

Compliance Today – August 2022 Oversight considerations for ethical artificial intelligence

By Nakis Urfi, JD, MPH, Shawn E. Marchese, CHC, and Dr. Keith Grimes

Nakis Urfi (nakis.urfi@babylonhealth.com) is Product Compliance Officer in Dallas, Texas, **Shawn E. Marchese** (shawn.marchese@babylonhealth.com) is Global Head of Compliance and Risk in Austin, Texas, and **Keith Grimes** (keith.grimes@babylonhealth.com) is Director of Clinical Product Management in London for Babylon.

- [linkedin.com/in/nurfi](https://www.linkedin.com/in/nurfi)
- [@nakisurfi](https://www.instagram.com/nakisurfi)
- [linkedin.com/in/shawn-e-marchese/](https://www.linkedin.com/in/shawn-e-marchese/)
- [@ShawnEMarchese](https://www.instagram.com/ShawnEMarchese)
- [linkedin.com/in/drkeithgrimes/](https://www.linkedin.com/in/drkeithgrimes/)
- [@keithgrimes](https://www.instagram.com/keithgrimes)

Imagine you are a leader at an enormously large healthcare organization, one where each functional area operates within its own silo: operations, client support, technology, etc. — especially technology. There is little transparency and a lack of communication between teams. Little to no documentation or shared frameworks. The services and the technology offered by the company are too large and complex for any single person or team to really understand how they all truly operate. Maybe you don't have to imagine this at all; maybe this is where you are today.

Suddenly, reports start to emerge that your company's technology is producing unintended consequences. Monitoring devices, software, and robots providing healthcare to individuals are not functioning as they're meant to, causing negative patient outcomes, inaccurate advice, and overall confusion. No one at your organization knows where to start to identify and address the root cause of the problem, because the technology is *everywhere* across the organization, housed in a variety of software systems and cloud solutions. Not only that, but the organization you work for is integrated with other health systems, and shutting down the service completely would have tremendous and unpredictable impacts to an incalculably large population.

A nightmare? A cautionary tale? Some new sci-fi horror movie on Netflix? Maybe.

But this is a potential scenario that *could occur in the near future*. There are currently insufficient regulations in place to prevent such an occurrence, and by the time this hypothetical scenario becomes a reality, regulations may still not be sufficient to mitigate the potential risks. With any luck, it won't happen. But if it does, it will likely trigger a host of new conversations and regulations on the part of industry thought leaders, elected



Nakis Urfi



Shawn E.
Marchese



Dr. Keith Grimes

officials, and regulators. Legislators will hold hearings to determine why it happened and what can be done to prevent its recurrence; industry will be called to account; and, finally, stricter oversight and regulation will be implemented that may have a chilling effect on innovation as a whole. Compliance professionals will study and discuss the case in conferences and magazine articles, discussing the many ways better oversight might have prevented the worst if only—if only—we had gotten involved sooner.

Why wait for the worst to happen?

What is artificial intelligence?

Before we get into questions of artificial intelligence (AI) oversight, it's worth defining AI. However, there is no single, universally accepted definition of "artificial intelligence"—nor indeed even of "intelligence." For our definition of AI, we will consider it as a set of advanced technologies that enable machines to carry out highly complex tasks effectively—tasks that would require intelligence if a person were to perform them. These include tasks such as decision-making, pattern recognition, machine learning, and natural language processing, which allow machines to operate more like the human mind.

AI is being used in more places than most people realize and can be used anywhere automation can occur. Media is full of advertisements for healthcare companies touting AI-powered capabilities.

AI can help bridge the gaps in care if used responsibly to meet the global demand of healthcare, which exceeds the supply of human brains available to meet it.^[1] AI is being used in various ways in healthcare today, including predictive analytics, risk analyses, health recommendations, claims processing, clinical documentation, revenue cycle management, and much more.

