ALGORITHMIC ACCOUNTABILITY BY THE BACK DOOR: THE SIGNIFICANT DATA FIDUCIARY DUE DILIGENCE MANDATE AS INDIA'S DE FACTO AI GOVERNANCE INSTRUMENT

Mr. Aryan Verma, University of Allahabad

ABSTRACT

India has currently decided not to establish a specific law for artificial intelligence. The India AI Governance Guidelines, released by the Ministry of Electronics and Information Technology on November 5, 2025, indicate this choice for a lighter, voluntary framework based on existing laws instead of a comprehensive statute like the one adopted by the European Union. Yet just nine days later, the same government announced the Digital Personal Data Protection Rules, 2025, whose Rule 13(3) imposes a binding duty on Significant Data Fiduciaries to ensure the algorithmic software they deploy does not pose a likely risk to the rights of Data Principals. This article argues that Rule 13(3), though framed as an incidental data protection obligation, functions in practice as India's first binding algorithmic accountability provision. It traces the provision's scope and enforcement design, places it alongside the voluntary AI Guidelines and the risk-tiered EU AI Act, and identifies four gaps in the current framework: the absence of risk classification, a narrow trigger confined to notified fiduciaries and personal data, a due diligence standard that speaks to the regulator rather than the individual, and the lack of any right for an affected person to know that a decision was automated. It closes with suggestions that could be implemented through delegated rule-making, without the need for fresh legislation.

Keywords: Digital Personal Data Protection Rules, 2025; Significant Data Fiduciary; algorithmic accountability; artificial intelligence governance; automated decision-making.

I. Introduction

Two regulatory documents issued within the same fortnight of November 2025 sit somewhat awkwardly beside each other. On 5 November, the Ministry of Electronics and Information Technology unveiled the India AI Governance Guidelines[1], prepared by a committee chaired by Professor Balaraman Ravindran of IIT Madras. The Guidelines state plainly that a separate AI statute is not presently required, and that existing sectoral regulators such as the Reserve Bank of India, SEBI and TRAI will oversee AI within their own domains through what the Guidelines call voluntary measures supported by techno-legal solutions.[2] At the launch, IT Secretary S. Krishnan was reported to have said that India did not intend to lead with regulation, and would rely on existing law wherever it already reached far enough.[3]

Then, on 14 November 2025, the Digital Personal Data Protection Rules, 2025 were notified[4], operationalising the Digital Personal Data Protection Act, 2023 more than two years after it received Presidential assent. Tucked inside Rule 13, which sets out the additional obligations of a Significant Data Fiduciary, is a clause that has attracted far less attention than it probably deserves. Rule 13(3) requires every notified Significant Data Fiduciary to observe due diligence to verify that the algorithmic software it deploys for hosting, displaying, uploading, modifying, publishing, transmitting, storing, updating or sharing personal data is not likely to pose a risk to the rights of Data Principals.[5] A failure to meet this standard is treated as a breach of Section 10 of the Act, which can attract a penalty of up to ₹150 crore.[6]

Read together, the two documents tell an interesting story. In the same fortnight that the government publicly and repeatedly said it did not intend to regulate AI through a dedicated law, it brought into force a rule that does exactly that, at least for a defined slice of the AI landscape, and did so almost without comment. This article makes a simple claim: Rule 13(3), whatever box it has been filed under, is presently the only binding, penalty-backed obligation in Indian law that speaks directly to the safety of algorithmic systems. Understanding it in those terms, rather than as a routine data protection compliance item, is the first step towards fixing

[1]Ministry of Electronics and Information Technology, Government of India, India AI Governance Guidelines (5 November 2025).
[2]ibid, Part II (Issues and Recommendations).
[3]Remarks attributed to IT Secretary S. Krishnan at the launch of the India AI Governance Guidelines, as reported in contemporaneous legal commentary, 12 November 2025.
[4]Digital Personal Data Protection Rules, 2025, notified vide G.S.R. 846(E), 13–14 November 2025.
[5]Digital Personal Data Protection Rules, 2025, r 13(3).
[6]Digital Personal Data Protection Act, 2023, Schedule, Item 4.

what is still missing from it.

