Regulating AI for Climate and Nature

On 18 March, the Centre for Climate Engagement hosted an event on Regulating AI for Climate and Nature, as part of the Cambridge Festival, Green Week at Hughes Hall and Cambridge Climate Week.

The event featured a keynote from Yu-Ting Kuo, Corporate Vice President at Microsoft, a presentation from Nick Scott, Manager of the Law for Climate Action Programme at CCE, and a panel discussion chaired by Professor Harro Van Asselt, Hatton Professor of Climate Law, featuring the following speakers:

- Lynn Dicks, Professor of Ecology at the Department of Zoology and the Conservation Research Institute.
- Dr Loïc Lannelongue, Assistant Research Professor at the Department of Public Health and Primary Care, and PI of the Cambridge Sustainable Computing Lab.
- Guilia Trojano, Senior Associate and AI expert at Hausfeld.
- Richard Turner, Professor of Machine Learning in the Department of Engineering at the University of Cambridge and Research Lead for AI for Weather Prediction at the Alan Turing Institute.

The CCE would like to thank our event partners Cambridge Zero, the Conservation Research Institute, Cambridge Institute of Sustainability Leadership, and the Cambridge Climate Society for their support in delivering this insightful discussion, takeaways of which are outlined below.

Industry Perspective: the AI-environment accountability gap

The rapid pace of AI development, and the growing dominance of major technology companies driving it, has brought the industry's environmental impacts into sharper focus. The material costs of AI, including energy consumption, natural resource use, and wider environmental externalities, have historically been outside the mainstream of both technical education and industry practice. In this context, bringing together expertise from across law, technology, policy, and the environmental sciences to consider how this AI-environment nexus can be better governed is both timely and necessary.

AI's environmental footprint

Understanding the full environmental impact of AI requires looking beyond carbon emissions alone. The hardware that underpins AI systems carries what is often termed an embodied footprint, which encompasses the extraction of raw materials, manufacturing processes, electronic waste, and transport emissions – and tends to account for the bulk of AI's overall environmental impact. Despite this, it receives comparatively limited attention in public debate, where focus centres more heavily on technologies' operational footprint: the energy consumed in running data centres. Calculations of an operational footprint are themselves highly variable, depending significantly on the proportion of renewable energy in a given country's grid.

Water is a critical dimension of AI's environmental impact. Data centres require water both indirectly though the electricity needed to power them (much of which depends on water in its generation) and directly for cooling. This presents a trade-off: low-power cooling systems tend to be more water-intensive, while low-water approaches typically consume more electricity. Technical solutions continue to be explored, but choices between them involve environmental trade-offs which are yet to receive mainstream focus in debates on AI's sustainability.

The local dimension of this has direct implications for how regulation should be designed and enforced. The environmental costs of AI, whether from resource extraction, water consumption or energy use, are felt by specific communities, ecosystems, and regions where infrastructure is built. Rather than relying solely on national or supranational frameworks, effective oversight may require empowering regional and municipal authorities to set binding standards on resource use, emissions, and land impact. This may also help to ensure that affected communities have a meaningful voice in regulatory processes, and that compliance is not assessed purely through corporate disclosures.

A more comprehensive accounting of AI's impacts across its lifecycle is a necessary foundation for any effective regulatory response. At present, the tendency to focus narrowly on carbon emissions specifically in the operational phase of these technologies means that a significant portion of environmental costs receives little scrutiny from regulators, businesses, and the public. Panellists explained that building lifecycle assessment into both AI regulation and related corporate reporting and auditing frameworks would enable more comprehensive scrutiny, and create stronger incentives for the industry to reduce environmental impacts at every stage.

Limitations of voluntary commitments

While some tech companies have made public environmental commitments, e.g. carbon neutrality by 2030, these targets remain largely voluntary and self-regulated. Reliance on self-regulation leaves ambition vulnerable to commercial pressures that could see goalposts shifted or transition efforts pushed back. Where climate targets have been set at a national level, governments risk falling short by placing too much reliance on major tech companies to fulfil their environmental commitments – promises that, without meaningful enforcement, carry no binding obligation to deliver.