Even closer to home for this audience, AI capabilities are making life easier for compliance teams, as well. AI can streamline compliance monitoring by helping mitigate the impact of false positives in exclusion screenings, enhance internal audit capabilities to make audits more data-driven and effective, and can even assist with regulatory change management—with natural language processing developing to a point where it can quickly identify changes in regulations. Finally, chatbots can help organizations answer simple questions and access documents such as policies and procedures.

Okay, then what is *ethical* AI?

What do we define as "ethical"? At its most basic, "ethics" (from the Greek word *ethos*, "custom, character") is a system of moral principles that affects how people make decisions and lead their lives. Another way to think of ethics is in the terms used by former U.S. Supreme Court Justice Potter Stewart: "Ethics is knowing the difference between what you have a right to do and what is right to do."

Currently, there are various entities defining and publishing principles for ethical AI, ranging from governments to the World Health Organization and other organizations; to universities, industry groups, and joint private and public collaborations; to—yes—individual companies. Most of these ethical AI principles cover common fundamental values that should be familiar to compliance professionals, such as protecting individual rights and autonomy, ensuring transparency and trust, and maintaining responsibility, among others.

But the regulatory landscape is currently not well developed in the field of AI, and current regulations may not fully oversee and prevent potential negative outcomes from its use. Furthermore, the technology is accelerating rapidly and will likely continue to outpace the development of adequate regulations for the foreseeable future. Therefore, it is incumbent upon those people who develop and use these AI technologies to be cognizant of the potential harms and unintended consequences that could arise.

Potential issues

Like many tools, AI can be used for both good and bad purposes. Many of the bad results may actually be unintended consequences far different than the initial intended purpose. A recent example may be found in the discussion about social media and the ill effects and the potential negative influence that we are beginning to understand come from its usage in so many facets of our society. Many of the key issues to be concerned about regarding the uses and results of AI are as follows.

Bias

One of the main issues that can arise from AI is bias. AI is only as good as the data that goes into it, but the available data in the US healthcare system is often inaccurate, filled with gaps, includes duplication, and may not be representative of a whole population. Using data collected mostly from individuals in high-income earning regions may not perform as well if the healthcare solution is to be used in lower-income settings. The AI models learn and develop from the data and, over time, increase the opportunities to amplify the bias in AI models and algorithms, which can lead to negative outcomes.

There are ways to correct bias, including identifying potential issues with data, AI models and algorithms, and developers. Once you understand your potential bias issues, you can impute and supplement data for missing data subsets in the populations to help address bias.

AI model drift

AI model drift occurs when changes in the data and other variables over time result in degradation of model performance. Without proper monitoring, these changes can affect the intended outcomes of the AI model in a negative manner and diverge quickly over time.

Trust and transparency

In any AI system, there is a need for trust, transparency, and explainability of the model and its results. Users should know the basics of how a given AI product is making decisions. Transparency should be a goal so that the public does not have to blindly trust a company's "black box" of opaque and proprietary algorithms with no method or attempt to explain how the AI works. In addition to transparency, traceability will help understand the genesis of the AI and the data that was used with its processes. Traceability supports the idea that companies should document how the AI was created and for what purpose in a way that explains why a system has particular dynamics or behaviors.

In addition, "human-in-the-loop" systems ensure that humans are able to provide direct feedback into AI models to supplement decisions and predictions with a low level of confidence. The presence of a human to review, validate, and make timely changes to algorithms demonstrably improves results while also allowing for the inspection and assessment of the data used by the algorithms. Having no human in the loop can open the door to potential issues that may not be detected and result in unintended outcomes.

User experience and safety

User experience is another potential issue. Much research and money are spent on understanding how to get people to engage and keep using their technology, but designs that are addictive and/or manipulative present ethical concerns. Companies employing AI should not engage in deceptive and manipulative practices. In addition, in healthcare, safety and accuracy must always be at the forefront of considerations for uses of AI in

order to avoid patient harm.

AI regulatory overview

The following is an overview of the various regulations and guidance emerging to address AI.

