QUANTIVE

POLICY REPORT

Does India Need a Dedicated Artificial Intelligence Law?

Evidence from Key Stakeholders

Does India Need a Dedicated Artificial Intelligence Law?

Evidence from Key Stakeholders

Prepared by

Quantive Advisory LLP
India
Email: info@quantiveadvisory.com, Web site: www.quantiveadvisory.com

Table of Contents

Acknowledgement *4*
Executive Summary *6*
1. Introduction *11*
2. Methodology *13*
3. Key Findings *15*
4. Decision Pathway for Government *26*
5. Recommendations and Implications for Government *30*
6. Conclusion *39*

Acknowledgement

Quantive Advisory LLP gratefully acknowledges the contributions of all individuals and institutions that supported this study. We are particularly thankful to the academics and technical experts, civil society organisations, think tanks, AI startups, sectoral deployers, industry bodies, legal professionals, former policymakers and regulatory practitioners who participated in the stakeholder consultations and shared their experience, concerns and perspectives on the future of AI governance in India.

Their inputs were essential in understanding the practical strengths and limitations of India's current regulatory framework, the challenges faced by startups and deploying organisations, and the safeguards required where AI systems affect rights, livelihoods, public services and institutional decision-making. The study benefited greatly from the openness and diversity of these perspectives.

We also acknowledge the contribution of researchers, policymakers, regulators and public-interest institutions whose published work informed the literature review and comparative analysis. In particular, the India AI Governance Guidelines and related policy initiatives provided an important foundation for assessing India's evolving approach to responsible, safe and innovation-enabling AI governance.

The analytical framing of this report also draws from the book *Governing Artificial Intelligence: From Technological Breakthroughs to Global Regulations and Market Power*, which examines AI as a socio-technical system shaped by law, markets, institutions and public policy.

We are grateful to the Quantive Advisory LLP team and associated researchers for their support in conducting interviews, reviewing documents, organising the evidence and preparing the report. Responsibility for the analysis, conclusions and recommendations contained in this publication rests entirely with Quantive Advisory LLP.

Finally, we thank all stakeholders who continue to contribute to India's rapidly evolving technology-policy discourse. It is hoped that this report will support informed discussion and constructive engagement among government, industry, academia, startups and civil society on the development of a balanced and future-ready AI-governance framework.

Executive Summary

Artificial intelligence (AI) is increasingly embedded in markets, public administration, healthcare, finance, education, employment and digital platforms. Its growing role in generating content, making predictions and supporting or automating decisions has intensified questions concerning accountability, transparency, fairness, safety, liability and human oversight.

India has so far adopted a flexible and sector-oriented approach to AI governance. Existing laws relating to information technology, data protection, consumer protection, competition, intellectual property and criminal conduct already apply to several AI-enabled activities. At the same time, these laws were not designed specifically for systems that learn continuously, generate or modify content, operate with varying degrees of autonomy and influence decisions affecting rights and livelihoods.

Against this background, the report examines whether India requires a dedicated AI law. The study adopts a qualitative, multi-stakeholder design combining a review of relevant literature with 40 semi-structured interviews across academia and technical experts, civil society and think tanks, AI startups, sectoral deployers and industry bodies, and legal, regulatory and policy experts.

The central finding is that India's policy choice should not be reduced to a binary question of whether to enact an AI law or continue without one. The more important issue is the design of a governance system that can address material harm while preserving innovation, investment, competition and domestic technological capability.

The evidence indicates that India does not face a complete legal vacuum. Many harmful outcomes—including fraudulent deepfakes, unlawful processing of personal data, misleading consumer practices, identity theft and copyright infringement—can already be addressed through existing laws. However, the present framework remains fragmented across statutes, institutions and sectoral mandates. It also does not consistently provide preventive safeguards for high-risk AI systems, such as impact assessments, bias and robustness testing, incident reporting, audit trails, human review and post-deployment monitoring.

The study therefore recommends a phased, risk-based and hybrid regulatory architecture. India should first strengthen the enforcement and clarification of existing laws, undertake a structured regulatory gap analysis and build technical and institutional capacity. Targeted amendments and sector-specific safeguards should then be introduced where current laws are inadequate. A limited framework AI law should be considered over the longer term only where evidence shows that statutory consolidation would improve regulatory certainty, public protection and enforcement.

A key recommendation is that India should regulate AI according to risk and impact rather than the mere presence of AI in a product or service. Routine productivity and creative tools should remain subject to general laws and voluntary standards. Chatbots, recommendation systems and synthetic media may require transparency, labelling and user-control obligations. High-risk uses in healthcare, credit, employment, welfare and critical infrastructure should be subject to stronger requirements, including impact assessments, testing, human oversight, auditability and redress.

Manipulative, unlawful or fundamentally rights-incompatible uses may require prohibition or tightly controlled authorisation.

The report also finds that liability should be allocated according to the function performed, degree of control, level of risk, foreseeability of harm and due diligence undertaken by each actor. Foundation-model developers, downstream developers, deployers, platforms, professional users and end users should not bear identical responsibility. At the same time, no actor exercising meaningful control over an AI system should remain outside the accountability framework.

