AI-ASSISTED CONTRACT REVIEW: LEGAL RISKS, DATA PRIVACY CHALLENGES AND THE NEED FOR RESPONSIBLE AI GOVERNANCE IN INDIA

Alpana Mishra, ILS Law College

ABSTRACT

Artificial Intelligence ("AI") is rapidly transforming commercial contract review. AI-enabled tools can assist legal professionals in summarising agreements, identifying contractual risks, comparing clauses against approved positions, extracting obligations and accelerating repetitive review activities. In increasingly technology-driven contracting environments, the use of approved AI tools is moving from experimentation towards integration into routine legal and commercial workflows.

The benefits of AI-assisted contract review, however, are accompanied by significant legal and governance concerns. Commercial agreements routinely contain confidential business information, personal data, intellectual property, pricing information, security requirements and commercially sensitive negotiated positions. Processing such information through AI systems may therefore raise questions relating to confidentiality, data protection, cybersecurity, accuracy, accountability and professional responsibility.

Against the backdrop of India's evolving digital regulatory framework, particularly the Digital Personal Data Protection Act, 2023 ("DPDP Act") and the Digital Personal Data Protection Rules, 2025 ("DPDP Rules"), this article examines the principal benefits and legal risks associated with AI-assisted contract review. It argues that organizations should move beyond broad statements of AI ethics and develop practical AI and data-privacy playbooks, risk-based review mechanisms and meaningful human oversight.

The objective should not be to prevent AI adoption in legal services. Rather, organizations should create governance structures that allow technological efficiency to operate within clearly defined legal, privacy and professional boundaries.

Keywords: Artificial Intelligence, Contract Review, Generative AI, Data Privacy, DPDP Act, AI Governance, Legal Technology, Contract Lifecycle Management, India.

1. Introduction: AI and the Changing Contracting Function

Artificial Intelligence is increasingly becoming part of everyday commercial contracting.

AI-enabled technologies can assist legal and contract-management professionals in reviewing agreements, summarising contractual provisions, identifying deviations from standard positions, extracting obligations and conducting preliminary risk assessments.

The attraction is understandable. Commercial contracts may contain dozens or even hundreds of pages addressing indemnification, limitation of liability, intellectual property, confidentiality, cybersecurity, data protection, termination, governing law, regulatory compliance and numerous commercial obligations. Reviewing these provisions manually requires significant professional time. AI can potentially perform some of the preliminary analytical work in seconds.

However, the speed at which an AI system can analyse a contract should not be confused with the reliability or legal sufficiency of its analysis. Contracts routinely contain confidential information, personal data, trade secrets, intellectual property, pricing arrangements and commercially sensitive negotiated positions. Introducing such information into an AI environment creates legal and governance questions extending considerably beyond productivity.

The central question for organizations should therefore not simply be whether AI can review contracts. It should be under what legal, privacy, technological and governance safeguards AI should be permitted to assist in contract review.

This distinction is important because responsible adoption of AI requires organizations to examine not merely what the technology can do, but what information it receives, how its output is verified, who remains accountable for decisions and what safeguards surround the process.

2. Benefits of AI-Assisted Contract Review

2.1 Efficiency and Scalability

Perhaps the most immediate benefit of AI-assisted contract review is efficiency. Depending

upon the technology being used, AI may assist professionals in summarising lengthy agreements, identifying relevant contractual provisions, extracting important dates and obligations, locating missing clauses, comparing contractual language with approved templates, identifying deviations from preferred positions and directing reviewers towards provisions requiring closer examination.

This can be particularly valuable for organizations processing substantial volumes of contracts. Importantly, the objective need not be replacement of the legal professional. A more sustainable model is AI-assisted legal review, under which technology performs appropriate extraction, comparison and preliminary analytical functions while professionals concentrate on interpretation, negotiation and consequential legal decisions.

2.2 Greater Consistency

Large organizations frequently have several professionals reviewing similar contractual provisions. Even where an established contracting playbook exists, reviewers may interpret risk positions differently.

An appropriately configured AI system may improve consistency by comparing provisions against predetermined parameters and systematically identifying deviations. However, consistency should not itself be confused with correctness. An AI system consistently applying an outdated or inappropriate contractual position merely produces consistent error. Human ownership of the underlying contractual playbook therefore remains essential.

2.3 Better Use of Institutional Knowledge

Organizations accumulate significant legal knowledge through previously negotiated contracts, approved clauses, fallback positions and historical risk decisions. AI can potentially make this institutional knowledge more accessible.

For example, rather than searching manually through numerous precedents, a professional may be able to identify previously approved contractual approaches more efficiently. This benefit, however, depends heavily upon the reliability and governance of the underlying information.

