UK AI COUNCIL

December 2022 Draft Memo

The UK Large Language Models opportunity

HMG should take action to build national and sovereign capability in Large Language Models, a type of Artificial Intelligence trained to analyse written language and generate convincing novel text. Without action to establish national capability, HMG risks handing the keys to a strategically significant technology to overseas companies, creating a technology gap that will expand over time and expose the UK to vulnerabilities across the public, private, and security sectors. The far reaching nature of LLMs means government needs to consider interventions in multiple parts of the ecosystem (commercial, academic, and otherwise) to deliver the scale of progress required. The AI Council stand ready to help to convene people around this important opportunity.

Rapid developments in 'Large Language Models' (LLMs) over the last 5 years are accelerating the capabilities of AI systems and represent a new exciting opportunity for the UK. LLMs are algorithms that are trained to analyse the structure and meaning of written language; they synthesise text, and can provide answers to user-posed questions with convincing detail and tone. The current generation of LLMs are typically trained on all the text data on the web; users can engage in dialogue with the LLM, asking questions and in return receiving novel answers generated by the model. Recent progress, as demonstrated by the success of systems such as GPT3, illustrate an accelerating trend in the capabilities of LLMs.

Widespread adoption of LLMs has potentially profound social and economic consequences. The accuracy and convincingness of these models represents a gear-change in AI capabilities. In education, they could act as personal tutors to enhance teaching and learning by providing accurate answers to student questions; in healthcare, they could synthesise complex patient information from diverse sources of text and data to support doctors in identifying treatment pathways; across the public sector, they could help citizens access the information and services they need, responding to user questions by synthesising large volumes of information from different websites and presenting a user-friendly answer. As the ability of these models to generate convincing, novel text increases, they may also be deployed to analyse genomic data (identifying the effects of genetic changes from nucleotide letter sequences) or environmental datasets (supporting sustainable living). If successfully developed in the UK, such systems could both alleviate pressures on public services nationally, and be sold internationally. If not carefully designed, deployed, stewarded, and regulated, these models can also: reproduce or reinforce harmful biases or online abuse; present false information; or mislead users, if interactions with an AI agent that generates convincingly 'human' answers are mistaken for interactions with another person.

Malicious use of LLMs could underpin new State threats; LLMs are also needed to enhance cybersecurity. LLMs are likely being deployed by hostile states to generate and propagate disinformation; engage UK and allied personnel in chat rooms; deliver malicious cyber payloads, and more. Further advances open the possibility of new threats that could materialise in 2-5 years, alongside unforeseen issues. For example: LLMs could be trained to scan code bases for vulnerabilities and generate code to exploit those vulnerabilities in real time, creating new cybersecurity risks; their ability to analyse all material ever written on the web could be leveraged to identify undercover agents, either internationally or in national policing and crime prevention efforts; LLMs could also be deployed in support of cultural conflict or polarisation. If these risks emerge, whether a Nation State has access to a cutting edge, national and sovereign LLM capability will have major consequences for national security. 'White hat' applications could also enhance defensive capabilities, for example detecting fake or malicious content.

Investment in LLMs is currently led by US and Chinese tech firms. The scale of investment and rate of progress is such that it is reasonable to expect the performance of these models to accelerate

significantly in the coming years. Open source LLMs are emerging, but these lack the advanced capabilities of state-of-the-art methods. Performance of these models – as with any model – will be limited by its initial design and training. Development of open source systems will be significantly slower than state-of-the-art or that of adversaries in a security context.

The UK is currently behind in LLM capabilities and there is a pressing need for national capability in LLM development and deployment. OpenAI, Meta, Google, Microsoft, NVIDIA, and other large companies are investing heavily in the development of LLM capabilities; VC interest in the field of generative AI is also growing. The concentration of technological capability in the US and Chinese private sector, the significance of LLMs as a tool for public service delivery, and the security implications of this technology create risks for the UK's technological sovereignty, if action is not taken to build national LLM capability. The scale of investment required to develop LLMs requires a critical mass of resources and expertise, in contrast to other AI methods which require distribution of capability. The UK has a thriving start-up scene, excellent academic expertise in AI, and links with partners in the US and Europe that could be leveraged in support of large-scale science, technology, and security collaborations. However, UK-based start-ups do not have access to the same long-term capital as in the US.Meanwhile academics do not have access to the compute power that tech giants possess. Further, the rapid development of national capability requires action across public and private sectors, creating an alignment challenge. The result is that the UK is behind in UK based LLM capability, creating a technological vulnerability.

Technologies available today could be deployed to support public service delivery. The LLM technologies already available today could be deployed for public benefit, leveraging public procurement to grow private sector capabilities and creating the conditions for the UK to reap wider economic benefits from nurturing technology leaders in a field with significant economic potential. Achieving the benefits of LLMs for public services requires a coordinated approach across the public and private sectors. Building private sector capability could be achieved through programmes focused on operational needs of the public sector (including defence and security) that are developed in conjunction with British founded, headquartered and owned applied AI specialist companies. This is a significantly concrete step change that enables the government to invest in foundational AI technologies in a practical fashion that will have returns to the taxpayer.

Progress in R&D is also necessary to shape the future of LLMs. The Alan Turing Institute has been engaging core partners from the public and private sector through its Foundation Model programme, alongside potential user groups. This work suggests a need for a large compute resource, an R&D effort focussed on the core computer science research to build large language models, training programmes (such as a Centre for Doctoral Training) to build capability, and mechanisms for engaging the start-up community. A roundtable with ARIA in January 2023 will further scope ideas in this space.

LLMs offer HMG an opportunity to leverage a new technology for strategic advantage and commercial benefit. Investment in the creation of a national and sovereign LLM would place the UK in a less vulnerable position in the scenario where the technology develops quickly, and would provide a foundational public service and security capability of potentially huge strategic value. It would also position the UK for international influence. Achieving this step-change in capabilities will involve investment in research, public procurement, and UK start ups to ensure the UK maintains ownership of critical national infrastructure, alongside efforts to collaborate internationally with allies that share the goal of enabling progress in trustworthy AI.

This requires further work to:

- prepare to position the public sector as an early and sophisticated customer of current LLM technologies;
- build private sector capability, using smart public procurement to incentivise private investment, develop public-private partnerships to support new collaborations, and foster competitive advantage in the longer-term;
- lead a conversation about the governance and regulation implications of LLMs in collaboration with international allies;
- bring together researchers, businesses and the public sector to drive progress in the next generation of LLM technologies, working across universities, the Alan Turing Institute, ARIA, HMG, and the private sector.

Recommendation

If the UK wishes to deliver on its ambition to be a world leader the field of AI, we should play to our strengths and move quickly to seize this opportunity. HMG should join up with industry and academia - to examine the options to build a national capability, to understand and articulate the risks of doing nothing, and to develop a clear plan of action to secure advantage in LLM.