


Artificial Footprints: the environmental impact of AI

A report from Register Dynamics

By Owain Jones

February 2024

R E G I S T E R
D Y N A M I C S



Artificial Intelligence is set to change the way we live. But have we thought about how it could change our planet?

AI technology is advancing rapidly, with multiple new products gaining widespread popularity in recent years, from personal assistants such as Siri and Alexa to more complex systems like ChatGPT. As AI becomes more prevalent, concerns around it have also grown. Substantial government and media attention has been devoted to understanding AI's potential impact on jobs, copyright and privacy. The EU even recently introduced the first piece of legislation to regulate AI [8], with a focus on protecting workers and personal privacy.

Meanwhile, relatively little attention has been given to another key aspect of the AI debate: How AI might contribute to - or help avert - climate change, resource depletion and environmental degradation in the future?

A small but growing body of research seeks to explore this question, and others. What are the environmental impacts of using, training and generating new AI models? Where does the majority of their environmental footprint come from? How sustainable is the hardware behind AI? Is it possible to use AI in an environmentally sustainable - or even beneficial - way?

Here we explore all of these questions, examining some of the current research in this area and making six key policy recommendations for promoting environmentally responsible use of AI.

Intensive schooling: the carbon costs of training AI

Building an AI model requires monumental computing power - and, therefore, energy use - to generate and train the new system. The computing power demand is largely driven by three things: the size of the model (in terms of parameters, the number of variables, or weights, that the model adjusts during training), the size of the training datasets, and tuning the hyperparameters of the model; with the latter often going underreported. So how exactly can you measure the environmental impact of creating a new AI model?

Strubell [27] devised a useful methodology for measuring the carbon emissions associated with training an AI model based on the power draw of the hardware used and the hours required. We applied their methodology to data from other sources to estimate the carbon emission for additional AI models, in order to give a broader picture of the range of emissions typically generated by training AI. Table 1 shows their results alongside our own additions. Note that these are estimates only, due to the difficulty of obtaining precise information on some of these models.

Emissions output is measured in kgCO₂e, or kilograms of Carbon Dioxide Equivalent, a unit used to measure carbon footprints in terms of the amount of CO₂ that would create the same level of global warming. [4]

Model	Hardware	Power (W)	Hours	kWh.PUE	CO2e (kg)
T2Tbig	P100x8	1515	84	201	87
ELMO	P100x3	518	336	275	119
BERTbase (V100)	V100x64	12,042	79	1507	652
BERTbase (TPU)	TPUv2x16	4000	96	607	261
BERTlarge	TPUv3x64	12,800	96	1941	840
NAS	P100x8	1515	274,120	656,347	284,000
NAS	TPUv2	250	32,623	12,900	5580
GPT-2	TPUv3x32	6400	168	1700	735
GPT-3	V100	300	3,100,000	1,474,000	638,000

Table 1 - Power requirements and CO2e emissions for AI models. Based on data and methodology from [27], with additional data from [17], [26], [28] and [22].

So what is the environmental impact of training AI? As illustrated above, there's no simple answer, with emissions ranging from 87kgCO2e (roughly equivalent to driving 220 miles in the average car [32]) to 638,000kgCO2e (equivalent to one person flying 10,500 miles from London to Sydney almost 142 times [20]).

To put this latter figure into further context, the average human is responsible for 5,000kgCO2e in a single year [24] - meaning it would take the average person over 127 years to generate the same level of emissions that it took to train GPT-3.

So why is there so much variation in energy consumption between different AI models?

The level of emissions produced by the AI training process is dependent on four key factors. The first and most obvious factor being the individual computational requirement of the individual system - the size of the model and of the training dataset. This can range from BERTbase, which took a mere 79 hours to train, to GPT-3, which required a whopping 3,100,000 hours - a little over 350 years - of total computing time. In the next section we'll delve more deeply into the reasons behind this and the trends that are emerging in computational demands.

The second key factor determining emissions from AI training is the Power Unit Equivalent (PUE), which is the effective power required for a server centre to produce one unit of computing power, due to additional power draws such as cooling. Strubell [27] estimates this at 1.58, but it's possible to reduce this by improving the efficiency and management of server centres. In fact, AI itself could be one tool for improving efficiency - more on this later.

