

# Transformation and Risk Podcast

## Episode 2

### AI and Generative AI in the Healthcare Industry

#### Hala Helm

Hello everyone, and welcome to the Healthcare Transformation and Risk podcast series. Here we explore the dynamic world of healthcare transformation and the accompanying risks. Join us as we speak with industry experts to uncover the challenges and strategies essential for healthcare systems on their transformation journey.

In this second episode, we'll focus on the challenges and impact of integrating artificial intelligence and generative artificial intelligence in the healthcare industry. I'm Hala Helm, Strategic Risk Advisor in the Marsh National Healthcare Practice. Today I'm joined by my colleagues, Jaymin Kim, Senior Vice President, Emerging Technologies at Marsh, and Kristin Rhebergen, Principal Digital Human Resources Advisory at Mercer.

Let's dive in. Many of our healthcare clients are excited about the integration of emerging AI and GAI technologies into existing services to drive improved efficiencies and outcomes. However, they understandably have questions and concerns about potential risks. Healthcare is already using artificial intelligence tools in multiple settings from administration to patient care. Just a few examples include applicant resume screening, medical imaging, chatbot responses to patient inquiries and other applications. So, Jaymin, as organizations begin to adopt even more sophisticated technologies, what should they be thinking about? Does the risk automatically increase when we adopt and deploy new technologies?

#### Jaymin Kim

So, as you alluded to Hala, the healthcare sector has been using various forms of artificial intelligence as well as other advanced technologies for that matter, for decades; like using AI to help identify patterns associated with specific diseases and large data sets

has been in practice because AI might be able to do that more efficiently than humans can. I would say the most recent addition to the various sophisticated technologies and use has been generative AI, which is a form of AI that can create original content across various domains like text, imaging and voice.

I think the first question that healthcare organizations should ask before deciding to deploy any new advanced technology is, what is the outcome we're looking for and how can this new technology help us optimally reach that outcome? Because sometimes it may not make sense to use the latest technology. It really depends on what your organization's goals are, your organization's risk appetite, and your ability to integrate these new technologies into existing or perhaps legacy infrastructures.

Risk doesn't automatically increase overall when we adopt and deploy these new technologies. In fact, some risks may get dialed up, whereas others might get dialed down. And so, it's really important to assess how your specific organization's risk profile changes. For example, some healthcare organizations are using generative AI to automate administrative tasks on the backend to help reduce burden on healthcare professionals, bring them up to do other higher value tasks, just to help reduce the risk of human error that might be made by, for example, overworked professionals and processing patient data.

But it may increase various data privacy and security risks as these technologies rely on massive volumes of large data sets. So the degree to which these risks get dialed up or down will largely depend on how the organization plans for resiliency and how well it implements risk mitigation controls appropriately.

I will say on the operational side of addressing generative AI risks, one of the things that I think all healthcare organizations should proactively be addressing is the fact that unlike other forms of AI technology that have been already used by the sector for a while, generative AI in the forms of, for example, ChatGPT and GenAI, and et cetera, is in the hands of every single individual practitioner who has access to the internet and a computer mobile device, whether the healthcare organization likes it or not, which means that organizations need to have a really clearly defined and well-communicated acceptable use policy and training guidelines for their professionals, as well as a way of enforcing said policy.

If you have no stance on whether it's acceptable to use, for example, the public version of ChatGPT, or if you say as an organization that practitioners shouldn't use this form of ChatGPT, but have no controls in place to enforce, you should probably as an organization be assuming that some professionals in your organization are using the public version of ChatGPT for their work courses.

## Hala Helm

You're so right, Jaymin. I mean, having worked in healthcare myself for a long time, I can recall when text became a thing and hospitals tried very hard to keep practitioners, doctors, and other providers from texting with little to no success. So really where we began to have success was when we started adopting that technology safely and responsibly instead of trying to avoid it. So, a very good point, and it takes me to another point around doctors in particular.

So one of the key challenges in this integration is the impact on the doctor-patient relationship and the role of healthcare professionals with AI and GAI becoming more prevalent in healthcare. It's really important to consider how this technology may affect the dynamics between doctors and patients and other healthcare providers and practitioners within healthcare organizations. So, Kristin, how should we be thinking about AI's impact on the employee experience in addition to how it'll be integrated into that day-to-day work?

