PUBLIC CONSULTATION ON ARTIFICIAL INTELLIGENCE AND SINGAPORE'S INTELLECTUAL PROPERTY REGIME

Consultation period: 26 Aug to 22 Oct 2026.

MINISTRY OF LAW
SINGAPORE

IPOS
INTELLECTUAL PROPERTY
OFFICE OF SINGAPORE

CONTENTS

Click on the tabs for quick access

PART I: EXECUTIVE SUMMARY 3

PART II: BACKGROUND 4

PART III: COPYRIGHT 7

PART IV: PATENTS 20

PART V: SUBMISSION OF COMMENTS 26

ANNEX A: INDUSTRY-DEVELOPED TECHNICAL MEASURES AND STANDARDS 27

ANNEX B: GLOSSARY 30

ENDNOTES 32

PART I: EXECUTIVE SUMMARY

AI is changing how people create, invent, research and do business. Singapore's IP regime is anchored on a careful balance: protecting the rights and incentives that support human creativity and investment, while enabling innovation for the public good. This consultation seeks views on how Singapore's IP laws and frameworks can remain fit for purpose, trusted and responsive to emerging developments as AI continues to advance.

Singapore government agencies have adopted a co-creation approach to our national policy on AI issues. Earlier engagements by MinLaw and IPOS with rights owners, technology companies, representatives from the creative industries, legal practitioners, academics and other stakeholders helped identify the issues covered in this paper. This public consultation now seeks broader views and perspectives from interested parties. The aim is to ensure thoughtful, carefully calibrated IP frameworks that can continue to support innovation while protecting the interests that underpin human creativity and ingenuity.

For **copyright**, the consultation focuses on three broad areas:

- How to foster greater certainty and accountability in AI training, including whether the scope and applications of the computational data analysis exception are sufficiently clear, how lawful access and rights-owner safeguards should operate in practice, and how the Government can support non-legally binding technical measures in a practical and industry-led way (see page 7).

- Copyright risk management in AI deployment and use, including how existing legal principles should apply when AI-generated output infringes copyright, how responsibility should be assessed among AI developers, deployers and end users, and what technical measures may help minimise output infringement risks in a proportionate and commercially viable way (see page 12).

- The nature of human creativity in AI-assisted works, focusing on how human contribution should be understood, recognised and evidenced where generative AI is used as a tool in creative and business workflows (see page 16).

Further details on the copyright consultation may be found at Part III (pages 7-19) of the paper.

For **patents**, the consultation focuses on two broad areas:

- How existing inventorship principles should apply across the spectrum of human-AI interactions in the inventive process, including problem formulation, selection from AI-generated outputs, and human modification of AI-generated technical solutions (see page 20).

- How the large-scale publication of AI-generated technical disclosures may affect the prior art landscape, including patent searches, examination, patentability assessments and incentives for innovation as AI-generated material becomes more publicly available (see page 24).

Further details on the patents consultation may be found at Part IV (pages 20-25) of the paper.

We look forward to receiving your feedback. **Part V** contains practical information on the submission of comments (see page 26).

PART II: BACKGROUND

01 Singapore's broader AI policy seeks to support innovation while maintaining trust, security, privacy and legitimate commercial interests. This includes work such as the National AI Strategy 2.0 and IMDA's guidance on AI governance, provenance, data transparency and liability issues. Singapore's IP policy should be considered in this broader context, with attention to both AI development and the confidence creators, inventors, businesses and investors need to participate in the innovation economy.

02 This consultation focuses on copyright and patents, where AI raises the most immediate policy questions. Copyright law has been at the centre of global litigation and debate over AI training, AI-generated output and human authorship. Patent law faces questions about how human contribution should be assessed where AI is used to generate, evaluate or refine technical solutions, and how AI-generated disclosures should be treated as prior art.

A. COPYRIGHT

Background leading up to this consultation

03 Singapore modernised its Copyright Act in 2021 to address digital-age technological and market developments, including the introduction of a computational data analysis ("**CDA**") exception[1] to support data-driven innovation, machine learning and Singapore's broader digital economy.

04 With the rapid acceleration of AI capabilities, what began as tools for prediction, classification and pattern recognition have developed into systems capable of generating text, images, audio, video, code and technical solutions. Increasingly, these systems are also taking multi-step action across digital environments.

05 The increasing use of generative AI has raised issues about whether copyright works could be used to train AI models and how the law should balance AI development with the legitimate interests of rights owners. AI outputs increasingly stand on their own, substitute for existing works, or closely resemble works on which a model may or may not have been trained. The policy questions have thus expanded beyond training to infringement and allocation of responsibility among AI developers, deployers and end users. As multimodal and agentic AI systems increasingly execute tasks across connected tools, applications and datasets, copyright risk will occur across training, deployment and use, rather than a single act of copying or a single point in the value chain.

06 A further question concerns the nature of human creativity when generative AI is used in content creation. AI may be a powerful tool for augmenting creativity, but the practical challenge is to identify when human choices – including prompting, grounding, selection, arrangement, editing and other creative interventions – are sufficiently reflected in the final work to justify copyright protection.

International developments and responses

07 International responses to the above questions have converged to varying degrees, with the AI training question reflecting the most gradation in treatment across jurisdictions. Singapore's approach sits within the global mainstream and continues to be guided by the defining features of our copyright regime: balanced yet robust protection, regulatory stability, futureproofing through technological neutrality, business certainty and collaborative policymaking.

Fig 1: Overview of global developments and Singapore's approach mapped to the AI lifecycle

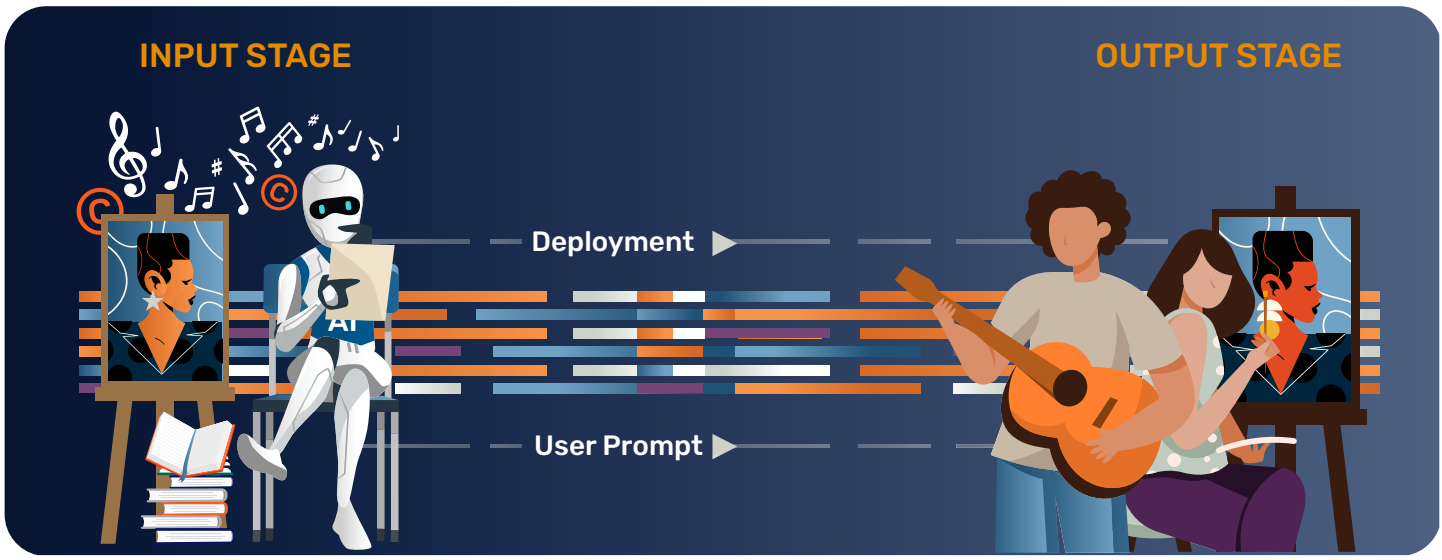

AI Training	Output Infringement	Authorship

Around the world

National positions range from more permissive approaches that allow text and data mining ("**TDM**") or machine learning subject to safeguards, to more restrictive approaches that rely more heavily on authorisation and licensing.[2] In parallel, market-based licensing arrangements have emerged as an important pathway for providing fair returns to rights owners and greater certainty to AI developers.[3]

Courts are still working through fact-specific infringement disputes involving AI-generated content.[4]

The position is more settled as international developments increasingly affirm a human-centric approach: autonomously generated AI output is generally not protected by copyright, while AI-assisted works may be protected where sufficient human creativity is reflected in the final expression.[5]

Singapore

The CDA exception permits AI training with safeguards, sharing common ground with jurisdictions that permit such activities through a TDM exception or fair use. The CDA exception and private licensing are complementary channels within a broader enterprise ecosystem, rather than competing solutions.

As with major jurisdictions, the Singapore courts may adapt established copyright principles to determine a case on its facts.

Singapore requires human authorship in the creative process and supports the use of generative AI as a tool to augment human creativity.

Co-creation in the copyright space

08 MinLaw and IPOS have approached these issues through engagement and co-creation. This includes a series of roundtable discussions in 2025 organised by the Singapore Academy of Law ("the **SAL Roundtable Discussions**") and supported by MinLaw and IPOS with a wide cross-section of industry, technical and legal representatives and experts. MCCY, MDDI, IMDA and NLB have also engaged with stakeholders in the media, literary arts, visual arts and performing arts communities, to better understand their views on the impact of AI. As a follow-up to these engagements, IPOS published a series of AI and copyright law explainers from October 2025.[6]

09 These engagements have surfaced a wide range of stakeholders' concerns. Ultimately, the copyright regime should strive to balance legitimate concerns on all sides. This consultation is a further step in that process.

B. PATENTS

Background leading up to this consultation

10 AI is increasingly used to support human research and innovation. For example, in drug discovery and electronic design automation, humans may define objectives and constraints while AI tools contribute to the generation, evaluation or refinement of possible technical solutions. Such uses of AI raise questions about how human contribution should be assessed. At the more contested end of the spectrum, Dr Stephen Thaler had presented his AI system named DABUS as having autonomously generated inventions and filed patent applications around the world in the name of DABUS as the inventor.

11 The increasing use of AI in the inventive process has also prompted discussion about the implications of AI-generated technical disclosures for the patent system – in particular, the prior art landscape in which future inventions are assessed. Generative AI can produce large volumes of technical disclosures at speed, some of which may be speculative, lack sufficient detail and accuracy to enable persons skilled in the art to carry out what is described, but some of which may nonetheless be publicly accessible and relevant to later patent applications. This raises questions about how such disclosures should be treated within existing prior art and patentability frameworks.

International developments and responses

12 Internationally, the question of AI inventorship has attracted significant attention, driven in part by the DABUS litigation. Patent offices and courts in major jurisdictions such as Japan, the United Kingdom and the United States have generally maintained that an inventor must be a natural person. Against this backdrop, the DABUS applications, which named an AI system as the inventor, were rejected in most patent offices, with subsequent court decisions generally affirming the requirement for human inventorship. IPOS similarly rejected the DABUS application on the basis that Singapore's Patents Act 1994 only recognises a natural person as an inventor. These developments have prompted debate on the level of human contribution required for inventorship and whether material AI involvement should be disclosed in patent applications.