Voluntary governance should remain an important first layer, particularly for low-risk applications and early-stage innovation. However, voluntary commitments cannot replace enforceable safeguards in high-risk contexts. India should therefore adopt a graduated compliance framework in which voluntary standards apply to low-risk systems, documented risk-management obligations apply to moderate-risk applications, and mandatory assessments and oversight apply to high-risk systems.

The study further highlights the need to protect startups from disproportionate compliance costs without exempting them from responsibility. Simplified documentation, phased timelines, common testing tools, subsidised evaluation infrastructure, regulatory sandboxes and technical support through the AI Safety Institute can make compliance more accessible. High-risk uses, however, should remain subject to appropriate safeguards regardless of firm size.

Human oversight and contestability emerge as non-negotiable principles for high-impact AI. Individuals affected by significant AI-assisted decisions should receive notice that AI was materially involved, understandable information about the basis of the decision, access to human review, the ability to challenge the outcome and timely grievance redressal.

Institutionally, the report does not recommend the immediate creation of a single powerful AI regulator. A whole-of-government model is likely to be more effective. MeitY should provide overall policy coordination; sectoral regulators should retain domain-specific enforcement authority; and institutions such as the AI Governance Group, Technology and Policy Expert Committee and AI Safety Institute should provide coordination, technical expertise, standards and testing support.

The proposed decision pathway has three stages:

- **Immediate term:** clarify how existing laws apply to AI, operationalise governance institutions, develop an India-specific risk framework, clarify liability, establish an AI incidents database and strengthen sectoral regulatory capacity.
- **Medium term:** introduce targeted amendments, establish common audit and testing standards, mandate safeguards for selected high-risk applications, pilot regulatory sandboxes and create simplified compliance mechanisms for startups.
- **Longer term:** assess whether regulatory fragmentation persists, conduct a Regulatory Impact Assessment, consult widely and consider a limited framework AI law where statutory consolidation is demonstrably necessary.

The report concludes that India should neither rush into a compliance-heavy AI Act nor rely indefinitely on fragmented rules and voluntary commitments. The most practical path is to strengthen the present framework, regulate according to risk, protect innovation and competition, and progressively introduce binding safeguards where evidence demonstrates clear and material harm.

Such an approach would provide government with flexibility, citizens with meaningful protection and industry with greater certainty and predictable compliance expectations. It would also enable India to develop a distinctive AI-governance model aligned with its constitutional values, developmental priorities, startup ecosystem and wider ambition to shape responsible AI governance for the Global South.

1. Introduction

Artificial intelligence (AI) has moved rapidly from a specialised technological tool to a foundational force shaping markets, public institutions, information systems and everyday decision-making. As AI becomes embedded in sectors such as finance, healthcare, education, employment, public administration and digital platforms, questions of transparency, accountability, fairness, safety and human oversight have become increasingly difficult to address through voluntary principles alone[1].

India has so far followed a largely flexible and sector-oriented approach to AI governance, relying on existing laws, regulatory guidance and emerging soft-law frameworks[2]. However, the growing use of generative AI, foundation models and autonomous systems has intensified the debate over whether this approach remains sufficient or whether India now requires a dedicated statutory framework[3].

This report examines that policy choice by bringing together stakeholder perspectives and a review of existing literature. Rather than treating the issue as a binary choice between introducing a comprehensive AI law and continuing with the existing framework, it assesses the relative merits of a phased and risk-based approach. In particular, the report evaluates where current laws remain adequate, where targeted amendments or sector-specific safeguards are required, and whether a limited framework law may eventually be necessary to improve coherence, accountability and enforcement. It also considers how such a framework can protect rights and

[1] https://unesdoc.unesco.org/ark:/48223/pf0000397817
[2] Governing Artificial Intelligence: From Technological Breakthroughs to Global Regulations and Market Power
[3] https://www.communicationstoday.co.in/india-moves-toward-a-dedicated-ai-law/

public interest without imposing disproportionate burdens on startups, innovation and responsible AI adoption.

2. Methodology

The study adopts a qualitative, multi-stakeholder research design comprising two principal components: a review of relevant literature and semi-structured interviews with key stakeholders.

Primary consultations were conducted with representatives from academia and technical experts, civil society and think tanks, AI startups, sectoral deployers and industry bodies, legal community, former policymakers, and regulatory practitioners. Table 1 presents the composition of the stakeholder sample, comprising 40 interviews across five respondent categories.

Table 1: Distribution of Stakeholder Interviews by Respondent Category

Stakeholder group	Number of interviews
Academia and technical experts	10
Civil society and think tanks	10
AI startups	10
Sectoral deployers and industry bodies	5
Legal, regulatory and policy experts	5
Total	**40**

The interviews examined stakeholder perspectives on existing regulatory gaps, the need for a dedicated AI law, preferred governance models, the

treatment of high-risk applications, accountability mechanisms, institutional capacity, startup-related concerns, and the appropriate balance between innovation and public-interest safeguards.

