Mapping AI risks for climate
Charlie Wilson & Felippa Amanta
28 May 2025

Ask what risk AI poses for climate, and there's a standard short answer. Energy-hungry data centres.

This is a narrow view of a much broader picture. Using a comprehensive AI risk taxonomy, we identify 30 ways in which the use and misuse of AI may jeopardise efforts to reduce carbon emissions. Our mapping of AI-climate risks ranges from adverse impacts on specific low-carbon technology deployment to a weakening of general climate policy and governance conditions. Our aim in setting out these AI risks for climate is to balance the prevailing 'AI for good' narrative and increase the prominence of climate change in AI governance frameworks.

Beneficial and Adverse Impacts of AI on Energy and Emissions

Energy-hungry AI models are driving data centre expansion from their current 1.5% of global electricity demand to a projected 3% by 2030 (IEA, 2025). Data centres to train and run AI models also explain tech companies' backsliding on their net-zero commitments. Consequently, tackling AI risks for climate has focused on efficiency standards and low-carbon power supplies for data centres (Wilson, Charlie, Fan, Yee Van, and Amanta, Felippa, 2025).

Data centres aside, the general narrative around AI's impact on climate is positive. AI for Good summits, the Global Partnership on AI, government strategies for harnessing AI to drive low-carbon growth - all rightly emphasise the numerous opportunities for AI to accelerate emission reduction efforts.

The global not-for-profit, Climate Change AI[i], documents over a hundred use cases for AI to strengthen climate science, accelerate scientific discovery (e.g. better battery catalysts), support integration of renewables into energy networks, and improve the efficiency of resource-consuming processes and systems throughout the economy (Rolnick *et al.*, 2022). AI also creates opportunities to progress not just on climate but also on wider sustainable development goals (Vinuesa *et al.*, 2020).

But … this is a one sided-story.

AI applications that reduce the cost, time, inconvenience, or effort of any given activity will induce more demand for that activity. This has trivial impacts on energy or other natural resources if the activity is online news, but potentially major impacts if the activity is self-driving cars. In extreme cases, a 'digital Jevons paradox' means that efficiency-enhancing AI results in demand for resources going up not down as scale growth outweighs productivity gains (Luccioni, Strubell and Crawford, 2025).

Beyond these direct implications for natural resource consumption, AI also impacts the social and governance institutions that underpin effective climate policy. Tackling climate change is a long-term, cross-cutting, multi-stakeholder, and transformative governance challenge. In democracies it is enabled by social trust, fairness (just transitions), cooperation, and consensus (Creutzig *et al.*, 2023; Wilkinson and Pickett, 2024). How may these elements of climate governance be affected by AI?

As one example, we know that our ever-more digitalised information and media environments (including the AI algorithms that manage our feeds) are associated with a well-informed citizenry, active civil society, and exposure to diverse opinions - all good for democratic institutions and so, by extension, for climate governance (Lorenz-Spreen *et al.*, 2023). But we also know that AI in digital media is associated with declining trust in institutions, open discrimination, misinformation, and political and social polarisation - all bad for democracy and so climate governance. AI in digital social networks further exacerbates these risks (Galaz *et al.*, 2025).

We used these examples of AI-climate risks as entry points to look more comprehensively at the many possible interactions between AI deployment and the building blocks of zero-emission pathways in line with Paris Agreement goals.

AI Risk Taxonomies

AI risk governance initiatives and ethics frameworks exhaustively map the potentials for AI to cause social or economic harms. This enables appropriate mitigation strategies to be put in place through industry commitments, multi-actor agreements like the Bletchley Declaration on AI Safety[ii] or the UN Global Digital Compact[iii], and regulations like the EU's AI Act[iv]. Risk mitigation strategies balance precautionary avoidance of harm with the openness for continued innovation.

Taxonomies like the MIT AI Risk Repository capture and define familiar dimensions of AI risk - misinformation, discrimination, privacy and security, malicious actors, job losses and other socioeconomic harms. These risks are widely discussed in many social and political contexts, but not specifically in relation to climate change. To the extent environmental harms are included, this is limited to … the energy footprint of data centres.

