FPOV's Response to the "Pause Giant AI Experiments" Open Letter

Introduction

The Future of Life Institute wrote an open letter titled "**Pause Giant AI Experiments**"[1] on March 28, 2023, as a means to "*call on all AI labs to immediately pause for at least six months the training of AI systems more powerful than GPT-4.*" The letter itself had an option to digitally add your name to the list and sign the letter.

Only a few days after its release, the letter had over 2,000 signatories. The list of signatories includes some heralded names in the echelon of technology innovation such as Elon Musk (CEO of SpaceX, Tesla, Twitter), Steve Wozniak (Co-founder of Apple), Emad Mostaque (CEO of Stability AI), Christopher Reardon (Head of Design at Meta), and many others in the world of academia.

No one at FPOV is signing this letter…

The last time we had a significant fracture in the scientific world, **by those that were the ones advancing the disruptive technology itself,** based on different ideas of moral and ethical responsibilities was the development of the first nuclear bomb as part of the Manhattan project in 1944.

Does this AI letter calling for a pause in development constitute a modern-day version of the Manhattan project? Was it written as a result of questioning the moral and ethical responsibility with a highly disruptive technology? What does the letter say and is the development of giant AI the disruptive equivalent of the unlocking of nuclear energy? Is this just a way to slow down the current leaders so others can figure out how to catch up?

In this **FPOV**iew, we will outline the issues addressed in the letter, deliver our response, and provide our outlook over the next 6 months.

We have been heartened that the discussion on the dangers of AI have been going on for some years now.

We believe that society needs to tread carefully with the development of AI. We have not had the best track record in the past with forecasting the direction technologies will go. An example of this is the rise in distracted driving caused by mobile device usage. This is killing a large amount of people every month and interestingly, no letter has been written about taking a pause on allowing mobile device usage in vehicles.

---

[1] Future of Life Institute. "Pause Giant AI Experiments: An Open Letter". https://futureoflife.org/open-letter/pause-giant-ai-experiments/

Six Months? - *April to October 2023*

The "Pause Giant AI Experiments" letter itself is asking for a six month pause. But really, what is going to change in six months? Most governments are not going to move fast enough to enact legislation in that time, and the technology companies in leading competitive positions are not going to slow down at all given that there is a financial advantage by continuing to develop more powerful systems.

Countries from around the world, especially China, must be amused at the idea of this pause and would sign the letter only in bad faith or cheer its advancement on from the sidelines. Not one of them will pause any AI effort being developed in their country because of the geopolitical risks in doing so.

Rather than use a letter that plays upon fear of an unknown, why not forecast where the technology is likely to be from a giant use case over the next six months. Will we reach a point, as the letter asks, that *nonhuman minds outnumber, outsmart, obsolete, and replace us? Will AI cause a loss of control over civilization?*

FPOView

Over the Next 6 Months

- Costs to maintain and scale the current AI systems will remain high.

- The margin for companies offering the use of AI systems remains small or are not profitable.

- Spending by large cap firms on AI systems will increase as they position AI as the new driver of growth and consolidate computing power.

- AI references in earnings calls and quarterly reports will continue to increase.

- Individual investors will continue to be left out of individual investment opportunities with AI based ETFs underperforming.

- Historical data to ingest by the systems will increase at least by 20-40%, which will provide an opportunity to increase AI system accuracy.

- Data science and the algorithms used within AI systems will allow for incremental efficiency improvements.

- The costs to re-train the systems based on the new data availability remains high and is a limiting factor in iterating for accuracy.

- Computational power for large models remains constrained. Nvidia makes most of the chips used for AI systems. Currently, they are reaching a plateau with chip manufacturing improvements slowing with per-transistor costs. (*Meaning they are not getting the same efficiency gains with each new chip.*)

- Proactive businesses will establish more defined transformation objectives through the technology in this first AI inclusion wave.

- Business using AI systems will find marginal gains in efficiency until a full strategy is developed with leadership using guideposts for adoption and detailing advanced investments in infrastructure.

The "Letter"