United States regulations

The United States does not have laws governing all aspects of AI, only certain areas and use cases. For example, there are various healthcare and privacy laws that may affect AI when AI implicates healthcare and privacy concerns. In addition, the Federal Trade Commission (FTC) has provided guidance warning companies to ensure that their algorithms and AI-based collections are not deceptive, biased, or unfair.^[2] The FTC's authority derives from Section 5 of the FTC Act, the Fair Credit Reporting Act, and Equal Credit Opportunity Act. There are biometric recognition laws that govern facial recognition technology that may be using AI. Also, Colorado passed a bill that prohibits insurers from using algorithms, external data sources, and predictive modeling systems to discriminate against people.^[3]

There are agencies and organizations that are also working on various initiatives in the AI space. The National Institute of Standards and Technology (NIST) is developing a voluntary framework for more ethical AI with its AI Risk Management Framework, which is expected to be released in January 2023.^[4] Additionally, NIST released Bias Guidance in March 2022, sharing three categories of potential bias—including systemic bias, statistical and computational biases, and human bias, and recommendations on how to address bias in AI.^[5] The Food and Drug Administration has also been exploring aspects of transparency and bias surrounding AI/machine learning-enabled medical devices.^[6] Finally, the Government Accountability Office has developed an AI Accountability Framework to address accountability challenges in AI by laying out key practices, questions, and audit procedures.^[7]

Global regulations

On the global stage, at least 60 countries have adopted some form of AI policy or regulation. Countries such as Brazil, China, Japan, and the UK are all developing various types of regulations, guidance, and strategies for their nations. Notably, the European Union (EU) shared a proposed regulation for the EU's approach to AI, which divides AI systems into categories of unacceptable-risk AI systems, high-risk AI systems, and limited- and minimal-risk AI systems.^[8] There are different requirements, including possible oversight obligations, based on each AI system's level of risk. The proposed regulation may go into effect as early as 2024.

Also, during its second meeting in 2022, the US-EU Trade and Technology Council released a statement that the EU and the US were committed to cooperating closely on certain key areas, such as technology standards, including cooperation on AI.^[9] Additionally, the International Organization for Standardization is working on various AI standards in hopes to gain widespread technology standardization that will lead to better transparency, and that focuses on various technical, societal, and ethical considerations.

Oversight considerations

Considering these risks, oversight should be considered to mitigate and control them. Fortunately, the oversight mechanisms that can mitigate risks in the ethical use of AI are concepts that should be familiar to any compliance professional and require structures and processes that should be in place at any established compliance program.

AI oversight strategy

Organizations should begin with developing an overall AI oversight strategy. As a first step, the organization needs to determine the appropriate stakeholders to help develop this strategy and documentation. A model for your AI oversight strategy document may be your company's compliance program document, but it should be augmented to include how to address AI-specific issues such as bias and explainability.

AI inventory

To begin a journey into addressing ethical AI, organizations should start by creating an AI inventory to identify all areas of the organization that are using AI—and all the uses. Speak with relevant stakeholders within the organization to gain an understanding of the various uses of AI within the organization; doing so will help with the creation of an inventory.

Risk assessments

Once there is an inventory, one can begin to assess the risks of the organization's AI uses. How do an organization's risk assessment methodologies contemplate technology risk? This should be specific to the organization's business and can include safety factors, operational impacts, reputational impacts, etc.

Governance

Governance mechanisms embedded throughout the organization should be leveraged to help mitigate risks of AI misuse. The organization needs to define a governance process for controls access, documentation and policies, and tracking and monitoring outcomes.

Top-down leadership and culture

What strategic positions and concerns have the organization's board and senior leadership articulated regarding the ethical use of AI? Is AI referenced in the organization's values? How an organization's leadership views business objectives and mission is an indicator of what the organization values. If the leadership values using AI and technology ethically and responsibly and makes statements that align with this view, it will drive the culture of the organization and establish the right attitude toward the ethical use of AI.