The interview data were subsequently coded and analysed thematically to identify recurring themes, areas of consensus and divergence, and stakeholder-specific priorities.

3. Key Findings

The stakeholder consultations and review of the existing literature indicate that India's policy choice should not be reduced to a binary question of whether to regulate or not regulate AI. The more relevant question is what form of governance can protect individuals and institutions from material harm while preserving India's ability to innovate, attract investment and build domestic AI capabilities.

India does not begin from a legal vacuum. Several existing laws already apply to AI-enabled activities, including those relating to information technology, data protection, consumer protection, competition, intellectual property and criminal conduct. At the same time, these laws were generally not designed for systems that generate content, make predictions, learn continuously, operate autonomously or influence decisions affecting rights and livelihoods. The resulting framework is functional in parts but fragmented as a whole.

The findings therefore point towards a phased and hybrid regulatory architecture. India should first strengthen the application of existing laws, identify specific legal gaps and build institutional capacity. Binding AI-specific obligations should be introduced where the level of risk justifies them, while low-risk and beneficial applications should remain subject to lighter governance. Over time, these measures could be consolidated into a concise framework law if fragmentation, emerging capabilities or enforcement gaps make such legislation necessary.

3.1 India has an AI-governance framework, but not yet a coherent AI-governance system

India has gradually developed an extensive policy architecture around responsible AI. This includes the Principles for Responsible AI, sector-specific guidance, the India AI Governance Guidelines, data-protection legislation and targeted rules concerning digital platforms and synthetic content. Sectoral bodies have also begun addressing AI in finance, competition, health and other fields.

These interventions demonstrate that the country is not waiting passively for a single AI law. However, they operate across different statutes, regulators and institutional mandates. There is no single, commonly applicable risk-classification system, no consolidated statement of obligations across the AI value chain and no unified process for dealing with AI-related grievances or incidents.

This creates uncertainty for both regulators and industry. A startup may have difficulty determining which rules apply to its model, while a citizen affected by an automated decision may not know which institution is responsible for redress. Government agencies may also interpret similar risks differently across sectors.

The immediate priority should therefore be regulatory consolidation and coordination, rather than legislation for its own sake.

3.2 Existing laws can address many AI harms, but they do not adequately cover AI-specific governance processes

The evidence suggests that many harmful outcomes associated with AI can already be dealt with under existing law. Fraudulent deepfakes, identity theft,

unlawful processing of personal data, deceptive consumer practices and copyright infringement do not become legally permissible merely because AI is involved.

However, existing laws are generally better at responding to harm after it has occurred than at preventing high-risk AI failures before deployment. They do not consistently answer questions such as:

- Whether an AI system should undergo an impact assessment;
- When bias and robustness testing should be mandatory;
- What documentation must be retained;
- Whether users should be informed when AI influences a decision;
- When human review must be available;
- Whether serious AI incidents must be reported;
- How model performance should be monitored after deployment.

This distinction is fundamental. Traditional laws regulate conduct and outcomes; AI governance must also regulate selected processes where the probability or consequences of failure are substantial.

The gaps are most serious where automated systems affect employment, credit, health, education, welfare, law enforcement, public infrastructure or other areas involving life, liberty or livelihood.

3.3 India should regulate risk and impact, not the mere use of AI

A central finding is that all AI systems should not be subjected to the same regulatory burden. The phrase “artificial intelligence” includes a very wide range of applications, from grammar tools and recommendation engines to medical diagnostics, automated credit decisions and autonomous systems.

A uniform law could therefore be both over-inclusive and under-inclusive. It could impose unnecessary obligations on low-risk applications while failing to distinguish the particular safeguards required in sectors such as healthcare, finance or criminal justice.

The India AI Governance Guidelines[4] support a balanced, flexible and evidence-led approach, including the development of an India-specific risk-assessment framework based on real-world harm. They identify major risks such as malicious use, bias, transparency failures, systemic risks, loss of control and national-security threats.

A practical Indian model should classify risk according to factors such as:

- The purpose of the system;
- The seriousness and probability of harm;
- The vulnerability of affected users;
- The degree of system autonomy;
- The scale of deployment;
- Whether the decision can be reversed or contested;
- Whether the system affects rights, essential services or critical infrastructure.

3.4 A dedicated law may eventually be required, but India should avoid copying the EU model wholesale

International practice demonstrates that there is no single accepted model of AI regulation[5]. The European Union (EU) has adopted a comprehensive

[4] https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc2025115685601.pdf
[5] Governing Artificial Intelligence: From Technological Breakthroughs to Global Regulations and Market Power

statutory framework, while the United Kingdom, Singapore and other jurisdictions rely more heavily on principles, sectoral regulators, voluntary tools and technical standards. Japan has chosen a framework law focused largely on institutional coordination and innovation rather than a compliance-heavy regime.

India's institutional capacity, development priorities, market structure and startup ecosystem differ from those of the European Union. A direct replication of the EU AI Act could create a substantial compliance burden without guaranteeing effective enforcement.

