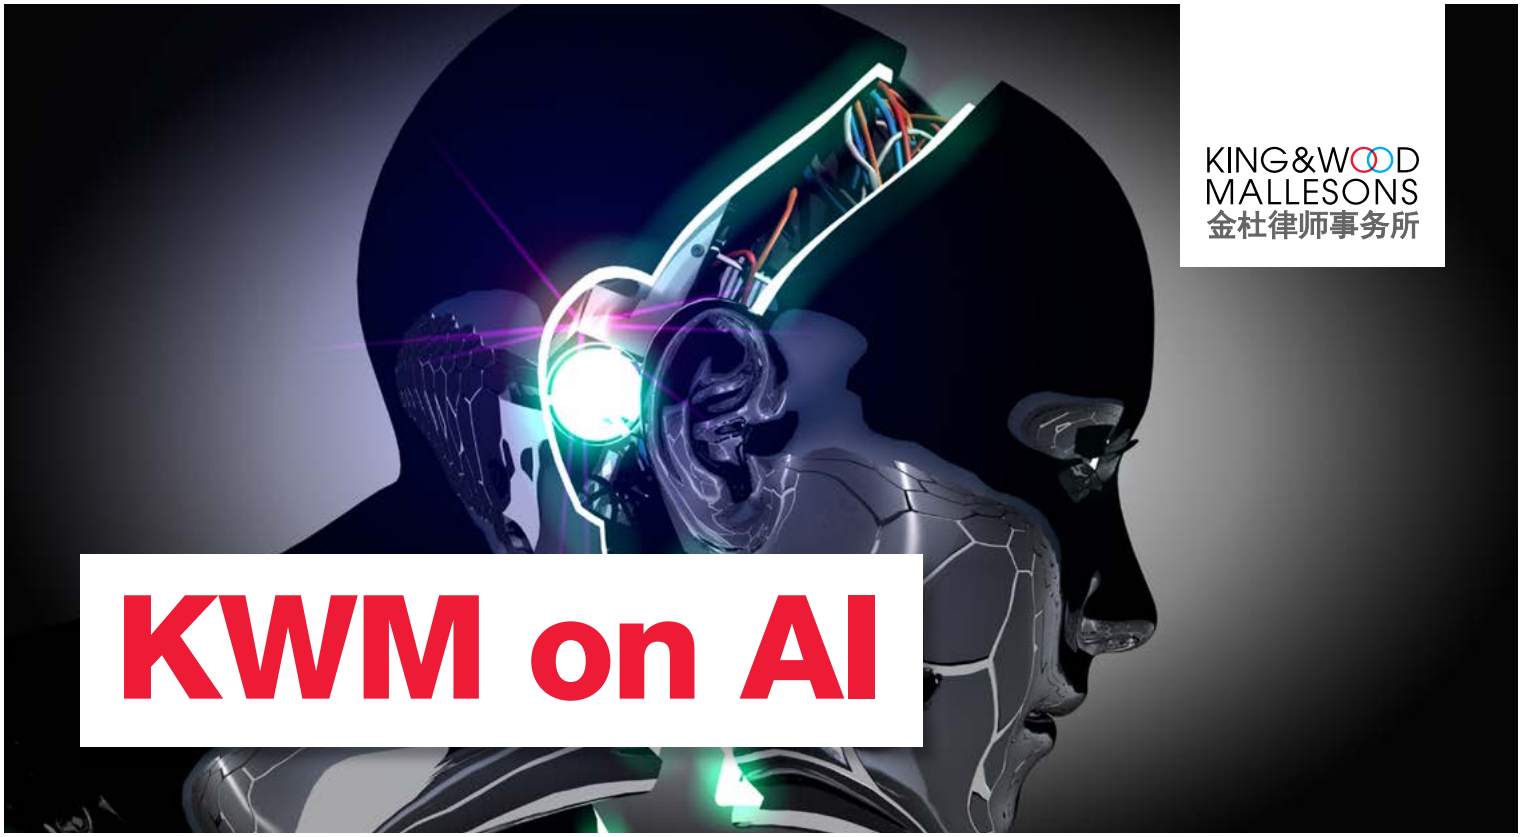

KING&WOOD MALLESONS
金杜律师事务所

KWM on AI

Artificial Intelligence or "AI" is the next great technology trend that promises to transform the way we do business and, indeed, function as a society. Whereas today we humans still tell machines what to do, the growing use of AI is ushering us towards a future in which machines figure things out for themselves. Clearly this will have significant legal, as well as societal, implications.

