FULL COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY
HEARING CHARTER

"Artificial Intelligence: Advancing Innovation Towards the National Interest"

Thursday, June 22, 2023
10:00 a.m.
2318 Rayburn House Office Building

Purpose: This hearing will discuss ways the federal government can advance artificial intelligence (AI) in a trustworthy and beneficial manner for all Americans. The hearing will provide an analysis of the technology's upsides and risks, as well as examine efforts in academia and industry to promote innovation, establish proper standards, and build the domestic AI workforce.

Witnesses

- Dr. Jason Matheny, President & CEO, RAND Corporation
- Dr. Shahin Farshchi, General Partner, Lux Capital
- Clement Delangue, Co-founder & CEO, HuggingFace
- Dr. Rumman Chowdhury, Responsible AI Fellow, Harvard University
- Dr. Dewey Murdick, Executive Director, Center for Security and Emerging Technology

Overarching questions

- How can Congress shape the development of artificial intelligence to maximize its benefits for the American public?
- What role does the federal government have in filling private sector R&D gaps to improve the trustworthiness of AI systems?
- What are the opportunities and risks of open-source AI systems?
- How can Congress invest in STEM education to develop the workforce needed to maintain US leadership in AI?
- How will AI change the demand for skilled labor and how can Congress prepare the U.S. workforce to adequately meet this shift?
- What are the consequences of falling behind China in fundamental research or commercialization of AI technologies?

BACKGROUND

History

Artificial intelligence (AI) refers to computer systems capable of performing tasks that typically require human intelligence, such as decision-making or content creation. Despite its recent popularity, it is not a new technology. "Narrow AI," or AI that targets a limited set of tasks, has been widely deployed for decades in various applications like automated warehouse robots, social media recommendation algorithms, and fraud detection in financial systems.

Although the term "artificial intelligence" was first coined in 1955, the field progressed slowly until the "machine learning" (ML) approach was popularized in the 2000s, a shift enabled by the proliferation of data on the Internet.[1] Unlike older AI systems which were pre-programmed to follow set rules, ML uses mathematical algorithms to learn patterns in data to make classifications or predictions.

In 2012, ML algorithms saw a dramatic leap in performance following advances in cloud computing, data storage, and "deep learning," a subfield of ML based on neural networks.[2] Neural networks use millions of nodes or "artificial neurons" to analyze large datasets that previous architectures could not handle. Training AI systems on advanced chips known as graphic processing units (GPUs) also led to large performance boosts.[3] Five years later in 2017, researchers at Google developed the "transformer" model, a new deep learning architecture that could take knowledge learned from one task (e.g., object recognition in images) and apply it to another task (e.g., generating new images).[4]

The ability of transformers to transfer 'knowledge' between tasks paved the way for models with "general" capabilities, resulting in an explosion of AI platforms that exhibited far greater reasoning and creativity than their predecessors. Recognizing this development, Stanford researchers popularized the term "foundation models" in 2021, highlighting these new models' foundational role for building next-generation AI applications.[5] Foundation models form the basis for "generative AI" — models that can create sophisticated writing, images, and other forms of content with minimal human input. For instance, GPT-4, the foundation model powering ChatGPT, can compose poems, score in the 90th percentile on the bar exam, and even write code to control robots.

The AI Supply Chain

Data: A key determinant of an AI system's performance is the quality of data the model is trained on. Using a car as an analogy, data is akin to oil and the model is like the engine. A high-quality dataset is a relevant, accurate, and comprehensive sample of the target domain.

---

[1] Leopold, G. (2016, June 16). *Proliferation of data driving machine learning*. Datanami.
[2] *AlexNet and ImageNet: The Birth of Deep Learning*. (n.d.). Pinecone.
[3] Fernández, P., et al. (2022). *The transformational role of GPU computing and deep learning in drug discovery*. Nature Machine Intelligence.
[4] Vaswani, A. (2017, June 12). *Attention Is All You Need*. arXiv.
[5] Bommasani, R. (2021, August 16). *On the Opportunities and Risks of Foundation Models*. arXiv.