## Kristin Rhebergen

I think Jaymin touched on a really great point in her initial comment about how generative AI can really be impactful. If we talk specifically about generative AI in helping to reduce some of that administrative burden, helping to free up time, so that you can focus on those interactions and that relationship and making sure that you're providing the best care and the most focused care to the patient population.

So if we think about that and take that and apply it to the employee experience, to how employees across healthcare organizations and employees in general are engaging with the platforms and solutions at work that they just need to do in order to be able to do things at work, to get stuff done. There's a couple of key points to think about here.

One, we've called this out already. AI is already being applied in the employee experience. We've got chatbots already helping to answer some questions, or we've got AI-powered search to help us find information. And what that is doing today is it is giving us basically links to other sites or links to other resources, but employees still have to go to those links and still have to parse out information and understand and discern how it pertains to that.

The power of generative AI in this area is to reduce that friction, to summarize the information and provide a personalized response back that really is applicable to the employee and their situation based on the question they asked. So rather than me having to take 10 minutes to go and type in a question about my leave policy and get the link to the knowledge article and read through it and be like, "I think this applies to me, I'm not really entirely sure". With generative AI there's the opportunity that the AI platform can do all of that for me

and turn a ten-minute interaction into a two to three-minute interaction.

In the long term, ideally, we would also see that generative AI working together in an agentic relationship, so agents talking to agents, if we think about on the smaller scale. I would then be able to take the information that the AI chatbot gave me about my leave policy, tell it that I want to initiate a leave beginning on September 1st, and it would then be able to help me complete that action in the necessary system so I don't have to go into another system to transact. I can just have a conversation with the generative AI chatbot. Now, there are risks to this and Jaymin highlighted some of those, and this is why it's not happening on a wide basis yet. We want to make sure before we deploy them that we can mitigate those risks and ensure that we're deploying ethically and in security.

## Hala Helm

Thanks so much for that insight. So as we begin to implement these types of systems, what do you think is Human Resources' role in helping to navigate this transformation?

## Kristin Rhebergen

It's a great question. And so very timely. As we take a step back and going again to what Jaymin said earlier, as we think about the outcome that we're trying to achieve, AI really has to be front and center in the conversations and in the planning and in the deployment. If our goal is to produce a more productive, a more effective, a more efficient, and yet still fulfilling workplace, we really need to be thinking about having the human at the center of it all, and not just thinking about how we can drive productivity, how we can turn the [widget] faster or produce [widgets] faster, if you will.

So we don't want to chase productivity gains when deploying AI or generative AI solutions at the expense of humanity. We want to make sure that we're serving our employee population and our patient population if we think about healthcare and not the other way around.

And so, as we look at this in the future, HR is really going to need to stand in as what we're calling the steward of humanity, making sure that they are a part of the team across the organization that is driving ethical use of AI, safeguarding the ESG experience and the values, and ensuring that the benefits realized by AI enhance the employee experience and don't detract from it. In many ways, HR is actually one of the first groups in organizations as we think about some of the more back office and administrative work, one of the first groups in organizations to start looking at and deploying generative AI. And it's a great way for HR to kind of lead the way, especially as we start talking

about deploying the platforms and the technologies and the solutions out across the entire workforce. It's a quick highlight. I do want to call out some great minds at Mercer really dove deeper into this topic recently in a white paper titled *The Future of Human Resources: Who will care for the human at work*, talking about how to navigate this transformation and this journey at large. And I encourage everybody to read it as they think about starting on their AI journey.

## Hala Helm

I really loved what you said about the human piece of this. I mean, as a former healthcare practitioner myself, people go into healthcare because they want to care for patients.

## Kristin Rhebergen

Yeah.

## Hala Helm

And healthcare has changed and evolved in a lot of ways, some good, some not so good. But one of the things that we really talk about in healthcare is burnout; burnout of healthcare practitioners, nurses, doctors, and I think a lot of that is because we have become less patient-centered and more focused on that productivity that you talked about. So anything that can bring us back to that core mission, I think has enormous potential to drive both patient and employee satisfaction.

So moving on, let's discuss the potential impact of AI and GAI on legal liability in healthcare. As these technologies become more advanced, it's important to consider how they may alter the legal landscape in terms of liability. With these technologies potentially making decisions and recommendations in healthcare, questions arise about who's responsible in case of errors or adverse outcomes, and it's essential for healthcare providers and legal professionals to work together to establish guidelines and frameworks that address these concerns and ensure accountability. So, Jaymin, how would you suggest healthcare organizations begin to proactively develop a framework to address this and other types of associated risks?

