Artificial Intelligence and Privacy Laws in Hong Kong: Symbiosis or Discord?

Cheung Tsz Him Ryan

3035398171

LLB IV HKU

Part A: Introduction

When the first wave of the Industrial Revolution hit the world, many relished in this new form of civilisation as it presented limitless opportunities. The steam engine and mass production became utopian means of achieving symbiosis with maximising a country's potential. Now, we are at the crux of a technological revolution that stands to be an upheaval of societal norms and raises unforeseen ethical quandaries. The notion of machines rising to power and forging a digital civilisation no longer seems like a castle in the air: would we be able to govern artificial intelligence ('AI') or would AI govern us?

This article will investigate whether the present landscape of privacy laws, as contained in the Personal Data (Privacy) Ordinance (Cap 486) ('PDPO') in Hong Kong present a symbiotic or a discordant environment for the development of AI. Part B of this article aims to explore the legal struggle between protection of data privacy on one hand, and the need for Big Data for the meaningful growth of AI on the other. This struggle is embodied across different aspects of AI, and these will be examined in turn: First, it will be shown how the present legal framework presents a *Jenga* conundrum to the development of AI, in that the removal of small data packets may jeopardise entire AI systems. Secondly, it will be argued that present privacy laws are not sufficient in regulating methods of data anonymisation for AI systems. Thirdly, AI's latent defects will be exposed when discussing its applications to algorithmic surveillance.

Part C of this article will explore how the legal landscape can be shaped to provide a symbiotic environment for AI to thrive. The importance of future AI-specific legislation will be highlighted as issues such as anonymity rules, strict liability, and principles of control have yet to be entrenched as law.

I. Scope of AI

As artificial intelligence becomes the new buzzword in spearheading the technological revolution, what it actually entails seems to be muddled with ideas of a mass upheaval of the status quo, as popularised by science fiction. To put it simply, AI refers to 'a system that can learn how to learn[1].' Learning is achieved through one of two ways: machine learning and pattern learning; the former functions by utilising algorithms to collect, process and adapt to data from the real world[2]. Pattern learning identifies regularities in data and automates the process of recognising trends for statistical analysis. Evidently, the effectiveness of both ways of learning hinge on the collection of big data, a process that generates 'high-volume, high-velocity and high-variety information assets'[3], positing a strong positive correlation between the amount of big data available and the quality of machine learning. It would then be the expectation that AI could generate new models to accurately map out data inputs and transform them into meaningful outputs. As the growth rate of data available becomes exponential, so would the improvement in AI's sophistication. With such an

[1] Francesco Corea, *Artificial Intelligence and Exponential Technologies: Business Models Evolution and New Investment Opportunities* (Cham: Springer International Publishing, 2017))
[2] Jordan Novet, "Everyone Keeps Talking about A.I.-Here's What It Really Is and Why It's so Hot Now," CNBC, June 23, 2017), https://www.cnbc.com/2017/06/17/what-is-artificial-intelligence.html).
[3] "Big Data," *Gartner Glossary*, accessed June 9 2020, and accessible at shorturl.at/twCI2

automated pathway, it becomes a difficult conundrum for legal professionals to identify the various gears in the machine and pinpoint regulatory risks.

Optimists would point to AI's harmlessness in its applications in gaming, translation, retail advertising and more through 'deep learning', a term coined by Igor Aizenberg in 2000 as a subsect of machine learning.[4] Deep learning delves into the mechanism that governs AI, and often entails 'an agent, action and reward.'[5] The agent, usually in the form of a robot or machine, interacts with its environment to observe a specific activity to respond to, with the hope of producing a beneficial reward. Through adhering to what was inputted, it engages in a series of trials and errors to maximise a cumulative reward. This bodes well with gaming as was the case with DeepMind's AlphaGo, where it utilised deep learning to play Go (a form of chess) and was able to beat the world champion in 2016.[6] OpenAI furthered this by engineering a robot arm (the aforementioned 'agent') to solve the Rubik's cube.[7] Intuitively, one would ponder about more intrusive AI uses, one akin to a machine so powerful that its decisions could jeopardise a person's livelihood. AtomNet, a deep learning system for structure-based rational drug designs[8], exemplifies this worry since it is used to predict novel candidate biomolecules to target diseases such as Ebola and multiple sclerosis.[9]