The third factor is the power draw of the hardware itself. If we compare the results of the BERTbase model to the BERTlarge model, the latter has just over three times the power use of the former, despite using four times the TPUs (Tensor Processing Units [34]). This is due to the latter using a newer version (TPUv3) with a lower power draw, illustrating how improvements in hardware can also reduce emissions from AI.

The final major factor is the carbon intensity of the grid. In his research, Strubell [27] assumes a carbon intensity of 433gCO₂/kWh based on 2018 estimates from the EPA. However, 2022 US grid carbon intensity was 376g/kWh [31], which would produce lower emissions than Strubell's estimates. Other locations would be even more efficient. For example, the EU's average grid carbon intensity in 2022 was 250g/kWh [7], whilst the UK's was 182g/kWh [21]. Hence, running an AI model on servers in the UK rather than the US could halve the carbon emissions from the model.

So, taking these factors into account, is it possible to train AI in an environmentally responsible way?

Fig. 1 and Table 2 compare the real-world emissions estimates above to the hypothetical emissions that would have occurred if PUE were reduced to 1.25 and if grid carbon were equivalent to that of the UK in 2022. Together, these data centre efficiency and carbon grid intensity measures would reduce emissions in each case by more than 50%.

Model	CO2e (kg) base	CO2e (kg) improved PUE	CO2e (kg) UK grid carbon	CO2e (kg) both
T2Tbig	87	68.73	36.54	28.8666
ELMO	119	94.01	49.98	39.4842
BERTbase	652	515.08	273.84	216.3336
BERTbase	261	206.19	109.62	86.5998
BERTlarge	840	663.6	352.8	278.712
NAS	284,000	224,360	119,280	94231.2
NAS	5580	4408.2	2343.6	1851.444
GPT-2	735	580.65	308.7	243.873
GPT-3	638,000	504,020	267,960	211,688.4

Table 2 - How reducing PUE (in this case from 1.58 to 1.25) and grid carbon (from 433gCO₂/kWh to 182gCO₂/kWh) affect emissions from AI.

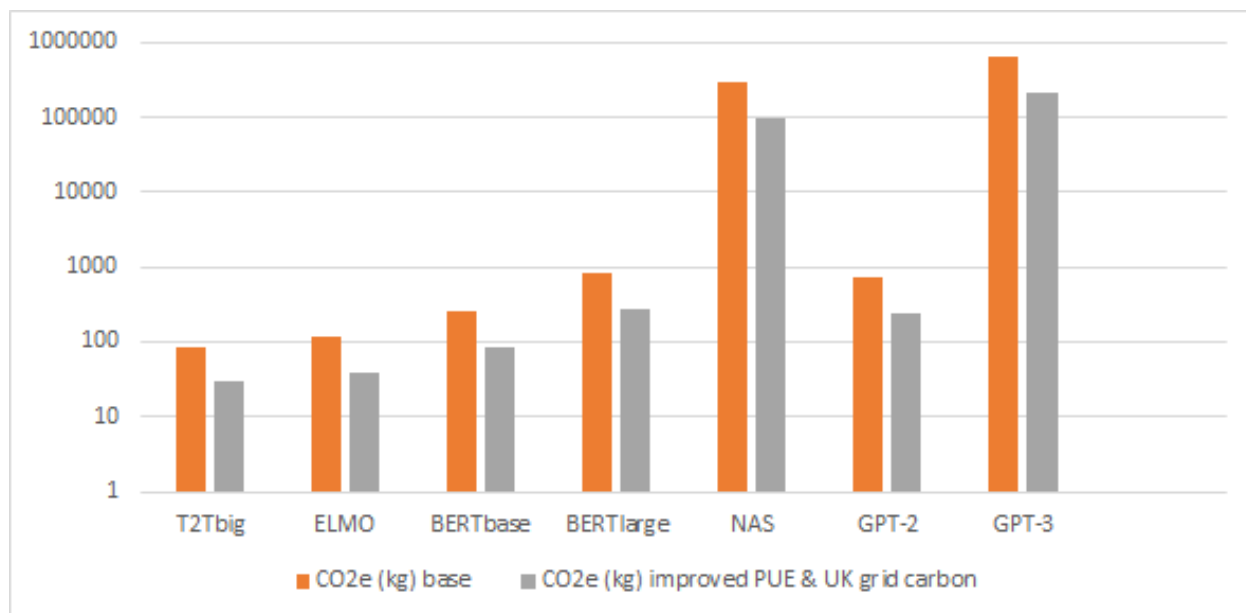


Figure 1 - Reducing PUE and grid carbon can drastically cut the emissions from AI, but large models still dwarf the emissions of smaller ones. The chart is on a logarithmic scale.