## Jaymin Kim

So at Marsh, we've constructed a risk framework specific to generative AI to help assess what's new here – if anything – from a risk perspective, with implications for who may be held accountable for what. With generative AI, what we're finding is that many of the risks, or in fact extensions of existing and familiar categories like data privacy and copyright infringement or technology misuse, and then potential errors, and of course the kind of professional liability that goes along with potentially misusing a technology.

I think there's three things that organizations can be doing to proactively address the changing legal landscape and what it means for what their organization may be held liable for. I think the first thing is organizations should be asking what risks come with using and deploying this new technology and which are most relevant for my organization? This will really depend on how your organization is using generative AI. From the emerging use cases that you're using generative AI for, you should be mapping out a risk management plan within your organization with clearly assigned roles and responsibilities. And this is all to ensure that the way that practitioners are using a generative AI is in a responsible way in alignment with existing guidelines. The same set of, for example, data privacy laws that come with HIPAA or GDPR are here to stay. It's essential to ensure that the ways in which you're deploying new technologies are still in alignment with these existing kinds of laws and regulations.

In the meantime, it's incumbent on organizations to have someone or some team that is responsible for keeping track of the latest emerging laws and regulations specific to artificial intelligence, whether it's at the federal level or the state level, or by the healthcare sector level, and map out based on how your organization is using generative AI, what are the potential ways that things could go wrong and how can the organization implement various kinds of risk management controls to ensure that your organization's usage of the technology remains compliant with the evolving landscape?

When it comes to various risk management controls, to Kristin's point earlier, this isn't just a technology topic. It's also about people. And so when we think about what organizations can be doing as they navigate the emerging legal landscape within their organizations, they have the power to put in technical controls as well as people, as well as process controls. So technical controls include things like looking at how do the AI models used by your organization process and store data, with what kinds of privacy and security controls, and are they compliant with the latest laws and regulations guiding the space? Or what's the data retention policy applied to your organization's usage of generative AI models?

People controls look at things like, how are colleagues trained on using these technologies, and how do we define the principles of providing human oversight? People controls include governance, so ensuring that there's a formalized multi-functional governance committee that includes not only the CSO's office, but also the CHR's office, as well as people who are in charge of deploying these tools in the field. The practitioners who are using these technologies themselves are best placed to provide that feedback loop in terms of what's working and not working practically on the ground.

As well as, of course, operationalizing processes to enable comprehensive monitoring of potential risks and corresponding liabilities, such as setting up a centralized repository and keeping track of emerging

use cases. Because again, what your organization may be liable for will directly depend on how your organization is using site technology.

## Jaymin Kim

But from a liability perspective, even if your particular healthcare organization does everything right, you may still be exposed to potential downstream or upstream risks. And so it's important to consider from a counterpart perspective, how are various vendors in your organization's ecosystem using generative AI? Do they have access to your data? And how are they handling your data? And if they're using generative AI or developing generative AI systems, what are their risk management practices? So it's important to conduct specific due diligence related to using and deploying and developing these models even among your vendor network.

And third, I briefly touched on how we're still seeing emerging regulatory guidelines and laws specific to artificial intelligence. Beyond the existing enterprise and security laws, we have various kinds of AI acts, for example, in the EU there's the EU AI Act. In the US we have the Biden Administration's Executive Order on AI. At the state level, there's been various levels of activity coming out of California and Colorado and New York and others.

I think it will be an evolving field. We're still in very early days. It's been about 18 months since we saw the birth of ChatGPT in the public domain. In the meantime, I think it's also going to be interesting to see, for example, are there new standards of care that arise around this space? For example, if a regulatory entity in the healthcare space, like the FDA, comes out and says going forward the best practice is going to be for healthcare organizations to use and deploy some kind of general AI model in these kinds of use cases. That's where we might see some of the responsibilities or liabilities change from being in the hands of the medical practitioner who's ultimate responsibility it has been to appropriately use technologies and provide that human expert guidance in knowing when to use and how to use over to perhaps a product liability standpoint where the onus is now more on the developers of these systems themselves.