Part B: Discord

I. Prohibitive Privacy Laws and AI: The *Jenga* Conundrum

A key element in AI is the usage of massive amounts of data, within a neural network. This means that where personal data is involved, it may be used over and over again in successive processes of the AI's algorithm. An example of this is the game Pokémon Go that was launched in 2016. It was aimed at providing players with an augmented reality to engage with the platform's virtual characters. Data such as where the players are going, their surrounding environment and their movements, are tracked and processed under the guise of enhancing the game's interactivity. At the same time, business corporations flocked to grasp the opportunity to gather a plethora of new consumer data, transforming performative data into a business goldmine.[10]

This presents two key issues. First, where a data subject revokes his consent for the AI system to use his data, there is potential for a substantial proportion of an AI system to be paralysed. This is so because, by nature of neural networks in deep learning, stopping the usage of small datasets will necessitate the halting of a large proportion of the entire system. Secondly, as an AI system

---

[4] Iman Raeesi Vanani, *Deep Learning for Opinion Mining* (Iran, 2019), at 40

[5] Tom Taulli, "Reinforcement Learning: The Next Big Thing For AI (Artificial Intelligence)?" *Forbes,* (June 5, 2020), accessible at https://www.forbes.com/sites/tomtaulli/2020/06/05/reinforcement-learning-the-next-big-thing-for-ai-artificial-intelligence/#2dd4460662ba

[6] *ibid*

[7] *ibid*

[8] Wallach, Izhar; Dzamba, Michael; Heifets, Abraham, *"AtomNet: A Deep Convolutional Neural Network for Bioactivity Prediction in Structure-based Drug Discovery".* 2015-10-09

[9] Caleb Garling, "Startup Harnesses Supercomputers to Seek Cures" *KQED,* (May 27, 2015), accessible at https://www.kqed.org/futureofyou/3461/startup-harnesses-supercomputers-to-seek-cures

[10] Kerr, Aphra, Marguerite Barry, and John D Kelleher. "Expectations of Artificial Intelligence and the Performativity of Ethics: Implications for Communication Governance." *Big Data & Society*, (January 2020). doi:10.1177/2053951720915939.

continues to learn, the purpose of successive processing of one's personal data may become different from the original purpose of data collection. Again, if no consent was sought for the new purpose, the successive processes of the AI may be rendered unusable and illegal.

These two issues expose a *Jenga* conundrum. It shows that where the law prohibits an AI system to continue to use a small packet of information, the whole system could be in danger of collapsing.

The Law on Consent

Under Data Protection Principles ("DPPs") 1 and 3 of the PDPO, consent is not a pre-requisite for collecting personal data, unless it is for a new purpose. DPP 1 requires the collection of personal data to be adequate and not excessive, and practicable steps to inform a data subject to the collection of data. In gist, it ensures that there is a right to exercise consent and to opt out from having one's personal data being processed.

This is in contrast to the position of the General Data Protection Regulation of the EU ('GDPR'), where Art 4(11) provides that consent of the data means 'any … informed and unambiguous indication of the data subject's wishes by which he or she … signifies agreement to the processing of personal'.[11] This is symbolic of an opt-in approach.

For these provisions, the right to withdraw consent can be exercised at any time.[12]

The Law on 'New Purpose'

DPP 3 of the PDPO requires prescribed consent if personal data collected was to be used for a new purpose. Under Schedule 1, s3(4):

'a new purpose, in relation to the use of personal data, means any purpose other than— (a) the purpose for which the data was to be used at the time of the collection of the data; or (b) a purpose directly related to the purpose referred to in paragraph (a)'

In furtherance to these generalities, the HKCFI has held *obiter* that it is legitimate to have regard to the '*reasonable expectations of the data subject*' when assessing original purposes of data collection and new purposes[13]. In *Ng Shek Wai*, the same court ruled that the data subject should reasonably expect that names of counsel in a public hearing (e.g. the Medical Council in the present case) will be released, should the public enquire about it. No new purpose was found.

In *Wing Lung Bank Ltd*[14], a data subject consented to the Bank using her personal data for promotional materials to be sent to her. On the basis of this consent, the Bank allowed a third party insurance company to call the data subject to promote insurance products. The insurance company

[11] Art 4(11) GDPR
[12] Art 7 GDPR, and also see s 2(3) PDPO where consent does not include any consent which has been withdrawn by notice in writing served on the person to whom the consent has been given.
[13] *Ng Shek Wai v Medical Council of Hong Kong* [2015] 3 HKC 455 at [52]
[14] *Wing Lung Bank LTd v Privacy Commissioner For Personal Data* [2010] 6 HKC 266

called under the disguise of representing the Bank, but notwithstanding this, the Commissioner found that promotion by a different service provider was in fact a new purpose.

