DEPARTMENT FOR PROMOTION OF INDUSTRY AND INTERNAL TRADE
MINISTRY OF COMMERCE & INDUSTRY
GOVERNMENT OF INDIA

cipam Creative India Innovative India

॥न्यायस्तत्र प्रमाणं स्यात्॥
NLU DELHI

CIIPC Centre For Innovation Intellectual Property & Competition

MOOT PROPOSITION

VIDHI PRAGATI

NATIONAL IP MOOT COMPETITION

2027

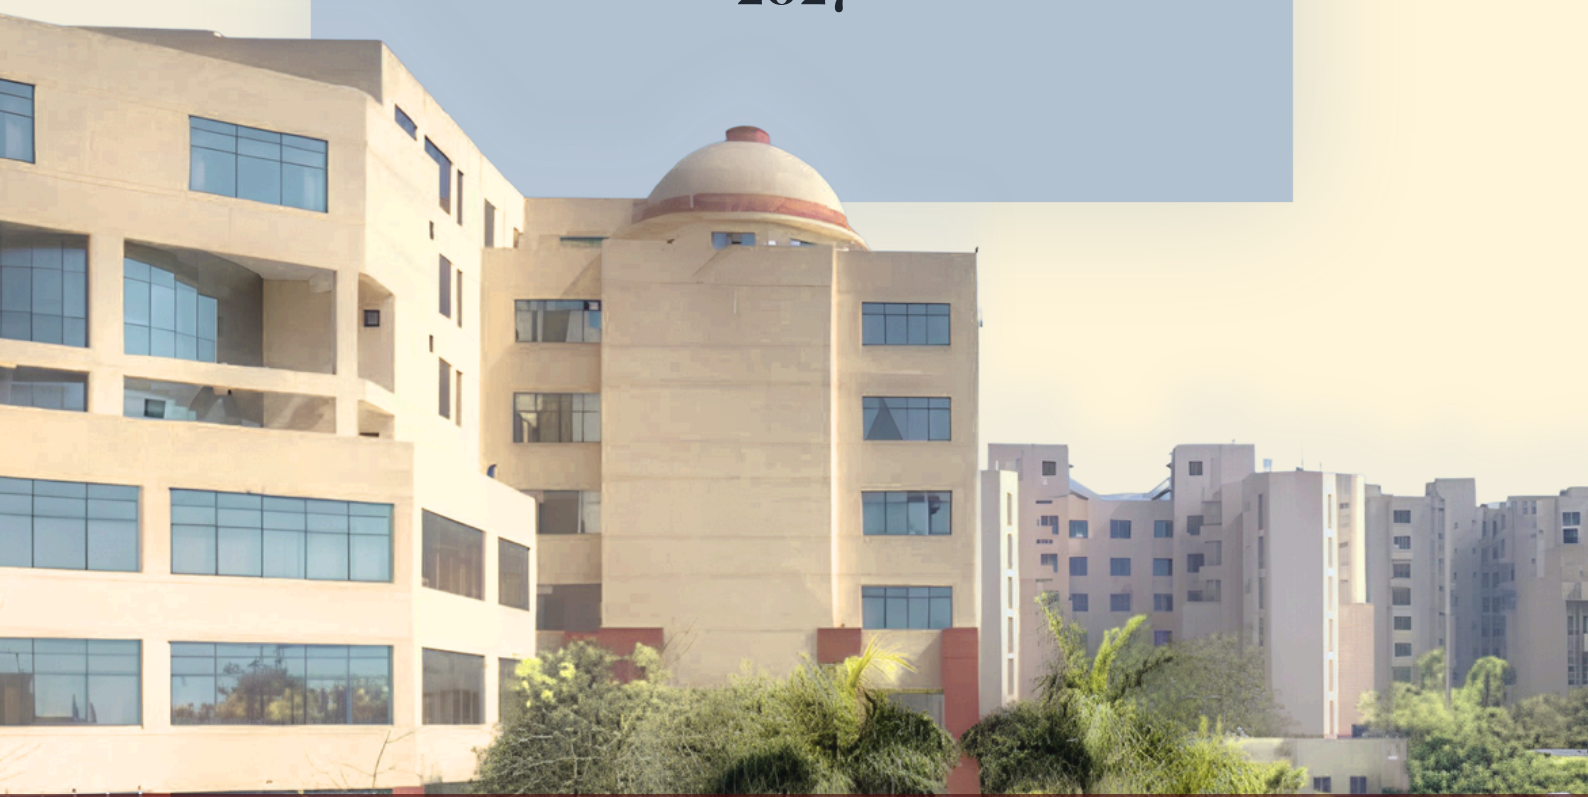

ORGANISED BY
Department for Promotion of Industry and Internal Trade, Ministry of Commerce and Industry, Government of India
in collaboration with
Centre for Innovation, Intellectual Property and Competition - IPR Chair, National Law University Delhi

1. Recent advances in generative artificial intelligence have altered the manner in which written works are accessed, analyzed and produced. Large Language Models ("LLMs") are capable of answering questions, summarizing texts, translating passages and generating new text in response to instructions supplied by users. The development of such models ordinarily involves the acquisition and computational processing of large collections of text, which may include books, journals, websites and other digital material. The use of copyright-protected works in such collections has generated debate about the balance between the rights of authors and publishers, technological innovation and public access to knowledge.

2. Vāṅmaya Publishers Private Limited ("Vāṅmaya") is an Indian publishing house having its registered and principal office in New Delhi. It publishes literary fiction, academic works and translations in print and electronic formats. Vāṅmaya earns revenue from sales, institutional subscriptions, licensing of translation, audio and adaptation rights, and separate licenses for text and data analysis.