Models: AI models are the underlying algorithms that are trained on data. Training advanced models with poor-quality data and training subpar models with high-quality data will result in inadequate performance. Different tasks require different types of models; for example, "large language models" (LLMs) have recently become popular for generative text applications such as ChatGPT.

Compute: As models have grown in algorithmic complexity and require increasing amounts of data, the energy and computational resources (processing power, memory, and storage capabilities) needed to train and run AI models efficiently has increased as well.[6] It is estimated that OpenAI's GPT-2 system (released in 2019) cost approximately $43K to train, GPT-3 (2020) cost $4M, and GPT-4 (2023) cost over $100M.[7] Access to compute is a significant bottleneck in the industry, as even the largest companies like OpenAI are struggling to find enough processors.[8] The advanced chips used for training today's most advanced AI systems are almost exclusively fabricated by TSMC in Taiwan,[9] though the CHIPS program to incentivize the construction of commercial fabrication facilities in the U.S. should help diversify this supply chain in the future.[10]

Workforce: Maintaining the U.S.'s continued leadership in AI requires a strong workforce capable of producing top research and applications in all levels of the AI stack—from consumer products down to microchips. This encompasses computer scientists, software and semiconductor engineers, cybersecurity experts, and workers in all industries who are trained to use AI tools to augment their performance.

Capital: Ensuring continuous capital flows into the sector is a crucial component for continued innovation, especially given the increasing costs of developing frontier AI systems. In 2022, private AI investments in the U.S. totaled $47.4 billion, which was 3.5 times higher than the next highest country, China ($13.4 billion).[11]

Open-source tools:
All software systems, including AI, fall on a spectrum of open to closed.[12] Many of the field's advancements can be credited to the open-source AI ecosystem, which refers to the community, projects, and resources in AI that are freely available and open for modification by anyone. Common open-source resources include code libraries, datasets, research papers, and pre-trained models—models that can be downloaded and used as-is without additional training. These resources foster innovation by lowering the financial costs and technical barriers to developing and deploying AI systems.

---

[6] *AI and Compute*. (n.d.). OpenAI.
[7] Knight, W. (n.d.). *OpenAI's CEO Says the Age of Giant AI Models is Already Over*. Wired.
[8] Seetharaman, D., & Dotan, T. (2023, May 29). *The AI Boom Runs on Chips, but It Can't Get Enough*. Wall Street Journal.
[9] Toews, R. (n.d.). *The Geopolitics of AI Chips Will Define the Future of AI*. Forbes.
[10] *Notice of Funding Opportunity: Commercial Fabrication Facilities*. NIST.
[11] Maslej, et al. (April 2023). "*The AI Index 2023 Annual Report*." Stanford Institute for Human-Centered AI.
[12] Solaiman, I. (n.d.). *Generative AI Systems Aren't Just Open or Closed Source*. Wired.

The AI research community has traditionally published most of its research open-source, even at corporations like OpenAI and Google. However, these companies have begun taking an increasingly closed approach in recent years as the capabilities of advanced models have increased exponentially, citing potential large-scale misuse as the primary reason.[13] For example, OpenAI released the open-source version of its flagship GPT-2 model in 2019 but has kept its subsequent versions GPT-3 and GPT-4 closed.

ECONOMIC IMPACT

Major players

Due to the high computational costs of conducting large-scale AI research, which has increased 300,000x since 2012 by OpenAI's estimation,[14] the majority of frontier AI R&D is conducted by a handful of well-capitalized entities in the private sector, including OpenAI, Anthropic, Google, Microsoft, and Meta. A recent PwC study estimates that AI will contribute $3.7 trillion to the North American economy in the next six years.[15]

Applications

AI is unique among emerging technologies in its wide-ranging applications in all sectors of the economy. Many experts have likened it to 'electricity' — just as electricity upended transportation, manufacturing, health care, agriculture, and more, AI is already doing the same. Below are a few examples of recent advances AI has made in key U.S. industries:

National security: AI is increasingly being applied to defense technologies, providing significant advancements and capabilities. In intelligence-gathering applications, AI-powered systems can detect patterns, identify potential threats, and provide real-time situational awareness by analyzing vast amounts of data from satellite imagery, sensor networks, and other sources. For example, Palantir, a defense contractor, recently launched a ChatGPT-style digital assistant that enables operators to efficiently deploy reconnaissance drones, devise tactical responses, and orchestrate enemy communication jamming.[16] AI is also being applied in autonomous systems such as drone swarms, enabling efficient monitoring of critical infrastructure and combat zones while reducing human risk.

