Generative AI: Navigating Privacy and Security Concerns in the U.S., EU and UK Webinar

Wed, Apr 26, 2023 3:35PM • 1:01:56

SUMMARY KEYWORDS
ai, generative, gpt, model, data, service, risks, employees, create, tools, laws, prompts, ftc, chat, privacy laws, aia, companies, privacy, relation, llm

SPEAKERS
Anna Westfelt, John Buyers, Frida Alim

Anna Westfelt 02:14
Hey, it's 9am so let's get started Hi everyone and welcome to a webinar on generative AI. My name is Anna Westfelt. I'm a partner at San Francisco and the head of Gunderson Dettmer's risk data privacy group. And I'm joined today by Frida Alim and associated data privacy group and John Buyers, a partner at Osborne Clark in London and the head of his firm's AI and machine learning team. We are going from Dettmer operate primarily in the venture backed company space, we represent 1000s of fast growing venture backed companies, as well as the venture capital firms that invest in them. And our data privacy group counsels clients on privacy, compliance risk and strategy. Anything from day to day privacy matters all the way through to mergers and IPOs. And John, do you want to give us an introduction to your firm?

John Buyers 03:18
Thank you very much, Anna. Yes, so I'm a partner at Osborne Clarke, which is an international law firm headquartered in London. And as Anna says, I lead the AI machine learning team, we're very much a full service law firm but our key markets are digitalization and technology. And We counsel clients on a wide issue a wide variety of issues relative to client's digital transformation journeys, including AI and privacy and data protection.

Anna Westfelt 03:52
Today's webinar is part of our generative AI series. And as I'm sure you all notice that the generative AI world continues to move at a very fast pace. Please check out our previous webinars a little keep in mind that there have been developments since we did those even though some of them are just a few weeks old. There are some links here, and you will get a copy of the slides with the links together with the recording after this webinar. Keep an eye out for future presentations in the series. We currently have plans for sessions covering generative AI issues and open source and in m&a and investment due diligence. Like all webinars in this series, this webinar will be recorded. So you will get the recording and you will get the slides in about a week. Feel free to submit questions today using the q&a

function and we will do our best to get to those during the session. And for those of you looking for CLE we will be providing two CLE codes during the presentation and there will be a silly form sent to attendees. Next slide please.

So before we dive in just a quick we'll recap to remind everyone of what generative AI is. It's the wonderful technology that gave us this image of the pope in a jacket. This was made using mid journey. But of course, not only that it is a type of AI that uses machine learning algorithms to create new and original content, often using vast amounts of data. We see generative AI products and services in text, where you can use it to write articles, scripts and poems.

For example, a very commonly used tool is chatty PT, which we'll be talking a lot about today. But we also see generative AI in images where you can create completely new images based on text prompts using for example, Valley or mid journey or another image generator, and sound where you can generate new music, sound effects, or voices. And that, of course, is something that has immense potential to be used for both good and bad, including deep fakes we'll be talking about a little bit today. If you're looking for a more detailed discussion of what generative AI is, and how large language models work, please refer to our earlier generative AI webinar linked on the previous slide where we go into more detail the technology right, next slide please. So today, we will assume that you are somewhat familiar with the most popular generative AI tools so that we can dive right in on our specific topics. I will start with going through with the current and anticipated US legal landscape looks like and John will cover the UK and EU legal landscape, including the upcoming EU AI act, Frida will run through the privacy risks, I will talk about cybersecurity and confidentiality as it relates to generative AI. And we will wrap up with a panel discussion on practical steps that you can take to prepare your company and mitigate your risk, including what you should think about when developing your own internal AI policy and when you're engaging vendors in the AI space. So we're really aiming to keep this presentation as practical as possible today. Next slide, please.

So on the US side, is all this regulated at this point? Well, in the US, justice, we don't have a comprehensive federal privacy law yet. We also don't have a comprehensive law governing the developmental use of AI. However, we do have a lot of regulation on the state and local level. We now have many state privacy laws regulating automated decision making. And that is very relevant to AI. California just introduced an AI law. Now we're keeping a close eye on because it will have a private right of action. It is fairly likely to pass there will probably be some amendments to it. But it is something that is worth keeping an eye on. And of course, we have laws regulating employees of AI. On the local and state levels in New York City, we have Illinois, we have Maryland with some fairly onerous laws. On the federal level, there have been a lot of attempts, passing voter privacy law that wouldn't have an impact on how you use AI. That's the American data privacy and protection act or Adva. It has not yet been passed, but it is also not dead. So it could we could see this Act passed this year.

