Podcast

We Get AI for Work: Establishing AI Policies and Governance (2)

By Joseph J. Lazzarotti & Eric J. Felsberg

December 11, 2024

Meet the Authors

Joseph J. Lazzarotti
Principal
908-795-5205
Joseph.Lazzarotti@jacksonlewis.com

Eric J. Felsberg
Principal
631-247-4640
Eric.Felsberg@jacksonlewis.com

Related Services

Artificial Intelligence & Automation

Details

December 11, 2024

Organizations are harnessing the benefits of using generative and traditional AI technologies to enhance productivity, streamline operations, and foster innovation. However, before employing these tools in the workplace, organizations must minimize potential risks and ensure the ethical and responsive use of AI.

Listen on Spotify
Listen on Apple Podcasts
YouTube

Transcript

INTRO

Organizations are harnessing the benefits of using generative and traditional AI technologies to enhance productivity, streamline operations, and foster innovation. However, before employing these tools in the workplace, organizations must minimize potential risks and ensure the ethical and responsive use of AI.

On this episode of We Get AI for Work*, we discuss the importance of developing comprehensive AI policies, the need for confidentiality, data privacy, and bias monitoring in AI outputs—and the potential pitfalls of not properly managing and vetting the appropriate use of AI technologies in the workplace.*

Today's co-hosts are Eric Felsberg, principal in Jackson Lewis' Long Island office, and Joe Lazzarotti, principal in the firm's Tampa office and co-leaders of the firm's AI Group.

Eric and Joe, given the importance of developing comprehensive AI policies, the question on everyone's mind today is: What factors should organizations consider when developing governance structures for the use of AI in the workplace, and how does that impact my organization?

CONTENT

Joseph J. Lazzarotti
Principal and Privacy, Data and Cybersecurity Co-Leader

Welcome, everyone. We're here with another episode of *We Get AI for Work*. As always, I have the pleasure of being with my good partner Eric Felsberg.

On one of the prior episodes, we talked a little bit about governance. It's a real

hot-button issue around AI and implementing AI in organizations. One of the things that we talked about is policies: How do we create policies around this technology? And one of the things we were talking about as well is maybe there's different types of policies — a handbook policy for employees, maybe an IT policy for development, maybe there's an assessment policy or a privacy policy. And what does that really look like? When do we need those?

If we're talking about policies generally — maybe more so a handbook policy for employees, I know there's a lot of interest there in terms of employers thinking about how we govern employees' use of this technology — let's dig into some of what you might expect to see. What are those provisions and what are some of the issues around those provisions? I know one of the things we always think about is: "Hey, are we using generative AI? Are we using more traditional AI? What is that policy covering?"

So, can you talk a little bit about when approaching a policy, what exactly is the technology and how might that impact the policy?

Eric J. Felsberg
Principal and Artificial Intelligence Co-Leader

Absolutely. There's one important thing to think about before we get into the specifics of a policy. It's pretty common that you and I will get a question from an employer that we're working with who says "Hey, do you have an AI policy template you can just send over to me?"

And you and I know our response as well: Broadly speaking, we have certain features we expect to see there, but the policy really should be specific to your particular organization and also, as we touched on this on the governance episode, how AI is going to be used in your particular organization. Is it going to be used, for example, to help select candidates for employment, as part of your applicant process? Or maybe it's going to be used as part of your core business model where you're looking to streamline and make more efficient some of the tests that you perform on an everyday basis.

So, yes, there's not really a template per se. But we'll talk a little bit about some of the features here. And certainly, one of things that you mentioned goes to the use case. Are we thinking about using more traditional artificial intelligence platforms where we're just taking data and we're trying to make predictions out of these data, identify trends? Or are we using the kind that everyone is really interested in currently — generative AI? Meaning, we're using an AI platform that generally speaking is going to create new content; that once we feed it information, it's going to create new content that didn't previously exist. To your point, the issues differ a little bit depending on what type of AI is being used and how it is being used.