In broad terms, we use "AI" to refer to technology that can do things that would usually require a human being. This does not necessarily mean that AI technologies will have human-like intelligence, but rather that they can replicate the results of human-like intelligence. Indeed, the first phase of AI technology will essentially be advanced data processing based on "machine learning" whereby computers will be trained to pick up patterns and correlations from a given data set in a way that allows them to replicate human decision-making, even though they are not using the same type of thought processes or reasoning that a human would. In this publication we look at the key building blocks of a machine learning system and identify some of the unique **legal** and **ethical** issues that they present.

Click on sections of the image to find out more. →

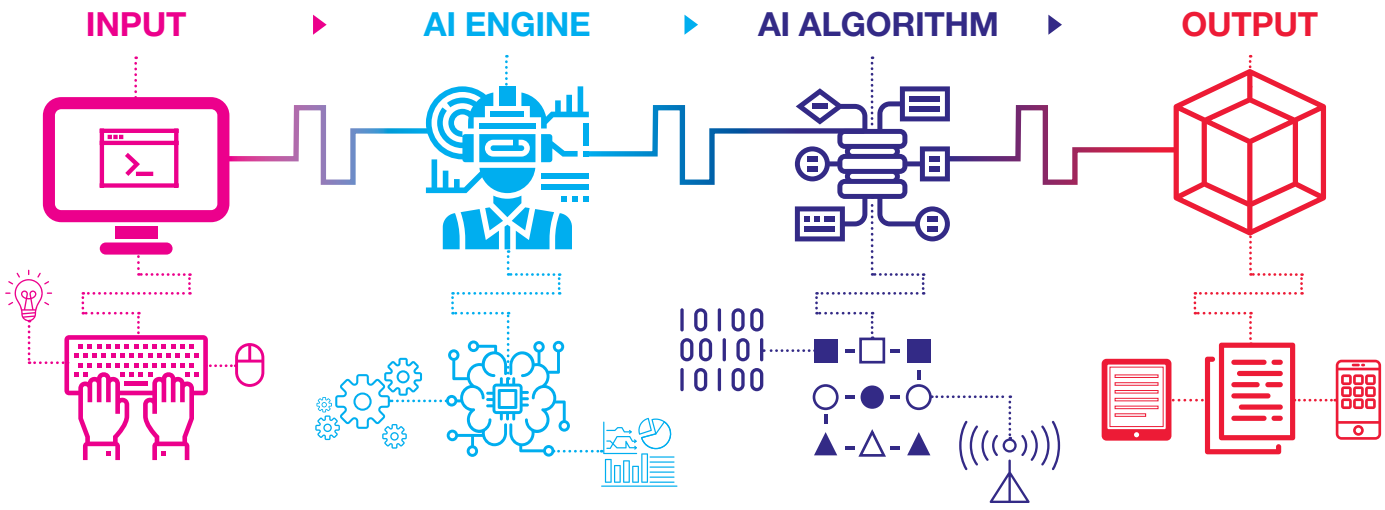

INPUT

What is it?

The input is the material that you will "feed" your AI system with. It is data that will both train the system and also enable the system to produce the output you are after. Your inputs may take many different forms, but remember that the old maxim "**rubbish in, rubbish out**" still applies in the world of AI. The more accurate, comprehensive and reliable your input data, the more accurate, comprehensive and reliable the work product produced by your AI will be. Accordingly, undertaking an initial "data cleansing" stage before feeding into the AI system can be a good investment for ensuring the overall success of your AI undertaking.

What legal issues does it present?