Biology: Google DeepMind's AlphaFold, first released in 2020, is an AI system that analyzes a protein's amino acid sequence to predict its 3D structure, which determines its function. The algorithm mapped nearly 98.5% of protein 3D structures when scientists only knew the structures for about 17% of the roughly 20,000 human proteins before AlphaFold.[17] This advancement has fueled new research in areas ranging from drug discovery to bioengineering that was previously impossible. A recent report by

[13] Chavez, P. (2023, May 30). *An AI Challenge: Balancing Open and Closed Systems.* CEPA.
[14] *AI and Compute*. (n.d.). OpenAI.
[15] *PwC's Global Artificial Intelligence Study: Sizing the Prize*. (n.d.). PwC.
[16] Sofrep. (2023, May 13). *Palantir Debuts Revolutionary Artificial Intelligence Platform for Military Decision-Making.* SOFREP.
[17] Toews, R. (2021, October 3). *AlphaFold Is The Most Important Achievement In AI—Ever*. Forbes.

Morgan Stanley estimated that the time and cost savings that AI could bring to the medicine creation process could result in an additional 50 novel therapies over a 10-year period.[18]

Designing advanced semiconductors: Some organizations are using reinforcement learning (RL)—a subfield within AI that uses a "trial-and-error" approach to learn and analyze unstructured environments—to design better and faster chips. In 2022, an Nvidia team was able to train an RL model to design chip circuits that were smaller, faster, and more efficient than the circuits designed by traditional electronic design automation tools. One of Nvidia's latest categories of chips, the Hopper GPU architecture, has over 13,000 instances of AI-designed circuits. Synopsys, a startup, has reached 100 commercial deployments of its AI-chip-design technology, which was able to reduce the size of one of its designs by 5%.[19] Since chips are produced in the tens or hundreds of millions, a 5% silicon area reduction can mean considerable cost savings.

Workforce

The widespread economic impact of AI will have large implications for the workforce. A new study by Goldman Sachs estimated that two-thirds of occupations could be partially automated by AI.[20] While some jobs are likely to be fully automated away, new jobs will also be created and other jobs will be augmented in ways that are difficult to predict. This has been the pattern for every new technology. For example, the introduction of ATMs in the 1970s was forecast to end the job of traditional bank tellers. While there was minor job loss, most bank teller jobs shifted primary responsibilities from counting money to managing client relationships.[21]

Data shows the demand for AI-related professional skills increasing across almost every American industrial sector. A 2023 Stanford report found that across every U.S. sector for which there is data (except for agriculture, forestry, fishing, and hunting), the number of AI-related job postings has increased on average from 1.7% in 2021 to 1.9% in 2022.[22] Employers in the United States are increasingly looking for workers with AI-related skills. In addition, companies are increasingly adopting AI—the proportion of companies using AI in 2022 has more than doubled since 2017.[23]

RISKS

Although AI systems have the potential to significantly improve our lives, they also have the potential to do significant harm. While risks to any type of information-based system also apply to AI (e.g., privacy, cybersecurity, and safety concerns), these systems also create a unique set of dangers that require special

[18] *How AI Could Speed Drug Discovery*. (n.d.). Morgan Stanley.
[19] Ward-Foxton, S. (2023, February 13). *AI-powered Chip Design Goes Mainstream*. EE Times.
[20] *Generative AI Could Raise Global Gdp by 7%*. (n.d.). Goldman Sachs.
[21] *Learning by Doing.* (n.d.). Yale University Press.
[22] Maslej, et al. (April 2023). "*The AI Index 2023 Annual Report*." Stanford Institute for Human-Centered AI.
[23] Ibid.

attention. AI systems used in the real-world are already sufficiently advanced to cause immediate harm, and given the rapid pace of technological development, long-term and structural risks are also present.