With generative AI, for example, when you're creating new content and you're going to rely on that output from the generative AI tool, we're going to talk in a little bit about ensuring accuracy of the output of some of these AI tools: Do you have a method in place for vetting the output that comes out of some of these AI tools, especially in the generative space? If it's more traditional, that still holds

true as well. How do we know that that particular AI tool or that algorithm is operating in a way that we expect it to? That it's performing the calculations correctly, that we feel comfortable relying on that particular output.

There's a few other kinds of nuances that will apply to each one, but it certainly is important to identify specifically the nature and type of AI that you're thinking about using. You really should be thinking about these issues before you actually use it. That may seem obvious, but it's often the case that we get these questions after it has been used for a while and maybe some harm has resulted, for example.

Lazzarotti

That's a great point because early on in the work that we've been doing in this area, both of us had this experience where clients would call and say: What is an AI notetaker?" And the reason they ask is because they got on the phone with a colleague and that technology was being employed. We can't really get into that specifically now — maybe we'll talk about that on a later episode in terms of an example of that technology.

One of the issues we expect to see, and not in all cases, is what uses are approved and how do employees go about suggesting or maybe introducing new technologies that the company might benefit from. What are you seeing there in terms of policies that say "Hey, employees, these are the things you should use, these you shouldn't"? — managing that process so that you don't have that surprise of "Wait a second, some employee is using this, we didn't even vet it." Can you speak to that a little bit?

Felsberg

Yes. This goes back to the point we made on previous episodes, as well as earlier on this one: An AI policy is not something that just kind of drops from the sky. We talked on a previous episode about this notion of a governance committee and how you want to think about the evaluation of these different AI tools.

The first thing you want to do is take inventory of all the AI tools that are out there and also try to figure out which AI tools that we currently don't have in our possession that we think we might want to use in the future to perform some aspect of our business, whether it's more administrative or more substantive part of our business. As part of that deliberation and evaluation of these different AI tools, you want to think about what we should be using this AI tool for: Is it a tool that should be used more broadly with different tasks? The idea being that each individual tool should be vetted for the specific use and it should not be used in a way that is inconsistent with the manner by which it's been vetted.

If we're using some sort of generative AI tool to create employment policies and we've kind of vetted that generative AI tool to help us do that, well, that's what the use has been and what it's been approved for because we've vetted it. What should not happen is using that same generative AI platform to create maybe a user manual, for example, that really impacts the core aspects of our business because that tool may not have been vetted for that particular use. When it's used in a manner that's inconsistent with how it's been vetted, that could lead to really

disastrous results. We don't know if what's coming out of that tool is reliable. Is it accurate? Have we taken steps to try to vet whether that output is accurate?

And so, again, we talked about a governance committee. There needs to be this function where the AI has been vetted. We fully understand what its intended use is and that the different stakeholders that are involved in this have been trained as to how this tool should be used. Any tool outside of that in a policy not only should be prohibited, but we should explain: "Here are the tools that we vetted. Here are the ones that are approved. Here are the approved uses."

We're going to talk about it again in a few minutes, you should also have a function in your policy so that if somebody does come across a new tool or a new use for an existing tool, they have an avenue to bring that up and maybe have that tool vetted for that particular use. We certainly wouldn't want to have folks using it in a manner inconsistent with what it's been approved for already.

Lazzarotti

Yes. There's a lot of flexibility there to structure it in a way that makes sense. Given everything that you're saying, for some organizations that don't have as many departments and employees, some of that might be more streamlined than maybe in other organizations where you have more defined roles. I imagine that certain employees in certain departments might have the approval to use certain tools where others do not. It's really thinking about, as you said, what those use cases are and how best to manage it in the organization and set that by policy.

We talked a little bit about governance already, but one other thing that also comes to mind from a policy perspective and what we expect to see is how do we manage confidentiality and data privacy and security. From my perspective, we go back to the same issues driving these concerns. Use cases are critical and what industry you're in is really important.

