



Evangelical
Lutheran Church
in America

CORPORATE
SOCIAL
RESPONSIBILITY

CSR

ARTIFICIAL INTELLIGENCE

Issue Paper

We need to ensure and safeguard a space for proper human control over the choices made by artificial intelligence programs: human dignity itself depends on it.

Pope Francis¹

BACKGROUND

Artificial intelligence (AI) and machine learning have been around for a long time, but in the past five years they have become nearly ubiquitous, thanks to faster computer chips and greater interest in their potential. There are potential benefits of AI, some of which we are already experiencing without even realizing we are being served by AI. There are also many concerns and worrying events that dictate caution and oversight when it comes to AI. Reactive and limited AI, whether through recommendations Netflix and Amazon give us, or through customer service chatbots, mitigate some concerns around privacy and data collection, but the rapidity of growth towards super intelligent AI systems raises the need to address the issue.

While there is no universal definition of AI, the World Intellectual Property Organization offers, via the UN Library: “AI is generally considered to be a discipline of computer science that is aimed at developing machines and systems that can carry out tasks considered to require human intelligence.”² The ELCA’s Information Technology department describes AI thus: “*Artificial Intelligence (AI) refers to the theory and development of computer systems that can perform tasks that typically require human intelligence, such as speech recognition, decision-making, and pattern identification. AI encompasses a broad spectrum of capabilities, from mimicking human actions and thought processes to acting and thinking rationally.*”³ It is classified with four descriptions: reactive, limited memory, theory of mind, and self-awareness.⁴ AI is developing so rapidly that what it can do and how it is being used as of this writing is bound to change within months, but what is described as AI comes with a range of impacts and capabilities, ranging from correcting grammar to “creating” artwork.

The benefit of AI use is that, once an AI system is trained, it’s able to sort through the onslaught of data we are collectively experiencing. Now that 90% of adults in the U.S. carry a smartphone⁵ and nearly every photo we have taken and conversation we have had is stored somewhere in the ether, sifting through too much information, rather than finding scarce information, becomes the challenge. AI’s proponents

Corporate Social Responsibility in the ELCA uses the tools of screening of investments, shareholder advocacy and community investing to work with corporations, calling them to ensure that people are treated fairly and with dignity and to create sustainable communities.

More at [ELCA.org/CSR](https://elca.org/CSR).

hope that on a large scale it will do things such as streamline medical service, choose diseases we are most likely to cure, and foster renewable energy. On a small scale, it could potentially increase workplace efficiency without replacing administrative workers.

The recent rapid pace of development and implementation of AI has presented complex ethical challenges and questions. The fundamental concerns around uses of AI raise questions of *transparency, responsibility, and scale*.

In terms of transparency, much of the work that goes into AI model learning is invisible, but vitally important. Should a user be notified any time AI is part of a process, no matter how insignificant that involvement might seem? What must be revealed about how a particular model of AI was trained? How can those whose data was used give their consent? Should we be concerned if consent has not been explicitly given?

Responsibility becomes vastly more complex. Who is responsible for the rights violations of citizens unlawfully targeted by law enforcement, or civilians targeted by the military? Are the writers of the initial programming responsible? Are those who oversaw the development of knowledge through AI responsible? Are those in management who made the decision to use AI responsible? Are the people closest to the actual effect responsible? If decisions such as allocating health care resources, or identifying targets for search by law enforcement, or choosing military targets are being informed by AI, or even made by AI, one error in the algorithm has manifold consequences for human rights, and the path to restitution of those rights is unclear.

The scale and pace of AI development, implementation and change make it close to impossible for legislators to formulate regulations that will safeguard human rights and protect vulnerable communities. As AI is put to use in novel situations, without sufficient oversight of its capabilities, unintended consequences arise. Legislators scramble to create boundaries and expectations after the fact. The Biden White House proposed a blueprint for an AI Bill of Rights that included five principles: safe and effective systems, protection from algorithmic discrimination, data privacy, notice and explanation, and availability of a human alternative. The Trump administration promptly rescinded that order, effectively removing roadblocks to AI development and announced that more resources would be developed to generate the power needed to continue to develop AI.

Legislation proposed by the EU details four levels of risk, from the unacceptable use of AI for real-time remote facial recognition, to heavily regulated high-risk uses such as insurance claims or credit-scoring systems, to limited risk such as chatbots. Even with uses such as real-time biometrics outlawed, rights groups such as Human Rights Watch have raised alarms⁶ about uses that will limit access to social goods as part of an effort to eliminate identity fraud, about failure to protect rights and about reinforcing of discrimination in the labor market. Many U.S. states have proposed legislation regulating the use of AI, with varying degrees of success on issues from deep fakes and images of child sexual abuse to artistic integrity.

