INTERNATIONAL JOURNAL FOR LEGAL RESEARCH AND ANALYSIS

Open Access, Refereed Journal Multi Disciplinary
Peer Reviewed



www.ijlra.com

ISSN: 2582-6433

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.

ISSN: 2582-6433

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IJLRA welcomes submissions from all branches of law, provided the work is original, unpublished, and submitted in accordance with the prescribed submission guidelines.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.

ISSN: 2582-6433

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.

ISSN: 2582-6433

ARTIFICIAL INTELLIGENCE AND CORPORATE SOCIAL RESPONSIBILITY: ETHICAL GOVERNANCE, ENVIRONMENTAL SUSTAINABILITY AND RESPONSIBLE INNOVATION

AUTHORED BY – BADARQUAH & SARAH MUMTAZ
B.A. LL.B. (Hons.), Amity University Kolkata

Abstract

The corporation of the algorithmic age no longer discharges its social obligations principally through philanthropy. It discharges or evades them through code. As artificial intelligence migrates from an operational curiosity to the central nervous system of corporate decision-making, Indian corporate law continues to conceptualise Corporate Social Responsibility ("CSR") as a quantum of profit to be spent, rather than a quality of conduct to be governed. This paper argues that the statutory architecture of CSR under Section 135 of the Companies Act, 2013, and the enumerated activities of Schedule VII, were designed for an industrial paradigm of externalities and are structurally incapable of capturing the distinctive harms of algorithmic enterprise: discriminatory automated decision-making, opaque "black-box" models, the erosion of informational privacy, and a prodigious and largely undisclosed environmental footprint.

It further contends that the fiduciary mandate of Section 166 — which commands directors to act for the benefit of employees, the community and the environment — already supplies a latent but unactivated basis for algorithmic accountability, and that the Securities and Exchange Board of India's Business Responsibility and Sustainability Reporting framework must be recalibrated to mandate disclosure of AI data ethics and automated human-resource systems. Reading the Companies Act alongside the National Guidelines on Responsible Business Conduct, the constitutional guarantees of equality and a wholesome environment, and the emerging transnational consensus reflected in the EU Artificial Intelligence Act, the OECD AI Principles and India's own AI Governance Guidelines of November 2025, the paper proposes a doctrinal and statutory model of "Responsible Innovation." That model would treat AI-induced environmental degradation as a core CSR liability, embed algorithmic impact assessments, mandatory human-in-the-loop safeguards and board-level AI ethics committees

ISSN: 2582-6433

into the governance fabric of the company, and thereby convert corporate engagement with AI from reactive compliance into proactive, constitutionally congruent stewardship ———

Keywords: corporate social responsibility, artificial intelligence governance, responsible innovation, algorithmic accountability, environmental sustainability

I. Introduction

The modern corporation has quietly undergone a transformation more consequential than any amendment to its memorandum. Where it once deliberated through committees of persons, it increasingly deliberates through systems of code: credit is extended or denied by a model, a résumé is ranked before a human ever reads it, a price is set by an optimisation routine, and the "discretion" of the firm is exercised, at scale and at speed, by statistical artefacts whose logic is frequently inscrutable even to those who deploy them. This migration of agency from the boardroom to the algorithm forces a question that Indian corporate law has not yet confronted squarely: if a company increasingly acts through artificial intelligence, through what law is that action to be rendered socially responsible? Corporate Social Responsibility in India is, uniquely among major jurisdictions, a creature of statute.

Section 135 of the Companies Act, 2013, married to the enumerated heads of Schedule VII, transformed an ethical aspiration into a measurable, auditable spending obligation.[1] That design — a mandated percentage of profit channelled into charitable heads — was a signal legislative achievement for an economy of factories, effluent and philanthropy. It is, however, a poor fit for an economy of inference. The harms that flow from a mis-specified hiring algorithm or a biased credit model are neither discharged by building a school nor captured by a line item of "amount spent." They are harms of conduct, distributed invisibly across data subjects, and they demand a conception of responsibility measured not in rupees disbursed but in governance exercised. This paper advances three connected claims.

