

Regulating Artificial Intelligence in Employment: How to Comply and Best Practices Transcript

Wed, Dec 14, 2022 2:28PM • 1:33:07

SUMMARY KEYWORDS

ai, law, tool, bias, companies, audit, applicant, employer, regulations, new york city, automated, data, required, draft regulations, questions, algorithmic, decision, employment, candidates, privacy

Natalie Pierce

Hello, everyone. And welcome to a discussion of a hot topic involving data privacy and employment law. And this has to do with a regulating, automated decision in employment. Listen, we know that we chose for you, those of you who are joining us live, we understand that the World Cup semifinal between France and Morocco is happening, now. We also know that some of you may be multitasking. So we have muted our attendees, but we invite you to submit questions through the chat in the next 75 minutes. And by the way, we'll be answering those questions at the end of the chat. In the next 75 minutes, we are going to be dealing with several questions.

And those areas will address five specific things we want to cover. The first being talking about how we got here, what have been some of the trends that we've seen in the last 10 years. And then second, we're going to talk about proliferation of the use of artificial intelligence in the workplace and some of the promises and concerns that that raises. Third, we're going to do a deep dive into a New York City Law which goes into effect next year. And then fourth, we're going to take a look at one of the components of that law which is the AI bias audit process and then fifth we're going to look, do a brief overview of some state, national and international trends that are really driving you movements in this area toward greater transparency explainability accountability.

And with that, I want to and by the way, we're also going to be looking at some solutions. With that, I want to introduce our fantastic panel, we're going to have everyone introduce themselves. And I'll go ahead and start with myself. My name is Natalie, I head of Gunderson's employment practice. And then let's go ahead and go next to Michael, and then to our data privacy team and then to our friends at holistic AI.

MEET THE PRESENTERS

Michael Hui

Hi, I'm Michael I am of counsel of Gunderson's Labor Employment Team.

Anna Westfelt

Hi, everyone. My name is Anna Westfelt and I'm a partner in Gunderson Dettmer's San Francisco office. I head of our data privacy group, and we also do a lot of work in the AI space together with Natalie.

Frida Alim

Hi, everyone. My name is Frida Alim. I'm a data privacy associated to Gunderson San Francisco office and help our startup and venture capital clients on all things data privacy and security.

Natalie Pierce

Great. And I'm not sure. Do we have our holistic

Emre Kazim

Hi, everyone. I'm Emre Kazim. I'm co founder and CEO of holistic AI.

Adriano Koshiyama

Yeah. I'm Adrian Koshiyama. I'm the co founder and CEO of holistic AI. Thanks for Natalie, for inviting us really pleased to be here.

Natalie Pierce

Great. Well, thank you so much for being here. And, Anna, let's go ahead and kick it off with you and sort of some trends and how it is that we got here.

OVER THE PAST 10 YEARS, AI AND MACHINE LEARNING HAVE REVOLUTIONIZED THE WORKPLACE, DRIVING EFFICIENCY AND INNOVATION

Ana Westfelt

Absolutely, thanks, Natalie. So this is a really interesting time for AI and machine learning both on the technology side and on the regulatory side. And with the exponential growth in computing power, we've seen them with incredible advancements in AI and machine learning technology in the past 10 years. So many of us attending this webinar today use smart speakers at home that have benefited from advancements in natural language processing.

Many of us have marveled over the incredible sophistication of deep fakes. And we have had a lot of fun with AI GIF image generators. Also, just in the last month, we've seen the chat TCP but providing answers to questions. And sometimes those answers are better than what humans can come up with.

So there are many ways that we use AI ourselves in our daily lives. But AI is of course, also used by companies and government agencies that we interact with, often on a daily basis. And that includes our workplace, which, of course will be the focus of today's webinar. I so all this AI that is I used by all of us on a daily basis is it's helpful. It's efficient, is driving innovation. But of course on the privacy side, we've seen over the last 10 years and increased attention on all the data that is collected to develop and train AI and machine learning technologies.

And just kind of looking at the basics. In order for a machine learning system to continue to evolve and improve, it needs new training data. The essence of the machine learning system is the ability to find and correlate patterns. And that generally requires a very large data set and continued infusion of new data.

We will talk today about how different laws tried to address the harms to individuals that can arise not just from how all these advances in AI and machine learning affects us in our daily lives, but also how the harms can result from how they're these systems are programmed. And also importantly from the underlying data set that is used to train the machine learning.

And at this time, we have a really high level of public awareness and concern over the use of AI in everyday life. But legislation has not kept up with the pace of development. And it has not kept up with this increase in awareness and concerns. However, as we will cover today, there was a lot going on in a regulatory space and we accepted expect to see some very impactful law coming into effect in the near future. So with that, I will hand over to Natalie who will talk about how how AI is used in today's workplace.

