AI

ISSUES IN ARTIFICIAL INTELLIGENCE

Blog Series by David Canton

Harrison Pensa TECHNOLOGY & PRIVACY LAW

You may download, save and/or print a copy of this document for your own use. This publication is for information only. It is not legal advice. Please consult a lawyer to obtain advice for your specific situation. Up to date versions of this publication and other Harrison Pensa publications are available for download at harrisonpensa.com.

Table of Contents

Consent issues with Sora text-to-video AI

21 February, 2024

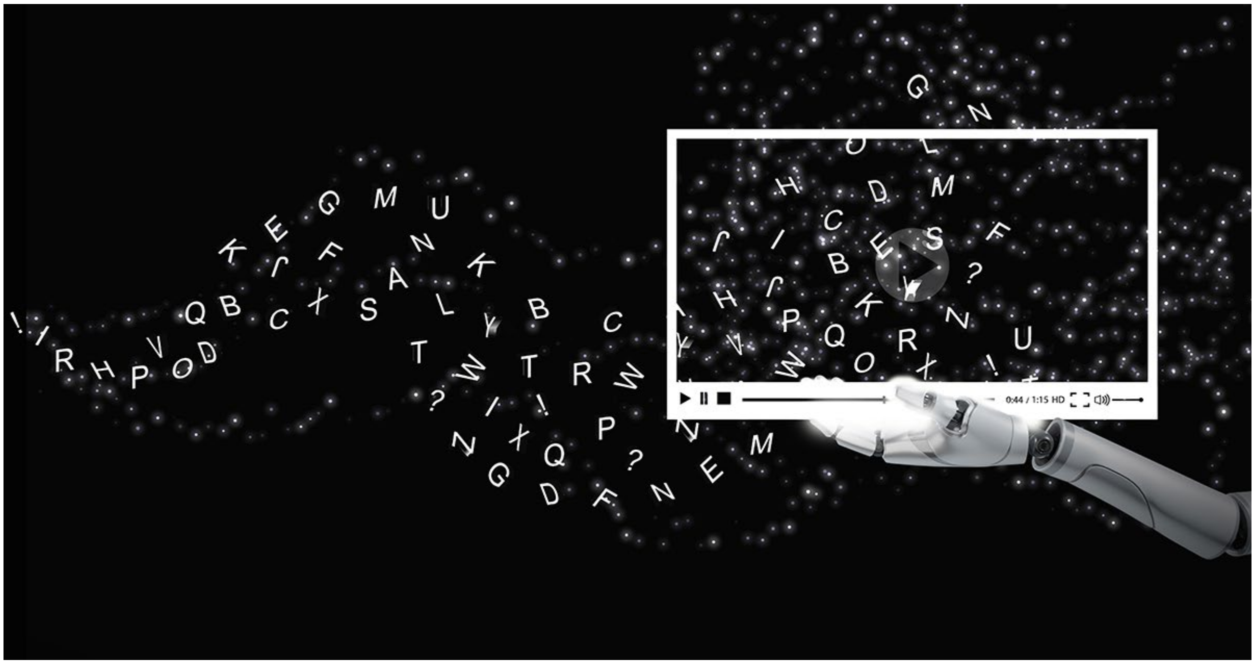

OpenAI just announced their Sora text-to-video diffusion model. It can generate realistic-looking one-minute videos based on simple text input. While it is not yet available for general use, the site has several examples. They are amazingly realistic. OpenAI admits that it can struggle to show a few things accurately, but it is orders of magnitude better than Dall-E was when it debuted just over a year ago.

This is simultaneously fascinating, compelling, disruptive, and terrifying.

Using this technology to create fake news or mislead people is going to be a real problem. Fundamentally this comes down to consent and deception.

AI fakes and consent

Consent in the context of needing permission to use someone’s likeness or creative works.

Deception in the context of misleading viewers that what they are seeing is real.

This technology is not evil per se — there are lots of legitimate uses for it. There are contexts that even fake video about a person can be a legitimate tool. For example, it would be fine for Harrison Pensa to create a video of me walking down the street reciting this blog post.

Provided that Harrison Pensa does that with my knowledge and permission, and it is identified as AI-generated and not really me.

We have written before about reporting during a time of misinformation, and whether there is a need for regulation to control AI replicas.

Nefarious purposes

Never underestimate the inclination of bad actors to use this tech for inappropriate or illegal purposes. That could range from fraud targeted at individuals, to false celebrity endorsements, to politicians trashing their opponents, to nation-states interfering in the politics of other countries.

OpenAI's website talks about safety steps they are working on to try to reduce misuse. That would include rejecting things like "… extreme violence, sexual content, hateful imagery, celebrity likeness, or the IP of others." And including watermarks to identify output as Sora AI-generated.

This type of platform control is important. But that alone won't be enough. Videos that might seem innocuous when generated could be used for misleading purposes. For example, a video of a crowd in a specific setting could be used on social media to claim an event happened that did not.

Social media and publishing platforms will need to be vigilant around this.

Emerging laws around AI may or may not be able to deal with these issues effectively.

We as viewers of this content have a role to play. We need to be skeptical that anything we see is true. But that is easier said than done when it looks realistic. Or when it rings true because of confirmation bias. And who has time to fact-check everything we see?

This leads to me ponder the irony of writing this post if it turns out that Sora is not what it is billed to be.

Image credit: ©OleCNX - stock.adobe.com

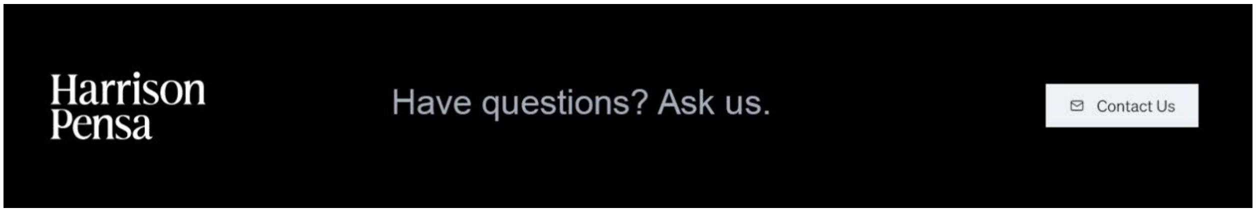

Regulation to control AI replicas

January 24, 2024

AI tools are available that allow the creation of voices and video that sound, look, and act like a real person based on relatively small amounts of actual audio or video recording. We have mentioned before issues around people using AI to make up deep fake stuff for nefarious proposes.

This technology can also be used to create new content that appears to be from live actors. For example, a movie studio could take video that it made of an actor for a movie and use AI to replace that actor in another movie. Without paying that actor for their appearance in the new movie.

Should a film studio (or anyone else for that matter) be able to take audio and images from actors or models and use it to create new audio or video that sounds or looks like the actor or model?

There are of course potential legal issues around that depending on the circumstances. Existing laws prevent others from using our images for commercial gain without our permission. There are also ethical and societal issues around it.

To some extent, it may depend on the language in actor, model, or photo agreements or releases that were signed when the original images, audio, video, or 3D body scans were taken. Those releases often allow very broad use. Some allow whoever took a picture or video of a person to do almost anything they want with it and alter it in any way they want.

This issue will inevitably be tested in court. It will be interesting to see to what extent courts will say it is acceptable because the wording is broad enough to allow it, and to what extent they will pause. After all, no one contemplated this kind of use when the releases were signed.

AI replicas

A major issue in the recent SAG-AFTRA (a labour union representing performers) strike was to give actors control over the ability of movie studios to replicate that actor’s voice and image without the actor’s consent. SAG-AFTRA recently announced an arrangement for union members to license digital replicas of their voices for video games.

A Bloomberg Law article titled “AI Threatens to Push Human Fashion Models Out of the Picture” takes a deep dive into the controversial issue of clothing models and AI modeling companies that do things like swap faces. Fashion companies can use AI models for a fraction of the cost of a human model.

As with many uses of AI, the question is where the boundary is between acceptable AI use around human images and unacceptable AI exploitation without compensation.

Current laws are relevant, but we should also ask what the right thing is to do. And flowing from that, is there a need for AI regulation to deal with this?

Gemerative AI Image: ©ink drop - stock.adobe.com

Use the EU Artificial Intelligence Act to guide your compliance

December 20, 2023

The EU has agreed on the terms of the EU AI Act that will regulate artificial intelligence. It is expected to take effect no sooner than 2025.

Its rules are based on a risk based system, with the strictest rules reserved for the highest risks. It also prohibits the use of AI to do certain things at all.

This Q&A from the European Commission is a worthwhile read.

It wasn't easy to get the act drafted, or without controversy. But anyone creating AI systems would be well advised to review the act now with a view to complying at least in spirit, even though it won't be law for a while. Even if a business does not operate in the EU, it may be a guidepost for eventual regulation and expectations elsewhere.

Here are 10 questions to ponder:

1. How well do the politicians really understand AI and its ramifications and how will this law will work in practice?

2. Does it strike the right balance between allowing innovation and protecting human rights?

3. Will the EU AI Act be looked upon for AI like the GDPR is for privacy — the toughest standard that international companies will adhere to?

4. Given that it will be a couple of years before the AI Act is in force, and how rapidly AI is advancing, how much of it will be outdated, and how many gaps will it have, before it is even in force?