From what we're seeing, some companies may already have a confidentiality and privacy and security policy. If they don't, they probably should. But you don't have to start from scratch in an AI policy all the time to establish basic principles around that. Then the question becomes: What data do we need? What data gets input into the tool? Who has access to that data and to the results from what you're trying to accomplish with your tool? So really building on the sensitivity of you may not be using a tool that is managed and maintained only on the company's information systems. It may be a third party's tool so that, in effect, if you're uploading data for the AI tool and adding to the prompt, depending on what you're trying to do, you may actually be disclosing confidential personal information to a third party and have not taken all the appropriate steps to do that.

So that's certainly an area of expectation in policies: Trying to manage the confidentiality, the data privacy issues, and then of course, cybersecurity and data security around the processing of data in the course of using the AI. Anything else you're seeing on that, Eric, for your practice?

Felsberg

For our listeners, when Joe is not in the AI space, he co-leads our data privacy practice group. We often have discussions about these issues even outside the context of AI.

One of things I know you've said to me in the past, which has stuck with me is "Wait, you don't know how these tools are storing data? You don't know who's managing some of these tools?" If you were to enter in proprietary company information or, let's say, you're trying to generate a document that is based on client information, for example, if as part of your business you have clients: This notion of putting into an AI tool either proprietary information, confidential company information or client data — it's tantamount to going out on the street potentially and just kind of stapling it to the telephone poles that are out there for anyone to read because you just don't know who has access to this thing. I always think about that, and that's something we try to convey to our clients: You want to have very strict parameters on the nature and type of information that may be entered into these tools.

To your point, Joe. I don't think I've ever seen a situation in a policy where we feel an employer has felt comfortable, a company has felt comfortable, to enter into that tool client information, confidential proprietary information, whatever else it may be, because of that concern. I think it's just absolutely critical that you have this in a policy because again, to people that are kind of using this casually, that don't spend their time thinking through these issues, they may not think about that. They may say "How great is this? I put in some client information. I get almost a completed memo back. The world is a wonderful place." What you don't know is who can see that information. So, I think about this too, like if you're ever sitting on a plane or a train where people have these privacy guards on their phone and laptops, but yet are entering information into like an AI platform. It's kind of funny in some respects, but it's certainly not a funny issue and something that you really should detail in a written policy. We'll keep coming back to this idea of training, but employees and users need to understand the significance and seriousness of this particular issue.

Lazzarotti

To build on that, there's a lot of considerations that go into what type of data we are talking about. Obviously, we were referring to personal information and others, but there's also intellectual property and copyright infringement issues that a lot of people are concerned about. I know we have an episode that's going to talk about that, so maybe we can kind of table that one for now, but it's certainly a significant issue from a policy perspective.

Whatever it is, accuracy of the data going in and the data coming out and trying to ensure that is critical. Can you talk about that, Eric, a little bit? I know there's a popular case that everybody likes to talk about, especially lawyers, because, of course, that would never happen to us: That story about the lawyer that wrote a brief and the AI had what's called a hallucination and made up a couple of cases — I'm paraphrasing — and submitted it to the court and the court kind of said "Hey, what's this? We never found these cases." But that can happen when you get an output. So, what do you think about in terms of policy there?

Felsberg

It's interesting. I know when AI, especially generative AI, first came onto the scene, there were a lot of stories about this idea of hallucination. It sounds a little funny, but it really is a serious issue where — and again, beyond the scope of this particular podcast — the way that the AI worked is that it produced a result that looked and felt terrific. In that legal case, it looked like a legal citation. It looked reliable, felt reliable, but nobody apparently vetted that. And the idea is that's where some of this generative AI potentially could be a bit nerve wracking because if you just rely blindly on the output, as has kind of been told in some of these stories that we've all heard about, and then that's relied upon and that particular output that's being relied upon is completely inaccurate with the hallucinations, completely made up, complete fiction. That could be really damaging.