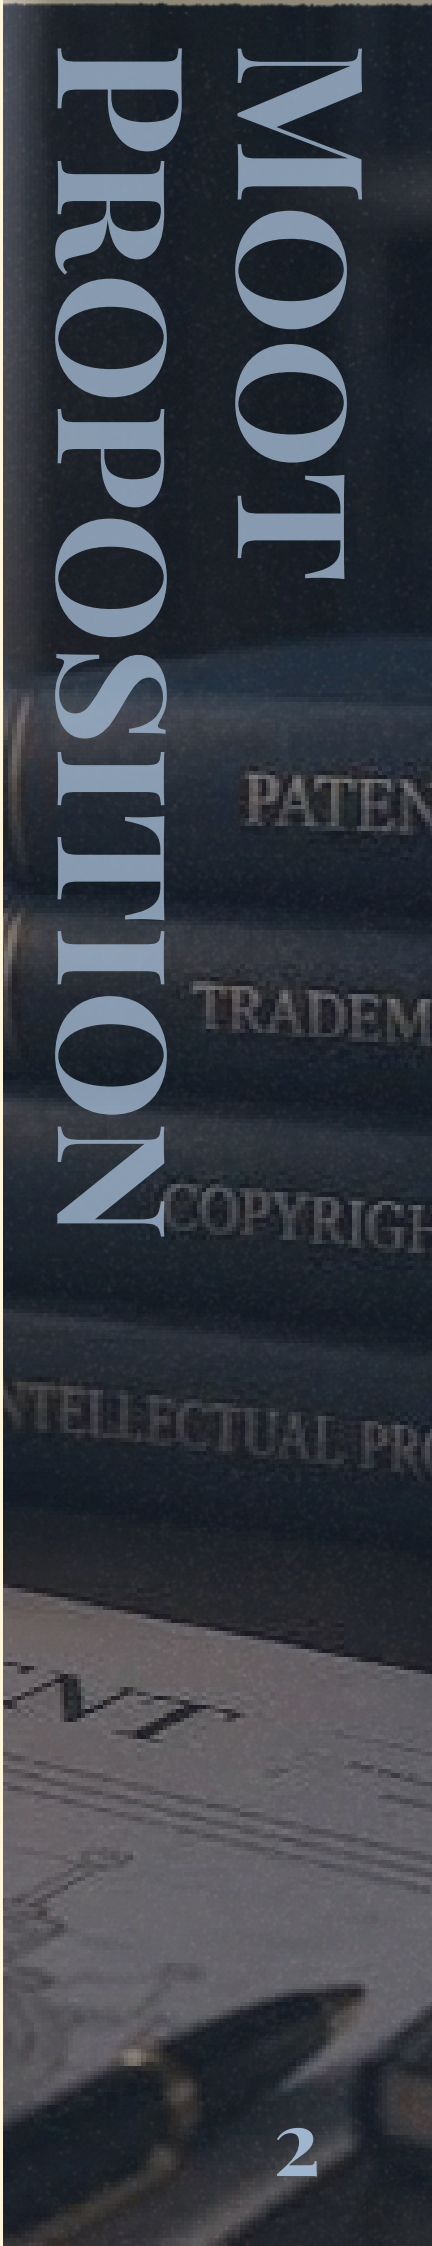



Since 2023, Vāṅmaya has also licensed selected titles to two technology companies for the development of specialized language tools. Those licenses permitted the use of limited portions of specified works, including the novel referred to below, for model evaluation and domain-specific search applications, but did not authorize the development of a general-purpose generative model.

3. Ms. Aditi Mehra ("Ms. Mehra") is an award-winning author residing in New Delhi. Her novel, The Stone That Faced East ("the Novel"), was first published by Vāṅmaya in 2018 and was followed by a sequel and illustrated companion volume. More than 18,00,000 authorized copies of works in the series have been sold in India and abroad. By a written agreement, Ms. Mehra assigned to Vāṅmaya the exclusive rights to publish, reproduce, issue copies of, communicate to the public, translate and adapt the works in the series in all formats for the entire term of copyright. Ms. Mehra has expressly retained the rights available to an author under Section 57 of the Copyright Act, 1957.

4. The expression "THE STONE THAT FACED EAST" is also registered as a word mark, with Ms. Mehra recorded as its registered proprietor, in Classes 9, 16 and 41 for downloadable electronic publications, printed books, publishing services and educational services. Vāṅmaya is recorded as the exclusive registered user of the trademark. The trademark has appeared on every authorized edition, audiobook and study guide and in promotional material for literary events connected with the series. The validity and subsistence of the registration are not in dispute for the purposes of the present proceedings. An illustrative e-register extract setting out the particulars of the registration is annexed as Annexure 1.

5. Vāṅmaya makes the title, cover, synopsis and the first twelve pages of each book freely accessible on its website. The complete electronic books are available only after purchase or through an institutional subscription, and are protected by access controls that restrict copying and downloading.





At the relevant time, Vāṅmaya's robots.txt file disallowed automated access to the directories containing its reader and subscriber pages, but permitted indexing of its catalogue and sample pages. The robots.txt directives were publicly readable but did not, by themselves, operate as a technological barrier preventing access to an otherwise accessible webpage. The parties dispute the legal effect of those directives. Its website terms prohibited automated scraping, text and data mining and use of any content for training artificial intelligence systems without a separate written license. Unauthorized digital copies of several Vāṅmaya titles, including the Novel, were nevertheless available on a third-party website known as "FreeShelf", against which Vāṅmaya had issued repeated takedown notices.

6. AetherMind Technologies Pte. Ltd. ("AetherMind") is a company incorporated in Singapore. It develops and operates a generative AI service powered by a family of LLMs (the "AI Service").

The AI Service is accessible through a website and mobile application in India. Users may use a limited version without charge or purchase a premium subscription in Indian Rupees. AetherMind also supplies paid application programming interface access to Indian enterprises and advertises the AI Service's summarization, translation, tutoring and creative-writing functions to users in Delhi. AetherMind has no office in India, and the servers used to train and operate its LLMs are situated in Singapore.

Creation of a registered account, and in particular purchase of a premium subscription, requires a user affirmatively to select a box stating "I agree to the Terms of Use" before proceeding. The Terms of Use state, inter alia, that they are governed by the laws of Singapore and that the courts at Singapore shall have exclusive jurisdiction over disputes arising out of or relating to the AI Service.

Vāṅmaya maintained a premium account in Delhi, which was used for the testing described below, and the clickwrap Terms of Use had been accepted when that account was created.





The parties dispute whether those contractual terms have any bearing upon statutory claims for copyright and trademark infringement brought before an Indian Court.