MIT's AI Risk Repository defines 7 domains with 24 subdomains of AI risk (Slattery *et al.*, 2024). (The repository itself is built from 1600+ risks extracted from 65 existing frameworks and classifications).

The left side of our AI-climate risk map sets out the main subdomains we found to be relevant for climate [see Figure].

Pathways To Zero-Emission Energy Systems

There are many different ways of representing what's needed to make progress towards climate targets. We set out a simple framework that focuses on zero-emission energy systems for two reasons. First, energy accounts for over 2/3 of global emissions. Second, information and communication technologies (ICTs) including AI most directly impact energy. Our framework could very simply extend to materials, food, land and other natural resources whose use is linked to emissions, as well as to carbon removal technologies with potential roles in net-zero emission pathways.

The right side of our AI-climate risk map sets out the main building blocks of zero-emission energy pathways [see Figure]. From top to bottom these building blocks range from proximate, direct drivers of emissions to ultimate, indirect drivers.

On the top right, reducing emissions means more efficient use of energy and more low or zero-carbon supply of energy (Allwood *et al.*, 2019). These changes in the energy system require low-carbon technology deployment such as heat pumps, electric vehicles, or solar panels. Households and firms shifting towards lower-carbon behaviours and business models are also part of energy system change.

Climate policy plays an important enabling role in driving emission reductions by setting incentives, aligning markets, and proscribing high-emission activity (Nachtigall *et al.*, 2024; Stechemesser *et al.*, 2024). In our framework we distinguish two particular elements of climate policy that may be impacted by AI. First, policymakers need the capacity and competence (time, skills, resources) to design, implement, monitor, and enforce climate policy (Luers *et al.*, 2024). Second, policymakers need action on climate to be salient, urgent, or otherwise important so it outcompetes other demands on their scarce resources.

This simple characterisation of climate policy for reducing emissions emphasises the role of state actors like government, ministries, and agencies. These in turn sit within a wider landscape of climate governance involving non-state actors, networks, and norms across multiple scales. In our framework we reduce this complexity to four particular elements that may be impacted by AI. First, climate governance needs political consensus behind the long-term strategic aim of transforming the energy system. Second, climate governance needs the urgency for policymakers to prioritise climate action to be reinforced by pressure from across civil society, business, and opinion leaders. Third, climate governance needs the legitimation of widespread social support for climate action given its distributional implications for winner and losers. Fourth, climate governance - like all governance - needs effective, well-functioning institutions including social trust, an informed and active citizenry, and a diverse information environment.

These select elements of climate policy and climate governance are the supporting scaffold for specific technological and behavioural changes in the energy system required to reduce emissions to zero. They are shown on the lower right side of our AI-climate risk map. It's important to reiterate this is a highly simplified representation of zero-emission energy pathways that emphasises what may be impacted by AI.

AI-Climate Risks: Mapping Potentially Harmful Impacts of AI for Zero-Emission Energy Systems

We used scientific literature, media reports, AI incident reports, policy and industry studies to identify links from any of the AI risk domains [left side of Figure] to any of the building blocks of zero-emission energy pathways [right side of Figure]. We compiled a database with an initial set of 30 AI-climate risks, supported by relevant citations.

Our initial efforts are informed but opportunistic. We do not claim that our AI-climate risk database is systematic. Indeed we expect there are many links we have missed, particularly those which are more speculative and lack specific climate examples.

For those 30 AI-climate risks in our database, each has both a general risk characterisation and a specific argument with an example of how the risk may cause harm to progress on climate mitigation.

To give an example [#A1 in Figure], 'disinformation' is a subdomain of AI risk within the 'malicious actors' domain. A general characterisation of this risk is that 'AI is used by malicious actors to create believable deepfakes for large-scale manipulation'. Such manipulation may be used to 'weaken support for low-carbon technology' (Berkebile-Weinberg *et al.*, 2025). This AI-climate risk is evidenced by a specific example: a fossil fuel-funded think tank in the US created a deepfake of dead whales to drum up opposition to an offshore wind project[v]. In our mapping, this information is condensed down into the AI risk - 'deepfakes for large-scale manipulation' - and the harmful climate impact, 'synthesised opposition to climate initiatives' [#A1 in Figure].

Each of the 30 AI-climate risks we have identified has a similar structure. We document details and examples in our database, and reduce these to summary labels in the figure.

