

Resource Paper on Artificial Intelligence

FCNL Policy Committee

This document supplements Section III.2.12: Artificial Intelligence in FCNL's The World We Seek: Statement of Legislative Policy. It includes a brief, non-technical overview of current AI, frequently asked questions (FAQ) about the Policy Committee process, a list of additional resources for consultation, and a list of definitions

Overview of Artificial Intelligence Today

Why Should We Care

Current innovations introduced by Artificial Intelligence (AI) dramatically exceed the capacity of governmental, social, and cultural institutions to respond wisely, effectively, and promptly. While AI has produced creative solutions to complex problems, it can also pose existential risk to humans, other species, and the environment. AI is already influencing – both overtly and covertly – all areas of policy work.

While there is not currently a FCNL staff member prepared to lobby on AI, and it is not likely one will be hired soon, having an AI policy provides needed guidance and permission for them to collaborate with other organizations who are similarly concerned about the need to regulate this rapidly evolving and potentially very dangerous technology.

What is Artificial Intelligence:

There are multiple ways to define artificial intelligence, from the simple to the more complicated. For our purposes, we define artificial intelligence as when machines (computers) perform tasks that in the past have required human intelligence. Examples are understanding speech, translating between languages, playing games, proving mathematical theorems, programming computers, and participating in conversations.

Generative AI is a “black box system” for generating text, images, videos, or code. This technique produced AI chatbots like ChatGPT, Claude, Copilot, and Gemini, which can answer questions, summarize text, help compose meeting minutes, draft project outlines, explain concepts, write poems or even books, and perform many other textual, image, video, or coding tasks. These applications are quickly improving.

Why AI Now?

Over the last five years, artificial intelligence has become increasingly ubiquitous. In our society, AI chatbots are integrated or optionally available in many of the software applications typically used, such as word processors and spreadsheets. Large numbers of people use chatbots to help them with their work or their schoolwork.

AI's profitability has created significant competition among major tech companies, who are investing billions in AI. As they, and supporting companies like chip makers, seek to dominate the market, AI tools are being developed and brought to market without any meaningful intervention or oversight. Given AI's potential for public harm and well as good, government regulation is imminently needed.

The Positive Side of AI

There are many positive uses of artificial intelligence. Chatbots can generate ideas, create outlines, summarize meeting results, write meeting minutes, and suggest ideas for new projects. With the help of new software, AI can also create new artwork, music, and videos. AI can assist with language translations and communication across cultures and species and serve as information experts on websites, helping users find information quickly and effectively.

On a limited basis, and with the user's double-checking, they can also conduct research or create computer code. Some of the most exciting possibilities exist in the field of medicine, where AI can assist with medical diagnoses or help determine the most effective approach in medical research. Developing new hypotheses is also a possible advantage in other, non-medical research. A team won the Nobel Prize for creating a neural network that can predict how proteins will fold, a problem that had stymied researchers for decades.

The Negative Side of AI

There are also opposing sides to artificial intelligence. Some of these are connected to AI hallucinations, where the AI fabricates information in response to a user's query. Chatbots can also lie or refuse to answer a question. Alternatively, results can be nonsensical, misleading, or wrong. Due to this tendency, users must review any results they receive from an AI tool. But how can the user determine that a plausible, confidently stated answer is actually wrong, particularly when the AI system does not cite sources for the information that it presents?

Bias:

AI systems include embedded bias because the data from which they learn already contains biases. This, in turn, increases the likelihood of bias in the results of AI systems. Consider the task of approving bank loans. If previous bank loan decisions were partially made based on racial, religious, or gender discrimination, AI system decisions might continue to incorporate that bias in making future decisions. Other areas potentially influenced are hiring decisions and criminal justice.

Another source of bias is intentional bias, where the developers of the system explicitly favor answers that agree with a given political perspective. Most versions of the chatbot Grok have been accused of having such intentional bias, in spite of Elon Musk claiming that he is building a neutral, truth-seeking AI. Grok has praised Hitler, denied the Holocaust, ranted about “white genocide” in South Africa, and called itself “MechaHitler.”

Energy and Environmental Impacts:

Even basic requests for information from AI chatbots require immense computing power to generate, thus necessitating the construction of enormous data centers. These data centers then consume a considerable amount of land, energy, and water, creating additional carbon emissions and environmental problems for society.

Privacy, Confidentiality, and Copyright:

Not all of the data or datasets used to train AI models have been appropriately licensed, nor have builders always sought permission to use authors’ materials to train AI models. Confidential information is not always removed from data before it is used for training purposes and can later appear in generated text or images. Many authors have belatedly discovered that their materials were used without their permission to train AI models.

