

Generative AI: Opportunities and Challenges for Students, Teachers, Professionals, Lawmakers and Policy-setters

Thank you Christoph, and friends, for allowing me to join today. Thank you again. For the next few minutes, we'll get to hopefully do some exchange and learn from each other. My background is actually not only studying law, but also studying computer science with a specialisation in AI when I was in university. So I feel very lucky, but I also feel the two disciplines are very much, interconnected, interrelated, in the sense that when you draft a good contract, you don't want syntax errors. You want to define terms properly, and you don't want to, you know, have infinite programming loops. You know, in some ways, I always have the urge to encode and decode something, find logic out of chaos and try to resolve it. So my childhood pastime includes how to, I still do programming these days, but includes how to understand, encode and decode some Chinese Feng Shui telling software, idea like Kei Mun Deon Gaap. Maybe you have heard about it. It's interesting, and I thought it's so much easier for people to put in the palm of their hands and work out, you know, the results from the algorithm. But innately, with all algorithms or programs, if you encode them, is it biased? Is it fair? Is it traceable? What is the logic behind it? If you don't trust me enough, you might then say, "No, I'm not going to use your program, Nick, for Fengshui". But you know, the same goes generally for everything we use.

We use AI every day, in, you know, trying to navigate my way to the university, in getting recommendations for my YouTube program or Netflix, or also just a spellcheck on my computer, I'm using AI. The computer is using the AI. We might not appreciate it. But there's so much of AI that is around us. And for lawyers, the legal profession, I think we can't afford not to embrace it. I actually teach Law and Technology at CityU, PCLL, LLB, LLM and JD. So let's go straight into it.

I think AI presents a huge paradigm shift for economic development. And I think as lawyers in the profession, apart from solving a problem at a time, we want to have a higher calling, higher purpose, which is to help economies and help, I think, bring about peaceful resolution of disputes and a win-win collaboration.

In history, we have the first industrial revolution, the second and the third, right? We have steam engines, we have electricity. We have, you know, the digital age. Now we are in the age of the AI revolution, and history tells us that if we don't seize the opportunity, we will fall behind. So as a lawyer, as a profession, as a community, as a school, as a country, nation, people, we must take heed of

warnings, but also the opportunities available to us. And I think it's not just a matter of understanding AI. Sorry, I should say it's not just a matter of adopting AI for our practice. It's also about how we contribute to the development of AI. Often, people say law lags behind technology. I like to think, "Well, true, but how can we be one step forward?" "How can we foster the development of AI?"

So when the Hong Kong government talked about, a recent consultation about how we would update our copyright ordinance, and in view of AI developments, my submission was, let's not beat around the bush. Let's go straight to it and pass an AI law that specifically deals with AI. I feel that a copyright law in itself cannot properly deal with ethical issues, because, for example, if you have a copyright ordinance, you could only really deal with what is copyrightable and what the exceptions are, fair use, how you use it, and how you compensate somebody. It doesn't really build an entire ecosystem. It sort of lays out some groundwork, and it follows on the International Berne Convention. Essentially, most people read and interpret it to mean a computer-generated AI, completed by computer with no human intervention, cannot have copyright to it. The idea of copyright is that it gives a monopoly to someone to use the material to make money out of it, but that goes against the idea of helping the society, you know, use a lot of different things, and foster the society. However, if you don't support someone with some monopoly over trademark copyright, then who's going to develop it? Right? So, yes, we give the society under certain conditions. The law gives a monopoly to people who create copyrights. But then, is that good enough? Where does it come from?

And the fact is, with AI, specifically Gen AI, even before we go to artificial-generated intelligence, we're really talking about what we're seeing a lot these days, why AI is in our mind these days, and its ability to give you an article that reads and looks like it's written by human. The Turing tests is about if someone stays behind the screen and communicates with you, can you tell that it's a computer or a person? If you cannot tell it's a computer, it passed the test. Then people tried to capture the imagination with playing chess. And next, I think, these days, people generally accept and agree that the article written by the computer is so beautiful, how is the computer now so smart? Is it really that smart? The generative AI? Honestly, it's not. It's still a baby at the moment. I think it has got a lot of room to grow, but also that means we have to get into it and get onto it right now. And I'll be looking forward to Ben's discussion later on. I think it's a strong case, but let us see what makes it

a strong case.