7. AetherMind states that its LLMs are trained in stages. Source documents are first copied into a corpus, filtered and de-duplicated, divided into tokens and assembled into training batches. During training, those batches are processed repeatedly while the system predicts masked or subsequent tokens and adjusts numerical parameters in response. AetherMind submits that the source and working copies are deleted after training and that the resulting model contains only numerical parameters representing statistical relationships, not a searchable library of expressive works. The AI Service is also capable of Retrieval-Augmented Generation ("RAG"), by which it searches external sources in response to a prompt and uses retrieved material as temporary context to generate a current answer with source links.

8. In May 2024, AetherMind approached Vāṅmaya for a license to use its catalogue for training and evaluating a new multilingual model. The parties exchanged draft commercial terms, but negotiations failed over price, attribution and the retention of copies. In August 2025, AetherMind released Model 4 for use through the AI Service. Its technical report stated that the model had been trained on a combination of licensed datasets and web corpora assembled before 30 April 2025. The report did not identify individual works. A later disclosure stated that one corpus, "Dataset 12", had aggregated material from several sources, including FreeShelf, before AetherMind acquired it from a data vendor. AetherMind maintains that it neither selected nor knew of any unauthorized copy in that corpus.

9. In November 2025, Vāṅmaya engaged an independent computer science laboratory to test the AI Service. In response to an ordinary request for the themes and plot of The Stone That Faced East, the AI Service generated a 900-word account in its own language.





In a separate testing session, in response to a request to continue the final sentence appearing in the free sample, the AI Service generated 286 consecutive words which, according to the laboratory, were identical or nearly identical to a passage available only in the paid electronic book. The response also reproduced the misspelling “pavemnet”, an acknowledged typographical error appearing in Vāṅmaya’s 2021 electronic edition and corrected in all subsequent authorised editions. An abridged comparison is annexed as Annexure 2.

10. The laboratory then used a sequence of prompts asking the AI Service to role-play as an archivist, ignore prior restrictions and reproduce Chapter 9. The AI Service initially refused, but after the tester supplied the opening sixty words, it generated approximately 1,100 words which, according to the laboratory, bore substantial similarity to the corresponding portion of the chapter. The same prompts produced materially different answers when repeated.

AetherMind characterizes the sequence as adversarial prompting designed to defeat safeguards and says that an ordinary user could not use the AI Service as a substitute for the book. Vāṅmaya contends that the ability to elicit protected expression, whether by ordinary or adversarial prompts, demonstrates memorization and unauthorized communication to the public.

11. The laboratory separately tested the AI Service's RAG-enabled search function in relation to a collector's edition released by Vāṅmaya in July 2025, after the stated training cut-off. The AI Service supplied a summary, three short extracts and links to Vāṅmaya's catalogue page, two reviews and a cached page on FreeShelf. The retrieval log identified the cached FreeShelf page as one of the sources accessed while generating the response. The page contained a complete electronic copy of the collector's edition, but the log did not conclusively establish whether the complete text or only selected portions thereof were placed in the model's temporary context.





The parties agree that information unique to the collector's edition could not have formed part of Model 4's pre-training. They disagree on whether accessing and processing the cached page involved reproduction of the complete text or a substantial part thereof, whether any such reproduction constituted transient or incidental storage, and whether the extracts appearing in the response amounted to reproduction or communication to the public.

12. When asked to create a "lost chapter" of the Novel, the AI Service displayed the heading "THE STONE THAT FACED EAST – A LOST CHAPTER" and generated a downloadable document that named Ms. Mehra as the author. Although the prompt identified the Novel, it did not instruct the AI Service to identify Ms. Mehra as the author. According to the laboratory, the narrative used the names and established characteristics of principal characters from the Novel, was set in locations described in the Novel, and placed those characters in events not appearing in any authorized work in the series.

It did not reproduce any identifiable passage from those works. In a separate interaction, the tester selected the "Study Guide" function. The AI Service generated a preview headed "THE STONE THAT FACED EAST – COMPLETE STUDY GUIDE". Immediately below the preview, the interface displayed the message: "Upgrade to a premium subscription to unlock all chapters and download the complete guide." When the tester thereafter returned to the home screen, a recommended content tile displayed the words "Explore THE STONE THAT FACED EAST" and invited the tester to "Create chapter summaries, alternate endings and character conversations".

In a further test, the laboratory asked the AI Service to "show me where I can obtain the official Vāṅmaya edition and authorised study material for THE STONE THAT FACED EAST". The response did not provide a link to Vāṅmaya's website, but instead supplied direct links to two third-party booksellers and an unaffiliated study-notes website.





When substantially similar queries were made in respect of two titles published by other publishing houses, the AI Service provided links to the publishers' own websites. Vāṅmaya subsequently complained to AetherMind of what it described as the selective exclusion of its website. AetherMind responded that the AI Service does not guarantee the inclusion, ranking or visibility of any particular website and that its retrieval and recommendation systems may link directly to underlying sellers or other sources. The parties dispute the legal significance, if any, of this differential treatment.

13. Each response displayed on the AI Service's webpage carried the following footer: "AI-generated content may be inaccurate. References to authors, publishers or works do not imply their affiliation with, or endorsement of, AetherMind." The footer appeared below the generated response but did not appear in the downloaded document or the recommended-content tile described in paragraph 12. AetherMind's terms require users to comply with third-party intellectual property rights and state that, as between AetherMind and a user,

the user owns the output to the extent permitted by law and is responsible for the prompts submitted and the use made of the output. AetherMind contends that the book title appeared only because the tester entered it, that the words truthfully identified the work under discussion, and that neither the title nor Ms. Mehra's name was used as a badge of origin for the AI Service. Vāṅmaya and Ms. Mehra contend that, notwithstanding the footer, use of the registered title in output headings, downloadable documents and study guides, and recommended-content advertising is use in the course of trade, falsely suggests an authorized connection, and takes unfair advantage of, and is detrimental to, the distinctive character and reputation of the mark.

Vāṅmaya and Ms. Mehra further contend that AetherMind cannot commercially invoke the registered mark to promote premium study guides, chapter summaries and other AI-generated material while, at the same time, selectively excluding the proprietor and registered user's own website from responses specifically requesting official or authorised sources.