13 The implications of AI-generated technical disclosures are comparatively less explored. As AI systems become increasingly capable of generating and publishing technical content at scale, patent systems will need to consider how such disclosures should be treated as prior art for assessing novelty and inventive step, and the broader policy implications on the state of the art.

Co-creation in the patents space

14 IPOS has engaged multinational technology companies, AI firms, legal academics and IP practitioners to understand how AI is being used in innovation and research. IPOS also conducts exchanges with fellow IP offices internationally to deepen our understanding of AI-related patent issues. Stakeholders generally view AI as a tool that augments, rather than replaces, human inventors, and have emphasised the need for clarity over the patentability of AI-related inventions. This feedback informed IPOS' October 2024 Supplemental Guidance[7] for the Examination of AI-related Patent Applications.

15 These discussions demonstrate that AI-related patent issues require a careful balance: recognising human contribution as the starting premise of inventorship, giving innovators confidence that AI-assisted inventions can be protected where patentability requirements are met, and ensuring that AI-generated disclosures do not undermine incentives for human and AI-assisted innovation.

16 Building on these engagements, broader views informed by real-world experiences would be required for further consideration on how Singapore's patent framework should address AI-assisted innovation.

A. Fostering greater certainty and accountability in AI training

17 Singapore's copyright regime provides a sound foundation for AI innovation and enterprise growth, combining legal exceptions with market-based licensing solutions to support the responsible use of copyright works for AI development. The CDA exception is central to this approach. It balances innovation and protection by allowing copyright works to be used for such purposes, subject to stringent safeguards to protect rights owners' interests.

18 The exception has undergirded Singapore's push towards data-driven innovation and transformation into an AI-powered economy. The exception, which expressly recognises the use of copyright works for machine learning, has been described as a "*game-changer*" in giving businesses and researchers confidence to innovate in an environment that is not fraught with litigation over the legality of using copyright works for AI training.[8] The exception's technological neutrality buttresses regulatory stability by avoiding the need for constant legislative recalibration in response to the rapid advancement of AI capabilities.

A Responsive Framework

19 Regulatory stability should not preclude continued review as technology evolves and matures. In the five years since the exception was introduced, Singapore has engaged stakeholders and, where appropriate, considered practical measures or refinements in response to feedback:

a. Reinforcing "lawful access" to support rights owners in controlling and monetising access to works for AI training: Singapore's 2024 public consultation affirmed that rights owners could rely on paywalls, password controls and other measures to control access to their works and derive value through charging for such access for AI training uses.

b. Highlighting commercial licensing as a complementary pathway: IPOS has highlighted the potential additional benefits of taking up private licences for AI uses and explained how the CDA exception complements such licences.[9]

c. Building mindshare to drive copyright compliance and accountability in AI training: IPOS's explainer on *Artificial Intelligence Training and Copyright Infringement in Singapore* (the "**AI Training Explainer**")[10] provides accessible explanations on how the exception operates and strengthened the exception's safeguards through practical guidance.

d. Seeding co-creation on industry priorities for refinement: Extensive engagements with key stakeholders, including via the SAL Roundtable Discussions,[11] have helped crystallise the core issues and scenarios where there is a shared need for greater certainty and refinement.

20 The use of copyright works for AI training is widely debated, with jurisdictions adopting varied approaches to balancing innovation, rights protection and accountability. Stakeholders differ on whether policy should emphasise licensing, exceptions or a combination of both. In Singapore's case where there is an established exception complemented by licensing, the next step in co-creation involves examining how this approach may be refined to further enhance certainty and accountability through Issues A(I) and A(II).

Issue A(I): Scope and applications of the CDA exception

21 This consultation seeks views on whether targeted refinements to the framework may further enhance certainty, transparency and accountability in the context of AI training. From the engagements and consultations, stakeholders' views have diverged particularly in two aspects of the exception's scope and applications.

Applicability of the exception to generative AI and other machine learning applications

22 A significant source of uncertainty appears to stem from the premise that since the exception predated the generative AI boom, the intention was not for it to apply to generative AI and other subsequent technological developments.[12] Another reason for uncertainty relates to the purpose-based limitation: for the exception to apply, copyright works can be used only for the purpose of CDA and no other purpose. Some have ventured that where generative AI is concerned, the use of works is not only for the purpose of CDA, but also for generating output, and therefore should not be covered by the exception.

23 The exception is expressly intended to apply to machine learning for all types of technologies, including generative AI – consistent with its technology-neutral design and policy intent to support AI and data-driven innovation in Singapore.[13] This aligns with a purposive interpretation that promotes this objective, as well as the established principle that a statute should be interpreted to take into account changes since its enactment, including technological developments and changes in scientific understanding, even if those changes could not have been reasonably contemplated or foreseen at the time of enactment.[14]

24 The purpose-based limitation therefore focuses on the training phase, where copyright works are used as information or data to improve a computer program's functions[15], and does not cover the eventual use of the AI model being trained. This paper clarifies that the CDA's policy intent is to cover a broader, rather than narrower, range of technological developments in AI training.

25 It should be noted that there could be infringement issues at the *output phase* where there may be further uses of copyright works (other than AI training) when such models are deployed and produce infringing output pursuant to users' prompts. The exception would not apply here as such uses do not involve CDA.

Applicability of the exception to commercial activities

Certain stakeholders have highlighted that the legislation does not explicitly state that the exception applies to commercial activities.

26 The exception itself is agnostic as to whether the use of copyright works is in commercial or non-commercial contexts. It is the *absence of conditions* in the legislation that permits the exception to apply in either of these contexts.[16] Hence, commercial CDA activities can benefit from the exception. This position and its rationale have been comprehensively covered in earlier policy materials,[17] the Parliamentary debates,[18] and public guidance and information resources such as the AI Training Explainer.

27 The legal position on the above two aspects is clear – the CDA exception may be used to develop generative AI models, and there is no express limitation on the purpose for which CDA is undertaken (including for commercial purposes). Nonetheless, we seek to understand if there is any need for greater clarity on these or other aspects relating to the exception's scope and applications.

28

Questions:

a. What existing and emerging applications of AI training would benefit from the exception?

b. Are the scope and applications of the exception sufficiently clear and if not, what statement(s) or illustration(s) should be added to the legislation for greater clarity?

Issue A(II): Legal obligations that safeguard rights owners' interests

29 Legal obligations that safeguard rights owners' interests are key to maintaining the balance struck by the exception. Clear obligations set defined boundaries for users and operate effectively as safeguards, preventing uncertainty from excusing non-compliance.

The legal measures (e.g., contractual obligations) and 30 technical measures (e.g., technological measures)[19] employed by rights owners to control access to or use of copyright works may give rise to legal obligations that attract civil and, in some cases, criminal liabilities if breached. The lawful access safeguard in the exception supports rights owners' use of such measures: users cannot rely on the exception if they unlawfully access copyright works for AI training purposes, including breaching a database's terms of use or circumventing a paywall or other access control measure. This is separate from any further liability that users may incur for breaching these obligations.

In special cases, users can lawfully access copyright works for AI training under the exception, despite such measures. These include cases where contract terms 31 are unenforceable or where technical measures are non-legally binding.

a. Unenforceable contract terms: Such terms have no legal effect and hence do not require users' compliance. In the Copyright Act, one category of such terms are those that purport to exclude or restrict the exception, whether directly or indirectly.[20] Such terms are void, ensuring that private contractual arrangements do not override the exception's benefits.[21]

b. Non-legally binding technical measures: Such measures do not create legal obligations on users. One example is the Robots Exclusion Protocol ("**REP**"), or "robots.txt", which website owners use to specify whether web crawlers are permitted to scrape content hosted on their site. This protocol relies on voluntary compliance and is not legally binding. It is also not a technical (access control) measure that Part 7, Division 4 of the Copyright Act makes unlawful to circumvent. Please see paragraphs 33 – 37 below for elaboration on this issue, as well as **Annex A – Industry-Developed Technical Measures and Standards.**

The perceived lack of certainty here arises from the highly fact-specific nature of the issue. Contract terms and technical measures take many forms and operate in ever-changing contexts. This consultation concentrates on improving clarity in the key scenarios and perspectives that have surfaced through stakeholder engagements. 32

Fig 2: Key scenarios considering whether the exception applies despite legal or technical measures

Lawful access

Is access in breach of contract unlawful? Currently, the exception provides an illustration that breach of a database's terms of use constitutes unlawful access.[22] There is some uncertainty as to the effectiveness of this illustration and whether it applies more generally to contractual breaches. Some stakeholders have recommended legislative clarification to elevate the illustration to a deeming provision to this effect: "*X is deemed not to have lawful access to the first copy if access to the first copy is obtained in breach of a contract…*".[23]

Is access in breach of other civil wrongs unlawful? Stakeholders have expressed different views, with SAL questioning whether the lawful access safeguard may be interpreted such that only criminal offences would be considered "unlawful". Such an interpretation would, however, substantially limit the scope of the safeguard because access in breach of laws that give rise to civil liability (such as a breach of contract[24] or a contravention of the prohibition against circumvention of access control measures[25]) would then not be considered unlawful.

Enforceability of contract terms

The next few scenarios ask whether certain contract terms are enforceable.

Is a term that controls the scope or manner of access enforceable? Stakeholders posit two specific scenarios – where a user accesses a non-public area of an online service, contrary to a term prohibiting such access, and where a user accesses an online service through means other than those provided by the service provider, contrary to a term prohibiting access through those means.[26]

Is a term that seeks to protect the security and stability of online services enforceable? For example, when a user uses automated means to copy materials from an online service, overburdening it and causing an outage that is in breach of a term that prohibits use of the service in any manner that could disable, overburden, damage or impair the service.[27]

Is a term that does not expressly reference the exception or any CDA-related uses enforceable? Commercial contracts often contain terms that shape whether and how copyright works can be used for AI training, even if they do not refer expressly to the exception or to CDA-related uses. This could include terms that permit only personal or non-commercial uses of copyright works (and remain silent on CDA), or terms that impose confidentiality obligations. A key issue is whether such terms purport to ***indirectly*** exclude or restrict the exception.

Accessing copyright works in breach of an enforceable contract term may constitute unlawful access (by way of breach of contract) and prevent the exception from applying.

Questions:

a. What technical or legal measures do rights owners employ to control access to or use of their copyright works for AI training?

b. Is it clear from the exception when the exception applies despite rights owners' use of any such measures?

   i. If not, should the issue be addressed by allowing case law to develop, legislative amendment, or non-legally binding instruments? Please explain the reasons for your preference.

   ii. If you propose legislative clarification or non-legally binding instruments, what form[28] should this take and what scenarios should be covered?

Please consider the above-mentioned scenarios (where applicable) and any other similar scenarios you may have encountered.

Issue A(III): Enhanced support for non-legally binding technical measures

33 Technical measures operate across different stages of the AI lifecycle and as mentioned in paragraph 31 above, not all technical measures are legally binding.

