International Center for Law & Economics

Comments of the International Center for Law & Economics

RE: Public Consultation on Artificial Intelligence and Singapore's Intellectual Property Regime

October 22, 2026

Authored by:

Kristian Stout, Director of Innovation Policy, ICLE
Sabrina Pekarovic, Innovation Fellow, ICLE
Giuseppe Colangelo, Senior Scholar, ICLE

I. Introduction

The International Center for Law & Economics (ICLE) thanks the Ministry of Law (MinLaw) and the Intellectual Property Office of Singapore (IPOS) for the opportunity to comment on the Public Consultation on Artificial Intelligence and Singapore's Intellectual Property Regime.

ICLE is a nonprofit, nonpartisan, global research and policy center based in Portland, Oregon, United States. It promotes law & economics methods and economic evidence to inform public policy. Its scholars have written extensively on copyright, technology, and innovation, including licensing copyrighted works for artificial intelligence (AI) training, exceptions for computational data analysis, and the treatment of AI outputs under competition principles. ICLE has also addressed these issues in jurisdictions considering reforms relevant to Singapore, including through comments to the European Commission on the Copyright in the Digital Single Market Directive and its application to AI.

Because Singapore's framework offers an influential model for copyright policy in the AI era, this consultation will help shape policy choices beyond Singapore's borders. We commend MinLaw and IPOS for seeking broad public input before proposing changes.

The Copyright Act 2021 already provides a foundation for AI training and the protection of rights holders' legitimate interests. Its computational data analysis (CDA) exception allows developers to train models using works they lawfully access.[1] But rights holders also retain tools to control and charge for access, including paywalls, contracts, and licenses. The exception does not authorize developers to circumvent those access controls.[2] The consultation correctly treats the CDA exception and commercial licensing as complementary channels. Together, they provide a sound basis for responsible AI development.[3]

Market evidence illustrates why both channels are important. Developers face substantial difficulty valuing each work's contribution to model training. Major licensing deals generally cover entire catalogs, with developers negotiating access through large rights holders and aggregators.[4] These arrangements can supply substantial volumes of data, but they cannot readily extend to every work that training uses.

Synthetic data—that is, data generated artificially—and licensed datasets can supplement training material and reduce some valuation difficulties. Their usefulness varies across training stages, however,

---

[1] Copyright Act 2021, §§ 243–244 (Sing.).

[2] Ministry of Law & Intellectual Property Office of Singapore, *Public Consultation on Artificial Intelligence and Singapore's IP Regime* ¶ 30 (Aug. 26, 2026) [hereinafter Public Consultation].

[3] *Id.* ¶ 17.

[4] Hamidah Oderinwale & Anna Kazlauskas, *The Economics of AI Training Data: A Research Agenda* (arXiv:2510.24990v3, 2026) (examining AI training-data licensing deals involving major rights holders from 2020 to 2025).

and they remain largely supplements to the extensive underlying data that training requires.[5] The CDA exception enables developers to use lawfully accessed works beyond the reach of licensing agreements.

There is an economic basis to this division. The CDA exception addresses uncertainty about input values *ex ante*—before training begins—and enables access at scale. It also gives developers greater legal certainty than frameworks that depend on case-by-case fair-use determinations.[6]

Licensing, meanwhile, allows parties to price the additional value of curated catalogs, platform feeds, and other organized datasets. Developers can pay for high-quality, vetted material and avoid some of the costs of collecting data from public sources.[7] The two channels perform distinct functions. The exception enables lawful training without individual negotiations, while licensing establishes prices for commercially valuable access and services. Requiring licenses for every training input would undermine the certainty that makes experimentation affordable, particularly for smaller developers.[8]

These comments address three questions. First, where should contractual control over access end and the CDA exception begin? Second, how do the exception's assumptions about valuing individual works align with a market that trades largely in catalogs and datasets? Third, who should bear the cost of preventing infringing outputs—a legally distinct question from the use of inputs for training?[9]

The economic analysis and available evidence support Singapore's current framework. Importing the European Union's (EU) transparency duties and rights-reservation mechanisms, which allow rights holders to opt out of certain training uses, would raise compliance costs while leaving central copyright questions unresolved.[10] These measures would add administrative and copyright burdens of the kind developers already face under uneven implementations of text-and-data-mining exceptions across EU member states.[11]

---

[5] *See* Feiyang Kang *et al.*, *Demystifying Synthetic Data in LLM Pre-training: A Systematic Study of Scaling Laws, Benefits, and Pitfalls*, at 1–3 (arXiv:2510.01631, 2025) (examining how synthetic data's benefits and limitations vary across training stages).