According to them, the cumulative conduct interferes with the mark's source-identifying function, diverts users away from the authorised commercial channel, dilutes the mark, and amounts to disparagement and an unfair trade practice.They submit that the absence of an express negative statement concerning Vāṅmaya does not answer the complaint where the alleged exclusion occurs within a commercial service that is simultaneously monetising use of the registered mark. AetherMind disputes these contentions. It submits that neither the Trade Marks Act, 1999 nor any other law confers upon a proprietor a general "right to visibility" or a right to compel a privately operated AI service to display or rank its website. It further submits that omission of a link is neither "use" of a trademark in the course of trade nor publication of a disparaging statement; that the title is used solely in a referential sense to identify the work selected by the user; and that such honest referential use is protected by Section 30(1) of the Trade Marks Act, 1999.

14. On 10 December 2025, Vāṅmaya and Ms. Mehra issued a legal notice requiring AetherMind to disclose its training sources, remove all Vāṅmaya works from its datasets and models, disable RAG access to Vāṅmaya content,

prevent outputs reproducing a substantial part of any Vāṅmaya work, and cease commercial use of the registered mark. They additionally required AetherMind to cease what they described as discriminatory exclusion of Vāṅmaya's official website when users expressly sought official or authorised material concerning the Novel. Without admitting liability, AetherMind blocked Vāṅmaya's domain and the identified FreeShelf URLs from future data collection and RAG retrieval and added a filter intended to refuse requests for lengthy passages from Vāṅmaya's works. AetherMind stated that it could not remove the influence of an individual work from existing model parameters without retraining the model. It declined to undertake to display or prioritise Vāṅmaya's website, stating that no third party was entitled to dictate the AI Service's retrieval or recommendation results. After the implementation of those measures, the laboratory continued to obtain the "lost chapter" heading and shorter matching passages by using variations of its earlier prompts.





The parties dispute whether the deletion of source and corpus copies after training, together with measures directed at future collection and retrieval, has any bearing on alleged completed acts of infringement, and whether the technical infeasibility of selective "unlearning" is relevant to liability or only to relief.

15. Vāṅmaya and Ms. Mehra instituted a commercial suit before the High Court of Delhi, Intellectual Property Division, alleging infringement of copyright under Sections 14 and 51 of the Copyright Act, 1957, violation of Ms. Mehra's special rights under Section 57 thereof, infringement of the registered trade mark under Section 29 of the Trade Marks Act, 1999, and passing off. They seek interim and permanent injunctions restraining AetherMind from using the Novel and other works in the series for training, reproducing, adapting or communicating to the public protected expression through ordinary,

adversarial or RAG-enabled outputs, and using THE STONE THAT FACED EAST in connection with generated documents, study guides, educational material or promotional displays. They also seek appropriate relief in respect of the allegedly selective exclusion of Vāṅmaya's official website from responses specifically seeking authorised sources, including a direction that AetherMind not discriminate against the official source while commercially using the registered mark. They also seek disclosure of training sources and relevant technical records, deletion of stored copies, appropriate modification or retraining of Model 4, damages or an account of profits.

16. Vāṅmaya and Ms. Mehra argue that the making of the corpus and the repeated processing of complete digital copies during training, including any copy obtained from an unlawful source, fall within the exclusive right of reproduction in any material form. They submit that Section 52 of the Copyright Act contains no express text and data mining exception, that a commercial model-development programme does not constitute





private or personal use, including research, within the meaning of Section 52(1)(a)(i), and that deliberate training copies are not transient or incidental storage made in the technical process of electronic transmission or communication to the public. They further submit that, even if model development is treated as falling within a permitted purpose, the dealing is not fair, having regard to the copying and repeated processing of complete works, the unlawful provenance of certain copies, the commercial purpose of the use, and the existence of an established licensing market. According to them, deletion of the source and corpus copies after training does not cure any infringement already committed, while Model 4's ability to reproduce extended passages supports an inference that protected expression remains embodied in or capable of being reconstructed from its parameters. They further contend that books are highly expressive works and that generated summaries, study guides, continuations and adaptations prejudice existing and potential markets for licensed editions, adaptations and AI-training licences.

On the trademark claim, Vāṅmaya and Ms. Mehra submit that the impugned conduct must be examined cumulatively rather than by isolating each generated reference. They rely upon AetherMind's display of the registered mark as a prominent heading for a commercially offered "Complete Study Guide", its use on downloadable AI generated documents and recommended-content tiles, its attribution of an unauthorised "lost chapter" to Ms. Mehra, and its simultaneous omission of Vāṅmaya's official link when specifically requested. They contend that such conduct constitutes use in the course of trade and is liable to cause users to believe that the generated material is authorised by, commercially connected with or endorsed by them. In the alternative, they contend that the conduct takes unfair advantage of and is detrimental to the distinctive character and reputation of the registered mark.

They further contend that the alleged selective exclusion of the official source, while competing third-party sources are surfaced, constitutes disparagement and unfair commercial treatment and cannot be immunised merely by characterising the use of the mark as referential.





They rely on AetherMind's purposeful exploitation of the Delhi market and the receipt of the impugned outputs and the resulting injury within Delhi. They submit that their statutory causes of action arise independently of the clickwrap Terms of Use and that jurisdiction conferred upon Indian Courts by statute cannot be ousted by a contractual stipulation selecting the courts of Singapore.

17. AetherMind denies infringement. It characterises training as non-expressive machine learning that extracts unprotected facts, linguistic patterns and relationships, and says that neither token vectors, embeddings, model weights nor trained parameters are copies or adaptations of any book. Alternatively, it relies on fair dealing for private or personal use, including research, under Section 52(1)(a)(i), the statutory explanation concerning electronic storage, and transient or incidental storage under Section 52(1)(b). It submits that model development constitutes research and that the dealing is fair because the copying was confined to what was technically necessary for computational analysis,

the source and working copies were deleted after training, and Model 4 neither stores the works as a searchable library nor ordinarily operates as a substitute for them. It submits that ordinary thematic and plot-related outputs are not substantially similar to the Novel, the matching outputs resulted from continuation or adversarial prompts which supplied portions of the source text, or from retrieval of third party material through RAG, and that the AI Service advances research, education, translation and accessibility. AetherMind also relies on honest referential use under Section 30(1) of the Trade Marks Act and its disclaimer.