5. Will it be ignored by companies creating AI that decide to just avoid selling their products in the EU?

6. Will Canada and other countries try to emulate the AI Act, or will they go their own way?

7. Should Canada slow down its AIDA draft AI legislation so it can be drafted in line with international standards? That might reduce the chances of becoming out of synch with the world and result in products being unavailable in Canada as a result of compliance costs.

8. The AIDA approach is to have a skeleton act that will be fleshed out by regulations. That approach has been criticized for putting policy decisions in the wrong place. Would a better approach be to pass an initial version of AIDA that covers some high risk situations or no-go zones similar to how the AI Act does, with a view to working on a more well thought out , more comprehensive, more internationally compliant approach on a longer term?

9. Is it a mistake to rush into regulating a complex cutting edge subject like AI, or is it too late if time is taken for politicians to thoroughly understand AI and the many aspects of possible risks and ramifications?

10. What role will voluntary codes of conduct play around AI and AI laws?

Image credit: ©Sergey Nivens – stock.adobe.com

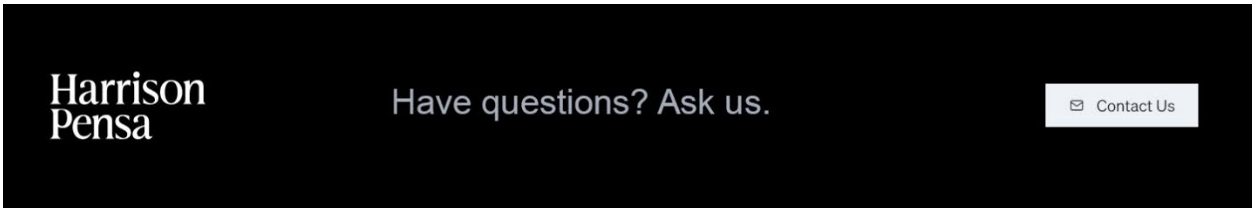

How the world is handling artificial intelligence rules

November 8, 2023

New rules around artificial intelligence seem to be coming at us almost as fast and furious as new AI products.

Many jurisdictions are scrambling to get something in place. While there may be some consensus on high-level AI issues that should attract regulation, there is controversy and uncertainty over the details, what is needed, what kinds of AI require higher scrutiny, and who they should apply to.

Getting AI regulation right is crucial. While there are serious issues to address, getting regulation wrong won't help anyone, can cause a lot of grief, and get in the way of innovation.

EU

The EU was one of the first to consider a law with its Artificial Intelligence Act. But it has been bogged down by disagreements amongst the lawmakers.

Canada

In Canada, we have proposed legislation in the form of the Artificial Intelligence and Data Act. It has little detail and will need either extensive changes or extensive regulations. As an interim measure, the government released a Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems. It is a non-binding voluntary code that was signed by 14 businesses that the government is encouraging everyone to follow.

G7

The G7 countries just agreed on “International Guiding Principles on Artificial Intelligence (AI) and a voluntary Code of Conduct for AI developers under the Hiroshima AI process.” It basically adopts the rules under the EU Artificial Intelligence Act.

Canada and the US are of course part of the G7.

Bletchley Declaration

At the AI Safety Summit in the UK, 28 countries, including Canada and the US, agreed to cooperate to identify AI safety risks and create risk-based AI policies.

U.S.

In the United States, President Biden just issued an Executive Order to address AI. It appears to be a blueprint for AI regulation by various government departments.

Reaction has been mixed.

This Mashable article titled “White House drops an AI regulation bombshell: 10 new mandates that’ll shake up the industryv” calls it “a thunderous executive order.”

Wired, on the other hand, says “Joe Biden’s Big AI Plan Sounds Scary — but Lacks Bite. Joe Biden’s new executive order is billed as the biggest governmental AI plan ever. Unless he can convince a dysfunctional US Congress and overseas rivals to play along, its effects will be limited.”

Practical Effects

So how does all this affect the AI industry and users of AI products? It depends.

The most serious and onerous aspects of these laws and codes apply to those doing cutting-edge and large AI projects. Or for AI that might pose high-risk to things like cybersecurity and critical infrastructure. The EU legislation classifies risk into 4 tiers depending on possible risk to health, safety, or fundamental human rights. The AIDA will apply to “High Impact AI systems”.

However, certain aspects of these laws and codes could affect how businesses and organizations use AI tools.

None of these are actually binding on anyone yet — other than the companies who have voluntarily declared they will follow them. But they are an indication of what laws will eventually require. They may also set the bar for consumer expectations, and perhaps even for governance and negligence tests. Businesses should at least pay attention to the nature of these rules and adopt them as best they can to the extent they might apply to what they are doing with AI.

Image credit: ©Alina – stock.adobe.com

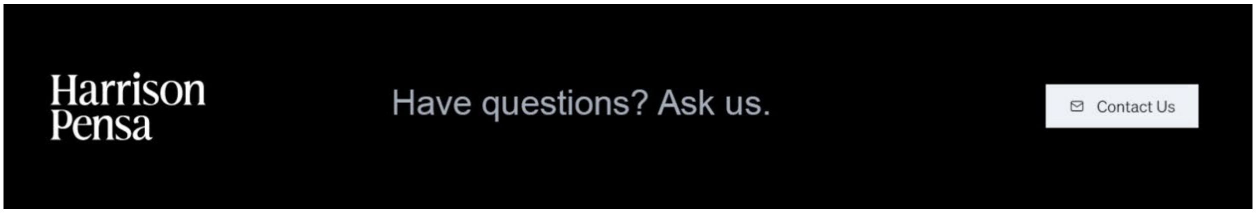

Remove AIDA from Bill C-27

October 4, 2023

Bill C-27, the proposed legislation that will replace the Personal Information Protection and Electronic Documents Act (PIPEDA), Canada's private sector privacy legislation, is slowly making its way through the legislative process. This post sets out where it is at, and what businesses need to know.

Bill C-27 includes:

- The Consumer Privacy Protection Act (CPPA) that will replace PIPEDA;
- The Personal Information and Data Protection Tribunal Act that would create a privacy tribunal as an enforcement mechanism in addition to the Privacy Commissioner; and,
- The Artificial Intelligence and Data Act (AIDA) that will regulate some commercial uses of artificial intelligence.

While the CPPA needs some changes before it is passed, it is the second attempt to draft a PIPEDA replacement and is reasonably close to the final form.

The AIDA on the other hand is a hastily put-together high-level shell without substance or detail. It needs a lot of work.

Bill C-27 recent developments

The Ministry just started hearings into Bill C-27 that got off to a bad start. The hearings were set to hear about 30 witnesses over several sessions. The Minister had promised amendments to the bill and was criticized for hearing from witnesses about the draft bill without knowing what the proposed changes were. Professor Michael Geist wrote, “This secretive, non-transparent approach is unfortunately consistent with the privacy and AI reform process.”

A group of 45 organizations and experts sent an open letter to the Industry, Science, and Economic Development (ISED) Minister. The letter sets out a number of concerns and wants the AIDA separated from the CPPA so the AIDA can be properly considered and drafted. They want ministries other than the ISED to be involved in the AIDA drafting process.

A few months ago, the government published The Artificial Intelligence and Data Act (AIDA) — Companion document. It is a high-level primer on the government’s approach to AI regulation.

The government just released a Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems. It is a non-binding voluntary code that was signed by 14 businesses, including Telus, Open Text, and Blackberry. The government refers to this code as “… a critical bridge between now and when [AIDA] would be coming into force.”

My take on all this

The government should disclose its proposed amendments to the CPPA so the witness comments are more relevant.

The AIDA should be separated from C-27. The CPPA is in a more mature state and should not be held up by the AIDA drafting and comment process. It is crucial to get AIDA right and pushing it through with the CPPA is not conducive to that.

What business should do

On the privacy side, be aware that when the CPPA replaces PIPEDA it will require a review of privacy practices to make sure they comply with the new rules. And that it will require internal documentation and policies in addition to current privacy policies.

If a business is providing AI products, it should become familiar with the AIDA companion document, and the voluntary code of conduct. While neither of those is enforceable law, they give an indication of where legislation is headed. Being able to state that you comply with the code may give some comfort to potential customers.

Image credit: ©Sono Creative – stock.adobe.com

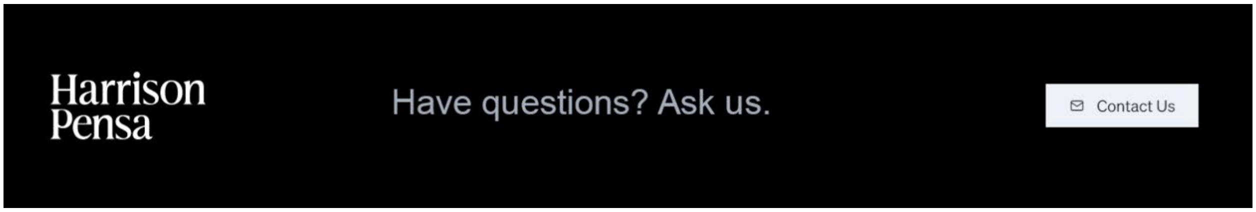

Privacy risks of AI for ID verification

August 23, 2023

The use of facial recognition software for ID verification is controversial for several reasons. A recent report by Calgary based legal advocacy organization Justice Centre for Constitutional Freedoms details privacy risks of using AI based technology for identity verification. It points out that while using facial recognition and other biometrics to authenticate people is tempting, it comes with some serious privacy issues.