The Jenga Conundrum

The legal parameters of consent and 'new purposes' will become prohibitive to the development of AI, due to the nature of machine learning from neural networks. Assume that a data user, e.g. a creator of an AI system, collects personal data to build a model that predicts changes in artistic styles between a community of artists. This AI system collects the images of an artist's creation and collates it with personality traits, psychological behaviours, and choices of aesthetics. Through deep learning, it continuously builds on every artists' collated data to develop its algorithm. Suppose that this AI system can now predict what potential creations would look like from each artist, what would happen if one or more data subjects withdraw consent for the use of their personal data?

It has been argued that while the efforts of the machine learning prior to the withdrawal of consent will still be valid, any subsequent learning would need to halted. This is due to the fact that the matrix of information has become so intertwined within the neural network of the AI, that a very substantial amount of computing will invariably have 'processed' the data set that is now non-consensual.[15]

This is detrimental to any further meaningful developments of AI, as machine learning cannot be facilitated without re-computing the AI system again with the non-consensual dataset removed. Furthermore, it has been noted that in some circumstances, a deletion of particular datasets may cause damage to the AI system's integrity.[16] There is an immense risk of a *Jenga*-esque collapse to the entire system.

Going back to the aforementioned AI artist-prediction system, what this may mean is that it's future predictions will be impaired, and no future refinements can be done. Another example that was raised by the Centre for Data Innovation (US) is how AI systems used for evaluating credit risks may be undermined by partial removal of personal datasets, with real world ramifications such as customers having unfair credit ratings assigned to them by the AI.[17]

On a similar vein, the 'new purpose' formulation in the PDPO poses a significant threat to automation of existing services with the aid of AI.

The present parameters of a 'new purposes' create a large barrier against the innovation of AI as new consent will be needed for any usage that differs from the original purpose. This means that 'companies cannot find serendipitous uses for data, *even when there are no privacy implications in doing so*.'[18] As a corollary, the cost and efficiency of developing AI systems will be significantly

[15] Matthew Humerick, "Taking AI Personally: How The E.U. Must Learn to Balance The Interests Of Personal Data Privacy & Artificial Intelligence." *Santa Clara High Technology Law Journal* 34, no. 4 (2018): 393-418 at 407
[16] *Ibid* at 408
[17] Nick Wallace and Daniel Castro, *The Impact of the EU's New Data Protection Regulation on AI,* (Centre for Data Innovation, 2018) at 13
[18] *Ibid* at 14

heightened if an organisation has to implement new systems of AI for data with existing consent – as new consent will need to be sought for the sole purpose of automation.[19]

It is difficult to see/predict whether the Commissioner, under the PDPO, will allow 'automation of data' or 'input of data in to a novel AI system' to be within the *Ng Shek Wai* formulation of 'reasonable expectation' of data use. It is also unclear whether automation of data for the same usage purpose will fall under such reasonable expectations. The formulation of a 'new purpose' under the PDPO may therefore lead to the risk of entire AI automated systems to be in breach of the ordinance.

The Royal Free Hospital ('the Hospital') example[20] from the UK shows how a prohibition on non-consensual repurposing of data may lead to a *Jenga-esque* collapse of an AI system. In 2017, the Hospital used personal data from 1.6 million patient records for a medical AI system (ran by Google's DeepMind) used for the advanced prediction of kidney injuries. The legal purpose for the initial collection of data was for the purpose of *inter alia* 'direct care'. However, the Information Commissioner's Office (UK) found that the input of data into the AI for automation purposes did not satisfy data protection principle one of the Act as it 'significantly differ[ed] from what data subjects might reasonably have expected to happen to their data'.[21] New consent was required to be sought from the 1.6 million patients and led to a *de facto* collapse of the entire system, as high costs and substantial time delays were incurred.

As Hong Kong organisations move towards data automation in different sectors, not least in medicine[22], this is a significant challenge that may prejudice imprudent data users. The present laws must require further facilitation.

These examples of AI systems running the risk of substantial failure in the face of repurposing or consent withdrawals shows that the law leaves room for a *Jenga* conundrum vis-à-vis AI where small obstacles may lead to a catastrophic collapse of the entire system.

II. Regulating Data Anonymisation

At present, the PDPO only extends to govern the use of personal data. As the data that are used to feed into AI's algorithms are often highly personal, AI developers have curated data anonymisation methods to protect users' privacy and to prevent a significant loss in data due to restrictions.