However, although that 50% reduction illustrates the value of improving data centre efficiency and grid carbon intensity, we can also see that it would still be far outweighed by the huge increase in emissions between GPT-2 and its subsequent iteration GPT-3. This is indicative of a wider trend in which small gains in efficiency are far outstripped by a shift towards more complex AI systems requiring more energy to run.

Insatiable demand: the exponential growth of computational power behind AI

It is clear that AI models generate vastly different amounts of emissions, and that some of the larger systems have a significant environmental footprint. But even more concerning is the astounding rate at which AI emissions as a whole are increasing. This is reflective of an overall trend towards larger and ever more complex models.

When OpenAI launched GPT-2 in 2019, the system was responsible for an estimated 735kgCO₂e – slightly less than an average-sized dog generates in a year [11]. But just a few months later, GPT-3 launched with estimated training emissions of 638,000kgCO₂e – far more than the average human generates in a lifetime. That's an increase of almost 87,000% - and this trend towards growth seems set to continue.

Information for GPT-4 is difficult to find at present, but it is indicated to have ten times the parameters of GPT-3 (1.8T compared to 175B), so we could reasonably expect another order of magnitude increase, which would take GPT-4 emissions into millions of kgCO₂e.

OpenAI is far from unique in this aspect. Fig. 1 shows the increase in computing power for a variety of Machine Learning models from 2012 to 2018. In total there was a 300,000 times increase in computing power used by these models over the six years.




Figure 2 - Increase in computing power for models, on a logarithmic scale [1].

Driving this growth in compute resources is a quest for ever greater accuracy, requiring increasingly complex models, which can improve existing use cases of AI and potentially open up new ones.

But is this greater accuracy enough to justify the environmental cost? Several papers have delved into the issues surrounding AI's increase in complexity [3, 26,27], with broadly similar findings. Improving the accuracy of a model is almost always a case of diminishing returns, with incremental increases in accuracy requiring exponential increases in computing power - and not necessarily improving user outcomes. In other words: when it comes to AI, bigger is not always better.

So if more complex models aren't necessarily delivering any more value, why are we still seeing such an obsession with accuracy? The issue here is one of culture. The current culture around AI is focused on accuracy of models, often ignoring other issues such as cost, efficiency and required computing power. In some cases, this might be out of a genuine scientific desire to push the limits of what AI

can do. But in other cases, it's a matter of organisations competing for prestige, attention and ultimately, business – improving accuracy is widely regarded as impressive in a way that improving efficiency usually isn't.

Schwartz [26] describes this pervasive attitude as 'Red AI', which simply seeks to improve accuracy regardless of the massive computational - and therefore environmental - cost. Schwartz instead encourages a focus on 'Green AI' research, aiming to yield novel results whilst taking into account computational cost.

In order to measure the efficiency of an AI model - and therefore how 'Red' or 'Green' it might be - Schwartz proposes examining the total number of floating point operations [] necessary to produce a result. This method directly computes the amount of work done and is hardware agnostic. This approach is therefore useful for comparing and designing the most efficient AI models and algorithms. Bender and Strubell [3, 27] also highlight the need for a greater focus on efficiency and computational cost. In addition, they suggest that retraining time and sensitivity to hyperparameters should be reported more widely in AI research, to make the environmental costs more transparent.

As the size of AI models increases, so does the need for these kinds of research measures and reporting standards in order to understand and mitigate the environmental impact of this technology. But at the same time as this research is becoming more crucial, it's also becoming more challenging. With the increased size of AI models comes increased running costs. Both Schwartz and Strubell [26,27] highlight how increased costs can create a barrier to academics seeking to conduct AI research and to reproducing the results of such large models, thus decreasing the transparency of such AI. Creating more equitable access to hardware and data resources could help to alleviate this and open up the cultural conversation around AI's environmental impact - a conversation that is more relevant than ever as AI becomes embedded in our daily lives.