[6] *See* Kristian Stout, Ben Sperry & Subiksha Ramakrishnan, *Generative AI: When Fair Use Becomes Unfair Competition*, 32 GEO. MASON L. REV. 779 (2026) (examining the economic foundations of copyright and AI training).

[7] Public Consultation, *supra* note 2, ¶¶ 19, 30.

[8] Stout *et al.*, *supra* note 6, at 779, 781–85 (examining the transaction costs of individualized licensing and the limits of fair use as a substitute for legal certainty).

[9] Public Consultation, *supra* note 2, ¶¶ 22–28, 41–50.

[10] Kristian Stout & Giuseppe Colangelo, *ICLE Comments to the European Commission on Copyright in the Digital Single Market Directive* (Int'l Ctr. for L. & Econ., June 25, 2026).

[11] João Pedro Quintais, *Generative AI, Copyright and the AI Act*, 56 COMPUT. L. & SEC. REV. 106107, at 5–7 (2025), https://doi.org/10.1016/j.clsr.2025.106107.

Singapore should first clarify the boundaries between the CDA exception, access controls, and commercial licensing through guidance. That experience should inform any later assessment of the need for additional obligations.

II. Valuation, Licensing, and Creator Compensation

The consultation's questions about rights-holder safeguards turn partly on a valuation problem. No accepted method currently measures the value of an individual copyrighted work within a training dataset.[12] A work's contribution depends on the model's design, the training objective, the other material available, and duplication across inputs.[13] Its importance to a human reader or viewer may reveal little about its contribution to a model trained on billions of inputs.

Timing also matters, as a work may add more value early in training than after the model has encountered similar material. Researchers continue to develop valuation methods, but most assess collections or subject areas rather than the individual works that a work-by-work licensing system would need to price.[14] The U.S. Copyright Office identifies the related practical obstacle that the cost of locating a rights holder and negotiating permission may exceed the value of the license.[15]

Research on orphan works—works whose rights holders cannot be identified or located—offers a parallel. Maurizio Borghi, Kris Erickson, and Marcella Favale find that locating rights holders for mass digitization entails substantial costs before bargaining or enforcement begins.[16] Even a reliable valuation method would leave those costs intact.

The Google Books litigation illustrates the difficulty of securing permissions at scale. Google had the resources and technology to digitize millions of books and agreed to fund a Book Rights Registry to locate rights holders and distribute payments under a proposed class-action settlement.[17] The court nevertheless rejected the settlement, partly because it would authorize future uses of works without their owners' affirmative consent, including orphan works. It concluded that arrangements for exploiting unclaimed books belonged more appropriately to Congress.[18]

---

[12] Oderinwale & Kazlauskas, *supra* note 4.

[13] *Id.* at 3, 11–12; *see also* Kang *et al.*, *supra* note 5, at 1–3.

[14] Oderinwale & Kazlauskas, *supra* note 4, at 3, 11–12.

[15] U.S. Copyright Office, *Copyright and Artificial Intelligence, Part 3: Generative AI Training* 71, 104 (pre-publication version May 2025).

[16] Maurizio Borghi, Kris Erickson & Marcella Favale, *With Enough Eyeballs All Searches Are Diligent: Mobilizing the Crowd in Copyright Clearance for Mass Digitization*, 16 CHI.-KENT J. INTELL. PROP. 135, 149–53 (2016), https://scholarship.kentlaw.iit.edu/ckjip/vol16/iss1/6.

[17] *Authors Guild v. Google, Inc.*, 770 F. Supp. 2d 666 (S.D.N.Y. 2011) (rejecting the amended settlement); *Authors Guild v. Google, Inc.*, 954 F. Supp. 2d 282 (S.D.N.Y. 2013), *aff'd*, 804 F.3d 202 (2d Cir. 2015), *cert. denied*, 578 U.S. 941 (2016).