Job Displacement:

Private corporations and other entities’ use of AI already leads to economic dislocation and job loss. AI devices are used to code, or chatbots to answer customer questions and perform even limited types of analyses. This can devalue and replace certain types of human labor and undermine the work of creators, possibly disrupting some industries and creating additional barriers to customer service.

Disinformation

AI also significantly enhances the ability of almost anyone to increase disinformation, due to the growing ease of creating text, images and videos with AI. Deepfakes (fake images, audio, and

video) make it appear that a person said or did something that they did not say or do. We are already seeing a number of these deepfakes used in the political arena against public figures, political candidates, and elected officials. This usage means an additional threat to our democracy. Adding to this issue is the complete lack of accountability in finding who is responsible for misinformation, plagiarism, or other problems.

Autonomous Weapons:

Not least among our concerns is the potential use of AI's power in creating autonomous weapon systems, which can be operated without significant human oversight. Autonomous weapons are already being used, and those powered by artificial intelligence can create a risk of loss of control, creating security risks. As Quakers, many view this as the main reason why artificial intelligence should be significantly restricted in today's world.

Threats to Children and Vulnerable Populations:

As access to AI systems is not currently regulated, completely open access to these systems has a growing potential for harm to children and other vulnerable populations. The problems with uneducated use among these populations alone strongly underscore the need for regulation of AI devices. But these are just one area of the increasing challenges to operating AI systems safely and ethically, as they become more autonomous and as humans lose more control over their operations.

Existential Threats:

AI agents pose existential threats if they become more powerful than humans, in part because they lack human values. The AI research community has increasingly affirmed the reality of such threats, and significant numbers of its leaders judge there is a 10% or greater risk that AI may in fact threaten the existence of our species (e.g. an AI could develop and release a virus that kills humanity). Experts have stressed, further, that no one has yet been able to provide solutions to the human values alignment problems or even to demonstrate that solutions are possible. *"How do we react when half the field can't agree on what risks are real?"* There is a serious debate among AI researchers as to whether to pause or halt AI development until such solutions can be developed.

Regulating AI

There has been some regulation already; The EU has already passed an AI Act in 2024, although its success in regulating AI is yet to be seen. In the United States, no federal legislation has passed.

The significant challenges in regulating artificial intelligence are that many systems are black boxes. In addition, regulating AI is like trying to stop a moving train that has already left the station. The ability of AI computing to create new connections that were not part of the original programming also poses unique challenges. As Yuval Noah Harari concludes (2024, p. 200), “[AI systems] can learn by themselves things that no human engineer programmed, and they can decide things that no human executive foresaw.” Finally, in attempting to regulate AI, we are up against the power of the world’s most powerful corporations, which often seek to create AI and AI-based products without any interference.

Beyond the technical challenges to regulation are the difficulties in integrating ethics and values into such complex computer systems that will be operating more and more autonomously. Several efforts have been made to outline the ethics needed in AI regulation: the Rome Call for AI Ethics (<https://www.romecall.org/the-call/>) (Renaissance Foundation, 2020) and the Mozilla Petition: Call for Trustworthy AI (<https://www.mozillafoundation.org/en/campaigns/sign-the-petition-and-support-mozillas-call-for-trustworthy-ai/>) (Mozilla Foundation, ND).

The pace of rapidly evolving technology necessitates that FCNL stand with other faith groups and people of conscience as we work together for peace and justice. Current innovations introduced by Artificial Intelligence (AI) dramatically exceed the capacity of governmental, social, and cultural institutions to respond wisely, effectively, and promptly. While AI has produced creative solutions to complex problems, it can also pose existential risk to humans, other species, and the environment.

The stakes are very high. AI is already embedded in many critical functions of our daily living. There is **no going back to ground zero**. Regulation, as we have traditionally developed and applied it, is not adequate to this task. Our challenge awaits us.

References

Harari, Y. N. (2024). *Nexus: A Brief History of Information Networks from the Stone Age to AI*. Random House.

Mozilla Foundation. (ND, ND ND). *Sign the Petition: Support Mozilla's Call for Trustworthy AI*. Sign the Petition. Retrieved March 25, 2025, from <https://www.mozillafoundation.org/en/campaigns/sign-the-petition-and-support-mozillas-call-for-trustworthy-ai/>

RenAIssance Foundation. (2020, February 8th). *The Call for AI Ethics*. The Call. Retrieved March 25, 2025, from <https://www.romecall.org/the-call/>

Bakker, Karen. *The Sounds of Life: How Digital Technology Is Bringing Us Closer to the Worlds of Animals and Plants*. Princeton University Press, 2024.