---

[19] *ibid*

[20] See generally, the letter from the Information Commissioner's Office (UK) to Sir David Sloman (Chief Executive, Royal Free NHS Foundation Trust), reference RFA0627721, titled "*provision of patient data to DeepMind".*

[21] *Ibid* at 5

[22] See "HKU Medical AI Laboratory Programme employs AI as new testing ground for more optimal and efficient clinical practice of eye diseases", June 10 2019, accessible at https://www.hku.hk/press/news_detail_19593.html It is also interesting to note the exemption on health related data as provided under s 59 PDPO, and whether such a case of automation akin to the Royal Free Hospital case in Hong Kong would be an 'application of … [data protection] provisions [that] would be likely to cause serious harm to the physical or mental health of the data subject' to merit exemption.

This article argues that current methods of data anonymisation is inadequate in rendering information as impersonal, as the identity of a person could still largely be *indirectly* ascertained. Discord arises between the law and anonymisation methods that proves to be another barrier in creating a symbiotic legal environment for AI development to thrive.

The Law

Personal data, as defined in s 2(1) PDPO, means any data—
'(a) relating directly or indirectly to a living individual;(b) from which it is practicable for the identity of the individual to be directly or indirectly ascertained; and (c) in a form in which access to or processing of the data is practicable.'

The provisions set out under s2(1) are incredibly vague, not least in relation to its minimal guidelines in the matter. This was somewhat improved by consolidations in various Hong Kong courts, such as in *Eastweek Publisher Ltd*[23] where the Court of Appeal held that a photograph of a person is personal data as it 'contains some of the most accurate information of the person.' In the Court's view, as the photograph depicts a data subject and 'satisfies the three requirements under s 2(1)', it is deemed as 'personal data'. In another case before the Privacy Commissioner, internet cookies (i.e. data package containing web-browsing data) were found to be personal data as, inferring from the facts of the case, it could be used to uniquely identify the individual as it contained the complainant's English name.[24]

Conversely, in the case of *Shi Tao*[25], the Administrative Appeals Board held that an IP address in the particular circumstances did not constitute as personal data. From the facts, the IP address only 'disclosed that the email was sent from a computer located at the address of a business entity, and the data and time of the transaction.' As a result, the decision was based on the IP address's inability to uniquely identify the actual person who sent the email from the IP address.

These cases show that the PDPO still relies on individual facets of the case rather than concretely defining the parameters for privacy protection. The bottom line remains to be that if an individual could be uniquely identified from a dataset, such dataset constitutes personal data.

Problems with data anonymisation

Data anonymisation refers to 'personal data being rendered anonymous in such a manner that the data subject is not or no longer identifiable'[26], and is done through de-identification and pseudonymisation.[27] AI complicates this as all its datasets are high-dimensional, increasing the interconnectedness of each data point.

---

[23] *Eastweek Publisher Ltd & Another v Privacy Commissioner for Personal Data* [2000] 1 HKC 692
[24] PCPD Case No.:2006C14, Case Note titled "Employee complained her employer logged in her computer collecting cookies without notifying her".
[25] *Shi Tao v Privacy Commissioner for Personal Data* [2008] 1 HKC 287
[26] Recital 26, GDPR
[27] Yves-Alexandre de Montjoye, Ali Farzanehfar, Julien Hendrickx and Luc Rocher, "Solving Artificial Intelligence's Privacy Problem", *Field Actions Science Reports*, Issue 17 (2017), 80-83

Two questions arise regarding the process of data anonymisation: whether it succeeds in protecting user privacy and if not, how the PDPO could be improved to regulate data collection.

De-identification is the process of separating personally identifiable information and the user's identity. One example is the k-anonymity model, where a dataset is said to be k-anonymous if no combination of user attributes is shared by fewer than *k* individuals.[28] It follows a general principle of generalising data then deleting raw and identifiable bits, ostensibly solving the issue of being able to link the data to a single living individual. An issue with k-anonymity is in instances where data is incredibly high dimensional, there are few points that could be easily isolated. Consequently, it becomes difficult to anonymise the data without an unacceptably high amount of information loss.[29] The concept of unicity was also raised as a risk metric for measuring the re-identifiability of high-dimensional anonymous data. In a study based on mobile phone metadata, it showed that just 4 points are sufficient to uniquely identify 95% of people in a dataset of 1.5 million individuals.[30] This is alarming as it merely takes 4 batches of information regarding where and when a user was to reveal their entire location history.