Statements from the "Pause Giant AI Experiments" Letter	FPOView
"AI labs [are] locked in an out-of-control race to develop and deploy ever more powerful digital minds that no one – not even their creators – can understand, predict, or reliably control."	Before asking about the race to create digital minds, it can be important to find a balance by asking questions like *how are we going to use the systems and who is going to use the systems* **along with** *what happens if an AI system is wrong and are the systems capable of AI hallucinations?* Then we must train people on how to use the systems effectively? For example: • Chat GPT-3 had an accuracy rate of between 70% - 80%. • Chat GPT-4 increased that to between 80% - 90%. • Medical AI systems for diagnosis are 90-94% accurate. • Automated driving systems are 99.999999% accurate. Provided that people understand when and how to use each system and what to expect, automated learning can be co-managed with a shared responsibility by developers and users.
"Should we automate away all the jobs, including the fulfilling ones?"	FPOV is running multiple assessments to identify the difference in the workplace "pre-AI" vs. "post-AI." Goldman Sachs believes that as many as 300 million jobs could be automated in this current wave of AI development. They stated that 18% of work globally could be computerized, with the effects felt more deeply in advanced economies than emerging markets. In the US, Goldman believes that approximately two-thirds of current jobs "are exposed to some degree of AI automation," and up to a quarter of all work could be done by AI completely. The US, other governments, and organizations themselves should be forecasting internal efficiencies, displacement, re-skilling, and unemployment. It will be important to determine if subsidies will be needed during a digital work shift. AI systems will likely be used in stages, first in a hybrid and co-working stage leading to a march toward more and more reliance on the technology. In all transitions prior to this (consider the industrial revolution as an example) there were increases in efficiencies, job opportunities, standard of living, and wages. There was also increased innovation that led to higher level of creativity and motivation. It is true that jobs will shift to new types of roles, those known and those soon to be identified, but this is not a valid reason to stop progress.

"AI labs and independent experts should use this pause to jointly develop and implement a set of shared safety protocols for advanced AI design and development that are rigorously audited and overseen by independent outside experts."

It is true that the risk related to the systems, based upon use, should be evaluated and appropriate and reasonable governance should be put into place. However, the letter calls for this to be fashioned at the governmental levels with the development of a new governing body that includes both government officials and AI developers.

It is unlikely that governments will be able to react that fast in any meaningful way. However, it will be important to establish governance through audits of data, algorithms, testing criteria, system oversight, and third party risk. It will also be important for senior leadership or boards to take responsibility for the way the technology is used in their own organization. Compliance will be necessary, like a SOX framework.

FPOV has already developed governance models that are effective in helping organizations comply with current and emerging pieces of legislation.

If we are going to ask questions, then what should we be asking?

1. *What is it we really hope to accomplish by slowing things down? Are we going to be any smarter in six months? When in the past have we tried this with success? Did we slow down the PC? The Internet? Cloud computing? What would have been the economic impact had we?*

FPOView

It has been shown throughout history that geopolitics and national interests will be prioritized over any scientific questioning of ethics. Therefore, pausing the development for any period of time within the US will create an opportunity for other nations with competing interests to catch up and or surpass the current capabilities.

During this current AI development phase, FPOV is forecasting growth against the level of reliance on technology, chips, rare earth minerals, etc. that the AI systems will need to continue expanding. We are working to determine if these developments reach the level of becoming a strategic national interest, what are the potential options to enhance or replace dwindling resources, and how we can ensure a sustained ability to compete in the presence of other global powers.

2. *Based on an actual forecast, how can we be better spend the next six months than implementing a moratorium?*

FPOView

We are better served spending the next six months identifying the specific problems we need to deal with and coming up with those plans and that legislation without interrupting commerce and progress.

Secondarily, to remain a leader in technological resources, forecasts will be needed to calculate the cost of scaling the AI systems to ensure adequate financial capacity. This should include the complimentary resources needed for system usage and the costs of workforce transition with specific milestones.

Businesses are best served by accelerating R&D activities to transition into full commercial use with the appropriate levels of governance.

Our View

The letter presents a case whereby, giant AI cannot be managed. However, using a thoughtful approach development can and should proceed alongside an increase in the development of reasonable governance.

AI systems will develop incrementally over the next 6-12 months for the reasons mentioned in the outlook provided in this document. What will be more important is to plan for when these innovations move from "radical innovations" or "breakthrough innovations."

Radical innovation is a type of innovation that combines the power of technology with a new business model. It is a concept that changes the relationship between customers and suppliers by displacing current products and services or by making new product categories. *Breakthrough innovation* is a specific significant technological advance that makes a large impact on the efficiency or cost of a given product, service, or process.

Timelines and events to prepare for include:

- The arrival of commercialized quantum computing on or about 2025. This will create a "breakthrough innovation" that increases AI capacity dramatically by removing the current chip plateau.
- The arrival of 6G on or about 2030. This will allow for a "radical innovation" for the IoT industry and allow for more data to be sent to the AI systems. It will increase the number of potential devices from 5G at 1 million connected devices per sq km to 10 million connected devices per sq km.

For more information on , see previous **FPOV**iews and information on AI including: Embracing Generative AI in Education and Artificial Intelligence is Already Transforming the Job Market. What Does this Mean for You? and Regulating AI: The European Union's AI Act.

Test your AI risk knowledge and get a free report with FPOV's AI Risk Review.

Follow us on Social Media for breaking updates and virtual events:
LinkedIn | Facebook | YouTube

Contact us for more information or for support in your AI development:
info@fpov.com.