

INTERNATIONAL JOURNAL FOR LEGAL RESEARCH AND ANALYSIS



Open Access, Refereed Journal Multi-Disciplinary
Peer Reviewed

www.ijlra.com

DISCLAIMER

No part of this publication may be reproduced, stored, transmitted, or distributed in any form or by any means, whether electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the Managing Editor of the *International Journal for Legal Research & Analysis (IJLRA)*.

The views, opinions, interpretations, and conclusions expressed in the articles published in this journal are solely those of the respective authors. They do not necessarily reflect the views of the Editorial Board, Editors, Reviewers, Advisors, or the Publisher of IJLRA.

Although every reasonable effort has been made to ensure the accuracy, authenticity, and proper citation of the content published in this journal, neither the Editorial Board nor IJLRA shall be held liable or responsible, in any manner whatsoever, for any loss, damage, or consequence arising from the use, reliance upon, or interpretation of the information contained in this publication.

The content published herein is intended solely for academic and informational purposes and shall not be construed as legal advice or professional opinion.

**Copyright © International Journal for Legal Research & Analysis.
All rights reserved.**

ABOUT US

The *International Journal for Legal Research & Analysis (IJLRA)* (ISSN: 2582-6433) is a peer-reviewed, academic, online journal published on a monthly basis. The journal aims to provide a comprehensive and interactive platform for the publication of original and high-quality legal research.

IJLRA publishes Short Articles, Long Articles, Research Papers, Case Comments, Book Reviews, Essays, and interdisciplinary studies in the field of law and allied disciplines. The journal seeks to promote critical analysis and informed discourse on contemporary legal, social, and policy issues.

The primary objective of IJLRA is to enhance academic engagement and scholarly dialogue among law students, researchers, academicians, legal professionals, and members of the Bar and Bench. The journal endeavours to establish itself as a credible and widely cited academic publication through the publication of original, well-researched, and analytically sound contributions.

IJLRA welcomes submissions from all branches of law, provided the work is original, unpublished, and submitted in accordance with the prescribed submission guidelines. All manuscripts are subject to a rigorous peer-review process to ensure academic quality, originality, and relevance.

Through its publications, the *International Journal for Legal Research & Analysis* aspires to contribute meaningfully to legal scholarship and the development of law as an instrument of justice and social progress.

PUBLICATION ETHICS, COPYRIGHT & AUTHOR RESPONSIBILITY STATEMENT

The *International Journal for Legal Research and Analysis (IJLRA)* is committed to upholding the highest standards of publication ethics and academic integrity. All manuscripts submitted to the journal must be original, unpublished, and free from plagiarism, data fabrication, falsification, or any form of unethical research or publication practice. Authors are solely responsible for the accuracy, originality, legality, and ethical compliance of their work and must ensure that all sources are properly cited and that necessary permissions for any third-party copyrighted material have been duly obtained prior to submission. Copyright in all published articles vests with IJLRA, unless otherwise expressly stated, and authors grant the journal the irrevocable right to publish, reproduce, distribute, and archive their work in print and electronic formats. The views and opinions expressed in the articles are those of the authors alone and do not reflect the views of the Editors, Editorial Board, Reviewers, or Publisher. IJLRA shall not be liable for any loss, damage, claim, or legal consequence arising from the use, reliance upon, or interpretation of the content published. By submitting a manuscript, the author(s) agree to fully indemnify and hold harmless the journal, its Editor-in-Chief, Editors, Editorial Board, Reviewers, Advisors, Publisher, and Management against any claims, liabilities, or legal proceedings arising out of plagiarism, copyright infringement, defamation, breach of confidentiality, or violation of third-party rights. The journal reserves the absolute right to reject, withdraw, retract, or remove any manuscript or published article in case of ethical or legal violations, without incurring any liability.

REGULATING ARTIFICIAL INTELLIGENCE IN INDIA: BALANCING INNOVATION WITH ACCOUNTABILITY AND FUNDAMENTAL RIGHTS

AUTHORED BY - SIMRAN

Abstract

Artificial Intelligence (AI) is rapidly transforming governance, finance, healthcare, content creation, law enforcement, and digital infrastructure. While India actively promotes AI innovation through initiatives like Digital India and IndiaAI, the absence of a dedicated legal framework poses significant risks to constitutional rights particularly privacy (Article 21), equality (Article 14), and free speech (Article 19). AI-driven surveillance, facial recognition, predictive policing, deepfake misinformation, and algorithmic bias can create an opaque, unaccountable governance ecosystem, where human dignity and due process are threatened. This paper examines the urgent need for rights-based AI regulation in India, critically analysing global regulatory models such as the EU AI Act and constitutional standards emerging from K.S. Puttaswamy (Privacy). It proposes a balanced regulatory architecture that protects innovation but ensures accountability, transparency, non-discrimination, human oversight, and data governance safeguards, advancing the principle that technological progress must not compromise constitutional morality.