## Hala Helm

I think you're exactly right. We are already seeing that there are some AI models that are more specific, more accurate than humans in some situations, or just serve as a really good adjunct to those human diagnostics. So I think that you're absolutely right, that we are going to see this evolving standard of care wherein at some point in time it will be the standard of care to use AI tools to augment physician professional judgment or other types of human interventions. So, great point.

Let's talk now about the first steps that organizations can take if they haven't already started their AI journey.

And I know a lot of healthcare organizations are very much on this path, but not every healthcare organization is an early adopter, and so some are just beginning to put their toes in the water again with some trepidation. So Kristin, what recommendations do you have for employers who are just starting this journey from an employer standpoint?

## Kristin Rhebergen

I think there's two key things here, and Jaymin called one of them out early on, about making sure that you know what the outcome you're looking for and trying to drive towards is really, really important. But if you keep that in mind, it's if you can optimize and leverage and we breakout AI versus generative AI, but if you can optimize and leverage other forms of AI that may already be embedded in your core platforms in an ethical and a secure manner, it may make the most sense for employers, for organizations to start there, right?

Why chase the shiny object and potentially spend more, if you can really optimize and take advantage of what you already have in place today? So that's the AI component, the more traditional deep learning or machine learning types of AI. But if we think about looking ahead to the application or starting the journey around generative AI and going down that path, is that what we're really seeing today and what we're actually counseling organizations to do is to take this year to educate, take this year to experiment, take this year to practice, to learn. And then next year think about enabling, think about deploying, think about embedding the solutions into your work and into everything that you're doing.

And we call this out because I say that I'm counseling organizations to do this, but I do also want to highlight that I'm really seeing organizations say 'we're not ready to deploy'. We need to take those foundational steps today which I think is really wonderful. And the reason I call this out is because it's so important for everyone from the C-suite down to your frontline worker to understand more about how generative AI can be used directly in their day-to-day work, beyond drafting an email or beyond summarizing a meeting. Those are amazing use cases, don't get me wrong. They really are. They're time saving. They're really great, but that's not what's going to push the needle, right? That's not what's going to drive that outcome the organizations are trying to push towards.

What we really are encouraging organizations to do is to focus on how they can apply generative AI, and use it in a way that fits into their working styles and roles. So this includes resetting on foundations, like providing lessons and education around things like prompting one-on-one and prompting two-on-one. How do we go from basic prompts to advanced prompts, doing the art of the possible and focused ideation sessions, and then getting hands-on practice in using the platforms to so that they know how it can drive effectiveness, drive efficiency, drive productivity, but also how and where they need to keep the human in the loop so that they're



not allowing the platform or the solution to make decisions, to do things for the human that should not be done.

And one of the great things that can come out of this education and experimentation is that people at the grassroots level, at that frontline level can really start to tell the story about the value that it could bring. So it's no longer this hypothetical kind of nebulous, cloudy thing, but they can actually say, this is how it's helping me do my work. This is how it's helping me be more productive. This is how it's helping me to be more human in my interactions on a day-to-day basis with my colleagues and the people that I'm caring for or supporting.

And finally, we recommend focusing on starting small to build momentum in the necessary digital behaviors and those stories of proofs really do help to communicate that with them, the 'what's in it for me'? And take that more mindful approach to rolling this new technology and these new solutions out, rather than going 'big bang', think about how you can embed it in a secure and an ethical and in a mindful manner.

And I also love Jaymin's comment about the importance of feedback loops and engaging with the frontline, because it's those people who are the ones who will be able to tell you what's working best and what's not working best. It's really important, again, that you focus on bottom-up education and experimentation as well as top-down experimentation and education.

## Hala Helm

Great advice. Jaymin, I know you touched on much of this in your last set of comments, but anything further to add around that?

## Jaymin Kim

I thought, Kristin, your points were great. The only thing I'll briefly add is that as organizations think about how to utilize various technologies like AI and generative AI with direct implications for their workforce strategy, I think something that goes hand in hand with that is, as certain roles get redeployed, and the workforce strategy for particular organizations evolve over time, it's important to think about the corresponding risk landscape as well.

Again, coming back to you, it's not that your overall risk picture suddenly increases or accelerates certain risks will get dialed up, whereas others will get dialed down. And so details will be important in terms of which functions within your organization are using which technologies for what use cases.