The most critical legal issue to consider in relation to input materials is whether you actually have adequate rights to use those materials for your AI process. This will depend on a number of factors. To the extent the material is capable of sustaining **copyright** (i.e. it is not simply raw data) it will be important to consider whether or not you own the copyright or, if the materials have been sourced from a third party, whether you have a licence from the copyright owner that permits use of those materials as part of the AI process. Apart from that, you will also need to consider whether any of the input materials may contain **confidential information** belonging to another party and, if so, whether the conditions under which you have obtained access to that information permit you to use and potentially disclose it for the purposes of the AI process. This may depend on a range of things, including whether you are using the AI for internal or external purposes, and whether or not your use of the AI will require some collaboration with a third party that will also require them to have access to the input materials. In addition to general duties of confidence, to the extent that the input materials may contain personal information (i.e. information that relates to an identified or identifiable individual) then you will need to consider potential **privacy implications**, including whether there is a legal basis for the use and disclosure of that information under applicable privacy laws. This may in turn require careful consideration of what you have told the individuals in question about your potential use of their information – as such, it may be appropriate to update your privacy collection notices and policies before embarking on any widespread use of AI technologies for your business.

What ethical issues does it present?

As flagged above, the quality of the AI output will be highly dependent upon the quality of your input materials. Any flaws within the input materials may be perpetuated or even magnified through the AI process. In particular, any **intrinsic bias or prejudice** that exists in source data may be reflected through the AI process and result in further bias or prejudice in the final outputs. This has been seen, for example, in the use of AI for recruitment and in law enforcement, where profiling using AI engines trained on historical data can lead to unfair and unjust outcomes (perpetuating historical biases against minority groups that are embedded in the data). Methods of gathering source data need to be carefully scrutinised in order to ensure that they do not undermine the work of the AI process and maintain the same or better level of quality than human decision makers could achieve.

AI ENGINE

What is it?

Unsurprisingly, given the name, the AI engine is the driving force behind your AI system and provides the data processing power that is required to analyse your input materials and produce a useful output. The engine will invariably consist of a software application running on a host platform and, unless your organisation has the capability to develop, support and run advanced software of its own, it's likely that you will be using a third party to provide one or more of these elements. Increasingly, this means that AI engines will be made available on an "as a service" basis. A range of machine learning platforms – such as Amazon Sagemaker, Azure Machine Learning and Google's Cloud Auto ML – are already available from the world's leading cloud technology providers, in order to help businesses develop, train and deploy machine learning solutions at scale.

What legal issues does it present?

Any third party that provides software or other system elements for your AI engine will naturally expect to retain ownership of the intellectual property in those elements. That is fair enough. However, you will want to take to **protect your interest in both the inputs fed into and outputs produced by** the system. In particular, you will want to be sure that you retain ownership and control over any input materials and that the third party is obliged to keep those inputs confidential and properly protected against unauthorised access or interference. You will also want to be sure that your contract with the third party system provider clearly transfers ownership of any intellectual property rights in the output produced by the AI engine to you and prohibits the system provider from further sharing or repurposing those outputs without your consent. When using any third party AI system, you should also think about **business continuity risks**. For example, is it clear what options you will have if your relationship with the third party runs into difficulty – will you be able to recover your data and easily transfer to an alternative AI engine without experiencing any material operational disruption?

You may want to test the third party's disaster recovery and business continuity procedures to ensure they're up to scratch, and think about whether there will be any formatting or other issues that may impede the transfer of your data to an alternative platform.

What ethical issues does it present?

Your selection of the AI engine may, of itself, not generate any specific ethical issues. However, ethical concerns should still be front of mind whenever you are implementing an AI system, so that appropriate ethical safeguards can be built into the system from the ground upwards following an "ethical by design" approach. While government and industry are still getting to grips with this area, one potential useful starting point is the **voluntary set of eight AI ethics principles** that the Australian government – in collaboration with business, academia and the community – has recently developed to encourage organisations using AI systems to deliver the best outcomes for Australia. The principles in this framework address: human, social and environmental wellbeing; human centred values; fairness; privacy protection and security, reliability and safety; transparency and explainability; contestability; and accountability. When selecting a new AI engine, it will be useful to consider how well it is adapted to address each of these principles as part of a process that has a sound ethical as well as legal and operational grounding. For more on this ethical framework, read: https://www.industry.gov.au/data-and-publications/building-australias-artificial-intelligence-capability/ai-ethics-framework/ai-ethics-principles

AI ALGORITHM

What is it?