34 At the AI training stage, non-legally binding technical measures include a range of tools, system design choices, and operational practices that enable rights owners to express and/or give effect to preferences regarding access to or use of their works. These may be supported by technical standards, which help technical measures operate effectively and strengthen industry-wide accountability, consistency and interoperability. Such standards can also enable industries to develop context-specific solutions. **Annex A** sets out an overview of common technical measures and supporting standards that rights owners have adopted to control access to or use of copyright works for AI training.

35 The effectiveness of such measures depends on widespread compliance and adoption by both rights owners and AI developers, deployers and other users. It is also shaped by practical aspects such as the non-binding nature and uneven adoption of the measures, and the varying levels of awareness, capability and resources among stakeholders. In this regard, WIPO urges collaboration on industry self-regulation and ethical standards that "*go beyond minimum legal compliance*".[29]

36 We see room to encourage greater respect in the use of copyright works for AI training – to enable not merely lawful, but responsible, ethical and trusted AI development. Singapore is aligned with other major jurisdictions in fostering such respect, including for users to build in non-binding technical measures when developing and deploying AI tools. Technical safeguards for AI should remain non-legally binding because technology and industry practices evolve faster than legislation can sensibly prescribe them.[30]

37 We seek views on how non-legally binding technical measures can be developed and adopted in ways that are practical, trusted and responsive to industry needs.

Questions:

a. How can the Government support rights owners in deploying non-legally binding technical measures to control access to and use of their copyright works for AI training, and encourage AI developers/deployers and other users to build in or respect such measures? What non-legally binding technical measures should it support?

B. Strengthening copyright risk management in AI deployment and use

38 Singapore's copyright regime provides rights owners with established protections that may be applied to AI output infringement. Stakeholders generally consider the existing framework to be robust in addressing output infringement, and Singapore continues to do well in international rankings for copyright protection.[31]

39 The inherently fact-based nature of infringement cases allows the Singapore courts to apply established legal principles to new technologies, actors and modes of use. However, there are fewer opportunities for the courts to publish decisions in emerging areas like AI, due to limited copyright litigation and dispute resolution options that facilitate out-of-court settlements.[32]

40 In addition to development by case law, the Government has adopted a proactive stance focused on actionable risk management strategies, giving enterprises and rights owners practical tools to assess exposure and address infringement risks upstream. For example, IPOS's explainer on *Artificial Intelligence Output and Copyright Infringement in Singapore* (the "**AI Output Infringement Explainer**")[33] explains how the existing legal framework can be applied to assess and mitigate the risk of AI-output infringement by various parties, distinguishes between lower- versus higher-risk conduct, and encourages the adoption of best practices when developing, deploying and using generative AI.

Fig 3: Extract of Risk-Based Assessment in AI Output Infringement Explainer

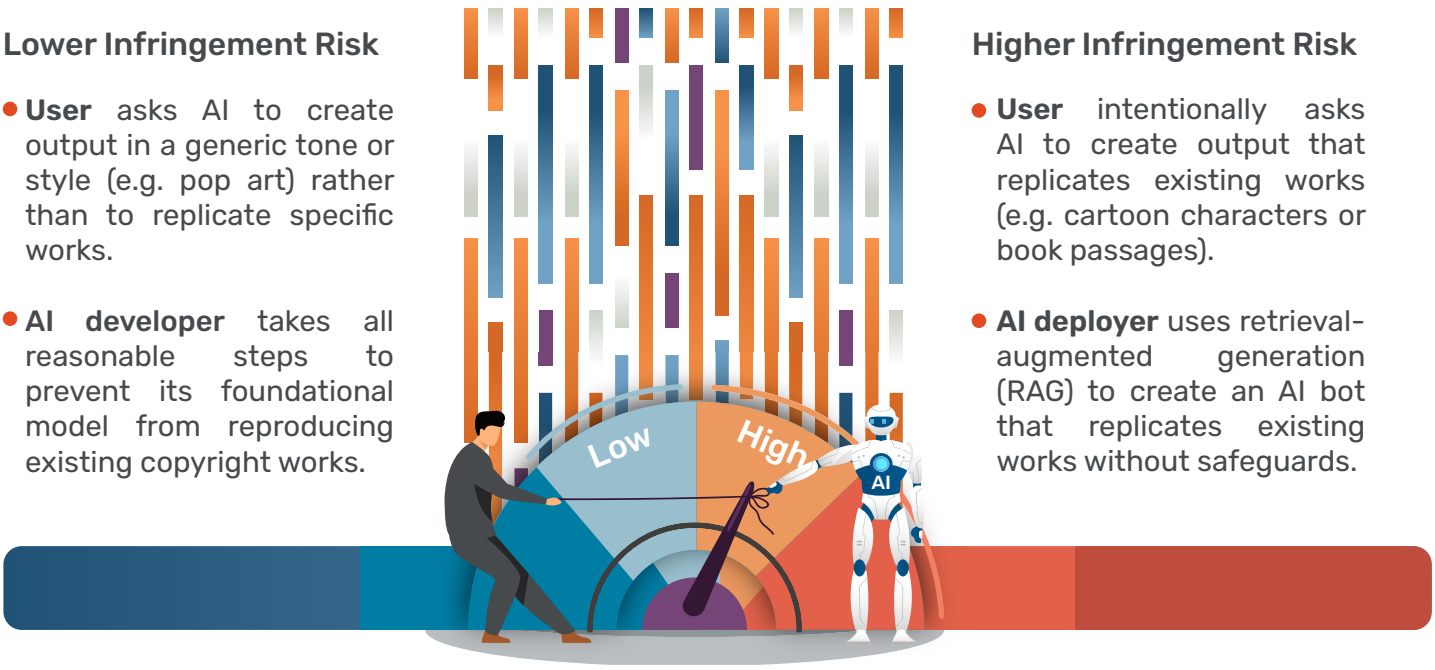

41 Singapore's approach sits within wider international efforts to ensure that AI innovation develops alongside robust copyright protection. Other major jurisdictions have similarly leveraged established copyright principles to respond to output infringement issues, and where necessary, supplemented those principles with guidance to help stakeholders manage risks and encourage responsible AI deployment and use.[34]

42 Against this backdrop, the questions in Issues B(I) and B(II) engage stakeholders on views that may inform the risk assessment and potential liability for copyright infringement in the output phase.

43 These questions focus exclusively on copyright infringement liability and are not intended to address broader issues of general transparency obligations and civil liability, such as liability and allocation of legal responsibility in the use of agentic AI, which are the subject of other workstreams by the Government.[35]

Issue B(I): Copyright infringement liability when AI generates infringing output

44 A risk-based approach must remain grounded in an accurate application of Singapore's copyright law. Depending on the facts, an AI developer, AI deployer and/or end user (i.e., prompter) may be potentially liable for copyright infringement by doing one or more of the following infringing acts:

a. Unauthorised reproduction: creation of output that substantially reproduces an existing work may infringe the rights owner's exclusive right of reproduction. The infringing output may be initiated by direct human instruction through end-user prompts or, in the case of agentic AI, executed through multi-step machine reasoning directed towards a human-defined objective. As AI models are trained on large amounts of data, there is a possibility of "memorisation" if a model is over-exposed to certain types of training data. This may result in output that is substantially similar to one or more existing works. However, such similarities may also arise even if the model has not been trained on the existing work(s), such as when end-user prompts are deliberately crafted to elicit output that substantially reproduces the work(s). It should be noted that insofar as the output only reproduces the general style of a work, there will be no infringement as copyright law protects the particular expression embodied in the work and not its underlying ideas, facts or stylistic elements.

b. Unauthorised communication to the public: if the output is an infringing copy that is also shared or disseminated online, such as on websites and social media platforms, this may result in a further infringement of the rights owner's exclusive right of communication to the public.

c. Authorising infringement: a party may be liable for authorising another party's infringing reproduction or communication, even if it does not carry out the infringing act itself. As set out by Singapore's Court of Appeal[36], liability for authorising copyright infringement is assessed holistically by considering the alleged authoriser's control and power to prevent the infringement, relationship with the actual infringer, reasonable preventive steps taken, and actual or constructive knowledge of the infringement or its likelihood.

45 The discussion of this issue in Singapore has focused on potential authorisation liability for AI developers and deployers, i.e. whether these parties have authorised an end user's infringing reproduction or merely acted as an intermediary or host by making available a service to the end user.[37]

46 We seek views on how existing legal principles should apply when AI-generated output infringes copyright, including how responsibility should be assessed among AI developers, deployers and end users. These views may help clarify questions around preventing, detecting and responding to infringement risks.

Questions:

a. When an AI output infringes the copyright in an existing work, when and to what extent should the following parties be liable under Singapore's copyright law?

   i. AI developer;
   ii. AI deployer;
   iii. end user (i.e. prompter); and/or
   iv. any other party not mentioned above.

b. How would your views differ in cases where the infringing output arises from the use of agentic AI?

Issue B(II): Technical measures to minimise infringement risks

47 Technical measures have played a significant role in influencing liability for copyright infringement in court cases globally. Beyond the use of such technical measures as voluntary risk-management tools, certain jurisdictions such as the EU and China have also imposed AI-specific regulatory obligations that mandate technical measures to mitigate risk of copyright infringement. The interaction between copyright liability and AI-specific regulatory obligations can be seen in China's *Ultraman* case.[38] In that case, an AI platform provider was liable for contributory copyright infringement by allowing users to generate clearly recognisable variations of the "Ultraman" character, and for contravening China's generative AI-related measures by failing to take preventive or filtering measures despite being capable of doing so.

48 These are important and evolving regulatory responses, whose implications and effectiveness will become clearer over time – as implementation experience accumulates across jurisdictions and through insights shared by stakeholders.

49 Minimising output infringement is a shared responsibility for parties that develop, deploy and use generative AI within their respective domains. Singapore does not mandate binding technical measures that target the generation of infringing output, but as explained in paragraph 40 above, encourages responsible, ethical and trusted AI development, deployment and use through resources such as the AI Output Infringement Explainer (Fig 3 above).[43]

50 The questions in this section seek views on how technical measures to minimise output infringement risks can continue to develop and be adopted in ways that support effective, proportionate and commercially viable practices based on a principled allocation of infringement liability. Given the potentially disproportionate impact on smaller or less well-resourced parties, we also seek to understand the extent to which any technical measures should be calibrated to the scale and capacity of AI developers, deployers and users.

Examples of voluntary technical measures influencing litigation outcomes:

- In the US case of *Kadrey v Meta Platforms, Inc.*, post-training "mitigations" implemented by Meta Platforms prevented their AI model from "memorising" and generating more than 50 tokens from the claimants' books. This led to the court's rejection of the claim that the AI model had been trained to "repackage" the works and contributed to the court's finding of fair use.[39]

- In the UK case of *Getty Images (US) Inc & Ors v Stability AI Ltd*,[40] Getty Images discontinued its output infringement claim after Stability AI blocked the prompts that had been allegedly used to generate the infringing output examples in the proceedings.