There is a lot of renewed interest in passing this now with developments in AI, there have been several attempts to pass an algorithmic Accountability Act. This Act is also not dead, so we're keeping an eye on it. The Biden administration also released a blueprint for an AI Bill of Rights, really giving an indication of how they think about how AI should be regulated. But this is really just a guideline that is not a binding law. So most of most of the action is really on the FTC level. The FTC is looking at making rules around AI, they are consulting on AI and rules. And they are very interested in enforcing in the AI space, they have released some statements on the risks of generative AI, they're making it very clear that in their view, you may be violating the FTC act if you are engaging in deceptive and unfair conduct relating to AI, and that really can be if you make or sell us a tool that is effectively designed to deceive even if that's not the intended or the sole purpose.

They have already issued a few enforcement's that required the deletion of an algorithmic model because they was created with unlawfully obtained data. That is really something to keep an eye on. Because of course that can be that can have enormous consequences not only of the developer of that model, but also the companies using it.

So really the takeaway here is that even though there isn't a federal AI law, there are many ways that existing legal frameworks can be applied in the AI world. And there's a lot of interest and motivation to enforce. We have the most aggressive Federal Trade Commission that we've seen in several decades. And we expect to see more enforcement this year, and perhaps even rulemaking from the FTC. So you really can't sit back and wait for a federal law to get passed. I will hand it over to John, who will talk us through what's going on the EU and the UK side.

John Buyers 10:31
Thank you very much, Anna. So if we could just leave it to the European and UK slide. I'm going to keep these comments understandably brief and high level and they're going to be we're going to be touching on them in a little bit more detail in FreeNAS content after this slide. But essentially, the position in the EU and the UK is slightly more foreign than the US. Although paradoxically, as Anna says the FTC has been making some great, very aggressive strides in terms of AI enforcement. So what exists at the moment?

Well, really, we've got the general data protection regulation, which is a uniform mechanism for the handling of personal data or personally identifiable information across the UK and the EU. And the position is the same in those two territories. And that has been forming the basis of some enforcement action in the EU, which we'll pick up at the end of this discussion. But basically, when you process personal data, in the EU and UK, you have to have a lawful basis on under which to do it. And typically, that is via direct consent of the individuals involved the data subjects or through a concept called legitimate interests. And as in the US, we also have a concept of automated decision making under Article 22 of the GDPR, which is where a machine is making a decision which has an impact, a lawful impact or similarly significant impact on an individual.

So essentially called a major impact on an individual in that situation, you need the direct consent of the individual in order to lawfully process that personal data. Clearly, so far as MLMs are concerned, they're no different. If they use if you're using personal information, personal data, you are still going to have to follow the principles of the GDPR. There are no exceptions there. And again, I'm glossing over really but the these rights tend to flow from something called a data protection impact assessment. That's the first step that you really do need to undertake and ensuring your GDPR compliance, which is undertaking a risk benefit analysis and analysis of your potential use case your data processing use case to ensure that the benefits to the data subjects outweighs their detriments. Why am I talking about the GDPR to a principally American audience? Well, we do need to understand that these measures have an the AIA will as well have extraterritorial impact. These will potentially touch American companies that are serving UK and European customers.

So far as data is concerned, and there are very significant penalties for failure to comply with their provisions, significant percentages of worldwide turnover as fines if you do get compliance issues wrong. So very quick skip through the GDPR. And say we'll go down to the enforcement points in a moment. But following on from the GDPR is the as I mentioned, the EU artificial intelligence that now this only applies to the European Union, because obviously the UK has split away from the EU following Brexit. What it does do is create a complementary measure. It's a measure which is very much something that is being created hand in hand with the GDPR. It's a top down unitary framework, which regulates what is called high risk AI. And these are AI use cases that fall into specific categories and therefore are subject to regulation and this could be and these have been identified as European by European regulators as being particularly important.

These are typically things like financial systems that provide credit insurance, underwriting systems, systems that get involved in employment selection and recruitment. For educational vocational systems to name a few. And there are some quite significant compliance steps that you need to undertake in relation to the AIA including making certain that you follow some quite prescriptive data governance steps in relation to the data that you use to power your AI system. And also in introducing specific degrees of transparency, and what's called logging by design to ensure that people can understand when a decision has been made by an AI system and what impact that has which are not of themselves. uncontroversial.

These are quite difficult obligations to achieve in the context of machine learning. So again, this is something that as US audience, you need to be mindful of, because when this comes into effect, probably in early 2024, this will also have extraterritorial effect. And if you get the compliance steps wrong, and you are providing high risk AI to the European market, then again, you will be potentially subject to some quite significant fines. And actually, if you're using an AI use case, which is deploying personal data, then there's a potential for a double whammy, you could be hit potentially for breach of the GDPR and the AIA, and those fines could be compounded. So they are incredibly significant. Not unsurprisingly, this progress in the AI has actually stalled because of our friends who are producing generative AI and Santa indicated this is a market that is moving incredibly quickly. And even though the AIA has a very nice logical structure, it specifically certainly in its earlier iterations has overlooked the concept of foundational AI, or AI MLMs, which are general purpose models, Mr, causing some concern in the European Parliament that has caused the progress to stall. So I would imagine that there's going to be some revision to that measure, which will enable it to sweep up specific concerns that have been raised in relation to MLMs and other foundational AI models. Very quickly, I'm conscious of time.