AI IN TODAY'S WORKPLACE

Natalie Pierce

Thanks, Anna. You know, I'll start with giving some statistics that I think help explain how it is that we are here and looking at more regulation of use of these technologies because I think many people will agree there is a concern that if the area is too regulated that it could very well stifle innovation and just the Equal Employment Opportunity Commission, I think, and Michael will talk about this more later. But the reason that they really are looking at least having some guardrails, some guidelines, and we have places like New York City that are now coming out and actually legislating the area, is because 83% of employers are currently using some form of automation to screen candidates to gauge employee performance. And when it comes to the fortune 500, it's actually 90% of employers that are using those technologies. And, and what we've seen is that that oftentimes, these tools can do a better job, they can do what is impossible, such as screening, 1000s of applications, they can do things like Unilever, for example, who uses a tool called biometrics, and has applicant engaged in games that are doing things like testing skills, risk, aptitude, and doing these measurements in a way that is linked directly to job skills. And tools like HireVue, or who are applying the same sorts of questions, criteria to applicants. And, you know, the question that I think we really need to ask when we look at the, at the regulation and and how far we go, really is Can humans do it better at the end of the day, because what studies have shown is that human biases, when it seeps into the interview process, you really have more of a match of 25% selection or accuracy of what the interviewer is seeing versus the predictability for jobs success. We're also seeing these tools used to do things like help employers measure productivity, to give them greater sense of security, because of the proliferation of remote workforces. And we also are certainly seeing how these tools are helping the automation of training. I think many of us are very cognizant of the fact that we have a tremendous skills gap coming. And so there are many types of skills that are not as prevalent in the workforce. And so the need to really find ways to upskill and, and automated tools are also helping in in that in that area as well.

So ultimately, you know, again, it is are these tools doing more harm than good? Are they helping to augment what humans can do in terms of their decision making, but the reality is that people are scared of what they don't understand 35% of global consumer consumers say that they simply don't trust artificial intelligence. And when it comes to who should be liable. the 77% of global consumers say that absolutely. The companies that are using these tools should be held liable. Okay, so, so what's the problem? AI has tremendous potential but the proliferation of AI in the workplace brings some problems such as artificial intelligence bias. This is a phenomena where an algorithm produces results that are systematically biased due to historical bias is found in data sets or models. This can lead to several questions. One can we trust AI, to will it be fair or simply perpetuate biases that are embedded in historical data that's used to train and models? Three, do we understand how the decisions are made? Are they transparent and explainable? Four what about accountability when things go wrong is the user responsible the developer both, as we've seen and will continue to see these concerns are fueling new regulations like the New York City Law, which we're about to dive into.

Now, when it comes to AI regulation in the workplace, we think that this is just the start of a growing call for more transparency, as we talked about explainability, fairness and accountability. The good news is that the FTC and other agencies and organizations provide several practical pieces of advice to help comply with the laws and public calls to action. And some of these are one watch out for discriminatory outcomes at various stages of the lifecycle. two Don't exaggerate what your algorithm can do. For example, if you're a developer Don't say things like it will provide 100% on bias hiring decisions. Three, tell the truth about how you use the data. And four the FTC, as general rule is that products should do more good than harm. And the five just remembering that a growing number of bodies will start holding those accountable when, when things do go wrong. I want to now turn it over to Michael and Frieda who are going to do that deep dive into what is this new New York City Law.

NYC'S AUTOMATED EMPLOYMENT DECISION TOOL LAW

Michael Hui

So let's deep dive time into New York City's automated employment decision tool law. So in December of 2021, the New York City Council enacted local law 144. And as Natalie mentioned, it's going to be effective next year, January 1 of 2023. So under local law 144. It's going to be unlawful in New York City for an employer or an employment agency to use an automated employment decision tool to screen a candidate or employee unless certain requirements are being met. In the next slide, we are going to discuss what those requirements and obligations look like. And to the extent you're wondering, you know, what the heck is an automated employment decision tool, we're going to cover that as well in subsequent slides. So under local law 144, there is an enforcement mechanism, it's going to authorize New York City's Corporation council to enforce this law by permitting it to file suit in a court of competent jurisdiction. Also, this, you know, claims under this law can be brought before the New York City Office of Administrative trials and hearings to seek civil penalties. And these civil penalties can range anywhere from \$375 to \$1,575. So, you know, what is the current state of local law 144. relatively recently, on November 4 of this year, the Department of Consumer and worker protection held a hearing to provide an opportunity for the public to weigh in on sort of implementation of this law. The hearing was very popular, and very well attended. And during this hearing, unfortunately, the

department didn't provide any further guidance on exactly when the proposed rules would be finalized. But currently, on the department's website, it has indicated that a second public hearing is going to be held to take further comments, and that enforcement of this law is going to be delayed until April 15 of 2023.

OBLIGATIONS UNDER NYC AEDT LAW

Okay, so I had talked about, you know, in New York City, it's going to be unlawful unless certain obligations are met for an employer to use an automated employment decision tool. So let me talk about what those obligations and what those requirements kind of look like. So an employer needs to complete an independent bias audit of the tool no more than a year prior to the use of the tool. Independent bias audit. The law defines it as an impartial evaluation by an independent auditor. I know what does that mean? So my colleague Frida, she'll get into it, but at a minimum, the bias audits going to need to include testing of the tools despair impact on race, ethnicity, or sex and disparate impact is sort of how discrimination can appear. That is not overt or directly applied to an individual, it's a little bit more insidious. Also, under local law 144 at an employer must publish on its website a summary of the bias audit results. So right now, the draft regulations state that the summary should include the date of the most recent by an audit, it should include the selection rates and the impact ratio for the categories being looked at as well as the distribution date of tool. And, you know, employers must also employers and employment agencies in New York must also give 10 business days notice to employees and candidates in New York City that the tool will be used. So the notice needs to include the job qualifications and characteristics, the tool is going to be measuring and assessing that candidate in that interview, for example, and I know there is probably some ambiguity in regarding you know, who is exactly and applicants. But to be clear, you know, this applies to New York City residents who have applied for a position. So according to the draft regulations, this is going to be an individual who has actually applied for specific employment position, and who has provided necessary information, something like a resume or application. The draft regulations also state that the notice provided can be done so through, for example, like in the career section of the website of the employer, or employment agency, so long as so long as it's conspicuous and clear, and can be posted in the job posting, the notice can be given through mail or email. And with the notice, the candidates must be permitted to request an alternative selection process or an accommodation. And finally, under local law 144 employers or employment agencies must make available when it's written upon written request when someone requests it in writing about the source, the type of data that's being collected by the tool, and the employer or employment agency also needs to provide its data retention policy, unless for some reason, the disclosure of such information is going to result in some sort of legal violation or it's, you know, interfere with a criminal investigation and once requested by an employer or an employment agency must provide that information within 30 days of the written requests.