Code of ethics and conduct

Organizations should consider including information about the ethical use of AI in their code of ethics and conduct. The code can list guiding principles that the organization intends to follow to ensure the organization and employees act with honesty and integrity, and that it has ethical responsibilities to all stakeholders.

Documentation

It is a good practice to have processes in place that document what the purpose of the AI technology is, how that purpose is being accomplished, and any sort of review process, among other aspects relevant to the organization.

Training

Regular training will help remind employees of the organization's values and responsibilities. Training can remind employees, especially employees handling the development of AI technology, that the end users' safety and well-being should be at the forefront of all technological advancements. Catchy phrases such as "Ethical

developers develop ethical technology” can help build culture; yet it is not just developers that need periodic reminders but also the vast array of other stakeholders who use, receive, and act on the data and functionality of the AI. This will help employees understand the bigger picture of their jobs, which include using powerful tools with great capabilities that come with inherent risks, and to have responsibility for their uses and outcomes.

Auditing and monitoring

Auditing and monitoring are powerful tools for oversight of ethical AI, just as they are powerful tools for compliance oversight. In fact, references to guidance and documents discussed earlier have recommendations on how to mitigate bias and make more ethical AI, which include more specific direction on how to audit and monitor AI and its outputs.

Frameworks, tool kits, and software development life cycle

Organizations can take advantage of the various frameworks and tool kits available online that identify and provide guidance and standards on how to increase adherence to ethical AI principles and reduce bias, unfairness, and lack of transparency. This can help an organization reduce risks and help its operations in the long run. Additionally, using an end-to-end software development life cycle will help the organization capture major steps and improve efficiencies with software development. This is another process to help ensure transparency with the organization’s AI development and also to reduce risk and address issues timely when they arise.

Communication and reporting

The AI regulatory landscape is changing, and it will be imperative that companies that develop and use AI technologies stay abreast of regulatory changes, including organizations with an international scope.

Considerations include how an organization’s hotline and other communication channels are doing. Do people feel comfortable reporting issues, especially possible unethical practices that may be arising in a certain area of your organization? Are AI risks and issues being brought up to the appropriate levels of leadership and the board? It may not be on the radar at the moment, but with the rapid adoption of AI in what seems to be almost all aspects of the healthcare industry, it may be something to be aware of for the future.

Transparency, collaboration, and knowledge sharing

What is being developed at the organization? Is it transparent, or is it being developed in what seems to be a secret lab? This area is complex and requires the expertise and collaboration of a variety of cross-functional disciplines. A good practice is to have technology teams share important technological updates throughout the organization.

AI ethics committee or forum

An organization may benefit from creating some type of body that meets to agree on principles and discuss questions related to the ethical use of AI. Different models will work for different organizations, from a formal AI ethics committee to a more informal ethics forum, panel, or a recurring agenda item, to an updated current standing meeting. There is no one-size-fits-all approach here. The purpose of the group meetings and the attendee list of multidisciplinary stakeholders should be decided early in the process. Regardless of the model chosen, the framework, governance mechanisms, and overall understanding of what is happening with AI development and usage can be discussed and captured here.

Conclusion

Although the US is lagging in developing regulations and guidance for AI, these will eventually come into effect, and when they do, organizations should be ready. However, variations between state and local governments with the federal government could complicate the AI regulatory landscape in the US as it develops. Organizations (particularly those based in or with a presence in other countries) should be watching trends in the global landscape and implementing oversight structures now.

Furthermore, ethics goes beyond regulation. Organizations and individuals have the responsibility to ensure their AI-embedded products and services are ethical and responsible. “When asked why companies aren’t mitigating all relevant risks...respondents in emerging economies are more likely than others to report that they are waiting until clearer regulations for risk mitigation are in place, and that they know from formal assessments that mitigation is more costly than the consequences of a risk-related incident.”^[10] As long as this is the case, there is always a chance that organizations will take unnecessary risks and face the consequences of negative incidents and delays, rather than implement strong risk mitigation processes now.