So AI is something we should look out for. And if we could help the law drafted in such a way to attract, imagine we pass a law that we attract developers to come to Hong Kong to develop AI, then, not only can we benefit coders and programmers, we can also benefit everybody else in industries who want to think about adopting AI. Then you have experts, more experts, software engineers on the ground, prompt engineers and so on and so forth. So attracting money, and attracting talent and innovation that it expresses itself, that is what I think should happen. Change is coming, whether you like it or not, AI is everywhere. You probably still listen to Spotify while sitting here. I'm kidding. But as lawyers, if we don't understand AI, we're missing out, not only missing out on productivity, but also missing out on the chance of being appointed for our jobs. For someone who studied AI, it's easy for me, in this day and age, to say, "Oh, you have an AI project? Look for me. I think I'm okay". It helps. So then, I was paid to fly to Europe to talk to Daniel Ek to investigate into whether it's worthwhile to invest in this music streaming company Spotify. So eventually, we stayed there for a few days and recommended investment, and we brought them to Asia. The rest is history. You know, living history, we're still using it a lot.

There are a lot of things to think about regarding Artificial Intelligence. A Goldman Sachs report suggests that it would replace a million full-time jobs, and Generative AI would replace 44% of legal jobs and so on. Is that worrying? Yeah, it looks worrying, but it's also creating a lot of jobs. So you want to look ahead to the future of work. How to get the mundane stuff out of the way? I mean, as a partner, I guess I try not to do that. I have some control over whether I do the mundane stuff or the more interesting stuff. But I would like to help my workforce and my colleagues develop, to be better persons and better lawyers by challenging themselves, challenging the mind. And boring stuff, let someone else, let the computer deal with it. So we should find gaps, places where we could be more productive. And so, it's going to increase global GDP by 7% annually.

AI has many different forms, and there are many different relevant companies and studies, but I guess the science of making things smart is one way of defining it. I think there's no general definition of why they accept it. But essentially, we all have our own way of thinking about it. One day, people will

talk about Matrix, I don't know if you've seen Matrix. You look probably too young, but Matrix has been made again, so I can say the Matrix. You know, the speed of the Skynet is basically the speed that we use now with some of the NVIDIA chips. So, are we there? Is Terminator here? It's not a scary story in reality. It's how we use it. But if we don't have the right ethical attitude to contain its use so that computers don't talk to themselves and we don't understand what they're saying. It happened before with Facebook, with Meta, and they stopped the project, they worried, and we perhaps should stop AI from making decisions about shooting missiles and starting World War III.

But with some of these names mentioned here, you might recognize this more than that, depending on where you live. But at the end of the day, though, some universities in Hong Kong, are developing large language models, studying the language, and preparing. Generally, it used to take two or three years to do it. Now it takes only a few months to have something reasonable. So I think these days it's more about sort of going deeper into specific domains. The University of Science and Technology, Hong Kong UST, for example, today, is investing in and working on a legal domain-specific LLM. So it's crunching a lot of law-specific documents. It's easier because Hong Kong is a bilingual jurisdiction where English and Chinese have the same value of precedence. So you can easily crunch the judgments, the law and contracts and so on. Then, that project is more public-facing, for public use, but the public could say, "Well, I want to have a divorce, and what are the issues I need to think about? I might need one. Can you help me?" And the computer is like a chatbot, talking you through, and it is more than just an internet search. It has studied the domain well enough to know what questions to ask and kind of expect the range of answers you might have, leading you to another question and answer.

I don't know when, at least when I studied law, you would have a "talk in a nutshell" and a "contract in a nutshell", you would use mind map and know the professor would, you know, ask this kind of question, then you know, you have to have a Boolean question. If A, then else; if A or B, then C, kind of concept. So you write it out. And you kind of, as a law student, probably have case ready to copy up to say "I cited this. I'm aware of this". And you kind of reverse-engineer the scoring scheme, right? So, artificial intelligence, how do transformers work? Transformer, of course, not the cartoon transformer. Transformer is really about encoding and, you know, enabling learning of grammar. Let's

have a look here. So essentially, if the user puts in “To sleep well”, then the computer will look into the past database; the database it has might be static, you know, it might not be. Then you'll be able to say the probability of the next word to be “you” is high, therefore, I'll give “you”. And then the next word “are” is high, then I give you “are”. But you might not know that.