II. A Deliberate Choice Against a Standalone AI Law

The AI Guidelines rest on seven guiding principles, described as Sutras, namely Trust, People First, Innovation over Restraint, Fairness and Equity, Accountability, Understandable by Design, and Safety, Resilience and Sustainability, alongside an action plan and practical guidance for industry.[7] Institutionally, the framework is built around an AI Governance Group chaired by the Principal Scientific Adviser, a Technology and Policy Expert Committee, and a proposed AI Safety Institute, none of which is given adjudicatory power over private conduct. Their stated tools are voluntary codes, algorithmic audits and techno-legal mechanisms.[8]

This is not an unreasonable choice, and it places India closer to Japan's light-touch AI statute and the United States' innovation-first posture than to the European Union's prescriptive, risk-tiered AI Act.[9] The Guidelines make this comparison themselves, noting that the EU model carries stricter obligations and steeper penalties for high-risk systems.[10] What the Guidelines propose instead is to lean on existing statutes, principally the Information Technology Act, the Bharatiya Nyaya Sanhita, the Consumer Protection Act and the DPDP Act, supplemented by targeted amendments rather than a new omnibus law.[11]

It is precisely this reliance on the DPDP Act that gives Rule 13(3) its weight. The Guidelines treat the Act mainly as background law relevant to AI through its consent requirements for training data. What they do not say, perhaps because the two instruments were drafted on separate tracks, is that the delegated legislation under that Act already contains a standalone obligation addressed to algorithmic software as such, not merely to the personal data such software happens to touch.

[7] India AI Governance Guidelines (n 2), Part I (the Seven Sutras) and Part III (Action Plan).
[8] ibid, Part II, on the proposed AI Governance Group, Technology and Policy Expert Committee and AI Safety Institute.
[9] For a comparative account, see commentary discussing Japan's Act on Promotion of Research, Development and Utilisation of AI-Related Technologies (May 2025) alongside the United States' voluntary, NIST-anchored approach, each contrasted with the European Union's AI Act.
[10] India AI Governance Guidelines (n 2).
[11] ibid, Part II, discussing amendments to the Information Technology Act, 2000, the Bharatiya Nyaya Sanhita, 2023, the Consumer Protection Act, 2019 and the DPDP Act, 2023 as the preferred route over a new omnibus statute.

III. Rule 13(3): What It Actually Requires

A. Who is covered

The obligation binds only entities notified as Significant Data Fiduciaries. Designation is left to the Central Government under Section 10(1) of the Act, having regard to the volume and sensitivity of personal data processed, the risk to Data Principals, and considerations such as sovereignty and public order.[12] There is no fixed self-assessment threshold; a fiduciary becomes an SDF only once notified.[13] Commentators expect the first cohort to include large social media platforms, e-commerce marketplaces, gaming intermediaries and recommendation-driven consumer platforms.[14]

B. What is required

Rule 13 imposes five obligations on a notified SDF: appointing a resident Data Protection Officer, appointing an independent data auditor, conducting an annual Data Protection Impact Assessment and audit, reporting significant findings to the Data Protection Board of India, and the algorithmic due diligence duty under Rule 13(3), alongside a conditional data localisation requirement in Rule 13(4).[15] The due diligence clause itself is activity based rather than technology based, covering software used for hosting, display, uploading, modification, publishing, transmission, storage, updating or sharing of personal data, a description wide enough to reach recommendation engines, ranking systems, fraud-detection models, targeted advertising tools and content moderation classifiers wherever personal data is part of the pipeline.[16] Notably, the obligation runs upward to the fiduciary's own governance and, eventually, to the Board. It does not, on its own words, require that an affected Data Principal be told that an algorithm was involved in a decision concerning her at all.