There are two fundamentally different ways to use facial recognition or other biometric based tech that knows who we are.

ID verification

The first is to verify that we are who we say we are. In other words, that this really is David trying to enter his locked office door or access his bank account. This kind of use has many issues that need to be considered and dealt with, but they can be addressed.

Identification

The second is to simply identify us. For example, a video feed of a public space that recognizes that David walked by this spot at 5:00 PM on Thursday. That is much more invasive than mere verification. Some countries have embraced this technology under the rationale of law enforcement being able to tell who committed a crime.

England is an example of a country that has embraced massive CCTV surveillance networks. That seems odd for a country that has one of the toughest private sector privacy laws in the world. It’s not the only country that seems to have a significant disconnect by being tough on private sector use of personal information, but insisting that the state needs massive surveillance powers.

Canadian perspective

In Canada, Privacy Commissioners have weighed in on this issue, banning Cambridge Analytica from scraping our images from the internet and selling those to police. They have also published their thoughts on how police use of facial recognition **should** be limited.

The problem

Facial recognition using artificial intelligence to identify people is notoriously inaccurate, especially when dealing with people who are not white. It has a serious embedded bias problem. Many jurisdictions have banned the use of it for that reason. The problem is the undercurrent is that it will be okay to use it once the bias /accuracy issue is sorted out.

The longer lasting problem, though, is not accuracy. Assuming the bias / accuracy problem can be fixed, the bigger problem is privacy.

Think of it this way. Police can only collect and retain our fingerprints if we are charged with a serious crime. So why should they be able to take our images or other biometric information from various sources without our knowledge or consent and use that to track our every move?

Image credit: ©Sergey Nivens – stock.adobe.com

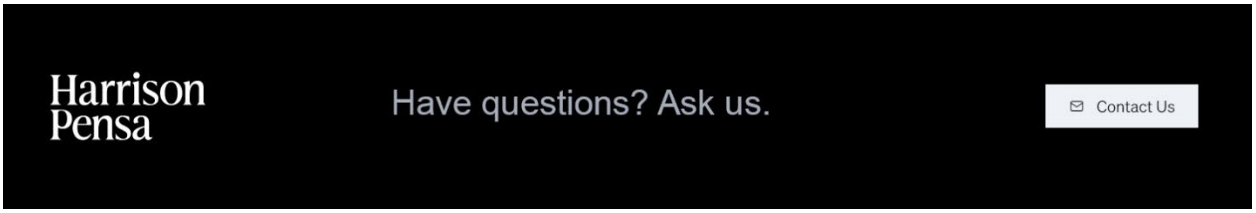

Legal, ethical issues of AI for recruiting

June 14, 2023

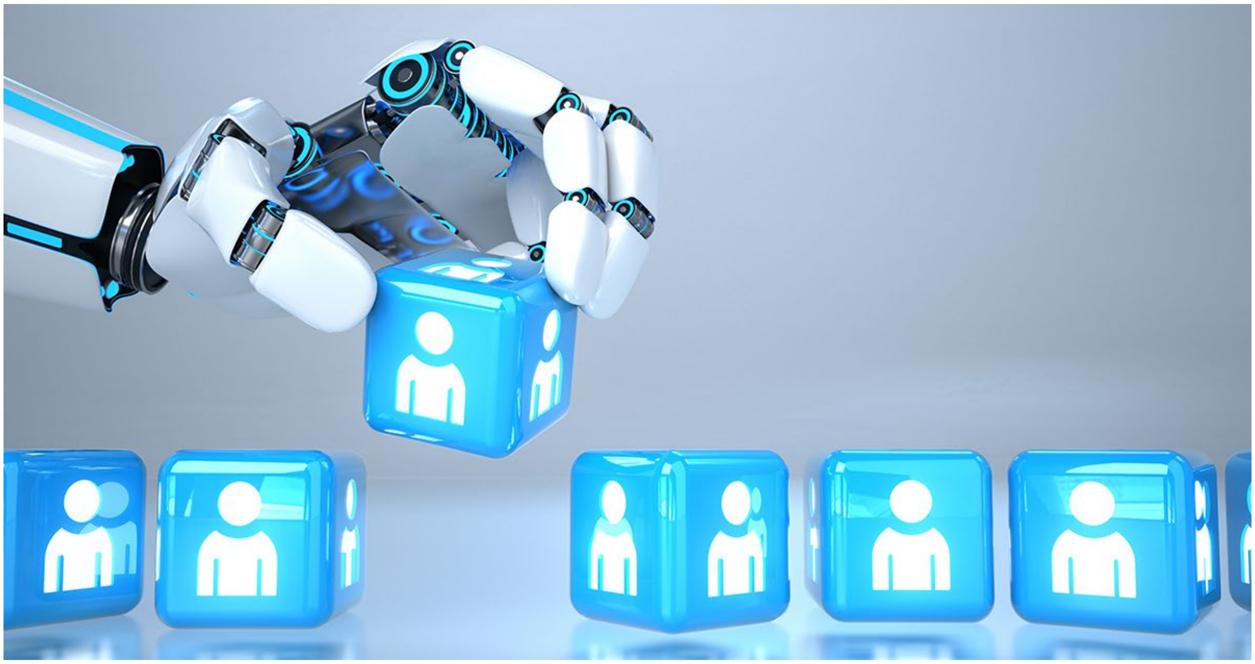

A recent survey of college students found that almost ½ of them are interested in using AI chatbots to write their resume and cover letter.

But almost the same percentage of human resources professionals surveyed said the use of AI on job applications would be a dealbreaker.

That's a bit of a head-scratcher. Generative AI tools like ChatGPT are notorious for making things up and providing output that is not true — often referred to as hallucinogenic output. Try it for yourself. Ask ChatGPT or another tool to write your CV and see what it serves up. When I did that on myself it sounded good — but got about one-third of it wrong.

Job Seeker AI tools

There are many AI-based tools that job seekers can use. Surely no one would use AI to help write their resume and include false information in it. (Or at least no more than would embellish their resumes on their own.)

So it's a mystery why so many HR professionals are so against it.

If job seekers use AI to research potential employers and get the details wrong, they do that at their own peril. Job seekers sometimes use AI tools to try to get by the automated tools employers use to compare resumes and job descriptions to reject candidates.

HR AI tools

The bigger issue is HR professionals using AI tools to research and select people they want to employ. It would have been interesting if the survey had asked those HR professionals if they would use AI tools in the hiring process.

Using AI tools to learn about candidates can be a problem. AI tools can get similar information as a Google search, but with a significant risk that it will serve up believable nonsense.

The use of Google and social media to look at job candidates is somewhat controversial. It can serve up information that one is not supposed to consider from a legal and human rights perspective. You can find the answers to questions you are not supposed to ask.

Doing AI chat searches on candidates may make that worse. Not only are you looking at information you are not supposed to consider, but that information may be wrong and biased.

Using AI tools to decide between candidates is also a problem. There are issues around embedded bias and algorithmic transparency when using AI that could lead to legal, human rights, and public relations consequences.

For example, a recent class action suit was commenced in California over alleged AI tool hiring discrimination. It is inevitable that AI tools will be used by both candidates and employers in the hiring process. The challenge is to use them in a fair and transparent way that is consistent with employment law principles.

Image credit: ©Alexander Limbach – stock.adobe.com

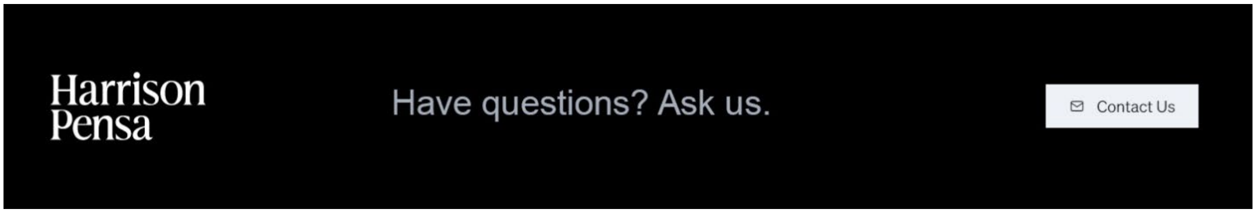

Impact of generative AI on fair use in creative works

June 1, 2023

Who owns creative works generated by artificial intelligence, and why does it matter?

AI can generate art, music, text, and software. These are often referred to as “creative works”. AI doesn’t create, it just generates its output based on a mashup of existing material. Knowing who owns that output can have significant ramifications. That controls who can exploit it for profit, who can stop others from copying and using it, and perhaps who is liable if there is something unlawful or actionable in it.

Human creators only

We start with the notion that ownership of intellectual property such as copyright and patents is based on it being created by a human. Courts and IP offices in various countries have routinely rejected attempts to claim AI or animals as authors.

For example, a U.S. court said a monkey that took pictures of itself after a photographer set up a camera could not be the author or owner of the pictures.