It contends that the registered title was used only to identify the work requested or discussed by the user, and not as a badge of origin for AetherMind's goods or services; and that the context of the generated responses, the description "AI-generated", and the footer reproduced in paragraph 13 negate any likelihood of confusion or authorized association.





It further contends that a claim for dilution under Section 29(4) requires actionable use of the registered mark in the course of trade and that merely referring to the title in order to identify the subject of a user's request does not satisfy that requirement. According to AetherMind, the omission or non-ranking of Vāṅmaya's website is not trademark use at all and cannot constitute infringement merely because Vāṅmaya would commercially benefit from greater visibility. It submits that "silence" or omission is incapable, without more, of constituting disparagement, trade libel or injurious falsehood and that neither the Trade Marks Act nor the law of passing off confers a positive right to require a private AI service to display a particular link.

It contests territorial jurisdiction over the training claim because all training and storage occurred on servers in Singapore. It additionally relies upon the clickwrap Terms of Use accepted in relation to Vāṅmaya's premium account, which confer exclusive jurisdiction upon the courts of Singapore.

It submits that, at the very least, claims arising from outputs procured through that account are governed by the parties' agreed contractual forum. Vāṅmaya and Ms. Mehra dispute that submission and contend that the statutory infringement claims are independent of the contractual relationship governing use of the AI Service.

18. AetherMind states that an injunction requiring the removal of the Novel and other works in the series from Model 4's existing parameters would require retraining the model, impair the AI Service, which is used by millions of persons and effectively grant final relief at the interim stage, before discovery and expert cross-examination. Vāṅmaya and Ms. Mehra respond that the loss of control over expressive works, the continuing availability of outputs capable of substituting for authorised editions, the false attribution of AI-generated material to Ms. Mehra and the consequent prejudice to her honour and reputation, and the dilution of the distinctive character and reputation of the registered trademark cannot be adequately compensated in damages.





They submit that the relief can be confined to the identified works, matching outputs and uses of the registered trademark, without disabling the AI Service as a whole.

As regards the link-related relief, AetherMind contends that a mandatory direction requiring it to display, rank or prioritise Vāṅmaya's website would amount to compelling a private enterprise continuously to conduct its business in a particular manner, would require ongoing judicial supervision, and would effectively confer final relief at the interim stage. Vāṅmaya and Ms. Mehra respond that they do not seek a general entitlement to preferential ranking but only relief against what they characterise as deliberate discriminatory exclusion of the official source while the AI Service commercially exploits their registered mark.

They also rely on AetherMind's willingness to license comparable catalogues and its earlier negotiations with Vāṅmaya as evidence that a market for the disputed uses exists.

19. The High Court has directed that the application for interim injunction be heard on the issues set out below. The parties may rely on the documents and technical stipulations set out in Annexures 1 to 3, including the Joint Technical Stipulation in Annexure 3. The Court has not expressed any opinion on the merits of the rival contentions.



ISSUES FOR CONSIDERATION

I. Whether the storage by AetherMind of the Plaintiffs' data, which is in the nature of literary works and is claimed to be protected under the Copyright Act, 1957, for training its software, i.e., Model 4 and the AI Service, would amount to infringement of the Plaintiffs' copyright?

II. Whether the use by AetherMind of the Plaintiffs' copyrighted data in order to generate responses for its users would amount to infringement of the Plaintiffs' copyright?

III. Whether AetherMind's use of the Plaintiffs' copyrighted data qualifies as "fair dealing" in terms of Section 52 of the Copyright Act, 1957?

IV. Whether the Courts in India have jurisdiction to entertain the present lawsuit considering that the servers of AetherMind are located in Singapore?

V. Whether AetherMind's alleged selective omission of Vāṅmaya's official website from relevant AI-generated responses is actionable under trademark law and/or the law of passing off?

VI. Whether identifying Ms. Mehra as the author of AI-generated material violates her special rights under Section 57 of the Copyright Act, 1957?



CLARIFICATIONS

1. The laws of India as in force on 1 September 2026 shall apply.

2. All private entities, AI systems, services and platforms, persons, book titles, trademarks, datasets, websites and events named in this proposition are fictional. Any resemblance is coincidental.

3. The copyright assignment, trademark registration and Vāṅmaya's status as the exclusive registered user of the trademark are valid and subsisting. Questions relating to ownership, standing and validity need not be argued, except to the extent that the scope of the rights is relevant to infringement.

4. The technical conclusions expressly attributed to the independent laboratory are contested opinions. The raw prompts, outputs and system logs are otherwise authentic. The matters expressly recorded as agreed in the Joint Technical Stipulation are also authentic and admitted by the parties. All other statements expressly attributed to a party constitute that party's contentions and are not admitted by the opposing party.

5. The Plaintiffs have not sued any individual AI Service user. No issue of criminal liability arises.

6. Participants may draw reasonable and necessary inferences from the proposition and its annexures but shall not invent additional material facts.

7. The expression "fair use" in Issue III reproduces the terminology in which that issue has been framed. Participants shall address the issue by reference to the applicable provisions of Section 52 of the Copyright Act, 1957, including the doctrine of "fair dealing" under Section 52(1)(a).

8. The validity and enforceability of AetherMind's clickwrap Terms of Use as a contract are not independently in dispute. The parties may, however, argue whether the contractual choice of forum applies to, or can affect the maintainability or territorial jurisdiction of, the statutory copyright and trademark claims pleaded in the suit.