Explainability and interpretability

Advanced AI systems are functionally black boxes, which means we cannot explain or interpret how they reach the decisions that they do. For instance, deep learning models, which most generative AI systems are built on, use thousands or millions of interconnected nodes to make complex calculations and arrive at an output. Humans can only observe the inputs and outputs of the system; what happens during calculation-time is a mystery.

False information

Generative AI systems often produce "hallucinations"—false information that seems credible. They can occur when users request information not in the training data or when models fail to "learn" the underlying dataset correctly. A malicious actor could use this flaw to create inaccurate or misleading information in disinformation campaigns. Below is an example AI-generated image of Mahatma Gandhi taking a selfie:[24]

Computational scarcity

Training AI systems requires a large amount of computational power. An analysis by OpenAI found that the amount of compute required to train the current most advanced models grew 300,000,000% from 2012 to 2017.[25] As a result, talent and cutting-edge innovation are increasingly concentrated in a handful of large companies that can afford the high computation costs. This Committee also finds that the

[24] Sukheja, B. (2023, March 21). *Artist Uses AI To Generate Selfies From The Past*. NDTV.
[25] *AI and Compute*. (n.d.). OpenAI

computing resources available to federal agencies are in scarce supply, oftentimes with 3-4x more demand for them than what is available.

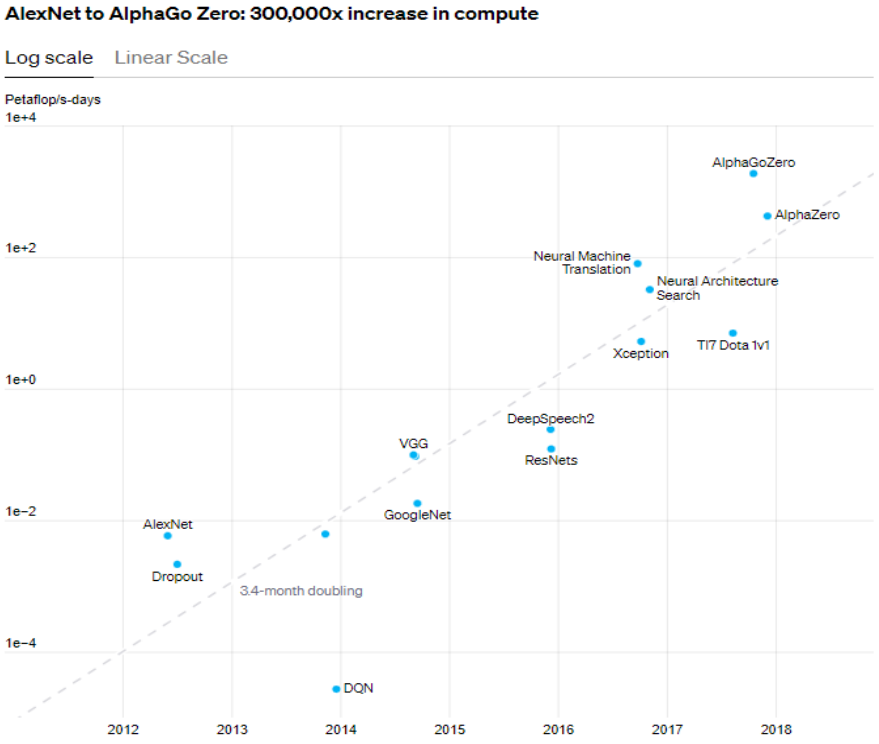

[26]

Harmful bias

There are three types of bias according to the National Institute of Standards and Technology (NIST).[27] First, systemic biases result when AI systems discriminate against certain groups while disadvantaging others. Second, statistical and computational biases arise from the AI system being trained on a dataset that is not representative of the population. These biases often arise when algorithms are trained on one type of data but cannot extrapolate beyond that. Lastly, human biases reflect systematic errors in human judgment. These biases are often implicit and tend to relate to how an individual or group perceives information (such as the output of an AI system) to make a decision or fill in missing or unknown information. Since AI systems are designed by humans, systematic bias is present across the entire AI lifecycle and in the use of AI applications once deployed.