The U.S. Copyright Office recently published a guide to the registration of works containing generative AI material. It confirms that AI-generated material is not subject to copyright, and sets out how to handle applications that have an AI component.

Creativity and effort threshold

At the same time, it is safe to say that if a human creates something using AI as a tool, and puts sufficient creativity and effort into it, that human would own the output. To some extent that is no different than creating an image using Photoshop.

The big questions are how much and what type of human input is required before the human owns it, how will the courts or legislators define the line, and will various countries be consistent?

There is case law in various countries on the subject of copyright ownership and infringement and what the thresholds are. That case law would form the basis for any decisions. Basically, courts require more than mere labour – it requires some level of skill, judgment, and creativity. (If you want to read more about copyright decisions on this issue, look at this Wikipedia article.)

The unknown is how these existing tests will be interpreted for works that are created using AI.

Advice to creators

If anyone wants to claim ownership of something they create using AI, they will have to prove that they put sufficient skill, judgment, and creativity into it. The challenge is that we don't yet know how much is sufficient. Creators would be well advised to keep a comprehensive record of their prompts, process, and as many versions of the output as possible from the initial prompt to the final product.

It may take some time to sort this out. And this ownership discussion ignores issues that might arise out of litigation on copyright issues AI tools are currently facing regarding the AI tool's use of existing material to generate its output. Various jurisdictions are working on legislation that will affect AI. That legislation may have something to say about this as well.

Image credit: ©EricSchumid - stock.adobe.c

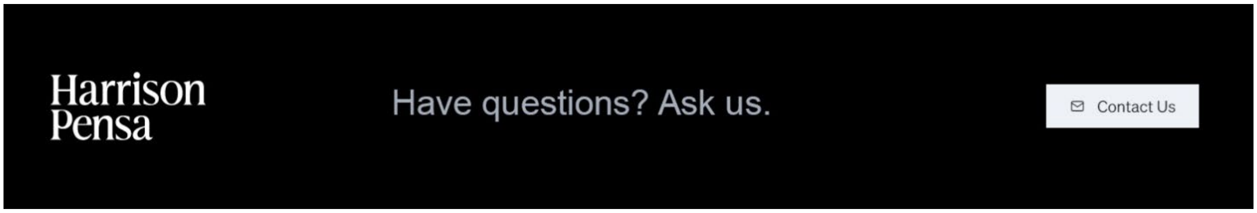

Artificial Intelligence caused these real-life issues

May 16, 2023

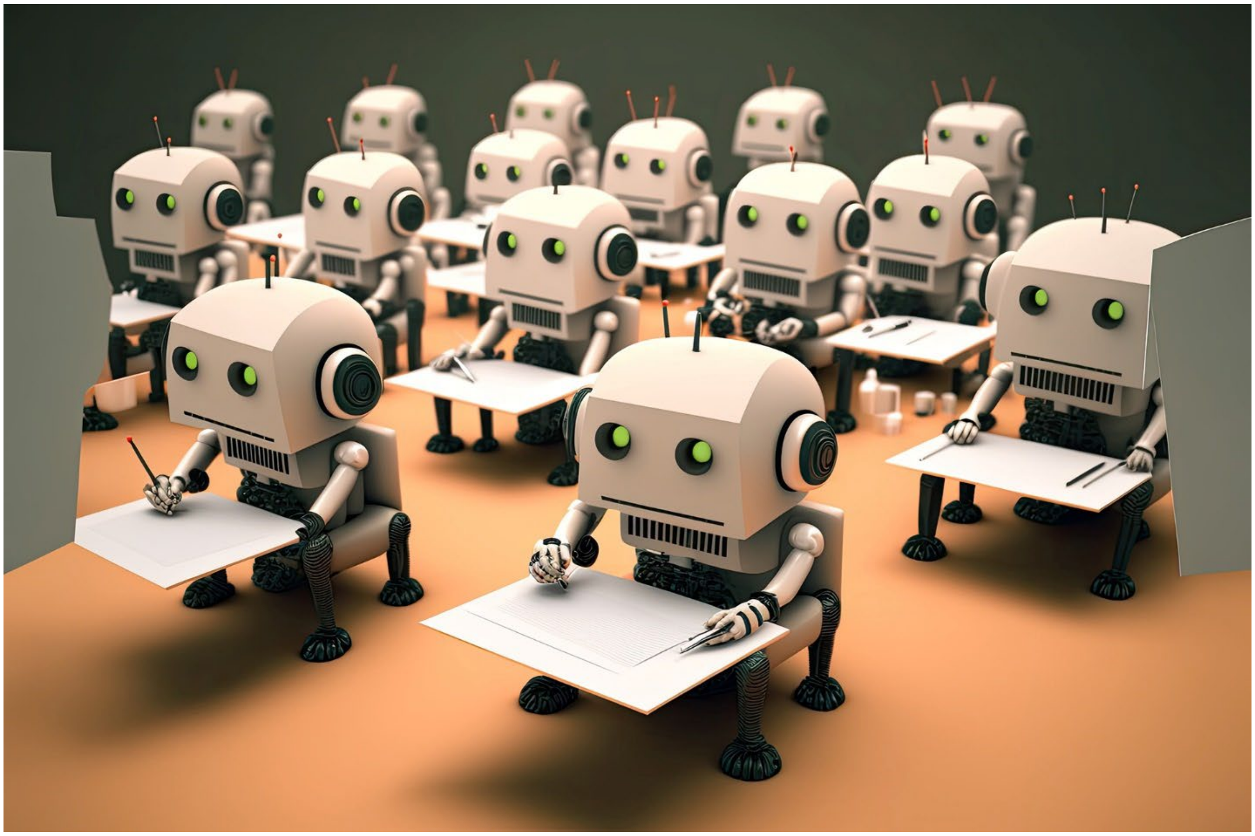

We have written about many legal and other issues around artificial intelligence. To show that these issues are not just theoretical, here are some examples of current actions.

Practicing law

We recently wrote about a class action lawsuit against DoNotPay, claiming that it uses AI to practice law without a license.

Copyright

A group of artists has sued Stability AI, Midjourney, and DeviantArt for copyright infringement. The California class action claims the services used their copyrighted images without permission.

Programmers have filed a class action against Microsoft, GitHub, and OpenAI claiming they violated copyright law by reproducing open-source code using AI.

Ownership

Courts and intellectual property registrars around the world have routinely refused to accept claims of authorship and ownership for copyright and patent purposes.

Privacy

Italy banned ChatGPT for privacy reasons. The Canadian Privacy Commissioner has launched an investigation into ChatGPT. The UK's Information Commissioner's Office published a post setting out 8 data privacy questions that should be asked about generative AI.

The Clearview AI facial recognition database has been banned in many countries, including Canada, for breaching privacy laws.

Issues include personal information in training databases, inaccurate output, and lack of consent to deal with personal information.

ChatGPT defamation

AI text output is notoriously inaccurate. ChatGPT said a Mayor in Australia was accused of bribery and sentenced to prison – which was not true. The mayor threatened OpenAI (the owner of ChatGPT) with a defamation lawsuit if they didn't fix it.

Fake interview

The family of former Formula One champion Michael Schumacher is reportedly launching legal action against a German magazine that purported to publish an interview with him. Michael suffered a brain injury in a 2013 skiing accident and has not been in public since. The purported interview quotes had been AI generated.

Just the start

These are examples that have been publicized. There are no doubt many more in the works. It will take some time to sort these issues out.

Image credit: ©Zephyrfoto - stock.adobe.com

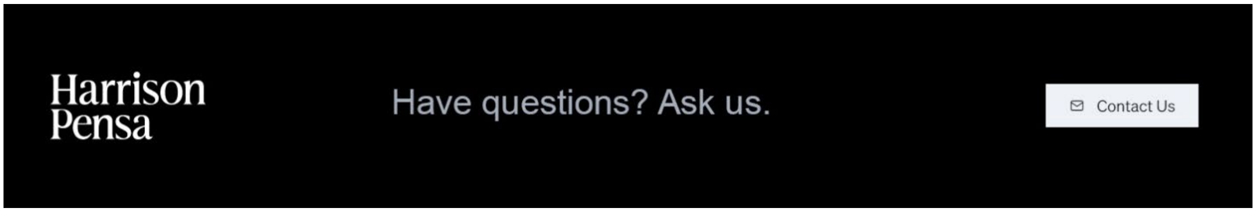

Do you need an AI policy?