Examples of regulatory obligations mandating technical measures:

- The EU General-Purpose AI Code of Practice was published in July 2025 as part of compliance with the EU AI Act's obligations. The Code of Practice provides measures that can be taken by general-purpose AI model providers to mitigate the risk of copyright-infringing outputs. These include implementing technical safeguards that prevent models from generating output that reproduces training content, and prohibiting infringing uses of a model in acceptable use policies, terms and conditions, or equivalent documents.[41]

- *China's National Standard of the People's Republic of China: Cybersecurity Technology - Basic Safety Requirements for Generative Artificial Intelligence Services* came into effect on 1 November 2025. This similarly required generative AI service providers to, amongst other obligations, (i) include technical solutions for addressing IP infringement risks in technical documentation; (ii) comply with specific IP protection requirements, such as establishing training data IP management strategies and rules, and communicating IP risks in the use of generated content in user service agreements; (iii) implement model output safety requirements, such as restricting infringing output and refusing to answer questions that obviously induce such output; (iv) suspend services to recalcitrant end-users.[42]

Questions:

a. What technical measures should AI developers, deployers and users adopt to minimise the risk of infringement arising from AI output? Where relevant, please identify: –

   i. any industry or international technical standards or norms; and

   ii. measures that would help minimise risks specific to agentic AI[44].

b. To what extent should any such measures be calibrated to the scale and capacity of AI developers, deployers and users? If so, how should the relevant threshold(s) be defined?

C. The nature of human creativity in AI-assisted works

51 Generative and agentic AI are changing creative processes. Works may now be produced through repeated interactions between humans and AI systems, including prompting, selection, editing and other human-directed choices, increasingly blurring the line between human and machine.

The Human-Centric Paradigm of Authorship

52 Around the world, regulatory responses and cases are converging on the centrality of human contribution to copyright protection of AI output. WIPO has affirmed copyright's design as a human-centric right,[45] and major jurisdictions such as the US, China, Korea, Japan, and the EU have upheld the requirement for human authorship. In a move that would bring the UK closer to these jurisdictions, the UK government has proposed abolishing its longstanding copyright protection for computer-generated authorial works that have no human authors, citing legal contradictions and minimal evidence that protection was being used or had significant economic effect.

53 With agreement on the need for human contribution, the global conversation is turning towards the nature, degree and evidence of human contribution when generative AI is used in the creative process. Notably, factors for assessing copyrightability and human contribution have since been developed by legal experts and authorities in several jurisdictions have issued guidance in this area.

Empowering the Man behind the Machine

54 Singapore's framework places human authorship at the centre of copyright protection, while allowing creators and businesses to use generative AI as a tool. As of 2025, more than 50% of Singapore's creative professionals have used some form of generative AI tools in their work,[46] even as creatives differ in their sentiments about generative AI usage.

Singapore's human-centric approach to authorship

Singapore's copyright legislation and courts have held steadfast to the requirement for human authorship for works (i.e., literary, dramatic, musical or artistic works) to receive copyright protection. As the Court of Appeal observed: "*…the objective of copyright law has been to encourage the creativity of* ***natural authors***".[47]

A work that is created with AI's involvement is only protected by copyright if it is **made by a human** and is an "***authorial creation that is causally connected with the engagement of the human intellect***". The courts will assess whether the human acts and decisions in the creative process demonstrate sufficient intellectual effort, and whether they are sufficiently reflected in the final work to justify copyright protection. In such cases, generative AI is used as a tool to augment human creativity and does not take over the role of its human operator. The prompter must apply intellectual effort or creativity, or exercise mental labour, skill or judgment in creating the final work. Effort (even intellectual) that is applied not towards the authorial creation but towards other ends will not count, such as most fact verification, data entry or purely mechanical tasks.

Issue C: Human Authorship of AI-assisted Works

55 The engagements to date provide a starting point for a deeper dialogue on the nature of creativity under the existing legal framework. Singapore's copyright regime provides protection in deserving cases where generative AI is used as a tool to augment human creativity. The question is thus how to achieve greater clarity and predictability in differentiating between protectable AI-assisted works and unprotectable AI-generated content, especially at the margins.

56 Views on this include:

a. On the Government's part, IPOS's explainer on *Artificial Intelligence and Copyright Protection in Singapore* helps creators understand the application of existing legal principles to generative AI output, with tips for demonstrating human input, choices, and judgment when using generative AI.[48]

b. For stakeholders, the SAL Roundtables Discussions discussed how the human authorship requirement can be demonstrated in machine-human creative processes. Participants recognised copyright protection along a spectrum: the greater the human creative contribution in shaping the work over which copyright is sought, the greater the prospects of obtaining copyright protection and the stronger the protection. For AI-assisted works that barely cross the threshold for protection, infringement will only be made out if such works are copied more substantially. Others will have more room to borrow from such works without infringing copyright – particularly where the copied elements are mostly attributable to the AI system's expressive choices rather than the human author's contribution. Some considered that this could help preserve a measure of comparative advantage for creators who continue to rely on more traditional, human-led modes of creation.

Fig 4: Identifying the human creative contribution in AI-assisted works

Pre-generation acts and decisions by the prompter

This entails context-setting, grounding, and prompting that help to manage the variability of the AI's output and ultimately reflect human control over the expressive elements of the final work over which copyright protection is sought. Such control changes the prompter from being a mere instructor to an author of a copyright-protected work. The prompter may demonstrate this through choices as to:

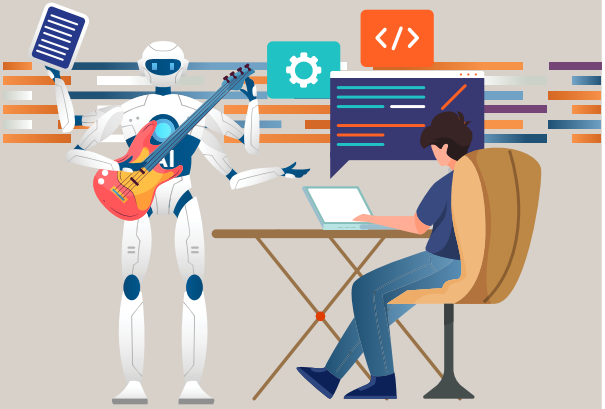

- the AI model or system used for generating the output, especially if a specific model or system possesses distinct stylistic or technical characteristics that can yield the expressive outcomes intended by the prompter;
- the detail and granularity of the prompts (e.g., stylistic or technical instructions, or negative prompts);
- the number of prompts (e.g., through iterative prompting, prompt chaining, or multiple generation attempts);
- the nature of the prompts (e.g., whether they were directed at shaping the expressive elements of the output);
- the source of the prompts (e.g., whether they originated from the prompter or were rewritten or modified by the AI system); and
- the adjustments of technical parameters, such as seeds, weights, or temperature settings to impose structural constraints and exert more precise control over the output.

The significance of prompting is still evolving globally. For example, in the United States, the Copyright Office does not recognise prompting alone as providing sufficient human contribution. South Korea has been less categorical in its denial of copyrightability; it has been observed that entering a prompt is unlikely to demonstrate controllability or predictability, and thus has a low chance of being recognised as a creative contribution. Conversely, in China, the Beijing Internet Court has found that multiple rounds of iterative prompting, the use of negative prompts, and continuous adjustment of output can evidence that the user exercised creative judgment and control over the resulting AI output.

- **Post-generation selection and arrangement of AI output**

This includes:

- choosing from multiple output options generated by the AI;
- choosing when to stop further generation or modification of the AI output; and
- selection and arrangement to combine multiple AI and/or human output into a compilation or broader copyright-protected work, such as a comic or multimedia composition.

Merely choosing one option from a set of near-identical alternatives may involve minimal creative judgment, whereas engaging in sustained iterative review across multiple AI-generated outputs to land on the option that best realises a particular expressive objective may involve more substantial creative judgment.

- **Post-generation editing**

This involves making original human-authored modifications to the AI output after generation, such as editing, redrawing, rewriting, retouching, transforming, or other creative intervention. Given the more direct nature of the human intervention on the final work product, major jurisdictions have shown greater willingness to recognise that such modifications may result in copyright protection of the final work product.

Notes:

1. *The above summary is based on the Government's stakeholder engagements, global legal commentaries, and guidance notes from other jurisdictions. It is intended to facilitate further feedback on the issue. The Government does not presently express any views on the possible considerations.*
2. *The prompter's actions may recur in an iterative process where one or more of the above steps are repeated to produce the final work over which copyright is sought. For example, an initial AI-generated output may be manually edited, fed back into the system for further generation, and yet further edited before the final work is produced.*

57 The questions in this section seek insights as to how human contribution should be understood and recognised as the Government charts its copyright policy in this age of AI-mediated creativity.

Questions:

a. What existing and emerging use cases of generative AI augment human creativity in the process of creating authorial works?

b. What acts or decisions by creators (i.e. human prompters) demonstrate human creativity when generative AI is involved in creating an authorial work?[49]

c. Should principles or factors be developed to assess the sufficiency of human creativity in such cases? If so, what should these be and should they be developed by the courts, prescribed in legislation, or provided in non-legally binding instruments?

d. How would your views to b and c differ in cases where the AI output is executed using an agentic AI system?

PART IV: PATENTS

58 The patent issues in Part IV are at an earlier stage of policy development than the copyright issues in Part III. While major jurisdictions including Singapore require a human inventor, advances in AI are giving rise to diverse forms of human-AI interactions in the inventive process. This Part seeks views from stakeholders, including enterprises, practitioners, inventors and researchers on how AI is affecting inventorship and patentability.

D. Assessing human contribution in AI-assisted inventions

59 Today, advances in AI are giving rise to a spectrum of human-AI interactions in the inventive process, ranging from:[50]

a. **AI-assisted inventions:** AI functions as an assistive tool, supporting human researchers who define the problem and conceive a solution; to

b. **AI-generated inventions:** AI autonomously identifies a problem and generates a novel solution with minimal to no human involvement.

Singapore's human-centric approach to inventorship

Under Singapore's patent law, only a human being can be named as an inventor. The Patents Act 1994 defines an "inventor" as "the actual deviser of the invention". Singapore courts have interpreted this to mean the natural person (human) who came up with the inventive concept(s),[51] i.e., "the heart" of the invention, or the "idea or principle…which entitles the inventor's achievement to be called inventive".

Generally, inventive concept(s) or contribution(s) include something in an invention that helps to solve a particular problem or answer a particular question in a new way, and may include the perfection or improvement of a solution.

To qualify as an inventor, a person must have formulated or contributed to the formulation of the inventive concept(s).[52]

To illustrate, this would include a person who formulated the inventive concept but had no capability to realise it into a functional product,[53] but would not include a person who merely did what was suggested to him/her by the actual inventor[54] or a person who contributed "unnecessary detail" or managerial and entrepreneurial contributions to an invention.[55]

Inventorship is a personal right that cannot be assigned,[56] though the ownership of a patent can be transferred. Patent applicants must file a statement identifying the inventor(s), failing which the application will be treated as abandoned.[57] It is possible to request to add or remove the inventor(s) in any patent granted or published application for a patent.[58] A patent may be revoked if it was granted to a person who was not entitled to be granted that patent.[59]

Issue D: How should human contribution be assessed for inventorship in AI-assisted inventions?