The scale of the risk is significant. Yu-Ting Kuo presented the case of Taiwan Semiconductor Manufacturing Company (TSMC) to illustrate the extent to which the energy demands of AI-related hardware manufacturing are placing pressure on national infrastructure and energy systems. As of the end of 2023, TSMC consumed approximately 9% of Taiwan's totally energy production,[1] with official forecasts suggesting this could rise to 24% by 2030.[2] This tension between industrial growth and environmental sustainability reflects a broader pattern for governments worldwide. Similar estimates exist for the US, with projections suggesting data centre energy consumption could rise to 13-27% of total domestic electricity by 2028.[3] Where industrial growth and technological leadership are national priorities – as in the EU[4] and the UK[5] – energy sustainability can be crowded out by pressure to attract investment and maintain market competitiveness. The practical response to mitigating and/or preventing this involves developing regulatory governance frameworks that specify what must be measured and reported – not to slow innovation, but to ensure businesses have the accountability structures they need to act with greater strategic clarity and consistency.

Communicating AI's impacts

One of the barriers to public understanding of AI's environmental impact is the abstract language often used to describe it. Yu-Ting Kuo used the analogy of "cloud computing" as an example – language that implies something intangible, obscuring the physical infrastructure beneath it. This is seen as a contributing factor to the lack of scrutiny of environmental trade-offs, reducing the pressure on companies to address them. Those working within the technology sector have a responsibility to communicate these trade-offs more transparently, not only to the public, but in how they factor environmental impacts into decisions on technology development and deployment. This is not yet common practice, and closing the gap will require a cultural shift as much as regulatory reform.

Understanding the AI-environment regulatory landscape

There is a wide range of ways in which the law is already governing at the intersection of AI and the environment, yet significant regulatory gaps remain. Academic literature specifically addressing AI regulation and climate change is relatively limited – a gap the CCE's discussion paper, *Regulating the AI-Climate Nexus,*[6] looks to address. In it, the CCE provides analysis of existing and proposed legislation. Author Nick Scott outlined key insights from the report across four broad categories of legislation at the AI-climate nexus:

- Legislation which specifically addresses the environmental impacts of AI
- Provisions in AI-focused legislation which address climate change or related issues
- Provisions in climate legislation which address AI.
- Legislation which indirectly addresses the AI-climate nexus

Across these categories, 'AI' covers a wide range of technologies with a wide range of potential environmental impacts. This diversity is not only a classification issue, but also a central challenge to designing regulation that is both comprehensive and precise. A rule designed for large-scale data centres, for instance, may be inadequate for the AI systems embedded in consumer devices, and vice versa. In this, there is an argument for effective regulation to be granular by design, in order to distinguish between specific AI use cases.

In terms of existing evidence, in some jurisdictions, data centres are subject to specific legal or policy measures, such as water efficiency requirements and renewable energy thresholds, as seen in a Special Action Plan in China[7] and proposed legislation in multiple US states. Planning, permitting, and environmental impact assessment processes have also been tested. In Ireland, litigants have pursued legal action to intervene in the planning approval of the multi-billion-euro Herbata data centre on the grounds that it "could breach Ireland's climate obligations"[8] and overpower the country's energy system. Analysis finds that the dozens of data centres surrounding Dublin are consuming more electricity than all of the urban homes in Ireland,[9] heightening fears of rolling blackouts, energy price increases, as well as significant disruption to national housebuilding targets. While this exemplifies litigants' growing willingness to test the boundaries of existing frameworks, litigation remains costly and reactive by nature, and questions remain around whether it is the most effective tool for closing the gap where legislation has not caught up.

Corporate responsibility and accountability

The mechanisms for holding both corporate providers and end user of AI products to account on their own reporting obligations remain inconsistent. Existing regulation such as the EU AI Act (Regulation (EU) 2024/1689)[10] – which is regarded as a leading example of AI policy that incorporates energy efficiency provisions – is yet to instil a whole-lifecycle approach.[11] This creates gaps in what can be regulated and measured at each stage of development and deployment, making it difficult to calculate environmental impacts consistently across a product's lifetime, and risks different methodologies being applied at different stages. Reporting obligations compound this problem – what must be reported, by who, why, and when remains uneven across jurisdictions, and moves such as the EU Omnibus Directive,[12] and the removal of a number of sustainability-related mandatory reporting requirements, risk deepening fragmentation further. The extent to which AI products and services and being used by oil and gas companies to expand extraction practices and enable fossil fuel production is another area of concern. The potential gap between stated values and real-world application reinforces the case for mandatory reporting requirements, specifically around the downstream uses of AI systems.