So there needs to be in the policy a very clearly laid out directive that if the organization decides that you're going to permit AI, that AI really should be a starting place. It's meant to come up with the general answer, if you will, if you're working on a particular issue. But it's not meant to be the final say. While it may set you on the right path to the potential answer, and it may be that the ultimate answer is the same thing that AI initially said, but there needs to be some human interaction there to independently vet that output.

Again, Joe and I are attorneys, and so this is near and dear to our heart: When you get a citation, go look up that citation to make sure that citation actually exists. And look, the one thing here, give credit where credit's due: The folks that work on these AI platforms. When these things first came out, this issue of hallucinations was a bit more prevalent. I don't see it as much now. But until this is completely nailed down, airtight, and we feel comfortable this never ever could happen — until that happens, we may not ever see that in our lifetime, we just don't know — there needs to be a mechanism in place where results are independently validated, separate and apart from the output of an AI tool. That should be laid out in a policy. And again, just going back to this idea I mentioned earlier — to the uninitiated, they may not even think about this — it's all of our jobs to make sure people understand how this stuff works and how much reliability they can afford to the output of some of these tools.

Lazzarotti

Reviewing the output of AI is important. We're still humans. We still could find some errors in it, even though it sometimes seems perfect. It's important to look at that. But there's a different iteration of that. Instead of a case citation or a case name or description being inaccurate, you can have a result that maybe has some bias embedded in it.

Eric, you've been a leader in the firm for a long time around doing analyses for affirmative action and government contract work. I know that as part of our team, you help clients with bias audits and examine a lot of the questions around this. It'd be really interesting to hear from a policy perspective: What do you recommend around policy to monitor for that kind of bias, which I think is a big issue. particularly in the HR space?

Felsberg

Yes, that's right. This is an issue that certainly everyone should be thinking about if you're using AI as part of your management, your personnel management process — so again, selecting candidates for employment, who you're to promote in some instances, who's likely to terminate and things of this nature, pay administration. If you're using these tools for anything like that, you have to think about this issue of bias.

When we think about bias, there could be situations — and what I'm about to say is a bit uncommon — where you may have an AI tool and algorithm that is affirmatively taking into account a protected characteristic when making a calculation and producing output. It seems kind of crazy that we have to think about these issues, but what if the algorithm is taking into account somebody's race when determining whether they would be a good candidate for employment? Now that seems probably obvious to all of us that that is not lawful, but what about the situation where the algorithm is taking into account facially completely neutral criteria, but for whatever reason, it's having a disproportionate impact on certain demographics. No intent. It's just having an impact. These tools and the output that you're relying on need to be monitored for impact against certain groups.

For those that practice in this area, that's known as a disparate impact analysis or even a disparate treatment where it's affirmatively taken into account. That may trigger some obligation either on your part because you are the user, but by extension, the developer, to have the tools validated — have an independent study as to how the tool is operating, which criteria variables it's considering, and how it's assigning value and how it's producing an output.

You see this in some of the regulations that are out there. For those of you that have employees or jobs in the City of New York, you know that the City of New York has a law dealing with automatic employment decision tools, AEDTs, that is clearly focused on this particular issue, meaning does it have a bias against certain demographics? This is a huge concern for the EEO agencies, both on the federal level and emerging on the state and local level across our country. It's a significant issue.

Now, do we want your employees out there all doing their independent bias audits? Well, absolutely not. But — and for a variety of reasons, which are probably beyond the scope of this particular discussion — certainly it's advisable, again, subject to your specific situation, to put this in your policy, to say these tools are going to be monitored for bias. We have to make sure that we do this. Before you embark on any sort of project using a new tool that hasn't been vetted, consult with, and again, we talked previously about governance committees or the Office of General Counsel or the head of human resources, so that we can ensure that whatever platforms that we're using, AI platforms, that they're either free of bias or, if we do see disparate kind of selection rates or recommendation rates, that the underlying technology has been validated. Really, really important in this area because again, this is a focus for lot of regulations that are out there right now.