It all adds up: emissions from AI use

So far, we have focused primarily on the impacts of training AI models, as this is reflective of where most of the literature has been focused to date. However, it is important to consider the use of AI. Single instances of AI use - known as 'inferences' - may individually have far smaller costs than training the model, but over time this can add up to be more than the initial training costs. Some studies estimate that 30-65% of emissions from AI could be attributed to inference [23, 35], which is likely to rise as AI is more widely used.




Figure 3: Tasks examined in [18]'s study, and the emissions they produce per 1,000 queries. The y-axis is logarithmic.

So when it comes to environmental costs, are all inferences alike? To answer this question, Luccioni [18] has taken a deep dive into AI inference, examining the differences in energy use between different tasks and between AI models, comparing both 'task specific' and 'general purpose'. Their findings summarised in Fig. 2 illustrate that different AI tasks use significantly different resources.

Two main takeaways are that generation tasks - i.e. having the AI generate something, such as images or text - cause significantly more emissions than classification-based tasks, and image-related tasks cause considerably more emissions than non-image based tasks. Image generation created far higher emissions than any other task studied. There is already significant debate around AI art and its ramifications for artists and copyright, and we can see above that its carbon footprint is another cause for concern. When it comes to AI's impact on the environment, it seems a picture really is worth a thousand words... if not more.




Figure 4: Mean model emissions for classification tasks for different datasets [18], showing the comparison between task-specific models and multi-purpose architectures. The y axis is in logarithmic scale, dot size is proportional to model size.

The study [18] also examined the differences between task-specific and general-purpose models, finding that the emissions from general-purpose AI are an order of magnitude higher than for task-specific AI when running for classification tasks. They find that for generative tasks there is a similar though smaller difference in emissions, with general-purpose AI generating 3 to 5 times more emissions than task-specific AI.

These findings would suggest that focusing on developing task-specific rather than general-purpose AI would be beneficial in reducing overall emissions from AI use. However, it may not be quite so simple. Other research [23] has suggested that large models which don't need to be retrained for different tasks - thus saving on the retraining costs - could help reduce emissions. More research into the tradeoffs between training, retraining and inference costs could be invaluable in finding the most energy-efficient solution.

Of course, the use of AI wouldn't be possible without hardware: servers, processing units and devices such as Alexa. Manufacturing this hardware typically requires rare earth metals obtained in labour-intensive mining processes which can have significant impacts on the local environment. Crawford & Joler [6] indicate that the mining of minerals and metals for AI can cause loss of biodiversity, contamination of soil and water, deforestation, erosion and high concentration of polluting substances.

We also see the intersection of environmental and sociopolitical issues, with resource extraction having knock-on impacts such as low-paid, dangerous labour and negative health impacts on the local population, often in poorer parts of the world.

It is, however, incredibly difficult to estimate the true impact of specific AI programs and devices due to the layers of supply chains and systems between the extracted resources and the finished product. And here we start to see another issue: that of complexity and lack of transparency.

Obscurity through complexity: the opacity issue

As we saw earlier when estimating the carbon emissions impact of AI programs, it was difficult to find precise information on the computing resources and hardware that went into these programs. This was also an issue when discussing the hardware resources that underpin AI. This highlights another issue with AI: lack of transparency and complexity, which together can make it difficult to examine the impacts of AI.

This opacity is partially due to a deliberate lack of transparency on the part of companies that make AI and other ICT technologies. In most cases, there is no incentive for companies to be upfront about their compute and energy use. If their energy use is seen as something that may put customers off they may even be incentivised to hide it to protect profits.

The lack of clarity is also due to the complexity of the supply chains that underpin AI and ICT products. This has been explored by studies such as Graham and Haarstad [10], pointing out that as commodities become more globalised, the production processes are broken into a complex set of supply chains and networks with geographically diverse nodes. Such complexity obscures the impacts of products from customers and makes it difficult even for customer watchdog organisations to perform analysis on commodity chains.

Crawford & Joler [6] also highlight this as an issue, particularly in metal supply chains, which they describe as a “zooming fractal of tens of thousands of suppliers, millions of km of shipped materials and hundreds of thousands of workers.”