First, that the statutory definition of CSR is substantively under-inclusive in the algorithmic age: Section 135 and Schedule VII speak the language of philanthropy, not of accountability, and the Securities and Exchange Board of India's ("SEBI") Business Responsibility and Sustainability Reporting ("BRSR") framework, for all its ambition, does not yet compel disclosure of the data-ethics and automated-decision practices that most affect stakeholders.[2]

1 Companies Act, 2013, § 135, No. 18 of 2013, India Code (2013); id. sch. VII.
2 See infra Part II.B.

ISSN: 2582-6433

Second, that the fiduciary architecture of the Act — principally the director's duty under Section 166(2) to act for the benefit of the company's members, employees, the community and the environment — already contains the normative resources to require algorithmic accountability, but that these resources lie dormant for want of interpretive and institutional activation.[3]

Third, that the environmental cost of AI — the electricity, water and hardware consumed in training and operating large-scale models — is a genuine corporate externality that the law should recognise as a core CSR liability rather than an invisible overhead. The argument proceeds in four substantive movements. Part II traces the evolution of CSR and exposes the statutory and disclosure gaps that AI reveals. Part III interrogates the fiduciary voids in corporate governance, situating algorithmic bias, automated decision-making and the tension between proprietary secrecy and transparency within the duties of directors. Part IV addresses the environmental paradox of a "clean" digital economy built on carbon- and water-intensive infrastructure, and reads it against Principle 6 of the National Guidelines on Responsible Business Conduct.

Part V synthesises these strands into a model of "Responsible Innovation," drawing on comparative instruments — the EU Artificial Intelligence Act, the OECD AI Principles, and India's AI Governance Guidelines of 2025 — and anchoring the proposal in the constitutional commitments to equality and sustainable development. Part VI concludes.

II. The Evolution Of Csr In The Algorithmic Age

The intellectual lineage of CSR runs from a shareholder-primacy orthodoxy — famously, that the social responsibility of business is to increase its profits[4] — toward a stakeholder conception in which the enterprise is accountable to the constituencies its activities affect. India's distinctive contribution was to statutorily codify that conception. Understanding why that codification now strains under the weight of AI requires a close reading of its component parts.

A. Deconstructing Section 135 and Schedule VII

Section 135 obliges every company which meets defined thresholds of net worth, turnover or net profit to constitute a CSR Committee and to spend, in every financial year, at least two per

3 Companies Act, 2013, § 166(2).

4 The locus classicus is Milton Friedman, A Friedman Doctrine—The Social Responsibility of Business Is to Increase Its Profits, N.Y. Times (Sept. 13, 1970). For the countervailing stakeholder tradition, see R. Edward Freeman, Strategic Management: A Stakeholder Approach (1984).

ISSN: 2582-6433

cent of the average net profits made during the three immediately preceding financial years on activities falling within its CSR policy.[5] The permissible objects of that expenditure are not unlimited. They are limited to the heads enumerated in Schedule VII. Eradicating hunger and poverty, promoting education, ensuring environmental sustainability, protecting national heritage, and the like.[6] Two structural features of this design are decisive for the present inquiry. The first is that Section 135 regulates the destination of money, not the manner of operation. It is a measure that looks at inputs and outputs: a sum is computed, a sum is spent, and compliance is substantially established. The Act is studiously indifferent to how the company earns the profit from which the two per cent is carved.

A firm may deploy a demonstrably discriminatory lending model throughout the year and nonetheless achieve impeccable CSR compliance by funding a rural health initiative. The law therefore separates the social consequences of the company's core business. This separation is especially problematic where AI is embedded in core business operations.

The second is that Schedule VII is an enumeration of philanthropic causes, not a taxonomy of corporate harms. Nowhere does it contemplate algorithmic accountability, data-protection stewardship, the auditing of automated systems, or the remediation of AI-induced bias. A company genuinely wishing to invest its CSR budget in, say, an independent audit of its own recruitment algorithm or a programme to retrain workers displaced by automation would find no secure statutory home for that expenditure, and might be told that such spending is a business expense rather than CSR.[7] The result is a definitional paradox: the most socially consequential activity of an algorithmically driven firm is, by design, the activity its CSR obligation cannot reach.

B. The BRSR Framework and the Disclosure Deficit

If Section 135 governs spending, the architecture of corporate disclosure is where the social conduct of the firm is meant to be rendered visible. SEBI's BRSR framework, introduced in 2021 and made mandatory for the top one thousand listed entities by market capitalisation, was a substantial advance: it aligned mandatory reporting with the nine principles of the National Guidelines on Responsible Business Conduct and required granular disclosure across

5 Companies Act, 2013, § 135(1), (5). The thresholds are a net worth of ₹500 crore or more, a turnover of ₹1000 crore or more, or a net profit of ₹5 crore or more, during the immediately preceding financial year. See also Companies (Corporate Social Responsibility Policy) Rules, 2014, Gazette of India, pt. II sec. 3(i) (Feb. 27, 2014).
6 Companies Act, 2013, sch. VII. The Schedule has been read liberally by the Ministry of Corporate Affairs but remains a closed enumeration of permissible activities, not a standard of conduct.
7 Expenditure that is undertaken in the normal course of business, or that benefits only the company's own employees in a manner that is not "CSR," is generally excluded. See Companies (CSR Policy) Rules, 2014, r. 2(1)(d) (defining CSR to exclude activities undertaken in pursuance of normal course of business).