Keywords

AI Regulation, Fundamental Rights, Algorithmic Bias, Privacy, Accountability, Digital India Act, EU AI Act, Human Oversight

Literature Review

Early AI discourse in India was dominated by economic and innovation policy (NITI Aayog's National Strategy for AI, 2018), presenting AI as a growth catalyst with limited legal discussion. In contrast, global scholarship especially in the EU and US has deeply engaged with AI ethics, algorithmic discrimination, and AI accountability models. The EU AI Act is the first comprehensive attempt at risk-based AI regulation, classifying AI systems into "unacceptable," "high-risk," and "low-risk" categories, imposing strict compliance and transparency in sensitive domains such as law enforcement and finance.

Recent Indian academic work including by Vidhi Centre for Legal Policy and CyberBRICS highlights AI's constitutional risks, especially surveillance capitalism, digital authoritarianism, and the weaponisation of deepfakes and misinformation. Law-and-technology scholars such as Dr. Usha Ramanathan and Gautam Bhatia argue that Indian AI governance must evolve not only through innovation policy but constitutional constraints, especially consent, fairness, and due process. Global literature also stresses the need for algorithmic audits, explainability, and human oversight, principles currently missing in India's Digital India strategy.

Research Methodology

This research follows a doctrinal, constitutional, and comparative methodology:

- Primary sources: Supreme Court judgments (Puttaswamy, Anuradha Bhasin¹), constitutional principles (Articles 14, 19, 21), Digital India Act (2024/2025 draft policy).
- Comparative study: EU AI Act (2024), OECD AI Principles, and US algorithmic governance debates.
- Secondary academic literature & policy papers on AI ethics, data governance, and digital rights.

The methodology is qualitative, rights-centred, and reform-oriented, evaluating AI regulation against constitutional morality, human dignity, and rule of law standards.

Introduction

Artificial Intelligence (AI) is no longer a futuristic concept it is now embedded in everyday decision-making processes, from facial recognition and content moderation to loan sanctions, predictive policing, hiring systems, and election campaigning. India's ambition to become a global AI hub through initiatives like IndiaAI Mission, Digital India Act (upcoming), and National Strategy for AI (NITI Aayog) reflects its commitment to innovation. However, unlike the European Union's legally enforceable AI Act, India's approach is still policy-heavy but law-light with no dedicated regulatory framework to safeguard fundamental rights.

While AI can bring efficiency, it also threatens constitutional guarantees under Articles 14 (equality), 19 (free speech), and 21 (privacy & dignity). Algorithmic opacity, data monopolies, bias, deepfake misinformation, and automated surveillance can undermine due process,

¹ *Anuradha Bhasin v. Union of India*, (2020) 3 SCC 637.

fairness, and democratic accountability. The Supreme Court's judgment in Justice K.S. Puttaswamy (2017)² elevated privacy and data protection as fundamental rights, implicitly demanding that AI regulation must pass the test of proportionality, necessity, and transparency. India now stands at a critical point: Can it regulate AI in a way that enables innovation while preserving constitutional morality? This research argues YES but only if AI regulation is made rights-centric, auditable, and human-overseen rather than techno-deterministic or purely market-driven.

Artificial Intelligence (AI) has become the defining technological force of the twenty-first century—one that is rapidly transforming governance, economy, and society. From facial recognition and predictive policing to automated credit scoring, AI now mediates crucial decisions that affect liberty, livelihood, and dignity. India's ambition to become a global AI innovation hub—through initiatives such as the *IndiaAI Mission*, *Digital India Programme*, and the forthcoming *Digital India Act*—reflects an optimistic vision of AI as a national growth multiplier. Yet this optimism exists alongside an alarming absence of a dedicated legal architecture that secures constitutional guarantees of privacy, equality, and free speech in an algorithmic age.