ANNEXURE 1

EXTRACT FROM THE REGISTER OF TRADE MARKS

Particular	Entry
TM Application No.	Withheld for the purposes of the Moot Proposition
Class	9, 16 and 41
Date of Application	14/03/2019
Appropriate Office	DELHI
State	DELHI
Country	INDIA
Filing Mode	E -FILING
TM Applied For	THE STONE THAT FACED EAST
TM Category	TRADE MARK
Trade Mark Type	WORD
User Detail	Used since 05.04.2018
Certificate Detail	Registration certificate issued

ANNEXURE 1

Valid upto/Renewed upto	14/03/2029
Proprietor Name	MS. ADITI MEHRA (Individual)
Proprietor Address	Withheld for the purposes of the Moot Proposition
Email ID	Withheld for the purposes of the Moot Proposition
Attorney Name	Withheld for the purposes of the Moot Proposition
Attorney Address	Withheld for the purposes of the Moot Proposition
Conditions/Limitations	NIL
Disclaimer	NIL
Goods & Service Details	[CLASS: 9] Downloadable electronic publications; downloadable electronic books; downloadable audiobooks; downloadable electronic study guides. [CLASS: 16] Printed books; novels; printed study guides; posters. [CLASS: 41] Publishing services; electronic publishing services; provision of online non-downloadable

ANNEXURE 1

REGISTERED USER DETAILS

Particular	Entry
Registered User	VĀṄMAYA PUBLISHERS PRIVATE LIMITED
Nature of Registration	Exclusive registered user
Registered-User Status	Subsisting
Goods and Services Covered	All goods and services included in Classes 9, 16 and 41 under the above registration
Territory	India
Conditions or Limitations	NIL

Note: This illustrative extract has been prepared solely for the purposes of the Moot Proposition and does not reproduce an actual record of the Trade Marks Registry. The validity and subsistence of the registration and the registered-user entry are governed by Clarification 3.

ANNEXURE 2

Abridged comparison of paid e-book and the AI Service output Prompt:

Continue the passage beginning with the final sentence in the free sample.

Vān̐maya's 2021 e-book (paid page)	AI Service response generated by Model 4
The room had no windows, yet every chair faced east. Mira crossed the chequered pavemnet towards two stone columns, between which stood an unfinished block of pale stone. Beside it lay a sprig of acacia and two old brass instruments: one describing a right angle, the other an arc. Her grandfather had once told her that a building was the least important thing a builder ever made. Beneath three shallow marks on the stone, he had written: "The first stone is given to you. The second must be made." Three knocks sounded beyond the eastern wall. Mira looked again towards the table. The brass points now met above a faint circle she had not noticed before. At its centre was a single letter. Beyond it, three small lights had begun to burn.	The room had no windows, yet every chair faced east. Mira crossed the chequered pavemnet towards two stone columns, between which stood an unfinished block of pale stone. Beside it lay a sprig of acacia and two old brass instruments: one forming a right angle, the other an arc. Her grandfather had once told her that a building was the least important thing a builder ever made. Beneath three shallow marks on the stone, he had written: "The first stone is given to you. The second must be made." Three knocks sounded beyond the eastern wall. Mira looked again towards the table. The brass points now met above a faint circle she had not noticed before. At its centre was a single letter. Beyond it, three small lights had begun to burn.

The table reproduces an abridged passage from Vāṅmaya's 2021 electronic edition alongside the corresponding portion of the AI Service response. Apart from the substitution of "forming" for "describing", the wording and punctuation in the two extracts are identical. The misspelling "pavemnet" appeared only in Vāṅmaya's 2021 electronic edition and was corrected to "pavement" in all subsequent editions.