This complexity means that it’s not only costly and time-consuming for a company to attempt to ensure they are ethically sourcing materials - it can also make it incredibly difficult. Global ICT company Intel ran into this problem in 2009 when they set out to ensure that all their supplies of minerals such as tantalum, tin, tungsten, gold and cobalt were sourced from ‘conflict-free’ sources. It wasn’t until 2013, four years later, that Intel finally understood its own suppliers and supply chains well enough to certify that all minerals used in their products came from ‘conflict-free’ sources [14].

That four year timeframe was for a company that was actively trying to understand and improve its supply chain from within. How much harder and more time-consuming would it be for external organisations, researchers, journalists or consumers to untangle these supply chains and understand what goes into products such as AI, especially in cases where companies are deliberately exhibiting a lack of transparency?

Cloud computing providers are particularly bad at providing carbon transparency. Whilst most now publish some metrics, the numbers they provide are often difficult to compare across providers and are impossible to verify. They also typically only consider the carbon emissions from a subset of sources such as energy use, whereas correctly accounting for environmental sustainability requires considering a broader range of issues. Initiatives such as the Circular Data Centre Compass [37] aim to remedy this, but more engagement from cloud providers is required.

In order to truly move towards environmentally sustainable AI, there will need to be far greater levels of openness and transparency from AI and ICT companies, sparked by either demand from consumers or - more likely - by legislation and enforcement.

How we use it: applications of AI

When it comes to AI inference and its potential impact on the environment, emissions from operational costs are not the whole story. It is also important to consider the ways in which AI technology can be utilised to directly help or harm the environment. Even if an AI is carbon-neutral, the ways in which it is being used may not be environmentally responsible.

How can we use AI to benefit the environment? Research has shown that AI can be used to save energy by improving the efficiency of systems, enabling them to derive more output from fewer resources. For example, testing by Google indicates that turning cooling control in their data centres over to AI could reduce energy use by 40% [15], whilst the IEA finds that digital technologies (including AI) could save energy in transport and buildings [13].

AI can also help the environment in other ways. For example, the World Bee Project is utilising AI to help reverse the decline in the bee population [19].

On the other hand, AI has the potential to be used with environmentally devastating consequences. Studies have shown that using AI to boost production could result in unsustainable consumption of resources, and could result in biodiversity loss if used to maximise agricultural yields without taking negative externalities into account [16]. AI could also be used to increase the production of oil and gas, and the amount of recoverable reserves [13], which would conflict with vital efforts to phase out fossil fuels needed to avert severe climate change, as highlighted at the recent CoP 28.

The potential for AI to harm the environment suggests that policy change will be needed to prevent it, but so far this hasn't been widely reflected in AI law. The EU's proposed AI act, for example, has several policies to minimise the use of AI in socially harmful ways or in ways that impinge on personal privacy and freedoms, but almost nothing to prevent AI being used in ways that impact the environment - despite one of their listed priorities being to ensure that AI used in the EU is environmentally friendly. This is an oversight that other jurisdictions seeking to enact AI laws can rectify.

At the government, organisational and individual level, it is important to consider how AI could be used to help and harm the environment, and to bring both aspects into the wider cultural debate around ethical use of AI.

No silver bullets: AI and renewables

AI's emissions come primarily through electricity use, which we can 'easily' decarbonise, so do they matter? If a gradual global shift towards renewable energy means that AI's power demand could one day be entirely met by renewables, is it worth worrying about AI's carbon footprint today?

Whilst we could (and should) meet electricity demand from renewables as soon as we can, it does not change the fact that currently there is still a significant amount of carbon in the electricity mix, and thus AI emissions will remain high in the short term. Working towards carbon-neutral AI does not excuse lack of action in dealing with today's problems.

Firstly, decarbonising the grid requires big investments in renewable capacity in order to replace fossil-fuel-based generation. This becomes significantly more difficult if overall demand continues to increase. China, for example, is a world leader in installed renewable capacity [12], yet its emissions have continued to grow almost every year [5] as demand continues to increase due to a growing economy and a more affluent populace. Growing demand makes it harder to decarbonise. It's crucial, therefore, to promote demand-side reduction in order to facilitate supply-side decarbonisation, and this should be strongly considered as part of the debate around AI.

Secondly, the move towards renewables has its own problems with decarbonisation. Manufacturing of batteries and solar cells consumes significant electricity and hence has an associated emissions cost whilst the grid still uses non-renewable sources. These devices also typically consume materials like lithium, the mining of which has negative local environmental effects and direct emissions as mentioned above [36]. So even if AI training and use consumed only energy from renewable sources, there would still be significant carbon emissions in the supply chain to consider.

