

Andreessen Horowitz—written evidence (LLM0114)

House of Lords Communications and Digital Select Committee inquiry: Large language models

Andreessen Horowitz (“a16z”) appreciates the opportunity to provide this letter during the House of Lords’ Communications and Digital Committee’s inquiry on Large Language Models. We welcome the collaborative discourse around AI governance in the UK, and the UK’s leadership in this space globally.

a16z is a venture capital firm that invests in seed, venture, and late-stage technology companies, focused on bio and healthcare, AI, consumer, crypto, enterprise, fintech, and games. a16z currently has more than £28 billion in committed capital under management across multiple funds, and we invest in startups that both build and rely on artificial intelligence technologies. Our funds typically have a 10-year time horizon, as we take a long-term view of our investments.

Why AI Can Make Everything We Care About Better

The most validated, core conclusion of social science across many decades and thousands of studies is that human intelligence makes a very broad range of life outcomes better. Smarter people have better outcomes in almost every domain of activity: academic achievement, job performance, occupational status, income, creativity, physical health, longevity, learning new skills, managing complex tasks, leadership, entrepreneurial success, conflict resolution, reading comprehension, financial decision making, understanding others’ perspectives, creative arts, parenting outcomes, and life satisfaction.

Further, human intelligence is the lever that we have used for millennia to create the world we live in today: science, technology, math, physics, chemistry, medicine, energy, construction, transportation, communication, art, music, culture, philosophy, ethics, morality. Without the application of intelligence on all these domains, we would all still be living in mud huts, scratching out a meager existence of subsistence farming.

Instead, we have used our intelligence to raise our standard of living on the order of 10,000X over the last 4,000 years. What AI offers us is the opportunity to profoundly augment human intelligence to make all of these outcomes of intelligence – and many others, from the creation of new medicines to ways to solve climate change to technologies to reach the stars – much, much better from here.

AI augmentation of human intelligence has already started (it is already

around us in the form of computer control systems of many kinds), is now rapidly escalating with AI Large Language Models (LLMs) like ChatGPT, and will accelerate very quickly from here – if we let it.

Protecting Open Source AI is a Necessity

We believe governments should not deviate from longstanding technology policy principles supporting open source computing that have been widely accepted and legally enshrined for decades, since the advent of the Internet. It is critical to realize that restricting the ability to develop open source software will undermine the competitive AI landscape and harm, rather than enhance, cyber-security.

Open source technology has played a pivotal role in fostering innovation, encouraging competition, and democratizing access to technology. For example, open source software, like Linux, underpins the vast majority of cloud computing and is used widely throughout the U.S. government. It has also been a critical component for creating a more secure technology stack. It is commonly recognized that the approach of “security through obscurity” has been a failure, and that the most secure code is a function of how broadly it is laid out for the world to see, share, and open to challenge. Open source technology is the reason the Internet has succeeded far beyond our imagination and it is in our national interest to protect it and encourage it in the AI era for three core reasons:

1. **Safety & Security:** The disclosure of code for public review, while seemingly counterintuitive, significantly enhances security. Consider open-source software as akin to a peer-reviewed publication. A large group of individuals meticulously examines each line of code, identifying potential flaws, vulnerabilities, and points of exploitation. This robust global community ensures that any weaknesses are promptly identified and fixed. Conversely, in proprietary systems, vulnerabilities may linger undetected, or worse, detected only by bad actors who go searching for these weaknesses in order to exploit them. In the open-source domain, these are promptly brought to light and addressed, thus cultivating a digital environment that is both innovative and inherently secure. Although advocates for AI safety guidelines often allude to the "black box" nature of AI models, where the logic behind their conclusions is not transparent, recent advancements in the AI sector have resolved this issue, thereby ensuring the integrity of open-source code models.
2. **Competition & Transparency:** The trajectory of AI's future should not be under the control of a select few corporate entities. Instead, it should be the result of open competition and encompass a set of global voices that allow diverse insights and ethical frameworks. Open-source AI is the catalyst for this competition and transparency. It guarantees that the algorithms shaping our society are transparent, accountable, and subject to modification by a diverse array of individuals with different

insights and areas of expertise, rather than an elite few. It also ensures that a broader set of participants beyond dominant technology incumbents will have a role in shaping this technology. By championing open source, we are safeguarding a future where technology is of the people, by the people, and reflective of humanity's collective wisdom. Regulation that allows several large tech corporations to capture the market will ensure only a select few own and operate a repository of human knowledge and culture. Open source ensures that the AI era is both inclusive and just.