Selected Definitions Related to Artificial Intelligence

Source: Many of these are selected from Article 3 of the European Union's Artificial Intelligence Act 2024

AI system means a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments; 2024)

Biometric data means personal data resulting from specific technical processing relating to the physical, physiological, or behavioural characteristics of a natural person, such as facial images or dactyloscopic data [information from fingerprints];

Deepfake refers to AI-generated or manipulated image, audio, or video content that resembles existing persons, objects, places, entities, or events and would falsely appear to a person to be authentic or truthful.

An emotion recognition system means an AI system for the purpose of identifying or inferring emotions or intentions of natural persons based on their biometric data

General-purpose AI model means an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications.

A general-purpose AI system refers to an AI system based on a general-purpose AI model, capable of serving various purposes both directly and through integration in other AI systems.

High-impact capabilities refer to those that match or exceed the capabilities of the most advanced general-purpose AI models.

Instructions for use means the information provided by the provider to inform the deployer of, in particular, an AI system's intended purpose and proper use;

Intended purpose means the use for which an AI system is intended by the provider, including the specific context and conditions of use, as specified in the information supplied by the provider in the instructions for use, promotional or sales materials, and statements, as well as in the technical documentation.

The performance of an AI system refers to its ability to achieve its intended purpose.

Post-market monitoring system refers to all activities carried out by providers of AI systems to collect and review experiences gained from the use of AI systems they place on the market or

put into service, to identify any need to apply necessary corrective or preventive actions immediately.

Reasonably foreseeable misuse refers to the use of an AI system in a manner that deviates from its intended purpose, but which may result from reasonably foreseeable human behavior or interaction with other systems, including other AI systems.

Risk means the combination of the probability of an occurrence of harm and the severity of that harm.

Safety component means a component of a product or of an AI system which fulfils a safety function for that product or AI system, or the failure or malfunctioning of which endangers the health and safety of persons or property; (Article 3 EU Artificial Intelligence Act 2024)

Systemic risk means a risk that is specific to the high-impact capabilities of general-purpose AI models, having a significant impact on the [national/international markets] due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain;

Training: Most modern AI systems have millions or billions of parameters that control how they operate. Training is inputting text, images, video, or audio to allow the system to learn values of those parameters that achieve the desired behavior.

Selected Resources

- Artificial Intelligence Legislation Tracker (<https://www.brennancenter.org/our-work/research-reports/artificial-intelligence-legislation-tracker>)
- AI For Nonprofits Resource Hub (<https://www.nten.org/learn/resource-hubs>(<https://www.nten.org/le/artificial-intelligence>)
- The Beginner's Guide to AI for Nonprofits (<https://www.daxko.com/insights/the-beginners-guide-to-ai-for-nonprofits>)
- Center for Democracy and Technology. 2025. Regulating Public-Sector AI: Emerging Trends in State Legislation. Available at <https://cdt.org/insights/regulating-public-sector-ai-emerging-trends-in-state-legislation/>
- Charting AI Regulations for Ethical Law Enforcement (<https://patimes.org/charting-ai-regulations-for-ethical-law-enforcement/>)
- *Coded Bias* (facial recognition, bias in algorithms). Documentary Film.
- Coleman, Jude. 2023. AI's Climate Impact Goes Beyond Its Emissions. *Scientific American*. Available at <https://www.scientificamerican.com/article/ais-climate-impact-goes-beyond-its-emissions/>
- Cox, Gray. 2023. *Smarter Planet or Wiser Earth? Dialogue and Collaboration in the Era of Artificial Intelligence* (also available as a PDF in Creative Commons at: www.smarterplanetorwiserearth.com). A video presentation of highlights is available as a West Falmouth Friends Second Sunday Forum on "A Quaker Approach to a Wiser Earth in the Era of Artificial Intelligence with Gray Cox" here: <https://www.youtube.com/watch?v=ICgrONK23RQ>
- Cox, Gray. A broader overview of the ongoing developments in AI, their social impacts, and ways Quakers might be led to frame and respond to them is provided in broadcast of the West Falmouth Friends Second Sunday Forum on "A Quaker Approach to a Wiser Earth in the Era of Artificial Intelligence with Gray Cox" available here: <https://www.youtube.com/watch?v=ICgrONK23RQ>
- Electronic Frontier Foundation (EFF). Transition Memo 2025. Pages 26-29. Available at <https://www.eff.org/press/releases/eff-sends-transition-memo-digital-policy-priorities-new-administration-and-congress>
- EPIC (Electronic Privacy Information Center) 2023. Outsourced and Automated: How AI Companies Have Taken Over Government Decision-Making. Available at <https://epic.org/new-epic-report-delves-into-state-ai-procurement/>.
- EPIC (Electronic Privacy Information Center). Generating Harms: Generate AI's Impact and Paths Forward. (I and II) Available at <https://epic.org/generating-harms/> and