Some of the harmful climate impacts have already been observed, like with the example of the deepfake dead whales.

Some of the harmful climate impacts are not yet directly observable but can be clearly inferred. An example is [#I3 in Figure]. The 'emotional amplification of worst case outcomes' by AI algorithms in online news platforms and a print media campaign catalysed resistance to proposed legislation for a gas boiler phase out in Germany. By halting or delaying the transition to heat pumps to decarbonise heat in buildings, this caused 'slower adoption of low-carbon technology' [#I3 in Figure].

Other harmful climate impacts in our database are evidenced and supported in scientific literature but are not yet observable specifically in relation to climate as they are part of a wider phenomenon. An example is #Q2 in our Figure. The 'job market and wage polarisation' effects of AI will worsen income inequality and so reduce the availability of capital among some households whose breadwinners will lose out (Acemoglu and Restrepo, 2020). Of the many welfare consequences, one will be to create 'financial constraints on low-carbon choices' [#Q2 in Figure]. Inequality poses much wider systemic challenges for governance which are picked up in other links #Q4-7 (Wilkinson and Pickett, 2024).

As two further examples, the top two links #S1 and #S2 in our Figure capture the most obvious and direct impacts of AI on energy systems: energy-hungry data centres [#S1] and induced demand for energy-intensive activity in buildings, industry, and transport sectors if AI makes those activities more cost or time efficient [#S2]. Both these links are documented in the IEA's recent report on energy and AI (IEA, 2025).

AI-Climate Risks: Interactions

Although our Figure shows unique links from one AI risk subdomain to one building block of zero-emission energy pathways, in reality different links interact. We try and capture these links in our database, but we don't show these in the Figure as it would turn into even more of a spaghetti.

An example is #D2 in our Figure. Biased training data for AI models marginalises data-poor groups. In an energy context this may include low-income households on non-smart prepayment meters. Using such models in policy assessments may then bias how resources are allocated to enable widespread participation in just energy transitions. This unique link shown in #D2 from the 'unfair discrimination' AI risk subdomain interacts with the 'power centralisation' AI risk which relates to the concentration of market power and political economic influence in the hands of a small number of AI providers. If they have commercial interests asymmetric to those of their model user, the distributional biases may be amplified. This interaction between AI risks is not shown in the Figure.

Another example is #P1 in our Figure. Data breaches from AI systems with vulnerabilities (one risk subdomain) can interact with malicious actors seeking to propagate disinformation, conduct surveillance, or otherwise exploit data (other risk subdomains). A pre-AI analogue of these interacting risks was the hacking of climate scientists' emails and the subsequent false representation of their content to undermine the integrity of climate science on warming trends observed historically (Leiserowitz *et al.*, 2013). Fast forward to today, and UK owners of Chinese electric vehicles are being cautioned from pairing their smartphone navigation apps to their cars due to concerns of data breaches and resulting surveillance.[vi]

AI-Climate Risks: The Full Picture

The full AI-climate risk mapping [Figure] is dense, complex, and diffuse.

First it's dense, with a large number of links that are just an initial set we expect to expand. Indeed we would welcome suggestions of risks we have missed, particularly if they are clearly evidenced already.

Second, it's complex in that no single AI risk domain maps onto a single zero-emission energy building block. However some links are more common, particularly those from socioeconomic AI risks (top left) to climate governance (bottom right).

Third, it's diffuse both in the variety of AI risks involved, and in the multiple levels of zero-emission energy pathways involved. This is not surprising as AI is a general purpose technology (a GPT!) with applications throughout the economy and society. Fully decarbonising our energy system is also a transformational challenge that similarly touches most economic and societal activity in some way. As both AI and energy transitions are systemic, links between them are too.

These characteristics of our AI-climate risk mapping make it harder to come up with clear mitigation strategies.

AI-Climate Risks: So What?

We had two main aims in developing our AI-climate risk map - pedagogy and policy.

First and foremost we simply wanted to communicate a story about potential AI risks for climate. It's intentionally very one-sided because we think the inverse account of AI opportunities for climate is being more widely told and heard. We hope our Figure achieves this pedagogical aim, even if it does need a fair bit of explanation! As we've shown, AI risks for climate are many, varied, and diffuse.