WHAT'S AN AEDT

Okay, so you heard about the obligation. So let me kind of get further insight on what exactly is an automated employment decision tool. So the regulation defines it as any computational process derived from machine learning, statistical modeling, data analytics, artificial intelligence, that essentially

generates a simplified output that's going to substantially assist or replace human discretion in screening candidates in New York City.

I will say at this time, it's a little bit unclear how much assistance is required to be considered substantial assistance. But with that said, what this law is really, you know, being focused on is things like algorithmic software or tools that can help identify, evaluate or select candidates in the interview screening process. You can think of it as, like AI, tools that filter resumes, based on keywords, gamified interviews that look at skills assessments, or you know, other skills testing through software. The draft regulations do make clear that automated employment decision tools do not include things like you know, a junk email filter, firewall, anti virus software, spreadsheets, underlying databases, things of that nature.

WHAT SHOULD A BIAS AUDIT COVER

Frida Alim

Thanks, Michael. So I'm going to talk about what the bias audit should cover under this law. So local law 144 requires that a business is conducted by an audit that looks at the ADTs disparate impact on individuals in protected categories such as race, ethnicity, or sex. Unfortunately, those are the only parameters that are actually set out in local law 144. But a couple months ago, we received draft regulations from the Department of Consumer and worker protection that have additional detail on how the buyer's audit should be conducted. So under the draft rights as we currently have them, again, these could be subject to change. The bias audit should calculate a selection rate, as well as an impact rate for each category that employers have to report on the EEO C's employment employer information report, which would include race, ethnicity, and sex. So the regulations provide that there are two ways of calculating the impact ratio, you could divide the selection rate by the selection rate of the most selected category. The other way to calculate the impact ratio is to divide the average score of individuals in a category, for example of race, ethnicity, or sex by the average score of the individual in the highest category. So note that the the term selection rate here means the rate at which individuals in a category are either selected to move forward in the hiring process or assigned to classification by the AEDT. So for example, if the AEDT selects 40, out of 100, Hispanic women, the selection rate for Hispanic woman would be 40%. So ultimately, what this will look like is essentially a table that you then have to publicly post on your website with the statistics that are required to be covered by this by the bias audit. And the idea here is that the candidate or the regulator could navigate to the company's website and see the impact that the use of the tool has on individuals in specific protected categories. And if you're curious about what this would look like, you can take a look at the draft regulations, which has an example of essentially what that table would contain and how it would look on your website.

THE GOOD, THE BAD, THE UNCLEAR

And next, I want to talk about what are kind of positive and negative aspects of this law and what's unclear, especially as we await final regulations from the regulator. So first, to tackle the good. From a consumer protection perspective, the law does deliver on providing accountability and transparency. You know, data privacy laws in the US have really been trending in that direction over the last few years. And part of the objective of the law is to make sure that candidates and employees have

meaningful information about what characteristics are evaluated by the AEDT. So as Michael mentioned, candidates and employees are entitled to a disclosure that states the job qualifications and characteristics that the tool has looked at in coming up with a recommendation. And the law is designed to introduce some accountability by requiring that covered entities that use AEDTS provide some insight on whether certain classes are disproportionately impacted by the AEDT. So as with any new law, there are also some challenging aspects of this regulation. So as we discussed earlier, some of the definitions such as you know, the key definition of an AEDT, are a little bit vague and could potentially capture some activities that perhaps New York City Council didn't actually mean to regulate. Another issue here is, is you know, especially for smaller companies, this law applies regardless of the size of the business. So in many of the privacy laws we've seen passed in the US in the last few years, legislators have drafted these laws so that they apply when a business hits certain thresholds such as a revenue threshold, we don't really have that in this law. Another issue is that as currently written in the regulations, if you publish just what set out to be included in the bias audit, it may not actually show the full picture of whether the ADT has a discriminatory outcome. There are some aspects of this law that are also a little bit unclear. So for example, the covered entity has to provide the candidate or the employee with instructions for how to request an alternative selection process or accommodation. But the regulations also say that the entity is not required to provide an alternative process an alternative selection process. So it's a little bit unclear how that notice is actually going to jive with company's practices. The law also states that, you know, specific information has to be disclosed, but it's silent on whether you can present additional information alongside the bias audit. For example, there are questions around whether a company can or should show that there are business or job related legitimate reasons for observed differences and the different treatment of different categories under the law. And that traditionally forms part of the disparate impact analysis, but it's not explicitly addressed in the regulations at this point. It's also unclear which candidates or employees should be included within the scope of the bias audit. So for example, should the bias audit only look at New York City candidates or is it acceptable for the bias audit to look at candidates globally. And now I will hand it over to Natalie to provide an overview of what businesses should start doing now to comply

NYC AEDT LAW- POTENTIAL SOLUTIONS

Natalie Pierce

Thank you, Michael. And thank you for you too. That was fantastic summary of New York City Law, we still have unanswered questions. But we are well poised, I think, to start to shift towards some of solutions. And so with that, we're very happy to be joined by a holistic AI. But let me just summarize this slide for you. Knowing if you're required to comply with the New York City Law, and if so, how do you get there? So here are the five steps outlined on this slide? One, are you a covered entity? Do you use automated tools to screen candidates for employment? Or employees for promotion in New York City? Two? If so, determine if your tools qualify as automated employment decision making tools. Three, if the answer to number one and two is yes, then you must use an independent auditor to conduct a bias audit of your ADT for publish the results of the bias audit and give notice to candidates and employees that you're using the ADTs. And then five, let them know how to request alternatives to using ADTs. So now it's it is with great pleasure that I'm going to turn it over to holistic AI to talk about conducting the audit and and what some of the reports might look alike. And I think it's over to you, Adriano.