RECOMMENDED by the
Advisory Committee
for Corporate Social
Responsibility, June 10,
2025

APPROVED by the Church
Council, October 2025
[CC25.10.42]

Benefits of AI

AI can be used to sort through massive amounts of data quickly and (often, but not always) accurately. It can be used to automate tedious tasks that require little thought or creativity. It could potentially lead to significant discoveries in scientific research,⁷ greater access to health care and education, and safer roadways. AI advocates promise that it can promote green energy by anticipating energy use and allocating resources more efficiently and can staff nursing homes more attentively and efficiently. AI is also increasingly integrated into our national defense system.⁸

Its potential benefits span a number of fields. In healthcare, generative AI is being used to develop new drugs, personalize treatment plans, and predict disease progression. The hope is to improve patient outcomes and streamline medical research. Possibly AI chatbots are commonplace in customer service at this point, providing even small organizations with customized customer service. AI implemented by manufacturing is expected to increase efficiency and strengthen quality control, while AI in banking can enhance fraud detection.

Risks of AI

Like any human endeavor, AI can fail to live up to its promises and can be abused and misused. Cigna Healthcare was recently sued⁹ when it was found to have used AI to deny more than 300,000 claims in less than two months. Though insurance companies are obliged legally to have doctors view case files and approve claims, doctors never even opened the files, which were handled at an average rate of one every 1.5 seconds. An AI app used by landlords was also implicated¹⁰ in rising rents and subject to a lawsuit from the Department of Justice for collusion.

While any new technology, especially one as ubiquitous as AI, is bound to experience hiccups, there are many factors in the process, training and development of AI that raise concerns.

Human rights: AI presents challenges to human rights on many levels. Discrimination based on race, class or gender that already exists can be codified and then magnified by AI,¹¹ particularly when resources are distributed dependent on judgments made by AI. Issues of privacy abound. As large language models are accumulating and testing knowledge, they have no way of distinguishing what should be considered private and what could be shared, and that information can be used negatively by many bad actors.¹² Copyright and compensation, particularly for images, raise problematics. Artists have raised important questions about copyright laws and the right to privacy violated being by training AI models. The laborers that help compile AI, something most users don't think about, are often poorly compensated and work under less-than-ideal conditions.¹³

Bias: The output of AI can be only as good as the data that goes into it. Skewed data produces skewed results. In using AI in its hiring practices, Amazon found that the hiring algorithm [favored](#) males for technical roles based on historical gender imbalances.¹⁴ Thanks to a data set that was faulty, recruitment was skewed based

on a history of discrimination, not merit. AI is assembled based on data with little to no discernment about the potential bias of the data, then encoded by a relatively small number of people who also have inherent biases, then potentially accepted as unbiased. As use of AI proliferates, it cannot simply be accepted at face value.

Environment: Proponents hope that in the long term, AI will help us become more energy-efficient by sorting through data faster and predicting where resources need to be allocated. In the short term, the training and inference phases of AI use up such a massive amount of data that tech giants refuse to reveal the full extent of that use, lumping data storage and AI model training into the same category. But looking at carbon emissions since 2019/2020, when tech giants began training large language models, is telling. Google, for example, reported that its greenhouse gas emissions had increased 48% from 2019 to 2024.¹⁵ Microsoft reported that its overall admissions increased 29% from 2020 to 2024.¹⁶ Training a model such as Generative Pre-trained Transformer 3 (or GPT-3) is estimated to use just under 1,300 megawatt hours (MWh) of electricity.¹⁷ This is roughly equivalent to the annual power consumption of 130 homes in the US.¹⁸ Training the more advanced GPT-4, meanwhile, is estimated to have used 50 times more electricity.¹⁹ AI requires energy to keep the temperature in computing facilities optimal, and massive amounts of fresh water in dry places to cool the facilities. The energy required to power AI development and data storage have prompted big tech companies to start looking to [nuclear power](#). They simply can't create enough power with our current capacity to generate power. One of the first moves the second Trump administration made was to announce a \$100 billion project to create new data centers to power AI. It may be that in the long run, AI will make energy use more efficient, but at the present rate of growth, AI alone could potentially use the amount of energy of a small country such as Ireland.²⁰

Scale: The scale of AI is both a benefit and a very challenging aspect. In theory, it could allow few people to accomplish much by eliminating the need for humans to perform predictable and repetitive tasks, at the risk of eliminating the need for human workers. But the scale at which it can operate also can cause significant disruption when things go wrong. Real estate corporation Zillow had to close a division and lay off 25% of its employees after its AI algorithm overestimated the price of housing.²¹

Lack of Transparency: Perhaps the most concerning part of addressing potential shortcomings in AI is its opacity. Users of AI don't always know they are using AI. Users of AI don't know how their AI was trained or on what data set. Decision-making can be hidden from the user. Programmers not versed in the field where the AI is being used may be effectively making decisions for experts simply because they are writing the algorithms. This is of particular concern when human life is at stake, as when AI is attached to weapons.