Unlike the European Union, which has enacted the world's first risk-based and rights-centric *AI Act (2024)*, India's current approach remains policy-driven but legally underdeveloped. While economic and strategic imperatives dominate the discourse, the constitutional implications of AI systems—especially regarding surveillance, bias, and opacity—are insufficiently addressed. This gap is no longer theoretical: AI systems are already being deployed in law enforcement, welfare distribution, and digital governance, often without judicial oversight or statutory safeguards. The risk is the emergence of a technocratic governance model, where opaque algorithms make or influence decisions without transparency, contestability, or human accountability.

From a constitutional standpoint, AI regulation in India must reconcile two imperatives:

- (1) The innovation imperative, which promotes AI development as a driver of economic growth and technological leadership; and**
- (2) The rights imperative, which mandates that technological progress remain subordinate to the Constitution's normative core—liberty, equality, and dignity.**

² *Justice K.S. Puttaswamy (Retd.) v. Union of India*, (2017) 10 SCC 1.

The Supreme Court's decision in *Justice K.S. Puttaswamy (Retd.) v. Union of India* (2017) provides the normative foundation for this reconciliation. By recognising privacy and informational self-determination as integral to life and liberty under Article 21, *Puttaswamy* articulated a structured proportionality test—legality, legitimate aim, necessity, and balancing—that now serves as a constitutional compass for regulating technology. Applying this framework to AI governance implies that no algorithmic system may restrict or shape rights without legal authority, necessity, and demonstrable safeguards for fairness, transparency, and human oversight.

This paper argues that India's forthcoming AI regulation must therefore move beyond voluntary ethics or innovation rhetoric towards a rights-based statutory model grounded in *Puttaswamy*-style proportionality. It advocates for a balanced architecture—one that fosters innovation through sandboxes and research carve-outs, yet enforces accountability through independent oversight, explainability, contestability, and algorithmic audits. In short, India must craft an AI regime that is not merely “Digital by Design” but “Constitutional by Design.”

Constitutional Risks of AI: Privacy, Equality & Free Speech

Artificial Intelligence, by its very design, transforms data into decisions. In doing so, it inescapably implicates the three foundational pillars of India's constitutional order — privacy (Article 21), equality (Article 14), and freedom of speech and expression (Article 19). When AI systems mediate governance, justice, employment, and information flows, they cease to be neutral tools; they become sites of constitutional power. Yet India's current AI ecosystem—marked by opaque algorithms, unregulated data flows, and expanding surveillance infrastructures—risks creating a digital state apparatus that operates outside traditional constitutional accountability.

The following subsections examine how unregulated AI directly threatens these constitutional guarantees and why a rights-based regulatory response is necessary.

1. Surveillance Capitalism and State Surveillance: Erosion of Privacy and Dignity (Article 21)

The proliferation of AI-driven surveillance systems—such as facial recognition technologies (FRTs), predictive policing algorithms, and data aggregation platforms—has fundamentally altered the citizen-state relationship. In cities like Delhi, Hyderabad, and Lucknow, real-time

FRT is already integrated into law enforcement and public safety networks without explicit statutory authority, judicial oversight, or sunset clauses.

This undermines the *Puttaswamy* mandate that any restriction on privacy must satisfy the four-fold proportionality test: (i) legality, (ii) legitimate aim, (iii) necessity/least restrictive means, and (iv) proportionality stricto sensu. Most AI-based surveillance systems fail at the very first step — legality — as they are deployed under vague executive circulars or internal police orders rather than a democratically enacted statute. Even if public safety is invoked as a legitimate aim, the absence of narrow tailoring, judicial warrants, and audit mechanisms means these systems fail the tests of necessity and balancing.

Moreover, mass surveillance through AI-powered analytics violates the right to informational self-determination, an essential facet of dignity and autonomy recognised in *Puttaswamy*. Continuous biometric tracking or behaviour profiling transforms citizens into data subjects under constant observation, producing a “chilling effect” on movement, association, and thought. As the Court warned in *Puttaswamy* and later reaffirmed in *Anuradha Bhasin v. Union of India* (2020), surveillance cannot become the new normal of governance. Without enforceable proportionality safeguards, AI surveillance risks converting a democratic state into a panoptic bureaucracy, where technology silently expands executive power beyond constitutional limits.