The AI engine may provide the data processing muscle for your AI system, but the brains will lie in the algorithm that tells the engine what to do. In a machine learning system, the algorithm may be preprogrammed or else may be developed or improved by the engine itself, from a combination of processing the data you feed into it and the feedback you provide on the outputs that are produced. Whatever its origins, the algorithm is the formula or the "secret sauce" that allows the AI system to produce an outcome that is tailored to your unique needs.

What legal issues does it present?

The key question in assessing legal issues associated with an AI algorithm is whether or not it is a **unique algorithm** generated by reference to your own unique circumstances (in which case you will have a stronger case to argue that you should control it), or whether it is a **generic algorithm** that has been produced by an AI technology provider and applied for many different customers in a similar position to you (in which case there will be a stronger case for the technology provider to retain control). More complex issues may arise where the AI algorithm learns from or is improved by the processing of your data inputs (or even develops in a way that incorporates some of that data, such as where inputs are used as exemplars or comparators within the algorithm) – in that scenario, even if the original algorithm was brought to the table by the technology provider, you may justifiably expect a degree of exclusivity (either temporary or permanent) to protect the added value generated from your data. In either case, an algorithm or formula of itself is unlikely to be something that can be protected by recognised system of proprietary rights such as copyright or patent. The best protection for this element of the AI system will likely come through the law of **confidentiality** or **trade secrets**. Accordingly, it will be essential to ensure that any contract with a third party provider of AI technology addresses these issues and makes clear what rights each party will have to access and reuse the algorithms that guide the AI engine. From a business continuity standpoint, it will also be important for you to be clear on any impediments to transferring the logic behind the relevant algorithm to a different AI platform if need be – that is, is it possible for you to neatly extract the logic behind the algorithm so that it can be transferred to a different platform? If not, then if you do at some stage need to change platforms you may need to effectively retrain your AI solution.

What ethical issues does it present?

The most critical ethical consideration in relation to the AI algorithm is the need to ensure that AI outputs can be adequately **explained** if necessary. In particular, if you use an AI system to make a decision that affects an individual data subject then you should consider how you would justify that decision if queried by the individual in question. In some cases you may be subject to a legal duty as well as an ethical duty to do this (e.g. the EU General Data Protection Regulation provides rights for data subjects to object to and contest automated decision making processes, including profiling). This may present some significant challenges where the AI algorithm is self-developed by the AI engine, rather than by a human author, and may potentially be beyond human understanding. In these circumstances, it will at least be necessary to identify the relevant data inputs and to be able to explain how those inputs are processed to generate an outcome by the AI engine. At a basic level, explainability is important to ensure that AI systems **do not perpetuate biases or prejudices** that would not be considered acceptable for a human decision-maker. For example, if your HR department uses an AI algorithm to filter candidate applications, then clearly it will be unacceptable for the algorithm to discriminate between applications simply based on gender or age. Explainability is also important for building trust in AI systems, particularly where they are to be deployed in business or safetycritical systems where stakeholders may be reluctant to give up control to technology that they are not able to understand.

OUTPUT

What is it?

The output is the result produced by the AI system – it is the end game, the reason why you are engaging the AI in the first place. It will likely take the form of some data, whether that be in the form of a decision, prediction, or answer to a question that you have presented to the AI system. The form of the output may depend on the way in which you have configured the AI system, but the output itself will be something that is purely the creation of the AI itself.

What legal issues does it present?