HOLISTIC AI

Adriano Koshiyama

Thanks, Natalie. And thanks Gunderson team for the invitation. Just very briefly, about holistic AI, we basically help companies adopt AI with confidence, we be doing work on this space on the research capacity. Over the last five years, we were very interested from the early days on AI ethics, responsible AI, this was all academic life. And nowadays, we are really way more commercial than we used to be. And more recently, with the whole New York City bias audits, legislation coming from last year, we've been helping companies, ranging from HR tech companies, to enterprise companies to conduct this bias audits. And the purpose of me here today is to explain to you how the process works, and show you some of glimpses of how some of the results and all the bits that we produce. Looks like. Could you the next slide, please?

BIAS AUDITING PROCESS

Yes. So what is the process about so usually is a process that takes a few weeks, it's actually getting closer and closer. Now it should be about two weeks now, for us to do a bias audit used to take four weeks when we started. But now we've been optimizing ourselves a lot. And technology plays a good role in enabling us to get to move much quicker than before.

Typical process involves a kickoff call. Usually it's run by a buyer solely to lead from our side, as well as a customer success. Lindsey, who might be in the call, she is going to be your key point of contact and the whole conversation. After a kickoff call when we basically set out an agenda to work with your team. Part of the agenda works on performative data extraction as the goddess and team have mentioned Frida in my Michael, very important is to calculate these impact ratios. It's a core element of the buyers audit. So that's where really ensuring that we collect the data, ensuring the data is in the correct format and templates. It just really speeds up the whole process of randomized audits. And we also try to collect some documents around how does that AEDT works, how the scoring is performed or the recommendation is performed. So then we can start to really draft this buyer's audit. After about one week, I know that sometimes it takes stress takes more time takes less time. But it broadly speaking from our experience, it's actually has been quite quick. We've discovered that a lot of companies in this space already have been conducting this in some shape or form internally. So it's quite straightforward in many cases. After this process is done, we try to calculate those metrics very quickly and report back to you with the impact ratios. So we can provide a feedback reports. And we sometimes there might be some gaps on the documentation that will be requested. Sometimes we produce a draft describing the system. And we'd like to bounce back, just to see whether there could be some clarifications, or maybe we maybe we have misrepresented parts of the products of just to get some feedback, and use us as a sounding board. Second week, most of the time, we spent on refining the results that we found on the first on the first week, sometimes the data was not on the correct shape. And maybe we miss the protected characteristic, maybe some impact ratio was a bit off because we lacked sample size. So you know, there is an element of back and forth normally happens, not a big problem. By the second or third week, we already have a good shape, a bias audit in a draft form, that you can circulate insights good with your counsel. And then we can start to go towards actually getting to a final form of the report after by the end of the second or the third week. In the meantime, the team is try to run some training

sessions. Some of you might be interested in learning more about neostigmine, New York City legislation, Lindsay tends to run sessions with your customer success team, because they are in the front line speaking with your own clients about it, so it's good for them to have some q&a. So we have all this collateral as well, helpful. Elements about the guidance itself, we tend to as was highlighted by Frida and Michael and notices, etc. We provide some guidance, but it's something that you have to work with your counsel, we recommend that very, very importantly, because notices are very important elements. And we've seen this with legislations like GDPR and even Illinois for example. Noticing simply use it as a way to claim and, and create damage. So

KEY DELIVERABLES

these are the in a nutshell, in about three weeks, we go from reading documentations getting the data to actually producing a buyer's audit getting go to the next slide. And then the next slide, we basically demonstrate the two key deliverables out of this process. First deliverable is internal buyers order reports. We're here, I'm just trying to highlight in a snapshot, the structure, there is an executive summary, we then highlight the methodology apply to do the BIOS audits, for example, how the the data process gathering looks like, what are the key elements of the new oxy legislation, and then how your system for example, or this organization qualifies under that. Then we describe your company, your service organization, we describe the systems, the ADTs that we had analyzed and then finally, demonstrate the impact ratio results. In some cases, and that's has Frida and Michael alluded to sometimes it's a bit of both if you want to demonstrate above and beyond compliance with, we also perform a controls review, it's generally useful for you to demonstrate that not only you got the impact ratios, right, but also you have good ways to prevent the text incorrect bias. It's optional. But it's always important to do this bit. And then finally, any recommendations or appendix we do, as you may imagine, is a good software company, lot of open source a lot of technologies that can be used to mitigate bias, as well. So if anything shows up during the process, we can we can point to you and provide some open source technologies that might be useful for this process. The second bit, which is part of the audits, and that's the public part is the summary for results webpage. And I think if we can go to the browser, if possible. I can briefly show more or less how the app works. I'm just running here a demo.