At the same time, indefinite reliance on fragmented rules and voluntary commitments would leave important questions unresolved. A dedicated law may therefore become useful if it performs a coordinating and enabling function, rather than attempting to prescribe every technical requirement within one statute.

A future Indian framework law could establish:

- Common definitions;
- National principles;
- A risk-based governance structure;
- Prohibited or restricted applications;
- Baseline rights of affected individuals;
- Allocation of responsibilities across the AI value chain;
- Powers for sectoral regulators;
- Institutional coordination;
- Sandboxes and innovation safeguards;

- Periodic review mechanisms.

Detailed rules should remain with sector-specific regulators that possess relevant domain expertise.

3.5 Liability must follow function, control, risk and due diligence

One of the clearest regulatory gaps concerns the allocation of responsibility when multiple actors contribute to the development and deployment of an AI system.

The AI value chain may involve:

- Providers of data and computing infrastructure;
- Foundation-model developers;
- Downstream application developers;
- Deployers;
- Digital platforms;
- Professional users;
- Individual end users.

A uniform liability rule would be impractical and potentially unfair. A model developer should not automatically bear full responsibility for every malicious downstream use, while a deploying organisation should not be able to avoid responsibility by attributing a harmful decision entirely to an external model.

The India AI Governance Guidelines recommend graded liability based on the function performed, the level of risk and whether reasonable due diligence was exercised.

A workable liability framework should consider:

- Which actor controlled the relevant part of the system;
- Whether the harm was reasonably foreseeable;
- Whether the model was modified or safeguards were removed;
- Whether the deployer selected an inappropriate high-risk use;
- Whether required testing and monitoring were conducted;
- Whether the actor responded appropriately after becoming aware of the harm.

3.6 Voluntary governance should remain the first layer, but not the final layer

Voluntary measures have an important place in a rapidly changing field. Codes of practice, red-teaming, impact assessments, technical standards, transparency reports, third-party audits and internal governance policies can be adopted and updated more quickly than legislation.

For startups and low-risk applications, such measures can encourage responsible innovation without imposing excessive legal costs. They can also generate evidence that helps regulators determine which practices should later become mandatory.

However, voluntary governance has clear limitations. It lacks legal enforceability, may be inconsistently applied and can allow firms to assess their own systems without independent scrutiny. It is therefore insufficient where AI decisions can significantly affect rights, safety or access to essential services.

The balanced approach is to use voluntary measures as part of a graduated compliance system:

- Voluntary standards for low-risk systems;
- Documented risk-management obligations for moderate-risk applications;
- Mandatory assessments and safeguards for high-risk uses;
- Legal prohibitions for unacceptable uses.

3.7 Regulation should protect startups from disproportionate compliance burdens without creating exemptions from responsibility

India's AI ambitions depend substantially on startups, research institutions and smaller technology companies. A compliance regime designed around the resources of major global firms could unintentionally reinforce market concentration.

Large firms can absorb the costs of legal advice, audit, documentation and technical testing. Smaller firms cannot always do so. Uniform obligations could therefore raise barriers to entry, discourage experimentation and increase dependence on the same large cloud and model providers that already occupy strategic positions across the AI ecosystem.

This concern is also consistent with the book's wider argument that AI governance must account for concentration in data, compute, infrastructure and distribution. Regulation that ignores these structural differences may strengthen incumbents while appearing formally neutral.

Nevertheless, startup status should not excuse harmful deployment. A small company using AI to make critical healthcare or credit decisions should still be subject to appropriate safeguards.

3.8 Human oversight and the right to contest AI decisions must be non-negotiable in high-impact contexts

The study finds broad normative support for retaining meaningful human authority over high-impact AI systems. Human oversight is not simply an ethical preference; it is essential for assigning responsibility, correcting errors and preserving public trust.

The book *Governing Artificial Intelligence: From Technological Breakthroughs to Global Regulations and Market Power*[6] also argues that human-in-the-loop governance should be treated as a structural requirement where AI affects people, resources or rights. It emphasises that AI may augment judgment but cannot replace responsibility, particularly in welfare, justice, healthcare and democratic systems.

A meaningful human-oversight requirement does not imply that every automated output must be manually approved. Depending on risk, it may involve:

- Human authorisation before a critical decision;
- Monitoring and escalation thresholds;
- The ability to override or suspend the system;
- Automated circuit breakers;

[6] https://www.bloomsbury.com/in/governing-artificial-intelligence-ai-9789369522262/

- Periodic human review;
- Accessible grievance and appeal channels;
- Audit trails that allow a decision to be reconstructed.

3.9 Institutional coordination is more urgent than the creation of a powerful new regulator

AI affects sectors with very different technical and legal characteristics. The Reserve Bank of India is better placed to understand AI-related financial risk; health regulators understand clinical safety; competition authorities understand market concentration; and MeitY remains central to digital policy and technology regulation.