How to improve: policy recommendations

Here we have examined the potential environmental impacts of AI, but how can we improve these things? Based on the issues we've discussed and recommendations from the literature, we have proposed a set of policy interventions that could help minimise the environmental impacts of AI.

1. Increase energy efficiency

Firstly, there is a need to increase energy efficiency - not just in the design of data centres, processors and other hardware underpinning AI, but also in the design of AI algorithms and programs. This will necessitate promoting a culture that prioritises efficiency, rather than just building the biggest and most accurate models possible. In other words, encouraging a culture of 'Green AI' over 'Red AI'. This may also require running AI programs on processors in locations and data centres that have more efficient running standards and higher proportions of renewable electricity.

2. Consider the entire tech ecosystem

Secondly, consider the entire tech ecosystem and all the systems that underpin AI and ICT hardware, promote ethical design standards which aim to minimise the impacts on the environment that these systems have, and where possible, promote the sourcing of hardware and materials from ethical sources.

3. Mandate transparency and choice

Thirdly, mandate transparency and accountability. Promote a culture of openness and responsibility, encourage and mandate more reporting of how, where and when AI is trained and hardware used in order to ensure that customers are as informed as possible.

Use procurement levers to disadvantage companies that do not do this, and work with social responsibility initiatives like the "B Corporation" mark to ensure environmentally sustainable AI use is included in their criteria.

Transparency alone will not be enough, there is also a need to ensure that customers have real choice (to move away from unethical companies) and methods to hold corporations to account for harmful practices (most likely through legislation).

4. Curb environmentally harmful AI use

Fourthly, curb the use of AI in environmentally harmful practices, for example, discourage the use of AI in maximising resource extraction. Ensure that where AI is used to maximise production and resource extraction, it accounts for negative externalities and does not simply seek to maximise efficiency or output regardless of negative external impacts.

5. Integrate climate policy into AI policy

There is a need to better integrate climate policy into all aspects of policy making, in this case tech and AI policy. Too often climate and environment is seen as a separate policy area that policy makers in other domains do not need to worry about. Instead, we need to promote the consideration of climate and environmental concerns in all aspects of policy making.

One example would be the Wellbeing of Future Generations (Wales) Act [9], which requires public bodies to consider the long-term impact of their decisions to prevent problems including climate change.

6. Improve public awareness

Ensure that the public are informed about the wider issues with AI, and how their use of AI can impact the environment. Educate around how different AI tasks can result in widely different related emissions. Indicate which companies are meeting ethical or environmental guidelines around AI, and make this information widely available.

Summary

AI is a rapidly developing and increasingly widespread technology which will impact both our lives and our planet. Although it is currently a small part of the overall conversation around AI, there have been a number of papers and articles discussing the potential impacts of AI on the environment and climate change.

Training some of the larger AI models can equate to more emissions than the average person generates in a lifetime, and this is a growing issue - we estimate that training GPT-3 caused almost 1000 times more emissions than GPT-2, and GPT-4 is likely to be even higher still.

AI models have been growing considerably in their computing power, growing 300,000 times larger in the six years from 2012-18, and therefore increasing electricity demand from data centres and ICT. The culture around AI is very much focused on 'bigger is better', getting the best accuracy at high computing cost. It would be beneficial to encourage a cultural shift towards developing more efficient models and algorithms.

Inference of AI models, when run millions of times, can also add up to cause substantial emissions - equalling and even exceeding the emissions from training the models. Image generation is by far the worst offender, generating an order of magnitude more emissions than other tasks. Task-specific models use less energy, but more research is needed to determine whether this saving is cancelled out by the additional training required. Resource use is also an issue, with AI hardware requiring minerals and metals that are often mined in environmentally (and socially) damaging ways.

Determining emissions from AI can be tricky due to difficulty in finding information on the exact hardware and the electricity mix behind AI models. Similarly, finding information on what is underpinning AI models and systems can be difficult, due to both a lack of transparency around AI and the inherent complexity of the global supply chains. This makes it difficult to estimate the impacts of AI and to identify AI systems and products that have a greater or lesser impact on the environment, making ethical consumer choice difficult.