## Hala Helm

These are all great points. So we've talked a lot about AI and GenAI, and they do obviously bring new

opportunities and new risks, but it's crucial for organizations to think beyond just these technologies to the overall tech transformation. Technology is moving so quickly and there are other technologies to consider beyond just AI and GenAI. So Jaymin, how would you encourage our clients to approach that exercise in thinking about risk management beyond AI into emerging technologies?

## Jaymin Kim

That's a great question, Hala. I think frequently we're talking about AI as though it's happening in a vacuum. And in reality, AI is being used in conjunction with various other emerging technologies. And so in order to address AI risk management, we actually have to reframe our thinking and take a look at the broader landscape of all the technologies of which AI is just one that we're using in order to improve, for example, patient outcomes.

In the healthcare sector, AI is being used in conjunction with other technologies like augmented reality, for example, where augmented reality technology is used to help with visualization tests and AI helps process large sets of data. And another example of the healthcare sector, AI is being used in conjunction with robotic process automation technology in order to automate some of the manual repetitive administrative tasks, freeing up the time of healthcare practitioners to focus more on other aspects of bettering patient outcomes. From a risk management standpoint, what that means is in order to address the risks that come with using AI, we can't just look at AI as the sole technology that organizations are using.

We have to look at the broader landscape of how AI is converging with other technologies like augmented reality, like RPA, like Haptics technology, or now with Neuralink brain-computer interface technology and more comprehensively assess, given we are using these technologies in conjunction with one another, what risks get dialed up versus down?

Ultimately, I think for our client organizations, I would advise starting with the outcome that we're trying to achieve, and then working backwards to identify, given this kind of patient outcome let's say, what combination of technologies including or excluding AI, are the right sets of technologies to use in order to most efficaciously achieve the outcome that we're trying to get to?

Earlier, Kristin raised the great point of how humans should be kept in the loop. And personally, I hear this phrase more frequently in the context of AI risk management. As AI is used in conjunction with various other technologies, we need to be very clear about what we mean by a human in the loop in order to help mitigate some of the risks that come with using AI and other technologies.

So for example, how do you define the loop or the scope of technological processes that we want to provide oversight to? Who is the right human given the

various technologies that we're using to provide that oversight? And what are the principles of oversight that we should be applying in practice? The answer will depend on whether we're just using AI in isolation, or whether we're using AI in conjunction with various other technologies. And so I would encourage our clients to look at not only AI, but also how AI is being used in everyday life and everyday practices within their healthcare organizations in order to be able to comprehensively map out the kinds of risks that they need to be thinking about when using AI.

## Hala Helm

So that's really critical I think, the idea that we're starting with the why and trying to understand what the purpose is, rather than chasing those shiny objects or the latest and greatest technology, especially in healthcare, where the why is really our mission. So with that, I think we're coming to the end of this. Kristin, any final thoughts that you want to share before we wrap up?

## Kristin Rhebergen

I think one thing to keep in mind too here is as we drive towards these outcomes, as we drive towards that why, a part of the consideration that everybody is going to have to have is what are the digital habits that we need across the board in order to be able to really optimize that human machine teaming? Keeping that human in the loop, making sure that we are recognizing how everything is coming together and where it makes the most sense for it to be a digitally enabled action versus a human enabled action or activity or interaction in order to drive those outcomes that we're trying to look at.

And I think we're kind of on this tipping point where, or on this precipice where digital habits have been something that people have talked about in terms of, 'oh, I use technology, or, oh, I do this'. But we're really starting to think about them in terms of what are the mindset components that you need to make sure that everyone in the workforce has as they approach the use of these tools and these platforms as well.

## Hala Helm

Thank you so much to my colleagues. This has been a great discussion. So in summary, proactive risk mitigation requires cross-functional leadership, a centralized process, and organizations should establish clear policies and guidelines for data privacy, security, and legal compliance while considering human factors and impacts. This will help mitigate potential risks and ensure a smooth integration of AI and GAI technologies.

And that wraps up our second episode of the Healthcare Transformation and Risk podcast series. Today we've taken a closer look at integrating artificial

intelligence and generative artificial intelligence in the healthcare industry, delving into the challenges organizations face and the strategies needed for a successful transformation.

We hope you found our discussion insightful and that it provided a deeper understanding of the complexities and solutions in this critical area of healthcare. Please join us in our upcoming episodes as we continue to explore more facets of healthcare transformation. And thank you so much for listening.

Copyright © 2024 Marsh LLC. All rights reserved.