Creating a single super-regulator could duplicate existing mandates and create institutional conflict. The stronger approach is a coordinated, whole-of-government architecture.

The India AI Governance Guidelines propose an AI Governance Group, supported by a Technology and Policy Expert Committee and an AI Safety Institute, while retaining enforcement powers with sectoral regulators.

3.10 AI governance should be supported by infrastructure, evidence and regulatory capacity, not law alone

Effective regulation depends on the government's ability to understand and test the systems it seeks to govern. Without adequate technical expertise, legal obligations may remain difficult to enforce or may be implemented through broad and uncertain restrictions.

India therefore needs:

- Representative evaluation datasets;
- Safety and bias-testing capacity;
- Standards for auditing and documentation;
- AI incident-reporting mechanisms;
- Content-authentication tools;
- Privacy-enhancing technologies;
- Trained regulators, judges and law-enforcement personnel;
- Horizon-scanning and scenario-planning capacity.

The India AI Governance Guidelines propose a phased action plan including institutional development, India-specific risk frameworks, an AI incidents database, liability rules, standards, sandboxes and capacity building.

4. Decision Pathway for Government

The findings point towards a sequenced approach to AI governance rather than an immediate shift to comprehensive legislation. India already has several laws, regulatory institutions and policy frameworks that apply to AI-related activities. However, gaps remain in areas such as risk classification, liability, automated decision-making, incident reporting, human oversight and institutional coordination.

Accordingly, the Government of India may consider a phased decision pathway that strengthens the existing framework in the immediate term, introduces targeted regulatory measures in the medium term and evaluates the need for a limited framework law over the longer term. This approach would allow governance arrangements to evolve alongside technological development, regulatory capacity and evidence of real-world harm.