Overall, the possible benefits of AI include many potential drawbacks that tech companies are often unwilling to discuss or regulate. Indeed, senior researchers at OpenAI expressed grave concerns in 2023 that have been ignored.²² There are also questions of meaning and relationship that religious organizations can and should ask, as Pope Francis did early in 2024. What does it mean when Meta creates

“friends” out of AI? Should we keep loved ones “alive” by feeding their memories into AI? The fact that we can do something does not mean we should do it.

ELCA SOCIAL TEACHING

ELCA social teaching touches on AI-related issues in different ways. Principally, while social teaching conveys that science and human reason are God-given and should be used to further human flourishing, it does not lead us to think of human knowledge as neutral or perfect. The ELCA social statement *Genetics, Faith and Responsibility* (2011) sets out some parameters as to how to assess scientific breakthroughs from a theological perspective. The [statement](#) notes “Knowledge and technology have never developed in a social vacuum, and genetic research and technology and their delivery are not socially neutral” When it comes to AI, that same social teaching calls us to careful moral scrutiny of technological developments.

Certain uses of AI, wherein a mistaken result might threaten human life (such as use in defense), would fall under the “precautionary principle” described in the statement. “Precaution comes into play when existing tools for risk assessment are overwhelmed by a high level of uncertainty and proposed actions may dramatically affect the integrity and limits of the earth or the existence of future generations. In such cases, the burden to demonstrate safety rests upon those who promote the novel action.” It would be incumbent upon proponents of AI to demonstrate safety when AI is used in defense.

Further sounding a note of caution, the ELCA social statement *For Peace in God’s World* (1995) focuses on humility in the face of decisions about war, acknowledging that “even our best intentions can produce harmful results. Our efforts must take account of the human tendency to dominate and destroy, and must recognize those ‘principalities’ and ‘powers’ (Ephesians 6:12, RSV) that cause strife in our world” ([p. 7](#)). The statement recognizes that human advancement implies both good and evil ([p. 8](#)) in “tension-filled interplay.” Decisions about war (inherently considered to be “mournful”) require “political wisdom and historical knowledge of the situation” ([p. 12](#)). We believe that rational and moral discernment are critical responsibilities of both political decision-makers, and individual combatants and machines are not capable of this discernment.

Concern for the vulnerable is also a significant thread in ELCA social teaching. The ELCA social message “Human Rights” (2017) expresses concern throughout that the rights and needs of vulnerable people be given particular attention. It lays out six categories of human rights. AI could be used to hamper or violate any number of these categories, for example, restricting freedom of thought and religious expression through virtual policing; inhibiting political, civil and economic participation through facial recognition technology; and denying people’s right to physical goods through inequitable algorithms. The message also specifically directs the ELCA to address human rights violations through corporate social responsibility and commits the ELCA to upholding human rights.

The healthcare social statement prioritizes equitable access and calls for services to those who need it most. Any use of AI that results in diversion of resources away from equitable access in health care is one that the ELCA is obligated by its social teaching to question.

ELCA SOCIAL INVESTMENT SCREENS

The military weapons and human rights social criteria investment screens both apply to this issue.

CORPORATE RESPONSE

Corporations have been quick to adopt new advances in AI. Sometimes the results are salutary, but some of the challenges in AI have led to negative impacts on public opinion. Rite Aid pharmacies, for example, were banned by the Federal Trade Commission from using facial recognition technology after their electronic surveillance repeatedly misidentified customers as shoplifters.²³

Corporations have begun to see the virtue of adopting codes of conduct, with the encouragement of the shareholder community. This is a good first step, but implementation of those codes of conduct across a large organization with fast-moving development is a significant challenge.

Big data users such as Microsoft and Google remain reluctant to reveal publicly the full data of the energy use of training AI. Microsoft quietly walked back predictions that were made about going carbon-neutral. Big tech aggregate AI and data storage together and are reluctant to disaggregate and reveal the full impact of their business activities.

RESOLUTION GUIDELINES FOR THE ELCA – GENERAL²⁴

1. We support requests for transparency reports that explain the company's use of AI in its business operations and the board's role in overseeing AI usage, and that set forth any ethical guidelines the company has adopted regarding its use of AI.
2. We support human rights impact assessments examining the actual and potential human rights impacts of a company's AI-driven targeted advertising policies and practices.
3. We support independent third-party reports on customer due diligence process to determine whether customers' use of its products or services with surveillance technology and AI capability or of its components that support autonomous military and police vehicles, contributes to human rights harms.
4. We support independent studies and reports to shareholders regarding:
 - i. The extent to which such technology may endanger, threaten or violate privacy and/ or civil rights, and unfairly or disproportionately target or

surveil people of color, immigrants and activists in the United States.