In Hong Kong, I studied overseas when I was 12, and I often got decent marks in English, even overseas, although my pronunciation might not be as native, because I think in Hong Kong, somehow the way we're taught in English is more formula-based, if I could say. Is that the phrase I should use? Maybe? Virginia says, “No, no, that's not the case”. But in some jurisdictions where every day, you wake up, everything you're speaking is in a particular language, say in English, you kind of picked it up and you are able to, you know, go with it, run with it, but without realizing how there're actually concepts and rules behind it. So the way it works out is that it isn't trying to work out a grammar rule. All this is doing is basically working out the probability of one word appearing after another word. That's all it does. So he has no concept of what's a preposition or what's a pronoun, and he has no idea. And if he says, “A is bigger than C”, if it comes up this way, you might ask, “How do you know it's bigger than C?” “I just know it”, like kids would say. But how would you know? “Because of the probability in all the books I read that A is bigger than C.” In most of the books I read A is bigger than C. You wouldn't be able to say, “Well, I read from somewhere, based on what Einstein said, A is bigger than B, and B is bigger than C, therefore, A is bigger than C.” It doesn't do that. So this kind of encoding is relatively easy. That's why it's so fast. It doesn't try to understand the context. There's no context to it. So that's why, you know, ChatGPT, the likes of this kind of generative AI, how it blossoms and blooms.

So you might pretrain something before you let people use it, and then you fine-train it, you know, with specific tasks, involving programmers or beta users trying to improve it and reinforcing the learning. So is learning a language, such as English and so on. When I was in a secondary school in the UK, I had some time in my hands. I was silly, rather than popping over to Europe to play over the weekend, apart from watching a musical, what I did was that I picked up a book on how to write a best-seller novel. I started trying to write, trying to be an author. And basically, you studied so many different books, like *To Kill A Mockingbird*. There are many different classics, and you're trying to work out what kind of protagonist you need, and what kind of plot you need. At the end, someone must

die. There must be some lovers. You know, there is something and some common ways that people can feel affinity and want to buy and it's a successful book. So, I was determined to turn it into a computer program as well. Okay, now I study this. You may punch in some details, you know, put in your name, and put in someone you hate, someone like your ex, or something, and then you come up with a book, and what kind of context you like. So I had something in the making of that, and I found it fun. It's about reverse engineering. What is the norm? What can one expect as right?

Also, sometimes, if you want to draft contracts, or sometimes seek advice, or maybe more simply put, a prospectus for an initial public offering for IPO, for listing, publicly listing companies, you need to have a book called prospectus. Then you need to verify everything. There are many statements. Everything will end with a full stop. And as a junior lawyer, some of you might have done this verification work. You turn every sentence into a question. "Ube is handsome. Is Ube that handsome?" You have to add the question mark to the answer. Then you have all this you signed off, and who I'll trust is a professor, so and so. It's not using my brain as a junior lawyer. So I didn't like that very much, but at least in a way, if you kind of know what the SFC wants to see in the book, then you kind of reverse engineer, and you could crawl effectively when people do the risk section. In the risk section, sometimes people just rip it off and copy it from some other recently listed companies in the same space. And then, of course, they did finetune it. Back then, they didn't have to help with Gen AI, so it's a little bit more work in a few more seconds. But people, lawyers, and professionals, are using technology a lot.

In terms of learning the language and learning the culture, AI translation isn't always accurate. Yesterday, I was at an international conference, and then someone said, "Salamat dating", but he missed the space in between, so it means something else, but it's a welcome dating, and it's meant to just say welcome. Anyway, I found that really interesting.

Here I want to talk about different large language models, as you're aware. And then one is a different large language model, some are for the general-use cases, and some are specific-use case models and foundation tools. There are different areas for lawyers to specialize in, or coders and other people to specialize in. There are lots of opportunities. I think with all this development, a lot of intellectual

property rights are involved and a lot of personal data sometimes are involved. So it generates a lot of opportunities for lawyers who want to contribute. And lawyers have to work with people in different disciplines, the English Department, for example, we could together build a lot. Large language models are not strong in logic or reasoning. It's just probability at the moment.

