
39. Tomašev, N. et al. AI for social good: unlocking the opportunity for positive impact. *Nat. Commun.* **11**, 2468 (2020).
40. Henderson, P. et al. Towards the systematic reporting of the energy and carbon footprints of machine learning. *J. Mach. Learn. Res.* **21**, 1–43 (2020).
41. Schmidt, V. et al. CodeCarbon: estimate and track carbon emissions from machine learning computing. *Zenodo* https://doi.org/10.5281/zenodo.4658424 (2021).
42. Anthony, L. F. W., Kanding, B. & Selvan, R. Carbontracker: tracking and predicting the carbon footprint of training deep learning models. Preprint at https://arxiv.org/abs/2007.03051 (2020).
43. Cai, E., Juan, D., Stamoulis, D. & Marculescu, D. NeuralPower: predict and deploy energy-efficient convolutional neural networks. In *9th Asian Conference on Machine Learning* (ACML, 2017).
44. Dodge, J., Gururangan, S., Card, D., Schwartz, R. & Smith, N. A. Show your work: improved reporting of experimental results. In *Proc. 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing* 2185–2194 (Association for Computational Linguistics, 2019); https://doi.org/10.18653/v1/D19-1224
45. Mattson, P. et al. (eds) *Proc. Machine Learning and Systems 2* 336–349 (MLSys, 2020).
46. Reddi, V. J. et al. MLPerf inference benchmark. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture* 446–459 (IEEE, 2020); https://doi.org/10.1109/ISCA45697.2020.00045
47. Hinton, G., Vinyals, O. & Dean, J. Distilling the knowledge in a neural network. In *NeurIPS Deep Learning Workshop* (NeurIPS, 2014).
48. Schaul, T., Quan, J., Antonoglou, I. & Silver, D. Prioritized experience replay. In *International Conference on Learning Representations* (ICLR, 2016).
49. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R. & Bengio, Y. Quantized neural networks: training neural networks with low precision weights and activations. *J. Mach. Learn. Res.* **18**, 1–30 (2018).
50. Pfeiffer, J. et al. Adapterhub: a framework for adapting transformers. In *Proc. 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* 46–54 (Association for Computational Linguistics, 2020).
51. Cai, H., Gan, C., Wang, T., Zhang, Z. & Han, S. Once-for-all: train one network and specialize it for efficient deployment. In *International Conference on Learning Representations* (ICLR, 2020).
52. Lepikhin, D. et al. GShard: scaling giant models with conditional computation and automatic sharding. In *International Conference on Learning Representations* (ICLR, 2021).
53. Hooker, S., Moorosi, N., Clark, G., Bengio, S. & Denton, E. Characterizing and mitigating bias in compact models. In *ICML Workshop on Human Interpretability in Machine Learning* (ICML, 2020).