ISSN: 2582-6433

environmental, social and governance dimensions.[8] SEBI subsequently introduced the "BRSR Core" which is a focused set of key performance indicators subject to reasonable assurance the objective of which is to lend credibility to what had been largely self-reported ESG data.[9] Yet for all its complexity, the BRSR framework exhibits a conspicuous lacuna: it does not dictate disclosure of the dimension of corporate conduct that AI has made most salient.

There is no required revelation of whether and how the entity uses automated decision-making in hiring, promotion, credit, pricing or dismissal; no mandated account of the data on which its models are trained or the attribution and consent-basis of that data; no obligation to report the results of bias testing or algorithmic audits; and no disaggregated disclosure of the energy and water consumed by the entity's digital infrastructure.[10] The consequence is that a listed entity may arrange consequential automated systems across its workforce and customer base while disclosing nothing about them, because the reporting taxonomy contains no field in which such practices must appear. What is not asked is not reported; what is not reported cannot be scrutinised by the market, by regulators, or by affected stakeholders. This disclosure deficit is not just a gap to be filled for completeness. Disclosure is the mechanism by which the Indian ESG regime alters private conduct into public accountability.

A reporting framework that is silent on AI therefore does more than omit a topic; it communicates, by its silence, that algorithmic governance is not a matter of corporate responsibility at all. Comparative developments cut sharply the other way: the European Union's Corporate Sustainability Reporting Directive develops the universe of mandatory sustainability information and entrenches a "double materiality" logic,[11] while the AI Act layers transparency and documentation duties onto high-risk systems.[12] India's disclosure architecture, if it is to remain credible, cannot treat the most transformative technology of the

8 Sec. & Exch. Bd. of India, Circular No. SEBI/HO/CFD/CMD-2/P/CIR/2021/562, Business Responsibility and Sustainability Reporting by Listed Entities (May 10, 2021). BRSR replaced the earlier Business Responsibility Report regime and became mandatory for the top 1000 listed entities by market capitalisation from the financial year 2022–23.

9 Sec. & Exch. Bd. of India, Circular No. SEBI/HO/CFD/CFD-SEC-2/P/CIR/2023/122, BRSR Core – Framework for Assurance and ESG Disclosures for Value Chain (July 12, 2023) (introducing a phased "reasonable assurance" requirement on a defined set of core ESG attributes).

10 The BRSR's Principle 3 (employee well-being) and Principle 5 (human rights) disclosures capture aspects of workforce treatment, and its Principle 6 disclosures capture energy and water at an entity level, but none is drafted to surface algorithmic human-resource systems or to isolate the footprint of AI workloads. Cf. Principle 9 (consumer responsibility), which addresses data privacy in general terms without reaching model governance.

11 Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 (Corporate Sustainability Reporting Directive), 2022 O.J. (L 322) 15, art. 19a (institutionalising the principle of "double materiality").

12 Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial Intelligence Act), 2024 O.J. (L, 12.7.2024), arts. 11, 13 (technical documentation and transparency obligations for high-risk AI systems).

ISSN: 2582-6433

corporate decade as beneath mention.

C. The Disruption of the Philanthropy Paradigm

Beyond rules and clear signals there is an idea problem. The charity view thinks that the place where a company makes money is separate from the place where it gives money back, to people. In that view CSR is a thought after business is done.. Using data predictions and automatic decisions breaks that split. When a computer chooses who to hire what rate a borrower must pay, which neighborhoods a service covers or which workers get the shifts the company's social impact is made right then. The social footprint is not a gift; it is a choice made from the start.

This text carries a doctrinal implication. A responsibility regime that engages after value creation— the Section 135 model— will always arrive too late to govern harms that are manufactured within value creation itself. The algorithmic firm therefore needs a conception of responsibility that's concurrent rather than consequent: responsibility must be embedded in the design, testing, deployment and monitoring of the systems through which the algorithmic firm acts not added as a charitable coda. The rest of this paper focuses on locating and where needed building, that conception, within Indian law.

III. Ethical Governance And Fiduciary Voids

If the CSR law gives the tool corporate governance law gives the right one. But it is not being used. The responsibilities of directors the ways the board checks on things and the rules about sharing information make up the system that keeps the company in check. This part says that the duty, in Section 166 already covers harm caused by algorithms but current practices have not dealt with the effects of AI. The law needs to balance keeping algorithms secret with making sure the public can see what they are doing.