Depending on the form, the output produced by an AI system may potentially be capable of sustaining **copyright**. While it is currently common for copyright works to be produced using software tools and other computer technology, human operators are generally involved in some part of the operation such that it is still possible to identify a human author for the purposes of establishing the subsistence of copyright. With higher levels of automation, as humans are progressively removed from the creation process, doubts may arise as to whether outputs of advanced AI systems will be protected by copyright as they will be too far divorced from any human author. In some cases, courts have already found that copyright does not subsist in certain computer-generated works because of the **absence of a human author**. A number of jurisdictions, such as NZ and the UK, are responding by changing their copyright laws to deal with this scenario and provide that the owner of the AI system will own the copyright in the work created by that system. For these reasons, even if there is some residual doubt as to what proprietary rights will be created, it will be important to ensure that appropriate copyright assignment provisions are present in all of your AI contracts to ensure ownership of any rights in the system outputs is passed to you. In addition, in case no such rights subsist, it will be critical to also ensure that your AI contracts also include robust confi-dentiality provisions that make clear all **outputs remain confidential** to you and may not be further used or disclosed without your permission. This will be particularly important where the outputs produced simply are not capable of sustaining copyright, which may well be the case where the outputs are simply comprised of raw data. Some middle ground may exist where a third party AI provider wishes to use aggregated or anonymised output data to form part of a broader research project. However, before contemplating this, you should carefully consider the risks of the information being re-identified and traced back to your organisation or to any individual data subjects

What ethical issues does it present?

As flagged above, the output produced by an AI system may in some cases be difficult for a human being to fully comprehend. Even where the AI algorithm is known, the output produced by the algorithm may appear wrong or flawed from a human standpoint, and can quickly become corrupted by being stoked with inappropriate input materials. For example, a few years ago a chatbot launched by Microsoft on Twitter to engage in conversation with other users was quickly subverted and used to propagate hateful and abusive comments. This is a reflection of the fact that advanced machine learning is still not capable of perfectly replicating human thinking and judgement. For this reason, it will be critical to ensure that AI outputs are **routinely reviewed and tested** by human decision-makers in order to ensure the results are sound. In addition, it will be important to ensure that AI decisions that affect an individual data subject or organisation are contestable in the sense that they can be **challenged and reviewed** by a human decision-maker. For example, where a bank rejects a loan application based on an automated AI assessment of the application, it should have a mechanism whereby the applicant can seek a review of that decision to test whether it is a fair reflection of credit risk. Indeed, this type of review process is already a legal requirement under some laws, such as under the EU General Data Protection Regulation for decisions based on personal data. Apart from addressing ethical concerns, establishing a review process will also help to **build trust** in AI systems in a way that maintains the social licence required to roll out of this type of technology.

For more on this, or any other issues relating to use of AI for your business, please contact one of the specialists in our national Tech Law team:

Melbourne

Michael Swinson
Partner, Melbourne
T +61 3 9643 4266
M +61 488 040 000
michael.swinson@au.kwm.com

Kate Creighton-Selvay
Partner, Melbourne
T +61 3 9643 4071
M +61 405 993 122
kate.creighton-selvay@au.kwm.com

Cheng Lim
Partner, Melbourne
T +61 3 9643 4193
M +61 419 357 172
cheng.lim@au.kwm.com

Sydney

Patrick Gunning
Partner, Sydney
T +61 2 9296 2170
M +61 418 297 018
patrick.gunning@au.kwm.com

Scott Bouvier
Partner, Sydney
T +61 2 9296 2472
M +61 407 894 528
scott.bouvier@au.kwm.com

Bryony Evans
Partner, Sydney
T +61 2 9296 2565
M +61 428 610 023
bryony.evans@au.kwm.com

Nicole Heller
Partner, Sydney
T +61 2 9296 2347
M +61 417 213 334
nicole.heller@au.kwm.com

Brisbane

Kirsten Bowe
Partner, Brisbane
T +61 7 3244 8206
M +61 409 460 861
kirsten.bowe@au.kwm.com

John Swinson
Partner, Brisbane
T +61 7 3244 8050
M +61 408 220 513
john.swinson@au.kwm.com

Canberra

Annabel Griffin
Partner, Canberra
T +61 2 6217 6075
M +61 408 847 519
annabel.griffin@au.kwm.com

Rachael Lewis
Partner, Canberra
T +61 2 6217 6074
M +61 448 056 645
rachael.lewis@au.kwm.com

1449 - 01/20