There is a companion measure to the AIA, which is the AI best directive in Europe, which will follow behind the artificial intelligence act. And what that does is it gives a private right of action to individuals who have been harmed under the AIA. It's a directive rather than a regulation, which means that each member state will have to implement their own laws to actually bring it into effect. So I wouldn't see any substantive action under this measure until at least the finalization of the core artificial intelligence act. So far as the UK is concerned, we're going in a completely different direction. So we have what's known as the AI governance framework. And the UK has chosen to take a much more pragmatic sector driven approach is not top down and unitary. And essentially, what that means is that regulators are going to have to take the lead in creating suitable laws based on what are defined by the UK government as the five pillars of responsible AI, which include safety and reliability, transparency, fairness, accountability, and contestability and redress. I wouldn't expect any certainly so far as the leading sectors are concerned in the UK with respect to any substantive process progress, but these are undertaken consultations at the moment until mid 2024 at the earliest.

So finally, I was just going to pick up on what has been happening in relation to large language models and principally in relation to open AI as Chat GPT. Many of you will have heard that the Italian regulator garantie has basically asked open AI to provide it with clarifications in relation to the way in which open AI operates, and has given it until I think 30th of May, to provide those clarifications and these really focusing on things like the legal basis upon which Chat GPT has been processing users data, principally the data that's been scraped in the algorithmic training phase of GPT. And their subsidiary issues in relation to transparency they'd like more transparency on the way in which that model works, and protections around information relating to miners.

So we wait to see how open AI will respond to that. But pending that Italian regulatory action, other European investigators or its regulators in Spain, France, Germany and Ireland have initiated initial inquiries and actually the edpb the European Data Protection Board is also set up a task force to look specifically at large language models. So there is some regulatory activity going on in continental Europe. Now. The UK regulator has, I guess, somewhat more pragmatically. Through Ico, The Information Commissioner's Office issued some guidelines on what you ought to be considering in relation to your use of large language models. So that was a very quick canter through the UK and EU position. And as I said, we'll probably pick up some more of those issues substantively when Frida is talking in the next section.

Frida Alim 21:03
Great, thanks, John. I'm going to be going through the privacy risks and concerns associated with using generative AI internally and incorporating it into your product. So start with we can discuss kind of internal use of generative AI. There have obviously been really rapid advancements in the AI space, and it's increasingly apparent that generative AI will be transformative across a range of industries and tasks. And of course, as with every new technology, it introduces some new risks. And with generative AI in particular, there's a bit of a clash between the way generative AI models operate and requirements under certain privacy laws.

It's also important to note that the types of risks that are triggered here also depend on how you're using the generative AI service and how the service itself operates. So at the outset, it is really important to understand how the model is trained, where the training data comes from, how inputs or prompts to the service will be used, and what type of service you're using. For example, are you relying on the direct consumer version of the service? Or are you using the enterprise version of the service because there can be differences in how the data is handled depending on which service you're using, as we'll discuss, it's also obviously really important to think through your policies on internal use of generative AI. You know, what types of use cases are permitted, and which types of use cases are prohibited, which we'll talk about later in the presentation. And with that, we can dive into some examples of privacy risks associated with generative AI services, and how to mitigate those risks. So if you are providing the generative AI service with personal data to train the model, you'll need to make sure that you have the right to use that personal data to train the model.

As John mentioned, under the GDPR, for example, you need to establish a legal basis for processing personal data. So if you're training the model using special categories of data, which are sensitive categories of data under the GDPR, like race, ethnicity, health data, etc, you may need consent of the data subject, which is one type of legal basis under the GDPR. Similarly, new state privacy laws have opt in or opt out requirements in relation to processing sensitive data for certain use cases. Importantly, if you fail to obtain any legally required consent, or if you're using personal data to train the model in a manner that's inconsistent with representations you've made to consumers, the FTC could consider this in the US to be an unfair deceptive act or practice, you know, in the EU, that could be a violation of the GDPR. And in the past, the FTC, as Anna mentioned, has required companies to discourage their algorithms, meaning destroy algorithms that are trained using ill gotten data so something to keep in mind. Note that if you're training based on public data, information relating to identified or identifiable individuals will still be considered personal data, even if it's published, publicly posted, and you'll need to make sure that you have a legal basis for using that data to train the model.

Another thing to notice several privacy laws have purpose and use limitations associated with processing of data, meaning companies should only be collecting and processing data that's adequate to fulfill a specific stated purpose, and that data should be limited to what's necessary to provide the service. Another thing to keep in mind here is that if you're using a generative AI service to create inferences about an individual, the inferences that are created by that generative AI service could themselves be considered personal data under certain privacy laws, meaning they're subject to kind of the rights that would attach to personal data under those and to data subjects to under those privacy laws. Another important privacy risks that's worth noting is it's possible that data contained within training data could be included in the output of the model in response to prompts for you from users.