60 Singapore's position — that only a natural person can be named as an inventor — is consistent with the approach taken in most jurisdictions. While it is accepted law that contribution to the inventive concept may take different forms and the assessment is heavily fact-dependent, this assessment becomes more complex when AI is involved in generating, evaluating or refining the technical solution. In such cases, it may be difficult to distinguish between human contributions that shaped the inventive concept and AI-generated outputs that merely informed, enabled or accelerated the inventive process.

61 The approach taken to assess the sufficiency of human contribution in the inventive process may have practical implications for R&D planning, investment decisions and IP management strategies. For example, existing R&D processes may need to be redesigned to incorporate explicit human-in-the-loop involvement, alongside increased documentation requirements to substantiate the nature and timing of human-AI interactions throughout the inventive process. Researchers and businesses may need to consider how inventorship should be attributed among researchers using AI tools, and navigate varying compliance requirements of different jurisdictions.

Fig 5: Human-AI interactions in the inventive process

To illustrate the potential implications, we set out below four hypothetical scenarios based on the stages of an AI-assisted drug discovery workflow.

Scenario A: Problem formulation and AI system configuration

A pharmaceutical R&D team seeks to identify candidate compounds capable of binding to a defined biological target implicated in a rare disease. The team deploys a general-purpose AI system configuring it with parameters they have defined: target binding affinity thresholds, molecular stability requirements, and acceptable toxicity profiles. The system autonomously generates and evaluates candidate structures – the eventual output is a finite set of 10 candidate molecular structures predicted to satisfy the parameters that might bind to the relevant biological target. The team did not design the underlying model, contribute training data, or intervene in the generative process. They receive the outputs and proceed to experimental validation.

Do the human acts of identifying the biological target, framing the design problem, and specifying the technical parameters constitute a contribution to the inventive concept?

Scenario B: Selection from AI-generated outputs

The same AI system is now used at greater scale, generating 10,000 candidate molecular structures ranked by predicted efficacy and safety. The research team reviews the ranked outputs, and drawing on their domain expertise, identifies a subset of 50 candidates for further consideration. They then select three compounds for synthesis and experimental validation. One of these compounds proves to be a promising lead candidate. The team's selection decisions were not guided by any additional computational analysis — they relied on scientific judgement about which candidates were most worth pursuing.

Does the exercise of expert judgement in selecting from AI-generated outputs constitute a contribution to the inventive concept?

Scenario C: Modification of AI-generated output

Building on Scenario B, the research team takes the lead candidate identified through AI-assisted screening and modifies its chemical structure. These modifications are made on the basis of the team's own understanding of structure-activity relationships, without further AI assistance. The modified compound demonstrates materially enhanced efficacy and pharmacokinetic properties compared to the AI-generated precursor, and is selected for development into a drug. The claimed invention is the modified compound and its therapeutic application.

Is human modification of an AI-generated output — where the modification is itself the product of independent human reasoning and expertise — sufficient to ground inventorship in the human researchers? Should the inventive concept extend to the AI-generated precursor that made the modification possible?

Scenario D: The role of AI system developers

In each of the preceding scenarios, the AI system was developed and trained by a separate team of engineers and scientists who designed its architecture, curated its training data, and optimised its outputs for molecular design tasks. These developers had no involvement in the specific research projects that produced the candidate compounds.

Should the developers or trainers of an AI system be regarded as having contributed to the inventive concept of downstream inventions generated using that system?

Note: These scenarios are intended to show different points along the spectrum of human involvement and AI use, and to facilitate discussion of how existing inventorship principles may apply in practice. They are not intended to represent legal conclusions or preferred policy outcomes.

62 AI continues to shape research and innovation. This warrants closer examination of how existing inventorship principles apply to AI-assisted inventions. At the same time, responses should be measured, having regard to evolving technological capabilities, changing innovation practices and the fundamental bargain underlying the patent system in which exclusive rights are granted in exchange for the disclosure of new knowledge. This will help ensure that Singapore's patent system remains adaptable to technological developments, continues to provide appropriate incentives for innovation, and supports a favourable environment for investment in new technologies.

Questions:

a. How is AI currently being used in the innovation and research processes in your field or organisation? Please describe, with examples whenever possible, the nature and extent of human involvement and AI involvement in these processes.

b. Singapore's current framework requires a human to have contributed to the inventive concept in order to qualify as an inventor.

   i. Is this adequate to determine inventorship given the human-AI interactions in innovation and research processes today?

   ii. If not, what aspects do you consider most in need of clarification?

c. In relation to AI-assisted inventions, what level or type of human contribution should be required for a person to qualify as an inventor? The following examples illustrate varying types of human contribution:

   i. a person who identified the problem;

   ii. a person who designed the AI model that came up with the solution;

   iii. a person who provided the training data for the AI model;

   iv. a person who formulated the prompts used with the AI model; and

   v. a person who selected from the AI model's outputs.

   For each of the Scenarios A-D in Fig 5 above, which human participant(s) should be treated as an inventor, if any, and why?

d. Where AI was used in the inventive process, what information should applicants be required to disclose about the AI model?

e. If clarification of the inventorship framework for AI-assisted inventions is needed, what form should this take? For example: (i) maintaining the current position and allowing the case law to develop; (ii) legal refinements; (iii) practical information resources or (iv) non-binding technical measures. Please explain the reasons for your preference.

E. The effect of AI-generated prior art on patentability

63 Prior art plays a central role in the patent system's balance between exclusivity and disclosure, by helping to determine whether the applicant has contributed something new and inventive beyond existing public knowledge in return for exclusive rights. In Singapore, the requirements of novelty[60] and inventive step[61] are assessed against the "state of the art", which generally comprises everything made available to the public before the priority date of the invention.[62] This includes published patents, scientific literature, products on the market and other publicly accessible information. In general, anything made available to the public before the priority date may constitute prior art. A claimed invention will not be patentable if it is not new, or if it lacks inventive step if obvious to a person skilled in the art, in light of the prior art. Prior art therefore provides the benchmark for deciding whether exclusive rights should be granted, while keeping existing public knowledge available for all to use.

Issue E: What is the effect of AI-generated prior art on patentability?

64 For over a decade, there have been discussions on computer-generated disclosures and their implications on patentability. They considered matters such as the increasing capabilities of computers to generate technical disclosures (including AI-generated defensive publications that generate large volumes of prior art disclosures) and whether such disclosures should qualify as prior art that may be cited against later inventions. Examples of defensive publication platforms include All Prior Art and All The Claims, which were developed with the express aim of generating and publishing large volumes of technical disclosures that may be relied upon as prior art against patent applications.[63]

65 Generative AI has brought these issues into sharper focus because it can produce technical descriptions at scale and speed. This may increase the volume of publicly available disclosures and affect later patent searches, applications and assessments of patentability.

66 The growing prevalence of AI-generated disclosures also gives rise to a number of practical considerations, for example:

a. Quality and reliability: AI may generate content that is speculative, inaccurate, internally inconsistent or technically unsound.

b. Enablement: AI-generated disclosures may lack the detail, accuracy or technical coherence needed to enable a person skilled in the art to carry out what is described.

c. Prior art search burden: AI may significantly increase the volume of searchable technical information, making it more difficult and resource-intensive for patent offices and other stakeholders to identify, assess and rely on the disclosures that are genuinely relevant.

67 We recognise the central role of prior art in ensuring that patent protection is reserved for inventions that make a genuine technical contribution beyond existing knowledge. At the same time, the growing prevalence of AI-generated disclosures raises important questions about how such content should be treated within the prior art framework, and whether existing approaches for assessing prior art remain adequate to uphold the underlying principles and effective operation of the patent system.

68 We seek the experiences of inventors, patent practitioners and other stakeholders in assessing and relying on AI-generated disclosures, including any challenges such disclosures may present for patent searching, examination and innovation activities. As the use of AI to generate and disseminate technical information evolves, we welcome views on any practical or policy considerations arising from AI-generated disclosures, as well as whether existing approaches to assessing prior art may require clarification or adaptation to ensure that the patent system continues to support both high-quality patents and innovation.

Questions:

a. Are you encountering AI-generated content when searching for existing technical information? How is this affecting R&D or patenting activities in your field?

b. Should AI-generated content be treated as prior art during the patent examination process? In your view, what are the factors or considerations for AI-generated content to be considered prior art? If it should be treated as prior art, should AI-generated content be treated under the same standards as human-generated content? In particular, how should public accessibility, enablement, operability and technical reliability be assessed where content is generated or published at scale? Should the same standards and presumptions apply?

c. How can a prior art disclosure be identified as AI-generated? Are there particular tools, processes, or disclosure requirements that would be helpful? How should the publication date, public accessibility, source and integrity of AI-generated content be established? Would particular metadata, record-keeping practices, tools or disclosure requirements be helpful?

d. Are you aware of instances where AI-generated content has been published strategically to prevent others from obtaining patents? Do you consider this a significant concern, and if so, how should the patent system respond?

F. Patents and other AI-related issues

69

Beyond the specific questions raised under Parts D and E, we would also welcome your views on the following broader questions.

Questions:

a. Patent protection is typically sought in multiple jurisdictions. Are there specific areas relating to the issues in this paper where you see divergence in the approaches between jurisdictions and what is the impact on your organisation?

b. Are there any other aspects of the relationship between AI and the patent system that you consider should be examined? Please describe the issue and, where possible, its practical significance for your organisation or field.

PART V: SUBMISSION OF COMMENTS

70 MinLaw and IPOS welcome views and comments on the above issues. Your views are important and will help us in keeping Singapore's intellectual property regime effective and fit for purpose. All submissions should be clearly and concisely written and should provide explanations for proposed revisions. Where feasible, parties should identify the specific section on which they are commenting and explain the basis for their proposals.

71 Feedback may be submitted via FormSG (at **go.gov.sg/ai-ip2026-copyright** for copyright-related feedback, and at **go.gov.sg/ai-ip2026-patents** for patents-related feedback), or via email to MinLaw (at **MLAW_Consultation@mlaw.gov.sg**). When providing your responses, please also include your name, contact number, and e-mail address, so that we may contact you if there are any follow-up questions.

72 We reserve the right to make public all or parts of any written submission, and to disclose the identity of the source. Commenting parties may request for confidentiality for any part of the submission that is believed to be proprietary, confidential or commercially sensitive. Any such information should be clearly marked and placed in a separate annex. If we grant confidential treatment, we will consider, but will not publicly disclose, the information. If we reject the request for confidential treatment, the information will not be considered. As far as possible, parties should limit any request for confidential treatment of information submitted. We will not accept any submission that requests confidential treatment of all, or a substantial part, of the submission.

73 Please submit any comments by **5pm on 22 Oct 2026**. Thank you.

ANNEX A: INDUSTRY-DEVELOPED TECHNICAL MEASURES AND STANDARDS

Technical measures operate across different stages of the AI lifecycle and function alongside legal, contractual, and policy mechanisms. Rights owners may choose to deploy such measures to safeguard their interests while AI developers, deployers and users may honour rights owners' choices or build in other measures to respect copyright and mitigate infringement risks.