A. Section 166 and the Latent Stakeholder Mandate

Section 166(2) of the Companies Act requires every director to act in faith to promote the objects of the company for the benefit of all members as a whole. It also requires directors to act in the interests of the company of its employees of its shareholders of the community and to protect the environment. In a single provision it pulls directors into serving employees, the community and the environment – the groups most exposed to harm. A director who permits a hiring model that systematically disadvantages a protected class is, on a reading not acting in the best interests of the company's employees and the community. A director who signs off on data-processing practices that risk the privacy of millions is, at least arguably breaching the duty. The difficulty is not in the words of the duty itself. In whether it can be enforced and what

ISSN: 2582-6433

it really means in practice. Section 166 has not led to legal precedent. Courts have not often stepped in to judge the substance of director decisions when they affect groups outside shareholders.

The strong principle of business-judgment deference in law means judges usually avoid questioning operational choices made by directors. Because of this Section 166 today acts like a suggestion than a real legal rule that can be enforced. This paper argues that artificial intelligence provides both the reason and the way to bring this duty to life. When the harm caused by a managed algorithm becomes predictable. Through the audit and impact assessment tools discussed in Part V. Then a director's failure to put in place proper oversight of such systems becomes a clear failure of the duty of care under Section 166(3). This kind of failure is similar, to the standing rule that directors must have information and reporting systems that are good enough to handle the risks the company faces. In short oversight of algorithms is not a duty. It is the existing duty of care applied to an important kind of risk.

B. The Regulatory Disconnect: Bias, Automation and Privacy

The negative externalities of corporate AI fall into three recurring families, and in each the gap between the harm and the governing law is stark. The first is algorithmic bias. The Models are trained on historical data inherit and frequently amplify the discriminatory patterns embedded in that data, producing outcomes that disadvantage individuals on grounds such as caste, religion, gender or disability without any discriminatory intent on the part of any human actor.[13] Indian anti-discrimination law, framed around state action and enumerated grounds, is poorly equipped to reach discrimination effected by a private firm's model, and offers the affected individual little by way of remedy or even detection. The second family concerns automated decision-making as such — the displacement of human judgment by systems whose reasoning is neither transparent nor contestable.

The person denied credit, filtered out of a hiring pool, or flagged by a fraud model is frequently unable to learn that an automated system decided the matter, let alone why, or to seek human reconsideration. Comparative law has begun to supply a vocabulary of rights here: the General Data Protection Regulation confers a qualified right not to be subject to solely automated decisions producing legal or similarly significant effects, together with rights to meaningful information about the logic involved.[14] India's Digital Personal Data Protection Act, 2023, by

13 The foundational legal treatment is Solon Barocas & Andrew D. Selbst, Big Data's Disparate Impact, 104 Calif. L. Rev. 671 (2016). For the broader social critique, see Cathy O'Neil, Weapons of Math Destruction (2016); Virginia Eubanks, Automating Inequality (2018); and Safiya Umoja Noble, Algorithms of Oppression (2018).

14 Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation), 2016 O.J. (L 119) 1, art. 22; see also id. arts. 13–15 (information and access rights

ISSN: 2582-6433

contrast, is largely silent on automated decision-making, omitting any analogue to Article 22 and thereby leaving a significant governance void at precisely the point where AI most directly touches individual rights.[15] The third family is informational privacy.

The Supreme Court's recognition of privacy as a fundamental right intrinsic to the guarantees of life and personal liberty supplies the constitutional backdrop,[16] but the corporate data practices that feed AI — mass collection, inference of sensitive attributes, repurposing of data beyond its original context — frequently outrun both the letter of data-protection law and the capacity of individuals to protect themselves. The common thread across all three families is that the governing legal regimes were drafted for a world of human decision-makers and discrete data transactions, and they struggle to discipline systems that decide continuously, opaquely and at population scale. This is the regulatory disconnect that corporate governance, properly activated, must help to close.

C. Proprietary Secrecy Versus Algorithmic Transparency

Any serious proposal to open algorithms to scrutiny collides with a powerful countervailing interest: the protection of proprietary models as trade secrets and intellectual property. A trained model, its architecture, its weights and its training corpus represent substantial investment and competitive advantage, and firms have strong and legitimate incentives to keep them confidential. The law has long protected such confidential business information,[17] and the public demand for "algorithmic transparency" cannot be satisfied by simply compelling disclosure of source code to the world, which would be both unworkable and self-defeating. The resolution lies in disaggregating "transparency" into its functional components. What stakeholders and regulators require is rarely the model's source code; it is assurance about the model's behaviour — that it has been tested for bias, that its decisions can be explained to those they affect, and that its performance is monitored over time.