[18] *Authors Guild v. Google, Inc.*, 770 F. Supp. 2d 666, 677–81 (S.D.N.Y. 2011) (rejecting the proposed settlement and concluding that Congress should address the exploitation of unclaimed works).

The proposed comprehensive digital library did not emerge from that settlement. Google's unlicensed search and snippet functions instead survived under fair use, a result the Supreme Court declined to review.[19] The episode illustrates how substantial resources and a willingness to pay can still leave major legal and practical barriers to assembling permissions. AI training datasets can encompass an even broader range of works.

Michael Heller describes the underlying coordination problem as the "tragedy of the anticommons." When many owners hold separate rights to exclude others, the cost of assembling permissions can prevent valuable uses of a resource.[20] AI training can present that problem on a vast scale, involving millions of rights holders, some unidentifiable, whose works may each make only a small contribution. Singapore's CDA exception helps avoid that gridlock.

These considerations inform both licensing policy and creator compensation. Markets generally price access to catalogs and datasets, while compulsory schemes would still need workable rates and distribution rules. How creators share in licensing revenue also depends on their contracts with publishers and other aggregators. The following subsections explain why Singapore should allow these commercial arrangements to develop before imposing new obligations.

A. Licensing Markets and the Limits of Compulsory Licensing

The first consequence concerns what licensing markets actually price. Large enterprise deals generally cover access to catalogs, archives, subscription feeds, and databases, rather than each work's contribution to training.[21] Major AI licensing arrangements commonly aggregate rights at this level.[22]

The CDA exception's lawful-access condition supports these transactions. Section 244 requires developers to respect controls over access, including paywalls, passwords, and applicable database terms.[23] Singapore's framework therefore supports commercial negotiations at the level markets can currently value—access to collections of works.

Bundled agreements further complicate valuation. As Matthew Sag and Peter K. Yu explain, AI partnerships often combine training rights with summaries, attribution, and hyperlinks. [24]

---

[19] *Authors Guild v. Google, Inc.*, 954 F. Supp. 2d 282 (S.D.N.Y. 2013), *aff'd*, 804 F.3d 202 (2d Cir. 2015), *cert. denied*, 578 U.S. 941 (2016).

[20] Michael A. Heller, *The Tragedy of the Anticommons: Property in the Transition from Marx to Markets*, 111 HARV. L. REV. 621 (1998); Michael A. Heller & Rebecca S. Eisenberg, *Can Patents Deter Innovation? The Anticommons in Biomedical Research*, 280 SCIENCE 698 (1998).

[21] *See* Oderinwale & Kazlauskas, *supra* note 4, at 7–8, 17–18; *see also* Public Consultation, *supra* note 2, ¶ 19 (discussing commercial licensing as a complementary channel).

[22] Oderinwale & Kazlauskas, *supra* note 4, at 7–8, 17–18 (cataloging licensing deals from 2020 to 2025 and their concentration among rights aggregators).

[23] Copyright Act 2021, § 244 (Sing.); Public Consultation, *supra* note 2, ¶ 30.

[24] Matthew Sag & Peter K. Yu, *The Globalization of Copyright Exceptions for AI Training*, 74 EMORY L.J. 1163, 1216–17 (2025), https://scholarlycommons.law.emory.edu/elj/vol74/iss5/4.

Confidential contract terms make it difficult to determine how much of a reported deal price reflects training rights alone.

The difficulty of valuing individual training inputs also complicates proposals for compulsory licensing or mandatory levies. These systems require administrators to set rates and decide how to distribute payments. Yet valuation research offers no dependable measure of an individual work's contribution to training. That gap complicates both pricing and distribution, even when a scheme pools royalties rather than assesses them work by work. Rate-setting also requires substantial administrative resources and can generate costly disputes even when values are measurable.[25]

The United Kingdom's review of copyright and AI found no case for legislating licensing terms and concluded that licensing should remain a matter for commercial negotiation.[26] The U.S. Copyright Office likewise concluded that compulsory licensing should await evidence of market failure.[27]

Singapore's consultation identifies no market failure that requires such intervention. Emerging markets for catalog licenses provide further reason to allow commercial arrangements to develop. Both the CDA framework and AI licensing markets remain young. More practical experience would help policymakers assess their performance before introducing costly and contentious legislative changes.