The quality of AI-assisted legal analysis cannot exceed the quality and appropriateness of the information, instructions and controls supporting it.

3. Legal Risks of AI-Assisted Contract Review

The benefits of AI should be considered alongside its risks. AI-assisted contract review potentially raises issues concerning confidentiality, personal data, intellectual property, cybersecurity, reliability of AI-generated analysis, third-party technology providers, retention of contractual information, accountability for incorrect recommendations and professional responsibility.

These risks do not mean that AI should not be used. Rather, they demonstrate why organizations require governance proportionate to the nature and consequence of each use case.

4. Confidentiality and Commercially Sensitive Information

Commercial contracts frequently contain information that an organization is legally or commercially obligated to protect. Such information may include pricing, intellectual property, customer information, business strategies, technical specifications, security requirements, trade secrets and negotiated contractual positions.

Submitting an agreement to an AI platform therefore involves more than simply asking whether the platform can accurately analyse the document. Organizations should understand whether contractual information is retained; whether submitted information may be used to train or improve AI models; where information is processed and stored; whether subprocessors receive access; what technical safeguards apply; how long information remains available; and how information is deleted.

These questions become particularly important when users have access to publicly available generative AI systems as well as enterprise-approved AI environments. The fact that an AI platform is widely available does not establish that it is appropriate for processing confidential contractual information.

This creates an important distinction between approved AI and uncontrolled AI use. Organizations adopting AI for legal functions should therefore clearly identify which systems may process contractual information and what categories of information may be submitted to those systems.

5. Data Privacy and the Indian Legal Framework

The privacy implications of AI-assisted contract review are equally important. These concerns must also be viewed in light of the constitutional recognition of privacy as a fundamental right in India.[8] The right to privacy has been recognized as an intrinsic part of the right to life and personal liberty under Article 21 of the Constitution.[9] Subsequent jurisprudence has reinforced the importance of proportionality and data-governance safeguards in digital ecosystems.[10] Contracts may contain personal data relating to employees, customers, vendors, directors, representatives and other individuals. Protection of such information is consistent with the broader constitutional recognition of informational privacy.[16] Processing those documents through an AI system may therefore involve processing digital personal data[15]. India's Digital Personal Data Protection Act, 2023 establishes the principal statutory framework for processing digital personal data.[1]

The Act applies to processing of digital personal data within India where such data is collected in digital form or subsequently digitised. It also has specified extra-territorial application where processing outside India relates to offering goods or services to Data Principals[13] within India.[2]

The DPDP Act places obligations upon Data Fiduciaries[14] concerning lawful processing and protection of personal data.[3]

Of particular relevance to organizations using technology providers, section 8 provides that a Data Fiduciary remains responsible for complying with the Act irrespective of an agreement to the contrary or a failure by a Data Principal to carry out specified duties.[4]

The Act also contemplates processing through Data Processors[12] on behalf of Data Fiduciaries. This has direct significance for AI-enabled contractual workflows. An organization cannot necessarily treat the involvement of an external technology platform as removing its own responsibility for the personal data being processed.[5]

5.1 Data Minimisation as a Practical AI Control

One useful governance question is whether the AI system actually needs all of the information contained in the contract. Suppose a contract reviewer requires AI assistance only to compare a limitation-of-liability provision against an approved contractual position. The names, telephone numbers, email addresses or signatures of individuals appearing elsewhere in the

agreement may have no relevance to that analysis.

Organizations should therefore consider whether information can be excluded, masked or redacted before it is processed. This is a practical example of how privacy principles can be translated into day-to-day AI governance.

5.2 Security Safeguards

The DPDP Act requires Data Fiduciaries to protect personal data in their possession or control, including processing undertaken on their behalf by a Data Processor, by taking reasonable security safeguards to prevent personal data breaches.[6]

The DPDP Rules, 2025 further develop the framework concerning reasonable security safeguards.[7] Organizations using AI platforms outside India should also consider statutory provisions relating to cross-border processing and transfers of personal data.[11]

AI governance therefore cannot operate independently of cybersecurity and privacy governance. An organization adopting AI-assisted contract review should understand the complete lifecycle of the information: contract, AI environment, processing, storage, AI output, human reviewer and retention or deletion. Mapping this lifecycle can expose risks that may not be apparent when an AI system is viewed merely as a productivity tool.

6. AI Hallucinations, Accuracy and Legal Risk

One of the most distinctive risks associated with generative AI is its ability to produce information that appears authoritative even when inaccurate. This is commonly described as an AI "hallucination." In legal contracting, however, the problem is broader than completely fabricated information.