54. *Greenhouse Gas Emissions Trajectories for the Information and Communication Technology Sector Compatible with the UNFCCC Paris Agreement* (International Telecommunication Union, 2020); http://handle.itu.int/11.1002/1000/14084
55. Malmodin, J. & Lundén, D. The energy and carbon footprint of the global ICT and E&M sectors 2010–2015. *Sustainability* **10**, 3027 (2018).
56. Masanet, E., Shehabi, A., Lei, N., Smith, S. & Koomey, J. Recalibrating global data center energy-use estimates. *Science* **367**, 984–986 (2020).
57. *Data Centres and Data Transmission Networks* (International Energy Agency, 2021); https://www.iea.org/reports/data-centres-and-data-transmission-networks
58. Montevecchi, F., Stickler, T., Hintemann, R. & Hinterholzer, S. *Energy-Efficient Cloud Computing Technologies and Policies for an Eco-Friendly Cloud Market* (2020); https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=71330
59. *Cisco Global Cloud Index: Forecast and Methodology, 2016–2021* (Cisco, 2018); https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.pdf
60. Compton, C. Cisco's Global Cloud Index study: acceleration of the multicloud era. *Cisco Blogs* https://blogs.cisco.com/news/acceleration-of-multicloud-era (2018).
61. Wu, C. et al. Machine learning at Facebook: understanding inference at the edge. In *2019 IEEE International Symposium on High Performance Computer Architecture* 331–344 (IEEE, 2019).
62. Koomey, J., Berard, S., Sanchez, M. & Wong, H. Implications of historical trends in the electrical efficiency of computing. *IEEE Ann. Hist. Comput.* **33**, 46–54 (2010).
63. Koomey, J. & Naffziger, S. Moore's law might be slowing down, but not energy efficiency. *IEEE Spectrum* **52**, 35 (2015).
64. *Facebook Sustainability Data 2020* (Facebook, 2021); https://sustainability.fb.com/wp-content/uploads/2021/06/2020_FB_Sustainability-Data.pdf
65. Naumov, M. et al. Deep learning training in Facebook data centers: design of scale-up and scale-out systems. Preprint at https://arxiv.org/abs/2003.09518 (2020).
66. Park, J. et al. Deep learning inference in Facebook data centers: characterization, performance optimizations and hardware implications. Preprint at https://arxiv.org/abs/1811.09886 (2018).
67. Shehabi, A. et al. *United States Data Center Energy Usage Report* (Lawrence Berkeley National Laboratory, 2016); https://eta.lbl.gov/publications/united-states-data-center-energy
68. Jouppi, N. P. et al. In-datacenter performance analysis of a tensor processing unit. In *Proc. 44th Annual International Symposium on Computer Architecture* 1–12 (ACM, 2017).
69. Radovanovic, A. Our data centers now work harder when the sun shines and wind blows. *The Keyword* https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows (2020).
70. Whitehead, B., Andrews, D. & Shah, A. The life cycle assessment of a UK data centre. *Int. J. Life Cycle Assess.* **20**, 332–349 (2015).
71. Masanet, E., Shehabi, A. & Koomey, J. Characteristics of low-carbon data centres. *Nat. Clim. Change* **3**, 627–630 (2013).
72. Hischier, R., Coroama, V. C., Schien, D. & Achachlouei, M. A. in *ICT Innovations for Sustainability* (eds Hilty, L. M. & Aebischer, B.) 171–189 (Springer, 2015).
73. André Barroso, L., Clidaras, J. & Hölzle, U. The datacenter as a computer: an introduction to the design of warehouse-scale machines. *Synth. Lect. Comput. Arch.* **8**, 1–154 (2013).
74. Gupta, U. et al. Chasing carbon: the elusive environmental footprint of computing. In *IEEE International Symposium on High-Performance Computer Architecture* 854–867 (IEEE, 2021).
75. Finer, M. et al. Combating deforestation: from satellite to intervention. *Science* **360**, 1303–1305 (2018).
76. Kulp, S. A. & Strauss, B. H. New elevation data triple estimates of global vulnerability to sea-level rise and coastal flooding. *Nat. Commun.* **10**, 4844 (2019).
77. Friederich, D., Kaack, L. H., Luccioni, A. & Steffen, B. Automated identification of climate risk disclosures in annual corporate reports. Preprint at https://arxiv.org/abs/2108.01415 (2021).
78. Liu, Y., Guo, B., Zou, X., Li, Y. & Shi, S. Machine learning assisted materials design and discovery for rechargeable batteries. *Energy Storage Mater*. **31**, 434–450 (2020).
79. Ahmed, R., Sreeram, V., Mishra, Y. & Arif, M. D. A review and evaluation of the state-of-the-art in PV solar power forecasting: techniques and optimization. *Renew. Sustain. Energy Rev.* **124**, 109792 (2020).
80. You, J., Li, X., Low, M., Lobell, D. & Ermon, S. Deep Gaussian process for crop yield prediction based on remote sensing data. In *Thirty-First AAAI Conference on Artificial Intelligence* 4559–4565 (ACM, 2017).
81. Toqué, F., Khouadjia, M., Come, E., Trepanier, M. & Oukhellou, L. Short & long term forecasting of multimodal transport passenger flows with machine learning methods. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems* 560–566 (IEEE, 2017).