- ii. The extent to which such technologies may be marketed and sold to authoritarian or repressive governments, including those identified by the U.S. Department of State Country Reports on Human Rights Practices.
 - iii. The potential loss of good will and other financial risks associated with these human rights issues.
5. We support audits of driver health and safety, evaluating the effects of performance metrics, policies, and procedures on driver health and safety across markets.
 6. We support calls for more quantitative and qualitative information on algorithmic systems. Exact disclosures are within management's discretion, but suggestions include: how a company uses algorithmic systems to target and deliver ads, error rates, and the impact these systems had on user speech and experiences. Management also has the discretion to consider using the recommendations and technical standards for algorithm and ad transparency put forward by the Mozilla Foundation and researchers at New York University.
 7. We support reports assessing the siting of data centers in countries of significant human rights concern, and the company's strategies for mitigating the related impacts.
 8. We support reports on customer due diligence process to determine whether customers' use of its products or services with surveillance technology and AI capability or of its components that support autonomous military and police vehicles, contributes to human rights harms.
 9. We support disclosing transition plans that result in new renewable energy capacity, or other actions that achieve actual emissions reductions at least equivalent to the energy demand associated with its expanded data center operations.
 10. We support reports, at a reasonable cost and omitting proprietary information, on the public health related costs and macroeconomic risks created by practices that limit or delay access to healthcare, particularly with regard to the use of AI.
 11. We support reports about quantitative water use reduction targets by data center location and report on practices implemented to secure social license to operate and reduce climate-related water risk. The report should be prepared at reasonable expense and omit proprietary information.
 12. We support data protection impact assessments on the company's health care service offerings that describe how the company is ensuring appropriate use of, and informed consent for collection of, patient data.
-

ENDNOTES

1. <https://www.vatican.va/content/francesco/en/speeches/2024/june/documents/20240614-g7-intelligenza-artificiale.html>
2. https://www.wipo.int/about-ip/en/artificial_intelligence/faq.html
3. Beyer, J. (2024, February 29). AI: Overview, Insights and Roadmap for Non-Profits. LinkedIn. <https://www.linkedin.com/pulse/ai-overview-insights-roadmap-non-profits-jon-beyer-rgdgc/>
4. For a description of each category, visit this link: <https://www.coursera.org/articles/types-of-ai>
5. <https://www.pewresearch.org/internet/2024/01/31/americans-use-of-mobile-technology-and-home-broadband/>
6. <https://www.hrw.org/news/2021/11/10/how-eus-flawed-artificial-intelligence-regulation-endangers-social-safety-net>
7. <https://www.nytimes.com/2024/12/23/science/ai-hallucinations-science.html>
8. <https://www.state.gov/bureau-of-arms-control-deterrence-and-stability/political-declaration-on-responsible-military-use-of-artificial-intelligence-and-autonomy>
9. <https://www.propublica.org/article/cigna-pxdx-medical-health-insurance-rejection-claims>
10. <https://www.npr.org/2024/08/23/nx-s1-5087586/realpage-rent-lawsuit-doj-real-estate-software-landlords-justice-department-price-fixing>
11. <https://www.brookings.edu/articles/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/>
12. <https://hai.stanford.edu/white-paper-rethinking-privacy-ai-era-policy-provocations-data-centric-world>
13. <https://time.com/6247678/openai-chatgpt-kenya-workers/>
14. <https://www.forbes.com/councils/forbestechcouncil/2023/09/25/ai-bias-in-recruitment-ethical-implications-and-transparency/>
15. <https://www.gstatic.com/gumdrop/sustainability/google-2024-environmental-report.pdf>
16. <https://blogs.microsoft.com/on-the-issues/2024/05/15/microsoft-environmental-sustainability-report-2024/>
17. <https://arxiv.org/pdf/2211.02001>
18. <https://www.theverge.com/24066646/ai-electricity-energy-watts-generative-consumption>
19. <https://www.economist.com/technology-quarterly/2024/01/29/data-centres-improved-greatly-in-energy-efficiency-as-they-grew-massively-larger>
20. [https://www.cell.com/joule/fulltext/S2542-4351\(23\)00365-3](https://www.cell.com/joule/fulltext/S2542-4351(23)00365-3)

21. <https://www.gsb.stanford.edu/insights/flip-flop-why-zillows-algorithmic-home-buying-venture-imploded>
22. <https://www.reuters.com/technology/sam-altmans-ouster-openai-was-precipitated-by-letter-board-about-ai-breakthrough-2023-11-22/>
23. <https://www.ftc.gov/news-events/news/press-releases/2023/12/rite-aid-banned-using-ai-facial-recognition-after-ftc-says-retailer-deployed-technology-without>
24. These guidelines may be used in proxy voting as well as to help determine resolutions to file and dialogues to support. Each resolution guideline should be looked at within the context of the entire resolution language and specific company situation.