Figure 1: Phased Decision Pathway for AI Governance in India

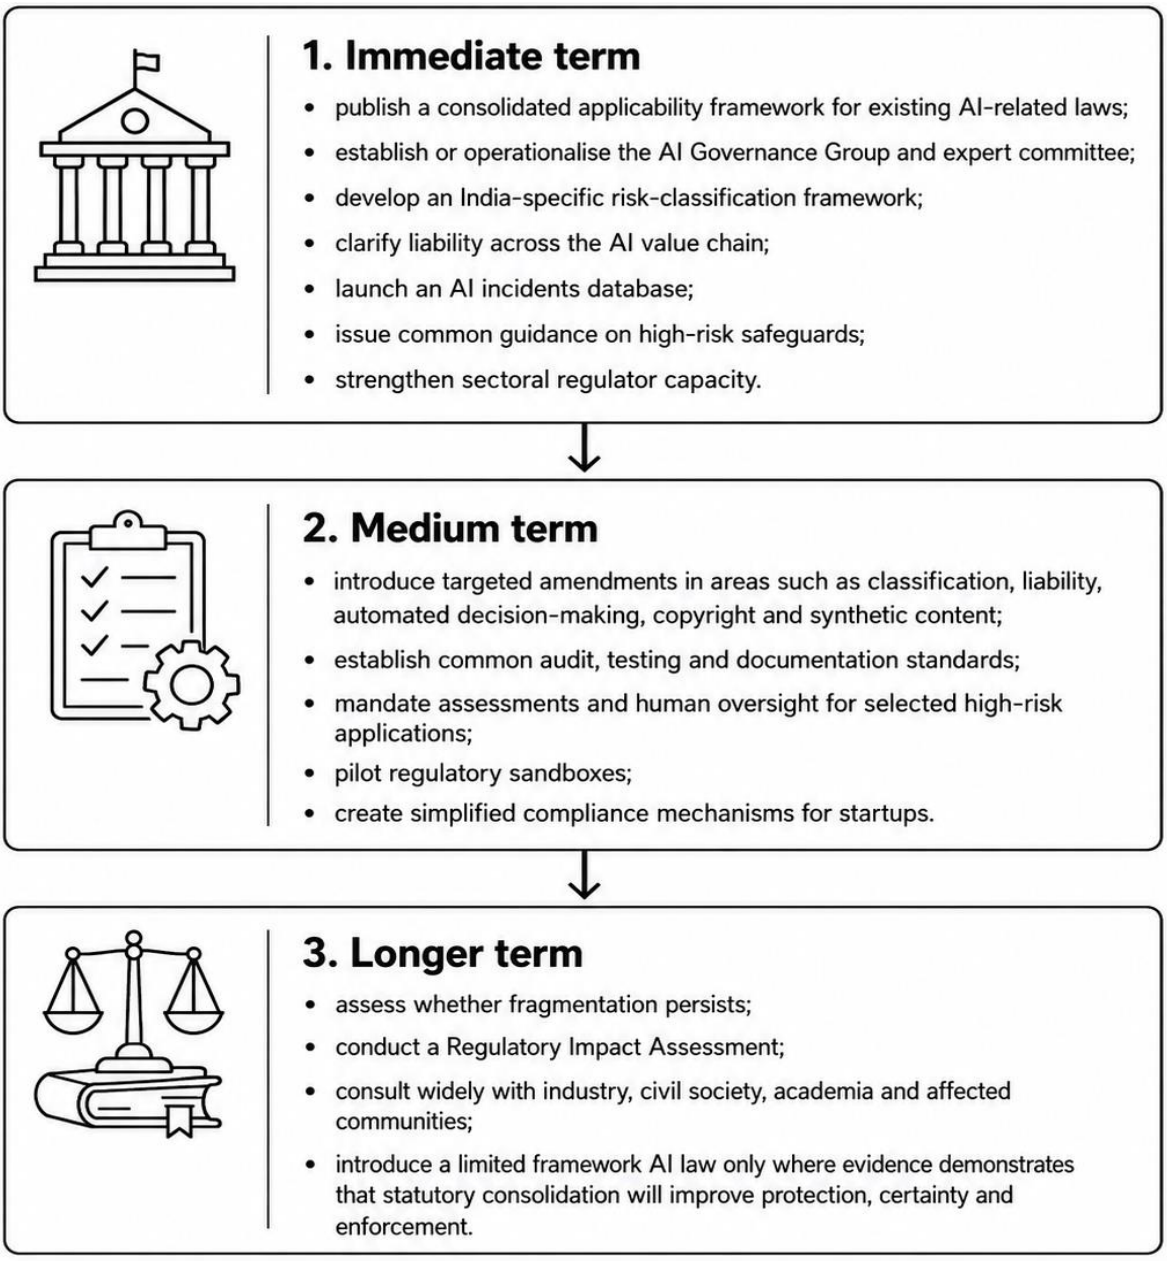

Source: Quantive's Analysis

Figure 1 presents a three-stage roadmap for the Government of India.

In the immediate term, the priority should be to improve clarity and institutional readiness. This includes publishing a consolidated framework

explaining how existing laws apply to AI systems, operationalising the AI Governance Group and expert committee, developing an India-specific risk-classification framework, clarifying liability across the AI value chain and launching a national AI incidents database. Common guidance on safeguards for high-risk systems and stronger capacity within sectoral regulators would help establish a coherent foundation for future intervention.

In the medium term, the government should act on the regulatory gaps identified through the initial review. This may involve targeted amendments relating to the classification of AI actors, liability, automated decision-making, copyright and synthetic content. Common standards for auditing, testing and documentation should also be developed. High-risk applications should gradually become subject to mandatory impact assessments, human oversight and redress mechanisms. Regulatory sandboxes and simplified compliance arrangements would allow startups and smaller firms to meet these obligations without facing disproportionate costs.

In the longer term, the government should assess whether fragmentation persists despite stronger enforcement, targeted amendments and institutional coordination. A Regulatory Impact Assessment should be undertaken to evaluate the likely economic, social and administrative consequences of a dedicated AI law. This process should include extensive consultation with industry, civil society, academia, startups and affected communities. A limited framework law should be considered only where the evidence demonstrates that statutory consolidation would materially improve regulatory certainty, accountability, public protection and enforcement.

The proposed pathway neither delays necessary governance nor assumes that comprehensive legislation is the only solution. It gives policymakers a practical sequence for intervention while providing industry with a predictable and proportionate route towards responsible compliance. It also allows India to build its regulatory architecture incrementally, drawing on evidence, institutional learning and changing technological capabilities rather than relying on a rigid one-time legislative response.

5. Recommendations and Implications for Government

The findings suggest that India should adopt a phased, risk-based and institutionally coordinated approach to AI governance. The immediate objective should not be to introduce an expansive and compliance-heavy AI law, but to improve regulatory coherence, clarify responsibilities, strengthen institutional capacity and introduce binding safeguards where the level of risk justifies them.

5.1 Publish a Consolidated AI Regulatory Applicability Framework

MeitY, in consultation with sectoral regulators and relevant ministries, should issue a consolidated AI regulatory applicability framework or master circular. This should clearly identify:

- The existing laws applicable to AI development and deployment;
- The regulators responsible for different sectors and types of harm;
- The obligations of model developers, deployers, platforms and professional users;
- The grievance, appellate and enforcement mechanisms available to affected individuals;
- The areas in which existing legislation does not adequately address AI-specific risks.

Such guidance would reduce regulatory uncertainty, prevent conflicting interpretations and help startups and companies understand their obligations without navigating multiple disconnected legal instruments. It would also enable the government to distinguish between gaps that require better enforcement and those that genuinely require legislative intervention.

5.2 Conduct a Structured Regulatory Gap Analysis

The government should undertake a systematic review of AI-related risks across priority sectors, including healthcare, finance, employment, education, welfare, law enforcement, telecommunications and critical infrastructure.

The review should assess:

- Whether existing laws adequately cover the identified harm;
- Whether regulators possess sufficient enforcement powers;
- Whether new procedural safeguards are required;
- Whether the risk can be addressed through standards, guidance or a targeted amendment;
- Whether statutory intervention is necessary.