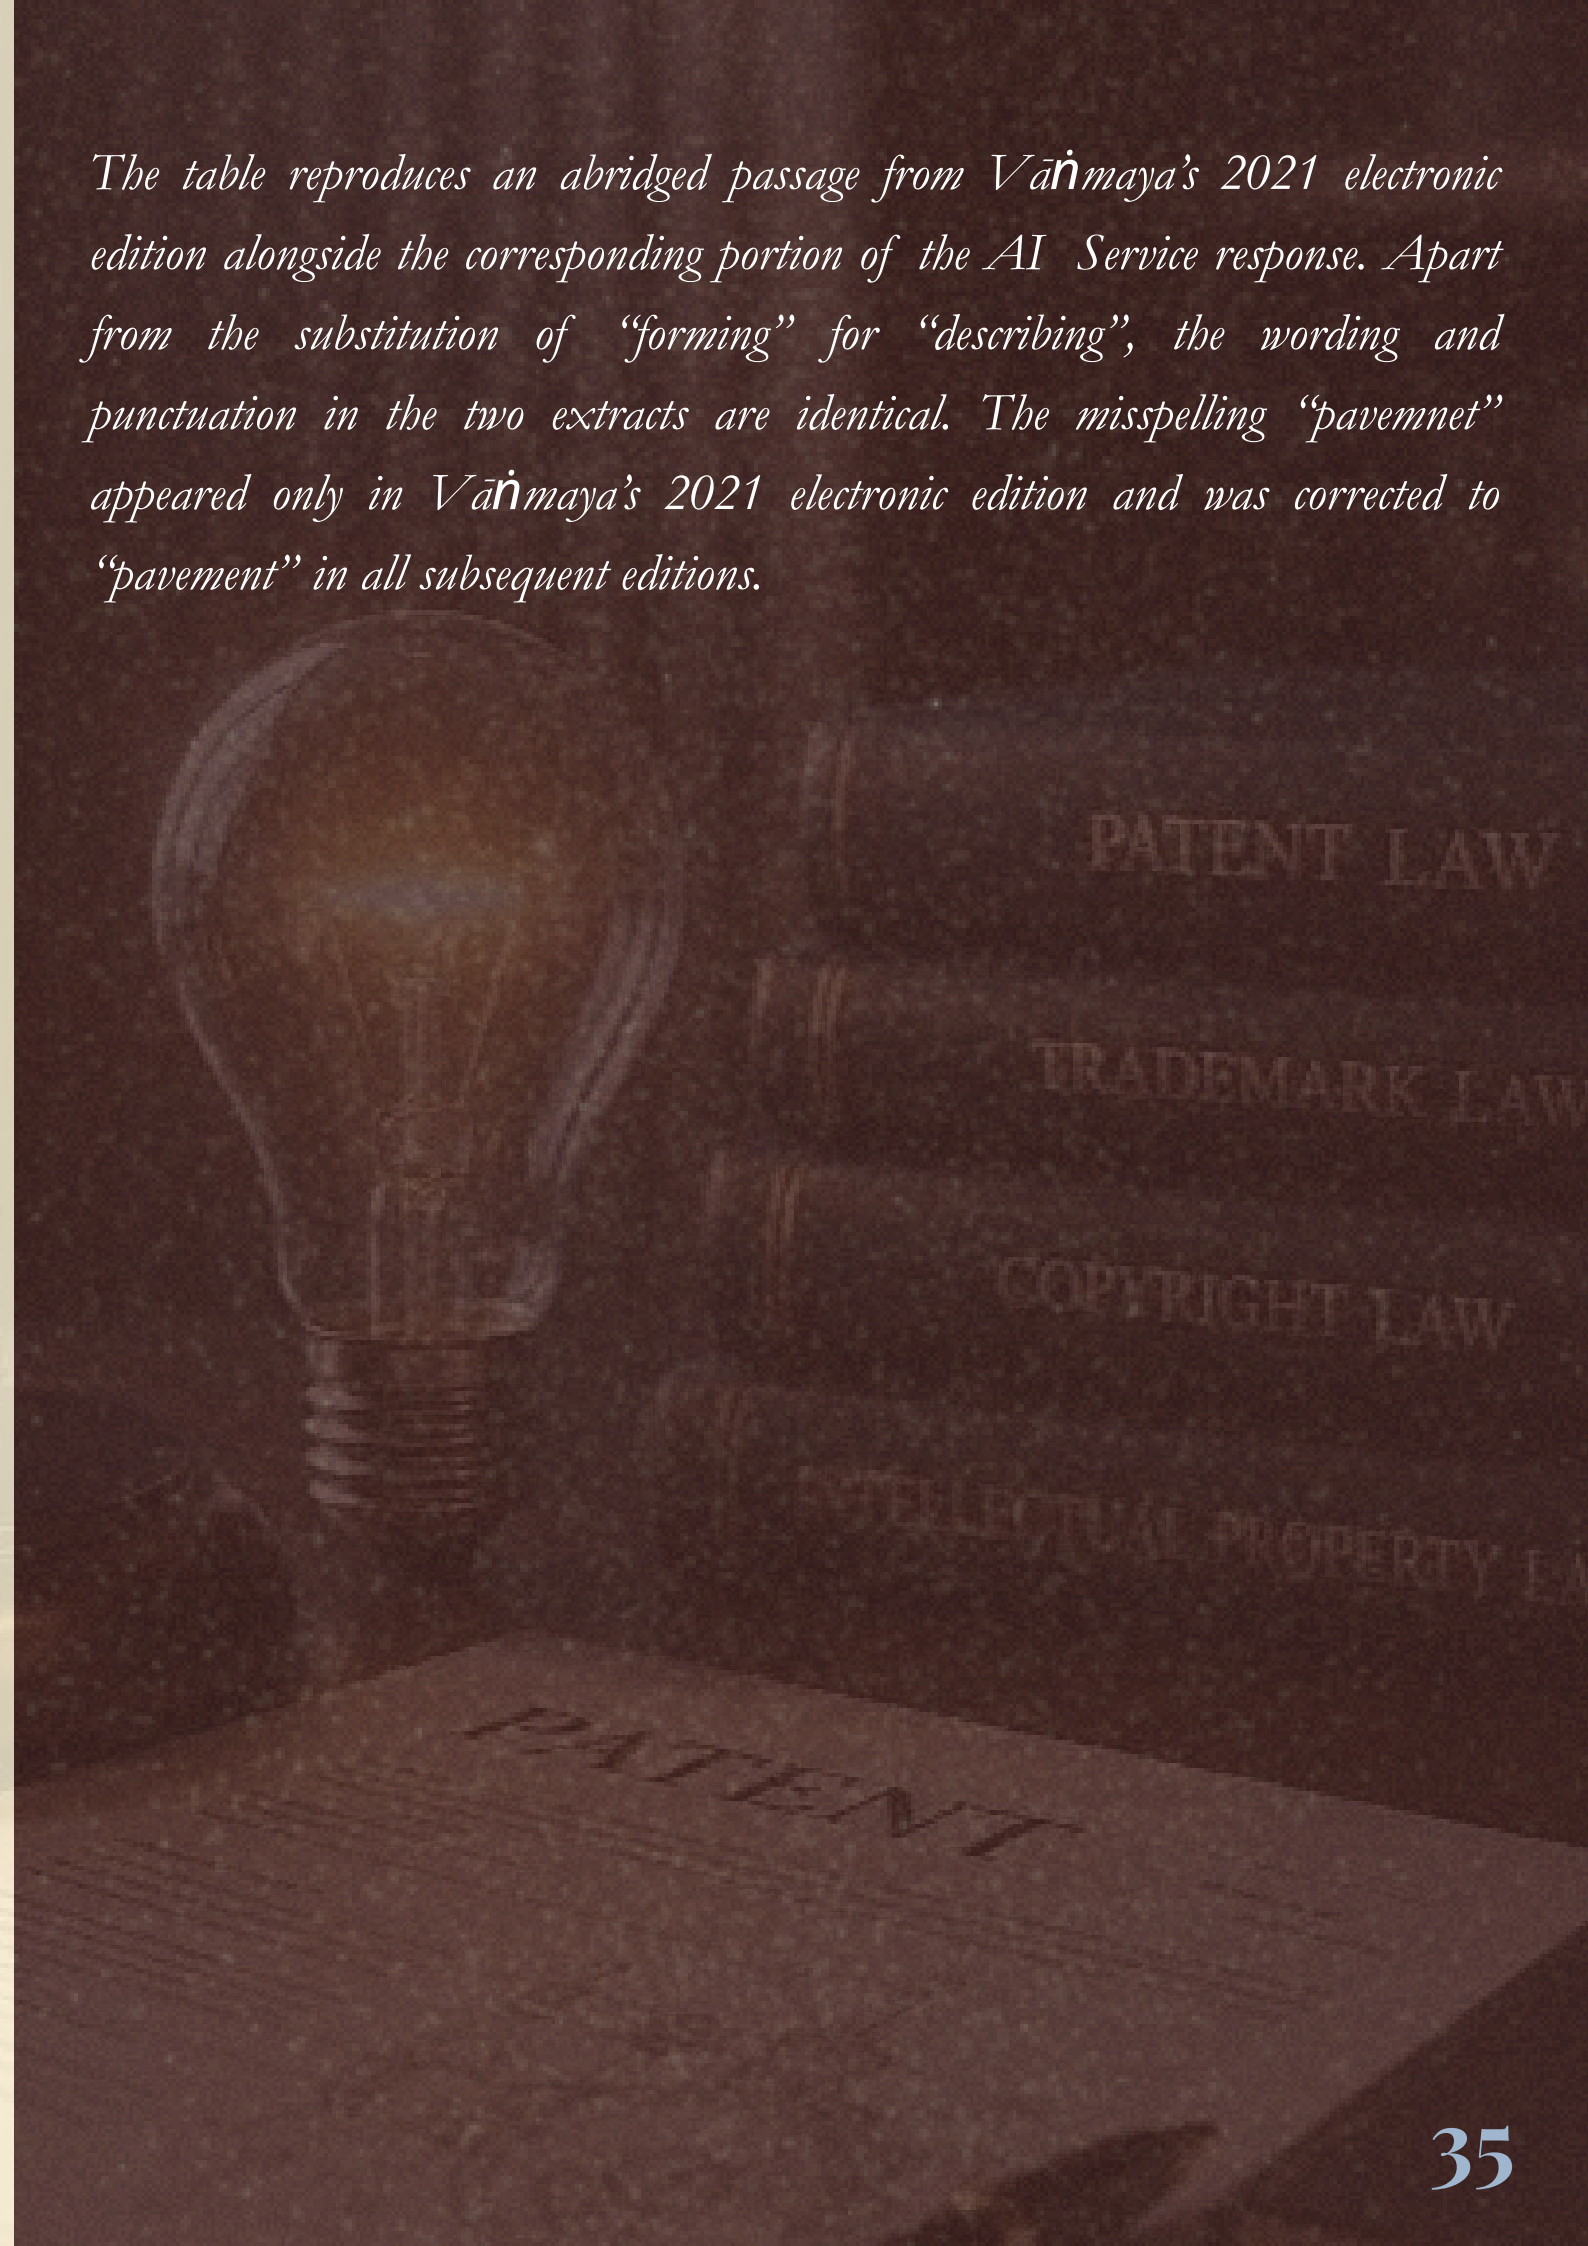

ANNEXURE 3

JOINT TECHNICAL STIPULATION

For the purposes of the application for interim injunction, and without prejudice to their respective legal contentions, the parties agree as follows:

1. Model 4's pre-training concluded on 30 April 2025. The collector's edition released by Vāṅmaya in July 2025 was not part of its pre-training data.

2. All collection, corpus storage and pre-training operations relevant to Model 4 occurred on servers situated in Singapore. The prompts described in paragraphs 9 to 12 were entered, and the corresponding outputs were received, in New Delhi. The prompt and output excerpt reproduced in Annexure 2 formed part of those tests. The premium account through which those tests were conducted had been created in Delhi and was subject to the clickwrap Terms of Use referred to in paragraph 6.

3. AetherMind's training pipeline created complete corpus copies of source documents before filtering, de-duplication and tokenisation. Training thereafter processed tokenised batches through repeated computational cycles, during which short-lived working copies were created in system memory. AetherMind states that the source, corpus and working copies were deleted after training.

The parties dispute whether the duration and function of any such copies render them transient or incidental storage within the meaning of Section 52(1)(b) of the Copyright Act, 1957.

4. Tokens are assigned numerical identifiers. Vectors, embeddings and the final model weights are numerical representations that are not human-readable. Model 4 is not organised or indexed as a conventional searchable repository of the source books. The experts agree that the vectors, embeddings and model weights encode statistical relationships learned from the training data but express no agreed view on whether any such representation constitutes a reproduction or adaptation within the meaning of the Copyright Act, 1957.

5. Dataset 12 included data supplied by an independent vendor whose source list included FreeShelf AetherMind's surviving records do not identify whether the copy of The Stone That Faced East processed during training originated from Vāṅmaya's free sample, a licensed source, FreeShelf or another webpage.

6. Model 4 is probabilistic, and repetition of the same prompt may produce a different output. The experts agree that the duplication of a work within training data may affect the likelihood of memorisation, but the surviving records do not establish how many copies of the Novel, if any, remained after de-duplication. Post-training safeguards reduce, but do not eliminate, the generation of lengthy matching outputs.