82. Evans, R. & Gao, J. DeepMind AI reduces Google data centre cooling bill by 40%. *DeepMind* https://www.deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40 (2016).
83. Roman, N. D., Bre, F., Fachinotti, V. D. & Lamberts, R. Application and characterization of metamodels based on artificial neural networks for building performance simulation: a systematic review. *Energy Build.* **217**, 109972 (2020).
84. Irrgang, C. et al. Towards neural Earth system modelling by integrating artificial intelligence in Earth system science. *Nat. Mach. Intell.* **3**, 667–674 (2021).
85. Jenssen, R. et al. Automatic autonomous vision-based power line inspection: a review of current status and the potential role of deep learning. *Int. J. Electr. Power Energy Syst.* **99**, 107–120 (2018).
86. Rudin, C. et al. Interpretable machine learning: fundamental principles and 10 grand challenges. *Stat. Surv.* **16**, 1–85 (2022).
87. Ghahramani, Z. Probabilistic machine learning and artificial intelligence. *Nature* **521**, 452–459 (2015).
88. Willard, J., Jia, X., Xu, S., Steinbach, M. & Kumar, V. Integrating physics-based modeling with machine learning: a survey. Preprint at https://arxiv.org/abs/2003.04919 (2020).
89. Zhuang, F. et al. A comprehensive survey on transfer learning. *Proc. IEEE* **109**, 43–76 (2020).
90. Adams-Progar, A., Fink, G.A.,Walker, E. & Llewellyn, D. in *Security and Privacy in Cyber-Physical Systems: Foundations, Principles and Applications* (eds Song, H. et al.) Ch. 18 (Wiley, 2017); https://doi.org/10.1002/9781119226079.ch18
91. Charles, H. et al. Meat consumption, health, and the environment. *Science* **361**, eaam5324 (2018).
92. Herweijer, C., Combes, B. & Gillham, J. *How AI Can Enable a Sustainable Future* (Microsoft & PWC, 2018); https://www.pwc.co.uk/services/sustainability-climate-change/insights/how-ai-future-can-enable-sustainable-future.html
93. *Climate AI: How Artificial Intelligence Can Power Your Climate Action Strategy* (Capgemini, 2020); https://www.capgemini.com/research/climate-ai/
94. Degot, C., Duranton, S., Frédeau, M. & Hutchinson, R. Reduce carbon and costs with the power of AI. *BCG* https://www.bcg.com/en-us/publications/2021/ai-to-reduce-carbon-emissions (2021).
95. Azevedo, I. M. L. Consumer end-use energy efficiency and rebound effects. *Annu. Rev. Environ. Resour.* **39**, 393–418 (2014).
96. Lange, S., Pohl, J. & Santarius, T. Digitalization and energy consumption. Does ICT reduce energy demand? *Ecol. Econ.* **176**, 106760 (2020).
97. Anderson, J. M. et al. *Autonomous Vehicle Technology: A Guide for Policymakers* (RAND Corporation, 2016); https://doi.org/10.7249/RR443-2
98. Creutzig, F. et al. Leveraging digitalization for sustainability in urban transport. *Glob. Sustain.* **2**, e14 (2019).
99. Wadud, Z., MacKenzie, D. & Leiby, P. Help or hindrance? The travel, energy and carbon impacts of highly automated vehicles. *Transport. Res. A* **86**, 1–18 (2016).
100. Chase, N., Maples, J. & Schipper, M. *Autonomous Vehicles: Uncertainties and Energy Implications* (EIA, 2018); https://www.eia.gov/outlooks/aeo/av.php
101. Arthur, W. B. Competing technologies, increasing returns, and lock-in by historical events. *Econ. J.* **99**, 116–131 (1989).
102. Cox, E., Royston, S. & Selby, J. *Impact of Non-Energy Policies on Energy Systems* (UK Energy Research Centre, 2016); https://ukerc.ac.uk/publications/impact-of-non-energy-policies-on-energy-systems/
103. Stilgoe, J., Owen, R. & Macnaghten, P. Developing a framework for responsible innovation. *Res. Policy* **42**, 1568–1580 (2013).
104. Jirotka, M., Grimpe, B., Stahl, B., Eden, G. & Hartswood, M. Responsible research and innovation in the digital age. *Commun. ACM* **60**, 62–68 (2017).
105. Itten, R. et al. Digital transformation-life cycle assessment of digital services, multifunctional devices and cloud computing. *Int. J. Life Cycle Assess.* **25**, 2093–2098 (2020).
106. Coroamă, V. C., Bergmark, P., Höjer, M. & Malmodin, J. A methodology for assessing the environmental effects induced by ICT services: part I. Single services. In *Proc. 7th International Conference on ICT for Sustainability* 36–45 (ACM, 2020).
107. Bergmark, P., Coroamă, V. C., Höjer, M. & Donovan, C. A Methodology for assessing the environmental effects induced by ICT services: part I. Multiple services and companies. In *Proc. 7th International Conference on ICT for Sustainability* 46–55 (ACM, 2020).
108. Haataja, M. & Bryson, J. J. What costs should we expect from the EU's AI Act? Preprint at *SocArXiv* https://osf.io/preprints/socarxiv/8nzb4 (2021).
109. Mytton, D. Hiding greenhouse gas emissions in the cloud. *Nat. Clim. Change* **10**, 701–701 (2020).
110. *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts* (European Commission, 2021); https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206