And in addition to training on an initial corpus of training data. Generative AI models can also be based on prompts and responses that are provided to users of the model. So this is important because if you're training a model on data that contains personal data, there is a risk that the personal data contained within the training data could itself be exposed to users of the service. Similarly, if you're using a pre to train generative AI model, you should understand whether data that's contained within prompts will be used for further training of the model, in which case, there's a risk that you know, any personal data included within a prompt could be exposed to other users of the service.

Another thing to note is that several privacy laws impose additional requirements for a company sells personal information. And the term sell is very broadly defined. So it's any disclosure of personal data to a third party for monetary or other valuable consideration. So if you are covered by the state privacy laws, you'll want to understand whether your disclosure of personal data, you know in prompts or in training data, data to a generative AI surface could be considered a sale or whether you can avoid this issue by putting in place the appropriate contractual language with that service provider which can help avoid a sale. Another thing to keep in mind is you may be required to flow down data subject requests if you are providing personal data to a generative AI service. So worth confirming whether that's feasible. As John mentioned, use of generative AI may trigger laws around automated decision making like you know, laws concerning profiling and automation and employment. Of course, it's very much depends on the use case.

And if you're using the generative AI service for use cases that trigger these laws. So if you are planning to use the generative AI service in any way to assist with hiring and promotion decision making, note that there are some specialized laws here that that could be triggered. And obviously, you should be considering whether or generative AI is, you know, the right tool in the first instance to be engaging in these types of activities. So how can you avoid some of the risks above, of course, the risks that we just discussed regarding leakage of personal data will be greatly mitigated if you avoid providing personal data and prompts and training data to the generative AI service. So it is really worth thinking about how you tend to use the service, and whether provision of personal data to the services can be avoided entirely. If you do need to provide personal data and prompts you may be able to select vendors that don't train on prompts, or that allow you to opt out from having prompts used for training. So for example, open AI recently announced they will not use data submitted by customers to their API to train or improve the model unless the user specifically opts into data, sharing data for this purpose. And they've also introduced the ability to turn off chat history and chat GPT.

So any conversations in the future that are started when chat history is disabled? Would it be used to train or improve the chat GPT models. So again, make sure that you understand you know, how data you provide to the services will be used under the terms of your agreement with that service. So with that, I'm going to chat about deploying generative AI and your services. You know, this carries a lot of the same risks that we discussed in the previous slide regarding internal use of generative AI. So again, if you're responsible for training a model and are using data that may include personal data to train the model, again, you should make sure that you have an appropriate legal basis to use that personal data for that purpose within the training set, or that you have the appropriate contractual rights to that data if you're using your customers data.

So for example, if you've incorporated a generative AI service into your platform, and you want to use prompts provided by your customers to train that model, you know, make sure that you have any rights required to train the model using that data and agreements with your customers or end users. There can also be issues associated with the accuracy of information that's included in responses provided by the generative AI model. You know, large language models may produce or quote hallucinate outputs that appear to be plausible, but are in fact incorrect. And there's some really interesting questions here around whether inaccurate outputs depending on how they're used, could violate the GDPR principles around accuracy of data and fairness of processing. So you'll need to think about the disclaimers that you provide and how the technology could or should be used. You may also want to consider parameters around how users can use the responses provided by the AI model, and whether any limitations should be implemented as responses that the CIC around responses that the generative AI model would provide freedom

John Buyers 29:23
If I could just add one point at least nations very briefly. And that's worth mentioning, because it's very much a feature of machine learning systems. And that is that what you'll tend to find with any machine learning model and not just an LLM, or chat GPT, but any other conversational foundational AI model is that typically, it's very difficult to get consistency and repeatability of contact of content, sorry. So you could ask a question multiple times in the context of particular use case and you may get an almost set you are going to get different answers to the same question that each time you asked the same question. And obviously, that is something we'll pick up in our open q&a At the end of this session. But it's great for a use case, which is based on creativity and creative writing, but could potentially be a problem if you're working on a use case where a definitive answer is required, such as a regulated professional services environment. And that's kind of an extension of the hallucination point.

Frida Alim 30:34
Yeah, that's a great point. Thank you, John. Another issue I wanted to raise is kind of disclosures or and this kind of relates to the issue John was discussing, discussing, you know, disclosures or explanations regarding how the model is used and justification as to why reliance on generative AI is justified for this specific use case. So again, you know, think about whether the output of generative AI is really appropriate for you know, the use case that you'd like to employ it for.