Pseudonymisation replaces personally identifiable information fields within a database with artificial identifies, or pseudonyms.[31] This process transforms personal data and these are aggregated to support AI machine learning. Latanya Sweeney's involvement with the US Supreme Court reviewing of the case of *IMS Health v Sorrell*[32] showed that "widespread aggregation of medical information threatens individual patient privacy", where there was no independent review system for IMS' de-identification process nor were there checks for vulnerabilities in data protection.[33] These factors combine to enable unauthorised re-identification of individuals via the Mosaic Effect. The Mosaic Effect occurs when a person is indirectly identifiable because information can be combined with other pieces of information, enabling the individual to be distinguished from others.

These offer disheartening examples of data anonymisation failing to protect user privacy that leaves it to be constituted as *personal data* under the PDPO. Present anonymisation methods are plainly insufficient to properly de-identify a person from his personal data. As such, many organisations which are contemplating the use of anonymisation methods for their AI systems may need to widen their collection of consent and detail clearly the potential purposes that the AI system may be used for.

---

[28] L. Sweeney, "k-anonymity: A model for protecting privacy", *International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems,* 10, no. 5 (2002), 557-570.

[29] Charu C Aggarwal, *On k-anonymity and the curse of dimensionality*, (In Proceedings of the 31st international conference on Very large data bases, 2005), pages 901–909, (accessible at https://hal.inria.fr/hal-01635002/document)

[30] Y. A. Montjoye *et al,* "Unique in the crowd: the privacy bounds of human mobility", *Scientific reports,* 3, (2013), at 1376

[31] Jan Lindquist, "Data science under GDPR with pseudonymization in the data pipeline", *Datvia*, (April 17, 2018), accessible at https://www.dativa.com/blogs/data-science-gdpr-pseudonymization-data-pipeline/

[32] 564 U.S. 552 (2011)

[33] Adam Tanner, *"Our Bodies, Our Data: How Companies Make Billions Selling Our Medical Records"* (Beacon, 2018)

Without the requisite consent, use of such personal data could be vulnerable to restrictions that would be highly detrimental to the development of AI machine learning, a facet of the *Jenga* conundrum outlined earlier in the essay.

III. The surveillance conundrum: AI's latent defects

Algorithmic surveillance is 'in literal terms, surveillance that makes use of automatic step-by-step instructions'.[34] In many instances, AI systems are used in algorithmic surveillance to engage masses of data in identifying threats to societies. This ranges from the Hubble deep-space telescope (for incoming planets and deep-space exploration)[35] to the Primsmatica/Cromatica movement-recognition system for the London Underground (for monitoring flow of people and suicide rates near the platforms).[36]

AI's role in algorithmic surveillance is to effectively use machine learning to identify patterns in large data sets to predict activity. An example of this in a law enforcement context is a pilot project at Berlin-Südkreuz, which was launched by the German Government to detect acts of violence and individuals in distress. It aimed to use facial recognition technologies on volunteers. By the end of the pilot project, the Government deemed the project to be a success, with an 80% accuracy of identifying participants, as well as general competency with detecting various criminal scenarios.[37] This article argues that if AI and algorithmic surveillance technologies were to be applied in a criminal setting in Hong Kong, the present legal rules will be overtly prohibitive against any meaningful local development or application. In any event, the latent defects in AI and algorithmic surveillance makes it an unideal technology to be applied locally on a large scale.

The Law

Under the PDPO, s 57 and s 58 exempts the use of personal data by the Government for purposes of *inter alia* security, international relations, and the prevention of crime from applications of DPP 3, 6 and s 18(1)(b) respectively. S 58A exempts the use of personal data from the entire PDPO when such data is contained in protected product or relevant records under the Interception of Communications and Surveillance Ordinance (Cap 589) ('ICSO').

It is at this juncture that we evaluate whether the law of covert surveillance in Hong Kong would allow the application of a system that is akin to the AI and algorithmic surveillance examples given above (in a criminal context).