These assurances can be delivered through qualified rather than public transparency: independent third-party audits conducted under confidentiality, disclosure of impact-assessment outcomes rather than underlying code, regulator access under protective conditions,

extending to the logic of automated processing).

15 Digital Personal Data Protection Act, 2023, No. 22 of 2023, India Code (2023). The Act establishes consent and purpose-limitation obligations but does not confer a right against solely automated decision-making or a right to an explanation of algorithmic logic.

16 Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 S.C.C. 1 (holding the right to privacy to be a fundamental right under Article 21 of the Constitution of India).

17 In India, trade-secret protection is a matter of common-law breach of confidence and contract rather than a dedicated statute; the obligation is reinforced by India's commitments concerning undisclosed information under the Agreement on Trade-Related Aspects of Intellectual Property Rights art. 39, Apr. 15, 1994, 1869 U.N.T.S. 299.

ISSN: 2582-6433

and explanations calibrated to the affected individual rather than the general public.[18] So conceived, the conflict between intellectual property and transparency is not irreconcilable; it is a design problem, solved by channelling scrutiny through accountable intermediaries rather than through wholesale publication. Part V operationalises this insight.

IV. Environmental Sustainability And The Tech Paradox

There is a comfortable assumption that the digital economy is a clean economy — that software, unlike smokestacks, imposes no ecological burden. The assumption is false, and increasingly so. The computational infrastructure on which artificial intelligence depends consumes electricity, water and raw materials on a scale that is now material to national and corporate climate commitments, yet which remains, for the most part, undisclosed. This Part sets out the footprint, reads it against the environmental obligations of Indian responsible-business law, and argues for its recognition as a core CSR liability.

A. The Ecological Footprint of Artificial Intelligence

The environmental cost of AI arises across three vectors. The first is energy. Training a large model is an extraordinarily electricity-intensive undertaking, and — contrary to common intuition — it is the inference phase, the continuous servicing of user queries, that accounts for the larger share of operational energy once a model is deployed.[19] Estimates for 2025 placed the carbon emissions attributable to AI activity in the order of eighty million tonnes of carbon dioxide, a figure comparable to the annual emissions of a major metropolis,[20] while global data-centre electricity consumption has been estimated at several hundred terawatt-hours and is projected to rise steeply through 2030.[21] The second vector is water. Data centres consume vast quantities of freshwater for cooling, both directly on site and indirectly through the thermoelectric generation of the power they draw.

Public estimates suggest that AI-related systems consumed hundreds of billions of litres of

18 The distinction between transparency to the public and accountability to a trusted auditor or regulator is central to the governance literature. See Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information (2015) (arguing for qualified disclosure regimes that preserve competitive confidentiality while enabling meaningful oversight).

19 Multiple independent assessments converge on the point that inference, not training, dominates lifecycle energy once a model is in large-scale deployment. For the influential early analysis of training-phase costs, see Emma Strubell, Ananya Ganesh & Andrew McCallum, Energy and Policy Considerations for Deep Learning in NLP, in Proc. 57th Ann. Meeting Ass'n for Computational Linguistics 3645 (2019).

20 See Alex de Vries-Gao, [study on AI energy and emissions], Patterns (Cell Press) (2025), https://doi.org/10.1016/j.patter.2025.101430 (estimating AI-attributable emissions of roughly 80 million tonnes of CO_2 for 2025). The author should verify the exact title, volume and figures of this source before citation.

21 See U.N. Env't Programme & related international assessments (2025) (estimating global data-centre electricity consumption on the order of 448 TWh in 2025, with projections exceeding 945 TWh by 2030). Exact figures and the issuing body should be confirmed.

ISSN: 2582-6433

water in 2025, and that the training of a single frontier model can carry a water footprint measured in hundreds of millions of litres.[22] This burden is acutely regressive in geographic terms: data centres are frequently sited in, and draw upon, regions already experiencing water stress, placing industrial cooling demand in direct competition with domestic and agricultural needs — a distributional concern of obvious salience in the Indian context. The third vector is the hardware lifecycle. The specialised processors on which AI depends embody energy-intensive manufacturing, the mining of critical minerals, and — at end of life — a growing stream of electronic waste whose disposal raises its own environmental and human-rights concerns. The footprint of AI is therefore not confined to the operation of a data centre; it extends backward into extraction and manufacture and forward into disposal, and a responsible accounting must span the whole.