Second, we wanted to contribute to policy and governance debates on how to mitigate AI-climate risks. These debates are rich and active for all sorts of other AI risks, but not (yet) for climate. What's our evidence for this claim?

Guidelines being developed for AI are founded on ethical principles. These principles emphasise transparency, justice and fairness, not causing harm, and so on. Sustainability - including a stable climate - is near the bottom of the list (Jobin, A., Ienca, M., and Vayena, E., 2019). It should be higher.

Regulatory frameworks for mitigating AI risks formalise these guidelines. For example, the EU's AI Act differentiates risks created by AI systems. Unacceptable risks are banned. High risks are regulated, for example, through precautionary assessment. Transparency is encouraged for more limited risks. Climate risks barely feature as a specific form of undesirable harm the framework is designed to avoid. They should be. In the EU this is particularly so because the EU's twin transition strategy entwines AI and climate as synergistic strategic goals.

Regulations are one of many possible forms of risk mitigation strategy. While the EU may take a more regulatory approach, other jurisdictions like the US rely more on industry commitments or self-regulation. Internationally, a variety of summits, declarations, compacts, agreements, and working processes involving both governments and industry articulate AI risk mitigation approaches. These can involve multilateral bodies like the UN, as well as civil society organisations.

Environmental assessments of AI impacts richly characterise and quantify the direct energy and emission footprint of ICT infrastructure. The IEA's recent landmark report goes further in estimating the indirect energy savings from efficiency-improving AI applications across the economy (IEA, 2025). The extent to which these will induce demand that mean net emissions go up not down is mentioned but not assessed beyond a hazy assumption in a graph on literally the last page of the report (p252, Fig 5.3). We should do better at measuring AI-climate impacts in the round.

Environmental assessments of AI impacts inform the importance of climate in ethical guidelines for AI governance. These in turn inform regulatory and other policy responses to mitigate AI risks.

AI-climate linkages should be given more attention throughout. Evidencing why is the second aim of our risk mapping. The scope and transformative potential for AI means its deployment should be aligned with climate stabilisation goals, or those goals will be pushed out of reach (Gaffney *et al.*, 2025).

AI-Climate Risks: Governance & Risk Mitigation

Going from mapping to mitigating AI-climate risks means going from what to who and how? Our mapping shows what risks there are. But which entities should mitigate each risk (who?) and with what levers (how?).

Who should be responsible for mitigating AI-climate risks?

In many if not most cases, the risk is the result of how AI systems are used. AI is just the most recent of a long line of digital, smart, connected, automated technologies deployed in our buildings, vehicles, factories, and cities. AI is an amplifier and an accelerator, not a prime cause. If you dropped AI, the climate risk would still be there. If AI-enabled optimisation indirectly leads to emission pressures in those use contexts, this is for 'conventional' climate policy to tackle downstream by making sure markets and processes are aligned towards zero-emission energy outcomes. The who? are therefore urban planners, transport system designers, energy market regulators, climate ministries, building designers, electricity network operators, and the many other actors with a stake in zero-emission energy pathways.

This is a 'weak' view of AI-climate risk governance. It argues that only those risks unique to AI should be accounted for as they will not be tackled elsewhere. For energy and emissions this mainly means data centres. For all the other more diffuse impacts of AI mediated by mis- and dis-information, cybersecurity threats, political and job market polarisation, and so on - these are not distinctively about climate but about AI more generally. So they should be left to general AI governance frameworks. A deep fake of a whale to breed opposition to offshore wind is like a deep fake of a politician to breed opposition to their re-election. They're both deep fakes.

In contrast, a 'strong' view of AI-climate governance sees risks as the result of how AI systems are developed as well as used. AI is a general purpose digital technology but uniquely so. AI risks can and should be mitigated at source. The AI industry is highly concentrated (particularly for generative AI), deeply resourced, and influential economically and politically. Intervening upstream represents a powerful leverage point for affecting change. The who? are primarily therefore tech companies and their financial backers, national governments, and international organisations including the UN family of institutions.