<https://epic.org/wp-content/uploads/2024/05/EPIC-Generative-AI-II-Report-May2024-1.pdf>

- Farid, Hany. 2025. This Month in Generative AI: Supercharged Frauds and Scams. *Content Authenticity Initiative*. October 2025. Accessed at <https://contentauthenticity.org/blog/october-2025-this-month-in-generative-ai-supercharged-frauds-and-scams>
- Friends of the Earth. 2024. Report: Artificial Intelligence- A Threat to Climate Change, Energy Usage, and Disinformation. Available at <https://foe.org/news/ai-threat-report/>.
- Harari, Yuval Noah, Nexus: A Brief History of Information Networks from the Stone Age to AI. New York, Random House, 2024.
- Khlaaf, Heidy, and West, Sarah Myers. 2025. The Rush to AI Threatens National Security. *New York Times*. Available at <https://www.nytimes.com/2025/01/27/opinion/ai-trump-military-national-security.html>
- Lucas, Julian. 2025. Tim Berners-Lee Invented the World Wide Web. Now He Wants to Save It. *New Yorker*. September 29, 2025. Accessed at <https://www.newyorker.com/magazine/2025/10/06/tim-berners-lee-invented-the-world-wide-web-now-he-wants-to-save-it>
- Media Justice. 2024. Exposing Harms, Reclaiming People Power. Available at <https://mediajustice.org/news/take-back-tech-2024-exposing-harms-reclaiming-people-power/>
- Mulligan, Scott J. 2025. A New York Legislator wants to pick up the pieces of the dead California AI bill. *MIT Technology Review* January 9th, 2025. Available at <https://www.technologyreview.com/2025/01/09/1109875/a-new-york-legislator-wants-to-pick-up-the-pieces-of-the-dead-california-ai-bill/>
- Mulligan, Scott J. 2024. There are more than 120 AI bills in Congress right now. *MIT Technology Review*, September 18th, 2024. Available at https://www.technologyreview.com/2024/09/18/1104015/here-are-all-the-ai-bills-in-congress-right-now/?gad_source=1&gclid=CjwKCAiAtNK8BhBBEiwA8wVt94XcGu2WNC9v_pmSREtyVMey0sLjv-uwFXajAy9ybj-iif-cq6ZO2RoCHSkQAvD_BwE
- National Telecommunications and Information Administration. 2024. Artificial Intelligence Accountability Policy. Available at <https://www.ntia.gov/issues/artificial-intelligence/ai-accountability-policy-report/overview>

- O'Donnell, James. 2025. Here's our forecast for AI this year. *MIT Technology Review* January 14th, 2025. Available at <https://www.technologyreview.com/2025/01/14/1109958/whats-next-for-ai-in-2025-2/>.
- <https://share.evernote.com/note/ca0fb8bd-2655-4b6e-9eed-7d41e221e7a3>
- OECD AI Policy Observatory. 2025. How the G7's new AI reporting framework could shape the future of AI governance. (a reporting framework developed from the G7 Hiroshima AI Process). Available at <https://oecd.ai/en/wonk/how-the-g7s-new-ai-reporting-framework-could-shape-the-future-of-ai-governance>.
- "A Grave New Urgency? Toward a Quaker Approach to Artificial Intelligence and Related Technologies" Quaker Institute for the Future. These are preliminary draft materials that are intended to help Friends begin the processes of educating ourselves on issues posed by Artificial Intelligence systems and related technologies, and help us begin to articulate Testimonies and Queries concerning them, as well as priorities for legislation and other forms of action that we may be called to advance. These drafts were developed by a Circle of Discernment in the Quaker Institute for the Future, whose attendees and participants have included Welling Hall, Scott Drysdale, Lois Yellowthunder, Gray Cox, Keith Runyan, Diane Keefe, Walter Long, Marge Pifer, and Phil Arcuni. These are drafts of documents that are very much works in progress. The materials are available at: https://docs.google.com/document/d/1686igI8CSG_-Y_eIVWBKM02vF7NWbA0R1-PkS2oNBVQ/edit?usp=sharing
- Reisner, Alex. 2025. The Unbelievable Scale of AI's Pirated-Books Problem. *The Atlantic*. <https://www.theatlantic.com/technology/archive/2025/03/libgen-meta-openai/682093/>
- Runyan, Keith. 2023. Keith Runyan's presentation to the Quaker Institute for the Future in 2023 provides a very accessible introduction to key concerns around safety and existential threats and ways they might be addressed: <https://www.youtube.com/watch?v=Nfg4HRMj7ag>
- Sustainable by Design – Advancing the Sustainability of AI (<https://blogs.microsoft.com/blog/2024/04/02/sustainable-by-design-advancing-the-sustainability-of-ai/>)
- Thompson, Stuart A. 2025. How Trump Is Using Fake Imagery to Attack Enemies and Rouse Supporters. *New York Times*. October 21, 2025. Accessed at https://www.nytimes.com/interactive/2025/10/21/business/media/trump-ai-truth-social-no-kings.html?unlocked_article_code=1.vE8.sJh6.xBfeFf80bJIe&smid=nytcore-ios-share&referringSource=articleShare

