How best to regulate AI

As AI can refer to many types of rapidly evolving technologies, it's better to focus on regulating specific applications of it

Alexander Woon

For The Straits Times

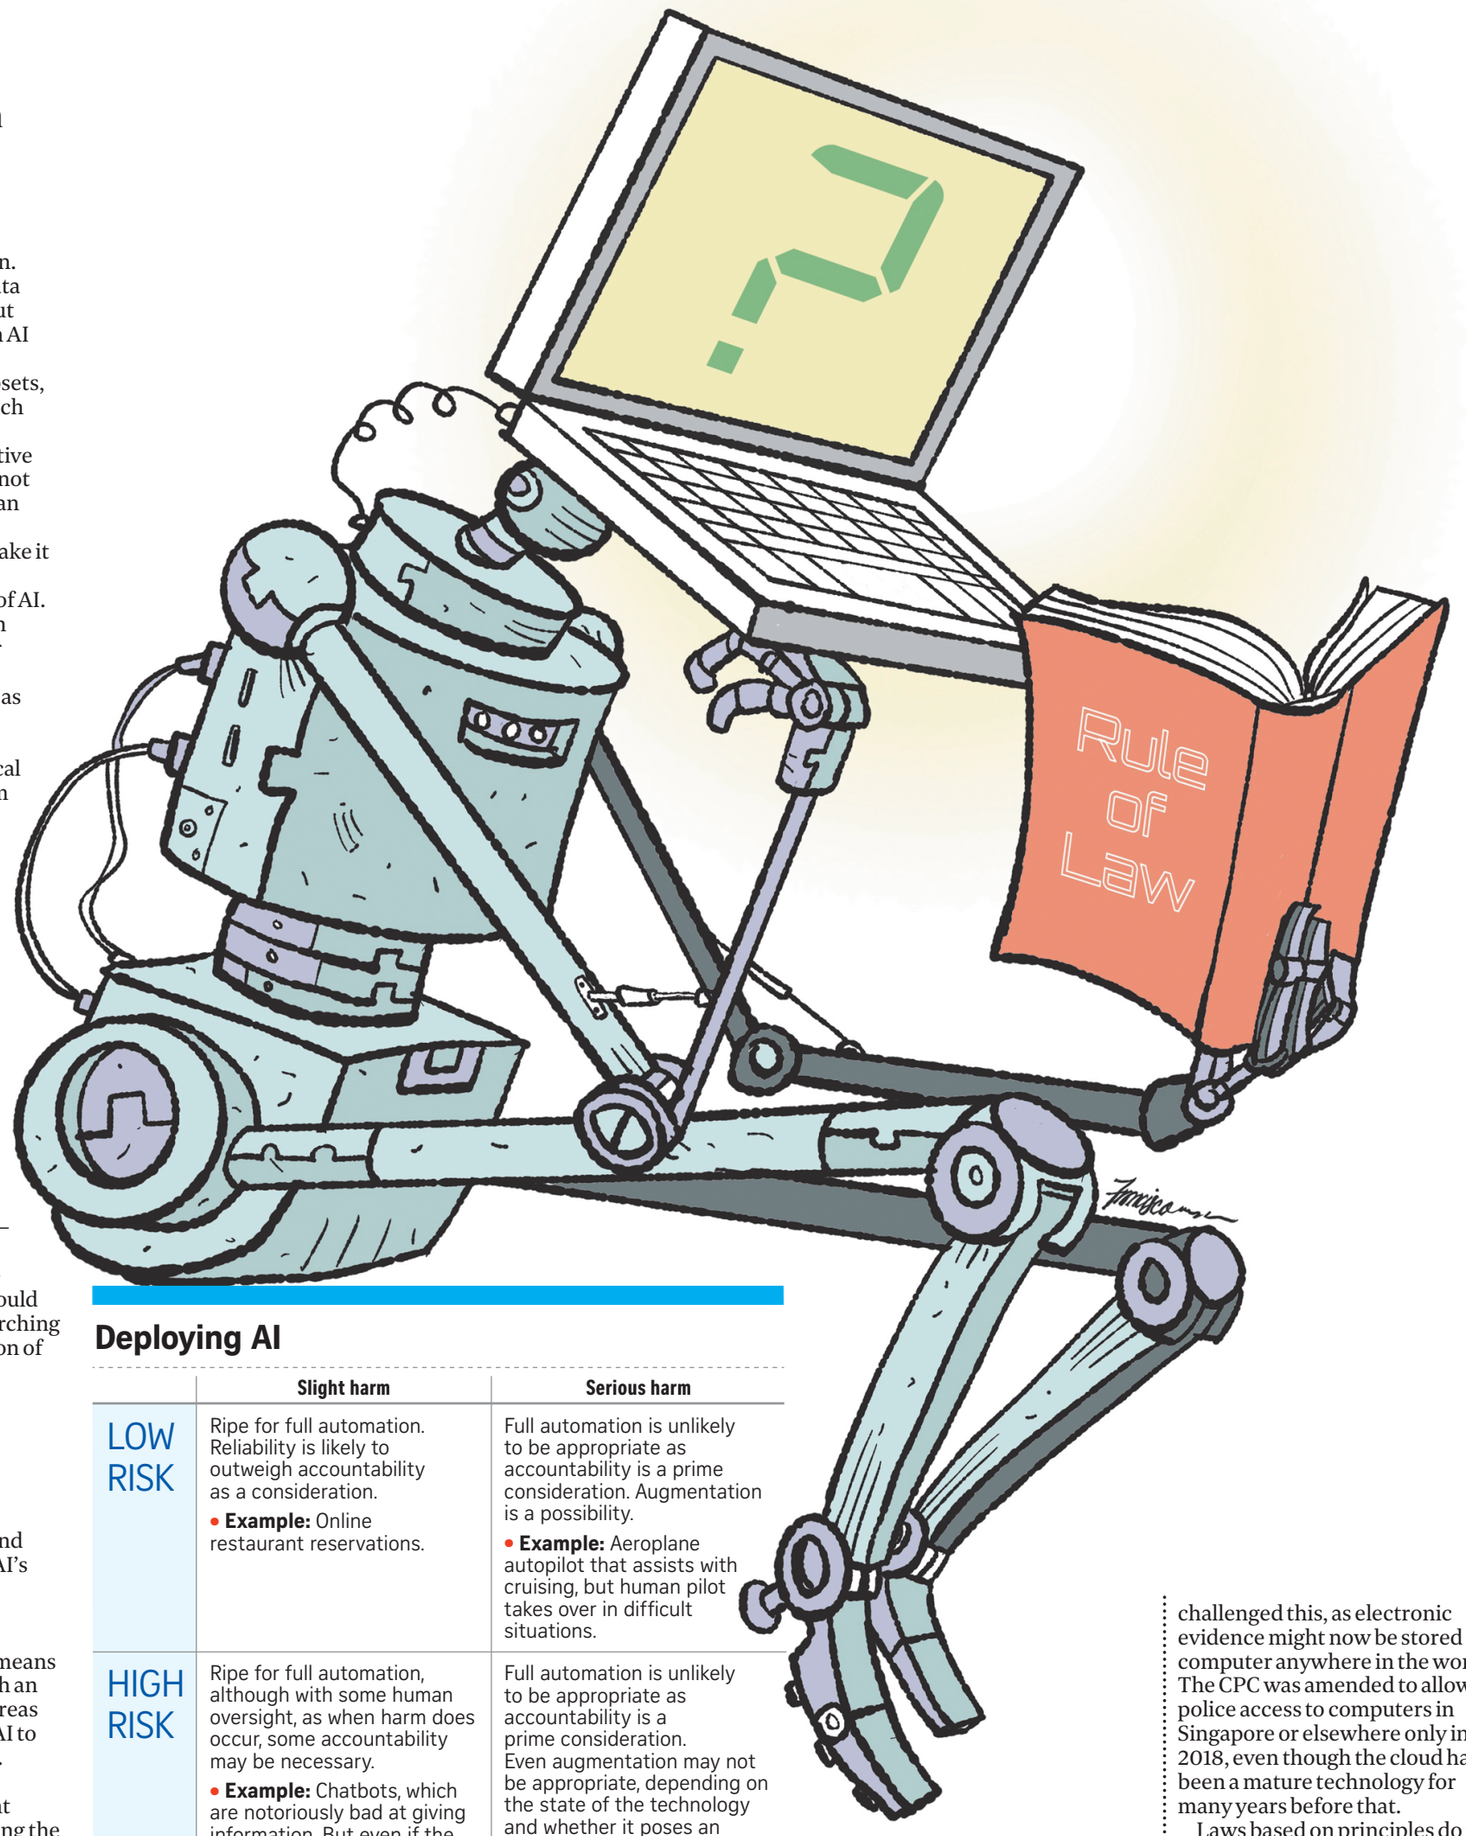

Artificial intelligence (AI) is everywhere now and there are growing calls for its regulation. There is, for example, the debate over whether AI should have its own legal personality and be subject to criminal law and punishment.