Don't use AI blindly. In some local cases, they are not mentioned on the slides, but in local cases, some junior judges wrote something, copying or plagiarizing another judge. It was called out. It was embarrassing. And sometimes people refer to a law case that doesn't even exist in the world, but ChatGPT says so. That inherent bias, if you only see the world through one lens, is different. So it's highly possible that if you use some Gen AI produced in the US based on US data, it may reflect local culture and norms, but it might be totally different from what you expect in Asia. For example, on Netflix, there's a movie called Coded Bias, and a person of colour complains about the computer not recognizing her as a human being because they just didn't have the bio data or facial recognition data to train the system. So, I think Hong Kong has a big role to play. We could access data from mainland China and overseas, and if we do it properly, this is a great place to provide unbiased LLM for the world. At Cyberport, we have put in a lot of money. The government put in 3 billion HKD, and it's about to be fully built out, and we will allow licensing. So, it is a sort of computer with supercomputing power and super AI computing power as a service. If your department or colleagues and friends are interested, come talk to us at Cyberport.

Misinformation is sometimes intentionally done, sometimes not. Some concepts are immoral or illegal, such as teaching you how to build a bomb, and so on. Traceability is also an issue I want to talk about. If A is bigger than C, it doesn't explain why A is bigger than C. It doesn't give you credit. When the credit is due, you don't know how you get it from. Then how can you rely on lawyers? How can you rely on a colleague who says, "This is the result"? If you're at the top of your game, people rely on it. It could be a bad thing. But if junior people want to double-check the stuff, you don't know where it came from, yeah, you might say it came from a Gen AI program. Next, confidentiality. Lawyers have to be careful not to breach clients' confidentiality. If you ask the program the question, "Oh, my client is this and faces that issue". Of course, you might be smart enough to say, "Take away the name of the client", at least, and then you put in the system. But then, if the system knows this question

and then gives you the answer, maybe you are the first one to ask and might not have a proper answer. Then you have iterative discussions with it. And then your opponent law firm, law team, or legal team asks, "What would the other side ask for some cases like this?". Then what if it comes out and blurts out, giving the client's confidential information off to a third party, like the AI model, is wrong. So how do you overcome it?

I think the way the Singapore government overcomes it is that they work with the developer and say, "Silo the system" so that the programming aspect is done elsewhere and it's prepared elsewhere, but I can use it on my premises, and the questions that are kept in my server don't go out, just like when you use the LeaveHomeSafe app. The app might ask some questions about you, but that data doesn't leave your device until you're found to be COVID-positive, and then you go to a government place and are willing to disclose it. Otherwise, it's kept on your device. Privacy-by-design, you know, those things are what we have to think about.

Okay, I might not have a lot of time to talk about different cases, but cases around the world generally share some common basis. First, Berne Convention is the copyright convention that a lot of countries rely on and model the local laws on. Generally, if something is not created with a lot human involvement or some human involvement, it is not copyrightable. So everyone else can copy that and its themes. But that's it? If you copy that, you're infringing someone else's stuff. You see a secondary infringement and so on. It could be. In a lot of these cases around the world, not just this one, people try to be the first mover in history and say, "My computer. With my computer, the computer-generated copyrightable work that is respected as a copyright". In most cases, they fail. If they succeed, it could be the judge taking the view, "I also want to make a first in history". But you know, generally it shouldn't be allowed, you know, because otherwise, how does that computer system defend itself if it owns something? How does it defend itself? Would things change when they become artificial intelligence with emotions and thinking like a human? Maybe things might change then. But if you just switch it off, then who owns it? Next of kin? Not sure. So there are a lot of things to think about. In Hong Kong and the US, there are different ways of dealing with copyright, although the Berne Convention underlies it. You know, the concepts expressively list what's allowed, what's not allowed, or something, are more specific in Hong Kong. This kind of exception in the US is more

general. It could be that the AI system is not eligible for copyright, as it is absent in human involvement. Zarya of the Dawn is only entitled to a copyright registration of only the text and the arrangement that is decided by the author. But, if the artwork has enough iterations of human involvement or prompt engineering, maybe that should be protected, too.