Lazzarotti

That makes a lot of sense, Eric. From the standpoint of any policy, we think about who's going to address questions about the policy. How do we enforce the policy? What is a violation? And who's going to be the point person to deal with that — and how do they? Is there a structure for managing that without getting it too complicated? But again, just understanding the application of policy, what are you thinking there in terms of approaching that as companies develop these policies?

Felsberg

There's a couple of things. One is you want to have an outlet for individuals to ask questions about the existing policy, existing tools. You also want to have an outlet for folks to propose new tools or new uses for existing tools. And then, lastly, an outlet for folks to report violations of the policy. Right now, how do they do that?

Well, the important thing is you don't want to make it difficult for the user. You don't want to have an employee trying to hunt around figuring out who the heck am I supposed to be reporting this to. This certainly should come out of some sort of governance committee. But again, that should not be left up to the employee to figure out — do I call the counsel or should I call HR? A lot of times, what will work with clients is to develop an email hotline, a telephone hotline, or a dedicated online form that can be submitted and that goes to this governance committee. Depending on the nature of the concern, the inquiry, then that committee is responsible for getting the particular stakeholders involved. If it's an inquiry around a potential violation of a policy, when it gets to the committee, maybe that's something that legal and HR need to focus on. If it's "Hey, I heard about this new notetaking technology and think it'd be really terrific for our engineers and their engineering meetings to do." Maybe that needs to go to IT compliance and the business leaders and et cetera, et cetera.

The important thing is you want to make this easy and seamless. Because what you don't want to have happen is that, if it's too complicated, the employee might just say: "Yeah, forget it. I'm just going to use it, right? Who cares?" And they use it. Now suddenly we have a problem later on. Whatever method you use, email, phone, whatever, that should be detailed in the policy. And again, I mentioned training a couple of times. Certainly, the outlet can be conveyed during a training session with employees so that they know exactly where to go when they have these issues.

Lazzarotti

That makes a lot of sense. We covered a number of things about policies. One thing you mentioned at the outset, Eric, was that this is a pretty rapidly developing area and it's going to have changes. Like any good policy, you want to have a mechanism to amend it and account for changes and think about how to communicate those changes and make the new policy effective.

Any other thoughts, Eric, on policies and procedures? I know we're going to be touching upon this in upcoming episodes in the context of specific AI

technologies. Generally, it's a really important topic, so any other thoughts there?

Felsberg

Like we said at the outset, this area is developing so rapidly and we're in this interesting time here where on the one hand it's exciting. We're at the forefront and seeing this really incredible technology really taking hold.

But in terms of developing policies, there should be a mention that this is a policy that is subject to frequent revision and update because you have to keep up with the developing technology. And it's probably a good idea to have in there that each time somebody embarks on a project where they're going to leverage AI that they consult the policy again. That's going to be the case for the foreseeable future until this technology settles up a little bit or settles down and becomes part of our everyday life. But right now, it's developing so rapidly these policies have to be revisited frequently. That would be my one last point that I would make.

Lazzarotti

All right, well, as always, my friend, good presenting with you. I hope this was helpful.

If any of the listeners have any questions or recommendations or thoughts, please reach out to us at ai@JacksonLewis.com. Until the next time, thank you so much.

Felsberg

Thank you, Joe.

OUTRO

Thank you for joining us on We get work™. Please tune into our next program where we will continue to tell you not only what's legal, but what is effective. We get work™ is available to stream and subscribe to on Apple Podcasts, Libsyn, SoundCloud, Spotify and YouTube. For more information on today's topic, our presenters and other Jackson Lewis resources, visit jacksonlewis.com.

As a reminder, this material is provided for informational purposes only. It is not intended to constitute legal advice, nor does it create a client-lawyer relationship between Jackson Lewis and any recipient.