So for example, if you're using a generative AI powered chat bot in your service, you'll want to think about how that chat bots being used and whether you need to disclose the fact that an individual is interacting with the chat bot. And the FTC has warned that you know, misleading consumers via doppelgangers, including chat bots could result in and has in fact in the past resulted in FTC enforcement action. The FTC is especially focused on AI tools at the moment, particularly those that can cause harm to children, teens and other at risk populations.

So if you are deploying a chat bot with those types of users, you should really carefully scrutinize the types of data that you collect via the Chat bot and outputs being produced by the Chat bot. And depending on the laws you're subject to and the jurisdictions in which you operate, you may be required to conduct a data protection impact assessment, or a similar assessment of the risks and safeguards that are associated with that service. And with that, I will hand it off to Anna to discuss cybersecurity and confidentiality issues.

Anna Westfelt 31:59
Right Thank you Frida. So, of course, privacy security, it was go together and we need to look at the implications for cybersecurity and for confidentiality when it comes to generative AI. And one of the main issues and this has actually been in the press a lot recently, associated data exposure loss. So when employees use generative AI tools, they often well, they are prone to inputting your organization's confidential information. We've already seen this with large organizations like Amazon and Samsung, which have found that their code has actually been out in the open because employees have inputted into chat GPT for bug fixes or doing some code analysis.

So this, this is a real problem. It's not just your code. We also see this in terms of personal information, employees may want to put your list of prospects or leads into captivity to learn more about them. That could become public information, or at least available to other chat GPT users if you don't use the right kind of opt in and opt out. Same if you put in information about your end users, there are not only privacy implications there, they're also confidentiality and proprietary information implications because this is typically very safely guarded information. Chat GPT is a great tool for analyzing code. But that also means if you put your code into GPT for analysis, it's possible that you are publishing your roadmap to hack your organization. If you are making clear what vulnerabilities there may be in your code, for example, if you're running a bug analysis, there are also a lot of risk associated with ingesting third party code and third party content from chat GPT. It really it really has been described as a game changer from a hackers perspective, and it really lowers the barrier of entry for hackers. If you ingest third party code, from chat GPT or another generative AI tool, you could be ingesting malicious elements and vulnerabilities. As I'm sure many of you have seen, chat GPT will agree a generating code.

There are some guardrails built into the system to avoid generating code that checks if it deems to be malicious or intended for hacking. But at this time is relatively easy to manipulate those guardrails and get around them. So chatty putty has been used to create some info stealer code and that code has been verified by threat researchers. This is this is again a very real risk and it's already happening. So you really have to make sure that there are several review layers if you're going to ingest third party code from any kind of generative AI tool, and that has to include a layer of human review. This is something you have to look at at on an organizational basis. That's up to you. GPT is also incredibly powerful for social engineering. And it really changes how you have to train your employees to spot these attacks. You can use it to write a really convincing phishing email, you can use it for what we call spear phishing, so very targeted social engineering attacks, because you can learn a lot about an individual using chatty btw.

And typically, the kind of phishing emails that we see produce a chat GPT are a lot more sophisticated than what you would typically see they don't have the typos. So the grammar mistakes that we often use to spot scams. In its simplest form, you can use it to write a really convincing looking password reset requests email from, for example, Microsoft. We have seen examples of Chat GPT being used to create very, very targeted attacks on high level management, trained on all the information that is that is out there in the training data set.

John Buyers 36:23
Yeah, it's like, at the point of this sort of interject because I think this is, this is all incredibly interesting stuff. And as with the large language model, foundational AI models, changing the paradigm for AI, the cybersecurity paradigm has also changed. So, in additional, in addition to social engineering attacks, and third party code, hacks, there's a new, fascinating class of cybersecurity hacking for LLM is called prompt injections. And that's where a bad actor essentially very artfully creates a prompt to a machine learning model to get it to do something which it wouldn't normally do. And these, these, these malicious prompts are actually rising in frequency. And they're very, very clever. And they can actually very, insidiously subvert the, the performance of these models.

Anna Westfelt 37:37

Yeah, that's really interesting. Thank you, John. And that really relates to my next point here is that we are really seeing a growing industry of AI detection software. So as you detect AI scams, you have to use AI. It is worth reviewing what your security tools are, to see what their capabilities are, in terms of detecting AI. This really is a growing industry, where we're seeing a lot of new technology coming out. And I think we'll see more and more sophisticated technology, but it's, it's almost like a race, the scans are getting the scans and attacks are getting much more sophisticated using generative AI.

So on the flip side, the detection software has to get a lot more sophisticated using generative AI. And of course, we see incredibly convincing deep fakes and voice impersonation, we have seen that being used for cyber attacks, where voice impersonation tools are used to train on, for example, high level management, employees public statements, so you just need a little bit of voice to create some really, really convincing AI generated voice prints that can be used for imposter attacks. So on the flip side, generative AI tools can be used to help cybersecurity as I mentioned, AI detection. Software is a really, really big thing right now. You can also use it to write Risk Management Policies, useful malware analysis, but again with a lot of guardrails around it. And it actually produces some really good results. When it comes to using it to write Risk Management Policies.