AI should be viewed as a tool, in my belief. Based on the technology available, it isn't a separate actor today. It isn't like another human. It is a tool. So if you think of it as a tool, how do you deal with it? China is quite advanced. I mean, I do see a lot of Western media saying, "Oh, the EU is very fast and the first to come up with legislation on AI and Gen AI". But don't forget the developments in China. The reality is, Gen AI in mainland China is exceptionally fast and you can see how it persists in being the guideline in China. AI, the ethical use of new generative AI, should be people-centered. It's people-centric, right? As a socialist economy, China always prioritizes people first, not just maximizing profit. There are similar thoughts in Hong Kong. The previous Commissioner has come up with guidance on ethical development use of AI. Feel free to read about that. At the international level, the UN has a recent initiative, from Xi Jinping's request or suggestions that's been adopted by a lot of countries, and that is to do with ethical development of AI. That is not that different.

These are highlights of work done since 2017. Our country has been very, very active in, you know, giving a good foundation of law to support the development of AI. The EU has its AI regulations and we definitely should read them and learn more about them. They talk about trying to identify and put things in different baskets. What will be an unacceptable risk if it happens? What is a high risk? And what is a limited risk? Therefore, it's a risk-based approach. That's quite good.

Then at the Law Society of Hong Kong, we have developed The Legal Ethical Code of Conduct regarding the deployment of AI, and we further developed it into an AI manifesto. So we've got over a dozen Law Associations from around the world coming to Hong Kong last month to sign off on this manifesto. We're trying to take this forward. I touched on this earlier. Next, legal application of AI. There are a lot of things we do the same. We use AI and we will use AI, but provided it's done ethically, traceable, and should not breach, you know, relevant rules. From day one, I think a lot of big law firms sent memos to colleagues, saying, "Don't use it because we fear you might breach client

confidentiality". Today, like our firm, we have dozens of programmers and a lot of AI experts. How could we, you know, adopt licensing or build AI tools to help our work? And if we don't use it, other firms will use it, and we will be less competitive. That's the situation in the field of profession. Later on, we will hear from Ben about, maybe the situation in education institutions and how we deal with that.

I'll briefly introduce E-discovery. Some of you might not be from the legal background. In common law jurisdictions, if we go to court against each other, then there's a process called discovery, where to be fair to the process, and for the interest of justice, each side must produce all relevant materials. Unlike a civil law jurisdiction, where if you can't find it or the judge doesn't ask for it, you don't have to produce it. In common law jurisdictions, you must produce it, except for privileged materials, like, you may have heard about, attorney-client privilege. Your communication with lawyers that you don't have to disclose. But if you have a document that is not in your favour, you still have to disclose it. When people disclose things, one of the strategies that people use is to overwhelm the other side. They will give me four container loads of documents and be like, you know, "Okay, be my guest. Read it". But they don't breach the disclosure requirement. It's document pump. They disclose it, but do I have the time and energy and does the client have the ability to pay for that low time cost to read it? Now AI will tell me the percentage likelihood of relevance, and you have to have contextual ability. If I look for murder or I want to find anything relevant to murder, you might also come up with manslaughter or, you know, accidental stepping. So it's able to think about it. So I think it's very helpful. And these days you could also go for videos and pictures. So you walk in, in milliseconds, it tells you which documents perhaps you should read first. But there is no substitute for proper training. AI helps E-discovery, saves a lot of money and makes legal research a lot faster. You know, a lot of people just punch in a case and say, "Tell me what is the reasoning of the case". But have you done more? Otherwise, you will lose the ability to decipher and think critically.

Another aspect is trial preparation. Sometimes AI is used to train witnesses. We are not allowed to coach witnesses. We are not allowed, but the computer has, you know, friendship with the witnesses sometimes. I heard about it, not in my cases, of course. In addition, you can ask chatbots questions like, "How do I prepare a case?" There are plenty of applications you can use in Beijing Internet Court,

Mainland China. Enforcement is made easy with a blockchain, for instance, you confiscate something that obviously is not stolen. Hong Kong has eBRAM and we raised 150 million in support from the government. So, the policy is very supportive. Thank you for the time, we overran a little bit. Thank you Christoph. Thank you. May I pass to Ben, thank you.