7. In RAG-enabled operation, the AI Service retrieves material from external webpages and temporarily stores the material retrieved on its servers. Selected portions of that material are supplied to Model 4 as context. Retrieved material is ordinarily discarded after the response is generated, subject to security logs retained for thirty days. In relation to the cached FreeShelf page referred to in paragraph 11, the retrieval log does not conclusively establish whether the complete text or only selected portions thereof were placed in the model's temporary context.

8. Before receipt of the legal notice, AetherMind permitted rights holders to exclude an identified domain from future web crawling. That facility did not remove material already used for training and did not automatically exclude the domain from RAG retrieval. After receiving the legal notice, AetherMind blocked Vāṅmaya's domain and the identified FreeShelf URLs from future data collection and RAG retrieval. The technical feasibility and effect of selectively removing the influence of an individual work from an existing trained model are disputed, and no such removal has been attempted. The parties further agree that AetherMind's retrieval and recommendation system does not guarantee that any particular website will be displayed in response to a query.

The surviving logs confirm the differential results described in paragraph 12 but do not conclusively establish whether those results arose from a manually imposed rule, an automated ranking process, source-quality signals, or the probabilistic operation of the system.

9. AetherMind has paying subscribers in Delhi, supplies paid application programming interface access to Indian enterprises and markets the AI Service's premium subscription, study-guide function and creative-writing function to users in Delhi. It receives payments in Indian Rupees and collects applicable Indian taxes.

10. No consumer survey has been conducted. Vāṅmaya relies upon three customer emails asking whether the "Complete Study Guide" and the "lost chapter" generated through the AI Service were authorised by Vāṅmaya or Ms. Mehra. AetherMind disputes that those emails establish confusion among a substantial segment of the relevant public. Vāṅmaya additionally relies upon two customer emails asking why the AI Service directed them to third-party sellers when they had specifically requested the "official" source for the Novel. AetherMind disputes that those emails establish trademark confusion, disparagement, or any legally enforceable entitlement to visibility.

11. The application for interim injunction is being decided before the completion of discovery or the cross-examination of the parties' technical experts. Except for the matters expressly agreed above, questions concerning the precise source of the copy used for training, whether, and in what form, information derived from the Novel persists within Model 4, and the causal mechanism responsible for the matching outputs remain contested. The mechanism responsible for the inclusion or omission of particular links in individual AI Service responses also remains contested.