While still an emerging area, a more standardised framework for independent auditing of AI's environmental impacts was identified as a critical accountability mechanism. This model could involve:

- **Standardised metrics**: agreement on what needs to be measured and how,
- **Scope definitions**: clarity on which entities are captured, given that AI's supply chain spans hardware manufacturers, cloud providers, and end users across multiple jurisdictions,
- **Mandatory disclosure**: guidelines on companies' legal requirement to provide auditors with access to relevant data,
- **Accreditation of auditors**: a framework for who is qualified to conduct these audits and what standards they are held to,
- **Enforcement mechanisms**: consequences for non-compliance or misreporting and/or incentive structures.

Although enforceability of an AI-specific environmental auditing regime at the global level is unlikely, in practice, a combination of mechanisms could be leveraged, including stronger regional regulation, internationally agreed disclosure standards, and market pressure. The role private industry has in shaping this, however, would need to be more clearly defined.

The case for proactive regulation

There is a case for tackling many of these issues ex ante – through clear, enforceable regulation – rather than waiting for them to be exposed through litigation. A fundamental concern is existing regulatory and government structures' inability to keep up with the speed and scale of technological change. Fit-for-purpose regulation will require not only new rules and requirements, but faster and more adaptive mechanisms for responding to technological change.

Building awareness and seizing opportunities

Regulation is one component to improving scrutiny of AI development and deployment, and one on which other forms of accountability can build. Greater public awareness of the socioenvironmental impacts of AI could become a meaningful force, shaping consumer behaviour, informing democratic participation in technology rollout, and creating the kinds of informed pressure that regulation alone cannot generate.

Enabling environmentally responsible uses of AI

AI's potential as a tool for environmental good is an equally important part of the debate. The same capabilities that make these technologies resource-intensive also make them powerful instruments for understanding complex systems. Panellists spoke of the use of AI technologies in their research, highlighting the potential to process data with more sophistication, greater speed, and at greater scales.

A principal concern for the immense potential of these technologies is the concentration of power that a select number of companies have over it; the benefits of technologies best-placed to address environmental challenges are unlikely to be distributed equitably when controlled by select actors. In the same respect, democratisation of access to AI technologies presents its own challenges. From an energy sustainability perspective, the proliferation of individual, personalised AI models may in aggregate consume considerably more energy than the centralised supercomputing resources they replace.

Panellists suggested there is a case for developing a policy mechanism analogous to green credits for AI, a resource allocations framework that actively incentivises environmentally efficient applications, which may also help address concerns to do with impact data auditing. Machine learning teams generally tend to report on the compute used to train their best-performing models, rarely accounting for the testing iterations and experimental workloads that enable them, despite this being where a substantial proportion of compute is consumed. Mandating fuller reporting across the entire development process would at minimum improve the traceability of these methods and improve standards of practice around waste – which may be particularly beneficial for evidencing more environmentally responsible AI use cases. Panel experts agreed that expecting new and emerging technologies to be energy-proportionate is not a radical position, and could serve as a practical point of alignment between climate and AI legislation as these develop.

Data transparency

Assessing AI's environmental impact fundamentally depends on access to data. Large tech companies have conventionally reported aggregated impact data, offering limited basis on which to distinguish the impact of any specific activity or workload. Despite growing evidence of consumer interest in understanding, and potentially choosing between, providers on the basis of efficiency, data transparency remains a significant barrier. France offers an instructive example: its electronic communications regulator, Arcep, holds legal powers to determine which metrics certain tech companies are required to report on, including data on their environmental footprint.[13,14]

Industry-wide, the metrics that do exist are not without limitations. Panellists spoke of the EU Energy Efficiency Directive,[15] which – despite being one of the main signals of regulatory progress in this area – imposes no mandatory requirement for data centres to meet specific renewable energy targets or efficiency thresholds. The most widely used measure, Power Usage Effectiveness (PUE), which divides total power supplied to a facility by the power consumed by its computing equipment, is susceptible to distortion. A data centre can improve its PUE by adding more computers and ensuring they are running, rather than by reducing the energy its operations actually consume. Additionally, data centres are routinely built to a larger capacity than what is currently needed in order to be future-proofed, meaning that significant embedded and operational costs are being incurred ahead of demand. As a standalone measure of environmental performance, current usage of PUE as a metric therefore has clear limits, and is an insufficient basis on which to assess the true environmental cost of data centre operations.