Basically, in this demo, you have three systems. So you have a company called holistic AI has, what is the AI is running a biased audit for itself. And holistic AI has three AEDTs. One is called candidate assessment. The other one's called candidate match and mobility. For example, there's three AEDTs a lot of the companies that we work with have one or two in one of them This three domains, broadly speaking, can you go to mobility? AEDT, is perfect. And it's just a dashboard, demonstrating some numbers about sample sizes, delivery scores, etc. It's not important, but you can see here on the left side, you have some collateral that get stored there, like the audit reports, the system description and deployment summary results. Can you click on deployments, please? Yes. So the deployments are basically a list of where this ADT has been deployed to do this, it is being used by CBOE, for example, CVS civic blahs, etc. So it's all dummy. No one is using this at the moment disclaimer. So if you click for example, in CVS, you have a summary of results page, that, for example, you could share with CVS with the key information that you would need, in summary results, ready for a description of the system, the key information, the mobility, use to make a scoring about a candidate's scroll down as you please. And then finally, the impact ratio is at the bottom, on a standalone form, and intersectional form that is

required by the legislation. And there's a left bar there, explain the system, the mobility, distribution, data, etc. And if you go back on top of IQ, a top of the of the page is perfect is going. There's a button there called Share. Let's see you excellent. You see the page is private. So no one can see it only you can see this page. But you can, for example, click on guests. And you'll be able to invite a guest, maybe someone from CVS that you'd like to share this results in private form for them. Or if you click on public, you get a URL. So now anyone in the words, can access this web page. And now it's public. And for example, you could share with your clients, and they could embed this on their own websites. So if you have 1000 clients, it saves a lot of time for you to be sharing a summary of results. Web pages, for example. That's pretty much from my from my side. That's part of the deliverables you get from this bias audit process.

Back to the slides should be final. Yep, the rest is some useful materials, feel free to visit our website, there's plenty of blogs and things for you to read them, and have fun. Thanks for the opportunity, Natalie. And I'm going to something thanks, Emre.

OTHER STATES REGULATING AUTOMATED DECISION-MAKING IN EMPLOYMENT

Natalie Pierce

Thank you, Adriana. And before we move on, you know, for even before we had the New York City Law, we really have been suggesting to employers the importance of monitoring these tools through the lifecycle. So to really think about doing things like pilot testing, to make sure that there are unintended consequences. Sometimes there's the very best intentions. And it may be to really mitigate against human biases to further diversify the workforce. And it's with that reason the tools are used, but it's oftentimes some unintended consequences and what we see happening, and so I guess the message is really think about doing these types of audits, even if you're not yet legally required to do these types of audits. And all of us who are participating on today's panel, we are of a mindset that what we're seeing in New York City, really is just the start, we can expect to see much more legislation in this area and without reason we thought for this final section, we want to take a step back and really look at other state laws, some national movement, some international moves in this space as well. So I will start and then I'll pass it on, but a brief To look at what's going on in other states, so Illinois, for example, has something called the artificial intelligence video interview act. So if, if a company uses AI tools to analyze video interviews, it requires the company explain how the AI works and get the applicant's consent. And while the Act does not provide a specific remedy, or private right of action, there, Illinois has a similar law. And that law went into effect in 2009. It's the biometric Privacy Information Protection Act. And what we've seen happen under that law is that there have been hundreds of class actions millions of dollars in settlements. And in fact, just within the last couple of months, we have the first case go to a jury trial that resulted in a \$220 million verdict against the company for using fingerprint scans of truck drivers without adequate notice, and consent. And so while we don't have we still don't have cases under the use of AI and video interview. We, we do have, for example, at least one company who's our who's now arguing that the AI video interview Act should free enter the biometric Privacy Act. But that issue has not been resolved, we do suspect that there will be more states who will be really looking to hold employers to account. And as we've seen through the Illinois example, when there is a remedy available, the damages can very quickly add up.

Maryland is has a law that's similar to Illinois companies who use facial recognition technology in job interviews, let's get applicants written consent. Now the fines in Maryland can be \$10,000 per violation, up to a total of \$600,000. Now, California, has a privacy rights act that will begin applying much of the chagrin of many employers will begin applying to employees and candidates in 2023. And Frida will be discussing that law in more detail. I'm not sure if next but further down in our presentation. All right over to you, Michael. Great.

GUIDANCE FROM EEOC & DOJ

Michael Hui

Thank you so much, Natalie. So I'm going to discuss about the guidance that the Equal Employment Opportunity Commission and Department of Justice has issued regarding the use of algorithmic decision making tools. The guidance really focuses on the intersectionality between these sorts of tools, and disabilities. In the guidance, they really focus on three primary areas where disabilities and algorithmic decision making tools can intersect. The first area is the arena of reasonable accommodations. The second one is what they call screen outs. And then the third one is focused on these tools potentially violating the ADA A's prohibition on asking disability related inquiries. So let me talk about the first one, reasonable accommodations. So this issue arises when an employer doesn't provide a reasonable accommodation that's necessary for a job applicant or employee to be rated fairly and accurately by the algorithm. So under federal law, there is an obligation to reasonably accommodate an employee with a disability, so that way they can do the job and this also applies to applicants. So for example, you know, let's say you have a job applicant that has some sort of manual dexterity issue that's caused by a disability and let's say they need to take some sort of knowledge based test that's being administered through computer which requires the use of a keyboard a mouse, you know, other sort of computer equipment. So, it is it might take them longer or they might have trouble answering these knowledge based questions even though they have the knowledge. So if these responses are for example, like times in which more often than not are, this kind of tests will will not accurately measure the applicants particular knowledge on whatever is being tested. And so the second part of the guidance is focused on screen screening outs. So this occurs when an employer relies on a tool and algorithmic decision making tool that effectively screens out an individual with a disability from the interview process, even though they can do the job with a reasonable accommodation. So essentially, the disability prevents the job applicant or the employee from meeting certain selection criteria set forth by the algorithmic decision making tool, and then the employee or the average applicant essentially, loses the opportunity as a result. So let me provide an example something kind of concrete to understand what that means. So let's say the guidance uses the example of like a chatbot. A chatbot, might be programmed to kind of reject certain applicants based on any gaps in their employment history. So if an applicant has a gap, their employment history due to a disability, or maybe they were seeking treatment for disability, then that chatbot my essentially function to screen them out from the interview process because of that disability. And now, there's sort of an example of screen out occurs, for example, like in video interviewing, you might have video interviewing software that kind of measure speech patterns determine sort of an individual and applicant's ability to solve problems. So if you have a individual with a speech impediment that rises to the level of a disability, it may cause you know, kind of differences in speech patterns. And that sort of video interviewing software might pick that up. And that could result in that applicant being effectively screened out and rejected because of the