An AI-generated conclusion can be legally dangerous even when it is partly correct. Consider an AI summary stating that a customer may terminate an agreement without cause. The complete provision might provide that the customer may exercise such a right only after an initial contractual period, after providing ninety days' notice and upon payment of an early termination charge. The AI summary is therefore not necessarily entirely false; it is legally incomplete.

In contractual interpretation, qualifications and exceptions can be as important as the principal obligation. Similar concerns arise where an AI tool overlooks an exception, incorrectly identifies a defined term, fails to connect provisions located in different sections, misinterprets an incorporated document, misses an amendment, incorrectly categorises contractual risk or recommends language inconsistent with applicable law.

This distinction is important because legal professionals should not regard AI-generated output as equivalent to a legal conclusion merely because it is presented confidently.

7. Client-Approved AI: Does Approval Transfer Responsibility?

An increasingly important practical situation arises where the use of a particular AI tool is approved, encouraged or required as part of a contractual or legal-services workflow. At first glance, client approval may appear to resolve the principal governance concern. It does not necessarily do so.

Approval of a technology answers whether the technology may be used. It does not automatically answer what information may be entered, for what purpose it may be used, which outputs require verification, which contractual risks require mandatory human review, who is accountable if the AI-generated recommendation is inaccurate, or what happens when the AI output conflicts with an approved legal playbook.

Accordingly, client or organizational approval of an AI platform should be regarded as the beginning of the governance analysis rather than its conclusion. Relevant stakeholders should define the permitted use case, not merely the permitted technology.

This distinction becomes particularly important where AI is integrated into routine contract review. Repeated use of an approved system can gradually create excessive reliance upon its outputs. An approved AI tool can still produce an incorrect legal analysis. Technology approval therefore should not be interpreted as delegation of professional judgment.

8. Why Organizations Need AI and Data-Privacy Playbooks

High-level AI policies are useful but frequently insufficient for professionals performing day-to-day contractual work. A policy might instruct employees to use AI responsibly or protect confidential information. A contract professional requires more specific guidance.

An AI and Data-Privacy Playbook can bridge the gap between broad principles and operational decisions. Such a playbook may address approved and prohibited AI technologies; permitted use cases; restricted categories of information; data-minimisation and redaction requirements; contractual clauses requiring mandatory human review; escalation thresholds; validation procedures; privacy requirements; cybersecurity requirements; documentation standards; incident-response procedures; and periodic review of AI-generated outcomes.

The objective of such a playbook should not simply be to identify prohibited activities. A mature governance framework should explain how professionals can use AI safely and effectively.

9. A Risk-Based Framework for AI-Assisted Contract Review

Treating every use of AI as equally risky would be impractical. Organizations should instead consider a risk-based model.

9.1 Lower-Risk Activities

Depending upon the information involved and the system being used, relatively lower-risk activities might include identifying defined terms, locating particular clauses, extracting dates, formatting information and summarising non-sensitive provisions.

9.2 Moderate-Risk Activities

Activities requiring greater oversight might include comparing clauses against templates, identifying missing contractual provisions, extracting contractual obligations, preliminary classification of deviations and summarising commercially significant provisions.

9.3 Higher-Risk Activities

Substantial human review should generally accompany AI-generated analysis relating to matters such as limitation of liability, indemnification, intellectual-property ownership, confidentiality, data protection, cybersecurity, regulatory compliance, termination rights, governing law and significant financial obligations.

The underlying governance principle is straightforward: as the potential legal or commercial consequence of an AI-generated output increases, the degree of required human review should

increase correspondingly. This approach enables organizations to obtain the benefits of AI without treating routine extraction and consequential legal judgment as equivalent activities.

10. Human Oversight and Professional Accountability

Discussion concerning AI and legal services is frequently framed around whether artificial intelligence will replace lawyers. That question overlooks the more immediate challenge. The relevant issue is how responsibilities should be allocated between technology and legal professionals.

AI can be particularly effective in processing volume, extracting information, comparing language and identifying patterns. Legal professionals contribute contextual interpretation, commercial judgment, negotiation strategy, legal reasoning and accountability. An effective contract-review framework should therefore combine these capabilities.

10.1 The Risk of Automation Bias

Human oversight itself must also be meaningful. Repeated exposure to generally accurate AI recommendations can create automation bias—a tendency to accept automated recommendations without sufficient independent scrutiny.

An organization may technically require human review while, in practice, reviewers routinely approve AI-generated conclusions without verifying the underlying contractual language. Such a process provides the appearance of human oversight without its substance.

Human verification should therefore be designed into the workflow. For higher-risk contractual provisions, the reviewer should have ready access to the original contractual language, the AI-generated analysis, the applicable contractual playbook and relevant escalation guidance. AI should assist the professional in reaching a decision, rather than become a substitute for the decision itself.