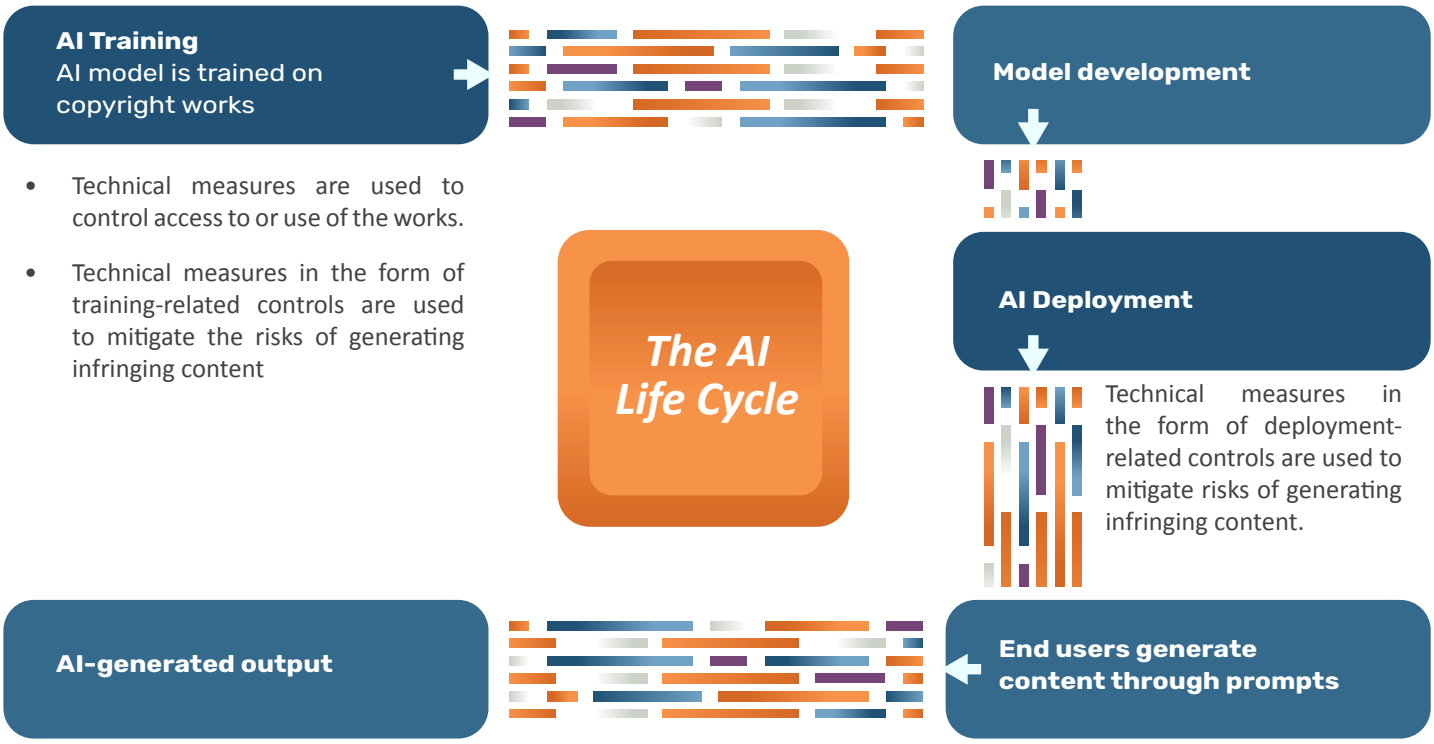

* Further details on these measures are set out below.

Technical standards help technical measures operate effectively and strengthen accountability, consistency and interoperability across the AI lifecycle while preserving flexibility for industry to develop context-specific solutions. These include international frameworks by various organisations and are available in the Organisation for Economic Co-operation and Development ("**OECD**")'s Catalogue of Tools and Metrics for Trustworthy AI:[64]

- the International Organization for Standardization ("**ISO**"), such as ISO/IEC 23894: 2023 *Guidance on Risk Management*; [65]
- the International Electrotechnical Commission ("**IEC**"); and
- the Internet Engineering Task Force ("**IETF**"), which is working on standardising building blocks that allow for rights owners to express preferences about how content is collected and processed for model development, deployment, and use.[66]

As the OECD observes, technical standards foster international cooperation by establishing common frameworks for responsible AI governance on a global scale.[67] Initiatives have arisen in multiple jurisdictions supporting the development, adoption, and promotion of international standards:

- EU's Copyright Infrastructure Task Force is an expert forum working on creating and maintaining an "Open Rights Data Framework".[68] This framework seeks to promote standardisation through improving the interoperability, reliability, and machine readability of copyright data.
- UK's AI Standards Hub provides a platform to facilitate knowledge sharing, coordination, and capacity building amongst different stakeholders.[69]

Technical Measures to Control Access to or Use of Copyright Works during the AI Training Process

Measures include a range of technological tools, system design choices, and operational practices that enable rights owners to express and sometimes enforce technical preferences regarding access to or use of their works.

Site-based tools

Site-based tools operate at a website or domain level. For example:

- the Robots Exclusion Protocol, which has been widely adopted as a voluntary baseline tool to control web scraping.
- website hosting platforms (e.g., Cloudflare[70] or Squarespace) enable website owners to block, rate-limit, or selectively allow AI web crawlers to access their content.

Unit-based tools

Unit-based tools operate at the individual file level. Some tools embed machine-readable metadata in the content itself to specify ownership, licensing and usage terms.

- For example, Adobe Content Authenticity allows creators to apply Content Credentials into each file, which are machine-readable metadata fields that specify ownership, tools used, and preferences around AI training and usage.[71]

Other tools operate to disincentivise AI training on web-scraped content without consent from rights owners.

- For example, the Glaze Project[72] introduced tools such as Glaze and Nightshade that incorporate changes to images that appear only to AI models to disrupt style mimicry or distort how those images are represented in the model.

Tools may be both site-based and unit-based, such as the EU's TDM Reservation Protocol, developed by the industry in response to the EU's AI Act's legislative requirements.[73] This web protocol provides a variety of techniques for rights owners to declare, via machine-readable formats, whether TDM activities can be carried out on their works.

AI developers may also complement the use of these tools by adopting consent-based datasets or incorporating rules within an AI model to ensure that rights owners' technical preferences are respected during model development.

Stability AI's use of Spawning API[74] in training Stable Diffusion V3 integrates opt-out preferences collected through partnerships with rights owners and content platforms, such as ArtStation and Shutterstock. Spawning API helps automate the identification and exclusion of non-consenting content, enabling developers to better align model training practices with rights owners' preferences.[75]

Technical Measures to Mitigate Risk of Generating Infringing Output

Measures to mitigate the risk of generating infringing output can be implemented during model development to reduce memorisation and verbatim reproduction of content, or during model deployment to detect, block, or modify potentially infringing output before it is delivered to users. These measures may work alongside platform-level guardrails, such as platform terms of use specifying that users must not use the model for copyright infringement.

Training-related controls include:

- Dataset curation to exclude the use of infringing data for model training. For example, using datasets which comply with access control measures or "copyright-clean" datasets like those provided in the KL3M Data Project[76].

- Anti-memorisation and deduplication training techniques, which reduce the model's tendency to memorise verbatim sequences. For example, Google DeepMind's VaultGemma model was trained with "differential privacy", a measure that adds calibrated noise into the model to prevent memorisation.[77]

- Post-training machine unlearning and finetuning, which removes specified sequences from the model to reduce the probability of exact reproduction without retraining. For example, Microsoft Research applied unlearning techniques to Meta's Llama2-7b model to replace specific expressions from copyright content with generic alternatives and fine-tuned the model on those relabelled terms.[78]

Deployment-related controls include:

- Output filtering and duplication detection mechanisms to suppress output that closely resembles known copyright works. For example, GitHub Copilot uses a duplicate detection filter to flag or block output that closely matches existing material in GitHub's repositories.[79]

- User-facing prompt controls that block or modify prompts seeking to reproduce copyright works. For example, Adobe Firefly automatically blocks prompts that are intended to generate output that may infringe on a third party's IP rights.[80]

ANNEX B: GLOSSARY

This glossary explains key technical terms used in this consultation paper. The terms include concepts relating to artificial intelligence (AI) and intellectual property (IP) and are intended to help readers who may not be familiar with all of the specialist terminology used in these areas. They should not be read as strict legal or technical definitions.

A

Access control measures
Technical measures in the form of technologies, devices or components that effectively control access to digital content, such as paywalls, password protection, authentication systems, encryption, subscription controls.

Agentic AI
AI systems, typically built on generative AI models, that can autonomously take multi-step actions towards achieving end-user-defined goal(s), such as planning, using tools, interacting with applications or carrying out tasks.

AI-assisted work
An authorial work that uses generative AI as a tool in the creative process, where the final expression reflects sufficient creative choices by a human creator.

AI deployer
One who puts an AI system to use or who makes it available for others to use, without necessarily having developed the system. Includes those who offer the system as a service to end users, or who integrate it into their own or others' services or workflows.

AI developer
One who builds, trains or substantially develops an AI model or system, whether from scratch or by modifying or fine-tuning an existing model or system.

AI-generated prior art
Technical information produced by AI that has been made publicly available before the priority date and may be considered when assessing whether a later invention is new or inventive.

AI model
A core component of an AI system, trained through machine learning to perform specific tasks such as identifying patterns, making predictions or recommendations, or generating content. Examples include forecasting models, large language models, and image generation models.

AI system
A machine-based system comprising one or more AI models and other technical components that support its processes, such as software and interfaces. Based on the input it receives, an AI system infers how to generate output such as predictions, content, recommendations, or decisions. AI systems vary in their levels of autonomy and adaptiveness after deployment. Examples of AI systems include chatbots, generative AI tools and applications, recommendation systems and agentic AI.

Authorial work
A literary, dramatic, musical or artistic work, as defined in section 9 of the Copyright Act 2021.

Authorship
The creation of an authorial work by a person, and the resulting status of that person as the work's creator (i.e., its "author"). Under Singapore law, an author must be a human being.

C

Computational data analysis
Refers to uses such as text and data mining, sentiment analysis, and machine learning. The Copyright Act defines it as including (a) using a computer program to identify, extract and analyse information or data from the work or recording; and (b) using the work or recording as an example of a type of information or data to improve the functioning of a computer program in relation to that type of information or data.

E

Enablement
Whether a disclosure in a patent specification gives enough information for a skilled person to make or perform the invention without undue difficulty.

End user / prompter
The person who interacts with or uses an AI system to do specific tasks, often by entering instructions or prompts and deciding how to use the AI output.

G

Generative AI
AI models or systems that generate output, such as text (including code), images, audio, or video, in response to end-user instructions or prompts.

I

Inventive concept
The core idea or technical contribution that makes an invention inventive, rather than merely routine or obvious.

Inventive step
A statutory requirement that the invention must not be obvious to a skilled person, considering what was already publicly known before the priority date. The presence of an inventive step is one of the three substantive requirements for patentability, alongside novelty and industrial applicability. See sections 13 and 15 of the Patents Act 1994 for more details.