speech impediment not because of a not lack of knowledge that's being tested. And finally, about sort of the ADA is prohibition about asking disability related questions. And algorithmic decision making tool might violate that if the assessment includes disability related inquiries, that might elicit or might directly ask information about an employee or applicants, disability. Next slide, please, Emri.

ANTICIPATED REGULATIONS IN CALIFORNIA

And then I'm going to talk about sort of sort of anticipated regulations that are going to be coming out in California regarding the use of automated decision systems. So on March 15, of 2022, the California Fair Employment and Housing Council, which is now known as California Civil Rights department, they publish draft modifications to existing employment anti discrimination regulations, that incorporates regulations that sort of determine the use of automated decision systems. So under these dropped regulations, an employer or a covered entity cannot use automated decision systems that screen out or tend to screen out an applicants or an employee based on a protected characteristic. In California, a protected characteristic is much broader than for example, what was being tested in the bias audit in local law 144. So under these draft regulations, there's not prohibition unless the employer or the covered entity can establish that, you know, the screen out is job related for the position, and it is consistent with some sort of business necessity of the company.

So, New York City uses the term automated employment decision tool. California calls it an automated decision system. And California defines it as a computational process. One that's basically derived from machine learning, statistics, other data processing or artificial intelligence techniques that screens, evaluates, categorizes, recommends, or otherwise makes a decision or facilitates human human decision making that impacts an employee or an applicant. So it's somewhat similar to New York City's definition. So notably, under these draft regulations, there can be liability not only on the company, but also on third party, third parties and administering the artificial the automated decision system tool, so this might include vendors, or other third parties responsible for these AI tools. Under the draft regulation, regulation, there are robust record keeping requirements as well expanded the requirements to keep such machine learning data to four years. So employers covered entities need to keep, for example, like datasets that's that are being used to train the automated decision system datasets that are provided by the individuals, the applicants, the employees, or data that's produced by the automated decision system.

STATE REGULATION OF GENERAL PURPOSE AI

I will also add in terms of record keeping the draft regulations also require persons or entities engaged in advertisements, sale provision, or use of the automated decision system to preserve that sort of assessment criteria that's being used by the tool for you know, each employer or a covered entity to whom the tools provided. Currently, the state of this is sort of still up in the air, there have been public workshops and review sessions on the draft regulations in California. But you know, as of right now, it has not been officially adopted and sort of adoption date is currently unknown. Frida I think I'm gonna hand it off to you.

Frida Alim

Great, so I'm just going to quickly talk about some state general purpose AI regulations that are coming into effect next year. So Colorado, California, Connecticut and Virginia have passed general purpose privacy legislation that have certain requirements regarding profiling or automated decision making. So these laws provide that a business that uses automated decision making that produces legal or other similarly significant effects have to provide certain information about the automated decision making to the data subject, allow the data subject to opt out of that automated decision making, and conducted data protection assessment that should identify and mitigate privacy risks that are associated with automated decision making. It's worth noting that most of these laws have exemptions for employee or candidate data. The exception here being California where as Natalie mentioned, the legislature actually failed to pass a bill that everyone expected would pass that would have extended the exemption. So starting on January 1, the California Privacy Rights Act is going to apply to employee and candidate data. And that includes these requirements around automated decision making. What's interesting is if you look at the definition of profiling, which kind of fits into the automated decision making term, it does suggest that if you're doing automated processing that evaluates performance at work, that's the type of activity that the legislature intended to specifically regulate. And I will hand it over to Anna to discuss federal regulation of AI.

FEDERAL REGULATION OF AI

Anna Westfelt

Thank you Freda. So I will cover what is going on the federal side in terms of regulating AI. And it's important to note that the US does not currently have a comprehensive federal AI law. However, there is a lot happening in the federal regulatory space at this time. And AI is in fact identified as a key technology focus area and the number of congressional bills, and then the government's strategic long term investment. So to start with, in October this year, we saw the release of the AI Bill of Rights or technosys, called a blueprint for an AI Bill of Rights. And it is a non binding completely voluntary blueprint to help guide the design, development and deployment of AI in a way that protects the American public. So the focus here is very much on the protection of individuals in the US. And this bill of rights provides five core principles to achieve these protections. And we see a lot of common themes that we've seen in other regulations here in terms of protecting individuals from unsafe or ineffective systems and providing explanations and notice and privacy and control and human alternatives. So the themes are the themes are familiar the Blueprint is not binding in any way. But we still think is really important because it really will guide future federal law impossibility efficacy approaches, and they were really influenced have government agencies for use AI. In particular, relevant to the focus of today's webinar. The Department of Labor has shown a lot of interest in this AI Bill of Rights and have provided some input on how AI affects workers when they don't have human oversight and don't have control over their own data. Next, it is worth mentioning the algorithmic Accountability Act, which was originally introduced in 2019. With a revised attempt earlier this year, this bill has not yet been passed by Congress, and at this time is uncertain if and when it will pass. However, it is important to consider this bill because, in our view, this sets the baseline for future federal AI legislation. And while it is an ambitious bill, it's also it's also mostly a broad high level proposal for companies to identify and mitigate social, ethical legal risks when deploying automated decision systems. So this bill, unlike a lot of other laws, really focuses on regulating automated decision systems rather than AI systems, as we see in the EU AI, which we'll talk about a little bit later. What I find really interesting with this bill is that it