B. Creator Compensation and Contractual Adjustment

The distribution of licensing revenue warrants further study. A review of major AI training-data deals announced between 2020 and 2025 finds that most agreements studied exclude original creators from compensation because publishers, platforms, and archives hold the relevant rights.[28]

This pattern alone does not establish a market failure. Training requires vast collections of works, which makes large catalog holders natural negotiating partners. Direct deals between individual creators and major AI developers may remain uncommon because the cost of negotiating each agreement can exceed the value of the creator's contribution to training.

Concerns about how creators share in licensing revenue therefore point first to their contracts with publishers and other aggregators. Many existing agreements predate commercial AI training and may not account for its value. As parties negotiate or renew contracts, they can address training uses explicitly through provisions that define the rights granted and allocate licensing revenue.

These adjustments are likely to use familiar contractual tools, including renewal provisions, revised royalty terms, and clear definitions of training rights. They do not inherently require new legislation.

---

[25] U.S. Copyright Office, *supra* note 15, at 95–97, 104–06 (explaining that rate-setting requires substantial administrative resources and that compulsory licenses warrant use only in exceptional cases).

[26] Dep't for Sci., Innovation & Tech. & Dep't for Culture, Media & Sport, *Report on Copyright and Artificial Intelligence* § G, ¶¶ 41, 63–64 (Mar. 2026) [hereinafter Report on Copyright and Artificial Intelligence].

[27] U.S. Copyright Office, *supra* note 15, at 104–06.

[28] Oderinwale & Kazlauskas, *supra* note 4, at 7–8, 17–18.

As creators negotiate compensation for training uses, aggregators' costs and obligations can change, informing the terms they negotiate with AI developers. Contractual adjustment can thus carry through the licensing chain.

Policymakers should give that process time to develop. Premature intervention could constrain contractual experimentation and introduce a compulsory licensing system that would be difficult to administer. Such a system would also likely rely on large aggregating bodies, much as music licensing relies on performing rights organizations. It could reproduce the existing distribution of revenue while adding compliance costs.

Aggregation also lowers the cost of licensing each work, helping new sellers enter the market.[29] Singapore should allow voluntary licensing arrangements to mature and assess their effects before concluding that they fail creators. The consultation itself should be commended for recognizing market-based licensing as a pathway to fair returns for rights holders.[30]

III. Access Terms and the CDA Exception

The consultation asks "whether the scope and applications of the exception [are] sufficiently clear."[31] The central uncertainty concerns how the CDA exception interacts with contractual terms. Sections 243 and 244 preserve rights holders' ability to control and charge for lawful access through paywalls, database restrictions, and licensing arrangements.[32] Section 187, meanwhile, voids contractual terms to the extent they directly or indirectly exclude or restrict uses the exception permits.[33]

The challenge lies in distinguishing legitimate access conditions from restrictions that undermine the exception. An expansive reading of "indirectly" could invalidate ordinary commercial terms and discourage the catalog-level licensing arrangements discussed in Section II.

Singapore should clarify this boundary through a functional distinction between conditions governing access and restrictions on computational analysis after lawful access. IPOS guidance, supported by practical examples, could preserve the exception's protections while giving parties greater confidence to negotiate. That approach would also allow policy to respond as licensing practices develop.

A. Distinguishing Access Terms from Restrictions on Training

The consultation's scenarios illustrate the uncertainty about which contract terms remain enforceable.[34] Does a personal-use or noncommercial-use restriction indirectly limit the CDA exception if it

---

[29] *Id.* (explaining that aggregation reduces transaction costs per work).

[30] Public Consultation, *supra* note 2, at 5 & fig. 1 (discussing market-based licensing as a pathway to fair returns for rights owners).

[31] *Id.* at 8, question A(I)(b).

[32] Copyright Act 2021, §§ 243–244 (Sing.); Public Consultation, *supra* note 2, ¶ 30.

[33] Copyright Act 2021, § 187 (Sing.).

[34] Public Consultation, *supra* note 2, ¶¶ 31–32.

says nothing about computational analysis? Does a confidentiality clause over a licensed catalog survive? What about a term specifying how a user may access the material? These provisions can determine whether a rights holder will offer access at all. Yet the consultation provides no clear framework for assessing them.