Inventorship
The legal identification of the person or persons who came up with the invention (see section 24 of the Patents Act 1994 for more details). Under Singapore law, an inventor must be a human being.

L

Lawful access
A safeguard for rights owners in the form of a legal requirement that users must satisfy to rely on the CDA exception for AI training. It requires accessing material in a way that is legally permitted, such as without breaching enforceable contract terms, bypassing access control measures, or contravening other applicable laws, see section 244(2)(d) of the Copyright Act 2021 for more details.

M

Machine learning
Used for AI development, where a computer analyses large datasets to learn patterns from the data and make inferences about new data. This is the foundation for most AI models.

Memorisation
A phenomenon where an AI system reproduces output that closely resembles the data that it was trained on.

N

Novelty
A statutory requirement that the invention must be new when compared with what was already publicly available before the priority date. Novelty is one of the three substantive requirements for patentability, alongside inventiveness (or an "inventive step") and industrial applicability. See sections 13 and 14 of the Patents Act 1994 for more details.

P

Patentability
Whether an invention qualifies for patent protection. The three substantive criteria for an invention to be patentable are that it must be new, inventive (see also: inventive step) and capable of industrial application (see sections 13-16 of the Patents Act 1994 for more details).

Prior art
Information that was publicly available before the priority date of a patent application. The existence of prior art affects whether an invention is new or inventive. See sections 14-15 of the Patents Act 1994 for more details.

Priority date
The priority date is the date used to assess whether an invention is new and involves an inventive step. It is usually the filing date of the patent application, but may be the filing date of an earlier application for the same invention if the relevant requirements are met (see section 17 of the Patents Act 1994 for more details).

R

Robots Exclusion Protocol / robots.txt
A commonly used web standard that allows website owners to indicate whether automated crawlers should access or scrape parts of their website. Compliance is generally voluntary unless supported by legal obligations.

T

Technical measures
Technological tools, settings, practices, or safeguards used in the development, deployment or use of AI models or systems, such as to control access to or use of content, signal or enforce preferences, reduce risks or shape behaviour. These may be supported by technical standards.

Technical standards
Common technical requirements, guidelines or frameworks that promote consistency, interoperability, and good practices in the development, deployment or use of AI models or systems.

Text and data mining
The automated analysis of large amounts of text, data or other information to identify patterns, trends, relationships or insights. Also includes uses beyond AI model training and machine learning, such as sentiment analysis, spam filtering, and the identification and extraction of specified information from large volumes of unstructured data.

ENDNOTES

[1] Copyright Act 2021, ss 243 and 244.

[2] Singapore and Japan are broadly permissive, expressly allowing AI training subject to safeguards. The US and EU take more intermediate approaches: US fair use has so far tended to favour developers, while the EU allows text and data mining subject to rights-owner opt outs. Australia remains the most restrictive, having declined to introduce an AI training exception and applying a relatively low infringement threshold.

[3] Examples of licensing deals include: OpenAI's deals with publishers e.g., Associated Press (July 2023), New York Times' deal with Amazon (May 2025).

[4] In Germany and China, rights owners have prevailed in case law against developers, where the court found that the model outputs were liable for infringement. In the UK, AI developer Stability prevailed against copyright owner Getty, where Getty had alleged secondary infringement (pending appeal). There are also cases in the US which are yet to be decided on this point.

[5] As the OECD notes, most jurisdictions agree that works generated autonomously by AI are not copyrightable; judicial and legislative developments across major jurisdictions, including the US and China, continue to reinforce a "*human-centric paradigm of creative authorship*", see Huijuan Peng & David Tan, "Copyright in the Generative AI Era: Reimagining Creativity and Fairness" (2025) 37 Singapore Academy of Law Journal 749 at [62].

[6] Available on IPOS's website at https://www.ipos.gov.sg/about-ip/copyright/copyright-resources/#8a212422e02608c37e8a1b6597b3c009 (accessed 21 August 2026).

[7]Intellectual Property Office of Singapore, "Supplemental Guidance for the Examination of AI-related Patent Applications" (11 October 2024) https://isomer-user-content.by.gov.sg/61/1889bce0-bd7f-4946-8879-4583d3a933a5/supplemental-guidance-for-the-examination-of-ai-related-patent-applications.pdf (accessed 25 August 2026).

[8] Espie Angelica A. de Leon, "Computational data analysis exception in Singapore's Copyright Act 2021 a game changer", Asia IP Law (7 December 2021) https://asiaiplaw.com/article/computational-data-analysis-exception-in-singapores-copyright-act-2021-a-game-changer (accessed 17 July 2026).

[9] Ministry of Law & Intellectual Property Office of Singapore, "Summary of key changes to prescribed exceptions in Part 6, Division 1 of the Copyright Regulations 2021" (19 December 2024), at [10(ii)], and Intellectual Property Office of Singapore, How does Singapore law treat the use of copyright works for AI training? (2026) https://isomer-user-content.by.gov.sg/61/7709dff4-3cdd-438d-b767-df5155aa943d/How%20does%20Singapore%20law%20treat%20the%20use%20of%20copyright%20works%20for%20AI%20training.pdf (accessed 17 July 2026).

[10] Intellectual Property Office of Singapore, "How does Singapore law treat the use of copyright works for AI training?" (2026) https://isomer-user-content.by.gov.sg/61/7709dff4-3cdd-438d-b767-df5155aa943d/How%20does%20Singapore%20law%20treat%20the%20use%20of%20copyright%20works%20for%20AI%20training.pdf (accessed 17 July 2026).

[11] Singapore Academy of Law, "Reform of the Computational Data Analysis Exception" (February 2026) https://sal.org.sg/wp-content/uploads/2026/02/Reform-of-the-Computational-Data-Analysis-Exception-12-Feb-2026.pdf (accessed 16 July 2026) ("**SAL CDA Report**").

[12] SAL CDA Report at [5] – [7].

[13] See also, the Parliamentary debates on the exception, which characterised the useful applications of CDA as "indispensable" to the digital economy and emphasised the exception's importance in supporting Singapore's Smart Nation initiatives, push towards data-driven innovation and efforts to grow its AI and technology sectors: Singapore Parliamentary Debates, Official Report (13 September 2021) vol 95, "Copyright Bill" (Edwin Tong Chun Fai, Second Minister for Law) http://sprs.parl.gov.sg/search/#/sprs3topic?reportid=bill-524 (accessed 20 July 2026).

[14] *Getty Images (US) Inc & Ors v Stability AI Ltd* [2025] EWHC 2863 (Ch) (4 November 2025). The "always speaking" principle is a settled principle in Singapore law, see for example, AAG v Estate of AAH, deceased [2010] 1 SLR 769 at [30].

[15] Copyright Act 2021, s 243(b) defines CDA as including the use of a work "…as an example of a type of information or data to improve the functioning of a computer program in relation to that type of information or data…".

[16] See s 244(2) and 244(4) of the Copyright Act 2021, which exhaustively provide for the conditions for the exception to apply. None relate to commerciality.

[17] Ministry of Law & Intellectual Property Office of Singapore, "Singapore Copyright Review Report" (17 January 2019) https://www.mlaw.gov.sg/files/news/public-consultations/2021/copyrightbill/Annex_A-Copyright_Report2019.pdf (accessed 21 July 2026) at pp 33 – 34.

[18] Singapore Parliamentary Debates, Official Report (13 September 2021) vol 95, "Copyright Bill" (Edwin Tong Chun Fai, Second Minister for Law) http://sprs.parl.gov.sg/search/#/sprs3topic?reportid=bill-524 (accessed 21 July 2026): "*…this permitted use…will apply in both commercial and non-commercial contexts.*"

[19] Including the access control measures considered in the above-mentioned public consultation in 2024. The terms "technological measure" and "access control measure" are as defined in Section 423 of the Copyright Act 2021.

[20] Section 187(1)(c) and 187(2) of the Copyright Act 2021.

[21] As explained in Parliament, the benefit of data analysis improves as the set of underlying data used is as complete as possible. Having specific databases locked out of the analysis because of contractual restrictions would make the results less useful, or worse, create bias or inaccuracies: Singapore Parliamentary Debates, Official Report (13 September 2021) vol 95, "Copyright Bill" (Edwin Tong Chun Fai, Second Minister for Law) http://sprs.parl.gov.sg/search/#/sprs3topic?reportid=bill-524 (accessed 21 July 2026).

[22] See illustration (b) of s 244(2)(d), Copyright Act 2021.

[23] In SAL's opinion, there is case law to suggest that breach of contract is not per se unlawful. The corollary of this is that accessing works for AI training in breach of contract would not be considered unlawful under the exception, despite the illustration currently suggesting otherwise. For more details, refer to SAL CDA Report at [11b] and [101] – [129].

[24] SAL CDA Report at [122]: "Breach of contract is a civil (or private) wrong, not a criminal (or public) wrong. If an obligor breaches a contract, the obligee may have a (private) cause of action for the breach; but the obligor is not exposed to a potential action by the state for that breach. As such, it does not appear to be correct to describe breach of contract as an unlawful activity".

[25] See s 425 (read with s 436), as well as illustration (a) of s 244(2)(d), Copyright Act 2021. SAL CDA Report, at Annexure 3, [5]: "It is not clear, however, whether such infringing activity is unlawful for the purpose of section 244(2)(d) of the Act. That is, there is an argument that all section 436 does is to establish a private cause of action (and associated remedies in section 438) which a rights owner can choose to avail itself of, but the infringing act is not unlawful unless the elements of the criminal offence are established."

[26] SAL CDA Report at [117(a) – (b)].

[27] SAL CDA Report at [117(c)].

[28] For example, legislative clarification could include additional statutory provisions, principles, factors, or illustrations.

[29] World Intellectual Property Organization, "Learning Machines: An Introduction to AI and IP for small and medium-sized enterprises" (WIPO Publication No. 1090/26) at 43.

[30] SAL CDA Report at [149] – [156].

[31] The SAL's 2026 Report on the Reform of Copyright Infringement in Light of Generative AI concludes that Singapore's copyright infringement laws are fundamentally sound and adaptable to new factual circumstances involving generative AI.

[32] For example, the WIPO Arbitration and Mediation Center offers a low-cost mediation scheme for parties to resolve intellectual property, technology and related commercial disputes without resorting to litigation. As the scheme provides a streamlined, informal and affordable way to settle disputes, it is particularly accessible for small and medium enterprises, startups, and individual creators. For more details, see WIPO, "WIPO Arbitration and Mediation Center – Mediation" https://www.wipo.int/amc/en/mediation/information.html (accessed 22 July 2026).

[33] Intellectual Property Office of Singapore, "How does Singapore law treat AI-generated content that may infringe copyright?" (2025) https://isomer-user-content.by.gov.sg/61/908918b6-a9f5-4494-a980-6f93ec6b3f13/ai-infographic-v2-how-does-singapore-law-treat-ai-generated-content-that-may-infringe-copyright.pdf (accessed 21 July 2026).