Whilst in the future, AI's growing electrical demand could be met from renewable sources, having an increasingly growing electricity demand will make grid decarbonisation significantly harder.

We proposed a number of policy recommendations, including:

- Promote a culture shift towards a focus on efficiency in AI models, as well as promoting energy efficiency in AI hardware.
- Consider the whole tech ecosystem and encourage sourcing the resources and products that underpin AI from ethical sources.
- Mandate transparency and real accountability in AI and ensure real choice for customers.
- Curb the use of AI in environmentally harmful practices, such as fossil fuel extraction.
- Better integrate climate policy into all aspects of policy making, including AI.
- Improve public awareness of the environmental issues around AI and which products are harmful.

References:

1. Amodei, Dario & Hernandez, Danny, AI and compute, 2018, OpenAI, <https://openai.com/research/ai-and-compute>
2. Bannour, Nesrine et.al.; Evaluating the carbon footprint of NLP methods: a survey and analysis of existing tools, 2021, Association for Computational Linguistics, <https://aclanthology.org/2021.sustainlp-1.2.pdf>
3. Bender, Emily M. et. al., On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?, 2021, Association for Computing Machinery, <https://dl.acm.org/doi/10.1145/3442188.3445922>
4. What does kg CO2e mean?, Carbon Cloud Community Knowledge, <https://knowledgebase.carboncloud.com/-kg-co2e> (accessed 2024)
5. Historical GHG Emissions, ClimateWatch, https://www.climatewatchdata.org/ghg-emissions?end_year=2020&gases=co2®ions=CHN§ors=energy&start_year=1990 (accessed 2024)
6. Crawford, Kate & Joler, Vladan, Anatomy of an AI System, 2018, <https://anatomyof.ai/>
7. Greenhouse gas emission intensity of electricity generation in Europe, 2023, European Environment Agency, <https://www.eea.europa.eu/en/analysis/indicators/greenhouse-gas-emission-intensity-of-1?activeAccordion=546a7c35-9188-4d23-94ee-005d97c26f2b>
8. EU AI Act: first regulation on artificial intelligence, 2023, European Parliament, <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
9. Well-being of Future Generations (Wales) Act 2015, 2023, Future Generations Commissioner for Wales, <https://www.futuregenerations.wales/about-us/future-generations-act/>
10. Graham, Mark & Haarstad, Håvard, Transparency and Development: Ethical Consumption Through Web 2.0 and the Internet of Things, 2011, Information Technologies & International Development Journal, <https://itidjournal.org/index.php/itid/article/download/693/693-1851-2-PB.pdf>
11. Han, Jessica, The Environmental Impact of Pets: Working Towards Sustainable Pet Ownership, 2023, earth.org, <https://earth.org/environmental-impact-of-pets/>
12. Howe, Colleen, Explainer: The numbers behind China's renewable energy boom, 2023, Reuters, <https://www.reuters.com/sustainability/climate-energy/numbers-behind-chinas-renewable-energy-boom-2023-11-15/>