March 15, 2023

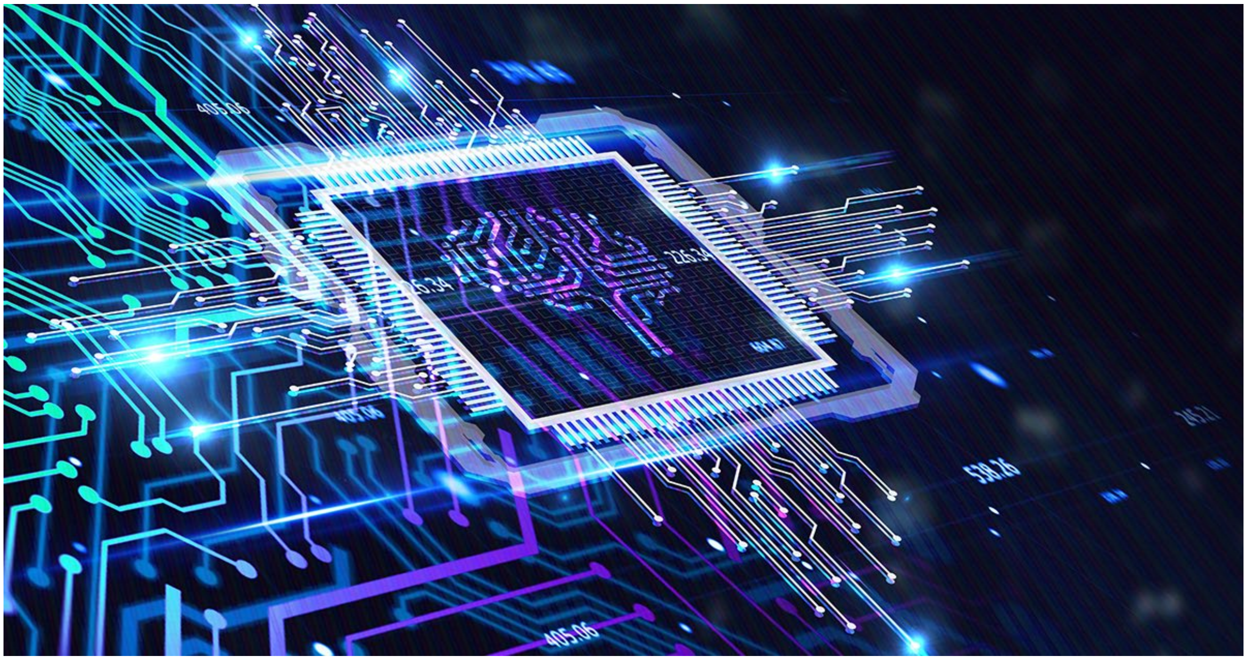

We are inundated with Artificial Intelligence (AI) tools. Examples include image generators such as Dall-E, Stable Diffusion, and Midjourney, and text generators such as ChatGPT, and Jasper. Google and Bing search tools also use AI to generate search results.

AI has been used for several years for things like recommendation engines, but only recently have AI tools been generally available for use by everyone.

It is inevitable that we will use AI tools in some fashion. But we need to keep in mind that AI carries several legal and ethical issues.

Depending on the entity, the nature of the AI tool, and how the output will be used, relying on them might cause issues.

Businesses and organizations should consider whether they should implement rules around the use of AI tools in an AI policy, or an addition to their existing technology use policy. Those rules might prevent certain activities or might put a vetting process in place before use.

The goal could be to set a framework for ethical AI use. Also, to control inappropriate rogue use of AI tools by employees who may not understand the risks and issues.

Some of these issues depend on how AI is being used. Using ChatGPT directly, for example, has more issues than if it is being used as a conversational interface layer to some other service you are using. It is important to understand that distinction and what the AI is actually doing in each situation.

AI policy issues to address

- Images and text created by AI tools are probably not protectable by copyright. So, you can’t stop others from copying and using what you ask AI to create. If you don’t want others to use what you create, that could be an issue.

- AI tools generate output based on a database of existing images or text. Some of them retain the queries you make and any document you enter.

- Inputting anything that is confidential, sensitive, or personal should be avoided.

- The output from AI tools is notorious for being incorrect. It also parrots the biases contained in the material in its database. If AI tools are to be used to make decisions, the tool should be vetted for accuracy and bias. AI output should be carefully reviewed before it is relied upon or published.

- If the business wants to formally use AI tools, it should vet the tools to make sure they can be used in a legal way and comply with published ethical standards.

Algorithmic transparency is desirable, and will no doubt be required at some point legally. Consider how transparent you should be with the public and your customers about how your AI makes decisions.

Using AI to get ideas or inspiration for images, articles, or other creative works may not be something that needs to be disclosed. But from an ethical perspective, disclosure may be desirable if it is being used to make decisions or to create published material. To some extent, it is no different than passing off someone else’s work as your own. See, for example, how Wired recently described how it will and won’t use generative AT technology.

Someone in the business should be accountable for the use of AI – similar to a privacy officer or CIO. That person would be the gatekeeper for AI issues and policies, vetting AI tools, and educating staff.

The last thing a business or organization needs is for an employee to use an AI tool in a way that embarrasses the business. Best to set the tone up front to prevent that from happening.

Image credit: ©putilov_denis - stock.adobe.com

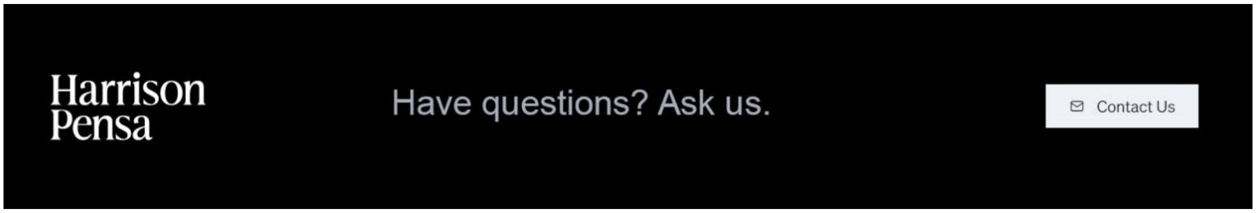

AI privacy:

Duty to protect your data

April 11, 2023

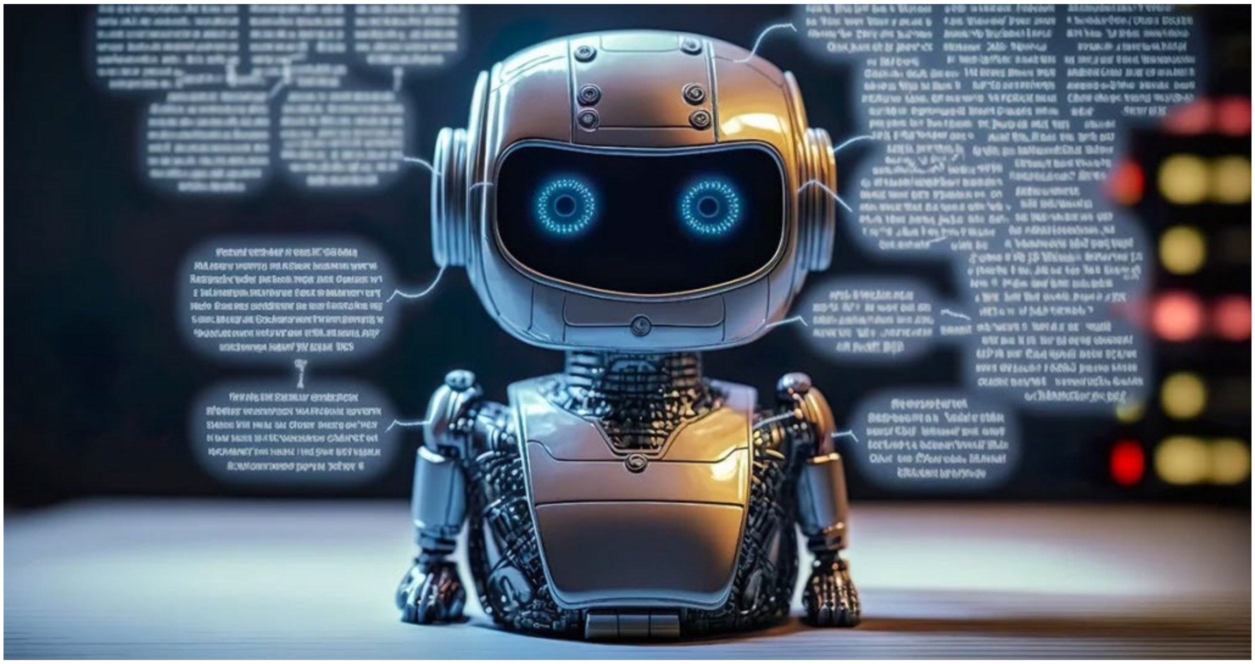

You may have seen that Italy banned ChatGPT for privacy reasons. So how does a generative AI tool have privacy issues?

Before we get into that, it is worth noting that this is not an issue peculiar to Italy. The Canadian Privacy Commissioner has launched an investigation into ChatGPT. The UK's Information Commissioner's Office published a post setting out 8 data privacy questions that should be asked about generative AI. Wired magazine says, "Many of the issues raised by the Italian regulator are likely to cut to the core of all development of machine learning and generative AI systems…"

The Italian regulator cites 4 issues. The first is a data breach that allegedly leaked subscriber payment information and user conversations. The second is that it doesn't adequately verify ages of users who are children. Those issues are not unique to generative AI tools.

The broader issues unique to generative AI are related to the fact that it uses machine learning to generate responses. It trains on data sets that have been put together by various means. Those data sets may contain personal information. Also, the data users input as prompts to generative AI tools may end up in the training database.

The privacy issue is that no one consented to their personal information being used in a data set for training purposes or to be published in AI-generated output. And the platform may not be obtaining proper consent to use data input by users. On top of that, the personal data in the training database and the output may be inaccurate.

In a recent blog, we asked if organizations need an AI policy, and pointed out that inputting anything confidential, sensitive, or personal should be avoided. Unless of course the platform expressly promises that it does not retain that information.