For example the I wanted to mention the role of the seaso This is something I really highlight to a lot of my clients is that you really need to involve your seaso or similar employee if you have one, when creating your generative AI policies and controls, because this is very much a CISO matter. You have to consider prohibiting certain inputs, such as code and personal information, because of all the rest that we have touched upon today. So that concludes the cybersecurity and confidentiality section. So we wanted to walk you through some practical tips. And then we are going to do a panel discussion and we will try to get to as many of the questions as we can I hear that we already have quite a few in the chat.

Frida Alim 40:11
Great and I will start us off with kind of practical tips and considerations for internal use of generative AI. So of course, at the outset really important to establish a policy around whether and how you know, generative AI can be used internally. So you should have this policy coverage general usage guidelines. So for example, what are permitted use cases of using generative AI, and what are prohibited use cases, it should also cover whether, you know, personal data may be provided to the generative AI service or whether certain types of confidential information should or shouldn't be provided to the service, it should also identify tools that have been approved for use by the company. In other words, you know, the specific generative AI services are tools, if any, that the company has sanctioned for internal use. And when you roll out the policy can also be beneficial to hold the training for employees regarding the policy and use of the generative AI service.

Beyond establishing a written policy, it's also prudent to make sure you're actually monitoring usage of these tools. And you know, one way to do this is monitoring who within the company is using the generative AI services on the network. You could also introduce a splash page with warnings that caution employees about the company's policy on generative AI when they're trying to navigate to certain web pages that have generative AI services. And the other thing to keep in mind when you're using these generative services AI services internally is you'll need to make sure that you are still complying with your obligations under Data privacy laws. So you know, again, just because the service is being used internally doesn't mean that you kind of shirk your data privacy responsibilities. And depending on the specific use case, that could include, you know, conducting an audit of the generative AI service, or conducting a data protection impact assessment. And so I'll hand it off to John to discuss product use of generative AI.

John Buyers 42:08
Thank you very much, Frida. And I'll keep again, I'll keep this relatively brief. But I think the first point that you need to bear in mind when you're considering deploying generative AI on the product side is to really understand the levels of service and the models and the ways in which they are actually being made available to you just get the nomenclature and the classification rights because there's a world of difference between and certainly in terms of open AI is offering the chat GPT offering GPT plus and GPT.

For business, they all offer varying degrees of opt ins and opt outs in relation to training prompts, or using prompts to train models, and use of chat history and models and varying degrees of protection. And they are not the same thing, they might use the same core model, which might be GPT, three or GPT 3.5. But they are different services. And to be fair to and there's a lot of misunderstanding of this to be fair to open AI, it has never positioned chat GPT the base public interface is anything other than a research tool. And you should not be contemplating use of that public base services, anything other than a research tool, you should be thinking about potentially a subscription service, or a service that is designed for business or through an API, which will give you more security.

For those of you who are in business, and are in highly regulated sectors or who are conservative about risk, but we'd still like to leverage the benefits and there are amazing benefits of using our LEMs then you might want to think about taking GPT in a Microsoft as your wrapper. And actually what the offering is here is a is a basically your own instance of the GPT model in an as your VPC which is essentially your virtual private environment where your data can be hosted. And where you have the benefit of an enterprise agreement with as your which will govern as per Microsoft as your SAS arrangement. What the obligations are between the parties in relation to that the issues that we've been discussing, such as privacy confidentiality, performance standards. So the as you're offering the enterprises you're offering may actually cost more money but it does provide you with the same GPT model, but with significantly more guardrails around it in a in an enterprise SAS environment to please my I would stress to you please understand what the levels of service are and what you're what you're getting for what you're paying, if that makes sense.

And the DPI a point really goes to what the way in which European regulation is, is heading both under the GDPR. And under the AIA, which is you really do need to think about responsible use of artificial intelligence, you really do need to make certain that you've evaluated very carefully that your use case is suitable for generative AI. And actually, the benefits to consumers outweighs the detriments of using generative AI. If that doesn't work, then you shouldn't be thinking about generative AI, you should be thinking about more traditional AI services, which are limited models that are single use models that will, for example, simply do facial recognition or vocal recognition or texture recognition or semantic analysis. In the traditional sense, you should not be looking at foundational AI. So that was that was that was my points on product use. I hand over I think it's to Anna. Yep.

Anna Westfelt 46:17
Thank you, John. So on the vendor management side, this really will have an impact on how you sign up vendors and your ongoing relationships with vendors. Even if you as an organization, decide that you are not going to allow any generative AI, it is just not for you, the risks are too great, and you should stay away from it. Your vendors are probably using some form of generative AI, or they are planning to they will change the service. We're seeing a lot of vendors kind of getting in on this and creating chat bots or other analytics tools using generative AI and their services. So you really have to understand how you let vendors are using generative AI and make sure that you understand what kind of data is going into that and how that data is being used. This also means that you may need to revisit your contracts with those vendors.