2. Algorithmic Bias and the Right to Equality (Article 14)

AI systems are often presented as “neutral” or “objective” decision-makers. Yet algorithmic decisions are only as fair as the data they are trained on. When historical or skewed datasets reflecting social hierarchies of caste, gender, religion, and region are encoded into algorithms, discrimination is no longer explicit—it becomes automated and concealed.

For example, AI-based hiring tools that disproportionately reject female applicants, or credit-scoring models that penalise minority neighbourhoods, replicate systemic bias while evading scrutiny under the guise of technical complexity. In law enforcement, predictive policing tools trained on biased arrest data risk reinforcing existing policing disparities, violating the equal protection guarantee under Article 14.

From a doctrinal standpoint, such systems offend both substantive and procedural equality.

Substantively, they produce discriminatory outcomes without rational justification. Procedurally, they deny individuals the right to be heard or to contest decisions, as algorithmic opacity (the “black box” problem) conceals the reasoning process. This violates the due process element of Article 21 and the “reasonableness” component of Article 14’s non-arbitrariness test (*E.P. Royappa v. State of Tamil Nadu, Maneka Gandhi v. Union of India*).

The constitutional remedy lies in algorithmic transparency and contestability: every consequential AI decision must be explainable, auditable, and subject to human review. Without these mechanisms, the right to equality risks being hollowed out by automated arbitrariness—what scholars term “algorithmic governance without accountability.”

3. Deepfakes, Misinformation, and Manipulation of Free Speech (Article 19)

AI-generated deepfakes and synthetic media now threaten the integrity of public discourse, elections, and journalistic truth. Manipulated videos and AI-enhanced misinformation can distort democratic deliberation, defame individuals, and erode citizens’ ability to form rational opinions. This raises complex tensions between protecting free expression (Article 19(1)(a)) and regulating harmful or deceptive content under Article 19(2).

The constitutional challenge lies in crafting responses that are proportionate and narrowly tailored. Blanket content takedowns, preemptive censorship, or algorithmic filtering by social media platforms risk creating a regulatory overreach, where legitimate speech is chilled alongside harmful content. The Supreme Court’s rulings in *Shreya Singhal v. Union of India* (2015) and *Anuradha Bhasin* emphasise that speech restrictions must be specific, transparent, and justified—not preventive or indefinite.

AI’s amplification of misinformation also introduces a new dimension: algorithmic amplification. Recommendation algorithms, designed to maximise engagement, often privilege sensational or polarising content. This raises questions about private platform accountability under constitutional values. While private corporations are not “State” actors, courts have increasingly recognised the horizontal application of rights in digital contexts (see *Kaushal Kishor v. Union of India*, 2023). Thus, platforms deploying generative or curatorial AI must adhere to constitutional reasonableness, ensuring that content moderation policies are transparent, non-discriminatory, and amenable to appeal.

The governance of AI-mediated speech therefore demands a dual proportionality framework:

- The State must ensure that content regulation is narrowly tailored to constitutional grounds; and
- Private platforms must design and deploy AI systems in a manner consistent with fairness, transparency, and accountability.

4. Due Process, Accountability, and the Rule of Law

Across these domains—surveillance, discrimination, and speech regulation—AI’s defining challenge to constitutionalism is the erosion of due process. Automated systems often produce decisions without prior notice, reasoning disclosure, or human oversight. This creates what legal theorists term “algorithmic opacity”: citizens are subject to decisions they cannot see, understand, or contest.

Article 21, as interpreted in *Maneka Gandhi* and *Puttaswamy*, mandates that any deprivation of life or liberty must follow “fair, just, and reasonable procedure.” If AI-driven systems determine access to welfare, employment, credit, or policing outcomes without procedural fairness, they effectively substitute algorithmic authority for constitutional legality. In a rule-of-law system, accountability cannot be outsourced to code.

Hence, the right to explanation, human oversight, and appeal must be treated not as policy preferences but as constitutional necessities—the digital analogues of natural justice. Only through such procedural guarantees can India prevent AI from becoming an instrument of unreviewable governance and preserve the integrity of its constitutional democracy.

India’s policy approach vs. Global ai regulation

A. India’s Current Approach Innovation-First, Regulation-Deferred

India currently regulates AI indirectly, through frameworks such as:

- Digital Personal Data Protection Act (DPDPA), 2023³: handles consent, but does not cover algorithmic bias, explainability, or AI accountability.
- Draft Digital India Act (2024/2025): concept note mentions “high-risk AI systems”, but no enforceable rights framework yet.