111. Hilbert, M. Big data for development: a review of promises and challenges. *Dev. Policy Rev.* **34**, 135–174 (2016).
112. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K. & Galstyan, A. A survey on bias and fairness in machine learning. in *ACM Computing Surveys* Vol. 54, 1–35 (ACM, 2021).
113. Bondi, E., Xu, L., Acosta-Navas, D. & Killian, J. A. Envisioning communities: a participatory approach towards AI for social good. In *Proc. 2021 AAAI/ACM Conference on AI, Ethics, and Society* 425–436 (ACM, 2021).
114. Pinch, T. J. & Bijker, W. E. The social construction of facts and artefacts: or how the sociology of science and the sociology of technology might benefit each other. *Soc. Stud. Sci.* **14**, 399–441 (1984).
115. Klein, H. K. & Lee Kleinman, D. The social construction of technology: structural considerations. *Sci. Technol. Human Values* **27**, 28–52 (2002).

Acknowledgements

We thank the Environmental Law Institute and Alfred P. Sloan Foundation for their support of this work. P.L.D. was also supported by a US Department of Energy Computational Science Graduate Fellowship (DE-FG02-97ER25308), the Center for Climate and Energy Decision Making through a cooperative agreement between the US National Science Foundation and Carnegie Mellon University (grant number SES-00949710) and the Siebel Scholars programme. D.R. was supported by the Canada CIFAR AI Chairs programme. We are grateful to W.Y. for assistance with the figures. We thank P. Bergmark, V. Coroama, D. Daniels, J. Dunietz, L. Klaaβen, S. Luccioni, J. Malmodin, D. Rejeski, C. Samaras, N. Schmid, T. S. Schmidt, M. Scheutz, S. Sewerin, B. Steffen and M. Voss for their input and comments on the manuscript. Opinions are G.K.'s own and do not reflect those of the OECD, the IEA or their member countries.

Author contributions

P.L.D., L.H.K. and D.R. conceived the idea for this manuscript. All authors wrote and edited the manuscript text and figures, with primary contributions from E.S. and G.K. to the section on computing-related impacts, from D.R., F.C. and L.H.K. to the sections on application-related impacts, from P.L.D. to the section on shaping ML's impacts and from L.H.K. to the introduction, roadmap for assessing impacts and overall conceptual framing.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence should be addressed to Lynn H. Kaack.

Peer review information *Nature Climate Change* thanks Lorenz Hilty, Christopher Irrgang and Yang Yu for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.