---

[34] Lucas D Introna and David Wood, "Picturing Algorithmic Surveillance: The Politics of Facial Recognition Systems", *Surveillance & Society,* 2, (2004), 177-198, at 181

[35] *Ibid*

[36] *Ibid*

[37] Anna V Eireiner, "Imminent dystopia? Media coverage of algorithmic surveillance at Berlin-Sudkreuz", *Internet Policy Review,* 9, no. 1 (2020)

Under the ICSO, it is likely that AI and algorithmic surveillance fall within the category of Type 1 surveillances[38]. Where law enforcement agencies desire to conduct Type 1 surveillances, permission before a panel judge would be required.[39] In granting permission for Type 1 surveillances to be conducted, it has to be proved that there is a circumstance in which the same objective cannot be achieved by a Type 2 surveillance (less intrusive surveillance techniques using optical or listening devices), or search warrants and court orders. The rationale behind this is to minimise intrusions to personal privacy.[40]

Further and in addition to the above, s 3 ICSO provides that covert surveillance would only be allowed for preventing or detecting *serious* crime[41] or protecting public security. Moreover, there needs to be a reasonable suspicion that any person under surveillance has been involved in that particular serious crime. Lastly, the considerations of necessity and proportionality has to be taken into account vis-à-vis the intrusiveness of the covert surveillance (e.g. people who will be affected by said surveillance).

AI and algorithmic surveillance and the ICSO regime

It would be certainly surprising if any large scale AI and algorithmic surveillance for purposes of preventing serious crimes will be permitted under the aforementioned ICSO regime.

An AI and algorithmic surveillance system, by nature, would only be meaningful if applied to a very broad data sample. As metaphorically confirmed by the former Director of the US National Security Agency, 'you need the haystack to find the needle'.[42] It is difficult to see how a substantial collection of data, which may very much contain a 'haystack' of people beyond suspicion, can be necessary and proportionate. This is especially so given the broad range of alternative options that law enforcement agencies already have for crime prevention (such as applications for warrants, search orders, and general police powers). Coupled with the immense level of intrusion for our citizen's privacy, and the potentially vast amount of people that does not attract reasonable suspicion, it would be unlikely that such surveillance would be proportionate.

Further, the recent 2019-2020 social unrest, as well as the impending application of the National Security Legislation[43], provides for another interesting hypothesis.

In the Code of Practice ('CoP') pursuant to ICSO, it is said that 'advocacy, protest or dissent … unless likely to be carried on by violent means, is **not** of itself regarded as a threat to public

[38] It is likely that AI and algorithmic surveillances will come under 'data surveillance devices' hence not falling within the purview of Type 2 surveillances which governs listening devices or optical surveillance devices. In any event, persons under large scale of algorithmic surveillances is unlikely to reasonably expect their words or activity to be heard or seen, thus outside of the provisions of Type 2 surveillances under s 2 ICSO.
[39] S 6 ICSO
[40] S 3(c)(ii) ICSO
[41] S 2(1) ICSO defines serious crime to be an offence punishable by a maximum sentence of not less than 7 years' imprisonment
[42] Barton Gellman and Ashkan Soltani, "NSA Collects Millions of E-mail Address Books Globally", *Washington Post,* October 14 2013, accessible at https://www.washingtonpost.com/world/national-security/nsacollects-millions-of-e-mail-address-books-globally/2013/10/14/8e58b5be-34f9-11e3-80c67e6dd8d22d8f_story.html.
[43] "China law requires Hong Kong to enact national security rules as soon as possible", *Reuters,* May 22 2020, accessible at https://www.reuters.com/article/us-china-parliament-hongkong-legislation-idUSKBN22Y0CG

security'. Furthermore, it explicitly says that applications for authorisation *must* comply with the statement made by the Secretary for Security in 2006 that the powers under this law *will not be used for investigation of criminal offences that are yet to be created under Art 23 of the Basic Law.* A point that is particularly thought-provoking is whether the Government will interpret the statement made by the Secretary in 2006 to include the provisions of the incoming National Security Legislation, or cast an exception so as to facilitate surveillance under such legislation.

Pitfalls of AI and algorithmic surveillance in general

In any event, there is a swathe of academic opinion that argues against the general application of AI and algorithmic surveillance. Most saliently, the technology doesn't seem to be ripe for vast application in a way that can ensure foreseeability and reliability. It has been shown that 'an inherent problem … is the base rate fallacy and the high likelihood of "false positives"'.[44] These may, in a criminal setting, lead to overtly unnecessary surveillance on the innocent and further exacerbates the disproportionality of this measure.

By the inherent nature of machine learning, the internalised surveillance algorithmic processes can even be 'opaque … to those who designed it'.[45] This alone caused academics to say that there is an 'unconquerable obstacle to the provision of foreseeability'.[46]

In sum, it is argued that the present laws in Hong Kong would prohibit any meaningful application of AI and algorithmic surveillance in a crime-prevention context. In any event, inherent pitfalls in these systems shows that much more needs to be done before such surveillance methods should be used.