3. Academia: The embrace of open source by the academic community has allowed researchers to operate outside of the restrictions imposed by proprietary platforms and tools. It provides researchers with an unbounded laboratory, free from the constraints of commercial licenses. Perhaps more importantly, the computing resources required to develop cutting-edge AI developments, like Large Language Models (LLMs), are substantial, and effectively out of reach of academic researchers today. In the absence of open-source models, academic researchers will lack the ability to meaningfully participate in the development of LLMs and other computationally intensive models.

Embracing AI is Critical for National Security

It's not hyperbole to assert that America's AI leadership is crucial, especially when viewed through the prism of the West's dynamic with China. As Beijing aggressively integrates AI into its military strategies, surveillance apparatus, and economic master plans, the U.K.'s place as a technological beacon isn't just innovative startups—it's about safeguarding national security. If the UK and aligned western governments let our AI momentum wane, we risk being outpaced in areas like cybersecurity, intelligence operations, and modern warfare, handing strategic advantages to a formidable competitor.

But this also has significant economic and ideological ramifications. AI is a foundational computing technology that will continue to transform all sectors of the economy and drive the creation of new industries and jobs in ways we cannot yet anticipate. The ability of the U.K. economy to disproportionately benefit from AI – as other countries have with microchips and the internet – depends critically on whether AI is developed by companies within the U.K. Additionally, the U.K.'s AI endeavors are intertwined with its democratic fabric, emphasizing individual freedoms, privacy, and an ethos of open innovation. In stark contrast, China's AI trajectory is heavily influenced by state control and surveillance priorities. If the U.K. leads the West by standing at the forefront of AI, we can drive global norms that prioritize these democratic values in the emerging AI-driven world. Overbearing regulations risk ceding western leadership to China reshaping the global tech ecosystem in a way that's less transparent and more authoritarian, with ripple effects that could redefine the internet's DNA for the next 20 years.

While the West currently holds a distinct edge over China, this advantage is not guaranteed. These monumental advancements, borne out of an ecosystem that championed unbridled AI R&D, now face an ironic twist. Just as we're unlocking AI's immense potential, there's a rising chorus suggesting we pump the brakes. Calls for industry restraint and AI licensing will pave the way for regulatory frameworks that incumbent behemoths will exploit, sidelining the very startups that inject fresh innovation into our tech landscape. As we navigate the rise of AI, we must ensure that the spirit of innovation that got us here isn't smothered by the very mechanisms intended to safeguard it.

Goals for AI Policy Outcomes

- Big AI companies should be allowed to build AI as fast and aggressively as they can – but not allowed to achieve regulatory capture, not allowed to establish a government-protected cartel that is insulated from market competition due to speculative claims of AI risk. This will maximize the technological and societal payoff from the amazing capabilities of these companies, which are jewels of modern capitalism.
- Startup AI companies should be allowed to build AI as fast and aggressively as they can. They should neither confront government-granted protection of big companies, nor should they receive government assistance. They should simply be allowed to compete. If and as startups don't succeed, their presence in the market will also continuously motivate big companies to be their best – our economies and societies win either way.
- Open source AI should be allowed to freely proliferate and compete with both big AI companies and startups. Development of open source code should continue to be unregulated - as it is today. Use of open source code by bad actors for illicit activity is already heavily regulated and criminally prohibited and those standards should apply to the use of open source AI. Even when open source does not beat companies, its widespread availability is a boon to students all over the world who want to learn how to build and use AI to become part of the technological future, and will ensure that AI is available to everyone who can benefit from it no matter who they are or how much money they have.
- To offset the risk of bad people doing bad things with AI, governments working in partnership with the private sector should vigorously engage in each area of potential risk, using AI to maximize society's defensive capabilities. This shouldn't be limited to AI-enabled risks but also more general problems such as malnutrition, disease, and climate. AI can be an incredibly powerful tool for solving problems, and we should embrace it as such.

- To prevent the risk of China achieving global AI dominance, we should use the full power of our private sector, our scientific establishment, and our governments in concert to drive Western AI to absolute global dominance, including ultimately inside China itself. We win, they lose.

Again, we deeply appreciate the Committee's inquiry into this important topic and the leadership of the UK in ensuring western leadership of AI LLM development.

1 December 2023