A single feature unites these vectors and frustrates their governance: opacity. Firms seldom disaggregate the energy, water and materials consumed by their AI workloads from their general operations, and the resulting absence of reliable, comparable data makes it difficult for regulators, investors or the public to hold any enterprise to account for its AI-specific footprint. The first task of environmental governance in this domain is thus not restriction but measurement and disclosure.

B. NGRBC Principle 6 and the Reconciliation Problem

India's National Guidelines on Responsible Business Conduct supply the normative anchor. Principle 6 enjoins that businesses should respect and make efforts to protect and restore the environment,[23] and the BRSR disclosures operationalise that principle by requiring reporting on energy consumption, water withdrawal and greenhouse-gas emissions. The reconciliation problem is that a technology firm may present itself, and report itself, as environmentally benign — it operates no factory, discharges no effluent — while its true ecological burden is displaced into the data centres, often third-party and offshore, on which its models run. The entity-level, facility-blind structure of present disclosure allows this displacement to pass unrecorded. Principle 6, read purposively, does not tolerate this.

Its injunction to "protect and restore" the environment is indifferent to whether the harm issues from a chimney or a server farm, and its reach extends, in the logic of responsible-business

22 See Shaolei Ren et al., Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models (2023) (quantifying the direct and indirect water footprint of training and inference); see also de Vries-Gao, supra note 23 (estimating AI-related water consumption on the order of 765 billion litres for 2025). Figures should be verified against the primary studies.

23 Ministry of Corp. Affairs, Gov't of India, National Guidelines on Responsible Business Conduct, princ. 6 (2019) ("Businesses should respect and make efforts to protect and restore the environment."). The NGRBC are the normative foundation on which SEBI's BRSR disclosures are built.

ISSN: 2582-6433

conduct, across the value chain rather than stopping at the factory gate. A tech-enterprise whose carbon- and water-intensive digital infrastructure is material to its environmental impact cannot credibly reconcile itself with Principle 6 while declining to measure, disclose and mitigate that impact. The guideline already demands what the disclosure framework does not yet ask; the task is to align the second with the first.

C. AI-Induced Degradation as a Core CSR Liability

The foregoing supports a reconceptualisation. The environmental cost of AI should not be treated as an ordinary operating expense, invisible on the social ledger; it should be recognised as a core CSR liability — a harm that the company is responsible for measuring, disclosing, mitigating and, where appropriate, restoring. Three mechanisms follow. First, measurement and disclosure: entities of material digital footprint should be required to report AI-attributable energy, water and emissions on a disaggregated basis, so that the ledger is no longer blank. Second, restoration: where an enterprise's AI operations impose a quantifiable ecological burden, targeted CSR allocation toward ecology-restoring activities within Schedule VII's environmental head offers a statutorily grounded route to internalising that burden.[24]

Third, "Green AI": the law should incentivise architectural and operational efficiency — model compression, carbon-aware scheduling of workloads, renewable-powered and water-efficient data centres — as a first-order design objective rather than an afterthought.[25] Each of the three principles that India's own AI governance framework associates with sustainability — and the global movement toward water replenishment and renewable-powered compute to which leading operators have begun to commit — points in the same direction: environmental responsibility must be designed into the AI enterprise, not donated around it.

V. Doctrinal Solutions And Policy Reset: Responsible Innovation

The diagnosis of the preceding Parts converges on a single prescription: Indian corporate law must move from a posture of reactive compliance — spend the percentage, file the report — to one of proactive technological governance. The organising idea is "Responsible Innovation," a principle, increasingly embedded in international instruments, that the capacity to anticipate, assess and mitigate the societal consequences of a technology is itself a condition of its

24 Companies Act, 2013, sch. VII cl. (iv) (ensuring environmental sustainability, ecological balance, conservation of natural resources). This head is capacious enough to accommodate water-replenishment, afforestation and renewable-transition expenditure tied to an entity's digital footprint.

25 The "Green AI" framing — prioritising computational and energy efficiency as a research and deployment value — is associated with Roy Schwartz, Jesse Dodge, Noah A. Smith & Oren Etzioni, Green AI, 63 Comm. ACM 54 (2020).

ISSN: 2582-6433

legitimate deployment.[26] This Part sets out how that principle can be mainstreamed into corporate law, the concrete protocols it entails, and the comparative and constitutional scaffolding on which it rests.