[12]Digital Personal Data Protection Act, 2023, s 10(1).
[13]SDF status takes effect from the date of Government notification; there is presently no published self-assessment threshold.
[14]Sector commentary anticipates large social media platforms, e-commerce marketplaces, online gaming intermediaries and search or recommendation driven consumer platforms as the likely first cohort of notified SDFs.
[15]Digital Personal Data Protection Rules, 2025, r 13(1)–(4), read with the Digital Personal Data Protection Act, 2023, s 10(2).
[16]Digital Personal Data Protection Rules, 2025, r 13(3).

C. How it is enforced

A breach of Section 10 obligations can draw a penalty of up to ₹150 crore per instance, determined by the Data Protection Board having regard to the nature and gravity of the breach and the fiduciary's compliance history.[17] This is a civil and administrative penalty rather than a criminal one, imposed by the Board with an appeal lying to the Telecom Disputes Settlement and Appellate Tribunal.[18] Unlike the voluntary codes proposed under the AI Guidelines, this obligation is capable of being formally adjudicated and appealed, which is what gives it real legal bite.

IV. Why This Is, in Substance, an AI Law

Three features together support treating Rule 13(3) as the operative core of India's algorithmic accountability regime, regardless of where it happens to sit in the statute book.

First, it is binding in a way nothing else currently is. The AI Guidelines, the Reserve Bank's FREE-AI framework[19] and similar sectoral guidance all work through voluntary codes and board-approved policies, enforceable only indirectly through general heads of law such as negligence or unfair trade practice. Rule 13(3) creates a direct statutory duty tied to the deployment of algorithmic software itself.

Second, it names its subject precisely. The DPDP Act's general obligations of consent, notice and purpose limitation would apply to personal data processing whether or not an algorithm was involved. Rule 13(3) instead isolates algorithmic software and asks a question that is unmistakably about system safety and fairness, not about data handling as such.

Third, it builds a reporting channel to a regulator. Audit findings under Rule 13, including presumably any adverse algorithmic due diligence findings, must reach the Data Protection Board, giving India the beginnings of a regulator-facing visibility mechanism for algorithmic harm well before the AI Guidelines' own proposed tiered incident reporting system[20] is even

[17]Digital Personal Data Protection Act, 2023, s 33 read with the Schedule, Item 4.
[18]Digital Personal Data Protection Act, 2023, ss 27–29, providing for appeal to the Telecom Disputes Settlement and Appellate Tribunal.
[19]Reserve Bank of India, Report of the Committee for developing a Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI), 13 August 2025.
[20]India AI Governance Guidelines (n 2), Part III, on the proposed tiered AI incident reporting mechanism.

built.

V. Four Gaps Worth Closing

A. No risk classification

Rule 13(3) applies the same undifferentiated standard to every algorithm an SDF runs, whether it decides what appears in a social feed or whether a loan application is approved. This is a marked departure from the EU AI Act's four-tier structure, in which obligations scale with the severity of the use case.[21] An undifferentiated rule risks two failures at once: light-touch checks on genuinely high-stakes systems, and unnecessary audit burden on low-stakes personalisation features.

B. A narrow trigger

Only notified SDFs are bound, and even then only where personal data is involved. A consequential but un-notified fintech lender owes no equivalent duty, and systems that operate on synthetic, aggregated or anonymised data fall outside the rule entirely, however significant their effects on real people. Generative AI outputs sit more naturally within MeitY's proposed amendments to the Intermediary Guidelines on synthetically generated content[22] than within the DPDP Rules at all.

C. A regulator-facing, not a citizen-facing, standard

Rule 13(3) is written as an exercise in internal governance: the fiduciary verifies, the auditor reports, the Board receives. Nothing in the text requires that a Data Principal be told that a decision affecting her was produced by an algorithm at all. Article 22 of the EU General Data Protection Regulation, by contrast, gives data subjects a right not to be subject to a solely automated decision with legal or similarly significant effects, along with a right to human intervention and to contest the outcome.[23] The DPDP Act's own rights of access, correction and grievance redressal contain no comparable right to an explanation of automated decision-making, which leaves a Data Principal with no textual foothold to raise a complaint on her own

[21]Regulation (EU) 2024/1689 (EU AI Act), Title III.
[22]MeitY's draft amendments (22 October 2025) to the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021, addressing the labelling of synthetically generated information.
[23]General Data Protection Regulation (EU) 2016/679, art 22.

initiative.