13. Digitisation and Energy, 2017, International Energy Agency, <https://iea.blob.core.windows.net/assets/b1e6600c-4e40-4d9c-809d-1d1724c763d5/DigitalizationandEnergy3.pdf>
14. Intel's Efforts to Achieve a Responsibly Sourced Mineral Supply Chain, 2022, intel, <https://www.intel.com/content/dam/www/central-libraries/us/en/documents/2023-05/responsible-minerals-white-paper.pdf>
15. Jones, Nicola, How to stop data centres gobbling up the world's electricity, 2018, nature, <http://complexityexplorer.s3.amazonaws.com/Computation+in+CS/SFI+1.4b.pdf>
16. Kanungo, Alokya, The Green Dilemma: Can AI Fulfill Its Potential Without Harming the Environment?, 2023, earth.org, <https://earth.org/the-green-dilemma-can-ai-fulfil-its-potential-without-harming-the-environment/>
17. Li, Chuan, OpenAI's GPT-3 Language Model: A Technical Overview, 2020, Lambda, <https://lambdalabs.com/blog/demystifying-gpt-3>
18. Luccioni, Alexandra Sasha, Jernite, Yacine & Strubell, Emma, Power Hungry Processing: Watts Driving the Cost of AI Deployment?, 2023, arXiv, <https://arxiv.org/pdf/2311.16863.pdf>
19. Marr, Bernard, How Artificial Intelligence, IoT And Big Data Can Save The Bees, 2021, <https://bernardmarr.com/how-artificial-intelligence-iot-and-big-data-can-save-the-bees/>
20. Calculate your emissions, myclimate, https://co2.myclimate.org/en/calculate_emissions (accessed 2024)
21. Britain's Electricity Explained: 2022 Review, 2023, NationalGrid, <https://www.nationalgrideso.com/news/britains-electricity-explained-2022-review>
22. NVIDIA Tesla V100 PCIe 16 GB, TechPowerUp, <https://www.techpowerup.com/gpu-specs/tesla-v100-pcie-16-gb.c2957>
23. Patterson, David et. al., The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink, 2022, arXiv, <https://arxiv.org/ftp/arxiv/papers/2204/2204.05149.pdf>
24. Ritchie, Hannah & Roser, Max, CO2 emissions, 2023, Our World in Data, <https://ourworldindata.org/co2-emissions>
25. Schlossberg, Tatiana, Flying Is Bad for the Planet. You Can Help Make It Better, 2017, New York Times, <https://www.nytimes.com/2017/07/27/climate/airplane-pollution-global-warming.html>
26. Schwartz, Roy, Dodge, Jesse, Smith, Noah A. & Etzioni, Oren, Green AI, 2020, Communications of the ACM Vol. 63 No. 12, <https://dl.acm.org/doi/pdf/10.1145/3381831>

27. Strubell, Emma, Ganesh, Ananya & McCallum, Andrew, Energy and Policy Considerations for Deep Learning in NLP, 2019, Association for Computational Linguistics, <https://aclanthology.org/P19-1355.pdf>
28. Teich, Paul, Tearing Apart Google's TPU 3.0 AI Coprocessor, 2018, The Next Platform, <https://www.nextplatform.com/2018/05/10/tearing-apart-googles-tpu-3-0-ai-coprocessor/>
29. Thomas, Samuel, AI revolution: what's the environmental impact?, 2023, Schroders, <https://www.schroders.com/en-gb/uk/intermediary/insights/ai-revolution-what-s-the-environmental-impact-/>
30. Energy Consumption of ICT, 2022, UK Parliament POST, <https://researchbriefings.files.parliament.uk/documents/POST-PN-0677/POST-PN-0677.pdf>
31. U.S. Energy-Related Carbon Dioxide Emissions, 2022, 2022, U.S. Energy Information Administration, <https://www.eia.gov/environment/emissions/carbon/>
32. Greenhouse Gas Emissions from a Typical Passenger Vehicle, 2023, U.S. Environmental Protection Agency, <https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle#driving>
33. FLOPS, 2024, wikipedia.org, <https://en.wikipedia.org/wiki/FLOPS>
34. Tensor Processing Unit, 2024, wikipedia.org, https://en.wikipedia.org/wiki/Tensor_Processing_Unit
35. Wu, Carole-Jean et. al., Sustainable AI: Environmental Implications, Challenges and Opportunities, 2022, arXiv, <https://arxiv.org/pdf/2111.00364.pdf>
36. Melin, Hans Eric, Circular Energy Storage, 2019, [Analysis of the climate impact of lithium-ion batteries and how to measure it.](#)
37. Andrews, Deborah, Circularity and Whole Systems Thinking, 2023. https://static.sched.com/hosted_files/stateofopencon2023/4e/D-Andrews-OpenCon-Slides.pdf and [CEDaCI Compass](#)

REGISTER DYNAMICS

Register Dynamics is a team of technical, data and business specialists working with public and private organisations to support them with their technical and data challenges. We are an independent data and technology consultancy based in the UK.

Our specialisms are in emerging technologies and our focus is in supporting organisations to drive their data agenda forwards. This includes implementing data and analytics strategies to boost data intelligence, building data capability that drives business value, establishing good data governance and nurturing a data-driven culture.

register-dynamics.co.uk | [@RegDyn](https://twitter.com/RegDyn)



About Owain Jones

Owain is a skilled data architect and technology with a passion for finding new ways to generate and use data to provide meaningful insights for business. He is an expert in data migration, the standardisation of data from various sources, database creation, data modelling and AI.

[Learn more about Owain](#)