A. Mainstreaming Responsible Innovation into Corporate Law

Two routes, interpretive and statutory, are available and should be pursued together. The interpretive route requires no legislative amendment: it reads the existing duties of Section 166 as already encompassing the obligation to govern the company's algorithmic systems with due care, and reads Schedule VII's environmental and other heads as capacious enough to accommodate AI-accountability and ecology-restoration expenditure. On this route, regulators and courts give present content to dormant text. The statutory route supplies what interpretation cannot: a clear mandate. It would amend the CSR and disclosure architecture to name algorithmic accountability as a dimension of responsible conduct, require board-level governance of AI risk, and compel disaggregated disclosure of AI data practices and environmental footprint. The two routes are complementary — interpretation can act immediately while legislation is debated, and legislation can entrench what interpretation begins.

B. Operational Protocols: Impact Assessments, Human-in-the-Loop and Ethics

Boards Responsible Innovation must be given operational teeth through three institutional protocols, each adaptable to the Indian corporate form. The first is the mandatory Algorithmic Impact Assessment ("AIA"): before deploying a consequential automated system, and periodically thereafter, the company should be required to document the system's purpose, data provenance, foreseeable risks of bias and error, tested performance across affected groups, and mitigation measures. The AIA is the AI analogue of the environmental impact assessment, and its outputs — not the underlying code — form the basis of qualified disclosure and audit.[27] The second protocol is the Human-in-the-Loop ("HITL") requirement: for decisions producing legal or similarly significant effects on individuals — hiring, dismissal, credit, access to essential services — the architecture must preserve meaningful human review and the affected person's ability to contest the outcome and obtain reconsideration.

HITL converts the abstract right against purely automated decision-making into an operational

26 The concept of "responsible research and innovation" is developed in, among others, Jack Stilgoe, Richard Owen & Phil Macnaghten, Developing a Framework for Responsible Innovation, 42 Res. Pol'y 1568 (2013), and has since been absorbed into AI-ethics frameworks worldwide. See also Luciano Floridi et al., AI4People—An Ethical Framework for a Good AI Society, 28 Minds & Machines 689 (2018).

27 The AIA model draws on, among others, the data-protection impact assessment of GDPR art. 35 and the conformity-assessment and risk-management architecture of the EU AI Act, arts. 9, 43. See Regulation (EU) 2024/1689, arts. 9, 43.

ISSN: 2582-6433

design constraint, and it maps directly onto the human-oversight obligations now emerging in comparative regulation.[28] The third protocol is institutional: the establishment of an internal AI Ethics Board or committee, ideally with independent membership and a reporting line to the board of directors, charged with overseeing the AIA process, reviewing high-risk deployments, and advising the board on algorithmic risk. Such a body translates the director's duty of care into a standing governance mechanism, discharging the oversight obligation that Section 166(3), read in light of established risk-oversight doctrine, increasingly implies. Together, the AIA, HITL and the ethics board convert Responsible Innovation from slogan into structure.

C. Comparative Harmonisation and Constitutional Anchoring

India does not legislate in a vacuum, and the architecture proposed here is consonant with the transnational consensus now crystallising. The EU Artificial Intelligence Act, whose principal obligations become applicable in August 2026, institutes a risk-tiered regime of prohibited practices, high-risk conformity assessment, transparency duties and human oversight,[29] and it supplies a mature template for the impact-assessment and oversight protocols advanced above. The OECD AI Principles — the first intergovernmental standard on AI, committing adherents to inclusive growth, human-centred values, transparency, robustness and accountability — furnish the normative vocabulary,[30] as does the UNESCO Recommendation on the Ethics of Artificial Intelligence.[31] Most significantly for a domestic audience, India has now articulated its own position.

The India AI Governance Guidelines released by the Ministry of Electronics and Information Technology in November 2025 — built upon the Reserve Bank of India's framework for the responsible and ethical enablement of AI and structured around seven guiding principles, or sutras, spanning trust, a people-first orientation, innovation over restraint, fairness and equity, accountability, being understandable by design, and safety, resilience and sustainability — adopt a deliberately enabling, principle-based posture that favours the adaptation of existing law over prescriptive new regulation.[32] This posture is a double-edged inheritance. It is nimble

28 Human oversight is a cornerstone obligation for high-risk systems under the EU framework. See Regulation (EU) 2024/1689, art. 14 (human oversight); cf. Regulation (EU) 2016/679, art. 22 (right not to be subject to solely automated decisions).

29 Regulation (EU) 2024/1689, arts. 5 (prohibited practices), 6 & annex III (high-risk classification), 9–15 (requirements for high-risk systems), 50 (transparency). Prohibited-practice provisions applied from 2 February 2025; the general body of obligations, including those for high-risk systems under Annex III and the transparency duties, applies from 2 August 2026, with certain product-related high-risk obligations following in 2027.