requires you to compare your automated decision system against the pre existing decision making process. And as we have touched on today, it's very possible that your automated decision system is going to be better than the human run process that you had previously. So this doesn't necessarily mean that kind of all automated decision systems are bad and biased, you could reach conclusion here that you actually have a better outcome with the automated decision systems. This law would require impact assessments before you deploy your automated decision system. And also ongoing assessments, which makes a lot of sense because AI continues to evolve and change. So just having an assessment before you launch, it doesn't really make sense in this space, and this is similar to the conformity assessment and the post market monitoring requirements that we see in the EU AI act. There's this kind of continuous obligation to make sure that you identify and mitigate the risks of your automated system. There, there are also some transparency obligations under the algorithmic Accountability Act that are that are pretty important if company if a company uses an automated decision system to make critical decisions, that expressly includes access to employment. So that is relevant for the context today, you have to make certain information available if you are using AI in the employment context under this law. Again, it is not an enacted law, but it is some these are obligations that we're likely to see if we end up with a federal AI law.

Interestingly, in comparison with EU AI Act, if you are subjected to the algorithmic Accountability Act, you have to do a, an impact assessment if you're using an automated decision system, whereas in the EU, you only have to make the equivalent assessment which is called conformity assessment, if you are using the deployment of high risk AI systems. This algorithmic Accountability Act only applies to large companies. And that definition is interesting. Interesting, because it is it is really very large companies only, you have to either have an annual turnover of over \$6 million, or have over \$250 million in equity value or process information of over 1 million users. At that excludes most small and medium enterprises, which I find really interesting, because those are often the kind of the hot AI companies that we're seeing right now, they would not meet this definition of a large company, they will not be subject to these obligations. I think the reality is this would actually apply to very few of the major AI companies that we see deploying AI that we use in our day to day lives. It also excludes governmental agencies, which is interesting because they are heavy users of AI.

And you would expect to see them regulated in some way, but probably not under this bill. It's also worth mentioning the American data privacy and protection act, the ADPPA. which is an omnibus federal bill with bipartisan support, which is really, really interesting that this happened this year, that it was introduced in June 2022. And however, this bill is completely stalled at this time over battle Tiller state preemption. So we don't expect it to pass in the current format. We don't expect this to happen this time soon. But like the algorithmic Accountability Act, we still consider it setting a federal baseline to just pay for privacy law. But it does have some provisions that are relevant for AI. It includes some national standards and safeguards for algorithms to be using algorithms to process personal information. There are some civil rights protections built in, which is unusual for a privacy law. And there are some requirements on algorithmic design evaluations. And, again, we see the same impact assessments. And this, this bill differs a lot to the Federal Trade Commission. And it does require companies to submit to sign evaluations and impact assessments to the Federal Trade Commission. And that that really leads us to what I think is the most interesting development in the in the federal space when it comes to AI. What the FTC is doing at this time, and FTC really now sees itself as a key

regulator of AI. They released a report in June this year regarding risks of AI. In August, they issued a formal notice of the rulemaking process by starting a public consultation and issuing a formal notice of the rulemaking process, which included a lot of questions that are relevant for AI questions around biometric information, automated decision making systems, concerns around algorithmic errors, algorithmic discrimination, this really shows the FTC has renewed interest in regulating AI and machine learning. And when considering the FTC, we also have to mention some recent settlements that really show the FTC focus on enforcing AI, where we had two recent settlements where the FTC took the view that the data was acquired in an unlawful way. And the sanction was not just deletion of that data, but it was deletion of the models and algorithms that were built on that thing that we've trained that data. So that really signifies the approach that the FCC will have, when it comes to regulating AI. And when it comes to enforcing the laws that they have, they have to enforce today and the laws that they will have to enforce in the future. So some common themes here in on the federal level. As we have seen, these laws really focus on privacy and equity by design at the start of the design process. So you really need to perform assessments before you put your AI product or your automated decision system out in the market. And then you have to continue to assess and evaluate the risks and the impacts on individuals. And we also see common theme for giving individuals agency control over how their data is used. We see things of human oversight of intervention, individuals should have a right to pull human into the loop. And again, it's not necessarily that the AI system is going to be worse than the human. But you need to have that possibility to get a human to intervene. If you think you're being treated unfairly. Or if you think there is some harm coming out of the use of AI systems. And then we see this common theme of impact assessments. And what I think is really interesting is that these are not checked the box exercises, these really have to be meaningful. And often you have to actually make them public or provide them to the regulators. And they have to be continuous. So that's where we will see a lot of help from technology, like we've seen from my friends at holistic AI today. It's not going to be a matter of creating a document that you find a way these really have to be live and ongoing. And and they really have to be thoughtfully performed in order to meet the requirements.

INTERNATIONAL REGULATION OF AI

And then, of course, some transparency requirements that will be really interesting to see in some cases where you have to make information about the use of AI systems public's, and then you have to provide it to the regulators. I think we're just going to learn a lot more about how companies use AI. The next slide, please. So, of course, we have to look at what's going on internationally.