D. Uncertain designation

Because SDF status depends entirely on a case-by-case government notification, and no published threshold presently exists, unlike the EU AI Act's quantifiable systemic-risk compute thresholds for general purpose models[24], entities and the public alike have little way of knowing in advance which algorithmic deployments are actually covered at any given time.

VI. A Brief Comparative Note

The comparison with the EU AI Act is useful not because India ought to copy it, but because it shows what Rule 13(3) is not. The European model is rights-facing and largely front-loaded: high-risk systems must clear a conformity assessment before deployment, and affected individuals have transparency and, in some cases, explanation rights. Rule 13(3) is fiduciary-facing and largely after the fact. The Reserve Bank's FREE-AI framework moves a little closer to a rights-facing model within finance by recommending explainability disclosures and board-approved product approval processes for AI-driven financial products[25], which suggests that a similar disclosure layer could be added to the DPDP Rules without abandoning India's broader preference for a lighter, sector-led approach.

VII. Suggestions

Each of the four gaps above could plausibly be addressed through delegated rule-making under the DPDP Act itself, consistent with the AI Guidelines' own preference for adapting existing frameworks rather than legislating afresh.

First, risk tiering by notification. MeitY and the Data Protection Board could use their existing powers to identify categories of algorithmic processing, such as those affecting credit, employment, insurance or access to public services, that call for a heightened due diligence standard beyond the general Rule 13(3) baseline.

Second, an interpretive circular connecting the two frameworks. The AI Governance Group

[24]EU AI Act (n 22), art 51, on systemic risk thresholds for general purpose AI models.
[25]Reserve Bank of India, FREE-AI Report (n 20), recommending explainability disclosures and board-approved product approval processes for AI-driven financial products.

could issue guidance explaining how the Seven Sutras, particularly Understandable by Design and Accountability, are meant to inform compliance with Rule 13(3), so that SDFs are not left to build their own due diligence methodology from a five-line rule.

Third, a qualified right to explanation. Sections 11 and 13 of the DPDP Act, or the Rules made under them, could be amended to give a Data Principal subject to a significant, solely automated decision the right to know that fact and to seek a review, on a scaled-down model of Article 22 GDPR that stops short of importing the EU's full conformity assessment machinery.

Fourth, aggregate transparency reporting. Alongside confidential reporting to the Board, SDFs could be asked to publish an annual, non-confidential summary of their algorithmic due diligence findings, in keeping with the AI Guidelines' own emphasis on transparency as a trust-building measure.

Fifth, published designation criteria. Even informal guidance on the factors that inform SDF notification would reduce the present uncertainty facing algorithm-heavy but not-yet-designated entities.

VIII. Conclusion

The debate over AI regulation in India has largely proceeded on the assumption that the country has opted for soft law alone, principles and sandboxes rather than the compulsion of a European-style statute. That assumption, while broadly true of AI-specific instruments, misses something: binding algorithmic accountability law already exists in India, sitting quietly inside Rule 13(3) of the DPDP Rules, 2025. Recognizing it for what it functionally is, the operative core of India's algorithmic accountability regime rather than a data protection footnote, changes the shape of the task ahead. The question is no longer whether India will choose hard law for AI at some future date. For algorithmic systems that touch personal data, it already has, whether by design or by the coincidence of two rule-making processes landing in the same month. What remains is to refine that existing core through risk tiering, a right to explanation, and clearer designation criteria, rather than to treat the matter as settled and look elsewhere for solutions India has, in a limited but real sense, already begun to build.