[34] See for example, UK Intellectual Property Office, "Copyright and Artificial Intelligence: Consultation" (December 2024) https://www.gov.uk/government/consultations/copyright-and-artificial-intelligence/copyright-and-artificial-intelligence (accessed 22 July 2026) at [161]: "…the copyright framework in relation to infringing outputs is reasonably clear and appears to be adequate", and Ministry of Culture, Sports and Tourism (Republic of Korea) & Korea Copyright Commission, "Guide to Preventing Copyright Disputes Related to Generative AI Outputs" (June 2025) https://www.copyright.or.kr/eng/doc/etc_pdf/Guide_to_Preventing_Copyright_Disputes_Related_to_Generative_AI_Outputs.pdf (accessed 22 July 2026) at 4: "No one would disagree that the criteria of access and substantial similarity, which are generally applied when determining copyright infringement, should also apply." It further provides stakeholder-specific guidance to rights holders, users of generative AI and generative AI operators.

[35] See for example, Infocomm Media Development Authority of Singapore, "Discussion Paper: Legal Responsibility for AI Agents" (May 2026) https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/agents-legal-responsibility.pdf (accessed 22 July 2026).

[36] *RecordTV Pte Ltd v MediaCorp TV Singapore Pte Ltd and others* [2011] 1 SLR 830 at [50].

[37] Singapore Academy of Law, "Reform of Copyright Infringement in Light of Generative AI" (March 2026) https://sal.org.sg/wp-content/uploads/2026/02/Reform-of-copyright-infringement-in-light-of-generative-AI.pdf (accessed 22 July 2026) at [78] – [87].

[38] *Shanghai Character Licence Administrative Co Ltd v Hangzhou Jellyfish Intelligent Technology Co Ltd* [2024] Zhe 01 Min Zhong No. 10332 (Hangzhou Intermediate People's Court, 20 December 2024).

[39] *Kadrey v Meta Platforms, Inc.* 788 F Supp 3d 1026 (N.D. Cal. Jun 25, 2025) at 12 – 13, and 18.

[40] *Getty Images (US) Inc & Ors v Stability AI Ltd* [2025] EWHC 2863 (Ch) at [9].

[41] See EU General-Purpose AI Code of Practice, Copyright Chapter, Measures 1.3 and 1.4, pp 5 – 6. This Code of Practice provides a voluntary means of demonstrating compliance with the EU AI Act's requirement for general purpose AI model providers to establish a policy to comply with EU copyright law, see Regulation (EU) 2024/1689 (AI Act), Art 53(1)(c).

[42] See for example, Center for Security and Emerging Technology, "National Standard of the People's Republic of China: Cybersecurity Technology - Basic Safety Requirements for Generative Artificial Intelligence Services" (28 May 2026) https://cset.georgetown.edu/wp-content/uploads/t0661_gen_ai_safety_standard_final_EN.pdf (accessed 22 July 2026) at [4.2.2], [5.2], [6.5], and Appendix B at [B2.1.2.2.2].

[43] Intellectual Property Office of Singapore, "How does Singapore law treat AI-generated content that may infringe copyright?" (2025) https://isomer-user-content.by.gov.sg/61/908918b6-a9f5-4494-a980-6f93ec6b3f13/ai-infographic-v2-how-does-singapore-law-treat-ai-generated-content-that-may-infringe-copyright.pdf (accessed 22 July 2026).

[44] These may include additional safeguards such as disclosure of risks (by the developer), prevention of risky behaviour (such as hacking), providing manual oversight features for the human user, including a human in the loop at appropriate decision-making points.

[45] World Intellectual Property Organization, "Generative AI" (2024) (WIPO Conversation on Intellectual Property (IP) and Frontier Technologies, WIPO Reference RN2024-2EN) https://www.wipo.int/publications/en/details.jsp?id=4750&plang=EN (accessed 14 July 2026) at p 7.

[46] SkillsFuture Singapore, "Skills Demand for the Future Economy Report 2025" (December 2024) https://jobsandskills.skillsfuture.gov.sg/sdfe-2025-published/assets/files/SSG_2024_12.pdf (accessed 14 July 2026) at p 51.

[47] *Asia Pacific Publishing Pte Ltd v Pioneers & Leaders* (Publishers) Pte Ltd [2011] 4 SLR 381 at [41] (emphasis added), [67], and [72].

[48] Intellectual Property Office of Singapore, How Does Singapore Law Treat AI-Generated Content? (2025) https://isomer-user-content.by.gov.sg/61/8f095ec1-521c-4b35-b0aa-ca99bd85a66d/How does Singapore law treat AI-generated content.pdf (accessed 14 July 2026)

[49] Whether or not generative AI (including agentic AI) is used as an intermediate step or to produce the final output.

[50] These are terminologies established by WIPO. See World Intellectual Property Organization, "AI Inventions" (2024) https://www.wipo.int/publications/en/details.jsp?id=4734 (accessed 18 July 2026) at pp 3 – 4.

[51] Patents Act 1994, s 2(1). *National University Hospital (Singapore) Pte Ltd v Cicada Cube Pte Ltd* [2017] SGHC 53 at [50]; *Energenics Pte Ltd v Musse Singapore Pte Ltd* [2013] SGHCR 21 ("*Energenics v Musse*") at [24(a)].

[52] *Cicada Cube Pte Ltd v National University Hospital (Singapore) Pte Ltd and another appeal* [2018] SGCA 52 ("***Cicada v NUH***") at [62] – [64].

[53] *Dien Ghin Electronic (S) Pte Ltd v Khek Tai Ting (trading as Soon Heng Digitax)* [2011] SGHC 36 at [13].

[54] *Cicada v NUH* at [86(b)].

[55] Such as the provision of money, facilities, materials and the like: *ibid*., at [64].

[56] *Energenics v Musse* at [24(b)].

[57] Patents Act 1994, s 24.

[58] Patents Rules, r 17(1).

[59] Patents Act 1994, s 80(1)(b).

[60] Patents Act 1994, s 14(1).

[61] Patents Act 1994.

[62] Patents Act 1994.

[63] All Prior Art is described as "*a project attempting to algorithmically create and publicly publish all possible new prior art, thereby making the published concepts not patent-able. The concept is to democratize ideas, provide an impetus for change in the patent system, and to preempt patent trolls. The system works by pulling text from the entire database of US issued and published (un-approved) patents and creating prior art from the patent language. While most inventions generated will be nonsensical, the cost to computationally create and publish millions of ideas is nearly zero – which allows for a higher probability of possible valid prior art." All The Claims is described to be "attempting the same thing, but with the use of claims and a more verbose alternative*". See Alexander Reben, "About", *All Prior Art* https://allpriorart.com/about/ (accessed 21 July 2026).

[64] OECD.AI, "Catalogue of Tools & Metrics for trustworthy AI" (2026) https://oecd.ai/en/catalogue/tools?toolTypeIds=6&page=1 (accessed 24 July 2026).

[65] ISO Library, "ISO 23894 Artificial Intelligence – Risk Management" (2023) https://iso-library.com/standard/23894/ (accessed 24 July 2026). This standard has been adopted by the UK as BS EN ISO/IEC 23894:2024, see British Standards Institution, "BS EN ISO/IEC 23894:2024 Information technology. Artificial intelligence. Guidance on risk management." (29 February 2024) https://knowledge.bsigroup.com/products/information-technology-artificial-intelligence-guidance-on-risk-management-1 (accessed 24 July 2026).

[66] Suresh Krishnan & Mark Nottingham, "IETF setting standards for AI preferences" (27 February 2025) https://www.ietf.org/blog/aipref-wg/ (accessed 24 July 2026).

[67] Paul Khullar, "AI Standards Hub makes its AI standards available on the OECD Catalogue of Tools and Metrics for trustworthy AI" (9 September 2024) https://oecd.ai/en/wonk/ai-standards-hub-oecd-catalogue-tools (accessed 24 July 2026).

[68] Ministry of Education and Culture (Finland), "Copyright Infrastructure Task Force" (2026) https://okm.fi/en/citf-en (accessed 24 July 2026).

[69] The Alan Turing Institute, "About the AI Standards Hub" (2024) https://aistandardshub.org/the-ai-standards-hub/ (accessed 24 July 2026).

[70] Cloudflare also allows website owners to opt to charge for access through a "pay per crawl" feature. For more details, see Will Allen & Simon Newton, "Introducing pay per crawl: Enabling content owners to charge AI crawlers for access" (1 July 2025) https://blog.cloudflare.com/introducing-pay-per-crawl/ (accessed 24 July 2026).

[71] Andy Parsons, "Adobe Content Authenticity, now in public beta, helps creators secure attribution" (24 April 2025) https://blog.adobe.com/en/publish/2025/04/24/adobe-content-authenticity-now-public-beta-helps-creators-secure-attribution (accessed 24 July 2026).

[72] University of Chicago, "About The Glaze Project" (2026) https://nightshade.cs.uchicago.edu/aboutus.html (accessed 24 July 2026).

[73] UK Intellectual Property Office, "Report on Copyright and Artificial Intelligence" (Presented to Parliament pursuant to s 136 of the Data (Use and Access) Act 2025, March 2026) at p 78. For more details, see European Digital Reading Lab, "The W3C TDM Reservation Protocol" (2026) https://www.edrlab.org/open-standards/tdmrep/ (accessed 24 July 2026), which describes the TDM Reservation Protocol as a rights reservation model including techniques such as embedding permissions or restrictions on TDM in webpages (e.g., HTML content) and files (e.g., EPUB and PDF content).

[74] An API, or Application Programming Interface, is a set of rules and protocols that allows different software applications to communicate and exchange data with each other.

[75] Laura Exline, "The Spawning Guide to Rights Reservations" (27 March 2024) https://spawning.substack.com/p/the-spawning-guide-to-rights-reservations (accessed 24 July 2026).

[76] Institute for the Advancement of Legal and Ethical AI, "KL3M Data Project" (15 April 2025) https://aleainstitute.ai/work/kl3m-data-project/ (accessed 24 July 2026).

[77] Amer Sinha & Ryan McKenna, "Vault Gemma: The world's most capable differentially private LLM" (Google Research, 12 September 2025) https://research.google/blog/vaultgemma-the-worlds-most-capable-differentially-private-llm/ (accessed 24 July 2026).

[78] Ronen Eldan & Mark Russinovich, "Who's Harry Potter? Making LLMs forget" (Microsoft Research, 4 October 2023) https://www.microsoft.com/en-us/research/articles/whos-harry-potter-making-llms-forget/ (accessed 24 July 2026).

[79] GitHub Learn, "GitHub Copilot" (2026) https://learn.github.com/learning-pathways/github-copilot (accessed 24 July 2026), and Derrick Mwiti, "GitHub Copilot Privacy: A Guide to Safeguards & Troubleshooting" (2 June 2026) https://www.datacamp.com/tutorial/github-copilot-privacy-and-troubleshooting (accessed 24 July 2026).

[80] Adobe, "Adobe Generative AI Product Specific Terms" (24 April 2025) https://wwwimages2.adobe.com/content/dam/cc/en/legal/servicetou/Adobe-Generative-AI-Product-Specific-Terms-en_US-20250424.pdf (accessed 6 August 2026).