Part C: Symbiosis

In view of the multifaceted discord between the present privacy laws in Hong Kong and the development of AI, this article contends that a symbiotic environment between the law and AI could be achieved if AI-specific provisions can be added to the PDPO.

The PCPD often brands the PDPO as a 'technology-neutral' and 'principle-based' creature inspired by the OECD Privacy Guidelines 1980 and the 1994 Law Reform Commission Report entitled "Reform of the Law Relating to the Protection of Personal Data".[47] The PDPO itself took effect from December 1996.

---

[44] Douwe Korff, "Technologies for the Use of Images: Automated Processes of Identification, Behavioural Analysis and Risk Detection Control at the Airports." Seminar presentation at Spanish Data Protection Agency Seminar, Madrid, June 9–11 2010. *See further* Bruce Schneier, "Terrorists, Data Mining, and the Base Rate Fallacy", *Schneier on Security,* July 10 2006.

[45] Maria H Murphy, "Algorithmic surveillance: the collection conundrum", *International Review of Law, Computers & Technology,* 31, no. 2, (2017), 225-242, at 230

[46] *ibid*

[47] "Ethical Accountability Framework for Hong Kong, China", *Report prepared for the Office of the Privacy Commissioner for Personal Data,* 2018, accessible at https://www.pcpd.org.hk/misc/files/Ethical_Accountability_Framework.pdf, at 1, and 17; *See further* "Reform of the Law Relating to the Protection of Personal Data (Topic 27)", *The Law Reform Commission of Hong Kong,* August 1994.

What is apparent from an examination of these documents is that no consideration was given to novel ideas such as 'big data' and 'artificial intelligence' – for the simple reason that such ideas were not prominent in the 90s. On a similar vein, the experts representing the European Commission[48] described that regimes 'designed with traditional … business models in mind' may be an 'inadequate match' for 'emerging digital technologies'.[49] In other words, the 'technology-neutral' element of the PDPO may, when applied to AI technologies, be analogous to 'fitting a square archaic peg into the hexagonal hole of modernity'.[50]

This article proposes the following additions to the PDPO to facilitate a symbiotic legal landscape to the development of AI:

Attribution of liability

Under the PDPO, a data user is a 'person' who controls the collection, holding, processing or use of the data.[51] At present, algorithms and AI have not been given legal personality under Hong Kong laws. Therefore, when amending the PDPO to facilitate AI development, clear parameters have to be set to identify whom liability is to be attributed to.

As laid down in the landmark Singaporean decision of *B2C2*[52], common law is chiefly concerned with two strands of automated data processing insofar as liability is concerned. First, deterministic automation, i.e. where the automation 'produces the exact same output when provided with the same input; and secondly, autonomous artificial intelligence, which could 'be said to have a mind of its own'.[53]

Case law has shown that where automation is deterministic, liability is attributed to the programmer of the system. Deterministic automation, as said by the Singapore Court, is like a kitchen blender – the output is the always same when provided with the same input. It has no mind of its own. In so suggesting, it held that where deterministic systems enter into automated contracts, it is the state of mind of the programmer which will be under scrutiny when evaluating liability.

Similarly, in *Yeung Sau Shing v Google*[54], Deputy Judge Marlene Ng held that there is a good arguable case that Google will be liable, as a publisher, for AutoComplete results generated by an automated algorithmic process that collects data 'precisely as [the programmers] intended'.[55] The lack of human input in the collection process of data did not affect the fact that the system performed as intended by the programmers.

---

[48] These experts refer to the New Technologies Formation of the Expert Group on Liability and New Technologies, set up by the European Commission and convened in June 2018

[49] "Liability for Artificial Intelligence and other emerging digital technologies", *Report from the Expert Group on Liability and New Technologies – New Technologies Formation,* (2019), catalogue number: DS-03-19-853-EN-N, at 28

[50] *Crookes v Newton* [2011] 3 SCR 269 at 110 (decision from the Supreme Court of Canada)

[51] S 2 PDPO

[52] *B2C2 Ltd v Quoine Pte Ltd* [2019] SGHC(I) 3

[53] *Ibid at* [206]

[54] *Dr Yeung Sau Shing Albert v Google Inc* [2014] HKCFI 404

[55] *Ibid* at [98]

Privacy legislation in Hong Kong, therefore, should follow established case law to state clearly that where there is deterministic automation (be it 'AI' or not), programmers should be liable for the system's breach of the PDPO if the system had acted in a way that was intended by the programmers.