This issue has the same root cause as complaints by artists that they didn't give AI generative art tools permission to use their works to learn on or to create new works that look like theirs.

Just because we publish our personal information and creative works on the internet or social media doesn't mean anyone can copy it for whatever they want.

Image credit: ©Atchariya - stock.adobe.com

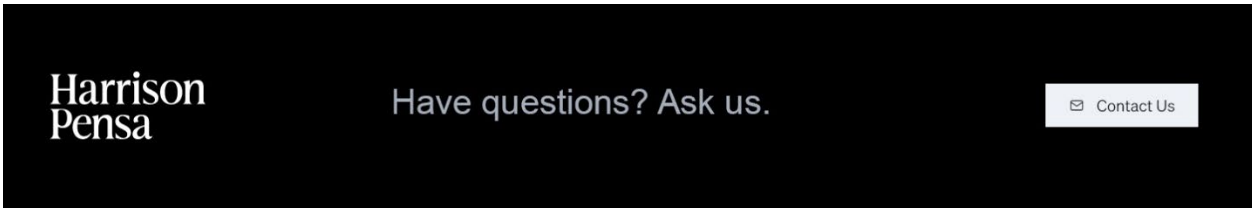

Privacy compliance when using facial recognition

April 26, 2023

The British Columbia Privacy Commissioner determined that four Canadian Tire stores in BC violated BC privacy laws for their use of facial recognition technology.

The stores had cameras that collected images of customers entering the stores. The system compared those to a database of individuals “representing persons of interest who had allegedly been involved in incidents at Canadian Tire stores in the same region.”

The stores removed the systems and deleted the databases when the privacy investigation started in 2021.

The decision points out that 130 privacy authorities around the world have “expressed significant concerns” around facial recognition technology.

Consent for facial recognition

Problems in this instance included that the stores didn’t notify customers it was doing this and didn’t obtain consent. Also, they didn’t demonstrate that the use of the facial recognition tech was reasonably necessary.

There is a high threshold to pass to convince privacy commissioners that using facial recognition technology — or any biometric information for that matter — is necessary. In this decision, the Commissioner went so far as to recommend that the BC government amend existing laws regarding the use of biometric information.

The decision recommended that, “The stores should build and maintain robust privacy management programs that guide internal practices and contracted services.” The need for privacy management programs and privacy impact assessments is becoming more apparent. They will be required under the new CPPA which will replace the Federal PIPEDA legislation. Privacy compliance will require more than just a privacy policy.

Image credit: ©ChaoticDesignStudio - stock.adobe.com

Alternate AI dystopia of ethics and originality

March 29, 2023

The typical dystopian AI concern is that it will become sentient and take over. But is there an alternate AI dystopia that is being overlooked?

I'm not trying to be alarmist, but it's an issue we should ponder. This might be to react appropriately to AI advances, to consider how AI should be regulated at law or approached from an ethical perspective, or perhaps for a dystopian movie script.

Don't get fooled by the hype around AI. It is not a passing fad. Bill Gates says artificial intelligence is "the most important advance in technology since the graphical user interface."

The alternate AI dystopia goes something like this:

- AI tools are here, they are not going away, and can be useful tools for many things. But it will be disruptive to many jobs and people.
- AI generative tools use existing material to "create" what we ask it to. That includes creating new images and text.
- But AI doesn't really create anything new. It only repackages what already exists.

- Creative people get frustrated because AI uses its existing material to generate its output, the AI output competes with them, and fewer people will see their original work.

- When more people use AI tools that summarize content instead of doing Google searches that link to one’s content, fewer people will find original material or know who wrote it. So, people stop creating or being thought leaders because their stuff just ends up in the amorphous AI ether. It feels like shouting into the dark abyss.

- At the same time, AI output is inaccurate and can perpetuate misinformation and biases.

- That flawed AI output is put back into the AI learning pool which reinforces the hallucination and misinformation — the ultimate confirmation bias.

- So, AI-produced information is increasingly unreliable and wrong. At the same time, people give up creating.

- It is a recipe for a downward spiral in creativity, truth, trust, civil discourse, problem-solving, and progress.

- So, AI wins and society loses — not because AI becomes sentient and takes over — but because people forget they are in control and give up.

No AI text tools were used to create this post. But the image for the post was created using an AI image tool to describe an alternate AI dystopia where artificial intelligence replaces creativity.

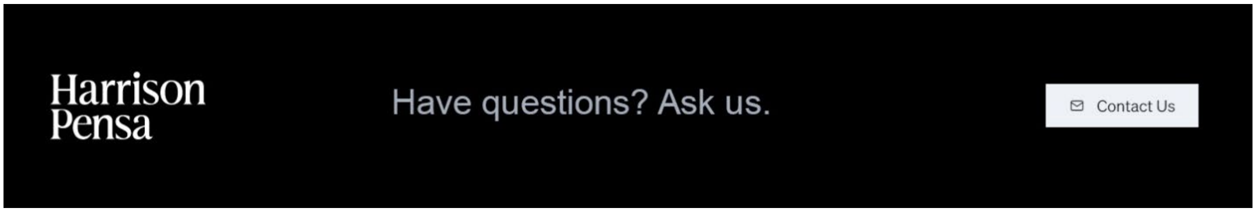
Have questions? Ask us.

Unreliable AI tools and machine learning

February 15, 2023

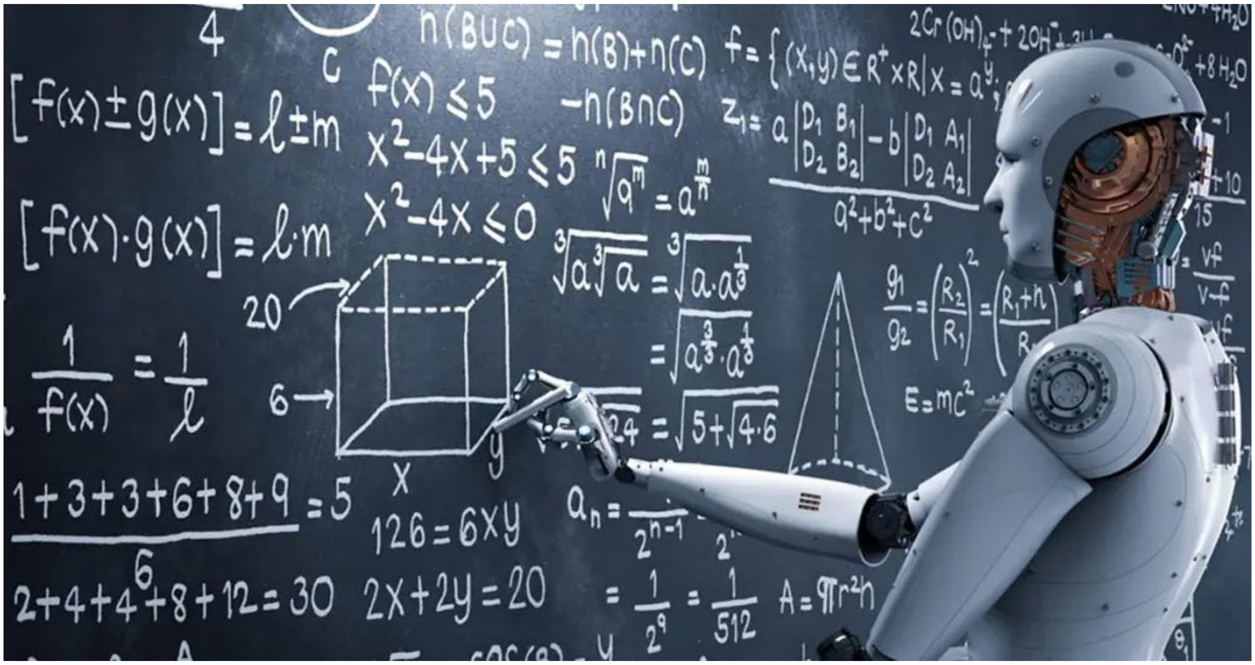

Artificial Intelligence and machine learning tools are coming at us so fast it’s like drinking from a fire hose. The dawn of AI and the analog-to-digital conversion that enabled it will be considered as significant as the Industrial Revolution. Both bring new opportunities, disruption, fear, uncertainty, and change. A significant difference, though, is the speed at which this is happening.

To quote from a Wired article, “artificial intelligence is here. It’s overhyped, poorly understood, and flawed but already core to our lives—and it’s only going to extend its reach.”

AI has been in development and in use for some time. It drives things like recommendation engines for online shopping and video providers, and facial recognition to unlock our phones. The recent launch of free publicly available tools like Dall-E to create images and ChatGPT to create text have brought AI into the public eye.

Microsoft and Google recently announced they are including AI in their search engines and other products. There is a huge amount of competition in AI-assisted search as well as AI content creation.

Machine-learning AI

AI in its current form is not intelligent. It is machine learning, meaning it looks at large databases of existing material and generates results in response to natural language queries based on what it learned. Dall-E, for example, learns from a database of images of art and photographs. ChatGPT learns from a database of written material.

It is more like a parrot than Data from Star Trek.