11. Recommendations for Responsible AI-Assisted Contracting

11.1 Establish Approved AI Environments

Professionals should clearly understand which systems may process contractual information and which may not.

11.2 Define Permitted Use Cases

Approving an AI platform is insufficient. Organizations should define what the platform may actually be used to do.

11.3 Minimise Data Exposure

Where information is unnecessary for the intended analysis, organizations should consider whether it can be excluded, masked or redacted.

11.4 Classify Contractual Risk

Routine information extraction and consequential legal recommendations should not be governed identically.

11.5 Preserve Meaningful Human Review

Professionals should verify AI-generated analysis before relying upon it for significant legal or commercial decisions.

11.6 Integrate Privacy and Cybersecurity

AI governance should involve legal, privacy, cybersecurity, procurement and technology stakeholders where appropriate.

11.7 Develop an Operational AI and Data-Privacy Playbook

Rules should be sufficiently specific to guide professionals encountering actual contracting scenarios.

11.8 Maintain Escalation Mechanisms

Professionals should know when an AI-assisted review requires specialist legal, privacy, cybersecurity or commercial input.

11.9 Periodically Test AI Performance

Organizations should consider periodically reviewing whether AI systems are correctly

identifying contractual risks and whether error patterns have emerged.

11.10 Preserve Accountability

AI may generate an analysis or recommendation. Responsibility for consequential contractual decisions should remain identifiable.

12. Conclusion

Artificial Intelligence has significant potential to transform commercial contract review in India. Used appropriately, AI can improve efficiency, consistency and access to institutional contractual knowledge. It can reduce repetitive review activities and allow legal professionals to concentrate on complex negotiations and higher-value legal analysis.

However, technological efficiency cannot be separated from legal responsibility. Contracts may contain confidential information, personal data, intellectual property and commercially significant obligations. Processing these documents through AI systems therefore raises questions involving confidentiality, privacy, cybersecurity, accuracy and accountability.

India's emerging digital data-protection framework makes these considerations increasingly significant. The appropriate response is neither unrestricted adoption nor complete avoidance of artificial intelligence.

Organizations should instead adopt a risk-based governance model supported by approved AI environments, defined use cases, data-minimisation practices, privacy and security safeguards, AI and data-privacy playbooks, escalation mechanisms and meaningful human oversight.

Importantly, even where the use of an AI system is approved or required within a contracting workflow, technology approval should not be confused with transfer of professional responsibility.

The future of contract review is therefore unlikely to involve a simple choice between lawyers and artificial intelligence. A more sustainable model is one in which AI enhances legal judgment without replacing responsibility for that judgment.

Endnotes

1. The Digital Personal Data Protection Act, 2023, Act No. 22 of 2023, enacted on 11 August 2023.

2. Digital Personal Data Protection Act, 2023, Section 3.

3. Digital Personal Data Protection Act, 2023, Sections 4–10.

4. Digital Personal Data Protection Act, 2023, Section 8(1).

5. Digital Personal Data Protection Act, 2023, Sections 2(i) and 8.

6. Digital Personal Data Protection Act, 2023, Section 8, including the obligation concerning reasonable security safeguards against personal data breaches.

7. Digital Personal Data Protection Rules, 2025, Ministry of Electronics and Information Technology, Government of India, notified on 14 November 2025. The Rules contain phased commencement provisions; applicability of individual requirements should be assessed by reference to their notified commencement dates.

8. Justice K.S. Puttaswamy (Retd.) & Anr. v. Union of India & Ors., (2017) 10 SCC 1.

9. Constitution of India, Article 21. In Justice K.S. Puttaswamy (Retd.) v. Union of India, the nine-judge Bench of the Supreme Court recognised privacy as a constitutionally protected right associated with liberty, dignity and autonomy.

10. K.S. Puttaswamy (Aadhaar) v. Union of India, (2019) 1 SCC 1.

11. Digital Personal Data Protection Act, 2023, Section 16, concerning processing of personal data outside India and the Central Government's power to restrict transfers to specified countries or territories.

12. Digital Personal Data Protection Act, 2023, Section 2(i), defining "Data Processor" as a person who processes personal data on behalf of a Data Fiduciary.

13. Digital Personal Data Protection Act, 2023, Section 2(j), defining "Data Principal" as the individual to whom the personal data relates, subject to statutory qualifications.

14. Digital Personal Data Protection Act, 2023, Section 2(k), defining a "Data Fiduciary" as a person who alone or in conjunction with others determines the purpose and means of processing personal data.

15. Digital Personal Data Protection Act, 2023, Section 2(t), defining "personal data" as data about an individual who is identifiable by or in relation to such data.

16. Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 SCC 1, recognising informational privacy within the broader constitutional protection of privacy.