From a GDPR perspective, it could be that you previously had the event vendors position that's pure processors, but if they start using the information in a different way, for their own purposes, for training their model for other customers, they may be controllers, and you have to decide if that's something you're comfortable with and what kind of restrictions you have on them using the data that you provide to them. So this is really going to change your vendor relationships, and it's something you have to pay attention to. That also relates to my next point here of your outbound contracts. So you have your public facing terms and your customer contracts. So you really have to make sure that your terms of service, if that's how you sign up your customers, that they are consistent with you use of generative AI tools. And same with the privacy policy that has to have accurate disclosure. So if you're using generative AI, that is probably something that you're going to need to update because even just a year ago, when we were updating for other loss, we didn't really have the generative AI tools available on the same basis that we do today.

Similarly, you have to keep an eye on your customer contracts when you're starting to see anti AI provisions. So it could be that your customers, especially if they are in a highly regulated industry, like financial services or health care, that they have provisions, basically saying that you are not allowed to use any AI tools. And you have to represent that you are not currently using any. That's where you have to make sure there isn't a disconnect between what your vendors do and what you agree in your customer contracts. Because you're going to get caught in the middle there. If you are in fact using AI tools you've been breached on that permission day one. So really pay very close attention to what your customer contracts say about the use of AI tools or whether they prohibit AI tools. So let's kick off the panel discussion. I am seeing some questions around the EUA. I act. I would like to kick that off with some questions for John. There.

We've had some questions about the extraterritorial effect of it and how it can actually apply to US companies and how that maps to the GDPR. And in relation to that, I would love it if you could talk a bit about why US companies should be concerned about it at this point, since it's actually not yet a law, as you mentioned, probably early 2024 is the earliest we'll see of being passed. Are you seeing US companies already taking steps to be compliant? And what do you think companies should be doing at this stage?

John Buyers 49:48
That's a great question. And thanks very much and taking the reverse order first. I mean, I think yes, I'm seeing a high degree of compliance that has been carried out by internationally facing US companies, which is very encouraging that I've really been motivated by the provisions of the AIA. Obviously, the AIA is not a law, as you say it's a draft law. But it's envisaged to have the same extraterritorial impact as the GDPR. So if you're marketing towards European consumers, if your business touches on European consumers or your AI will impact them will affect them, potentially deprive them of rights, then you will be caught by the AIA. So it's very much a consideration, if you consider the European market, which is a very large market to be one of your target markets. Now, it may not be for all US companies.

But certainly the larger US companies will be very interested in that in that. What do you need to be doing to ensure compliance while it? Again, the measure is in a state of somewhat in a state of flux, because of the disruptive impact of generative AI, which is, I think, potentially causing European legislators to revisit this high risk structure, I don't think that's going to be completely thrown away. I think that it's going to stay but it's going to be significantly amended to allow for foundational AI models. And there's been some movement in the latest revisions of the act to to cover that you should, frankly speaking, be undertaking the AI equivalent of a DPA, which is going to be called a fundamental rights impact assessment, which is the same exercise, which is objectively justifying that your use case in the European Union so far as artificial intelligence is concerned outweighs in terms of benefits the detriments to the users.

So you simply can't implement and you won't be able to implement in Europe and nice to have use case for generative AI, you're going to have to provide something which is substantive in terms of your reasoning for using generative AI. And actually the bar for generative AI, I think is going to be higher than some types of more traditional artificial intelligence, which are more defined as a use case, simply because it is multi purpose. And it is drawing the attention of several large regulators, the regulators of several large European countries at the moment, who are very concerned about this wide reach that the technology has. So hopefully, that answers the question, but hopefully sufficient for the questioner.

Anna Westfelt 52:44
Yeah, that's great. Thank you, John. And we did get a question about who owns the content generated by AI that I can address quickly. It's really up to the terms of the tool, the AI tool that you use. So for example, chat GPT states that you own the input and the output, but that is as between you and open AI. So that doesn't mean necessarily that you have intellectual property rights in what is created. It could be infringing someone else's intellectual property. So for a more detailed discussion of the ownership issues, please check out our introductory generative AI webinar, there is a link in the slides that you will receive, we go into much more detail there on kind of ownership issues. But it really is something you have to look at the terms for each tool to see who owns what is between you and the tool vendor who owns the output. So Frida, can you talk a bit about what could happen if a company is using a generative AI service that has been training that has been trained using unlawfully obtained data? So for example, where the generative AI hadn't obtained individual's consent, if the applicable laws actually required consent?