30 Org. for Econ. Co-operation & Dev., Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449 (adopted May 22, 2019; amended 2024). India is an adherent to the OECD AI Principles.

31 U.N. Educ., Sci. & Cultural Org., Recommendation on the Ethics of Artificial Intelligence (Nov. 23, 2021).

32 Ministry of Electronics & Info. Tech., Gov't of India, India AI Governance Guidelines (Nov. 5, 2025). The Guidelines comprise seven guiding principles ("sutras"), six pillars of recommendations, and an action plan, and designate MeitY as the nodal ministry coordinating a proposed AI Governance Group. They build on the Report

ISSN: 2582-6433

and innovation-friendly, but its reliance on existing law to carry the governance burden only sharpens the argument of this paper: if existing law is to do the work, the corporate-law instruments examined here — Sections 135 and 166, the BRSR framework and the NGRBC — must be interpreted and amended to bear it.

The MeitY Guidelines themselves acknowledge that their principles are to serve as interpretive aids for sectoral regulators and courts in applying statutory obligations already in force; SEBI and the Ministry of Corporate Affairs are precisely such regulators, and corporate law is precisely such a field. The deepest anchor, however, is constitutional. Algorithmic discrimination implicates the guarantee of equality before the law and the equal protection of the laws under Article 14, whose settled doctrine forbids arbitrary and unreasonable classification — a standard that a biased or inscrutable model may fail as surely as any arbitrary administrative act.[33] The environmental burden of AI implicates the right to a wholesome environment read into the right to life, and the principle of sustainable development that Indian courts have woven into the constitutional fabric.[34] And informational privacy, the interest most directly threatened by data-hungry models, is now a recognised fundamental right.

Harmonising corporate AI governance with these constitutional commitments is therefore not an act of policy generosity; it is the alignment of private technological power with the basic law of the Republic.

VI. Conclusion

The main argument of the paper is that the corporation's adoption of artificial intelligence has gone beyond the legal framework by which India regulates corporate responsibility. Since corporate social responsibility (CSR) was set up as a statutory tax on profits and included a list of philanthropic activities, it was designed for a world in which a company's social impact occurred as a result of its business operations. It is only with the introduction of artificial intelligence that this social impact has moved upstream and become part of the design of the systems through which the company now operates, thereby revealing the philanthropic model

of the Subcommittee on 'AI Governance Guidelines Development' (Jan. 2025) and the Reserve Bank of India's FREE-AI framework (2025). See also NITI Aayog, National Strategy for Artificial Intelligence #AIForAll (2018); NITI Aayog, Responsible AI for All (2021).

33 India Const. art. 14. The guarantee against arbitrariness, developed from E.P. Royappa v. State of Tamil Nadu, (1974) 4 S.C.C. 3, and Maneka Gandhi v. Union of India, (1978) 1 S.C.C. 248, supplies a constitutional vocabulary for challenging arbitrary algorithmic classification, at least where state or state-adjacent action is engaged.

34 The right to a wholesome environment is an aspect of Article 21. See Subhash Kumar v. State of Bihar, (1991) 1 S.C.C. 598. Sustainable development, the precautionary principle and the polluter-pays principle were received into Indian law in Vellore Citizens' Welfare Forum v. Union of India, (1996) 5 S.C.C. 647. See also India Const. arts. 48A, 51A(g) (directive and fundamental-duty provisions on protection of the environment).

ISSN: 2582-6433

as necessary yet fundamentally inadequate. It is not necessary to come up with a response out of nothing. The duty of directors under Section 166 already includes a requirement to act faithfully towards employees, the community and the environment.

The NGRBC's Principle 6 already calls for the protection and restoration of the environment; and the Constitution already secures equality, a healthy environment and privacy. What is lacking, however, is a unifying principle and the necessary institutional arrangements that will link these obligations to the actual circumstances of algorithmic business such as the disaggregated disclosure of AI data practices and environmental impact, algorithmic impact assessments, human-in-the-loop safeguards, and board-level ethics oversight and the acknowledgement, both in interpretation and in statute, that governing one's algorithms is now an integral part of what it means to be a responsible company.

To regard AI-caused harm as a fundamental aspect of a company's corporate social responsibility and to treat Responsible Innovation as a prerequisite for the legitimate use of such technologies is to ask Indian corporate law to carry out something it has done in the past. that is, to convert a constitutional and ethical aim into a binding standard of corporate behaviour. The other option is letting the most significant technology of the present age function in an area which the law has never before sought to regulate is one that Indian corporate jurisprudence neither has the means to make nor, in a way that is consistent with its basic principles, can justify.