³ Digital Personal Data Protection Act, 2023, No. 22 of 2023.

- NITI Aayog’s National Strategy for AI (2018): promotes “AI for All,” but focuses on economic scalability more than rights & ethics.
- No mandatory AI audits, transparency, or liability laws yet exist.

In short policy enthusiasm exists, but legal enforceability is missing.

B. EU AI Act A Global Gold Standard

The EU AI Act (2024) is the world’s first binding AI law based on a risk-classification model:

AI Category	Regulatory Status
Unacceptable AI	Banned (e.g., social scoring, emotion recognition in workplace/schools)
High-Risk AI	Allowed but heavily regulated — requires audits, human oversight, transparency
Low/Minimal Risk AI	Largely unregulated — innovation-friendly

Key principle: the higher the risk to fundamental rights, the stricter the law.

C. OECD & UNESCO AI Principles

- Promote transparency, explainability, accountability, and human-in-the-loop oversight.
- Ethical AI must be auditable, bias-tested, and avoid harm to democracy and dignity.

D. India’s Lag No Legally Binding AI Risk Classification Yet

India lacks a formal mechanism to:

- Classify AI systems by risk to rights
- Enforce pre-deployment audit or certification
- Hold corporations or govt legally liable for AI harm

Result: AI enters public systems (policing, welfare, banking) without constitutional filtering.

The Case For Rights-Based Ai Regulation In India (Puttaswamy + Proportionality)

A rights-centric AI statute should take Puttaswamy’s four-part proportionality as its backbone: legality, legitimate aim, necessity/least-restrictive means, and strict balancing. Translated to AI governance:

- Legality: No high-risk AI deployment (esp. in State functions) without clear statutory authority, not mere executive circulars.

- Legitimate Aim: Objectives must be precise (e.g., fraud detection), not open-ended governance optimisation.
- Necessity / LRM: The deploying authority must show why non-AI or less-intrusive AI (smaller models, narrower features, edge processing, stronger de-identification) would not suffice.
- Balancing: Demonstrable rights-safeguards (human oversight, contestability, audits) must outweigh rights-risks (bias, opacity, chilling effects).

Crucially, proportionality should be front-loaded (pre-deployment risk assessments, sandboxing) and back-loaded (ongoing audits, incident reporting, redress).

A Model Indian Ai Statute: Core Architecture

1) Risk Classification (EU-style, India-adapted)

- Unacceptable-risk AI (ban): social scoring by the State; real-time remote biometric identification in public spaces (save judge-approved exigencies); emotion-recognition for workplace/schools; manipulative systems targeting children.
- High-risk AI (licence + audits): systems in law enforcement, welfare eligibility, credit scoring, hiring, education, healthcare, elections, and critical infrastructure.
- Limited/Minimal risk: notice/transparency duties; voluntary codes.

2) Algorithmic Impact Assessment (AIA) & Data Protection Impact Assessment (DPIA)

Before any high-risk deployment: publish (with necessary redactions) an AIA/DPIA covering purpose, data sources, potential bias, security, and mitigations, with an independent reviewer's opinion.

3) Independent AI Authority (AIAI)

A multi-member, security-cleared regulator with quasi-judicial powers to:

- accredit auditors, certify high-risk systems, issue use-stop orders, levy penalties, and mandate model retraining;
- run a public registry of certified high-risk AI; publish annual transparency statistics.

4) Human Oversight & Contestability

- Right to human review for consequential decisions (credit, welfare, policing flags, employment).

- Explainability-on-demand: meaningful rationale of key factors; adverse-action notices when AI significantly contributed to the outcome.
- Appeal & correction routes with time-bound remedies.

5) Fairness, Non-Discrimination, and Testing

- Mandatory pre- and post-deployment bias testing on protected attributes (sex, caste, religion, disability, etc.).
- Representative datasets; document data provenance; prohibit proxy discrimination unless justified and narrowly tailored.
- Model cards / data sheets documenting training data, intended use, and known limitations.

6) Data Governance & Security

- Align with DPDPA, 2023: purpose limitation, minimisation, retention controls; prohibit function-creep without renewed legal basis.
- Robust security: secure enclaves, access controls, audit logs, incident reporting within strict timelines.