How should AI-climate risks be mitigated? Risk governance is not synonymous with regulation. Policymakers are reluctant to stifle innovation through overly rigid legal constraints. Guardrails can be put in place through other means including industry-led processes with oversight and transparency. Competitive differentiation is already seeing some AI companies pledging anything from more energy-efficient models to more ethically-led deployment.

The weak and strong views of AI-climate governance place differing emphases on mitigation through climate policy and mitigation through AI policy respectively. The actors involved correspondingly differ: energy companies vs tech companies, climate ministries vs science and innovation departments, households and firms using energy vs households and firms using AI.

These are just initial ideas. A next stage in our work is to build out a more nuanced AI-climate risk mitigation framework for each of the risks we show in our map.

This goes way beyond data centres. It also means tackling AI risks of "exacerbating social injustice, eroding social stability, and weakening our shared understanding of reality" because these undermine effective climate policy and governance (Gaffney *et al.*, 2025).

References

Acemoglu, D. and Restrepo, P. (2020) 'Robots and Jobs: Evidence from US Labor Markets', *Journal of Political Economy*, 128(6), pp. 2188–2244. Available at: https://doi.org/10.1086/705716.

Allwood, J.M. *et al.* (2019) *Absolute Zero: Delivering the UK's climate change commitment with incremental changes to today's technologies*. Cambridge, UK: University of Cambridge. Available at: doi.org/10.17863/CAM.46075.

Berkebile-Weinberg, M. *et al.* (2025) 'Internet image search outputs propagate climate change sentiment and impact policy support', *Nature Climate Change*, 15(1), pp. 44–50. Available at: https://doi.org/10.1038/s41558-024-02178-w.

Creutzig, F. *et al.* (2023) 'Designing a virtuous cycle: Quality of governance, effective climate change mitigation, and just outcomes support each other', *Global Environmental Change*, 82, p. 102726. Available at: https://doi.org/10.1016/j.gloenvcha.2023.102726.

Gaffney, O. *et al.* (2025) ‘The Earth alignment principle for artificial intelligence’, *Nature Sustainability*, 8(5), pp. 467–469. Available at: https://doi.org/10.1038/s41893-025-01536-6.

Galaz, V. *et al.* (2025) ‘Artificial intelligence, digital social networks, and climate emotions’, *npj Climate Action*, 4(1), p. 23. Available at: https://doi.org/10.1038/s44168-025-00225-6.

IEA (2025) *Energy and AI*. Paris, France: International Energy Agency (IEA).

Jobin, A., Ienca, M., and Vayena, E. (2019) ‘The global landscape of AI ethics guidelines’, *Nature Machine Intelligence*, 1, pp. 389–399. Available at: https://doi.org/10.1038/s42256-019-0088-2.

Leiserowitz, A.A. *et al.* (2013) ‘Climategate, Public Opinion, and the Loss of Trust’, *American Behavioral Scientist*, 57(6), pp. 818–837. Available at: https://doi.org/10.1177/0002764212458272.

Lorenz-Spreen, P. *et al.* (2023) ‘A systematic review of worldwide causal and correlational evidence on digital media and democracy’, *Nature Human Behaviour*, 7(1), pp. 74–101. Available at: https://doi.org/10.1038/s41562-022-01460-1.

Luccioni, A.S., Strubell, E. and Crawford, K. (2025) ‘From Efficiency Gains to Rebound Effects: The Problem of Jevons’ Paradox in AI’s Polarized Environmental Debate’. arXiv. Available at: https://doi.org/10.48550/arXiv.2501.16548.

Luers, A. *et al.* (2024) ‘Will AI accelerate or delay the race to net-zero emissions?’, *Nature*, 628, pp. 718–720. Available at: https://doi.org/10.1038/d41586-024-01137-x.

Nachtigall, D. *et al.* (2024) ‘The Climate Actions and Policies Measurement Framework: A Database to Monitor and Assess Countries’ Mitigation Action’, *Environmental and Resource Economics*, 87(1), pp. 191–217. Available at: https://doi.org/10.1007/s10640-023-00821-2.

Rolnick, D. *et al.* (2022) ‘Tackling Climate Change with Machine Learning’, *ACM Comput. Surv.*, 55(2), p. Article 42. Available at: https://doi.org/10.1145/3485128.