Long-term risks

A growing base of AI experts fear the technology can one day supersede human interests. The Future for Life Institute notably released a letter calling on "AI labs to immediately pause for at least 6 months the

[26] *AI and Compute*. (n.d.). OpenAI

[27] Reva Schwartz et al., "*Towards a Standard for Identifying and Managing Bias in Artificial Intelligence*," NIST, March 2022

training of AI systems more powerful than GPT-4."[28] The letter garnered signatures from Elon Musk, Steve Wozniak (Apple co-founder), and other technology researchers and executives.

THE GLOBAL RACE TO LEAD IN ARTIFICIAL INTELLIGENCE

There is a global consensus that AI will be transformative across the scientific, economic, and defense realms. For that reason, there is a global race among great powers to lead in both fundamental AI research and commercial applications. The U.S. remains the leader in fundamental research and consistently produces cutting-edge AI applications such as ChatGPT before other nations, but adversarial nations like China have taken the lead in narrow fields such as facial recognition.

China

China is the second global leader in private AI investments behind the U.S., totaling $13.4 billion behind the U.S.'s $47.4 billion.[29] Although this private investment gap reflects the U.S.'s lead in R&D, China is closing the gap through AI industrial policy which invests billions through state-financed investment funds, designates "national AI champions," and provides preferential tax treatment to grow AI startups.[30] [31] By many metrics, China has caught up or surpassed the U.S. in research and commercial capabilities. For instance, nine of the top ten universities ranked by number of AI papers published in 2021 were Chinese (the 10th was the Massachusetts Institute of Technology). Chinese-published papers also received nearly the same share of citations as US researchers (22% vs. 24%) and China installed more automated industrial robots than the rest of the world combined in 2021.[32] China also officially implemented a K-12 AI curriculum and is on track to produce nearly 2x more STEM PhDs as the U.S. by 2025. A recent report by CSET found that U.S. investors played a key role in fueling China's AI rise, accounting for nearly one-fifth of all investments in Chinese AI companies from 2015 to 2021 which totaled $40.2 billion, or 37% of the total amount raised during the six-year period.[33]

United Kingdom

In March 2023, the UK announced a £3.5 billion investment in advancing science and technology.[34] The package includes £900 million for a new supercomputer with a portion dedicated to an AI Research Resource, similar to the U.S. National AI Research Resource (NAIRR) proposed by the NAIRR Task Force in January. One month later, the UK announced an additional £100 million for foundation model research and commercialization.[35] The UK has also declared a "pro-innovation approach to AI

---

[28] *Pause Giant AI Experiments: An Open Letter*. (n.d.). Future of Life Institute.
[29] Maslej, et al. (April 2023). "*The AI Index 2023 Annual Report*." Stanford Institute for Human-Centered AI.
[30] *Understanding Chinese Government Guidance Funds*.. (n.d.). *Center for Security and Emerging Technology*.
[31] *China Creates National New Generation Artificial Intelligence Innovation and Development Pilot Zones.* (n.d.). Center for Security and Emerging Technology.
[32] Maslej, et al. (April 2023). "*The AI Index 2023 Annual Report*." Stanford Institute for Human-Centered AI.
[33] Emily S. Weinstein and Ngor Luong, "*U.S. Outbound Investment into Chinese AI Companies*" (Center for Security and Emerging Technology, February 2023).
[34] *Government Commits Up to £3.5 Billion to Future of Tech and Science*. (n.d.). GOV.UK.
[35] *Initial £100 Million for Expert Taskforce to Help Uk Build and Adopt Next Generation of Safe AI*. (n.d.). GOV.UK.

regulation"[36] and designated an official technology diplomat to the U.S., who recently visited Silicon Valley executives in June 2023.[37]

European Union

The European Parliament, a main legislative branch of the E.U., passed the draft version of the E.U. AI Act on June 13th, 2023. The legislation mandates use and development requirements based on classifying AI systems by risk. The bill also introduces constraints on broad applications and processes, such as strongly curtailing uses of facial recognition software, and requires companies to publish summaries of copyrighted material used for training generative AI systems.

