Response to the UK's Future of Compute Review: A Missed Opportunity to Lead in Compute Governance

Jess Whittlestone (Centre for Long-Term Resilience)
Shahar Avin (Centre for the Study of Existential Risk)
Lennart Heim (Centre for the Governance of AI)
Markus Anderljung (Centre for the Governance of AI)
Girish Sastry (OpenAI)

We are pleased to see the publication of the UK's Future of Compute Review, making a number of recommendations for investing in a strong compute ecosystem in the UK, including the development of a UK AI Research Resource to address the compute divide between academia and industry (an idea several of us supported in evidence submissions to the review).

However, we also believe there is a significant missed opportunity: **the review does not address how to ensure that compute is used responsibly or how its usage might be governed**, especially in the context of frontier AI development.

In the National AI Strategy, the UK has already shown that it is committed to creating a governance environment that enables safe and responsible development of AI. Compute is currently a key ingredient in the development of cutting-edge AI systems. The governance of compute therefore offers a promising avenue for AI governance, enabling the UK to steer global development of this technology in line with its values.

The most advanced AI systems, developed with large-scale compute, tend to have emergent capabilities that are difficult to predict in advance, and are not necessarily beneficial. Recent frontier advances such as OpenAI's ChatGPT and, in particular, Microsoft's Bing/Sydney have demonstrated deceptive and harmful behaviour, as well as potential for misuse, capturing public attention in a way that no previous advances in AI have. There is growing expert consensus that compute-intensive AI systems will continue to advance

rapidly and unpredictably in coming years, with potential risks to individual safety, economic stability, and national security. With increasing public attention and recognition of these risks, now is the time for the UK government to demonstrate leadership.

Governing the compute used to develop frontier AI systems provides a clear mechanism through which to oversee and intervene in the highest-risk areas of AI development and deployment, while leaving all other areas of AI development unencumbered. In contrast to other inputs to AI progress such as data, algorithms, and talent, compute hardware has many features which make it a good governance target: it is centralised, easily quantified, and hard to duplicate.

While the final Future of Compute review acknowledges that "compute-intensive applications of AI pose novel risks", this is only a short paragraph in a long report, with no associated practical recommendations for addressing the risks. Elsewhere, the report emphasises the importance of ensuring the security of compute infrastructure, acknowledging that "a cultural change towards the adoption of a risk-based approach is required to make best use of available resources". We strongly agree, and would underline that there are many more productive steps the UK could take to adopt a risk-based approach above and beyond focusing on compute security.

Our high-level recommendation is that the UK Government should explore ways to govern high-risk uses of compute, particularly in frontier AI development. Ideas to explore include:

1. **A tiered access approach to compute provision** via the proposed UK AI Research Resource, where access to larger amounts of compute comes with additional requirements: to demonstrate responsible use or subject systems to external review or scrutiny. For more details see recommendation (2) of CLTR's submission and recommendation (4) of GovAI's submission to the Future of Compute Review's Call for Evidence.

2. **Requiring AI companies to report, or possibly in the future apply for a license for, training runs above a certain (very high) threshold**. Such a reporting regime could give government oversight over and knowledge of particularly high-stakes AI

development at the frontier of capabilities, while leaving all other economically beneficial progress unencumbered. A first step could be a voluntary reporting pilot with companies particularly committed to responsible development, of which we believe there would be several. For more details see recommendation (2) of CLTR's submission.

3. **Requiring compute providers to have "Know Your Customer (KYC)" processes** around the use of very large amounts of compute for AI development, including potentially checking customers against blacklists, or investigating the risk that their compute provision aids human rights abuses. This is analogous to requirements imposed on banks to know who their customers are, to thwart tax evasion and money laundering. This would complement efforts to ensure compute security, recognising that misuse can come from many sources. We would only expect this to apply to a handful of customers and so wouldn't be overly burdensome on providers.

4. **Facilitating academic access to large pre-trained models to address the compute divide**. The compute review suggests the creation of a national computing resource. This is partly valuable because it allows academic researchers to scrutinise the world's frontier models, helping them keep frontier AI development accountable to the public interest. However, that goal may be more directly achieved by facilitating academic access to large pre-trained models already developed by private frontier AI labs. We recommend that the UK AI Research Resource provides not only compute, but also API access to frontier models. For more details see GovAI's research post on Compute Funds and Pre-trained Models and recommendations (2) and (3) of GovAI's submission.

Governing compute usage in these ways needn't hinder, and can actually support, the UK's ambition to be a science and technology superpower. By focusing governance only on the small number of cases where compute is driving high-risk AI development, it will be easier for the UK to move fast in all other areas, knowing the most severe harms are being effectively managed. As the review itself recognises, "the compute required for AI is distinct from that of more traditional uses", relying largely on more specialised AI accelerators, such as graphical processing units (GPUs), Graphcore's IPUs, or Google's TPUs. This means that

the kind of governance we are proposing need not even touch most of the economically beneficial uses of compute that the review discusses. It need not even touch most uses of compute in AI research - only that which is being used to develop and deploy increasingly general-purpose, large-scale AI models at leading companies.

The UK has an opportunity to demonstrate international leadership here. We believe that it will become increasingly clear in the coming years that some form of compute governance is essential to making the most of technological opportunities while mitigating the largest risks they might pose. The idea that large amounts of compute should come with large responsibility is receiving increasing attention among AI policy experts and leading AI companies - in a recent blog post, OpenAI CEO Sam Altman stated "we think it is important that major world governments have insight about training runs above a certain scale." The UK has shown that it is ahead of the curve in recognising the importance of compute for future economies, and simultaneously in committing to establish a pro-innovation regulatory regime for AI. Compute governance sits squarely at the intersection of these two areas, and it would be a shame not to explore it.