Section 187 reflects a legislative judgment about which contractual restrictions should yield to the CDA exception. Although the Copyright Act 2021 predates the generative AI boom and illustrates computational analysis through image classification, its rationale extends beyond that technology.[35] The government has confirmed that the exception covers machine learning across technologies, including generative AI, consistent with interpreting legislation to accommodate technological change.[36]

The underlying economics also remain relevant. Requiring individual licenses for dispersed training inputs can make valuable uses prohibitively expensive. That problem may be greater for frontier models, which draw on more works without making individual contributions easier to value.

Section 187 addresses this problem by preventing parties from contracting out of the exception. Without that protection, rights holders could require developers to negotiate permission for uses the statute otherwise permits. A clause expressly prohibiting AI training presents the clearest case for applying this statutory protection.

The difficulty lies in the provision's treatment of terms that "indirectly" restrict the exception. An expansive interpretation could reach ordinary conditions governing access, including terms that support the licensing markets the Act otherwise accommodates.

Paywalls, logins, rate limits, and personal-use licenses do not necessarily prohibit computational analysis. They may instead define the conditions under which a user can obtain material. Section 244 makes lawful access a condition of the CDA exception.[37] The two provisions therefore require a coherent boundary. Section 244 governs entitlement to access, while section 187 limits contractual restrictions on uses the exception permits.

Without clearer guidance, parties cannot reliably predict whether an access condition will later count as an indirect restriction. The consultation describes that assessment as "highly fact-specific."[38] The consequences also vary across transactions. Uncertainty about standard access terms can affect many users, while uncertainty about a negotiated catalog license can undermine the terms that make a substantial commercial agreement worthwhile.

---

[35] *Id.* ¶¶ 22–24.

[36] *Id.* ¶ 23.

[37] Copyright Act 2021, § 244 (Sing.).

[38] Public Consultation, *supra* note 2, ¶ 32.

1. The Costs of Contractual Uncertainty

Rules that prevent parties from varying statutory entitlements can also discourage useful agreements tailored to different works and commercial needs.[39] The valuation problem discussed in Section II illustrates the stakes. No accepted method measures an individual work's contribution to AI training.[40] That contribution depends on the model, training objective, existing dataset, and overlap with other inputs.[41] The U.S. Copyright Office likewise observes that identifying a rights holder and negotiating permission may cost more than the license is worth.[42]

These conditions favor agreements covering collections of works. Catalog holders can price access without assigning a separate value to every input. But those agreements require confidence that ordinary commercial provisions—including confidentiality obligations, limits on permitted activities, and access-method requirements—will remain enforceable.

Uncertainty about section 187 can raise the price of a license *ex ante* or prevent an agreement altogether. The United Kingdom's copyright and AI review found no case for legislative intervention in licensing terms.[43] Similar restraint should guide the interpretation of rules that could invalidate those terms.

2. A Functional Boundary Between Access and Use

Guidance should distinguish terms governing access from terms restricting computational analysis after lawful access. Genuine access conditions should remain enforceable. Restrictions on subsequent use should stand only to the extent they respect the CDA exception.

The inquiry should focus on what a term does. Does it govern whether and how a user may obtain material, or does it prohibit computational analysis of copies the user has lawfully obtained? Access-method requirements and server-load limits ordinarily fall in the first category. They govern delivery and protect the systems that provide access.

Personal-use restrictions can perform both functions. A provision defining who may obtain access can serve as an access condition. To the extent it also prohibits computational analysis of lawfully obtained copies, it must yield to the exception.

Confidentiality obligations ordinarily support commercial access and should generally remain enforceable. Parties price those obligations into their agreement. An incidental effect on training

---

[39] *See generally* Stout *et al.*, *supra* note 6.

[40] Oderinwale & Kazlauskas, *supra* note 4.

[41] *Id.* at 1–3.

[42] U.S. Copyright Office, *supra* note 15, at 104.

[43] Report on Copyright and Artificial Intelligence, *supra* note 26, § G, ¶¶ 41, 63–64.

should not alone invalidate a genuine confidentiality provision, although labeling a training prohibition "confidentiality" should not shield it from section 187.

This approach would preserve the statutory protection while giving parties greater confidence to negotiate. Outright exclusions of the CDA exception would fail. Genuine access conditions would stand. Terms combining both functions would remain enforceable only to the extent consistent with the statute. Although some cases would still require close analysis, a functional framework would reduce reliance on uncertain judgments about a drafter's intent.