But AI is currently very far from reaching sapience, the ability to think and reflect in the manner of a person. Killer robots like HAL 9000 from Arthur C. Clarke's novel 2001: A Space Odyssey remain strictly science fiction.

Some over-enthusiastic proposals are therefore merely a manifestation of the Frankenstein complex, a term coined by science-fiction author Isaac Asimov to describe humanity's instinctive fear of being overtaken by its own creations.

Attempting to regulate AI in the same way that we regulate human behaviour is doomed to failure because AI is not "intelligent" in the way that humans are.

To paraphrase Professor Pedro Domingos of the University of Washington, the problem is not that AI is too smart and is going to take over the world; the problem is that it is too dumb and it already has.

It is clear that AI does need to be regulated. The question is how to do this effectively.

WHY AN OMNIBUS LAW WON'T WORK

It may not be possible to regulate AI in general. It is not possible to comprehensively map all the uses of AI, especially as the technology is constantly and rapidly evolving.

An omnibus AI regulation framework might miss out on certain niche use cases and be unduly restrictive, imposing a one-size-fits-all policy on a technology that has a myriad of different applications.

Instead, we might focus on regulating specific applications of the technology. To give an analogy, we do not have an omnibus "Internet law". The Internet's reach is too pervasive and its uses too multifarious to encapsulate in a single regulatory framework. For instance, online scamming is dealt with by the Penal Code; electronic commerce by the Electronic Transactions Act; Internet service providers are regulated under the Telecommunications Act and the Broadcasting Act, and so on.

A similar model could be applied to the regulation of AI, for three main reasons:

- There is no clear definition of AI.
- Appropriate regulation depends on the nature of the precise use case.
- The law should, as far as possible, be technology neutral.

GENERAL AND NARROW A.I.

There is no clear definition of AI as it can refer to many types of technologies, and there is no global consensus on which technologies constitute proper AI.

First, we must distinguish between general and narrow AI. General AI is the stuff of science fiction: an AI capable of fully autonomous functioning that is able to perform a wide variety of tasks – like Data from Star Trek.

There is no need to devote too much attention to general AI as it does not currently exist and it is unclear whether it is even possible for it to exist.

What we need to deal with urgently is narrow AI, designed to accomplish very specific tasks, such as picking stocks or hiring candidates for jobs.

Second, there are many types of narrow AI, and they do not all work the same way. Principally, we need to distinguish between rules-based and machine learning systems.

In simple terms, rules-based AI is just a complicated collection of "if-then" statements. It has no ability to learn or change, merely a complicated logical decision tree that simulates intelligence. If something happens outside the scope of its programmed rules, the AI is unable to adapt.

In contrast, machine learning systems are able, to varying degrees, to learn on their own. Such AI is trained on large data sets. It is fundamentally about pattern recognition. Modern AI research tends to focus on machine learning and its subsets, including deep learning, which relies on multiple layers of algorithms to create an iterative learning process, one that is not always explainable to a human being.

Third, these differences make it difficult to formulate an appropriate legal definition of AI.

For example, the European Commission earlier this year published a proposed AI regulation, which defines AI as including not only machine learning and rules-based approaches, but also statistical approaches. An AI is a system developed using one of the techniques specified above which can, with a human-defined objective, produce certain output.

In contrast, Singapore's Personal Data Protection Commission's Model AI Governance Framework defines AI as a set of technologies that seek to simulate human traits, such as reasoning and problem solving. Production of an output is optional.

Perhaps it is time to simply cut the Gordian knot – pursuing a usage-focused regulation regime avoids the problem entirely, as there would then be no need for an overarching definition of AI. The definition of AI could be scoped for each specific use case.

APPROPRIATE REGULATION DEPENDS ON USE CASE

The nature and extent of appropriate regulation depend heavily on the nature of the AI's use case.

First, we must distinguish between automation and augmentation. Automating means replacing a human being with an autonomous AI system, whereas augmentation means using AI to enhance human capabilities.

For example, there are AI systems designed to augment human drivers, like those using the Global Positioning System that suggest routes to the driver, and systems designed to automate driving entirely, like Tesla's Autopilot.

Implementation of AI is not a binary yes/no decision. In law, for example, there is an instinctive distrust of allowing AI to make judicial decisions. But AI need not replace human judges completely – judges could be assisted by AI that helps to collate case law or extrapolate sentencing trends, for example.

The AI would merely be a research tool, much like current legal database search tools, while the final decision would remain with a human judge.

Second, when regulating a specific application of AI, we should ask whether the concern is reliability or accountability. Reliability means the rate at which a correct decision is made. Accountability means the ability to take responsibility when a wrong decision is made.