Where existing law is sufficient, the priority should be timely enforcement and clear guidance. Where gaps exist, targeted amendments or sector-specific rules should be preferred over a broad law that may impose unnecessary obligations across the entire AI ecosystem.

5.3 Adopt an India-Specific Risk Classification Framework

India should regulate AI according to the seriousness and context of risk rather than the mere presence of AI in a product or service. A four-tier model could provide a practical starting point:

Risk category	Illustrative applications	Recommended treatment
Minimal risk	Routine productivity tools and creative assistance	Existing laws and voluntary standards
Transparency risk	Chatbots, recommendation systems and synthetic media	Disclosure, labelling, user controls and grievance mechanisms
High risk	Healthcare, credit, employment, welfare and critical infrastructure	Impact assessments, safety testing, audit trails, human oversight and redress
Unacceptable risk	Manipulative, unlawful or fundamentally rights-incompatible uses	Prohibition or tightly controlled authorisation

The classification should consider the purpose of the system, degree of autonomy, scale of deployment, vulnerability of affected users, reversibility of harm and impact on life, liberty, livelihood or essential services.

This approach would protect the public while avoiding the unnecessary classification of every AI-enabled product as a high-risk system.

5.4 Follow a Phased Legislative Pathway

The policy question should not be framed as a choice between immediately adopting an AI Act and taking no regulatory action. The preferred sequence should be:

- Enforce existing laws consistently;
- Clarify their application to AI systems;
- Identify and address regulatory gaps through targeted amendments;

- Strengthen institutions, standards and testing capacity;
- Consider a framework statute when evidence demonstrates that consolidation is necessary.

A future framework law should be concise, technology-neutral and enabling. It could establish common definitions, core principles, prohibited uses, baseline rights, graded responsibilities and institutional coordination, while leaving detailed sectoral obligations to specialised regulators.

5.5 Establish a Graded Responsibility and Liability Matrix

The government should develop a statutory or regulatory responsibility matrix covering the principal actors in the AI value chain, including:

- Foundation-model developers;
- Downstream developers;
- Deployers;
- Digital platforms;
- Professional users;
- End users.

Liability should reflect:

- The function performed by each actor;
- The degree of control over the relevant system or decision;
- The foreseeability of harm;
- The level of risk;
- Whether safeguards were removed or bypassed;

- Whether reasonable due diligence, testing and monitoring were undertaken;
- How the actor responded after becoming aware of the harm.

This would avoid both extremes: automatically imposing liability on the original developer for every downstream misuse, and allowing deployers or platforms to avoid accountability by attributing decisions entirely to third-party models.

5.6 Create a Graduated Compliance Framework

Voluntary governance should remain an important first layer, especially for low-risk systems and early-stage innovation. However, voluntary commitments should not replace enforceable safeguards in high-risk contexts.

The government could adopt the following graduated approach:

- Voluntary codes and standards for low-risk systems;
- Documented internal risk-management processes for moderate-risk applications;
- Mandatory assessments, auditability and oversight for high-risk systems;
- Prohibition or prior authorisation for unacceptable uses.

Compliant firms could be encouraged through:

- Certification and public recognition;
- Preference in government procurement;
- Access to regulatory sandboxes;

- Technical assistance;
- Reduced reporting burdens;
- Access to common testing infrastructure.

At the same time, baseline obligations such as grievance redressal, serious-incident reporting, human oversight and high-risk impact assessments should become mandatory as standards and institutional capacity mature.

5.7 Protect Startups Through Proportionate Compliance

India should distinguish between risk proportionality and company-size exemption. Smaller firms should not be exempted from safeguards merely because of their size, particularly where they operate in high-risk sectors. However, the compliance method should remain affordable and practicable.

Support measures should include:

- Simplified documentation for low-risk applications;
- Phased compliance timelines;
- Standardised templates and regulatory toolkits;
- Shared testing and auditing tools;
- Subsidised access to evaluation datasets and computing resources;
- Regulatory sandboxes;
- Technical assistance through the AI Safety Institute;
- Safe harbours for supervised research and good-faith experimentation.

This would prevent regulation from becoming a barrier to entry while maintaining meaningful protection in high-impact use cases.

5.8 Guarantee Rights for Individuals Affected by High-Impact AI

Any future legal or regulatory framework should provide clear protections for individuals affected by significant AI-assisted or automated decisions.

These should include:

- Notice that AI was materially involved;
- Accessible and understandable information about the basis of the decision;
- The right to request human review;
- The ability to challenge or appeal the outcome;
- Timely and accessible grievance redressal;
- Protection against retaliation for reporting harm.

These rights would strengthen public trust and provide industry with a standardised compliance expectation across high-risk sectors.

5.9 Operationalise a Whole-of-Government Institutional Architecture

India's AI-governance system should combine central coordination with sector-specific expertise. The institutional architecture could include:

- **MeitY:** nodal ministry and overall policy coordinator;
- **AI Governance Group:** cross-government coordination and strategic decision-making;
- **Technology and Policy Expert Committee:** multidisciplinary technical, legal and policy advice;
- **AI Safety Institute:** testing, standards, safety research and regulatory support;

- **Sectoral regulators:** domain-specific rules, supervision and enforcement;
- **Standards bodies:** benchmarks, certification methods and audit protocols;
- **Civil society, academia and startups:** structured consultation, evidence generation and independent scrutiny.