As for an autonomous AI system that 'has a mind of its own', the attribution of liability may be inherently difficult, as some algorithmic processes can even be 'opaque … to those who designed it'.[56] This presents problems in ascertaining the causal nexus between a mistake generated by the autonomous system, and the fault of the programmer.

Experts from the European Commission have argued that the prudent measure would be to impose strict liability on programmers for liabilities incurred by their autonomous AI creations.[57] This article wholly agrees insofar as liabilities under the PDPO is concerned. Not only would this increase awareness and due diligence by programmers, this would guarantee that victims of the breach of privacy would be entitled redress.

In short, it is proposed that the PDPO include provisions which attribute liability incurred by deterministic and autonomous systems under the PDPO to programmers. For the former, the intention of the programmer would have to be ascertained, for the latter, strict liability should be imposed.

Anonymisation

Understanding the pitfalls of previous anonymisation efforts would enable the PDPO to effectively target stakeholders and uphold data privacy protection. Whilst stakeholders could take the form of AI algorithms, control mechanisms could be put in place to ensure no single individual can be identified.

Mechanisms could be added as an extension to DPP 6 PDPO regarding access to personal data. Access control could be tightened through privacy-enhancing technologies (PET) that allow datasets to be used a privacy conscientious way.[58] Taking a step further from listing what the data subject is entitled to doing, it could create an audit system to ensure that any interaction with the data is recorded and accessible, hopefully mitigating the risk of exploitation of personal data.

A successful example of access control coupled with an audit system is Google's DeepMind, which trained machine learning algorithms on individual-level health data records from the NHS.[59] Their 'Verifiable Data Audit' increases accountability on all levels from whoever has access to the dataset.

Implementing a tailored version of the right to be forgotten from the GDPR to the PDPO by focusing on anonymisation efforts rather than outright erasure, it could minimise the risk of large

[56] *See n. 45*
[57] *See n. 49* at 39
[58] *See n. 27*
[59] Suleyman, M., Laurie, B, "Trust, confidence and Verifiable Data Audit, DeepMind Blog", accessible at https://deepmind.com/blong/trust-confidence-verifiable-data-audit

data loss. This removes the need for data removal and allows for abundant neural network circulation in AI to aid its machine learning. As AI thrives from being able to access large amounts of data, the quality of AI would not be sacrificed as a result.

In staying aligned with the rapid pace of data growth, the PDPO could prevent future breaches of privacy instead of passively waiting for the next data collapse.

Additional exemptions and regimes for automated systems

Part B of this article saw an exposure of a *Jenga* conundrum under the present formulation of a 'new purpose' for the use of data under the PDPO.

To remedy this lacuna, it is proposed that an additional exemption is provided under Part 8 (Exemptions) for AI systems. This exemption should clearly provide that where there is consent for an existing purpose in using personal data, the automation of such existing purpose should be exempt from the provisions under DPP 3.

This will create an immensely symbiotic environment for organisations to structure their datasets in an automated way (often done by AI as seen from examples above) and allow the development of such systems to flourish.

Furthermore, an additional regime can be provided under the ICSO for potential applications for AI and algorithmic surveillance systems in the future. At present, the ordinance is mainly aimed at regulating authorisation for surveillance on persons who have already attracted reasonable suspicion[60].

This is different from the main aims of algorithmic surveillance, which seeks to find out who law enforcement agencies should be suspicious about in the first place. In other words, the algorithmic surveillance should be a treated as a separate type of surveillance compared to the regimes under the ICSO. Hence, legislation should aim to regulate algorithmic surveillance with separate safeguards to ensure that the right of privacy as guaranteed under the Basic Law and the Bill of Rights Ordinance.

Conclusion

Societal improvement is buttressed in staunch belief of a legal system. As AI continues to strengthen its neural networks and ability to learn, likewise, we must strive to enhance our adaptability through the PDPO.

Whilst discord may arise between PDPO's prohibitive parameters and the development of AI, where tightened regulations may cause AI to crumble under the *Jenga* conundrum, hope is not lost. There is room for the law to be changed with regards to progressions in AI advancements, this article has proved that in engaging headfirst with AI's intricacies, the law could foster a symbiotic environment for AI to thrive in Hong Kong. Some AI technologies may require alterations in terms of its use in algorithmic surveillances and anonymisation methods, but this should not hinder the

[60] *See* s 3(b) ICSO

PDPO as it could improve alongside AI with proposals outlined in Part C. Enabling the PDPO to be on the forefront in this race to maximising AI's utility would render what was once a discordant environment into one that is appropriately symbiotic.