Russia

While Russia has used AI-enabled autonomous weapons in Syria and Ukraine, it lags far behind the world in research and commercial output. A recent study by Stanford found that Russia only had 3 authors on significant machine learning papers in 2022, compared to the US's 285 and China's 49. The study also found Russia produced one 'significant' system in 2022 compared to the US's 16, the UK's 8, and China's 3.[38]

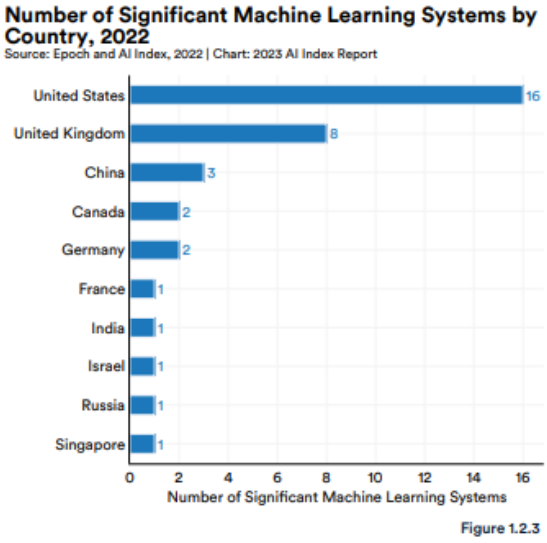

Figure 1.2.3

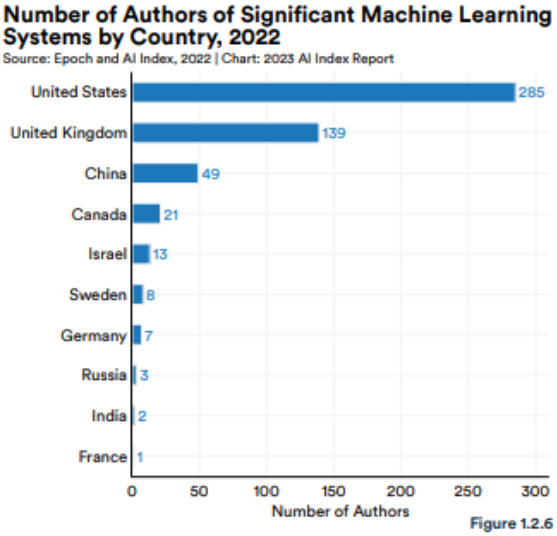

Figure 1.2.6

FURTHER READING

Artificial Intelligence Basics

- CRS - Artificial Intelligence: Background, Selected Issues, and Policy Considerations
- CRS - Overview of Artificial Intelligence
- IBM - Artificial Intelligence Basics

[36] *AI Regulation: A Pro-innovation Approach*. (n.d.). GOV.UK.
[37] Bordelon, B. (2023, June 14). *The British Diplomat Trying to Win Over the U.S. Tech Industry*. POLITICO.
[38] Maslej, et al. (April 2023). "*The AI Index 2023 Annual Report*." Stanford Institute for Human-Centered AI.

Generative AI

- GAO - Science & Tech Spotlight: Generative AI
- CRS - Generative Artificial Intelligence and Data Privacy
- CRS - Generative Artificial Intelligence: Overview, Issues, and Questions for Congress
- Stephen Wolfram - What is ChatGPT doing and why does it work?
- a16z - Who owns the generative AI platform?

Trustworthy AI and risks

- NIST - AI Risk Management Framework
- Google, OpenAI, Berkeley, Stanford - Concrete Problems in AI Safety
- Center for Strategic and International Studies - The Path to Trustworthy AI

National Security

- CRS - Artificial Intelligence and National Security
- CRS - Defense Primer: Emerging Technologies
- CRS - Deep Fakes and National Security
- The National Security Commission on Artificial Intelligence Report

AI in China

- Center for Security and Emerging Technology (CSET) - China AI-Brain Research
- CSET - Guidelines for National New Generation Artificial Intelligence Innovation and Development Pilot Zone
- Dr. Jeffrey Ding - China's AI National Team
- Tech in Asia - Chinese voice recognition firm iFlytek gets $407m from state investors
- CCP Draft Guidance for Generative AI