7) Public Procurement & Vendor Liability

- Government contracts must include audit access, security warranties, indemnities, and source transparency (at least to the regulator).
- Joint liability for deployer and vendor in high-risk contexts unless due diligence and compliance are proved.

8) Research, Sandboxes, and Innovation

- Regulatory sandboxes with ethics guardrails; safe-harbours for bona fide research using privacy-preserving techniques (synthetic data, federated learning).

Sector-Specific Safeguards

A. Law Enforcement & National Security

- Judge-issued warrants for identity-search via facial recognition; strict necessity for any real-time deployment; default to post-facto forensic use.
- Prohibit dragnet watchlists; mandate audit trails and discovery in criminal trials so defendants can challenge algorithmic evidence (accuracy, bias, error rates).

B. Welfare & Financial Inclusion

- Algorithmic eligibility tools must disclose features used and error rates; provide manual override and appeal.

- For credit scoring/insurance, require adverse-action notices with concise explainability.

C. Elections & Information Integrity

- Label synthetic media; require platforms to implement traceable provenance (watermarking/content credentials) where feasible; rapid notice-and-contest for deepfake harms balanced by free-speech proportionality.

D. Employment & Education

- Ban emotion recognition; require bias-tested assessments; offer appeal and human review of automated grading/hiring.

Accountability, Enforcement & Liability

- Civil penalties scaled to turnover for negligent deployments causing rights harm; enhanced penalties for knowing/reckless violations.
- Private rights of action for individuals harmed by unlawful AI use; class/representative actions for systemic harms.
- Incident response: mandatory reporting of significant failures (security, bias spikes, wrongful denials/arrests).
- Whistle-blower protections for engineers and auditors reporting concealed risks.

Key Components

1. Accountability

- **Who is responsible?** Both the AI system deployer (e.g., government agency, company) and the AI developer/vendor can be held accountable.
- **Internal accountability mechanisms:** Organizations must maintain audit trails, document AI decisions, and implement internal compliance programs.
- **Human-in-the-loop responsibility:** Even automated systems require human oversight to ensure ethical and lawful outcomes.

2. Enforcement

- **Regulatory oversight:** An independent AI regulatory authority (like your proposed AIAI) monitors AI systems, audits deployments, and issues corrective measures.
- **Penalties for non-compliance:**
 - Civil fines proportional to the scale of harm or turnover.
 - Enhanced penalties for intentional or reckless violations.

- **Incident reporting:** Mandatory reporting of AI failures, bias spikes, security breaches, or wrongful decisions to the regulator.
- **Whistleblower protection:** Engineers, auditors, or employees who report hidden risks are legally protected from retaliation.

3. Liability

- **Civil liability:** Individuals harmed by AI (e.g., wrongful denial of credit, misidentification by facial recognition) can sue for damages.
- **Joint liability:** Deployers and vendors may share liability unless they demonstrate due diligence and regulatory compliance.
- **Sector-specific rules:** Stricter liability for high-risk areas like law enforcement, healthcare, elections, and financial services.
- **Corrective action:** Regulators can mandate retraining of AI models, suspension of deployment, or compensation to affected individuals.

Purpose in AI Regulation

- Ensures compliance with constitutional rights (Articles 14, 19, 21).
- Creates deterrence against careless or harmful AI deployment.
- Balances innovation with public safety and accountability, so AI can be developed responsibly without undermining trust.

Anticipating Objections

“Regulation will chill innovation.”

A risk-tiered model targets high-risk deployments, leaving low-risk uses largely free. Clear rules reduce uncertainty, encouraging responsible innovation.

“Explainability degrades accuracy.”

Require fit-for-purpose explanations: global model cards + local reason codes for decisions. Use post-hoc techniques where native interpretability is infeasible, coupled with human review.

“Security needs secrecy.”

Secrecy about targets can coexist with ex ante warrants, independent audits, and aggregate transparency just as in other sensitive domains.

Conclusion

India can and should craft an AI statute that champions innovation while constitutionalising accountability. Anchoring regulation in Puttaswamy’s proportionality, adopting an EU-style risk taxonomy, mandating audits, human oversight, contestability, explainability, and erecting an independent AI regulator would align AI deployment with Articles 14, 19, and 21. With pragmatic sandboxes, clear liability, and sector-specific guardrails, India can avoid techno-solutionism and rights-blind deployments, building an AI ecosystem that is innovative, trustworthy, and fundamentally constitutional.