As this will both impact what US legislation ends up looking like and also some of these laws have such a broad extraterritorial effect that they will directly apply to US companies, even if these companies don't have an establishment in the EU. So really, We have to look at the international law for many different reasons. And in the EU a lot of focus right now is on the proposed to use artificial intelligence act. But before we get to that Frida will be covered how the GDPR already regulates.

Frida

Yeah, absolutely. And as Anna mentioned, the GDPR has really set the tone for privacy regulation in the US, which is why we see a lot of US laws that are coming into force next year, essentially mirroring GDPR requirements around profiling and automated decision making. So as we discussed earlier, you

know, in these US states, consumers have a right to opt out, it's the same under the GDPR. So controllers that engage in automated decision making, or profiling that produces legal or similarly significant effects have must allow the data subject to opt out of that decision making. And regulators have said that, for example, e-recruiting practices that don't involve any human intervention are examples of decisions that produce legal or similarly significant effects. The GDPR also requires that data subjects be given meaningful information about the logic involved in the decision making, and the significance of the consequences to the individual of having their information assessed using that kind of automated decision making system. In the EU. We also have the Digital Services Act, which recently entered into force. And that act is not specifically directed at the kind of employment sphere. But it does generally require that certain platforms outline their recommendation, and their algorithmic recommendation systems and their terms of service. So consumers there are entitled to information about how content is recommended to them.

Anna Westfelt

Great, thanks, Rita. So as, I should back one slide, please. We will just briefly mentioned the EU AI act. It has not yet been passed, it is going through the legislative process. And we have seen a lot of amendments to it. But we're definitely seeing some important themes in this act that will be very relevant to you as companies and to get your attention here penalties under this Act are really, really severe. They are much higher than under the GDPR. For certain violations, we could be fined up to 30 million euros or 6% of global annual turnover, whichever is higher. We could also potentially see double deference if you also have a GDPR violation. We don't know yet if that is how this will work in practice, but it is a possibility. So definitely an asset you have to pay attention to if you're in the AI space, it will apply to companies in the US. If the AI system in question was used to process data about EU citizens, you don't even need an establishment in the EU. Similar to the GDPR has a very broad extraterritorial effect. This law is really based on a risk based approach. And the higher the perceived risk restricted the rules, and a category categorized as am systems according to this, when certain AI use cases classified as prohibited, and therefore outright banned, and others tears are high, limited and minimal. And to tie back to our focus today, AI systems used to implement contexts, for example, for recruiting and performance evaluations would most of the time be considered high risk AI systems, it would be separate subject to heavy compliance applications. And while most of the focus in this act is on the providers, so it's really the developer of AI system. If you are an employer using an AI system, you also have to make sure that the system meets all the requirements under the Act, including that it has undergone a conformity assessment and that it has received the required Compliance Certification. So that's even if you just purchase a service from a third party and deploy it in an employment context.

Natalie Pierce

to say thank you so much. And I'm freed up for this I want to just remind people we're gonna have Emre closes out with our final five minutes left in today's webinars, just remind people that much of what we've discussed today in our panel is in client alert, that is available online. And you can always reach out to any of us with any questions. But with that, I'd love to have Emre close this out with our final five minutes. Thank you. I hope my sound is okay. Please do not be if it's not.

WHERE ARE WE HEADED

Emre Kazim

Yeah, so I think it's just maybe just to round up is just to talk about that we're really we're fully seeing the thin edge of a much larger, thin edge of a thin end of a much larger edge. Wedge rather, of legislation and activity taking place regarding trustworthy AI and the regulation of AI. So I think it's very important for us to think about this as more broadly and the GDPR became the global gold standard for privacy and the EU AI Act is very much in the same ilk as the GDPR legislation for privacy and how we expect to have this global outreach but more generally the way that we are reading the New York City legislation and having the conversation with the people drafting, the people who have drafted and pushed for the legislation really at the core of what they are interested in is accountability and transparency so we know with great interest how close the legislation passed in Colorado regarding if you will the assessments of the algorithms used in insurance was to the New York City Legislation and generally we expect there to be a kind of cascading effect. So at the core there is probably a push for transparency. We see this more generally in customer transparency in data, people making informed decisions. Accountability, which we can start to call more company's and individuals and the important part there is ownership of the risks, and evidencing and documenting how decisions were made and to be able to defend them. And then consent. Do we have the right to opt-out, in a much broader sense. We saw this in the first wave in the terms of the cookie consent in GDR where something pops up on your computer website and you click on it, you are probably going to see similar things emerge. You will be made aware that you are working and interacting with an algorithm. Whether data is being processed in this way more generally so we expect the culprit the problem of consent emerge more and more . And human oversight, as mentioned it may not even be in terms of the algorithm system itself but it may be more generally over who owns the risk and what decisions have been made. AI employment we see that the use targets in recruitment side but actually we see the use in algorithms across the human resources when someone joins the company their performance assessment and the calculation of their wage or bonuses, the recommendations for trainings, for potentially alternative roles, etc. algorithms are being used across a company are done with algorithms, so its being used across the business and we can anticipate that area just to continue. What we are seeing now is just the beginning of it. And as mentioned, I think it was briefly mentioned, but really we are seeing this as something a lot more than just filling in a form and then filing it away. But actually thinking about the capacities, the technologies that can help with the governance in transforming the compliance function within a company thinking more broadly in monitoring those systems.

I want to say thank you also for arranging this and im very happy to be here

Natalie Pierce

Thanks for joining us.