AI systems are obviously incapable of accountability as they are not legal or moral agents. But they may very well be more reliable than humans for certain tasks. With self-driving cars, there is some evidence that AI systems are actually safer drivers than humans, but if something goes wrong, then who is to be held accountable?

A balance needs to be found between reliability and accountability, but exactly where that balance is struck will depend on the precise nature of the use case.

Third, appropriate regulation of AI is a function of risk management. It is a matter of policy as to when and to what extent AI should be deployed. This may be illustrated by a matrix comprising two axes: the risk of harm occurring and the severity of the potential harm. The matrix is a heuristic that allows for general classification and application of a broad approach to regulation (see table).

Deploying AI

	Slight harm	Serious harm
LOW RISK	Ripe for full automation. Reliability is likely to outweigh accountability as a consideration. • **Example:** Online restaurant reservations.	Full automation is unlikely to be appropriate as accountability is a prime consideration. Augmentation is a possibility. • **Example:** Aeroplane autopilot that assists with cruising, but human pilot takes over in difficult situations.
HIGH RISK	Ripe for full automation, although with some human oversight, as when harm does occur, some accountability may be necessary. • **Example:** Chatbots, which are notoriously bad at giving information. But even if the information given is incorrect, the harm to the user is very slight.	Full automation is unlikely to be appropriate as accountability is a prime consideration. Even augmentation may not be appropriate, depending on the state of the technology and whether it poses an acceptable risk. • **Example:** Complex litigation still needs to be done by trained lawyers. Even output generated by legal technology may need to be double-checked manually, as any mistake could result in severe consequences for the client.

Source: ALEXANDER WOON STRAITS TIMES GRAPHICS

Fourth, AI does not operate in isolation. Machine learning AI, for example, relies on big data, and therefore any regulation of AI alone would be incomplete if it does not also consider data collection and usage.

To give another example, regulation of self-driving cars must take into account the regulation of human drivers. Many accidents involving self-driving cars are the result of a human counterparty's error. It might make more sense to view this as an issue of traffic regulation and not an issue of AI regulation. The solution may very well be to regulate the human drivers and not the AI.

LAW SHOULD BE TECHNOLOGY NEUTRAL

Finally, the law should strive to be technology neutral. Technology evolves much faster than the law. If laws are drafted to be very specific about the technology used, they will become obsolete when the technology inevitably changes.

For example, Singapore's Criminal Procedure Code (CPC) was previously drafted with a server-based paradigm in mind, in which police officers could access computers suspected in connection with an offence only if they were located in Singapore.

The advent of cloud computing challenged this, as electronic evidence might now be stored in a computer anywhere in the world. The CPC was amended to allow for police access to computers in Singapore or elsewhere only in 2018, even though the cloud had been a mature technology for many years before that.

Laws based on principles do not require similar amendment every time the technology changes. For example, the definition of the offence of cheating in the Penal Code has remained remarkably stable since 1987. This is because it defines cheating broadly as deception causing some form of damage to the victim, without the necessity for specifying any particular means or method for carrying out the cheating. This makes it robust to technological change.

Similarly, when it comes to regulation of AI, perhaps we should simply think about how to make fundamental principles of law technology neutral so that they can be applied to any new technology.

For example, when it comes to AI and the law of armed conflict, debates about whether autonomous weapons should be allowed might be reducible to questions about the existing foundational principles of distinction and proportionality – that is, whether autonomous weapons are as capable as human operators of adequately picking only military targets, and attacking them with such force that prevents unacceptable collateral damage.

AI needs to be regulated, but regulation should not necessarily be about AI. It is not so much the technology that matters but its use and potential impact on society. A focus on fundamental principles of justice – what is right or wrong, harmful and helpful – may be a better way to address these issues than to get fixated on the technology itself.

stopinion@sph.com.sg

• Alexander Woon is a lecturer at the School of Law, Singapore University of Social Sciences, and Of Counsel with RHTLaw Asia LLP.

OPINION

What Spider-Man looks like to AI

Over the last few months, mobile apps featuring art generated by artificial intelligence (AI) have been lighting up social media.

The latest app is Wombo Dream (left), which allows people to key in words to prompt the AI to generate a picture. The Straits Times looks at how the Wombo Dream AI visualises Spider-Man, Squid Game and Pokemon, and explores the recent history of other text-to-image generators.

There are other applications of AI beyond the realm of art, including simulation of tennis matches, or in fields such as customer service and urban planning.

This raises issues of economics and ethics, such as: How can governments manage potential job losses arising from AI and automation? When an algorithm is used to create a product, who owns and profits from it – the person who wrote the algorithm, or the person who put it to use?

str.sg/wTSp