cohere

Date:	October 22, 2023
Subject:	*U.S. Senate AI Insight Forum: Innovation - Cohere's Written Submissions*

Cohere is a frontier foundational model developer at the forefront of generative AI model innovation

Cohere is the leading enterprise-focused foundation model developer worldwide, providing developers and businesses access to natural language processing (NLP) AI capabilities powered by large language models (LLMs). Cohere's state-of-the-art foundational models, all trained in-house, compete with the world's other leading generative AI developers (e.g. OpenAI, Google, Meta, and Anthropic), while differentiating itself on independence - unlike the others, Cohere has not taken multi-billion dollar funding from one of the tech giants - data security, deployment flexibility - with the ability to easily deploy in multiple cloud environments such as GCP, Azure, AWS, and OCI, as well as on-premise - and customization capabilities through expert finetuning. Cohere is the only LLM model developer to sign *both* the White House's Updated Voluntary Commitments and the Canadian Federal Government's Voluntary Artificial Intelligence Code of Conduct.

Cohere is at the forefront of developing advanced LLMs that are deployed responsibly, emphasizing safety and accountability through:

- *Cohere's industry-leading machine learning team*: Cohere's research team is world-renowned, having contributed to the development of the transformer architecture[1], sentence transformers for semantic search[2], dynamic adversarial data collection and red teaming, and retrieval augmented generation[3].
- *Investing in research to better support responsible model development and deployment*: Cohere For AI (C4AI), Cohere's non-profit research lab, is dedicated to contributing fundamental research that explores the unknown in machine learning. C4AI's research focus includes model evaluations, model interpretability, traceability, and fairness. C4AI also supports a growing external ecosystem of scholarship on these issues via both a dedicated scholar program to support researchers working on responsible AI as well as a full-time research staff contributing peer reviewed scholarship to academic venues.

  Over the last year, C4AI has launched 15 open science initiatives, supported the publication of more than 25 papers, and established a global community of over 1,500 independent researchers across 107 countries. C4AI also democratizes access to resources for the research community through the Research Grant Program. To help

[1] Prior to founding Cohere, Aidan Gomez, Cohere's CEO, was part of the Google Brain research team, and one of the co-authors of the groundbreaking academic paper, that introduced the "Transformers," the neural network architecture that underlies Cohere's product offerings, as well as systems such as GPT-4 and Anthropic's Claude model.

[2] Nils Reimers, Director of ML, built the sentence transformer - sBERT.

[3] Patrick Lewis co-authored the seminal *Retrieval Augmented Generation* paper.

promote diversity and preserve linguistic and cultural heritage, Cohere has committed to large scale open science initiatives like *AYA* that improve multilingual model performance. An up-to-date list of C4AI's research publications is available here.

With a focus on enterprises, Cohere is exposed to a wide range of use cases in AI deployment, and a diverse array of organizations implementing AI in their everyday business activities. This unique lens provides Cohere with a close-up view of how enterprises (and government agencies) are adopting AI to drive innovation within their own organizations and in their industries. We welcome the opportunity to share these insights with policymakers at this Forum and offer ourselves as an experienced industry resource to US policymakers seeking to understand the ways in which AI is impacting businesses more broadly but also within specific domains (e.g., healthcare, financial, educational, retail and public sectors).

Encouraging a diverse AI ecosystem is critical to ensuring innovation in AI

As an independent and cloud agnostic LLM developer, Cohere welcomes this Forum's focus on how US policymakers can maintain a broad and diverse foundation model development ecosystem and promote market competition, to ensure AI companies continue to invest in innovation.

To promote innovation and competition, take advantage of the tremendous economic and productivity gains the adoption of AI technologies are likely to unleash, and ensure a beneficial integration and adoption of advanced AI systems across the economic spectrum, US policymakers should consider:

- Encouraging federal agencies and departments to procure AI technologies from diverse industry participants, especially emerging technology innovators, rather than incumbents from previous technology development cycles. As with every previous technology cycle, the innovators are rarely the incumbents - with the dawn of the internet in the mid-90s, Google wasn't even founded until 1998, and few of the tech incumbents of those days are even around anymore - think about Nokia, Lucent, Nortel - let alone members of the "trillion-dollar club" today. The Transformer Architecture was invented at Google Brain, but all 7 authors of the seminal "*Attention is all you need*" paper have since left Google to found startups, including Cohere.
- Developing pilot programs that incentivize public sector adoption of emerging technology innovators' products and services, and requiring public sector bodies to disclose how their AI procurement supports a diverse and vibrant AI-ecosystem. Such mechanisms would further prevent regulatory capture that usually place larger burdens on newer entrants.
- Providing incentives (e.g., tax-credits) to businesses that procure AI solutions from emerging market participants to encourage private-sector adoption of offerings from those who may become, with such support, the next generation of innovators, creating economic value and jobs in the United States.
- Adopting a sector-specific, principles-based approach to AI policy, and ensuring that appropriate regulations, compliance obligations and enforcement is performance-based, not technologically prescriptive, and proportional to scale and sophistication.

Driving innovation and fostering responsible development

Cohere believes that generative AI will profoundly transform our society by increasing the productivity of individuals and organizations, and spurring a myriad of new innovations, products, and services across the global economy. We also recognize that there are significant risks associated with AI systems that can have broad social impact, such as bias, fairness, reliability, toxicity, privacy and security, and other limitations.

However, spending our time and resources stoking existential fear of AI has been a distraction - one might argue a convenient distraction for incumbents - from the many very real risks similar to those we faced with social media, that have the potential to be exacerbated by AI.

As AI and large language model (LLM) technology moves from a consumer novelty to a core part of everyday business for millions of Americans and others around the world, there are a set of challenges that need to be addressed. These include protecting sensitive data, mitigating bias and misinformation, and knowing when to keep humans in the loop for oversight.[4] We believe the US government can help re-calibrate the current AI-safety discussion by encouraging innovation around imminent and medium-term AI risk-identification and mitigation by:

- Supporting and funding outcomes-based research that addresses issues related to imminent AI risk development and deployment.
- Incentivizing and supporting greater coordination between academic researchers and industry AI-labs throughout the technology development lifecycle. For example, currently C4AI supports academic researchers and paid scholar programs to do responsible AI research. However, larger-scale programs and support from the government would be welcomed. Other possible mechanisms that should be considered include joint research programs and national compute policies that support open science initiatives.
- Fostering standardized protocols for how foundation models are assessed for risk. There is a possible role for the US government to play in fostering further efforts towards holistic evaluation of model quality along specific dimensions (for example, accuracy, robustness, fairness, and efficiency). In particular, an effort on benchmarks which can be applied in a standardized and replicative manner are crucial for model deployment. Eventually, there may be a role for regulatory agencies to play in formally evaluating model quality as a part of regulatory oversight or enforcement, but the space is still immature and requires deeper investment and attention by a broad range of stakeholders spanning government, industry, and academia.

[4] We strongly believe that conversations overwhelmingly centering on the existential risk that systems pose is a distraction from the set of challenges that need to be addressed to foster responsible adoption of Generative AI technology. You can read more in a recent piece authored by Cohere's CEO, Aidan Gomez.

Encouraging a technological-centric approach to addressing AI-risks

Cohere believes that technological solutions will be instrumental in mitigating both current and future risks associated with AI-adoption and use, including issues of transparency, toxicity, fairness, and hallucinations. Consequently, the US government should ensure its research and incentives programs are appropriately encouraging such a technological-based approach to AI risk evaluation and mitigation.

Cohere's own product-development experience is instructive. Recently, Cohere released Coral, an "enterprise-knowledge assistant" designed specifically for business use. Coral enables knowledge workers across industries to receive responses to requests specific to their sectors, based on their proprietary company data. This offering expands well beyond publicly available generative AI tools, improving the accuracy and relevance of generative AI responses by adding data that was not used in its original training through Retrieval-Augmented Generation (RAG). RAG also enables Coral to mitigate the hallucination problems by providing grounded responses that have citations to proprietary, internal company data as well as publicly available sources enabling users to verify the veracity of Coral's response.

Conclusion

We look forward to engaging members of the U.S. Senate AI Insight Forum and discussing the most imminent and critical risks of generative AI, effective strategies and standards needed to ensure we maximize innovation while mitigating potential harm, and watching out for the emerging dominance of the 'Big Tech' oligopoly, to address the challenges ahead with scaled adoption of this technology.