Nobody asked if our images, information, writing, and art could be used by AI to learn about and create potentially competing works. Scraping online images of our faces for use in facial recognition tools has been banned for privacy reasons. Despite our willing publication of photos of our faces, we didn’t consent to them being used to create an online police lineup. Might similar logic be applied to AI use of our content for learning purposes?

Much art and writing are derivative in that they are influenced by what we have seen and read before. That’s what style is all about. But laws and ethical rules kick in when that morphs into copyright violations and plagiarism. Are these AI tools nothing more than “high-tech plagiarism”?

Are AI copyright and plagiarism issues like privacy and practical obscurity? Before the web, many things were considered public information. But putting documents online that, for example, used to be only in paper records in a courthouse that few would take the time and effort to find and look at makes them available to anyone with a browser. Practical obscurity effectively hid from view things that were considered public. But the notion of those being public had to be reconsidered when they were easy to access. In that context, what should private or confidential information really mean? Do we have to similarly rethink the notions of ethics, copyright, and derivative works when AI is creating things based on pre-existing works?

AI ownership

Attempts to attribute ownership, authorship, or invention for patents and copyright to AI or non-human creators have not been successful. Does that mean anything generated by an AI tool is in the public domain and free for anyone to use or copy? Should we have some sort of AI watermark or disclosure obligation when publishing AI-generated works? How much human manipulation and skill must be added to AI-generated art or text before that human is considered the creator, author, or owner?

AI output sounds confident and convincing. But it is notoriously inaccurate and unreliable. Even Google produced a wrong answer when they debuted their product. Is AI text output destined to be obtuse, specious, and inaccurate? Are we going down the path of lowering writing to the standard of a supermarket tabloid?

We already have a huge problem with misinformation, confirmation bias, and the lack of a desire to base decisions and opinions on facts. Is AI going to make it worse? How do we prevent this? How do we tell what is nonsense?

Is the confident misinformation AI presents the nature of the beast, or is it because the material it learns is of dubious quality? After all, garbage in, garbage out. Can the misinformation issue be corrected? AI output can read like link bait articles found on the web or social media that have no substance, don’t provide what they promised, or rely on sensationalism to attract attention. Perhaps AI has learned what humans, unfortunately, do well?

Lots of questions

How do we deal with malicious uses of AI, such as those that use it to create malware or use it to impersonate for fraudulent purposes?

How do we deal with the issues of AI bias and transparency?

Will AI replace [insert your job here]?

How will AI change how I do my job, and how fast? How do we learn the new skill of prompt drafting?

AI tools will support, supplant, or replace things people do. At what point will it become negligent for humans to perform certain tasks or make certain decisions without using AI tools that can do them better, more accurately, and faster?

Will AI ever become sentient? If so, will it be like Data from Star Trek or Skynet from the Terminator? How do we control and direct that?

There is a myriad of AI ethical frameworks. The EU and Canada are contemplating legislation to govern AI. Legislators don’t understand most tech/internet/social media issues, and IMHO gets most legislation that touches it horribly wrong. So how can they possibly legislate AI issues in a way that prevents harm, but doesn’t get in the way of progress?

Impact on web search

Will our online material have to focus less on SEO (search engine optimization) and more on being optimized for AI queries?

In a future where AI chat supplants search and web traffic goes down, how do you promote yourself and your products? How does social media fit into this? Will AI chat tools effectively unite social media platforms?

Now if only these questions could be answered at the same speed AI is developing.

Image credit: ©phonlamaiphoto - stock.adobe.com

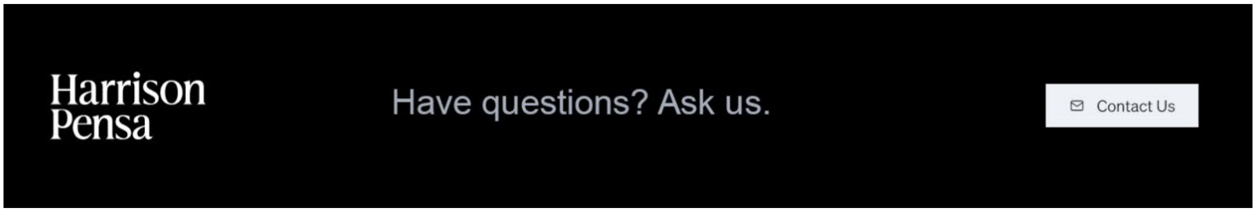

Copyright controversy of AI art apps

November 9, 2022

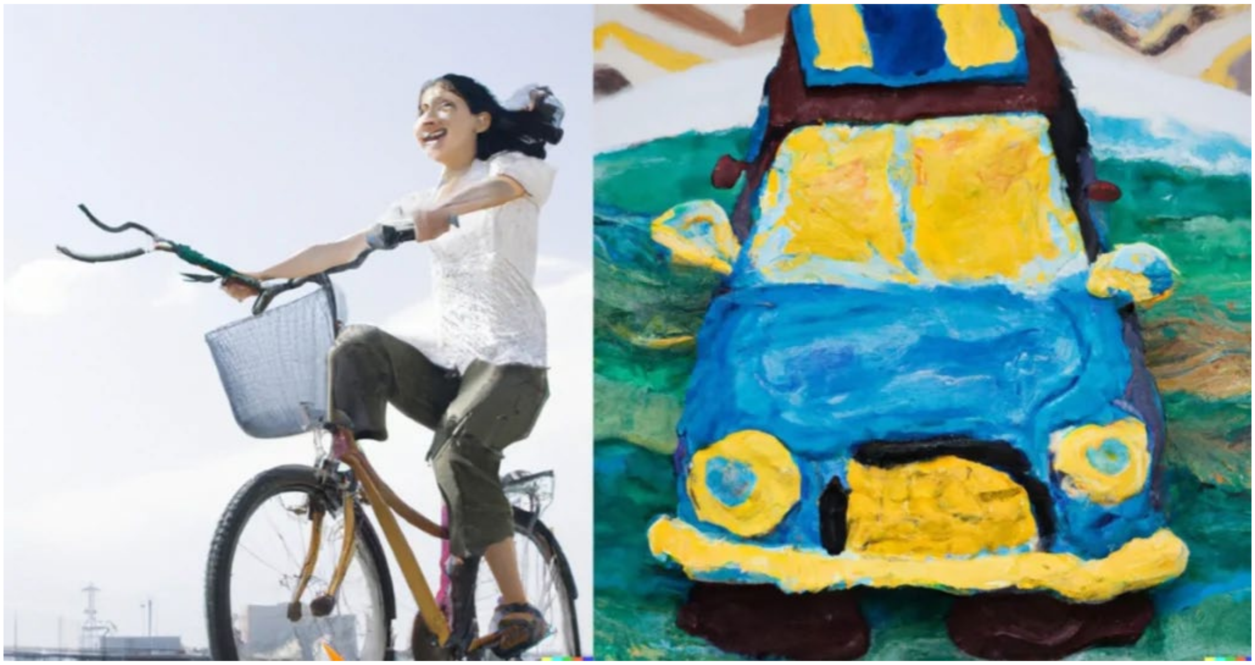

AI art generator tools are becoming readily available; Dall-E and Stable Diffusion, for example. Microsoft is even adding it to its Office suite. So, coming to a PowerPoint near you.

These artificial intelligence apps generate art based on a text-to-image model. If you type “woman flying a bicycle” into Dall-E, it will generate that image. The image is generated based on images the tool has learned from.

But you can also type in “car on a cake in the style of Van Gogh”. It will then generate images that look like it was created by that artist.

Copyright risk

AI generated art is controversial for many reasons.

On the legal side, there are potential copyright issues depending on what images were used to start with and what instructions were given. Some art communities and image sites have banned AI generated art either because of the copyright risk or deference to artists and photographers.

Shutterstock was one of those. But Shutterstock announced it will sell AI generated images that learned from images in Shutterstock archives. They will compensate those artists, but one artist describes it as a “sewer water leak into the drinking supply”.

Who owns AI art?

Authorship can be an issue. Courts have been reluctant to recognize authorship and ownership claims not based on human efforts, such as AI generated authorship of patents, and monkey selfies. How much human effort is needed to reflect human authorship and ownership?

Disruptive tech

It is common for new tech to rile up those who fear it will replace them or existing business models, similar to how the camera concerned artists and videotape concerned movie makers. Those didn’t hurt the existing players. On the other hand, the development of the car put a quick end to those in the horse and buggy trade.

So, which one is AI generated art? Will it put artists out of work? To what extent does it infringe on artists’ rights?

Shelly Palmer says: “AI will absolutely replace millions of commercial artists, and it will do so in the most damaging way.”

Work in progress

Anyone who has tried to create an image using one of these tools knows they are rather crude. Getting the images often shown as examples is not easy and can take time and effort to get right. That will probably change, similar to how digital cameras evolved from a curiosity to replacing film in a short time.

Having AI generated images based on the style of an existing artist is a tough one. Many artists are known for their particular style, and the paintings they create are recognizable as theirs. Understandably, they are not thrilled with the ability of anyone to generate AI based art that looks like theirs – for free. A suggestion has been made that AI tools should not be able to do that for living artists.

It may take some time to sort all this out, but that is small solace to anyone concerned that it will negatively affect their livelihood.