A single super-regulator may create duplication and conflict with existing bodies. A coordinated networked model would preserve sectoral expertise while ensuring national consistency.

5.10 Build Regulatory and Technical Capacity Before Imposing Complex Obligations

Regulatory ambition must be matched by implementation capacity. Before introducing complex compliance requirements, the government should invest in:

- Representative Indian evaluation datasets;
- Safety, fairness and bias-testing infrastructure;
- Common audit and documentation standards;
- AI incident-reporting systems;
- Content-authentication and provenance tools;
- Privacy-enhancing technologies;
- Training for regulators, judges, prosecutors and law-enforcement personnel;
- Horizon-scanning and scenario-planning capabilities.

The government should also create a national AI incidents database with standard reporting protocols and confidentiality safeguards. Evidence from actual incidents can inform future risk classifications, amendments and enforcement priorities.

The recommended approach is neither regulatory delay nor premature overregulation. It is a sequenced governance strategy that begins with clarity, coordination and targeted intervention and develops into stronger statutory safeguards where real-world evidence demonstrates the need.

This approach would allow the government to:

- Protect citizens and vulnerable groups;
- Reduce uncertainty for responsible firms;
- Preserve innovation and startup participation;
- Avoid strengthening market concentration through excessive compliance costs;
- Improve the enforceability of AI-related rules;
- Create a governance system capable of adapting to technological change.

For industry, it would provide predictable obligations, differentiated treatment based on risk and access to practical compliance support. For policymakers, it would create a credible pathway from existing fragmented regulation towards a coherent and future-ready AI-governance framework.

6. Conclusion

The stakeholder consultations and review of the existing literature indicate that India's AI-governance challenge cannot be resolved through a simple choice between immediate comprehensive legislation and continued regulatory restraint. India already possesses a substantial legal and policy foundation that can address many AI-related harms. However, this framework remains fragmented across statutes, regulators and sector-specific instruments, while important gaps persist in areas such as algorithmic accountability, automated decision-making, liability, incident reporting, human oversight and redress.

The study therefore finds that India should pursue a phased, risk-based and hybrid regulatory approach. Existing laws should be applied more consistently and clarified through consolidated guidance. Specific gaps should be addressed through targeted legal amendments and sectoral rules. Binding obligations should be reserved primarily for high-risk and rights-affecting applications, while low-risk and socially beneficial uses should remain subject to lighter, flexible governance.

A future dedicated AI law may eventually be useful, but its purpose should be to create coherence rather than excessive control. Such a law should establish common definitions, baseline principles, a national risk framework, clear responsibilities across the AI value chain, minimum rights for affected individuals, institutional coordination and powers for sectoral regulators. It should avoid prescribing rigid technical requirements that may quickly become outdated or imposing uniform compliance costs that disadvantage startups and smaller Indian firms.

The findings also suggest that accountability should follow the actual function performed, the degree of control exercised, the level of risk involved and the due diligence undertaken by each actor. Developers, deployers, platforms and users should not carry identical responsibility, but no actor exercising meaningful control over an AI system should remain outside the accountability framework.

Human oversight, transparency and contestability should form the non-negotiable core of governance in high-impact contexts. Where AI affects health, employment, credit, welfare, education, law enforcement or essential public services, individuals should know that AI was involved, have access to understandable information, and be able to seek human review and effective redress.

At the institutional level, India's immediate requirement is not necessarily a single powerful AI regulator. A coordinated whole-of-government structure is likely to be more practical. MeitY can provide overall policy direction, sectoral regulators can retain domain-specific enforcement powers, and institutions such as the AI Governance Group, Technology and Policy Expert Committee and AI Safety Institute can provide coordination, technical expertise, testing support and standards development.

Finally, law alone will not ensure effective governance. India must simultaneously invest in regulatory capacity, evaluation datasets, safety testing, algorithmic auditing, incident reporting, content authentication, privacy-enhancing technologies and training for regulators, judges and law-enforcement agencies. Without such capabilities, even well-designed legislation may remain difficult to implement.

The central conclusion of the study is therefore that India should neither rush into a compliance-heavy AI Act nor rely indefinitely on fragmented rules and voluntary commitments. The most practical path is to strengthen the present framework, regulate according to risk, protect innovation and competition, and progressively introduce binding safeguards where evidence demonstrates clear and material harm.

Such an approach would give government the flexibility to respond to emerging risks while providing industry with greater certainty, predictable compliance expectations and space for responsible innovation. It would also enable India to develop a distinctive governance model—one that is aligned with its constitutional values, developmental priorities, startup ecosystem and ambition to shape AI governance for the Global South.

QUANTIVE

India

Email: info@quantiveadvisory.com

Website: www.quantiveadvisory.com