B. Clarifying the Boundary Through Guidance

The consultation asks whether case law, legislative amendment, or nonbinding guidance should clarify the boundary between access conditions and restrictions on use.[44] Litigation takes time, and the consultation acknowledges that Singapore's limited copyright litigation may leave these questions unresolved for years.[45] Legislative amendment could fix a boundary before licensing practice has sufficiently developed.

IPOS should instead publish illustrative guidance explaining how section 187 applies to common contract terms. It should distinguish access conditions from restrictions on computational analysis and explain when a term "indirectly" restricts the CDA exception.

Comparative research on copyright and user contracts identifies model licenses and codes of conduct as useful responses when contractual terms conflict with statutory exceptions.[46] IPOS guidance could serve a similar purpose by helping parties draft enforceable access conditions that preserve the CDA exception.[47]

Guidance would also allow IPOS to respond as licensing arrangements and data-valuation methods develop.[48] It would preserve the distinction between legal obligations and the voluntary technical measures the consultation considers separately.[49] Those measures include the Robots Exclusion Protocol, through which website operators communicate automated-access preferences. The consultation appropriately proposes keeping such voluntary measures outside legislation.[50]

---

[44] Public Consultation, *supra* note 2, question A(II)(b)(i).

[45] *Id.* ¶ 39.

[46] Estelle Derclaye & Marcella Favale, *Copyright and Contract Law: Regulating User Contracts: The State of the Art and a Research Agenda*, 18 J. INTELL. PROP. L. 65, 138–39 (2010), https://digitalcommons.law.uga.edu/jipl/vol18/iss1/4.

[47] *Id.*

[48] Oderinwale & Kazlauskas, *supra* note 4, at 11–12.

[49] Public Consultation, *supra* note 2, ¶¶ 33–37.

[50] *Id.* ¶ 31.

The same measured approach should govern the contract boundary. Singapore should first clarify how its existing rules apply, allowing commercial practice and further evidence to inform any later legislative changes.

IV. Responsibility for Infringing Outputs

The consultation appropriately distinguishes the use of copyrighted inputs for training from infringement through downstream outputs.[51] Some courts have resisted treating that distinction as a bright-line rule.[52] Liability for infringing outputs should nevertheless turn on each party's conduct, knowledge, and practical ability to prevent infringement at reasonable cost. Placing obligations on those best positioned to distinguish lawful uses from infringement encourages effective precautions while preserving the benefits of general-purpose technology.[53]

Singapore's existing law provides a suitable framework. In *RecordTV*, the Court of Appeal assessed the alleged authorizer's control, relationship with the infringer, reasonable preventive steps, and actual or constructive knowledge—what it knew or reasonably should have known. The court emphasized that these factors require consideration together. The ability to prevent infringement by shutting down a service did not itself establish a duty to do so. It ultimately held that RecordTV had not authorized its users to infringe MediaCorp's copyright.[54]

Supplying a general-purpose AI tool capable of lawful uses should likewise be insufficient, by itself, to establish liability for users' infringing outputs. Liability becomes more plausible where a developer deliberately designs a product to reproduce protected expression without authorization or knowingly facilitates infringement in circumstances satisfying the authorization test. A developer also remains responsible for its own infringing acts.

A firm deploying an AI system may exercise substantial control through its choice of data sources, system instructions, and product features. Retrieval-augmented generation (RAG), which allows a model to draw on external material when answering requests, illustrates the point. A RAG system deliberately configured to deliver unauthorized copies presents a stronger case for liability than a general research assistant. Deliberately avoiding knowledge of users' infringement, or continuing to facilitate it after receiving credible evidence, should also inform the assessment.[55]

Reasonable safeguards should weigh strongly against authorization liability. Depending on the service, these might include controls on prompts seeking unauthorized reproduction, limits on access

---

[51] Kristian Stout, Geoffrey A. Manne & Emily Corbeille, *ICLE Comments on Artificial Intelligence and Copyright* § III.B, at 21–24 (Int'l Ctr. for L. & Econ., Oct. 30, 2023).

[52] Stout *et al.*, *supra* note 6, at 794–95 (discussing *Kadrey v. Meta*).