Frida Alim 54:12
Yeah, absolutely from a US perspective, we've actually seen this play out in a few enforcement actions that we mentioned where the company is training an AI model based on ill gotten data, kind of classic example here is, you know, scraping photos from the internet and using biometric data contained in those photos to create, you know, a facial recognition model. And in one of those cases, the FTC required algorithmic disgorgement, meaning that, you know, the algorithm that was trained on, you know, the data that was obtained without consent, which was the standard required under certain laws, or, you know, in contravention of what that company had told to users then had to destroy that algorithm.

So, the risk here is that, you know, if you are creating a model and you're training it based on data, you know, you didn't have a right to or you're training it in a way that, you know, it's not consistent with what you've told individuals, you may be required by the FTC to destroy that algorithm as that may be viewed as an unfair deceptive act or practice. And if you are incorporating, you know, a third party service that trained based on ill gotten data and you know, their algorithm is required to be destroyed, you're kind of put in an awkward position there, if you are reliant on that algorithm to provide your services. So just a bit of risk there. If you're, you know, relying on a generative AI service without understanding how it was trained, or, you know, if you're relying on a service that perhaps didn't have, you know, the appropriate legal basis to train that model.

Anna Westfelt 55:47
Right, thank you Frida. Frida do you have some question?

Frida Alim 55:59
Yeah, I do so and we've heard a lot of risks highlighted during these this presentation. You know, should organizations be completely banning their employees from using generative AI services? Or, you know, are there some use cases here where there's there's lower risk to the company? And what can companies do to mitigate their risks associated with these services?

Anna Westfelt 56:21
Yeah, absolutely. I personally don't believe in a complete ban, because it is pretty easy for employees to just use your own personal devices, even if you block all the generative AI tools that you can think of your employees will probably use them. And sometimes it can be for fairly benign things, performing research, or maybe just using it to draft a sales email that they could use as a starting point. So I really believe in a balanced approach.

And really, the important part here is to communicate clearly with your employees, what your expectations are, with respect to the use of generative AI. And what the guardrails are. And as with all policies is really important to not just kind of push out a policy, have employees sign it when you onboard, posted with the OSHA policies, and then forget about it, you have to make sure that the policy is actually followed, that you monitor use that employees are made aware that you're monitoring their use of generative AI. And you need to do kind of continued education throughout your organization to make sure that this really becomes something that is very familiar to employees how they can and cannot use chatty putty. So really, I think the worst thing you could do is create a really onerous policy, have people sign up to it. And then for the policy away in a drawer, I think this becomes much more of an operational issue where you have to make sure that that generative AI is used responsibly throughout your organization. And some of the other guardrails that were mentioned can be really useful, like having a splash page that really reminds employees about what information they can and cannot input into a generative AI tool.

If you are in a particularly sensitive information is sensitive industry, then you really have to be careful. And you probably have to train your employees even more and put some additional guardrails around it. But there is a lot that you can do to make sure that you can use generative AI responsibly. I also think, just with the incredible popularity of these tools, and there are a lot of really powerful tools available. There are some really positive aspects of generative AI. And I think there's a risk that organizations fall behind if they don't let their employees get the benefit of some of those tools. So again, it really comes back to balance and responsible use.

John Buyers 58:50
I would I would concur with that completely. And I think if any in the audience have not yet used generative AI then they should sit seriously think about trying it out. Because frankly, it is a bewitching experience. I think being subsumed in that richness of interaction with an LLM is an experience unlike anything I've ever had. It's extraordinary and quite powerful. But actually I think it's a very simple message that and I would agree entirely with Anna it's really apply common sense to your usage of MLMs and and realize that they are more suited because there are inherently forms of machine learning. So they suffer from the same vulnerabilities. But they are much better at other use cases than they are others and as a starting point where you're looking to have something create some creative text to get your creative juices flowing. I think that's a fantastic use case there for that technology. However, if you're in a business where you are in such as Anna and I are in the business of providing regulated professional advice to clients where there are legal consequences.

If you don't get the advice, right, then you really shouldn't be using an LLM in an in an unrestricted context to provide advice to your end clients. And that would go for advice industries, such as financial services, and not just in relation to regulation. If, for example, the use case is supported by a contract. So if you're obligated under a contract to provide certain kinds of advice or outputs, then common sense would dictate that you shouldn't be using a large language model to do what a human could otherwise do on a more accurate basis. So I would go back to different horses for different courses essentially.

Anna Westfelt 1:00:52
Yeah, absolutely. Thank you, John, that was really, really helpful. And while we still have you all, I'm going to give you the Cle codes. The first one is 1092. And the next one is 6436. So that's 1092. And the next one is 6436. And a certificate for CLE will be emailed to attendees after this webinar. And I think that means we are just about out of time. So apologies to those whose questions we can get to feel free to follow up with us after the webinar and keep an eye out for invites to future webinars, we will enjoy talking about generative AI it is such an interesting field and things change daily it seems like so

thank you everyone for joining. Thank you Frida and John for joining the panel. And we look forward to speaking with you or more in the future.

John Buyers 1:01:51
Thank you very much