- Tiku, Natasha, and De Vynck, Gerritt. *Washington Post*. 2025. Google drops pledge not to use AI for weapons or surveillance. Available at <https://www.washingtonpost.com/technology/2025/02/04/google-ai-policies-weapons-harm/>
- Witt, Stephen. 2025. The A.I. Prompt That Could End the World. *New York Times*. October 10, 2025. Accessed at https://www.nytimes.com/2025/10/10/opinion/ai-destruction-technology-future.html?unlocked_article_code=1.sU8.qbxf.WkMPfgA-pVJ3&smid=nytcare-ios-share&referringSource=articleShare

Frequently Asked Questions (FAQs)

We are making this FAQ available because these questions came up during listening sessions and the committee's deliberations. This includes both substantive questions about AI and about the Policy Committee's policy development process and how staff use the FCNL policy statement.

FCNL Process

1. Why is the proposed AI section longer than most of the other sections in the *World We Seek*? This new proposed section is longer than most others because AI is a new concept and technology to many, and requires additional explanation and context.
2. How did the proposed AI section end up being placed in the Part III of *The World We Seek*, rather than the others? The Policy Committee first considered integrating language throughout *The World We Seek* document. Ultimately, the committee is proposing the new language for Part III: We seek a community where every person's potential may be fulfilled. We decided much of AI's impact would be upon individuals and society, so it belonged there.
3. How will FCNL staff use this to lobby or advocate on the AI issue? FCNL staff can use the policy guidance being proposed here to join the many new coalitions being formed to address artificial intelligence, including new interfaith coalitions, and to join letters on the topic.
4. Will another staff person be hired to advocate on the AI issue? It is not likely a new staff person would be hired to advocate just on the AI issue, and currently, there is not a staff person who could lobby on the topic. It is more likely that FCNL staff would use this language to guide their participation in coalitions or letters.
5. If important legislation develops on AI that requires advocacy, could FCNL engage in advocacy, even if there is no dedicated FCNL lobbyist on the topic? Yes, if approved, this policy statement opens the way for FCNL to take advocacy action as the opportunity arises.

6. Will the proposed statement be added to this Congress' legislative priorities? At this time, it is being proposed as an addition to *The World We Seek*, and not to the legislative priorities of the current Congress.
7. Does this statement imply that FCNL will lobby states with active AI legislation? At this time, FCNL primarily lobbies the U.S. Congress.
8. Why are we not just saying AI is bad and it should be abolished and not allowed to continue? The issues surrounding artificial intelligence and its usage are complex and require a great deal more nuance around multiple AI sub-issues.

Artificial Intelligence Questions

1. How would we use AI to help us discern truth? While AI can be used to process information, gather and sort information, and generate ideas, Quaker discernment of the truth is centered around group discernment of the Spirit's leadings. It's not clear at this point how AI would assist.
2. What are deepfakes, and why are they important when talking about AI? Deepfakes are images or videos that might look real but are, in fact, images manipulated or generated by artificial intelligence. Due to advances in software, it is quite easy for most people to create deepfakes, and they can be used to manipulate public opinion, elections, marketplaces, and other social processes. Deepfakes have already been responsible for harmful outcomes to individuals, particularly children and other vulnerable populations, and the potential for additional harm is tremendous.
3. What are some ways to safeguard truth, knowledge, and wisdom that undergird human civilization? Strategies for preserving truth, knowledge, and insight in the area of technologies like artificial intelligence should focus on governmental regulation, as is already happening in the European Union and some American states, education around AI, training on responsible uses of the technology, and empowerment of all voices in the community to take part in the creation of knowledge and the making of decisions.
4. How do we work to advocate for government action that safeguards the truth, knowledge, and wisdom undergirding human civilization? Educating ourselves on AI

and surrounding issues and helping to propose reasonable and effective policy laws and regulations. Global cooperation is critical.