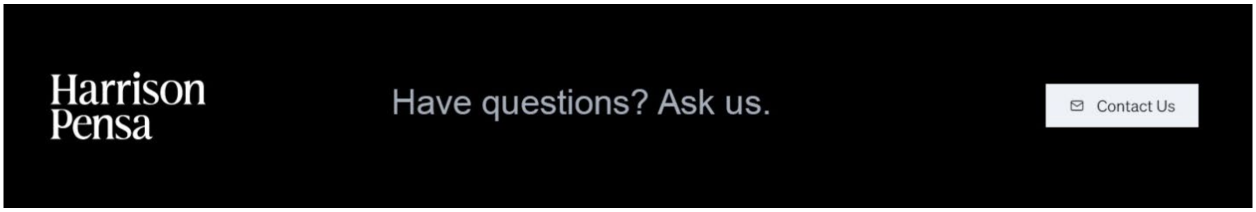

What privacy Bill C-27 means for business

June 23, 2022

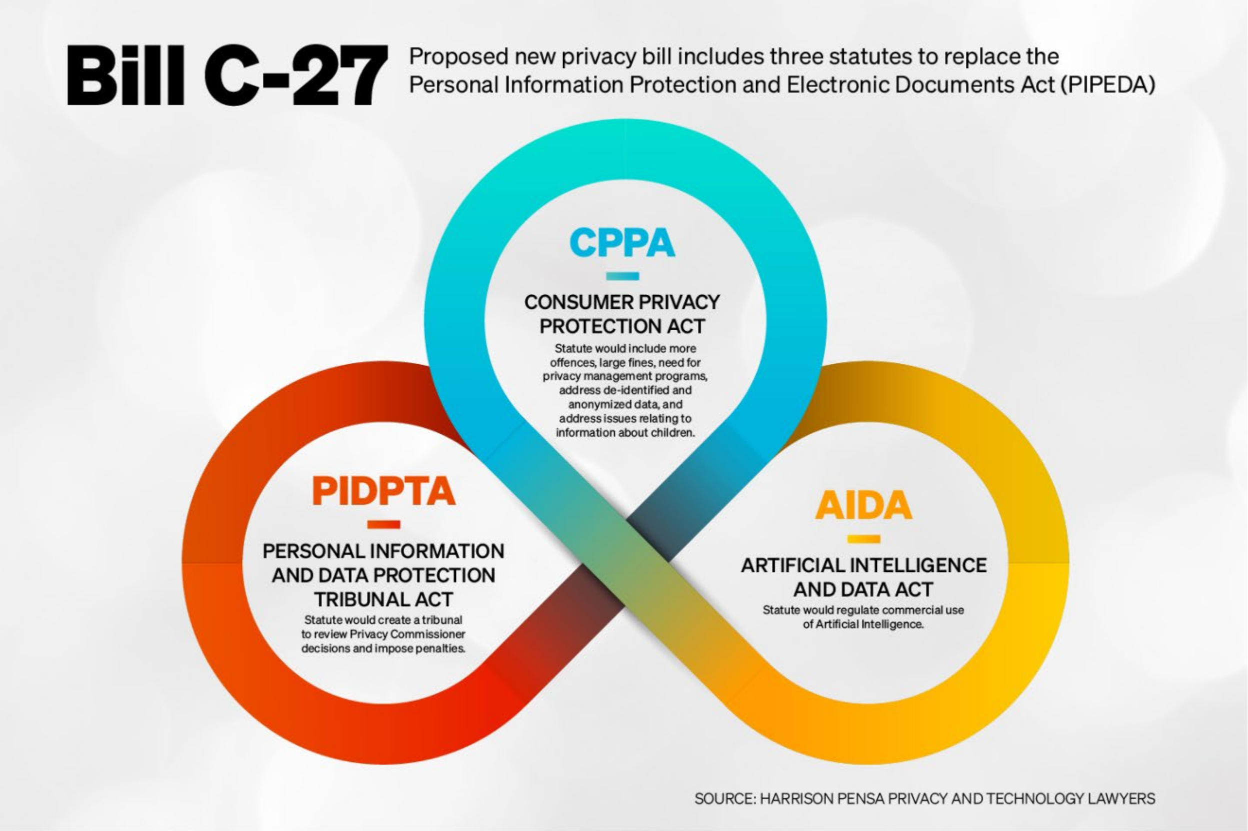

Bill C-27 Proposed new privacy bill includes three statutes to replace the Personal Information Protection and Electronic Documents Act (PIPEDA)

CPPA
CONSUMER PRIVACY PROTECTION ACT
Statute would include more offences, large fines, need for privacy management programs, address de-identified and anonymized data, and address issues relating to information about children.
PIDPTA
PERSONAL INFORMATION AND DATA PROTECTION TRIBUNAL ACT
Statute would create a tribunal to review Privacy Commissioner decisions and impose penalties.
AIDA
ARTIFICIAL INTELLIGENCE AND DATA ACT
Statute would regulate commercial use of Artificial Intelligence.

SOURCE: HARRISON PENSA PRIVACY AND TECHNOLOGY LAWYERS

The Federal government has tabled a new privacy Bill C-27 to replace the Personal Information Protection and Electronic Documents Act (PIPEDA), the legislation that governs the commercial use of personal information in most provinces, including Ontario.

The official title of Bill C-27 is “An Act to enact the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act, and the Artificial Intelligence and Data Act and to make consequential and related amendments to other Acts.”

It will take a while to fully digest the lengthy bill. Here are my initial observations.

- Bill C-27 would enact three separate statutes.
- The Consumer Privacy Protection Act (CPPA) replaces PIPEDA. The Personal Information and Data Protection Tribunal Act would create a tribunal to review Privacy Commissioner decisions and impose penalties.

Those parts are similar to the former Bill C-11 (not to be confused with the current Bill C-11, the controversial Online Streaming Act) that was proposed but not passed before the last Federal election.

- The Artificial Intelligence and Data Act (AIDA) would regulate the commercial use of Artificial Intelligence. This is completely new.

- The outgoing Privacy Commissioner didn’t feel that the former Bill C-11 went far enough. We now have an interim Commissioner and a proposed new Commissioner who will take over in the near future. It will be interesting to see the perspective of the new Commissioners.

- The CPPA will include new order-making powers for the Commissioner.

- It will include more offences and potentially large administrative monetary penalties (aka fines) of millions of dollars.

- It introduces the need for businesses to have a privacy management program (in addition to the normal public-facing privacy policy) that sets out how the business will comply with the Act, including processes to deal with complaints and staff training. The Commissioner has the right to review that.

- The CPPA addresses issues of de-identified and anonymized data.

- It addresses issues relating to information about children.

- PIPEDA obligations extend only to the entity individuals deal with directly, and not to service providers to that entity. (The European Union’s General Data Protection Regulation — GDPR — calls them Data Controller and Data Processor.) It is up to that entity to impose privacy obligations on the service provider contractually. But the CPPA has language that imposes direct obligations on service providers under the Act for some things, such as security safeguards.

- AIDA would regulate “high impact” AI systems. The definition of “high impact” will be included in regulations that have yet to be drafted.

- The AIDA would require the publication of plain language explanations of how AI is used and what it does. Regulations that have yet to be drafted will contain much of the detail.

- The AIDA contemplates a new senior official called the “Artificial Intelligence and Data Commissioner.”

- Offences under AIDA carry significant eight-figure penalties.

It will take some time before this becomes law, and we may see changes before it’s passed. The regulations that have yet to be drafted will be a crucial part.

There will no doubt be criticism from those who feel it does not go far enough, and from those who think parts are impractical.

No matter how it shakes out, all businesses subject to the act will have to review their privacy policies and procedures to make sure they are compliant. Most will have to revise or replace written policies and procedures they now have, especially in light of the new privacy management program obligations. They may also need to change processes and systems to comply with new obligations and rights of individuals.

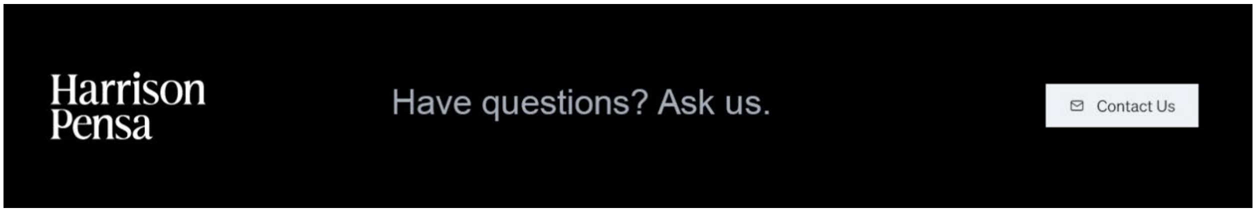

Sign Up:

Tech Law Weekly

Get top news stories from around the world delivered to your inbox so you're up to date on the latest in technology, privacy law, artificial intelligence law, and intellectual property matters.

David Canton is a business lawyer and trademark agent at Harrison Pensa with a practice focusing on technology, privacy law, AI law, technology companies and intellectual property.

Harrison Pensa LLP
130 Dufferin Avenue, Suite 1101
London, ON Canada N6A 5R2

harrisonpensa.com

Harrison Pensa TECHNOLOGY & PRIVACY LAW

hptechlaw@harrisonpensa.com