Slattery, P. *et al.* (2024) ‘The AI Risk Repository: A Comprehensive Meta-Review, Database, and Taxonomy of Risks From Artificial Intelligence’. arXiv. Available at: https://doi.org/10.48550/arXiv.2408.12622.

Stechemesser, A. *et al.* (2024) ‘Climate policies that achieved major emission reductions: Global evidence from two decades’, *Science*, 385(6711), pp. 884–892. Available at: https://doi.org/doi:10.1126/science.adl6547.

Vinuesa, R. *et al.* (2020) ‘The role of artificial intelligence in achieving the Sustainable Development Goals’, *Nature Communications*, 11(1), p. 233. Available at: https://doi.org/10.1038/s41467-019-14108-y.

Wilkinson, R.G. and Pickett, K.E. (2024) ‘Why the world cannot afford the rich’, *Nature*, 627(8003), pp. 268–270. Available at: https://doi.org/10.1038/d41586-024-00723-3.

Wilson, Charlie, Fan, Yee Van, and Amanta, Felippa (2025) ‘AI’s indirect impacts on climate outweigh concerns over its direct energy footprint’, *Oxford Energy Forum*, Artificial Intelligence and its Implications for Electricity System(145), pp. 10–15.

Figure: AI-Climate Risk Map.

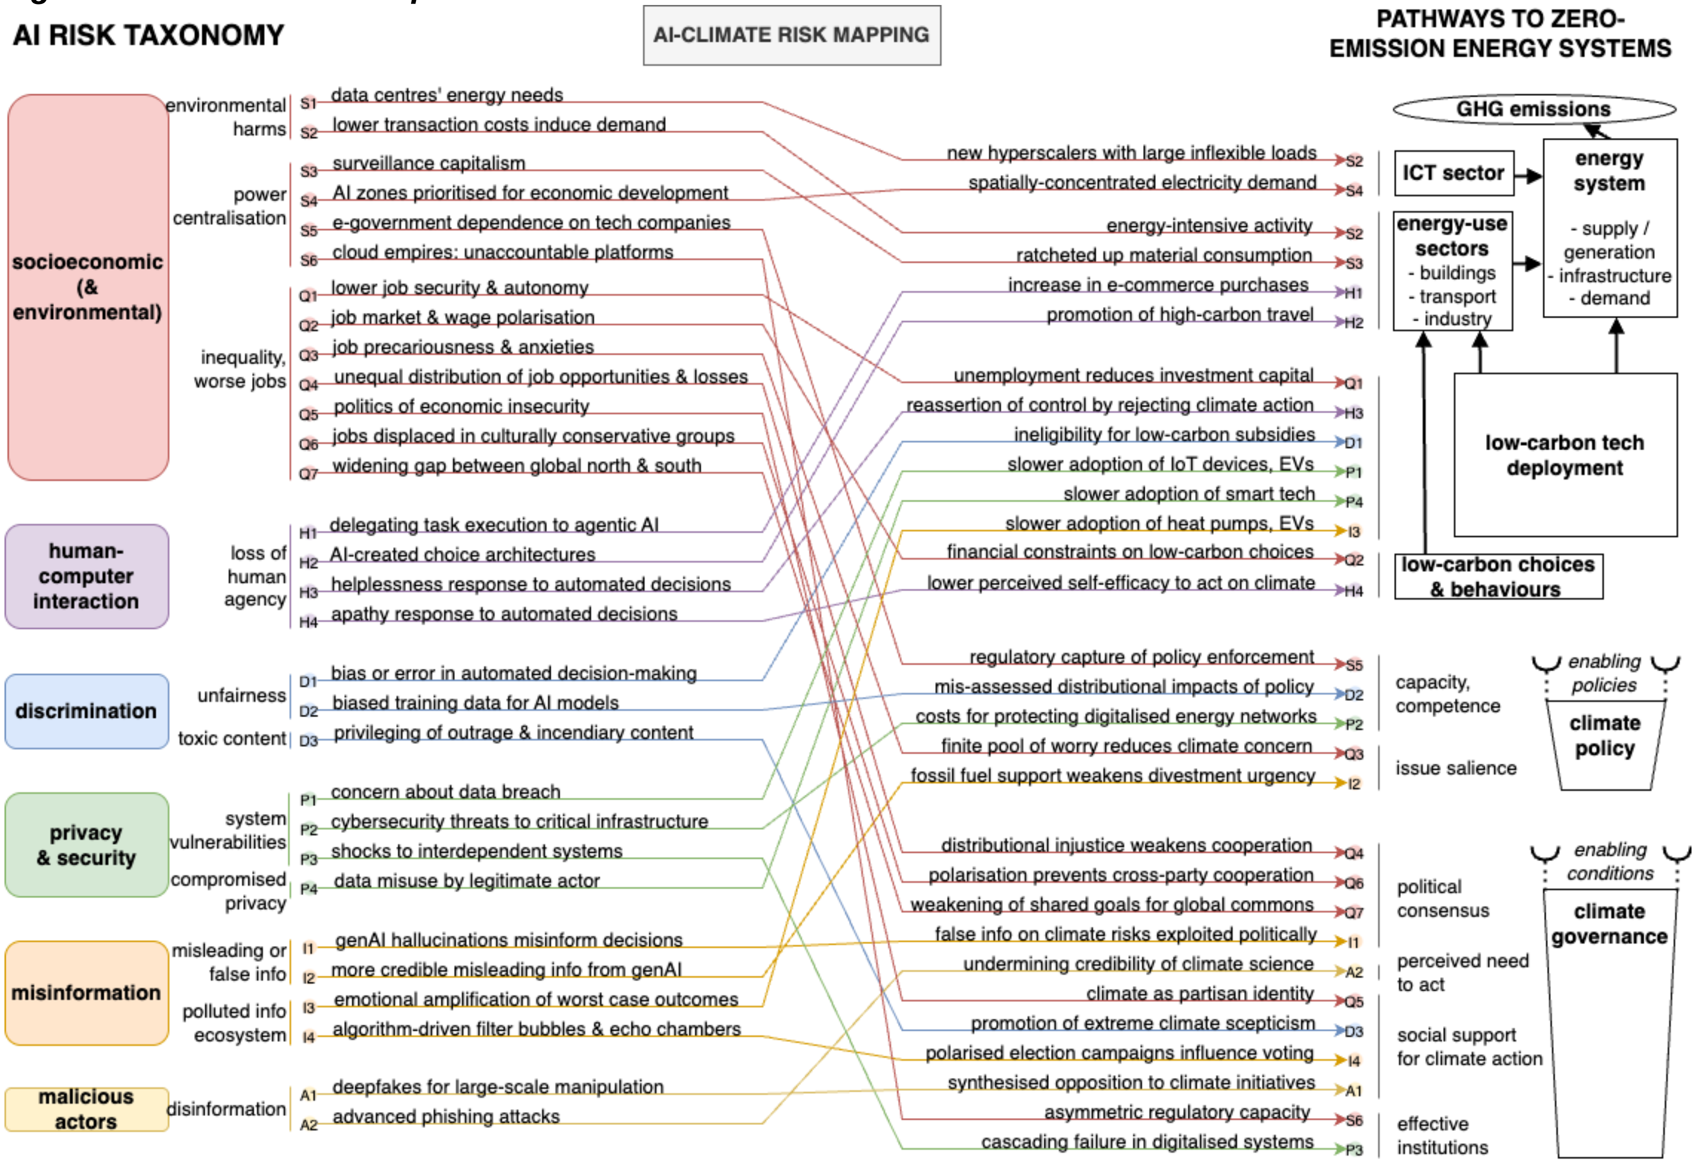

Endnotes

[i] https://www.climatechange.ai
[ii] https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023
[iii] https://unglobalcompact.org
[iv] https://artificialintelligenceact.eu/the-act/
[v] Taft, M. ‘Leading denier think tank uses AI image of dead whale and wind turbines’. 16 March 2023. Gizmodo. https://gizmodo.com/climate-denier-newsletter-ai-image-dead-whale-wind-farm-1850234135
[vi] Milmo, D. ‘Source of data: are electric cars vulnerable to cyber spies and hackers?’. The Guardian. 29 April 2025. https://www.theguardian.com/environment/2025/apr/29/source-of-data-are-electric-cars-vulnerable-to-cyber-spies-and-hackers