Looking ahead, the rise of agentic AI systems – capable of acting autonomously and interacting with one another – is likely to drive an exponential increase in this consumption demand and related environmental footprint. Panellists cautioned that existing data reporting frameworks are not designed to anticipate or respond to this. Without transparency frameworks that are able to scale alongside this trajectory, regulatory responses risk being perpetually behind the curve.

AI's impact on policymaking processes

Policymakers are increasingly reliant on projections, models, and evidence that may themselves be generated by AI systems. The use of AI in generating evidence for policymaking introduces a number of risks. From an evidence synthesis perspective, chief among these are questions of bias, both in the data on which models are trained and in the outputs they produce, and a lack of transparency around the sources and processes that shape these outputs. On a practical level, AI technologies remain limited in their ability to generalise unseen conditions and project into the future, which has serious implications for climate science and the policy decisions that rely on it.

Panellists highlighted how the current tendency to embed AI across as many applications as possible, rather than deploying it more strategically, is problematic not only from a sustainability standpoint (i.e. leading to unnecessary compute, energy, and resources expenditure), but also in terms of outcomes. A more targeted AI policy approach would be to identify the societal problems that present the most immediate and/or severe risks in order to direct AI development towards them, rather than letting the technologies' development determine how and for what purpose they are used. It is also important to challenge assumptions of near-infinite data, compute, and storage, as these become increasingly difficult to sustain. The concentration of AI expertise has had a clear role in shaping the field's culture and priorities in ways that are not always well-suited to the broader societal contexts in which these technologies are now being deployed and embedded within. Meaningful reform to governance frameworks must move beyond regulating AI as it exists, towards actively shaping the conditions under which it is developed.

Conclusion

The relationship between AI and the environment is an evidently complex policy issue – the scale and scope of AI's environmental risks are exposing the limits of existing legal and regulatory frameworks; voluntary commitments are proving insufficient at preventing environmental externalities; and the development of technologies is outpacing the regulatory response. AI's expanding physical footprint is placing real pressure on energy systems, natural resources, and climate commitments at a time when the data needed to fully account for these impacts remains limited. Yet these technologies have genuine potential as tools for understanding and responding to environmental challenges. Capturing the opportunity, while holding its environmental consequences to account, will depend on greater cross-disciplinary engagement – from between AI developers and climate scientists, between technologists and policymakers, to between industry and the public. Despite the current fragmented legislative context, some patterns are beginning to take shape across jurisdictions, with multiple pieces of legislation introducing broader environmental principles embedded in AI legislation, procedural safeguards, and renewable energy requirements for data centres. What is needed may not necessarily be a wholesale new regime, but a more deliberate and joined-up approach to defining AI's environmental impact and standardising how it is measured, reported, and mitigated. There is shared optimism that stronger regulation could indeed drive innovation that leads to better outcomes for industry, for policymakers, and for nature and climate.

The CCE is committed to working with partners – researchers, practitioners, funders and networks – interested in contributing expertise, co-developing projects or exploring aligned opportunities in this area.

---

[1] 'TSMC to Consumer Three Nuclear Reactors' Worth of Power', CommonWealth Magazine (2024). Link.
[2] 'TSMC could account for 24% of Taiwan's electricity consumption by 2030', Data Centre Dynamics (2024). Link.
[3] 'How Much Electricity Does a Data Center Use? Complete 2025 Analysis', IAEI Magazine (2026). Link.
[4] The AI Continent Action Plan, European Commission (2025). Link.
[5] AI Growth Zones, UK Government (2025). Link.
[6] 'Regulating the AI-Climate Nexus: Trends, emerging issues, and ways forward', Centre for Climate Engagement (2025). Link.

[7] Global AI Governance Action Plan, Ministry of Foreign Affairs People's Republic of China (2025). Link
[8] 'ClientEarth lawyers intervene on Irish data centre that could send dangerous signal for Europe', ClientEarth (2025). Link.
[9] 'The AI race is already taking a toll. Irelan's massive data centers are a cautionary tale', Fast Company (2024). Link.
[10] Regulation (EU) 2024/1689 of the European Parliament, European Commission (2024). Link.
[11] 'The EU AI Act and environmental protection: the case for a missed opportunity', Heinrich Boll Stiftung (2024). Link.
[12] 'Omnibus Package', European Commission (2025). Link.
[13] 'After digital technologies, how do we manage AI's environmental footprint?', Pollutec (2025). Link.
[14] 'Digital sustainability', Arcep (2025). Link.
[15] Energy Efficiency Directive, European Commission (2023). Link.