There is a multitude of benefits that can stem from implementing AI oversight mechanisms. In addition to enabling smoother operations and processes within a business, employees will feel better about their employer knowing that the organization has made a commitment toward ethical use of AI. There can be positive societal and governance impacts that can be added to ESG strategies. Additionally, trust with stakeholders, such as investors, customers and patients, employees, regulators, and the public as a whole, can be built when transparent and explainable technologies that focus on the safety of the users are used. This will provide a competitive advantage in the marketplace as trends move toward sustainable practices that involve ethics and social responsibility, and investors are interested in profitable businesses that mitigate looming risks. Additionally, jumping ahead with best practices should enable an organization to be better prepared for the upcoming regulatory demands in the industry.

This is a complex area that is evolving rapidly, and as in many areas of the industry, innovation and regulation are frequently at odds when it comes to the use of AI. There is an opportunity for organizations, including compliance professionals, to get ahead of the potential problems and regulations that will inevitably arise from the future advanced development of AI. Now, before there are clear regulatory requirements to do so, organizations can take initiative and proactively address ethical use of AI, and compliance programs are well positioned to lead in this area.

Takeaways

- Ethical artificial intelligence (AI) follows defined principles regarding fundamental values, such as protecting individual rights and autonomy, ensuring transparency and trust, and maintaining responsibility, among other values.
- AI has great potential to both enhance our healthcare system and cause unintended issues and harm.
- AI regulations are lagging behind the rapid innovation of AI and its never-ending applications.
- Developing an organization-wide strategy on how to approach AI through various oversight mechanisms can benefit organizations in a multitude of ways.
- Compliance professionals have an opportunity and are uniquely positioned to implement effective ethical AI oversight mechanisms at their organizations.

- 1** HCCA-SCCE, “Marchese, Urfi and Grimes on Ethical AI in Healthcare,” podcast, *The Compliance and Ethics Blog*, February 22, 2022, <https://www.complianceandethics.org/marchese-urfi-and-grimes-on-ethical-ai-in-healthcarepodcast/>.
- 2** Elisa Jillson, “Aiming for truth, fairness, and equity in your company’s use of AI,” *Business Blog*, Federal Trade Commission, April 19, 2021, <https://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai>.
- 3** S.B. 21-169, 2021 Leg., Reg. Sess. (Colo. 2021).
- 4** “AI Risk Management Framework: Initial Draft,” National Institute of Standards and Technology, March 17, 2022, <https://www.nist.gov/system/files/documents/2022/03/17/AI-RMF-1stdraft.pdf>.
- 5** Reva Schwartz et al., “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence,” NIST Special Publication 1270, March 2022, <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf>.
- 6** U.S. Food and Drug Administration, “Virtual Public Workshop – Transparency of Artificial Intelligence/Machine Learning-enabled Medical Devices,” October 14, 2021, <https://www.fda.gov/medical-devices/workshops-conferences-medical-devices/virtual-public-workshop-transparency-artificial-intelligencemachine-learning-enabled-medical-devices>.
- 7** U.S. Government Accountability Office, “Our New Framework for Artificial Intelligence (video),” WatchBlog, July 21, 2021, <https://www.gao.gov/blog/our-new-framework-artificial-intelligence-video>.
- 8** Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, COM (2021) 206 final (April 21, 2021), <https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1623335154975&uri=CELEX%3A52021PC0206>.
- 9** U.S. Department of Commerce, “U.S.-EU Joint Statement of the Trade and Technology Council,” news release, May 16, 2022, <https://www.commerce.gov/news/press-releases/2022/05/us-eu-joint-statement-trade-and-technology-council>.
- 10** McKinsey Analytics, *The state of AI in 2021*, December 2021, <https://www.mckinsey.com/~media/McKinsey/Business%20Functions/McKinsey%20Analytics/Our%20Insights/survey-The-state-of-AI-in-2021.pdf>.

This publication is only available to members. To view all documents, please log in or become a member.

[Become a Member Login](#)