[53] Stout & Colangelo, *supra* note 10, § II.

[54] *RecordTV Pte Ltd v. MediaCorp TV Singapore Pte Ltd*, [2010] SGCA 43, ¶¶ 50, 59–62 (Sing.).

[55] *Id.* ¶¶ 50, 56–59; *see also* Geoffrey A. Manne & Kristian Stout, *Passive Resistance and Contributory Liability*, TRUTH ON THE MKT. (June 22, 2015).

to protected material, and procedures for responding to credible infringement notices. Such precautions would not, of course, excuse a provider's own direct infringement.[56]

Users often control the conduct most directly connected to infringement. They choose their requests, may direct the reproduction of existing works, and decide how to use or distribute the results. Liability should ordinarily follow that infringing conduct. Intentional wrongdoing presents a particularly clear case, but infringement and culpability remain distinct questions.[57] Singapore's protection against damages for innocent infringement reflects that distinction. It addresses what a user knew or reasonably could have known without eliminating other remedies.[58]

The same principles should apply when users employ AI agents—software that carries out tasks on their behalf. A person should not escape responsibility for directing unauthorized copying merely because software performs the intermediate steps. Applying the *RecordTV* framework requires examining who sets the agent's objectives, selects its sources, grants access and publication permissions, and can review or stop its actions.

A user who instructs an agent to reproduce and distribute protected works presents a strong case for responsibility. Unexpected conduct outside the user's instructions should not, by itself, establish authorization liability. Where the deployer retains the relevant control and the user supplies only a broad objective, the deployer may bear greater responsibility. The assessment should follow the parties' actual conduct and control.[59]

Copyright policy should place obligations where they can most effectively reduce infringement, with due regard for the cost of restricting lawful activity.[60]

V. Conclusion

Singapore should clarify the boundaries of its copyright framework while preserving its core. The consultation confirms that the CDA exception covers generative AI, commercial development, and training.[61] The economic rationale remains sound. As training draws on more works, the costs of identifying rights holders, valuing individual contributions, and negotiating separate permissions become more consequential.

The exception and commercial licensing serve complementary functions. The exception enables training on lawfully accessed material at scale, while licensing markets price access to catalogs,

---

[56] *RecordTV*, [2010] SGCA 43, ¶¶ 50, 60–62.

[57] Stout *et al.*, *supra* note 51, § III.B.1, at 21–22; Copyright Act 2021, § 306 (Sing.).

[58] Copyright Act 2021, § 306 (Sing.).

[59] *RecordTV*, [2010] SGCA 43, ¶¶ 50, 59–62.

[60] Kristian Stout & Geoffrey A. Manne, *The Law & Economics of Online Intermediary Liability* 1–2 (Int'l Ctr. for L. & Econ., Oct. 2021).

[61] Public Consultation, *supra* note 2, ¶¶ 22–28.

curated datasets, and related services.[62] Those markets also allow creators and aggregators to negotiate how they share revenue from training uses. Policymakers should give these arrangements time to develop before imposing compulsory licensing, levies, or additional compliance obligations.

IPOS should clarify the contract boundary through practical guidance. Genuine access conditions should remain enforceable, while restrictions on computational analysis of lawfully obtained copies should yield to the CDA exception. Guidance can help parties distinguish these functions and adapt as licensing practices mature.[63]

Responsibility for infringing outputs should follow conduct, knowledge, and practical control under the Court of Appeal's existing authorization framework. Reasonable safeguards should weigh against authorization liability, while leaving responsibility for direct infringement intact. The same principles should govern developers, deployers, users, and those directing AI agents.[64]

These recommendations require no legislative amendment. Guidance, existing legal principles, and evidence from developing markets provide a proportionate response to the consultation's concerns.[65] Singapore should make its framework more predictable and allow experience to establish whether further intervention is necessary.

---

[62] Oderinwale & Kazlauskas, *supra* note 4, at 7–8, 17–18; U.S. Copyright Office, *supra* note 15, at 104–06.

[63] Copyright Act 2021, §§ 187, 243–244 (Sing.); Public Consultation, *supra* note 2, ¶¶ 30–32.

[64] *RecordTV*, [2010] SGCA 43, ¶¶ 50, 59–62.

[65] Public Consultation, *supra* note 2, ¶¶ 41–50.